

A stylized illustration of a human brain, split vertically to show the left and right hemispheres. The brain is rendered in shades of purple and blue, with intricate, swirling patterns that suggest neural activity or complex thought processes. The background is a solid yellow color. The title and author's name are centered over the brain.

# Serotonin in the Bayesian brain

Filip Novický

**DONDERS**  
SERIES

**RADBOD  
UNIVERSITY  
PRESS**

Radboud  
Dissertation  
Series



# Serotonin in the Bayesian brain

Proefschrift ter verkrijging van de graad van doctor  
aan de Radboud Universiteit Nijmegen  
op gezag van de rector magnificus prof. dr. J.M. Sanders,  
volgens besluit van het college voor promoties  
in het openbaar te verdedigen op

donderdag 4 december 2025  
om 16.30 uur precies

door

**Filip Novický**

geboren op 21 december 1997  
te Pilsen, Tsjechië

## **Serotonin in the Bayesian brain**

Filip Novický

### **Radboud Dissertation Series**

ISSN: 2950-2772 (Online); 2950-2780 (Print)

Published by RADBOUD UNIVERSITY PRESS

Postbus 9100, 6500 HA Nijmegen, The Netherlands

[www.radbouduniversitypress.nl](http://www.radbouduniversitypress.nl)

Design: Filip Novický

Cover: Filip Novický

Printing: DPN Rikken/Pumbo

ISBN: 9789465151793

DOI: 10.54195/9789465151793

Free download at: <https://doi.org/10.54195/9789465151793>

© 2025 Filip Novický

**RADBOUD  
UNIVERSITY  
PRESS**

This is an Open Access book published under the terms of Creative Commons Attribution-Noncommercial-NoDerivatives International license (CC BY-NC-ND 4.0). This license allows reusers to copy and distribute the material in any medium or format in unadapted form only, for noncommercial purposes only, and only so long as attribution is given to the creator, see <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

# Serotonin in the Bayesian brain

Dissertation to obtain the degree of doctor  
from Radboud University Nijmegen  
on the authority of the Rector Magnificus prof. dr. J.M. Sanders,  
according to the decision of the Doctorate Board  
to be defended in public on

Thursday, December 4, 2025  
at 4:30pm

by

**Filip Novický**

born on December 21, 1997  
in Pilsen, Czech Republic

**PhD supervisors:**

Prof. dr. F.P. Battaglia

Prof. dr. J.R. Homberg

**PhD co-supervisors:**

Dr. F. Zeldenrust

Dr. P.L. Lanillos Pradas

**Manuscript Committee:**

Prof. dr. F.P. de Lange

Prof. dr. ir. J.H.P. Kwisthout

Prof. dr. J. Limanowski; University of Greifswald, Germany

# Contents

<b>1</b>	<b>General Introduction</b>	<b>8</b>
1.1	Serotonin, a neurotransmitter . . . . .	10
1.1.1	Serotonergic Modulation of Exploration . . . . .	12
1.2	Overview of psychedelics research . . . . .	13
1.2.1	History of psychedelics use . . . . .	13
1.2.2	Serotonin and psychedelics . . . . .	15
1.2.3	Information theory perspective on psychedelics . . . . .	16
1.3	Active Inference and the Free Energy Principle . . . . .	17
1.3.1	Mathematical background . . . . .	18
1.3.2	Expected Free Energy . . . . .	21
1.3.3	Precision-Weighted Inference in Biological Systems . . . . .	22
1.3.4	Unified approach towards psychedelics . . . . .	25
1.4	Summary . . . . .	26
	Bibliography . . . . .	28
<b>2</b>	<b>Active inference, whisking, and serotonin</b>	<b>37</b>
2.1	Introduction . . . . .	38
2.2	Methods . . . . .	40
2.2.1	Active Inference . . . . .	41
2.3	Results . . . . .	48
2.3.1	Face validity . . . . .	48
2.3.2	Serotonergic manipulations . . . . .	49
2.3.3	Temporal analysis of the whisking-amplitude switch . . . . .	56
2.4	Discussion . . . . .	58

2.5	Data Availability Statement . . . . .	63
2.6	Funding . . . . .	64
	Bibliography . . . . .	65
2.7	Appendix . . . . .	72
2.7.1	Initial hidden states . . . . .	72
2.7.2	Likelihood matrices of how data are generated . . . . .	72
2.7.3	Transition Matrices . . . . .	72
2.7.4	Habitual vector . . . . .	73
<b>3</b>	<b>Robotic active tactile sensing</b>	<b>74</b>
3.1	Introduction . . . . .	75
3.2	Methods . . . . .	77
3.2.1	Agent Model . . . . .	78
3.2.2	Active inference . . . . .	78
3.2.3	Precision modulation . . . . .	80
3.2.4	Model tailored to robotic tactile active sensing . . . . .	82
3.2.5	Experimental design . . . . .	83
3.2.6	Implementation . . . . .	83
3.3	Results . . . . .	85
3.3.1	Simulations . . . . .	85
3.3.2	Validation with humanoid . . . . .	87
3.4	Discussion . . . . .	88
3.4.1	Limitation . . . . .	93
3.5	Conclusion . . . . .	93
3.6	Acknowledgement . . . . .	94
	Bibliography . . . . .	95
3.7	Appendix . . . . .	99
3.7.1	Python scripts . . . . .	99
3.7.2	Experiment with different sensory stimulation . . . . .	99
3.7.3	Experiment videos . . . . .	100
<b>4</b>	<b>Precision not Prediction</b>	<b>102</b>
4.1	Introduction . . . . .	103
4.1.1	Previous research on limb illusions and related computational models . . . . .	105



4.1.2	A unified body-ownership model based on precision . . . . .	106
4.2	Methods . . . . .	109
4.2.1	Problem formulation and parameters . . . . .	109
4.2.2	Maximum posterior precision body-ownership model . . . . .	110
4.2.3	Evaluative measures for the model . . . . .	115
4.3	Results . . . . .	116
4.3.1	Body-ownership illusion: model selection based on the un- certainty minimization . . . . .	117
4.3.2	Estimating noise through precision adaptation . . . . .	117
4.3.3	Validating the model via breaking precision adaptation . . . . .	118
4.3.4	Dependence of the illusion on inter-hand distance . . . . .	119
4.3.5	Proprioceptive drift in competing models . . . . .	120
4.4	Discussion . . . . .	124
4.4.1	Potential neural underpinnings of the model . . . . .	125
4.4.2	Alignment with other uncertainty-based models . . . . .	125
4.4.3	Testing our model and limitations . . . . .	126
4.4.4	Future work . . . . .	128
	Bibliography . . . . .	130
4.5	Appendix . . . . .	136
4.5.1	Related computational models of body illusions . . . . .	136
4.5.2	Precision parametrization . . . . .	138
4.5.3	Code and parameters used for simulations . . . . .	138
<b>5</b>	<b>Reclaiming saliency</b>	<b>140</b>
5.1	Introduction . . . . .	141
5.2	Attention and salience in neuroscience . . . . .	143
5.2.1	Attention as neural gain control . . . . .	144
5.2.2	Salience as uncertainty minimisation . . . . .	145
5.2.3	Rhythmic coupling of attention and salience . . . . .	146
5.3	Proposed precision-modulated account . . . . .	148
5.4	Precision-based attention for Robotics . . . . .	149
5.4.1	Previous brain-inspired attention models in robotics . . . . .	151
5.4.2	Precision-modulated perception . . . . .	152
5.4.3	Precision-modulated action . . . . .	162

5.5	Concluding remarks . . . . .	168
	Bibliography . . . . .	171
<b>6</b>	<b>Psilocybin Accelerates EEG Microstate Transitions and Elevates Approximate Entropy</b>	<b>182</b>
6.1	Introduction . . . . .	183
6.2	Materials and methods . . . . .	185
6.2.1	Participants and procedure . . . . .	187
6.2.2	EEG Data Acquisition . . . . .	187
6.2.3	Preprocessing . . . . .	187
6.2.4	Analysis Methods . . . . .	188
6.2.5	Statistical analyses . . . . .	193
6.3	Results . . . . .	193
6.3.1	Decomposing the EEG signal into microstates . . . . .	194
6.3.2	No significant difference between MBCT and non-MBCT groups . . . . .	194
6.3.3	Sub-conditions reveal changing patterns of attentional states	195
6.3.4	The effects of psilocybin . . . . .	196
6.3.5	Transition and Coverage . . . . .	197
6.3.6	Approximate entropy . . . . .	199
6.4	Discussion . . . . .	199
6.4.1	Limitations . . . . .	203
6.4.2	Future directions . . . . .	204
6.5	Conclusion . . . . .	205
	Bibliography . . . . .	206
6.6	Supplementary materials . . . . .	213
6.6.1	Demographics of participants . . . . .	213
6.6.2	Details on Approximate Entropy . . . . .	213
6.6.3	Statistical comparisons . . . . .	215
<b>7</b>	<b>General Discussion</b>	<b>221</b>
7.1	Complementary serotonergic receptors . . . . .	222
7.2	Precision and body ownership . . . . .	225
7.3	Serotonin and body ownership . . . . .	226
7.4	Other computational models of serotonin . . . . .	227

7.5	Limitations . . . . .	229
7.6	Conclusion and future directions . . . . .	230
	Bibliography . . . . .	231
<b>8</b>	<b>Appendices</b>	<b>236</b>
8.1	Summary . . . . .	236
8.2	Samenvatting . . . . .	238
8.3	Research Data Management . . . . .	240
8.3	Research Data Management Plan . . . . .	240
8.4	Curriculum Vitae . . . . .	241
8.5	Portfolio . . . . .	242
8.6	List of Publications . . . . .	243
8.7	Acknowledgment . . . . .	243
8.8	Donders Graduate School . . . . .	246

# Chapter 1

## General Introduction

The brain faces a fundamental challenge: how to make sense of ambiguous sensory information while maintaining adaptive behaviour in an uncertain world. This challenge becomes particularly apparent when studying neuromodulatory systems, such as serotonin, which appear to globally influence information processing without being directly responsible for any single function. Understanding how neuromodulators shape perception and action requires theoretical frameworks that can bridge multiple levels of analysis—from molecular mechanisms to behaviour.

Neuroscience has recently witnessed significant theoretical developments suggesting that the brain employs a form of Bayesian inference to solve this challenge (Clark, 2013; Hohwy, 2013; Clark, 2015). Bayesian inference is the process of inverting a probabilistic generative model of how data are generated. The idea is that, if one understands how measurable signals are caused by unobserved states of the world, it should be possible to reverse engineer the likely causes from those measurements. This comes with particularly interesting interpretations of the brain function, and also perception per se, as perception is conceptualised as an inference or ‘the best guess’ as to the causes of sensory input (Hohwy, 2013). There are several interpretations of how Bayesian computation is achievable in the brain. For instance, predictive processing (PP) proposes that the brain’s function is to maintain a hierarchical model of the world by generating predictions about sensory input (Clark, 2015). These predictions are compared with actual sensory signals, creating prediction errors that update the brain’s beliefs. Active inference goes be-

yond this by explaining how organisms act to test their predictions, while standard Bayesian inference only describes how they update beliefs from observations.

No matter which of these Bayesian theories one selects, all of them provide a strong mathematical underpinning of how the brain could achieve its sophisticated behavioural and perceptual output. However, how such Bayesian interpretations of brain functions are implemented in the brain, and how these models can be linked to specific neural recordings, is still an open question.

As the brain is a complex organ with many neurons, brain regions, blood vessels and neurotransmitters, all of them affect and define the structure or the function of the brain, it will take many years to come to provide a strong Bayesian model that will well integrate the Bayesian inference and the biological structure of the brain. This thesis aims to help with this project by providing a possible homologue of serotonin under the Bayesian brain hypothesis. It is said that 'serotonin is involved in everything but responsible for nothing' (Carhart-Harris and Nutt, 2017; Müller and Homberg, 2015).

Nevertheless, this thesis investigates how serotonin shapes the way organisms explore and perceive their environment. To understand this, we employ the concept of "precision", which regulates how much weight or confidence the brain assigns to different sources of information. Consider this: when you are in a noisy room, your brain reduces the precision (confidence) it assigns to auditory information; when you are in the dark, it reduces visual precision. This dynamic weighting is crucial for adaptive behaviour. We propose that serotonin acts as a key regulator of this precision-weighting process <sup>1</sup>. To test this hypothesis, we examined serotonergic function through three complementary approaches: observing exploratory behaviour in animals, implementing similar mechanisms in a humanoid robot, and analyzing the effects of psychedelics (which act on serotonin receptors) in humans. In addition, the thesis continues on extending how precision can be generally used during body perception and attention.

The introduction below is structured as follows: First, the basics of serotonergic research and the core knowledge about this neurotransmitter are provided, establishing the neurobiological foundation necessary to understand the subsequent sections. The review on psychedelics—pharmacological agents that specifically

---

<sup>1</sup>Note that throughout the text, precision represents the confidence or certainty assigned by an agent and not the external precision of the environmental sensory input, unless stated otherwise.


target the serotonergic system—follows, offering unique insights into serotonin’s functional roles. Lastly, the technical section on active inference and the free energy principle describes the computational framework that will be used throughout this thesis to model and interpret serotonergic function, particularly the hypothesis that serotonin modulates sensory precision.

## 1.1 Serotonin, a neurotransmitter

Serotonin, also known as 5-hydroxytryptamine (5-HT), is a monoamine neurotransmitter that plays crucial roles both in the central nervous system (CNS) as well as peripheral nervous system (PNS) Julius (1991). Originally identified for its vasoconstrictive properties in blood serum, serotonin has emerged as one of the most extensively studied neurotransmitters due to its involvement in diverse physiological processes and pathological conditions Rapport et al. (1949); Berger et al. (2009). What is interesting about serotonergic neurons when it comes to the organization of neural structure, their presence comprises only about 0.0001% of all neurons, but they form 1 out of 500 axonal connections(Jacobs and Azmitia, 1992).

The molecular structure of serotonin consists of an indole ring structure with a hydroxyl group and an amino group. This is visualized in Figure 1.1. Although extensively studied and associated with many functions, the molecule appears to have a complex role, as it is not solely responsible for any single function (Julius, 1991). This complexity might be related to the multiple receptor subtypes, classified as 5HT1 through 5HT7, each with distinct molecular structures, signaling mechanisms, and tissue distributions Hoyer et al. (1994); Pytliak et al. (2011). Except for 5HT3 receptor, which functions as a ligand gated ion channel, all the receptors belong to the G protein-coupled receptor group (Julius, 1991; Hen, 1993).

Within the CNS, most of the serotonergic neurons are found to originate from the raphe nuclei located in the brainstem. These neurons project extensively throughout the brain, where they regulate a wide array of functions (Lesch and Waider, 2012). Among many, these include mood, cognition, sleep, appetite, and pain perception (Jacobs and Azmitia, 1992). In peripheral tissues, serotonin modulates various physiological processes, including gastrointestinal motility, platelet aggregation, and vascular tone (Berger et al., 2009). One can see from the list of functions that



**Figure 1.1:** A molecular structure of serotonin.

we are dealing with a molecule with diverse physiological roles. In addition to this list, dysregulation of serotonergic signaling has been implicated in numerous psychological conditions, including major depressive disorder, anxiety disorders, and schizophrenia (Meltzer, 2012). Moreover, the serotonin system represents a major psychotherapeutic target, with drugs modulating serotonergic neurotransmission being among the most prescribed medications worldwide (Wong et al., 2005).

In addition to the studies above, there are also other ways to study the function of serotonin. Therefore, next to studying the effects of activating different serotonin receptors, another approach to better understand serotonergic function is to develop genetic models that manipulate key components of the serotonin system. Particularly relevant to this thesis are knockout models targeting the serotonin transporter (SERT) and tryptophan hydroxylase 2 (TPH2). SERT is responsible for the serotonergic reuptake from the synaptic cleft, effectively terminating serotonergic signaling (Demchyshyn et al., 1994). SERT knockout animals ( $SERT^{-/-}$ ) consequently exhibit chronically elevated extracellular serotonin levels throughout development and adulthood (Bengel et al., 1998).  $SERT^{-/-}$  rodents typically exhibit heightened anxiety-like behaviours, reduced exploratory activity, and enhanced sensitivity to environmental stimuli (Homberg, 2012). On the other hand, TPH2 knockout animals are characterized by a nearly complete absence of serotonin in the central nervous system, as ( $TPH2^{-/-}$ ) is the rate-limiting enzyme in brain serotonin synthesis (Gutknecht et al., 2012). Thus, these genetic models display distinctive phenotypes that illuminate serotonin's role in information processing.

Studies of serotonin transporter (SERT) polymorphisms have revealed that variations in serotonergic function are associated with differences in environmental sensitivity and exploratory behaviour (Caspi et al., 2010).  $SERT^{-/-}$  animals typically show enhanced sensitivity to both positive and negative environmental inputs (Homberg et al., 2016). This differential susceptibility suggests that serotonin helps calibrate how organisms sample and learn from their environment (Belsky and Pluess, 2009).

### 1.1.1 Serotonergic Modulation of Exploration

Although the list of serotonergic functions is already extensive, this thesis focuses on the less studied modulatory power of serotonin, which is regulating exploratory behaviour and environmental sensitivity. Research has shown that serotonin influences how biological agents process and respond to environmental stimuli across multiple timescales (Homberg et al., 2016). At a basic level, serotonin affects sensory gain control, modulating the signal-to-noise ratio in interpreting sensory information (Palacios et al., 1991). This mechanism may help organisms appropriately tune their sensitivity to environmental cues during exploration.

Recent work has begun to bridge cellular and behavioural levels by examining how serotonergic circuits implement exploration. Optogenetic studies have shown that activation of specific serotonergic pathways can promote exploratory behaviour, while others favor exploitation of current resources (Lottem et al., 2018; Doya et al., 2021). This circuit-level organization may provide a flexible mechanism for regulating the exploration-exploitation trade-off based on environmental conditions and internal states.

Computational models incorporating serotonergic function have helped formalize these findings on exploration (Daw et al., 2002). Such models typically cast the functional role of serotonin as a modulator of either the variability of action selection or the learning rate from environmental feedback (Cools et al., 2011). However, more recent active inference accounts suggest serotonin—specifically found to be activated by psychedelic drugs—may fundamentally regulate the precision of sensory evidence relative to prior beliefs, thereby controlling the acquisition of new information (Carhart-Harris and Friston, 2019).

In conclusion, the serotonin system plays a role in many functions of the brain.



A part of this thesis specifically addresses the serotonergic role in exploration and regulating environmental sensitivity.

## 1.2 Overview of psychedelics research

Having reviewed the complex role of serotonin in neural function, we now turn to examine a class of compounds that primarily act through the serotonergic system. Psychedelics are a class of psychoactive compounds that primarily act as 5HT<sub>2A</sub> agonists, producing profound alterations in perception of the world (Carhart-Harris and Friston, 2019). These substances, which include compounds such as psilocybin, lysergic acid diethylamide (LSD), and dimethyltryptamine (DMT) are characterized by their ability to alter sensory processing and brain activity, such as default mode network activity (Carhart-Harris et al., 2014a). Unlike other psychoactive substances, classical psychedelics are distinguished by their non-addictive properties and their capacity to induce non-ordinary states of consciousness while maintaining physiological safety under controlled conditions (Nichols et al., 2017).

This part of the introduction thus provides a detailed look into the current state of classical psychedelics research, starting from traditional practices to modern scientific applications.

### 1.2.1 History of psychedelics use

The history of consuming psychedelic substances is thought to be deeply rooted in human culture, with evidence spanning diverse civilizations and continents. Archaeological findings suggest the use of psychedelic compounds in spiritual and healing practices, and it is assumed that early humans were surrounded by psychedelics for more than a million years (Winkelman, 2019). However, the first documented use of psychedelics dates back thousands of years, with substances like peyote having a documented use from over 5,700 years ago (El-Seedi et al., 2005). These traditional practices were integrated into religious ceremonies, healing rituals, and rites of passage, which shows sophisticated early understanding of the effects of these compounds into the society.

Although the study of psychedelics has already emerged in the 19th century,

the modern scientific era of psychedelic research is usually considered from the 1940s onwards, marked by Albert Hofmann’s discovery of LSD’s psychoactive properties (Sessa, 2016). This discovery catalyzed an intensive scientific interest in psychedelics’ therapeutic applications. Throughout the 1950s and early 1960s, numerous clinical trials investigated LSD as well as psilocybin in treating various psychological conditions.

Initially, researchers and doctors thought about psychedelics from the contemporary theories at that time, specifically from psychoanalysis. Researchers like Stanislav Grof integrated psychedelic experiences into the psychoanalytic framework, proposing that these compounds could accelerate therapeutic processes of addiction (Abramson, 1966), depression (Savage et al., 1964) or even dystonia (Stewart et al., 2020) by making ‘unconscious material—the jargon of the psychoanalysts—more accessible (Grof and Halifax, 1977).

As research methodologies evolved, a more cognitive perspective emerged in the 1990s and 2000s with the revival of psychedelic research. This approach focused on understanding the effects of psychedelics on specific mental processes and brain functions, such as illusion or mystical experiences and neuroplasticity through rigorous experimental protocols (Vollenweider et al., 1998; Vollenweider and Kometer, 2010; Strassman, 2000). This was an important shift, as the studies on psychedelics before they were banned in the early 1970s were often criticized for poor sample sizes and unreliable results (for instance, see Stewart et al. (2020)). These studies found that psychedelics consistently enhance emotional openness, reduce cognitive rigidity, and promote novel thought patterns, suggesting fundamental alterations in information processing (Strassman, 2000; Carhart-Harris and Friston, 2019).

The most recent perspective in this field has come with the emergence of computational approaches, integrating modern neuroscience with sophisticated mathematical models to understand how psychedelics alter information processing in the brain. Contemporary researchers employ tools from network science or statistical physics to characterize the brain’s altered states under psychedelic influence. These approaches have led to novel theories such as the entropic brain hypothesis and the REBUS model (Carhart-Harris et al., 2014b, 2018; Carhart-Harris and Friston, 2019). These theories will be discussed in Section 1.3.4.

### 1.2.2 Serotonin and psychedelics

Throughout the recent decades of psychedelics research, scientific rigour regarding the research into psychedelic action has resulted in several tangible observations that have advanced an understanding of their neurochemical mechanisms. A crucial breakthrough came with the identification of the 5HT<sub>2A</sub> receptor as the primary target for classical psychedelics (Glennon et al., 1984; Hintzen and Passie, 2010). This finding was further supported via pharmacological blocking studies, where researchers demonstrated that ketanserin, a selective 5HT<sub>2A</sub> antagonist, prevents the subjective effects of psychedelics (Vollenweider et al., 1998). This direct causal relationship between receptor binding and psychedelic effects established a clear molecular basis for their action. For instance, the 5HT<sub>2A</sub> receptors have been found in the dendrites in layer 5 (Weber and Andrade, 2010). These neurons are particularly significant as they are known to generate neural alpha oscillations (around 8-12Hz) (Silva et al., 1991; Sun and Dan, 2009).

Alpha oscillations are quite important in the psychedelics research community, as this activity is thought to play a crucial role of how the information propagates in the brain by implementing selective inhibition, effectively filtering out irrelevant sensory information while facilitating the processing of task-relevant inputs (Jensen and Mazaheri, 2010). Further studies in humans have shown that psilocybin administration leads to a marked decrease in alpha waves (Carhart-Harris et al., 2014b; Valle et al., 2016; Pallavicini et al., 2019), suggesting a disruption of this filtering mechanism. This reduction in alpha power may explain the enhanced sensory processing often reported during psychedelic states (i.e., perceiving vivid colors, intricate geometric patterns, or inducing synesthesia), as the usual constraints on sensory information flow are relaxed (Carhart-Harris et al., 2014b).

Neuroimaging studies on humans have revealed significant alterations in key brain networks under psychedelic influences, particularly affecting the Default Mode Network (DMN), putamen, and thalamus (Carhart-Harris et al., 2012). The laminar organization of cortical connections is likely significant in these network-level effects, as psychedelics primarily target 5HT<sub>2A</sub> receptors on layer 5 pyramidal neurons which have specific projection patterns to thalamic nuclei (Shipp, 2007). Layer-specific cortical connections form distinct ascending and descending pathways that normally regulate information flow between cortical areas and subcortical structures, providing an anatomical basis for the network alterations. The DMN

holds special significance in this context, as it is widely considered to represent the highest hierarchical level in cortical processing (Carhart-Harris and Friston, 2010; Smallwood et al., 2021). This network’s interaction with sensory input streams has also been carefully mapped (Sepulcre et al., 2012).

### 1.2.3 Information theory perspective on psychedelics

One of the emerging and quite important hypotheses focuses on the more entropic state of the brain on psychedelics. In a series of theories by Carhart-Harris and his colleagues, the idea of an entropic brain under psychedelics was introduced (Carhart-Harris et al., 2014b, 2018). In these papers, the core idea was that psychedelics—with its metabolites attaching to the 5HT<sub>2A</sub> receptors in the brain—increase disordered behaviour of neuronal activity. This coincides with an increase in the brain’s criticality (Beggs and Plenz, 2003; Beggs, 2022). The concept of criticality in neural systems describes a state where neural activity patterns exhibit scale-free behaviour, characterized by power law distributions of activity in both spatial and temporal domains. At this critical point, the brain demonstrates optimal computational capabilities: ordered enough to maintain stable information processing, yet sufficiently flexible to respond to new inputs. This state maximizes several information-theoretic measures including dynamic range, information transmission, and information capacity (Shew and Plenz, 2013). The disruption of the critical state in various neurological conditions suggests that maintenance of criticality is essential for normal brain function, providing a framework for understanding how the brain achieves its balance between stability and adaptability (Meisel, 2020).

Recent evidence supports the aforementioned theoretical proposal, with studies showing that psychedelics tune the brain even closer to criticality than is evident in normal waking consciousness, particularly demonstrated through LSD (Atasoy et al., 2017; Varley et al., 2020; Herzog et al., 2023). These findings align with broader research on critical connectivity in neural systems, suggesting that psychedelics may optimize the brain’s capacity for information processing by maintaining it at a critical point between order and chaos. This optimization at criticality could explain the enhanced perceptual sensitivity observed in the psychedelic state (Carhart-Harris et al., 2014a; Carhart-Harris, 2018).

The goal of this section is to give a background on the current understanding and progress of psychedelic research. The goal here is to unite this field within one unified approach: active inference. To do that, we first have to provide the introduction to this theory, which follows in the next Section. At the end of this section, in 1.3.4, we provide a comprehensive way of approaching psychedelics from the active inference perspective.

### 1.3 Active Inference and the Free Energy Principle

Serotonin research is a highly interdisciplinary field that draws from neurophysiology, behavioural neuroscience, and psychopharmacology. To articulate hypotheses that speak to each of these areas, it is useful to appeal to a theoretical framework that is applicable across these domains. Active Inference is one such approach that bridges different perspectives. Crucially, the mathematical formulations that underwrite Active Inference have a dual interpretation as descriptions of both perceptual inferences and the physiological processes that support those inferences. Each of these is ultimately interpretable in terms of the actions they cause, i.e., the movement of parts of one's body. The benefit of this theoretical and mathematical framing for this thesis is (1) that it allows one to link together concepts like synaptic efficacy, attention, and neuromodulation in a common language; and (2) that this framing is explicitly embodied. The second of these is something that will be particularly relevant when dealing with both the manipulation of parts of the body to gain information about the world and with the distinction between self and other implicit in the body ownership illusions dealt with in Chapter 4. Active Inference does not give us any answers as to the role of serotonin in the brain, but it does allow one to ask questions and pose hypotheses. Each hypothesis then has implications for what we would measure across multiple empirical domains.

The free energy principle (FEP) and its corollary theory active inference follow from quite a simple observation, which is that humans and other living systems maintain their physical structure (Maturana and Varela, 2012) and thus combat the second law of thermodynamics—the tendency for entropy to increase over time. Since cognitive processes are ultimately realized through physical systems (i.e., neural activity in the brain), they too must operate within the same thermodynamic constraints. This suggests that cognitive functions may be understood as specialized

processes that help living systems maintain their physical integrity (i.e., survival). Simply put, preservation of self over time is the key imperative for a living system and cognition emerges as an adaptive mechanism that supports this imperative.

The FEP postulates that living systems must, by their very nature, resist the universe’s tendency toward disorganization (Friston and Stephan, 2007). A living organism that fails to maintain its essential functions or parameters within viable bounds will ultimately be unable to survive. Simply, the system would disintegrate. This builds on a theory where life itself can be understood as a process of continuous prediction and self-organization. In other words, the system is organized by and for minimizing free energy (a measure of the difference between an organism’s model of the world and reality), as the better prediction about the world leads to a more sustained functioning in it (e.g., having a good prediction of what will happen after falling off a cliff makes one not to fall). The question is, how is this realized in the biological body?

Active inference aims to explain how creatures resist a tendency to dispersion (i.e., increases in entropy) (Parr et al., 2022). Living systems actively probe and shape the environment by acting on it rather than passively reacting to it. This leads to the construction of internal models that help them predict and navigate their world. This process manifests across scales, from bacteria to human cognition, as has been argued in (Kirchhoff et al., 2018), providing a framework across species and essentially uniting life under one umbrella, trying to answer Schrodinger’s question ‘what is life?’ (Schrodinger, 1946). The revolutionary idea of this framework is that action and perception are accounted in terms of optimisation of the same quantity. The mathematics underlying FEP in active inference assumes that living systems are, in essence, performing Bayesian inference (Clark, 2013; Hohwy, 2013).

### 1.3.1 Mathematical background

Before examining the mathematical details, it is important to note that active inference can be formulated in both discrete-time and continuous-time frameworks (Lanillos et al., 2021). Continuous-time models utilize differential equations to represent the dynamics of belief updating and can capture the smooth, ongoing nature of biological processes. In contrast, discrete-time models represent updating at specific time points and often use categorical probability distributions (i.e., a

distribution assigning probabilities to a finite set of discrete outcomes), making them mathematically more tractable and computationally less demanding. This section explains the discrete state models of active inference, as they are mathematically more straightforward to comprehend compared to the continuous models, and in the end the core idea of active inference is preserved (Da Costa et al., 2020).

In the FEP interpretation of the world, an organism observes the world via observations  $o$ , such as visual, tactile or auditory stimulation, but has no access to the (hidden) states of the world  $s$  that cause these experiences. The organism’s aim is thus to reconstruct the hidden states by their observable consequences of the state. To do a proper reconstruction, one needs to know the prior, or the previous knowledge about the inferred state. This is either through previous experience or genetic coding. In addition, a likelihood is also necessary, as it describes a relationship between a given observation to the inferred state, for all possible states. The mismatch between observation and prediction is known as prediction error in the literature (Clark, 2013). Prediction error is often used as a proxy for free energy, a tractable measure to perform Bayesian inference (i.e., a bound on how two distributions are alike). We can write the free energy as the discrepancy between the internal model and its environment, and how close is the agent to fulfill their preferences (or prior expectations) (Da Costa et al., 2020):

$$F = \underbrace{D_{KL}[Q(s)||P(s)]}_{\text{Complexity}} - \underbrace{E_{Q(s)}[\ln P(o|s)]}_{\text{Accuracy}} \quad (1.1)$$

where  $Q(s)$  represents approximate posterior beliefs about the hidden states  $s$ , and  $P(s)$  represents prior beliefs about those states. The free energy  $F$  comprises two main terms: a complexity term  $D_{KL}[Q(s)||P(s)]$  that quantifies the Kullback-Leibler divergence between posterior and prior beliefs, and an accuracy term  $E_{Q(s)}[\ln P(o|s)]$  that measures how well the model explains observed sensory data  $o$ , for its expected (or average) value  $E$ . Minimising the complexity term prevents the posterior from diverging too far from the prior. This is because the bigger the difference between  $Q(s)$  and  $P(s)$ , the higher the complexity value and thus the higher the free energy (Da Costa et al., 2020). Minimising the accuracy term, expressed as an expected log-likelihood under the approximate posterior, ensures the model can adequately explain sensory observations. Meaning, the beliefs about states that can explain the observations the best, will maximize the accuracy value

while incurring a minimal complexity cost, leading to a smaller free energy in the end.

For a different interpretation of the FEP for a better clarity of what the principle represents, one can reconstruct the equation above in order to comprehend what this principle is about in its core. This is possible by scrutinizing the Kullback-Leibler divergence  $D_{KL}$  in Equations (1.2) and (1.3):

$$D_{KL}[Q(s)||P(s|o)] = D_{KL}[Q(s)||P(s)] - E_{Q(s)}[\ln P(o|s)] + \ln P(o) \quad (1.2)$$

$$D_{KL}[Q(s)||P(s)] - E_{Q(s)}[\ln P(o|s)] = D_{KL}[Q(s)||P(s|o)] - \ln P(o) \quad (1.3)$$

Hence, an alternative formulation of the free energy principle expresses  $F$  in terms of the posterior divergence and surprisal (defined as  $-\ln P(o)$ ), where  $D_{KL}[Q(s)||P(s|o)]$  represents the divergence between the approximate posterior  $Q(s)$  and the true posterior  $P(s|o)$ :

$$F = \underbrace{D_{KL}[Q(s)||P(s|o)]}_{\text{Posterior divergence}} - \underbrace{\ln P(o)}_{\text{Surprisal}} \quad (1.4)$$

$$F \geq -\ln P(o) \quad (1.5)$$

The KL divergence in Equation (1.4) is by definition non-negative ( $\geq 0$ ) because of Jensen's inequality. The second term,  $\ln P(o)$ , represents surprisal (or self-information), which quantifies the unexpectedness of observations. This decomposition reveals that minimizing free energy is equivalent to minimizing the divergence between approximate and true posteriors while simultaneously minimizing surprisal. From here, it is also apparent why the free energy is considered a lower bound on surprisal, as highlighted in the Equation (1.5). This is because if the posterior divergence in Equation (1.4) is equal to zero, then free energy is equal to surprisal. And since the first term can only be positive, the free energy will always be the upper bound of surprisal. Importantly, this formulation is mathematically equivalent to the previous expression involving complexity and accuracy terms, but it emphasizes



different aspects of the optimization problem. While the complexity-accuracy decomposition in Equation (1.1) highlights the trade-off between model complexity and fit, this posterior-surprisal decomposition in Equation (1.4) emphasizes how organisms attempt to minimize the unexpectedness of their sensory inputs while maintaining accurate internal models.

### 1.3.2 Expected Free Energy

A key innovation of active inference is its treatment of action selection (Friston, 2009). Instead of generating movement from the free energy per se, an agent may consider all possible actions and select the one that will minimize the free energy expected at a future time, given the current beliefs. The selection of future actions refers to planning. It is the plans that depend upon expected free energy. Mathematically, this is done through the minimization of expected free energy  $G$ , which extends the free energy principle into the future (Da Costa et al., 2020):

$$G(\pi) = \underbrace{D_{KL}[Q(s_\tau|\pi)||P(s_\tau)]}_{\text{Risk}} + \underbrace{E_{Q(s_\tau|\pi)P(o_\tau|s_\tau)}[\ln P(o_\tau|s_\tau)]}_{\text{Ambiguity}} \quad (1.6)$$

In this formulation,  $G(\pi)$  reflects the expected free energy for a particular policy (i.e., a set of actions)  $\pi$ . This quantity comprises two fundamental terms: risk and ambiguity. The risk term, expressed as  $D_{KL}[Q(s_\tau|\pi)||P(s_\tau)]$ , quantifies the divergence between predicted and preferred states, where  $Q(s_\tau|\pi)$  represents beliefs about future states at time point  $\tau$  under a given policy, and  $P(s_\tau)$  encodes prior preferences about future states at the same time point. The ambiguity term,  $E_{Q(s_\tau|\pi)P(o_\tau|s_\tau)}[\ln P(o_\tau|s_\tau)]$  represents the expected uncertainty in observations given future states. Uncertainty in this context refers to the variability or unpredictability in the mapping between states and observations.

This formulation elegantly captures two fundamental aspects of decision making. First, it accounts for exploratory behaviour through the minimization of ambiguity, driving organisms to seek out situations where they can make precise predictions about sensory inputs. Second, it explains goal-directed behaviour through the minimization of risk, ensuring that selected actions align with preferred outcomes. This dual optimization creates a natural balance between exploratory and exploitative behaviours, providing a principled mathematical framework for understanding

how organisms make decisions in uncertain conditions while maintaining their homeostatic objectives. It is worth noting that this framework naturally accommodates the existence of multiple locally optimal solutions in the expected free energy landscape. Rather than requiring a single globally optimal solution, the theory incorporates precision parameters that effectively modulate how deterministically an organism commits to particular actions or what is perceived.

### 1.3.3 Precision-Weighted Inference in Biological Systems

A complete account of biological self-organization under the Bayesian inference view requires understanding how systems weight the reliability of different information sources. In probabilistic inference, precision refers to the inverse variance of a distribution. In other words, higher precision indicates greater certainty and reliability of the information and vice versa. Precision weighting is thus a mechanism by which the brain assigns different levels of confidence to various sources of information, effectively determining how much influence each source should have on belief updating. For instance, when sensory inputs are deemed more precise (reliable), they exert greater influence on posterior beliefs. Conversely, when prior beliefs are considered more precise, they dominate the inference process. This weighting of information based on its reliability is essential for optimal Bayesian inference, as it allows organisms to adaptively respond to changing environmental conditions and internal states. To understand how precision weighting is achieved mathematically, it is important to understand the generative model construction first. In active inference, we begin by specifying a generative model that describes how observations are generated from hidden states:

$$p(o, s | \pi) = p(\pi) \prod_{t=1}^T p(o_t | s_t) p(s_t | s_{t-1}, \pi) \quad (1.7)$$

where the joint probability of observations  $o$  and states  $s$  under policy  $\pi$  is decomposed into a product of likelihood terms  $p(o_t | s_t)$  and state transitions  $p(s_t | s_{t-1}, \pi)$ . In active inference, the former is formalized as matrix A, while the latter as matrix B. Note that this formulation is for discrete state space models.

This generative model specifies the probabilistic relationship between states and observations. However, to perform inference under this model, biological systems

must account for the reliability or precision of different information sources. In the active inference framework, this is achieved by introducing precision parameters that modulate the influence of different terms in the model.

When we perform inference under this generative model, the precision parameters enter through the softmax function, which converts log-probabilities into normalized probability distributions. The precision parameters  $(\omega, \zeta, \gamma)$  act as inverse temperature parameters that control how sharply peaked these distributions are (Novicky et al., 2024).

Here,  $B$  represents the state transition matrix encoding  $p(s_t|s_{t-1})$ , and  $A$  represents the likelihood matrix encoding  $p(o_t|s_t)$ , as explained above. The precision parameters modulate these relationships as follows:

$$p(s_t|s_{t-1}) = \sigma(\omega \cdot \ln B(s_t|s_{t-1})) \quad (1.8)$$

$$p(o_t|s_t) = \sigma(\zeta \cdot \ln A(o_t|s_t)) \quad (1.9)$$

$$Q(\pi) = \sigma(-\gamma \cdot G(\pi)) \quad (1.10)$$

These equations show how precision parameters transform the log-probabilities from the generative model into working distributions for inference. Higher precision values  $(\omega, \zeta, \gamma)$  lead to more deterministic mappings, while lower values create more uniform distributions. This mathematical formulation allows biological systems to flexibly weight different information sources based on their estimated reliability.

This mathematical framework reveals that precision has an important role, which varies according to the distribution with which it is associated. If a prior is held with high precision, it is resistant to change, and posterior probabilities will be very similar to priors. In contrast, if a likelihood has very high precision, this may support a large belief update from prior to posterior following sensory observations.

Once we consider planning and action selection, and score them based upon prior and (anticipated) posterior beliefs (as in Equation 1.6), the influence of prior and likelihood precisions has a profound impact on which policies might lead to the greatest risk or ambiguity. In other words, selecting a course of action based

upon our current beliefs means behaviour is highly sensitive to our beliefs and to the precisions that help to determine them.

At the computational level, precision weights the relative influence of different prediction errors during belief updating, effectively modulating the balance between prior beliefs and incoming sensory evidence. This precision weighting mechanism explains how organisms selectively process relevant information while ignoring irrelevant sensory inputs (Parr and Friston, 2019).

### 1.3.3.1 Precision in Attention and Body Perception

Above we explored how precision—defined as the inverse of variance in probability distributions—operates within the FEP framework. We showed how an agent’s estimate of precision can influence action selection to minimize expected free energy. Here, we will examine how this precision-weighting mechanism provides a unifying framework for understanding both attentional processes and body perception phenomena.

Within the active inference framework, both attention and perceptual illusions can be understood through the precision weighting mechanisms described in the previous section. The framework suggests that attention—conceptualized as the process by which certain sensory channels or spatial locations are selectively prioritized—effectively modulates the precision of prediction errors, determining how much influence different sensory channels have on belief updating (Feldman and Friston, 2010; Brown et al., 2013; Kanai et al., 2015; Kok et al., 2012). For instance, when attending to visual input, the precision of visual prediction errors for a specific region of the visual field increases relative to others, leading to the enhanced influence of this particular space on perceptual inference (Hillyard et al., 1998).

Body ownership illusions like the Rubber Hand Illusion (RHI)—where a participant perceives a rubber hand as their own when both hands are simultaneously stroked and the real hand is covered, creating the illusion that the touch is felt at the rubber hand—illustrate how precision weighting shapes our experience of embodiment. Under the active inference framework, the brain must arbitrate between competing models of sensory input, i.e., whether touch sensations arise from one’s actual hand or a fake hand (Samad et al., 2015; Limanowski and Friston, 2020; Apps and Tsakiris, 2014). This arbitration depends critically on the relative

precision of different sensory channels and prior beliefs. Thus, when a person sees that the rubber hand is receiving tactile input, they can perceive this as if the real hand is getting the tactile input, leading to the illusion of ownership over the fake hand (Suzuki et al., 2013; Tsakiris, 2010).

These phenomena demonstrate that precision weighting might serve as a key mechanism for both selective attention and the resolution of multisensory conflicts. By dynamically adjusting the precision of different sensory channels and prior beliefs, the brain can optimize its inferences about the state of the world and the body. This precision-based account aligns with empirical findings about the role of top-down predictions and bottom-up prediction errors in both attentional selection and perceptual illusions (Clark, 2013; Adams et al., 2013; Seth, 2014).

### 1.3.4 Unified approach towards psychedelics

After providing a background both to the entropic brain hypothesis and the free energy principle, which share their foundation in information theory, particularly in their utilization of concepts derived from Shannon entropy, we can integrate them together. The integration of these two has been labeled as the Reduced Beliefs Under Psychedelics (REBUS) theory (Carhart-Harris and Friston, 2019), highlighting that psychedelics disrupt the hierarchical processes in the brain, such that the priors are relaxed. This results in that the new incoming sensory input is considered more important (i.e., has more weight), which means that the signal can propagate deeper into the neuronal hierarchical architecture. This theoretical synthesis is particularly noteworthy as it brings together two seemingly distinct approaches to understanding brain function: one focused on the psychedelic state and altered consciousness, and the other offering a unifying theory of brain organization and function.

The REBUS theory becomes especially important when considering the role of precision-weighting in predictive processing. It suggests that, in normal waking consciousness, the brain maintains a delicate balance between the estimated precision assigned to top-down predictions and that of bottom-up sensory information (Clark, 2013). Psychedelics appears to fundamentally alter this balance by reducing the precision-weighting of prior beliefs, thereby relaxing the constraints that usually shape our perception and cognition. In addition, this mechanism also explains

increased cognitive flexibility and enhanced plasticity observed under psychedelics (van Elk and Yaden, 2022; Calder and Hasler, 2023).

## 1.4 Summary

In summary, the theoretical background above has outlined the basics of serotonin, psychedelics, and active inference. It positions the reader to better grasp the goal of this thesis, which is to demonstrate that serotonergic modulation can be modelled using the sensory precision parameter, along with expanding the theories of precision on the neuronal function in general.

Chapter 2 examines how serotonin modulates active sensing behaviour through precision weighting in an exploratory task, providing a concrete example of how this neurotransmitter may fundamentally act to regulate the precision of sensory inputs and habitual expectations. This is investigated to establish a neurobiological basis for computational theories of precision in a well-controlled animal model.

Chapter 3 translates these biological insights into robotics, showing how precision-based active inference can guide autonomous behaviour in artificial systems. This demonstrates the broad applicability of precision as an organizing principle across biological and artificial agents.

Chapter 4 explores how precision optimization can explain both perception and action in body ownership illusions, specifically modelling the rubber hand illusion through precision-weighted inference between competing models of sensory input. This chapter is pursued in order to extend our understanding of precision's role from basic active sensory processing to more complex phenomena of embodied cognition.

Chapter 5 examines how rhythmic precision modulation may underlie both attention and action, providing a mechanistic account of how precision connects perception and behaviour. This investigation helps bridge discrete computational models with continuous neural dynamics to create a more biologically realistic framework.

Finally, Chapter 6 analyzes how psychedelics, specifically psilocybin, may work by modulating precision in neural communication, supporting the REBUS theory highlighting their modulatory effects on the brain function. We study this to test whether the active inference perspective on serotonergic function as precision

modulator can explain the effects of psychedelics. This strengthens the position of serotonin as a precision neuromodulator and integrates our computational framework with broader serotonergic function across the brain and substances.

Together, these chapters aim to demonstrate that precision weighting provides a common computational framework that can unify our understanding of seemingly disparate phenomena united under the serotonergic modulation. By examining how precision is fit in a model to replicate findings in the serotonergic research across these different contexts, this thesis aims to establish serotonin as a biological homologue of precision within active inference theory, while extending these insights to theories of embodiment with a special focus on precision in generative models.

# Bibliography

- Abramson, H. A. (1966). Lsd in psychotherapy and alcoholism. *American journal of psychotherapy*, 20(3):415–438.
- Adams, R. A., Shipp, S., and Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Structure and Function*, 218:611–643.
- Apps, M. A. and Tsakiris, M. (2014). The free-energy self: a predictive coding account of self-recognition. *Neuroscience & Biobehavioral Reviews*, 41:85–97.
- Atasoy, S., Roseman, L., Kaelen, M., Kringelbach, M. L., Deco, G., and Carhart-Harris, R. L. (2017). Connectome-harmonic decomposition of human brain activity reveals dynamical repertoire re-organization under lsd. *Scientific reports*, 7(1):17661.
- Beggs, J. M. (2022). *The cortex and the critical point: understanding the power of emergence*. MIT Press.
- Beggs, J. M. and Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *Journal of neuroscience*, 23(35):11167–11177.
- Belsky, J. and Pluess, M. (2009). Beyond diathesis stress: differential susceptibility to environmental influences. *Psychological bulletin*, 135(6):885.
- Bengel, D., Murphy, D. L., Andrews, A. M., Wichems, C. H., Feltner, D., Heils, A., Mössner, R., Westphal, H., and Lesch, K.-P. (1998). Altered brain serotonin homeostasis and locomotor insensitivity to 3, 4-methylenedioxymethamphetamine (“ecstasy”) in serotonin transporter-deficient mic. *Molecular pharmacology*, 53(4):649–655.



- Berger, M., Gray, J. A., and Roth, B. L. (2009). The expanded biology of serotonin. *Annual Review of Medicine*, 60:355–366.
- Brown, H., Adams, R. A., Parees, I., Edwards, M., and Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive processing*, 14:411–427.
- Calder, A. E. and Hasler, G. (2023). Towards an understanding of psychedelic-induced neuroplasticity. *Neuropsychopharmacology*, 48(1):104–112.
- Carhart-Harris, R. L. (2018). The entropic brain-revisited. *Neuropharmacology*, 142:167–178.
- Carhart-Harris, R. L., Erritzoe, D., Williams, T., Stone, J. M., Reed, L. J., Colasanti, A., Tyacke, R. J., Leech, R., Malizia, A. L., Murphy, K., et al. (2012). Neural correlates of the psychedelic state as determined by fmri studies with psilocybin. *Proceedings of the National Academy of Sciences*, 109(6):2138–2143.
- Carhart-Harris, R. L. and Friston, K. J. (2010). The default-mode, ego-functions and free-energy: a neurobiological account of freudian ideas. *Brain*, 133(4):1265–1283.
- Carhart-Harris, R. L. and Friston, K. J. (2019). Rebus and the anarchic brain: toward a unified model of the brain action of psychedelics. *Pharmacological reviews*, 71(3):316–344.
- Carhart-Harris, R. L., Leech, R., Hellyer, P. J., Shanahan, M., Feilding, A., Tagliazucchi, E., Chialvo, D. R., and Nutt, D. (2014a). The entropic brain: a theory of conscious states informed by neuroimaging research with psychedelic drugs. *Frontiers in human neuroscience*, 8:55875.
- Carhart-Harris, R. L., Leech, R., Hellyer, P. J., Shanahan, M., Feilding, A., Tagliazucchi, E., Chialvo, D. R., and Nutt, D. (2014b). The entropic brain: a theory of conscious states informed by neuroimaging research with psychedelic drugs. *Frontiers in human neuroscience*, 8:20.
- Carhart-Harris, R. L. and Nutt, D. J. (2017). Serotonin and brain function: a tale of two receptors. *Journal of psychopharmacology*, 31(9):1091–1120.

- Carhart-Harris, R. L., Roseman, L., Haijen, E., Erritzoe, D., Watts, R., Branchi, I., and Kaelen, M. (2018). Psychedelics and the essential importance of context. *Journal of Psychopharmacology*, 32(7):725–731.
- Caspi, A., Hariri, A. R., Holmes, A., Uher, R., and Moffitt, T. E. (2010). Genetic sensitivity to the environment: the case of the serotonin transporter gene and its implications for studying complex diseases and traits. *American journal of Psychiatry*, 167(5):509–527.
- Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, 36(3):181–204.
- Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Cools, R., Nakamura, K., and Daw, N. D. (2011). Serotonin and dopamine: unifying affective, activational, and decision functions. *Neuropsychopharmacology*, 36(1):98–113.
- Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., and Friston, K. (2020). Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology*, 99:102447.
- Daw, N. D., Kakade, S., and Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural networks*, 15(4-6):603–616.
- Demchyshyn, L. L., Pristupa, Z. B., Sugamori, K. S., Barker, E. L., Blakely, R. D., Wolfgang, W. J., Forte, M. A., and Niznik, H. B. (1994). Cloning, expression, and localization of a chloride-facilitated, cocaine-sensitive serotonin transporter from drosophila melanogaster. *Proceedings of the National Academy of Sciences*, 91(11):5158–5162.
- Doya, K., Miyazaki, K. W., and Miyazaki, K. (2021). Serotonergic modulation of cognitive computations. *Current Opinion in Behavioral Sciences*, 38:116–123.
- El-Seedi, H. R., De Smet, P. A., Beck, O., Possnert, G., and Bruhn, J. G. (2005). Prehistoric peyote use: alkaloid analysis and radiocarbon dating of archaeological specimens of lophophora from texas. *Journal of ethnopharmacology*, 101(1-3):238–242.

- Feldman, H. and Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4:215.
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in cognitive sciences*, 13(7):293–301.
- Friston, K. J. and Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, 159:417–458.
- Glennon, R. A., Titeler, M., and McKenney, J. (1984). Evidence for 5-HT<sub>2</sub> involvement in the mechanism of action of hallucinogenic agents. *Life sciences*, 35(25):2505–2511.
- Grof, S. and Halifax, J. (1977). *The human encounter with death*. EP Dutton New York.
- Gutknecht, L., Araragi, N., Merker, S., Waider, J., Sommerlandt, F. M., Mlinar, B., Baccini, G., Mayer, U., Proft, F., Hamon, M., et al. (2012). Impacts of brain serotonin deficiency following tph2 inactivation on development and raphe neuron serotonergic specification.
- Hen, R. (1993). Structural and functional conservation of serotonin receptors throughout evolution. *Comparative molecular neurobiology*, pages 266–278.
- Herzog, R., Mediano, P. A., Rosas, F. E., Lodder, P., Carhart-Harris, R., Perl, Y. S., Tagliazucchi, E., and Cofre, R. (2023). A whole-brain model of the neural entropy increase elicited by psychedelic drugs. *Scientific reports*, 13(1):6244.
- Hillyard, S. A., Vogel, E. K., and Luck, S. J. (1998). Sensory gain control (amplification) as a mechanism of selective attention: electrophysiological and neuroimaging evidence. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 353(1373):1257–1270.
- Hintzen, A. and Passie, T. (2010). *The pharmacology of LSD*. OUP Oxford.
- Hohwy, J. (2013). *The predictive mind*. Oxford University Press.
- Homberg, J. R. (2012). Serotonin and decision making processes. *Neuroscience Biobehavioral Reviews*, 36(1):218–236.

- Homberg, J. R., Kyzar, E. J., Nguyen, M., Norton, W. H., Pittman, J., Poudel, M. K., Gaikwad, S., Nakamura, S., Koshiba, M., Yamanouchi, H., et al. (2016). Understanding autism and other neurodevelopmental disorders through experimental translational neurobehavioral models. *Neuroscience & Biobehavioral Reviews*, 65:292–312.
- Hoyer, D., Clarke, D. E., Fozard, J. R., Hartig, P. R., Martin, G. R., Mylecharane, E. J., Saxena, P. R., and Humphrey, P. (1994). International union of pharmacology classification of receptors for 5-hydroxytryptamine (serotonin). *Pharmacological reviews*, 46(2):157–203.
- Jacobs, B. L. and Azmitia, E. C. (1992). Structure and function of the brain serotonin system. *Physiological reviews*, 72(1):165–229.
- Jensen, O. and Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: gating by inhibition. *Frontiers in human neuroscience*, 4:186.
- Julius, D. (1991). Molecular biology of serotonin receptors. *Annual review of neuroscience*, 14(1):335–360.
- Kanai, R., Komura, Y., Shipp, S., and Friston, K. (2015). Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1668):20140169.
- Kirchhoff, M., Parr, T., Palacios, E., Friston, K., and Kiverstein, J. (2018). The markov blankets of life: autonomy, active inference and the free energy principle. *Journal of The royal society interface*, 15(138):20170792.
- Kok, P., Rahnev, D., Jehee, J. F., Lau, H. C., and De Lange, F. P. (2012). Attention reverses the effect of prediction in silencing sensory signals. *Cerebral cortex*, 22(9):2197–2206.
- Lanillos, P., Franklin, S., Maselli, A., and Franklin, D. W. (2021). Active strategies for multisensory conflict suppression in the virtual hand illusion. *Scientific Reports*, 11(1):22844.
- Lesch, K.-P. and Waider, J. (2012). Serotonin in the modulation of neural plasticity and networks: implications for neurodevelopmental disorders. *Neuron*, 76(1):175–191.

- Limanowski, J. and Friston, K. (2020). Active inference under visuo-proprioceptive conflict: Simulation and empirical results. *Scientific reports*, 10(1):4010.
- Lottem, E., Banerjee, D., Vertechi, P., Sarra, D., Lohuis, M. o., and Mainen, Z. F. (2018). Activation of serotonin neurons promotes active persistence in a probabilistic foraging task. *Nature communications*, 9(1):1000.
- Maturana, H. R. and Varela, F. J. (2012). *Autopoiesis and cognition: The realization of the living*, volume 42. Springer Science & Business Media.
- Meisel, C. (2020). Antiepileptic drugs induce subcritical dynamics in human cortical networks. *Proceedings of the National Academy of Sciences*, 117(20):11118–11125.
- Meltzer, H. Y. (2012). Serotonergic mechanisms as targets for existing and novel antipsychotics. *Current antipsychotics*, pages 87–124.
- Müller, C. P. and Homberg, J. R. (2015). The role of serotonin in drug use and addiction. *Behavioural brain research*, 277:146–192.
- Nichols, D. E., Johnson, M. W., and Nichols, C. D. (2017). Psychedelics as medicines: an emerging new paradigm. *Clinical Pharmacology & Therapeutics*, 101(2):209–219.
- Novicky, F., Parr, T., Friston, K., Mirza, M. B., and Sajid, N. (2024). Bistable perception, precision and neuromodulation. *Cerebral Cortex*, 34(1):bhad401.
- Palacios, J., Waeber, C., Mengod, G., and Pompeiano, M. (1991). Molecular neuroanatomy of 5-HT receptors. *Serotonin: molecular biology, receptors and functional effects*, pages 5–20.
- Pallavicini, C., Vilas, M. G., Villarreal, M., Zamberlan, F., Muthukumaraswamy, S., Nutt, D., Carhart-Harris, R., and Tagliazucchi, E. (2019). Spectral signatures of serotonergic psychedelics and glutamatergic dissociatives. *Neuroimage*, 200:281–291.
- Parr, T. and Friston, K. J. (2019). Attention or salience? *Current opinion in psychology*, 29:1–5.
- Parr, T., Pezzulo, G., and Friston, K. J. (2022). *Active inference: the free energy principle in mind, brain, and behavior*. MIT Press.

- Pytliak, M., Vargová, V., Mechírová, V., and Felsöci, M. (2011). Serotonin receptors-from molecular biology to clinical applications. *Physiological research*, 60(1):15.
- Rapport, M. M., Green, A. A., and Page, I. H. (1949). Serum vasoconstrictor (serotonin). *J Biol Chem*, 176(297):1243–1251.
- Samad, M., Chung, A. J., and Shams, L. (2015). Perception of body ownership is driven by bayesian sensory inference. *PloS one*, 10(2):e0117178.
- Savage, C., Savage, E., Fadiman, J., and Harman, W. (1964). Lsd: Therapeutic effects of the psychedelic experience. *Psychological Reports*, 14(1):111–120.
- Schrodinger, E. (1946). What is life?: the physical aspect of the living cell.
- Sepulcre, J., Sabuncu, M. R., Yeo, T. B., Liu, H., and Johnson, K. A. (2012). Stepwise connectivity of the modal cortex reveals the multimodal organization of the human brain. *Journal of Neuroscience*, 32(31):10649–10661.
- Sessa, B. (2016). The history of psychedelics in medicine. *Handbuch psychoaktive substanzen*, pages 1–26.
- Seth, A. K. (2014). The cybernetic bayesian brain. In *Open mind*. Open MIND. Frankfurt am Main: MIND Group.
- Shew, W. L. and Plenz, D. (2013). The functional benefits of criticality in the cortex. *The neuroscientist*, 19(1):88–100.
- Shipp, S. (2007). Structure and function of the cerebral cortex. *Current Biology*, 17(12):R443–R449.
- Silva, L. R., Amitai, Y., and Connors, B. W. (1991). Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. *Science*, 251(4992):432–435.
- Smallwood, J., Bernhardt, B. C., Leech, R., Bzdok, D., Jefferies, E., and Margulies, D. S. (2021). The default mode network in cognition: a topographical perspective. *Nature reviews neuroscience*, 22(8):503–513.
- Stewart, B., Dean, J. G., Koek, A., Chua, J., Wabl, R., Martin, K., Davoodian, N., Becker, C., Himedan, M., Kim, A., et al. (2020). Psychedelic-assisted therapy

- for functional neurological disorders: a theoretical framework and review of prior reports. *Pharmacology Research & Perspectives*, 8(6):e00688.
- Strassman, R. (2000). *DMT: The spirit molecule: A doctor's revolutionary research into the biology of near-death and mystical experiences*. Simon and Schuster.
- Sun, W. and Dan, Y. (2009). Layer-specific network oscillation and spatiotemporal receptive field in the visual cortex. *Proceedings of the National Academy of Sciences*, 106(42):17986–17991.
- Suzuki, K., Garfinkel, S. N., Critchley, H. D., and Seth, A. K. (2013). Multi-sensory integration across exteroceptive and interoceptive domains modulates self-experience in the rubber-hand illusion. *Neuropsychologia*, 51(13):2909–2917.
- Tsakiris, M. (2010). My body in the brain: a neurocognitive model of body-ownership. *Neuropsychologia*, 48(3):703–712.
- Valle, M., Maqueda, A. E., Rabella, M., Rodríguez-Pujadas, A., Antonijoan, R. M., Romero, S., Alonso, J. F., Mañanas, M. À., Barker, S., Friedlander, P., et al. (2016). Inhibition of alpha oscillations through serotonin-2a receptor activation underlies the visual effects of ayahuasca in humans. *European Neuropsychopharmacology*, 26(7):1161–1175.
- van Elk, M. and Yaden, D. B. (2022). Pharmacological, neural, and psychological mechanisms underlying psychedelics: A critical review. *Neuroscience & Biobehavioral Reviews*, 140:104793.
- Varley, T. F., Carhart-Harris, R., Roseman, L., Menon, D. K., and Stamatakis, E. A. (2020). Serotonergic psychedelics lsd & psilocybin increase the fractal dimension of cortical brain activity in spatial and temporal domains. *Neuroimage*, 220:117049.
- Vollenweider, F. X. and Kometer, M. (2010). The neurobiology of psychedelic drugs: implications for the treatment of mood disorders. *Nature Reviews Neuroscience*, 11(9):642–651.
- Vollenweider, F. X., Vollenweider-Scherpenhuyzen, M. F., Bäbler, A., Vogel, H., and Hell, D. (1998). Psilocybin induces schizophrenia-like psychosis in humans via a serotonin-2 agonist action. *Neuroreport*, 9(17):3897–3902.

