


Master of Science Thesis

The Effect of Prior Knowledge on the Sense of Agency

Using Expectation-Maximization to Reproduce the Moving Rubber-Hand Illusion

Meilin Tan
2020



THE EFFECT OF PRIOR KNOWLEDGE ON SENSE OF AGENCY

USING EXPECTATION-MAXIMIZATION TO REPRODUCE THE MOVING
RUBBER-HAND ILLUSION

by

Meilin TAN

Master of Science Thesis

for the degree of Master of Science in Mechanical Engineering
at Delft University of Technology
to be defended on December 17, 2020 at 14:00

Supervisors: Prof. dr. ir. Martijn Wisse
Dr. Arkady Zgonnikov
Thesis committee: Prof. dr. ir. Martijn Wisse
Dr. Arkady Zgonnikov
Dr. Ksander de Winkel

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

ABSTRACT

When humans make inferences that go beyond limited, noisy, or ambiguous input data, background knowledge is necessary to make generalizations. Such inferences are important for designing intelligent artificial agents. Bayesian inference, a statistical method commonly used as a cognitive model of the brain, formalizes the integration of background knowledge, i.e. "priors", and sensory evidence. I research the effect of priors on *Sense of Agency* (SoA) in an artificial agent. In humans, SoA is the subjective experience of control over actions and their consequence. In robotics, developing an agent with SoA is a popular challenge and the first step towards the artificial self.

This thesis designs an artificial agent where the same prior knowledge improves SoA in unambiguous environments and induces incorrect, illusory SoA in noisy, ambiguous environments. First, I provide a comprehensive overview of the role of priors in Bayesian inference and the computational principles of a SoA. Second, I define the "point mass *moving rubber-hand illusion* (mRHI)", a simulation for an artificial agent that simplifies the human mRHI experiment, where participants gain illusory SoA over a rubber-hand. Third, I use two parameter estimators, *ordinary least squares* (OLS) and *Expectation-Maximization* (EM), to calculate SoA in the point mass mRHI.

I found that SoA requires a prior belief about the causal relationship between outcome and action. In the point mass mRHI, the agent has to identify which of the three point masses it can control with its force input. The agent has the prior knowledge that there is a causal relationship between its actions and the states of one point mass. Though with OLS the agent has correct and incorrect SoA, some of the results not comparable to the results from human mRHI experiment. A likely cause is that this simplest replication of the mRHI models SoA as a binary variable and does not use Bayesian inference. The (partially) Bayesian parameter estimator EM calculates SoA as a (continuous) posterior probability and finds results similar to the human mRHI experiment. Comparing an OLS algorithm without prior to the EM algorithm, I found that in unambiguous environments the Bayesian prior improves the agent's general SoA and SoA over time but does not improve the SoA for the point mass that the agent does control. In the ambiguous environments, the prior generally motivates the agent towards a SoA over the incorrect mass. However, in a noisy but slightly less ambiguous environment, the prior improves the agent's SoA. In short, the prior improves a SoA when the agent needs to make inferences that go beyond the data but can induce an incorrect SoA in a noisy and ambiguous environment.

CONTENTS

Abstract	iii
List of Figures	ix
List of Tables	xiii
Acronyms	xvii
Acknowledgements	xix
1 Introduction	1
1.1 Motivation	1
1.1.1 Bayesian Inference Unites Prior Knowledge and Sense of Agency . .	2
1.1.2 The Effect of Priors on Inferences	2
1.2 Problem Statement	4
1.2.1 Addressing Two Research Challenges of Priors in Agents	4
1.2.2 Research Goal	5
2 Background	9
2.1 The Influence of the Prior on Bayesian Inferences	10
2.1.1 Statistics: Theoretical Role of the Prior	10
2.1.2 Cognitive Science: Models and Empirical Evidence	12
2.1.3 Neuroscience: Bayesian Inference and Priors in the Brain	15
2.1.4 Robotics: Bayesian Inference and Priors in Artificial Agents	15
2.1.5 Summary: the Influence of the Prior on Bayesian Inference	17
2.2 Principles Underlying the Sense of Agency	19
2.2.1 General Overview Sense of Agency (SoA)	19
2.2.2 Three Underlying Computational Principles of SoA	21
2.2.3 Conclusion: Connecting Principles of SoA to Priors	23
2.3 Bringing Together Priors, the SoA, and the mRHI	24
2.3.1 The mRHI	24
2.3.2 Simple Bayesian Explanation of SoA in the mRHI	25
2.4 Conclusion	28
3 Experimental Design	29
3.1 The Point Mass mRHI: the Mathematical Model	30
3.1.1 The Agent's Model: Initial Explanation	30
3.1.2 The Plant: Generating the Environmental States	31
3.1.3 The Agent's Model: Agency over the Observed States	32
3.1.4 Uniting the Plant and the Agent's Model	34

3.2	Point Mass mRHI as an Analogy of Human mRHI.	35
3.2.1	Comparison of the Experimental Set-Up.	35
3.2.2	Computational Principles of SoA in the Point Mass mRHI	35
3.3	Experiments	38
3.3.1	Parameters of the Point Mass mRHI	38
3.3.2	Performance Measures.	40
3.4	Conclusion	43
4	The Simplest Point Mass mRHI	45
4.1	The Ordinary Least-Squares Algorithm for Agency	46
4.2	Implementation of the Prior in OLS.	48
4.2.1	Sense of Agency is Determined with P	48
4.2.2	OLS Algorithm	49
4.3	Results	51
4.3.1	Experiment 1: Correct SoA in Unambiguous Environments	51
4.3.2	Experiment 2: Incorrect SoA in Ambiguous Environments.	54
4.4	Discussion	57
4.4.1	Comparison of the SoA in Human and the Agent.	57
4.4.2	Improvements to the Point Mass mRHI	58
4.5	Conclusion	60
5	The Point Mass mRHI with Expectation-Maximization	61
5.1	Theory: Expectation-Maximization for SoA.	62
5.1.1	The Observed Sensory Data	62
5.1.2	Expectation-Maximization Algorithm Overview	63
5.1.3	The Gaussian Mixture Model.	64
5.1.4	E-step	65
5.1.5	M-step	66
5.1.6	Defining Agency and SoA	66
5.1.7	Summary: EM Algorithm	68
5.2	The EM Algorithm Motivation	69
5.2.1	Computational Advantages of EM	69
5.2.2	Similarity Between Cognition and the EM Algorithm.	69
5.3	Experiment 1A: EM in a Static Environment.	71
5.3.1	Method: EM Algorithm for the Point Mass mRHI in Python	71
5.3.2	Results and Discussion.	75
5.4	EM in a Dynamic Environment	77
5.4.1	Extra Hidden Causes	77
5.4.2	Method: Implementation of the EM algorithm.	80
5.4.3	Results	84
5.5	Discussion: Evaluation of the Point Mass mRHI.	99
5.5.1	Three SoA Computational Principles.	99
5.5.2	Improved SoA in Unambiguous Environments.	100
5.5.3	Incorrect SoA in Ambiguous Environments	101
5.5.4	General Issues in the EM Algorithm	101
5.6	Conclusion	103

6	Conclusions	105
6.1	Summary	106
6.1.1	Answering the Research Questions.	106
6.1.2	Addressing the Research Challenges of Priors	107
6.2	Recommendations	108
6.2.1	Extensions of the Point Mass mRHI	108
6.2.2	Future Directions for Priors in Artificial Agents	109
6.2.3	Future Directions for SoA in Artificial Agents.	110
A	Measures of Sense of Agency in Experimental Investigation	123
B	Generating Discrete-Time States With the Plant	125
C	Results from Point Mass mRHI with OLS	127
C.1	Experiment 1: Unambiguous Environments	128
C.2	Experiment 2: Ambiguous Environments	129
D	Results form Point Mass mRHI with EM	131
D.1	Experiment 1: Unambiguous Environments	132
D.2	Experiment 2: Ambiguous Environments	133
D.2.1	Table of the Performance Measures	133
D.2.2	Estimation of the Parameters the Clusters	133
E	The Influence of the Initial Covariance Matrix on SoA	137

LIST OF FIGURES

1.1	This figure provides a simplified overview of the agent’s inferential process for <i>Sense of Agency</i> (SoA).	6
2.1	This schematic shows three cases when the posterior belief shifts towards the prior belief.	11
2.2	This figure provides a Bayesian account of the binocular rivalry illusion, where the prior plays a key role.	14
2.3	This figure provides an overview of the connection between agency, SoA, and the three computational principles of SoA.	23
2.4	This figure shows the different experimental conditions of the <i>rubber-hand illusion</i> (RHI).	24
2.5	This figure shows the two experimental conditions in which the agent uses Bayesian inference to determine over which hand it has agency.	25
2.6	A schematic of the (hypothetical) posterior probabilities after some time for the participants hand, $p(A_1 \mathbf{s}, \mathbf{p}, \mathbf{v}_1)$, rubber hand 1 $p(A_2 \mathbf{s}, \mathbf{p}, \mathbf{v}_2)$, and rubber hand 2 $p(A_3 \mathbf{s}, \mathbf{p}, \mathbf{v}_3)$.	27
3.1	This figure illustrates the environment that the agent observes.	30
3.2	This block diagram shows that the plant and the agent’s model are connected to determine SoA.	34
3.3	This figure compares the point mass mRHI to the human mRHI experiment.	36
4.1	This flowchart shows how the point mass mRHI OLS algorithm connects to the computational principles of SoA and the agent’s prior.	50
4.2	This flowchart of an agent in the point mass mRHI that only uses OLS to finds its SoA.	50
4.3	These plots show that, for the static environment (experiment 1A), the agent’s prior knowledge does not significantly affect the estimation of the agency parameter \hat{b}_j because the SoA did not improve.	52
4.4	These plots show that, for the dynamic environment (experiment 1B), the agent estimation of the agency matrix $\hat{\mathbf{B}}$ is closer to the true value of the matrix \mathbf{B} with the prior than without the prior.	53
4.5	These plots show that, for experiment 1C, the agent’s prior improves estimation of the agency parameter \hat{b}_j through an improved SoA.	53
4.6	This plot shows, for the dynamic environment (experiment 1B), that the agent’s estimation of the agency parameter \hat{b}_1 is the same with and without prior.	54

4.7	These plots show that, for experiment 2A, the agent's prior knowledge motivates an incorrect estimation of the agency parameter \hat{b}_j meaning the SoA is also incorrect.	55
4.8	These plots show that in this ambiguous environment (experiment 2B) the agent's prior knowledge motivates an incorrect estimation of the agency parameter \hat{b}_j meaning the SoA is also incorrect.	56
4.9	These plots show that for the ambiguous asynchronous environment (experiment 2C) the agent's prior knowledge motivates an incorrect estimation of the agency parameter \hat{b}_j meaning the SoA is also incorrect.	56
5.1	This histogram shows that partial data \mathbf{X}_B can be approximated by two Gaussian functions.	63
5.2	This figure illustrates the probability distributions of the two clusters that the agent uses to describe all observed data points b_{ij}	72
5.3	This figure shows that the <i>Expectation-Maximization</i> (EM) algorithm correctly clusters the data into agency and no-agency clusters for the static environment defined by experiment 1A.	76
5.4	This figure shows the matrix size of the observed data, \mathbf{X} , that the agent uses to gain SoA over one of the three masses in the point mass mRHI. \mathbf{X}_B is dimension 0 and \mathbf{X}_c is dimension 1.	78
5.5	This plot shows how the agent could identify four hidden causes based on the observations b_{ij} and the perceived error $\hat{\mathbf{x}}$ in experiment 1B.	79
5.6	The flowchart connects SoA, the agent's prior, and the EM algorithm for the point mass mRHI.	85
5.7	This plot shows that the agent has a correct SoA over the observations influenced by the agent's input force, F_{int} , for trial 3 in experiment 1C.	87
5.8	These plots show that prior on the SoA parameter \hat{b}_j in the point mass mRHI promotes a correct SoA in experiment 1B.	88
5.9	These plots, like Figure 5.8, show that prior on the SoA parameter \hat{b}_j in the point mass mRHI promotes a correct SoA in a dynamic environment with noise (experiment 1C).	89
5.10	This figure shows that for the dynamic environment (experiment 1B) the EM agent correctly clusters the data according to their hidden causes, despite the incorrect initial conditions of the parameters.	91
5.11	This shows that in a dynamic environment with noise (experiment 1C) the agent also correctly has a correct SoA at the end of the simulation Figure b.	92
5.12	This figure shows that an agent using the EM algorithm without the prior knowledge that it can control one mass has less correct SoA over the observations from $j = 1$	93
5.13	This figure shows that the prior steers the agent away from the correct SoA for a synchronous, ambiguous environment.	95
5.14	This figure shows that in this ambiguous environment (experiment 2B), the prior induces an incorrect SoA over mass m_2 with parameter \hat{b}_{j3}	95
5.15	This figure shows that in this ambiguous environment (experiment 2C), the prior induces a correct SoA over mass m_1 with parameter \hat{b}_{j1}	96

5.16	This figure shows that in an ambiguous environment (experiment 2A) with the EM algorithm, the agent without the prior knowledge that it can control one mass has a less incorrect SoA over the observations from $j = 3$ but also less correct SoA over the observations from $j = 1$	97
5.17	This figure shows that in an ambiguous environment (experiment 2B) with the EM algorithm, the agent without the prior knowledge that it can control one mass has a less incorrect SoA over the observations from $j = 3$	97
5.18	This figure shows that in an asynchronous, ambiguous environment (experiment 2C) with the EM algorithm, the agent without the prior knowledge has a less correct SoA over the observations from $j = 1$	98
D.1	This figure shows that in a synchronous ambiguous environment (experiment 2A) the agent gains an incorrectly SoA over the data from $j = 3$	134
D.2	This figure shows that in the synchronous ambiguous environment (experiment 2B) the agent incorrectly clusters the data according to their hidden causes with the EM algorithm.	135
D.3	This figure shows that in this trial of experiment 2C the agent correctly clusters the data according to their hidden causes with the EM algorithm.	135
E.1	This figure shows that the estimation error of \hat{b}_{i1} is dependent on the initial covariance matrix, Σ , of the agency cluster.	138

LIST OF TABLES

3.1	Overview of the time and noise parameters of the point mass mRHI. . . .	38
3.2	Overview of the parameters per experimental condition (1A, 1B, and 1C) of the unambiguous environment, meaning the agent should have a correct SoA mass m_1	40
3.3	Overview of the parameters per experimental condition (2A, 2B, and 2C) of the ambiguous environment, meaning the agent should have a incorrect SoA over mass m_2	41
4.1	This table shows that for unambiguous environments the agent with a prior has a correct SoA over mass m_1	51
4.2	This table shows that, for the ambiguous environments (experiment 2) of the point mass mRHI, the agent using the OLS algorithm with a "prior" has a more incorrect SoA than the agent without any prior knowledge.	54
5.1	This table connects agency and SoA to the variables of the EM algorithm. .	67
5.2	This table shows that for unambiguous environments the agent with a prior generally has a correct SoA over mass m_1	86
5.3	This table shows that on average for the 1000 trials the EM algorithm (with prior) improves agent's estimation of the agency parameter \hat{b}_1 with time. The OLS algorithm (no prior) is more accurate throughout the simulation but the EM algorithm the agent is more accurate in the final time window of the simulation.	92
5.4	This table shows that for ambiguous environments the agent with a prior generally has an incorrect SoA over mass m_2 in the synchronous conditions (2A and 2B) and more correct SoA over mass m_1 in the asynchronous condition (2C).	94
C.1	This table compares the performance of the agent in the point mass mRHI using the OLS algorithm with and without a "prior" per experiment number of experiment 1.	128
C.2	This table compares the performance of the agent in the point mass mRHI using the OLS algorithm with and without a "prior" per experiment number of experiment 2.	129
D.1	This table shows for unambiguous environments the agent's agency improves with a prior w.r.t. without a prior.	132
D.2	This table shows for ambiguous environments the agent's agency deteriorates with a prior w.r.t. without a prior.	133

LISTINGS

5.1	Import libraries necessary.	71
5.2	Initial conditions of the EM algorithm.	73
5.3	Iteration loop of EM algorithm.	73
5.4	Expectation (E) step of the EM algorithm.	74
5.5	Maximization (M) step of the EM algorithm.	75
5.6	Python code for the part of the for loop in which the agent determines which cluster the agent has agency at each time step.	83

ACRONYMS

DEM	<i>Dynamic Expectation Maximization</i>
EM	<i>Expectation-Maximization</i>
mRHI	<i>moving rubber-hand illusion</i>
OLS	<i>ordinary least squares</i>
RHI	<i>rubber-hand illusion</i>
SoA	<i>Sense of Agency</i>
VB	<i>Variational Bayes</i>

ACKNOWLEDGEMENTS

I would like to thank both of my supervisors Martijn and Arkady for their guidance during my thesis. All our meetings, even though they were on Zoom, were very constructive. Martijn thank you for helping me to keep the bigger picture of my thesis in mind. With your critical questions and feedback, I always know what next steps to take. Also, thank you Arkady for being equally as involved in my thesis even after joining later. Your feedback, knowledge on Bayesian inference, and Python tips have been very helpful.

Also, I would like to thank my family and friends for supporting me and making my time as a student in Delft unforgettable.

Delft University of Technology
December 7, 2020

Meilin Tan

1

INTRODUCTION

1.1. MOTIVATION

IN recent decades, research has focused on building artificial agents that achieve intelligence similar to human intelligence. A central challenge in attaining human intelligence is understanding how we can build causal models, make generalizations, and abstractions in the presence of limited, noisy, or ambiguous input data [1]. If input data is missing, more abstract background knowledge is necessary for humans to make generalizations. Psychologists term this knowledge “constraints”, artificial intelligence researchers “inductive biases”, and statisticians “priors”. In any case, all fields recognize the importance of prior knowledge to make inferences that go beyond the data [1][2].

Despite the recognized importance of prior knowledge, artificial agents currently cannot use prior knowledge to make inferences that go beyond data like humans [3]. Prior knowledge required to learn a new task is often acquired from a previous related task. In search of achieving human-like learning, many learning algorithms successfully take their inspiration from biological systems, such as Hebbian learning [4], reinforcement learning [5], and deep learning [6]. Though the performance of agents may exceed human performance in specific tasks, even in seemingly simple tasks humans still outperform artificial agents. For instance, humans need fewer examples of handwritten characters than the best algorithms (e.g. [7]) to recognize and generate new examples of the handwritten characters [3]. Thus, to enable higher-level cognition in artificial agents, agents should use prior knowledge similar to humans.

In the context of robotics, using prior knowledge is an important step to developing an agent that reproduces the human experience of self. This is a popular challenge for roboticists today (e.g. [8], [9], [10], and [11]). In humans, the emergence of self-consciousness in its minimal form, or the *minimal-self*, is a key step in cognitive development. Gallagher describes the minimal-self as the “consciousness of oneself as an immediate subject of experience, unextended in time” [12]. The minimal-self is characterized by two aspects: the *sense of ownership*, or the feeling of mineness that we have over our body parts, and *sense of agency*, or the feeling of control of our own body parts.

In robots, skills such as self-other distinction and agency are valuable because these improve machine adaptability, human-robot interaction, and test the computational models for embodiment to improve understanding about these mechanisms [13][10]. However, reproducing the artificial self is difficult because agents do not inherently build a minimal self as humans do in the first months of life [14]. As a result, investigation on the artificial self is still incomplete [15]. I aim to explore the role of prior knowledge in the development of the artificial self. To limit the scope of this thesis, I will focus on *Sense of Agency* (SoA).

1.1.1. BAYESIAN INFERENCE UNITES PRIOR KNOWLEDGE AND SENSE OF AGENCY

Bayesian inference is a suitable model to research the role of prior knowledge in SoA in an agent because it provides a mathematical framework for combining prior knowledge and sensory evidence. According to Bayesian inference, the brain makes inferences or hypotheses about the hidden causes of uncertain sensory evidence using prior knowledge [16]. The agent also learns from observations by transforming a probability distribution that describes prior knowledge (defined before observing the data) into a probability distribution describing the inferences (defined after observing the data) [17]. Another advantage of Bayesian models is that they model uncertainty, which allows an understanding of the confidence in observations and predictions. Ghahramani argues that "intelligence relies on understanding and acting in an imperfectly sensed and uncertain world", meaning Bayesian models are possibly key for the development of intelligence in artificial agents [17].

Bayesian inference is applicable for reproducing a SoA in particular because theoretical models of SoA are generally Bayesian (e.g. [18] and [19]). However, literature still lacks computational principles that can reproduce a SoA [20]. From a Bayesian perspective, the SoA is an inferential process where the agent finds the probability that its current state is a result of their own action, given their prior knowledge and sensory evidence. In other words, the subjective feeling arises from the predictability of our own body compared to the unpredictability of others [21]. Computational models are an important tool for quantitatively testing the proposed principles that underlie a SoA. As of yet, only three computational models exist (i.e. [22], [21], and [20]), all of which use Bayesian mathematics.

1.1.2. THE EFFECT OF PRIORS ON INFERENCES

The role of prior knowledge in Bayesian inference is that it can promote convergence to both correct and incorrect SoA. Two fields of research explain why both aspects are an important function of the prior. While computational fields view the prior as helpful because it promotes correct inferences, in neuroscience the prior is important because it explains incorrect inferences in humans. In computational fields (e.g. statistics) a correctly chosen prior helps avoid overly rapid convergence to incorrect inferences and promotes more rapid convergence to correct inferences [23]. A prior also helps the agent avoid the trap of letting a few data points steer the agent away from the true conclusion [24][23]. Thus, inferences are a result of the constant interaction between priors and sensory observations, meaning a correct prior should promote the convergence to a correct

SoA.

On the other hand, in neuroscience, the prior in Bayesian inference plays a fundamental role in understanding human behavior because the prior explains why humans perform seemingly suboptimal behaviors. Different conclusions in the same environments are a consequence of applying the same rules of optimal statistical inference over different (suboptimal) priors [2]. The susceptibility of humans to sensory illusions is a fundamental example of seemingly suboptimal behaviors. Illusions are considered a result of artificial scenarios that violate prior beliefs [25][26]. In Hohwy's book, *The Predictive Mind*, he states "it would be strange if the brain had evolved a highly sophisticated inferential process to deal with a perceptual situation it encounters mainly in artificial laboratory settings" [27]. Therefore, illusions can reveal the constraints of the brain that have developed to efficiently format representations to sufficiently describe the external environment [28]. Bayesian inference can explain how an incorrect SoA results from the interaction between a prior and ambiguous and/or noisy sensory evidence.

1.2. PROBLEM STATEMENT

COMPUTATIONAL fields and neuroscience explain the general key role of priors in Bayesian inference. However, the implementation of priors in artificial agents poses two research challenges of a different nature. The first challenge pertains to all agents applying Bayesian inference including biological agents (such as humans), while the second is specifically a disadvantage for designers of artificial agents. I will explain both challenges and explain how this thesis will use SoA as a tool to investigate these challenges while also advancing research of SoA in artificial agents.

1.2.1. ADDRESSING TWO RESEARCH CHALLENGES OF PRIORS IN AGENTS

First, the strength of priors in Bayesian inference according to neuroscience (i.e. their ability to clarify suboptimal behavior), poses a challenge for the performance of artificial agents that apply Bayesian inference. In the presence of noisy and ambiguous sensory evidence, the prior can negatively influence the agent's inferential process. Depending on the application, incorrect inferences in an artificial agent, such as illusions, can lead to fatal mistakes [29]. Understanding the role of the prior in an illusion will therefore be useful to help understand the limitations of an artificial agent.

Box 1.1: The rubber hand illusions

The static rubber hand illusion (RHI). In the original (static) RHI, Botvinick and Cohen [30] show that humans can experience a rubber hand as their hand. In the experiment, the participant can see the rubber hand while their hand is hidden. To induce the illusion, the experimenter simultaneously strokes the participant's hand and rubber hand with a brush.

The moving rubber hand illusion (mRHI). Instead of synchronous visuo-tactile stimulation, the rubber hand and participant's hand move simultaneously because the hands are attached with a connective stick in the mRHI by Kalckert and Ehrsson [31][32]. The mRHI exists in two forms: in the (1) active mRHI the participant moves the rubber hand with their own input force, whereas in the (2) passive mRHI the experimenter moves both the rubber hand and the participant's hand by moving the stick [32].

To investigate this first research challenge, I reproduce the *moving rubber-hand illusion* (mRHI) (see Box 1.1) in an artificial agent because cognitive scientists use this experimental paradigm to research SoA [32]. In the RHI and passive mRHI experiment, participants experience a sense of ownership over the rubber hand. When participants move their own hand in the active mRHI, participants experience a *Sense of Agency* over the rubber hand. The mRHI has not yet been reproduced in an artificial agent. In this thesis, I design a simplified version of the active mRHI experiment for an artificial agent that is termed the "point mass mRHI". The SoA in the agent is measured using a parameter that reflects the relationship between the agent's actions and the sensory evidence from the environment. A correct estimation of the parameter translates to **correct SoA**, while an incorrect estimation is an **incorrect, illusory SoA**. The agent needs a parameter estimator that also allows separation of the data into "agency" and "no-agency" groups. I investigate whether a prior allows an agent to experience an incorrect, illusory, SoA in

some environments but an improved SoA in other environments.

The second research challenge for implementing a prior is the challenge for designers of artificial agents to translate subjective prior beliefs into a mathematical prior that can be included in an algorithm. The predictions of a Bayesian model often depend intensely on the prior distribution chosen by the designer [23]. As a result, within the field of statistics and engineering, Bayesian approaches are criticized for their necessity of a prior, arguing that the data should “speak for itself”. On the other hand, cognitive models try to formalize how human inferences use their prior knowledge. Choosing the appropriate prior distribution is important to capture the background knowledge humans have. However, the Bayesian framework does not explain how designers should select prior probability distributions. In cognitive modeling, the practical difficulties of how to define a prior often result in non-informative priors, which restricts the potential of cognitive models [33]. Conversely, the risk of incorrectly translating human background knowledge to an informative prior probability distribution is that the artificial agent makes incorrect conclusions. Therefore, researchers need to understand how to correctly translate subjective priors into mathematical priors.

To investigate this second research challenge, this thesis researches what prior beliefs allow humans to experience the mRHI and a SoA. In contrast to humans, there is no *a priori* developed SoA in robots [34]. Therefore, I identify the underlying principles of SoA that translate to a SoA computational model. One such principle is that the prior belief encodes the causal relationship between intended actions and predicted sensory outcomes. In the point mass mRHI, the agent’s prior knowledge is that it has agency over its hand. This thesis estimates the SoA parameter with two parameter estimators. First, as an elementary investigation into SoA I use the simple (and non-Bayesian) *ordinary least squares* (OLS). Second, for the more human-like model of SoA, the agent uses *Expectation-Maximization* (EM). Here, the agent can apply the prior causal belief in a Bayesian manner to test whether the same prior induces both mRHI and a SoA depending on the agent’s sensory evidence.

1.2.2. RESEARCH GOAL

The motivation of this thesis is to show that understanding priors is an important prerequisite in advancing the intelligence of artificial agents to human-level intelligence. While showing that implementing prior knowledge in an artificial agent improves SoA, this thesis will also research the two challenges. Hence, the **research goal** of this thesis is: *design an artificial agent with a Sense of Agency (SoA) where the same prior knowledge (i) improves SoA in normal environments and (ii) result in incorrect, illusory SoA in noisy, ambiguous environments.*

RESEARCH QUESTIONS AND THESIS OUTLINE

To achieve this research goal, I compose three research questions which I answer in the following four chapters of this thesis. Figure 1.1 is a simplified overview of the agent’s inferential process for SoA. The agent uses its sensory evidence from the environment and prior knowledge to make inferences about agency, which then updates the agent’s prior knowledge for the next time steps. Also, Figure 1.1 shows which research questions answer which part of the agent’s inferential process.

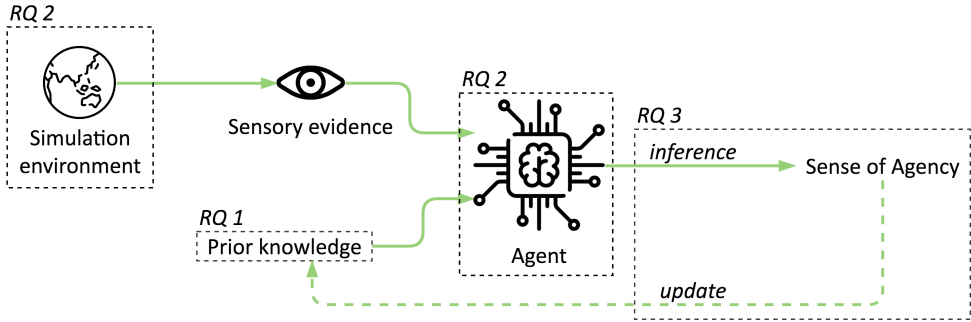


Figure 1.1. This figure provides a simplified overview of the agent's inferential process for *Sense of Agency* (SoA). Each part of the process is covered by a research question (RQ). The three research questions are as follows: *RQ 1* What prior knowledge enables a SoA in humans in the mRHI? *RQ 2* In what artificial environment, similar to the mRHI, should an agent experience a correct and incorrect SoA? and *RQ 3* What is the effect of the prior on a SoA in the point mass mRHI? Is this effect comparable to a SoA in human? For a further explanation of the research questions, see the main text.

RQ 1. What prior knowledge enables a sense of agency in humans in the moving rubber hand illusion? This initial research question provides the background knowledge necessary to design an agent with SoA. I present literature that discusses the Bayesian priors, SoA, and the mRHI.

Chapter 2 answers the first research question in three parts. First, I answer the question: *what is the role of the prior in Bayesian inference?* This question is answered according to four relevant research fields: statistics, cognitive science, neuroscience, and robotics. Using these four different research fields, establish the theoretical, cognitive, neural, and cognitive role of the prior. Each field highlights the value of the prior but also encounters at least one of the two aforementioned research challenges. Second, I answer the question: *what are the computational principles underlying a SoA?* Current literature does not agree upon one computational model of SoA because of the range of properties that influence a SoA. I summarize three underlying computational principles which interact to produce a SoA: prediction error, prior belief about causality, and precision of sensory signals. Third, I bring together the answers from the two previous sub-question to give a Bayesian explanation of how prior beliefs and the environment interact to produce correct SoA and incorrect SoA in the mRHI experiment.

RQ 2. In what artificial environment, similar to the moving rubber hand illusion, should an agent experience a correct and incorrect sense of agency? This second research question (mathematically) defines the model of the environment in which the agent should experience correct and incorrect SoA. This simulation, the "point mass mRHI", is a simplified version of the mRHI experimental paradigm by Kalckert and Ehrsson [31] (see Box 1.1).

Chapter 3 answers this research question in three parts: mathematically defining the point mass mRHI, comparing the point mass mRHI to human mRHI and SoA, and establishing the parameters that define the experiments that the artificial agent performs. First, in the point mass mRHI, the agent observes three point masses that are connected

to spring-damper systems. The agent's prior knowledge is that it can control only one of the point masses. Thus, the agent's goal is to discover which mass it has agency over by exerting a force on this mass. Mathematically, the agent's agency is described by a matrix that relates the agent's force input to the observed sensory signals. The prior defines that only one element of the agency matrix can be non-zero. The agent's SoA is determined by the values of the agency matrix. Second, the point mass mRHI is a simplified version of the experimental set-up of the mRHI because one mass represents the agent's real hand and the other two masses represent the rubber hands. Also, the agent uses the three computational principles of SoA to gain a SoA. Finally, I define the environments (i.e. experimental conditions) for which the artificial agent should experience correct SoA over its "real hand" and incorrect SoA over a "rubber hand" according to the results of the human mRHI experiments. These environmental conditions are defined by the parameters of the point mass and spring-damper systems.

RQ 3. What is the effect of the prior on a sense of agency in the point mass moving rubber hand illusion? Is this effect comparable to a sense of agency in human? If the answer to the research question is that effect of the prior, in the point mass mRHI, can result in both an improvement of SoA and incorrect, illusory SoA, then the research goal of this thesis is achieved. Furthermore, I compare the effect of the prior to results from human SoA and mRHI experiments to justify the point mass mRHI as a (simplified) model of SoA.

To gain SoA, the agent makes inferences about the agency matrix using a parameter estimator. The agency matrix describes the relationship between the agent's input force and the predicted sensory consequences. This thesis tests two parameter estimator to answer the third research question. First, Chapter 4 implements the point mass mRHI with *ordinary least squares* (OLS). This statistical method estimates the relationship by minimizing the sum of the squares over the difference between the observed and predicted values. This thesis analyzes whether this simplest estimator is sufficient for improved and incorrect SoA in the agent. This preliminary implementation of the point mass mRHI reveals what properties are necessary for an algorithm that allows the agent to experience a SoA.

In Chapter 5, I reproduce the point mass mRHI with the system identification method *Expectation-Maximization* (EM). EM is a popular (partially) Bayesian inference algorithm and appropriate because it can define SoA as a Bayesian inferential process. With the experimental conditions from Chapter 3, I find that the prior can promote both improved and incorrect, illusory SoA depending on the agent's sensory evidence. Chapter 6 concludes this thesis and defines recommendations for improvements to the EM algorithm as a model for SoA and topics for future research.

2

BACKGROUND: THE EFFECT OF PRIORS ON SENSE OF AGENCY

This chapter answers the first research question: *What prior knowledge enables a Sense of Agency in humans in the moving rubber-hand illusion?* At the end of this chapter the reader should have an understanding of the main concepts of this thesis, priors, Bayesian inference, SoA, and the *moving rubber-hand illusion* (mRHI), through the discussion of existing literature.

I answer this research question in three parts. Sec. 2.1 answers *What is the role of the prior in Bayesian inference?* according to four relevant research fields: statistics, cognitive science, neuroscience, and robotics. Each field highlights the value of the prior, but also encounters at least one of the two aforementioned research challenges. Next, Sec. 2.2 answers the question: *what are the computational principles underlying a SoA?* I look at empirical experiments such as the mRHI, theoretical, and computational models to conclude three computational principles that enable a SoA in humans. Sec. 2.3 brings together the two previous subsections to provide a Bayesian explanation of how prior beliefs interact with the environment to improve correct and enable incorrect SoA for a participant in the mRHI experiment. This section answers the overarching research question of this chapter. Finally, Sec. 2.4 concludes this chapter.

2.1. THE INFLUENCE OF THE PRIOR ON BAYESIAN INFERENCE

IN the introduction (Chapter 1), I establish that prior knowledge plays a key role in making inferences that go beyond the data [1], meaning that prior knowledge adjusts the inferences made by agents. But, *how* does the prior influence Bayesian Inference? I will explain the role of the priors according to four different fields of study. First, the field of statistics explains the theoretical role of the prior. Next, cognitive science researches the role of Bayesian priors in human cognition by developing Bayesian models of cognition that are supported by empirical evidence. Third, neuroscience researches evidence for Bayesian inference and priors in the brain. Lastly, engineering fields show the effects of implementing priors with Bayesian inference in artificial agents.

Each field researches the role of the prior in Bayesian inference but also encounters the two research challenges of different nature identified in the introduction (Chapter 1). Therefore, I will also explain the effect of the first research challenge, incorrect conclusions (specifically, illusions), and the second research challenge, the translation of subjective prior beliefs into a mathematical prior, on Bayesian inference for each research field.

2.1.1. STATISTICS: THEORETICAL ROLE OF THE PRIOR

Bayesian inference originates from the field of statistics and mathematically defines how to combine prior knowledge (or, in Bayesian terms “the prior”) with sensory information to make inferences about the underlying structure (hidden causes) that gave rise to the sensory data. Using the new prior knowledge of the hidden causes, an agent can make predictions about how new data is generated. Thus, statistics explains the theoretical role of prior knowledge in inference and Bayesian inference in the EM algorithm (Chapter 5).

THE THEORY BEHIND BAYESIAN INFERENCE

The Bayesian inferential process grows out of a formula called Bayes’ theorem (see Box 2.1). The inferences (posterior probability in Bayesian term) are updated according to Bayes’ theorem as more sensory evidence becomes available. This statistical model of measurement generation is also termed the generative model.

With Bayes theorem, there are three ways how the prior probability adjusts the posterior probability. Figure 2.1 illustrates the interaction of the Gaussian probability distribution of the prior belief and likelihood to cause an incorrect posterior belief. The width of the probability distributions is a function of their variance (or, the inverse of precision). In Figure 2.1, the prior moves the posterior away from the correct conclusion in three ways:

- ① Decreasing the precision of the likelihood shifts the posterior belief towards the prior probability. A decrease in sensory precision results from too much noise or uncertainty in sensory evidence, which causes an incorrect conclusion.
- ② Increasing the prior precision shifts the posterior towards the prior. An increase in prior precision occurs when the agent is certain about its prior knowledge, which means the agent is more likely to ignore or reject contradictory sensory information.

- ③ Shifting the sensory evidence distribution towards the prior probability also shifts the posterior belief towards the prior belief. This occurs when the environment seems to align with the prior.

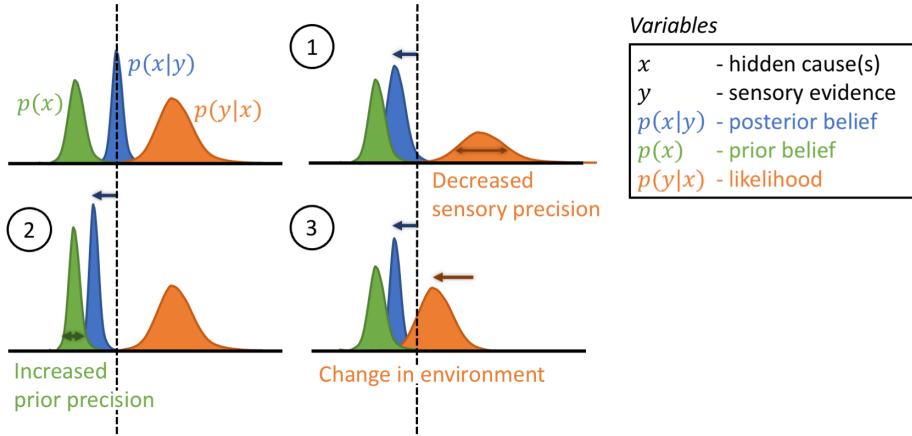


Figure 2.1. This schematic shows three cases when the posterior belief shifts towards the prior belief. The Gaussian probability distributions represent prior beliefs, posterior beliefs, and the likelihood of some sensory evidence as functions of the hidden cause, where the width of the probabilities corresponds to their variance (the inverse of precision). The dotted line corresponds to the posterior expectation. The relative variance or mean biases the posterior belief to the prior. See the main text for an explanation of the three cases.

Box 2.1: Bayes theorem

Bayes theorem indicates how an agent should revise their belief in a hypothesis given the observed data. Mathematically, the agent finds the probability that the sensory evidence, y , is a result of the hidden cause, x , by calculating the *posterior* probability, $p(x|y)$ with

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)}. \quad (2.1)$$

The *prior* probability, $p(x)$, is the probability that summarizes previous knowledge before receiving sensory information, or the prevalence of that hidden cause. The *likelihood*, $p(y|x)$, is the evidence provided by the senses to update the hypothesis. The *marginal likelihood*, $p(y)$, is the probability of observing the sensory evidence independent of whether the hidden causes occur.

In all three cases, the mean and relative variance of the prior steers the posterior towards a conclusion because the prior contains information that is not covered by the likelihood. This allows the prior to solve ambiguities that the sensory data cannot. Though Figure 2.1 only illustrates how the prior can move the posterior belief away from the cor-

rect conclusions, prior knowledge can facilitate the agent in making more accurate conclusions. According to Bayesian inference, both the relative mean *and* variance of the prior probability can shift the agent's posterior probabilities to both more accurate or incorrect SoA.

2.1.2. COGNITIVE SCIENCE: MODELS AND EMPIRICAL EVIDENCE

The field of cognitive science applies the statistics of Bayesian inference to model how prior knowledge allows people to learn from limited sensory data and interact with this data to guide generalization and action [23]. To establish that a Bayesian prior may affect SoA, I explain the role of Bayesian priors in both correct and incorrect inferences according to cognitive science research.