- Weber, E. T. and Andrade, R. (2010). Htr2a gene and 5-HT<sub>2A</sub> receptor expression in the cerebral cortex studied using genetically modified mice. *Frontiers in neuroscience*, 4:36.
- Winkelman, M. (2019). Introduction: Evidence for entheogen use in prehistory and world religions. *Journal of Psychedelic Studies*, 3(2):43–62.
- Wong, D. T., Perry, K. W., and Bymaster, F. P. (2005). The discovery of fluoxetine hydrochloride (prozac). *Nature reviews Drug discovery*, 4(9):764–774.



## Chapter 2

# Active inference, whisking, and serotonin

Based on:

Novický, F., Zeldenrust, F., Lanillos, P., Homberg, J. R., Friston, K., & Parr, T. Active inference, whisking, and serotonin. *Manuscript in preparation.*

### Abstract

Serotonin is a neuromodulator in the brain that facilitates the interaction between the organism and its environment. However, the precise computational mechanism behind this exchange is not well understood. Here, we aimed to elucidate the computational role of serotonergic neuromodulation on sentient behaviour. Specifically, we focused on tactile active perception, i.e., whisking, as it has previously been used to study serotonin's function in perception. Taking advantage of the open-source data collected on whisking in the gap crossing task, we asked whether an active inference model of whisking could predict the behaviour observed following serotonergic modulation. We studied two conditions in a synthetic rat: a rat with normal levels of serotonin and a rat lacking the serotonin transporter ( $SERT^{-/-}$ ). We tested two hypotheses: that serotonin either affects the precision of sensory (tactile) data or the precision of habits. We conclude that a reduction of sensory precision was sufficient to explain  $SERT^{-/-}$  rat behaviour.

## Keywords

Serotonin • active sensing • precision modulation • active inference

## 2.1 Introduction

One of the most important roles of cognitive systems is to select the information that allows them to learn about their environments. This selection is necessitated by the need to move, and the different sensory inputs solicited by different movements. The perceptions informed by these movements reciprocally influence behaviour, and the ensuing action-perception cycle (Fuster, 2004): a core aspect of active sensing. One of the neuromodulators implicated in active sensing is serotonin or 5-Hydroxytryptamine (5-HT) (Azarfar et al., 2018; Lottem et al., 2018). Serotonin is a monoamine neurotransmitter with complex and multifaceted biological functions. Along with active sensing, these include the modulation of mood, cognition, reward, learning, memory, behavioural inhibition, social interactions and numerous physiological processes (Cools et al., 2008).

The effect of serotonin in active sensing has been investigated in the literature in tactile perception tasks using whiskers (Azarfar et al., 2018). Whisking in rodents offers an interesting framework to study active perception and its serotonergic alterations, as whiskers are actively targeted to explore the environment (Voigts et al., 2008) and subsequently guide behaviour, optimising both perception and action selection. This has been studied using the gap-crossing task (Azarfar et al., 2018; Celikel and Sakmann, 2007; Voigts et al., 2008) where whisking animals cross a gap of variable size between two platforms, relying completely on the tactile information actively generated via their whiskers. The changes in whisking behaviour, while crossing the platform, offer a measure of the way in which sensory data influences perception, and the way in which perception influences (information seeking) behaviour. We choose this paradigm, as it allows one to study rats with high or normal levels of serotonin, and hence to study the function of serotonin on this kind of exploratory behaviour (Azarfar et al., 2019). Specifically, the rats lacking the serotonin transporter (*SERT*<sup>-/-</sup>) take more time to switch from a large to small whisking amplitude during the gap-crossing paradigm.

But what is the specific function of serotonin that can achieve this prolonged

switch in behaviour? Research suggests that serotonin mediates environmental sensitivity (Homberg et al., 2016). In particular, this neurotransmitter appears to reverse sensory gain (i.e., flatten neural responses generated by a stimulus) across several modalities, including the tactile domain (Marquez and Chacron, 2020), which consequently implies a failure to inhibit irrelevant sensory stimulation (Farid et al., 2000; Kohnomi et al., 2008; Feifel et al., 2003; Shilling et al., 2004; Sipes and Geyer, 1994, 1997; Varty et al., 1999). From the predictive coding perspective (Brown et al., 2013; Clark, 2013; Feldman and Friston, 2010), these effects of serotonin on sensory information can be realized in two possible ways: (i), the weights (i.e., the precision) on sensory information are decreased so that they have less effect on action selection, or (ii) sensory data are not informative about action selection due to a precise habitual prior, which leads to flattened responses. The latter mechanism can be intuited from experiments where increased serotonin is associated with prolonged waiting for a reward, which might be explained as increasing the habitual (or learned) behaviour that leads to a reward (Miyazaki et al., 2018).

To test these hypotheses, we appeal to analogies between the gap-crossing paradigm and paradigms in active vision, but for tactile data. Active vision is sometimes investigated using the active inference framework (Friston et al., 2017a; Parr et al., 2022), which assumes that creatures behave based upon internal models of their world, and act to minimise the discrepancy (i.e., surprise) between the predictions of their models and the sensory data sampled from the world. This touches upon the so-called ‘dark room problem’ (Friston et al., 2012): i.e., if creatures act to minimise surprise, why do they not simply stay in a dark room, where they will always observe darkness, and therefore never be surprised? The answer to this apparent paradox (Friston et al., 2015; Parr and Friston, 2017) comes from the prior belief that creatures act to resolve their uncertainty (i.e., expected surprise) about the world. Under this prior, it would be very surprising to find oneself in a dark room, as sensory samples would be ambiguous and resolve no uncertainty. In other words, it would be all noise with no signal. In contrast, being in the light provides a much more precise signal to noise ratio, and therefore offers greater potential for uncertainty reduction. Analogously, in the tactile domain, the presence of an external object (when palpated with a whisker) provides informative data while the absence of an object offers only ambiguous data. In this sense,

whisking in the region of an external object is the tactile analogue of turning on a light-switch in the visual domain. In active inference, whisking behaviour therefore depends upon the precision either associated with tactile sensations or with habitual priors that underwrite whisking. The former reflects the confidence that can be placed in tactile data, while the latter reflects the strength of habits, where a high habitual precision favours learned action sequences that guide behaviour in a way that cannot be revised by sensory signals. In other words, a strong prior over habits results in a consistent behaviour; namely a consistent whisking amplitude in the case of this study.

In order to model the role of serotonin, we adopt the same approach as previous computational models that treat neuromodulatory transmitters as encoding precision via synaptic gain modulation (Parr and Friston, 2017). The relative precision of sensory and prior signals can have a profound influence on action selection. Previous computational studies have examined the roles of other modulators in a similar fashion, including acetylcholine (Moran et al., 2013), noradrenaline (Vincent et al., 2019), and dopamine (Schwartenbeck et al., 2015). Hence, here we ask whether serotonin modulates one (or both) of the hypothesised precision terms. In this work, we show that the serotonergic modulation may be better understood as encoding sensory (i.e., tactile) precision as it modulates the switch between a large and small amplitude whisking, independently of the habitual precision.

In what follows, we describe a computational model, based on the active inference framework as described in (Friston et al., 2017a; Smith et al., 2022), that accounts for whisking as active sensing in the cross-gap paradigm formalized as a (partially observed) Markov Decision Process (POMDP). Our approach is complementary to continuous time models of whisking using active inference (Mannella et al., 2021), but focuses on a discrete time formulation. This lets us deal in a tractable way the evaluation of alternative action plans without the need for complex function approximators thus, facilitating the analysis of the active sensing behaviour.

## 2.2 Methods

This section provides an overview of the active inference model used to simulate rodent behaviour in the gap-crossing paradigm (Azarfar et al., 2018; Celikel and Sakmann, 2007). The description is minimally technical and focuses on the in-

tuitions behind active inference. For a detailed derivation of the active inference equations we refer to (Da Costa et al., 2020; Friston et al., 2017a).

## 2.2.1 Active Inference

### 2.2.1.1 Variational Free Energy

Active inference is a corollary theory of the free energy principle that aims to unify neuroscientific findings under one relatively straightforward principle (Friston, 2010). It is built upon the idea that any living system self-organizes by minimizing its surprisal (i.e., negative log of the probability of sensory outcomes under some model). Evaluating surprisal over time, this measure corresponds to Shannon’s entropy of observable outcomes. Furthermore, minimising surprisal corresponds to maximizing model evidence (Hohwy, 2016). However, calculating surprisal for large scale problems is computationally intractable (Da Costa et al., 2020; Friston et al., 2017a; Sajid et al., 2021). Alternatively, under the variational inference approach (Beal, 2003; Blei et al., 2017), we can compute an upper bound on surprisal (i.e., Variational Free Energy), such that minimizing free energy implicitly minimizes surprisal (Da Costa et al., 2020). Equation (1) provides a definition of free energy ( $F$ ), expressed in terms of the prior beliefs ( $P(s)$ ) agents hold about the states of their world, their beliefs about how those states map to observable sensory outcomes ( $P(o|s)$ ), and the (approximate) posterior beliefs (in the form of a variational distribution  $Q(s)$ ) they optimise, once they have updated their priors to posteriors:

$$F = D_{KL} \underbrace{[Q(s)||P(s)]}_{\text{Complexity}} - E_{Q(s)} \underbrace{[\ln P(o|s)]}_{\text{Accuracy}} \quad (2.1)$$

Equation (1) expresses free energy as a functional<sup>1</sup> of the probability distributions outlined above. The notation DKL means a Kullback-Leibler divergence, which reaches its minimum (of zero) when the two distributions that form its arguments approximate one another. The notation  $EQ(s)$  means the expected value (or average) under the distribution  $Q(s)$ . The divergence term has been named Complexity, as its minimisation ensures deviations between prior and posterior probabilities are penalised, precluding overfitting of models to data (Jefferys and

---

<sup>1</sup>A functional means a function of functions.

Berger, 1992).

The second term, Accuracy, is an expected log likelihood under the approximated posterior distribution. If only the first (complexity) term were minimized, the agent would completely ignore the environment it finds itself in and would follow priors only. The accuracy term embeds the agent in the world via its sensory receptors, as this measure relies on soliciting predicted outcomes from the environment. The free energy quantity is then a combination of these two terms. Equation (2) uses an alternative (but equivalent) arrangement of the free energy, along with the non-negativity of the KL-Divergence, to demonstrate the bounding of surprisal.

$$F = D_{KL} \underbrace{[Q(s)||P(s|o)]}_{\geq 0} - \ln P(o) \geq \underbrace{-\ln P(o)}_{\text{Surprisal}} \quad (2.2)$$


While variational free energy is useful in drawing inferences about states, in order to address prospective planning (computing the best actions in the future), we must supplement this formulation with the concept of the expected free energy. We turn to this in the next subsection.

### 2.2.1.2 Expected Free Energy

The expected free energy  $G(\pi)$  resembles the free energy  $F$ , however it evaluates a sequence of actions (i.e., policy  $\pi$ ) by taking expectations under beliefs about future time-steps, and outcomes that have yet to happen. It can be expressed as follows:

$$G(\pi, \tau) = D_{KL} \underbrace{[Q(s_\tau, A|\pi)||P(s_\tau, A)]}_{\text{Risk}} + E_{Q(s_\tau, A|\pi)P(o_\tau|s_\tau, A)} \underbrace{[\ln P(o_\tau|s_\tau, A)]}_{\text{Ambiguity}} \quad (2.3)$$

where,  $\tau$  represents the time-step about which beliefs are held, and  $A$  is the likelihood matrix of outcomes given hidden states. The term Risk is effectively the expected Complexity described above. It measures the deviation from the current belief distribution given the selected policy  $Q(s_\tau, A|\pi)$ , and the real posterior  $P(s_\tau, A)$ , but now we assume that prior beliefs encode preferences about the states held by the agent being simulated. Risk minimisation is achieved by selecting those policies ( $\pi$ ), that can deliver the agent to its preferred states (or outcomes). For this



**Figure 2.1: A generative model for the gap-crossing task.** This schematic illustrates a plausible model that a rodent might use to predict its sensory data during a whisking task. The model has two hidden factors (i.e., two sorts of hidden states): Location and Whisker Position. The former is split into Near and Far from the platform. The latter is separated into full and partial retraction (FR and PR, respectively), and full and partial protraction (FP and PP, respectively). Note that partial retraction and protraction hidden states are duplicated in order to provide sine wave-like behaviour and a different frequency for large and small amplitudes. This is elucidated in the likelihood distribution in the middle-right matrix for the proprioceptive outcomes. Different combinations of these two hidden factors provide different likelihoods for the sensory data (here comprising tactile and proprioceptive modalities), where the far location provides a completely uncertain environment (analogous to a dark room), while the near location differs in the partial and full protraction hidden states with the former providing the most precise outcome. These likelihood probabilities are illustrated graphically by the matrices in the middle-left, which depict high probabilities as black, low probabilities as white, and intermediate probabilities as grey. Each cell shows the probability that, given the hidden state combination indicated by the column, the outcome indicated by the row will be observed. (Continue on the next page.)

**Figure 2.1:** (Continued) These distributions are modulated by the  $\zeta$  parameter (tactile precision) and this representation shows high  $\zeta$  (e.g., 1). If  $\zeta$  were low (e.g., 0.1), the likelihood distribution in the near condition would be completely uniform and look the same as the far condition. Lastly, there are two policies that either result in a large (1) or small (2) amplitude of whisking, coloured in purple and yellow, respectively. The former is assigned a higher prior probability of selection and is therefore linked to habitual behaviour as shown in the upper-left probability distribution. This prior probability is modulated by the parameter  $\rho$ . The mathematical parameterisation of the model and all parameters are described in the Appendix.

reason, risk minimisation is normally associated with exploitative behaviour. The second Ambiguity term follows from the formulation of the term Accuracy; however, the likelihood is expected under a predictive posterior distribution. Ambiguity minimisation is typically associated with exploratory behaviour, as the agent searches for the least ambiguous outcomes to aid in uncertainty resolution that ultimately minimises the free energy. Once the uncertainty has been resolved, the agent will naturally switch to exploitation because there is no further uncertainty to be resolved. Finally, Equation (4) explains how the final behaviour (action plan  $\pi^*$ ) is selected, where the vector  $E$  specifies habits, i.e., a prior preference for certain policies over others.

$$\pi^* = \arg \max Q(\pi) = \arg \max \tau[\rho \ln E(\pi) - G(\pi)] \quad (2.4)$$

$Q(\pi)$  expresses posterior beliefs about how to act in terms of a softmax (normalised exponential) function of two log probabilities (i.e., potentials) that reflect habitual ( $E$ ) and goal directed ( $G$ ) motivational drives<sup>2</sup>. The relative influence of the former is modulated by a precision parameter  $\rho$ . This parameter affects the overall confidence in the policy selection process and is explained in Section 2.2.1.4 in more detail.


### 2.2.1.3 Generative model of the gap-crossing task

Figure 1 describes the states and outcomes we consider in our model, as an abstraction of the graph representation of active sensing formalized as a POMDP and described in Figure 2. When the animal whisks with a large amplitude, the whisker cycles through the fully and partially protracted, and partially and fully

---

<sup>2</sup>Note that in a discrete state space formulation the softmax normalizes the value of each policy to furnish a probability mass function.





**Figure 2.2:** POMDP structure: On the left side, the panels provide equations that follow the realization of the agent under the active inference framework with discrete state and time variables. The right side pictures the factor graph of the Markov Decision Process of the agent. Adapted with permission from (Friston et al., 2017b).

retracted whisking positions. In the small amplitude case, the whisker goes through only partially protracted and retracted positions. These hidden states are discrete, and the agent decides the next action (policy) at every position. By discretizing the whisker motion to six positions (with 2 of them being duplicated)—where each next position is reached from the previously connected position with a fixed time step—small amplitudes lead to a faster whisking frequency than large amplitudes ( $\sim 10\text{Hz}$  and  $\sim 20\text{Hz}$ , respectively), as has been observed empirically (Berg and Kleinfeld, 2003; Knutsen et al., 2006; Mitchinson et al., 2011; Sofroniew et al., 2014; Voigts et al., 2008).

We assume that habitual priors favour the large amplitudes due to the more frequent use of this behaviour. Moreover, the whisking agent can find itself either in the far or near position relative to the platform, which is set before the trial and defines the condition. There are two outcome modalities, proprioceptive and tactile, where the former has a precise mapping to its corresponding ‘whisker position’ hidden state, while the latter depends on the location hidden state. In the far condition, all tactile outcomes are uniform, while in the near condition, the edge outcome is completely precise for the partially protracted condition, and the fully protracted position shows a precise mapping of touching the object, but imprecise in terms of which part of the object was touched. This is motivated as the whisker bends over the object with the fully protracted position. Consequently, the animal’s generative model is realized with the following equation:

$$P(o, s, \pi | A, B, E, \zeta, \rho) = \underbrace{P(\pi | A, B, E)}_{\text{Priors}} \prod_{t=1}^T \underbrace{P(o_t | s_t, A, \zeta)}_{\text{Likelihood}} \underbrace{P(s_t | s_{t-1}, \pi, B, E, \rho)}_{\text{Transition}} \quad (2.5)$$

Mathematically, a POMDP is simply a joint distribution over several variables, including (observable) outcomes, hidden states, policies, and sometimes additional model parameters. Expanding the joint probability (i.e. the generative model), we can see that it comprises the conditional probabilities of each policy, and the current, but also future, outcomes given hidden states and their related likelihood matrix, and the prospective hidden states, given the hidden states inferred previously, the current policy, and the transition matrix. The vector  $E$  defines the prior probability for policy selection, independently of the expected free energy  $G$  as shown in Equation 4. For a more complete understanding of the model and active inference, see (Da Costa et al., 2020; Smith et al., 2022).

**Implementation** Here, we describe how our simulated rodent collects data and subsequently generates behaviour in four steps:

1. The agent draws inferences about hidden states by forming beliefs about these states that minimise the free energy the most, based on the given outcomes. In this step, the agent infers its location (near or far) and the whisker position (FP, PP, RP, or FR) based on the tactile (edge, surface, nothing) and proprioceptive outcomes (FP, PP, RP, or FR).
2. After the agent infers its most likely position in the environment, it projects its beliefs into the future in order to select a policy that has the minimal expected free energy. This can either be the large or small amplitude whisking.
3. The policy with the least expected free energy is selected and an action is performed.
4. After the action is performed, new outcomes are generated that need to be evaluated and thus the first step repeats.

#### 2.2.1.4 Serotonergic manipulation as precision

As discussed in the introduction, we hypothesise that the function of serotonin can be understood as encoding or modulating the precision of tactile data, and/or via the habitual prior. In the context of a likelihood distribution, precision can be thought of as the signal-to-noise ratio associated with the outcome-generating process. In keeping with previous publications, we use the  $\zeta$  symbol to indicate the likelihood precision (Parr and Friston, 2017). It is modelled as:

$$A = \sigma(\zeta \cdot \log(A + e^{-8})) \quad (2.6)$$

The  $\exp(-8)$  added to the  $A$  is to ensure there are no values equal to zero inside the logarithm (i.e., to preclude numerical overflow), where the bold  $A$  represents the likelihood process of how data are generated, while  $A$  is the model's representation of this generation. For the description of  $A$  see the Appendix. Conversely, the precision of the habitual prior  $\rho$  modulates the strength of the preferred policy. This means that the stronger the habitual precision is, the more likely the agent is to select the larger amplitude, as the initial distribution of the habits is specified to prefer this action: see Equation 4.

#### 2.2.1.5 Simulations of Local Field Potential and Time Frequency Response

The LFP data are simulated based upon prediction errors, or neuronal encoding of hidden states. Specifically, they can be derived from normalized firing rates that are generated from the Bayesian model average for every iteration of message passing (16 iterations were used). In the same fashion, a time frequency response is simulated. The equation describing the underlying belief updating can be written as follows:

$$\begin{aligned} \dot{v}_{\pi,\tau} &= \epsilon_{\pi,\tau}, \quad s_{\pi,\tau} = \sigma(v_{\pi,\tau}), \\ \epsilon_{\pi,\tau} &= \sigma(\ln B_{\pi,\tau-1} s_{\pi,\tau-1} + \ln B_{\pi,\tau} \cdot s_{\pi,\tau+1} + \zeta \ln A \cdot o_\tau + \rho \ln E_{\pi,\tau} - \ln s_{\pi,\tau}) \end{aligned} \quad (2.7)$$

Where  $\epsilon$  represents a prediction error, which determines the rate of change of a variable  $v$  which plays a role analogous to a membrane potential. This potential is

converted via the softmax function to the expected states, which can be regarded as neuronal firing rates, the LFP, and the (filtered) rate of change of the potential. These electrophysiological responses can be expressed in the frequency domain via a Fourier transform to illustrate time-frequency responses. Heuristically, a higher prediction error indicates faster neural updating and greater activity in higher frequency bands. More details can be found in the simulation script (see Data Availability Statement).

In summary, the active inference framework is built on three pillars: (i) perception minimises an agent’s free energy by finding the best fit between observed outcomes and the inferred hidden state; (ii) The agent performs actions in order to resolve uncertainty under preferred outcomes; and (iii) all of this is embedded within a generative model that attempts to explain sensory outcomes (such as touch) in terms of their causes (that it was this object or event that caused it). The first is achieved by Bayesian belief updating that minimises variational free energy, the second entails planning by evaluating the expected free energy of plausible policies, and the last is realised with the joint probability distribution that forms a generative model formalized as a POMDP (Equation 5).

## 2.3 Results

### 2.3.1 Face validity

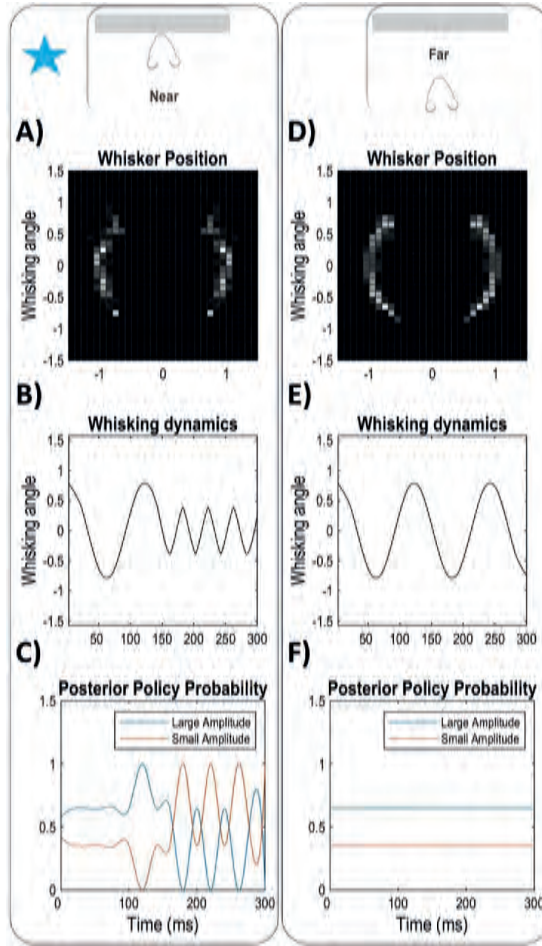
First, we explored the difference between whisking near the platform and far from it, for wild-type animals (i.e., a control group) as shown in Figure 3. This replicates the classical behaviour of an active sensing transition, where a wild-type rodent adapts to the presence of an object by switching to a smaller whisking amplitude (Celikel and Sakmann, 2007; Voigts et al., 2008). These two behaviours were simulated using exactly the same generative model, with the same prior beliefs. However, the ‘true’ location hidden state was set to either near or far to ensure the outcomes obtained by our synthetic rodent differed between the two conditions.

In both situations, the habitual prior favouring large amplitude whisking was the dominant behavioural influence at first. However, as the rat in the near condition obtained sensory data consistent with the presence of the platform (from which high quality sensory data were available), the ambiguity aversion built into the

expected free energy (see Equation 3) began to have a greater influence, prompting smaller amplitude whisking that ensured precise, high-quality sensory outcomes were obtained. We know from the literature (Berg and Kleinfeld, 2003; Knutsen et al., 2006; Mitchinson et al., 2011; Sofroniew et al., 2014; Voigts et al., 2008) that rats increase their whisking frequency from 10Hz to 20Hz (as indicated in our Figure 3B,C,E, and F) when whisking amplitude is decreased during the time they explore an environment, and vice versa: namely, that they decrease the whisking frequency with large amplitude whisking. This justifies why the whisking frequency increases with the amplitude switch in our figure, which is triggered by the posterior policy probability (Figure 3C,F). Empirically, the whisking amplitude changes approximately at the third cycle (Voigts et al., 2015), which is consistent with our simulation (Figure 3B). We plot the posterior policy probability in Figure 3C and 3F to see how the model performs the switch. Note that the reason the policy probability fluctuates is due to the same selection of outcomes, as both policies lead to the same action from the PR to PP positions and vice versa (see Figure 1). Thus the agent selects the large amplitude action due to the prior E (see Equation 3), but the small whisking frequency is not influenced by this. In addition, the policy selection is assumed to be implemented in the striatum (Parr and Friston, 2018): Figure 3 shows a hypothetical difference in neural responses in this region for whisking near and far from the object, where being near the object increases the activity.

### 2.3.2 Serotonergic manipulations

We were particularly interested in how the transition between large amplitude and low frequency whisking to high frequency and small amplitude whisking changes with levels of serotonin. We therefore sought to replicate the behaviour from (Azarfar et al., 2019) using our active inference whisking model. The authors of this paper showed that the switch from free whisking to active sensing (i.e., from large to small amplitude) is delayed in rats with higher levels of serotonin ( $SERT^{-/-}$ ) in comparison to the wild-type animals. To replicate this, we examined the influence of different parameter modulations on behaviour, to determine whether any of these computational manipulations could be a plausible explanation for the behavioural changes in serotonin transporter knockout mice and suggest what kind of behaviour



**Figure 2.3:** A comparison of whisking when an agent is near to or far from the platform to resemble the standardized behaviour of the wild type animals. Both habitual prior and tactile precision levels were kept constant at 1.5 and 1, respectively. A) A 2D histogram plot of the whisker tip position. The point  $[0,0]$  is the tip of the nose. This takes the positions implied by the whisking trajectory and adds a small amount of Gaussian noise. B) The whisking angle as a function of time. This is constructed by assigning the discrete whisker protractions from the simulation to locations in continuous space and interpolating between these locations to create a simulated trajectory. In this panel, the agent is near the platform which activates the somatosensory signals transmitted through the whiskers. Eventually, this leads to a behavioural switch. C) Posterior policy probability. This plot shows the measure that causes the switch. The reason for the fluctuation of both policies is due to the same consequence of selecting one or the other (see Figure 1). The figure has been interpolated based on discrete responses which provides a continuous look. Figures D), E), and F) represent the same template as figures A), B), C), respectively, but in the far condition. The blue star indicates where this simulation occurs in Figure 6.

should appear in animals that lack serotonin. This group is known as tryptophan hydroxylase 2 ( $TPH2^{-/-}$ ) (Kaplan et al., 2016). In what follows, we examine two candidate explanations.

### 2.3.2.1 Tactile precision

The first manipulation we investigated was the precision associated with the likelihood distribution for tactile outcomes (see Figure 1, the left middle panel). This is particularly relevant to the point made in the introduction about the role of serotonin in gating the influence of sensory data on neuronal activity, as tactile whisking data can be thought of as consequent upon where a rodent chooses to place its whiskers. The results are shown in Figure 4, where the tactile precision labels represent low to high serotonergic activity. Figures 4A and B show that with increases in the tactile precision in the near condition, the behaviour shows a faster switch from a large to small whisking amplitude, thus decreases in tactile precision seem to replicate the effect of knocking out the serotonin transporter, which shows a negative correlation with the serotonergic levels in the brain (see Section 2.2.1.4). The reason for this is that, as precision decreases, it takes longer to confidently infer that the platform is nearby. It is only when this has been inferred that the difference in the ambiguity term between large amplitude and small amplitude whisking positions can be fully appreciated by the rat.

An interesting comparison would be with rats that lack TPH2 ( $TPH2^{-/-}$ ), the enzyme synthesizing serotonin in the brain, to the wild-type condition. We incorporated hypotheses about this group in our simulations as follows, where with higher tactile precision that might be associated with low serotonin shows a faster switch to the small amplitude. Hence, this provides a testable hypothesis that  $TPH2^{-/-}$  animals will have the opposite behavioural effect on the whisking amplitude than  $SERT^{-/-}$ . Figure 4D-F shows the same precision modulation but now the agent is in the far condition as opposed to the near condition and hence not touching the platform. This figure shows that in the far condition the agent does not dramatically modulate the behavioural switches, as the distribution remains flat in all conditions and the only tactile outcome sampled is the ‘nothing’ outcome (as described in Figure 1). However, when the tactile precision increases there is a higher chance of performing a switch as shown on the right side of Figure 4E, suggesting that  $TPH2^{-/-}$  rats should show a faster whisking frequency and

smaller amplitude on average in free whisking trials. This is addressed in Figure 6.


As discussed in the methods (2.2.1.5), an advantage of using the active inference framework is that we can appeal to the associated process theories (Friston et al., 2017a)—i.e., theories about the electrophysiological manifestations of belief-updating—to validate this behavioural findings using other experimental measures. Figures 4C and 4F show simulated local field potential (LFP) signals that are superimposed on the time-frequency response for the whiskers hidden factor in the near and far conditions, respectively. These are generated from the (filtered) rates of belief updating (assuming that posterior expectations are represented by average firing rates in neural populations). In Figure 4C, during the exploration of an object, there is an apparent increase of the gamma frequency (30Hz+) with the increase of serotonin, meaning *TPH2*<sup>-/-</sup> animals show the lowest levels of gamma, while the *SERT*<sup>-/-</sup> shows the highest. One interpretation for this is that the fast belief-updating for the *TPH2*<sup>-/-</sup> mice occurs almost instantaneously at the beginning of the trial, with relatively little subsequent belief-updating.

The key thing to observe is that most belief-updating occurs prior to the switch in behaviour (as the switch depends upon the inference that the object is nearby), meaning there is more sustained high-frequency depolarisation when the behavioural switch is delayed. For the far condition (Figure 4F), the gamma frequency remains approximately the same for all groups. However, note the small decrease in the high frequency response following the temporary switch to the low amplitude whisking in the high-precision condition, and the restoration of high frequency activity on switching back to the large amplitude whisking. This reflects a brief false inference that the object is near, so that fast belief-updating is followed by a short period with relatively little belief-updating. The LFP signals themselves do not bear any clear qualitatively distinct patterns. The differences between these signals are clearest in their time-frequency representation.

### 2.3.2.2 Strength of habits

An alternative explanation for observed whisking behaviour is a modulation of ‘habits’  $E$  (see Equation 4), i.e., the prior probability of policies, as stated in the introduction. The baseline of the model assumes the preference for large amplitude whisking, as rodents more often use this strategy. Plausibly, delays in changing from the habitual to an alternative policy could be a function of the strength of this



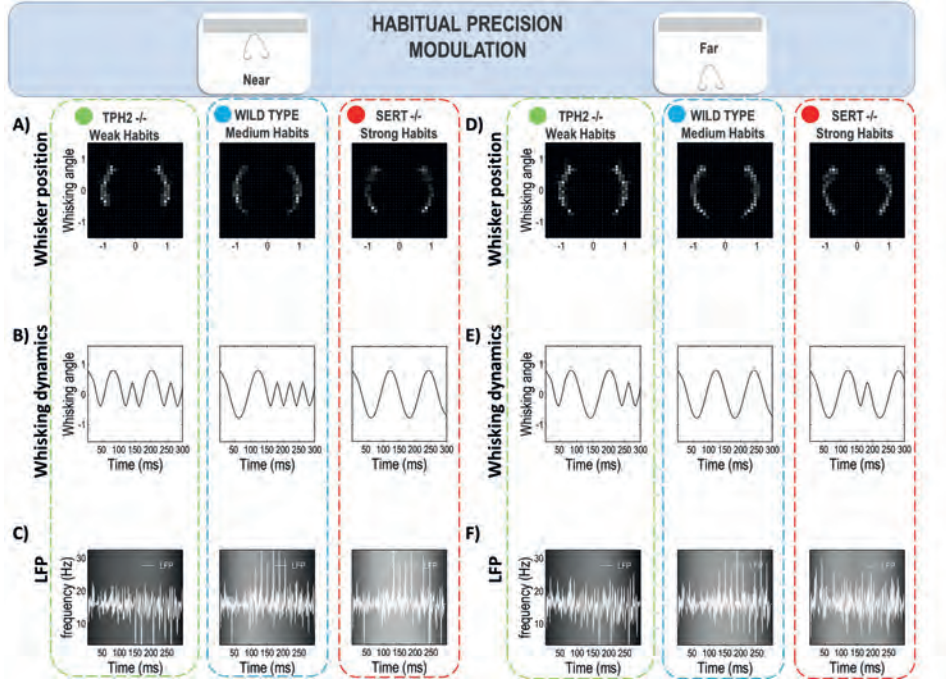


**Figure 2.4:** A modulation of tactile precision in a single simulation using three distinct  $\zeta$  values (0.01, 0.5, and 2) divided into near (A-C) and far (D-F) conditions, where the habitual prior precision ( $\rho$ ) remained constant at 0.5. The three tactile precision terms represent  $SERT^{-/-}$ , wild-type, and  $TPH2^{-/-}$  animals. A) Whisker position distribution throughout the whole trial. The y-axis represents the whisking position, where 1 is a fully protracted whisker, and -1 represents the fully retracted position. The y-axis represents the angular distance from the tip of the nose that is arbitrarily added. B) Whisking dynamics showing the temporal change of the whisker position. C) LFP data superimposed on the time-frequency response graph. Figures D-F represent the same template as in A-C. The triangle symbols represent where these simulations are plotted in Figure 6.

habit. Modulating this parameter, we observe the same behaviour as obtained from empirical data and our simulations based upon tactile precision (c.f. Figure 4 and 5). This makes sense, as the increase of the habitual prior over the large amplitude whisking policy motivates the agent to stick with this behaviour. Thus, both the modulation of tactile precision and of habits could explain how serotonin might modulate the behaviour when the agent is next to the platform, with stronger habits corresponding to raised serotonin levels. Specifically, we found that, at high levels of habitual precision, simulated behaviour resembles the behaviour of *SERT*<sup>-/-</sup> animals, as they show no switch towards the smaller and faster amplitude. On the other hand, low habitual precision should resemble the behaviour of *TPH2*<sup>-/-</sup> animals, and for this group we showed a faster switch in the near condition, in comparison to the wild-type animals, and also inconsistent and thus smaller and more frequent whisking in the far condition. These are the same findings as in the tactile precision modulation, but with the assignments of high and low precision to the genetic variants reversed (Figure 4).

Both time frequency responses in the near conditions for both tested precision terms (Figures 4C and 5C) reveal the same predictions about the serotonergic levels in the brain, as there is an apparent increase of the gamma frequency with the increase of serotonin. However, there is a tendency for a decreased gamma frequency for the *TPH2*<sup>-/-</sup> group in the far condition in comparison to the Wild type animals (Figure 5F), which provides a testable difference, as this result differs from the tactile precision modulation hypothesis.

The key result here is that the behavioural whisking switches depend on modulating habitual or tactile precision, which is—in both cases—accompanied by an increase of gamma frequency in the near condition. The time-frequency response can specifically help to dissect these two hypotheses by focusing on the gamma frequency in the free whisking trials, where if the habitual precision hypothesis is correct, one should observe a decrease in the gamma frequency in the barrel cortex of the *TPH2*<sup>-/-</sup> group. The LFP signals do not show any obvious differences between the two hypotheses.



**Figure 2.5:** A modulation of the habitual precision in order to achieve the *SERT<sup>-/-</sup>* behaviour. We used three  $\rho$  values: 0.01, 1, and 2, where the  $\zeta$  value remained constant at 0.5. The format is the same as in Figure 4. *SERT<sup>-/-</sup>* behaviour is reproduced here via high  $\rho$  value in comparison to low  $\rho$  values related to the wild type rats. The low  $\rho$  value is suggested to represent the *TPH2<sup>-/-</sup>* behaviour. The circle symbols on the top of the simulations indicate where these simulations occur in Figure 6.

### 2.3.3 Temporal analysis of the whisking-amplitude switch

Above, we showed that modulating tactile and habitual precision terms can reproduce the serotonergic modulation of active sensing found empirically (Azarfar et al., 2019). However, the modulations above were conditioned on each other, as in both cases the non-modulated precision term was kept constant. Therefore, we sought to analyse whether both parameters can reproduce the *SERT*<sup>-/-</sup> behaviour independently or whether they interact. To consider the influence of both hyperparameters,  $\zeta$  and  $\rho$ , we looked at the time delay of the switch between the large and small amplitude, i.e., when the active sensing (i.e., small amplitude whisking) strategy is selected relative to the beginning of the trial. We ran 32 simulations of the combination of 7  $\zeta$  values and 10  $\rho$  values (i.e., 70 models simulated 32 times per each) that are shown in Figure 6. We analysed both the average and median time delay of the switch, in both the far and near conditions. In the near condition Figures 6A and 6B (referring to the average and median, respectively) show that only the tactile precision parameter can reproduce the non-switching behaviour independently of the habitual precision. In other words, although there is an interaction between the two parameters, the qualitative pattern of the switching behaviour is consistent for the tactile precision, whatever value the habitual precision takes. In contrast low tactile precision (0.01) consistently shows no or rare switches in all combinations with habitual precision, implying the hypothesised association between habitual precision and serotonin is conditional upon specific values of the tactile precision. Conversely, the high habitual precision that was previously associated with the *SERT*<sup>-/-</sup> behaviour does not reproduce a consistent result in the combination with high tactile precision (i.e., switches are observed when  $\zeta \geq 1.5$ ). We arbitrarily increased the  $\rho$  values up to 100 to ensure that this is not due to a limited range of values. This  $\rho$  value shows a confident prior probability (converging to 1) of selecting the large amplitude action. This suggests that serotonin is more parsimoniously understood as the tactile precision neuromodulator, as even very high habitual precision cannot achieve the non-switching behaviour in the combination with high tactile precision ( $\zeta \geq 1.5$ ). Of the two hypotheses, the tactile precision hypothesis requires fewer assumptions. However, the simpler explanation is only the best explanation when it is also able to accurately explain data—this highlights that we cannot definitively adjudicate between the two hypotheses without resorting to experiment. One such experiment

might examine the electrophysiological responses of  $TPH2^{-/-}$  mice preceding switches in the far condition, as outlined above. The simulations from Figures 3-5 are linked to Figure 6 via their related symbols. Interestingly, this analysis also shows that the combination of both parameters modulates the timing of the switch between policy one and policy two, or free whisking and active sensing, which speaks for the complexity of precision in Bayesian inference models (Clark, 2013).

We also tested the same behaviour in the far condition (Figure 6C-F). This analysis shows that with the decrease of habitual precision and the increase of tactile precision, more random whisking behaviour emerges, which is associated with  $TPH2^{-/-}$  as shown in Figures 4 and 5. The reason the low habitual precision shows the most significant difference is due to initially setting higher habitual prior precision over the larger amplitude policy, as discussed above. With the increase of tactile precision, the drive for ambiguity resolution is enhanced (as the difference in ambiguity is greater between the two behavioural policies). This means that any possibility that the object might be present justifies the ambiguity resolving small-amplitude whisking behaviour—which would otherwise be compensated for by the habitual drive to large-amplitude whisking. This causes the inconsistent behaviour. Figures 6C and D show average and median time delay of the switch when only one small amplitude whisk suffices to be considered a switch, and Figures 6E and F show the same analysis but when at least two small whisks were observed. Comparing these figures shows that although the switches to smaller amplitude can happen in the far condition, the agents do not stay with it for too long. The median time delay shows that most of the time, the animals do not use small whisks at all during the far condition, and if so, they small-whisk only once. This provides a testable hypothesis on how fluctuant the whisking amplitude is in the free whisking trials.

In Figure 6 we have included the simulations of Figures 3-5 and the differences in the  $SERT^{-/-}$ ; wild-type; and  $TPH2^{-/-}$  groups. Increasing the number of simulations shows the same behaviour as reported above. This means that  $SERT^{-/-}$  rats show a consistent selection of the larger amplitude throughout both conditions in comparison to the wild-type animals, and is associated with low tactile precision independently on habitual precision. The wild-type condition is associated with a faster switch to low amplitude whisking when near the object in comparison to the  $SERT^{-/-}$  animals but a slower switch in comparison to


$TPH2^{-/-}$  animals where the selection of the smaller amplitude in the far condition happens rarely or not at all. The wild-type group behaviour is only achieved via a combination of medium levels of tactile precision and habitual prior precision or where one precision is high, while the other remains low. This can be associated with balanced neuromodulatory release.

## 2.4 Discussion

In this study we identified two precision parameters—tactile and habitual—that, under our generative model, modulated active sensing as shown in (Azarfar et al., 2018). Based on the temporal analyses, we found support only for the tactile precision contribution as being able to replicate the behaviour of  $SERT^{-/-}$  rats independently of the other precision (Figures 6A and B); although the habitual precision hypothesis could also reproduce the behaviour at specific values of the tactile precision. Although serotonin can modulate both precision terms, for instance due to the neurotransmitter’s several functions and receptors, we argue that for the whisking tasks only one precision term modulation suffices.

Along with  $SERT^{-/-}$  animals, we generated behaviour that can hypothetically be linked to animals lacking serotonin, i.e.,  $TPH2^{-/-}$ . The latter group, in comparison to the wild-type animals, shows a faster switch to the smaller amplitude during object detection trials and in the far condition (i.e., during free whisking) this group shows a faster frequency and a smaller amplitude on average, again in comparison to the wild type. Moreover, we simulated the LFP data along with time-frequency response for all three groups of rats with the wild-type animals as a control group. Both with the increase of habitual precision and the decrease of tactile precision (representing the increase of serotonin), we were able to show that there is an increase of the gamma frequency in the near condition (i.e., during object exploration). This suggests that the higher levels of serotonin are associated with an increase in the gamma frequency and vice versa, hypothetically at the barrel cortex as the signals come from the whiskers hidden factor. We also identified specific differences of the model’s predictions for both precision terms, and observed that there is one testable hypothesis to dissect between these two precision terms, which suggests that if the habitual precision hypothesis is correct, one should observe a decrease in the gamma frequency of the  $TPH2^{-/-}$  group during free





**Figure 2.6:** A comparison of the time delay of the behavioural switch throughout 32 trials between the large and small amplitude for the combination of 7  $\zeta$  values representing tactile precision and 10  $\rho$  values representing habitual precision. The time delay has been converted to milliseconds up to 300ms. A) An average time of the occurrence of the switch in the near condition. B) A median time of the occurrence of the switch in the near condition. The dashed red boxes represent the conditions that fail to reproduce the *SERT*<sup>-/-</sup> behaviour when the values of both precision terms are high. C) A representation when an agent selected a small whisking amplitude on average at least for one whisk cycle in the far condition D). A median representation of the same measurement as in C). E). The same as in C) but here an agent had to show two small whisk cycles happening right after each other to consider a switch. F) Same as in E) but for the median time delay. Symbols shown on the left side represent the simulations in Figures 3-5.

whisking in comparison to the wild type animals. Altogether, our simulations support the hypothesis that serotonin modulates tactile precision and provide testable measurements to scrutinize this via measuring time-frequency responses during free whisking tasks in  $TPH2^{-/-}$ .

Previous findings have provided a strong case for attributing the modulation of sensory inputs to serotonin across several modalities, which suggests that our conclusions can be generalized from tactile precision to sensory precision (Baisley et al., 2012; Farid et al., 2000; Geyer and Braff, 1987; Jacobs and Fornal, 1999; Lottem et al., 2016; Marquez and Chacron, 2020; Quednow et al., 2009). Specifically, previous studies suggest that higher levels of serotonin dampen sensory precision across several modalities for relevant stimuli, while it increases the responses to irrelevant stimuli (i.e., inhibits sensory gating). This is in accordance with our observation that increased levels of serotonin found in  $SERT^{-/-}$  rats attenuates the precision assigned to tactile information. This would give an animal the opportunity to attend to more (task-irrelevant) sensory cues. Thus, while sensory precision and thereby goal-directed behaviour is decreased in  $SERT^{-/-}$  rats relative to wild-type controls (Nonkes et al., 2010), it could foster increased adaptability when environmental circumstances change. As such,  $SERT^{-/-}$  rats show a slower learning rate and increased cognitive and behavioural flexibility (Nonkes et al., 2013), as well as reduced latent inhibition (Nonkes et al., 2012b,a), which interestingly correlates with decreased sensory gating (Jones et al., 2016). The slower learning rate naturally follows from the decreased fidelity of sensory data, as the creature cannot properly update its own beliefs. More technically, sensory precision plays the role of a learning or inference rate; enabling the faster accumulation of precise sensory information or, conversely, slowing down the assimilation of imprecise sensory input and, implicitly, assigning more precision to prior beliefs.

The idea that serotonin may decrease sensory precision has interesting consequences for some clinical conditions in which serotonergic signalling is impaired. Clinical depression is often treated with serotonin reuptake inhibitors (SSRI)—a pharmacological manipulation that increases extracellular serotonin levels. If this diminishes the precision that patients afford sensations, the action of these medications may be to help rescue people from overly precise prior beliefs about the state of the world, or about themselves. These may correspond to the unhelpful thinking styles (and their behavioural responses) that cognitive behavioural



therapy is designed to address (Williams and Garland, 2002)(Williams & Garland, 2002), and might explain the apparent benefits of combining psychological therapies and serotonergic therapies (Carhart-Harris et al., 2016; Cuijpers et al., 2020) (Carhart-Harris et al., 2016; Cuijpers et al., 2020).

In terms of active inference, the modulation of prior precision can explain the effects of decreased sensory gain as a by-product of a selection of one common policy or action. This is because as the agent has more confidence in what actions to select, even before assimilating sensory stimuli; this bottom-up sensory information then becomes redundant and can be explained away more easily (Friston, 2005). Additionally, Miyazaki et al. (2018) observed that when the chance of receiving reward was sufficiently high, the mice with increased serotonin levels via optogenetic activation waited for a reward longer. A similar finding has been observed in (Lottem et al., 2018), where mice with optogenetically activated serotonergic regions showed the same pattern of behaviour in a foraging task. Although these behaviours are intuitively explained via an increase in habitual prior precision over the rewarding policies, we showed here that the selection of one consistent behaviour is better understood in terms of an attenuation of sensory precision. Thus, rather than having prolonged waiting due to the habitual behaviour, it may be caused by reducing sensory precision, and impeding belief updating based upon sensory cues. In addition, Lottem et al. (2016) reported that serotonin decreases top-down activity while bottom-up sensory signals remain unaffected. This is consistent with the attenuation of sensory precision, in the sense that higher levels of processing (e.g., policy evaluation) are less sensitive to bottom-up influences.

Focusing on LFP data generated from our model, the tactile precision modulation showed that the SERT-/- group, associated with low tactile precision, should show a decrease in the LFP signal during the far condition or during the free whisking trial. Interestingly, this is consistent with LFP findings from the barrel cortex, where it has been shown that the ingestion of SSRI suppresses such responses in vivo (Akhmetshina et al., 2016). However, the opposite was observed in (Miceli et al., 2017). This suggests that there may be different responses to serotonin depending upon different factors, including the brain region in which it acts, the predominant receptor subtypes in different tissue samples, and the chronicity of exposure to serotonin (which may influence sensitivity to serotonin via up or down-regulation of receptors). We have provided hypotheses that disambiguate

between the habitual and tactile precision modulatory effects of serotonin: thus, the LFP signals in *TPH2*<sup>-/-</sup> could be tested to see whether serotonin is really better understood as the tactile precision modulator. If this is so, then the increase of LFP signals in free whisking task should be observed. If the influence of habits is indeed playing a role, the opposite should be observed.

It has previously been suggested that serotonin can work as a modulator of sensory precision (Carhart-Harris and Friston, 2019). This modulation is related to psychedelics and thus to 5HT2A receptor. There is evidence from normalization models that 5HT2A flattens neural responses (Azimi et al., 2018), which is in accordance with the assumption that serotonin decreases sensory precision and thus synaptic gain, as normalization models loosely effect the same function (Reynolds and Heeger, 2009). This suggests that by blocking 5HT2A in *SERT*<sup>-/-</sup> rats (for instance via ketanserin), their whisking amplitude should be unchanged or faster in comparison to the wild type rats. Among other things, (Azarfar et al., 2018; Miceli et al., 2017) observed that *SERT*<sup>-/-</sup> rats crossed the gap faster than their wild-type peers. If we assume that recognition of the platform is necessary for crossing the gap quickly, then this is apparently contradictory to the idea that the *SERT*<sup>-/-</sup> rats find themselves in the hypothetical dark room (see introduction), since they can recognize the second platform faster. This limitation might call for a more complex (i.e., hierarchical) model that could analyse the sensory information in higher levels of the hierarchical model and thus to infer a platform based on the touch. Another self-evident limitation of our current model is that it suffers from having only one whisker and thus the link to the integrative function of several whiskers (Suchkov et al., 2018) is missing. However, it has been shown that mice with a deprived number of whiskers show unaffected performance in the gap-crossing task (Papaioannou et al., 2013). Additionally, since crossing a gap can be seen as decision making, it might be interesting to compare different sorts of behaviour with manipulations of serotonin. For instance, Guo et al. (2020) showed that *SERT*<sup>-/-</sup> rats show slower decision making when categorizing new stimuli.

To conclude, we have demonstrated that the reduction in whisking amplitude, on encountering an informative stimulus, associated with the *SERT*<sup>-/-</sup> behaviour (Azarfar et al., 2019) can be replicated in silico through appealing to ambiguity minimisation (or information gain) principles; i.e., a modulation of tactile precision and the habitual precision behaviour. Assuming that serotonin selectively modulates

one kind of precision, we showed that changes in tactile precision provides a simpler account than the habitual prior precision, as the former can reproduce the qualitative  $SERT^{-/-}$  behaviour independently of the latter, while the modulation of latter depends on the former. Due to the similar effects of serotonin on tactile information and visual, odour, and auditory modalities (Marquez and Chacron, 2020) tactile precision can be generalized to sensory precision. This however needs to be further tested. Additionally, we have specified hypotheses on how to test this assumption via recording of LFP signals, and have also hypothesised behavioural response patterns for  $TPH2^{-/-}$  rats. This pattern is associated with a faster decrease in whisking amplitude, in comparison to the wild-type animals when exploring an object. Analysing the responses of  $TPH2^{-/-}$  rats can also proceed a free whisking paradigm, where we suggest that the amplitude should be smaller on average in comparison to the wild type animals but also to the  $SERT^{-/-}$ . However, it is still uncertain which serotonergic receptor could mediate changes in sensory precision. Based on the previous literature, it has been suggested that psychedelics, known to activate the 5HT2A receptor, might share the same computational function observed here (i.e., decrease precision) (Carhart-Harris and Friston, 2019), thus future work should aim to scrutinize 5HT2A as a potential mediator of precision control. In  $SERT^{-/-}$  rodents, this can be tested via ingesting ketanserin, a 5HT2A receptor antagonist, that blocks the effects of psychedelics (Preller et al., 2017).

## 2.5 Data Availability Statement

The model developed and tested here can be found at [github.com/filipnovicky/whisking\\_model](https://github.com/filipnovicky/whisking_model). Also, use the SPM toolbox to run this simulation [www.fil.ion.ucl.ac.uk/spm/software/spm12/](http://www.fil.ion.ucl.ac.uk/spm/software/spm12/). From this toolbox, we used `SPM_MDP_VB_X.m` as a standard routine and for the LFP data we used `SPM_MDP_VB_LFP.m`. The used generative model was the one specified in this article. We invite readers to reproduce and also customize the results.

## 2.6 Funding

FN is funded by the ITN Serotonin and Beyond (953327). TP is affiliated with The Wellcome Centre for Human Neuroimaging which is supported by core funding (203147/Z/16/Z). PL is partially supported by the Spikeference project (HBP, SGA3, Grant agreement ID: 945539).

# Bibliography

- Akhmetshina, D., Zakharov, A., Vinokurova, D., Nasretdinov, A., Valeeva, G., and Khazipov, R. (2016). The serotonin reuptake inhibitor citalopram suppresses activity in the neonatal rat barrel cortex in vivo. *Brain Research Bulletin*, 124:48–54.
- Azarfar, A., Zhang, Y., Alishbayli, A., Miceli, S., Kepser, L., van der Wielen, D., van de Moosdijk, M., Homberg, J., Schubert, D., Proville, R., and Celikel, T. (2018). An open-source high-speed infrared videography database to study the principles of active sensing in freely navigating rodents. *Gigascience*, 7(12).
- Azarfar, A., Zhang, Y., Alishbayli, A., Schubert, D., Homberg, J. R., and Celikel, T. (2019). Serotonergic development of active sensing. *bioRxiv*, page 762534.
- Azimi, Z., Spoida, K., Barzan, R., Wollenweber, P., Mark, M. D., Herlitze, S., and Jancke, D. (2018). Subtraction and division of visual cortical population responses by the serotonergic system. *bioRxiv*, page 444943.
- Baisley, S. K., Fallace, K. L., Rajbhandari, A. K., and Bakshi, V. P. (2012). Mutual independence of 5-HT<sub>2</sub> and  $\alpha$ 1 noradrenergic receptors in mediating deficits in sensorimotor gating. *Psychopharmacology (Berl)*, 220(3):465–479.
- Beal, M. J. (2003). *Variational algorithms for approximate Bayesian inference*. University of London, University College London (United Kingdom).
- Berg, R. W. and Kleinfeld, D. (2003). Rhythmic whisking by rat: retraction as well as protraction of the vibrissae is under active muscular control. *Journal of Neurophysiology*, 89(1):104–117.

- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877.
- Brown, H., Adams, R. A., Parees, I., Edwards, M., and Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cogn Process*, 14(4):411–427.
- Carhart-Harris, R. L., Bolstridge, M., Rucker, J., Day, C. M. J., Erritzoe, D., Kaalen, M., Bloomfield, M., Rickard, J. A., Forbes, B., Feilding, A., Taylor, D., Pilling, S., Curran, V. H., and Nutt, D. J. (2016). Psilocybin with psychological support for treatment-resistant depression: an open-label feasibility study. *The Lancet Psychiatry*, 3(7):619–627.
- Carhart-Harris, R. L. and Friston, K. (2019). Rebus and the anarchic brain: toward a unified model of the brain action of psychedelics. *Pharmacological Reviews*, 71(3):316–344.
- Celikel, T. and Sakmann, B. (2007). Sensory integration across space and in time for decision making in the somatosensory system of rodents. *Proceedings of the National Academy of Sciences*, 104(4):1395–1400.
- Clark, A. (2013). The many faces of precision (replies to commentaries on “whatever next? neural prediction, situated agents, and the future of cognitive science”). *Frontiers in psychology*, 4:270.
- Cools, R., Roberts, A. C., and Robbins, T. W. (2008). Serotonergic regulation of emotional and behavioural control processes. *Trends in Cognitive Sciences*, 12(1):31–40.
- Cuijpers, P., Noma, H., Karyotaki, E., Vinkers, C. H., Cipriani, A., and Furukawa, T. A. (2020). A network meta-analysis of the effects of psychotherapies, pharmacotherapies and their combination in the treatment of adult depression. *World Psychiatry*, 19(1):92–107.
- Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., and Friston, K. (2020). Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology*, 99:102447.

- Farid, M., Martinez, Z. A., Geyer, M. A., and Swerdlow, N. R. (2000). Regulation of sensorimotor gating of the startle reflex by serotonin 2a receptors: Ontogeny and strain differences. *Neuropsychopharmacology*, 23(6):623–632.
- Feifel, D., Melendez, G., and Shilling, P. D. (2003). A systemically administered neurotensin agonist blocks disruption of prepulse inhibition produced by a serotonin-2a agonist. *Neuropsychopharmacology*, 28(4):651–653.
- Feldman, H. and Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Front Hum Neurosci*, 4:215.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456):815–836.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138.
- Friston, K., Fitzgerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017a). Active inference: A process theory. *Neural Computation*, 29(1):1–49.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., and Pezzulo, G. (2015). Active inference and epistemic value. *Cogn Neurosci*, 6(4):187–214.
- Friston, K., Thornton, C., and Clark, A. (2012). Free-energy minimization and the dark-room problem. *Frontiers in Psychology*, 3.
- Friston, K. J., Parr, T., and De Vries, B. (2017b). The graphical brain: Belief propagation and active inference. *Network Neuroscience*, 1(4):381–414.
- Fuster, J. M. (2004). Upper processing stages of the perception–action cycle. *Trends in Cognitive Sciences*, 8(4):143–145.
- Geyer, M. A. and Braff, D. L. (1987). Startle habituation and sensorimotor gating in schizophrenia and related animal models. *Schizophr Bull*, 13(4):643–668.
- Guo, C. C.-G., Minda, J. P., and Homberg, J. (2020). Serotonin transporter knockout in rats modulates category learning. *bioRxiv*, page 2020.11.09.373886.
- Hohwy, J. (2016). The self-evidencing brain. *Noûs*, 50(2):259–285.

- Homberg, J. R., Kyzar, E. J., Nguyen, M., Norton, W. H., Pittman, J., Poudel, M. K., Gaikwad, S., Nakamura, S., Koshiba, M., Yamanouchi, H., Scattoni, M. L., Ullman, J. F., Diamond, D. M., Kaluyeva, A. A., Parker, M. O., Klimenko, V. M., Apryatin, S. A., Brown, R. E., Song, C., Gainetdinov, R. R., Gottesman, I., and Kalueff, A. V. (2016). Understanding autism and other neurodevelopmental disorders through experimental translational neurobehavioral models. *Neurosci Biobehav Rev*, 65:292–312.
- Jacobs, B. L. and Fornal, C. A. (1999). Activity of serotonergic neurons in behaving animals. *Neuropsychopharmacology*, 21(2 Suppl):9s–15s.
- Jefferys, W. H. and Berger, J. O. (1992). Ockham’s razor and bayesian analysis. *American Scientist*, 80(1):64–72.
- Jones, L. A., Hills, P. J., Dick, K. M., Jones, S. P., and Bright, P. (2016). Cognitive mechanisms associated with auditory sensory gating. *Brain Cogn*, 102:33–45.
- Kaplan, K., Echert, A. E., Massat, B., Puissant, M. M., Palygin, O., Geurts, A. M., and Hodges, M. R. (2016). Chronic central serotonin depletion attenuates ventilation and body temperature in young but not adult tph2 knockout rats. *Journal of Applied Physiology*, 120(9):1070–1081.
- Knutsen, P. M., Pietr, M., and Ahissar, E. (2006). Haptic object localization in the vibrissal system: behavior and performance. *Journal of Neuroscience*, 26(33):8451–8464.
- Kohnomi, S., Suemaru, K., Kawasaki, H., and Araki, H. (2008). Effect of aripiprazole on 5-HT<sub>2</sub> receptor-mediated wet-dog shake responses and disruption of prepulse inhibition in rats. *J Pharmacol Sci*, 106(4):645–650.
- Lottem, E., Banerjee, D., Vertechi, P., Sarra, D., Lohuis, M. O., and Mainen, Z. F. (2018). Activation of serotonin neurons promotes active persistence in a probabilistic foraging task. *Nature Communications*, 9(1).
- Lottem, E., Lorincz, M. L., and Mainen, Z. F. (2016). Optogenetic activation of dorsal raphe serotonin neurons rapidly inhibits spontaneous but not odor-evoked activity in olfactory cortex. *Journal of Neuroscience*, 36(1):7–18.



- Mannella, F., Maggiore, F., Baltieri, M., and Pezzulo, G. (2021). Active inference through whiskers. *Neural Networks*, 144:428–437.
- Marquez, M. M. and Chacron, M. J. (2020). Serotonin and sensory processing. In *Handbook of Behavioral Neuroscience*, volume 31, pages 449–459. Elsevier.
- Miceli, S., Nadif Kasri, N., Joosten, J., Huang, C., Kepser, L., Proville, R., Selten, M. M., van Eijs, F., Azarfar, A., and Homberg, J. R. (2017). Reduced inhibition within layer iv of sert knockout rat barrel cortex is associated with faster sensory integration. *Cerebral Cortex*, 27(2):933–949.
- Mitchinson, B., Grant, R. A., Arkley, K., Rankov, V., Perkon, I., and Prescott, T. J. (2011). Active vibrissal sensing in rodents and marsupials. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1581):3037–3048.
- Miyazaki, K., Miyazaki, K. W., Yamanaka, A., Tokuda, T., Tanaka, K. F., and Doya, K. (2018). Reward probability and timing uncertainty alter the effect of dorsal raphe serotonin neurons on patience. *Nature Communications*, 9(1):1–11.
- Moran, R. J., Campo, P., Symmonds, M., Stephan, K. E., Dolan, R. J., and Friston, K. J. (2013). Free energy, precision and learning: the role of cholinergic neuromodulation. *The Journal of Neuroscience*, 33(19):8227–8236.
- Nonkes, L. J., de Pooter, M., and Homberg, J. R. (2012a). Behavioural therapy based on distraction alleviates impaired fear extinction in male serotonin transporter knockout rats. *Journal of Psychiatry and Neuroscience*, 37(4):224–230.
- Nonkes, L. J., Maes, J. H., and Homberg, J. R. (2013). Improved cognitive flexibility in serotonin transporter knockout rats is unchanged following chronic cocaine self-administration. *Addiction biology*, 18(3):434–440.
- Nonkes, L. J., Tomson, K., Mærtin, A., Dederen, J., Maes, J. R., and Homberg, J. (2010). Orbitofrontal cortex and amygdalar over-activity is associated with an inability to use the value of expected outcomes to guide behaviour in serotonin transporter knockout rats. *Neurobiology of Learning and Memory*, 94(1):65–72.
- Nonkes, L. J., van de Vondervoort, I. I., de Leeuw, M. J., Wijlaars, L. P., Maes, J. H., and Homberg, J. R. (2012b). Serotonin transporter knockout rats show

- improved strategy set-shifting and reduced latent inhibition. *Learning & Memory*, 19(5):190–193.
- Papaioannou, S., Brigham, L., and Krieger, P. (2013). Sensory deprivation during early development causes an increased exploratory behavior in a whisker-dependent decision task. *Brain and Behavior*, 3(1):24–34.
- Parr, T. and Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of the Royal Society Interface*, 14(136):20170376.
- Parr, T. and Friston, K. J. (2018). The anatomy of inference: generative models and brain structure. *Frontiers in Computational Neuroscience*, page 90.
- Parr, T., Pezzulo, G., and Friston, K. J. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. The MIT Press.
- Preller, K. H., Herdener, M., Pokorny, T., Planzer, A., Kraehenmann, R., Stämpfli, P., Liechti, M. E., Seifritz, E., and Vollenweider, F. X. (2017). The fabric of meaning and subjective effects in lsd-induced states depend on serotonin 2a receptor activation. *Current Biology*, 27(3):451–457.
- Quednow, B. B., Schmechtig, A., Ettinger, U., Petrovsky, N., Collier, D. A., Vollenweider, F. X., Wagner, M., and Kumari, V. (2009). Sensorimotor gating depends on polymorphisms of the serotonin-2a receptor and catechol-o-methyltransferase, but not on neuregulin-1 arg38gln genotype: a replication study. *Biological Psychiatry*, 66(6):614–620.
- Reynolds, J. H. and Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61(2):168–185.
- Sajid, N., Ball, P. J., Parr, T., and Friston, K. J. (2021). Active inference: demystified and compared. *Neural Computation*, 33(3):674–712.
- Schwartenbeck, P., Fitzgerald, T. H. B., Mathys, C., Dolan, R., and Friston, K. (2015). The dopaminergic midbrain encodes the expected certainty about desired outcomes. *Cerebral Cortex*, 25(10):3434–3445.
- Shilling, P., Melendez, G., Priebe, K., Richelson, E., and Feifel, D. (2004). Neureotensin agonists block the prepulse inhibition deficits produced by a 5-HT<sub>2A</sub> and an  $\alpha$ 1 agonist. *Psychopharmacology*, 175(3):353–359.

- Sipes, T. and Geyer, M. (1994). Multiple serotonin receptor subtypes modulate prepulse inhibition of the startle response in rats. *Neuropharmacology*, 33(3-4):441–448.
- Sipes, T. E. and Geyer, M. A. (1997). Doi disrupts prepulse inhibition of startle in rats via 5-HT<sub>2A</sub> receptors in the ventral pallidum. *Brain Research*, 761(1):97–104.
- Smith, R., Friston, K. J., and Whyte, C. J. (2022). A step-by-step tutorial on active inference and its application to empirical data. *Journal of Mathematical Psychology*, 107:102632.
- Sofroniew, N. J., Cohen, J. D., Lee, A. K., and Svoboda, K. (2014). Natural whisker-guided behavior by head-fixed mice in tactile virtual reality. *Journal of Neuroscience*, 34(29):9537–9550.
- Suchkov, D., Sharipzyanova, L., and Minlebaev, M. (2018). Horizontal synchronization of neuronal activity in the barrel cortex of the neonatal rat by spindle-burst oscillations. *Frontiers in Cellular Neuroscience*, 12:5.
- Varty, G. B., Bakshi, V. P., and Geyer, M. A. (1999). M100907, a serotonin 5-HT<sub>2A</sub> receptor antagonist and putative antipsychotic, blocks dizocilpine-induced prepulse inhibition deficits in sprague-dawley and wistar rats. *Neuropsychopharmacology*, 20(4):311–321.
- Vincent, P., Parr, T., Benrimoh, D., and Friston, K. J. (2019). With an eye on uncertainty: Modelling pupillary responses to environmental volatility. *PLoS Computational Biology*, 15(7):e1007126.
- Voigts, J., Herman, D. H., and Celikel, T. (2015). Tactile object localization by anticipatory whisker motion. *Journal of Neurophysiology*, 113(2):620–632.
- Voigts, J., Sakmann, B., and Celikel, T. (2008). Unsupervised whisker tracking in unrestrained behaving animals. *Journal of Neurophysiology*, 100(1):504–515.
- Williams, C. and Garland, A. (2002). Identifying and challenging unhelpful thinking. *Advances in Psychiatric Treatment*, 8(5):377–386.

## 2.7 Appendix

Here we describe the specific parameterisation of the whisking model depicted in Figure 1.

### 2.7.1 Initial hidden states

The agent starts with pre-set beliefs about the initial position. Specifically, for Figures 3-5 the agent believed to start in a fully protracted position, while for Figure 6 the agent started in either fully protracted or partially protracted positions. The agent always believed in starting in the far condition and had to infer it is next to the platform, if it was so.

### 2.7.2 Likelihood matrices of how data are generated

The tactile data were generated via a random seed initialisation with the following likelihood matrices for the near condition shown in Equation (8):

$$A_{near} = \begin{bmatrix} 0 & 0 & 1/3 & 1/3 & 1/3 & 0 \\ 0.5 & 0 & 1/3 & 1/3 & 1/3 & 0 \\ 0.5 & 1 & 1/3 & 1/3 & 1/3 & 1 \end{bmatrix} \quad (2.8)$$

And for the far condition shown in Equation (9):

$$A_{far} = \begin{bmatrix} 1/3 & 1/3 & 1/3 & 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 & 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 & 1/3 & 1/3 & 1/3 \end{bmatrix} \quad (2.9)$$

Where columns and rows represent whisker position with the duplicated partial retraction (5th column) and partial protraction (6th column) as the last two columns. Specifically, the sequence of columns is FP, PP, PR, FR, PR, PP. The rows represent ‘Nothing’, ‘Surface’, and ‘Edge’ from top to down. The template is essentially the same as in Figure 1.

### 2.7.3 Transition Matrices

The transition matrix for the far and near condition was set as an identity matrix without an agent’s control over it. The equation is shown below:

$$B_{condition} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (2.10)$$

The transition matrix for the whisker position was set in such a way to allow the agent to move in a sinusoidal wave and for that the partial protraction and retraction states had to be duplicated, otherwise the agent would switch from the full retraction to full protraction position. All the action showed precise mapping of probability 1.

$$B_{Large\ amplitude} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \quad (2.11)$$

Where the rows represent the whisking positions in the next time step, while the columns represent the current time step, going from the full protraction to full retraction and then back to partial protraction. The small amplitude matrix is shown below:

$$B_{Small\ amplitude} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (2.12)$$

Both these matrices are defined as controllable states by the agent.

#### 2.7.4 Habitual vector

The agent started with the preference over the large amplitude that is then modulated via the precision term  $\rho$  as in Equation 4. The initial E vector is defined as  $E = [0.6 \ 0.4]$ .

## Chapter 3

# Robotic active tactile sensing inspired by serotonergic modulation using active inference

Based on:

Novicky, F., Offergeld, J., Janssen, S., & Lanillos, P. (2023). Robotic active tactile sensing inspired by serotonergic modulation using active inference. In F. Meder, A. Hunt, L. Margheri, A. Mura, & B. Mazzolai (Eds.), *Biomimetic and biohybrid systems* (Living Machines 2023, Lecture Notes in Computer Science, Vol. 14157). Springer, Cham. [https://doi.org/10.1007/978-3-031-38857-6\\_3](https://doi.org/10.1007/978-3-031-38857-6_3)

## Abstract

When faced with uncertainty in the world, biological agents actively sense the environment to acquire the most informative input to fulfil their tasks. Actions are performed to adjust bodily sensors to maximize the collected information, which is usually known as active sensing. For instance, rodents continuously adjust the speed and amplitude of whisking to better identify objects and body location in space, which ultimately regulates navigation. Whilst, the internal mechanism that drives active sensing in humans is still under research, recent evidence points towards neuromodulators, such as serotonin, that influence whether the habitual behaviour is preferred over sensor adjustments to trigger exploration. Here, we present an active tactile-sensing model for a robot inspired by the serotonergic function viewed from the uncertainty minimization perspective. To mechanistically explain this neuromodulatory function, we associated it with precision parameters regulating habitual behaviour and tactile encoding based on previous findings. We qualitatively evaluated the model using an experiment inspired by the gap-crossing paradigm but tailored to a humanoid with tactile sensing. Behavioural switch timing results show the strong dependencies between active sensing and precision regulation. Ultimately, this work discusses how the neural microcircuitry regulates active sensing, hence opening future research of such neuromodulatory processes translated to robotics active sensing and perception.

## Keywords

Active Sensing • Active Perception • Active Inference • Serotonin • Precision • Neuro-robotics

## 3.1 Introduction

Perception is not passive. Biological agents actively explore the world by selecting such actions that inform them about the world, given their embodied constraints. Whilst this active sensing process is not yet fully understood (Bajcsy et al., 2018; Meera et al., 2022), recent findings in neuroscience point to neuromodulators, such as serotonin, to be strongly involved. In humans, serotonin has been suggested to be involved in the sensitivity of sensory processing (Homberg et al., 2016) and perceptual decision making—by tracking

sensory uncertainty (Bang et al., 2020). But regarding active sensing a revealing study comes from exploratory functions of rodents (Azarfar et al., 2019), where it has been shown that while wild type rodents, with normal levels of serotonin, show a switch from a large to a small amplitude of their whiskers when they touch an object—to acquire more information—, rodents with altered levels of serotonin (i.e., knock-out serotonin mice, SERT-/-) do not present this behaviour. This may mean that SERT-/- rodents have their active sensing mechanism turned off. We suggest that the changes in the exploratory behaviour can be understood as an uncertainty minimization strategy that is triggered by switching to a more efficient mode of collecting data under higher uncertainty.

We hypothesize that abnormal active sensing behaviour can be accounted for through aberrant precision modulation (i.e., broadly inverse covariance) in two ways: modulating incoming sensory input (via reducing sensory precision) or modulating habitual behaviours (increasing habitual precision). The sensory precision dictates how accurately a stimulus will be encoded in the brain and how much information it provides for an apt inference. Habitual precision, on the other hand, modulates habits that are selected a priori without considering current input (Miller et al., 2019). Habits have the advantages that they do not require the complex computations of policy selection. Hence, we suggest that proper precision modulation is at the core of producing active sensing behaviour.

Here, we present a robotic active perception computational model inspired by the serotonin-based active sensing observed in rodents. Our work follows from previous research on active sensing where – from the computational point of view – this behaviour relates to uncertainty minimization or information gain maximization (Yang et al., 2016; Martinez-Hernandez et al., 2017; Parr et al., 2021b). Our model thus extends on these theories via providing precision parameters as a neuromodulatory function that regulates this behaviour—see (Parr and Friston, 2017; Parr et al., 2022).

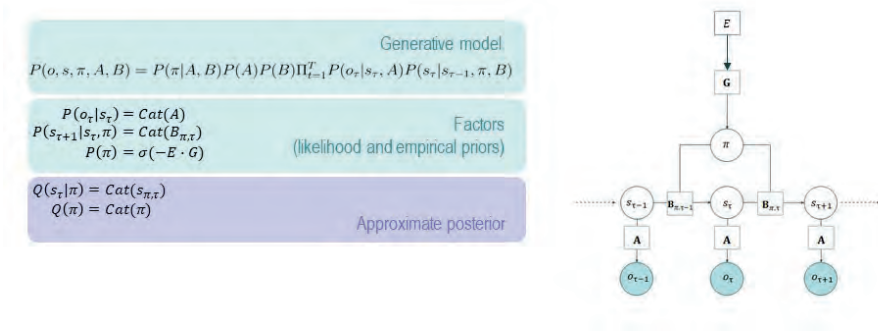
While rodents active sensing differs from human tactile sensing, there are some relevant connections, such as how tactile information is obtained through habitual versus exploratory behaviours. Disentangling this mechanism may provide a new direction in tactile active perception, particularly in robotics applications, such as navigation (Rasouli et al., 2020) and object recognition (Kaboli et al., 2017). Besides, our model can be specifically used to translate its theoretical findings into animal or human behaviour via the deployment of the model in a humanoid robot.

Hence, to validate the model, we deployed the model on a humanoid robot and designed an exploration experiment inspired by rodent navigation via the so-called gap-crossing paradigm (Voigts et al., 2008). The humanoid robot relies only on somatosensory (tactile) data with a broad movement amplitude (i.e., randomly moving robot’s hands in the air) to maximize the range of where the stimulation can come from. Active sensing is



then considered when a small range of movement around a more informative stimulus is selected. This is consistent with previous computational models that understand active sensing as information gain behaviour (Yang et al., 2016; Parr et al., 2021b). We found that optimal levels of precision lead to more informative behaviour, while dysregulations of these move the agent to suboptimal behaviour where information gained from an object are missed, causing an agent to be ignorant about the environment. Thus, indicating the strong connection between precision modulation and optimal active sensing behaviours. The specific contribution of our model lies in applying the uncertainty minimization perspective in active tactile-sensing, and its demonstration in a humanoid.

## 3.2 Methods



**Figure 3.1:** General factor graph of the POMDP behind the discrete active inference model of the agent. On the left panel, the equations specify the generative model. The generative model is thus a joint probability of outcomes, states, and model parameters. The right panel shows the graphical representation of the same generative model. This figure has been adapted from (Friston et al., 2017b). The terms  $o$ ,  $s$ ,  $\pi$  and  $\tau$  refer to outcomes, hidden states, policies, and time-steps respectively. Model parameters are labeled as  $A$ ,  $B$  and  $E$  which relates to likelihood and transition matrices, and a habitual vector, respectively. Lastly,  $\text{Cat}(\cdot)$  specifies a categorical distribution, and  $\sigma$  is a softmax function.

Our model is grounded on the Active Inference (Friston et al., 2017a) (AIF) mathematical framework—see (Lanillos et al., 2021) for a survey on robotic applications. Particularly, we used the discrete state space AIF to model the robot’s behaviour. Thus, we first briefly describe active inference and then the model used for active sensing.

### 3.2.1 Agent Model

We modelled the agent dynamics and action selection using the discrete Partially-Observable Markov Decision Process (POMDP) formalism. The POMDP structure provides a mapping from outcomes to hidden states, and across time-steps, referring to likelihood and transition functions. The structure is specified in Figure 3.1. The agent has a hidden state  $s$  that evolves through time depending on the transition function  $B$  and is updated given the observation  $o$  at each time step. The agent has a mapping (likelihood) of observations given the state.

### 3.2.2 Active inference

Action selection is modelled under the AIF mathematical framework, a corollary theory of the free energy principle. The free energy principle assumes that all living systems follow a simple process of self-organization which is to minimize the surprisal value (Friston, 2010). Variational free energy ( $\mathcal{F}$ ) then works as an upper-bound on this value (Friston et al., 2017a):

$$\mathcal{F} \geq -\ln P(o) \quad (3.1)$$

where  $\mathcal{F}$  is the free energy,  $o$  denotes outcomes or input data and  $P(o)$  represents the surprisal in probabilistic terms. This equation states that minimizing the free energy minimizes the surprisal as well (Oliver et al., 2021). The reason we do not minimize the surprisal directly is due to its intractability when scaling beyond very low-dimensional distributions (Da Costa et al., 2020). The variational free energy bound combats this mathematical pit via the use of variational inference (Blei et al., 2017), an approximate Bayesian inference that uses an approximate distribution  $Q$  of the generative model's distribution  $P$ . The variational free energy can be defined as:

$$\mathcal{F} = \underbrace{D_{KL}[Q(s)||P(s)]}_{Complexity} - \underbrace{E_{Q(s)}(\ln P(o|s))}_{Accuracy} \quad (3.2)$$

where  $s$  denotes the hidden state of the system,  $D_{KL}$  is the Kullback–Leibler divergence that measures the distance between two distributions and  $E_{Q(s)}$  is the expectation of the observation likelihood given the approximate distribution. Eq. 3.2 shows that the free energy is at its minimum in case where both the complexity and accuracy term are equal to zero. The complexity term compares the approximate posterior distribution  $Q(s)$

with the model’s prior distribution  $P(s)$ . This value is minimized when the model’s prior knowledge only explains the new stimulation forming the posterior. The accuracy term assures the agent of inferring those (hidden) states that can explain the stimuli  $o$  the most. It is computed by averaging over the posterior approximate distribution  $Q(s)$  the likelihood function of the model. Without the complexity term, the accuracy would lead to observations overfitting, and without the accuracy term the agent would ignore the outside world and focus on the model’s beliefs without updating them. Hence, free energy is a balance of exploitation (i.e., fulfilling our priors) and exploration (i.e., inferring the most informative hidden states).

### 3.2.2.1 Action selection.

We define the agent’s action selection as computing the optimal policy  $\pi$ , given the combination of the habitual behaviours  $E^1$  (i.e. a behaviour occurring automatically, without a complex inference) and the optimization of the Expected Free Energy ( $G$ ) of executing a plan of actions in the future<sup>2</sup>:

$$\pi = \arg \max Q(\pi) = \arg \max \sigma[\rho \ln E(\pi) - G(\pi)] \quad (3.3)$$

where  $\rho$  is the precision parameter that modulates the interplay between habitual policies and new computed optimal policies, given the model and new observations. The  $\sigma$  represents a softmax function that in a discrete state representation transforms the objective values into probabilities. Thus, the policy with the higher probability is selected.

To enable the agent to plan actions ahead (Lanillos et al., 2021), we used the Expected Free Energy (EFE) objective function  $G$ , which predicts the free energy value in the next time-point (Kaplan and Friston, 2018):

$$G(\pi) = \underbrace{D_{KL}[Q(s_\tau|\pi)||P(s_\tau)]}_{Risk} - \underbrace{E_{Q(s_\tau|\pi)P(o_\tau|s_\tau)}[\ln P(o_\tau|s_\tau, A)]}_{Ambiguity} \quad (3.4)$$

where  $\pi$  is a policy/plan (i.e. a sequence of actions),  $\tau$  is a time point, and  $A$  represents the likelihood function of the model. The terms risk and ambiguity closely follow the terms complexity and accuracy from Equation 3.2, respectively. A relevant difference

---

<sup>1</sup> $E$  can be provided by the designer or learned and encodes the agent’s preferences given by evolution.

<sup>2</sup>Note that the EFE is a generalization of the variational free energy conditioned on an action plan or policy.

between the free energy  $F$  and  $G$  is that the latter is sampled from the predictive posterior distribution. This means that the agent predicts the consequences of its actions and calculates the deviance from surprise based on these. The risk calculates the posterior value of a hidden state in the next time-step, given a policy selected at a given moment that is compared to the prior beliefs of the agent. This assures a selection of such a policy that leads to fulfilling the priors. On the other hand, the ambiguity term functions similarly to the accuracy term mentioned above, but the likelihood is predictive (over the next time step) based on the averaged approximate distribution of an expected hidden state given a policy  $Q(s_\tau|\pi)$ . This means that the ambiguity term selects those policies that are most informative.

### 3.2.3 Precision modulation

As described in the introduction, we suggest that active sensing is modulated by neurotransmitters, such as serotonin, which we assimilate to sensory precision or habitual precision. Sensory precision modulates the likelihood function of the model, while the habitual precision modulates policy selection (through EFE computation).

#### 3.2.3.1 Likelihood modulation.


Sensory precision  $\zeta$  modulates the impact of the sensory input.

$$A = \sigma(\zeta \ln(\mathbf{A} + \epsilon)) \quad (3.5)$$

where  $\sigma$  represents a softmax function, and  $\epsilon$  is some arbitrary small number that prevents a mathematical error due to taking a logarithm of 0. The bold matrix  $\mathbf{A}$  denotes a distribution shown in Figure 3.2 (shown in the middle panel) and represents the likelihood mapping between states and observations. Since only those columns linked to tactile sensing have a non-uniform distribution, the equation above then implicitly modulates only these columns, as the inverse temperature parameter  $\zeta$  cannot modulate the uniform distribution.

#### 3.2.3.2 Policy selection modulation.

The habitual precision parameter  $\rho$  is described in Equation 3.3 of action selection. The  $E$  vector prespecifies a preference over the habitual policy, and thus the precision parameter can only minimize or maximize the difference, but not shift it around and make a habitual preference for a small movement policy. High  $\rho$  thus indicates a strong preference for the broad range of movement, while low  $\rho$  makes the habitual preferences disappear.



**Figure 3.2:** *The generative model of active sensing in the Nao humanoid.* The model contains two hidden factors: object appearance and arm position. The former indicates where the object will appear in the environment, defined by the arm position. There are three objects that can be linked to three distinct arm positions, and also a case when there is no object which is what the agent believes at the beginning of every trial. Next, there are 5 arm positions, where the middle three positions are duplicated due to the sinewave-like behaviour (i.e., going back and forth). There is only one outcome modality, touch, which takes a boolean value of either touched or untouched. The likelihood matrices depend on what context (i.e., object appearance) is inferred. We defined the likelihood function  $A$  in an uncertainty-minimization environment, where the motivation for the agent to explore an object (where the stimulation is sensed) is defined by a uniform likelihood distribution except for the part where the hand can sense an object. In the matrix, the black colour relates to high probability, while white boxes are of low probability. The grey boxes represent uncertain distribution. Lastly, we specified two policies. The first one, which has a higher habitual value, is related to a complete hand circle which is a logical way of exploring the (uncertain) world as fully as possible. The second policy provides some nuances on how to escape the former policy and stick with the most certain outcome related to a given arm position. This policy has a lower habitual value and therefore its selection depends on the uncertainty minimization behaviour. These actions are then derived from transition matrices  $B$  specified in Figure 3.1.

### 3.2.4 Model tailored to robotic tactile active sensing

Our active sensing model is described in Figure 3.2. Inspired by the gap-crossing paradigm in rodents (Voigts et al., 2008), we designed the robotic model with two hidden factors, which describe the state: object appearance (or context) and arm position. The former hidden factor is split into four states. Three states (indicated with black, grey, and white boxes) reflect three different positions of where the object can be detected. The fourth hidden state assumes that no object can be detected. For the hidden factor arm position, we specify five different arm positions. Note the three middle positions are duplicated so we can provide a back and forth movement in a discrete space, as much as in whisking movement. Therefore, the arm positions are represented by eight different states. There is one outcome modality in the model, which is associated with touch. The outcome is then either being touched or not. Following the outcome modalities and hidden states, we can provide likelihood functions, specifying the probability of being touched, given a specific context and arm position. If the agent infers that there is no stimulus around, the likelihood function is completely uniform and thus uninformative<sup>3</sup>. This is provided as a prior for the model at the beginning of every trial. Next, the three positions in which an object can appear are associated with distinct arm positions (i.e. an agent can only infer that there is an object in a position when the arm is in that same position). These likelihood distributions then provide a precise likelihood distribution only over the adequate association of the object’s appearance and arm position. Note that these exact distributions are then modulated with the  $\zeta$  precision parameter provided in Equation 3.5. Lastly, the policies (or actions) are split into two. The former one (indicated in yellow colour in the figure) is related to the large movement amplitude. This policy has a higher habitual value in the vector  $E$ , as we assume that the bigger hand movements should be more intuitive. The small amplitude of the hand movement is indicated by the purple colour. This policy assures the possibility of switching to the smaller, more precise, movements, that are then related to active sensing strategies. The behaviour for the switch is to find the most certainly interpretable outcome and stick with it, hence uncertainty minimizing behaviour.

---

<sup>3</sup>It might seem counter-intuitive that without any stimulation, the agent does not recognize that nothing has been touched, but rather increases the uncertainty of the world. This is motivated by the so-called dark room problem (Friston et al., 2012; Baltieri and Buckley, 2019) that resolves a problem in the free energy that states that, if agents minimize the uncertainty of the world, why not stay in completely dark environments which should provide completely predictable stimuli (i.e., always dark)? The answer is that perceiving darkness is highly ambiguous, therefore we treat no stimulation in the same way and build likelihood matrices in such a way that increases uncertainty over sensing no-stimulation signals while the stimulation outcomes provide higher certainty of causes.

### 3.2.5 Experimental design


The model was validated in the humanoid Nao robot (Nao, 2017) using an object exploration experiment using one limb, inspired by the rodents navigation paradigm. The goal of the robot is to explore the environment looking for objects. Once the robot infers there is an object in a given position by sensing the object with the touch sensor, we expect a switch in behaviour from broad movement amplitude to small movement amplitude as it will start minimizing the uncertainty over a more certain input. In other words, the big movement amplitude is used when the environment is more uncertain and all inputs are equally uncertain. The object provides more certainty to the agent, hence it makes sense that the agent will stay around more certain inputs with a small movement amplitude.

To test distinct tactile exploratory processes under precision modulation, we conducted experiments in which a Nao robot moves its arm back and forth as hard-coded based on our model (see Figure 3.2). The experiments contain 40 time-steps (with one exception of 80 time-steps), where a time step is simply defined by an arm position. Hence, during a single experiment, the robot moves its arm 40 times. Each time-step takes approximately 2.5 seconds (communication speed between the robot and the host computer and the robot physically moving between two states), which means that running an experiment takes approximately one minute and forty-five seconds. The Nao robot will always start with the first arm position. After the 8th time-step, we introduce an object (the researcher’s hand) at a specific arm position.

### 3.2.6 Implementation

To implement the model that we described in section 3.2.1, we used the open-source ‘pymdp’ package (Heins et al., 2022). This software implements active inference agents in discrete time, state, and action spaces. To translate the model into the physical Nao robot we used the NAOqi driver. Thus, we developed the interface between the NAOqi SDK and the pymdp algorithm to control the robot. Since the NAOqi software only runs on Python 2.7 and the python model needs to be implemented in Python 3 because of the pymdp library, we implemented a socket module to communicate observations from the physical robot to the model and the actions from the model to the robot.

Furthermore, the acting and sensing computations of the robot were subdivided into parallel threads to allow the robot to sample faster in the environment when moving the arm down and up in a particular state, thereby increasing the probability that a touch is observed. The acting thread was responsible for moving the robot’s arm and for interacting with the model script. To move the robot’s arm, we used the motion



**Figure 3.3:** *Illustrations of the Nao robot used for the experiments.* a) A particular arm position of the Nao robot in the experiment. The red dots indicate all states of the robot, where the most right arm position is referred to as arm position 1 and the most left arm position is referred to as arm position 5. The darker red line indicates the trajectory between arm positions and the dots indicate a movement down and up again to sense the environment. b) The location of the touch sensor on the back of the robot's right hand indicated by a red circle. The touch sensor is a capacitive sensor that can distinguish between two different states: touched and not touched.

module of the Nao robot (ALMotion) to access the joints of the robot's right shoulder. In accordance with the model, five arm positions have to be specified and the three middle positions have to be duplicated, so the robot can move back and forth between states in discrete space (see Figure 3.3a)<sup>4</sup>. When the right hand arrives in a new state, it will slightly move down (0.1 radians vertically) in order to sense whether there is an object in the state. Afterward, it will move up again and continue to the next state.

The sensing thread was responsible for registering whether the touch sensor of the robot is touched and updating the variables in the script accordingly. To register a touch, we used a touch sensor on the back of the robot's right hand. In particular, we only used the middle touch sensor on the back of the robot's hand, which is a capacitive sensor with a binary output (touched or not touched) (see Figure 3.3b for the location of this touch sensor in red). The touch sensor is always facing down (angle is 1.74 radians) (see Figure 3.3a for the positioning of the hand in a particular state). By accessing the memory of the Nao robot (ALMemory) and reading out the value for the back touch sensor on the right hand, we could infer whether the sensor was touched (value is 1) or not (value is 0)

<sup>4</sup>Note that the movement of the robot's arm is only horizontal. Therefore, we could fix the joint position of the shoulder pitch to 0.1 radians. For the shoulder roll, we chose to evenly space the joint positions between 0.2 and -1.0 radians, such that the shoulder roll has to move 0.3 radians between states.



<sup>5</sup>. Because reading from memory is very costly, we introduced locks in our solution to only read from memory when the robot’s hand moves up and down in a state and not when moving between different states.

In general, the communication between the two scripts proceeded as follows: first, the acting thread read whether the robot is touched in the previous state and sent this observation to the model. The model subsequently computed the next action for the robot to execute and sent this action to the acting thread. This thread then moved to the next state, based on the action. Afterward, it moved the hand down and up again in this next state, while the sensing thread simultaneously read from memory whether or not the touch sensor on the robot was touched. If the sensor was touched, the thread would update the shared variable storing the touch value accordingly, such that this observation was available for the acting thread in the next time step. Detailed information on the software implementation can be found in Figure 3.7, Appendix 3.7.1.

### 3.3 Results

First, based on observations from the simulations, we evaluated behavioural and perceptual differences depending on the precision modulation parameter ( $\zeta$  and  $\rho$ ) values that are the hypothesized homologues of serotonergic neuromodulation and that affects active sensing behaviour. Second, based on the simulations we introduced a qualitative analysis studying the robot NAO behaviour with low and high  $\zeta$  values, combined with low and high  $\rho$  values.

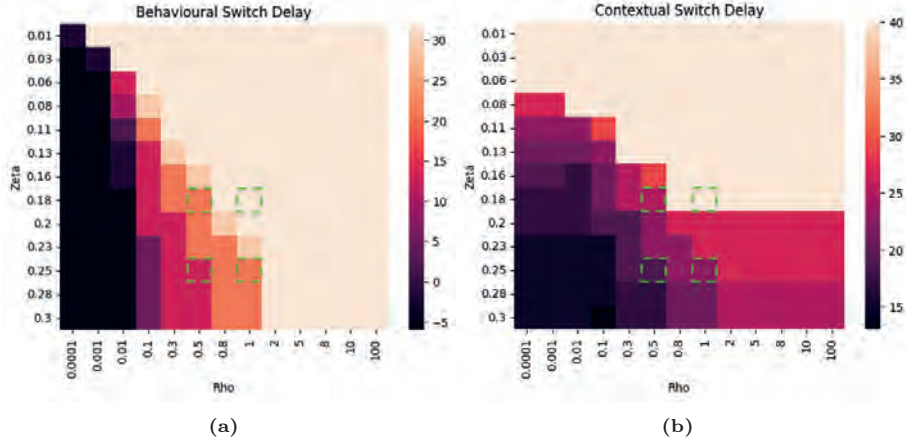
#### 3.3.1 Simulations

We analyzed the experiment in simulation for 13 different  $\zeta$  values and 13 different  $\rho$  values, selected to capture the whole range of potential behaviours. All of the experiments followed the structure as described above, meaning 40 time-steps per experiment, where after the 8th time step we simulated object appearance at arm position 4. We thus ran 169 (13x13) experiments, 8 times each. Note that for the simulated experiments, the starting arm position is randomized. These experiments show numerical differences in tactile exploratory behaviour under distinct precision modulation. Specifically, we focused on *i*) the time when the (behavioural) switch from a large to small hand movement amplitude; and *ii*) the switch of the object appearance posterior occurs after introducing the object.

Figure 3.4a represents the time it takes to switch from a large to a small hand

---

<sup>5</sup>We also used the speech module of the Nao robot (ALTextToSpeech) to have an auditory cue which could indicate whether the sensor was touched.



**Figure 3.4:** Numerical analyses of the active sensing behaviour. Figure (a) represents a time delay of behaviour switch. Figure (b) shows a time delay of inferred contextual switch. Based on these analyses, we decided to simulate four precision combination in the humanoid that are represented in the dashed green squares in both heatmaps.

movement amplitude from the moment an object has been touched. Here, we can see that the increase of  $\zeta$  and decrease of  $\rho$  gradually speeds up the switch. Figure 3.4b represents the same problem, but for inferring the correct context (i.e. that an object is located at arm position 4). This has been calculated as a time difference from an object has been touched to the point where context 4 (a context where an object is believed to appear to arm position 4) has a higher posterior probability than context 1 (the initial context where no object is believed to be in front of the humanoid). As shown in the figure above, all cases show a gradual change in the same manner for both precision parameters. Based on these analyses, we selected four different cases where behaviour is distinct from each other. This is highlighted in dashed green boxes of Figure 3.4, which were then applied in the humanoid. Based on this finding, we also conclude that there is a strong binding of lower sensory precision with stronger habitual behaviour, at least in tactile detection and exploratory domain, as these two appear to modulate the behaviour in the same manner. This is because the effects of these two precision terms provide the same function when it comes to a temporal delay of active sensing. Given the methodological structure of (Azarfar et al., 2019) that compared animals with high and normal levels of serotonin, the upper-left and bottom-right green dashed boxes can be regarded as a behaviour from the wild-type animals, providing the baseline model. The other two boxes then demonstrate a behaviour of agents with higher levels of serotonin.

### 3.3.2 Validation with humanoid

We evaluated the model on the physical Nao robot. Videos of the experiments can be found in Appendix 3.7.3. Figure 3.5 shows the results after running the four experiments in the Nao robot. In these simulations, an object appeared after the 8th time-step at arm position 4 (which can also be inferred from the plots representing the observation and arm position). Compared to Figure 3.4, it seems apparent that the simulated results with the robotics application are consistent. First of all, the plots of the arm position (upper-right panels) show a trend that increasing  $\zeta$  or decreasing  $\rho$  leads to a faster switch in policy from a large to a small movement amplitude. Specifically, the switch can happen up to 8 time-steps faster (representing one full arm swing). This can also be seen in the plots of the policy posterior in the bottom-left panels (i.e. the value for the small amplitude is higher than the value for the large amplitude at an earlier time-step). Next, the plots representing the context posterior show that with a higher  $\zeta$  or lower  $\rho$ , the correct context can be inferred four time-steps faster (i.e. half an arm swing).

In addition to testing whether the simulation results when varying  $\zeta$  and  $\rho$  can be validated in experiments with the Nao robot, we also conducted an experiment to see whether behaviour is similar when the robot is touched in different arm positions. Results between touching the robot in arm position 3 (see Figure 3.6a) and in arm position 4 (see Figure 3.5c) under the same precision parameters show that this is indeed the case. Namely, the switch between policies occurs at the same time (considering it takes one time-step longer to see the change in policy for Figure 3.6a) and also the context posteriors for the respective arm position (i.e. 'Touched right' (red line) in Figure 3.5c and 'Touched middle' (yellow line) in Figure 3.6a) seem to increase with a similar trend. The difference after time-step 32 is due to a missed sensory stimulation. The missed sensory stimulation also provides an interesting case where the consequences of unreliable sensors affect active sensing. This then leads to a longer exploration and a sudden drop in the context posterior. This indicates a loss of confidence in the agent's beliefs about understanding the newly acquired structure of the environment, causing it to explore for longer.

Lastly, we conducted a longer experiment with 80 time-steps (see Figure 3.6b), where we switched from touching the robot in arm position 2 to arm position 3 after 30 time-steps (indicated by the red line). As can be seen in the plot for the context posterior (bottom-right panel), the robot was able to correctly infer the change in the context (i.e. the context of being touched in arm position 2 goes down after 30 time-steps and the context of being touched in arm position 3 goes up) and adapted its behaviour accordingly (i.e. after time-step 30, the robot switches back to large movement amplitude and after time-step 45, the robot switches to a small movement amplitude around arm position 3).



These results suggest that the model can quickly adapt to new information, as a switch in arm position leads to a switch in behaviour and posterior probabilities. Doing so, we just touched the surface of the robotics application for the neuroscientific research, but such extensions can inform neuroscience and psychology research more about the assumptions originating from these fields, such as adaptive behavioural switching in an environment where objects are not stationary, something that the gap-crossing paradigm does not consider.


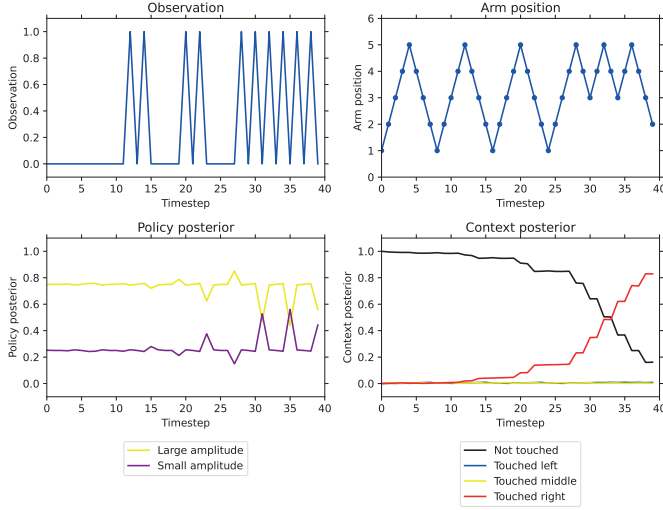
### 3.4 Discussion

We studied the behaviour of a humanoid robot with serotonergic depletion using the active inference framework (Adams et al., 2013; Friston et al., 2017a; Da Costa et al., 2020) – where we associated serotonin with a modulation of sensory or habitual precision.

Although this work was inspired by rodent behaviour with serotonin alteration (Azarfar et al., 2019), humans also show a similar aspect of switching from a coarse-grained action execution to a more fine-grained one – once an object or environment is being explored. The more detailed behaviour provides more information about the texture or shape, or even motion of the object (Ryan et al., 2021). In our case, we focused on exploratory strategies in the tactile domain which extends current neuroscientific and psychological research focusing on exploration mostly in vision (Najemnik and Geisler, 2005, 2008; Renninger et al., 2007). We also extend this by focusing on subtle changes in neuromodulation – specifically serotonin – and show that to fully comprehend how we are able to resolve uncertainty in the world via adjusting our sensors, neuromodulatory systems also have to be investigated.

There are, however, other computational models focusing on the function of serotonin. Most of them are developed under the reinforcement learning framework focusing on assertive behaviour (Daw et al., 2002; Doya, 2002; Dayan and Huys, 2008; Cools et al., 2011). These models thus analyze the exploitative behaviour, i.e., how to behave to collect the least punishment. Our model – where we modelled (consequences of) altered levels of serotonin via habitual and sensory precision in order to bias exploratory behaviour – seems to be compatible with previous models as well. In particular, (Dayan and Huys, 2008) suggested that modulating the reward value – in order to promote assertive behaviour linked to serotonin (Dayan and Huys, 2009) – limits the exploration of cases where the reward is expected to be lower. Hence our speculation is that the consequence of an altered level of serotonin (especially from birth which is the case of SERT-/- mice) can lead to changes in sensory and/or habitual precision. Aligning exploratory and exploitative models of serotonin is thus an interesting field for future research.

(a)  $\zeta = 0.18, \rho = 0.5$ (b)  $\zeta = 0.18, \rho = 1.0$

(c)  $\zeta = 0.25, \rho = 0.5$ (d)  $\zeta = 0.25, \rho = 1.0$ 



**Figure 3.5:** *Application of the generative model in a humanoid.* This figure represents the experiments conducted in the Nao robot for different values of  $\zeta$  and  $\rho$ , where an object appears in arm position 4 after the 8th time-step. For each experiment, the top-left panels labeled as 'Observation' represent whether or not the robot is touched for every time-step. Here, a value of 1 means that a sensor was triggered, while a value of 0 says that no touch was sensed. The top-right panels 'Arm Position' represent the (discrete) position of the hand per time-step. (Continue on the next page.)

**Figure 3.5:** (Continued) There are five hand positions that can be reached. These graphs nicely show the switch from the large amplitude to the small one. The bottom-left panels represent the 'Policy Posterior' for the large amplitude (yellow) and small amplitude (purple). Lastly, the bottom-right panels 'Contextual posterior' represent the posterior probabilities for four different contexts in which the agent can find itself. It starts with the preference over 'Not touched' and eventually switches to 'Touched right', which is the context where a hand was touched at position 4 as specified above.

Although we linked our analysis to serotonin specifically, this simulation can also test the tactile exploratory differences in people with Autism Spectrum Disorder (ASD). This is theoretically possible, as previously it has been suggested that ASD is associated with aberrant sensory precision (Van Boxtel and Lu, 2013; Lawson et al., 2014; Palmer et al., 2017). These theories relate ASD to an increase in sensory precision. In this work, as we specifically tested the sensory precision parameter  $\zeta$ , our analysis can also be used to explore ASD abnormalities in the tactile exploratory domain. Based on our results and following the theories, we hypothesize that the time of switching to a detailed exploratory hand movement will happen faster in comparison to a typical group. Such associations can also deepen the connection between serotonergic alterations in ASD (Muller et al., 2016).

In our work, we specifically focus on how active sensing is modulated via serotonin which extends on previous assumptions of active sensing models (Yang et al., 2016; Martinez-Hernandez et al., 2017; Bajcsy et al., 2018; Parr et al., 2021b). In addition to this, we specifically motivate our model via the dark-room problem (Friston et al., 2012) which states that having no information (i.e., not sensing any object) increases the uncertainty of the world rather than decreasing it via holding some understanding (or a related hidden state) that is certain about such a situation. Although this approach is motivated by philosophy (Sun and Firestone, 2020) and is thus speculative, we nevertheless show that we can use it as a motivation to create a model of active sensing. Thus, our modelling contributes to the field of neurorobotics that aims to strengthen both robotics applications and create better replicas of biological agents (Krichmar, 2018; Tani and White, 2022) via the application of (purposefully built) uncertain environments to trigger active sensing.

Lastly, our work has touched upon the concept of embodiment in distinct living systems but also distinct architectures of robots (Glenberg, 2010; Lux et al., 2021). We highlight the analogies of different bodies (c.f., humans and rodents) that are exploited by agents in order to fulfil the same task, which is to collect the most informative stimuli from the world to resolve a model's uncertainty of it. Surely, different bodies provide different strategies, but as we have shown here, some strategies can be united and driven

(a)  $\zeta = 0.25, \rho = 0.5$ (b)  $\zeta = 0.5, \rho = 0.5$ 

**Figure 3.6:** *Additional simulations in the robot.* The template of results is the same as in Figure 3.5. Figure 6a represents an interesting case when an object appeared at the arm position 3; and one stimulation at the 32nd time-step was missed, leading to more exploration compared to Figure 5c. In addition, this led to a sudden drop in context posterior. Figure 6b is slightly different, as there are 80 time-steps instead of 40 as in previous experiments. The red lines here specify when a context was changed (i.e., when a researcher’s hand position appeared at a different position).



very similarly, such as exploring an environment using whiskers or hand movement. In robotics, there are several distinct architectures of robots to fulfil distinct tasks (Rubio et al., 2019). Although this is outside the scope of this paper, we hope to motivate future researchers to play with the concept of embodiment and see what strategies are universal across several bodies, and what changes lead to different neural encodings, for instance.

### 3.4.1 Limitation

Here we consider several limitations of our work that are worth mentioning. First, the difference between hand movement and animal whisking is obviously distinct, and the fact that both are used for tactile exploration can be misleading from the neuroscientific point of view, as these processes are encoded with different importance (i.e. it is less important to lose one whisker but more tragic to lose a human arm). For instance, the speed of whisking during object exploration is up to 25Hz (Berg and Kleinfeld, 2003), which is a speed that cannot be reached by a hand. However, the model is simplified enough to justify the use of both models for both species independently, and the link should be further tested to suggest whether this generalization is possible. Furthermore, while the model works with discrete actions—which can be action primitives—its extension to continuous control could be interesting for robot exploration tasks. The models combining discrete and continuous domains have been shown in (Parr et al., 2021a, 2022). As a final practical consideration, we acknowledge that the touch sensors on the Nao robot are not suited for object exploration, as reliable stimulation of the touch sensor only occurred when using a researcher’s hand (see also Appendix 3.7.2).

## 3.5 Conclusion

We described an active sensing model inspired by serotonergic modulation under the active inference framework. Particularly, we designed a generative model to perform tactile exploration in a humanoid robot. While the presented robotic model is still far away from human tactile exploration we suggested some possible relations. Specifically, we have shown that to resolve uncertainty in the environment, the influence of high sensory and/or low habitual precision on the inference process leads to a faster appearance of an active sensing behaviour. Moreover, our findings also contribute to scrutinizing the vast functions of serotonin (Berger et al., 2009) that in our case are associated with the aforementioned parameters. Next, we highlight that it might be a complex task to dissociate the effect of habitual behaviour with imprecise sensory estimation at least in active sensing. Whether this is a universal criterion of behaviour and thus extends to

other behavioural processes such as social cognition (Lockwood and Klein-Flügge, 2021) or reward learning (Lindström et al., 2021), and whether these are regulated via serotonin in the same manner, should be further tested.

## 3.6 Acknowledgement

FN was supported by the EU H2020 MSCA ITN project ‘Serotonin and Beyond’ (N 953327). We thank Tansu Celikel and Thomas Parr for their helpful discussions.

# Bibliography

- (2017). Nao software 1.14.5 documentation.
- Adams, R. A., Shipp, S., and Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Structure and Function*, 218:611–643.
- Azarfar, A., Zhang, Y., Alishbayli, A., Schubert, D., Homberg, J. R., and Celikel, T. (2019). Serotonergic development of active sensing. *BioRxiv*, page 762534.
- Bajcsy, R., Aloimonos, Y., and Tsotsos, J. K. (2018). Revisiting active perception. *Autonomous Robots*, 42:177–196.
- Baltieri, M. and Buckley, C. L. (2019). The dark room problem in predictive processing and active inference, a legacy of cognitivism? In *Artificial Life Conference Proceedings*, pages 40–47. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info . . . .
- Bang, D., Kishida, K. T., Lohrenz, T., White, J. P., Laxton, A. W., Tatter, S. B., Fleming, S. M., and Montague, P. R. (2020). Sub-second dopamine and serotonin signaling in human striatum during perceptual decision-making. *Neuron*, 108(5):999–1010.
- Berg, R. W. and Kleinfeld, D. (2003). Rhythmic whisking by rat: retraction as well as protraction of the vibrissae is under active muscular control. *Journal of neurophysiology*, 89(1):104–117.
- Berger, M., Gray, J. A., and Roth, B. L. (2009). The expanded biology of serotonin. *Annual review of medicine*, 60:355–366.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877.

- Cools, R., Nakamura, K., and Daw, N. D. (2011). Serotonin and dopamine: unifying affective, activational, and decision functions. *Neuropsychopharmacology*, 36(1):98–113.
- Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., and Friston, K. (2020). Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology*, 99:102447.
- Daw, N. D., Kakade, S., and Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural networks*, 15(4-6):603–616.
- Dayan, P. and Huys, Q. J. (2009). Serotonin in affective control. *Annual review of neuroscience*, 32:95–126.
- Dayan, P. and Huys, Q. J. M. (2008). Serotonin, inhibition, and negative mood. *PLoS computational biology*, 4(2):e4.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural networks*, 15(4-6):495–506.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2):127–138.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017a). Active inference: a process theory. *Neural computation*, 29(1):1–49.
- Friston, K., Thornton, C., and Clark, A. (2012). Free-energy minimization and the dark-room problem. *Frontiers in psychology*, page 130.
- Friston, K. J., Parr, T., and de Vries, B. (2017b). The graphical brain: belief propagation and active inference. *Network neuroscience*, 1(4):381–414.
- Glenberg, A. M. (2010). Embodiment as a unifying perspective for psychology. *Wiley interdisciplinary reviews: Cognitive science*, 1(4):586–596.
- Heins, C., Millidge, B., Demekas, D., Klein, B., Friston, K., Couzin, I., and Tschantz, A. (2022). pymdp: A python library for active inference in discrete state spaces. *arXiv preprint arXiv:2201.03904*.
- Homberg, J. R., Schubert, D., Asan, E., and Aron, E. N. (2016). Sensory processing sensitivity and serotonin gene variance: Insights into mechanisms shaping environmental sensitivity. *Neuroscience & Biobehavioral Reviews*, 71:472–483.
- Kaboli, M., Feng, D., Yao, K., Lanillos, P., and Cheng, G. (2017). A tactile-based framework for active object learning and discrimination using multimodal robotic skin. *IEEE Robotics and Automation Letters*, 2(4):2143–2150.

- Kaplan, R. and Friston, K. J. (2018). Planning and navigation as active inference. *Biological cybernetics*, 112(4):323–343.
- Krichmar, J. L. (2018). Neurorobotics—a thriving community and a promising pathway toward intelligent cognitive robots. *Frontiers in neurorobotics*, 12:42.
- Lanillos, P., Meo, C., Pezzato, C., Meera, A. A., Baioumy, M., Ohata, W., Tschantz, A., Millidge, B., Wisse, M., Buckley, C. L., et al. (2021). Active inference in robotics and artificial agents: Survey and challenges. *arXiv preprint arXiv:2112.01871*.
- Lawson, R. P., Rees, G., and Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in human neuroscience*, 8:302.
- Lindström, B., Bellander, M., Schultner, D. T., Chang, A., Tobler, P. N., and Amodio, D. M. (2021). A computational reward learning account of social media engagement. *Nature communications*, 12(1):1311.
- Lockwood, P. L. and Klein-Flügge, M. C. (2021). Computational modelling of social cognition and behaviour—a reinforcement learning primer. *Social Cognitive and Affective Neuroscience*, 16(8):761–771.
- Lux, V., Non, A. L., Pexman, P. M., Stadler, W., Weber, L. A., and Krüger, M. (2021). A developmental framework for embodiment research: The next step toward integrating concepts and methods. *Frontiers in Systems Neuroscience*, 15:672740.
- Martinez-Hernandez, U., Dodd, T. J., Evans, M. H., Prescott, T. J., and Lepora, N. F. (2017). Active sensorimotor control for tactile exploration. *Robotics and Autonomous Systems*, 87:15–27.
- Meera, A. A., Novicky, F., Parr, T., Friston, K., Lanillos, P., and Sajid, N. (2022). Reclaiming saliency: rhythmic precision-modulated action and perception. *arXiv preprint arXiv:2203.12652*.
- Miller, K. J., Shenhav, A., and Ludvig, E. A. (2019). Habits without values. *Psychological review*, 126(2):292.
- Muller, C. L., Anacker, A. M., and Veenstra-VanderWeele, J. (2016). The serotonin system in autism spectrum disorder: From biomarker to animal models. *Neuroscience*, 321:24–41.
- Najemnik, J. and Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434(7031):387–391.

- Najemnik, J. and Geisler, W. S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of Vision*, 8(3):4–4.
- Oliver, G., Lanillos, P., and Cheng, G. (2021). An empirical study of active inference on a humanoid robot. *IEEE Transactions on Cognitive and Developmental Systems*, 14(2):462–471.
- Palmer, C. J., Lawson, R. P., and Hohwy, J. (2017). Bayesian approaches to autism: Towards volatility, action, and behavior. *Psychological bulletin*, 143(5):521.
- Parr, T. and Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of the Royal Society Interface*, 14(136):20170376.
- Parr, T., Limanowski, J., Rawji, V., and Friston, K. (2021a). The computational neurology of movement under active inference. *Brain*, 144(6):1799–1818.
- Parr, T., Pezzulo, G., and Friston, K. J. (2022). *Active inference: the free energy principle in mind, brain, and behavior*. MIT Press.
- Parr, T., Sajid, N., Da Costa, L., Mirza, M. B., and Friston, K. J. (2021b). Generative models for active vision. *Frontiers in Neurorobotics*, 15:651432.
- Rasouli, A., Lanillos, P., Cheng, G., and Tsotsos, J. K. (2020). Attention-based active visual search for mobile robots. *Autonomous Robots*, 44(2):131–146.
- Renninger, L. W., Verghese, P., and Coughlan, J. (2007). Where to look next? eye movements reduce local uncertainty. *Journal of vision*, 7(3):6–6.
- Rubio, F., Valero, F., and Llopi-Albert, C. (2019). A review of mobile robots: Concepts, methods, theoretical framework, and applications. *International Journal of Advanced Robotic Systems*, 16(2):1729881419839596.
- Ryan, C. P., Bettelani, G. C., Ciotti, S., Parise, C., Moscatelli, A., and Bianchi, M. (2021). The interaction between motion and texture in the sense of touch. *Journal of Neurophysiology*, 126(4):1375–1390.
- Sun, Z. and Firestone, C. (2020). The dark room problem. *Trends in Cognitive Sciences*, 24(5):346–348.
- Tani, J. and White, J. (2022). Cognitive neurorobotics and self in the shared world, a focused review of ongoing research. *Adaptive Behavior*, 30(1):81–100.

Van Boxtel, J. J. and Lu, H. (2013). A predictive coding perspective on autism spectrum disorders.

Voigts, J., Sakmann, B., and Celikel, T. (2008). Unsupervised whisker tracking in unrestrained behaving animals. *Journal of neurophysiology*, 100(1):504–515.

Yang, S. C.-H., Wolpert, D. M., and Lengyel, M. (2016). Theoretical perspectives on active sensing. *Current opinion in behavioral sciences*, 11:100–108.

## 3.7 Appendix

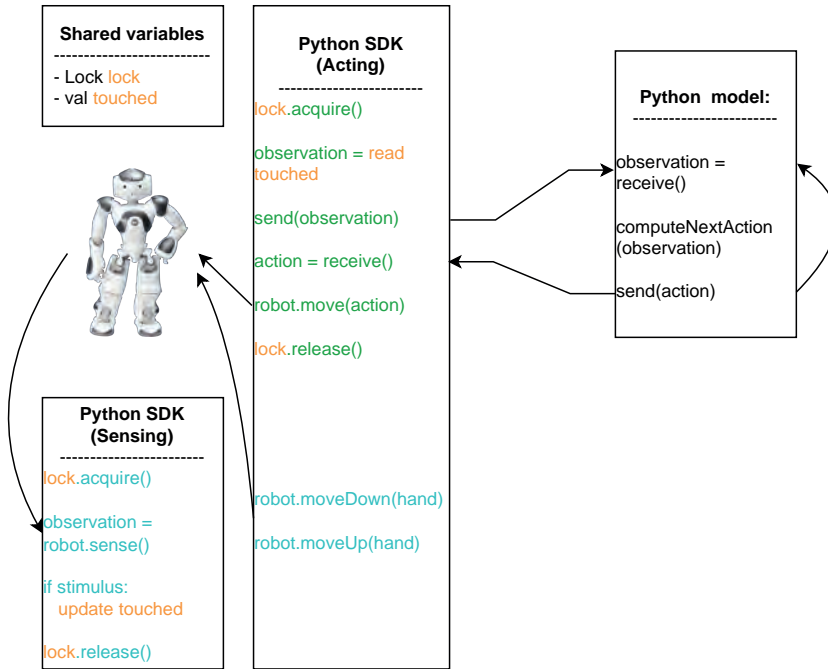
### 3.7.1 Python scripts

The scripts that we used in order to simulate the experiment and run the experiment in the Nao robot can be found on the following GitHub page: <https://github.com/SimonJanssen1/Serotonin-based-active-sensing.git>

In order to generate the heatmaps in section 3.3.1, we used a Jupyter Notebooks file in which we iterated through different combinations of  $\zeta$  and  $\rho$  parameters, while simulating a touch in arm position 4. We also used the scripts as explained in section 3.2.6 (see also Figure 3.7 for a visual description of the software architecture) in order to conduct the experiment with the Nao robot. Next, we simulated the script communicating with the Nao robot to test whether the architecture is working correctly. In this script, the general architecture is the same with an acting and sensing thread and some shared variables. However, instead of communicating with the robot to sense the environment and execute actions, the script simulates these behaviours. For simulating touch, we defined three options: 1) simulating touch at specific time-steps, which could be indicated by means of a list, 2) simulating touch in a specific arm position, or 3) simulating touch randomly, where a probability of touch could be specified.