BAYESIAN INFERENCE IN HUMAN COGNITION

Before I investigate the role of the Bayesian prior in human cognition, it is necessary to establish that Bayesian inference can explain aspects of human cognition. A recurrent claim by cognitive scientists is that Bayesian inference provides a unifying framework for explaining human inferences in different settings (e.g. [2], [35], and [36]). This statement is supported by a range of human behaviors that researchers find consistent with Bayesian inference including decision making [37], cue combination [38], motor control [39], feedback response to force perturbations [40], and causal learning [41]. Previous studies also try to understand atypical human behavior such as neurological disorders with Bayes. For example, the incorrect balance of prior precision in Bayesian inference might explain symptoms of psychosis [42]. Adams et al. [42] reproduce one such symptom with a computational model, whose task is the visual pursuit of a target. Reducing the precision of high-level prediction errors (i.e. priors) causes both impaired smooth pursuit eye movements during occlusion of a target and the contradictory improvement of responses to unpredictable changes in the target direction. Such eye movements are some of the most reproducible signs of someone who has schizophrenia. Thus, as stated by Parr et al. [16] Bayesian inference "captures many different types of behavior, including apparently sub-optimal behavior".

Despite Bayesian inference being one of the most prominent theories of cognition and behavior, Bayesian models also receive several elaborate critiques. Jones and Love [43] find Bayesian inference conceptually trivial because it is just a "vote-counting" technique with unclear implications. Bowers and Davis' [44] main critique is that Bayesian models are unfalsifiable due to the flexibility of priors and likelihoods. Also, they find that Bayesian theories are rarely better at predicting data than other non-Bayesian models. I highlight these critiques so the reader has a more balanced view of Bayesian cognitive models. However, as there are a wealth of Bayesian models of human cognition including models for SoA and the RHI, I continue with Bayesian models as a probable model of human cognition.

THE EFFECT OF THE PRIOR ON HUMAN COGNITION

If Bayesian inference explains various behaviors, then it is probable that prior knowledge is applied in a Bayesian manner (Box 2.1). Studies show human integration of prior information is consistent with Bayesian inference in various tasks such as visual movement [45], force estimation [40], and motor adaptation [46]. Bayesian inference enables

an understanding of the prior assumptions that a human has, providing “a window into the mind” [17].

In many cases, prior knowledge improves conclusions that humans make, or, as previously stated, allows humans to make inferences that go beyond the data. For example, in a rapid pointing task, Tassinari et al. [47] find that participants optimally integrate uncertain stimulus information with the prior knowledge of the target according to the Bayesian model. This work shows participants rely more on prior information as sensory data of the target location becomes uncertain, resulting in more accurate pointing.

Besides the role of the prior in improving inferences, the prior also induces seemingly sub-optimal behaviors in humans. The *complete class theorem* [48] captures the idea that both optimal and sub-optimal behavior result from Bayesian inference, while behavior remains Bayes optimal. Therefore, any behavior can be computationally detailed to quantify the prior beliefs that subjects bring to any environment [49]. This theorem is particularly important to the field of computational psychiatry because the prior holds the key to understanding neurological disorders, which are sub-optimal behaviors. Another type of sub-optimal behavior that is more relevant to this thesis, is perceptual illusions. Illusions are hypothesized to be a result of the mean and/or variance of the prior probability shifting the posterior belief away from the likelihood. The theory behind Bayesian inference helps us understand how the mRHI occurs, which allows us to understand how prior knowledge can induce both correct and incorrect SoA.

To give the reader an initial understanding of how prior beliefs induce incorrect conclusions, I will explain how the visual phenomenon binocular rivalry occurs. Hohwy uses Bayesian inference to hypothesize that binocular rivalry result from the relative precisions of the likelihoods and the prior beliefs (case 1 and case 2 in Figure 2.1) of the different possible hidden causes (hypotheses) [25]. In binocular rivalry, each eye is presented with a different stimulus such as a face and a house. The subjective experience is an alternation between the house and the face. Figure 2.2 shows that there are three possible hidden causes (hypotheses) for what you are seeing: a face, a house, or a face-house. Seeing either a face or a house has a low likelihood while seeing a face-house has a higher likelihood. However, from prior knowledge you know that faces and houses exist (high probability), but face-houses do not (low probability). Therefore, the increased prior precision of the face or house hypothesis relative to the face-house hypothesis causes the brain to infer an alternation between seeing a face and a house. Similar to Hohwy’s hypothesis of binocular rivalry [25], the prior plays a key role in determining the posterior belief which causes the mRHI.

While Hohwy’s Bayesian model of binocular rivalry provides a *qualitative* understanding of the visual phenomenon, the strength of *quantitative* computational models is that they allow comparison between the artificial agent’s behavior and human behavior. It is unknown whether an agent behaving according to Howhy’s Bayesian explanation experiences the binocular rivalry. Parr et al.[50] reproduce binocular rivalry and Troxler fading with a computational model to find what prior belief activates the phenomena. This supports Hohwy’s binocular rivalry hypothesis [25]. The other reproduced visual phenomenon, Troxler fading, is an illusion that the artificial agent experiences with the prior belief that the world changes over time with a certain probability [51]. In another computational model, Brown et al.[26] define prior beliefs about the spatial



<i>Selection as driven by priors</i>	
Input: I	
Hypotheses	F+H: "It's a face-house" H: "It's a house" F: "It's a face"
Likelihoods	$P(I/F) = P(I/H) < P(I/F+H)$
Priors	$P(F) > P(H) \gg P(F+H)$
Perceptual inference	$P(F/I) > P(H/I) > P(F+H/I)$ 

Figure 2.2. This figure provides a Bayesian account of the binocular rivalry illusion, where the prior plays a key role. Reprinted from [25]. In binocular rivalry, one stimulus is selected rather than a blend of the stimuli. None of the hypotheses have both high likelihood and high prior probability. Therefore, if the “face-house” hypothesis has a low likelihood, the hypothesis with the highest prior probability wins .

gradients of illuminance and reflectance that are often observed in visual scenes. When presented with the stimuli, the computational model perceives the Cornsweet illusion and Mach bands. Also, Raman et al. [52] find filling in the blind-spot is a consequence of learning different natural visual scene. These computational models show the key role that priors can have in instigating incorrect conclusions, hinting that the prior plays an important role in the mRHI.

OBSTACLES IN FOR BAYESIAN MODEL REPRODUCING HUMAN COGNITION

Even though Bayesian inference helps cognitive scientists model the effect of priors on cognition, a challenge lies in defining a prior distribution. This research challenge, introduced in Chapter 1, is an obstacle for cognitive scientists. The source of this challenge is that the prior can be defined by any type of probability distribution, meaning that defining the prior requires more estimations than defining variables in deterministic models. An approach to solving this problem is suggested in [45], where Stocker and Simoncelli “reverse-engineer” the shape of the prior distribution of human visual speed perception through psychophysical experiments. However, such experiments cannot be done for all human behavior and it is especially more complex for higher-level cognition than human perception. On the other hand, Chater et al. [23] frame the challenge of prior definition as an advantage because the prior forces modelers to clearly state their assumptions about the shape and content of the agent’s prior knowledge.

Assumptions about the priors are often made using theoretical considerations or measurements from the agent’s environment [45], but these approximations might cause a discrepancy between the model’s behavior and human behavior. For example, Samad et al. [53] develop a Bayesian causal inference model of the RHI. Given the distance from the real body part and visuotactile delay (of the paintbrush, see Box 1.1), the model com-

putes the likelihood that an artificial body part belongs to the agent. Though results are qualitatively consistent with empirical findings, the model's proprioceptive drift is larger than the usual findings during the RHI. Schürmann et al. [54] hypothesize that the proprioceptive drift overestimation can be attributed to the chosen prior knowledge distribution. The uniform prior distribution in Samad et al. [53] is common in Bayesian data analysis, but are not necessarily a plausible quantification of prior knowledge in cognitive modeling [54]. Adjusting the prior to an informed distribution in the Bayesian causal inference model, Schürmann et al. [54] find more adequate predictions of proprioceptive drift in the rubber foot illusion. The explicit definition of the prior distribution (in [53]) allows other researchers (in [54]) to adjust the prior to find results similar to human behavior.

2.1.3. NEUROSCIENCE: BAYESIAN INFERENCE AND PRIORS IN THE BRAIN

Though there is ample evidence for Bayesian inference and priors in human behavior, less research is devoted to understanding how neural circuitry may be implementing Bayesian inference [2]. Strengthening the link between Bayesian inference and neural function improves the (biological) plausibility of Bayesian inference as a model for cognition. Pouget et al. [55] highlight that though neurons represent probability distributions, neuroscientists have to still understand to "what extent the brain uses them". As all neural circuits share similar features, these circuits may have one computational principle: probabilistic inference. Bayesian predictive coding is one way in which neuroscience tries to understand how the brain uses Bayesian inference [56]. Predictive coding is based on the idea that instead of representing the input directly, it is more efficient to represent the prediction error [57]. The strength of predictive coding is that it describes neural responses, which complements the strength of Bayesian inference (that it describes behavior). Prominent experimental support of combined Bayesian predictive coding comes from the visual cortex [58][52]. Recent applications of Bayesian predictive coding are the Free-Energy Principle [59] and its extension, Active Inference [60]. Friston describes the models as a "unified brain theory", which extends predictive coding [61].

The influence of the prior on neural mechanisms supports the possibility of Bayesian predictive coding in the brain. Hansen et al. [62] find that prior knowledge increases visual cortical activity when sensory evidence is difficult to perceive and decreases visual cortical activity when sensory evidence is easy to perceive but difficult to interpret (ambiguous). The effect of prior knowledge on the brain is dependent on the environment, which is consistent with Bayesian inference where the effect of prior knowledge is dependent on its distribution *relative* to the sensory evidence (recall Sec. 2.1.1). Another study that supports Bayesian priors in the brain finds that priors change the stimulus representations in the early visual cortex [63]. This suggests that probabilistic inference is also present in the early sensory regions and not just in higher-order neural areas. These results support that perception is a process of Bayesian inference, where top-down and bottom-up information is integrated at every level of the cortical hierarchy [59].

2.1.4. ROBOTICS: BAYESIAN INFERENCE AND PRIORS IN ARTIFICIAL AGENTS

In robotics, Bayesian approaches are used in a range of algorithms (e.g. from Bayesian reinforcement learning [24] to deep Active Inference [64]) to perform a range of tasks (e.g.

from sensory anomaly detection [65] to robot self/other distinction [66]). I highlight the strengths and weaknesses of Bayesian inference and priors in artificial agents because the agent will use Bayesian inference to achieve SoA (in Chapter 5).

2

ADVANTAGES AND DISADVANTAGES OF BAYESIAN INFERENCE IN ROBOTS

There are two other practical advantages of using Bayesian inference. First, Bayesian inference models uncertainty which is relevant for all artificial agents that interact with the real-world. Uncertainty in robotics can arise from five different factors: environments, sensors, actuation, models, and computation [67]. The explicit representation of uncertainties means robots can define the precision of incoming information and its degree of belief in this data. Second, robotics researchers often aspire to human performance. As cognitive science explores the possibility for Bayesian inference as a computational model of human cognition [1], robots might achieve human performance if the robots use such cognitive models.

The computational disadvantage of Bayesian inference is that in practice Bayesian approaches are no longer fully Bayesian. This problem exists because finding the posterior involves the marginalization of all the variables in the model except for the variables of interest. This becomes an intractable problem even in simple cases (see Box 2.2). To solve this issue, approximate integration algorithms are developed [17] such as Markov chain Monte Carlo (MCMC) methods, variational approximations, and expectation propagation. The former approximate integration method provides a guarantee of producing accurate results, while the latter two algorithms can be computationally competitive; the approximate inference methods are characterized by a speed-accuracy trade-off [68]. Also, the use of approximations weakens the argument that Bayesian inference provides a coherent way of representing beliefs. The coherence argument follows from Cox axioms [69], which states that the only coherent way of manipulating the 'degrees of belief' (ranging from 'impossible' to 'absolutely certain') is to have them satisfy the rules of probability, such as Bayes rule. Despite these disadvantages, these approximation methods are promising for advancing Bayesian research in robotics.

ADVANTAGES AND DISADVANTAGES OF BAYESIAN PRIOR IN ROBOTS

The effect of the prior in Bayesian inference in robotics can be separated into the advantages and disadvantages because the prior can both positively and negatively influence computation (i.e. improved or incorrect inferences).

The computational advantage of the prior is that it enables a robot to make inferences that go beyond the data [1], which I also discuss in the theoretical explanation of Bayesian inference according to statistics (see Sec. 2.1.1). This implies that the prior should allow the robot to make *better* conclusions than without prior knowledge. To strengthen this statement, Chappell et al.[70] find that when priors are non-informative, the parameter estimates obtained from Variational Bayes (a method providing quick Bayesian inference by estimating the parameters of a factorized approximation to the posterior distribution) are the same as those found through nonlinear least-squares. With an informative prior, the parameters are estimated with greater confidence especially in the presence of poor data. Also, Ghavamzadeh et al.[24] describe the effect of the prior as "regularization". By defining a prior on, for example, the model parameters,

a few data points can no longer move the robot away from the true parameters and the prior prevents overly rapid convergence (as stated in the introduction).

The computational disadvantage of the prior is that again the robot designers have to define the prior. Different from cognitive scientists, in the field of robotics the priors do not have to reflect human prior knowledge, but rather should result in the desired behavior of the robot. Incorrect choice of the prior might lead to the undesired behavior of the robot. Beal categorizes assigning priors into three schools of thought (see pg. 27 of [71]). *Subjective priors* “encapsulate prior knowledge as fully as possible” through experimental data or expert knowledge [71], whereas *objective priors* (or, non-informative priors) try to minimize the influence on the posterior distribution. Third, *hierarchical priors* put priors on the prior to build a hierarchical Bayesian model. Modelers can continue to put prior on prior until the top level of priors is vague enough. These different types of priors highlight how the subjectivity of the prior is an obstacle for the implementation of Bayesian inference in robots. In this thesis, I will define a subjective prior because the goal is to reproduce human SoA in an agent.

Box 2.2: Bayesian inference is an intractable problem

To calculate the denominator of Bayes theorem, the sum rule states that the marginal probability of y is obtained integrating (or summing if the variables are discrete) the joint over x . Therefore, that Bayes theorem calculates the posterior probability, $p(x|y)$ with

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} = \frac{p(y|x)p(x)}{\int p(y|x)p(x)dx}. \quad (2.2)$$

Calculating the marginal integral is analytically intractable or in the discrete case the number of calculations may grow exponentially with the number of states.

2.1.5. SUMMARY: THE INFLUENCE OF THE PRIOR ON BAYESIAN INFERENCE

To summarize, an overview follows of the answers per research field to the question: *how does the prior influence Bayesian Inference?*

STATISTICS

- Assuming Gaussian distributions, the relative mean and precision of the probability distributions of the *prior* and the *likelihood* determines the *posterior* (or, inference) according to Bayes theorem.

COGNITIVE SCIENCE

- Prior knowledge allows humans to make inferences that go beyond the sensory data.
- Using Bayesian inference, the complete class theorem explains *seemingly* sub-optimal behavior results from incorrect priors for that environment (e.g. illusions) [48].

- The discrepancy between a computational model and human behavior might result from defining an inaccurate prior through assumptions made using theoretical considerations or measurements from the agent's environment [45].

NEUROSCIENCE

- Empirical evidence exists that supports that prior knowledge adjusts neural activity in the brain according to Bayesian inference [62][63].

ROBOTICS

- The prior can improve the results in robots and has a "regularization" effect [24].
- Three schools of thought exist on how the prior should be defined [71], highlighting that the subjectivity of the prior is an obstacle for implementing Bayesian inference robotics and a reason that the prior can motivate incorrect inferences.

2.2. PRINCIPLES UNDERLYING THE SENSE OF AGENCY

RECENTLY, SoA has become a popular topic of study across various disciplines such as psychology, cognitive sciences, neuroscience, and robotics. This section identifies the computational principles that underlie SoA (for a definition see Box 2.3) to understand how to reproduce SoA in an agent. Though SoA was initially linked only to motor control, SoA is likely also a result of intentional aspects outside of our body. The lack of one coherent model of SoA means the answer to this question is not straightforward. I will present this answer in two parts: (1) a general overview of research on SoA and (2) three main principles that adjust SoA.

Box 2.3: Definition of Sense of Agency

Agency is the *objective* measure of control. *Sense of Agency* (SoA) is the *subjective* feeling “*I am the one generating this action*” [12][18][72]. Therefore, SoA is a result of voluntary actions [73] and allows an agent to distinguish between their own actions and actions generated by others [12][74].

2.2.1. GENERAL OVERVIEW SENSE OF AGENCY (SOA)

SoA is typically an unconscious experience, making it difficult to measure in experiments [74]. Unlike for conscious experiences, such as vision, experimenters have to develop paradigms through which to detect SoA.¹ I will present both theoretical and computational models of SoA with experimental investigations that support the models.

THEORIES: UNDERSTANDING HOW SOA EMERGES

In the search for understanding how SoA arises and what factors influence the subjective feeling, researchers have developed theoretical models. Two influential models of SoA, “the comparator model” [75] and the “theory of apparent mental causation” [76] provide competing views of SoA, which are more recently brought together in the “cue integration theory” [18].

According to the traditional model of motor control, *the comparator model* [75][77], SoA occurs when there is a “match” between the internal motor signals and actual sensory outcome. This model assumes that intended actions are accompanied by an efferent copy, which computes the expected sensory consequence. Comparing the efferent copy to the actual feedback allows the agent to distinguish between internally and externally generated sensations. Consistent with this theory, internal motor signals are important in SoA. To highlight a relevant example, in the mRHI, Kalckert and Ehrsson [32] find that participants experience a greater SoA in the active condition, where their hand and rubber-hand are moved by their own actions (i.e. voluntary actions), as opposed to the passive condition, where both hands are moved by the experimenter (i.e. involuntary action).

In contrast, the *theory of apparent mental causation* [76] rejects that internal cues, the motor signals, play such a strong role in the SoA. Rather, this theory hypothesizes

¹See Appendix A for a discussion of the different measures of SoA in experimental investigations and how robotics uses these measures to further research on the artificial self.

that SoA occurs from the relationship between the thought and the action, which emphasizes the role of external cues [78]. Quoting from Moore [74], “if our intention [belief about an action] to act happens before we act, is consistent with the action, and is the only plausible cause of the action, then we feel as though we have caused the action”. In this postdictive process, the agency is inferred more after the movement [76] and influenced by higher-level causal beliefs [21]. To illustrate, Wegner et al. [76] use priming to manipulate the prior conscious thought about an effect. This external cue (the prime) generated a false SoA as a result of movements that participants had not performed.

Recently, much SoA research agrees that both internal and external cues contribute to the SoA [18][79]. The *cue integration theory* of agency unites both theories by stating that various sources of information are responsible for the SoA and the reliability of the sources of information determines their contribution to SoA [18][78]. To support the influence of weighting and integration of cues, Moore et al. [18] found that external cues (that influence prior beliefs such as priming) have more effect on intentional binding if internal (predictive) cues are missing. Also, in line with the Bayesian framework, Moore et al. [78] support that SoA is based on the integration of cues weighted according to their precision. The extension of the cue integration theory to the Bayesian framework allows the mathematical incorporation of the prior. However, the origin of the prior is a “recurring theoretical problem” [19].

COMPUTATIONAL MODELS: PUTTING THE THEORETICAL MODELS TO PRACTICE

Despite the substantial amounts of literature debating the theoretical model of SoA, only three computational models exist which try to formalize the computational principles underlying SoA. Computational models support an easier adaptation of the principles for SoA in an artificial agent. These three models use Bayesian, predictive coding, and/or Active Inference (i.e. influential models of the brain in cognitive science and neuroscience) to identify the computational principles that might underlie SoA.

Legaspi and Toyoizumi [20] support the cue integration theory with a Bayesian model. The SoA is defined by the confidence in the causal estimate (CCE), which is high when (i) sensory signals are consistent with the action-outcome effect, (ii) the causal prior probability is high, and (iii) the sensory inputs are precise. The model reproduces the empirical results from two studies on intentional binding (i.e. [73] and [80]). Intentional binding is the “perceived compression of the time interval between voluntary action and its outcome” [19].² This confirms the factors contributing to the CCE are relevant to SoA, meaning SoA emerges when there is a causal relationship between voluntary actions and the sensory signals and by the precision-dependent integration of different agency cues. Drawbacks are that the model does not use spatial but only temporal aspects of action and the model is specifically designed to reproduce the empirical results from [73] and [80] (i.e. it cannot be easily generalized other actions of agents).

Second, the aforementioned Active Inference model [22] introduces the role of action in SoA. In Active Inference, intentions, or prior beliefs about actions, are referred to as the *control* states, whereas *actions* are the physical states of the real world. SoA emerges as a result of minimizing the divergence between the predicted outcomes of available policies for action and the prior beliefs about future sensory states. The prior

²See Appendix A for a more elaborate explanation of intentional binding.

belief is defined by the future states that an agent will occupy given the current states. This shows how the agent's prior belief of agency plays an important role in decreasing the prediction error. The prediction error can be minimized through perceptual and active inference. In perceptual inference, the agent updates its expectations of agency while passively receiving sensory input, whereas in active inference the agent selectively samples sensory input to satisfy beliefs about sensory input in future states, given certain actions [81][72]. A drawback of the Active Inference model is that there is no measure of SoA in the model.

The third model from Kahl and Kopp[21] use principles of Predictive Coding [58] and Active Inference [60] to model the human sensorimotor system, where the prior belief is updated through a hierarchy of different motor beliefs. For each higher level of the hierarchy of the model, the Hierarchical Predictive Belief Update model (HPBU), increases in abstractions over descriptions of movement [21]. The SoA is updated through two processes. The predictive process is based on the causes of the actions and uses forward models to predict sensed action-outcomes. The postdictive process observes action outcomes and then applies higher-level causal beliefs and inference to determine an explanation. An interesting addition that the model makes is a Kalman filter that estimates the accumulation of SoA over time. The filter depends on the current free-energy and precision of the sequences of motor acts. Thus, the model suggests that at the core of a SoA is "the predictability of our own body".

2.2.2. THREE UNDERLYING COMPUTATIONAL PRINCIPLES OF SOA

Based on the previously presented theoretical and computational models, I conclude three main computational principles that increase SoA in humans. I discuss the influence of these three computational mechanisms on SoA according to current robotics research.

1. LOW PREDICTION ERROR

People experience a stronger SoA when their intended actions caused sensory outcomes that were predicted by their prior knowledge (lower *prediction error*), compared to when the effects were not predicted by the prior knowledge (higher prediction error) [82]. The SoA is minimized through perceptual and/or active inference. The role of *perceptual inference* in SoA is to distinguish between sensations that it can and cannot control (i.e. low vs. high prediction error) [72]. The role of active inference in SoA is to actively choose actions that minimize the prediction error [22][72]. Thus, the prior belief of agency can affect the prediction error in three ways:

- 1 A correct choice of agency should result in a decrease in prediction error [72].
- 2 A correct choice of no-agency should not affect the prediction error minimization because the agent cannot control the object to reduce the prediction error [72].
- 3 An incorrect choice of agency (i.e. the intended action does not cause the predicted sensory outcome) should increase the prediction error, which, if large enough, causes an adjustment in the prior belief of the hidden causes [83].

Robotics research shows that the prediction error affects sensory attenuation which is thought to be an implicit measure of SoA [8][9][84]. To illustrate, Schillaci et al.[8] apply a forward model to a robot to distinguish between self-produced and externally

produced signals. Using the robot's ego-noise, the robot attenuates signals with a small prediction error (self-produced) compared to signals with a large prediction error (externally produced). Similarly, Lang et al.[9] use sensory attenuation in a robot to suppress the predicted visual consequences of an intended movement, which allows for the sense of object permanence. The sense of object permanence is correlated to SoA in developmental psychology. Calculating the prediction error of consequences of self-generated actions is argued to be the first step in understanding the self in an artificial agent[85].

2. STRONG PRIOR BELIEF THAT INTENDED ACTION CAUSES THE OUTCOME

The role of the prior belief on SoA is defining the causal relationship between intended action and sensory outcome [19]. Both prior belief of *causality* and the *intention* of action are necessary conditions for SoA. The strength of the prior belief influences the SoA as follows: an agent acting on an environment with a more precise priors (smaller variance) experiences a stronger SoA because it is less likely to change its prior beliefs and more likely to change the environment [20].

To illustrate the importance of causal prior beliefs on SoA in robots, Ohata and Tani [11] implement a Predictive-coding-inspired Variational Recurrent Neural Network (PV-RNN) model to research the effect of the precision of the prior belief on the imitative interaction between a robot and a human. The researchers postulate that a robot that imitates a human has to adapt its internal states, therefore, having less SoA than a robot that acts egocentrically (changing the environment). To adjust the weight of the prior belief, the researchers adjust the tightness of the complexity term. The complexity is the Kullback-Leibler divergence between the approximate posterior and prior, meaning that by increasing the tightness of this term the importance of the prior decreases. With the decreased importance of the prior, the agent imitates the human counterpart more meaning a lower SoA.

3. PRECISE SENSORY SIGNALS

An agent acting on an environment that receives less *precise* sensory signal (larger variance) will have a lower SoA. Also, Perrykad et al.[72] find that precision of the sensory signals influences the selected action. In an uncertain environment, participants were more likely to adjust their actions to achieve a SoA. Perrykad et al. interpret this as active inference because the participants face accumulating prediction error meaning the participants will not achieve a SoA under the current chosen action. Moreover, the Bayesian model by Legaspi and Toyoiuzumi[20] predicts that with a small uncertainty of the sensory input signals, even involuntary causal action result in strong SoA. However, this prediction awaits empirical confirmation.

To my knowledge, the influence of the precision of sensory signals on SoA is not tested in robots. However, the influence of an incorrect sensory signal is tested in Kahl and Kopp's[21] SoA computational model. First, the artificial agent learns the visual and proprioceptive spatial and temporal dynamics of the handwritten digits from 0 to 9. Then, the agent has to reproduce the handwritten digits while receiving the correct proprioceptive feedback with the visual feedback of either the same or a different digit. The agents receiving incorrect feedback perceive a much lower SoA than when receiving the correct feedback.

2.2.3. CONCLUSION: CONNECTING PRINCIPLES OF SOA TO PRIORS

In conclusion, the three computational principles determine the agent's subjective SoA that the agent *believes* describes its agency over an action (see Figure 2.3). However, the agent's belief about its agency (SoA) does not always align with the agent's true agency. Therefore, the three computational principles do not influence the agency but the agent's SoA. I will use these three computational principles in the agent to determine when the agent has SoA (starting in Chapter 3).

While the three computational principles influence SoA, the principles are also connected because the prediction error calculates the difference between the predicted sensory outcome of the intended action and the observed sensory signal (see Figure 2.3). When a sensory signal has more noise or the agent makes an inaccurate prior prediction, then the prediction error increases, meaning a decrease in SoA. However, I keep the computational principles separated into three principles rather than just one for the prediction error. The reason is that the prediction error does not encode the relative precisions of the sensory signals and the predicted outcome. The precisions are important to connect the principles to Bayesian inference. Another reason is that the prediction error does not include the causality of the prior belief.

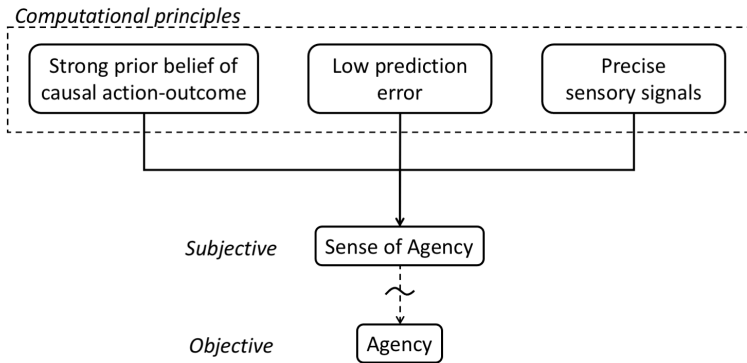


Figure 2.3. This figure provides an overview of the connection between agency, SoA, and the three computational principles of SoA.

2.3. BRINGING TOGETHER PRIORS, THE SOA, AND THE MRHI

IN this section, I answer the main research question of this chapter: *What prior knowledge enables a sense of agency in humans in the moving rubber-hand illusion?* I can already answer this question based on the previous subsections. The prior enables SoA in two ways: (i) the correctness of the prediction by the prior belief (i.e. prediction error) and (ii) the strength of the causal prior belief. To understand the implication of this prior in the context of the mRHI, this section provides a Bayesian explanation for both improved and incorrect SoA.

2.3.1. THE MRHI

To understand all the components of the mRHI, I elaborate on the mRHI explanation from the introduction. Figure 2.4 shows the experimental set-up for different experimental conditions of the RHI [32]. Figure 2.4A shows the overall experimental set-up, where Figure 2.4B and Figure 2.4C shows the active and passive mRHI, respectively. In the active mRHI the participant moves a finger of the rubber-hand with the input force from their finger. In the passive mRHI, the experimenter moves both the finger from the rubber hand and the participant's finger with the connective stick.




Figure 2.4. This figure shows the general A and three B-D different experimental conditions of the *moving rubber-hand illusion* (mRHI). Reprinted from [32]. A Experimental setup with the occluding cloth as used in the experiment. B Active *moving rubber-hand illusion* (mRHI): the participant taps the finger. C Passive mRHI: the experimenter moves both the rubber hand's and participant's fingers by moving the connecting stick. D The RHI: the experimenter strokes both fingers with a small brush (not discussed in this thesis).

Results from the mRHI experiment research two of the computational principles of SoA³. First, SoA occurs when there is a low prediction error. Participants only experience

³Another interesting result is that participants experience a SoA regardless of whether the rubber-hand is in an

SoA over the rubber hand when there is temporal congruence between the hand and the rubber hand [31]. Second, SoA disappears when the movement is passively induced by the experimenter (Figure 2.4C). Thus, the participant must have the intention to move, which is likely linked to their understanding of the causal origin of actions. The third computational principle, the precision of the sensory signal, is involved in the process of experiencing mRHI. However, the influence of the precision of sensory signals is not used as an independent variable in the mRHI.

2.3.2. SIMPLE BAYESIAN EXPLANATION OF SOA IN THE MRHI

I translate the original experimental set-up for the mRHI to a virtual environment where two rubber-hands and an agent's real hand are on a table. This virtual environment will become the "point mass mRHI" starting in Chapter 3.

While the two rubber hands are moving, the agent uses Bayesian inference to determine over which hand it has agency. Two experimental conditions result in correct and incorrect (illusory) SoA. Figure 2.5 shows the agent's view of the hands for the two experimental conditions: the normal condition (a) and the mRHI condition (b).




Figure 2.5. This figure shows the two experimental conditions in which the agent uses Bayesian inference to determine over which hand it has agency. In the first condition, the agent sees all three hands ($h = 1, 2, 3$) while in the second condition the agent's hand ($h = 1$) is hidden under the table (similar to the mRHI). The agent receives three cues: a visual v_h , proprioceptive p , and sensorimotor s , where the former cue is external and the latter two are internal cues.

The participant determines its SoA over each hand using the posterior probability from Bayes' theorem (similar to the Bayesian cue integration framework described in

anatomically plausible position [31]. This highlights the difference between a SoA and a sense of ownership. The feeling of a sense of ownership increases when the rubber-hand is in an anatomically plausible position and the movements are synchronous but do not require the movement of the rubber-hand to be a result of voluntary control [32].

[19]). The hidden cause of the sensory evidence, A_h , is agency over each hand, $h = (1, 2, 3)$. The sensory evidence is the internal and external cues that an participant receives about the states (i.e. position, velocity, acceleration) of each hand. For simplicity, I assume that the participant receives three cues: a visual \mathbf{v}_h , proprioceptive \mathbf{p} , and sensorimotor \mathbf{s} , where the former cue is external and the latter two are internal cues. Using Bayes theorem, SoA, or the posterior probability that the agent has agency over hand h is

$$p(A_h|\mathbf{s}, \mathbf{p}, \mathbf{v}_h) = \frac{p(\mathbf{s}, \mathbf{p}, \mathbf{v}_h|A_h)p(A_h)}{p(\mathbf{s}, \mathbf{p}, \mathbf{v}_h)}, \quad (2.3)$$

where the likelihood, $p(\mathbf{s}, \mathbf{p}, \mathbf{v}_h|A_h)$, describes the probability of the sensory evidence from the cues occurring simultaneously given that agent has agency over that hand, $h = (1, 2, 3)$. The prior probability, $p(A_h)$, describes the prior belief of the hidden cause and $p(\mathbf{s}, \mathbf{p}, \mathbf{v}_h)$, or the marginal likelihood, describes the probability of observing the states of the hand independent of whether the agent has agency over the hand.

In line with Friston et al. [22], the prior of agency, $p(A_h)$, is based on the prior beliefs about how the hand will move in the future. In other words, $p(A_h)$, is the prior knowledge on the causal relationship between intended action and outcome, which is one of the computational principles of SoA (Sec. 2.2.2). Assuming the agent does not have SoA over any of the hands before the experiment, the agent easily gain a SoA by shortly moving their hand in a certain direction with its prior knowledge of the causal relationship.

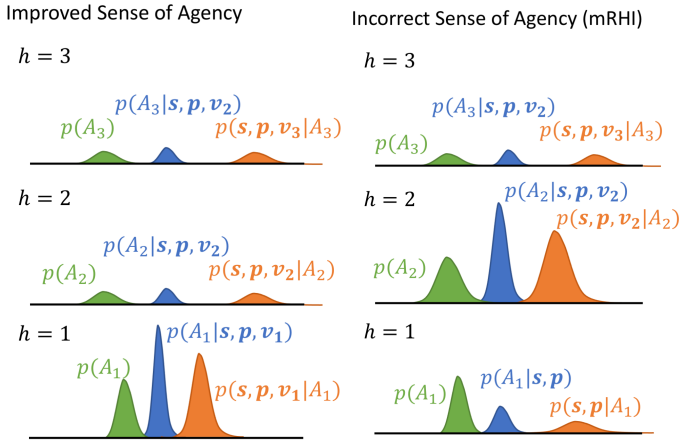
HOW THE PRIOR IMPROVES SENSE OF AGENCY

To illustrate how the prior improves SoA in Figure 2.5a, I have drawn the agent's hypothetical posterior probabilities for the agency over each hand in Figure 2.6a. Initially, the agent moves their hand and only knows the causal relationship between action and outcome. The agent observes a movement in their hand ($h = 1$) that corresponds to the prior belief of the outcome given the intended action (low prediction error). Therefore, the likelihood, $P(\mathbf{s}, \mathbf{p}, \mathbf{v}_1|A_1)$, is higher than the likelihood of both rubber hands. This increases the posterior belief that the agent has agency over the real hand $P(A_1|\mathbf{s}, \mathbf{p}, \mathbf{v}_1)$. In the next time steps, the agent's prior belief $P(A_1)$ is continually updated with the posterior belief from the previous time step $P(A_1|\mathbf{s}, \mathbf{p}, \mathbf{v}_1)$, meaning the agent has increasingly more SoA over their hand.

The prior improves the agent's SoA at two moments. First, at the beginning of the experiment, the agent would not know that their muscle input causes some sensory output. The agent would have to learn, e.g. through correlation of input and output, that they can control their hand. Second, continuously updating the prior belief means that later estimations of SoA become even more accurate. The strong prior belief of agency will allow the agent to believe that they can control their hand without even moving their hand (similar to the effect of the primed prior in [76]); the agent makes an inference that goes beyond the data.

HOW THE PRIOR PROMOTES INCORRECT SENSE OF AGENCY

To illustrate how the prior knowledge steers the posterior towards incorrect SoA in Figure 2.5b, I have drawn the agent's hypothetical posterior beliefs for the SoA over each



(a) The participant is subject to the first experimental condition as demonstrated in Figure 2.5a. The posterior probabilities stated in Figure 2.5a show that the participant has SoA over its bilities show that the participant has agency over the first rubber hand, $h = 2$.
 (b) The participant is subject to the second experimental condition as demonstrated in Figure 2.5b. The posterior probabilities stated in Figure 2.5b show that the participant has agency over the first rubber hand, $h = 2$.

Figure 2.6. A schematic of the (hypothetical) posterior probabilities after some time for the participants hand, $p(A_1|s, p, v_1)$, rubber hand 1 $p(A_2|s, p, v_2)$, and rubber hand 2 $p(A_3|s, p, v_3)$

hand in Figure 2.6b. For the mRHI experimental condition, I make two adjustments to the virtual environment. First, the agent places their hand under the table. Then, the agent can no longer use the sensory evidence from their vision, v_1 to determine the hand position. The agent can only estimate the position of the hand with the less precise sensory signals proprioception and sensorimotor predictions. This decreases the precision of the marginal likelihood for the real hand, $p(s, p|A_1)$. Second, the first rubber hand ($h = 2$) now moves synchronously with the agent's hand. With the prior knowledge that there is a causal relationship, the agent concludes that they have agency over the rubber hand (lowest prediction error).

2.4. CONCLUSION

In conclusion, I will briefly answer the main research question: *What prior knowledge enables a sense of agency in humans in the moving rubber-hand illusion?* In short, the prior knowledge of the causal relationship between intended actions and sensory outcomes. When this prior belief aligns with observed sensory signals (low prediction error) the SoA increases. Independent of the environment, the strength of this prior belief is also correlated to the strength of the SoA.

The more extensive answer to this question requires an understanding of the main concepts of this thesis: priors, Bayesian inference, SoA, and the *moving rubber-hand illusion* (mRHI). This is addressed by two sub-questions. First, *how does the prior influence Bayesian inference?* In four different research fields (statistics, cognitive science, neuroscience, and robotics), the prior allows (Bayesian) inferences that go beyond the data, but also induces sub-optimal behavior such as illusions. Computational models of the brain and robotics research find that the subjectivity of the prior is an obstacle for the implementation of Bayesian inference models. Second, I answer *what (computational) principles induce a SoA?* Though SoA is a complex process, I conclude three main principles underlying SoA: (i) the prediction error between intended actions and sensory outcomes, (ii) the prior belief of the causal action-outcome relationship, and (iii) the (precision of) sensory signals influence SoA. These principles align with Bayesian inference. The content of this chapter provides the foundation for understanding how to test SoA in an artificial agent with the mRHI.

3

EXPERIMENTAL DESIGN: DEFINING THE POINT MASS MRHI

This chapter answers the research question: *In what artificial environment, similar to the moving rubber hand illusion, should an agent experience correct and incorrect sense of agency?* To answer the question, the chapter translates the knowledge on priors and SoA from the Chapter 2 into the point mass mRHI. This chapter is named “Experimental Design” because Chapter 4 and 5 use the simulation defined in this chapter to show improved and incorrect SoA in an artificial agent.

The first section (Sec. 3.1) describes the model of the artificial environment, i.e. the “point mass mRHI”, in mathematical detail. The model simplifies Kalckert and Ehrsson’s[31][32] mRHI experiment for an artificial agent that observes three (point) masses connected to spring-damper systems. Next, Sec. 3.2 compares the point mass mRHI to the human mRHI experiment. This comparison explains, first, how the point mass mRHI is an analogy of the human mRHI experiment and, second, how the point mass mRHI captures the underlying computational principles of SoA defined in Sec. 2.2.2. Third, Sec. 3.3, mathematically answers the research question because I define the parameters of experimental conditions in which the agent should experience correct and incorrect SoA over one of the masses. Finally, Sec. 3.4 concludes this chapter.

3.1. THE POINT MASS MRHI: THE MATHEMATICAL MODEL

To explore the effect of prior knowledge on SoA in an artificial agent, I simplify the mRHI experiment by Kalckert and Ehrsson[31][32] to the point mass mRHI. This section describes the point mass mRHI in mathematical detail and defines the matrix that will describe agency and SoA.

3.1.1. THE AGENT'S MODEL: INITIAL EXPLANATION

The goal of the agent in the point mass mRHI is to achieve a SoA over one of the three point masses (from now on masses). Figure 3.1 illustrates that each mass is attached to a spring-damper system. The agent achieves SoA by learning on which mass it applies its force input, F_{int} , and how it controls this mass with force F_{int} . In the point mass mRHI, the agent can control the first mass, m_1 , because this is the mass on which it applies its input force F_{int} (see Figure 3.1). However, the agent does not know that it applies force F_{int} on mass m_1 . By estimating the agency matrix the agent will try to uncover how its force input adjusts the observed states of the masses. To prevent the agent from trying to gain agency over all three masses, the agent's prior knowledge is that it knows that it can have agency over exactly *one* mass.