### 3.7.2 Experiment with different sensory stimulation

In addition to the experiments with the Nao robot presented in section 3.3.2, we conducted another experiment where we touched the robot in arm position 4 with a metal object instead of using a researcher’s hand. Because the robot has capacitive sensors on the back of the right hand, it should still be able to sense a touch when using a metal object. However, Figure 3.8 shows that this is not the case. Namely, the robot only registers a touch four times, whereas we expected it to register a touch every time the arm was at




**Figure 3.7:** *The software architecture of the Nao robot to conduct the experiment.* The script that interacts with the robot consists of two different threads that are responsible for acting and sensing, respectively. The acting thread reads an observation from the shared variable storing the touch value, which it sends to the python model. The python model subsequently computes the next action for the robot, given the observation, and sends this action back to the acting thread, which uses it to move to the next state. When in the new state, the robot moves its arm down and up again, during which the sensing thread is also active and records whether or not the touch sensor of the robot is touched. If the robot's sensor is touched, the sensing thread will make sure to update the shared variable storing the touch value accordingly, such that the acting thread can send the observation to the python model in the next time step.

position 4. These results thus show that the back touch sensor on the hand of the Nao robot was not very sensitive to touch with materials other than a hand.

### 3.7.3 Experiment videos

During the experiments, we also recorded the humanoid's behaviour. These results can be found under the following link: <https://drive.google.com/drive/folders/1QFmy1nHtQ1mWxmKPZa-kbgj84FBJ9Q6n?usp=sharing>.





**Figure 3.8:** *Stimulating the Nao’s sensors using a metal object.* These are the plots for the experiment conducted in the Nao robot with  $\zeta = 0.5$  and  $\rho = 0.5$ , where the robot is touched with a metal object in arm position 4. For the experiment, the plots show for every time-step whether or not the robot is touched (top left), which position the hand is in (top right), which policy is preferred (bottom left), and what context the robot believes it is in (bottom right). As can be seen in the observation plot, touching the robot is often not registered and results in unreliable behaviour.

## Chapter 4

# Precision not Prediction: Body-ownership illusion as a consequence of online precision adaptation under bayesian inference

Based on:

Novický, F., Meera, A. A., Zeldenrust, F., & Lanillos, P. (2024). Precision not prediction: Body-ownership illusion as a consequence of online precision adaptation under Bayesian inference. *bioRxiv*. <https://doi.org/10.1101/2024.09.04.611162>

## Abstract

Humans can experience body ownership of new (external) body parts, for instance, via visuotactile stimulation. While there are models that capture the influence of such body illusions in body localization and recalibration, the computational mechanism that drives the experience of body ownership of external limbs is still not well understood and under discussion. Here, we describe a mathematical model of the dynamics of this phenomenon via uncertainty minimization. Using the Rubber Hand Illusion (RHI) as a proxy, we show that to properly estimate one's arm position, an agent needs to infer the least uncertain world model that explains the observed reality through online adaptation of the signals' relevance, i.e., its precision parameters (the inverse variance of the prediction error signal). Our computational model describes that the illusion is triggered when the sensory precision estimate quickly adapts to account for the increase of sensory noise during the physical stimulation of the rubber hand due to the occlusion of the real hand. This adaptation produces a change in the uncertainty of the body position estimates, yielding a switch of the perceived reality: the "rubber hand is the agent's hand" becomes the most plausible model (i.e., it has the least posterior uncertainty). Overall, our theoretical account, along with the numerical simulations provided, suggests that while the perceptual drifts in body localization may be driven by prediction error minimization, body-ownership illusions may be a consequence of estimating the signals' precision, i.e., the uncertainty associated with the prediction error.

## Keywords

Body ownership • Rubber hand illusion • Precision • Bayesian inference • Sensory adaptation

## 4.1 Introduction

Humans navigate the environment and interact with it through their bodies. This happens with a precise balance between internal representations of our body learned over our lifetime (Peviani et al., 2024) and sensorimotor-related processes, such as sensory signals arriving from the periphery caused by the localization of our body in space (Ehrsson, 2020; Tsakiris and Haggard, 2005). Remarkably, these processes are "flexible", probably to handle uncertainty and conflicts in the sensory input and to adapt to bodily changes and changes in the external world. Strong experimental evidence reveals that body perception can be easily manipulated (Ehrsson, 2012; Kilteni et al., 2015; Hide et al., 2021). Body-

ownership, i.e., “perceptual experience of body or body part as one’s own” (Ehrsson, 2020), has been associated with information integration in the ventral premotor cortex (Ehrsson et al., 2004), right posterior insula, right frontal operculum as well as posterior parietal cortex junction (Chancel et al., 2022; Gentile et al., 2013; Limanowski and Blankenburg, 2016). However, the exact computational mechanisms behind and its intricate neural underpinnings are still not well understood. By developing computational models that replicate perceptual and motor effects of humans when subjected to multisensory conflicts, such as body illusions (Ehrsson, 2022; Lanillos et al., 2021), some inner characteristics of these hidden mechanisms may be revealed (Kilteni et al., 2015; Hinz et al., 2018). In this sense, one of the most intriguing scientific questions is how we can experience the embodiment of external objects as a part of our own body.

To study this, we focused on the mechanism behind the body ownership illusion and its temporal dynamics (Finotti et al., 2023). While the research literature already provides computational accounts to explain body mislocalization as an effect of body illusions (Kilteni et al., 2015) and its dynamics (Hinz et al., 2018; Lanillos et al., 2021), there are no mathematical accounts that properly describe and explain the dynamics of body-ownership illusions. This is going beyond static Bayesian Causal Inference (BCI) (Samad et al., 2015; Chancel and Ehrsson, 2023) and showing the temporal dynamics of the mechanism behind how an external object becomes part of your body during the illusion. Here, we provide a novel mathematical account that explains the effect of body-ownership illusions as an indirect consequence of uncertainty minimization. In particular, our model describes that perceptual drifts in body localization are driven by prediction error minimization, and body-ownership illusions are a consequence of online estimation of the signals’ precision. Our model shows that the illusion is triggered when the sensory precision estimate quickly adapts to account for the increase of sensory noise during the physical stimulation of the fake body part due to the occlusion of the real body part. This adaptation produces a change in the uncertainty of the body localization estimates, yielding a change in the perceived reality: the ‘rubber hand is the agent’s hand’ becomes the most plausible model—it is the hypothesis or model that has the least posterior uncertainty. Importantly, our model aligns with previous conceptual accounts of causal inference (Samad et al., 2015; Hinz et al., 2018) of body perception and the recent findings regarding the relation between the emergence of the illusion and visual noise (Chancel et al., 2022; Chancel and Ehrsson, 2023).

To introduce the proposed model, we first summarize previous research on limb illusions and related computational models that are the backdrop of our body ownership model based on precision.

### 4.1.1 Previous research on limb illusions and related computational models

There is an extensive body of literature on limb-body illusions and manipulations—see (Kilteni et al., 2015) and (Ehrsson, 2022) for a review. Here, we briefly summarize the basic findings that point out two well-known phenomena: *i*) body-ownership of external/fake limbs and *ii*) limb mislocalization (or perceptual drift).

Evidence of the body-ownership of external/fake limbs shows that the majority of tested subjects—not all of them—perceive a fake limb as their own after a short period of a specific type of visuotactile stimulation (Botvinick and Cohen, 1998; Kilteni et al., 2015). The most well-studied body illusion, the Rubber Hand Illusion (RHI) (Botvinick and Cohen, 1998), can be seen as a phenomenon where people hold false beliefs about the causes of tactile stimulation. Specifically, the participant sees a rubber hand being stroked (e.g., with a brush) at the same time as the real hand, which is occluded. In less than one minute, the participant has the subjective experience of owning the fake hand as part of the body (Botvinick and Cohen, 1998). This illusion comes with appearance and positional constraints, and its intensity varies depending on the subject (Lanfranco et al., 2023). The illusion can be measured using subjective questionnaires. Objective measures, such as skin conductance, were introduced to evaluate the body’s reaction to the process of the fake limb embodiment (Armel and Ramachandran, 2003; Guterstam et al., 2019; Fan et al., 2021) and confirmed that external hands also produce physical responses.

The second well-known phenomenon is perceptual drift, which can be seen as a mislocalization phenomenon due to sensory integration (Tosi et al., 2023). Participants localize their own limb slightly drifted towards the fake limb (Tsakiris and Haggard, 2005; Guterstam et al., 2013). By measuring the perceived location of the hand before and after the illusion is induced, researchers have suggested that the closer participants estimate their real hand to the position of the rubber hand, the stronger the illusion is (Tosi et al., 2023). While studies used this measure to elucidate the effects of the ownership effect (Botvinick and Cohen, 1998; Costantini and Haggard, 2007), other studies highlighted that the correlation between the perceived ownership and the proprioceptive drift does not necessarily emerge (Holmes et al., 2004, 2006; Rohde et al., 2011, 2013; Abdulkarim and Ehrsson, 2016), questioning the use of perceptual drift to investigate the feeling of body ownership, and thus suggesting that perceptual drift and body-ownership illusion are a consequence of two interrelated but different processes.

Computational models of body perception should be able to account (at least) for the two well-known phenomena mentioned above. There are already candidate computational models for the mislocalization effect based on the optimal integration of the multisensory

signals (Samad et al., 2015; Hinz et al., 2018; Lanillos et al., 2021; Rood et al., 2020; Maselli et al., 2022). The rationale is that the brain optimally integrates tactile, visual and proprioceptive cues to make them fit into an internal model of the body (learned through experience). In the presence of multisensory conflicts induced by the new visuotactile input, the part of the internal model that is used to explain body localization and that is adaptive forces the system to incorrectly estimate the limb in the space. The models differ in the mathematical framework used to integrate the signals (Kilteni et al., 2015): optimal signal integration (Ernst and Banks, 2002), Bayesian Causal Inference (Samad et al., 2015), Free Energy Minimization (Hinz et al., 2018; Lanillos et al., 2021), and synaptic plasticity (e.g., spike-timing-dependent plasticity) (Zhao et al., 2023).

The computational mechanism behind the body ownership illusion is more elusive than the mechanism behind the mislocalization effect. One of the most influential proposed mechanisms, from a conceptual point of view, is that the brain, during the body illusion experiment, decides whether the information comes from the fake limb or the own limb (Samad et al., 2015; Kilteni et al., 2015). This means that the subject decides between two models of the perceived reality. This may explain why the increase of the number of presented cues (visuotactile synchronization, coherent poses, active movement) that support the fake limb model results in a stronger illusion. Thus, according to this computational model, causal inference plays a strong role in the body illusion. Still, the computational process and its dynamics behind the perceptual experience where the participant starts feeling a fake limb as its own remains unknown. Bayesian Causal Inference, as presented in (Samad et al., 2015), describes body ownership as a static probability computation, where no explanation is given on the inner functional mechanism producing the dynamics of the perceptual drifts nor the illusion (Finotti et al., 2023). Importantly, in recent experiments, the notion of uncertainty as a driving mechanism due to (visual) noise induction (Chancel et al., 2022; Chancel and Ehrsson, 2023) is gaining strength. It aligns with a hypothesis that body ownership works as a perceptual inference process, where uncertainty monitoring modulates the experience. (Lanillos et al., 2021). For a more comprehensive summary of computational models of body illusions, see the Appendix 4.5.1.

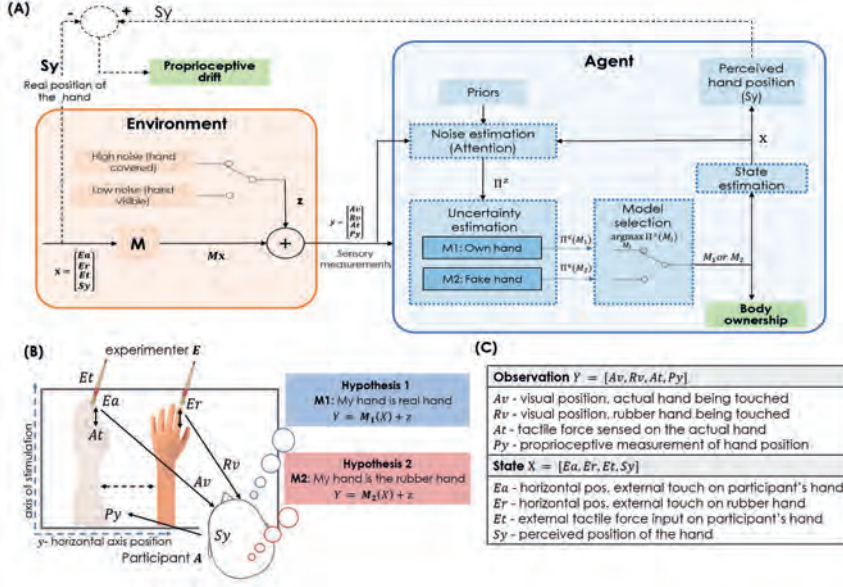
#### 4.1.2 A unified body-ownership model based on precision

The models described in the previous section vary significantly in their approaches to investigating specific phenomena observed in limb embodiment illusions. However, none have yet fully explained the mechanism and dynamics behind the ownership illusion and its relation to multisensory integration and localization drifts. We address this

gap by presenting a novel mathematical description of the mechanism that may trigger the feeling of owning a fake limb. Our model builds upon previous seminal Bayesian modeling accounts (i.e., Predictive Processing (Lanillos et al., 2021) and Bayesian Causal Inference (Samad et al., 2015)), but focuses on the role of the precision (inverse variance) adaptation during body illusions. In Bayesian models, the sensory precision parameter determines the weight of prediction errors arising from different sensory modalities. These prediction errors are, in turn, used in attention and uncertainty processing in the brain (Parr and Friston, 2017, 2019). Prior work suggests that precision optimization may be critical for resolving conflicts between competing models of sensory input (Chancel et al., 2022). However, existing computational accounts have not fully explored how dynamic precision adaptation could explain the temporal evolution of body ownership illusions. Understanding this process could reveal how the neuronal processes actively adjust the relative influence of different sensory signals to resolve uncertainty when faced with conflicting bodily information.

Our hypothesis posits that body illusions are a consequence of the brain aiming to minimize uncertainty by (online) adapting to the precision of prediction errors. Neurobiologically, this precision parameter can be related to the postsynaptic gain, which influences the encoding of prediction errors (Mumford, 1992; Feldman and Friston, 2010; Adams et al., 2013; Friston, 2023; Anil Meera et al., 2022). Interestingly, aberrant precision has been implicated in the causes of hallucinations (Brown et al., 2013; Jeganathan and Breakspear, 2021; Feeney et al., 2017; Jardri and Denève, 2013; Adams et al., 2013), and these precisions are thought to be modulated via NMDA receptors and neuromodulators such as dopamine, acetylcholine, serotonin, and noradrenaline (Adams et al., 2013; Parr and Friston, 2017; Parr et al., 2018; Novicky et al., 2023b,a). Essentially, as precision reflects the uncertainty of the prediction error, it may also function as a model of the experimental results observed when introducing visual noise in a body illusion experiment, where sensory uncertainty was found to be essential for the ownership illusion (Chancel et al., 2022).

We propose a computational model of body ownership illusion – evaluated in the RHI paradigm – that uses precision and its adaptation to sensory noise levels to explain the perceptual switch of feeling a fake hand as part of the body. The model schematic, the paradigm and used variables are described in Figure 4.1A,B and C, respectively. Similarly to the Bayesian Causal Inference approach (Samad et al., 2015), the brain decides between two different hypotheses or models  $M$  of reality. However, our approach will base its decision on each model’s posterior uncertainty. Model  $M_1$  explains the observed data as “my hand is the real hand”, while  $M_2$  explains the data as “my hand is the rubber hand”. The induced noise from covering the real hand prompts the system



**Figure 4.1: Maximum posterior precision body-ownership model.** (A) Our proposed model consists of two main blocks: the environment and the agent. The state of the environment  $x$  produces the output  $y$  that is measured by the agent as sensory measurements, along with the observation noise  $z$ . The visual noise is high when the real hand is covered, and low when the real hand is visible. The agent estimates the noise via attentional mechanisms, i.e., precision priors, which leads to estimating the uncertainty of the  $M_1$  and  $M_2$  models. Through model selection, the agent selects the best suitable model, producing the body-ownership illusion. Using the selected model, the state estimation is performed, resulting in the state estimate  $X$  that represents  $x$ . The proprioceptive drift is the difference between the perceived hand position  $Sy$  and the real position of the hand  $Sy$  (that the agent does not have access to). Note that all variables in **bold** belong to the generative process (such as  $Sy$ ), while non-bold variables belong to the generative model (such as  $Sy$ ). (B) Experiment formalization. The brain is testing two competing models (i.e., hypotheses):  $M_1$  stands for believing that “my hand is the real hand”, while  $M_2$  indicates that “my hand is the rubber hand”. These two internal models maps (as a generative model) the states  $X$  into the observations  $Y$ . (C) Notation used for the observations (sensory measurements) and the states.

to adapt. This adaptation makes the posterior estimates more uncertain. Consequently, hypothesis/model selection switches to the most probable one, which is to own the fake limb, thus producing the illusion.

Our computational model suggests that: *i*) body mislocalizations (e.g., proprioceptive drift) result from body state estimation, where prediction errors are minimized, and *ii*) body ownership illusions, i.e., perception of a fake limb/hand as one’s own, result from uncertainty minimization, where precision adaptation to overcome induced noise triggers a model selection switch.



## 4.2 Methods

We introduce a mathematical model for the body-ownership illusion, which is drawn from the Bayesian brain hypothesis (Doya et al., 2007) and grounded on the free energy principle (FEP) (Friston et al., 2017). This principle proposes that the brain generates cognition and behaviour following a Bayesian optimization strategy. It postulates that living beings show resilience over time, which is achieved via minimizing entropy (the second law of thermodynamics) to preserve their equilibrium state (i.e., to self-organize) (Friston, 2010; Parr et al., 2022). To do this, living systems embed a generative internal model that explains the world (i.e., perception) and adjusts to the environment (by changing their belief or exerting actions) via minimizing the free energy functional (an upper bound on sensory surprisal) (Friston, 2009). We build on this theory to describe body illusions being driven by uncertainty minimization, i.e., the aim to obtain a maximum posterior precision or minimum posterior uncertainty. We first present the model description and notation and then detail each computational sub-process.

### 4.2.1 Problem formulation and parameters

Figure (4.1A) shows the model schematic for the body-ownership illusion based on posterior precision. It consists of the environment and the agent. The environment comprises an experimenter (E) that provides the participant (or the agent) with sensory stimuli—see Fig. (4.1C) for the notation used. The experimenter strokes the real hand (which is at position  $\mathbf{Ea}$ ) and/or the rubber hand (which is at position  $\mathbf{Er}$ ) locations. Thereby, the experimenter generates a force ( $\mathbf{Et}$ ) on the real hand, which the agent (A) can feel. Moreover, the agent can collect the visual signals  $\mathbf{Av}$  and  $\mathbf{Rv}$  that represent the position of the real hand and the rubber hand, respectively. The agent also feels a tactile sensation  $\mathbf{At}$  on its real hand. The  $y$  coordinate position of the real hand  $\mathbf{Sy}$  is observed by the agent as  $\mathbf{Py}$  (proprioceptive measurement). Stacking these variables together we have: i) the environment state  $\mathbf{X} = [\mathbf{Ea} \ \mathbf{Er} \ \mathbf{Et} \ \mathbf{Sy}]^T$  and ii) the measurements of the agent as  $\mathbf{Y} = [\mathbf{Av} \ \mathbf{Rv} \ \mathbf{At} \ \mathbf{Py}]^T$ .

We assume that the environment is a generative process<sup>1</sup> that produces the sensory measurements from the state variables. We define this mapping as a linear function  $\mathbf{M}$  between the state  $\mathbf{X}$  and output (measurements)  $\mathbf{Y}$  with measurement noise  $\mathbf{z}$  as:

---

<sup>1</sup>In the FEP literature, the generative process refers to the real world process that the agent has only access through the senses, and the generative model refers to the agent model approximation of the reality learnt by experience.

$$\mathbf{Y} = \mathbf{M}\mathbf{X} + \mathbf{z} \rightarrow \begin{bmatrix} \mathbf{A}\mathbf{v} \\ \mathbf{R}\mathbf{v} \\ \mathbf{A}\mathbf{t} \\ \mathbf{P}\mathbf{y} \end{bmatrix} = \mathbf{M} \begin{bmatrix} \mathbf{E}\mathbf{a} \\ \mathbf{E}\mathbf{r} \\ \mathbf{E}\mathbf{t} \\ \mathbf{S}\mathbf{y} \end{bmatrix} + \mathbf{z} \quad (4.1)$$

The agent processes the measurement  $\mathbf{Y}$  to make sense of the world through estimation. However, the experimenter can manipulate  $\mathbf{Y}$  to make the agent fall into the body illusion. We model this process within the RHI paradigm using the following components, as shown in Figure 4.1A: i) noise estimation (attention), ii) state estimation (body localization), iii) posterior uncertainty estimation, and iv) model selection (body ownership). The next section will explain each of these computational blocks in detail.

## 4.2.2 Maximum posterior precision body-ownership model

In our proposed approach, the agent selects the model that best minimises its uncertainty in estimation. This implies that the agent estimates, at each time step, both the state and the noise in the environment given the competing models or hypotheses of the perceived reality. For instance, in the RHI, the two models are “my hand is my real hand” and “my hand is the rubber hand”. In the following subsections, all subprocesses of the agent, depicted in Figure 4.1A, are detailed.

### 4.2.2.1 Competing models of perceived reality

We assume that the brain can maintain different hypotheses of the perceived reality, where each of them defines how sensory information is generated. In the classical RHI, this accounts for two hypotheses or competing models of reality ( $M_1$  and  $M_2$ ). However, this approach can be extended to any number of hypotheses.  $M_1$  accounts for accurate body estimation neglecting the rubber hand, and  $M_2$  considers that the rubber hand is our hand. We further model the participant’s sensitivity to experience the illusion by means of how both models are integrated. We define body ownership existing on a continuous spectrum parameterized by  $\alpha$ . This spectrum spans between two extreme cases: a strongly non-illusory model ( $\alpha = 10$ ) where the agent maintains accurate body ownership and a strongly illusory model ( $\alpha = -10$ ) where the agent is susceptible to incorporating external objects into their body representation. Mathematically, this is defined as the weighted combination of both non-illusory and illusory models:

$$S_M = \frac{1}{1 + e^{-\alpha}} M_1 + \frac{e^{-\alpha}}{1 + e^{-\alpha}} M_2 \quad (4.2)$$

where  $M_1$  represents the mapping for accurate body perception:

$$Y = M_1 X + z \rightarrow \begin{bmatrix} Av \\ Rv \\ At \\ Py \end{bmatrix} = \begin{bmatrix} 0.8 & 0 & 0.2 & 0 \\ 0 & 1 & 0 & 0 \\ -0.3 & 0 & 0.7 & 0 \\ Ay - Ry & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} Ea \\ Er \\ Et \\ Sy \end{bmatrix} + z, \quad (4.3)$$

and  $M_2$  represents how sensory information will be generated if the rubber hand is our hand:

$$Y = M_2 X + z \rightarrow \begin{bmatrix} Av \\ Rv \\ At \\ Py \end{bmatrix} = \begin{bmatrix} 0.8 & 0 & 0.2 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & -0.3 & 0.7 & 0 \\ 0 & Ay - Ry & 0 & 1 \end{bmatrix} \begin{bmatrix} Ea \\ Er \\ Et \\ Sy \end{bmatrix} + z. \quad (4.4)$$

---

**Algorithm 1** Model selection by uncertainty minimization

---

**Require:** Models  $M = \{M_1, M_2\}$

**for**  $t = 0 : dt : T$  **do**

$\mathbf{Y} = \text{DataGeneration}()$

$\triangleright$  Environment generates the data

**for**  $M_i \in M$  **do**

$\triangleright$  Iterate through possible models

$\Delta\lambda = \text{NoiseEstimation}(\mathbf{Y}, \lambda, M_i)$

$\triangleright$  Compute  $\Delta\lambda$  by minimizing  $F$

$\lambda_{t+1}(M_i) = \lambda_t(M_i) + \Delta\lambda$

$\triangleright$  Update the noise parameter

$\Delta x = \text{StateEstimation}(\mathbf{Y}, x, M_i)$

$\triangleright$  Compute  $\Delta x$  by minimizing  $F$

$x_{t+1}(M_i) = x_t(M_i) + \Delta x$

$\triangleright$  Update the state estimation

$\Pi^\lambda(M_i) = \text{EvaluatePrecision}()$

$\triangleright$  Precision on  $\lambda$

$\Pi^x(M_i) = \text{EvaluatePrecision}()$

$\triangleright$  Precision on  $X$

$\Sigma^x(M_i) = (\Pi^x(M_i))^{-1}$

$\triangleright$  Compute uncertainty in state estimation

**end for**

$M^* = \arg \min_{M_i} \text{trace}(\sqrt{\Sigma^x(M_i)})$

$\triangleright$  Model selection

$x_{t+1} = x_{t+1}(M^*)$

$\triangleright$  Select the state

$\lambda_{t+1} = \lambda_{t+1}(M^*)$

$\triangleright$  Select the noise

**end for**

---

The key differences between these mappings (or generative models) lie in the processing of tactile information ( $At$ ): while in both models tactile sensation depends on the proprioceptive estimation ( $Py$ ), in  $M_1$ , tactile sensation receives input from the real hand stimulation ( $Ea$ ), while in  $M_2$  it depends on the rubber hand stimulation ( $Er$ ).  $M_1$  maintains an accurate position estimation factoring in hand separation (parametrized as the distance between the hand position  $Ay$  and the rubber hand position  $Ry$ ), while  $M_2$  allows information integration using the rubber hand position. Thus, both models use the hand separation parameter ( $Ay - Ry$ ), but they apply it differently -  $M_1$  applies it to the

real hand position while  $M_2$  applies it to the rubber hand position. This mathematical difference captures the essence of the illusion, where the proprioceptive signals become linked to the rubber hand rather than the real hand. Note that all variables highlighted in **bold text** belong to the generative process (environment), and the non-bold variables belong to the generative model (agent).

The inclusion of tactile force ( $Et$ ) in our model represents the continuous nature of the stroking stimulation of the experimenter during the RHI. While  $Ea$  and  $Er$  represent the spatial locations where stroking occurs on the real and rubber hands, respectively,  $Et$  captures the temporal dynamics of the applied force. Having visual and tactile signals allows us to model both the spatial and temporal aspects of visuotactile integration. The tactile force parameter enables our model to capture how the strength and timing of tactile stimulation could influence the illusion. Although this implementation is a simplification of the complex spatiotemporal patterns involved in actual RHI experiments, it provides a tractable way to model continuous sensory stimulation.

The  $\alpha$  parameter of Eq. 4.2 determines the relative weighting between the two running hypotheses, with more negative values of  $\alpha$  leading to stronger illusory effects and more positive values maintaining accurate body representation. This formulation allows us to capture both the gradual onset of the illusion and individual variations in susceptibility while maintaining the key computational features that enable both accurate body representation and illusory perception when appropriate. A participant with  $\alpha = 10$  will consider the information from the experimenter touching the real hand and also the force applied to the real hand as the most important, compared to the other extreme participant ( $\alpha = -10$ ), which values the rubber hand as the main source of these stimuli, slightly ignoring the force  $Et$ .

#### 4.2.2.2 Model selection by uncertainty minimization

To determine the most plausible hypothesis, our agent selects the one with higher posterior precision. First, we describe the overall mechanism of model selection<sup>2</sup> that produces the body ownership illusion—depicted as a pseudocode in algorithm 1—and then we detail the different computations, such as noise precision estimation, body state estimation and uncertainty estimation. The code and the parameters for replication of the results and figures can be found in Appendix 4.5.3.

For all competing models  $M_i$ , we compute the posterior precision of states  $\Pi^X(M_i)$ , along with the estimated state ( $X$ ). We define the model selection criterion, using the posterior precision of state estimation, for the competing models as:

---

<sup>2</sup>See Discussion section for the reason of using model selection instead of model averaging, which is usually followed in computational neuroscience.

$$M = \operatorname{argmin}_{M_i} \operatorname{trace}[(\Pi^X M_i)^{-\frac{1}{2}}]. \quad (4.5)$$

The core idea is that the agent selects the model that minimises the uncertainty in state estimation, mathematically described as the trace (sum of diagonal elements) of the square root of the posterior covariance matrix (inverse of precision matrix,  $\Sigma^X = (\Pi^X)^{-1}$ ) of state estimation. According to the proposed model, the RHI will begin at the moment where  $M_2$  has lower uncertainty, compared to  $M_1$ . Once the RHI starts, the distinction between the actual hand and the rubber hand is hypothesised to begin to fade from the agent's point of view, resulting in a flow of information from external inputs on the rubber hand ( $Er$ ) to the sensory perception about the actual hand ( $At$  and  $Py$ ). This can be observed by comparing the differences in the last two rows of equations (4.3) and (4.4): As the agent switches from model  $M_1$  to model  $M_2$  the RHI begins.

**Free parameters:** The core free parameter of the model is the estimate of the noise parameter  $\lambda$ , which is learned dynamically through precision adaptation. This parameter captures the key phenomena. Our model has other parameters: *i*) The prior noise precision parameter  $P^\lambda$  for each sensory input; *ii*) The model matrices  $M_1$  and  $M_2$  that define how states map to observations; *iii*) The prior mean and precision ( $\eta^\lambda$  and  $P^\lambda$ ) for  $\lambda$  that helps to embed prior information for precision adaptation; and *iv*) The learning rates  $k$  for state and precision estimation.

#### 4.2.2.3 Role of precision in the computational model

In our computational model, we consider three sources of variance, i.e. three precision parameters that need to be estimated: *i*) noise precision  $\Pi^z$ , that is parameterized using  $\lambda$  (details in appendix 4.5.2), *ii*) prior precision on estimates  $P^\lambda$ , and *iii*) posterior precision on estimates  $\Pi^X$  and  $\Pi^\lambda$ . While the noise precision ( $\Pi^z$ ) is learned online by our model using the agent's prior precision ( $P^\lambda$ ), the posterior precision ( $\Pi^X$ ) is computed alongside the state estimates  $X$ . The noise precision learning contributes to the identification of the noise levels in the sensory data, while the posterior precision computation contributes to the decision (model selection) of which model to rely on ( $M_1$  or  $M_2$ ). The following sections will detail the mathematical formulations behind these processes.

#### 4.2.2.4 Joint State and Noise Precision Estimation

We model body state estimation as inferring the internal states of the body from the observations  $Y$ , given a model of reality. The agent keeps track of the states  $X$ , the activities of the experimenter  $Ea$ ,  $Er$ ,  $Et$  and the perceived hand position  $Sy$ . This estimate is

computed in the agent by approximate Bayesian inference through the minimization of the free energy (Meera and Wisse, 2020). The free energy of the agent that tries to predict the outputs  $\mathbf{Y}$ , under state  $X$  and model  $M$ , and learns (or estimates) the measurement noise parameter  $\lambda$ , with prior mean and precision  $\eta^\lambda$  and  $P^\lambda$  respectively, can be written down as (see (Meera and Lanillos, 2023) for full derivation):

$$F = \frac{1}{2}(\lambda - \eta^\lambda)^T P^\lambda (\lambda - \eta^\lambda) + \frac{1}{2}(\mathbf{Y} - MX)^T \Pi^z (\mathbf{Y} - MX) - \log |\Pi^z|, \quad (4.6)$$

where the operation  $^T$  denoted the transpose. Here, the noise precision  $\Pi^z$  is parametrized using  $\lambda$  (see appendix 4.5.2). The agent performs state estimation under a selected model using free energy minimization:

$$X(t + \Delta t) = X(t) + \left( e^{-k \frac{\partial^2 F}{\partial X^2} \Delta t} - I \right) \left( \frac{\partial^2 F}{\partial X^2} \right)^{-1} \frac{\partial F}{\partial X}, \quad (4.7)$$

where  $I$  denotes the identity matrix. The update rule uses the first two gradients of  $F$  in Eq 4.7, which can be calculated by differentiating Eq (4.6) with respect to  $X$ . This mathematical routine updates  $X$  such that it minimises  $F$  in Eq 4.6 with respect to  $X$ . The agent performs online noise precision learning by updating  $\lambda$  such that it minimises the free energy, using its first two free energy gradients:

$$\lambda(t + \Delta t) = \lambda(t) + \left( e^{-k \frac{\partial^2 F}{\partial \lambda^2} \Delta t} - I \right) \left( \frac{\partial^2 F}{\partial \lambda^2} \right)^{-1} \frac{\partial F}{\partial \lambda}. \quad (4.8)$$

The gradients of  $F$  can be calculated by differentiating Eq (4.6) with respect to  $\lambda$ . Note that the update of  $\lambda$  changes the learned observation noise  $\Pi^z$  as per Eq (4.10) in Appendix. The updated parameter  $\lambda$  is used to compute  $\Pi^z$  in Eq (4.6). This is further used in the state estimation in Eq (4.7) in the next time step. Therefore, the noise parameter  $\lambda$  directly influences the state estimation. In this particular setting when the noise has a large amplitude, the agent can perceive illusory body state estimation.

#### 4.2.2.5 Uncertainty in estimation

According to FEP, the posterior precision of the estimate (the second-order statistic reflecting the confidence about the estimates (Meera and Lanillos, 2024)) is obtained using the second-order gradient of  $F$ . We use this to evaluate the uncertainty about the estimates by taking the inverse of the posterior precision of the state estimates as (Meera and Lanillos, 2023):

$$\Sigma^X = (\Pi^X)^{-1} = \left( \frac{\partial^2 F}{\partial X^2} \right)^{-1}, \quad (4.9)$$

where  $\Sigma^X$  denotes the agent’s uncertainty in making predictions about the states, given the observed data. The agent computes the precision for both the competing models  $\Pi^X(M_1)$  and  $\Pi^X(M_2)$  to decide. It is used directly in the model selection criterion given in Eq (4.5). The core idea is that the agent decides to believe in a model that best minimizes (or resolves) its uncertainty in estimation. The hypothesis is that once the visual signals of the hand ( $Av$ ) are highly noisy (hand covered), the illusory condition ( $\text{trace}(\sqrt{\Sigma^X(M_1)}) > \text{trace}(\sqrt{\Sigma^X(M_2)})$ ) becomes true and the agent switches its model selection from a non-illusory model ( $M_1$ ) to an illusory model ( $M_2$ ), resulting in the experience of RHI.

### 4.2.3 Evaluative measures for the model

Now that it has been explained how the agent uses the uncertainty in state estimation for model selection, we conclude with a description of how to model hand coverage, the illusion sensitivity and measure proprioceptive drift. Note that these measures are not modeled as part of the agent in Figure 4.1, but are used as fitted parameters for evaluations and comparisons.

#### 4.2.3.1 Modeling hand coverage

We model the real hand coverage by making the visual signal very noisy. The measurement noise  $\mathbf{z}$  is composed of four components ( $\mathbf{z}^{Av}, \mathbf{z}^{Rv}, \mathbf{z}^{At}, \mathbf{z}^{Py}$ ). When the real hand is covered, the visual signal  $\mathbf{A}v$  obstruction is modeled by adding a high noise  $\mathbf{z}^{Av}$ . Analogously, for no coverage of the real hand, we use a low noise  $\mathbf{z}^{Av}$  on the signal. In broad terms, the noise level  $\mathbf{z}^{Av}$  modulates the error in the visual information input. This modeling, although controversial, is grounded on the idea that darkness is completely uncertain (Friston et al., 2012) and hence, that the absence of visual information invokes a strong noise component—see discussion (Sec. 4.4.3) for further explanation.

#### 4.2.3.2 Individual differences through illusion sensitivity

There is a high variability between individual susceptibility to the body illusion. Moreover, depending on the manipulation (e.g., temporal synchrony), the body illusion sensitivity varies (Lanfranco et al., 2023). To account for the agent’s sensitivity to experience the illusion, we modulate the  $\alpha$  parameter from Eq 4.2. Technically, reliance on prior versus posterior evidence of each participant is modeled through parameter  $\alpha$ . Intuitively, the higher the  $\alpha$ , the less susceptible the agent is to experience the illusion. Thus, Eq (4.2) captures the variability of body ownership sensitivity, which also influences the experienced proprioceptive drift.

### 4.2.3.3 Measuring the proprioceptive drift

During the RHI experiment, the agent’s state estimate of its hand position  $Sy$  may be disrupted, due to the high noise amplitude in the visual data ( $\mathbf{z}^{\mathbf{A}^v}$ ), leading to the experience of a proprioceptive drift, where the perceived hand position  $Sy$  shifts from the real hand position  $Ay$ , towards the rubber hand position  $Ry$ . The agent cannot report the proprioceptive drift during the RHI because it is an illusion, and the agent does not have access to the true signal  $\mathbf{S}y$ . The proprioceptive drift in this paper is computed by taking the difference between the real position of the hand  $\mathbf{S}y$  and the agent’s estimated hand position  $Sy$  as shown in Figure 4.1.

## 4.3 Results

Here, we show using simulations how the proposed precision-based model may explain the mechanism and dynamics behind body ownership illusions and their relation to multisensory integration and localization drifts. To this end, we simulated the RHI paradigm, being one of the most paradigmatic body perceptual manipulations, where a participant experiences the illusion that an isolated fake hand becomes an agent’s hand. The parameters used to generate the results are described in Appendix 4.5.3. We describe the following results:

1. We validate the proposed model in a simulated RHI experiment (Fig. 4.2), where we show how posterior uncertainty minimization under the two competing models of reality yields to a perceptual switch that produces the illusion of owning the fake hand.
2. We analyse the precision adaptation mechanism of the proposed model, which is underneath the perceptual switch (Fig. 4.3). The introduction of visual noise by covering the hand is one of the factors producing the effect. The online adaptation to the sensory noise through precision estimation appears as an essential property to experience the illusion.
3. To further validate our model, we show that breaking the noise precision learning – where the model does not adapt to the sensory noise – leads to a state where switching the current model of reality to a different one is implausible, hence, no illusion can occur (Fig. 4.4).
4. We assess the dependence of the body ownership illusion (model switch) on the distance between the rubber hand and the participant’s own hand (Fig. 4.5)
5. We evaluate how the proprioceptive drift depends on the selected competing model



(Fig. 4.6), where  $M_1$  produces noisy estimations but not drift, while when selecting  $M_2$ , the model produces the experimentally observed proprioceptive drift.


6. We validate the effect of the inter-hand distance on the proprioceptive drift dynamics, given that the illusion is happening (Fig. 4.7).
7. Finally, using the modeling of artificial participants with individual differences, we explore the effect of the participant's susceptibility to experience the RHI ( $\alpha$ ) and the inter-hand distance ( $Ry$ ) on the generated proprioceptive drift (Fig. 4.8). The results show a similarity between the profiles found in a recent experiment in humans and macaques (Fang et al., 2019) and our model.

### 4.3.1 Body-ownership illusion: model selection based on the uncertainty minimization

Figure 4.2 describes the proposed model behaviour, as described in algorithm (1), as a function of uncertainty estimation when the variance of the sensory input increases (Fig. 4.2). First, the posterior uncertainty ( $(\Pi^X)^{-1}$ )—when both hands are visible and being stroked by the researcher (Figure (4.2, panel 1)—favors the non-illusory hypothesis ( $M_1$ ). Once the hand becomes covered (Figure 4.2 panel 2), there is an immediate increase in  $M_1$  uncertainty, which is due to the increase of the noise estimate of the real hand ( $\lambda$ ). After a few moments, there is a switch between the two models regarding posterior uncertainty (Figure 4.2 panel 3). This leads to the participant believing that they own the rubber hand, as the  $M_2$  hypothesis now provides a lower uncertainty compared to  $M_1$ . As this is happening, we can see that in model  $M_2$ , there is also an observable proprioceptive drift driving away from  $P_y=0$  towards the location of  $Ry=0.8$  (red line, bottom panel). Therefore, this simulation shows the two main effects of body-ownership illusions and the potential computational mechanisms responsible for body ownership (minimizing uncertainty) and proprioceptive drift or mislocalization (state estimation as multisensory integration).

### 4.3.2 Estimating noise through precision adaptation

We evaluated the reliability of our model when estimating the sensory noise. Figure (4.3) shows that both proprioceptive (purple line) and visual (blue line) signals from the real hand become noisy after the hand is covered at the 15th second. This is because covering the real hand results in noisier visual and proprioceptive signals. The coloured lines in Figure 4.3B follow the real noise levels (the dashed lines) in all signals, showing that model is capable of correctly estimating the precision levels of all sensory signals from

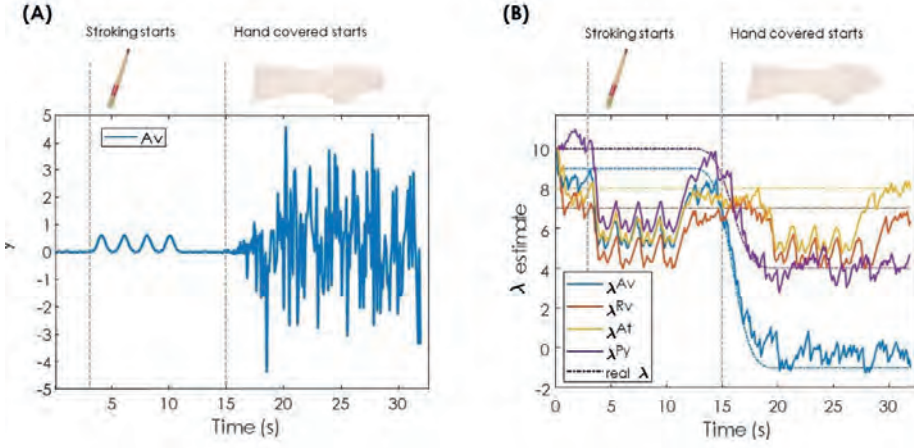


**Figure 4.2: Body-ownership perceptual switch.** The experiment sequence and its results in our model: 1) Both hands are being simultaneously stroked; 2) The real hand is covered and the strokes continue; 3) The participant starts experiencing the RHI. In the bottom graph, the competition of the two models and their uncertainty during the experiment are shown. The middle panel shows the estimated posterior state uncertainty  $((\Pi^X)^{-1})$  of both models and the bottom panel the perceived position of the hand ( $Sy$ ) given the participant believes in model  $M_1$  (blue) or in model  $M_2$  (red). After the hand is covered, there is a switch between these two models, as  $M_2$  starts being more precise (lower uncertainty). This then leads to believing that the rubber hand is the participant’s hand, thus leading to a proprioceptive drift depicted in the bottom graph. Note that the blue line in this panel represents the real proprioceptive signal  $Sy$ , whereas the red line represents the estimated  $Sy$  under the assumption of  $M_2$ . In this simulation, we used rubber hand position  $Ry = 0.8$ .

data.

### 4.3.3 Validating the model via breaking precision adaptation

We hypothesised that precision learning/adaptation is a necessary precondition for model switching. In other words precision learning is crucial for the model switching, and if this is broken (by providing priors that inhibit learning), the switch will not occur. To test this, we disrupted the noise precision learning by biasing it with a high prior precision on noise. We defined  $(P^\lambda = \text{diag}([e^{10}, e^{-4}, e^{-4}, e^{-4}]))$  for the signal  $Av$ . In Figure (4.4),




**Figure 4.3: Precision adaptation (attention).** (A) Visualization of visual signal from the real hand  $Av$  signal in blue. It is highly noisy after the real hand is covered (at time = 15 seconds). This represents the highly uncertain nature of the visual data for the hand position once it is covered. (B) Noises estimate block from our model in Fig. 4.1, given the sensory measurements. The colored plots track the true noise parameter  $\lambda$  (in dashed coloured lines), showing that our model is capable of correctly estimating the precision levels of all sensory signals. In other words, the uncertainty of the sensory signals is fully captured by our model in real time. Simulation parameters can be found in Appendix 4.5.3.

the effects of this manipulation are shown. Figure (4.4A) shows the improperly learned precision signals of the input  $Av$  due to the high prior on  $\lambda$ . Consequently, Figure (4.4B) demonstrates that the posterior state uncertainty  $(\Pi^X)^{-1}$  of  $M_1$  remains consistently lower than  $M_2$ , regardless of hand coverage in our earlier findings in Figure (4.2). These results provide face-validation for the necessity of precision learning/adaptation in achieving the body-ownership illusion under our proposed model.

#### 4.3.4 Dependence of the illusion on inter-hand distance

We further examined how the illusory perceptual switch depends on the distance between the real and the rubber hand ( $|Sy - Ry|$ ). Empirically, the RHI breaks down when the fake hand is too far from the real one, as the brain struggles to integrate highly discrepant multisensory information. Our model captures this effect naturally (Figure 4.5). Because  $M_1$  has a lower uncertainty with an increased inter-hand distance, the illusory state cannot emerge. Figure 4.5A shows that the non-illusory model  $M_1$  is always more certain about the states for the time frames when the is not covered, independent of the distance  $Ry$ .

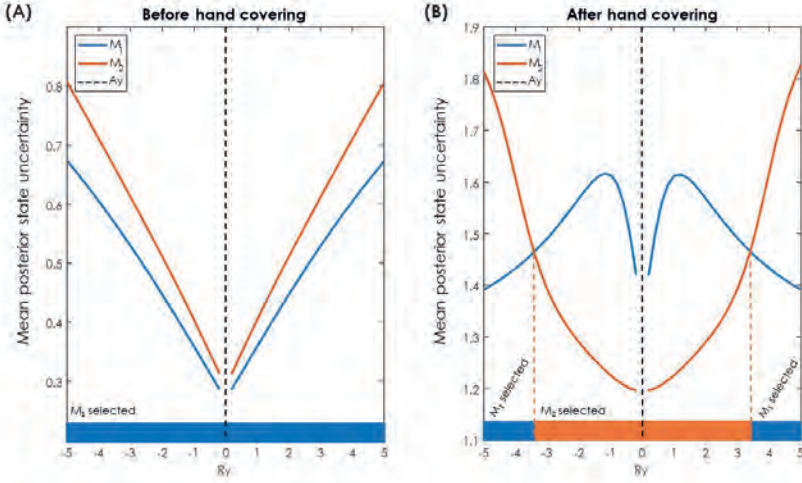


**Figure 4.4: Breaking precision adaptation leads to no illusion.** (A) This figure is similar to Figure 4.3B, except for the fact that the agent is now biased towards believing that the signal  $A_v$  is highly precise by using a higher prior precision ( $P^\lambda = \text{diag}([e^{10}, e^{-4}, e^{-4}, e^{-4}])$ ) for the signal  $A_v$ ), although it is as noisy as in Figure 4.3A. Note that  $\lambda^{Av}$  (in blue) stays high, around 10 instead of tracking the real  $\lambda$  (in black). (B) If an agent is (incorrectly) overconfident over the visual signal  $A_v$ , it will lead to a behaviour where the two competing models will not switch—the red curve always stays above the blue one (compare with the switching in Figure 4.2). Simulation parameters can be found in Appendix 4.5.3.

Figure 4.5B shows the model’s posterior state uncertainty for the time frame when the hand is covered. When the rubber hand is too far from real hand (when  $|Ry| > 3$ ), the non-illusory model becomes more certain (blue curve below the red curve), implying that the RHI will not happen when the inter-hand distance is too high. Notice that Figure 4.5A and B are completely symmetrical, highlighting that the model behaves in the same manner with negative and positive values.

### 4.3.5 Proprioceptive drift in competing models


We have shown that our model can explain the RHI as a result of adaptation to the precision of sensory stimuli. In this section, we will now show how our model explains proprioceptive drift in more detail than in Figure 4.2, where it is mentioned only briefly as an independent mechanism that happens alongside model selection. Here, we show that proprioceptive drift occurs only in the illusory model  $M_2$  (Figure 4.6B) but not in the non-illusory model  $M_1$  (Figure 4.6A). These plots represent a single trial of the experiment described in Figure 4.2. In Figure 4.6, the red lines that represent the estimated perception of the signal coincide with the blue lines that represent the real values of the experimenter’s stimulation. Therefore, the agent correctly estimates the location of the touch of the real hand ( $Ea$ ) (Figure 4.6 first panel), of the fake hand



**Figure 4.5: Model switch is dependent on the inter-hand distance.** The vertical axis of panel A represents the mean posterior state uncertainty of all 4 signals ( $Ea$ ,  $Er$ ,  $Et$ ,  $Sy$ ) averaged when hand is uncovered (from the 4th to the 11th second), while in panel B, the average when hand is covered is from 20th second to the 30th second. These averages are then plotted for both  $M_1$  and  $M_2$  for different inter-hand distances. The horizontal axis represents different positions of the rubber hand ( $Ry = -5 : 0.2 : 5$ ). Notice that the value of  $Ry = Ay = 0$  was removed, as it would mean that the fake hand is literally lying on the participant’s real hand. The colors on the bottom indicate whether the model remains at  $M_1$  or switches to  $M_2$  under a given  $Ry$ . Simulation parameters can be found in Appendix 4.5.3.

( $Er$ ) (Figure 4.6 second panel), and the tactile sensation ( $Et$ ) (Figure 4.6 third panel). However, as shown in Figure 4.6B, even during the first half of the experiment, when both hands are visible, the  $M_2$  model produces a small deviation in the estimated hand position  $Sy$  from its true value (Figure 4.6B in the bottom panel). Here, the red curve deviates from the blue line towards the dashed horizontal line ( $Ry$ ). This plot highlights that the model is able to generate the proprioceptive drift.

The proprioceptive drift also occurs before the hand is covered, though participants would not experience this drift since the non-illusory model  $M_1$  has a lower posterior state uncertainty compared to  $M_2$  (see Figure 4.6A last panel for comparison, and Figure 4.2 for model selection). This demonstrates that the proprioceptive drift exists within the mathematical structure of the  $M_2$  model even before it becomes the selected model. Moreover, we show how the process of model selection is independent from the mechanism of drift computation. Only when  $M_2$  becomes the selected model after hand occlusion does this drift become part of the agent’s experienced hand position.



**Figure 4.6: Replicating the proprioceptive drift.** (A) State estimates of all the variables under the non-illusory model  $M_1$ . (B) The same state estimates when the illusory model  $M_2$  has been chosen. Both figures show state estimates of all input data  $Ea$ ,  $Er$ ,  $Et$ , and  $Sy$ . Only in the illusory model on the right, the proprioceptive drift ( $Sy$ ) and the related touch ( $Et$ ) are observed. Also note the increase of noise in  $Ea$  in under both models when the hand is covered. The distance between the rubber hand and real hand is 0.8 in this case. The distance dependence on the proprioceptive drift is evaluated in Figure 4.8.

#### 4.3.5.1 Dependence of proprioceptive drift on inter-hand distance

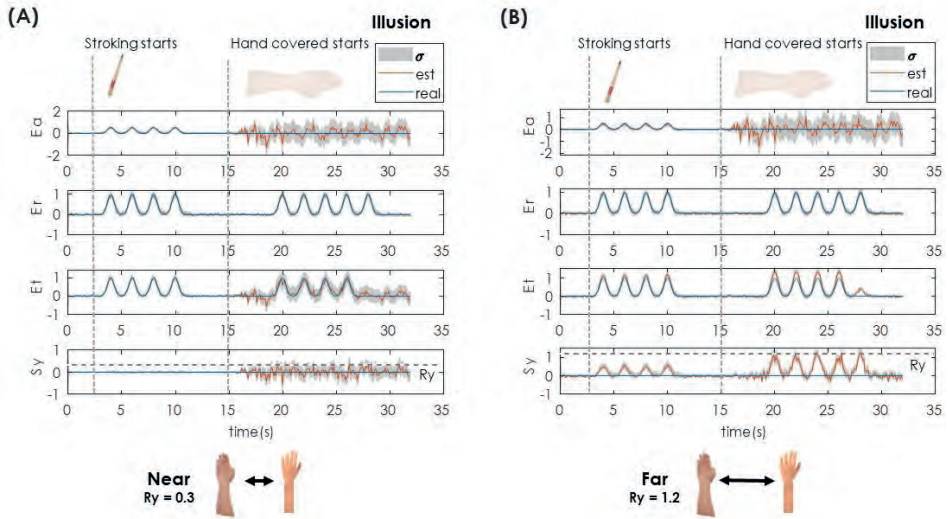
The dependence of the proprioceptive drift on the inter-hand distance has been observed empirically in many experiments (Preston, 2013; Kalckert et al., 2019). Here, we also validate our model with this logic. In Figure (4.7), we compare two illusory models where the distance of the two (real and rubber) hands is increased. In Figure (4.7A), the inter-hand distance is smaller ( $Ry = 0.3$ ) compared to the one in Figure (4.7B, ( $Ry=1.2$ )). In the latter figure, the estimate in  $Er$  is incorrect, which leads to a smaller and shorter proprioceptive drift (measured as the difference between the real and estimated values of  $Ry$ , bottom panels) compared to Figure (4.7A).

Next to the inter-hand distance, we also we modulated the strength of both the illusory and non-illusory models via the parameter  $\alpha$  in the softmax function in Eq (4.2) – see methods for more detail. By doing so, we assume that we modulate the subjective susceptibility of the competing models. We vary these two parameters, the inter-hand distance  $Ry$  and the susceptibility  $\alpha$ , independently, to show that the proprioceptive drift is increasing in amplitude when slowly increasing the distance between the two hands. The speed of the amplitude increase becomes smaller with higher distance (Figure




(4.8)). In this figure, a higher  $\alpha$  value is associated with the agent being less susceptible to the illusion and vice versa. We plot the maximum value of the proprioceptive drift for simplification, which is represented on the vertical axis, while the distance between the subject's and the rubber hand is on the horizontal axis. In addition, we show that modulating the subjective strength of the illusory model results in different strengths of the proprioceptive drift, as shown in (Fang et al., 2019).

In the introduction, we touched upon the discussion of whether proprioceptive drift is a good indicator of the RHI. The results, from our perspective, also link the difference between the selection of the illusory model and feeling the drift, and not selecting the illusory model and still feeling the drift, ultimately supporting the case of two distinct computations for proprioceptive drift and illusion per se. Here, we therefore argue that these two are the result of separate but related computational mechanisms (Holmes et al., 2004, 2006; Rohde et al., 2011, 2013) – state estimation and precision maximisation.



**Figure 4.7: Inter-hand distance influence on the preceptual drift.** The same template as in Figure 4.6 was used. Here, however, both figures represent the state estimates under the illusory model, where we varied the interhand distance. **(A)** Near ( $R_y = 0.3$ ). Effect of  $M_2$  when the inter-hand distance is small ( $R_y = 0.3$ ). **(B)** Far ( $R_y = 1.2$ ). Effect of  $M_2$  when inter-hand distance is increased. This shows that with the increase of the distance between the hands, the proprioceptive drift becomes stronger (or increases in amplitude).



**Figure 4.8: Dependence of the inter-hand distance and the participant’s susceptibility to experience the illusions on the proprioceptive drift.** Inspired by (Fang et al., 2019), we show a similar trend from previous empirical findings where the modulation of the inter-hand distance and different subjective experience to the illusion affects the proprioceptive drift. Every line in the plot can be related to individual differences when transiting from  $M_1$  non-illusory to  $M_2$  illusory models (i.e., susceptibility to experience the illusion). This is modulated through the parameter  $\alpha$ . Here,  $\alpha = 10$  and  $\alpha = -10$  are equal to models from equations (4.3) and (4.4), respectively. We measured the proprioceptive drift as the average of  $S_y$  between the 17th and the 30th second. Note that the dashed lines represent the distances for near and far conditions from Figure (4.7).

## 4.4 Discussion

We have proposed a novel theoretical model that can explain body illusion dynamics and evaluated it through a simulated RHI experiment. We built a general model that selects between competing interpretations of body representations and switches between them by maximizing the precision (minimizing the uncertainty) of the state estimation (Fig. 4.2). The model was face-validated on previous experimental results via replicating the proprioceptive drift (Fig. 4.6), its dependence on the inter-hand distance (Fig. 4.7), and the illusion dependency on the inter-hand distance (Fig. 4.5). We showed that the model can reproduce the observed differences in proprioceptive drift among participants due to the subjective susceptibility towards illusion—modulated by the  $\alpha$  parameter (Figure



4.8). Overall, our model provides *i*) useful insights into the process dynamics during the RHI, going beyond previous static computational models, and *ii*) a unified mathematical account that can distinguish between the ownership illusion and the perceptual drift.

In the following subsections, we discuss how this model’s predictions may contribute to future research in body ownership: its neural underpinnings, its relation with other recent body-ownership models, and how to empirically test the model with the challenges and limitations of the current implementation.

#### 4.4.1 Potential neural underpinnings of the model

The current work introduces the model and does not explicitly embed the computations in a specific brain region. However, some neurobiological clues exist regarding where such computations could be performed in the brain. Recent neuroimaging evidence points to a network of regions involved in body ownership illusions, particularly the premotor cortex and posterior parietal cortex (IPS) (Chancel et al., 2022; Gentile et al., 2013). The posterior parietal cortex is especially relevant to our precision-based model, as it plays a key role in multisensory integration and uncertainty processing. Recently, Chancel et al. (2020) (Chancel and Ehrsson, 2020) directly applied a causal inference model to fMRI data, identifying the posterior parietal cortex as crucial for resolving uncertainty during body ownership decisions. Additionally, the premotor cortex shows consistent activation during body ownership illusions (Ehrsson, 2020), suggesting its involvement in maintaining coherent body representations under sensory uncertainty.

#### 4.4.2 Alignment with other uncertainty-based models

There is a strong connection between our model and other models built on Bayesian Causal Inference (Samad et al., 2015). However, there are important differences. For instance, Chancel and colleagues (Chancel et al., 2022) defined a model where the implementation of uncertainty relies on visual noise manipulation via augmented reality with three discrete noise levels. However, our model treats uncertainty as a continuous variable that dynamically adapts through precision estimation. This fundamental difference in uncertainty handling reflects distinct theoretical approaches to the decision-making process. In (Chancel et al., 2022), a fixed decision criterion is used for each noise level to study how participants infer common cause probability, while our model employs online precision adaptation, where the system actively learns about and adjusts to sensory uncertainty. These models should be viewed as complementary rather than competing accounts. Relevantly, (Chancel et al., 2022) provides rigorous psychophysical evidence for the role of Bayesian inference in body ownership, demonstrating how increased visual

noise facilitates ownership illusions. This aligns with our model’s predictions about the relationship between sensory uncertainty and integration. However, our model extends this by suggesting how the brain might continuously update its estimates of sensory precision rather than simply responding to externally manipulated uncertainty levels. Their careful experimental paradigm validates key principles that our more mechanistic model assumes, particularly regarding how the brain weighs and integrates multisensory information under varying levels of uncertainty.

The combination of Chancel et al’s experimental rigor (Chancel and Ehrsson, 2020; Chancel et al., 2022) with our more mechanistic model of precision estimation suggests promising directions for future research. Their findings about the relationship between synchrony detection and ownership judgments could be extended to test conditions of continuously varying uncertainty, as predicted by our model. Similarly, our model’s predictions about dynamic precision estimation could be tested using modified versions of their psychophysical paradigms. Together, these approaches could provide a more complete understanding of how the brain determines and maintains body ownership through precision-weighted integration of multisensory signals.

### 4.4.3 Testing our model and limitations

Our computational model of body ownership based on precision adaptation makes several testable predictions that could be empirically validated. Here, we outline experimental approaches that could confirm or challenge key aspects of our model, including *i*) predictions about dynamic proprioceptive drift, *ii*) the mechanisms of model selection versus averaging, *iii*) the distinction between uncertain and absent signals, and *iv*) the temporal dynamics of the illusion. These proposed experimental tests would help validate our theoretical framework and potentially refine our understanding of the computational mechanisms underlying body ownership illusions.