Though the agent has an incomplete model of the environment, the agent does know (i) the states of the three masses, (ii) the parameters of each mass, and (iii) the force that it applies. The agent observes the states of the three masses as depicted in Figure 3.1. The states of each ($s = 1, 2, 3$) mass, are the displacement x_{m_s} , velocity \dot{x}_{m_s} , and acceleration \ddot{x}_{m_s} with respect to the equilibrium position. Second, the agent knows its input force F_{int} . Third, the agent knows the parameters of each mass. The parameters are the spring and damper constants, k_s and c_s , and the mass m_s .

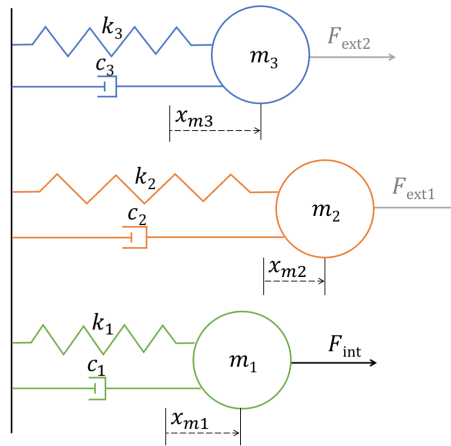


Figure 3.1. This figure illustrates the environment that the agent observes. The environment consists of three masses ($s = 1, 2, 3$) attached to spring-damper systems, where m_s is the mass, k_s is the spring constant, c_s is the damper constant, and x_{m_s} is the distance from the equilibrium position to the mass. The agent applies force F_{int} , while the external forces F_{ext1} and F_{ext2} are applied by the environment.

3.1.2. THE PLANT: GENERATING THE ENVIRONMENTAL STATES

The first part of the mathematical model of the point mass mRHI is the plant. The plant generates the states of the three masses, which describe the environment that the agent observes. The states are defined by an equation of motion for each 2nd order mass-spring-damper system excited with an input signal:

$$\ddot{x}_{m_1}(t) = \frac{1}{m_1}(F_{\text{int}}(t) - c_1 \dot{x}_{m_1}(t) - k_1 x_{m_1}(t)), \quad (3.1)$$

$$\ddot{x}_{m_2}(t) = \frac{1}{m_2}(F_{\text{ext1}}(t) - c_2 \dot{x}_{m_2}(t) - k_2 x_{m_2}(t)), \quad (3.2)$$

$$\ddot{x}_{m_3}(t) = \frac{1}{m_3}(F_{\text{ext2}}(t) - c_3 \dot{x}_{m_3}(t) - k_3 x_{m_3}(t)). \quad (3.3)$$

For clarity, the equations of motion are rewritten as the state-space equation. Solving this state-space equations with forward Euler generates the matrix for the states, $\mathbf{x}(t)$, and derivatives of the states, $\dot{\mathbf{x}}(t)$ for the duration of the simulation, T^1 . The state-space equations of the plant, indicated by subscript p , are as follows,

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}_p \mathbf{u}_p(t) + \boldsymbol{\omega}(t) \quad (3.4)$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \mathbf{z}(t). \quad (3.5)$$

The states of the three masses (m_1 , m_2 , and m_3) are defined by,

$$\mathbf{x}(t) = \begin{bmatrix} x_{m_1}(t) \\ \dot{x}_{m_1}(t) \\ x_{m_2}(t) \\ \dot{x}_{m_2}(t) \\ x_{m_3}(t) \\ \dot{x}_{m_3}(t) \end{bmatrix} = \begin{bmatrix} x_1(t) \\ x_2(t) \\ x_3(t) \\ x_4(t) \\ x_5(t) \\ x_6(t) \end{bmatrix}. \quad (3.6)$$

The input of the plant is,

$$\mathbf{u}_p(t) = \begin{bmatrix} 0 \\ F_{\text{int}}(t) \\ 0 \\ F_{\text{ext1}}(t) \\ 0 \\ F_{\text{ext2}}(t) \end{bmatrix}. \quad (3.7)$$

¹As the simulation is implemented in Python, the continuous state-space equations for the plant must be discretized to generate the states of the masses. Appendix B details how to find the discrete-time states.

The state matrix, \mathbf{A} , and output matrix, \mathbf{B} , are defined by,

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ -\frac{k_1}{m_1} & -\frac{c_1}{m_1} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -\frac{k_2}{m_2} & -\frac{c_2}{m_2} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & -\frac{k_3}{m_3} & -\frac{c_3}{m_3} \end{bmatrix} \quad (3.8)$$

$$\mathbf{B}_p = \begin{bmatrix} 0 \\ \frac{1}{m_1} \\ 0 \\ \frac{1}{m_2} \\ 0 \\ \frac{1}{m_3} \end{bmatrix}. \quad (3.9)$$

The agent only observes the position of each mass and then infers the motion states of the plant using the agent's own state-space equation (see eq. 3.11). Therefore, the output matrix is defined by

$$\mathbf{C} = [1 \ 0 \ 1 \ 0 \ 1 \ 0]. \quad (3.10)$$

The vectors $\boldsymbol{\omega}(t)$ and $\mathbf{z}(t)$ are the process and measurement noise of the system, respectively.²

3.1.3. THE AGENT'S MODEL: AGENCY OVER THE OBSERVED STATES

The second part of the mathematical model of the point mass mRHI is the agent's incomplete state-space model of the environment. The agent's state-space equation is

$$\dot{\hat{\mathbf{x}}}(t) = \mathbf{A}\hat{\mathbf{x}}(t) + \hat{\mathbf{B}}u_a(t) \quad (3.11)$$

$$\hat{\mathbf{y}}(t) = \mathbf{C}\hat{\mathbf{x}}(t). \quad (3.12)$$

THE AGENCY MATRIX AND SENSE OF AGENCY

In this state-space model, the agent determines its agency by estimating matrix $\hat{\mathbf{B}}$. The matrix $\hat{\mathbf{B}}$ is the *agency matrix* because this matrix describes the causal relationship between the agent's input, u_a , and the predicted states, $\hat{\mathbf{x}}$, and their derivatives, $\dot{\hat{\mathbf{x}}}$. Matrix $\hat{\mathbf{B}}$ is unknown to the agent at the start of simulation and is estimated by the agent with a parameter estimator (OLS in Chapter 4 and EM in Chapter 5). The hat on the variables in eq. 3.11 and 3.12 represents the agent's belief about that variable. Thus, matrix $\hat{\mathbf{B}}$ is the belief the agent has about its true agency, \mathbf{B} .

The agent's SoA is not described by a parameter of the agent's state-space equation because it is a subjective feeling that scales between the feeling of no-agency and agency.

²Instead of describing these noises with white noise, which is commonly assumed, $\boldsymbol{\omega}(t)$ and $\mathbf{z}(t)$ are calculated using normalized zero-mean Gaussian. This definition of noise is based on Active Inference where noises are assumed to be correlated at different time steps [86]. This correlation is a result of the dynamic processes in real-world systems generating the noise. Different from white noise, which is uncorrelated over time, correlated noise can be calculated by any differentiable (auto)correlation function.

The quantification of SoA depends on the parameter estimator. In short, with the OLS algorithm (Chapter 4), the agent's SoA is binary. As the agent knows that it can only have agency over one mass, the agent's SoA chooses which of the elements ($j = 1, \dots, 6$) of the agency matrix, $\hat{\mathbf{B}}$, is non-zero. The agent gains SoA over the observed state that has the lowest prediction error (a computational principles of SoA, Sec. 2.2.2). The EM algorithm (Chapter 5) uses a similar process to determine SoA but experiences SoA on a continuous scale. Both chapters will elaborate on the quantification of SoA.

If the agent correctly has a SoA over mass m_1 and correctly estimates this agency parameter, \hat{b}_j , then the estimated agency matrix,

$$\hat{\mathbf{B}} = \begin{bmatrix} \hat{b}_1 \\ \hat{b}_2 \\ \hat{b}_3 \\ \hat{b}_4 \\ \hat{b}_5 \\ \hat{b}_6 \end{bmatrix}, \quad (3.13)$$

should converge to the true agency matrix,

$$\mathbf{B} = \begin{bmatrix} 0 \\ \frac{1}{m_1} \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (3.14)$$

DIFFERENCES BETWEEN THE AGENT AND THE PLANT

Aside from the unknown agency matrix $\hat{\mathbf{B}}$, there are two more differences between the agent's (eq. 3.11 and 3.12) and plant's (eq. 3.4 and 3.5) state-space equations. First, the only input to the agent's model of the system is the force F_{int} because this is the only force that the agent itself applies,

$$u_a(t) = F_{\text{int}}(t). \quad (3.15)$$

Second, for simplicity, the noise is omitted meaning the agent does not model the process noise in the point mass mRHI. As a result of the unmodeled noise and estimated agency matrix $\hat{\mathbf{B}}$, the states are also predictions made by the agent, or,

$$\hat{\mathbf{x}} = \begin{bmatrix} \hat{x}_{m_1}(t) \\ \hat{\dot{x}}_{m_1}(t) \\ \hat{x}_{m_2}(t) \\ \hat{\dot{x}}_{m_2}(t) \\ \hat{x}_{m_3}(t) \\ \hat{\dot{x}}_{m_3}(t) \end{bmatrix} = \begin{bmatrix} \hat{x}_1(t) \\ \hat{x}_2(t) \\ \hat{x}_3(t) \\ \hat{x}_4(t) \\ \hat{x}_5(t) \\ \hat{x}_6(t) \end{bmatrix}. \quad (3.16)$$

3.1.4. UNITING THE PLANT AND THE AGENT'S MODEL

The agent gains SoA over the mass whose sensory consequence, y_j , is closest to the agent's predicted sensory consequence, \hat{y}_j . This means that the agent uses its perceived prediction error, ε_y , to update its agency matrix $\hat{\mathbf{B}}$. However, in the point mass mRHI, the agent uses the prediction error of the state derivatives, $\varepsilon_{\dot{x}}$, rather than the sensory observations, ε_y , because the agent can more directly observe the effects of its input force, F_{int} .

Figure 3.2 illustrates both state-spaces and their connection to each other through the prediction error. The choice of the non-zero element j of agency matrix, $\hat{\mathbf{B}}$, is based on the lowest j error of matrix $\varepsilon_{\dot{x}}$. The process of calculating the updated agency matrix, $\hat{\mathbf{B}}$, depends on the choice of parameter estimator. Chapter 4 and 5 detail how the prior belief is imposed on the matrix.

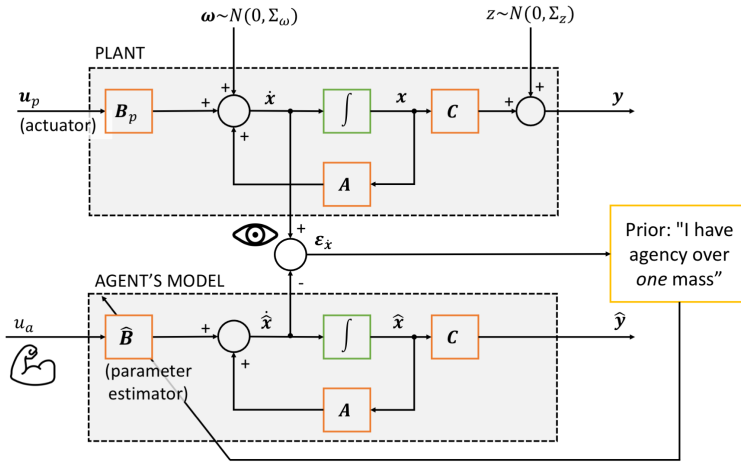


Figure 3.2. This block diagram shows that the plant and the agent's model are connected to determine SoA. The plant produces the agent's sensory evidence by defining the states and output of the three masses, x and y , with the state-space equation (eq. 3.4 and eq. 3.5). The agent estimates the agency matrix, $\hat{\mathbf{B}}$, to predict the motions states \hat{x} with the agent's state-space equation (eq. 3.11). As a result of the agent's prior knowledge (i.e. that it has agency over exactly *one* mass), the agent gains a SoA over the mass with the lowest prediction error, $\varepsilon_{\dot{x}}$.

3.2. POINT MASS mRHI AS AN ANALOGY OF HUMAN mRHI

USING the mathematical model from the previous section, this section discusses how the point mass mRHI is an analogy of the human mRHI experiment. The analogy is compared in two ways. First, I compare the experimental set-ups. Second, I analyze how the agent's model in the point mass mRHI reflects the three underlying computational principles of SoA.

3.2.1. COMPARISON OF THE EXPERIMENTAL SET-UP

The setup of the mRHI experiment by Kalckert and Ehrsson [32] is similar to the point mass mRHI. Figure 3.3 compares the two experimental setups. Mass m_1 represents the participant's real hand because, like the real hand, the agent applies a force directly to this mass. The other two masses, m_2 and m_3 represent the rubber hand. This highlights one obvious difference between the two experimental setups; the point mass mRHI has two rubber hands. The extra rubber hand allows testing of more experimental conditions for the mRHI.

Another similarity is that between the experimental set-ups is that the participants from the human mRHI and the agent in the point mass mRHI both know their force input F_{int} . The human mRHI experiment shows it is important that the agent knows its force input, F_{int} because the participants only have a SoA in the active condition (i.e. when the participant moves their hand) [31][32]. In the passive condition (i.e. the experimenter moves the hand), the participants do not know the action that causes the rubber hand and their own hand to move.

3.2.2. COMPUTATIONAL PRINCIPLES OF SoA IN THE POINT MASS mRHI

To support the point mass mRHI as a simplification of the human mRHI experiment (and SoA), I will identify the three computational mechanisms of SoA in the point mass mRHI.

THE PRECISION OF SENSORY SIGNALS

In the human mRHI, reducing the precision of the sensory signals is the first step towards the illusory SoA over the rubber-hand. The participant's sensory precision of the estimation of their hand position decreases by placing the hand under the table. This highlights a less apparent, but crucial, difference between the mRHI and the simulation. The participant (unconsciously) estimates the states of their invisible hand, \mathbf{x}_{hand} , with their proprioceptive signals, while the agent uses its "vision" to detect the states of m_1 , x_{m_1} . The visual system has greater precision than the proprioceptive sense in estimating hand position [87]. To accommodate this increased uncertainty in the states of the hand in the human mRHI, I increase the noise on the states of mass m_1 (i.e. ω). Increasing the noise on the states of the agent's "real hand" is the first step towards the agent attributing an incorrect SoA to mass m_1 .

LOW PREDICTION ERROR

In the mRHI experiments, SoA over the rubber hand only occurs when there is temporal congruence between the real and rubber hand because of the low prediction error between the predicted sensory outcomes of the real hand and observed outcomes of the

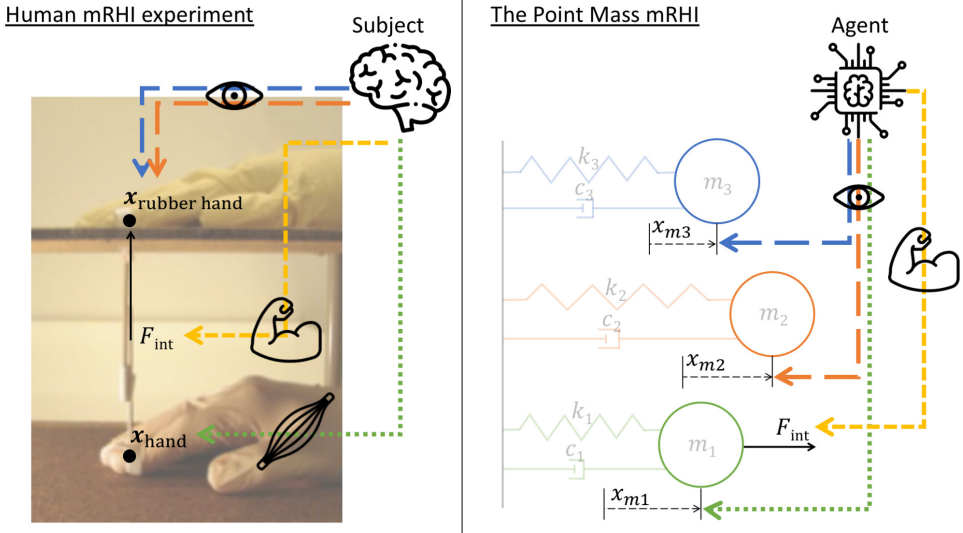


Figure 3.3. This figure compares the point mass mRHI to the human mRHI experiment (image from [32]). The first mass, m_1 represents the “real hand”, while masses m_2 and m_3 are the “rubber hands”. In the human mRHI experiment (left), the participant visually observes the states of the rubber hand $x_{\text{rubber hand}}$, uses proprioception to estimate the state of their hand x_{hand} , and knows the force input F_{int} applied to their hand. In the point mass mRHI (right), the agent observes the states of the three masses x_{m1} , x_{m2} , and x_{m3} and knows the force input F_{int} .

rubber-hand. In the point mass mRHI, the agent only gains SoA over the mass with the lowest prediction error, $\epsilon_{\dot{x}}$.

The process of how the agent uses the prediction error to gain a SoA is based on two sub-process that influence SoA in humans (e.g. [21]). In the predictive process, the agent predicts the motion states \dot{x} with its knowledge of its model (see Figure 3.2). The agent uses its SoA from a previous time step to make these predictions. Second, the postdictive process uses the prediction error, $\epsilon_{\dot{x}}$, to inform the prior belief which mass has the lowest prediction error. This is the mass that the agent actually should have SoA over and the agent keeps the corresponding element of the \hat{B} matrix as non-zero. With the updated \hat{B} , the agent can again perform the predictive process. By distinguishing between sensations that it can and cannot control, the agent lowers its prediction error.

STRONG PRIOR BELIEF THAT INTENDED ACTION CAUSES OUTCOME

In the human mRHI, the participants have strong prior beliefs that the actions of their hand cause sensory consequences [88]. These causal prior beliefs are important for inducing an illusory SoA over the rubber hand. As causality is encoded by the agency matrix, \hat{B} , the prior belief is imposed on this matrix.³ The prior belief defines that there is causality for only one mass. Though this prior knowledge does not capture the complexity of human prior knowledge, this prior does capture that participants of the mRHI have the prior belief that there is a causal relationship only for their own hand.

³How the prior belief is imposed on the matrix is dependent on the algorithms used in Chapter 4 and 5.

Two differences should be noted between the agent's and the participant's prior. First, different to humans, the agent does not know how F_{int} relates to the states of the masses before the mRHI experiment starts (i.e. does not know $\hat{\mathbf{B}}$). I assume this difference has a negligible effect because of the simplicity of the environment that the agent tries to estimate. Second, the agent's prior assumes that the agent will only feel agency over one mass at a time. In the human mRHI experiments, Kalckert and Ehrsson[31][32] did not research whether humans maintained agency over their own hand while also having a SoA over the rubber hand. Assuming that the SoA over the real hand decreases as a result of the decreased sensory precision [20][72], placing the hand under the table will likely result in a lower SoA over the real hand. Therefore, this assumption is acceptable for this simplified model of the mRHI experiment.

3.3. EXPERIMENTS

To generate the sensory evidence for which the agent either experiences an improved, correct or illusory, incorrect SoA, this section defines the parameters of the plant's mathematical model (see Sec. 3.1). This section also defines the performance measures that I will use to compare an agent's correct or incorrect SoA with and without a prior.

3.3.1. PARAMETERS OF THE POINT MASS MRHI

The generation of the states of the three masses requires a definition of the parameters and the input signal. The states generated by the plant are the environment that the agent "senses" and uses to estimate $\hat{\mathbf{B}}$ in Chapter 4 and 5. The agent will be subject to two sets of environments of which the agent will perform 1000 trials each. In the first set of *unambiguous* environments (experiment 1), the agent should gain a correct SoA. In the second set of *ambiguous* environments (experiment 2), the agent should experience the mRHI, which is an incorrect SoA.

CONSTANT PARAMETERS: TIME AND NOISE

Before I elaborate on the difference between the parameters of experiments 1 and 2, I define the time and noise parameters that remain constant in both experiments (see Table 3.1). The simulation time is chosen such that it is long enough to determine whether the agent does not adjust its SoA. The time window is the number of passed steps that the agent stores in its "memory" to avoid a continuous increase in the size of the data set. The agent should have just enough data to infer its SoA over a mass. Therefore, the time window is a tenth of the entire simulation time. The covariance of the states and observations, Σ_ω and Σ_z , and the roughness parameter γ generate the smooth Gaussian noise, ω and z . These values are set to introduce some noise in the observations of the states of the masses, without the states losing their predictability. The covariance matrix of the state noise, Σ_ω , is increased in some experimental conditions of experiment 1 and 2.

Table 3.1. Overview of the time and noise parameters of the point mass mRHI.

Parameter name	Parameter sign	Parameter value
Simulation time	T	20 s
Sampling period	h	0.01 s
Time window	δ_N	200
Covariance of the observation noise	Σ_z	[0.01]
Covariance of the state noise	Σ_ω	$\begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}$
Roughness parameter	γ	16

EXPERIMENT 1: UNAMBIGUOUS ENVIRONMENTS FOR A CORRECT SENSE OF AGENCY

In experiment 1, the agent should gain a correct SoA for the three experimental conditions. Table 3.2 defines the parameters for each experimental condition. Since the point mass mRHI is a simplification of the human mRHI, the goal is not to emulate the exact experimental conditions of the mRHI. The agent should experience a correct SoA for

all parameters. This requires that the agent can correctly predict the states of mass m_1 . Therefore, the parameters of experiment 1 must fulfill three conditions. First, the movement of the other masses, m_2 and m_3 , must be sufficiently asynchronous with mass m_1 . Second, the magnitude of the noise should be small relative to the magnitude of the states. Third, the combination of parameters should avoid resonance in the three masses. I handpicked parameter values that fulfill these conditions, meaning the agent should find a correct SoA in all experiments.

The parameters (m_s, k_s, c_s) are inspired by the properties of the human arm. The mass of each s point masses, m_s , varies between 1.5 kg and 2.5 kg because this is a plausible range for the hand and lower arm [89]. The combination of the spring and damping constant and the input force are chosen such that the resulting movement of each mass is plausible for the human body. This means that, for example, the amplitude of the resulting position of the mass is no further away than around 0.5 m from the equilibrium position to represent a plausible maximum distance between the elbow and the hand. For SoA the agent must have actions to distinguish from others. The agent applies a large enough sinusoidal force on mass m_1 , $F_{\text{int}} = 2 \cos(4t_p) + 3$, to ensure that the mass remains excited throughout the simulation.

The difference between the three experimental conditions of the unambiguous environment is the input forces, u_p , and the state noise, ω .

- **Experiment 1A: a static environment.** The agent observes a static environment because mass m_1 is the only mass on which a force is applied. The static environment is the simplest environment where an agent should have a SoA.
- **Experiment 1B: an asynchronous dynamic environment.** The environment of the agent is dynamic, meaning the agent now has to distinguish its own force from the sinusoidal forcing functions applied to both mass m_2 and m_3 . The only difference between mass m_2 mass m_1 is the 750 ms shift in the input force.
- **Experiment 1C: an asynchronous dynamic environment with noise.** This experimental condition is the first step towards inducing the mRHI by increasing the noise of the acceleration state of mass m_1 to resemble the participant's hidden hand. The noise on the acceleration state increases in noise because the agent only controls this state with its input force. The large time lag of the force input of mass m_2 (750 ms) means the agent should not experience the mRHI.

The covariance of the state noise of mass m_1 in experiment 1C increases to,

$$\Sigma_w = \begin{bmatrix} 0.01 & 0 \\ 0 & 1 \end{bmatrix}. \quad (3.17)$$

EXPERIMENT 2: AMBIGUOUS ENVIRONMENTS FOR AN INCORRECT SENSE OF AGENCY

To generate the ambiguous for experiment 2, there is an increased state noise for mass m_1 (defined by eq. 3.17) and (near) temporal congruence between the observations of mass m_1 and m_2 . These experimental conditions will reproduce the experimental conditions from the human mRHI experiment (i.e. [31] and [32]).

Table 3.2. Overview of the parameters per experimental condition (1A, 1B, and 1C) of the unambiguous environment, meaning the agent should have a correct SoA mass m_1 .

Mass number	Mass, kg	Spring constant, Nm ⁻¹	Damping constant, Ns m ⁻¹	Input force, N	Covariance of the state noise Σ_w
Experiment 1A: static environment					
m_1	2.5	6	2	$2 \cos(4t_p) + 3$	diag(0.01, 0.01)
m_2	2.5	6	2	0	diag(0.1, 0.1)
m_3	2	4	4	0	diag(0.1, 0.1)
Experiment 1B: an asynchronous dynamic environment					
m_1	2.5	6	2	$2 \cos(4t_p) + 3$	diag(0.01, 0.01)
m_2	2.5	6	2	$2 \cos(4t_p + 3) + 3$	diag(0.01, 0.01)
m_3	2	4	4	$\sin(2t_p)$	diag(0.01, 0.01)
Experiment 1C: an asynchronous dynamic environment with noise					
m_1	2.5	6	2	$2 \cos(4t_p) + 3$	diag(0.01, 1)
m_2	2.5	6	2	$2 \cos(4t_p + 3) + 3$	diag(0.01, 0.01)
m_3	2	4	4	$\sin(2t_p)$	diag(0.01, 0.01)

There are three ambiguous experimental conditions in which the agent should incorrectly conclude that it has agency. Table 3.3 defines the parameters of each experimental condition. The experimental conditions are as follows:

- **Experiment 2A: the active synchronous mRHI.** This experimental condition reproduces the “active synchronous” condition from mRHI experiment [32]. The plant generated states of mass m_1 and m_2 are the same because the parameters and the input force are equal. The difference is the increased state noise on mass m_1 .
- **Experiment 2B: the active synchronous mRHI with a different mass.** The second experiment condition is identical to the first except for a 10% in the mass of mass m_2 . This increases the amplitude of the position for m_2 , which has not been tested in mRHI experiments (such as [31], [32], and [90]).
- **Experiment 2C: the active asynchronous mRHI.** This experimental condition reproduces the “active asynchronous” condition from mRHI experiment [32]. Thus, the parameters of mass m_1 and m_2 are identical and the sinusoidal force on mass m_1 is shifted by 500 ms. This should decrease the (incorrect) SoA over the mass m_2 because a lag of 500 ms was sufficient to significantly reduce the mRHI in the human mRHI experiment [31][32]. However, the agent in this conditions should still have more incorrect SoA over mass m_2 than in experiment 1.

3.3.2. PERFORMANCE MEASURES

In all experiments, the goal is to quantify whether the prior improves and/or promotes illusory SoA. To compare the results of an agent with prior to a prior-less agent in the point mass mRHI (in Chapter 4 and 5), I define five performance measures.

Table 3.3. Overview of the parameters per experimental condition (2A, 2B, and 2C) of the ambiguous environment, meaning the agent should have a incorrect SoA over mass m_2 .

Mass number	Mass, kg	Spring constant, Nm ⁻¹	Damping constant, Ns m ⁻¹	Input force, N	Covariance of the state noise Σ_w
Experiment 2A: active synchronous mRHI					
m_1	2.5	6	2	$2 \cos(4t_p) + 3$	diag(0.01, 1)
m_2	2.5	6	2	$2 \cos(4t_p) + 3$	diag(0.01, 0.01)
m_3	2	4	4	$\sin(2t_p)$	diag(0.01, 0.01)
Experiment 2B: the active synchronous mRHI with a different mass					
m_1	2.5	6	2	$2 \cos(4t_p) + 3$	diag(0.01, 1)
m_2	2.25	6	2	$2 \cos(4t_p) + 3$	diag(0.01, 0.01)
m_3	2	4	4	$\sin(2t_p)$	diag(0.01, 0.01)
Experiment 2C: the active asynchronous mRHI					
m_1	2.5	6	2	$2 \cos(4t_p) + 3$	diag(0.01, 1)
m_2	2.25	6	2	$2 \cos(4t_p + 2) + 3$	diag(0.01, 0.01)
m_3	2	4	4	$\sin(2t_p)$	diag(0.01, 0.01)

The first measure quantifies the agent's SoA throughout all 1000 trials. As SoA is the degree that the agent believes to have agency over a mass (ranging from no agency to complete agency), the agent's SoA will be defined by the total amount of time that the agent has a SoA over each mass. The time of SoA is calculated by multiplying the agent's SoA over one of the observed states by the sampling period h .

As the agent could have a correct SoA, while incorrectly predicting the states or incorrectly estimating the agency matrix, $\hat{\mathbf{B}}$, I define four other performance measures. The second and third measure calculate whether the agent's estimation of the agency matrix, $\hat{\mathbf{B}}$, is accurate. Despite the agent's have a correct SoA over mass m_1 , the agent's estimated agency may still be inaccurate. A first measure that determines whether the agent's SoA is based on an accurate belief of agency, $\hat{\mathbf{B}}$, and is calculated by finding the difference between the mean of each state j of the estimated agency matrix, $\hat{\mathbf{B}}$, or $\mu_{\hat{b}_j}$, to the true agency matrix, \mathbf{B} (from eq. 3.14). This finds the error in the SoA estimation after all time steps, N , with,

$$\sum_j^k |\mu_{\hat{b}_j} - b_j|, \quad (\text{agency accuracy}). \quad (3.18)$$

The second measure focuses on whether the agent's SoA is based on an accurate belief of agency only for mass m_1 . I calculate the agent's prediction error for the motion states of mass m_1 . An agent with an accurate agency matrix, $\hat{\mathbf{B}}$, should have a low prediction error for mass m_1 . This measure calculates the mean error over all time steps between the predicted motion states, \hat{x}_1 and \hat{x}_2 , and the observed motion states, \hat{x}_1 and \hat{x}_2 , or

$$\sum_j^2 \mu_{\varepsilon_{\hat{x}_j}}, \quad (\text{state accuracy}). \quad (3.19)$$

The fourth measure calculates the variance of the agent's agency matrix, $\hat{\mathbf{B}}$ because

this reflects how the agent perceives the precision of the sensory signals. A lower sensory precision, or a higher variance in the agency matrix, $\hat{\mathbf{B}}$, should result in lower SoA. I calculate the agent's perceived sensory precision this is as follows,

$$\sum_j^k \sigma_{\hat{b}_j}^2, \quad (\text{perceived sensory precision}). \quad (3.20)$$

The fifth measure investigates whether an agent with prior performs prediction error minimization when compared to an agent without a prior. An agent only accumulates a prediction error over objects over which the agent has a SoA [72]. Therefore, the agent's accumulated error only adds the prediction error, $\varepsilon_{\hat{x}_j}$, of the states over which the agent believes it has a SoA, a .

$$\mu_{\varepsilon_{\hat{x}_a}}, \quad (\text{perceived error}). \quad (3.21)$$

I use these performance measures in the results of Chapter 4 and 5.

3.4. CONCLUSION

In conclusion, this chapter answers the research question: *In what artificial environment, similar to the moving rubber hand illusion, should an agent experience a correct and incorrect sense of agency?* In short, I mathematically define an analogy of the human mRHI experiment, termed “the point mass mRHI”. When I increase the noise on the states of mass m_1 and the three masses in the point mass mRHI move (near) synchronously, then the agent should experience an incorrect, illusory SoA. In all other cases, the agent should experience a correct SoA.

To elaborate, the first part of this chapter (Sec 3.1) establishes the mathematical model of the point mass mRHI. In this model, an artificial agent observes three masses connected to spring-damper systems (see Figure 3.1). The agent has an incomplete model of the environment because, though it knows its input force, the agent does not know on which mass it applies this force. Therefore, the agent estimates the matrix $\hat{\mathbf{B}}$ (see eq. 3.13), which describes the agent’s agency over each mass. The agent’s prior knowledge updates agency matrix $\hat{\mathbf{B}}$ because the agent knows it can only control one mass, i.e. only one element of $\hat{\mathbf{B}}$ is non-zero. The quantification of the SoA is dependent on what element of $\hat{\mathbf{B}}$ is non-zero and will further be specified later in this thesis. Next, Sec 3.2 shows that the point mass mRHI is a simplification of the human mRHI experiment. One mass represents the participant’s hand, while the two other masses represent the rubber-hand (see Figure 3.3). Like in human SoA, the agent’s SoA is dependent on the prediction error described by $\epsilon_{\hat{x}}$, causality described by $\hat{\mathbf{B}}$, and precision of the sensory data defined by ω . Sec. 3.3 defines the parameters for which the agent should experience a correct (experiment 1) and incorrect (experiment 2) SoA based on the experimental conditions of the human mRHI experiment in [32].

Chapter 4 and 5 implement “point mass mRHI” by estimating the agency matrix, $\hat{\mathbf{B}}$, with OLS and EM, respectively. The quantification of the agent’s SoA is dependent on the parameter estimator.

4

THE SIMPLEST POINT MASS MRHI WITH ORDINARY LEAST SQUARES

This chapter is the first step to answering the final research questions: *What is the effect of the prior on a sense of agency in the point mass moving rubber-hand illusion? Is this effect comparable to a sense of agency in human?* In this point mass mRHI, the agent uses the parameter estimator *ordinary least squares* (OLS) to estimate the agency matrix, $\hat{\mathbf{B}}$, which describes the relationship between its force inputs and the predicted sensory consequences. The goal of this chapter is to provide an elementary explanation of the effect of prior beliefs on a SoA in an artificial agent with this straightforward parameter estimator. Also, to understand how to improve the point mass mRHI so that the model can become a more correct model of SoA.

First, Sec. 4.1, provides a theoretical explanation of OLS as a parameter estimator for the agency matrix, $\hat{\mathbf{B}}$. Sec. 4.2 focuses on the implementation of the point mass mRHI with OLS. I explain how the agent uses prior knowledge to gain SoA. Mathematically, this means its choice of the non-zero element of agency matrix $\hat{\mathbf{B}}$. Then I explain the complete algorithm (see Figure 4.1) that includes the prior knowledge and OLS to gain SoA. Next, Sec. 4.3 presents the results the algorithm for the experimental conditions defined by the experimental design (Sec. 3.4). The results show that with the prior belief the agent has an improved SoA in experiment 1 and an incorrect SoA in experiment 2. In Sec. 4.4, I discuss whether the results reflect the three underlying computational principles of SoA (causal prior belief, prediction error, sensory signals). Chapter 5 uses these improvements as a starting point for the EM algorithm.

4.1. THE ORDINARY LEAST-SQUARES ALGORITHM FOR AGENCY

In the point mass mRHI, the agent's goal is to estimate its agency matrix $\hat{\mathbf{B}}$ to gain a SoA over one of the masses. With OLS, the agent estimates the unknown agency matrix $\hat{\mathbf{B}}$ before applying the prior by minimizing the differences between the observations and predictions.

The agent's state-space model (defined in eq. 3.11) describes the states of the masses for one time step. To use OLS, I extend the state-space model to describe the agent's model for the N passed time steps. However, the agent will not store the matrices over *all* passed time steps but rather will forget time steps that occurred too long ago to avoid that the amount of stored data approaches infinity with time. The number of steps over which the agent calculates the $\hat{\mathbf{B}}$ matrix is termed the "time window", or δ_N . Therefore, the equation for the motion states is given by,

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \hat{\mathbf{B}}\mathbf{u}_a, \quad (4.1)$$

where, the matrices of the state-space that change size with each time step i ($i = 0, \dots, N-1$),

$$\begin{aligned} \mathbf{x} &= [\mathbf{x}_p \quad \mathbf{x}_{p+1} \quad \dots \quad \mathbf{x}_{N-1}] = (x_{ij}) \in \mathbb{R}^{n \times (N-p)}, \\ \mathbf{u}_a &= [(u_a)_p \quad (u_a)_{p+1} \quad \dots \quad (u_a)_{N-1}] = (u_a) \in \mathbb{R}^{1 \times (N-p)}. \end{aligned}$$

The variable p describes the last time step that the agent remembers with,

$$p = \begin{cases} 0 & \text{if } (N-1) \leq \delta_N \\ (N-1) - \delta_N & \text{if } (N-1) > \delta_N \end{cases}$$

In the general OLS in matrix form, the variable, \mathbf{Y} , is a vector of the past $N-p$ ($i = p, \dots, N-1$) time-steps, that depends on N observations, \mathbf{X} , and j ($j = 0, \dots, n-1$) unknown parameters, $\boldsymbol{\beta}$. Mathematically [91],

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (4.2)$$

where $\boldsymbol{\varepsilon}$ is a vector of the $N-p$ unobserved disturbances.

Different to the general OLS matrix form, the agent in the point mass mRHI observes $n = 6$ states at each time step. Nonetheless, the variables from the agent's state-space can be equated to the variables from the OLS matrix form as follows,

$$\mathbf{Y} = \begin{bmatrix} [\dot{\mathbf{x}}_p - \mathbf{A}\mathbf{x}_p]^T \\ [\dot{\mathbf{x}}_{p+1} - \mathbf{A}\mathbf{x}_{p+1}]^T \\ \vdots \\ [\dot{\mathbf{x}}_{N-1} - \mathbf{A}\mathbf{x}_{N-1}]^T \end{bmatrix} \in \mathbb{R}^{(N-p) \times n}, \quad (4.3)$$

$$\mathbf{X} = \mathbf{u}_a^T \in \mathbb{R}^{(N-p) \times 1}, \quad (4.4)$$

$$\boldsymbol{\beta} = \hat{\mathbf{B}}^T \in \mathbb{R}^{1 \times n}. \quad (4.5)$$

To estimate the unknown parameter, $\boldsymbol{\beta}$, at each time step (incrementally) over the passed time steps within time window δ_N , OLS solves for the unknown parameter in

eq. 4.2 by minimizing the sum of squared errors, ϵ . The formula for β calculates the psuedo-inverse of \mathbf{X} ,

$$\beta = (\mathbf{X}' \cdot \mathbf{X})^{-1} \mathbf{X}' \mathbf{Y}. \quad (4.6)$$

By filling in eq. 4.3-4.5 in eq. 4.6, the agent solves for β which is the estimation of the agent's agency, $\hat{\mathbf{B}}$, *before* applying the prior.

4.2. IMPLEMENTATION OF THE PRIOR IN OLS

WITH the OLS estimated agency matrix $\hat{\mathbf{B}}$, the agent can apply its prior knowledge to gain SoA over one mass and then update $\hat{\mathbf{B}}$. First, this section explains how the agent uses its prior knowledge and SoA to update the agency matrix $\hat{\mathbf{B}}$. Second, the overall algorithm is presented.

4.2.1. SENSE OF AGENCY IS DETERMINED WITH P

The prior knowledge in the OLS algorithm is not a Bayesian prior because the agent does not combine the agent's prior knowledge and the sensory evidence in a Bayesian way. However, "the prior" in the OLS algorithm does communicate to the agent that only one of the elements of the agency matrix, $\hat{\mathbf{B}}$, can be non-zero. The agent cannot directly recognize this information from the collected sensory evidence. Therefore, for this simplest point mass mRHI, this non-Bayesian prior is sufficient.

The agent's SoA with the OLS algorithm is binary, meaning the agent either does or does not have SoA over an observed state. The SoA arises based on which observed states best fulfill two of the computational principles of SoA: a causal prior belief of action-outcome and low prediction error. The principle of the precision of sensory signals does not influence the choice of SoA because OLS does not describe the distribution of data. Combining the agent's prior belief with the principles means the agent's SoA over one mass through the j th unknown parameter is defined by,

$$\text{SoA} = \begin{cases} \text{yes} & \text{if } \hat{b}_j \neq 0 \text{ and } \operatorname{argmin}_{j \in k} \varepsilon_{\dot{x}_j} \\ \text{no} & \text{otherwise} \end{cases} \quad (4.7)$$

The first requirement, $\hat{b}_j \neq 0$, is a simplification of the principle that humans have strong prior beliefs about the causal relationship between actions and outcomes. Since each \hat{b}_j describes the causal relationship between the force input, \mathbf{u}_a , and motion states, $\dot{\mathbf{x}}$, when $\hat{b}_j \approx 0$ then there is no causal relationship. This requirement is a simplification of the computational principle as this simplification only encodes that there is a causal relationship and not *how* this relationship looks or the strength of this relationship. OLS finds $\hat{b}_j \approx 0$ in two cases. First, for the observed velocity states of each mass because the force does not directly control the velocity of a mass. Thus, for all even-numbered $j = (0, 2, 4)$, $\hat{b}_j \approx 0$. Second, the agency parameter is $\hat{b}_j \approx 0$ when a mass is not influenced by an input force. In this case, the agent can accurately predict the states of that mass with $\hat{b}_j \approx 0$.