**Dynamic proprioceptive drift prediction** A key prediction from our precision-based model concerns the temporal dynamics of body ownership illusions. Unlike static models that focus only on end states, our approach captures moment-to-moment changes in the proprioceptive drift (PD). Our model thus predicts that the PD should dynamically vary with each brush stroke during the RHI induction rather than being a static effect measured only after the illusion. While current experimental paradigms typically measure PD only before and after the induction period, this prediction suggests the need for continuous measurement techniques to capture potential stroke-by-stroke variations in perceived hand position. Such dynamic measurements could provide crucial validation of our model and offer new insights into the temporal dynamics of body ownership illusions.

There are already several methodologies that can be directly implemented in the model (Riemer et al., 2015; Kharrat et al., 2022), awaiting further tests of the precision-based theory developed here.

**Model selection versus model averaging** Although we used model selection as the mechanism behind for reporting body ownership, we also offered a way to perform model averaging through the weighted combination to capture individual differences and the illusion sensitivity. While model averaging is the most widely accepted computational account to combine different uncertain models in the brain, we took inspiration from the affordance competition hypothesis (ACH) that states that there are at least two competing neural correlates for distinct actions in the parieto-dorsal stream (Cisek and Kalaska, 2005; Cisek, 2007; Sakreida et al., 2016). Once a decision is made, only one such neural activation of hand position can remain. This leads to the final action, referred to as the winner-take-all algorithm. The ACH slightly resembles our modeling effort, but there are differences. Our approach does not consider a ‘simple’ hand movement but rather a more complicated model that combines body ownership via an intricate network of multisensory integration. Thus, different models of body ownership compete rather than actions that are generated from these models. Ultimately, though, we can consider a similar switch in neuronal activity of brain regions as characterized by the ACH showing similar neuronal patterns.

**Uncertain signals versus no signals** One of the assumptions of our body-ownership illusion model states that although the real hand is “invisible”, it still produces a signal that is highly noisy – from the agent’s perspective – rather than not producing a signal at all. This plays a role in the switch between the two models and can be used to study different levels of perceptual manipulation where, instead of fully hiding the hand, the degree of visibility is manipulated (similarly to (Chancel et al., 2022)). We emphasized that the primary driving mechanism stems from the adaptation to the changing precision of the sensory input signal that triggers the switch between hypotheses or models of reality. Assuming high noise of an invisible hand rather than no signal is an alternative way of modeling input signals, which is, however backed up by logical arguments regarding the free energy principle called the dark room problem (Friston et al., 2012). Simply, if the goal of an agent is to resolve all uncertainty, how come all agents do not live in a dark room, where everything is completely expectable? The resolution is that darkness is completely uncertain.

**Time scales, proprioceptive drift and onset illusion** In our simulations (Figures 4.6 and 4.8) the proprioceptive drift appears to follow the touch of the stimulation and dissipates when the strokes no longer occur. This, in empirical settings, would result in a participant perceiving the proprioceptive drift only during an episode of touch. A more stable belief of the proprioceptive drift can be obtained using temporal integration, operating over multiple timescales, such as implementing a slow-updating belief state that retains information over 10-30 seconds rather than the rapid updates (sub-second) currently implemented in our model, for instance, involving exponentially weighted averaging of the proprioceptive position estimate across multiple touch episodes, similar to integrating over the signal across a longer period (as in (Hinz et al., 2018)).

Another issue related to time scales within the simulation is that our model predicts an unrealistically rapid onset of the ownership illusion and proprioceptive drift. While empirical studies show the rubber hand illusion typically requires 10-30 seconds to develop (Chancel and Ehrsson, 2020; Ehrsson et al., 2004; Kalckert and Ehrsson, 2017), our model shows effects within 2-3 seconds of hand occlusion. This discrepancy likely stems from our simplified implementation of precision adaptation dynamics. In reality, the brain may require longer periods of sustained multisensory evidence to adjust precision weights and switch between competing models of body ownership. Incorporating slower learning rates or additional constraints on precision adaptation will better match the empirically observed temporal dynamics. Additionally, rather than implementing a binary switch between hypotheses, a more gradual accumulation of evidence through probabilistic model averaging might better capture the progressive emergence of the illusion reported in experimental studies.

**Temporal synchronicity** While we model synchronicity generation of the signals within the participant’s sensitivity (Lanfranco et al., 2023) to the body illusion through the  $\alpha$  parameter, one of the main limitations of our model is that the simulation does not account for deviations in the visuo-tactile temporal synchronization of the strokes, a well-studied aspect of the RHI. Thus, explicit temporal synchronisation of the signals should be incorporated to allow fair comparison with human data experiments.

#### 4.4.4 Future work

In summary, future research should focus on the experimental validation of the model to assess the dynamics of switching between competing models depending on the posterior uncertainty (precision), as proposed in our study. In this context, researchers could create scenarios where multiple states of body ownership are simultaneously induced, allowing only one to predominate. The introduction of different temporal scales and explicit

temporal synchronicity could refine the model to align with recent human experimental data, allowing for more direct empirical testing. Furthermore, the generalisation of the proposed model for any competing hypotheses would enable the study of probabilistic transition dynamics between different competing perceptions of reality. These modifications could potentially explain phenomena such as double-limb and supernumerary limb illusions (Guterstam et al., 2013; Fan et al., 2021), providing valuable insights into the mechanisms of body ownership.

## Acknowledgments

FN and FZ have received funding from, Serotonin & Beyond project, the European Union's Horizon 2020 Marie Skłodowska-Curie Actions under grant agreement No. 953327. AM is funded by the EU Metatool project (Grant agreement 101070940) under the EIC Pathfinder program. The research has been partially funded by the DEEPSELF project (467045002 DFG SPP The Active Self).

# Bibliography

- Abdulkarim, Z. and Ehrsson, H. H. (2016). No causal link between changes in hand position sense and feeling of limb ownership in the rubber hand illusion. *Attention, Perception, & Psychophysics*, 78:707–720.
- Abdulkarim, Z., Hayatou, Z., and Ehrsson, H. H. (2021). Sustained rubber hand illusion after the end of visuotactile stimulation with a similar time course for the reduction of subjective ownership and proprioceptive drift. *Experimental brain research*, 239(12):3471–3486.
- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., and Friston, K. J. (2013). The computational anatomy of psychosis. *Frontiers in psychiatry*, 4:47.
- Anil Meera, A., Novicky, F., Parr, T., Friston, K., Lanillos, P., and Sajid, N. (2022). Reclaiming saliency: Rhythmic precision-modulated action and perception. *Frontiers in Neurorobotics*, 16:896229.
- Armell, K. C. and Ramachandran, V. S. (2003). Projecting sensations to external objects: evidence from skin conductance response. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1523):1499–1506.
- Botvinick, M. and Cohen, J. (1998). Rubber hands ‘feel’ touch that eyes see. *Nature*, 391(6669):756–756.
- Brown, H., Adams, R. A., Parees, I., Edwards, M., and Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive processing*, 14:411–427.
- Chancel, M. and Ehrsson, H. H. (2020). Which hand is mine? discriminating body ownership perception in a two-alternative forced-choice task. *Attention, Perception, & Psychophysics*, 82(8):4058–4083.

- Chancel, M. and Ehrsson, H. H. (2023). Proprioceptive uncertainty promotes the rubber hand illusion. *Cortex*, 165:70–85.
- Chancel, M., Ehrsson, H. H., and Ma, W. J. (2022). Uncertainty-based inference of a common cause for body ownership. *Elife*, 11:e77221.
- Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485):1585–1599.
- Cisek, P. and Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron*, 45(5):801–814.
- Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Costantini, M. and Haggard, P. (2007). The rubber hand illusion: sensitivity and reference frame for body ownership. *Consciousness and cognition*, 16(2):229–240.
- Doya, K., Ishii, S., Pouget, A., and Rao, R. P. (2007). *Bayesian brain: Probabilistic approaches to neural coding*. MIT press.
- Ehrsson, H. H. (2012). The Concept of Body Ownership and Its Relation to Multisensory Integration. In *The New Handbook of Multisensory Processing*. The MIT Press.
- Ehrsson, H. H. (2020). Multisensory processes in body ownership. *Multisensory perception*, pages 179–200.
- Ehrsson, H. H. (2022). Bodily illusions. *The Routledge Handbook of Bodily Awareness*.
- Ehrsson, H. H., Spence, C., and Passingham, R. E. (2004). That’s my hand! activity in premotor cortex reflects feeling of ownership of a limb. *Science*, 305(5685):875–877.
- Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433.
- Fan, C., Coppi, S., and Ehrsson, H. H. (2021). The supernumerary rubber hand illusion revisited: Perceived duplication of limbs and visuotactile events. *Journal of Experimental Psychology: Human Perception and Performance*, 47(6):810.
- Fang, W., Li, J., Qi, G., Li, S., Sigman, M., and Wang, L. (2019). Statistical inference of body representation in the macaque brain. *Proceedings of the National Academy of Sciences*, 116(40):20151–20157.

- Feeney, E. J., Groman, S. M., Taylor, J. R., and Corlett, P. R. (2017). Explaining delusions: reducing uncertainty through basic and computational neuroscience. *Schizophrenia bulletin*, 43(2):263–272.
- Feldman, H. and Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in human neuroscience*, 4:215.
- Finotti, G., Garofalo, S., Costantini, M., and Proffitt, D. R. (2023). Temporal dynamics of the rubber hand illusion. *Scientific Reports*, 13(1):7526.
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends in cognitive sciences*, 13(7):293–301.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2):127–138.
- Friston, K. (2023). Computational psychiatry: from synapses to sentience. *Molecular psychiatry*, 28(1):256–268.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017). Active inference: a process theory. *Neural computation*, 29(1):1–49.
- Friston, K., Thornton, C., and Clark, A. (2012). Free-energy minimization and the dark-room problem. *Frontiers in psychology*, page 130.
- Gentile, G., Guterstam, A., Brozzoli, C., and Ehrsson, H. H. (2013). Disintegration of multisensory signals from the real hand reduces default limb self-attribution: an fmri study. *Journal of Neuroscience*, 33(33):13350–13366.
- Guterstam, A., Gentile, G., and Ehrsson, H. H. (2013). The invisible hand illusion: multisensory integration leads to the embodiment of a discrete volume of empty space. *Journal of cognitive neuroscience*, 25(7):1078–1099.
- Guterstam, A., Larsson, D. E., Zeberg, H., and Ehrsson, H. H. (2019). Multisensory correlations—not tactile expectations—determine the sense of body ownership. *PLoS One*, 14(2):e0213265.
- Hide, M., Ito, Y., Kuroda, N., Kanda, M., and Teramoto, W. (2021). Multisensory integration involved in the body perception of community-dwelling older adults. *Scientific Reports*, 11(1):1581.



- Hinz, N.-A., Lanillos, P., Mueller, H., and Cheng, G. (2018). Drifting perceptual patterns suggest prediction errors fusion rather than hypothesis selection: replicating the rubber-hand illusion on a robot. In *2018 Joint IEEE 8th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, pages 125–132. IEEE.
- Hohwy, J. (2013). *The predictive mind*. OUP Oxford.
- Holmes, N. P., Crozier, G., and Spence, C. (2004). When mirrors lie: “visual capture” of arm position impairs reaching performance. *Cognitive, Affective, & Behavioral Neuroscience*, 4(2):193–200.
- Holmes, N. P., Snijders, H. J., and Spence, C. (2006). Reaching with alien limbs: Visual exposure to prosthetic hands in a mirror biases proprioception without accompanying illusions of ownership. *Perception & psychophysics*, 68(4):685–701.
- Jardri, R. and Denève, S. (2013). Computational models of hallucinations. *The neuroscience of hallucinations*, pages 289–313.
- Jeganathan, J. and Breakspear, M. (2021). An active inference perspective on the negative symptoms of schizophrenia. *The Lancet Psychiatry*, 8(8):732–738.
- Kalckert, A. and Ehrsson, H. (2017). The onset time of the ownership sensation in the moving rubber hand illusion. *Frontiers in psychology*, 8:344.
- Kalckert, A., Perera, A. T.-M., Ganesan, Y., and Tan, E. (2019). Rubber hands in space: the role of distance and relative position in the rubber hand illusion. *Experimental Brain Research*, 237(7):1821–1832.
- Kharrat, I., Achour, I., Trabelsi, J. J., Trigui, M., Thabet, W., Mnejja, M., Hammami, B., Chakroun, A., and Charfeddine, I. (2022). Prediction of difficulty in direct laryngoscopy. *Scientific Reports*, 12(1):10722.
- Kiltner, K., Maselli, A., Kording, K. P., and Slater, M. (2015). Over my fake body: body ownership illusions for studying the multisensory basis of own-body perception. *Frontiers in human neuroscience*, 9:141.
- Lanfranco, R. C., Chancel, M., and Ehrsson, H. H. (2023). Quantifying body ownership information processing and perceptual bias in the rubber hand illusion. *Cognition*, 238:105491.

- Lanillos, P., Franklin, S., Maselli, A., and Franklin, D. W. (2021). Active strategies for multisensory conflict suppression in the virtual hand illusion. *Scientific Reports*, 11(1):22844.
- Limanowski, J. and Blankenburg, F. (2016). Integration of visual and proprioceptive limb position information in human posterior parietal, premotor, and extrastriate cortex. *Journal of Neuroscience*, 36(9):2582–2589.
- Limanowski, J. and Friston, K. (2020). Active inference under visuo-proprioceptive conflict: Simulation and empirical results. *Scientific reports*, 10(1):4010.
- Maselli, A., Lanillos, P., and Pezzulo, G. (2022). Active inference unifies intentional and conflict-resolution imperatives of motor control. *PLoS computational biology*, 18(6):e1010095.
- Maselli, A. and Slater, M. (2013). The building blocks of the full body ownership illusion. *Frontiers in human neuroscience*, 7:83.
- Meera, A. A. and Lanillos, P. (2023). Adaptive noise covariance estimation under colored noise using dynamic expectation maximization. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 165–171. IEEE.
- Meera, A. A. and Lanillos, P. (2024). Confidence-aware decision-making and control for tool selection. *arXiv preprint arXiv:2403.03808*.
- Meera, A. A. and Wisse, M. (2020). Free energy principle based state and input observer design for linear systems with colored noise. In *2020 American Control Conference (ACC)*, pages 5052–5058. IEEE.
- Mumford, D. (1992). On the computational architecture of the neocortex: II the role of cortico-cortical loops. *Biological cybernetics*, 66(3):241–251.
- Novicky, F., Offergeld, J., Janssen, S., and Lanillos, P. (2023a). Robotic active tactile sensing inspired by serotonergic modulation using active inference. In *Conference on Biomimetic and Biohybrid Systems*, pages 33–55. Springer.
- Novicky, F., Parr, T., Friston, K., Mirza, M. B., and Sajid, N. (2023b). Bistable perception, precision and neuromodulation. *Cerebral Cortex*, 34(1):bhad401.
- Parr, T., Benrimoh, D. A., Vincent, P., and Friston, K. J. (2018). Precision and false perceptual inference. *Frontiers in integrative neuroscience*, 12:39.

- Parr, T. and Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of the Royal Society Interface*, 14(136):20170376.
- Parr, T. and Friston, K. J. (2019). Attention or salience? *Current opinion in psychology*, 29:1–5.
- Parr, T., Pezzulo, G., and Friston, K. J. (2022). *Active inference: the free energy principle in mind, brain, and behavior*. MIT Press.
- Peviani, V. C., Miller, L. E., and Medendorp, W. P. (2024). Biases in hand perception are driven by somatosensory computations, not a distorted hand model. *Current Biology*, 34(10):2238–2246.
- Preston, C. (2013). The role of distance from the body and distance from the real hand in ownership and disownership during the rubber hand illusion. *Acta psychologica*, 142(2):177–183.
- Rao, R. P. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87.
- Reynolds, J. H. and Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61(2):168–185.
- Reynolds, J. H., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of v4 neurons. *Neuron*, 26(3):703–714.
- Riemer, M., Bublatzky, F., Trojan, J., and Alpers, G. W. (2015). Defensive activation during the rubber hand illusion: Ownership versus proprioceptive drift. *Biological psychology*, 109:86–92.
- Rohde, M., Di Luca, M., and Ernst, M. O. (2011). The rubber hand illusion: feeling of ownership and proprioceptive drift do not go hand in hand. *PloS one*, 6(6):e21659.
- Rohde, M., Wold, A., Karnath, H.-O., and Ernst, M. O. (2013). The human touch: skin temperature during the rubber hand illusion in manual and automated stroking procedures. *PloS one*, 8(11):e80688.
- Rood, T., Gerven, M. v., and Lanillos, P. (2020). A deep active inference model of the rubber-hand illusion. In *International Workshop on Active Inference*, pages 84–91. Springer.

- Sakreida, K., Effnert, I., Thill, S., Menz, M. M., Jirak, D., Eickhoff, C. R., Ziemke, T., Eickhoff, S. B., Borghi, A. M., and Binkofski, F. (2016). Affordance processing in segregated parieto-frontal dorsal stream sub-pathways. *Neuroscience & Biobehavioral Reviews*, 69:89–112.
- Samad, M., Chung, A. J., and Shams, L. (2015). Perception of body ownership is driven by bayesian sensory inference. *PloS one*, 10(2):e0117178.
- Sancaktar, C., van Gerven, M. A., and Lanillos, P. (2020). End-to-end pixel-based deep active inference for body perception and action. In *2020 Joint IEEE 10th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, pages 1–8. IEEE.
- Slater, M., Pérez Marcos, D., Ehrsson, H., and Sanchez-Vives, M. V. (2009). Inducing illusory ownership of a virtual body. *Frontiers in neuroscience*, page 29.
- Tosi, G., Montesana, B., and Romano, D. (2023). The correlation between proprioceptive drift and subjective embodiment during the rubber hand illusion: A meta-analytic approach. *Quarterly Journal of Experimental Psychology*, page 17470218231156849.
- Tsakiris, M. and Haggard, P. (2005). The rubber hand illusion revisited: visuotactile integration and self-attribution. *Journal of experimental psychology: Human perception and performance*, 31(1):80.
- Zhao, Y., Lu, E., and Zeng, Y. (2023). Brain-inspired bodily self-perception model for robot rubber hand illusion. *Patterns*, 4(12).

## 4.5 Appendix

### 4.5.1 Related computational models of body illusions

This section discusses previous computational models of body illusions and how they approached the topic as a problem of multisensory conflicts, especially the RHI (Botvinick and Cohen, 1998; Hinz et al., 2018; Limanowski and Friston, 2020) and its virtual reality (VR) version (Slater et al., 2009; Lanillos et al., 2021). One of the first seminal works on modeling RHI comes from (Samad et al., 2015), where a model driven by Bayesian inference is proposed, focusing on the synchronization of stimuli in time. The authors went beyond the optimal integration model (Fang et al., 2019), by introducing causal inference, where two potential models of reality are presented: either the tactile information is

coming from the subject own hand or from the fake hand. However, this account presents body ownership as a static probability computation, where no explanation is given on the inner functional mechanism producing the dynamics of the perceptual drifts nor the illusion (Finotti et al., 2023). Bayesian causal inference was revisited in (Chancel et al., 2022), focusing on the impact of the input noise to the emergence of the RHI.

Another line of research inspired by the predictive processing approach to perception (Rao and Ballard, 1999; Hohwy, 2013; Clark, 2015), yielded in a set of computational models of body illusions that follow the prediction error minimization as the inner mechanism. (Hinz et al., 2018) aimed to investigate how the distance between the rubber and real hands modulates the proprioceptive drift. The work compared the data obtained from a human RHI experiment and a humanoid RHI experiment. It used the learnt generative model of the sensory signals (proprioceptive, visual and visuotactile) to dynamically estimate the location of the real hand. Mathematically, the model used approximated Bayesian inference (i.e., free energy minimization) to infer the body pose given the input signals. Whilst, this work supports the separation of body-ownership illusion and proprioceptive drift, as these seem to be split into two interconnected processes, as suggested by (Abdulkarim and Ehrsson, 2016; Abdulkarim et al., 2021), it only provided the computation for replicating the proprioceptive drift and not the body-ownership. In a follow-up model (Lanillos et al., 2021; Maselli et al., 2022), the focus shifted from the passive perceptual process to the active component of the illusion. The model and human evidence suggests that participants actively move their body position (or generate forces) to minimize the discrepancy between proprioceptive signals coming from the real hand, and visual signals coming from the rubber hand. As an aside result, the model explained that perceptual drifts are not necessarily linked to tactile stimulation, but is achievable with visual and proprioceptive signals in VR (Maselli and Slater, 2013).

The models described above assumed that the location of the hand is already processed. The work by (Rood et al., 2020) introduced deep neural networks to allow the body illusion model to use large-scale visual inputs. The architecture implemented the deep active inference framework for continuous time variables (Sancaktar et al., 2020). This model was able to account for the perceptual and active drifts dynamics during the illusion using synthetic images generated from VR. Furthermore, it introduced the synchronicity of the input signals as a variable. However, this model did not contribute to the biological underpinnings of body ownership. There are other connectivity based models, such as (Zhao et al., 2023), where the neural network connections are inspired by brain functional segregation.

### 4.5.2 Precision parametrization

Precision learning follows the online learning of measurement noise precision (or the inverse noise covariance matrix)  $\Pi^z$ . Intuitively, this refers to the brain’s attentional mechanism of gauging the right amount of noise in the environment so as to make inferences about the world under high noise (uncertainties)—see (Reynolds et al., 2000; Reynolds and Heeger, 2009). We postulate that an accurate attentional mechanism (online precision learning) is critical to experiencing the RHI. We use the FEP for online precision learning. This follows two steps: modeling and learning. The noise precision  $\Pi^z$  is modeled through an exponential parametrization with respect to parameter  $\lambda = [\lambda^1 \ \lambda^2 \ \dots \ \lambda^m]^T$  (where  $m$  is the number of observations) as:

$$\Pi^z = \begin{bmatrix} e^{\lambda^1} & 0 & \dots & \\ 0 & e^{\lambda^2} & 0 & \dots \\ 0 & 0 & e^{\lambda^3} & \dots \\ \dots & \dots & \dots & \dots \end{bmatrix}_{m \times m}. \quad (4.10)$$

This parametrization was chosen to ensure that  $\Pi^z$  always remains positive definite ( $\Pi^z \succ O$ ) throughout the estimation. Intuitively, this ensures the noise covariance to be a positive real number.

### 4.5.3 Code and parameters used for simulations

The MATLAB code <sup>3</sup> for the simulation is available: [https://github.com/ajitham123/precision\\_RHI/](https://github.com/ajitham123/precision_RHI/). This section lists the parameters that were used in our simulations. The generative process uses the model given in Eq (4.3). The same model is used as  $M_1$  with  $Ay = 0$  and  $Ry = 0.8$ , while the model in Eq (4.4) is used as  $M_2$ . The  $\alpha$  in Eq (4.2) is set to 10 for a low sensitivity to the illusion, and to -10 for a high sensitivity to the illusion. The generative process is simulated for 32 seconds with a sampling time of 0.1s. The observation noise for each signal is parametrized using  $\lambda = [\lambda^{Av} \ \lambda^{Rv} \ \lambda^{At} \ \lambda^{Py}]^T = [9, 7, 8, 10]$  for hand visible, and  $\lambda = [-1, 7, 8, 4]$  for hand covered (see Figure (4.3) for the generated data). The causes (hand strokes)  $\mathbf{X}$  (plotted in blue in Figure (4.6)) comprises of i) four Gaussian bumps centered around  $t = 4, 6, 8, 10$  seconds for hand visible case and ii) five Gaussian bumps centered around  $t = 20, 22, 24, 26, 28$  seconds for hand covered case. The magnitude of these bumps for **Ea**, **Er**, **Et**, **Sy** are 0.5, 1, 1, 0, respectively (when stroked). The estimation starts with priors on noise parameters (in Eq (4.6)) as  $\eta^\lambda = [10, 10, 10, 10]$  with prior precision  $P^\lambda = \text{diag}([e^{-4}, e^{-4}, e^{-4}, e^{-4}])$  where  $\text{diag}(\cdot)$  converts an array into a diagonal matrix. All the numbers described here were arbitrarily selected to start the

<sup>3</sup>The code will be made public upon paper acceptance

estimation from low noise levels with low confidence. For breaking the precision learning in figure 4.4, a higher prior precision ( $P^\lambda = \text{diag}([e^{10}, e^{-4}, e^{-4}, e^{-4}])$ ) for the signal  $Av$  was used. The state estimation starts at  $t = 0$  with a unit precision. All the learning rates in Section 4.2.2.4 are set to  $k = 4$  for a quick learning. All figures are generated using these parameters as the basic reference for the simulation, unless another value is mentioned for a few variables within the text.

## Chapter 5

# Reclaiming saliency: Rhythmic precision-modulated action and perception

Based on:

Meera, A. A., Novicky, F., Parr, T., Friston, K., Lanillos, P., & Sajid, N. (2022). Reclaiming saliency: Rhythmic precision-modulated action and perception. *Frontiers in Neurorobotics*, 16. <https://doi.org/10.3389/fnbot.2022.850512>



## Abstract

Computational models of visual attention in artificial intelligence and robotics have been inspired by the concept of a saliency map. These models account for the mutual information between the (current) visual information and its estimated causes. However, they fail to consider the circular causality between perception and action. In other words, they do not consider where to sample next, given current beliefs. Here, we reclaim salience as an active inference process that relies on two basic principles: uncertainty minimisation and rhythmic scheduling. For this, we make a distinction between attention and salience. Briefly, we associate attention with precision control, i.e., the confidence with which beliefs can be updated given sampled sensory data, and salience with uncertainty minimisation that underwrites the selection of future sensory data. Using this, we propose a new account of attention based on rhythmic precision-modulation and discuss its potential in robotics, providing numerical experiments that showcase its advantages for state and noise estimation, system identification and action selection for informative path planning.

## Keywords

Attention • Saliency • Free-energy principle • Active inference • Precision • Brain-inspired robotics • Cognitive robotics

## 5.1 Introduction

Attention is a fundamental cognitive ability that determines which events from the environment, and the body, are preferentially processed (Itti and Koch, 2001). For example, the motor system directs the visual sensory stream by orienting the fovea centralis (i.e., the retinal region of highest visual acuity) towards points of interest within the visual scene. Thus, the confidence with which the causes of sampled visual information are inferred is constrained by the physical structure of the eye – and eye movements are necessary to minimise uncertainty about visual percepts (Ahnelt, 1998). In neuroscience, this can be attributed to two distinct, but highly interdependent attentional processes: (i) attentional gain mechanisms reliant on estimating the sensory precision of current data (Feldman and Friston, 2010; Yang et al., 2016a), and (ii) attentional salience that involves actively engaging with the sensorium to sample appropriate future data (Parr and Friston, 2019; Lengyel et al., 2016). Here we refer to perceptual-related salience, i.e., processing of low-level visual information (Santangelo, 2015). Put simply, we formalise the fundamental difference between attention — as optimising perceptual processing —

and salience as optimising the sampling of what is processed. This highlights the dynamic, circular nature with which biological agents acquire, and process, sensory information.

Understanding the computational mechanisms that undergird these two attentional phenomena is pertinent for deploying apt models of (visual) perception in artificial agents (Klink et al., 2014; Mousavi et al., 2016; Atrey et al., 2019) and robots (Frintrop and Jensfelt, 2008; Begum and Karay, 2010; Ferreira and Dias, 2014; Lanillos et al., 2015a). Previous computational models of visual attention, used in artificial intelligence and robotics, have been inspired (and limited) by the feature integration theory proposed by (Treisman and Gelade, 1980) and the concept of a saliency map (Tsotsos et al., 1995; Itti and Koch, 2001; Borji and Itti, 2012). Briefly, a saliency map is a static two-dimensional ‘image’ that encodes stimulus relevance, e.g., the importance of particular region. These maps are then used to isolate relevant information for control (e.g., to direct foveation of the maximum valued region). Accordingly, computational models reliant on this formulation do not consider the circular-dependence between action selection and cue relevance – and simply use these static saliency maps to guide action.

In this article, we adopt a first principles account to disambiguate the computational mechanisms that underpin attention and salience (Parr and Friston, 2019) and provide a new account of attention. Specifically, our formulation can be effectively implemented for robotic systems and facilitates both state-estimation and action selection. For this, we associate attention with precision control, i.e., the confidence with which beliefs can be updated given (current) sampled sensory data. Salience is associated with uncertainty minimisation that influences the selection of future sensory data. This formulation speaks to a computational distinction between action selection (i.e., where to look next) and visual sampling (i.e., what information is being processed). Importantly, recent evidence demonstrates the rhythmic nature of these processes via a theta-cycle coupling that fluctuates between high and low precision—as unpacked in Sec 5.2. From a robotics perspective, resolving uncertainty about states of affair speaks to a form of Bayesian optimality, in which decisions are made to maximise expected information gain (Lindley, 1956; Friston et al., 2021; Sajid et al., 2021a). The duality between attention and salience is important for resolving uncertainty and enabling active perception. Significantly, it addresses an important challenge for defining autonomous robotics systems that can balance optimally between data assimilation (i.e., confidently perceiving current observations) and exploratory behaviour to maximise information gain (Bajcsy et al., 2018).

In what follows, we review the neuroscience of attention and salience (Sec. 5.2) to develop a novel (computational) account of attention based on precision-modulation that underwrites perception and action (Sec. 5.3). Next, we face-validate our formulation within a robotics context using numerical experiments (Sec. 5.4). The robotics

implementation instantiates a free energy principle (FEP) approach to information processing (Friston, 2010). This allows us to modulate the (appropriate) precision parameters to solve relevant robotics challenges in perception and control; namely, state-estimation (Sec. 5.4.2.2), system identification (Sec. 5.4.2.3), planning (Sec. 5.4.3), and active perception (Sec. 5.4.3.3). We conclude with a discussion of the requisite steps for instantiating a full-fledged computational model of precision-modulated attention – and its implications in a robotics setting.

## 5.2 Attention and salience in neuroscience

Our interactions with the world are guided by efficient gathering and processing of sensory information. The quality of these acquired sensory data is reflected in attentional resources that select sensations which influence our beliefs about the (current and future) states of affairs (Lengyel et al., 2016; Yang et al., 2016b). This selection is often related to gain control, i.e., an increase of neural spikes when an object is attended to. However, gain control only accounts for half the story because we can only attend to those objects that are within our visual field. Accordingly, if a salient object is outside the centre of our visual field, we orient the fovea to points of interest. This involves two separate, but often conflated, processes: attention and salience – where the former relates to processing current visual data, and the latter to ensuring the agent samples salient data in the future (Parr and Friston, 2019). That these two processes are strongly coupled is exemplified by the pre-motor theory of attention (Rizzolatti et al., 1987), which highlights the close relationship between overt saccadic sampling of the visual field and the covert deployment of attention in the absence of eye movements. Specifically, it posits that covert attention<sup>1</sup> is realised via processes that are generated by particular eye movements but inhibits the action itself. In this sense, it does not distinguish between covert and overt<sup>2</sup> types of attention.

From a first principles (Bayesian) account, it is necessary to separate between attention and salience because they speak to different optimisation processes. Explicitly, attention as a precision-dependent (neural) gain control mechanism that facilitates optimisation of the *current* sampled sensory data (Feldman and Friston, 2010; Desimone, 1996). Conversely, salience is associated with selection of *future* data that reduces uncertainty (Mirza et al., 2016a; Friston et al., 2015a; Parr and Friston, 2019). Put simply, it is possible to optimise attention in the absence of eye movements and active vision, whereas salience is necessary to optimise the deployment of eye movements. In what follows, we formalise this distinction

---

<sup>1</sup>Covert attention is where saccadic eye movements do not occur.

<sup>2</sup>Overt attention deals with how an agent tracks the object with eye movements

with a particular focus on visual attention (Kanwisher and Wojciulik, 2000), and discuss recent findings that speak to a rhythmic coupling that underwrites periodic deployment of gain control and saccades, via modulation of distinct precision parameters.

### 5.2.1 Attention as neural gain control

Neural gain control can be regarded as an amplifier of neural communication during attention tasks (Eldar et al., 2013; Reynolds et al., 2000). Computationally, this is analogous to modulating a precision term, or the inverse temperature parameter (Feldman and Friston, 2010; Parr and Friston, 2017a). For this reason, we refer to precision and gain control interchangeably. An increase in gain amplifies the postsynaptic responses of neurons to their pre-synaptic input. Thus, gain control rests on synaptic modulation that can emphasise — or preferentially select — a particular type of sensory data. From a Bayesian perspective (Rao, 2005; Spratling, 2008; Parr et al., 2018), this speaks to the confidence with which beliefs can be updated given sampled sensory data (i.e., optimal state estimation) – under a generative model (Whiteley and Sahani, 2008; Parr et al., 2018). For example, affording high precision to certain sensory inputs would lead to confident Bayesian belief updating. However, low precision reduces the influence of sensory input by attenuating the precision of the likelihood, relative to a prior belief, and current observations would do little to resolve ensuing uncertainty. Thus, sampled visual data (from different areas) can be predicted with varying levels of precision, where attention accentuates sensory precision. The deployment of precision or attention is influenced by competition between stimuli (i.e., which sensory data to sample) and prior beliefs. Interestingly, casting attention as precision or, equivalently, synaptic gain offers a coherency between biased competition (Desimone, 1996), predictive coding (Spratling, 2008) and generic active inference schemes (Feldman and Friston, 2010; Parr et al., 2018; Brown et al., 2013; Kanai et al., 2015).

Naturally, gain control is accompanied by neuronal variability, i.e., sharpened neural responses for the same task over time. Consistent with gain control, these fluctuations in neural responses across trials can be explained by precision engineered message passing (Clark, 2013) via *(i)* normalization models (Reynolds and Heeger, 2009; Ruff and Cohen, 2016), *(ii)* temperature parameter manipulation (Mirza et al., 2019; Parr et al., 2018; Parr and Friston, 2017a; Parr et al., 2019; Feldman and Friston, 2010), or *(iii)* introduction of (conjugate hyper-)priors that are either pre-specified (Sajid et al., 2021b, 2020) or optimised using uninformed priors (Friston et al., 2003; Anil Meera and Wisse, 2021). Recently, these approaches have been used to simulate attention by accentuating predictions about a given visual stimulus (Reynolds and Heeger, 2009; Feldman and Friston, 2010; Ruff and

Cohen, 2016). For example, normalization models propose that every neuronal response is normalized within its neuronal ensemble (i.e., the surrounding neuronal responses) (Heeger, 1992; Louie and Glimcher, 2019). Thus, to amplify the neuronal response of particular neuron, the neuronal pool has to be inhibited such that that particular neuron has a sharper evoked response (Schmitz and Duncan, 2018). Importantly, these (superficially distinct) formulations simulate similar functions using different procedures to accentuate responses over a particular neuronal pool for a given neuron or a group of neurons. This introduces shifts in precision to produce attentional gain and the precision of neuronal encoding.

### 5.2.2 Salience as uncertainty minimisation

In the neurosciences, (visual) salience refers to the ‘significance’ of particular objects in the environment. Salience often implicates the superior colliculus, a region that encodes eye movements (White et al., 2017). This makes intuitive sense, as the superior colliculus plays a role in generation of eye movements – being an integral part of the brainstem oculomotor network (Raybourn and Keller, 1977) – and salient objects provide information that is best resolved in the centre of the visual field, thus motivating eye movements to that location. For this reason, our understanding of salience is a quintessentially action-driving phenomenon (Parr and Friston, 2019). Mathematically, salience has been defined as Bayesian surprise (Itti and Koch, 2001; Itti and Baldi, 2009), intrinsic motivation (Oudeyer and Kaplan, 2009), and subsequently, epistemic value under active inference (Mirza et al., 2016b; Parr et al., 2018). Active inference – a Bayesian account of perception and action (Friston et al., 2017a; Da Costa et al., 2020) – stipulates that action selection is determined by uncertainty minimisation. Formally, uncertainty minimisation speaks to minimisation of an expected free energy functional over future trajectories (Da Costa et al., 2020; Sajid et al., 2021a). This action selection objective can be decomposed into epistemic and extrinsic value, where the former pertains to exploratory drives that encourage resolution of uncertainty by sampling salient observations, e.g., only checking one’s watch when one does not know the time. However, after checking the watch there is little epistemic value in looking at it again. Generally, the tendency to seek out new locations – once uncertainty has been resolved at the current fixation point – is called inhibition of return (Klein, 2000).

From an active inference perspective, this phenomenon is prevalent because a recent action has already resolved the uncertainty about the time and checking again would offer nothing more in terms of information gain (Parr and Friston, 2019). Accordingly, salience involves seeking sensory data that have a predictable, uncertainty reducing, effect on

current beliefs about states of affairs in the world (Parr et al., 2018; Mirza et al., 2016b). Thus salience contends with beliefs about data that must be acquired and the precision of beliefs about policies (i.e., action trajectories) that dictate it. Formally, this emerges from the imperative to maximise the amount of information gained regarding beliefs, from observing the environment. Happily, prior studies have made the connection between eye movements, salience, and precision manipulation (Friston et al., 2011; Brown et al., 2013; Crevecoeur and Kording, 2017). This connection emerges from planning strategies that allow the agent to minimise uncertainty by garnering the right kind of data.

Next, we consider recent findings on how the coupling of these two mechanisms, attention and salience, may be realised in the brain.

### 5.2.3 Rhythmic coupling of attention and salience

To illustrate the coupling between attention and salience, we turn to a recent rhythmic theory of attention. The theory proposes that coupling of saccades, during sampling of visual information, happens at neuronal and behavioural theta oscillations; a frequency of 3-8Hz (Fiebelkorn and Kastner, 2019, 2021). This frequency simultaneously allows for: (i) a systematic integration of visual samples with action, and (ii) a temporal schedule to disengage and search the environment for more relevant information.


Given that gain control is related to increased sensory precision, we can accordingly relate saccadic eye movements to the decreased precision. This introduces saccadic suppression, a phenomenon that decreases visual gain during eye movements (Crevecoeur and Kording, 2017). This phenomenon was described by Helmholtz who observed that externally initiated eye movements (e.g., when oneself gently presses a side of an eye) eludes the saccadic suppression that accompanies normal eye movements – and we see the world shift, because optic flow is not attenuated (Helmholtz, 1925). An interesting consequence of this is that, as eye movements happen periodically (Rucci et al., 2018; Benedetto et al., 2020), there must be a periodic switch between high and low sensory precision, with high precision (or enhanced gain) during fixations and low precision (or suppressed gain) during saccades. Interestingly, it has been shown that rather than having action resetting the neural periodicity, it is better understood as something that aligns within an already existing rhythm (Tomassini et al., 2017; Hogendoorn, 2016). Additionally, the rhythmicity of higher and lower fidelity of sensory sampling has been shown to fluctuate rhythmically around 3Hz (Benedetto and Morrone, 2017), suggesting that action emerges rhythmically when visual precision is low (Hogendoorn, 2016), triggering salience.

Building upon this, we hypothesise that theta rhythms generated in the fronto-parietal network (Fiebelkorn et al., 2018; Helfrich et al., 2018; Fiebelkorn and Kastner, 2020)

couples saccades with saccadic suppression causing the switches between visual sampling and saccadic shifting. This introduces a diachronic aspect to the belief updating process (Friston et al., 2020; Parr and Pezzulo, 2021; Sajid et al., 2022); i.e., sequential fluctuations between attending to current data (perception) and seeking new data (action). This supports empirical findings that both eye movements (Sommer and Wurtz, 2006) and filtering irrelevant information (Nakajima et al., 2019; Phillips et al., 2016; Fiebelkorn and Kastner, 2020) are initiated in this cortical network. Interestingly, both eye movements and visual filtering then propagate to sub-cortical regions, i.e., the superior colliculus—for saliency map composition (White et al., 2017)—and the thalamus—for gain control (Kanai et al., 2015; Fiebelkorn et al., 2019), respectively. Furthermore, this is consistent with recent findings that the periodicity of neural responses are important for understanding the relation of motor responses and sensory information – i.e., perception-action coupling (Benedetto et al., 2020). Importantly, theta rhythms also speak to the speed (i.e., the temporal schedule) with which visual information is sampled from the environment (Busch and VanRullen, 2010; Dugué et al., 2015, 2016; Helfrich et al., 2018). Meaning visual information is not sampled continuously, as our visual experiences would suggest, but rather it is made of successive discrete samples (VanRullen, 2016; Parr et al., 2021).

The prefrontal theta rhythm has been associated with working memory (WM), a process that holds compressed information about the previously observed stimuli, in the sense that measured power in this frequency range using electroencephalography increases during tasks that place demands on WM (Hsieh and Ranganath, 2014; Pomper and Ansorge, 2021; Brzezicka et al., 2019; Balestrieri et al., 2021; Axmacher et al., 2010; Peters et al., 2020; Köster et al., 2018). The implication is that the neural processes that underwrite WM may depend upon temporal cycles with periods similar to that of perceptual sampling. Importantly, this cognitive process is influenced by how salient a particular stimulus was (Santangelo and Macaluso, 2013; Santangelo et al., 2015; Fine and Minnery, 2009). Moreover, WM has been implicated with attentional mechanisms (Gazzaley and Nobre, 2012; Knudsen, 2007; Oberauer, 2019; Panichello and Buschman, 2021; Peters et al., 2020). This is aligned with our account where we illustrate a rhythmic coupling between salience and attention.

In summary, the computations that underwrite attention and active vision are coupled and exhibit circular causality. Briefly, selective attention and sensory attenuation optimize the processing of sensory samples and which particular visual percepts are inferred. In turn, this determines appropriateness of future eye movements (or actions) and shapes which prior stimuli are encoded into the agent’s working memory. Interestingly, the close functional (and computational) link between the two mechanisms endorses the pre-motor



**Figure 5.1:** A graphical illustration of the precision-modulated account of perception and action. Saliency and attention are computed based upon beliefs (assumed to be) encoded in parts of the fronto-parietal network and realised in distinct brain regions: superior colliculus (SC) for perception as inference and thalamus for planning as inference, respectively. To deploy attentional processes efficiently, these two mechanisms have to be aligned, which is done rhythmically, hypothetically in theta frequency. This coupling enables the saccadic suppression phenomenon through fluctuations in precision (on an arbitrary scale). When precision is low (i.e., the trough of the theta rhythm), the saccade emerges. Note that there might be distinct processes inhibiting the action (e.g., covert attention), and (despite a decline in precision) saccades might not emerge in every theta cycle. On the other hand, high precision facilitates confident inferences about the causes of visual data. Under this account, thalamus is used for initiating gain control (or visual sampling in general) by providing stronger sensory input, while superior colliculus dictates next saccades, that lead to most informative fixation positions.

theory of attention.

### 5.3 Proposed precision-modulated account of attention and saliency

Here, we introduce our precision-modulated account of perception and action. A graphical illustration is provided in Figure 5.1. For this, we turn to attention and salient action selection which have their roots in biological processes relevant for acquiring task-relevant



information. Under an active inference account, this attention influences (posterior) state estimation and can be associated with increased precision of belief updating and gain control—described in Sec. 5.2.1. Furthermore, this is distinct from salience despite interdependent neuronal composition and computations.

Further alignment between the two constructs can be revealed by considering the temporal scheduling between movement (i.e., action) and perception for uncertainty resolution (Parr and Friston, 2019). We postulate that this perception-action coupling is best understood as a periodic fluctuation between minimising uncertainty and precision control. Subsequently, action is deployed to reduce uncertainty. Such an alignment specifies what stimulus is selected and under what level of precision it is processed. . Parr and Friston (2019) hypothesise that action alignment with precision is due to the eye structure that provides precise information in the fovea and requires the agent to foveate the most informative stimulus. We extend this by considering the periodic deployment of gain control with saccades (Hogendoorn, 2016; Fiebelkorn and Kastner, 2019; Nakayama and Motoyoshi, 2019; Tomassini et al., 2017; Benedetto and Morrone, 2017).

Accordingly, our formulation defines attention as precision control and salience as uncertainty minimisation supported by discrete sampling of visual information at a theta rhythm. This synchronises perception and action together in an oscillatory fashion (Hogendoorn, 2016). Importantly, a Bayesian formulation of this can be realised as precision manipulation over particular model parameters. We reserve further details for Sec. 5.4.

**Summary** Based upon our review, we propose a precision-modulated account of attention and salience, emphasising the diachronic realisation of action and perception. In the following sections, we investigate a realisation of this model for a robotic system.

## 5.4 Precision-based attention for Robotics

The previous section introduced a conceptual account to explain the computational mechanisms that undergird attention based on neuroscience findings. We focused on reclaiming saliency as an active process that relies on neural gain control, uncertainty minimisation and structured scheduling. Here, we describe how we can mathematically realise some of these mechanisms in the context of well-known challenges in robotics. Enabling robots with this type of attention may be crucial to filter the sensory signals and internal variables that are relevant to estimate the robot/world state and complete any task. More importantly, the active component of salience (i.e., behaviour) is essential to interact with the world—as argued in active perception approaches (Bajcsy et al., 2018).

We revisit the standard view of attention in robotics by introducing sensory precision (inverse variance) as the driving mechanism for modulating both perception and

action (Clark, 2013; Friston et al., 2011). Although saliency was originally described to underwrite behaviour, most models used in robotics, strongly biased by computer vision approaches, focus on computing the most relevant region of an image (Borji and Itti, 2012)—mainly computing human fixation maps—relegating action to a secondary process. Illustratively, state-of-the-art deep learning saliency models—as shown in the MIT saliency benchmark (Bylinskii et al., 2019)—do not have the action as an output. Conversely, the active perception approach properly defines the action as an essential process of active sensing to gather the relevant information. Our proposed model, based on precision modulated action and perception coupling (*i*) place attention as essential for state-estimation and system identification and (*ii*) and reclaims saliency as a driver for information-seeking behaviour, as proposed in early works (Tsotsos et al., 1995), but goes beyond human fixation maps for both improving the model of the environment (exploration) and solving the task (exploitation).

**Table 5.1:** Robotics applications and their precision realisations.

Task	Application	Precision manipulation	Section
Perception	State & input estimation	Noise precision modelling $\tilde{\Pi}$	5.4.2.2
	System Identification	Posterior parameter precision learning $\Pi^\theta$	5.4.2.3
	Exploration-exploitation in learning	Prior parameter precision modelling $P^\theta$	5.4.2.4
	Noise estimation	Noise precision learning $\tilde{\Pi}$	5.4.2.5
Action	Informative Path Planning (IPP)	Precision optimisation (of map)	5.4.3.2
Active perception	IPP with action-perception cycle	Precision modulation	5.4.3.3

In what follows, we highlight the key role of precision by reviewing relevant brain-inspired attention models deployed in robotics (Sec. 5.4.1). We propose precision-modulated attentional mechanisms for robots in three contexts - perception (Sec. 5.4.2), action (Sec. 5.4.3) and active perception (Sec. 5.4.3.3). The precision-modulated perception is formalised for a robotics setting; via (*i*) state estimation (i.e., estimating the hidden states of a dynamic system from sensory signals – Sec. 5.4.2.2), and (*ii*) system identification (i.e., estimating the parameters of the dynamic system from sensory signals – Sec. 5.4.2.3). Next, we show that precision-modulated action can be realised through precision optimisation (planning future actions – Sec. 5.4.3.2) and discuss practical considerations for coupling with precision-modulated perception (precision based active perception - Sec. 5.4.3.3). Table 5.1 summarises our proposed precision manipulations to solve relevant problems in robot perception and action. Table 5.2 provides the definitions of precision within our mechanism.

**Table 5.2:** Precision parameters that are manipulated in Sec.5.4.2

Term	Symbol	Definition
Sensory precision	$\Pi^z$	Inverse covariance of sensory noise $\mathbf{z}$ (Eqn. 5.1).
Prior parameter precision	$P^\theta$	The robot’s confidence on its prior parameters $\eta^\theta$ .
Noise precision	$\tilde{\Pi}$	The inverse covariance of all noises (Eqn. 5.5).
Posterior parameter precision	$\Pi^\theta$	The robot’s confidence on its parameter estimates.

### 5.4.1 Previous brain-inspired attention models in robotics

Brain-inspired attention has been mainly addressed in robotics from a ‘passive’ visual saliency perspective, e.g., which pixels of the image are the most relevant. This saliency map is then generally used to foveate the most salient region. This approach was strongly influenced by early computational models of visual attention (Tsotsos et al., 1995; Itti and Koch, 2001). The first models deployed in robots were bottom-up, where the sensory input was transformed into an array of values that represents the importance (or salience) of each cue. Thus, the robot was able to identify which region of the scene has to look at, independently of the task performed—see Borji and Itti (2012) for a review on visual saliency. These models have also been useful for acquiring meaningful visual features in applications, such as object recognition (Orabona et al., 2005; Frintrop, 2006), localisation, mapping and navigation (Frintrop and Jensfelt, 2008; Kim and Eustice, 2013; Roberts et al., 2012). Saliency computation was usually employed as a helper for the selection of the relevant characteristics of the environment to be encoded. Thus, reducing the information needed to process.

More refined methods of visual attention employed top-down modulation, where the context, task or goal bias the relevance of the visual input. These methods were used, for instance, to identify humans using motion patterns (Butko et al., 2008; Morén et al., 2008). A few works also focused on object/target search applications, where top-down and bottom-up saliency attention were used to find objects or people in a search and rescue scenario (Rasouli et al., 2020).

Attention has also been considered in human-robot interaction and social robotics applications (Ferreira and Dias, 2014), mainly for scene or task understanding (Ude et al., 2005; Kragic et al., 2005; Lanillos et al., 2016), and gaze estimation (Shon et al., 2005) and generation (Lanillos et al., 2015a). For instance, computing where the human is looking at and where the robot should look at or which object should be grasped. Furthermore, multi-sensory and 3D saliency computation has also been investigated (Lanillos et al., 2015b). Finally, more complex attention behaviours, particularly designed for social robotics and based on human non-verbal communication, such as joint attention, have

also been addressed. Here the robot and the human share the attention of one object through meaningful saccades, i.e., head/eye movements (Kaplan and Hafner, 2006; Nagai et al., 2003; Lanillos et al., 2015a).

Although attention mechanisms have been widely investigated in robotics, specially to model visual cognition (Begum and Karray, 2010; Kragic et al., 2005), the majority of the works have treated attention as an extra feature that can help the visual processing, instead of a crucial component needed for the proper functioning of the cognitive abilities of the robot (Lanillos and Cheng, 2018a). Furthermore, these methods had the tendency to leave the action generation out of the attention process. One of the reasons for not including saliency computation, in robotic systems, is that the majority of the models only output ‘human-fixation map’ predictions, given a static image. Saliency computation introduces extra computational complexity, which can be finessed by visual segmentation algorithms (e.g., line detectors in autonomous navigation). However, it does not resolve uncertainty nor select actions that maximise information gain in the future. In essence, the incomplete view of attention models that output human-fixation maps has arguably obscured the huge potential of neuroscience-inspired attentional mechanisms for robotics.

Our proposed model of attention, based on precision modulation, abandons the current robotics narrow view of attention and saliency by explicitly modelling attention within state estimation, learning and control. Thus, placing attentional processes at the core of the robot computation and not as an extra add-on. In the following sections, we describe the realisation of our precision-based attention formulation in robotics using common practical applications as the backbone motif.

### 5.4.2 Precision-modulated perception

We formalise precision-modulated perception from a first principles Bayesian perspective – explicitly the free energy principle approach proposed by Friston et al. (2011). Practically, this entails optimising precision parameters over (particular) model parameters.

Through numerical examples show how our model is able to perform accurate state estimation (Bos et al., 2021) and stable parameter learning (Meera and Wisse, 2021a,b). To illustrate the approach, we first introduce a dynamic system modelled as a linear state space system in robotics (Sec. 5.4.2.1)—we used this formulation in all our numerical experiments. We briefly review the formal terminologies for a robotics context to appropriately situate our precision-based mechanism for perception. Explicitly, we introduce: precision modelling (by adapting a known form of the precision matrix), precision learning (by learning the full precision matrix), and precision optimisation (use precision as an objective function during learning). As a reminder, precision modelling is associated with

(instantaneous) gain control and precision learning (at slower time scales) is associated with optimising that control.

#### 5.4.2.1 Precision for state space models

A linear dynamic system can be modelled using the following state space equations (boldface notation denotes components of the real system and non-boldface notation its estimates):

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{w}, \quad \mathbf{y} = \mathbf{C}\mathbf{x} + \mathbf{z}. \quad (5.1)$$

where  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{C}$  are constant matrices defining the system parameters,  $\mathbf{x} \in \mathbb{R}^n$  is the system state (usually an unobserved variable),  $\mathbf{u} \in \mathbb{R}^r$  is the input or control actions,  $\mathbf{y} \in \mathbb{R}^m$  is the output or the sensory measurements,  $\mathbf{w} \in \mathbb{R}^n$  is the process noise with precision  $\mathbf{\Pi}^{\mathbf{w}}$  (or inverse variance  $\mathbf{\Sigma}^{\mathbf{w}-1}$ ), and  $\mathbf{z} \in \mathbb{R}^m$  is the measurement noise with precision  $\mathbf{\Pi}^{\mathbf{z}}$ .

For instance, we can describe a mass-spring damper system (depicted in Fig. 5.2b) using state space equations. A mass ( $m = 1.4\text{kg}$ ) is attached to a spring with elasticity constant ( $k = 0.8\text{N/m}$ ), and a damper with a damping coefficient ( $b = 0.4\text{Ns/m}$ ). When a force ( $u(t) = e^{-0.25(t-12)^2}$ ) is applied on the mass, it displaces  $x$  from its equilibrium point. The linear dynamics of this system is given by:

$$\begin{bmatrix} \dot{x} \\ \ddot{x} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\frac{k}{m} & -\frac{b}{m} \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{m} \end{bmatrix} u, \quad y = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix}. \quad (5.2)$$

Note that Eq. (5.2) is equivalent to Eq. (5.1) with parameters  $\mathbf{A} = \begin{bmatrix} 0 & 1 \\ -\frac{k}{m} & -\frac{b}{m} \end{bmatrix}$ ,  $\mathbf{B} = \begin{bmatrix} 0 \\ \frac{1}{m} \end{bmatrix}^T$  and  $\mathbf{C} = \begin{bmatrix} 1 & 0 \end{bmatrix}$ , and state  $\mathbf{x} = \begin{bmatrix} x, \dot{x} \end{bmatrix}^T$ .

Now we introduce attention as precision modulation assuming that the robotic goal is to minimise the prediction error (Friston et al., 2011; Lanillos and Cheng, 2018b; Meera and Wisse, 2020), i.e., to refine its model of the environment and perform accurate state estimation, given the information available. In other words, the robot has to estimate  $\mathbf{x}$  and  $\mathbf{u}$  from input prior  $\eta^u$  with a prior precision of  $P^u$ , given the measurements  $\mathbf{y}$ , parameters  $\mathbf{A}$ ,  $\mathbf{B}$ ,  $\mathbf{C}$  and noise precision  $\mathbf{\Pi}^{\mathbf{w}}$  and  $\mathbf{\Pi}^{\mathbf{z}}$ . Formally, the prediction error  $\tilde{\epsilon}$  of

the sensory measurements  $\tilde{\epsilon}^y$ , control input reference  $\tilde{\epsilon}^u$  and state  $\tilde{\epsilon}^x$  are:

$$\tilde{\epsilon} = \begin{bmatrix} \tilde{\epsilon}^y \\ \tilde{\epsilon}^u \\ \tilde{\epsilon}^x \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{y}} - \tilde{\mathbf{C}}\tilde{x} \\ \tilde{u} - \tilde{\eta}^u \\ D^x\tilde{x} - \tilde{\mathbf{A}}\tilde{x} - \tilde{\mathbf{B}}\tilde{u} \end{bmatrix} \begin{cases} \text{sensory prediction error} \\ \text{control input prediction error} \\ \text{state prediction error} \end{cases} \quad (5.3)$$

Note that  $\tilde{\epsilon}^y = \tilde{\mathbf{y}} - \tilde{\mathbf{C}}\tilde{x}$  is the difference between the observed measurement and the predicted sensory input given the state<sup>3</sup>. Here  $D^x$  performs the (block) derivative operation, which is equivalent to shifting up all the components in generalised coordinates by one block.

We can estimate the state and input using the Dynamic Expectation Maximisation (DEM) algorithm (Friston et al., 2008; Meera and Wisse, 2020) that optimises a free energy variational bound  $\mathcal{F}$  to be tractable<sup>4</sup>. This is:

$$X = \begin{bmatrix} \tilde{x} \\ \tilde{u} \end{bmatrix} = \arg \max_X \mathcal{F} = \arg \max_X -\frac{1}{2}\tilde{\epsilon}^T \tilde{\Pi} \tilde{\epsilon} \quad (5.4)$$

Crucially,  $\tilde{\Pi}$  is the generalised noise precision that modulates the contribution of each prediction error to the estimation of the state and the computation of the action. Thus,  $\tilde{\Pi}$  is equivalent to attentional gain. For instance, we can model the precision matrix to attend to the most informative signal derivatives in  $\tilde{\mathbf{y}}$ . Concisely, the precision  $\tilde{\Pi}$  has the following form:

$$\tilde{\Pi} = \begin{bmatrix} S \otimes \Pi^z & 0 & 0 \\ 0 & S \otimes P^u & 0 \\ 0 & 0 & S \otimes \Pi^w \end{bmatrix}, \quad (5.5)$$

where  $S$  is the smoothness matrix. In Sec. 5.4.2.2, we show that modelling the precision matrix  $\tilde{\Pi}$  using the  $S$  matrix improves the estimation quality.

The full free energy functional (time integral of free energy  $\bar{\mathcal{F}} = \int \mathcal{F} dt$  at optimal precision) that the robot optimises to perform state-estimation and system identification

---

<sup>3</sup>The tilde over the variable refers to the generalised coordinates, i.e., the variable includes all temporal derivatives. Thus,  $\tilde{\epsilon}$  is the combined prediction error of outputs, inputs and states. For example, the generalised output  $\tilde{y}$  is given by  $\tilde{\mathbf{y}} = [y, y', y'' \dots]^T$ , where the prime operator denotes the derivatives. We use generalised coordinates (Friston et al., 2010) for achieving accurate state and input estimation during the presence of (coloured) noise by modelling the time dependent quantities  $(x, v, y, w, z)$  in generalised coordinates. This involves keeping track of the evolution of the trajectory of the probability distributions of states, instead of just their point estimates. Here the coloured noise  $w$  and  $z$  are modelled as a white noise convoluted with a Gaussian kernel. The use of generalised coordinates has recently shown to outperform classical approaches under coloured noise on real quadrotor flight (Bos et al., 2021)

<sup>4</sup>Note that this expression of the variational free energy is using the Laplace and mean-field approximations commonly used in the FEP literature

is described in Eq. (5.6)—for readability we omitted the details of the derivation of this cost function, and we refer to (Anil Meera and Wisse, 2021) for further details.

$$\begin{aligned}
\bar{\mathcal{F}} = & -\frac{1}{2} \sum_t \left[ \underbrace{\tilde{\epsilon}^{yT} \tilde{\Pi}^z \tilde{\epsilon}^y + \tilde{\epsilon}^{uT} P^{\tilde{u}} \tilde{\epsilon}^u + \tilde{\epsilon}^{xT} \tilde{\Pi}^w \tilde{\epsilon}^x}_{\text{precision weighed prediction error}} \right] \\
& - \frac{1}{2} \left[ \underbrace{\epsilon^{\theta T} P^\theta \epsilon^\theta + \epsilon^{\lambda T} P^\lambda \epsilon^\lambda}_{\text{prior precision weighed prediction error of } \theta \text{ and } \lambda} \right] \\
& + \underbrace{\frac{1}{2} n_t \ln |\Sigma^X|}_{\text{state and input entropy}} + \underbrace{\frac{1}{2} n^t [\ln |\tilde{\Pi}^z| + \ln |P^{\tilde{v}}| + \ln |\tilde{\Pi}^w|]}_{\text{noise entropy}} \\
& + \underbrace{\frac{1}{2} \ln |\Sigma^\theta P^\theta|}_{\text{parameter entropy}} + \underbrace{\frac{1}{2} \ln |\Sigma^\lambda P^\lambda|}_{\text{hyperparameter entropy}}
\end{aligned} \tag{5.6}$$

Here  $\epsilon^\theta = \theta - \eta^\theta$ ,  $\epsilon^\lambda = \lambda - \eta^\lambda$  are the prediction errors of parameters and hyperparameters<sup>5</sup>.  $\bar{\mathcal{F}}$  consist of two main components: i) precision weighed prediction errors and ii) precision-based entropy. The dominant role of precision — in the free energy objective — is reflected in how modulating these precision parameters can have a profound influence perception and behaviour. The theoretical guarantees for stable estimation (Meera and Wisse, 2021b), and its application on real robots (Lanillos et al., 2021) make this formulation very appealing to robotic systems.