The second requirement, $\operatorname{argmin}_{j \in k} \varepsilon_{\dot{x}_j}$, describes the principle that humans only experience a SoA when the prediction error is low. This means that for the non-zero \hat{b}_j , the agent only has a SoA over the observed motion state, \dot{x}_j , that is best predicted by $\hat{\dot{x}}_j$. For the purpose of the point mass mRHI, the prediction error of the state derivative, $\dot{\mathbf{x}}$, is assumed to be a function of the current time step, i , and the past time window, δ_N . At time step $i = N$, the prediction error is

$$\varepsilon_{\dot{x}_j} = \sum_{i=p}^{N-1} |\hat{\dot{x}}_{ij} - \dot{x}_{ij}|, \quad (4.8)$$

where \hat{x}_{ij} and x_{ij} are the predicted and observed j th state derivatives at time step i . Calculating the prediction error over the entire time window avoids that the agent's SoA becomes too susceptible to noise.

The agent's binary SoA described in eq. 4.7 updates the $\hat{\mathbf{B}}$ matrix. At each time step, the agent chooses the non-zero \hat{b}_j with,

$$\hat{b}_j = \begin{cases} \hat{b}_j & \text{if } \hat{b}_j \neq 0 \text{ and } \arg \min_{j \in k \in \mathcal{E} \hat{x}_j} \\ 0 & \end{cases} \quad (4.9)$$

where \hat{b}_j is the j th element of $\hat{\mathbf{B}}$.

4.2.2. OLS ALGORITHM

I define an OLS algorithm with and without prior knowledge so that I can compare these algorithms to each other.

WITH PRIOR KNOWLEDGE

The flowchart in Figure 4.1 depicts the algorithm of the agent that determines its SoA with OLS and the prior.¹ The left column describes the qualitative requirements for a SoA over an observed state. The prior stipulates that the agent can control only one mass. According to underlying computational principles of SoA, SoA requires on the existence of causality and a low prediction error. The right column of Figure 4.1 translates the requirements for SoA from the left column and estimates $\hat{\mathbf{B}}$ with OLS. In the algorithm, first the agent observes a prediction error (eq. 4.8) in the motion states, $\hat{\mathbf{x}}$, and predicts $\hat{\mathbf{B}}$ using OLS (eq. 4.6). With this predicted $\hat{\mathbf{B}}$ the agent eliminates the agency parameters with which it will not have agency by finding the $\hat{b}_j \approx 0$. With the remaining non-zero agency parameters \hat{b}_j , the agent determines the j th element for which the prediction error is the minimum. The agency matrix, $\hat{\mathbf{B}}$, is updated such that only this \hat{b}_j remains non-zero. This updated matrix $\hat{\mathbf{B}}$ predicts the motion states at the next time step, $\hat{\mathbf{x}}_{i+1}$.

WITHOUT PRIOR KNOWLEDGE

The flowchart in Figure 4.2 depicts the algorithm where the agent finds its SoA with only OLS. I assume that this agent knows there must be a causal relationship for a SoA because the agent estimates the causal agency matrix, $\hat{\mathbf{B}}$. Therefore, the agent has SoA over all of the agency parameters that OLS estimates as $\hat{b}_j \neq 0$. Unlike the agent with prior knowledge, this agent does not know that it has SoA over the only one mass. As a result, the agent does use the prediction error of the motion states to gain a SoA over one mass.

¹The code of this simplest point mass mRHI with OLS is available on GitHub at https://github.com/mtan1503/point-mass-mrhi/blob/main/ols-algorithm/agent_ols.py.

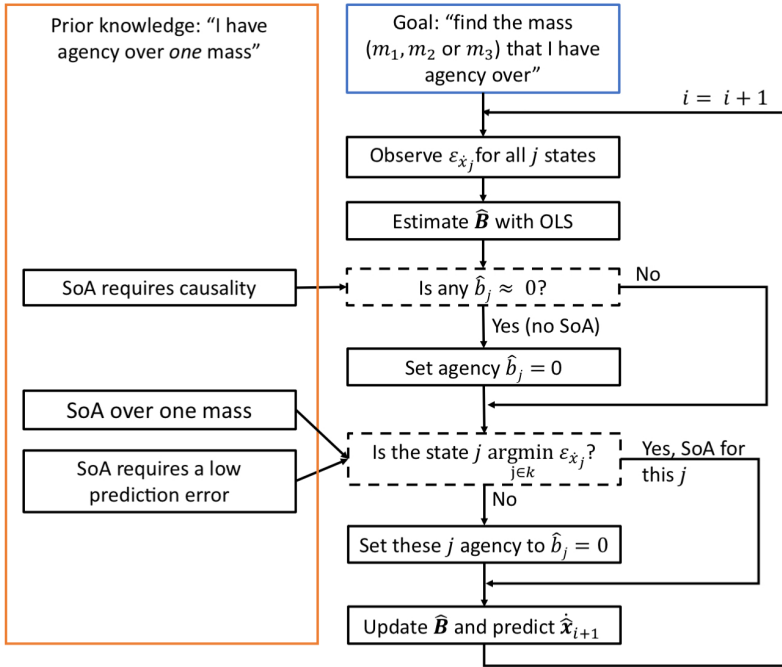


Figure 4.1. This flowchart shows how the point mass mRHI OLS algorithm connects to the computational principles of SoA and the agent's prior. The left column describes the prior knowledge that determines the agent's SoA: the prior knowledge (i) that it only controls one mass, (ii) that the relationship must be causal, and (iii) the prediction error must be low. The latter two requirements are based on the underlying principles of human SoA (see Sec. 2.2.2). The right column describes the algorithm with mathematical translation of the requirements.

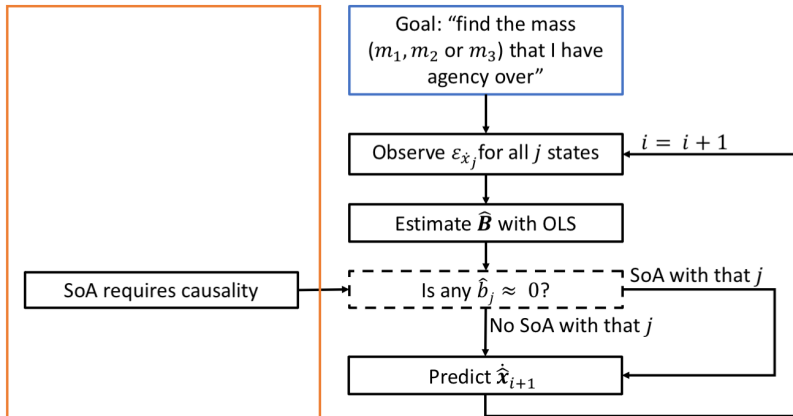


Figure 4.2. This flowchart of an agent in the point mass mRHI that only uses OLS to find its SoA. The left column describes that determine the agent knows the relationship must be causal, and (iii) the prediction error must be low. The right column describes the OLS algorithm.

4.3. RESULTS

THIS section compares the results of an agent using the OLS algorithm with prior knowledge (see Figure 4.1) to the results an agent using an OLS algorithm without prior knowledge (see Figure 4.2). First, I will present the results from the unambiguous environment where the prior should improve the SoA. Second, I will present the results from the ambiguous environment where the prior motivate incorrect SoA.

4.3.1. EXPERIMENT 1: CORRECT SOA IN UNAMBIGUOUS ENVIRONMENTS

For the unambiguous environments (experiment 1), Table 4.1 shows the percentage of time that the agent has a SoA over each state j of the estimated $\hat{\mathbf{B}}$ out of the 1000 20 s trials per experimental condition. The agent with prior correctly has SoA over mass m_1 100% of the time by estimating the agency parameter \hat{b}_1 as non-zero. Conversely, the agent without the prior has SoA over all masses that have a force applied to them. In the static environment (experiment 1A), this means that the agent with and without prior have the same performance. However, in dynamic environments (experiment 1B and 1C) the agent without a prior also has SoA over all odd j parameters, i.e. \hat{b}_1 , \hat{b}_3 , and \hat{b}_5 . Thus, the agent's SoA over the three masses improves with the prior.

Table 4.1. This table shows that for unambiguous environments the agent with a prior has a correct SoA over mass m_1 . The table shows the percentage of time that the agent has a SoA over each observed state per experimental condition (exp. nr) of experiment 1. The agent performs 1000 20 s trials of each experimental condition. The "prior" only allows one non-zero element in the agency matrix, $\hat{\mathbf{B}}$, meaning when $\hat{b}_j > 0$ the agent has a SoA over that observed state for that time step.

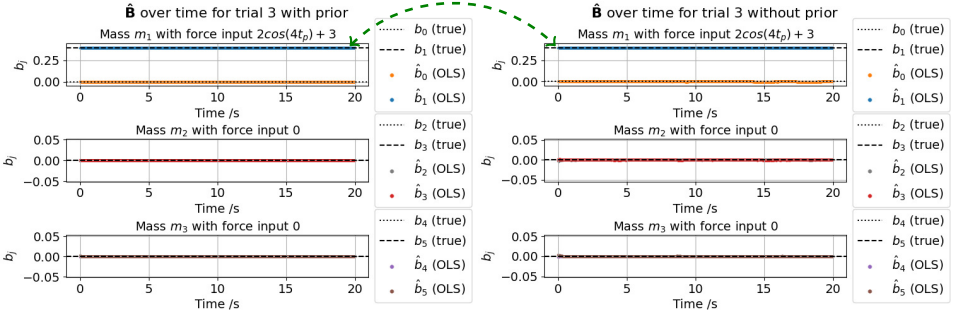
Exp. nr.	% of time that $\hat{b}_0 > 0$	% of time that $\hat{b}_1 > 0$	% of time that $\hat{b}_2 > 0$	% of time that $\hat{b}_3 > 0$	% of time that $\hat{b}_4 > 0$	% of time that $\hat{b}_5 > 0$
<i>With prior</i>						
1A	0.00	100	0.00	0.00	0.00	0.00
1B	0.00	100	0.00	0.00	0.00	0.00
1C	0.00	100	0.00	0.00	0.00	0.00
<i>Without prior</i>						
1A	0.00	100	0.00	0.00	0.00	0.00
1B,1C	0.00	100	0.00	100	0.00	92.8

To visualize the estimation of the agency matrix, $\hat{\mathbf{B}}$, Figures 4.3 to 4.5 compares the OLS estimation of the agency matrix $\hat{\mathbf{B}}$ with and without prior over time for one trial (that is representative of the 1000 trials) of the unambiguous environments of experiment 1.² The only non-zero parameter \hat{b}_j for each experimental condition should be \hat{b}_1 as indicated by the arrow between each set of figures. Also, the dotted lines in each figure represent the true value of the agency matrix \mathbf{B} that the agent tries to estimate (see eq. 3.14).

For the static environment (experiment 1), Figure 4.3a and 4.3b shows that the agent only estimates \hat{b}_1 as non-zero meaning the agent does not need the prior to identify the

²For brevity, I only present the results of the estimation accuracy of the agency matrix $\hat{\mathbf{B}}$ in this section. For the numerical results for the agency accuracy, perceived precision, state accuracy, and perceived error for each environmental condition of experiment 1, see Appendix C.1.

mass that it has agency over because the agent's input force, F_{int} , is the only force acting in the environment. The only (barely observable) difference between the estimated agency matrix $\hat{\mathbf{B}}$ is all the parameters besides \hat{b}_1 are set to 0 in Figure 4.3a, whereas these estimations fluctuate around 0 for the agent without a prior (Figure 4.3b).



(a) OLS estimated \hat{b}_j over time with prior (see Figure 4.1 for (b) OLS estimated \hat{b}_j over time without prior (see Figure 4.2 for algorithm).

Figure 4.3. These plots show that, for the static environment (experiment 1A), the agent's prior knowledge does not significantly affect the estimation of the agency parameter \hat{b}_j because the SoA did not improve (Table 4.1). Both agents estimate the agency matrix, $\hat{\mathbf{b}}_j$, over time window $\delta_N = 200$, with OLS. However, the prior in Figure a outlines that the agent can control *one* mass. Therefore, the only non-zero parameter of the true values b_j , represented by the black dashed lines, is b_1 (see eq. 3.14). Parameter b_1 connects the agent's input force, F_{int} , to the states of mass m_1 .

For the dynamic environment in experiment 1B, the only non-zero \hat{b}_j in Figure 4.4a is \hat{b}_1 meaning the agent with the prior correctly has SoA over mass m_1 . Without the prior in Figure 4.4b, the agent tries to predict the observations of mass m_2 and m_3 with fluctuating agency parameters, \hat{b}_3 and \hat{b}_5 . When there are more forces affecting the states of the environment, the prior helps the agent discriminate between its own actions and external actions. Here, the prior improves the estimation of the agency matrix $\hat{\mathbf{B}}$ and the SoA.

Next, experiment 1C increases the state noise for mass m_1 . Figure 4.5 illustrates that despite the noise the agent with the prior correctly has SoA over mass m_1 from the start of time, while the agent without the prior tries to gain agency over all three masses. Also, the increase in noise barely influences the accuracy of the agency estimation since the estimated \hat{b}_1 remains close to the true value b_1 (the variance in \hat{b}_1 increases from 4.42×10^5 in experiment 1B to 5.31×10^5 in 1C). Figure 4.5b explains the decrease in the SoA in the agent without prior. In the bottom plot, we see the estimated \hat{b}_5 fluctuates around 0. Therefore, some of the time when $\hat{b}_5 \approx 0$ the agent without prior does not have a SoA over this mass.

To illustrate the effect of the prior on the agency parameter with which the agent control mass m_1 , Figure 4.6 plots parameter \hat{b}_1 over time from the dynamic environment (experiment 1B). Though the prior improves the overall estimation of the matrix $\hat{\mathbf{B}}$, the prior does not improve agency parameter \hat{b}_1 . The estimation of parameter \hat{b}_1 with and without prior are equal for all experiments of the point mass MRHI.³ The equal \hat{b}_1 is due

³For numerical results see Appendix C.1, where the accuracy of the states of mass m_1 are equal in the prior

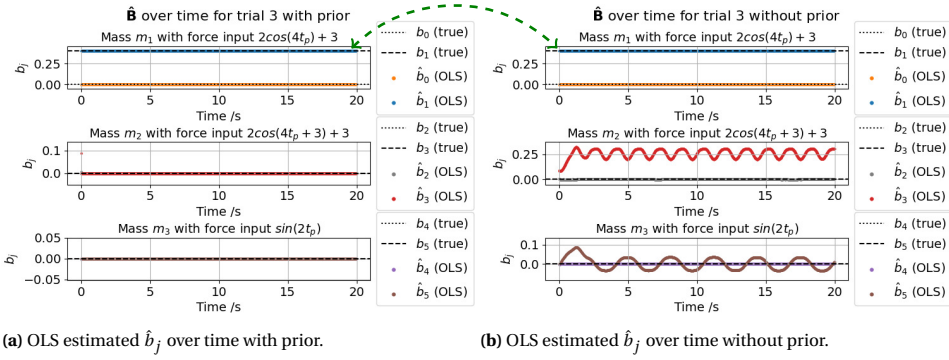


Figure 4.4. These plots show that, for experiment 1B, the agent estimation of the agency matrix $\hat{\mathbf{B}}$ is closer to the true value of the matrix \mathbf{B} with (Figure a) the prior than without the prior (Figure b). Thus, the prior improves SoA and thereby agency. This is a continuation of the plots of experiment 1A in Figure 4.3.

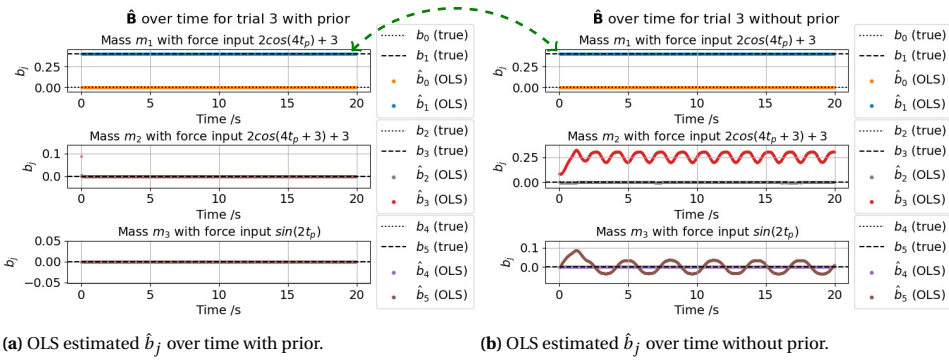


Figure 4.5. These plots show that, for experiment 1C, the agent's prior improves estimation of the agency parameter \hat{b}_j through an improved SoA. This is a continuation of the plots of experiment 1B in Figure 4.4.

to the agent's prior belief only adjusting the agency parameters that the agent does *not* have agency over. Thus, applying this prior does not improve the agent's SoA for the mass which it does have agency over (i.e. mass m_1).

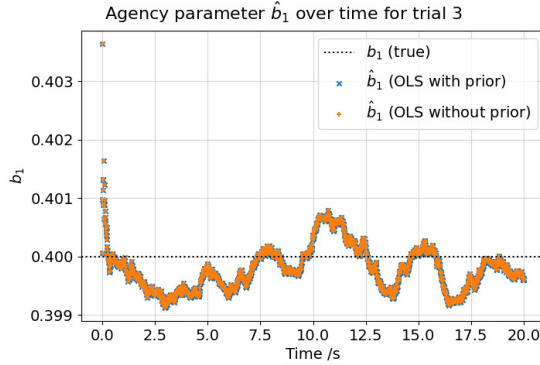


Figure 4.6. This figure shows, for the dynamic environment (experiment 1B), that the agent's estimation of the agency parameter \hat{b}_1 is the same with and without prior. The parameter \hat{b}_1 is the parameter with which the agent can control mass m_1 and the dotted line represents the true value of b_1 .

4.3.2. EXPERIMENT 2: INCORRECT SOA IN AMBIGUOUS ENVIRONMENTS

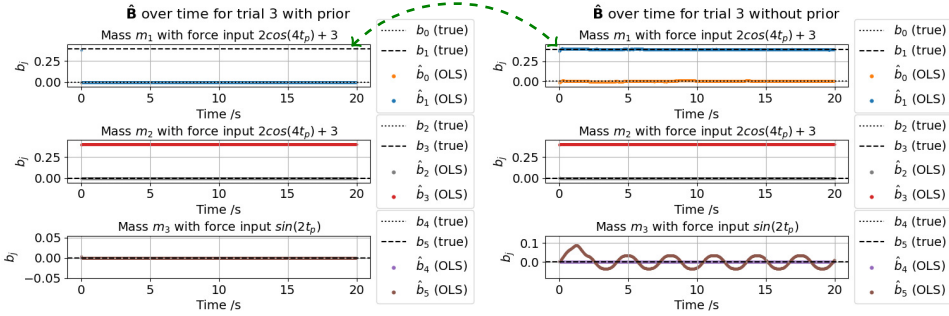
For an ambiguous environment (experiment 2), Table 4.2 shows the percentage of time that the agent has a SoA over each state j of the estimated $\hat{\mathbf{B}}$ out of the 1000 20 s trials. The agent with prior generally incorrectly believes that it has a SoA over mass m_2 with \hat{b}_3 most of the time (89.8%-94.7%), while the rest of the time the agent correctly believes has SoA over mass m_2 with \hat{b}_4 . Without prior, the agent's SoA is the same as in the ambiguous environment. The prior, requires the agent to have a SoA over only one of the observed states, induces an incorrect SoA in the agent. The incorrect SoA occurs when the prediction error of the states of mass m_1 is larger than that of mass m_2 .

Table 4.2. This table shows that, for the ambiguous environments (experiment 2) of the point mass mRHI, the agent using the OLS algorithm with a "prior" has a more incorrect SoA than the agent without any prior knowledge. This table is a continuation of Table 4.1 that describes experiment 1.

Exp. nr.	% of time that $\hat{b}_0 > 0$	% of time that $\hat{b}_1 > 0$	% of time that $\hat{b}_2 > 0$	% of time that $\hat{b}_3 > 0$	% of time that $\hat{b}_4 > 0$	% of time that $\hat{b}_5 > 0$
2A	0.00	5.35	0.00	94.7	0.00	0.00
2B	0.00	6.85	0.00	93.2	0.00	0.00
2C	0.00	10.2	0.00	89.8	0.00	0.00
<i>Without prior</i>						
2A,2B,2C	0.00	100	0.00	100	0.00	92.8

To visualize the incorrect estimation of the agency matrix, $\hat{\mathbf{B}}$, Figure 4.7 compares and no prior condition because the estimation of \hat{b}_1 is equal.

the OLS estimated matrix $\hat{\mathbf{B}}$ with and without prior over time for one trial of the active synchronous condition (experiment 2A).⁴ Again, the arrow indicates the location where parameter \hat{b}_j should be non-zero. The left plot (Figure 4.7a) shows that with a prior the agent incorrectly chooses parameter \hat{b}_3 as non-zero, while the right plot shows that the agent without a prior chooses parameters \hat{b}_1 , \hat{b}_3 , and \hat{b}_5 , as non-zero (Figure 4.7b). Thus, the prior promotes an incorrect SoA because without the prior the agent tries to gain agency over all three masses.



(a) OLS estimated \hat{b}_j over time with prior (see Figure 4.1 for (b) OLS estimated \hat{b}_j over time without prior (see Figure 4.2 for algorithm).

Figure 4.7. These plots show that, for experiment 2A, the agent's prior knowledge motivates an incorrect estimation of the agency parameter \hat{b}_j meaning the SoA is also incorrect. Both agents estimate the agency parameters, \hat{b}_j , over time window $\delta_N = 200$, with OLS. The true values of the agency parameter b_j , represented by the black dashed lines, show that only non-zero parameter should be b_1 (see equation 3.14) because this parameter connects the agent's input force, F_{int} , to the states of mass m_1 .

Similarly, Figure 4.8 shows that for experiment 2B the agent has an incorrect SoA over mass m_2 . Experiment 2B is also the active synchronous condition but with a decrease in mass for mass m_2 . Comparing the estimated agency parameter, $\hat{\mathbf{B}}$, from experiment 2A (see Figure 4.8a) to 2B, shows two differences. First, the estimated \hat{b}_3 is higher in experiment 2B because of the lower mass m_2 . Second, looking closely, the estimation of \hat{b}_3 fluctuates more in experiment 2B than 2A. Based on this trial, the decrease in mass decreases the precision of the agency estimation \hat{b}_3 .

In the active asynchronous environment (experiment 1C), I introduce a time lag introduced in the force input of mass m_2 (w.r.t. the agent's force input, F_{int}). Figure 4.9a shows that, different from all previous experiments, the agent's estimation of the non-zero agency parameter, \hat{b}_3 , fluctuates. In the OLS algorithm with prior, the agent does not have any prior belief about the required precision of the agency parameter. The agent believes that by choosing \hat{b}_3 as non-zero its perceived error is lower than with \hat{b}_1 , even though this estimation would be constant (see Figure 4.9b). Despite the time-lag which makes the environment less ambiguous, the prior still motivates the agent to have an incorrect SoA over mass m_2 and the agency matrix, $\hat{\mathbf{B}}$.

⁴For brevity, I only present the results of the estimation accuracy of the agency matrix $\hat{\mathbf{B}}$ in this section. For the numerical results of the agency accuracy, perceived precision, state accuracy, and perceived error for each environmental condition of experiment 2, see Appendix C.2.

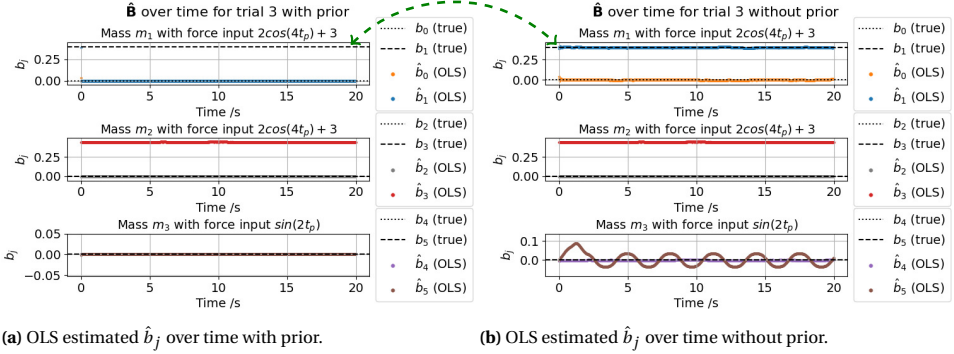


Figure 4.8. These plots show that in this ambiguous environment (experiment 2B) the agent's prior knowledge motivates an incorrect estimation of the agency parameter \hat{b}_j meaning the SoA is also incorrect. The difference between experiment 2A and 2B is that the mass of point mass m_2 is decreased by 10%. These plots are a continuation of Figure 4.7.

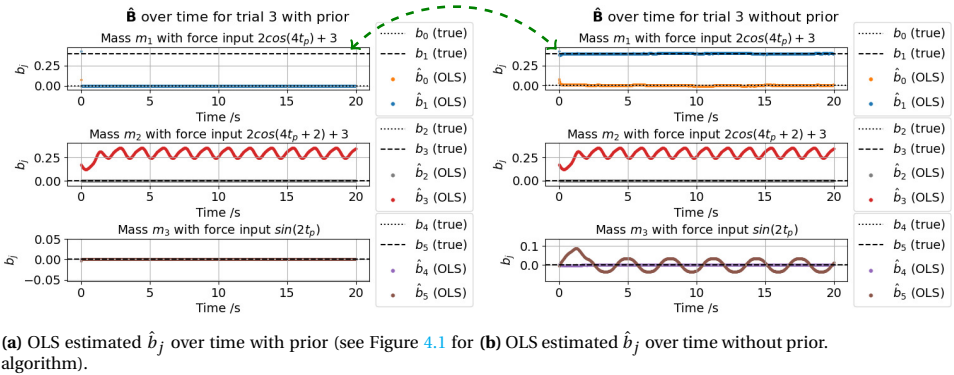


Figure 4.9. These plots show that for the ambiguous asynchronous environment (experiment 2C) the agent's prior knowledge motivates an incorrect estimation of the agency parameter \hat{b}_j meaning the SoA is also incorrect. The difference between experiment 2A and 2C is a 500 ms lag in the force input of mass m_2 w.r.t. the agent's force input, F_{int} . The plots are a continuation of Figure 4.7.

4.4. DISCUSSION

IN this OLS implementation of the point mass mRHI, the agent's prior knowledge allows the agent have a correct SoA in "normal" environments (experiment 1) and an incorrect SoA in ambiguous environments (experiment 2). On the surface, the results seem to match the findings of the human mRHI experiment by Kalckert and Ehrsson, where humans only experience an incorrect SoA over the rubber-hand when the rubber hand and their own hand move synchronously [31][32]. Therefore, this chapter achieves the goal of using OLS as a simple system identification method for an initial exploration of the effect of the prior in a simulation of the human mRHI.

4.4.1. COMPARISON OF THE SOA IN HUMAN AND THE AGENT

However, the simplicity of the algorithm causes the results to lack similarity to some aspects of human SoA, which impedes me from making generalizations about the influence of priors on SoA from the point mass mRHI. I will discuss two main differences between SoA in humans and the agent. These differences will highlight that the computational principles of SoA are too simplified in this point mass mRHI. Based on this discussion, formulate improvements for the point mass mRHI that should increase its similarity to the human mRHI. Chapter 5 implements these improvements.

THE AGENT'S "PRIOR" DOES NOT IMPROVE SOA

First, the results show that the prior does not improve the estimation of the agency parameter \hat{b}_2 . This diverges from human SoA, where the prior belief of strong causality can improve SoA [79][20] and even override other explanations of agency [76][73]. The current prior belief does not capture all aspects of the SoA principle that SoA arises when there is a *strong prior belief that action causes outcome*.

The agent does implement a part of the computational principle that a strong prior belief that action causes outcome induces a SoA. In the point mass mRHI, the prior is imposed on the correct parameter because the agency matrix, $\hat{\mathbf{B}}$, encodes the causality from the agent's known actions, F_{int} , to its outcomes. By enforcing the prior on the agency matrix, the agent separates the observations into the belief of causality and no causality based on how well the agent can predict the outcomes based on its actions. This aligns with the idea that self-other distinction arises from a SoA [77][21]. Thus, the current prior does describe that the action causes the outcome.

However, how the agent enforces the prior belief is incorrect because it does not encode the strength of the belief. Bayesian inference can define the strength of a prior meaning that a higher prior probability of agency results in a higher SoA [18]. Bayesian priors have two valuable characteristics that the current prior misses. First, a Bayesian prior is updated by the posterior making it more accurate as time progresses. The current prior knowledge is applied at every time step and the agent does not remember the agency/no-agency conclusion from the previous time steps. Therefore, the prior knowledge is constant despite the agent collecting more data. Second, there is no interaction between the prior in the OLS algorithm and the sensory evidence to adjust the SoA.

TOO MUCH INCORRECT SoA IN THE ACTIVE ASYNCHRONOUS CONDITION

The second difference between the simplest point mass and the human mRHI is that in the active asynchronous condition (experiment 2C), the agent's incorrect SoA over mass m_2 only decreased by 5% w.r.t. the active synchronous condition (experiment 2A). In the point mass mRHI, I expect a larger decrease in incorrect SoA over mass m_2 because for the agent without SoA the mean prediction error for motion state \hat{x}_3 (6.97×10^{-1}) is larger than for \hat{x}_1 (1.21×10^{-1}). In the active asynchronous condition of the human mRHI experiment, the participant's SoA decreases by 50% (from 2.7 to 1.3, where 1 means no illusory SoA) [32]. Therefore, similar to the asynchronous condition in the human mRHI experiment, the agent should experience a larger decrease in SoA to correctly minimize its *prediction error*.

The point mass mRHI does follow some aspects of the computational principle that the SoA arises with low prediction error because the agent does minimize its prediction error to some extent. OLS minimizes the sum of squares of the difference between the observed variables of the environment and the predicted values from the linear function. Then, the agent only gains SoA over the mass with the lowest prediction error. This indicates SoA is dependent on the magnitude of the prediction error and agrees with empirical evidence where the minimization of the prediction error explains shifts in the experience of agency. For example, a symptom of schizophrenia is the misattribution of agency to an external agent [78]. A previous study finds a higher signal-to-noise ratio in schizophrenics for signals that contribute to SoA, which increases the likelihood of a larger prediction error [92]. The increase in prediction error explains the abnormal updating of beliefs about the environment [93]. In both humans and the agent, prediction error influences the choice of SoA.

However, the agent's minimization of the prediction error is incorrect because the agent does not switch its SoA during a trial despite that switching SoA will lower the prediction error. The agent has an incorrect SoA over mass m_2 89.8% of the time. SoA does not switch during a trial because the agent predicts the motion state, \hat{x} , with its prior updated agency matrix, $\hat{\mathbf{B}}$. For the agent to switch its SoA, the agent's prediction error for the mass over which it currently has SoA over must be much larger than the prediction error of the states of the other mass. This is inconsistent with studies that find a longer interval between action and feedback results in a weaker SoA [82][94]. In a comparable study in humans, Perrykkad et al.[72] quantify prediction error in humans with eye-tracking and find that humans change their hypothesis of agency over a square on a computer screen when the prediction error is too high. More correlation between the prediction error on the SoA should make the results of experiment 2C more comparable to the asynchronous active mRHI condition of Kalckert and Ehrsson's experiment [31][32].

4.4.2. IMPROVEMENTS TO THE POINT MASS MRHI

Based on the discussion of the simplest point mass mRHI, I will formulate three improvements to the simulation that should increase the similarity between the agent and human SoA.

First, in the future point mass mRHI, I will *quantify SoA as the posterior probability of agency* so that it is a continuous variable. Empirical research theorizes that human

SoA is a continuous experience (e.g. through questionnaires in the mRHI [32]). Moreover, SoA as a posterior probability of agency aligns with the Bayesian cue integration [19] (Sec. 2.3) and the Bayesian computational models of SoA [21][20]. Another reason for a continuous SoA is that the magnitude of the three computational principles of SoA should directly affect the magnitude of the agent's SoA. For example, predictive processing frameworks postulate that the degree of prediction error correlates with the magnitude of SoA [95][96][34]. Continuous SoA should result in more flexible switching of agency during a trial. This might lead to an increase of correct SoA in the active asynchronous condition (experiment 2C) of the point mass mRHI.

Another improvement to the point mass mRHI, is changing the prior belief to a *Bayesian prior* to encode the strength of the causal prior belief. Sec. 2.1 specified the general advantages of a Bayesian prior. The main advantage for the point mass mRHI is that it is more intuitive to translate the general prior that the agent can have agency over one mass into a Bayesian prior. This prior then adjusts during the experiment to improve future inferences, which should lead to an improved estimation of the agency parameter \hat{b}_1 .

A final adjustment to the point mass mRHI involves the third computational principle of SoA, the precision of the sensory signals. In the current point mass mRHI, the agent has to choose which mass it has agency over in the first time step. However, the uncertainty on the states means that the agent might be unsure about which mass it has agency over. The agent should be able to continuously assign the incoming sensory data to whether it believes it does or does not have agency over that data. To illustrate, in the first few time steps of the simulation, the agent might be 70% certain that the states of mass m_1 are a result of its input force, leaving 30% possibility that the agent has agency over mass m_2 . Allowing the agent to continuously attribute each data point to the agent's own agency means the precision of the sensory signals has a more direct effect on agency, which should result in a more continuous SoA.

4.5. CONCLUSION

THIS chapter is the first step to answering the research question: *What is the effect of the prior on a sense of agency in the point mass moving rubber-hand illusion? Is this effect comparable to a sense of agency in human?* In the simplest point mass mRHI, the agent has correct SoA in certain environments (experiment 1) but incorrect SoA in environments where sensory data is noisy and ambiguous (experiment 2). This OLS algorithm achieves its goal as the simplest implementation of the mRHI for an initial exploration of the effect of the prior in a simulation of the human mRHI.

However, the simplicity of the algorithm impedes me from generalizing the influence of priors on SoA from the point mass mRHI. In particular, the prior does not improve the SoA and the behavior of the agent in the active asynchronous (experiment 2C) condition does not align with the human mRHI experiment. This finding is likely a result of the three underlying computational principles of SoA that are too simplified in the simulation. Therefore, I formulate three improvements that should increase the similarity between the agent and human SoA for the point mass mRHI:

- Quantify SoA as a continuous variable that is defined by the posterior probability of agency.
- Define the prior knowledge that the agent only can have agency over one mass as a Bayesian prior.
- Allowing the agent to assign a probability that it has agency over a certain observed state.

Though I could continue to use the OLS algorithm by extending it to include all the features that should increase its similarity to humans, the algorithm will feel fabricated. In Chapter 5, I will present EM as a Bayesian parameter estimation tool for the point mass mRHI which more inherently incorporates the improvements from this chapter.

5

THE POINT MASS mRHI WITH EXPECTATION-MAXIMIZATION

This chapter is the second step out of two for answering: *What is the influence of the prior on SoA in the point mass mRHI?* This chapter reproduces the point mass mRHI with the system identification method *Expectation-Maximization* (EM). EM is a popular (partially) Bayesian inference algorithm that I will show is appropriate for the point mass mRHI simulation. Using the experimental conditions defined in Chapter 3, this chapter concludes that the prior in some ways improves the SoA in unambiguous environments and produces incorrect SoA in ambiguous environments.

First, Sec. 5.1, explains the theory of EM for estimating the agency matrix, $\hat{\mathbf{B}}$. Next (Sec. 5.2), I elaborate on the motivation for the EM algorithm from a computational and human perspective. Sec. 5.3, tests the EM algorithm in a static environment (experiment 1A). After explaining the implementation of this theory in Python, I present the results. The following section (Sec. 5.4) extends the EM algorithm to dynamic environments to perform experiment 1B, 1C, 2A, 2B, and 2C. In the discussion of this chapter, I highlight the benefits and downsides of the prior. Also, I examine what gaps exist in the EM implementation of the point mass mRHI as an analogy for the human mRHI and SoA.

5.1. THEORY: EXPECTATION-MAXIMIZATION FOR SOA

IN the point mass mRHI, the agent's goal is to determine the effect of its input force, F_{int} , on the states, \mathbf{x} , by estimating the agency matrix $\hat{\mathbf{B}}$. The agent in the simplest point mass mRHI (Chapter 4) used OLS to estimate matrix $\hat{\mathbf{B}}$, while in this chapter the agent will estimate this matrix with EM. The advantage of EM is that it implements the improvements defined in Sec. 4.4 that increase the similarity between the agent's and human SoA. Before I elaborate on why EM is a suitable algorithm (Sec. 5.2), this section presents the theory.

5.1.1. THE OBSERVED SENSORY DATA

To gain a SoA, the agent must first observe some sensory data. Recall that in the simplest point mass mRHI (Chapter 4), the agent uses its observations of the states, \mathbf{x} , and its derivatives, $\dot{\mathbf{x}}$, over the time window δ_N to estimate agency matrix $\hat{\mathbf{B}}$ with OLS. As simplicity facilitates understanding, this section uses the static environment in experiment 1A for an initial implementation of EM. Sec. 5.4 extends the EM algorithm to the dynamic environments (for experiment 1B to 2C).

In a static environment, it was sufficient for the OLS algorithm to determine the agent's binary SoA through the requirement that the j th unknown agency parameter should fulfill $\hat{b}_j > 0$ ¹. Therefore, for the EM algorithm, the agent will observe both the states, \mathbf{x} , and derivatives of the states, $\dot{\mathbf{x}}$. The agent summarizes all these *observations* in one vector by calculating an incremental value of the agency matrix, or \mathbf{B}_i with,

$$\mathbf{B}_i = \frac{\dot{\mathbf{x}}_i - \mathbf{A}\mathbf{x}_i}{(u_a)_i}, \quad (5.1)$$

This is different to the OLS algorithm because OLS estimates the matrix $\hat{\mathbf{B}}$ by taking the average over all past observations. Formulations of the EM algorithm exist that estimate the parameters A, B, C , and D of a non-linear state-space model given the measured input u and output signals y (e.g. [97]). However, our agent knows A, C , and D . Therefore, I prepare the data by calculating \mathbf{B}_i incrementally to serve as the agent's observations instead of output data \mathbf{y} .

The matrix \mathbf{X}_B stores the n ($j = 0, \dots, n-1$) observations of \mathbf{B}_i for all $N-p$ ($i = p, \dots, N-1$) past time steps. Thus, matrix \mathbf{X}_B is defined by,

$$\mathbf{X}_B = \begin{pmatrix} \mathbf{B}_p \\ \mathbf{B}_{p+1} \\ \vdots \\ \mathbf{B}_{N-1} \end{pmatrix} = \begin{pmatrix} b_{p,0} & b_{p,1} & \cdots & b_{p,n-1} \\ b_{p+1,0} & b_{p+1,1} & \cdots & b_{p+1,n-1} \\ \vdots & \vdots & \ddots & \vdots \\ b_{N-1,0} & b_{N-1,1} & \cdots & b_{N-1,n-1} \end{pmatrix} = (b_{ij}) \in \mathbb{R}^{(N-p) \times n} \quad (5.2)$$

where

$$p = \begin{cases} 0 & \text{if } (N-1) \leq \delta_N \\ (N-1) - \delta_N & \text{if } (N-1) > \delta_N \end{cases}$$

¹This contrasts with the dynamic environment that requires the prediction error of the motions states $\epsilon_{\dot{\mathbf{x}}}$ to determine SoA

At each time step, the agent observes partial data \mathbf{X}_B , which generates two probability distributions of b_{ij} that can each be estimated with a Gaussian function. The histogram in Figure 5.1 visualizes the Gaussian probability distributions by plotting the observed partial data \mathbf{X}_B for the final time step. The two ($k = \{0, 1\}$) Gaussians, can be described by their means, μ_k , variances σ_k^2 , and the amount of data in each distribution π_k . This unknown set of parameters are summarized with $\theta = \{\mu_0, \mu_1, \sigma_0^2, \sigma_1^2, \pi_0, \pi_1\}$. The EM algorithm can determine the parameters, θ , when the number of distributions is known before-hand.

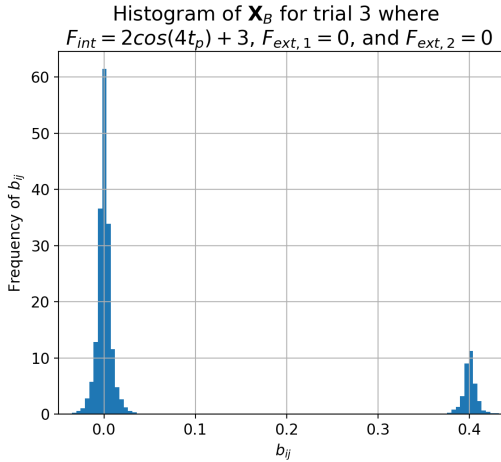


Figure 5.1. This histogram shows that partial data \mathbf{X}_B can be approximated by two Gaussian functions. This graph plots the observation matrix \mathbf{X}_B that is a collection observations \mathbf{B}_i in the final time window, δ_N , at $T = 20$ s in a static environment (experiment 1A).

5.1.2. EXPECTATION-MAXIMIZATION ALGORITHM OVERVIEW

To describe the observed data, \mathbf{X}_B in Figure 5.1, with two Gaussian functions, the agent needs to estimate the unknown set of parameters, θ , that later estimates matrix $\hat{\mathbf{B}}$. The parameters, θ , can be estimated by maximizing a likelihood function. Mathematically, this means we find the maximum likelihood estimate for the set of parameters θ , which maximize the probability of the observed data, or $P(\mathbf{X}_B | \theta)$. The maximum likelihood function is conveniently written as a log-likelihood,

$$\ell(\theta) = \ln P(\mathbf{X}_B | \theta). \quad (5.3)$$

However, this optimization problem is analytically intractable if there is hidden data. In the point mass mRHL, the data hidden to the agent is whether or not an observation is a result of the agent's force input, F_{int} . Recall in the Bayesian explanation of SoA (Sec. 2.3.2), the hidden cause of the sensory observations is the agency, A_h . The EM algorithm[98] solves the likelihood function by assuming the existence of some set of hidden causes, \mathbf{Z} . Therefore, the output of the EM algorithm is the estimated set of hidden causes, \mathbf{Z} , for each observation and the set of unknown causes θ for the entire set of observations, \mathbf{X}_B .

The EM algorithm starts with an estimation of the initial conditions for the parameters θ . Then at each time step, the agent observes \mathbf{B}_i and performs the EM algorithm until convergence of the log-likelihood function, $\ell(\theta)$. Each iteration of the EM algorithm [98] first estimates the expected value for the latent variables, \mathbf{Z} , given the observed data, \mathbf{X}_B , and the current estimate of the model parameters, θ (E-step). Second, the algorithm maximizes the likelihood of the model parameters, θ , assuming that the latent variables, \mathbf{Z} , are known (M-step). When $\ell(\theta)$ converges, the agent can update its agency parameter, $\hat{\mathbf{B}}$. See [98] for convergence properties and derivation of the EM algorithm.

There is a fundamental difference between the EM algorithm and the OLS algorithm from Chapter 4. In OLS, the agent observes data with which it estimates agency matrix, $\hat{\mathbf{B}}$. I extended the OLS algorithm to then determine whether the agent can use \hat{b}_j to control a mass or not. In contrast, the EM algorithm because this inherently clusters data according to the hidden cause (agency or no-agency) and then finds the parameters of each cluster. Clustering analysis is the grouping of the unlabeled observations such that the observations within a group are more similar to each other than those in another group. In the point mass mRHI, the agent wants to cluster the observations based on whether it has a SoA over these observations. With the parameters that describe the observations over which the agent has a SoA (i.e. the agency cluster), the agent can estimate $\hat{\mathbf{B}}$. Another difference to OLS is that EM performs soft clustering, which allows each observation to belong to several clusters with a fractional degree.

Next, I elaborate on the mathematical formulation of the log-likelihood function, followed by the E- and then the M-steps.

5.1.3. THE GAUSSIAN MIXTURE MODEL

The log-likelihood defines the formula that maximizes the likelihood of observing all \mathbf{X}_B given model parameters θ , or $P(\mathbf{X}_B|\theta)$ (eq. 5.3). A *Gaussian Mixture Model* finds an optimal way to group Gaussian distributed observations into K ($k = \{0, \dots, K-1\}$) clusters according to the parameters, θ .

Assume the agent observes $b_{p,0}, \dots, b_{N-1,n-1}$ and that each b_{ij} results from one of the K mixture components. Each observation b_{ij} has a label $Z_{ij} \in \{0, \dots, K-1\}$ that indicates the hidden cause of b_{ij} . In a static environment, there are two possible hidden causes meaning two clusters ($K = 2$). The latent variable Z_{ij} is either 0 or 1 depending on whether b_{ij} is a result of the agent's input force, F_{int} , or an external force (agency or no-agency). The agent has to uncover the latent variable Z_{ij} as the agent does not know which mass it controls.

For optimal grouping, the agent wants to maximize the probability of observing all data points, b_{ij} . To find this marginal probability, $P(b_{ij})$, I apply the law of total probability,

$$P(b_{ij}) = \sum_{k=0}^{K-1} P(Z_{ij} = k) P(b_{ij} | Z_{ij} = k), \quad (5.4)$$

where $P(Z_{ij} = k)$ is the prior probability that b_{ij} belongs to the k th cluster and $P(b_{ij} | Z_{ij} = k)$ represents the distribution of b_{ij} assuming it comes from cluster k .

The two terms that make up the marginal probability are known. First, the probability $P(Z_{ij} = k)$ describes the fraction of data that belongs to each cluster k and is termed

the mixture proportion, $\pi_k = P(Z_{ij} = k)$. The mixture proportions are non-negative and sum to one $\sum_{k=1}^K \pi_k = 1$. Second, assuming that a data point belongs to cluster k means $P(b_{ij} | Z_{ij} = k)$ can be formulated as a Gaussian probability density function, or

$$f(b_{ij} | \mu_k, \sigma_k) = \frac{1}{\sigma_k \sqrt{2\pi}} e^{-(b_{ij} - \mu_k)^2 / 2\sigma_k^2}. \quad (5.5)$$

Therefore, the marginal probability of b_{ij} can be reformulated with

$$P(b_{ij}) = \sum_{k=0}^{K-1} \pi_k f(b_{ij} | \mu_k, \sigma_k^2). \quad (5.6)$$

Eq. 5.6 defines a Gaussian Mixture Model and shows the model's dependence on the generally unknown parameters $\theta = \{\mu_0, \dots, \mu_{K-1}, \sigma_0^2, \dots, \sigma_{K-1}^2, \pi_0, \dots, \pi_{K-1}\}$. To find the optimal values of the unknown parameters, θ , an objective function maximizes the likelihood value (or, the joint probability) of observing all data b_{ij} and is defined by

$$P(\mathbf{X}_B | \theta) = \prod_{i=p}^{N-1} \prod_{j=0}^{n-1} P(b_{ij}) = \prod_{i=p}^{N-1} \prod_{j=0}^{n-1} \sum_{k=0}^{K-1} \pi_k f(b_{ij} | \mu_k, \sigma_k^2). \quad (5.7)$$

This equation simplifies to a sum by taking the log-likelihood such that,

$$\ell(\theta) = \ln P(\mathbf{X}_B | \theta) = \sum_{i=p}^{N-1} \sum_{j=0}^{n-1} \ln \sum_{k=0}^{K-1} \pi_k f(b_{ij} | \mu_k, \sigma_k^2). \quad (5.8)$$

This log-likelihood function is the same function as at the start of the section (eq. 5.3), but extended for the Gaussian Mixture Model. Finding the optimal parameters θ is computationally intractable because of the latent variables. This is where the EM algorithm is necessary.

5.1.4. E-STEP

The E-step calculates the distributions of each hidden cause, Z_{ij} , given the observed data, \mathbf{X}_B , and the current estimate of the parameters, θ , with Bayes' theorem (previously formulated in eq. 2.1). The posterior probability is calculated with,

$$P(Z_{ij} = k | b_{ij}) = \frac{P(b_{ij} | Z_{ij} = k) P(Z_{ij} = k)}{P(b_{ij})}. \quad (5.9)$$

From earlier derivations, the terms in this equation are known (see eq. 5.4-5.6). The posterior probability is also termed the responsibility, $r_i(Z_{ij} = k)$, because it describes how responsible a hidden cause is for an observation. Different to the typical EM where there is only one set of observations, the agent in the point mass mRHI collects data from $n = 6$ different sources (the number states of the state-space system). The agent knows that eventually it will gain agency over one of the $n = 6$ observations, by estimating the unknown parameters of the agency cluster, θ . Therefore, the probability that an observation is a result of a hidden cause varies per observation source, $j = 0, \dots, n-1$.

The responsibility at time step i , $\mathbf{r}_i(Z_{ij} = k)$, or for brevity \mathbf{r}_{ik} , is of size $1 \times n$, and is calculated with,

$$P(Z_{ij} = k | \mathbf{B}_i) = \mathbf{r}_i(Z_{ij} = k) = \frac{\boldsymbol{\pi}_k f(\mathbf{B}_i | \mu_k, \sigma_k^2)}{\sum_{k=0}^{K-1} \boldsymbol{\pi}_k f(\mathbf{B}_i | \mu_k, \sigma_k^2)} \quad (5.10)$$

5.1.5. M-STEP

The maximization step updates the estimations of the means, μ_k , variances, σ_k^2 , and mixture proportions, $\boldsymbol{\pi}_k$. For the equations of these parameters, I differentiate $\ell(\boldsymbol{\theta})$ with respect to each parameters and equate it to 0. After differentiation, this parameters are

$$\mu_k = \frac{\sum_{i=p}^{N-1} (\mathbf{r}_{ik} \cdot \mathbf{B}_i)}{N_k}, \quad (5.11)$$

$$\sigma_k^2 = \frac{\sum_{i=p}^{N-1} \mathbf{r}_{ik} (\mathbf{B}_i - \mu_k)^2}{N_k}, \quad (5.12)$$

$$\boldsymbol{\pi}_k = \frac{\sum_{i=p}^{N-1} \mathbf{r}_{ik}}{N}. \quad (5.13)$$

where N_k can be described as the number of observations assigned to cluster k , and calculated with,

$$N_k = \sum_{i=p}^{N-1} \sum_{j=1}^n r_{ij}(Z_{ij} = k).$$

The mean, μ_k , and variance, σ_k^2 , of each cluster k , are calculated over all n sets of observations of the partial data set \mathbf{x}_B . However, the mixture proportion, $\boldsymbol{\pi}_k$, is calculated separately for each j . Therefore, μ_k and σ_k^2 are scalars, whereas the mixture proportion is a vector, $\boldsymbol{\pi}_k \in \mathbf{M}_{1 \times n}(\mathbb{R})$. Also, the mixture proportion, or the prior, finds the mean of all passed \mathbf{r}_{ik} . Thus, the new prior is the mean of the past posterior probabilities causing the integration of inferences from past time steps with current inferences.

5.1.6. DEFINING AGENCY AND SOA

With this theoretical understanding of EM, I will link EM to agency and SoA. Connecting the variables to the point mass mRHI is a starting point for motivating why the EM is an appropriate algorithm for the point mass mRHI (in Sec. 5.2).

The unknown parameters of the agency cluster, $\boldsymbol{\theta}_0 = \{\mu_0, \sigma_0^2, \boldsymbol{\pi}_0\}$, describe the agent's estimation of agency. The mean μ_0 describes the weighted mean value of the observations belonging to a cluster. The mixture proportion, $\boldsymbol{\pi}_0$, describes which j th element of the observed b_{ij} the mean describes. Therefore, the mixture proportion is the agent's belief that each j set of observations is a result of F_{int} . By multiplying these values the

agent calculates the agency matrix $\hat{\mathbf{B}}$,²

$$\hat{\mathbf{B}} = \mu_0 \cdot \boldsymbol{\pi}_0. \quad (5.14)$$

Through Bayes theorem, I link SoA to the variables of the EM algorithm. Recall, the SoA is defined by the posterior probability that the agent has agency over hand h (see Sec. 2.3.2), or

$$P(A_h | \mathbf{s}, \mathbf{p}, \mathbf{v}_h) = \frac{P(\mathbf{s}, \mathbf{p}, \mathbf{v}_h | A_h) P(A_h)}{P(\mathbf{s}, \mathbf{p}, \mathbf{v}_h)}, \quad (5.15)$$

where the hidden cause of the sensory evidence, A_h , is agency over each hand and the sensory observation is collected from three sources: visual \mathbf{v}_h , proprioceptive \mathbf{p} , and sensorimotor \mathbf{s} . This aligns with Bayes theorem in the E-step (eq. 5.9) where the hidden cause of the sensory evidence, ($Z_{ij} = k$), also describes agency and the sensory observations are b_{ij} . Thus, SoA is described by the responsibility for the hidden cause of agency, $r_i(Z_{ij} = 0)$.

Table 5.1 lists the variables of the EM algorithm that describe the observations that result from the agent's force input, F_{int} . In the right column of the table, I connect the variables of the EM algorithm to agency and SoA based on the comparison between the two equations for Bayesian inference (eq. 5.9 and eq. 5.15).

Table 5.1. This table connects agency and SoA to the variables of the EM algorithm. These variables describe the (cluster of) observations that result from the agent's input force, F_{int} . The table provides a description of the meaning of each variable. In the right column the table links each variable to agency through the comparison between two Bayes equation for SoA in eq. 5.15 and for the responsibility in eq. 5.10.

EM variables	Variable name	General description	Connection to agency
$Z_{ij} = 0$	Latent variable	Hidden cause for cluster $k = 0$	Hidden cause is agency
μ_0	Mean	Mean of observations for $Z_{ij} = 0$	Magnitude of agency
σ_0	Variance	Inverse of precision of observations for $Z_{ij} = 0$	Precision of agency
$\boldsymbol{\pi}_0$	Mixture proportion	Prior probability of observing \mathbf{B}_i given the latent variable $Z_{ij} = 0$	Prior probability of agency
$f(b_{ij} \mu_0, \sigma_0^2)$	Gaussian function	Likelihood of observing \mathbf{B}_i given $Z_{ij} = 0$	Likelihood of agency
$r_i(Z_{ij} = 0)$	Responsibility	Posterior probability that each observation \mathbf{B}_i is a result of $Z_{ij} = 0$	Sense of agency

²Note that if the agent is not yet certain about the observations that it has agency over, then this automatically results in an incorrect agency matrix because the agent is still converging to the correct value of the mixture proportion of the agency matrix, $\boldsymbol{\pi}_0$. The agency matrix, $\hat{\mathbf{B}}$, can be incorrect even if the mean of the cluster being correct. In Sec. 5.3, the advantage of formulating the equation in this way is that it motivates the agent to converge to the correct mixture proportion to minimize the prediction error, $\boldsymbol{\epsilon}_{\hat{\mathbf{x}}_{ij}}$.

5.1.7. SUMMARY: EM ALGORITHM

I summarize the complete EM algorithm for an agent in a static environment in Algorithm 1 using the equations described in this section. This pseudocode runs the E and M-step until the log-likelihood converges to find the optimal parameters, θ , of the clusters and the hidden cause of the observations, Z_{ij} . At the end of each time step, the agent calculates the agency matrix $\hat{\mathbf{B}}$ with the parameters from the agency cluster $k = 0$.

Algorithm 1 The one dimensional Expectation-Maximization (EM) algorithm.

Input: \mathbf{B}_i
Initialize: $\theta^{(0)} = \{\mu_0, \mu_1, \sigma_0^2, \sigma_1^2, \pi_0, \pi_1\}$ and $\ell_0(\theta)$
for $i = 0$ **to** $N - 1$ **do**
 Observe \mathbf{B}_i
 while $\Delta\ell(\theta) \gg 0$ **do**
 $\ell_1(\theta) = \ell_0(\theta)$
 $\mathbf{r}_{ik} \leftarrow P(Z_{ij} = k \mid \mathbf{B}_i)$, where $\text{SoA} \leftarrow \mathbf{r}_{i0}$ ▷ E-step (eq. 5.10)
 $\mu_k \leftarrow \mu_k(\mathbf{r}_{ik}, \mathbf{B}_i)$ ▷ M-step (eq. 5.11-5.13)
 $\sigma_k^2 \leftarrow \sigma_k^2(\mathbf{r}_{ik}, \mathbf{B}_i, \mu_k)$
 $\pi_k \leftarrow \pi_k(\mathbf{r}_{ik})$, where $P(A_h) \leftarrow \pi_0$
 $\ell_0(\theta)$ ▷ Log-likelihood (eq. 5.3)
 $\Delta\ell(\theta) = |\ell_0(\theta) - \ell_1(\theta)|$
 end while
 $\hat{\mathbf{B}} = \mu_0 \cdot \pi_0$ ▷ Agency matrix (eq. 5.14)
end for
Output: $\theta = \{\mu_0, \mu_1, \sigma_0^2, \sigma_1^2, \pi_0, \pi_1\}$ and $\hat{\mathbf{B}}$

5.2. THE EM ALGORITHM MOTIVATION

THE main motivation for choosing the EM algorithm as system identification method for the point mass mRHI is that, like humans, the agent's SoA is described by a continuous posterior probability, the responsibility of the agency cluster r_{i0} . There are other advantages of the EM algorithm: its computational advantages and its similarity to human cognition.

5.2.1. COMPUTATIONAL ADVANTAGES OF EM

The EM algorithm has several computational advantages that make it a popular technique for determining the parameters of a mixture when the number of clusters is given *a priori*. First, the EM algorithm is an efficient way of solving for parameters of a Gaussian Mixture Model when there is hidden data [99]. Since the agent's observation data, \mathbf{X}_B , can be approximated by Gaussian distributions, the EM algorithm is an appropriate algorithm to describe agency and SoA. Also, the EM algorithm is guaranteed to converge as a result of the log-likelihood optimization [98]. However, it is important to note that convergence to a globally optimal solution is not guaranteed. Third, in contrast to OLS, the EM algorithm models noise because the parameters describe the data with probability distributions. Modeling noise is advantageous for an agent to determine the reliability of the observed data. A fourth computational advantage is that the EM algorithm is partially Bayesian. This means that EM uses the posterior probability from the previous iteration as a prior in the next iteration to make more informed decisions. The EM algorithm is not fully Bayesian because only the estimation of the hidden variable, Z_{ij} is Bayesian, while the estimations of the unknown parameters, θ , are point estimates.

5.2.2. SIMILARITY BETWEEN COGNITION AND THE EM ALGORITHM

The simplest point mass mRHI using OLS lacks enough similarity to the underlying computational principles of SoA (see discussion in Sec. 4.4) to generalize the simulation to the human mRHI experiment. The goal of changing the parameter estimator to EM is to increase the similarity between the point mass and human mRHI while maintaining the simplicity of the point mass simulation. Therefore, I will discuss the similarities between the EM algorithm and human cognition.

An advantage for the similarity between EM and human cognition is the use of Bayesian inference (E-step). Besides that the variables of EM can now be framed in terms of agency and SoA (see Table 5.1), there are three more specific advantages for using Bayes theorem (eq. 5.10) in the point mass mRHI. First, the SoA calculated with the responsibility term, $r_i(Z_{ij} = 0)$, is continuous because the assignment of data to the agency cluster is continuous. This is different from the OLS algorithm, where the agency assignment is binary. In the human mRHI, SoA over the rubber-hand is measured on a continuous scale [31][32]. Second, Bayes theorem incorporates information from the previous time step to update the responsibility term in the current time step. The agency prior, i.e. mixture proportion π_0 , updates with each time step meaning the agent's prior knowledge adjusts while gathering more data. Therefore, the responsibility $r_i(Z_{ij} = 0)$, or the agent's SoA, should improve with time. The prior also increases similarity to human SoA because it includes the predictive aspect of SoA (in addition to the postdictive process, see Sec. 2.2). Third, the variance of $Z_{ij} = 0$, σ_0^2 , specifies the precision of the agency sensory evidence,

which is one of the computational principles that influence the magnitude of the SoA. Higher precision (lower variance) correlates to a stronger SoA.

Another advantage of the EM algorithm as a model for SoA is that it is analogous to inference and learning in humans. The expectation step corresponds to inference, while the maximization step corresponds to learning [59]. Therefore, in the point mass mRHI, the E-step corresponds to the inference of SoA and the M-step is learning agency. According to Friston, it is remarkable that both inference and learning are driven by the minimization of one principle: the log-likelihood [59].³ Comparing human performance in perceptual and motor tasks to ideal observers, such as maximum-likelihood, suggests that this might be suitable frameworks for understanding aspects of human cognition [47][100].

A final advantage is that EM is related to Friston's promising brain theory Active Inference because it is an extension of one of the theories within Active Inference, *Dynamic Expectation Maximization* (DEM). As the naming suggests, DEM [101] is an extension of EM [98] that applies to dynamic systems. Due to the complexity and minimal implementations of DEM in artificial agents, I limit this research to the implementation with EM. Nonetheless, EM is a forerunner of DEM meaning it incorporates some parts of this novel system identification method.

³In Active Inference, this is the minimization of surprise about sensory inputs (i.e. deviations from predictions).

5.3. EXPERIMENT 1A: EM IN A STATIC ENVIRONMENT

INITIALLY, I test the EM algorithm (from Algorithm 1) in the static environment defined by experiment 1A as the first step towards formulating an EM algorithm for the dynamic point mass mRHI simulations. This section first describes the implementation in Python and then discusses the results of this algorithm.

5.3.1. METHOD: EM ALGORITHM FOR THE POINT MASS MRHI IN PYTHON

While detailing the steps in Python 3, I will explain how I apply the general equations from the EM algorithm to the point mass mRHI⁴. The implementation of the EM algorithm in Python is important because it defines the initial conditions and explains how to enforce the non-zero prior in the M-step. These are specific for the static environment.

INITIALIZATION

To simplify some calculations and allow plotting, first, I import the following libraries.

Listing 5.1: Import libraries necessary.

```
1 import numpy as np
2 from scipy import stats
```

Next, the all N observations \mathbf{B}_i , stored in a .txt file are imported to the script. In reality, the observations \mathbf{B}_i should be calculated at each time step during the algorithm but for simplicity, I do this beforehand.

INITIAL CONDITIONS

The EM algorithm starts with initial guesses for the unknown parameters θ . Eventually, the estimated parameters, $\hat{\theta}$, will approximate the partial data, \mathbf{X}_B , as accurately as possible. The agent's estimation should look something like Figure 5.2. The initial conditions are the first step towards finding these probability distributions and the first moment that the agent's prior knowledge is enforced. For clarity, I separate the initial conditions into four categories.

The number of clusters. Since EM requires setting the number of hidden causes (or, clusters) *a priori*, the agent must guess this number. For this initial EM algorithm, the agent estimates two hidden causes: agency and no-agency. Therefore, two clusters are necessary to describe the data, $K = 2$.

The properties of the agency cluster. The agent knows that when the hidden cause of observations is its force F_{int} , i.e. $Z_{ij} = 0$, then the mean of those observations must be some positive number, or $\mu_0 > 0$. If the mean is $\mu_0 = 0$, then the agent has no control over these observations. Since the agent's prior specifies that it has agency over one of the three masses, the agent knows that its SoA should be strong for the agency observations. Therefore, the agent assumes the variance is some small number, $\sigma_0^2 \leq 10^{-3}$. For the initial implementation of the EM algorithm, the agent sets $\mu_0 = 0.1$ and $\sigma_0^2 = 10^{-3}$.⁵

⁴The source code for the EM algorithm in a static environment is available on GitHub https://github.com/mtan1503/point-mass-mrhi/blob/main/em-algorithm/agent_em-static.py

⁵Since the EM algorithm is known to be strongly influenced by the initial conditions, I test different initial values for the variance. I find that between a range of variances including 10^{-3} , the results are comparable.




Figure 5.2. This figure illustrates the probability distributions of the two clusters that the agent uses to describe all observed data points b_{ij} . This agent believes the b_{ij} that belong to the no-agency cluster, $Z_{ij} = 1$, are normally distributed around a mean $\mu_1 = 0$ with some small variance $\sigma_1^2 \leq 10^{-3}$, while the b_{ij} values that belong to the agency cluster, $Z_{ij} = 0$, are some positive value with a mean, $\mu_0 > 0$ and variance, $\sigma_0^2 \leq 10^{-3}$.

5

The properties of the no-agency cluster. The agent does not know what type of environment it will encounter before the simulation starts. Therefore, the agent only knows that $\mu_1 \gtrsim 0$ and $\sigma_1^2 \geq \sigma_0^2$. The agent expects a higher variance in the no-agency cluster because the agent's agency is described by a constant value. What the agent does know is that if it finds "the agency combination" of properties, i.e. $\mu_k > 0$ and $\sigma_k \leq 10^{-3}$, then this cluster of observed data is a result of the hidden cause agency ($Z_{ij} = 0$). Therefore, the initial conditions of the no-agency cluster are guessed for the agent. For the purpose of this initial implementation of the EM algorithm, the mean and variance are guessed at $\mu_1 = 0$ and $\sigma_1^2 = 10^{-3}$.

The mixture proportions of the clusters. The agent knows that it has agency over only one out of the n observations of matrix \mathbf{B}_i because the prior specifies that the agent will gain agency by finding one of n parameters of \mathbf{B} as non-zero. Therefore, the mixture proportions, i.e. prior, of the agency ($Z_{ij} = 0$) and no-agency ($Z_{ij} = 1$) hidden causes, are defined by,

$$\begin{aligned}\pi_0 &= \frac{1}{n} = \frac{1}{6}, \\ \pi_1 &= \frac{n-1}{n} = \frac{5}{6}.\end{aligned}$$

For the initial condition, the mixture proportion across all n sets of observations is,

$$\boldsymbol{\pi} = \begin{pmatrix} 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 5/6 & 5/6 & 5/6 & 5/6 & 5/6 & 5/6 \end{pmatrix}.$$

Listing 5.2 shows the definition of these initial conditions in Python.

See Appendix E.

Listing 5.2: Initial conditions of the EM algorithm.

```

1 n_n = 6 # number of states
2 n_k = 2 # number of clusters
3 mu = np.zeros((n_k,1))
4 mu[0] = 0.1 # mean of agent
5 mu[1] = 0 # mean of environment
6
7 sigma = np.zeros((n_k,1))
8 sigma[0] = np.sqrt(10**(-3)) # standard deviation of agent
9 sigma[1] = np.sqrt(10**(-3)) # standard deviation of environment
10
11 pi = np.zeros((n_k,n_n))
12 pi[0,:] = 1/n_n # prior of agent
13 pi[1,:] = (n_n-1)/n_n # prior of environment

```

Initially, the agent does not know which column of \mathbf{X}_B belongs to the agency cluster; therefore, each column has an equal probability of belonging to the agency cluster, $\boldsymbol{\pi}_0$. As time progresses, the mixture proportions per j observation should converge to 1 for either $\boldsymbol{\pi}_0$ or $\boldsymbol{\pi}_1$ because each column of \mathbf{X}_B belongs either to the agency or no-agency cluster. The estimation of the mixture proportions should converge to the true mixture proportion,

$$\boldsymbol{\pi}^* = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}. \quad (5.16)$$

THE EM ALGORITHM

The overall EM algorithm is implemented in Python in Listing 5.3. The agent performs the parameter optimization in two loops. The first larger loop (line 11-40) is the time step, i , where the agent is at in the simulation. At each time step, the agent executes the EM algorithm to cluster the incoming sensory data (\mathbf{B}_i). The second loop (line 21-38) is within the first loop and specifies the number of iterations of the EM algorithm. The number depends on whether the log-likelihood, $\ell(\theta)$, converges before the maximum number of iterations is met. Since the agent performs EM at each time step, the number of maximum iterations is set to a low number to reduce calculation time. In this case, this maximum is set to 5, but the algorithm generally converges after 1 or 2 iterations.

Listing 5.3: Iteration loop of EM algorithm.

```

1 # initialize matrices for storing data
2 r_i = np.zeros((N,n_n,n_k,trials)) # responsibility
3 B_hat = np.zeros(shape=B_i.shape) # estimation of B matrix
4
5 # convergence criteria
6 max_iter = 5 # maximum number of iteration steps allowed
7 tol = 0.001 # difference in the log-likelihood
8 steps_window = delta_N
9
10 # for every time step
11 for i in t_steps:
12     # set window boundary
13     d_Nt = steps_window
14     if i-d_Nt < 0:
15         d_Nt = i

```

```

16 # log-likelihood with the current parameters
17 ll_new = log_likelihood(pi, mu, sigma, X_B[i-delta_N:i+1,:])
18
19
20 # run EM until convergence criteria are met
21 for it in range(max_iter):
22     ll_old = ll_new
23
24     # E-step
25     r_i[i,:,:] = e_step(pi, mu, sigma, X_B[i,:])
26     # M-step
27     mu, sigma_temp, pi = m_step(r_i[i-d_Nt:i+1,:,:], X_B[i-d_Nt:i
28 +1,:])
29
30     # only calculate the std after some time steps passed
31     if i>10:
32         sigma = sigma_temp
33
34     # check exit condition
35     ll_new = log_likelihood(pi, mu, sigma, X_B[i-delta_N:i+1,:])
36     diff_ll = abs(ll_old-ll_new)
37     if (diff_ll<tol) or np.isnan(diff_ll)==True:
38         break
39
40 # estimate B_hat using agency mu and pi
41 B_hat[i,:] = pi[0,:]*mu[0]

```

5

One adjustment I make to the EM algorithm is adding line 30 and 31 to the algorithm. In the first few time steps, the agent has a limited amount of data. Therefore, initially, the agent's estimation of the standard deviation of each cluster, σ_k , is too small and does not reflect the noise in the data. The overly precise σ_k will limit the agent from adjusting the mean of each cluster, μ_k , to a correct value. Thus, I set all σ_k to the initial values for the first 10 time steps.

After the EM algorithm at the end of each time step, the agent estimates its agency matrix, $\hat{\mathbf{B}}$ (line 41).

E-STEP

The expectation step (in line 29 of Listing 5.3) finds the responsibility term for each cluster, $r_i(Z_i = k)$ (see eq. 5.10). Listing 5.4 details the function for the E-step. In Python, the likelihood or Gaussian function $f(\mathbf{B}_i | \mu_k, \sigma_k^2)$ is calculated with the `stats.norm` function from the `scipy` library. The prior, π_k , is known from the M-step. To calculate the marginal likelihood, $P(\mathbf{B}_i)$, which is the denominator of the responsibility equation I apply the law of total probability (eq. 5.4). The responsibility is then calculated per cluster k .

Listing 5.4: Expectation (E) step of the EM algorithm.

```

1 def e_step(pi_k, mu_k, sigma_k, B_i):
2     """The expectation step: compute the probability each data point is
3     a result of cluster k, P(k|x_i)"""
4     # find the likelihood, P(B_i|k)
5     f_k = stats.norm(loc=mu_k, scale=sigma_k).pdf(B_i)
6     # find the marginal likelihood, P(x_i)
7     r_evidence = np.sum(pi_k * f_k, axis=0)
8     # find the posterior P(k|B_i) or responsibility

```



```

8     r_ik = [pi_k[k,:] * f_k[k,:] / r_evidence for k in range(len(mu_k))
9     ]
10    r_ik = np.array(r_ik).T
11    return r_ik

```

M-STEP

In the maximization step (line 31 of Listing 5.3), the agent updates the estimations of the means, μ_k , variances, σ_k^2 , and mixture proportion π_k . Listing 5.5 details these functions in Python.

Enforcing the agent's prior. To enforce the agent's prior knowledge that it has agency over *exactly* one of the n observations, I adjust the calculation of the mixture proportion from the general EM formula. The denominator of mixture proportion π_k is equal to the number of data points in $r_i(Z_i = k)$, which is typically equal to the number of time steps passed (see eq. 5.13). However, in the point mass mRHI, the number of time steps passed, N , is equal to only $1/n$ of the total number of observable data points from \mathbf{X}_B . Since the agent knows that it can gain agency over one mass, the agent knows that it can control exactly $1/n$ of the data points in \mathbf{X}_B . Therefore, to determine the mixture proportions of the agency cluster, π_0 , the denominator is set equal to the length of one column (i.e. the number of time steps passed within the time window). The mixture proportions must obey the criteria $\sum_{k=0}^{K-1} \pi_k = 1$, meaning the mixture proportion of the environment is $\pi_1 = 1 - \pi_0$.

Listing 5.5: Maximization (M) step of the EM algorithm.

```

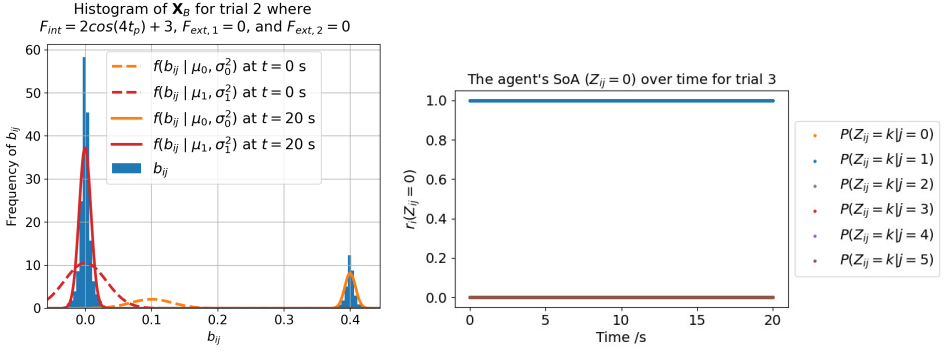
1 def m_step(r_ik, x):
2     """The maximization step: update parameter values"""
3     # sum of r_ik (posterior) over i, resulting shape Nk [Kx1]
4     Nk = np.sum(np.sum(r_ik,axis=1),axis=0)
5     # update mean
6     mu_k = [np.sum(r_ik[:, :, k]*x)/Nk[k] for k in range(r_ik.shape[2])]
7     mu_k = np.array([mu_k]).T
8     # update std
9     var_k = [np.sum(r_ik[:, :, k]*(x-mu_k[k,])**2)/Nk[k] for k in range(
10    r_ik.shape[2])]
11    sigma_k = np.array([np.sqrt(var_k)]).T
12    # update mixture proportions of the clusters (should actually be
13    mean but this allows us to enforce prior)
14    pi_k = np.sum(r_ik,axis=0).T
15    # enforce prior on agent
16    pi_k[0,:] = pi_k[0,]/Nk[0]
17    pi_k[1,:] = 1-pi_k[0,:]
18    return mu_k, sigma_k, pi_k

```

5.3.2. RESULTS AND DISCUSSION

For experiment 1A (see Sec. 3.3), the EM algorithm successfully attributes the observed \mathbf{B}_i to the agency ($Z_{ij} = 0$) and no-agency ($Z_{ij} = 1$) hidden causes at each time-step. The agent has SoA over mass m_1 100% of the time. Figure 5.3a plots the probability distributions of each cluster before and after the EM algorithm, which shows that the agent can move the initial incorrect estimation of the agency mean, μ_0 , to the correct location of

the data point which is generated by the agent's input force. Next, Figure 5.3b shows that the agent can continuously assign the probability that b_{ij} belongs to the agent cluster. The responsibility term, $r_i(Z_{ij} = 0)$, the agent's SoA, shows that throughout the simulation the agent has SoA over mass m_1 because $r_i(Z_{ij} = 0)$ is greater than 0 only for the 2nd set of observations ($j = 2$) of matrix \mathbf{B}_i .



(a) The histogram plots the value b_{ij} values against their frequency. The four probability distributions of the agent, $f(b_{ij} | \mu_0, \sigma_0^2)$, and the environment, $f(b_{ij} | \mu_1, \sigma_1^2)$, plot the EM estimated probability distribution before and after all time steps.

(b) The graphs plots the agent's SoA over time. SoA is described by the responsibility term, $r_i(Z_{ij} = 0)$. This describes probability that the data point, b_{ij} , is a result of the hidden cause of agency ($Z_{ij} = 0$) over time for each state j .

Figure 5.3. This figure shows that the *Expectation-Maximization* (EM) algorithm correctly clusters the data into agency and no-agency clusters for the static environment defined by experiment 1A.

Comparing the agent's estimation of the agency matrix $\hat{\mathbf{B}}$ from the EM algorithm to the OLS algorithm, shows similar accuracy (1.42×10^{-6} vs. 8.56×10^{-6}) and precision (5.88×10^{-8} vs. 1.49×10^{-7}). Also, the estimation of the agency parameter that exerts control of mass m_1 , i.e. \hat{b}_{i2} , is similar between EM and OLS (error is 1.41×10^{-6} vs. 3.32×10^{-6}). This initial implementation of the EM algorithm shows that the agent gains a correct SoA over the observations from mass m_1 by separating the observed \mathbf{B}_i into agency and no-agency clusters. Moreover, this section furnishes a greater understanding of the EM algorithm before Sec. 5.4 extends it to the more complex dynamic environments.

5.4. EM IN A DYNAMIC ENVIRONMENT

IN this section, I extend the EM algorithm from the static to the dynamic environment. The goal is to show that using a prior in an EM algorithm to determine the agency parameter results in improved SoA under normal environments (experiment 1), and is susceptible to illusory SoA in ambiguous environments (experiment 2).

5.4.1. EXTRA HIDDEN CAUSES

In a dynamic environment, there are more hidden causes, Z_{ij} , influencing the observed data because of the external forces acting on both masses m_2 and m_3 . This requires two changes to the EM algorithm. The first change is that the extra hidden causes of the sensory evidence influences the number of clusters, K . In the static environment EM algorithm, the hidden causes of the data are either the agent's input force, F_{int} , or no force input. However, in dynamic environments, the states of mass m_2 and m_3 are affected by external forces $F_{\text{ext}1}$ and $F_{\text{ext}2}$, respectively. The agent has to recognize which observations are a result of its input force amidst the four possible hidden causes.

The second change is that the agent needs an extra set of observations to determine the hidden causes of the observed data. In the static environment, the agent only uses the observed b_{ij} to cluster the data (see Figure 5.3a) because all other observations are $b_{ij} \approx 0$. However, with the extra external forces in the dynamic environment, these observations will be larger than 0 ($b_{ij} > 0$), meaning the observed b_{ij} of mass m_2 or m_3 might overlap with the observations of mass m_1 . The extra set of observations that the agent uses to gain a SoA is the prediction error of the motion states $\dot{\mathbf{x}}$. Again, the prediction error, $\varepsilon_{\dot{x}_{ij}}$, describes the difference between the predicted $\hat{\dot{x}}_{ij}$ (from the estimated agency matrix $\hat{\mathbf{B}}$) and the observed \dot{x}_{ij} . The prediction error was successful for correctly estimating $\hat{\mathbf{B}}$ with OLS in unambiguous environments (see Sec. 4.3) and is one of the underlying principles of SoA (see Sec. 2.2.2). This means that the EM algorithm has to be extended to its two-dimensional form.

CHANGING THE MATRIX SIZES

As a result of the increase to two-dimensional observations and $K = 4$ hidden causes, I have to adjust the sizes of the matrices.

The second set of observations for the n ($j = 0, \dots, n-1$) states and $N-p$ ($i = p, \dots, N-1$) past time steps is described by,

$$\mathbf{X}_\varepsilon = \begin{pmatrix} \varepsilon_{\dot{x}_{p,0}} & \varepsilon_{\dot{x}_{p,1}} & \cdots & \varepsilon_{\dot{x}_{p,n-1}} \\ \varepsilon_{\dot{x}_{p+1,0}} & \varepsilon_{\dot{x}_{p+1,1}} & \cdots & \varepsilon_{\dot{x}_{p+1,n-1}} \\ \vdots & \vdots & \ddots & \vdots \\ \varepsilon_{\dot{x}_{N-1,0}} & \varepsilon_{\dot{x}_{N-1,1}} & \cdots & \varepsilon_{\dot{x}_{N-1,n-1}} \end{pmatrix} = \left(\varepsilon_{\dot{x}_{ij}} \right) \in \mathbb{R}^{(N-p) \times n} \quad (5.17)$$

where,

$$p = \begin{cases} 0 & \text{if } (N-1) \leq \delta_N \\ (N-1) - \delta_N & \text{if } (N-1) > \delta_N \end{cases}$$

The overall matrix, \mathbf{X} , stores all observations by stacking \mathbf{X}_B (see eq. 5.2) and \mathbf{X}_ε (see eq. 5.17) and is visualized in Figure 5.4.