Note that we can manipulate three kinds of precision within the state space formulation: i) prior precision ( $P^{\tilde{u}}, P^\theta, P^\lambda$ ), ii) conditional precision on estimates ( $\Pi^X, \Pi^\theta, \Pi^\lambda$ ) and iii) noise precision ( $\Pi^z, \Pi^w$ ). Therefore, to learn the correct parameter values  $\theta$ , we i) learn the parameter precision  $\Pi^\theta$ , ii) model the prior parameter precision  $P^\theta$ , and iii) learn the noise precision  $\Pi^w$  and  $\Pi^z$  (parameterised using  $\lambda$ ).


#### 5.4.2.2 State and input estimation

State estimation is the process of estimating the unobserved states of a real system from (noisy) measurements. Here, we show how we can achieve accurate estimation through precision modulation in a linear time invariant system under the influence of coloured noise (Meera and Wisse, 2020). State estimation in the presence of coloured noise is inherently challenging, owing to the non-white nature of the noise, which is often ignored

<sup>5</sup>System identification involves the estimation of system parameters (denoted by  $\theta$ , e.g., vectorised  $\mathbf{A}$ ), given  $\mathbf{y}, \mathbf{u}$ , by starting from a parameter prior of  $\eta^\theta$  with prior precision  $P^\theta$ , and a prior on noise hyper-parameter  $\eta^\lambda$  with a prior precision of  $P^\lambda$ . Note that we parametrise noise precision ( $\Pi^w$  and  $\Pi^z$ ) using  $\lambda \in \mathbb{R}^{2 \times 1} = \begin{bmatrix} \lambda^z \\ \lambda^w \end{bmatrix}$  as an exponential relation (e.g.,  $\Pi^w(\lambda^w) = \exp(\lambda^w) I^{n \times n}$ ).

in conventional approaches, such as the Kalman Filter (Welch et al., 1995).

Figure 5.2 summarises a numerical example that shows how one can use precision modulation to focus on the less noisy derivatives (lower derivatives) of measurements, relative to imprecise higher derivatives. Thus, enabling the robot to use the most informative data for state and input estimation, while discarding imprecise input. Figure 5.2b depicts the mass-spring damper system used. The numerical results show that the quality of the estimation increases as the embedding ordering increases but the lack of information in the higher order derivatives of the sensory input do not affect the final performance due to the precision modulation. The higher order derivatives (Fig. 5.2a) are less precise than the lower derivatives, thereby reflecting the loss of information in higher derivatives. The state and input estimation was performed using the optimisation framework described in the previous section. The quality of estimation is shown in Fig. 5.2c, where the input estimation using six derivatives (blue curve) is closer to the real input (yellow curve) than when compared to the estimation using only one derivative (red curve). The quality of the estimation reports the sum of squared error (SSE) in the estimation of states and inputs with respect to the embedding order (number of signal derivatives considered).



**Figure 5.2:** An illustration of an attention mechanism in state and input estimation. The quality of the estimation improves as the embedding order (number of derivatives) of generalised coordinates are increased. However, the imprecise information in the higher order derivatives of the sensory input  $\mathbf{y}$  does not affect the final performance of the observer because of attentional selection, which selectively weighs the importance afforded to each derivative, in the free energy optimisation scheme.

To obtain accurate state estimation by optimising the precision parameters, we recall



that the precision weights the prediction errors. From Eq. (5.3), the structural form of  $\tilde{\Pi}$  is mainly dictated by the smoothness matrix  $S$ , which establishes the interdependence between the components of the variable expressed in generalised coordinates (e.g., the dependence between  $\mathbf{y}$ ,  $\mathbf{y}'$  and  $\mathbf{y}''$  in  $\tilde{\mathbf{y}}$ ). For instance, the  $S$  matrix for a Gaussian kernel is as follows (Meera and Wisse, 2022):

$$S = \begin{bmatrix} \frac{35}{16} & 0 & \frac{35}{8}s^2 & 0 & \frac{7}{4}s^4 & 0 & \frac{1}{6}s^6 \\ 0 & \frac{35}{4}s^2 & 0 & 7s^4 & 0 & s^6 & 0 \\ \frac{35}{8}s^2 & 0 & \frac{77}{4}s^4 & 0 & \frac{19}{2}s^6 & 0 & s^8 \\ 0 & 7s^4 & 0 & 8s^6 & 0 & \frac{4}{3}s^8 & 0 \\ \frac{7}{4}s^4 & 0 & \frac{19}{2}s^6 & 0 & \frac{17}{3}s^8 & 0 & \frac{2}{3}s^{10} \\ 0 & s^6 & 0 & \frac{4}{3}s^8 & 0 & \frac{4}{15}s^{10} & 0 \\ \frac{1}{6}s^6 & 0 & s^8 & 0 & \frac{2}{3}s^{10} & 0 & \frac{4}{45}s^{12} \end{bmatrix}, \quad (5.7)$$


where  $s$  is the kernel width of the Gaussian filter that is assumed to be responsible for serial correlations in measurement or state noise. Here, the order of generalised coordinates (number of derivatives under consideration) is taken as six ( $S \in \mathbb{R}^{7 \times 7}$ ). For practical robotics applications, the measurement frequency is high, resulting in  $0 < s < 1$ . It can be observed that the diagonal elements of  $S$  decreases because  $s < 1$ , resulting in a higher attention (or weighting) on the prediction errors from the lower derivatives when compared to the higher derivatives. The higher the noise colour (i.e.,  $s$  increases), the higher the weight given to the higher state derivatives (last diagonal elements of  $S$  increases). This reflects the fact that smooth fluctuations have more information content in their higher derivatives. Having established the potential importance of precision weighting in state estimation, we now turn to the estimation (i.e., learning) of precision in any given context.

### 5.4.2.3 System identification

This section shows how to optimise system identification by means of precision learning (Anil Meera and Wisse, 2021; Meera and Wisse, 2021b). Specifically, we show how to fuse prior knowledge about the dynamic model with the data to recover unknown parameters of the system through an attention mechanism. This involves the learning of the 1) parameters and 2) noise precisions. Our model ‘turns’ the attention to the least precise parameters and uses the data to update those parameters to increase their precision. Hence, allowing faster parameter learning.

For the sake of clarity, we use again the mass-spring-damper system as the driving example (Sec. 5.4.2.1). We formalise system identification as evaluating the unknown

parameters  $k$ ,  $m$  and  $b$ , given the input  $\mathbf{u}$ , the output  $\mathbf{y}$ , and the general form of the linear system in Eq. (5.2).



**Figure 5.3:** The schematic of the robot's attention mechanism for learning the least precise parameters of a given generative model of a mass-spring-damper system.

Figure 5.3 depicts the process of learning unknown parameters (dotted boxes denote the processes inside the robot brain). The robot measures its position  $x(t)$  using its sensors (e.g., vision or range sensor). We assume that the robot has observed the behaviour of a mass-spring-damper system before or a model is provided by the expert designer. However, some of the parameters are unknown. The robot can reuse the prior learned model of the system to relearn the new system. This can be realised by setting a high prior precision on the known parameters and a low prior precision on the unknown parameters. By means of precision learning, the robot uses the sensory signals to learn the parameter precision  $\Pi^\theta$ , thereby improving the confidence in the parameter estimates  $\theta$ . This directs the robot's attention towards the refinement of the parameters with least precision as they are the most uncertain. The requisite parameter learning proceeds by the gradient ascent of the free energy functional given in Eq. (5.6). The parameter precision learning proceeds by tracking the negative curvature of  $\bar{\mathcal{F}}$  as  $\Pi^\theta = -\frac{\partial^2 \bar{\mathcal{F}}}{\partial \theta^2}$  (Anil Meera and Wisse, 2021).

The learning process – by means of variational free energy optimisation (maximisation) – is shown in Fig. 5.3b. The learning involves two parallel processes: precision learning (Fig. 5.3a), and parameter learning (Fig. 5.3c). Precision learning comprises of parameter precision learning (top graph) – i.e., identifying the precision of an approximate posterior

density for the parameters being estimated – and noise precision learning (bottom graph). The high prior precision on the known system parameters (0 and 1), and low prior precision on the unknown system parameters ( $-\frac{k}{m}$ ,  $-\frac{b}{m}$  and  $\frac{1}{m}$ , highlighted in blue) directs attention towards learning the unknown parameters and their precision. Note that in Fig. 5.3a, the precision on the three unknown parameters start from a low prior precision of  $P^\theta = 1$  and increase with each iteration, whereas the precision of known parameters (0 and 1) remains a constant ( $3.3 \times 10^6$ ). The noise precisions are learned simultaneously, which starts from a low prior precision of  $P^{\lambda^w} = P^{\lambda^z} = 1$  and finally converges to the true noise precision (dotted black line). Both precisions are used to learn the three parameters of the system (Fig. 5.3b), which starts from randomly selected values within the range  $[-2, 2]$  and finally converges to the true parameter values of the system ( $\theta_3 = -\frac{k}{m} = -0.5714$ ,  $\theta_4 = -\frac{b}{m} = -0.2857$  and  $\theta_6 = \frac{1}{m} = 0.7143$ ), denoted by black dotted lines. From an attentional perspective, the lower plot in (Fig. 5.3a) is particularly significant here. This is because the robot discovers the data are more informative than initially assumed, thereby leading to an increase in its estimate of the precision of the data-generating process. This means that the robot is not only using the data to optimise its beliefs about states and parameters (system identification), it is also using these data to optimise the way in which it assimilates these data.

In summary, precision-based attention, in the form of precision learning, helps the robot to accurately learn unknown parameters by fusing prior knowledge with new incoming data (sensory measurements), and attending to the least precise parameters.


#### 5.4.2.4 Precision-modulated exploration and exploitation in system identification

Exploration and exploitation in the parameter space can be advantageous to robots during system identification. Precision-based attention—here the prior precision—allows a graceful balance between the two, mediated by the prior precision<sup>6</sup>. A very high prior precision encourages exploitation and biases the robot towards believing its priors, while a low prior precision encourages exploration and makes the robot sensitive to new information.

We use again the mass-spring-damper system example but with a different prior parameter precision  $P^\theta$ . The prior parameters are initialised at random and learned using optimisation. Figure 5.4b shows the increase in parameter estimation error (SSE) as the prior parameter precision  $P^\theta$  increases until it finally saturates. The bottom left

---

<sup>6</sup>Note that here we are using exploration and exploitation not in terms of behaviour but for parameter learning. Exploration means adapting the parameter to a different (unexplored) value and exploitation means keeping that value



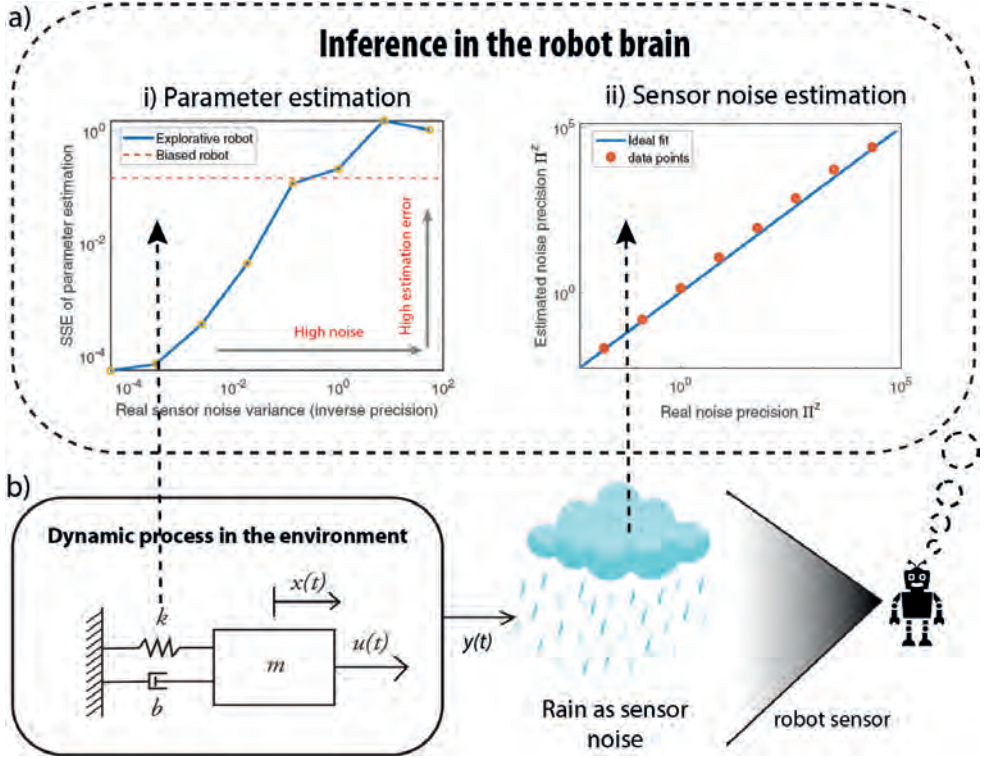
**Figure 5.4:** Precision-based attention allows exploration and exploitation balanced model learning mediated by the prior precisions on the parameters  $P^\theta$ . The higher the  $P^\theta$ , the higher the attention on prior parameters  $\eta^\theta$  and the lower the attention on the sensory signals while learning.

region (circled in red) indicates the region where the prior precision is low, encouraging exploration with high attention on the sensory signals for learning the model. This region over-exposes the robot to its sensory signals by neglecting the prior parameters. The top right region (circled in red) indicates the biased robot where the prior precision is high, encouraging the robot to exploit its prior beliefs by retaining high attention on prior parameters. This regime biases the robot into being confident about its priors and disregarding new information from the sensory signals. Between those extreme regimes (blue curve) the prior precision balances the exploration-exploitation trade-off. Figure 5.4a describes how increased attention to sensory signals helped the robot to recover from poor initial estimates of parameter values and converge towards the correct values (dotted black line). Conversely, in Fig. 5.4c, high attention on prior parameters did not help the robot to learn the correct parameter values.

These results establish that prior precision modelling allows balanced exploration and exploitation of parameter space during system identification. Although the results show that an over-exposed robot provides better parameter learning, we show – in the next section – that this is not always be the case.

#### 5.4.2.5 Noise estimation

In real-world applications, sensory measurements are often highly noisy and unpredictable. Furthermore, the robot does not have access to the noise levels. Thus, it needs to learn the noise precision ( $\Pi^z$ ) for accurate estimation and robust control. Precision-based attention enables this learning. In what follows, we show how one can estimate  $\Pi^z$  using



**Figure 5.5:** Simulations demonstrating how a biased robot could be advantageous, especially while learning in a highly noisy environment. As the sensor noise increases, the quality of parameter estimation deteriorates to a point where an explorative robot generates higher parameter estimation errors than when compared to the biased robot that relies on its prior parameters. However, the sensor noise estimation is accurate even for high noise environments, demonstrating the success of the attention mechanism using the noise precision learning.

noise precision learning and that biasing the robot to prior beliefs can be advantageous in highly noisy environments.

Consider again the mass-spring-damper system in Figure 5.5b, where heavy rain-fall/snow corrupts visual sensory signals. We evaluate the parameter estimation error under different noise conditions, using different levels of noise variances (inverse precision). For an over-exposed robot (only attending to sensory measurements), left plot of Fig. 5.5a, the estimation error increases as the noise strength increases, to a point where the error surpasses the error from a prior-biased robot. This shows that a robot, confident in its prior model, assigns low attention to sensory signals and outperforms an over-exposed robot that assigns high attention to sensory signals, in a highly noisy environment. The

right plot of Fig. 5.5a shows the quality of noise precision learning for an over-exposed robot. It can be seen that all the data points in red lie close to the blue line, indicating that the estimated noise precision is close to the real noise precision. Therefore, the robot is capable of recovering the correct sensory noise levels even when the environment is extremely noisy, where accurate parameter estimation is difficult.

These numerical results show that attention mechanism — by means of noise precision learning — allows the estimation of the noise levels in the environment and thereby protects against over-fitting or overconfident parameter estimation.

**Summary.** We have shown how precision-based attention—through precision modelling and learning— yields to accurate robot state estimation, parameter identification and sensory noise estimation. In the next section, we discuss how action is generated in this framework.


### 5.4.3 Precision-modulated action

Selecting the optimal sequence of actions to fulfil a task is essential for robotics (LaValle, 2006). One of the most prominent challenges is to ensure robust behaviour given the uncertainty emerging from a highly complex and dynamic real world, where the robots have to operate on. A proper attention system should provide action plans that resolve uncertainty and maximise information gain. For instance, it may minimise the information entropy, thereby encouraging repeated sensory measurements (observations) on high uncertainty sensory information.

Saliency, which in neuroscience is sometimes identified as Bayesian surprise (i.e., divergence between prior and posterior), describes which information is relevant to process. We go one step further by defining the saliency map as the epistemic value of a particular action Friston et al. (2015b). Thus, the (expected) divergence now becomes the mutual information under a particular action or plan. This makes the saliency map more sophisticated because it is an explicit measure of the reduction in uncertainty or mutual information associated with a particular action (i.e., active sampling), and more pragmatic because it tells you where to sample data next, given current Bayesian beliefs.

We first describe a precision representation usually used in information gathering problems and then how to directly generate action plans through precision optimisation. Afterwards, we discuss the realisation of the full-fledged model presented in the neuroscience section for active perception. We use the informative path planning (IPP) problem, described in Fig. 5.6, as an illustrative example to drive intuitions.






**Figure 5.6:** IPP problem for localising human victims in an urban search and rescue scenario (Meera et al., 2019). A UAV, in a realistic simulation environment, plans a finite look-ahead path to minimise the uncertainty of its human occupancy map (e.g., modelled as a Gaussian process) of the world. The planned path is then executed, during which the UAV flies and captures images at a constant measurement frequency. After the data acquisition is complete, a human detection algorithm is executed to detect all the humans on the images. These detections are then fused into the UAV’s human location map. The cycle is repeated until the uncertainty of the map is completely resolved (this usually implies enough area coverage and repeated measurements on uncertain locations). The ground truth of the human occupancy map and the UAV belief is shown in (c) and (b) respectively. The final map approaches the ground truth and all the seven humans on the ground are correctly detected.


#### 5.4.3.1 Precision maps as saliency

One of the popular approaches in information gathering problems is to model the information map as a distribution (e.g., using Gaussian processes (Hitz et al., 2017)). This is widely used in applications, such as a target search, coverage and navigation. The robot keeps track of an occupancy map and the associated uncertainty map (covariance matrix or inverse precision). While the occupancy map records the presence of the target on the map, the uncertainty map records the quality of those observations. The goal of the robot is to learn the distribution using some learning algorithm (Marchant and Ramos, 2014). A popular strategy is to plan the robot path such that it minimises the uncertainty of the map in future (Popović et al., 2017). In Sec. 5.4.3.2, we will show how we can use the map precision to perform active perception, i.e., optimise the robot path for maximal information gain. Optimising the map precision drives the robot towards an exploratory behaviour.

### 5.4.3.2 Precision optimisation for action planning



**Figure 5.7:** Finding humans with unmanned air vehicles (UAVs): an informative path planning (IPP) approach (Anil Meera, 2018). The simulation environment on the left consists of a tall building at the centre, surrounded by seven humans lying on the floor. The goal of the UAV is to compute the action sequence that allows maximum information gathering, i.e., the humans location uncertainty is minimised. On the right is the final occupancy map coloured with the probability of finding a human at that location. It can be observed that all humans on the simulation environment were correctly detected by the robot.




**Figure 5.8:** Variance map of the probability distribution of people location (Fig. 5.7) – inverse precision of human occupancy map. The plot sequence shows the reduction of map uncertainty (inverse precision) after measurements (Anil Meera, 2018).

To introduce precision-based saliency we use an exemplary application of search and rescue. The goal is to find all humans using an unmanned air vehicle (UAV) (Lanillos, 2013; Lanillos et al., 2014; Rasouli et al., 2020; Meera et al., 2019). We use precision



for two purposes: i) precision optimisation for action planning (plan flight path) and ii) precision learning for map refinement. In contrast to previous models of action selection within active inference in robotics (Oliver et al., 2021; Lanillos et al., 2021) here precision explicitly drives the agent behaviour. Figure 5.7 describes the scenario in simulation. The seven human targets on the ground are correctly identified by the UAV. We can formalise the solution as the UAV actions (next flight path) that minimise the future uncertainties of the human occupancy map. In our precision-based attention scheme, this objective is equivalent to maximising the posterior precision of the map. Figure 5.8 shows the reduction in map uncertainty after subsequent assimilation of the measurements (camera images from the UAV, processed by a human detector). The map (and precision) is learned using a recursive Kalman Filter by fusing the human detector outcome onto the map (and precision). The algorithm drives the UAV towards the least explored regions in the environment, defined by the precision map.



**Figure 5.9:** The human occupancy map (probability to find humans at every location of the environment) at four time instances during the UAV flight showing ambiguity resolution. The ambiguity arising from imprecise sensor measurements (false positive) is resolved through repeated measurements at the same location. The plot sequence shows how the assimilation of the measurements updates the probability of the people being in each location of the map (Meera et al., 2019).

Furthermore, Fig. 5.9 shows an example of uncertainty resolution under false positives. In this case, human targets are moved to the bottom half of the map. The first measurement provides a wrong human detection with high uncertainty. However, after repeated measurements at the same location in the map the algorithm was capable of resolving this ambiguity, to finally learn the correct ground truth map. Hence, the sought behaviour is to take actions that encourage repeated measurements at uncertain locations for reducing uncertainty.

Although the IPP example illustrates how to generate control actions through precision optimisation, the task, by construction, is constrained to explicitly reduce uncertainty. This is similar to the description of visual search described in (Friston et al., 2012), where the location was chosen to maximise information gain. Information gain (i.e., the Bayesian surprise expected following an action) is a key part of the expected free energy functional


that underwrite action selection in active inference. In brief, expected free energy can be decomposed into two parts the first corresponds to the information gain above (a.k.a., epistemic value or affordance). The second corresponds to the expected log evidence or marginal likelihood of sensory samples (a.k.a., pragmatic value). When this likelihood is read as a prior preference, it contextualises the imperative to reduce uncertainty by including a goal-directed, imperative. For example, in the search paradigm above, we could have formulated the problem in terms of reducing uncertainty about whether each location was occupied by a human or not. We could have then equipped the agent with prior preferences for observing humans.

In principle, this would have produced searching behaviour until uncertainty had been resolved about the scene; after which, the robot would seek out humans; simply because, these are its preferred outcomes. In thinking about how this kind of neuroscience inspired or biomimetic approach could be implemented in robotics, one has to consider carefully, the precision afforded sensory inputs (i.e., the likelihood of sensory data, given its latent causes) – and how this changes during robotic flight and periods of data gathering. This brings us back to the precision modulation and the temporal scheduling of searching and securing data. In the final section, we conclude with a brief discussion of how this might be implemented in future applications.

### 5.4.3.3 Precision-based active perception

In this section, we discuss the realisation of a biomimetic brain-inspired model in relation to existing solutions in robotics in the context of path-planning. Figure 5.10 compares our proposed precision-modulated attention model—from Fig. 5.1—with the action-perception loop widely used in robotics. By analogy with eye saccades to the next visual sample, the UAV flies (action) over the environment to assimilate sensory data for an informed scene construction (perception). Once the flight time of the UAV is exhausted (similar to saccade window of the eye), the action is complete, after which the map is updated, and the next flight path is planned.

In standard applications of active inference, the information gain is supplemented with expected log preferences to provide a complete expected free energy functional (Sajid et al., 2021a). This accommodates the two kinds of uncertainty that actions and choices typically reduce. The first kind of uncertainty is inherent in unknowns in the environment. This is the information gain we have focused on above. The second kind of uncertainty corresponds to expected surprise, where surprise rests upon a priori expected or preferred outcomes. As noted above, equipping robots with both epistemic and pragmatic aspects to their action selection or planning could produce realistic and useful behaviour that automatically resolves the exploration-exploitation dilemma. This follows because the



**Figure 5.10:** Precision-modulated attention model adapted to the action-perception loop in robotics. Each cycle consists of two steps: 1) action (planning and execution of a finite-time look ahead of the robot path for data collection) and 2) perception (learning using the collected data). This scheduling, using a finite time look-ahead plan, is quite common in real applications and of particular importance when processing is computationally expensive, e.g., slow rate of classification, non-scalable data fusion algorithms, Exponential planners, etc. However, the benefits of incorporating ‘optimal’ scheduled loop driven by precision should be further studied.

expected free energy contains the optical mixture of epistemic (information-seeking) and pragmatic (i.e., preference seeking) components. Usually, after a period of exploration, the preference seeking components predominate because uncertainty has been resolved. Although expected free energy provides a fairly universal objective function for sentient behaviour, it does not specify how to deploy behaviour and sensory processing optimally. This brings us to the precision modulation model, inspired by neuroscientific considerations of attention and salience.

Hence, there are key differences between biological and robotic implementations of the search behaviour. First, the use of oscillatory precision to modulate visual sampling and movement cycles, as opposed to arbitrary discrete action and perception steps currently used in robotics. Second, precision modulation influences both state estimation and action following the same uncertainty reduction principle. Importantly, our salience formulation speaks to selecting future data that reduces this uncertainty. For instance, we have shown—in the information gathering IPP example described in the previous subsection—that by optimising precision we also optimise behaviour.

We argue the potential need and the advantages of realising precision based temporal scheduling, as described the our brain-inspired model, for two practically relevant test cases: (i) learning dynamic models and (ii) information seeking applications.

In Section 5.4.2.4, we have shown how the exploration-exploitation trade-off can be mediated by the prior parameter precision during learning. However, the accuracy-precision curve (Fig. 5.4b) is often practically unavailable due to unknown true parameters values, challenging the modelling of prior precision. An alternative would be to use a precision based temporal scheduling mechanism to alternate between exploration and exploitation by means of a varying  $P^\theta$  (similar to Fig. 5.10) during learning, such that system identification is neither biased nor over exposed to sensory measurements. In Fig. 5.5a, we showed how noise levels influence estimation accuracy, and how biasing the robot by modelling  $P^\theta$  can be beneficial for highly noisy environments. A precision based temporal scheduling mechanism by means of a varying  $P^\theta$  could provide a balanced solution between a biased robot (that exploits its model) and an exploratory one.

Furthermore, temporal scheduling, in the same way that eye saccades are generated, can be adapted for information gathering applications, such as target search, simultaneous localization and mapping, environment monitoring, etc. For instance, introducing precision-modulation scheduling for solving the IPP, and scheduling perception (map learning) and action (UAV flight). Precision modulation will switch between action and perception: when the precision is high, perception occurs (c.f., visual sampling), and when the precision is low, action occurs (c.f., eye movements). This switch, which is often implemented in the robotics literature using a budget for flight time, will be now dictated by precision dynamics.

In short, we have sketched the basis for a future realisation of precision-based active perception, where the robot computes the actions to minimise the expected uncertainty. While most attentional mechanisms in robotics are limited to providing a ‘saliency’ map highlighting the most relevant features, our attention mechanism proposes a general scheduling mechanism with action in the loop with perception, both driven by precision.

## 5.5 Concluding remarks

We have considered attention and salience as two distinct processes that rest upon oscillatory precision control processes. Accordingly, they require particular temporal considerations: attention to reliably estimate latent states from current sensory data and salience for uncertainty reduction regarding future data samples. This formulation addresses visual search from a first principles (Bayesian) account of how these mechanisms might manifest – and the circular causality that undergirds them via a rhythmic theta-coupling. Crucially, we have revisited the definition of salience from the visual neurosciences; where it is read as Bayesian surprise (i.e., the Kullback Leibler divergence between prior and posterior beliefs). We took this one step further and defined salience as

the expected Bayesian surprise (i.e., epistemic value) of a particular action (e.g., sampling this set of data) (Friston et al., 2017b; Sajid et al., 2021a). Formulating salience as the expected divergence renders it the mutual information under a particular action (or action trajectory) (Friston et al., 2021), – and highlights its role in encoding working memory (Parr and Friston, 2017b). For brevity, our narrative was centred around visual attention and its realisation via eye movements. However, this model does not strictly need to be limited to visual information processing, because it addresses sensorimotor and auditory processing in general. This means it explains how action and perception can be coupled in other sensory modalities. For instance, (Tomassini et al., 2017) showed that visual information is coupled with finger movements at a theta rhythm.

The point of contact with the robotics use of salience emerges because the co-variation between a particular parameterisation and the inputs is a measure of the mutual information between the data and its estimated causes. In this sense, both definitions of salience reflect the mutual information – or information about a particular representation of a (latent) cause – afforded by an observation or consequence. However, our formulation is more sophisticated. Briefly, because it is an explicit measure of the reduction in uncertainty (i.e., mutual information) associated with a particular action (i.e., active sampling) and specifies where to sample data next, given current Bayesian beliefs. These processes (attention and salience) are a consequence of precision of beliefs over distinct model parameters. Explicitly, attention contends with precision over the causes of (current) outcomes and salience contends with beliefs about the data that has to be acquired and precision over beliefs about actions that dictate it. Since both processes can be linked via precision manipulation, the crucial thing is the precision that differentiates whether the agent acquires new information (under high precision) or resolves uncertainty by moving (low precision).

The focus of this work has been to illustrate the importance of optimising precision at various places in generative models used for data assimilation, system identification and active sensing. A key point – implicit in these demonstrations – rests upon the mean field approximation used in all applications. Crucially, this means that getting the precision right matters, because updating posterior estimates of states, parameters and precisions all depend upon each other. This may be particularly prescient for making the most sense of samples that maximises information gain. In other words, although attention and salience are separable optimisation processes, they depend upon each other during active sensing. This was the focus of our final numerical studies of action planning.

To face-validate our formulation, we evaluated precision-modulated attentional processes in the robotic domain. We presented numerical examples to show how precision manipulation underwrites accurate state and noise estimation (e.g., selecting relevant

information), as well as allowing system identification (e.g., learning unknown parameters of the dynamics). We also showed how one can use precision-based optimisation to solve interesting problems; like the informative path planning in search and rescue scenarios. Thus, in contrast to previous uses of attention in robotics, we placed attention and saliency as integral processes for efficient gathering and processing of sensory information. Accordingly, ‘attention’ is not only about filtering the current flow of information from the sensors but performing those actions that minimise expected uncertainty. Still, the full potential of our proposal has yet to be realised, as the precision-based attention should be able to account for prior preferences beyond the IPP problem (e.g., localising people using UAVs). Finally, we briefly considered the realisation of temporal scheduling for information gathering tasks, opening up interesting lines of research to provide robots with biologically plausible attention.

## Author Contributions

AAM and FN are responsible for the novel account and its translation to robotics. All authors contributed to conception and design of the work. AAM, FN, PL and NS wrote the manuscript. All authors contributed to manuscript revision, read, and approved the submitted version.

## Funding

The open access publication of the manuscript is funded by the TU Delft Library. F.N is funded by the Serotonin & Beyond project (953327). P.L. is partially supported by Spikeference project, Human Brain Project Specific Grant Agreement 3 (ID: 945539). N.S. is funded by the Medical Research Council (MR/S502522/1) and 2021–2022 Microsoft PhD Fellowship. K.F. is supported by funding for the Wellcome Centre for Human Neuroimaging (Ref: 205103/Z/16/Z) and a Canada-UK Artificial Intelligence Initiative (Ref: ES/T01279X/1).

# Bibliography

- Ahnelt, P. (1998). The photoreceptor mosaic. *Eye*, 12(3):531–540.
- Anil Meera, A. (2018). Informative path planning for search and rescue using a uav. *Masters thesis, TU Delft*.
- Anil Meera, A. and Wisse, M. (2021). Dynamic expectation maximization algorithm for estimation of linear systems with colored noise. *Entropy*, 23(10):1306.
- Atrey, A., Clary, K., and Jensen, D. (2019). Exploratory not explanatory: Counterfactual analysis of saliency maps for deep reinforcement learning. *arXiv preprint arXiv:1912.05743*.
- Axmacher, N., Henseler, M. M., Jensen, O., Weinreich, I., Elger, C. E., and Fell, J. (2010). Cross-frequency coupling supports multi-item working memory in the human hippocampus. *Proceedings of the National Academy of Sciences*, 107(7):3228–3233.
- Bajcsy, R., Aloimonos, Y., and Tsotsos, J. K. (2018). Revisiting active perception. *Autonomous Robots*, 42(2):177–196.
- Balestrieri, E., Ronconi, L., and Melcher, D. (2021). Shared resources between visual attention and visual working memory are allocated through rhythmic sampling. *European Journal of Neuroscience*.
- Begum, M. and Karray, F. (2010). Visual attention for robotic cognition: A survey. *IEEE Transactions on Autonomous Mental Development*, 3(1):92–105.
- Benedetto, A. and Morrone, M. C. (2017). Saccadic suppression is embedded within extended oscillatory modulation of sensitivity. *Journal of Neuroscience*, 37(13):3661–3670.

- Benedetto, A., Morrone, M. C., and Tomassini, A. (2020). The common rhythm of action and perception. *Journal of cognitive neuroscience*, 32(2):187–200.
- Borji, A. and Itti, L. (2012). State-of-the-art in visual attention modeling. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):185–207.
- Bos, F., Meera, A. A., Benders, D., and Wisse, M. (2021). Free energy principle for state and input estimation of a quadcopter flying in wind. *arXiv preprint arXiv:2109.12052*.
- Brown, H., Adams, R. A., Parees, I., Edwards, M., and Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive processing*, 14(4):411–427.
- Brzezicka, A., Kamiński, J., Reed, C. M., Chung, J. M., Mamelak, A. N., and Rutishauser, U. (2019). Working memory load-related theta power decreases in dorsolateral prefrontal cortex predict individual differences in performance. *Journal of cognitive neuroscience*, 31(9):1290–1307.
- Busch, N. A. and VanRullen, R. (2010). Spontaneous eeg oscillations reveal periodic sampling of visual attention. *Proceedings of the National Academy of Sciences*, 107(37):16048–16053.
- Butko, N. J., Zhang, L., Cottrell, G. W., and Movellan, J. R. (2008). Visual saliency model for robot cameras. In *2008 IEEE International Conference on Robotics and Automation*, pages 2398–2403. IEEE.
- Bylinskii, Z., Judd, T., Borji, A., Itti, L., Durand, F., Oliva, A., and Torralba, A. (2019). Mit saliency benchmark. <http://saliency.mit.edu/>.
- Clark, A. (2013). The many faces of precision (replies to commentaries on “whatever next? neural prediction, situated agents, and the future of cognitive science”). *Frontiers in Psychology*, 4.
- Crevecœur, F. and Kording, K. P. (2017). Saccadic suppression as a perceptual consequence of efficient sensorimotor estimation. *Elife*, 6:e25073.
- Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., and Friston, K. (2020). Active inference on discrete state-spaces: a synthesis. *Journal of Mathematical Psychology*, 99:102447.
- Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences*, 93(24):13494–13499.



- Dugué, L., Marque, P., and VanRullen, R. (2015). Theta oscillations modulate attentional search performance periodically. *Journal of cognitive neuroscience*, 27(5):945–958.
- Dugué, L., Roberts, M., and Carrasco, M. (2016). Attention reorients periodically. *Current Biology*, 26(12):1595–1601.
- Eldar, E., Cohen, J. D., and Niv, Y. (2013). The effects of neural gain on attention and learning. *Nature neuroscience*, 16(8):1146–1153.
- Feldman, H. and Friston, K. (2010). Attention, uncertainty, and free-energy. *Frontiers in human neuroscience*, 4:215.
- Ferreira, J. F. and Dias, J. (2014). Attentional mechanisms for socially interactive robots—a survey. *IEEE Transactions on Autonomous Mental Development*, 6(2):110–125.
- Fiebelkorn, I. C. and Kastner, S. (2019). A rhythmic theory of attention. *Trends in cognitive sciences*, 23(2):87–101.
- Fiebelkorn, I. C. and Kastner, S. (2020). Functional specialization in the attention network. *Annual review of psychology*, 71:221–249.
- Fiebelkorn, I. C. and Kastner, S. (2021). Spike timing in the attention network predicts behavioral outcome prior to target selection. *Neuron*, 109(1):177–188.
- Fiebelkorn, I. C., Pinsk, M. A., and Kastner, S. (2018). A dynamic interplay within the frontoparietal network underlies rhythmic spatial attention. *Neuron*, 99(4):842–853.
- Fiebelkorn, I. C., Pinsk, M. A., and Kastner, S. (2019). The mediodorsal pulvinar coordinates the macaque fronto-parietal network during rhythmic spatial attention. *Nature communications*, 10(1):1–15.
- Fine, M. S. and Minnery, B. S. (2009). Visual salience affects performance in a working memory task. *Journal of Neuroscience*, 29(25):8016–8021.
- Frintrop, S. (2006). *VOCUS: A visual attention system for object detection and goal-directed search*, volume 3899. Springer.
- Frintrop, S. and Jensfelt, P. (2008). Attentional landmarks and active gaze control for visual slam. *IEEE Transactions on Robotics*, 24(5):1054–1065.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2):127–138.

- Friston, K., Adams, R., Perrinet, L., and Breakspear, M. (2012). Perceptions as hypotheses: saccades as experiments. *Frontiers in psychology*, 3:151.
- Friston, K., Da Costa, L., Hafner, D., Hesp, C., and Parr, T. (2021). Sophisticated inference. *Neural Computation*, 33(3):713–763.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017a). Active inference: a process theory. *Neural computation*, 29(1):1–49.
- Friston, K., Mattout, J., and Kilner, J. (2011). Action understanding and active inference. *Biological cybernetics*, 104(1):137–160.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., and Pezzulo, G. (2015a). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4):187–214. PMID: 25689102.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., and Pezzulo, G. (2015b). Active inference and epistemic value. *Cognitive neuroscience*, 6(4):187–214.
- Friston, K., Stephan, K., Li, B., and Daunizeau, J. (2010). Generalised filtering. *Mathematical Problems in Engineering*, 2010.
- Friston, K. J., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, 19(4):1273–1302.
- Friston, K. J., Lin, M., Frith, C. D., Pezzulo, G., Hobson, J. A., and Ondobaka, S. (2017b). Active inference, curiosity and insight. *Neural computation*, 29(10):2633–2683.
- Friston, K. J., Parr, T., Yufik, Y., Sajid, N., Price, C. J., and Holmes, E. (2020). Generative models, linguistic communication and active inference. *Neuroscience Biobehavioral Reviews*, 118:42–64.
- Friston, K. J., Trujillo-Barreto, N., and Daunizeau, J. (2008). Dem: a variational treatment of dynamic systems. *Neuroimage*, 41(3):849–885.
- Gazzaley, A. and Nobre, A. C. (2012). Top-down modulation: bridging selective attention and working memory. *Trends in cognitive sciences*, 16(2):129–135.
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual neuroscience*, 9(2):181–197.
- Helfrich, R. F., Fiebelkorn, I. C., Szczepanski, S. M., Lin, J. J., Parvizi, J., Knight, R. T., and Kastner, S. (2018). Neural mechanisms of sustained attention are rhythmic. *Neuron*, 99(4):854–865.

- Helmholtz, H. v. (1925). Handbook of physiological optics southall. *JPC The Optical Society of America*, 3.
- Hitz, G., Galceran, E., Garneau, M.-È., Pomerleau, F., and Siegwart, R. (2017). Adaptive continuous-space informative path planning for online environmental monitoring. *Journal of Field Robotics*, 34(8):1427–1449.
- Hogendoorn, H. (2016). Voluntary saccadic eye movements ride the attentional rhythm. *Journal of cognitive neuroscience*, 28(10):1625–1635.
- Hsieh, L.-T. and Ranganath, C. (2014). Frontal midline theta oscillations during working memory maintenance and episodic encoding and retrieval. *Neuroimage*, 85:721–729.
- Itti, L. and Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision research*, 49(10):1295–1306.
- Itti, L. and Koch, C. (2001). Computational modelling of visual attention. *Nature reviews neuroscience*, 2(3):194–203.
- Kanai, R., Komura, Y., Shipp, S., and Friston, K. (2015). Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1668):20140169.
- Kanwisher, N. and Wojciulik, E. (2000). Visual attention: insights from brain imaging. *Nature reviews neuroscience*, 1(2):91–100.
- Kaplan, F. and Hafner, V. V. (2006). The challenges of joint attention. *Interaction Studies*, 7(2):135–169.
- Kim, A. and Eustice, R. M. (2013). Real-time visual slam for autonomous underwater hull inspection using visual saliency. *IEEE Transactions on Robotics*, 29(3):719–733.
- Klein, R. M. (2000). Inhibition of return. *Trends in cognitive sciences*, 4(4):138–147.
- Klink, P. C., Jentgens, P., and Lorteije, J. A. (2014). Priority maps explain the roles of value, attention, and salience in goal-oriented behavior. *Journal of Neuroscience*, 34(42):13867–13869.
- Knudsen, E. I. (2007). Fundamental components of attention. *Annu. Rev. Neurosci.*, 30:57–78.
- Köster, M., Finger, H., Graetz, S., Kater, M., and Gruber, T. (2018). Theta-gamma coupling binds visual perceptual features in an associative memory task. *Scientific reports*, 8(1):1–9.

- Kragic, D., Björkman, M., Christensen, H. I., and Eklundh, J.-O. (2005). Vision for robotic object manipulation in domestic settings. *Robotics and autonomous Systems*, 52(1):85–100.
- Lanillos, P. (2013). *Minimum time search of moving targets in uncertain environments*. PhD thesis, PhD thesis.
- Lanillos, P. and Cheng, G. (2018a). Active attention applications in robotics. In *International Workshop on Active vision, Attention, and Learning, IEEE Developmental Learning and Epigenetic Robotics (ICDL-Epirob)*. IEEE.
- Lanillos, P. and Cheng, G. (2018b). Adaptive robot body learning and estimation through predictive coding. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4083–4090. IEEE.
- Lanillos, P., Dean-Leon, E., and Cheng, G. (2016). Yielding self-perception in robots through sensorimotor contingencies. *IEEE Transactions on Cognitive and Developmental Systems*, 9(2):100–112.
- Lanillos, P., Ferreira, J. F., and Dias, J. (2015a). Designing an artificial attention system for social robots. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4171–4178. IEEE.
- Lanillos, P., Ferreira, J. F., and Dias, J. (2015b). Multisensory 3d saliency for artificial attention systems.
- Lanillos, P., Gan, S. K., Besada-Portas, E., Pajares, G., and Sukkarieh, S. (2014). Multi-uav target search using decentralized gradient-based negotiation with expected observation. *Information Sciences*, 282:92–110.
- Lanillos, P., Meo, C., Pezzato, C., Meera, A. A., Baoumy, M., Ohata, W., Tschantz, A., Millidge, B., Wisse, M., Buckley, C. L., et al. (2021). Active inference in robotics and artificial agents: Survey and challenges. *arXiv preprint arXiv:2112.01871*.
- LaValle, S. M. (2006). *Planning algorithms*. Cambridge university press.
- Lengyel, M., Yang, S. C.-H., and Wolpert, D. M. (2016). Active sensing in the categorization of visual patterns. *eLife*.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005.

- Louie, K. and Glimcher, P. W. (2019). Normalization principles in computational neuroscience. In *Oxford Research Encyclopedia of Neuroscience*.
- Marchant, R. and Ramos, F. (2014). Bayesian optimisation for informative continuous path planning. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6136–6143. IEEE.
- Meera, A. A., Popović, M., Millane, A., and Siegwart, R. (2019). Obstacle-aware adaptive informative path planning for uav-based target search. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 718–724. IEEE.
- Meera, A. A. and Wisse, M. (2020). Free energy principle based state and input observer design for linear systems with colored noise. In *2020 American Control Conference (ACC)*, pages 5052–5058. IEEE.
- Meera, A. A. and Wisse, M. (2021a). A brain inspired learning algorithm for the perception of a quadrotor in wind. *arXiv preprint arXiv:2109.11971*.
- Meera, A. A. and Wisse, M. (2021b). On the convergence of dem’s linear parameter estimator. In *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, pages 692–700, Cham. Springer International Publishing.
- Meera, A. A. and Wisse, M. (2022). Free energy principle for the noise smoothness estimation of linear systems with colored noise. *arXiv preprint arXiv:2204.01796*.
- Mirza, M. B., Adams, R. A., Friston, K., and Parr, T. (2019). Introducing a bayesian model of selective attention based on active inference. *Scientific reports*, 9(1):1–22.
- Mirza, M. B., Adams, R. A., Mathys, C. D., and Friston, K. J. (2016a). Scene construction, visual foraging, and active inference. *Frontiers in Computational Neuroscience*, 10.
- Mirza, M. B., Adams, R. A., Mathys, C. D., and Friston, K. J. (2016b). Scene construction, visual foraging, and active inference. *Frontiers in computational neuroscience*, 10:56.
- Morén, J., UDE, A., Koene, A., and Cheng, G. (2008). Biologically based top-down attention modulation for humanoid interactions. *International Journal of Humanoid Robotics*, 5(01):3–24.
- Mousavi, S., Schukat, M., Howley, E., Borji, A., and Mozayani, N. (2016). Learning to predict where to look in interactive environments using deep recurrent q-learning. *arXiv preprint arXiv:1612.05753*.

- Nagai, Y., Hosoda, K., Morita, A., and Asada, M. (2003). A constructive model for the development of joint attention. *Connection Science*, 15(4):211–229.
- Nakajima, M., Schmitt, L. I., and Halassa, M. M. (2019). Prefrontal cortex regulates sensory filtering through a basal ganglia-to-thalamus pathway. *Neuron*, 103(3):445–458.
- Nakayama, R. and Motoyoshi, I. (2019). Attention periodically binds visual features as single events depending on neural oscillations phase-locked to action. *Journal of Neuroscience*, 39(21):4153–4161.
- Oberauer, K. (2019). Working memory and attention—a conceptual analysis and review. *Journal of cognition*, 2(1).
- Oliver, G., Lanillos, P., and Cheng, G. (2021). An empirical study of active inference on a humanoid robot. *IEEE Transactions on Cognitive and Developmental Systems*.
- Orabona, F., Metta, G., and Sandini, G. (2005). Object-based visual attention: a model for a behaving robot. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)-Workshops*, pages 89–89. IEEE.
- Oudeyer, P.-Y. and Kaplan, F. (2009). What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, 1:6.
- Panichello, M. F. and Buschman, T. J. (2021). Shared mechanisms underlie the control of working memory and attention. *Nature*, 592(7855):601–605.
- Parr, T., Benrimoh, D. A., Vincent, P., and Friston, K. J. (2018). Precision and false perceptual inference. *Frontiers in integrative neuroscience*, 12:39.
- Parr, T., Corcoran, A. W., Friston, K. J., and Hohwy, J. (2019). Perceptual awareness and active inference. *Neuroscience of consciousness*, 2019(1):niz012.
- Parr, T. and Friston, K. J. (2017a). Uncertainty, epistemics and active inference. *Journal of The Royal Society Interface*, 14(136):20170376.
- Parr, T. and Friston, K. J. (2017b). Working memory, attention, and salience in active inference. *Scientific reports*, 7(1):1–21.
- Parr, T. and Friston, K. J. (2019). Attention or salience? *Current opinion in psychology*, 29:1–5.
- Parr, T. and Pezzulo, G. (2021). Understanding, explanation, and active inference. *Frontiers in Systems Neuroscience*, 15.

- Parr, T., Sajid, N., Da Costa, L., Mirza, M. B., and Friston, K. J. (2021). Generative models for active vision. *Frontiers in Neurorobotics*, 15:34.
- Peters, B., Kaiser, J., Rahm, B., and Bledowski, C. (2020). Object-based attention prioritizes working memory contents at a theta rhythm. *Journal of Experimental Psychology: General*.
- Phillips, J. M., Kambi, N. A., and Saalmann, Y. B. (2016). A subcortical pathway for rapid, goal-driven, attentional filtering. *Trends in neurosciences*, 39(2):49–51.
- Pomper, U. and Ansorge, U. (2021). Theta-rhythmic oscillation of working memory performance. *Psychological Science*, page 09567976211013045.
- Popović, M., Vidal-Calleja, T., Hitz, G., Sa, I., Siegwart, R., and Nieto, J. (2017). Multiresolution mapping and informative path planning for uav-based terrain monitoring. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1382–1388. IEEE.
- Rao, R. P. (2005). Bayesian inference and attentional modulation in the visual cortex. *Neuroreport*, 16(16):1843–1848.
- Rasouli, A., Lanillos, P., Cheng, G., and Tsotsos, J. K. (2020). Attention-based active visual search for mobile robots. *Autonomous Robots*, 44(2):131–146.
- Raybourn, M. S. and Keller, E. L. (1977). Colliculoreticular organization in primate oculomotor system. *Journal of Neurophysiology*, 40(4):861–878.
- Reynolds, J. H. and Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61(2):168–185.
- Reynolds, J. H., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of v4 neurons. *Neuron*, 26(3):703–714.
- Rizzolatti, G., Riggio, L., Dascola, I., and Umiltà, C. (1987). Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia*, 25(1):31–40.
- Roberts, R., Ta, D.-N., Straub, J., Ok, K., and Dellaert, F. (2012). Saliency detection and model-based tracking: a two part vision system for small robot navigation in forested environment. In *Unmanned Systems Technology XIV*, volume 8387, page 83870S. International Society for Optics and Photonics.

- Rucci, M., Ahissar, E., and Burr, D. (2018). Temporal coding of visual space. *Trends in cognitive sciences*, 22(10):883–895.
- Ruff, D. A. and Cohen, M. R. (2016). Stimulus dependence of correlated variability across cortical areas. *Journal of Neuroscience*, 36(28):7546–7556.
- Sajid, N., Da Costa, L., Parr, T., and Friston, K. (2021a). Active inference, bayesian optimal design, and expected utility. *arXiv preprint arXiv:2110.04074*.
- Sajid, N., Faccio, F., Da Costa, L., Parr, T., Schmidhuber, J., and Friston, K. (2021b). Bayesian brains and the r\'enyi divergence. *arXiv preprint arXiv:2107.05438*.
- Sajid, N., Friston, K. J., Ekert, J. O., Price, C. J., and W Green, D. (2020). Neuromodulatory control and language recovery in bilingual aphasia: An active inference approach. *Behavioral Sciences*, 10(10):161.
- Sajid, N., Holmes, E., Costa, L. D., Price, C., and Friston, K. (2022). A mixed generative model of auditory word repetition. *bioRxiv*.
- Santangelo, V. (2015). Forced to remember: when memory is biased by salient information. *Behavioural brain research*, 283:1–10.
- Santangelo, V., Di Francesco, S. A., Mastroberardino, S., and Macaluso, E. (2015). Parietal cortex integrates contextual and saliency signals during the encoding of natural scenes in working memory. *Human Brain Mapping*, 36(12):5003–5017.
- Santangelo, V. and Macaluso, E. (2013). Visual salience improves spatial working memory via enhanced parieto-temporal functional connectivity. *Journal of Neuroscience*, 33(9):4110–4117.
- Schmitz, T. W. and Duncan, J. (2018). Normalization and the cholinergic microcircuit: a unified basis for attention. *Trends in cognitive sciences*, 22(5):422–437.
- Shon, A. P., Grimes, D. B., Baker, C. L., Hoffman, M. W., Zhou, S., and Rao, R. P. (2005). Probabilistic gaze imitation and saliency learning in a robotic head. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 2865–2870. IEEE.
- Sommer, M. A. and Wurtz, R. H. (2006). Influence of the thalamus on spatial visual processing in frontal cortex. *Nature*, 444(7117):374–377.
- Spratling, M. W. (2008). Predictive coding as a model of biased competition in visual attention. *Vision research*, 48(12):1391–1408.



- Tomassini, A., Ambrogioni, L., Medendorp, W. P., and Maris, E. (2017). Theta oscillations locked to intended actions rhythmically modulate perception. *Elife*, 6:e25618.
- Treisman, A. M. and Gelade, G. (1980). A feature-integration theory of attention. *Cognitive psychology*, 12(1):97–136.
- Tsotsos, J. K., Culhane, S. M., Wai, W. Y. K., Lai, Y., Davis, N., and Nuflo, F. (1995). Modeling visual attention via selective tuning. *Artificial intelligence*, 78(1-2):507–545.
- Ude, A., Wyart, V., Lin, L.-H., and Cheng, G. (2005). Distributed visual attention on a humanoid robot. In *5th IEEE-RAS International Conference on Humanoid Robots, 2005.*, pages 381–386. IEEE.
- VanRullen, R. (2016). Perceptual rhythms. *Stevens Handbook of Experimental Psychology*, 1.
- Welch, G., Bishop, G., et al. (1995). An introduction to the kalman filter.
- White, B. J., Berg, D. J., Kan, J. Y., Marino, R. A., Itti, L., and Munoz, D. P. (2017). Superior colliculus neurons encode a visual saliency map during free viewing of natural dynamic video. *Nature communications*, 8(1):1–9.
- Whiteley, L. and Sahani, M. (2008). Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *Journal of vision*, 8(3):2–2.
- Yang, S. C.-H., Lengyel, M., and Wolpert, D. M. (2016a). Active sensing in the categorization of visual patterns. *Elife*, 5:e12215.
- Yang, S. C.-H., Wolpert, D. M., and Lengyel, M. (2016b). Theoretical perspectives on active sensing. *Current opinion in behavioral sciences*, 11:100–108.

## Chapter 6

# Psilocybin Accelerates EEG Microstate Transitions and Elevates Approximate Entropy

Based on:

Novický, F., Stoliker, D., Razi, A., & Zeldenrust, F. Psilocybin accelerates EEG microstate transitions and elevates approximate entropy. *Manuscript in preparation*.

### Abstract

Although the therapeutic potential of psilocybin has been documented, its effects on brain function remain incompletely understood. The relaxed beliefs under psychedelics (REBUS) theory proposes that psychedelics work by relaxing prior beliefs and increasing neural entropy, but this has primarily been tested using fMRI rather than EEG, which offers superior temporal resolution. This study investigated how psilocybin affects the spatiotemporal dynamics of the brain on the millisecond scale, examining whether its effects are modulated by mindfulness training and different cognitive states. Using EEG microstate and approximate entropy analyzes, we compared people who completed an 8-week Mindfulness-Based Cognitive Therapy program ( $N = 33$ ) with controls ( $N = 30$ ) in four conditions: video watching, resting state, meditation and music listening. The results showed that the 19 mg dose of psilocybin significantly altered brain dynamics,

decreasing the duration of the microstate while increasing the rates of occurrence and the complexity of the signal. Mindfulness training showed no significant effect on these changes. While brain activity patterns primarily distinguished between eyes-open and eyes-closed cognitive states, psilocybin notably diminished the typical neural differences between passive rest and attentional states (meditation and music). These findings support the prediction of the REBUS theory of increased and neural entropy and under psychedelics and suggest that psilocybin creates a more dynamic and less constrained brain state, particularly when visual input is reduced. The combination of microstate and entropy analyses provides complementary insights into how psychedelics affect both the temporal organization and complexity of neural activity.

## Keywords

psilocybin • resting state • EEG microstates • approximate entropy • mindfulness-based cognitive therapy

## 6.1 Introduction

Psilocybin, the primary psychoactive compound found in 'magic mushrooms', has emerged as a powerful tool for investigating fundamental brain dynamics and information processing. Acting primarily as a serotonin 5-HT<sub>2A</sub> receptor agonist, psilocybin induces profound alterations in neural activity and network organization (Nichols, 2016; Vollenweider and Preller, 2020). While its therapeutic potential in treating various psychiatric disorders has been well-documented ((Carhart-Harris et al., 2021; Goldberg et al., 2020; Griffiths et al., 2018), the precise mechanisms by which psilocybin modulates the activity of large-scale neural networks and information flow in the brain remain incompletely understood.

The relaxed beliefs under psychedelics (REBUS) theory, proposed by (Carhart-Harris and Friston, 2019), provides a compelling framework for understanding the neurophysiological and cognitive effects of psychedelics like psilocybin. This theory posits that psychedelics work by relaxing the precision weighting of prior beliefs, thereby altering the brain's predictive processing and allowing for more flexible, "bottom-up" information flow. According to this theory, the brain should exhibit increased randomness and variability in its neural activity under the influence of psychedelics due to the lower precision of top-down predictions, leading to an increase in signal diversity and complexity.

The REBUS theory has gained substantial empirical support across multiple studies using various methodologies. For instance, (Tagliazucchi et al., 2014) demonstrated an increased global functional connectivity in the brain after psilocybin administration using

fMRI which increases the variability of neural activity. (Herzog et al., 2023) developed whole-brain models showing an increased neural entropy under psychedelics, while (Savino and Nichols, 2022) found similar entropy increases in rat prefrontal cortex during LSD administration. Recent work by (Shinozuka et al., 2024) revealed that LSD flattens the hierarchy of directed information flow in whole-brain dynamics, and (Viol et al., 2019) characterized network changes in functional brain connectivity induced by ayahuasca using entropy-degree diagrams. Additional support comes from (McCulloch et al., 2023), (Mediano et al., 2024), (Singleton et al., 2021), and (Siegel et al., 2024), who consistently found increased entropy under psychedelics using various analytical approaches. However, these studies predominantly relied on either animal models or fMRI measurements, leaving a gap in our understanding of these effects through EEG analysis, which offers millisecond-level temporal resolution in measuring neural activity. EEG data can be particularly analyzed with two complementary analytical approaches: microstate analysis and approximate entropy (ApEn). Microstate analysis is a clustering method specifically designed for EEG data that characterizes the brain’s electrical activity as a sequence of quasi-stable spatial configurations of the scalp potential field (Pascual-Marqui et al., 1995). These microstates, typically lasting tens of milliseconds (ZanESCO, 2024), represent the coordinated activity of large-scale neural networks and have been linked to various cognitive processes such as attention, memory, and decision making (Michel and Koenig, 2018; Khanna et al., 2015). By analyzing transitions between specific microstates, we can directly examine how psilocybin affects the spatiotemporal organization of brain activity, with increased transition rates and shorter microstate durations potentially indicating a higher entropy in neural dynamics.

Approximate entropy (ApEn) analysis provides a complementary perspective by quantifying the regularity or predictability of time series data (Pincus, 1991; Cannard and Delorme, 2022). When applied to EEG signals, ApEn can detect increases in signal complexity that may reflect the relaxation of prior beliefs proposed by the REBUS theory. Higher ApEn values indicate more complex and irregular signals, potentially capturing the increased bottom-up information flow and reduced hierarchical constraints that REBUS theory predicts should occur under psychedelics.