Figure 5.4. This figure shows the matrix size of the observed data, \mathbf{X} , that the agent uses to gain SoA over one of the three masses in the point mass mRHI. \mathbf{X}_B is dimension 0 and \mathbf{X}_E is dimension 1.

The extra dimension for the observations means the dimensions of the unknown parameters θ also increase. The unknown parameters at each time step i ,

$$\theta^{(i)} = \{\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\pi}\},$$

which consists of the matrices for the means, $\boldsymbol{\mu}$, covariances, $\boldsymbol{\Sigma}$, and mixture proportions, $\boldsymbol{\pi}$. The parameters are described by,

$$\begin{aligned} \boldsymbol{\mu} &= \begin{pmatrix} \mu_{0,0} & \mu_{0,1} & \cdots & \mu_{0,K-1} \\ \mu_{1,0} & \mu_{1,1} & \cdots & \mu_{1,K-1} \end{pmatrix} \\ \boldsymbol{\Sigma} &= (\Sigma_0 \quad \Sigma_1 \quad \cdots \quad \Sigma_{K-1}) \\ \boldsymbol{\pi} &= \begin{pmatrix} \pi_{0,0} & \pi_{0,1} & \cdots & \pi_{0,n-1} \\ \pi_{1,0} & \pi_{1,1} & \cdots & \pi_{1,n-1} \\ \vdots & \vdots & \ddots & \vdots \\ \pi_{K-1,0} & \pi_{K-1,1} & \cdots & \pi_{K-1,n-1} \end{pmatrix}. \end{aligned}$$

The mean matrix $\boldsymbol{\mu}$, describes the mean for the two dimensions (i.e. b_{ij} and $\varepsilon_{\dot{x}_{ij}}$) per cluster, k . Second, the covariance matrix, $\boldsymbol{\Sigma}$, is a collection of covariance matrices for each cluster k . The covariance matrix for a cluster, Σ_k , is a 2×2 matrix that measures the correlation between the two dimensions of the observations, $\text{cov}(b_{ij}, \varepsilon_{\dot{x}_{ij}})$. The covariance matrix is calculated with,

$$\Sigma_k = \begin{pmatrix} \sigma_{xx}^2 & \sigma_{xy}^2 \\ \sigma_{yx}^2 & \sigma_{yy}^2 \end{pmatrix},$$

where the x and y represent the dimensions of the observations, b_{ij} and $\varepsilon_{\dot{x}_{ij}}$, respectively. The final unknown parameter, the mixture proportion matrix, $\boldsymbol{\pi}$, describes the prior probability that an each j observation belongs to a cluster, k .

CONNECTING THE OBSERVATIONS TO THE HIDDEN CAUSES

To visualize how the EM algorithm will estimate the parameters of the clusters, $\theta^{(i)}$, to best describe the observations, \mathbf{X} , Figure 5.5 plots the incrementally observed b_{ij} against the error of the predicted motion state $\hat{\mathbf{x}}^6$ for experiment 1B.

⁶For explanatory purpose, this figure plots the prediction of the motion states with the OLS estimation of $\hat{\mathbf{B}}$ without priors (see Sec. 4.4).

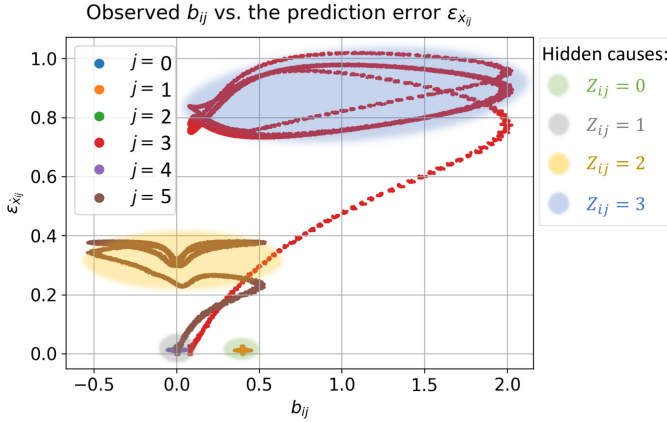


Figure 5.5. This plot shows how the agent could identify four hidden causes based on the observations b_{ij} and the perceived error \hat{x} in experiment 1B if it uses the EM algorithm. The motion states are predicted with an OLS algorithm without priors (see Sec. 4.2). Each set of dots represents a different, j , observation. The colored ovals are hypothetical clusters, which I added to illustrate how the agent should cluster the data according to the four hidden causes, Z_{ij} that affect the observations.

Figure 5.5 shows that the $n = 6$ sets of observations can be separated according to the four hidden causes. The ovals are the hypothetical clusters that the agent would estimate to separate the observations according to their hidden causes. I will elaborate on how the hidden causes relate to the observed data points. The first hidden cause, $Z_{ij} = 0$, is the agent's input force, F_{int} , and affects observations from $j = 1$. Since the agent knows this input force, the observed data points have a small distribution and a low error. The second hidden cause, $Z_{ij} = 1$, influences observations that are not affected by a force. Therefore, this hidden cause describes all observations b_{ij} which are $b_{ij} \approx 0$. In the static EM algorithm (Sec. 5.3), these observations were described the no-agency cluster. The third and fourth hidden causes, $Z_{ij} = 2$ and $Z_{ij} = 3$, are the external forces, $F_{\text{ext}1}$ and $F_{\text{ext}2}$. Different from the first two hidden causes, the observations have a high variance in both dimensions and a high prediction error. The reason is that these forces are unknown to the agent meaning the agent "observes" b_{ij} so that the agent can explain the states, \mathbf{x}_i and motion states, $\dot{\mathbf{x}}$, with its known input force, F_{int} (see eq. 5.1). To compensate for the unknown forces, the agent has to constantly adjust the observed b_{ij} . These data points generate a repeating pattern because all forces are sinusoidal.

Also, Figure 5.5 shows that even in the presence of more hidden causes, the agent should recognize which of the hidden causes is the agent's input force, F_{int} . We, as readers, know that the agent controls the data observed from $j = 1$ meaning hidden cause $Z_{ij} = 0$ affects the agent's observations. The agent does not know this but does know that the observations affected by its force input, F_{int} , should follow the three computational principles of SoA. First, the prediction error the observed motion states (i.e. $\varepsilon_{\dot{\mathbf{x}}_{ij}} \approx 0$) is low. Second, there is a causal relationship between actions and sensory consequences (i.e. $b_{ij} \neq 0$). Third, the sensory precision is high because the variance of the observations of $j = 1$ is low. Thus, despite the extra hidden causes, the agent should gain SoA

over the correct set of observations in Figure 5.5.

5.4.2. METHOD: IMPLEMENTATION OF THE EM ALGORITHM

To incorporate the extra dimension and the increased number of hidden causes, I adjust the Python code detailed in Sec. 5.3.1. The main changes to the code occur before and after the EM algorithm. However, an important change to the EM algorithm itself is the calculation of the mixture proportion, $\boldsymbol{\pi}$, in the M-step. I change this calculation so that the Bayesian prior still satisfies the agent's prior knowledge that its actions only have a causal effect on one set of observations⁷.

SETTING THE INITIAL CONDITIONS

Before the execution of the EM algorithm, I must adjust the initial conditions. The essence of the initial conditions is that the agent knows (i) the agency cluster is the only that can be described by low variance and positive observations for b_{ij} and (ii) the agent only has agency over 1 out of n observations per time step. For clarity, I separate the initial conditions into four categories.

5

The number of clusters. Since EM requires setting the number of clusters *a priori*, the agent must guess the number of hidden causes. The agent does not know the number of clusters before the simulation starts but does know the number of unknown parameters in the matrix \mathbf{B} . In the agency matrix, \mathbf{B} , the odd-numbered j elements are always $b_j = 0$, while the even-numbered j might be a number other than 0. Assuming that the agent knows that the maximum number of clusters, K , is dependent on the number of elements n , the agent knows the maximum number of clusters is,

$$K = \frac{n}{2} + 1 = 4.$$

The properties of the agency cluster. Like in the static environment, the agent knows that when observations are influenced by its own agency, i.e. $Z_{ij} = 0$, then the mean of those observations for the b_{ij} dimension must be some positive number, or $\mu_{0,0} > 0$. For the second dimension, the agent knows that it can only have SoA when its prediction error $\varepsilon_{\hat{x}_{ij}}$, is low (≈ 0). Therefore, the initial condition of the mean of the agency cluster is set to

$$\boldsymbol{\mu}_{:,0}^{(0)} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

The variance of two-dimensional observed data is described in a covariance matrix. If the agent gains agency over the correct mass, mass m_1 , then the non-diagonal entries of the matrix must be approximately 0 meaning the observed agency values b_{ij} are uncorrelated to the prediction errors $\varepsilon_{\hat{x}_{ij}}$. No correlation may seem unlikely since the predicted motion states, \hat{x}_{ij} , are calculated with the estimated agency parameter \hat{b}_{ij} .

⁷The entire code can be found on GitHub at https://github.com/mtan1503/point-mass-mrhi/blob/main/em-algorithm/agent_em-dyn.py.

However, if the agent correctly estimates agency parameter \hat{b}_{ij} , then the prediction error $\varepsilon_{\dot{x}_{ij}}$ is only a result of the noise on the motion state \dot{x}_{ij} . Therefore, the prediction error $\varepsilon_{\dot{x}_{ij}}$ is uncorrelated to the observations b_{ij} . For the agency cluster, the variance in b_{ij} and $\varepsilon_{\dot{x}_{ij}}$ is low, and therefore the initial covariance matrix is set to

$$\Sigma_0^{(0)} = \begin{pmatrix} 10^{-3} & 0 \\ 0 & 10^{-3} \end{pmatrix}.$$

The properties of the no-agency cluster(s). When there are four clusters, then there is a maximum of three remaining clusters describing the environment. One cluster describes the characteristics of a static environment ($Z_{ij} = 1$). The properties of the static environment are the same as the agent's cluster, except for that mean of the b_{ij} dimension should be around 0. Therefore, the mean of the no-agency cluster that describes static states are set to

$$\boldsymbol{\mu}_{:,1}^{(0)} = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

and the covariance matrix is $\Sigma_1^{(0)} = \Sigma_0^{(0)}$.

For the two remaining clusters, the agent only knows that it will not gain agency over these states. As a result, the prediction error, $\varepsilon_{\dot{x}_{ij}}$, is high and the variance for both dimensions is higher than the variance of the agency cluster. I guess the initial mean values for the no-agency clusters with

$$\boldsymbol{\mu}_{:,2:4}^{(0)} = \begin{pmatrix} 0.5 & 0.5 & 0.5 \\ 0.5 & 1.0 & 1.5 \end{pmatrix}.$$

The variance for both dimensions of these two no-agency clusters are larger than the variance of the agency cluster. Especially, the variance of the $b_{i,j}$ values is large because the agent does not know the applied external force. Therefore, I define the covariance matrices for the initial condition as follows,

$$\Sigma_2^{(0)} = \Sigma_3^{(0)} = \Sigma_4^{(0)} = \begin{pmatrix} 10^{-1} & 0 \\ 0 & 10^{-2} \end{pmatrix}.$$

The mixture proportions of the clusters. The initial condition of the mixture proportion of the agency cluster remains the same because the agent knows the probability that an observation is result of agency is 1 out of n . The probability that the observation per j is affected by a no-agency clusters is the remaining $n - 1$ out of n . Since the mixture proportions must obey $\sum_{k=0}^{K-1} \pi_k = 1$ and the agent estimates that all no-agency clusters have an equally likely probability of occurring for each j , the initial condition of the mixture proportion is,

$$\boldsymbol{\pi}^{(0)} = \begin{pmatrix} 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 5/18 & 5/18 & 5/18 & 5/18 & 5/18 & 5/18 \\ 5/18 & 5/18 & 5/18 & 5/18 & 5/18 & 5/18 \\ 5/18 & 5/18 & 5/18 & 5/18 & 5/18 & 5/18 \end{pmatrix}.$$

M-STEP: ENFORCING THE PRIOR WITH THE MIXTURE PROPORTION.

Aside from the change in matrix sizes, the only change for the dynamic environment within the E and M steps of the algorithm (w.r.t. Listing 5.3, 5.4, and 5.5 in Sec. 5.3) is the calculation of the mixture proportion is different in the dynamic environment. Recall, that in the static environment algorithm I enforce the prior that the agent has agency over one mass by adjusting the equation of the mixture proportion for the agent, $\boldsymbol{\pi}_0$. Instead of dividing the agent's sum of the responsibilities, $\sum_{i=p}^{N-1} \mathbf{r}_{i0}$, by the total number of observations, the equation for the mixture proportion of the agent is

$$\boldsymbol{\pi}_0 = \frac{\sum_{i=p}^{N-1} \mathbf{r}_{i0}}{\sum_{j=0}^{n-1} \sum_{i=p}^{N-1} \mathbf{r}_{ij}(Z_{ij} = 0)}. \quad (5.18)$$

The responsibility term in the numerator, \mathbf{r}_{i0} , is a vector of the current $N - p$ observations summed the passed time steps within the time window, δ_N . The responsibility term in the denominator, $\mathbf{r}_{ij}(Z_{ij} = 0)$, is additionally summed over the number of unknown parameters n . In a static environment, the mixture proportion of the no agency cluster is calculated with $\pi_1 = 1 - \pi_0$. However, with more than two clusters, this means that, $1 - \boldsymbol{\pi}_0 = \sum_{k=0}^{K-1} \boldsymbol{\pi}_k$. The numerator of each mixture proportion $\boldsymbol{\pi}_k$ remains the same, or $\sum_{i=p}^{N-1} \mathbf{r}_{ik}$. I assume the denominator of all no-agency clusters, x , are equal, therefore defined by,

$$1 - \boldsymbol{\pi}_0 = \frac{\sum_{k=0}^{K-1} \left(\sum_{i=p}^{N-1} \mathbf{r}_{ik} \right)}{x}.$$

Solving this equation I find that the denominator, x , is,

$$x = \frac{\sum_{k=0}^{K-1} \sum_{i=p}^{N-1} \mathbf{r}_{ik}}{1 - \pi_0}.$$

Replacing the denominator of the general equation for the mixture proportion (see eq. 5.13) with the denominator defined by x gives the mixture proportion for the remaining (no-agency) k clusters,

$$\boldsymbol{\pi}_k = \frac{\sum_{i=p}^{N-1} \mathbf{r}_{ik} (1 - \boldsymbol{\pi}_0)}{\sum_{k=0}^{K-1} \sum_{i=p}^{N-1} \mathbf{r}_{ik}}. \quad (5.19)$$

AFTER THE E- AND M- STEPS

After performing the E and M steps, the agent determines which k cluster it has agency over, calculates its agency parameter, $\hat{\mathbf{B}}$, and predicts the motion states, $\hat{\mathbf{x}}$.

The postdictive agency process. The agent uses the computational principles of SoA and prior knowledge to gain SoA over a cluster. This process can be compared to the postdictive process modulating SoA in humans [19]. In the initial conditions, I state that the agency cluster is cluster $k = 0$. However, since the agent initially guesses the mean of the agency cluster, $\boldsymbol{\mu}_0$, the estimated distribution might be closer to the data points that

are not a result of the agent's input force, F_{int} . Therefore, the agent can adjust which of the k clusters the agent has a SoA over.

To show how the agent finds its agency cluster, Listing 5.6 implements the process in Python. The agent's goal is to find the number of the agency cluster, k_{agent} . The prior knowledge states that the agent can only have agency over one of the observed states (lines 33 and 37). Based on the computational principles of SoA, the agent knows that for observation b_{ij} , the mean, $\mu_{0,\cdot}$, must be greater than 0 (line 29), while the variance (line 34-36) and prediction error (line 38-40) must be small. These steps allow the agent to change its agency cluster to contain the observations that the agent believes are influenced by its input force, F_{int} .

Listing 5.6: Python code for the part of the for loop in which the agent determines which cluster the agent has agency at each time step.

```

1 # run all time steps
2 for i in t_steps:
3     ...
4     ...
5     ...
25 # postdictive process of finding SoA (w prediction error and
    causality)
26 if i>10:
27     i_k = np.arange(0,n_k) # range of all k
28     mask = np.ones(len(i_k), dtype=bool) # boolean mask
29     static = np.where(mu[:,0]<0.1)[0] # static k
30     mask[static] = False # adjust agency k
31     k_agent = i_k[mask] # remove from agency k
32     # while more than one possible agency cluster
33     while sum(mask)>1:
34         max_cov = k_agent[np.argmax(cov[k_agent,0,0])]
35         mask[max_cov] = False # remove max. cov. from agency k
36         k_agent = i_k[mask] # adjust agency k
37         if sum(mask)==1: break
38         max_eps = k_agent[np.argmax(mu[k_agent,1])]
39         mask[max_eps] = False # remove max. error from agency k
40         k_agent = i_k[mask] # adjust agency k
41     k_agent = k_agent[0] # choose integer
42     k_env = i_k[~mask] # no-agency k are opposite of
    agency k
43     ...

```

Predicting the motion states. Another change in the EM algorithm for the dynamic environment is that the agent has to predict the motion states, $\hat{\mathbf{x}}$, so that it can observe a prediction error in the next time step. Therefore, like in the static environment, the agent first predicts $\hat{\mathbf{B}}$ with the mean and variance of the agency cluster. However, as the agency cluster can change during a trial of the point mass mRHI, the agency cluster is described by $k = a$, then when n is the number of states,

$$\hat{\mathbf{B}} = \mu_{0,a} \cdot \boldsymbol{\pi}_{a,*}^T \quad (5.20)$$

$$= \mu_{0,a} \cdot [\pi_{a,0} \quad \pi_{a,1} \quad \cdots \quad \pi_{a,n-1}]^T \quad (5.21)$$

where $\mu_{0,a}$ is the estimated mean of the agency cluster in the b_{ij} dimension and $\pi_{a,*}$ is the agency mixture proportion per state.

With matrix $\hat{\mathbf{B}}$, the agent can predict the states, $\hat{\mathbf{x}}$, with forward Euler⁸ and the motion states, $\hat{\dot{\mathbf{x}}}$, with agent's state-space equation (see eq. 4.1 in Sec. 3.1). At the start of the next time step, the agent uses the predicted $\hat{\mathbf{x}}$, to perceive the prediction error, $\varepsilon_{\hat{\mathbf{x}}}$.⁹

SUMMARY: THE ALGORITHM IN A DYNAMIC ENVIRONMENT

To summarize, I make two main adjustments to the EM algorithm: an increase in the number of clusters to $K = 4$ and including the prediction error as a second set of observations. These adjustments then require altering the EM algorithm from the static environment in Sec. 5.3.1 in the following ways: (1) changing the initial conditions, (2) changing the equation for the mixture proportions of the no-agency clusters, (3) the postdictive process of checking which cluster describes SoA, and (4) predicting the motion states to observe a prediction error.

The flowchart in Figure 5.6 connects SoA, the prior knowledge of the agent, and the previously described (Python) algorithm. The right column contains the simplified steps of this (EM) algorithm. The agent starts with the initial conditions. Then at each time step, the agent observes b_{ij} and senses the prediction error $\varepsilon_{\hat{\mathbf{x}}_{ij}}$. After performing the EM algorithm to determine the hidden causes, Z_{ij} , and the unknown parameters, $\theta^{(i)}$, the agent checks whether it has SoA over the correct cluster. Finally, the agent updates its agency parameter, $\hat{\mathbf{B}}$, to predict the motion states of the future time steps, $\hat{\dot{\mathbf{x}}}$. The left column is the "higher-level" description of the prior and the three principles of a SoA. The underlying computational principles of SoA (see Sec. 2.2.2) specify that SoA occurs when there is (i) a low prediction error between the agent's intended actions and sensory consequences, (ii) a strong prior belief that action causes outcome (i.e. causality), and (iii) a precise sensory signal. The prior specifies that the agent can only gain agency over one out of the three mass. The middle column translates the descriptive prior and principles to a mathematical formulation that I can implement in the algorithm.

5

5.4.3. RESULTS

In this section, I will present the results that show if the same prior knowledge (i) improves SoA in unambiguous environments and (ii) induces incorrect, illusory SoA in noisy, ambiguous environments. This requires a comparison to an algorithm without prior. Since Bayesian inference and clustering are a part of the EM algorithm, I cannot define an EM algorithm that contains no prior. Therefore, I compare the estimations of $\hat{\mathbf{B}}$ from the EM algorithm to the estimations from the no prior OLS algorithm. This OLS algorithm only relies on its sensory data to estimate $\hat{\mathbf{B}}$ (see Figure 4.2 in Sec. 4.2).

The tables presented in the results summarize the results from all 1000 trials of each experimental condition. Each figure shows one trial of the simulation per experimental condition. This trial represents a "typical" performance of the EM algorithm with and without prior.

⁸See Appendix B for the Forward Euler equations in discrete-time.

⁹The process of calculating the prediction error is the same as for the OLS algorithm in Sec. 4.2.

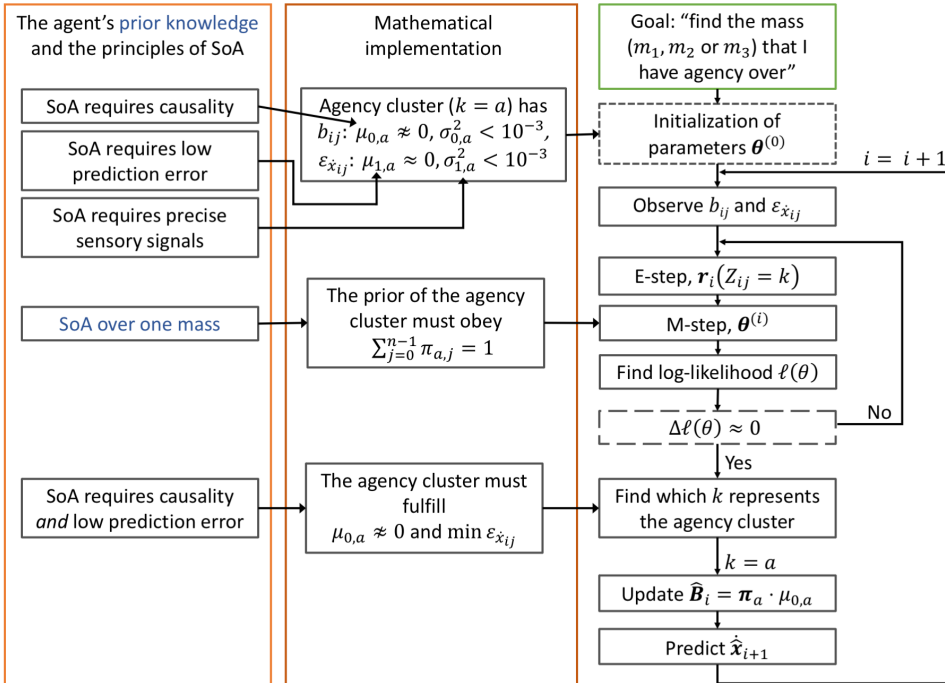


Figure 5.6. The flowchart connects SoA, the agent's prior, and the EM algorithm for the point mass mRHI. The left column is the "higher-level" description of the prior and the three principles underlying SoA. The middle column describes how the prior and these principles are translated to a mathematical formulation. The right column contains the simplified steps of the (EM) algorithm. See the text for a more elaborate explanation.

EXPERIMENT 1: IMPROVED SOA IN UNAMBIGUOUS ENVIRONMENTS

For the unambiguous environments (experiment 1), Table 5.2 shows the percentage of time that the agent has a SoA over each state j of the estimated $\hat{\mathbf{B}}$ out of the 1000 20 s trials per experimental condition. For EM, the SoA time is calculated by multiplying the agent's SoA, r_{ik} , with the length of a time step h . To determine if the prior improves a SoA, I compare the EM algorithm to the OLS algorithm without prior.

The agent correctly has a SoA over mass m_1 by estimating the agency parameter \hat{b}_1 as non-zero 100% of the time in the static environment (experiment 1A). These SoA results are the same with and without the prior. In a static environment, the agent does not need the prior to identify the mass that it has agency over. Therefore, the results of the unambiguous environmental conditions will focus on the dynamic environments (experiment 1B and 1C).

In the first dynamic environment (experiment 1B), the agent's correct SoA over mass m_1 decreases slightly to 99.9%. Adding noise (in experiment 1C) further decreases this SoA to 98.1%. Conversely, the agent without the prior has a correct SoA over mass m_1 100% of the time for all three experimental conditions. This suggests the prior and the EM algorithm decreases the accuracy of the agent's SoA for the dynamic environments. On the other hand, the agent's incorrect SoA over the other masses in the environment is low for the agent with prior (less than 1.99%), whereas the agent without the prior has SoA over all the masses that have a force applied to them (i.e. \hat{b}_1 , \hat{b}_3 , and \hat{b}_5). Based on these results, I would conclude that with EM, the agent's SoA improves for the masses that it does not have agency, mass m_2 and m_3 .

Table 5.2. This table shows that for unambiguous environments the agent with a prior generally has a correct SoA over mass m_1 . The table shows the percentage of time that the agent has a SoA over each observed state per experimental condition (exp. nr.) (see Sec. 3.3) of experiment 1. The agent with the prior uses the EM algorithm (see Figure 5.6) and the agent without the prior uses an OLS algorithm (see Figure 4.1). Each experimental condition is run for 1000 trials for a duration of 20 s. The time of SoA is calculated with, $r_{ik} \cdot h$, where h is the time step length h and r_{ik} is the SoA.

Exp. nr.	% of time that $\hat{b}_0 > 0$	% of time that $\hat{b}_1 > 0$	% of time that $\hat{b}_2 > 0$	% of time that $\hat{b}_3 > 0$	% of time that $\hat{b}_4 > 0$	% of time that $\hat{b}_5 > 0$
<i>With prior (EM)</i>						
1A	0.00	100	0.00	0.00	0.00	0.00
1B	0.00	99.9	0.00	0.00876	0.00	0.0375
1C	0.00	98.1	0.00	0.374	0.00	1.06
<i>Without prior (OLS)</i>						
1A	0.00	100	0.00	0.00	0.00	0.00
1B,1C	0.00	100	0.00	100	0.00	92.8

Figure 5.7 shows the agent's SoA over time for trial 3 of the dynamic environment with noise (experiment 1C). The dots indicate the SoA, or responsibility r_{ik} , that the agent has over each j observation set. Initially the agent has a SoA over the observations from $j = 3$, which are observations from mass m_2 . After the first time steps, the agent has 100% SoA over the observations influenced by the agent's force input, F_{int} (i.e. $j = 1$) and 0% SoA over the other masses.

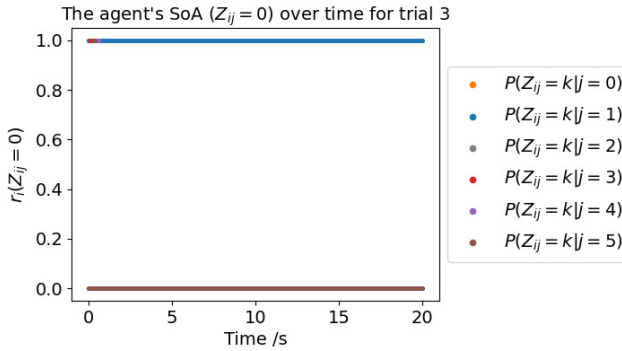


Figure 5.7. This plot shows that the agent has a correct SoA over the observations influenced by the agent's input force, F_{int} , for trial 3 in experiment 1B. SoA is described by the responsibility term, $r_i(Z_{ij} = 0)$. This describes probability that an observation, b_{ij} and $\varepsilon_{\dot{x}_{ij}}$, is a result of the hidden cause of agency ($Z_{ij} = 0$) over time for each j observation.

The results for the percent of time of SoA does not allow me to conclude whether the SoA *improves* as a result of the agent's prior. Overall the prior improves the SoA, but for mass m_1 , specifically, the SoA does not improve (Table 5.2). Therefore, to further show whether the SoA improves as a result of the prior, I compare the agent's SoA in three additional ways.

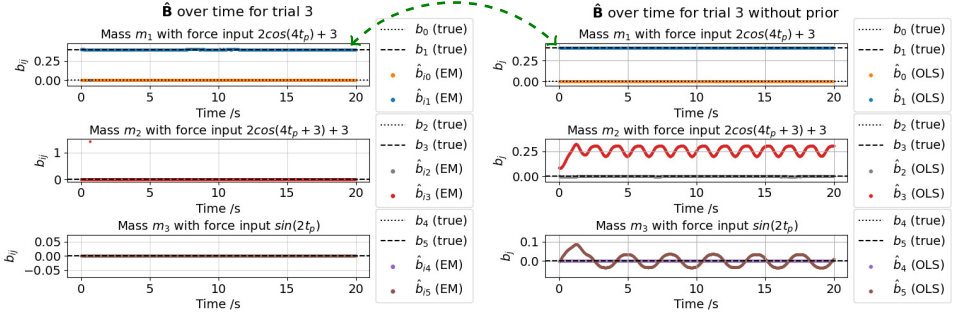
Estimation of the agency matrix, $\hat{\mathbf{B}}$. Besides the percentage of time of agency, four other performance measures indicate whether the agent's SoA has improved (see Sec. 3.3.2).¹⁰ These performance measures are relevant because though the agent's accuracy of the estimated agency matrix, $\hat{\mathbf{B}}$, and the error in the motion states, $\varepsilon_{\dot{x}}$, should be correlated to the agent's SoA, this does not have to be the case.

To understand how correct the agent's estimation of the agency matrix was in the dynamic environment of experiment 1B, Figure 5.8 plots each unknown j parameter in matrix $\hat{\mathbf{B}}$. The arrow points to the parameter \hat{b}_{i1} , which the agent should estimate as non-zero. The EM estimation (i.e. with prior) of \hat{b}_{ij} in Figure 5.8a correctly chooses the agency parameter \hat{b}_{i1} as non-zero. Without the prior in Figure 5.8b does not estimate any of the agency parameters as 0. Therefore, the prior improves the agent's SoA for the masses which it cannot control (i.e. mass m_2 and m_3) because the agent no longer has incorrect SoA over these masses.

Figure 5.8c provides a zoomed-in view of the estimated agency parameter \hat{b}_{i1} with (EM) and without prior (OLS). Both estimations are close to the true value of b_1 but the estimation without prior is closer to the real value of b_1 . According to the numerical values OLS estimation is slightly more accurate (see Appendix C.1). However, notice that the scale of the y-axis of Figure 5.8c is small, which indicates that the difference in accuracy between the prior and no prior is not significant. For experiment 1B, overall

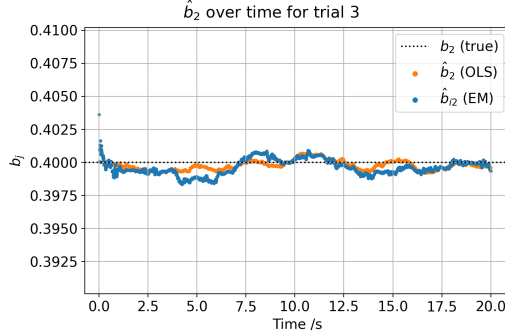
¹⁰See Appendix D.1 for the numerical values of the four performance measures per the experimental condition of experiment 1.

prior improves the SoA but the prior does not make a difference for the agent's SoA over the mass that it controls, i.e. mass m_1 .



(a) The EM estimated \hat{b}_j over time (with prior).

(b) The OLS estimated \hat{b}_j over time (without prior).

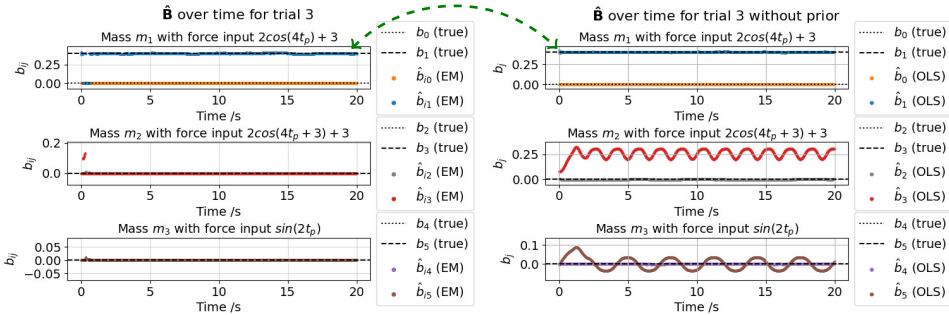


(c) EM and OLS estimated \hat{b}_2 over time.

Figure 5.8. These plots show that prior on the SoA parameter \hat{b}_j in the point mass MRHI promotes a correct SoA in experiment 1B. I plot the estimated agency parameter, \hat{b}_{ij} , over time window $\delta_N = 200$ steps, where the black dashed lines represent the true b_j (see equation 3.14). The parameter b_1 should be the only non-zero estimation element of $\hat{\mathbf{B}}$ because it connects the agent's input force to the states of mass m_1 . In Figure a, the agent estimates \hat{b}_j with EM (as described in Figure 5.6), whereas in the right figure (b) the agent estimates its agency parameter OLS and without prior knowledge. The bottom figure shows that in experiment 1B the estimations are relatively similar for EM, \hat{b}_{i1} , and OLS, \hat{b}_1 .


In the next experimental condition, I increase the state noise of motion state \mathbf{x}_1 (experiment 1C). Like Figure 5.8, Figure 5.9 compares the estimation of the agency matrix, $\hat{\mathbf{B}}$ with and without prior. Again, with the prior in Figure 5.9a the agent correctly estimates the agency parameter \hat{b}_{i1} as non-zero and without the prior in Figure 5.9b the agent estimates multiple agency parameters as non-zero.

Figure 5.9c shows that the increased noise on the states, ω , increases the variance of both the estimates of the agency parameter b_1 with (\hat{b}_{i1}) and without (\hat{b}_1) prior. However, the EM estimation fluctuates more than the prior-less OLS estimation. Thus, in experiment 1C, the prior does not improve the estimation of agency parameter \hat{b}_{i1} .



(a) EM estimated \hat{b}_j over time (with prior).

(b) OLS estimated \hat{b}_j over time (without prior).



(c) EM and OLS estimated \hat{b}_2 over time.

Figure 5.9. These plots, like Figure 5.8, show that prior promotes a correct SoA in a dynamic environment with noise (experiment 1C). The difference between experiment 1B and 1C is the increased noise on the states, ω . Figure 5.9c zooms in to the OLS and EM estimation of b_1 to show that the increase in state noise increases the variance in the estimation of b_1 , which decreases the SoA.

SoA throughout the trial. Though prior does not improve the estimation of the agency parameter \hat{b}_{i1} when compared to the OLS estimation in unambiguous environments, the prior should improve the agent's SoA over time as the prior constantly updates. The agent starts with an incorrect estimation of the unknown parameters, $\theta^{(0)}$ which through the EM iterations moves towards a more correct estimation.

Figure 5.10 plots the two-dimensional Gaussian distributions for each k cluster in the dynamic environment (experiment 1B) on top of the observations from that time window, δ_N . Box 5.1 gives a more detailed explanation of the meaning of the contours and data points in the plot. Figure 5.10a plots the distributions with the initial conditions, $\theta^{(0)}$, and the data points that the agent will encounter in the first δ_N time steps of the simulation. The observed data points from $j = 3$ and $j = 5$ are not Gaussian distributed because the agent does not the input force to mass m_2 and m_3 (as explained in Sec. 5.4.1). When comparing Figure 5.10a to 5.10b, the agent correctly clusters the data into agency and five no-agency clusters. Moreover, the agent has moved the green contours towards the observations for $j = 1$ meaning that the parameters of the bivariate distribution try to describe the unknown parameter that controls mass m_1 , i.e. \hat{b}_1 . This means that the agent's changes from an incorrect SoA at $t = 0$ s to a correct SoA at $t = 20$ s.

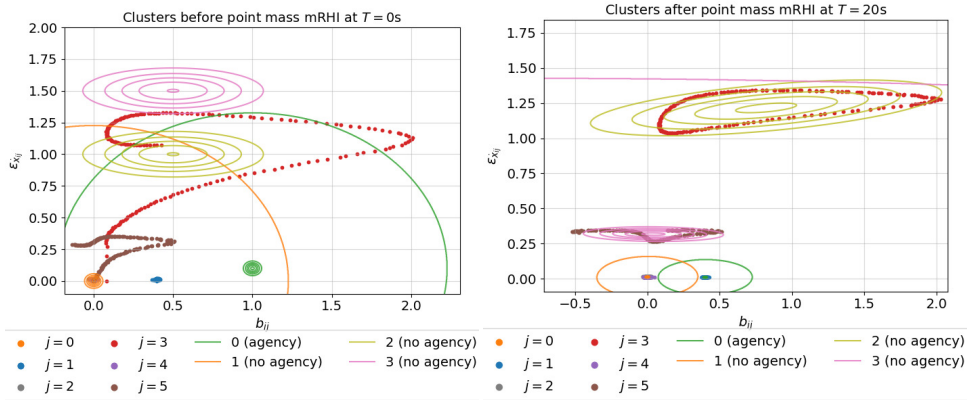
5

Box 5.1: EM: bivariate probability distributions describing observations

This box explains the figures of the clusters and the observations, such as Figure 5.10. These figures are a continuation of Figure 5.5. However, instead of plotting a hypothetical distribution on the data points, the agent estimates K probability distributions that best describe the data points depicted in the figures. The data points represent the two dimensions of the agent's observations (i.e. $\varepsilon_{\dot{x}_{ij}}$ and \hat{b}_{ij}). The contours represent a bivariate Gaussian probability distribution. The means, μ , define the center of the probability distribution. The covariance matrices, Σ define the spacing of the contours. As a result, the closer the contours are to each other, the smaller values of the covariance matrix. A data point closer to the center of the contours has the highest probability of being a result of the hidden cause that the cluster tries to describe.

In the figures of this thesis, the green contours describe the agency cluster whereas the other contours describe the no agency clusters. If the agent can correctly identify four different hidden causes, then the contours should overlay data points that are all a result of the same hidden cause, Z_{ij} . If the agent correctly identifies the observations from $j = 1$ as a result of its input force, F_{int} , then the green, agency contours should overlay the blue data points ($j = 1$).

Figure 5.11 plots the agent's estimation of the parameters of the clusters for the dynamic environment with noise (experiment 1C). A notable difference between the initial time window of experiment 1B (see Figure 5.10a) and this experiment in Figure 5.11a is that the data points of $j = 1$ have a wider variance in the both axes of the plot and a higher prediction error $\varepsilon_{\dot{x}_{ij}}$. This is a result of the increase in the uncertainty of the motion states, ω . Despite the increased error, the agent can separate the agency cluster from the no-agency cluster. For the final time window, Figure 5.11b shows that the agent's correct SoA over mass m_1 also decreased the prediction error, $\varepsilon_{\dot{x}_{i1}}$. Thus, the



(a) The initial time window of data and clusters before of the simulation ($T = 0s$) (b) The final time window of data and clusters at the end of the simulation ($T = 20s$).

Figure 5.10. This figure shows that for the dynamic environment (experiment 1B) the (EM) agent correctly clusters the data according to their hidden causes, despite the incorrect initial conditions of the parameters, $\theta^{(0)}$ describe the clusters in Figure a. The lines represent the probability density distributions of the agency and no-agency clusters. The dots represents the data for the prediction error, $\epsilon_{x_{ij}}$, and observed b_{ij} for each j . After the simulation, in Figure b, the agent correctly believes that the hidden cause of the blue points ($j = 1$), is its own force input F_{int} because the agency probability density distribution (green lines) tries to describe these data points. Here, the agent correctly identifies that it has agency over mass m_1 .


5

agent's changes from an incorrect SoA at $t = 0s$ to a correct SoA at $t = 20s$.

Table 5.3 compares the mean estimation of \hat{b}_{i1} over the entire simulation to the final time window of the simulation for both of the dynamic and unambiguous environments (experiment 1B and 1C). First, I compare the left columns to the right columns for both the agency accuracy and variance and find that with the EM algorithm, the estimation of \hat{b}_{i1} is more accurate and precise at the end of the simulation than over the entire simulation. The agent without a prior (with the OLS algorithm) does not significantly adjust its accuracy or variance of the estimation of parameter \hat{b}_1 . The relative lower accuracy of the EM algorithm over the entire simulation is due to the (incorrectly) guessed initial conditions, $\theta^{(0)}$. At the start of the simulation, the agent first has to adjust its parameters, $\theta^{(i)}$, to find more optimal parameters at the end of the simulation. This suggests that the prior improves the agent's SoA over time.

Second, with Table 5.3, I compare the accuracy and variance of the prior to the no prior algorithm. Over the entire simulation, the agent without prior has lower accuracy and variance than the agent with prior. This is highlighted by the green cell colors. However, for both experimental conditions at the end of the simulation, there is no significant difference between the agent with and without the prior. Thus, at the end of the unambiguous experimental conditions, the SoA of the agent with prior and without prior is approximately equally correct.

Comparing the effect of the prior on SoA to a less informative prior. In the M-step, I adjusted the EM algorithm to enforce the prior knowledge that the agent can only control



(a) The initial time window of data and clusters before of the simulation ($T = 0s$) (b) The final time window of data and clusters at the end of the simulation ($T = 20s$).

Figure 5.11. This figure is a continuation of Figure 5.10 and shows that in a dynamic environment with noise (experiment IC) the agent also correctly has a correct SoA at the end of the simulation Figure b, despite the incorrect clustering of data in the initial conditions in Figure a. A difference between this figure and Figure 5.10 is the increased uncertainty on the motion states, ω , which increases the error $\varepsilon_{\hat{x}_{ij}}$ and the variance of b_{ij} of mass m_1 (i.e. $j = 1$ and $j = 1$). Therefore, the probability distribution of the agency function has a higher mean, μ , in the prediction error dimension and a larger covariance, Σ , in the b_{ij} dimension than in experiment 1B.

Table 5.3. This table shows that on average for the 1000 trials the EM algorithm (with prior) improves agent's estimation of the agency parameter \hat{b}_1 with time. The OLS algorithm (no prior) is more accurate throughout the simulation but the EM algorithm the agent is more accurate in the final time window of the simulation.

Exp. nr.		Agency accuracy, $ \mu_{\hat{b}_1} - b_1 $		Agency variance, $\sigma_{\hat{b}_1}^2$	
		$t = [0s, 20s]$	$t = [18s, 20s]$	$t = [0s, 20s]$	$t = [18s, 20s]$
1B	No prior	3.08×10^{-6}	3.79×10^{-6}	1.50×10^{-7}	1.07×10^{-7}
	Prior	5.45×10^{-4}	2.43×10^{-6}	2.16×10^{-4}	1.11×10^{-6}
1C	No prior	9.73×10^{-5}	3.23×10^{-4}	1.61×10^{-5}	1.54×10^{-5}
	Prior	1.07×10^{-2}	2.31×10^{-4}	3.27×10^{-3}	4.55×10^{-5}

one mass (see eq. 5.18). To understand the effect of giving the agent this prior knowledge, I compare the agent's SoA with this calculation of the mixture proportion (eq. 5.18) to the SoA of the agent with regular calculation of the mixture proportion, π_k , (as defined by [98], see eq. 5.13). I term the latter calculation of the mixture proportion, the "less informative prior"

With this less informative prior, π_k , the agent's correct SoA over mass m_1 decreases from 99.9% to 91.9% for experiment 1B. With the introduction of noise in experiment 1C, a larger decrease in SoA over mass m_1 happens, from 98.1% to 68.1% for experiment 1C. To illustrate, Figure 5.12 plots the agent's SoA for one trial of the dynamic environment with noise (experiment 1C) with the regular calculation of the mixture proportion and the prior knowledge calculation of the mixture proportion. The left plot (Figure 5.12b) shows that the agent without the prior knowledge needs around 10 s to gain a constant SoA over the set of observations from $j = 1$. The agent that knows it must control one of the masses (Figure 5.12a), only has an incorrect SoA in the first few time steps. Thus, the agent with the adjusted prior has an improved SoA because generally, the agent has a complete SoA over the observations from mass m_1 from the first time step.




Figure 5.12. This figure shows that an agent using the EM algorithm without the prior knowledge that it can control one mass has a less correct SoA over the observations from $j = 1$. Probability $P(Z_{ij} = k | j = \{0, \dots, n-1\})$ describes the agent's SoA over the observations produced by its own input force, where $Z_{ij} = k$ depends on the k cluster that the agent believes is the hidden cause Z_{ij} .

EXPERIMENT 2: INCORRECT SoA IN AMBIGUOUS ENVIRONMENTS

For the ambiguous environments (experiment 2), Table 5.4 shows the percentage of time that the agent has a SoA over each state j of the estimated $\hat{\mathbf{B}}$ out of the 1000 20 s trials per experimental condition. To determine if the prior induces an incorrect SoA, I compare the EM algorithm to the OLS algorithm without prior.

For the synchronous ambiguous environments (experiment 2A and 2B), the agent only has a correct SoA over mass m_1 by estimating the agency parameter \hat{b}_1 as non-zero 6.46% and 8.97% of the time. Most of the remaining time the agent has an incorrect SoA over mass m_2 with \hat{b}_3 . Shifting the input force of mass m_2 by 500 ms defines the asynchronous environment (experiment 2C), which results in the agent's correct SoA over mass m_1 increasing to 68.6%. All three experiments contrasts with the no prior OLS algorithm which has a SoA over \hat{b}_1 100% of the time. Thus, the prior decreases the agent's correct SoA of mass m_1 .

Table 5.4. This table shows that for ambiguous environments the agent with a prior generally has an incorrect SoA over mass m_2 in the synchronous conditions (2A and 2B) and more correct SoA over mass m_1 in the asynchronous condition (2C). The table is a continuation of Table 5.2 but for the ambiguous environments defined by experiment 2 (see Sec. 3.3).

Exp. nr.	% of time that $\hat{b}_0 > 0$	% of time that $\hat{b}_1 > 0$	% of time that $\hat{b}_2 > 0$	% of time that $\hat{b}_3 > 0$	% of time that $\hat{b}_4 > 0$	% of time that $\hat{b}_5 > 0$
<i>With prior (EM)</i>						
2A	0.0415	6.46	0.00	88.7	0.00	0.296
2B	0.0507	8.97	0.00	95.5	0.00	0.422
2C	0.621	68.6	0.00	18.4	0.00	9.87
<i>Without prior (OLS)</i>						
2A,2B,2C	0.00	100	0.00	100	0.00	92.8

Overall the prior promotes an incorrect SoA over mass m_2 . As it is clear the prior promotes an incorrect SoA and for brevity, I will not discuss the figures of the clusters (e.g. Figure D.3) for the unambiguous environments. However, these figures help explain the differences in the SoA time between the experimental conditions of an incorrect SoA. See Appendix D.2.2 for this explanation. I will show the results of the estimation of the agency matrix, $\hat{\mathbf{B}}$, and the effect of adjusting the mixture proportion, π_k to further understand the effect of this incorrect SoA.

Estimation of the agent matrix, $\hat{\mathbf{B}}$. Aside from the percentage of time of SoA, four other performance measures indicate whether the agent's SoA has improved (see Sec. 3.3.2.¹¹ I will show the effect of the incorrect SoA as a result of the prior on the estimations of the agency matrix $\hat{\mathbf{B}}$.

Figure 5.13a plots the EM estimation of the agency matrix, $\hat{\mathbf{B}}$, for the synchronous ambiguous environment (experiment 2A). Though the agent has an incorrect SoA over mass m_2 in the final time step, the agent takes longer to gain agency over one mass than in the unambiguous environments (experiment 1). In the first approximately 5 s the agent has a SoA over the observations from both mass m_1 and m_2 . Without a prior in Figure 5.13b, the agent's estimation of $\hat{\mathbf{B}}$ does not take time to converge to an estimation for $\hat{\mathbf{B}}$ because the prior does not try to "steer" the agent away from its incorrect agency estimations. In the end, the agent with the prior sets the \hat{b}_{i1} to 0 because this explains the sensory evidence while reducing the agent's prediction error. For this condition, the prior steers the agent away from the correct conclusion.

In the next experimental condition, the mass of mass m_2 decreases. Figure 5.14 plots the estimation of the j parameters \hat{b}_{ij} for both the agent using a prior and not using a prior in experiment 2B. The agent with a prior (see Figure 5.14a) set parameter \hat{b}_{i1} to 0 at around 1 s rather than the 6 s in experiment 2A. Despite the decrease of mass m_2 , the agent still incorrectly concludes agency over mass m_2 with parameter \hat{b}_{i1} at the end of the simulation. Without a prior the agent has SoA over all masses (see Figure 5.14b), in-

¹¹See Appendix D.2.1 for the numerical values of the four performance measures per experimental condition of experiment 2.

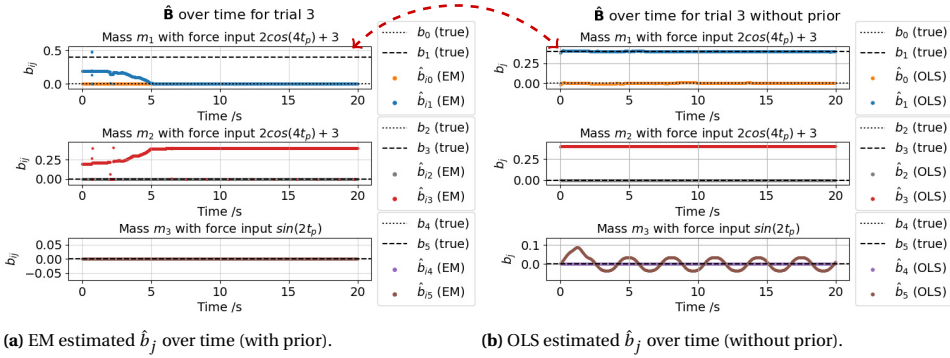


Figure 5.13. This figure shows that the prior steers the agent away from the correct SoA for a synchronous, ambiguous environment. Like Figure 5.9, this figure plots the estimation of the SoA parameter \hat{b}_j for the agent with a prior and without a prior. The only difference between experiment 2A and experiment 1C is that the sinusoidal force input to mass m_1 and m_2 is now synchronous.

cluding the correct mass m_1 . In this ambiguous environment, the prior steers the agent away from the correct SoA.

5

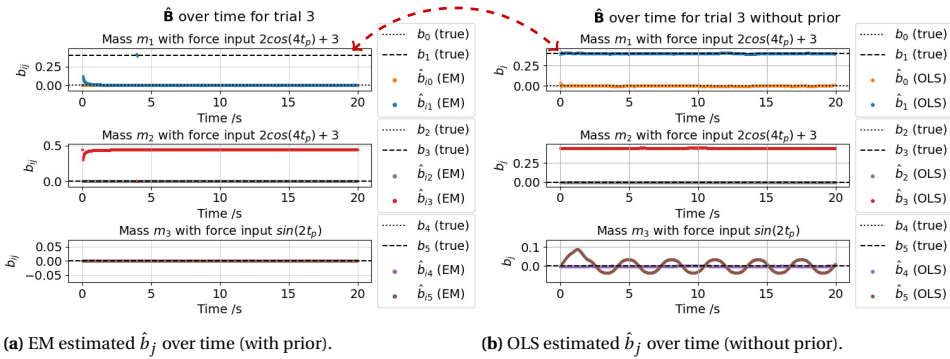


Figure 5.14. This figure shows that in this ambiguous environment (experiment 2B), the prior induces an incorrect SoA over mass m_2 with parameter \hat{b}_{13} . The difference between experiment 2A and experiment 2B is a 10% decrease in mass m_2 with respect to mass m_1 . The decreased mass increases the agent's estimation of \hat{b}_{13} and decreases time during which $\hat{b}_1 = 0$ is.

The final experiment condition of the ambiguous environments (experiment 2C) introduces a time lag of 500 ms in the force input of mass m_2 w.r.t. to the force input of mass m_1 , i.e. F_{int} . Figure 5.15a shows that the agent with the prior estimates the odd j elements of $\hat{\mathbf{B}}$ as non-zero at some point of the simulation. The agent without prior estimates \hat{b}_3 with a varying function to compensate for the 500 ms lag (Figure 5.15b). The reason for the agent with the prior eventually gaining a correct SoA over mass m_1 is because of the higher variance of the observations of \hat{b}_{13} . The agent knows that a higher variance is proportional to a lower SoA because this indicates an imprecise sensory signal. Therefore, the prior motivates the agent to gain a correct SoA over mass m_1 during the trial.

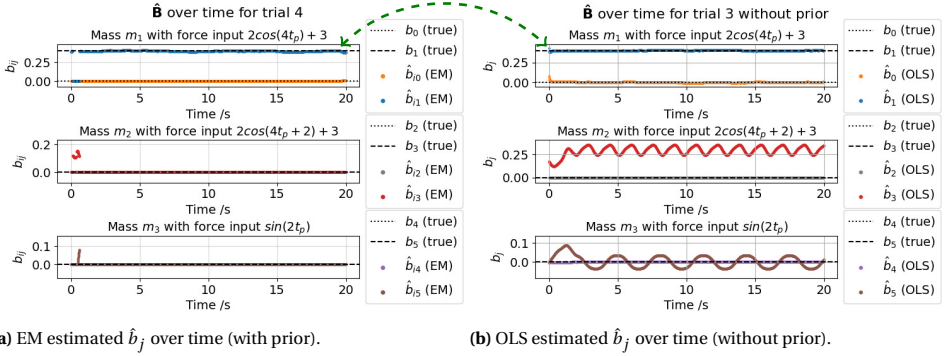
(a) EM estimated \hat{b}_j over time (with prior).(b) OLS estimated \hat{b}_j over time (without prior).

Figure 5.15. This figure shows that in this ambiguous environment (experiment 2C), the prior induces a correct SoA over mass m_1 with parameter \hat{b}_{j1} . The difference between experiment 2A and experiment 2C is a 0.5 s shift of the force input of mass m_2 with respect to mass m_1 . In the initial 5 s this makes the agent unsure over which mass it has a SoA. However, after 5 s the agent correctly has SoA over mass m_1 .

5

Comparing the effect of the prior on SoA to a less informative prior. To show the effect of adjusting the mixture proportion, π_k , on the agent's SoA over time in ambiguous environments, I compare the results of the EM algorithm to the results with a less informative prior (see eq. 5.13).

For the synchronous ambiguous environment (experiment 2A), the changing the prior, π_k , to the less informative prior, decreases the agent's incorrect SoA over mass m_2 from 88.7% to 42.5%. However, the agent's correct SoA over mass m_1 remains the same at around 8%. Figure 5.16 compares the agent's SoA during one trial for an agent with a mixture proportion, π_k , that is updated according to the normal EM (Figure 5.16b) and the SoA over an agent that knows it can control exactly one mass (Figure 5.16a). Figure 5.16a shows that initially the agent is unsure about which set of observations it has agency over. As the agent collects more observations and updates its prior, π_k , the agent becomes more certain that it has agency over mass m_2 . This shows that the constraint on the mixture proportion increases the agent's incorrect SoA over mass m_2 through the observations from $j = 3$.

For the synchronous ambiguous environment with lower mass (experiment 2B), the change in the calculation of the prior, π_k , decreased the agent's incorrect SoA over mass m_2 from 95.5% to 40.2%. Also, the agent's correct SoA over mass m_1 decreases from around 8.97% to 2.10%. These results are similar to the results from the previous experimental condition. However, in experiment 2B, the prior seems to motivate even bigger differences because Figure 5.17a shows that with prior the agent almost immediately has a SoA over mass m_2 . This contrasts with Figure 5.17b where the agent only shortly has SoA over the observations $j = 3$ from mass m_2 . Thus, the constraint on the prior, π_k , increases the agent's incorrect SoA.

For the asynchronous ambiguous environment (experiment 2C), changing in the the prior, π_k , to the less informative prior, the agent's correct SoA over mass m_1 decreases from 68.6% to 54.9%. Also, the agent's incorrect SoA over mass m_2 remains decreases from around 18.4% to 4.05%. In trial 3 of experiment 2C, the agent with the prior knowl-




Figure 5.16. This figure shows that in an ambiguous environment (experiment 2A) with the EM algorithm, the agent without the prior knowledge that it can control one mass has a less incorrect SoA over the observations from $j = 3$ but also less correct SoA over the observations from $j = 1$. The difference in the figures is an adjustment of the calculation of the prior, π_k , from eq. 5.18 in Figure a to eq. 5.13 in Figure b. Probability $P(Z_{ij} = k | j = \{0, \dots, n - 1\})$ describes the agent's SoA over the observations produced by its own input force, where $Z_{ij} = k$ depends on the k cluster that the agent believes is the hidden cause Z_{ij} .

5

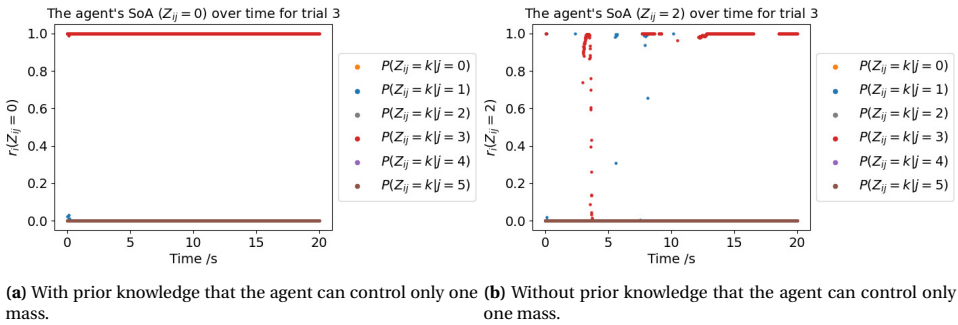



Figure 5.17. This figure shows that in an ambiguous environment (experiment 2B) with the EM algorithm, the agent without the prior knowledge that it can control one mass has a less incorrect SoA over the observations from $j = 3$. This plot is a continuation of Figure 5.16 but in experiment 2B the mass of m_2 decreases by 0.25 kg w.r.t. to the mass of m_1 .

edge that it controls one mass has a correct SoA over mass m_1 100% of the time (see Figure 5.18a). Conversely, Figure 5.18b shows that the agent without this prior knowledge needs more time to infer that it has a SoA over mass m_1 . This suggests that the constraint on the mixture proportion, i.e. the extra prior knowledge, improves the agent's SoA by allowing it to more quickly come to the correct conclusion.



(a) With prior knowledge that the agent can control only one mass. (b) Without prior knowledge that the agent can control only one mass.

5

Figure 5.18. This figure shows that in an asynchronous, ambiguous environment (experiment 2C) with the EM algorithm, the agent without the prior knowledge has a less correct SoA over the observations from $j = 1$. This plot is a continuation of Figure 5.16 but in experiment 2C force input to mass m_2 has a time-lag of 500 ms w.r.t. to the force input to mass of m_1 .

5.5. DISCUSSION: EVALUATION OF THE POINT MASS MRHI

IN this chapter, I presented the EM algorithm as a method for achieving SoA in an agent in the point mass mRHI. First, I will discuss how this EM algorithm incorporates the computational principles of a SoA. Second, I will discuss whether the prior improved SoA in ambiguous environments. Next, I will discuss whether the prior induced incorrect SoA in the point mass mRHI and how this connects to the point mass mRHI and the human mRHI experiment. Finally, I will discuss the computational issues of the EM algorithm.

5.5.1. THREE SOA COMPUTATIONAL PRINCIPLES

In the EM implementation of the point mass mRHI, SoA works according to the three computational principles described in Sec. 2.2.2. This formalizes this EM algorithm as a model of SoA in the point mass mRHI environment. However, though the point mass mRHI shows how prior knowledge, the prediction error, and sensory evidence interact to form a SoA, this analogy does not capture the complexity of all of the processes underlying a SoA in humans.

The first computational principle correlates the strength of the prior belief that action causes outcome to the strength of the SoA. As SoA, or responsibility r_{ia} , is calculated with the prior π_k , it is evident that when π_k approaches 1 (stronger causality) then the SoA also increases. This aligns with the SoA computational model by Legaspi and Toyoizumi[20], where causality increases the intentional binding effects, which are an indication of SoA. However, in their model, the causal prior is set before-hand, whereas, in the point mass mRHI, the agent determines the prior through the EM algorithm.

Second, a SoA is correlated to a lower prediction error. This computational principle is reflected in the asynchronous ambiguous environment (experiment 2C) because the difference in prediction error for the observations of mass m_1 and m_2 is smaller than in the other experimental conditions. With the OLS algorithm in Chapter 4, I showed that the agent generally believed it had SoA over mass m_2 despite the larger prediction error. However, with the EM algorithm, the agent has a correct SoA over mass m_1 around 70% of the time. This indicates that the agent tries to lower its prediction error through its SoA.

The third computational principle correlates the strength of SoA to the precision of the sensory signals. The only difference between experiment 1B to 1C is a decrease in the precision of the sensory signals of mass m_1 . This results in a slight decrease in the agent's correct SoA over mass m_1 from around 100% to 98%. Therefore, the sensory signals influence the strength of the SoA in the point mass mRHI.

5.5.2. IMPROVED SOA IN UNAMBIGUOUS ENVIRONMENTS

To summarize the results of the unambiguous environments (experiment 1), I list the main findings from the results on whether the agent's prior improves SoA:

- The prior does *not* improve the agent's SoA over the mass which the agent should have a SoA over, i.e. mass m_1 .
- The prior *does* improve the agent's SoA over the mass which the agent should not have a SoA over, i.e. mass m_2 and m_3 .
- The prior *does* improve the agent's SoA during the simulation.
- The prior *does* improve the agent's SoA when compared to a less informative prior.

The lack of improvement in the agent's SoA over mass m_1 (see Table 5.2) is probably a consequence of the simplicity of the point mass mRHI. The agent using OLS to estimate the agency matrix, \mathbf{B} , does not learn any parameters but only makes statistical inferences to calculate the matrix. In the point mass mRHI, letting the data "speak for itself" is sufficient to outperform an agent with prior knowledge for the prediction of the states of mass m_1 . Experiment 1 does not utilize the computational advantage of the prior (see Sec. 2.1.4). In the asynchronous ambiguous environment (experiment 2C), the environment became slightly more complex because the sensory evidence was less apparent for the agent to choose its SoA. With the prior, the agent eventually gains a correct SoA over mass m_1 . Therefore, I predict that the prior improves the agent's SoA in a more complex environment when the agent needs to prior to make inferences about its SoA that go beyond the data.

Nonetheless, the prior does improve the agent's SoA over the mass which the agent should not have a SoA over (see Table 5.2) because the sensory evidence does not inform the agent that it only has agency over one mass. The prior allows the agent to separate between its internal model and the external states of the world. An agent that understands the difference between objects that it can and cannot control is more important when the agent has self-defined input forces.

In addition, the prior does improve the agent's SoA and agency estimation, $\hat{\mathbf{B}}$, over time (see Figure 5.8 and 5.9), which is a result of two aspects of the EM algorithm. First, with Bayesian inference, the prior uses the sensory evidence that the agent collects to update the agent's SoA at a future time step. Therefore, collecting more data improves the prior, which improves the SoA. Second, the prior is needed to improve the incorrect initial conditions of the unknown parameters, $\theta^{(0)}$. These initial parameters negatively influence the estimations of the agency matrix $\hat{\mathbf{B}}$ at the beginning of the simulation. Accordingly, if the agent would know the correct parameters of the agency clusters before the simulation starts, then the agent's estimation accuracy of parameter \hat{b}_{i1} would increase. Also, an agent with correct initial conditions is more comparable to participants in the mRHI because the participants have SoA over their hand before the experiment. However, knowing the initial conditions is less relevant for an artificial agent because in a more complicated system the parameters of the robot are often unknown. This highlights an obstacle to reproducing human-like behavior in an agent. Therefore, the agent needs a Bayesian prior that updates with time that can eventually (when enough data is collected) improve the agent's SoA.

The final way in which the prior knowledge improves is in comparison to the less informative prior (see Figure 5.12). This comparison is important to show that adjusting the mixture proportion, π_k , leads to more accurate SoA because, again, this gives the agent extra knowledge about the environment that it cannot infer from its sensory data. Adding more information in the prior improves the agent's SoA.

5.5.3. INCORRECT SOA IN AMBIGUOUS ENVIRONMENTS

In the EM algorithm, the prior induces an incorrect SoA when the environment is ambiguous. The agent in the synchronous ambiguous environments (experiment 2A and 2B) has the most incorrect SoA over mass m_2 with \hat{b}_3 (around 89% and 96% of the time). This is reflective of Kalckert and Ehrsson's results in the synchronous active mRHI condition, where the agent have a SoA over the rubber-hand [31][31]. Also, different from the OLS algorithm, the agent does switch its SoA during a trial. This aligns with the human mRHI because the humans adjust their SoA to include the rubber hand [31][32].

The ambiguity of the environment decreases by introducing a time lag in the force input to mass m_2 (experiment 2C), which in turn decreases the agent's incorrect SoA over mass m_2 to approximately 18% of the time. These results appear to align with the active asynchronous condition of the human mRHI experiment [32] and with a continuous hand movement experiment [88]. In both experiments, introducing a time lag significantly reduced the SoA over the participants over their hand. This suggests that the point mass mRHI is a sufficient initial analogy for the human mRHI experiment.

To further compare a SoA with the mRHI in artificial agents, some more aspects of SoA in humans should be quantified. The question of quantification of human behavior is a problem during different moments of the simulation, such as:

- When the participants place their hands under the table, how much extra noise do participants perceive in the estimation of their hand location?
- How much (if any) SoA does a subject experience over their own hand while experiencing SoA over the rubber-hand?
- How does the SoA shift to the rubber-hand? Is this a smooth transition or do the participants suddenly experience incorrect SoA when subject to the synchronous active mRHI experimental condition?

These questions are ideas for future work on SoA in humans.

5.5.4. GENERAL ISSUES IN THE EM ALGORITHM

Till now, the discussion and results mainly focus on the connection of the EM algorithm to human SoA and the mRHI, but I would also like to highlight several computational issues of the EM algorithm. First, when the agent does not have SoA over a j set of observations, the corresponding prediction error, $\varepsilon_{\hat{x}}$, does not produce Gaussian data. The prediction error is then dependent on the noise and the difference between the incorrect prediction of \hat{x} and the observed \hat{x} , which is not Gaussian. The EM algorithm tries to approximate the prediction error with a Gaussian distribution. Though this is not theoretically correct, the prediction error data is only used to separate the different clusters

and the agent does not use the mean and variance of this dimension to estimate the agency parameter, $\hat{\mathbf{B}}$.

Second, I set the number of iterations per time step for the EM algorithm to an arbitrary maximum of 5. This maximum is only enforced in the initial time steps of the algorithm because later the agent seems to minimize the log-likelihood function within around one to two iterations. However, I do not know whether this is comparable to human behavior. Moreover, if I would extend this simulation to the real world, then I do not know whether the iterations are fast enough for the robot to choose the mass that it has agency over at each time step. Online variations of the EM algorithm exist that might be more suitable for this application (e.g. [102]).

Third, the number of clusters of the EM algorithm is fixed meaning that it cannot change if the number of hidden causes changes. The EM algorithm requires that the number of clusters is set *a priori*. A quick fix to this problem is setting the number of hidden causes to a high number. Then the EM algorithm leaves some of the clusters as empty, meaning it does not describe any of the observations. However, for a more accurate representation of the hidden causes of the environment, the agent should be able to adjust the number of clusters, K , during the simulation.

Next, the results of the EM algorithm are sensitive to the initial conditions. Though this does demonstrate the importance of prior knowledge to make correct inferences when the initial values are set to numbers that are far from the true values the agent can no longer cluster the incoming data. This is different from humans who observe and learn from their environments to try to recover from their incorrect prior knowledge.

Finally, a disadvantage of the EM algorithm is that it is not fully Bayesian because EM computes the point estimates of posterior distributions of the parameters, $\theta^{(t)}$. It would be advantageous to describe the prior, π_k , with some distribution rather than the uniform prior which currently describes the agent's prior knowledge. The current prior in this EM algorithm is a uniform prior per j observation as it only describes a probability that the agent has agency for each j observation. As stated, more information in the prior could lead to more improvement of the SoA. By describing the prior with a distribution, the agent can encode the confidence in their prior by encoding a precision. Also, a fully Bayesian approach might be a more suitable approach to reproducing the human mRHI experiment and SoA. I will discuss this in the recommendations for the thesis in Sec. 6.2.

5.6. CONCLUSION

THIS chapter answers the research question: *What is the effect of the prior on a sense of agency in the point mass moving rubber-hand illusion? Is this effect comparable to a sense of agency in human?* The prior, or the mixture proportion of the agency cluster π_a , defines the probability that data is a result of the agent's input force F_{int} . The prior that the agent can only have agency over one mass improves some aspects of the agent's correct SoA over mass m_1 in unambiguous environments (experiment 1) while promoting an incorrect SoA over mass m_2 in ambiguous environments (experiment 2). For a prior to further improve SoA in experiment 1, the environment that the agent is in should have more limited or noisy data. The results from experiment 2 indicate that an incorrect SoA could be comparable to the illusory SoA in the human mRHI experiment.

In summary, this chapter formalizes SoA with the EM algorithm where the artificial agent's SoA is the responsibility term, $r_i(Z_{ij} = k)$. The advantage of the EM algorithm is that it inherently clusters data according to the hidden causes of the data. This is consistent with the causal nature of SoA. Also, as the agent performs Bayesian inference the SoA can be quantified as a probability. This aligns with Moore et al.'s Bayesian cue integration framework [19]. The results for both the static (in Sec. 5.4) and dynamic environments (in Sec. 5.5) show that the artificial SoA is dependent on the three principles that literature believes brings about SoA (see Sec. 2.2.2): the precision of the sensory signals, the prediction error of the motion states, and the causality between action and outcome. By simplifying Kalckert and Ehrsson's mRHI [31][32] to a point mass mRHI, I find that the prior can improve the SoA in unambiguous environments (experiment 1) but also can promote incorrect conclusions in ambiguous environments (experiment 2).

6

CONCLUSIONS

This is the concluding chapter of this thesis. First, this chapter summarizes the answers to the three research questions presented in the introduction. Then this chapter presents several recommendations for further research on priors and *Sense of Agency* (SoA) in artificial agents.

6.1. SUMMARY

THE research goal of this thesis is as follows: *design an artificial agent with a SoA where the same prior knowledge (i) improves SoA in unambiguous environments and (ii) result in incorrect, illusory SoA in noisy, ambiguous environments.* An artificial agent in the “point mass mRHI” with a prior improved certain aspects of its SoA, but also experienced illusory SoA. The agent’s prior is that the agent can have agency over one out of three point masses.

6.1.1. ANSWERING THE RESEARCH QUESTIONS

To achieve the research goal, I answer three research questions (RQ) in this thesis.

RQ 1: What prior knowledge enables a sense of agency in humans in the moving rubber hand illusion? SoA is enabled by prior knowledge of the causal relationship between intended actions and sensory outcomes. When this predicted sensory signal aligns with the observed sensory signal (low prediction error) the SoA increases. Thus, prior knowledge in humans plays an important role in the main computational principles I identify as necessary for a SoA to emerge: (1) prior belief of causal action-outcomes, (2) magnitude of the prediction error, and (3) precision of sensory signals. Also, different research fields (statistics, cognitive science, neuroscience, and robotics) underline the prior as a possible influential tool in inferences. In Bayesian inference, the prior is responsible for allowing inferences that go beyond the data, but also inducing sub-optimal behavior such as illusions. Thus, the prior causal belief of action-outcome should both improve SoA and induce illusory SoA depending on the environment.

RQ 2: In what artificial environment, similar to the moving rubber hand illusion, should an agent experience a correct and incorrect sense of agency? Like the human mRHI experiment (i.e. [31] and [32]), the “point mass mRHI” tests both correct and incorrect SoA in an artificial agent. In this simulation, the agent has agency over (i.e. applies a force on) one of three point masses that are each connected to spring-damper systems. The agent’s prior knowledge is that it knows there is a causal relationship between its force input and the states of *one* mass but does not know which mass. In unambiguous environments (defined by high sensory precision and asynchronous movements of the masses), the agent should experience a correct SoA. According to the Kalckert and Ehrsson’s mRHI experiments [31][32], when two masses move synchronously *and* there is an increased noise on the mass the agent should experience an incorrect (illusory) SoA. To gain agency over a point mass, the agent estimates the agency matrix $\hat{\mathbf{B}}$ from its state-space equation with a parameter estimator.

RQ 3: What is the effect of the prior on a sense of agency in the point mass moving rubber hand illusion? Is this effect comparable to a sense of agency in human? The prior forces the agent to distinguish between internal and external states by predicting the hidden causes of the sensory signals. In unambiguous environments, the prior motivates the agent towards a correct SoA and in some aspects improves the agent’s SoA. Conversely, in ambiguous environments, the prior induces an illusory SoA over the in-

correct. These findings are comparable to the human mRHI experiment (i.e. [31] and [32]).

I found this answer in two steps with two parameter estimators. The first parameter estimator, OLS, is as a simple system identification algorithm, which allowed an initial exploration of the simulation. I termed this simulation the “simplest point mass mRHI”. Though I concluded that the OLS algorithm is not sufficient to capture human SoA and mRHI, this preliminary implementation reveals the necessary properties for the agent to experience a more realistic SoA. The main finding is that SoA should be described by a continuous posterior probability rather than a discrete state. Second, I used EM in the point mass mRHI to estimate the agent’s agency parameter and quantify SoA. EM is a popular (partially) Bayesian inference algorithm. Aside from the computational advantages of EM, EM is suitable for the point mass mRHI because the steps correspond to Bayesian inference and learning in humans. Also, similar to how humans cluster the causes of sensations into “a result of my own actions” or “a result of external factors”, the algorithm inherently clusters observations according to predictions about their hidden causes (i.e. agency or no agency). Though the agent using EM does not outperform the agent using OLS, the agent’s SoA is more similar to human SoA. Moreover, I find that the prior improves the agent’s SoA during the simulation, and setting a correct prior improves the SoA. In the synchronous mRHI experimental condition, the prior motivates the agent to gain a SoA over the “rubber-hand”.

6.1.2. ADDRESSING THE RESEARCH CHALLENGES OF PRIORS

With the research goal formulated in this thesis, I addressed two research challenges of a different nature formulated in the introduction. The first challenge is the general challenge of Bayesian inference that priors can generate incorrect conclusions that might cause (fatal) mistakes in an artificial agent. This research showed that the same prior that motivates a correct SoA can lead to illusory SoA over the incorrect object in ambiguous environments. The second challenge, for designers of artificial agents, is to translate subjective prior beliefs into a mathematical prior. I had to acquire a deeper understanding of SoA before I could formulate a quantitative prior that reflects the prior knowledge humans use for a SoA. Still, the prior in the point mass mRHI is a simplification of human prior knowledge. Also, the prior of this agent might be hard to generalize to other more complex artificial agents.

6.2. RECOMMENDATIONS

IN this section, I will describe recommendations for further research. These recommendations can be grouped into three categories: direct extensions for the “point mass mRHI”, future directions for research on priors in artificial agents, and future directions for SoA in artificial agents.

6.2.1. EXTENSIONS OF THE POINT MASS MRHI

Extensions of the point mass mRHI use the current mathematical model as defined in Chapter 3 to research further aspects of SoA or priors that are outside of the scope of this thesis.

COMPARISON OF THE POINT MASS MRHI TO A HUMAN EXPERIMENT

The point mass mRHI is presented as a simplification of the human mRHI experiment from Kalckert and Ehrsson [31][32]. However, to increase the similarity between the point mass mRHI and the human experience of SoA, future research could perform a point mass mRHI for humans. For the human point mass mRHI, the experiment could consist of participants viewing three point masses attached to spring-damper systems on a computer screen. By exerting a “force” using a joystick, the participants can control one of the three masses. The participant’s goal is to minimize the prediction error between the mass that they can control and some target state (which could be fixed or moving). Similar to the simulation in this thesis, the experimental design can include two independent variables. First, whether the participants are given the prior knowledge that they can control only one of the three masses with their force (but do not know which mass). The second independent variable is the simulation environment which could be either unambiguous or ambiguous (in this thesis experiment 1 and experiment 2). To measure the participant’s SoA, the experiment could use a questionnaire and/or temporal binding.

The advantage of comparing such an experiment to the point mass mRHI is that it removes the prior knowledge that participants have about agency over of their arm in the human mRHI experiment. If the effect of the prior and the different simulation environments on the results are similar in both human and artificial agent point mass mRHI, then this would further support the point mass mRHI as the simplest mRHI experiment.

CHANGING THE SYSTEM IDENTIFICATION ALGORITHM

The EM algorithm has some general limitations as a parameter estimator, but more importantly for this thesis EM has some specific limitations in the context of priors and SoA (see Sec. 5.5). *Variational Bayes* (VB) might solve these limitations. Variational methods provide an approximation for solving intractable integrals that arise in Bayesian inference problems (see Sec. 2.1) [71], which make it popular in Machine Learning. EM and VB are similar in that both use an alternative iterative procedure to converge to the optimal parameter values. Also, both calculate the posterior distribution of a set of latent variables, \mathbf{Z} , that determine the hidden cause of an observations (i.e. agency or no-agency) [98][103]. In the context of this thesis, the most important difference between the two is that EM computes point estimates of posterior distributions of the parameters, $\theta^{(t)} = (\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$, whereas VB fits distributions to these parameters and then finds

the parameters of this distribution, i.e. the hyperparameters. This means the unknown parameters, θ , also become latent variables.

The addition of distributions over the parameters in VB is advantageous for the point mass mRHI because of two reasons. First, all estimations of the variables of the algorithm become Bayesian, which is an advantage because Bayes models the uncertainty in the estimation of certain parameters. Second, in the context of this thesis, the hyperparameters can better encode the prior knowledge the agent has about the parameters. Specifically, by setting the hyperparameters of the “agency” cluster. In the EM algorithm, the mixture proportion (i.e. prior belief of agency) is a flat prior. With VB, I can define a distribution for the prior. Also, I could set the distribution hyperparameter of the mean parameter is wide, while the distribution of the variance parameter is narrow. This captures that the agent knows that it has agency over one of the point masses but does not know how it can control this mass. Though this may improve the point mass mRHI as an analogy of the human mRHI, it also makes the simulation more complicated which may defeat the purpose of the simple analogy.

THE EFFECT OF ADJUSTING THE ENVIRONMENT

Another interesting (minor) extension of the point mass mRHI is adjusting the environment during one trial of the simulation. The agent’s environment would change from ambiguous to unambiguous, or vice versa. This extension researches the effect of the agent’s strengthened SoA over one mas, which becomes the agent’s prior knowledge when it enters a new environment. I hypothesize that if the agent has SoA over a mass in one environment, then the agent will take longer to experience a SoA over another mass in the second environment compared to when the agent is directly exposed to the second environment. This is a result of the agent already having learned *how* to describe its SoA over a mass, i.e. having a stronger prior knowledge.

6.2.2. FUTURE DIRECTIONS FOR PRIORS IN ARTIFICIAL AGENTS

Empirical priors are an interesting future method for formulating the agent’s prior. Empirical priors are priors that are learned from data. This might seem to conflict with the advantages of priors discussed in this thesis, but are beneficial (i) when trying to build a cognitive model of human behavior and (ii) when implementing Bayesian hierarchical models. Cognitive scientists often try to reproduce results from an empirical study using their cognitive model. This way they can understand the underlying principles of the behavior researched in the study. It is common for cognitive models to use vague or non-informative priors [33] because defining priors is challenging. To illustrate the advantage of empirical priors, Schürmann et al. [54] use empirical data from the rubber foot illusion to define an informative prior (see Sec. 2.1). The results of their model are more comparison to results from human experiments than when using non-informative prior. Though the goal of the current point mass mRHI was not to as a cognitive model, formulating empirical priors will increase the similarity between the point mass mRHI and human mRHI.

The second advantage of empirical priors is their implementation within hierarchical models. In hierarchical models with empirical priors, estimates at one level act as priors on a lower level while the priors are also updated through inference [59]. Design-

ers of artificial agents can then only define the top-level prior, which represents the most abstract level of information about the world, or allow the hierarchical system to define top-level priors. The lower-level priors are updated by empirical evidence. Moreover, a hierarchical model incorporates causality through the connections between the levels of hierarchy [59]. This can encode the SoA prior belief about the causal relationship between action and outcome in the point mass mRHI. Recall, Kahl and Kopp[21] model SoA with a three-level hierarchy, where motor beliefs are abstractions for motor actions (see Sec. 2.2). This agrees with much of SoA research highlights the interaction of higher-level prior beliefs about actions and outcomes and lower-level sensory signals [104][72].