Together, the methods described above offer distinct yet complementary windows into how psilocybin affects brain dynamics: microstate analysis reveals changes in large-scale network coordination, while ApEn quantifies shifts in signal complexity at a more granular level. This dual approach allows us to test REBUS theory’s predictions about increased neural entropy and flexibility using temporal dynamics captured by EEG, complementing existing fMRI findings about spatial network reorganization.

Building on these analytical approaches, we investigate the acute effects of psilocybin

on resting state EEG across multiple experimental conditions aims to elucidate psilocybin's impact on large-scale brain network dynamics and signal complexity (Tagliazucchi et al., 2014; Lebedev et al., 2016). This study extends previous research by investigating the acute effects of psilocybin on resting state EEG using both microstate and approximate entropy analyses. In addition to this, the study compares individuals who followed a Mindfulness-Based Cognitive Therapy (MBCT) program with those who did not, allowing for an exploration of potential interactions between psilocybin effects and mindfulness practices (Kuyken et al., 2016b; Teasdale et al., 2000). More importantly, though, the analysis encompasses four distinct resting state conditions: watching a video (eyes open), closed eyes, meditation, and listening to music. This multi-faceted approach enables a more detailed understanding of psilocybin's effects across various contexts and sensory inputs (Carhart-Harris et al., 2014; Schartner et al., 2017; Barrett et al., 2018).


Building upon the REBUS theory and previous research on psychedelic alteration of brain activity (Carhart-Harris and Friston, 2019), this study explores several intriguing questions regarding psilocybin's effects on brain function:

1. How does prior MBCT training influence brain activity patterns and their modulation by psilocybin on microstates?
2. In what ways do psilocybin's effects differ across various resting state conditions (eyes open, closed eyes, meditation, and music listening)?
3. What changes in microstate dynamics and approximate entropy occur following psilocybin administration, and how do these reflect alterations in large-scale network coordination and signal complexity?

By addressing these questions, this study aims to provide a more comprehensive understanding of how psilocybin modulates brain function. The exploration of these aspects may offer valuable insights into the complex interplay between psychedelic substances, mindfulness practices, and varied cognitive states.

## 6.2 Materials and methods

This section explains the processing of PsiConnect data from (Novelli et al., 2025), the analyses used in this study and their statistical comparisons. For clarity, the reader can follow Figure 6.1. The links to code repository and the PsiConnect dataset is available in Software note.



**Figure 6.1: Methodological workflow of the study.** The pipeline consists of three main stages: (1) EEG processing, including filtering and downsampling of the PsiConnect dataset; (2) Analysis methods, comprising microstate analysis and approximate entropy calculations; and (3) Statistical comparisons examining differences between MBCT/non-MBCT groups, experimental sub-conditions, and psilocybin/control conditions. Each analysis branch shows the sequential processing steps applied to the data and their interconnections.

### 6.2.1 Participants and procedure

EEG datasets from the PsiConnect project were used for this study. The PsiConnect dataset consists of recordings from 63 participants that were assigned to two groups: MBCT (N=33, average age = 35.52; SD = 10.64; range = 20 — 54; 14 women, 1 non-binary) and non-MBCT (N = 30, average age = 40.03; SD = 10.85; range 19 — 53; 16 women). The assignment was non-randomized and balanced. The MBCT group followed a routine of meditating daily (around 20 minutes) and attended a meditation session lasting around 70 minutes every week. This lasted 8 weeks and followed the Mindfulness-Based Cognitive Practice (MBCT) guideline (Williams et al., 2011). For more detailed information about the participants, please see Table S1. The non-MBCT did not do any of this.

EEG recordings were conducted at two time points: a few weeks before psilocybin intake (labeled as control) and directly after psilocybin intake (labeled as psilocybin). The intake consisted of 19mg of psilocybin administered by a trained doctor. The EEG recording happened roughly after 75 to 120 minutes after the psilocybin intake. During both control and psilocybin EEG measurements, there were four consecutive recordings happening in the following order for each participant — watching a video, resting state (with the only condition that the eyes were kept closed), guided meditation, and listening to music. Note that in a few cases, the order was changed due to a malfunctioning of equipment. The first three recordings lasted for 5 minutes, while the last one – music – lasted for 7 minutes.

### 6.2.2 EEG Data Acquisition

A BrainAmp MR Plus EEG device with 64 channels was used in the Psiconnect project with the 10-20 system. This system means that the distances between adjacent electrodes are 10% and 20% of the total front-back and right-left distance of the skull, respectively (HH, 1958). The sampling rate was 500Hz, with the electrode *Fz*—located in the frontal midline—set as the acquisition reference. Participants were seated comfortably in an upright position during the measurements. Prior to the analysis, the participants gave their ethical consent to participating in the study.

### 6.2.3 Preprocessing

The datasets were preprocessed using EEGLAB (Delorme and Makeig, 2004) following these steps:

1. Re-referencing to average and bandpass filtering (1-40 Hz)<sup>1</sup>
2. Automated artifact removal using the CleanRawData and Artifact Subspace Reconstruction functions in EEGLAB
3. Independent Component Analysis (ICA) for further artifact detection and removal using ICLabels (Pion-Tonachini et al., 2019)
4. Visual quality check using MICROSTATELAB’s k-means clustering (12 clusters) (Nagabhushan Kalburgi et al., 2023). Datasets with >7% unexplained variance underwent verification and manual cleaning of overlooked artifacts
5. Downsampling to a sampling rate of 125 Hz for Approximate Entropy analysis

Datasets with excessive electrode (more than 25%) or time segment removal (more than 66%) were discarded. Final dataset counts for each condition are presented in Table 6.1. Note that this cleaning procedure mostly affected the video sub-condition, particularly in the psilocybin condition.

Condition	Group	Video	Resting State	Meditation	Music
Control	Non-MBCT	24	28	29	29
	MBCT	29	31	31	31
Psilocybin	Non-MBCT	17	25	26	24
	MBCT	22	27	28	28
Shared subjects	Non-MBCT	15	23	25	23
	MBCT	19	25	26	26

**Table 6.1:** Final number of recordings for each subgroup (MBCT = trained meditators; non-MBCT = no training in meditation), in both psilocybin and control conditions, for all four sub-conditions. The last row indicates how many participants had both EEG recordings for psilocybin and control available after cleaning.

## 6.2.4 Analysis Methods

### 6.2.4.1 EEG microstate analysis

**Global Field Power** EEG microstate analysis starts with computing the Global Field Power (GFP) from the pre-processed EEG recording as follows:

<sup>1</sup>Average referencing involves taking the mean voltage of all electrodes at each time point and subtracting it from each individual electrode’s voltage. Note that this process is particularly helpful when dealing with topographical maps, as this method keeps the distribution of individual electrodes relative to each other (Murray et al., 2008; Dien, 1998).



$$GFP_j = \sqrt{\frac{\sum_{i=1}^{n_{ch}} X_{i,j}^2}{n_{ch}}}. \quad (6.1)$$

Here,  $j$  is the index of a given time point,  $i$  is the channel index, so  $X_{i,j}^2$  represents the squared value of a recorded channel value  $i$  at the time-step  $j$ . The final sum is then divided by the number of channels ( $n_{ch}$ ). Once the GFP is obtained, one also needs to obtain the peaks of these. These peaks are considered only above a given threshold, usually 2 standard deviations. The peaks are detected using a sign function:

$$\text{GFP Peak} = \{X_j | \text{sgn}(\delta_j) - \text{sgn}(\delta_{j-1}) = -2\} \quad (6.2)$$

where  $\delta_j = GFP_{j+1} - GFP_j$ . Thus, the value GFP Peak stores the peak time-points of the Global Field Power  $GFP$  in  $X_j$ . In other words,  $\delta$  is positive if the GFP signal goes up and negative if it goes down. To get to -2,  $\delta_j$  needs to be negative and  $\delta_{j-1}$  positive, so one goes from a rise to a decline, which is indeed the top of a peak. Ultimately, the GFP peak values are then used in the clustering algorithm, because they represent time points where the electric field strength is maximal, providing the highest signal-to-noise patterns of brain activity (Murray et al., 2008). Additionally, using only GFP peaks rather than all time points reduces the computational load considerably (Jajcay and Hlinka, 2023).

**Atomize and Agglomerate Hierarchical clustering** To identify the dominant spatial patterns at the GFP peaks obtained above, we employ the Atomize and Agglomerate Hierarchical Clustering (AAHC) method. Unlike the traditional modified k-means approach (Pascual-Marqui et al., 1995), AAHC provides a deterministic solution, thus yielding fully consistent results for every run (Murray et al., 2008; Von Wegner et al., 2018). This is because AAHC creates a topographic map for every GFP Peak value. Next, it starts removing the topographic map with the lowest global explained variance (Von Wegner et al., 2018), one per iteration. Once this is done, the components of the removed cluster are re-assigned to one of the remaining clusters. This is repeated until a final number of pre-selected clusters remain (Murray et al., 2008; Von Wegner et al., 2018), which in the case of this study is four for consistency with previous research on microstates (Michel and Koenig, 2018). Once the clusters have been found, they are fitted back to the original data, where only one microstate is assigned to every time step, which is determined by the highest explained variance out of all microstates. For more technical details on this process, please see (Von Wegner et al., 2018).

The final result of applying the procedure described above to our data is shown in Figure (6.2, top row). These topographic maps are used as a reference for all the clusters in all individual datasets. The deterministic output of AAHC establishes a rigour, as one


can rely on running the same analysis on the same dataset and always get the same final topographic maps, which is not necessarily the case with a stochastic modified k-means method, which can be problematic especially with large datasets such as PsiConnect, although it is worth highlighting that it has been previously shown that both methods provide negligible differences when compared (Jajcay and Hlinka, 2023).

The analysis was conducted using four microstates, following the established tradition in EEG microstate research (Michel and Koenig, 2018; Khanna et al., 2015). While the optimal number of microstates can vary depending on the specific dataset and research question, four microstates have been consistently found to explain a significant portion of the variance in resting-state EEG data across numerous studies (Britz et al., 2010; Koenig et al., 2002). This approach allows for comparability with existing literature and strikes a balance between data reduction and physiological interpretability.


**Shared Variance between individual and universal microstates** To ensure data quality, we quantified how well individual participants' microstate topographies matched the reference topographies shown in Figure 6.2. This similarity was measured using shared variance, calculated as the squared spatial correlation coefficient between the voltage distributions of an individual's microstate and its corresponding reference microstate (Nagabhushan Kalburgi et al., 2023). If an individual's microstate topography shared less than 60% variance with its reference, that microstate's statistics were excluded from subsequent analyses for that participant. This quality control step resulted in excluding approximately 20% of all microstates from further analysis. Importantly, this exclusion was applied only after the clustering procedure was complete and the microstates were used to explain the overall variance in the EEG signals. The distributions of shared variance values for each microstate type are shown in Figure 6.3. Microstates A and B required more frequent exclusion (21.12% and 22.38% of cases respectively) compared to Microstates C and D (13.52% and 12.59% respectively), suggesting that the latter two patterns were more consistently expressed across participants.

#### 6.2.4.2 Approximate Entropy

Quantifying the complexity and regularity of the EEG signals was done with an Approximate Entropy (ApEn) algorithm (Pincus, 1991). ApEn is a statistical measure that quantifies the unpredictability of fluctuations in a time series. Lower ApEn values indicate more regular and predictable signals, while higher values suggest more complex and irregular signals. The ApEn function processes the EEG data in sliding windows (5s in our study) and computes the mean ApEn for each channel. The code is available in the repository (see Software Note), and the simplified pseudoalgorithm can be found in



**Figure 6.2: Topographies for the final four microstates:** The first row shows the reference microstates, created by averaging across all datasets. These reference topographies were used as templates to ensure consistent labeling of microstates across individual datasets. The second and third rows display the average topographies for control and psilocybin conditions respectively, demonstrating their consistency with the reference templates. Red coloring indicates positive voltage relative to the reference, while blue indicates negative voltage. Note that these condition-specific averages are shown for visualization only and were not used for microstate labeling. The consistency between all three rows validates our microstate classification approach. Red and blue colors show positive and negative voltages, respectively. Identical topographies with opposite polarities (reversed red/blue patterns) are considered the same microstate, as they represent the same neural configuration.



**Figure 6.3: Distributions of shared variance between individual topographies with their reference:** The red dashed lines in all plots highlight a percentile at the 60% of the shared variance between individual dataset microstates and reference microstate. This is a threshold line and all information related to the individual cluster below 60% was removed from the analysis.

the Supplementary Materials.

### 6.2.5 Statistical analyses

30 variables from the MICROSTATELAB plugin in EEGLAB by (Nagabhushan Kalburgi et al., 2023) were analyzed — apart from Total time of the dataset after removing artifacts and Total explained variance of the individual dataset by AAHC. This consisted of data from Occurrence – the frequency of a microstate appearing per second, Duration – the average time a microstate remains stable before transitioning to another microstate, Coverage – the percentage of total analysis time occupied by a specific microstate, and mean GFP (see Equation 6.1) from all four microstates. Next, for both occurrence and duration, a summation of all microstates was also provided (i.e., the sum of all occurrence and duration values across all 4 microstates.) The other variables consisted of transitions among all individual microstates. All comparisons can be found in the Supplementary materials in section 'Statistical comparisons'.

#### 6.2.5.1 Assumption of the normality distribution

The choice of parametric or non-parametric tests for comparing continuous distributions depends on whether a given dataset is normally distributed. To test this, the Shapiro-Wilk test for normality was applied to all datasets that were compared using the microstate analysis. For a given statistical comparison, the parametric relative or individual sample tests were applied if both measured distributions resulted in a non-significant value (p-value  $>0.05$ ) of the Shapiro-Wilk test. If this condition was not met, the non-parametric Mann-Whitney U or Wilcoxon Rank-Sum tests were applied instead. All these tests were implemented using the SciPy toolbox (Virtanen et al., 2020). An alternative would be to use an ANOVA, which would provide a correction of p-values with it, but this was resolved by including the False Discovery rate (FDR) correction (Benjamini and Hochberg, 1995). Its precise application is found the Supplementary materials in section 'Statistical comparisons'.

## 6.3 Results

This study investigated the acute effects of psilocybin on resting state EEG activity across different mental states. We examined whether prior mindfulness training (MBCT) influences psilocybin's effects on brain activity, how these effects differ across cognitive states (video, resting state, guided meditation, and music), and how psilocybin alters the brain's microstate dynamics and signal complexity.

In this section, we first describe the microstate analysis approach, then examine MBCT vs non-MBCT group differences, compare the sub-conditions, and finally analyze psilocybin's effects on brain dynamics through both microstate and entropy analyses.

### 6.3.1 Decomposing the EEG signal into microstates

Before examining psilocybin's effects, we first needed to establish whether our microstate analysis could reliably capture distinct brain states across conditions. To assess this, we examined how well our four microstate clusters could explain the variance in the EEG data across different conditions.

The following topographies were identified for this study: right-frontal left-posterior (microstate A), left-frontal right-posterior (microstate B), midline fronto-occipital (microstate C) and a maximum at the parieto-occipital region (microstate D). These are shown in Figure 6.2.

In order to assess the performance of the AAHC clustering method in extracting the basis topographies for representing the EEG data, the Total Explained Variance was analyzed, as it represents how well the 4 clusters can reconstruct represent the entire raw dataset. Note that this analysis has been done with all the intact microstates, including those that did not meet the criterion of sharing at least 60% of variance with the clusters (Figure 6.3). The total explained variance when using the four clusters to reconstruct the EEG signal in the psilocybin and control conditions was 66.6% and 65.8%, respectively. For the sub-conditions ( where psilocybin and control were grouped together), it was 59.89% for video, 68.51% for resting state, 67.27% for music and 68.19% for meditation. Lastly, for the groups the results are 65.98% for MBCT and 66.56% for the non-MBCT group. Apart from the video sub-condition, which showed lower explained variance (59.89%) compared to other conditions ( 68%), the explained variance was consistent across conditions. The reduced performance in the video condition can be attributed to increased artifacts from eye movements and blinks - the same artifacts that led to more data rejection during preprocessing. The high and consistent explained variance across remaining conditions demonstrates that our four-microstate decomposition provides a robust representation of the EEG dynamics across experimental conditions.

### 6.3.2 No significant difference between MBCT and non-MBCT groups

Having validated our microstate decomposition approach, we next examined whether mindfulness training influences brain dynamics under psilocybin by comparing MBCT and non-MBCT groups across all experimental conditions.

For both psilocybin and control conditions, the MBCT and non-MBCT groups were compared per sub-conditions. This meant that there were  $4 \times 30 = 120$  tests for the control condition, and the same amount for the psilocybin condition. The statistical analysis was performed either by a non-parametric Mann-Whitney U test (when the distributions were not considered normally distributed based on the Shapiro-Wilk test) or by parametric independent t-test. After the FDR correction for both psilocybin and control, the results indicated that there were no significant differences between MBCT and non-MBCT for any of the sub-conditions or measured variables (Table S2). Hence, for the following analyses, the MBCT and non-MBCT were grouped together. This increased the number of participants per a group, which naturally leads to more statistical power.

### 6.3.3 Sub-conditions reveal changing patterns of attentional states

Having established that MBCT and non-MBCT groups can be pooled for analysis, we next examined how different cognitive states - video watching, resting state, meditation, and music listening - influence brain dynamics under psilocybin and control conditions.

The following analysis examined differences between sub-conditions within both control and psilocybin conditions, comparing video watching, resting state, meditation, and music listening. Simplified results are provided in Table 6.2, and the extended ones can be found in Table S3. The most striking differences emerged between eyes-open (video) and eyes-closed conditions, but subtle yet important distinctions also appeared in how psilocybin affected the relationship between passive rest and more directed attentional states.

In both control and psilocybin conditions, the video condition showed markedly different patterns from all other conditions. Of the significant differences found, 46 out of 59 in the control condition and 34 out of 36 in the psilocybin condition involved comparisons between video and other states. Table 6.2 summarizes these consistent video-related effects, showing a decreased Global Field Power (GFP) and duration but increased occurrence across all microstates. However, this difference might be taken with caution, as the video sub-condition had more artifacts which resulted in a smaller number of participants, and the overall explained variance using the microstate analysis was also smaller, compared to other sub-conditions.

A more nuanced pattern emerged when examining the eyes-closed conditions. In the control condition, significant differences in mean GFP in all microstates distinguished the resting state from both meditation and music listening (8 significant comparisons), suggesting neurophysiologically distinct states between passive rest and directed atten-



tion, such as meditating and listening to music. Intriguingly, under psilocybin, these distinctions largely disappeared, with only two significant differences for microstate A in the GFP difference when comparing the music sub-condition with the resting state. This convergence of resting state with meditation and music listening under psilocybin suggests that these traditionally distinct mental states become more similar at a neurophysiological level between active and passive conditions. While in the control condition, the passive resting state shows distinct GFP patterns from meditation and music listening, psilocybin appears to reduce these differences, resulting in more similar neural signatures across all eyes-closed conditions. This finding is particularly intriguing as it suggests psilocybin may alter the usual boundaries between different attentional states, though the precise nature—whether psilocybin makes the resting state more active or the active states more passive—is unclear.

Microstate A	Microstate B	Microstate C	Microstate D
Mean GFP ↓	Mean GFP ↓	Mean GFP ↓	Mean GFP ↓
Duration ↓	Duration ↓	Duration ↓	Duration ↓
Occurrence ↑	Occurrence ↑	Occurrence ↑	Occurrence ↑

**Table 6.2:** Significant differences for both psilocybin and control conditions, found in all comparisons between video and the rest of the sub-conditions. Arrows indicate whether the average mean for video increased (↑) or decreased (↓) compared to all remaining conditions.

### 6.3.4 The effects of psilocybin

The analysis revealed distinct patterns in how psilocybin modulates brain activity across different cognitive states. Here, we examine how psilocybin alters three key aspects of brain dynamics: microstate occurrence rates, duration patterns, and Global Field Power. By analyzing these metrics across different experimental conditions, we aim to characterize how psilocybin influences the temporal organization and strength of neural activity patterns.

When examining microstate occurrence (Figure 6.4A), a clear pattern emerged: during eyes-closed conditions (resting state, meditation, and music), psilocybin significantly increased the frequency of all microstates. This suggests a more dynamic brain state with more frequent transitions between neural configurations. In contrast, during the video condition (eyes open), only microstate D showed increased occurrence, indicating that visual input may partially constrain psilocybin’s effects on brain dynamics.

The duration patterns (Figure 6.4B) showed a complementary effect: psilocybin consistently decreased the duration of all microstates during eyes-closed conditions, while



in the video condition, only microstate B showed reduced duration. This shortened duration coupled with increased occurrence suggests that psilocybin creates a more rapidly shifting neural landscape when visual input is missing.

Global Field Power analysis (Figure 6.4C) revealed state-specific effects of psilocybin. During resting state, microstates A, B, and D showed a significantly reduced GFP under psilocybin, while during music and meditation, only microstates B and D exhibited this reduction. Notably, while the video condition showed no significant GFP differences between psilocybin and control conditions, it consistently displayed a lower overall GFP compared to eyes-closed conditions. This lower GFP during video watching likely reflects the brain’s engagement with structured visual input, which may override some of psilocybin’s typical effects on neural dynamics.

These findings collectively suggest that psilocybin’s primary effect is to create a more dynamic and rapidly shifting brain state when visual input is reduced. The drug appears to promote a pattern of shorter, more frequent microstate transitions, potentially reflecting a more flexible and less constrained neural system.


### 6.3.5 Transition and Coverage

Here, we examined the transitions between different microstates and their relative prominence (coverage) to understand whether psilocybin affects the sequential organization of brain states and their overall expression. This analysis complemented our previous findings by focusing on the directional flow between brain states rather than their individual transition characteristics.

Our analysis of (directed) microstate transitions revealed surprisingly few significant differences across conditions. Out of 240 total transition comparisons, only 5 showed statistically significant changes. Under the influence of psilocybin, we observed an increase in transitions from microstate B to D and a decrease from C to A. In the control condition, when comparing video and meditation sub-conditions, we found increases in transitions from A to B and B to A, along with a decrease from D to C (see Tables S4 and S3).

However, these transition findings should be interpreted with caution. Recent research by (Kleinert et al., 2024) has demonstrated that directed microstate transitions generally show low reliability as a metric. This methodological limitation suggests that while the observed transitions may represent real phenomena, additional studies would be needed to confirm their reliability and functional significance.

Similarly, coverage analysis—the percentage of total recording time dominated by each microstate—showed minimal differences across conditions, with only one significant result out of 80 comparisons. This single difference appeared when comparing video



**Figure 6.4: Effects of psilocybin on EEG microstate dynamics across different conditions.** A) Microstate occurrence rate (number of times per second) for each microstate class (A-D). B) Mean duration (in milliseconds) of each microstate class. C) Global Field Power (GFP) for each microstate class, representing the strength of the scalp electric field. Conditions shown are: video watching, eyes closed resting state, meditation, and music listening. Significant differences between psilocybin and control conditions are indicated by asterisks: \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$ . Note the differential effects of psilocybin across microstate classes and experimental conditions.

and meditation sub-conditions in the control condition, where the video sub-condition showed a decreased coverage (Table S3). The scarcity of significant coverage differences suggests that the relative distribution of individual microstates remains largely stable across both psychedelic states and different cognitive contexts. This stability in coverage, contrasted with the clear changes in occurrence and duration reported earlier, indicates that psilocybin may alter the temporal dynamics of brain states while preserving their overall proportional expression. In other words, psilocybin does not target a specific brain region to fire more or less, but rather only functions by modulating the connectivity.

### 6.3.6 Approximate entropy


To strengthen the claim that the increased microstate transitions and their shorter duration highlights an increase in the brain's dynamic repertoire was done by using the approximate entropy analysis (Pincus, 1991). This analysis revealed significant changes in signal complexity across experimental conditions (Figure 6.5). In the comparison between psilocybin and control conditions, psilocybin consistently increased ApEn in the resting state, meditation, and music conditions. In agreement with the results described above, the video condition showed no significant differences, possibly due to the strong influence of visual input and artifacts by blinking.

Notably, the eyes-closed conditions exhibited the largest increase in ApEn under psilocybin compared to the control condition, suggesting that psilocybin has the strongest effect on the complexity of brain signals when visual input is minimized.

The video sub-condition showed the least difference in ApEn between psilocybin and control conditions, but again it has to be highlighted that this sub-condition suffered from more artifacts, reduced explained variance and a smaller number of datasets (after cleaning).

## 6.4 Discussion

Our analysis of psilocybin's effects on brain dynamics revealed three key findings. First, MBCT training showed no significant effect on neural dynamics within either drug condition across all experimental states. Second, when comparing cognitive states, the most prominent distinction emerged between eyes-open (video) and eyes-closed conditions, with an intriguing observation that psilocybin diminished the typical neural differences between passive rest and directed attention states (meditation and music) found in the control condition. The finding that the microstate coverage has not changed after the psilocybin administration points in a direction that the core effect of psilocybin



**Figure 6.5: EEG Topomap Visualization: ApEn differences and p-values across activities:** Topographical maps showing differences in Approximate Entropy (ApEn, left; calculated as psilocybin minus control) and statistical significance (right) between psilocybin and control conditions. Rows A-E represent: (A) Overall (i.e., all sub-conditions), (B) Meditation, (C) Video, (D) Resting State, and (E) Music. ApEn maps use a purple (-0.15) to yellow (0.15) scale, with yellow indicating a higher complexity under psilocybin. Significance maps use white to red, with darker red showing a higher significance.

is not to activate a specific brain region, but rather modulate the connectivity among regions. Importantly, it needs to be highlighted that the explained variance and number of datasets after the cleaning procedure of the video sub-condition was lower, thus the results on this analysis should be taken with more scepticism. Finally, both microstate and approximate entropy analyses provided converging evidence that psilocybin increases the brain’s dynamic repertoire during eyes-closed states, characterized by more frequent transitions between neural configurations, shorter microstate durations, and higher signal complexity. All the findings are described in detail below.

This study reports no difference between the MBCT and non-MBCT groups. These groups differed in that the MBCT group followed the mindfulness-based cognitive therapy (MBCT). They did this based on the book by (Williams et al., 2011): participants meditated daily for 8 weeks, approximately 20 minutes each day, and attended a weekly hour-long session with a trainer. MBCT’s effect on the modulation of neuronal and cognitive states seems to be undefined yet, as there are both observations that MBCT has no effect, or some significant observations were found (Szumska et al., 2021; Qiu et al., 2024; Yang et al., 2022; MacCoon et al., 2014; Kuyken et al., 2016a; Zimmerman et al., 2019). A potential explanation for the lack of effect in the current study is that the 8-week MBCT program may not have been sufficient to induce lasting changes in resting-state brain dynamics detectable by our EEG measures. (Gotink et al., 2016) suggest that more prolonged meditation practice might be necessary to observe robust changes in brain function. However, it needs to be highlighted that there could still be a potential modulation of this training across eyes-closed sub-conditions, which this study did not address. Hypothetically, the MBCT training can still have an influence on the separation of active perception—meditation and music sub-conditions—and the passive resting state.

When comparing all sub-conditions – video, music, resting state and meditation – using the microstate analysis, two key patterns emerged. First, the most prominent distinction was between eyes-open (video) and eyes-closed conditions, independent of drug consumption. Second, while the control condition showed distinct GFP patterns between resting state and attentional tasks (meditation and music), these differences largely disappeared under psilocybin. However, the primary influence on whole-brain activity measured by EEG remains whether there is visual input or not, rather than the type of cognitive engagement. This is further supported by the analysis on approximate entropy that found no difference within the video sub-condition when comparing psilocybin with control conditions (Figure 6.5C). Although the limitation of the video sub-condition was highlighted several times in this study, the fundamental difference between eyes opened and eyes closed is well-established, as it is one of the most readily observable phenomena

in raw EEG signals, with the frequency in the visual cortex changing to alpha waves (i.e., 8-13Hz) when eyes are shut (Karakaş et al., 2007). The convergence of resting state with meditation and music conditions under psilocybin points to a fundamental reorganization of attentional states. While sober brain activity shows clear neurophysiological distinctions between passive rest and active engagement with internal experiences (meditation) or external stimuli (music), psilocybin appears to dissolve these boundaries.

The main finding of this study is the effect of psilocybin on brain activity. The microstate analysis found that—for all sub-conditions with closed eyes—Global Field Power (GFP) for B and D microstates, and occurrence and duration for all microstates, significantly differed, such that occurrence increased, duration decreased and GFP decreased under psychedelics (Figure 6.4). The approximate entropy analysis found that most electrodes increase ApEn after the consumption of psychedelics (Figure 6.5). These results suggest that consuming psilocybin disorganizes brain connectivity, aligning with the REBUS theory proposed by (Carhart-Harris and Friston, 2019). However, what cannot be answered with our method is whether the brain becomes more chaotic, or whether there is a higher occurrence of more diverse, but organized, brain activity. The methods implemented only suggest that there is more entropy in time series and more spatiotemporal diversity, not whether these are chaotic or organized.

Decentralized activity produces a smaller EEG signal compared to grouped activity in a single region (Schartner et al., 2017). Thus, the smaller GFP hypothetically indicates less centralized activity. Therefore, the smaller GFP found in two out of the four microstates, further supports the claim that brain activity becomes more decentralized with psilocybin consumption. It should also be noted that for the video comparison, there were only two significant effects: an increased occurrence of the D microstate and a decreased duration of the B microstate. Since psychedelics are mostly known for how they affect visual perception (Kometer et al., 2013), it is difficult to argue that psilocybin would have no effect on brain activity when eyes are open. However, the reason for insignificant results is most likely due to a low explained variance for the video sub-condition compared to the rest, and the smaller number of participants after the cleaning procedure. This leads to taking the results and interpretations on the video sub-condition with more caution.

This study tested several comparisons within the coverage of individual microstates and their directed microstate transitions. Apart from the significant increase in the transition from B to D and a decreased transition from C to A in the resting state during the comparison of psilocybin intake and control, all comparisons of coverage and other connectivity between microstates was insignificant. The reason why coverage did not change could be indicate that psilocybin does not modulate particular brain region's activity locally, but rather these regions that are affected modulate the global neuronal

dynamics. For directional microstate transitions, it has been previously reported that this measure has a low reliability (Kleinert et al., 2024), thus the significant results on directional microstate transition—described in Table S4 but also Table S3—should await further support.

### 6.4.1 Limitations

A significant limitation of this study, is the fixed order of experimental conditions, with the video sub-condition always occurring first, followed by eyes closed resting state, meditation, and music listening. This non-randomized design introduces potential confounds that may impact the interpretation of the results. First, the primacy effect (Murdock Jr, 1962) might have influenced participants' attention and alertness, potentially leading to overestimation of differences between the video and subsequent sub-conditions. Second, carry-over effects (Greenwald, 1976) from one condition to the next cannot be ruled out, particularly considering the temporal dynamics of psilocybin's effects (Carhart-Harris et al., 2011). For instance, the acute effects of psilocybin may have intensified over time, potentially conflating drug effects with condition-specific effects in later trials. Moreover, fatigue or habituation effects (Groves and Thompson, 1970) could have influenced participants' brain states in later conditions, further complicating the interpretation of condition-specific results. (Maneshi et al., 2012) demonstrated that the order of resting-state conditions can significantly impact EEG microstate parameters, underscoring the importance of counterbalancing in future studies. To address this limitation, future research should employ a randomized or counterbalanced design for condition order, or utilize a between-subjects design where different groups undergo different condition orders.

The next limitation—and more severe one when it comes to the video sub-condition—is in the EEG cleaning procedure. From the number of remaining datasets, it is apparent that the video sub-condition contained more artifacts compared to the other sub-conditions. This is expected, as during video watching there are stronger artifacts from blinking and eye movements (Jung et al., 2000). Nevertheless, this led to more removed independent components which could affect the final analysis. This is corroborated by the decreased explained variance using the EEG microstate analysis for the video sub-condition. Additionally, it is surprising that both analyses comparing sub-conditions between video and the rest, and also testing differences between control and psilocybin conditions, led to almost all clusters in GFP, occurrence, and duration differing significantly.

The limitation of the EEG microstates analysis is in the selection of clusters. This study followed the tradition of using 4 clusters (Koenig et al., 1999; Michel and Koenig, 2018) and did not include any sophisticated selection of the 'best' number of clusters,



as this is still a topic of debate and different methods produce different results (see (Nagabhushan Kalburgi et al., 2023)). It is therefore possible that a different number of clusters would lead to different results.

The selection of statistical test forms a consecutive limitation. The automated selection of either performing the non-parametric or parametric test using the p-value from the Shapiro-Wilk test done by (Virtanen et al., 2020) can influence the result. First, one can criticize that there are some datasets above 50 participants, thus for these, this study should have implemented the Kolmogorov-Smirnov normality test instead (Massey Jr, 1951). The reason not to do so was: *i*) the largest number of participants in a distribution was only 60; and *ii*) it provides a universal approach of using only one normality test for the number of datasets implemented here. The next problem with the automated normality assumption is that normality should be visually checked (Ghasemi and Zahediasl, 2012). This is, however, quite difficult for the number of distributions implemented in this study, hence the reason for automation. Lastly, the number of statistical comparisons is quite high, which with the False Discovery Rate (FDR) correction might lead to false negatives (Benjamini and Hochberg, 1995).

### 6.4.2 Future directions

There are significant limits to the analyses performed here that can—and should—be overcome. For the microstates analysis, one can apply the convolutional non-negative matrix factorization method (Mackevicius et al., 2019) to capture approximate general spatiotemporal patterns of brain activity (MacDowell and Buschman, 2020). This is important, as the assumption that microstates are stable has been disputed (Mishra et al., 2020). This novel method can thus create a better temporal resolution of evolving neuronal activity. Regarding ApEn, the issue is that only individual electrodes are considered. A new method could be developed that considers the complexity of a signal across several electrodes, rather than within single electrodes, as the current approach misses valuable information about how distinct parts of the brain are wired together and could potentially show stable connectivity. In conclusion, there are multiple ways to improve the methods in neurosciences by focusing on the spatiotemporal layouts of data, since this is also what devices for humans—EEG, ECoG, fMRI, PET etc.—are able to capture (i.e., measurements of the brain activity across time).



## 6.5 Conclusion

To sum up, this study provides compelling evidence that psilocybin increases the brain's dynamic repertoire, particularly during eyes-closed states, characterized by more frequent transitions between neural configurations and higher temporal complexity. The dual methodological approach—combining microstate analysis with entropy measures—revealed that while mindfulness training had no detectable effect on brain dynamics, the primary determinant of neural activity patterns was the presence or absence of visual input. Notably, psilocybin appeared to dissolve the typical neurophysiological boundaries between passive rest and directed attention state, but only when eyes were closed. These findings align with the REBUS theory's predictions of increased neural entropy under psychedelics, though our methods cannot definitively determine whether this represents a shift toward chaos or toward a more diverse but organized pattern of brain activity. This fundamental reorganization of brain dynamics under psilocybin advances our understanding of how psychedelics affect neural information processing at the millisecond scale.

## Software note

To produce the figures, obtain statistical analyses, and apply the ApEn algorithm see the repository on GitHub: [github.com/filipnovicky/Entropic\\_psilocybin](https://github.com/filipnovicky/Entropic_psilocybin). The PsiConnect dataset is available here after the publication of (Novelli et al., 2025)

## Acknowledgment

This work benefited from the guidance of Adeel Razi's lab regarding the PsiConnect dataset.

## Funding

This work was funded by the Serotonin & Beyond programme (953327 to F.N.)

## Competing interests

Authors declare no competing interests.

# Bibliography

- Barrett, F. S., Carbonaro, T. M., Hurwitz, E., Johnson, M. W., and Griffiths, R. R. (2018). Double-blind comparison of the two hallucinogens psilocybin and dextromethorphan: effects on cognition. *Psychopharmacology*, 235:2915–2927.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300.
- Britz, J., Van De Ville, D., and Michel, C. M. (2010). Bold correlates of eeg topography reveal rapid resting-state network dynamics. *Neuroimage*, 52(4):1162–1170.
- Cannard, C. and Delorme, A. (2022). An open-source eeglab plugin for computing entropy-based measures on meeg signals.
- Carhart-Harris, R., Giribaldi, B., Watts, R., Baker-Jones, M., Murphy-Beiner, A., Murphy, R., Martell, J., Blemings, A., Erritzoe, D., and Nutt, D. J. (2021). Trial of psilocybin versus escitalopram for depression. *New England Journal of Medicine*, 384(15):1402–1411.
- Carhart-Harris, R. L., Erritzoe, D., Williams, T., Stone, J. M., Reed, L. J., Colasanti, A., Tyacke, R. J., Leech, R., Malizia, A. L., Murphy, K., et al. (2011). Neural correlates of the psychedelic state as determined by fmri studies with psilocybin. *Proceedings of the National Academy of Sciences*, 108(5):2138–2143.
- Carhart-Harris, R. L. and Friston, K. J. (2019). Rebus and the anarchic brain: toward a unified model of the brain action of psychedelics. *Pharmacological reviews*, 71(3):316–344.
- Carhart-Harris, R. L., Leech, R., Hellyer, P. J., Shanahan, M., Feilding, A., Tagliazucchi, E., Chialvo, D. R., and Nutt, D. (2014). The entropic brain: a theory of conscious

- states informed by neuroimaging research with psychedelic drugs. *Frontiers in Human Neuroscience*, 8:20.
- Delorme, A. and Makeig, S. (2004). Eeglab: an open source toolbox for analysis of single-trial eeg dynamics including independent component analysis. *Journal of neuroscience methods*, 134(1):9–21.
- Dien, J. (1998). Issues in the application of the average reference: Review, critiques, and recommendations. *Behavior Research Methods, Instruments, & Computers*, 30(1):34–43.
- Ghasemi, A. and Zahediasl, S. (2012). Normality tests for statistical analysis: a guide for non-statisticians. *International journal of endocrinology and metabolism*, 10(2):486.
- Goldberg, S. B., Pace, B. T., Nicholas, C. R., Raison, C. L., and Hutson, P. R. (2020). The experimental effects of psilocybin on symptoms of anxiety and depression: A meta-analysis. *Psychiatry Research*, 284:112749.
- Gotink, R. A., Meijboom, R., Vernooij, M. W., Smits, M., and Hunink, M. M. (2016). 8-week mindfulness based stress reduction induces brain changes similar to traditional long-term meditation practice—a systematic review. *Brain and cognition*, 108:32–41.
- Greenwald, A. G. (1976). Within-subjects designs: To use or not to use? *Psychological Bulletin*, 83(2):314.
- Griffiths, R. R., Johnson, M. W., Richards, W. A., Richards, B. D., Jesse, R., MacLean, K. A., Barrett, F. S., Cosimano, M. P., and Klinedinst, M. A. (2018). Psilocybin-occasioned mystical-type experience in combination with meditation and other spiritual practices produces enduring positive changes in psychological functioning and in trait measures of prosocial attitudes and behaviors. *Journal of psychopharmacology*, 32(1):49–69.
- Groves, P. M. and Thompson, R. F. (1970). Habituation: a dual-process theory. *Psychological review*, 77(5):419.
- Herzog, R., Mediano, P. A., Rosas, F. E., Lodder, P., Carhart-Harris, R., Perl, Y. S., Tagliazucchi, E., and Cofre, R. (2023). A whole-brain model of the neural entropy increase elicited by psychedelic drugs. *Scientific reports*, 13(1):6244.
- HH, J. (1958). The ten-twenty electrode system of the international federation. *Electroenceph clin Neurophysiol*, 10:367–380.
- Jajcay, N. and Hlinka, J. (2023). Towards a dynamical understanding of microstate analysis of m/eeg data. *NeuroImage*, 281:120371.

- Jung, T.-P., Makeig, S., Humphries, C., Lee, T.-W., McKeown, M. J., Iragui, V., and Sejnowski, T. J. (2000). Removing electroencephalographic artifacts by blind source separation. *Psychophysiology*, 37(2):163–178.
- Karakaş, S., Çakmak, E. D., Bekçi, B., and Aydın, H. (2007). Oscillatory responses representing differential auditory processing in sleep. *International Journal of Psychophysiology*, 65(1):40–50.
- Khanna, A., Pascual-Leone, A., Michel, C. M., and Farzan, F. (2015). Microstates in resting-state eeg: current status and future directions. *Neuroscience Biobehavioral Reviews*, 49:105–113.
- Kleinert, T., Koenig, T., Nash, K., and Wascher, E. (2024). On the reliability of the eeg microstate approach. *Brain topography*, 37(2):271–286.
- Koenig, T., Lehmann, D., Merlo, M. C., Kochi, K., Hell, D., and Koukkou, M. (1999). A deviant eeg brain microstate in acute, neuroleptic-naïve schizophrenics at rest. *European archives of psychiatry and clinical neuroscience*, 249:205–211.
- Koenig, T., Prichep, L., Lehmann, D., Sosa, P. V., Braeker, E., Kleinlogel, H., Isenhardt, R., and John, E. R. (2002). Millisecond by millisecond, year by year: normative eeg microstates and developmental stages. *NeuroImage*, 16(1):41–48.
- Kometer, M., Schmidt, A., Bachmann, R., Studerus, E., Seifritz, E., and Vollenweider, F. X. (2013). Psilocybin biases facial recognition, goal-directed behavior, and mood state toward positive relative to negative emotions through different serotonergic subreceptors. *Biological psychiatry*, 74(8):636–644.
- Kuyken, W., Hayes, R., Barrett, B., Byng, R., Dalgleish, T., Kessler, D., Lewis, G., Watkins, E., Brejcha, C., Cardy, J., et al. (2016a). Effectiveness and cost-effectiveness of mindfulness-based cognitive therapy compared with maintenance antidepressant treatment in the prevention of depressive relapse or recurrence (prevent): a randomised controlled trial. *The Lancet*, 386(9988):63–73.
- Kuyken, W., Warren, F. C., Taylor, R. S., Whalley, B., Crane, C., Bondolfi, G., Hayes, R., Huijbers, M., Ma, H., Schweizer, S., et al. (2016b). Efficacy of mindfulness-based cognitive therapy in prevention of depressive relapse: an individual patient data meta-analysis from randomized trials. *JAMA psychiatry*, 73(6):565–574.
- Lebedev, A. V., Kaelen, M., Lövdén, M., Nilsson, J., Feilding, A., Nutt, D. J., and Carhart-Harris, R. L. (2016). Lsd-induced entropic brain activity predicts subsequent personality change. *Human brain mapping*, 37(9):3203–3213.

- MacCoon, D. G., MacLean, K. A., Davidson, R. J., Saron, C. D., and Lutz, A. (2014). No sustained attention differences in a longitudinal randomized trial comparing mindfulness based stress reduction versus active control. *PloS one*, 9(6):e97551.
- MacDowell, C. J. and Buschman, T. J. (2020). Low-dimensional spatiotemporal dynamics underlie cortex-wide neural activity. *Current Biology*, 30(14):2665–2680.
- Mackevicius, E. L., Bahle, A. H., Williams, A. H., Gu, S., Denisenko, N. I., Goldman, M. S., and Fee, M. S. (2019). Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *Elife*, 8:e38471.
- Maneshi, M., Moeller, F., Fahoum, F., Gotman, J., and Grova, C. (2012). Resting-state connectivity of the sustained attention network correlates with disease duration in idiopathic generalized epilepsy. *PloS one*, 7(12):e50359.
- Massey Jr, F. J. (1951). The kolmogorov-smirnov test for goodness of fit. *Journal of the American statistical Association*, 46(253):68–78.
- McCulloch, D. E.-W., Olsen, A. S., Ozenne, B., Stenbaek, D. S., Armand, S., Madsen, M. K., Knudsen, G. M., and Fisher, P. M. (2023). Navigating the chaos of psychedelic neuroimaging: A multi-metric evaluation of acute psilocybin effects on brain entropy. *medRxiv*, pages 2023–07.
- Mediano, P. A., Rosas, F. E., Timmermann, C., Roseman, L., Nutt, D. J., Feilding, A., Kaelen, M., Kringelbach, M. L., Barrett, A. B., Seth, A. K., et al. (2024). Effects of external stimulation on psychedelic state neurodynamics. *ACS Chemical Neuroscience*, 15(3):462–471.
- Michel, C. M. and Koenig, T. (2018). Eeg microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: a review. *Neuroimage*, 180:577–593.
- Mishra, A., Englitz, B., and Cohen, M. X. (2020). Eeg microstates as a continuous phenomenon. *Neuroimage*, 208:116454.
- Murdock Jr, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, 64(5):482.
- Murray, M. M., Brunet, D., and Michel, C. M. (2008). Topographic erp analyses: a step-by-step tutorial review. *Brain topography*, 20(4):249–264.
- Nagabhushan Kalburgi, S., Kleinert, T., Aryan, D., Nash, K., Schiller, B., and Koenig, T. (2023). Microstatelab: the eeglab toolbox for resting-state microstate analysis. *Brain topography*, pages 1–25.

- Nichols, D. E. (2016). Psychedelics. *Pharmacological reviews*, 68(2):264–355.
- Novelli, L., Stoliker, D., Barta, T., Greaves, M. D., Chopra, S., Jackson, J., Kwee, J., Williams, M., and Razi, A. (2025). Psiconnect: A multimodal neuroimaging study of psilocybin-induced changes in brain and behaviour. *bioRxiv*, pages 2025–04.
- Pascual-Marqui, R. D., Michel, C. M., and Lehmann, D. (1995). Segmentation of brain electrical activity into microstates: model estimation and validation. *IEEE Transactions on Biomedical Engineering*, 42(7):658–665.
- Pincus, S. M. (1991). Approximate entropy as a measure of system complexity. *Proceedings of the National Academy of Sciences*, 88(6):2297–2301.
- Pion-Tonachini, L., Kreutz-Delgado, K., and Makeig, S. (2019). Iclabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, 198:181–197.
- Qiu, J., Gong, Y., Zhang, X., and Mao, W. (2024). Effectiveness of mindfulness-based cognitive therapy on depressive symptoms, brain potential, and neuroimmunoinflammatory factors in depressed patients. *Clinical neuropharmacology*, 47(4):128–133.
- Savino, A. and Nichols, C. D. (2022). Lysergic acid diethylamide induces increased signalling entropy in rats’ prefrontal cortex. *Journal of neurochemistry*, 162(1):9–23.
- Schartner, M. M., Carhart-Harris, R. L., Barrett, A. B., Seth, A. K., and Muthukumaraswamy, S. D. (2017). Increased spontaneous meg signal diversity for psychoactive doses of ketamine, lsd and psilocybin. *Scientific reports*, 7(1):1–12.
- Shinozuka, K., Tewarie, P., Luppi, A., Lynn, C., Roseman, L., Muthukumaraswamy, S. D., Nutt, D. J., Carhart-Harris, R., Deco, G., and Kringelbach, M. L. (2024). Lsd flattens the hierarchy of directed information flow in fast whole-brain dynamics. *bioRxiv*, pages 2024–04.
- Siegel, J. S., Subramanian, S., Perry, D., Kay, B. P., Gordon, E. M., Laumann, T. O., Reneau, T. R., Metcalf, N. V., Chacko, R. V., Gratton, C., et al. (2024). Psilocybin desynchronizes the human brain. *Nature*, 632(8023):131–138.
- Singleton, S. P., Luppi, A. I., Carhart-Harris, R. L., Cruzat, J., Roseman, L., Deco, G., Kringelbach, M. L., Stamatakis, E. A., and Kuceyeski, A. (2021). Lsd flattens the brain’s energy landscape: evidence from receptor-informed network control theory. *bioRxiv*, pages 2021–05.

- Szumska, I., Gola, M., Rusanowska, M., Krajewska, M., Żygierewicz, J., Krejtz, I., Nezlek, J. B., and Holas, P. (2021). Mindfulness-based cognitive therapy reduces clinical symptoms, but do not change frontal alpha asymmetry in people with major depression disorder. *International Journal of Neuroscience*, 131(5):453–461.
- Tagliazucchi, E., Carhart-Harris, R., Leech, R., Nutt, D., and Chialvo, D. R. (2014). Enhanced repertoire of brain dynamical states during the psychedelic experience. *Human brain mapping*, 35(11):5442–5456.
- Teasdale, J. D., Segal, Z. V., Williams, J. M. G., Ridgeway, V. A., Soulsby, J. M., and Lau, M. A. (2000). Prevention of relapse/recurrence in major depression by mindfulness-based cognitive therapy. *Journal of consulting and clinical psychology*, 68(4):615.
- Viol, A., Palhano-Fontes, F., Onias, H., de Araujo, D. B., Hövel, P., and Viswanathan, G. M. (2019). Characterizing complex networks using entropy-degree diagrams: unveiling changes in functional brain connectivity induced by ayahuasca. *Entropy*, 21(2):128.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.
- Vollenweider, F. X. and Preller, K. H. (2020). The neurobiology of psychedelic drugs: implications for the treatment of mood disorders. *Nature Reviews Neuroscience*, 21(11):611–624.
- Von Wegner, F., Knaut, P., and Laufs, H. (2018). Eeg microstate sequences from different clustering algorithms are information-theoretically invariant. *Frontiers in computational neuroscience*, 12:70.
- Williams, M. J., McManus, F., Muse, K., and Williams, J. M. G. (2011). Mindfulness-based cognitive therapy for severe health anxiety (hypochondriasis): An interpretative phenomenological analysis of patients’ experiences. *British Journal of Clinical Psychology*, 50(4):379–397.

- Yang, J., Lei, D., Suo, X., Tallman, M. J., Qin, K., Li, W., Bruns, K. M., Blom, T. J., Duran, L. R. P., Cotton, S., et al. (2022). A preliminary study of the effects of mindfulness-based cognitive therapy on structural brain networks in mood-dysregulated youth with a familial risk for bipolar disorder. *Early Intervention in Psychiatry*, 16(9):1011–1019.
- ZanESCO, A. P. (2024). Normative temporal dynamics of resting eeg microstates. *Brain Topography*, 37(2):243–264.
- Zimmerman, B., Finnegan, M., Paul, S., Schmidt, S., Tai, Y., Roth, K., Chen, Y., and Husain, F. T. (2019). Functional brain changes during mindfulness-based cognitive therapy associated with tinnitus severity. *Frontiers in Neuroscience*, 13:747.



## 6.6 Supplementary materials

### 6.6.1 Demographics of participants

Table S1 provides a detailed description of participants from the PsiConnect dataset.

### 6.6.2 Details on Approximate Entropy

Below the pseudocode for analyzing the Approximate Entropy of the psychedelics influence is provided:

---

**Algorithm S1** Approximate Entropy Calculation
 

---

```

1: procedure APEN(datasets,  $m$ ,  $r$ ,  $f_s$ , windowSize)
2:   Input: datasets (list of EEG data matrices),  $m = 2$ ,  $r = 0.2$ ,  $f_s = 125$ ,
      window Size = 5
3:   Initialize all ApEn Values
4:   for each dataset in datasets do
5:     for each window in dataset do
6:       Normalize window data
7:       for  $i = 1$  to 2 do
8:         Initialize template matrices of length  $m + i - 1$ 
9:         Compute Chebyshev distances between templates
10:        Count matches within tolerance  $r \cdot STD_{electrode}$ 
11:        Calculate logarithmic likelihoods
12:      end for
13:      Compute window ApEn as the difference of likelihoods
14:    end for
15:    Average ApEn values across windows for each channel
16:    Append matched ApEn values to all ApEn Values
17:  end for
18:  return all ApEn Values
19: end procedure

```

---

Where  $m$  is the embedding dimension (set to 2),  $r$  is the tolerance (0.2 times the standard deviation of each window),  $f_s$  is the sampling frequency in Hz (125Hz in this study), and *window Size* is the size of the sliding window in seconds (default: 5s). Except for the sampling frequency, these parameters are selected by a researcher, and follow from previous literature (Pincus, 1991). A template vector is a subsequence of the time series data, representing  $m$  consecutive time points across all EEG channels for a given window. The template vector is then repeated for  $m + 1$  time points. For visualization of what exactly ApEn is doing, the reader can visit this interactive website:

Characteristic	Med. (N=33)	Non-med. (N=30)
Age: Mean (SD)	35.52(1064)	40.03(1085)
Range	20 - 54	19 - 53
<b>Gender</b>		
Female	14 (42.42%)	16 (53.33%)
Male	18 (54.55%)	14 (46.67%)
Non-binary	1 (3.04%)	0 (0.00%)
<b>Ethnicity</b>		
Arab	0 (0.00%)	1 (3.33%)
Asian	7 (21.21%)	6 (20.00%)
Asian/White	1 (3.03%)	0 (0.00%)
Hispanic	0 (0.00%)	2 (6.67%)
White	25 (75.76%)	21 (70.00%)
<b>Handedness</b>		
Ambidextrous	0 (0.00%)	1 (3.33%)
Left	2 (6.06%)	2 (6.67%)
Right	31 (93.94%)	27 (90.00%)
<b>Highest education</b>		
Highschool	0 (0.00%)	1 (3.33%)
Postgrad	10 (30.30%)	14 (46.67%)
Postgrad (incomp.)	6 (18.18%)	2 (6.67%)
Tertiary dipl/cert	0 (0.00%)	2 (6.67%)
Tertiary (incomp.)	0 (0.00%)	2 (6.67%)
Trade/tech	1 (3.03%)	1 (3.33%)
Trade & tertiary	0 (0.00%)	1 (3.33%)
Undergrad	10 (30.30%)	5 (16.67%)
Undergrad (incomp.)	5 (15.15%)	4 (13.33%)
<b>Psych. Use (3y)</b>		
No	30 (90.91%)	29 (96.67%)
Yes	3 (9.09%)	1 (3.33%)
<b>Med. Exp. (3y)</b>		
No	23 (69.70%)	17 (56.67%)
Yes	10 (30.30%)	13 (43.33%)

Table S1: Detailed description of the participants.

<https://github.com/filipnovicky/ApEn-Visualization>.

The logarithmic likelihood calculation transforms the multiplicative probabilities into additive quantities, making the calculations more numerically stable and better suited for handling very small probability values often encountered in EEG analysis. By comparing sequences of length  $m$  with  $m + 1$ , we can quantify how predictable the signal is - if patterns matching at length  $m$  continue to match at length  $m + 1$ , this indicates greater regularity in the signal. The difference between these likelihoods provides a measure of the signal's complexity, with larger differences suggesting more complex or chaotic signals and smaller differences indicating more regular patterns. To identify matching patterns, the algorithm employs Chebyshev distance between template vectors, calculated as the maximum absolute difference across any time point between two templates. Lastly, the algorithm returns a vector containing the mean ApEn value for each EEG channel, averaged across all windows. The code and the final output are included in the repository.

### 6.6.3 Statistical comparisons

Table S2 shows p-values for comparing participants who underwent the MBCT training with those who did not. The False Discovery Rate (FDR) correction was applied both within the control and psilocybin conditions independently, hence the FDR corrected 120 tests twice.

Table S3 compares sub-conditions among each other, both within control condition and psilocybin condition. In here, FDR was applied to all statistical comparisons, thus 360 tests underwent this procedure.

Lastly, Table S4 shows FDR values for comparison of drug conditions. This analysis used the FDR correction for all 120 tests.

Microstate	Metric	Control				Psilocybin			
		video	meditation	music	resting state	video	meditation	music	resting state
A	Mean Occurrence	1.0000	1.0000	1.0000	1.0000	0.8089	0.7309	0.9781	0.7309
A	Mean Duration	1.0000	1.0000	1.0000	1.0000	0.9892	0.7392	0.8189	0.7309
A	Mean GFP	1.0000	1.0000	1.0000	1.0000	0.8479	0.8189	0.9892	0.8391
A	Coverage	1.0000	1.0000	1.0000	1.0000	0.9907	0.7580	0.9892	0.9892
B	Mean Occurrence	1.0000	1.0000	1.0000	1.0000	0.9892	0.9892	0.9984	0.9892
B	Mean Duration	1.0000	1.0000	1.0000	1.0000	0.8391	0.7309	0.8189	0.9781
B	Mean GFP	1.0000	1.0000	1.0000	1.0000	0.9727	0.7392	0.9103	0.8189
B	Coverage	1.0000	1.0000	1.0000	1.0000	0.8089	0.3049	0.7309	0.8370
C	Mean Occurrence	1.0000	1.0000	1.0000	1.0000	0.9955	0.7392	0.9781	0.7309
C	Mean Duration	1.0000	1.0000	1.0000	1.0000	0.9781	0.8132	0.9892	0.8396
C	Mean GFP	1.0000	1.0000	1.0000	1.0000	0.9892	0.8189	0.8633	0.8189
C	Coverage	1.0000	1.0000	1.0000	1.0000	0.9892	0.8370	0.9892	0.7309
D	Mean Occurrence	1.0000	1.0000	1.0000	1.0000	0.9892	0.7309	0.7309	0.8396
D	Mean Duration	1.0000	1.0000	1.0000	1.0000	0.9892	0.8370	1.0000	0.8370
D	Mean GFP	1.0000	1.0000	1.0000	1.0000	0.9955	0.8189	0.9892	0.8370
D	Coverage	1.0000	1.0000	1.0000	1.0000	0.9907	0.7580	0.7309	0.8370
Microstate Transitions									
-	A→B	1.0000	1.0000	1.0000	1.0000	0.8370	0.8396	0.8370	0.7309
-	A→C	1.0000	1.0000	1.0000	1.0000	0.7309	0.7309	0.9892	0.7392
-	A→D	1.0000	1.0000	1.0000	1.0000	0.8089	0.7309	0.7309	0.8479
-	B→A	1.0000	1.0000	1.0000	1.0000	0.8370	0.8089	0.9815	0.7309
-	B→C	0.6253	1.0000	1.0000	1.0000	0.8396	0.3013	0.7309	0.7392
-	B→D	1.0000	1.0000	1.0000	1.0000	0.8132	0.8089	0.8370	0.7309
-	C→A	1.0000	1.0000	1.0000	1.0000	0.7309	0.9586	0.9892	0.8089
-	C→B	1.0000	1.0000	1.0000	1.0000	0.8189	0.3013	0.8089	0.9103
-	C→D	1.0000	1.0000	1.0000	1.0000	0.9892	0.7309	0.8089	0.7309
-	D→A	1.0000	1.0000	1.0000	1.0000	0.8189	0.7309	0.7309	0.8089
-	D→B	1.0000	1.0000	1.0000	1.0000	0.7309	0.8370	0.9733	0.7309
-	D→C	1.0000	1.0000	1.0000	1.0000	0.8633	0.7309	0.9892	0.7309
All Groups									
-	Mean Occurrence for All	1.0000	1.0000	1.0000	1.0000	0.9781	0.9907	0.9892	0.9955
-	Mean Duration for All	1.0000	1.0000	1.0000	1.0000	0.9781	0.9907	0.9781	0.9892

(Continued on next page)

(Continued from previous page)

Group	Metric	Control				PsiLocybin			
		video	meditation	music	resting state	video	meditation	music	resting state

**Table S2: FDR corrected p-values of all the comparisons within four sub-conditions of MBCT and non-MBCT groups:** For this comparison of MBCT or non-MBCT, either non-parametric Mann-Whitney U (MWU) or parametric independent t-test were used for the comparisons, where the choice depended on how normal a given distribution is, which was measured by the Shapiro-Wilk test. Metrics for Control and PsiLocybin have been FDR corrected separately. Note that the reason why there are so many 1 values in the control is not because the datasets are identical but because of the FDR correction.

Group	Metric	Control						Psilocybin							
		V vs Me	V vs Mu	V vs RS	RS vs Me	Mu vs RS	Mu vs Me	V vs Me	V vs Mu	V vs RS	RS vs Me	Mu vs RS	Mu vs Me		
A	Mean Occurrence	>0.0001↑	>0.0001↑	>0.0001↑	0.9861	0.2907	0.3021	0.0008↑	0.0011↑	0.0187↑	0.1575	0.1987	0.6902		
	Mean Duration	>0.0001↓	>0.0001↓	>0.0001↓	0.0267↑	0.0020↓	0.6000	>0.0001↓	>0.0001↓	>0.0001↓	0.2718	0.8040	0.6854		
	Mean GFP	>0.0001↓	>0.0001↓	>0.0001↓	>0.0001↑	>0.0001↓	0.9861	0.0094↓	0.0139↓	0.0060↓	0.3881	0.0024↑	0.3930		
	Coverage	0.3331	0.1588	0.5220	0.5607	0.5366	0.8319	0.5220	0.4229	0.9008	0.8475	0.5607	0.7644		
B	Mean Occurrence	>0.0001↑	>0.0001↑	>0.0001↑	0.4229	0.0647	0.0594	0.0006↑	0.0172↑	0.0809	0.0496↑	0.2609	0.1548		
	Mean Duration	>0.0001↓	>0.0001↓	>0.0001↓	0.5452	0.5079	0.4399	>0.0001↓	>0.0001↓	0.0006↓	0.8319	0.5386	0.7644		
	Mean GFP	>0.0001↓	>0.0001↓	0.0007↑	0.0007↑	0.0172↓	0.8783	>0.0001↓	>0.0001↓	0.0002↓	0.3279	0.7644	0.6124		
	Coverage	0.5078	0.5400	0.2907	0.7427	0.5146	0.7645	0.8058	0.2410	0.2771	0.3061	0.8319	0.4680		
C	Mean Occurrence	>0.0001↑	>0.0001↑	>0.0001↑	0.5508	0.2712	0.9072	0.0001↑	0.0016↑	0.0001↑	0.2907	0.6293	0.7644		
	Mean Duration	>0.0001↓	>0.0001↓	>0.0001↓	0.8040	0.2314	0.2907	>0.0001↓	0.0001↓	0.0007↓	0.4474	0.8058	0.4705		
	Mean GFP	>0.0001↓	>0.0001↓	>0.0001↓	>0.0001↑	>0.0001↓	0.3707	0.0001↓	>0.0001↓	>0.0001↓	0.0947	0.1646	0.7923		
	Coverage	0.5166	0.6551	0.5836	0.8752	0.9400	0.6142	0.6008	0.8765	0.7640	0.9854	0.9235	0.8058		
D	Mean Occurrence	>0.0001↑	>0.0001↑	>0.0001↑	0.4774	0.0854	0.5386	>0.0001↑	>0.0001↑	>0.0001↑	0.9639	0.9427	0.8040		
	Mean Duration	>0.0001↓	>0.0001↓	>0.0001↓	0.4418	0.0125↓	0.0606	>0.0001↓	>0.0001↓	0.0002↓	0.3707	0.6551	0.5451		
	Mean GFP	>0.0001↓	>0.0001↓	>0.0001↓	>0.0001↑	>0.0001↓	0.2496	>0.0001↓	>0.0001↓	>0.0001↓	0.2454	0.6124	0.8494		
	Coverage	0.0015↓	0.3661	0.2503	0.8040	0.9455	0.5220	0.8494	0.8319	0.7644	0.7644	0.9280	0.7842		
Microstate Transitions															
All Groups	A→B	0.0222↑	0.0699	0.0666	0.5400	0.8319	0.3279	0.7948	0.9861	0.4090	0.1485	0.4090	0.1485	0.7644	0.1157
	A→C	0.5366	0.9008	0.6922	0.8123	0.6125	0.4874	0.8495	0.6944	0.8333	0.8708	0.8333	0.8708	0.7644	0.9978
	A→D	0.3908	0.8058	0.8378	0.3144	0.3507	0.7644	0.3220	0.2144	0.2771	0.3794	0.2771	0.3794	0.4091	0.8233
	B→A	0.0149↑	0.3020	0.1136	0.4942	0.9400	0.5950	0.3218	0.7782	0.3661	0.2161	0.3661	0.2161	0.6945	0.4091
	B→C	0.7645	0.9749	0.3692	0.5679	0.8058	0.8059	0.5452	0.2907	0.3165	0.1646	0.3165	0.1646	0.3678	0.3556
	B→D	0.3297	0.9520	0.3370	0.9149	0.3997	0.7645	0.9535	0.9209	0.3592	0.1848	0.3592	0.1848	0.3386	0.9377
	C→A	0.8054	0.9028	0.3278	0.5200	0.8319	0.5453	0.9234	0.775	0.9678	0.9639	0.9678	0.9639	0.3603	0.9303
	C→B	0.9025	0.9723	0.3275	0.2927	0.5346	0.8494	0.5332	0.7832	0.7369	0.7644	0.7369	0.7644	0.6009	0.9978
	C→D	0.1322	0.6551	0.2699	0.9078	0.8319	0.5158	0.8319	0.775	0.3003	0.4117	0.3003	0.4117	0.9817	0.6098
	D→A	0.8764	0.8764	0.9639	0.5643	0.8319	0.9073	0.3279	0.0664	0.3001	0.9881	0.3001	0.9881	0.8708	0.5386
	D→B	0.5494	0.8282	0.6878	0.6008	0.8494	0.4774	0.5220	0.9861	0.9639	0.9250	0.9639	0.9250	0.7644	0.3912
	D→C	0.0030↓	0.2245	0.0873	0.6728	0.9360	0.4427	0.7645	0.6944	0.9303	0.5678	0.9303	0.5678	0.4091	0.4724

**Table S3: FDR corrected p-values of sub-condition comparisons among each other both in control and psilocybin:** This table provides all statistical comparisons (calculated either by non-parametric Wilcoxon or parametric t-test, both used for dependent samples. Here, all the values ( $30 \cdot 6 \cdot 2 = 360$ ) were FDR corrected. V = Video; Me = Meditation; Mu = Music; RS = Resting State.

Microstate	Metric	Video	Music	Meditation	Resting state
A	Mean Occurrence	0.3020	<b>0.0057</b> ↑	<b>0.0011</b> ↑	<b>0.0001</b> ↑
A	Mean Duration	0.0517	<b>0.0001</b> ↓	<b>0.0001</b> ↓	> <b>0.0001</b> ↓
A	Mean GFP	0.5342	0.0963	0.0572	<b>0.0011</b> ↓
A	Coverage	0.9968	0.8433	0.8643	0.2503
B	Mean Occurrence	0.6217	<b>0.0033</b> ↑	> <b>0.0001</b> ↑	<b>0.0002</b> ↑
B	Mean Duration	<b>0.0401</b> ↓	<b>0.0056</b> ↓	<b>0.0001</b> ↓	<b>0.0013</b> ↓
B	Mean GFP	0.6217	<b>0.0401</b> ↓	<b>0.0072</b> ↓	<b>0.0023</b> ↓
B	Coverage	0.1966	0.9615	0.8643	0.0504
C	Mean Occurrence	0.0761	<b>0.0003</b> ↑	<b>0.0001</b> ↑	<b>0.0002</b> ↑
C	Mean Duration	0.7545	<b>0.0002</b> ↓	<b>0.0001</b> ↓	> <b>0.0001</b> ↓
C	Mean GFP	0.8643	0.0651	0.0945	0.0651
C	Coverage	0.2052	0.8643	0.8718	0.8643
D	Mean Occurrence	<b>0.0059</b> ↑	<b>0.0001</b> ↑	<b>0.0007</b> ↑	> <b>0.0001</b> ↑
D	Mean Duration	0.5634	> <b>0.0001</b> ↓	<b>0.0002</b> ↓	> <b>0.0001</b> ↓
D	Mean GFP	0.9192	<b>0.0187</b> ↓	<b>0.0288</b> ↓	<b>0.0002</b> ↓
D	Coverage	0.1451	0.8433	0.8643	0.8835
Microstate Transitions					
	A→B	0.3998	0.6217	0.9060	0.3354
	A→C	0.9615	0.9192	0.6785	0.5195
	A→D	0.3317	0.6217	0.9615	0.3791
	B→A	0.1589	0.6217	0.8643	0.8835
	B→C	0.9916	0.8643	0.8433	0.2503
	B→D	0.3998	0.9236	0.9442	<b>0.0245</b> ↑
	C→A	0.8835	0.7688	0.9615	<b>0.0256</b> ↓
	C→B	0.9940	0.9060	0.8433	0.8643
	C→D	0.1577	0.8433	0.8835	0.9192
	D→A	0.3998	0.9060	0.8896	0.1966
	D→B	0.8643	0.6939	0.5586	0.1308
	D→C	0.2671	0.9236	0.8918	0.9060
All Groups					
	Mean Occurrence for All	0.0963	<b>0.0033</b> ↑	<b>0.0009</b> ↑	<b>0.0108</b> ↑
	Mean Duration for All	0.0676	<b>0.0055</b> ↓	<b>0.0020</b> ↓	<b>0.0055</b> ↓

(Continued on next page)

(Continued from previous page)

Group	Metric	VIDEO	MUSIC	GM	RS
<b>Table S4: FDR corrected p-values for psilocybin vs control data comparison across sub-conditions:</b> Using either non-parametric or parametric statistical tests for dependent samples, this table shows the final p-values of such comparisons, between control and psilocybin conditions.					



# Chapter 7

## General Discussion

In this thesis, we explored how neuromodulators shape perception and behaviour through computational mechanisms. Particularly, we aim to establish serotonin as a biological modulator of precision estimates articulated within the active inference framework, demonstrating how this neurotransmitter influences neural processing during exploration, perception, and psychedelic use. We first investigated serotonergic effects on active sensing in rodents (Chapter 2), showing that chronically elevated serotonin levels lead to reduced estimated sensory precision, suppressed exploratory behaviour, and greater reliance on prior beliefs. We then extended this computational framework to humanoid robotics (Chapter 3), before exploring how precision modulation explains body ownership illusions (Chapter 4) and attention mechanisms (Chapter 5). Finally, we examined how psilocybin's effects on neural dynamics (Chapter 6) align with our precision-based account, providing converging evidence that serotonin fundamentally regulates the balance between sensory evidence and prior expectations across diverse cognitive contexts.

In chapter 2, we investigated how serotonergic signaling affects active sensing by comparing whisking behaviour in  $SERT^{-/-}$  and wild-type animals. Two competing computational hypotheses about how this altered serotonergic signaling might affect precision parameters within an active inference framework were tested: either through modulation of prior precision or sensory precision. Our simulations supported the hypothesis that  $SERT^{-/-}$  animals exhibit reduced sensory precision, resulting in less exploratory whisking behaviour and greater reliance on prior beliefs. This finding suggests that heightened serotonergic tone diminishes the influence of incoming sensory information on perceptual inference, resulting in a computational phenotype that favors existing models over new sensory evidence. The same theoretical principle was further demonstrated

in Chapter 3, where the same model was deployed in a humanoid. However, instead of exploring using whiskers, the humanoid robot explored the space using a hand. This work thus aimed to showcase that the demonstrated serotonergic function in exploration might be considered in other species, such as humans.

The computational account of precision modulation in active sensing connects to the work from Chapter 5 on rhythmic attention through the embodied nature of inference (Fiebelkorn and Kastner, 2019). When organisms actively sample their environment through movement (e.g., whisking or eye movement), they must differentiate self-generated sensations from external stimuli to avoid misinterpretation. In this Chapter, we speculated that the brain accomplishes this through rhythmically switching levels of sensory precision, effectively implementing sensory attenuation during self-generated movement. This mechanism represents a specific application we identified in attentional processes, such as sensory attenuation (Brown et al., 2013) and gain control (Hillyard et al., 1998). In both contexts, precision acts as a control parameter that dynamically allocates computational resources between prior beliefs and sensory evidence, with serotonin playing a key modulatory role in this balancing act.

The REBUS theory has been tested and finally supported in Chapter 6. This study focused on analyzing the influence of psilocybin on the spatiotemporal organization of brain activity. Using EEG microstates analysis (Michel and Koenig, 2018) and quantifying approximate entropy (Pincus, 1991), this study reported that the brain is more entropic and dynamic after the psilocybin consumption. Faster EEG microstates switches and elevated approximate entropy can be interpreted as a decrease in prior precision, thus increasing the flow of new information into the brain with more importance.

## 7.1 Complementary serotonergic receptors

Serotonin exerts its diverse effects through at least 14 distinct receptor subtypes, grouped into seven families (5HT1 through 5HT7), each with unique signaling mechanisms and functional roles (Barnes and Sharp, 1999; Hoyer et al., 2002). These receptors differ not only in their molecular structure and downstream signaling pathways, but also in their anatomical distribution and physiological effects (Pytliak et al., 2011). The 5HT1 family, which includes subtypes like 5HT1A, generally couples to inhibitory G-proteins ( $G_i/G_o$ ), reducing cellular excitability and cyclic AMP levels (Raymond et al., 2001). In contrast, the 5HT2 family, including 5HT2A receptors, typically activates excitatory  $G_{q/11}$  proteins, leading to increased neuronal activity through phospholipase C activation (Berg et al., 2005). This fundamental difference in signaling mechanisms suggests that different serotonin receptors may have opposing effects on neural computation, including

precision modulation.

Understanding these receptor-specific mechanisms is crucial for reconciling the seemingly contradictory findings presented in this thesis. While both chapters 2 and 6 implicate serotonin in precision modulation within the active inference framework, they reveal opposite directional effects: increased prior precision in  $SERT^{-/-}$  animals versus decreased prior precision following psilocybin administration. This apparent paradox may be resolved by considering which specific serotonin receptors mediate these effects and how their distinct signaling properties translate to computational parameters like precision weighting.

Chapters 2 and 6, although targeting the same neuromodulator and studied via the lens of active inference, differed in the study design. First, one study is in animals, the other in humans. Second, the latter work was studied while controlling the psychedelics consumption dose, thus targeting the 5HT2A receptor. However, the former affected the serotonergic system from animals' birth (as the  $SERT^{-/-}$  are genetically mutated). Lastly, humans were in a calm, passive state, while animals moved and explored freely.

These chapters highlight that serotonin modulates prior precision, but in different directions. Meaning, the psilocybin consumption decreases the priors, while the priors increase in the  $SERT^{-/-}$  animals. In this section, we will speculate how these seemingly paradoxical results can be reconciled by distinct receptors. First of all, a stimulation of distinct serotonergic receptors show distinct outputs (Pytliak et al., 2011), questioning the universality of the computational approaches to serotonin specifically (Doya et al., 2021; Azmitia, 2020; Daw et al., 2002; Crockett et al., 2009). By only focusing on the two most described and most influential serotonergic receptors we know of so far—5HT1A and 5HT2A—we can unite the two chapters and their opposite findings.

The REBUS theory has made a strong case for explaining the 5HT2A receptors as a biological mechanism of relaxing the priors by decreasing its precision. This is clearly aligned with our findings in the psilocybin study (Chapter 6), where we observed increased brain entropy and dynamics. Meanwhile, Chapter 2 demonstrated that  $SERT^{-/-}$  animals exhibit reduced exploratory whisking behaviour that could be modeled computationally as decreased sensory precision, leading to increased reliance on priors, which is an effect that operates in the opposite direction.

Carhart-Harris and Nutt (2017) proposed a model where 5HT1A and 5HT2A receptors serve complementary but opposing functions: 5HT1A receptors are involved in passive coping by mediating stress reduction and promoting resilience, while 5HT2A receptors facilitate active coping through enhanced plasticity and adaptability. According to this theory, 5HT1A receptor signaling represents the brain's default response to moderate stress, while 5HT2A signaling enables major behavioural changes when significant adaptation is

required.

Importantly, extensive research has demonstrated that 5HT1A receptors show pronounced changes in  $SERT^{-/-}$  animals. These receptors exhibit marked desensitization with a significantly reduced sensitivity to agonist activation, compared to wildtype controls Araragi et al. (2013). This desensitization is more severe and consistent than changes observed in other receptor subtypes, where 5HT2A and 5HT2C receptors show variable, region-specific alterations Li et al. (2003). Furthermore,  $SERT^{-/-}$  mice display enhanced 5HT1A receptor-dependent feedback control, with several-fold greater disinhibition of dorsal raphe neurons when 5HT1A receptors are blocked Soiza-Reilly et al. (2015).

Since we can confidently attribute psilocybin's effects to 5HT2A receptor activation, and these effects operate in the opposite direction to what we observed in  $SERT^{-/-}$  animals, this suggests that exploratory behaviour in  $SERT^{-/-}$  animals may be primarily influenced by 5HT1A receptor mechanisms. While SERT knockout affects all serotonergic receptors due to globally increased extracellular serotonin, the predominant desensitization of 5HT1A receptors combined with the opposing precision effects we observed computationally align with the opposing functions attributed to 5HT1A versus 5HT2A receptors in Carhart-Harris and Nutt (2017) model. Additionally, it was found that the 5HT1A receptor is the most abundant in the mammalian brain, thus both in humans and rodents Barnes and Sharp (1999). This further supports the possibility that 5HT1A mediated effects might predominate in shaping the exploratory behaviour of  $SERT^{-/-}$  animals. Our reasoning provides a mathematical underpinning of the theory of the two serotonergic receptors by Carhart-Harris and Nutt (2017).

It is crucial to note that this receptor-specific interpretation is a theoretical framework rather than an established mechanism. The  $SERT^{-/-}$  model is far more complex than can be captured by single receptor dynamics. These animals experience lifelong serotonin elevation from birth, leading to extensive developmental alterations including changes in glutamate/GABA signaling and compensatory adaptations across multiple receptor systems Homberg et al. (2010); Ansorge et al. (2004). While studies have documented altered 5HT1A receptor expression and function in  $SERT^{-/-}$  rats (including receptor desensitization and reduced autoreceptor function) Fabre et al. (2000); Li et al. (1999), and have shown that 5HT1A receptors are more severely affected than other serotonin receptor subtypes Holmes et al. (2003), no research has definitively attributed their behavioural phenotype to any single receptor subtype. The computational patterns we observe provide one possible lens for organizing these complex effects, but should not be taken as evidence for simple receptor-specific mechanisms. Rather, they highlight how precision-based frameworks might help generate testable hypotheses about the diverse manifestations of serotonergic dysfunction, for instance about the function of 5HT1A in

SERT<sup>-/-</sup>.

Having strengthened the case that serotonin is central to regulating the sensory precision parameter, how can this idea be united with Chapters 4 and 5, where we demonstrated the function of precision in body ownership and account for the limitations of precision in action-perception cycle?

## 7.2 Precision and body ownership

Precision is so well integrated in the neuronal function, and is essential to regulate uncertainty in the environment. In addition, it has been reported to be particularly relevant when treating illusions (Brown et al., 2013; Deneve and Jardri, 2016). This provides a motivation for investigating precision estimation in bodily illusions, specifically.

As a proof of principle that the precision is fundamental in body illusion, we used the rubber hand illusion (RHI) as a proxy (Botvinick and Cohen, 1998) in Chapter 4, showing that RHI can be simulated by precision learning and its modulation. In addition, we also built an account of a rhythmic limit of neuronal architecture to the perception-action cycle using a proxy of attention in Chapter 5. This cycle is argued to be essential for correct body movement and estimation. The question then is, how are these cognitive abilities modulated by serotonin, since precision is behind both serotonin and body ownership?

First of all, Influential computational neuroscience accounts, such as active inference, place precision as an essential component for the correct body estimation, and proprioception in general (Seth and Friston, 2016; Limanowski, 2022; Brown et al., 2013; Hinz et al., 2018). These theories suggest that the brain maintains the body model, where the goal is to continuously integrate multisensory inputs. Here, precision control serves as a key mechanism for balancing different sensory inputs (Seth and Friston, 2016). Specifically, precision weighting allows the brain to contextually modulate the processing of visual versus proprioceptive information in a top-down manner (Limanowski, 2022). A crucial aspect of this process is that proprioceptive information, despite its fundamental importance for body perception and action control (Sakamoto et al., 1989), can be selectively attenuated to facilitate the integration of novel or conflicting visual feedback (Lanillos et al., 2021). This has been demonstrated both in static scenarios like the rubber hand illusion, where proprioceptive attenuation helps resolve visuo-proprioceptive conflicts (Botvinick and Cohen, 1998), and in dynamic contexts like visuomotor adaptation, where it aids in learning new sensorimotor connections (Bernier et al., 2009). The neuronal implementation of this precision control appears to operate through modulation of control gain in sensory cortices, particularly affecting activity in somatosensory areas during periods of sensory conflict (Zeller et al., 2016). This conceptualization provides a mechanistic

explanation for how the brain can maintain a coherent body representation while still allowing for the flexibility needed to adapt to changing circumstances or incorporate tools and prosthetics into the body schema (Maravita and Iriki, 2004).

The framework builds on the Chapter 4, where we used the RHI experiment to show that the precision is central to the correct body estimation, and Chapter 5 which argued for a perception-action cycle for attention, modulated by rhythmic attention. Given the support for the importance of precision in body estimation, and demonstrating that serotonin modelled as sensory precision explains empirical findings on this neuromodulator, the next section looks at the connection between serotonin and body estimation.

### 7.3 Serotonin and body ownership

The relationship between precision weighting and body ownership described above provides some indications for serotonin’s role from the early demonstrations, showing that serotonin follows the precision computation (i.e., the physiological action of serotonin is the biological manifestation of precision modulation). This connection manifests itself across multiple domains, from basic postural control to targeting more complex cognition that may result in self-dissolution. Both cases offer unique insights into how serotonergic signaling shapes our sense of embodiment.

Serotonin has been suggested to play a role in postural control, with studies showing that serotonergic neurons in the raphe nuclei project extensively to motor neurons in the spinal cord and influence muscle tone regulation (Tierney and Mangiamele, 2001; Harris-Warrick and Kravitz, 1984; Jacobs et al., 2002; Hornung, 2003; Ghosh and Pearse, 2015). This regulatory function is further evidenced by the serotonin syndrome, where excessive serotonergic activity leads to, apart from the most recognized increase of body temperature, muscle rigidity (Boyer and Shannon, 2005; Simon et al., 2024). Furthermore, developmental studies demonstrate serotonin’s critical role in postural maturation (Pflieger et al., 2002).

The link between the serotonergic function and body ownership becomes particularly evident in cases of dystonia, where abnormal muscle contractions lead to unusual postures (Smit et al., 2016). Interestingly, classical psychedelics have been reported to alleviate certain forms of dystonia and other motor disorders (Stewart et al., 2020), although it should be noted that dopaminergic drugs are often used in clinics for these disorders, as they are most effective (Calne, 1981). This therapeutic effect of serotonin might be explained through the precision-modulating framework and the REBUS theory. Speculatively, psychedelics could therefore help reset abnormal sensorimotor patterns that maintain dystonic postures by altering the precision weighting of proprioceptive signals.

Perhaps the most striking evidence for serotonin’s role in body ownership comes from the phenomenon of ego dissolution during psychedelic experiences. When classical psychedelics activate 5HT<sub>2A</sub> receptors, they can induce strong alterations in the sense of self and body ownership (Nour et al., 2016), characterized by a dissolution of the boundary between self and environment. These experiences can be understood as an extreme case of precision modulation, where psychedelics decrease the precision of prior beliefs about body ownership and self-boundaries, temporarily disrupting the normal hierarchical processing that maintains our sense of embodied self (Millière, 2022). This mechanism aligns well with the precision account developed earlier in this thesis. Just as precision modulation helps resolve sensory conflicts in the rubber hand illusion (Chapter 4), serotonergic modulation of precision appears to regulate the broader integration of multisensory signals that constitute our sense of bodily self. The dramatic alterations in self-perception induced through serotonergic mechanisms further supports the argument that serotonin functions as a key precision modulator in perceptual hierarchies.

This basic function demonstrates serotonin’s involvement in proprioceptive precision—the accuracy with which we sense and maintain our body’s position in space. Alterations in serotonergic function can lead to postural instability and motor control issues, suggesting that serotonin behaves as the precision of proprioceptive signals used for maintaining balance and posture, in line with the general function of precision on proprioception (Seth and Friston, 2016; Limanowski, 2022).

## 7.4 Other computational models of serotonin

In this thesis, we argue that the main function of serotonin is the precision adaptation in an active inference framework. However, there are several other computational works suggesting alternative homologues for serotonin compared to our view of a precision modulator. First, Miyazaki et al. (2018, 2020) suggested that the stimulation of serotonergic neurons leads to an overestimation of the prior probability of reward delivery, or also in confidence of reward acquisition. This work is derived from a different background looking at the reward mechanism instead of the exploratory one which was motivated from our end. Notably, such reward-based frameworks would be unable to reproduce the effects observed in our whisking simulations, which operate in a reward-free setting and demonstrate serotonin’s role in sensory processing independent of reward contingencies. This means that the models do not necessarily need to argue for opposite, just for distinct functions for distinct goals. Nevertheless, as confidence of reward acquisition can simply be seen as a precision of probability of reward acquisition, it supports the view that serotonin modulates precision across different domains. As the sensory precision will

modulate the exploratory mechanisms, the reward mechanism could exploit the precision of acquiring a given goal.

In addition to the reward-based approach, Dayan and Huys (2008) proposed a computational model linking serotonin's role in behavioural inhibition to mood disorders. Their model conceptualized trains of thought as actions that navigate through different states, with serotonin enabling inhibition of thought patterns predicted to lead to negative outcomes. They demonstrated that when serotonin levels are high, thoughts leading to negative outcomes are adaptively pruned, resulting in an optimistic bias. However, when modelling sudden serotonin depletion, this protective mechanism fails, leading to greater exposure to negative outcomes and large negative prediction errors.

While the inhibition-based model focuses primarily on serotonin's role in behavioural inhibition, it could potentially be integrated with our account of serotonin as a sensory precision modulator. In Dayan and Huys (2008), serotonin helps inhibit actions associated with negative outcomes. We propose this could be reinterpreted through precision: when encountering situations previously associated with aversive experiences, an agent with high serotonin levels might assign lower precision to new sensory information about these negative experiences and higher precision to prior expectations of avoiding harm. This precision balance would effectively prune thought patterns leading to negative states. Conversely, low serotonin might lead to assigning high precision to negative sensory information while maintaining low precision for optimistic priors, resulting in rumination on negative experiences. This recasting of the inhibition-based model in precision terms could potentially unify seemingly disparate accounts of serotonergic function.

A third computational account by Lottem et al. (2018) challenged the behavioural inhibition theory through a carefully designed foraging task. In traditional paradigms, waiting for rewards is often conflated with behavioural inhibition. To address this, they developed a task where mice had to actively nose-poke at reward ports to obtain water. In this 'probabilistically declining schedule', each successive nose-poke had a decreasing probability of yielding a reward, eventually reaching zero. This design required mice to decide when to abandon a depleting reward port in favor of an alternative one. The behavioural inhibition hypothesis, which suggests serotonin primarily suppresses actions, would predict that activating serotonergic neurons should decrease active nose-poking behaviour and cause mice to abandon ports sooner. However, using optogenetic activation of dorsal raphe serotonin neurons, Lottem et al. (2018) found the opposite: stimulation increased the number of nose-pokes mice would perform before giving up on a reward port.

Their results demonstrated that serotonin promotes active persistence, or in other words the continuation of effortful behaviour despite diminishing returns, rather than



passive waiting or inhibition. This behaviour was well-captured by a proportional hazards model, a statistical framework commonly used in survival analysis. In this model, the ‘hazard rate’ represents the probability of switching from one reward port to another at any given moment. Serotonin activation reduced this hazard rate, meaning it decreased the probability that mice would abandon their current action. This study thus demonstrated that the serotonergic function in waiting and patience cannot be reduced to simple behavioural inhibition. The findings align well with treating serotonin as a sensory precision modulator, as by increasing the weight given to prior beliefs about reward availability, serotonin could promote behavioural persistence even in the face of reward omissions. However, the exact computational framework linking these perspectives remains to be developed formally.

In conclusion, we support the precision-modulation framework over alternative models, as it offers a more comprehensive explanation for diverse effects of serotonin across behavioural contexts. However, the limitations particularly arise when addressing receptor-specific effects, and due to its largely theoretical focus. To test this model, future studies should directly measure precision-weighting during serotonergic manipulation, perform causal experiments independently manipulating both serotonin and precision, and design tasks that dissociate sensory and prior precision to determine if receptor-specific serotonergic agents affect these components as predicted. While the precision-modulation framework is testable through its specific predictions, only future empirical studies will ultimately determine whether the modelling efforts are confirmed or falsified.

## 7.5 Limitations

Despite the robust theoretical framework that the Free Energy Principle (FEP) provides for understanding serotonergic function as precision modulation, several limitations of this approach should be acknowledged. A notable criticism of the FEP, which emerges implicitly throughout this thesis, concerns its quite complex mathematical abstraction and the challenges in translating these theories into falsifiable empirical predictions. While our applications of precision modulation to whisking behaviour, body ownership, and psychedelic states demonstrate the principle’s explanatory power, the gap between computational formalism and measurable neural mechanisms remains substantial. The flexibility of FEP in explaining diverse phenomena could be viewed as both a strength and a potential weakness, raising questions about whether the framework risks becoming unfalsifiable when parameters can be adjusted to fit observed behaviours post-hoc.

Additionally, the thesis reveals a tension between the universality claimed by the FEP and the receptor-specific effects observed in serotonergic function. As discussed

in above, the opposing effects of different serotonin receptors (5HT1A versus 5HT2A) on precision weighting illustrate that biological reality may be more complex than a unified computational principle suggests. This complexity is particularly evident when attempting to reconcile seemingly contradictory findings from SERT knockout animals and psilocybin studies. While the precision-modulation framework provides a valuable theoretical background for understanding these observations, it will require further refinement to accommodate the nuanced, receptor-specific mechanisms that underlie serotonergic function in different behavioural and pharmacological contexts.

## 7.6 Conclusion and future directions

By treating biological brain processes as computational mechanisms, it was demonstrated that serotonin might play a critical role as a precision parameter. This was studied both in human and animal contexts, strengthened with theoretical work on attention and illusion. We aligned our finding of serotonin with the REBUS theory stating that psychedelics relax the weight of priors by modulating the precision parameter in the brain. It was also speculated on its influence in modulating both attention and body ownership via ego dissolution and motor disorders. However, this link is only indirect, and has to be strengthened with future studies specifically aiming to build a better understanding of serotonin as a precision modulator in attention and body ownership. Although this thesis has built a starting path by demonstrating that serotonin acts as a sensory precision parameter, either with the whisking and humanoid project, or with the psychedelics research, there is still more to be done to establish clear link between serotonin, precision, and body ownership and attention. Future computational modelling experiments could directly test these implications.

# Bibliography

- Ansorge, M. S., Zhou, M., Lira, A., Hen, R., and Gingrich, J. A. (2004). Early-life blockade of the 5-ht transporter alters emotional behavior in adult mice. *Science*, 306(5697):879–881.
- Araragi, N., Mlinar, B., Baccini, G., Gutknecht, L., Lesch, K.-P., and Corradetti, R. (2013). Conservation of 5-ht<sub>1A</sub> receptor-mediated autoinhibition of serotonin (5-ht) neurons in mice with altered 5-ht homeostasis. *Frontiers in pharmacology*, 4:97.
- Azmitia, E. C. (2020). Evolution of serotonin: sunlight to suicide. In *Handbook of behavioral neuroscience*, volume 31, pages 3–22. Elsevier.
- Barnes, N. M. and Sharp, T. (1999). A review of central 5-ht receptors and their function. *Neuropharmacology*, 38(8):1083–1152.
- Berg, K. A., Harvey, J. A., Spampinato, U., and Clarke, W. P. (2005). Physiological relevance of constitutive activity of 5-ht<sub>2A</sub> and 5-ht<sub>2C</sub> receptors. *Trends in pharmacological sciences*, 26(12):625–630.
- Bernier, P.-M., Burle, B., Vidal, F., Hasbroucq, T., and Blouin, J. (2009). Direct evidence for cortical suppression of somatosensory afferents during visuomotor adaptation. *Cerebral Cortex*, 19(9):2106–2113.
- Botvinick, M. and Cohen, J. (1998). Rubber hands ‘feel’ touch that eyes see. *Nature*, 391(6669):756–756.
- Boyer, E. W. and Shannon, M. (2005). The serotonin syndrome. *New England Journal of Medicine*, 352(11):1112–1120.
- Brown, H., Adams, R. A., Parees, I., Edwards, M., and Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive processing*, 14:411–427.

- Calne, D. B. (1981). Dopamine receptors in movement disorders. In *Movement disorders*, pages 348–355. Elsevier.
- Carhart-Harris, R. L. and Nutt, D. J. (2017). Serotonin and brain function: a tale of two receptors. *Journal of psychopharmacology*, 31(9):1091–1120.
- Crockett, M. J., Clark, L., and Robbins, T. W. (2009). Reconciling the role of serotonin in behavioral inhibition and aversion: acute tryptophan depletion abolishes punishment-induced inhibition in humans. *Journal of Neuroscience*, 29(38):11993–11999.
- Daw, N. D., Kakade, S., and Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural networks*, 15(4-6):603–616.
- Dayan, P. and Huys, Q. J. M. (2008). Serotonin, inhibition, and negative mood. *PLoS computational biology*, 4(2):e4.
- Deneve, S. and Jardri, R. (2016). Circular inference: mistaken belief, misplaced trust. *Current Opinion in Behavioral Sciences*, 11:40–48.
- Doya, K., Miyazaki, K. W., and Miyazaki, K. (2021). Serotonergic modulation of cognitive computations. *Current Opinion in Behavioral Sciences*, 38:116–123.
- Fabre, V., Beaufour, C., Evrard, A., Rioux, A., Hanoun, N., Lesch, K., Murphy, D., Lanfumey, L., Hamon, M., and Martres, M.-P. (2000). Altered expression and functions of serotonin 5-HT<sub>1A</sub> and 5-HT<sub>1B</sub> receptors in knock-out mice lacking the 5-HT transporter. *European Journal of Neuroscience*, 12(7):2299–2310.
- Fiebelkorn, I. C. and Kastner, S. (2019). A rhythmic theory of attention. *Trends in cognitive sciences*, 23(2):87–101.
- Ghosh, M. and Pearse, D. D. (2015). The role of the serotonergic system in locomotor recovery after spinal cord injury. *Frontiers in neural circuits*, 8:151.
- Harris-Warrick, R. M. and Kravitz, E. A. (1984). Cellular mechanisms for modulation of posture by octopamine and serotonin in the lobster. *Journal of Neuroscience*, 4(8):1976–1993.
- Hillyard, S. A., Vogel, E. K., and Luck, S. J. (1998). Sensory gain control (amplification) as a mechanism of selective attention: electrophysiological and neuroimaging evidence. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 353(1373):1257–1270.

- Hinz, N.-A., Lanillos, P., Mueller, H., and Cheng, G. (2018). Drifting perceptual patterns suggest prediction errors fusion rather than hypothesis selection: replicating the rubber-hand illusion on a robot. In *2018 Joint IEEE 8th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, pages 125–132. IEEE.
- Holmes, A., Yang, R. J., Lesch, K.-P., Crawley, J. N., and Murphy, D. L. (2003). Mice lacking the serotonin transporter exhibit 5-HT<sub>1A</sub> receptor-mediated abnormalities in tests for anxiety-like behavior. *Neuropsychopharmacology*, 28(12):2077–2088.
- Homberg, J. R., Schubert, D., and Gaspar, P. (2010). New perspectives on the neurodevelopmental effects of SSRIs. *Trends in pharmacological sciences*, 31(2):60–65.
- Hornung, J.-P. (2003). The human raphe nuclei and the serotonergic system. *Journal of chemical neuroanatomy*, 26(4):331–343.
- Hoyer, D., Hannon, J. P., and Martin, G. R. (2002). Molecular, pharmacological and functional diversity of 5-HT receptors. *Pharmacology Biochemistry and Behavior*, 71(4):533–554.
- Jacobs, B. L., Martin-Cora, F. J., and Fornal, C. A. (2002). Activity of medullary serotonergic neurons in freely moving animals. *Brain Research Reviews*, 40(1-3):45–52.
- Lanillos, P., Franklin, S., Maselli, A., and Franklin, D. W. (2021). Active strategies for multisensory conflict suppression in the virtual hand illusion. *Scientific Reports*, 11(1):22844.
- Li, Q., Wichems, C., Heils, A., Van de Kar, L. D., Lesch, K.-P., and Murphy, D. L. (1999). Reduction of 5-hydroxytryptamine (5-HT) <sub>1A</sub>-mediated temperature and neuroendocrine responses and 5-HT<sub>1A</sub> binding sites in 5-HT transporter knockout mice. *The Journal of pharmacology and experimental therapeutics*, 291(3):999–1007.
- Li, Q., Wichems, C. H., Ma, L., Van de Kar, L. D., Garcia, F., and Murphy, D. L. (2003). Brain region-specific alterations of 5-HT<sub>2A</sub> and 5-HT<sub>2C</sub> receptors in serotonin transporter knockout mice. *Journal of Neurochemistry*, 84(6):1256–1265.
- Limanowski, J. (2022). Precision control for a flexible body representation. *Neuroscience & Biobehavioral Reviews*, 134:104401.
- Lottem, E., Banerjee, D., Vertechi, P., Sarra, D., Lohuis, M. O., and Mainen, Z. F. (2018). Activation of serotonin neurons promotes active persistence in a probabilistic foraging task. *Nature communications*, 9(1):1000.

- Maravita, A. and Iriki, A. (2004). Tools for the body (schema). *Trends in cognitive sciences*, 8(2):79–86.
- Michel, C. M. and Koenig, T. (2018). Eeg microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: a review. *Neuroimage*, 180:577–593.
- Millière, R. (2022). Drug-induced alterations of bodily awareness. In *The Routledge Handbook of Bodily Awareness*, pages 503–521. Routledge.
- Miyazaki, K., Miyazaki, K. W., Sivori, G., Yamanaka, A., Tanaka, K. F., and Doya, K. (2020). Serotonergic projections to the orbitofrontal and medial prefrontal cortices differentially modulate waiting for future rewards. *Science advances*, 6(48):eabc7246.
- Miyazaki, K., Miyazaki, K. W., Yamanaka, A., Tokuda, T., Tanaka, K. F., and Doya, K. (2018). Reward probability and timing uncertainty alter the effect of dorsal raphe serotonin neurons on patience. *Nature communications*, 9(1):2048.
- Nour, M. M., Evans, L., Nutt, D., and Carhart-Harris, R. L. (2016). Ego-dissolution and psychedelics: validation of the ego-dissolution inventory (edi). *Frontiers in human neuroscience*, 10:269.
- Pfieger, J.-F., Clarac, F., and Vinay, L. (2002). Postural modifications and neuronal excitability changes induced by a short-term serotonin depletion during neonatal development in the rat. *Journal of Neuroscience*, 22(12):5108–5117.
- Pincus, S. M. (1991). Approximate entropy as a measure of system complexity. *Proceedings of the national academy of sciences*, 88(6):2297–2301.
- Pytliak, M., Vargová, V., Mechírová, V., and Felsöci, M. (2011). Serotonin receptors-from molecular biology to clinical applications. *Physiological research*, 60(1):15.
- Raymond, J. R., Mukhin, Y. V., Gelasco, A., Turner, J., Collinsworth, G., Gettys, T. W., Grewal, J. S., and Garnovskaya, M. N. (2001). Multiplicity of mechanisms of serotonin receptor signal transduction. *Pharmacology & therapeutics*, 92(2-3):179–212.
- Sakamoto, T., Arissian, K., and Asanuma, H. (1989). Functional role of the sensory cortex in learning motor skills in cats. *Brain research*, 503(2):258–264.
- Seth, A. K. and Friston, K. J. (2016). Active interoceptive inference and the emotional brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1708):20160007.

- Simon, L. V., Torrico, T. J., and Keenaghan, M. (2024). Serotonin syndrome. In *StatPearls [Internet]*. StatPearls Publishing.
- Smit, M., Bartels, A., van Faassen, M., Kuiper, A., Niezen-Koning, K., Kema, I., Dierckx, R., de Koning, T., and Tijssen, M. (2016). Serotonergic perturbations in dystonia disorders—a systematic review. *Neuroscience & Biobehavioral Reviews*, 65:264–275.
- Soiza-Reilly, M., Goodfellow, N. M., Lambe, E. K., and Commons, K. G. (2015). Enhanced 5-HT<sub>1A</sub> receptor-dependent feedback control over dorsal raphe serotonin neurons in the ser<sup>t</sup> knockout mouse. *Neuropharmacology*, 89:185–192.
- Stewart, B., Dean, J. G., Koek, A., Chua, J., Wabl, R., Martin, K., Davoodian, N., Becker, C., Himedan, M., Kim, A., et al. (2020). Psychedelic-assisted therapy for functional neurological disorders: a theoretical framework and review of prior reports. *Pharmacology Research & Perspectives*, 8(6):e00688.
- Tierney, A. and Mangiamele, L. (2001). Effects of serotonin and serotonin analogs on posture and agonistic behavior in crayfish. *Journal of Comparative Physiology A*, 187:757–767.
- Zeller, D., Friston, K. J., and Classen, J. (2016). Dynamic causal modeling of touch-evoked potentials in the rubber hand illusion. *Neuroimage*, 138:266–273.

# Chapter 8

## Appendices

### 8.1 Summary

While theoretical models have increasingly embraced Bayesian frameworks—mathematical methods that combine prior knowledge with new evidence to make optimal predictions—to explain neural processing, there remains a significant gap in understanding how these mathematical principles are implemented in the biological brain. Specifically, we lack clarity on how neuromodulators like serotonin translate abstract computational parameters into neural activity. This thesis addresses this knowledge gap by proposing that serotonin serves as a biological homologue of precision modulation in predictive processing. This is tested and corroborated within experiments focusing on exploratory behavior, perceptual illusions, and altered states of neuronal activity induced by psychedelics.

The investigation of serotonergic modulation in rodent whisking behavior reveals how precision weighting may serve as a key mechanism by which this neurotransmitter system influences sensory processing or active sensing. Using an active inference framework, the research demonstrates that serotonin may fundamentally act to modulate the precision of sensory inputs and/or prior habits, thereby regulating exploratory behavior and environmental sampling.

The robotics work translates these biological insights into physical artificial systems, showing how precision-based active inference can guide autonomous behavior. By implementing precision-weighted sensory processing in a humanoid robot, this study demonstrates how this framework can enable adaptive sensorimotor control and efficient information seeking behavior in artificial agents.

The theoretical work on attention and body ownership illusions builds on the theme



of precision-control in embodied systems. The thesis then directly demonstrates how precision modulation can explain both perception and action. Through mathematical modeling of the rubber hand illusion, the research shows how the brain might arbitrate between competing models of sensory input via precision-weighted inference, providing a novel account of how body ownership experiences emerge.

The analysis of psilocybin's effects on neural dynamics demonstrates that this compound, which work as a serotonergic agent, may work by elevating the chaotic responses in the brain, supporting the theory that psychedelics decrease the precision of hierarchical neural communication. Through EEG microstate analysis and approximate entropy measures, the results show that psilocybin increases the transition rate and complexity of brain states, suggesting a fundamental reorganization of neural dynamics through precision modulation.

The thesis concludes with a discussion exploring how the serotonergic system may influence the integration of proprioceptive signals leading to a sense of body, and further evaluating alternative computational frameworks for understanding serotonergic function across different neural and behavioral domains.

## 8.2 Samenvatting

Terwijl theoretische modellen steeds meer Bayesiaanse raamwerken omarmen om neurale verwerking te verklaren, blijft er een grote kloof bestaan in het begrijpen hoe deze wiskundige principes worden geïmplementeerd in het biologische brein. Specifiek ontbreekt het ons aan duidelijkheid over hoe neuromodulators zoals serotonine abstracte computationele parameters vertalen in neurale activiteit. In dit proefschrift wordt dit vraagstuk aangepakt door voor te stellen dat serotonine dient als een biologische homoloog van precisie-modulatie in voorspellende verwerking. Dit wordt getest en bevestigd in experimenten die zich richten op exploratief gedrag, perceptuele illusies en veranderde toestanden van neuronale activiteit veroorzaakt door psychedelica.

Het onderzoek naar serotonerge modulatie in het fluistergedrag van knaagdieren onthult hoe precisieafweging kan dienen als een belangrijk mechanisme waarmee dit neurotransmittersysteem sensorische verwerking of actieve detectie beïnvloedt. Met behulp van een raamwerk voor actieve inferentie toont het onderzoek aan dat serotonine fundamenteel kan werken om de precisie van zintuiglijke input en/of voorafgaande gewoonten te moduleren, waardoor exploratief gedrag en het onderzoeken van de omgeving worden gereguleerd.

Het roboticawerk vertaalt deze biologische inzichten naar fysieke kunstmatige systemen en laat zien hoe op precisie gebaseerde actieve inferentie autonoom gedrag kan sturen. Door het implementeren van precisie-gewogen sensorische verwerking in een humanoïde robot, laat dit onderzoek zien hoe dit raamwerk adaptieve sensorimotorische controle en efficiënt informatiezoekend gedrag in kunstmatige agenten mogelijk kan maken.

Het theoretische werk over aandacht en lichaamseigendom illusies bouwt voort op het thema van precisie-controle in belichaamde systemen. Het proefschrift toont vervolgens direct aan hoe precisie-modulatie zowel perceptie als actie kan verklaren. Door middel van wiskundige modellering van de rubberen hand illusie laat het onderzoek zien hoe de hersenen zouden kunnen arbitrerende tussen concurrerende modellen van zintuiglijke input via precisie-gewogen inferentie, wat een nieuwe verklaring biedt voor hoe lichaamseigendom ervaringen ontstaan.

De analyse van de effecten van psilocybine op neurale dynamica toont aan dat deze verbinding, die werkt als een serotonerge stof, mogelijk werkt door de chaotische reacties in de hersenen te verhogen, wat de theorie ondersteunt dat psychedelica de precisie van hiërarchische neurale communicatie verminderen. Door middel van EEG microstatus analyse en benaderende entropiematen laten de resultaten zien dat psilocybine de overgangssnelheid en complexiteit van hersentoestanden verhoogt, wat wijst op een fundamentele reorganisatie van neurale dynamiek door precisie-modulatie.

Het proefschrift sluit af met een discussie waarin wordt onderzocht hoe het serotonerge systeem de integratie van proprioceptieve signalen kan beïnvloeden, wat leidt tot een gevoel van lichaam, en waarin alternatieve computationele raamwerken voor het begrijpen van serotonerge functie in verschillende neurale en gedragsdomeinen verder worden geëvalueerd.

## 8.3 Research Data Management

### Data Management

Mathematical models and simulation algorithms developed for Chapters 2-5 operate independently of any subject data and focus purely on theoretical neuroscientific concepts. The research in Chapter 6 utilized high-density EEG data from a 64-channel BrainAmp MR Plus system, acquired from the PSICONNECT study (ACTRN12621001375842) at Monash University. This dataset captured neural activity across multiple cognitive states recorded both before and after psilocybin administration. Data ownership and usage rights were established through formal agreements with Monash University as the primary data custodian, with access permissions secured through official data sharing protocols.

### Ethical Approval

The original data collection received full ethical clearance from Monash University Human Research Ethics Committee and maintains registration with the Australian New Zealand Clinical Trials Registry (ACTRN12621001375842).

All participants provided written informed consent for data collection and future research applications, with consent documentation explicitly covering potential data sharing for scientific research purposes. All data were fully anonymized by the original research team with participant identifiers removed and replaced with anonymous codes, ensuring privacy safeguards were maintained throughout all stages of data handling and analysis.

### Findability and Accessibility

The source dataset originates from research conducted at Monash Biomedical Imaging, Australia, spanning February through December 2022. Dataset archival follows Monash University's institutional data management framework, with future public release planned following primary publication of the PSICONNECT findings. Data access requests should be directed to the PSICONNECT principal investigators: Devon Stoliker ([devon.stoliker@monash.edu](mailto:devon.stoliker@monash.edu)) and Adeel Razi ([adeel.razi@monash.edu](mailto:adeel.razi@monash.edu)).

## 8.4 Curriculum Vitae

Filip Novický was born in the Czech Republic and completed his Bachelor's degree in Psychology at Masaryk University in 2020. A European exchange programme at Bilkent University in 2019 sparked his interest in neuroscience and motivated him to pursue studies abroad. He moved to the Netherlands to pursue a Master of Science degree in Cognitive Neuroscience at Maastricht University, completing it in 2021. As part of his master's programme, he undertook an internship at University College London, where he joined Karl Friston's lab under the supervision of Berk Mirza and Noor Sajid, developing a bistable perception model using the active inference framework.

Filip is a PhD candidate in Computational Neuroscience at the Donders Institute, Radboud University, based in Nijmegen, Netherlands. His doctoral research is supported by the Marie Curie Serotonin & Beyond programme. In this PhD, he focused on Bayesian modeling of brain dynamics and neuroimaging analyses. His work is supervised by Fleur Zeldenrust, Judith Homberg, Pablo Lanillos, and Thomas Parr. His research interests center on the intersection of computational modeling, neurobiology, and robotics, with expertise in Bayesian inference, specifically precision-weighted perception.

Filip's research experience includes internships at the University of Florence (2023), where he worked under Alessandro Scaglione and Francesco Resta applying convolutional non-negative matrix factorization methods to interpret spatiotemporal patterns of brain activity in anesthetized mice, and at Monash University (2024) under Adeel Razi's supervision, studying the influence of psychedelics on brain dynamics using EEG microstates and approximate entropy measures. He has published first-author publications in venues including *Frontiers in Neurorobotics* and *Cerebral Cortex*, spanning from theoretical frameworks linking serotonin to computational models of perception to practical applications in robotics inspired by neural mechanisms.

Filip is engaged in the scientific community through conference presentations across Europe and teaching responsibilities at Radboud University, including co-supervising master's students and leading weekly tutorials for undergraduate biology students in mathematics. His technical skills encompass programming in Python and MATLAB, and proficiency in neuroimaging analysis packages including SPM12, EEGLAB, and MNE.

8.5 Portfolio

Courses	Organizer	Year
Introduction Day	Radboud University	2021
Human-Robot Interaction	Radboud University	2022
Scientific Integrity Course	Radboud University	2025
Mathematics for Biologists (Teaching)	Radboud University	2023, 2025

External lectures and conferences

Title	Role	Year
European Conference on Visual Perception (Netherlands)	Poster Presentation	2022
3rd International Workshop on Active Inference (France)	Verbal Presentation	2022
Serotonin & Beyond Conference (France)	Verbal Presentation	2022
Dutch Neuroscience Meeting (Netherlands)	Poster Presentation	2023
Living Machines Conference (Italy)	Verbal Presentation	2023
Serotonin & Beyond Conference (Germany)	Verbal Presentation	2023
Serotonin & Beyond Conference (Italy)	Verbal Presentation	2024
BRAINNET Conference (Sweden)	Verbal Presentation	2025
Collective Dynamics and Information Processing in Neural Systems (Italy)	Poster Presentation	2025
Computational Neuroscience Meeting (Italy)	Poster Presentation	2025

Research experience and internships

Institution	Supervisor	Year	Duration
University College London	Berk Mirza, Noor Sajid	2021	6 months
University of Florence	Alessandro Scaglione, Francesco Resta	2023	3 months
Monash University	Adeel Razi	2024	6 months

## 8.6 List of Publications

### Published

- Meera, A. A., Novicky, F., Parr, T., Friston, K., Lanillos, P., & Sajid, N. (2022). Reclaiming saliency: Rhythmic precision-modulated action and perception. *Frontiers in Neurobotics*, 16. <https://doi.org/10.3389/fnbot.2022.850512>
- Novicky, F., Offergeld, J., Janssen, S., & Lanillos, P. (2023). Robotic active tactile sensing inspired by serotonergic modulation using active inference. In F. Meder, A. Hunt, L. Margheri, A. Mura, & B. Mazzolai (Eds.), *Biomimetic and biohybrid systems (Living Machines 2023)*. *Lecture Notes in Computer Science* (Vol. 14157). Springer, Cham. [https://doi.org/10.1007/978-3-031-38857-6\\_3](https://doi.org/10.1007/978-3-031-38857-6_3)
- Novicky, F., Parr, T., Friston, K., Mirza, M. B., Sajid, N., & Berk, M. (2024). Bistable perception, precision and neuromodulation. *Cerebral Cortex*, 34(1), bhad401. <https://doi.org/10.1093/cercor/bhad401>

### Preprint

- Novický, F., Meera, A. A., Zeldenrust, F., & Lanillos, P. (2024). Precision not prediction: Body-ownership illusion as a consequence of online precision adaptation under Bayesian inference. *bioRxiv*. <https://doi.org/10.1101/2024.09.04.611162>

### In preparation

- Novický, F., Stoliker, D., Razi, A., & Zeldenrust, F. Psilocybin accelerates EEG microstate transitions and elevates approximate entropy. *Manuscript in preparation*.
- Novický, F., Friston, K., Zeldenrust, F., Lanillos, P., Parr, T., & Homberg, J. Precision, embodiment, and serotonin: Reconciling neurobiology with computational frameworks. *Manuscript in preparation*.
- Novický, F., Zeldenrust, F., Lanillos, P., Homberg, J. R., Friston, K., & Parr, T. Active inference, whisking, and serotonin. *Manuscript in preparation*.

## 8.7 Acknowledgment

Reflecting on this PhD journey, I am very grateful to the many remarkable people who have shaped both my research and personal growth along the way. Each person mentioned here has contributed something unique and invaluable, making this work not just possible, but truly meaningful.

In the beginning of my PhD, which started as a roller coaster with several mishaps, Fleur took on another PhD student and helped me establish my new PhD trajectory during this period. I owe her a lot for letting me explore different sides in science. For me, as a trained psychologist, it was quite novel and encouraging to understand her physics approach towards neuroscience, which

in the end broadened my perspective of what this field can be about. In addition, her eagerness towards building a great lab has worked out quite well, and thanks to her I could be a part of a nice and young lab.

Pablo's insights were also quite considerable, since thanks to him I could get involved in a bit more technical niche of robotics. Working with him on several projects and seeing his efforts to include me in conferences is something I appreciate a lot.

Judith's biological knowledge on the anatomical level was always very appreciated, which was particularly useful when setting up models. Mostly, though, the fact that she started the entire Serotonin & Beyond network, that I had the pleasure to be part of, is admirable and I benefited a lot both from seeing small details of running such an organization, while being shielded from much of the administrative burden.

Lastly, it was a great pleasure to have Tom as an external supervisor on board. His insights and wide, yet integrated, knowledge of the neuroscientific findings is something I have always admired. I have learned a lot from him, and the aim to write academic texts with such clarity is something I will always strive for, with the goal of reaching his level one day.

I am particularly grateful for the opportunity to spend part of my PhD in Florence and Melbourne, experiences that enriched both my research and personal perspective. In Italy, Alessandro and Francesco welcomed me into their lab and introduced me to fascinating new analytical approaches, while also allowing me to discover the beauty of Florence. Their mentorship opened up new methodological horizons that taught me how to see and approach neuronal data. My time in Australia was equally transformative, thanks to Adeel's guidance in analyzing a compelling psilocybin dataset. In Australia, I could also spend some (or more) time in Jakob's lab, full of philosophers, which often ended up in quite engaging discussions.

People I could directly work with made my projects so much easier. It was amazing to work with Ajith, a great scientist I could collaborate with on two projects. As a very knowledgeable roboticist, his insights in technical understanding are truly astonishing to me. Thanks to Simon and Joshua, also technically minded coworkers, I had a great time working on our small robotics project, where all of us learned a lot.

I would also like to acknowledge all my great lab mates. Cheers to the musician Nicolas, whose talks about the scientific endeavor were both a source of excitement, and a way to ventilate our frustration at the same time. Mostly, though, he could teach me to play on guitar, which I won't forget, as it is stuck in my fingers now. Tousif, the best world tourist I've ever met, could come up with great events by himself – quite literally – and provided not one, but two coffee machines for the lab, which is beyond appreciation. Farhad's easygoing nature and all the funny teasing made everybody around him laugh, creating some great moments that will stick in my mind.

I feel very fortunate to have friends who, throughout the years we lived in different countries, are still very close to me. Starting with my German twin Katharina, who was that patient to listen to all my problems during the PhD struggles, and did not complain at all. My friend Walter gave me lots of inspiration, both through his reading recommendations and his questions about science, which were quite motivating. He would show me that philosophy of science is particularly great to know of, especially if you work as a scientist. I am quite lucky that I could have Rao around, almost throughout the entire PhD whose energetic and easygoing nature made life in



Nijmegen more engaging. Lastly, thanks to Daniel, whose philosophical night talks during our bachelor's studies actually started my scientific career, though he would never take credit for that. Though, I appreciate him even more for his sense of humor which is just perfect.

Finally, I want to acknowledge my mother, whose endless support has been the foundation of my entire life-long academic journey. Her dedication allowed me to pursue my studies with complete focus, as she skillfully managed countless details and challenges that I have only come to fully appreciate now. While there will be many future occasions to properly express what your guidance has meant to me, I want to recognize here that this achievement is as much yours as it is mine. Thank you for everything.

## 8.8 Donders Graduate School

For a successful research Institute, it is vital to train the next generation of scientists. To achieve this goal, the Donders Institute for Brain, Cognition and Behaviour established the Donders Graduate School in 2009. The mission of the Donders Graduate School is to guide our graduates to become skilled academics who are equipped for a wide range of professions. To achieve this, we do our utmost to ensure that our PhD candidates receive support and supervision of the highest quality.

Since 2009, the Donders Graduate School has grown into a vibrant community of highly talented national and international PhD candidates, with over 500 PhD candidates enrolled. Their backgrounds cover a wide range of disciplines, from physics to psychology, medicine to psycholinguistics, and biology to artificial intelligence. Similarly, their interdisciplinary research covers genetic, molecular, and cellular processes at one end and computational, system-level neuroscience with cognitive and behavioural analysis at the other end. We ask all PhD candidates within the Donders Graduate School to publish their PhD thesis in the Donders Thesis Series. This series currently includes over 600 PhD theses from our PhD graduates and thereby provides a comprehensive overview of the diverse types of research performed at the Donders Institute. A complete overview of the Donders Thesis Series can be found on our website: <https://www.ru.nl/donders/donders-series>

The Donders Graduate School tracks the careers of our PhD graduates carefully. In general, the PhD graduates end up at high-quality positions in different sectors, for a complete overview see <https://www.ru.nl/donders/destination-our-former-phd>. A large proportion of our PhD alumni continue in academia (>50%). Most of them first work as a postdoc before growing into more senior research positions. They work at top institutes worldwide, such as University of Oxford, University of Cambridge, Stanford University, Princeton University, UCL London, MPI Leipzig, Karolinska Institute, UC Berkeley, EPFL Lausanne, and many others. In addition, a large group of PhD graduates continue in clinical positions, sometimes combining it with academic research. Clinical positions can be divided into medical doctors, for instance, in genetics, geriatrics, psychiatry, or neurology, and in psychologists, for instance as healthcare psychologist, clinical neuropsychologist, or clinical psychologist. Furthermore, there are PhD graduates who continue to work as researchers outside academia, for instance at non-profit or government organizations, or in pharmaceutical companies. There are also PhD graduates who work in education, such as teachers in high school, or as lecturers in higher education. Others continue in a wide range of positions, such as policy advisors, project managers, consultants, data scientists, web- or software developers, business owners, regulatory affairs specialists, engineers, managers, or IT architects. As such, the career paths of Donders PhD graduates span a broad range of sectors and professions, but the common factor is that they almost all have become successful professionals.

For more information on the Donders Graduate School, as well as past and upcoming defences please visit: <http://www.ru.nl/donders/graduate-school/phd/>



**DONDERS**  
INSTITUTE



Radboud University



Radboudumc