6.2.3. FUTURE DIRECTIONS FOR SOA IN ARTIFICIAL AGENTS

Finally, I have two recommendations for researching SoA in artificial agents that do not involve priors.

THE ROLE OF ACTION IN MINIMIZING THE PREDICTION ERROR

A promising area of future research for SoA in artificial agents is to use the action as a way to minimize the prediction error and, therefore, achieve a SoA. Recently, Perrykkad et al.[72] researched how humans employ actions to minimize prediction error in (uncertain) environments. An interesting finding in the context of artificial agents is that when the participants report a SoA, there was more prediction error minimization. This metric could be used to quantify SoA in an artificial agent. Allowing an agent to act on an environment to minimize its prediction error is more challenging than having an agent perceive and make choices based on perception (as was done in this thesis).

ACTIVE INFERENCE AS A MODEL FOR SOA

Many of the previously mentioned interesting directions for future research, i.e. VB, empirical priors, hierarchical models, and the role of action in minimizing the prediction error, are all a part of Friston's Active Inference [105][22]. Active Inference provides another interesting way to research priors and SoA in an artificial agent. Friston formulates Active Inference as a model for SoA [22] but does not test whether an Active Inference artificial agent acts similar to humans in SoA experiments. Robotics research with Active Inference shows promising results for using prediction error minimization to bring a robot arm to a target state [106][107]. Nonetheless, Active Inference implementations in artificial agents and robots are in their early stages. Therefore, it might take some time before the implementation of Active Inference in robots has developed enough to research SoA.

As a concluding remark, I would like to leave the reader with a question. Though a benefit of artificial agents as a research tool for human cognition is the access to all of the internal states, when can we conclude that a robot is experiencing something subjective like SoA?

BIBLIOGRAPHY

- [1] J. B. Tenenbaum, C. Kemp, T. L. Griffiths, and N. D. Goodman, “How to Grow a Mind: Statistics, Structure, and Abstraction”, *Science*, vol. 331, no. 6022, pp. 1279–1285, Mar. 2011, ISSN: 0036-8075. DOI: [10.1126/science.1192788](https://doi.org/10.1126/science.1192788). [Online]. Available: <https://www.sciencemag.org/lookup/doi/10.1126/science.1192788>.
- [2] T. L. Griffiths, N. Chater, C. Kemp, A. Perfors, and J. B. Tenenbaum, “Probabilistic models of cognition: exploring representations and inductive biases”, *Trends in Cognitive Sciences*, vol. 14, no. 8, pp. 357–364, Aug. 2010, ISSN: 13646613. DOI: [10.1016/j.tics.2010.05.004](https://doi.org/10.1016/j.tics.2010.05.004). [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1364661310001129>.
- [3] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, “Building machines that learn and think like people”, *Behavioral and Brain Sciences*, vol. 40, 2017, ISSN: 14691825. DOI: [10.1017/S0140525X16001837](https://doi.org/10.1017/S0140525X16001837). [Online]. Available: <http://cims.nyu.edu/~brenden/http://www.mit.edu/~tomeru/http://web.mit.edu/cocosci/josh.htmlhttp://gershmanlab.webfactual.com/index.html>.
- [4] D. O. Hebb, “The Organization of Behavior; A Neuropsychological Theory”, *The American Journal of Psychology*, 1949, ISSN: 00029556. DOI: [10.2307/1418888](https://doi.org/10.2307/1418888). [Online]. Available: <https://psycnet.apa.org/record/1950-02200-000>.
- [5] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 2018.
- [6] J. Schmidhuber, “Deep Learning in Neural Networks: An Overview”, *Neural Networks*, vol. 61, pp. 85–117, Apr. 2014, ISSN: 18792782. DOI: [10.1016/j.neunet.2014.09.003](https://doi.org/10.1016/j.neunet.2014.09.003). [Online]. Available: <http://www.idsia.ch/~CB%9Cjuergen/DeepLearning80ct2014.texCompleteBIBTEXfile%20http://arxiv.org/abs/1404.7828%20http://dx.doi.org/10.1016/j.neunet.2014.09.003>.
- [7] D. Ciregan, U. Meier, and J. Schmidhuber, “Multi-column deep neural networks for image classification”, in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Feb. 2012, pp. 3642–3649, ISBN: 9781467312264. DOI: [10.1109/CVPR.2012.6248110](https://doi.org/10.1109/CVPR.2012.6248110). [Online]. Available: <http://arxiv.org/abs/1202.2745>.
- [8] G. Schillaci, C.-N. Ritter, V. V. Hafner, and B. Lara, “Body Representations for Robot Ego-Noise Modelling and Prediction. Towards the Development of a Sense of Agency in Artificial Agents”, 2016, pp. 390–397. DOI: [10.7551/978-0-262-33936-0-ch065](https://doi.org/10.7551/978-0-262-33936-0-ch065).

- [9] C. Lang, G. Schillaci, and V. V. Hafner, “A deep convolutional neural network model for sense of agency and object permanence in robots”, in *2018 Joint IEEE 8th International Conference on Development and Learning and Epigenetic Robotics, ICDL-EpiRob 2018*, Institute of Electrical and Electronics Engineers Inc., Sep. 2018, pp. 257–262, ISBN: 9781538661109. DOI: [10.1109/DEVLRN.2018.8761015](https://doi.org/10.1109/DEVLRN.2018.8761015).
- [10] P. Lanillos, J. Pages, and G. Cheng, “Robot self/other distinction: active inference meets neural networks learning in a mirror”, Apr. 2020, ISSN: 09226389. DOI: [10.3233/FAIA200372](https://doi.org/10.3233/FAIA200372). [Online]. Available: <http://arxiv.org/abs/2004.05473>.
- [11] W. Ohata and J. Tani, “Investigation of the Sense of Agency in Social Cognition, Based on Frameworks of Predictive Coding and Active Inference: A Simulation Study on Multimodal Imitative Interaction”, *Frontiers in Neurorobotics*, vol. 14, p. 61, Sep. 2020, ISSN: 1662-5218. DOI: [10.3389/fnbot.2020.00061](https://doi.org/10.3389/fnbot.2020.00061). [Online]. Available: <https://www.frontiersin.org/article/10.3389/fnbot.2020.00061/full>.
- [12] S. Gallagher, *Philosophical conceptions of the self: Implications for cognitive science*, Jan. 2000. DOI: [10.1016/S1364-6613\(99\)01417-5](https://doi.org/10.1016/S1364-6613(99)01417-5).
- [13] N.-A. Hinz, P. Lanillos, H. Mueller, and G. Cheng, “Drifting perceptual patterns suggest prediction errors fusion rather than hypothesis selection: replicating the rubber-hand illusion on a robot”, pp. 125–132, 2019. DOI: [10.1109/devlrn.2018.8761005](https://doi.org/10.1109/devlrn.2018.8761005).
- [14] P. Rochat, “What is it Like to be a Newborn?”, in *The Oxford Handbook of the Self*, 2011, ISBN: 9780191724848. DOI: [10.1093/oxfordhb/9780199548019.003.0003](https://doi.org/10.1093/oxfordhb/9780199548019.003.0003).
- [15] Y. K. Georgie, G. Schillaci, and V. V. Hafner, “An interdisciplinary overview of developmental indices and behavioral measures of the minimal self”, *2019 Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics, ICDL-EpiRob 2019*, pp. 129–136, Jun. 2019. DOI: [10.1109/DEVLRN.2019.8850703](https://doi.org/10.1109/DEVLRN.2019.8850703). [Online]. Available: <http://arxiv.org/abs/1907.00709>.
- [16] T. Parr, G. Rees, and K. J. Friston, “Computational neuropsychology and bayesian inference”, *Frontiers in Human Neuroscience*, vol. 12, no. February, pp. 1–14, 2018, ISSN: 16625161. DOI: [10.3389/fnhum.2018.00061](https://doi.org/10.3389/fnhum.2018.00061).
- [17] Z. Ghahramani, *Probabilistic machine learning and artificial intelligence*, May 2015. DOI: [10.1038/nature14541](https://doi.org/10.1038/nature14541). [Online]. Available: <https://www.nature.com/articles/nature14541>.
- [18] J. W. Moore, D. M. Wegner, and P. Haggard, “Modulating the sense of agency with external cues”, *Consciousness and Cognition*, vol. 18, no. 4, pp. 1056–1064, Dec. 2009, ISSN: 10538100. DOI: [10.1016/j.concog.2009.05.004](https://doi.org/10.1016/j.concog.2009.05.004). [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/19515577/> <https://linkinghub.elsevier.com/retrieve/pii/S105381000900083X>.
- [19] J. W. Moore and S. S. Obhi, “Intentional binding and the sense of agency: A review”, *Consciousness and Cognition*, vol. 21, no. 1, pp. 546–561, Mar. 2012, ISSN: 10538100. DOI: [10.1016/j.concog.2011.12.002](https://doi.org/10.1016/j.concog.2011.12.002). [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1053810011002881>.

- [20] R. Legaspi and T. Toyoizumi, “A Bayesian psychophysics model of sense of agency”, *Nature Communications*, vol. 10, no. 1, pp. 1–11, 2019, ISSN: 20411723. DOI: [10.1038/s41467-019-12170-0](https://doi.org/10.1038/s41467-019-12170-0). [Online]. Available: <http://dx.doi.org/10.1038/s41467-019-12170-0>.
- [21] S. Kahl and S. Kopp, “A predictive processing model of perception and action for self-other distinction”, *Frontiers in Psychology*, vol. 9, no. DEC, p. 2421, Oct. 2018, ISSN: 1664-1078. DOI: [10.3389/fpsyg.2018.02421](https://doi.org/10.3389/fpsyg.2018.02421). [Online]. Available: <https://www.frontiersin.org/article/10.3389/fpsyg.2018.02421/full%20http://arxiv.org/abs/1810.09879%20http://dx.doi.org/10.3389/fpsyg.2018.02421>.
- [22] K. Friston, P. Schwartenbeck, T. FitzGerald, M. Moutoussis, T. Behrens, and R. J. Dolan, “The anatomy of choice: active inference and agency”, *Frontiers in Human Neuroscience*, vol. 7, no. SEP, Sep. 2013, ISSN: 1662-5161. DOI: [10.3389/fnhum.2013.00598](https://doi.org/10.3389/fnhum.2013.00598). [Online]. Available: <http://journal.frontiersin.org/article/10.3389/fnhum.2013.00598/abstract>.
- [23] N. Chater, M. Oaksford, U. Hahn, and E. Heit, “Bayesian models of cognition”, *Wiley Interdisciplinary Reviews: Cognitive Science*, vol. 1, no. 6, pp. 811–823, 2010, ISSN: 19395078. DOI: [10.1002/wcs.79](https://doi.org/10.1002/wcs.79).
- [24] M. Ghavamzadeh, S. Mannor, J. Pineau, and A. Tamar, *Bayesian reinforcement learning: A survey*, 2015. DOI: [10.1561/22000000049](https://doi.org/10.1561/22000000049).
- [25] J. Hohwy, A. Roepstorff, and K. Friston, “Predictive coding explains binocular rivalry: An epistemological review”, *Cognition*, vol. 108, no. 3, pp. 687–701, Sep. 2008, ISSN: 00100277. DOI: [10.1016/j.cognition.2008.05.010](https://doi.org/10.1016/j.cognition.2008.05.010). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010027708001327>.
- [26] H. Brown and K. J. Friston, “Free-energy and illusions: The Cornsweet effect”, *Frontiers in Psychology*, vol. 3, no. FEB, pp. 1–13, 2012, ISSN: 16641078. DOI: [10.3389/fpsyg.2012.00043](https://doi.org/10.3389/fpsyg.2012.00043).
- [27] J. Hohwy, *The Predictive Mind*. Oxford University Press, Nov. 2013, ISBN: 9780199682737. DOI: [10.1093/acprof:oso/9780199682737.001.0001](https://doi.org/10.1093/acprof:oso/9780199682737.001.0001). [Online]. Available: <http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199682737.001.0001/acprof-9780199682737>.
- [28] D. M. Eagleman, “Visual illusions and neurobiology”, *Nature Reviews Neuroscience*, vol. 2, no. 12, pp. 920–926, Dec. 2001, ISSN: 1471-003X. DOI: [10.1038/35104092](https://doi.org/10.1038/35104092). [Online]. Available: <http://www.nature.com/articles/35104092>.
- [29] E. Watanabe, A. Kitaoka, K. Sakamoto, M. Yasugi, and K. Tanaka, “Illusory motion reproduced by deep neural networks trained for prediction”, *Frontiers in Psychology*, vol. 9, no. MAR, pp. 1–12, 2018, ISSN: 16641078. DOI: [10.3389/fpsyg.2018.00345](https://doi.org/10.3389/fpsyg.2018.00345).
- [30] M. Botvinick and J. Cohen, “Rubber hands ‘feel’ touch that eyes see”, *Nature*, vol. 391, no. 6669, pp. 756–756, Feb. 1998, ISSN: 0028-0836. DOI: [10.1038/35784](https://doi.org/10.1038/35784). [Online]. Available: <http://www.nature.com/articles/35784>.

- [31] A. Kalckert and H. H. Ehrsson, “Moving a Rubber Hand that Feels Like Your Own: A Dissociation of Ownership and Agency”, *Frontiers in Human Neuroscience*, vol. 6, no. Article 40, Mar. 2012, ISSN: 1662-5161. DOI: [10.3389/fnhum.2012.00040](https://doi.org/10.3389/fnhum.2012.00040). [Online]. Available: <http://journal.frontiersin.org/article/10.3389/fnhum.2012.00040/abstract>.
- [32] —, “The moving rubber hand illusion revisited: Comparing movements and visuotactile stimulation to induce illusory ownership”, *Consciousness and Cognition*, vol. 26, no. 1, pp. 117–132, May 2014, ISSN: 10902376. DOI: [10.1016/j.concog.2014.02.003](https://doi.org/10.1016/j.concog.2014.02.003).
- [33] M. D. Lee and W. Vanpaemel, “Determining informative priors for cognitive models”, *Psychonomic Bulletin and Review*, vol. 25, no. 1, pp. 114–127, 2018, ISSN: 15315320. DOI: [10.3758/s13423-017-1238-3](https://doi.org/10.3758/s13423-017-1238-3).
- [34] L. Zaadnoordijk, T. R. Besold, and S. Hunnius, “A match does not make a sense: on the sufficiency of the comparator model for explaining the sense of agency”, *Neuroscience of Consciousness*, vol. 2019, no. 1, Jan. 2019, ISSN: 2057-2107. DOI: [10.1093/nc/niz006](https://doi.org/10.1093/nc/niz006). [Online]. Available: <https://orcid.org/0000-0002-8002-0049%20https://academic.oup.com/nc/article/doi/10.1093/nc/niz006/5488527>.
- [35] A. Clark, “Whatever next? Predictive brains, situated agents, and the future of cognitive science”, *Behavioral and Brain Sciences*, vol. 36, no. 3, pp. 181–204, 2013, ISSN: 14691825. DOI: [10.1017/S0140525X12000477](https://doi.org/10.1017/S0140525X12000477).
- [36] K. Friston and S. Kiebel, “Predictive coding under the free-energy principle”, *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 364, no. 1521, pp. 1211–1221, 2009, ISSN: 14712970. DOI: [10.1098/rstb.2008.0300](https://doi.org/10.1098/rstb.2008.0300).
- [37] J. M. Beck, W. J. Ma, R. Kiani, T. Hanks, A. K. Churchland, J. Roitman, M. N. Shadlen, P. E. Latham, and A. Pouget, “Probabilistic Population Codes for Bayesian Decision Making”, *Neuron*, vol. 60, no. 6, pp. 1142–1152, Dec. 2008, ISSN: 08966273. DOI: [10.1016/j.neuron.2008.09.021](https://doi.org/10.1016/j.neuron.2008.09.021).
- [38] M. O. Ernst and M. S. Banks, “Humans integrate visual and haptic information in a statistically optimal fashion.”, *Nature*, vol. 415, no. 6870, pp. 429–33, Jan. 2002, ISSN: 0028-0836. DOI: [10.1038/415429a](https://doi.org/10.1038/415429a). [Online]. Available: www.nature.com/20http://www.nature.com/articles/415429a%20http://www.ncbi.nlm.nih.gov/pubmed/11807554.
- [39] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan, “An internal model for sensorimotor integration”, *Science*, vol. 269, no. 5232, pp. 1880–1882, Sep. 1995, ISSN: 00368075. DOI: [10.1126/science.7569931](https://doi.org/10.1126/science.7569931).
- [40] K. P. Körding, S. P. Ku, and D. M. Wolpert, “Bayesian integration in force estimation”, *Journal of Neurophysiology*, vol. 92, no. 5, pp. 3161–3165, Nov. 2004, ISSN: 00223077. DOI: [10.1152/jn.00275.2004](https://doi.org/10.1152/jn.00275.2004). [Online]. Available: <https://www.physiology.org/doi/10.1152/jn.00275.2004>.
- [41] T. L. Griffiths and J. B. Tenenbaum, “Theory-Based Causal Induction”, *Psychological Review*, 2009, ISSN: 0033295X. DOI: [10.1037/a0017201](https://doi.org/10.1037/a0017201).

- [42] R. A. Adams, K. E. Stephan, H. R. Brown, C. D. Frith, and K. J. Friston, “The Computational Anatomy of Psychosis”, *Frontiers in Psychiatry*, vol. 4, no. May, pp. 1–26, 2013, ISSN: 1664-0640. DOI: [10.3389/fpsy.2013.00047](https://doi.org/10.3389/fpsy.2013.00047).
- [43] M. Jones and B. C. Love, “Bayesian fundamentalism or enlightenment? on the explanatory status and theoretical contributions of bayesian models of cognition”, *Behavioral and Brain Sciences*, vol. 34, no. 4, pp. 169–188, 2011, ISSN: 0140525X. DOI: [10.1017/S0140525X10003134](https://doi.org/10.1017/S0140525X10003134). [Online]. Available: <http://love..>
- [44] J. S. Bowers and C. J. Davis, “Bayesian just-so stories in psychology and neuroscience”, *Psychological Bulletin*, vol. 138, no. 3, pp. 389–414, May 2012, ISSN: 00332909. DOI: [10.1037/a0026450](https://doi.org/10.1037/a0026450). [Online]. Available: [/record/2012-10601-001](http://record/2012-10601-001).
- [45] A. A. Stocker and E. P. Simoncelli, “Noise characteristics and prior expectations in human visual speed perception”, *Nature Neuroscience*, vol. 9, no. 4, pp. 578–585, 2006, ISSN: 10976256. DOI: [10.1038/nn1669](https://doi.org/10.1038/nn1669).
- [46] M. Berniker and K. Kording, “Estimating the sources of motor errors for adaptation and generalization”, *Nature Neuroscience*, 2008, ISSN: 10976256. DOI: [10.1038/nn.2229](https://doi.org/10.1038/nn.2229).
- [47] H. Tassinari, T. E. Hudson, and M. S. Landy, “Combining priors and noisy visual cues in a rapid pointing task”, *Journal of Neuroscience*, 2006, ISSN: 02706474. DOI: [10.1523/JNEUROSCI.2779-06.2006](https://doi.org/10.1523/JNEUROSCI.2779-06.2006).
- [48] A. Wald, “An Essentially Complete Class of Admissible Decision Functions”, *The Annals of Mathematical Statistics*, vol. 18, no. 4, pp. 549–555, 1947, ISSN: 0003-4851. DOI: [10.1214/aoms/1177730345](https://doi.org/10.1214/aoms/1177730345).
- [49] K. Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, J. O’Doherty, and G. Pezzulo, *Active inference and learning*, Sep. 2016. DOI: [10.1016/j.neubiorev.2016.06.022](https://doi.org/10.1016/j.neubiorev.2016.06.022).
- [50] T. Parr, A. W. Corcoran, K. J. Friston, and J. Hohwy, “Perceptual awareness and active inference”, *Neuroscience of Consciousness*, vol. 2019, no. 1, pp. 1–15, 2019, ISSN: 2057-2107. DOI: [10.1093/nc/niz012](https://doi.org/10.1093/nc/niz012).
- [51] T. Parr and K. J. Friston, “Generalised free energy and active inference”, *Biological Cybernetics*, vol. 113, no. 5, pp. 495–513, 2019, ISSN: 14320770. DOI: [10.1007/s00422-019-00805-w](https://doi.org/10.1007/s00422-019-00805-w). [Online]. Available: <https://doi.org/10.1007/s00422-019-00805-w>.
- [52] R. Raman and S. Sarkar, “Predictive Coding: A Possible Explanation of Filling-In at the Blind Spot”, *PLOS ONE*, vol. 11, no. 3, M. J. Chacron, Ed., e0151194, Mar. 2016, ISSN: 1932-6203. DOI: [10.1371/journal.pone.0151194](https://doi.org/10.1371/journal.pone.0151194). [Online]. Available: <https://dx.plos.org/10.1371/journal.pone.0151194>.

- [53] M. Samad, A. J. Chung, and L. Shams, "Perception of Body Ownership Is Driven by Bayesian Sensory Inference", *PLoS ONE*, vol. 10, no. 2, M. A. Lebedev, Ed., e0117178, Feb. 2015, ISSN: 1932-6203. DOI: [10.1371/journal.pone.0117178](https://doi.org/10.1371/journal.pone.0117178). [Online]. Available: <https://dx.plos.org/10.1371/journal.pone.0117178>; <http://www.ncbi.nlm.nih.gov/pubmed/25658822>; <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4320053>.
- [54] T. Schürmann, J. Vogt, O. Christ, and P. Beckerle, "The Bayesian causal inference model benefits from an informed prior to predict proprioceptive drift in the rubber foot illusion", *Cognitive Processing*, vol. 20, no. 4, pp. 447–457, Nov. 2019, ISSN: 1612-4782. DOI: [10.1007/s10339-019-00928-9](https://doi.org/10.1007/s10339-019-00928-9). [Online]. Available: <https://doi.org/10.1007/s10339-019-00928-9>; <http://link.springer.com/10.1007/s10339-019-00928-9>.
- [55] A. Pouget, J. M. Beck, W. J. Ma, and P. E. Latham, *Probabilistic brains: Knowns and unknowns*, 2013. DOI: [10.1038/nn.3495](https://doi.org/10.1038/nn.3495). [Online]. Available: <http://www.nature.com/reprints/index.html>.
- [56] L. Aitchison and M. Lengyel, "With or without you: predictive coding and Bayesian inference in the brain", *Current Opinion in Neurobiology*, vol. 46, pp. 219–227, 2017, ISSN: 18736882. DOI: [10.1016/j.conb.2017.08.010](https://doi.org/10.1016/j.conb.2017.08.010).
- [57] C. Mead, "Neuromorphic Electronic Systems", *Proceedings of the IEEE*, 1990, ISSN: 15582256. DOI: [10.1109/5.58356](https://doi.org/10.1109/5.58356).
- [58] R. P. N. Rao and D. H. Ballard, "Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects", *Nature Neuroscience*, vol. 2, no. 1, pp. 79–87, 1999, ISSN: 10976256. DOI: [10.1038/4580](https://doi.org/10.1038/4580).
- [59] K. Friston, "A theory of cortical responses", *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 360, no. 1456, pp. 815–836, 2005, ISSN: 09628436. DOI: [10.1098/rstb.2005.1622](https://doi.org/10.1098/rstb.2005.1622).
- [60] K. J. Friston, J. Daunizeau, J. Kilner, and S. J. Kiebel, "Action and behavior: A free-energy formulation", *Biological Cybernetics*, vol. 102, no. 3, pp. 227–260, 2010, ISSN: 03401200. DOI: [10.1007/s00422-010-0364-z](https://doi.org/10.1007/s00422-010-0364-z).
- [61] R. Bogacz, "A tutorial on the free-energy framework for modelling perception and learning", *Journal of Mathematical Psychology*, vol. 76, pp. 198–211, Feb. 2017, ISSN: 00222496. DOI: [10.1016/j.jmp.2015.11.003](https://doi.org/10.1016/j.jmp.2015.11.003). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0022249615000759>; <https://linkinghub.elsevier.com/retrieve/pii/S0022249615000759>.
- [62] K. A. Hansen, S. F. Hillenbrand, and L. G. Ungerleider, "Effects of prior knowledge on decisions made under perceptual vs. Categorical uncertainty", *Frontiers in Neuroscience*, vol. 6, no. NOV, p. 163, Nov. 2012, ISSN: 16624548. DOI: [10.3389/fnins.2012.00163](https://doi.org/10.3389/fnins.2012.00163). [Online]. Available: <http://journal.frontiersin.org/article/10.3389/fnins.2012.00163/abstract>.

- [63] P. Kok, G. J. Brouwer, M. A. J. van Gerven, and F. P. de Lange, “Prior Expectations Bias Sensory Representations in Visual Cortex”, *Journal of Neuroscience*, vol. 33, no. 41, pp. 16 275–16 284, Oct. 2013, ISSN: 0270-6474. DOI: [10.1523/JNEUROSCI.0742-13.2013](https://doi.org/10.1523/JNEUROSCI.0742-13.2013). [Online]. Available: <http://surfer.nmr.mgh.harvard.edu/http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.0742-13.2013>.
- [64] T. Rood, M. van Gerven, and P. Lanillos, “A deep active inference model of the rubber-hand illusion”, 2020. [Online]. Available: <http://arxiv.org/abs/2008.07408>.
- [65] M. Castellano-Quero, J.-A. Fernández-Madrigal, and A. García-Cerezo, “Improving Bayesian inference efficiency for sensory anomaly detection and recovery in mobile robots”, *Expert Systems with Applications*, vol. 163, p. 113 755, Jan. 2021, ISSN: 09574174. DOI: [10.1016/j.eswa.2020.113755](https://doi.org/10.1016/j.eswa.2020.113755). [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0957417420305790>.
- [66] P. Lanillos and G. Cheng, “Active inference with function learning for robot body perception”, 2018. [Online]. Available: www.selfception.eu.
- [67] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005, ISBN: 9780262201629. [Online]. Available: <https://mitpress.mit.edu/books/probabilistic-robotics>.
- [68] C. M. Bishop, “A New Framework for Machine Learning”, in *IEEE World Congress on Computational Intelligence*, Heidelberg: Springer, Jun. 2008, pp. 1–24. [Online]. Available: <http://research.microsoft.com/%E2%88%BCmbishop>.
- [69] R. T. Cox, “Probability, Frequency and Reasonable Expectation”, *American Journal of Physics*, 1946, ISSN: 0002-9505. DOI: [10.1119/1.1990764](https://doi.org/10.1119/1.1990764).
- [70] M. A. Chappell, A. R. Groves, B. Whitcher, and M. W. Woolrich, “Variational Bayesian inference for a nonlinear forward model”, *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 223–236, 2009, ISSN: 1053587X. DOI: [10.1109/TSP.2008.2005752](https://doi.org/10.1109/TSP.2008.2005752).
- [71] M. J. Beal, “Variational Algorithms for Approximate Bayesian Inference”, University College London, London, Tech. Rep., 2003.
- [72] K. Perrykkad, R. P. Lawson, S. D. Jamadar, and J. Hohwy, “The Effect of Uncertainty on Prediction Error in the Action-Perception Loop”, *bioRxiv*, Jun. 2020. DOI: [10.1101/2020.06.22.166108](https://doi.org/10.1101/2020.06.22.166108). [Online]. Available: <https://doi.org/10.1101/2020.06.22.166108><https://www.biorxiv.org/content/10.1101/2020.06.22.166108v1><https://www.biorxiv.org/content/10.1101/2020.06.22.166108v1.abstract>.
- [73] P. Haggard, S. Clark, and J. Kalogeras, “Voluntary action and conscious awareness”, *Nature Neuroscience*, vol. 5, no. 4, pp. 382–385, Apr. 2002, ISSN: 1097-6256. DOI: [10.1038/nn827](https://doi.org/10.1038/nn827). [Online]. Available: <http://www.nature.com/articles/nn827>.

- [74] J. W. Moore, "What Is the Sense of Agency and Why Does it Matter?", *Frontiers in Psychology*, vol. 7, no. AUG, Aug. 2016, ISSN: 1664-1078. DOI: [10.3389/fpsyg.2016.01272](https://doi.org/10.3389/fpsyg.2016.01272). [Online]. Available: [/pmc/articles/PMC5002400/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5002400/%20http://journal.frontiersin.org/Article/10.3389/fpsyg.2016.01272/abstract](https://pmc/articles/PMC5002400/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5002400/%20http://journal.frontiersin.org/Article/10.3389/fpsyg.2016.01272/abstract).
- [75] C. D. Frith, S.-J. Blakemore, and D. M. Wolpert, "Abnormalities in the awareness and control of action", *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 355, no. 1404, pp. 1771–1788, Dec. 2000, ISSN: 0962-8436. DOI: [10.1098/rstb.2000.0734](https://doi.org/10.1098/rstb.2000.0734). [Online]. Available: [/pmc/articles/PMC1692910/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1692910/%20https://royalsocietypublishing.org/doi/10.1098/rstb.2000.0734](https://pmc/articles/PMC1692910/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1692910/%20https://royalsocietypublishing.org/doi/10.1098/rstb.2000.0734).
- [76] D. M. Wegner and T. Wheatley, "Apparent mental causation: Sources of the experience of will", *American Psychologist*, vol. 54, no. 7, pp. 480–492, 1999, ISSN: 0003066X. DOI: [10.1037/0003-066X.54.7.480](https://doi.org/10.1037/0003-066X.54.7.480). [Online]. Available: [/record/1999-05760-003](https://record/1999-05760-003).
- [77] V. Chambon, N. Sidarus, and P. Haggard, "From action intentions to action effects: how does the sense of agency come about?", *Frontiers in Human Neuroscience*, vol. 8, no. MAY, May 2014, ISSN: 1662-5161. DOI: [10.3389/fnhum.2014.00320](https://doi.org/10.3389/fnhum.2014.00320). [Online]. Available: [/pmc/articles/PMC4030148/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4030148/%20http://journal.frontiersin.org/article/10.3389/fnhum.2014.00320/abstract](https://pmc/articles/PMC4030148/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4030148/%20http://journal.frontiersin.org/article/10.3389/fnhum.2014.00320/abstract).
- [78] J. W. Moore and P. C. Fletcher, "Sense of agency in health and disease: A review of cue integration approaches", *Consciousness and Cognition*, vol. 21, no. 1, pp. 59–68, Mar. 2012, ISSN: 10538100. DOI: [10.1016/j.concog.2011.08.010](https://doi.org/10.1016/j.concog.2011.08.010).
- [79] M. Synofzik, G. Vosgerau, and A. Newen, "Beyond the comparator model: a multifactorial two-step account of agency", *Consciousness and Cognition*, vol. 17, no. 1, pp. 219–39, Mar. 2008, ISSN: 1090-2376. DOI: [10.1016/j.concog.2007.03.010](https://doi.org/10.1016/j.concog.2007.03.010). [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/17482480>.
- [80] N. Wolpe, P. Haggard, H. R. Siebner, and J. B. Rowe, "Cue integration and the perception of action in intentional binding", *Experimental Brain Research*, 2013, ISSN: 00144819. DOI: [10.1007/s00221-013-3419-2](https://doi.org/10.1007/s00221-013-3419-2).
- [81] K. Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, and G. Pezzulo, "Active inference: A process theory", *Neural Computation*, 2017, ISSN: 1530888X. DOI: [10.1162/NECO{a}_{00912](https://doi.org/10.1162/NECO{a}_{00912).
- [82] A. Sato and A. Yasuda, "Illusion of sense of self-agency: discrepancy between the predicted and actual sensory consequences of actions modulates the sense of self-agency, but not the sense of self-ownership", *Cognition*, vol. 94, no. 3, pp. 241–255, Jan. 2005, ISSN: 00100277. DOI: [10.1016/j.cognition.2004.04.003](https://doi.org/10.1016/j.cognition.2004.04.003). [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S001002770400099X>.

- [83] N. Braun, S. Debener, N. Spychala, E. Bongartz, P. Sörös, H. H. Müller, and A. Philipsen, *The senses of agency and ownership: A review*, Apr. 2018. DOI: [10.3389/fpsyg.2018.00535](https://doi.org/10.3389/fpsyg.2018.00535). [Online]. Available: www.frontiersin.org.
- [84] S. Bechtle, G. Schillaci, and V. V. Hafner, “On the sense of agency and of object permanence in robots”, in *2016 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, IEEE, Sep. 2016, pp. 166–171, ISBN: 978-1-5090-5069-7. DOI: [10.1109/DEVLRN.2016.7846812](https://doi.org/10.1109/DEVLRN.2016.7846812). [Online]. Available: <http://ieeexplore.ieee.org/document/7846812/>.
- [85] V. V. Hafner, P. Loviken, A. Pico Villalpando, and G. Schillaci, *Prerequisites for an Artificial Self*, 2020. DOI: [10.3389/fnbot.2020.00005](https://doi.org/10.3389/fnbot.2020.00005). [Online]. Available: www.frontiersin.org.
- [86] K. Friston, “Hierarchical Models in the Brain”, *PLoS Computational Biology*, vol. 4, no. 11, O. Sporns, Ed., e1000211, Nov. 2008, ISSN: 1553-7358. DOI: [10.1371/journal.pcbi.1000211](https://doi.org/10.1371/journal.pcbi.1000211). [Online]. Available: <https://dx.plos.org/10.1371/journal.pcbi.1000211>.
- [87] R. J. Van Beers, A. C. Sittig, and J. J. Denier Van Der Gon, “The precision of proprioceptive position sense”, *Experimental Brain Research*, 1998, ISSN: 00144819. DOI: [10.1007/s002210050525](https://doi.org/10.1007/s002210050525).
- [88] O. Nomura and Y. Miyake, “Sense of agency on continuous hand movement with lagged visual feedback”, in *2017 56th Annual Conference of the Society of Instrument and Control Engineers of Japan, SICE 2017*, vol. 2017-Novem, Institute of Electrical and Electronics Engineers Inc., Nov. 2017, pp. 970–975, ISBN: 9784907764579. DOI: [10.23919/SICE.2017.8105458](https://doi.org/10.23919/SICE.2017.8105458).
- [89] S. Plagenhoef, F. Gaynor Evans, and T. Abdelnour, “Anatomical Data for Analyzing Human Motion”, *Research Quarterly for Exercise and Sport*, vol. 54, no. 2, pp. 169–178, 1983, ISSN: 21683824. DOI: [10.1080/02701367.1983.10605290](https://doi.org/10.1080/02701367.1983.10605290). [Online]. Available: <https://shapeamerica.tandfonline.com/doi/abs/10.1080/02701367.1983.10605290>.
- [90] A. Marotta, F. Bombieri, M. Zampini, F. Schena, C. Dallocchio, M. Fiorio, and M. Tinazzi, “The moving rubber hand illusion reveals that explicit sense of agency for tapping movements is preserved in functional movement disorders”, *Frontiers in Human Neuroscience*, vol. 11, p. 291, Jun. 2017, ISSN: 16625161. DOI: [10.3389/fnhum.2017.00291](https://doi.org/10.3389/fnhum.2017.00291). [Online]. Available: <http://journal.frontiersin.org/article/10.3389/fnhum.2017.00291/full>.
- [91] C. Heij, P. d. Boer, P. H. Franses, T. Kloek, H. K. v. Dijk, and A. a. t. E. U. i. Rotterdam, *Econometric Methods with Applications in Business and Economics*. Oxford University Press, 2004, p. 816, ISBN: 0191608408. [Online]. Available: https://www.researchgate.net/profile/Herman_Van_Dijk2/publication/227467912_Econometric_Methods_with_Applications_in_Business_and_Economics/links/5695a71208ae3ad8e33d8d93.pdf.
- [92] C. Frith, “The self in action: Lessons from delusions of control”, *Consciousness and Cognition*, vol. 14, no. 4, pp. 752–770, Dec. 2005, ISSN: 10902376. DOI: [10.1016/j.concog.2005.04.002](https://doi.org/10.1016/j.concog.2005.04.002).

- [93] P. C. Fletcher and C. D. Frith, “Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia”, *Nature Reviews Neuroscience*, vol. 10, no. 1, pp. 48–58, Jan. 2009, ISSN: 1471-003X. DOI: [10.1038/nrn2536](https://doi.org/10.1038/nrn2536). [Online]. Available: <http://www.nature.com/articles/nrn2536>.
- [94] W. Wen, A. Yamashita, and H. Asama, “The sense of agency during continuous action: Performance is more important than action-Feedback association”, *PLoS ONE*, vol. 10, no. 4, Apr. 2015, ISSN: 19326203. DOI: [10.1371/journal.pone.0125226](https://doi.org/10.1371/journal.pone.0125226). [Online]. Available: [/pmc/articles/PMC4404253/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4404253/](https://pubmed.ncbi.nlm.nih.gov/20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4404253/).
- [95] K. Friston, “The free-energy principle: A unified brain theory?”, *Nature Reviews Neuroscience*, vol. 11, no. 2, pp. 127–138, 2010, ISSN: 1471003X. DOI: [10.1038/nrn2787](https://doi.org/10.1038/nrn2787). [Online]. Available: <http://dx.doi.org/10.1038/nrn2787>.
- [96] M. A. Apps and M. Tsakiris, “The free-energy self: A predictive coding account of self-recognition”, *Neuroscience and Biobehavioral Reviews*, vol. 41, pp. 85–97, 2014, ISSN: 18737528. DOI: [10.1016/j.neubiorev.2013.01.029](https://doi.org/10.1016/j.neubiorev.2013.01.029). [Online]. Available: <http://dx.doi.org/10.1016/j.neubiorev.2013.01.029>.
- [97] Z. Ghahramani and S. T. Roweis, “Learning nonlinear dynamical systems using an EM algorithm”, in *Advances in Neural Information Processing Systems*, 1999, pp. 431–437, ISBN: 0262112450. [Online]. Available: <http://www.gatsby.ucl.ac.uk/>.
- [98] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum Likelihood from Incomplete Data Via the EM Algorithm”, *Journal of the Royal Statistical Society: Series B (Methodological)*, 1977. DOI: [10.1111/j.2517-6161.1977.tb01600.x](https://doi.org/10.1111/j.2517-6161.1977.tb01600.x).
- [99] J. A. Bilmes, “A Gentle Tutorial of the EM Algorithm and its Application”, *International Computer Science Institute*, vol. 4, no. 510, p. 126, 1998.
- [100] P. W. Battaglia, R. A. Jacobs, and R. N. Aslin, “Bayesian integration of visual and auditory signals for spatial localization”, *Journal of the Optical Society of America A*, 2003, ISSN: 1084-7529. DOI: [10.1364/josaa.20.001391](https://doi.org/10.1364/josaa.20.001391).
- [101] K. J. Friston, N. Trujillo-Barreto, and J. Daunizeau, “DEM: A variational treatment of dynamic systems”, *NeuroImage*, vol. 41, no. 3, pp. 849–885, Jul. 2008, ISSN: 10538119. DOI: [10.1016/j.neuroimage.2008.02.054](https://doi.org/10.1016/j.neuroimage.2008.02.054).
- [102] O. Cappé, “Online Expectation-Maximisation”, *Mixtures: Estimation and Applications*, pp. 31–53, Nov. 2010. DOI: [10.1002/9781119995678.ch2](https://doi.org/10.1002/9781119995678.ch2). [Online]. Available: <http://doi.wiley.com/10.1002/9781119995678.ch2%20http://arxiv.org/abs/1011.1745>.
- [103] C. W. Fox and S. J. Roberts, *A Tutorial on Variational Bayesian Inference*, 2012. DOI: [10.1007/s10462-011-9236-8](https://doi.org/10.1007/s10462-011-9236-8). [Online]. Available: <https://www.researchgate.net/publication/226071015>.

- [104] N. Braun, J. D. Thorne, H. Hildebrandt, and S. Debener, “Interplay of agency and ownership: The intentional binding and rubber hand illusion Paradigm Combined”, *PLoS ONE*, vol. 9, no. 11, M. Costantini, Ed., e111967, Nov. 2014, ISSN: 19326203. DOI: [10.1371/journal.pone.0111967](https://doi.org/10.1371/journal.pone.0111967). [Online]. Available: <https://dx.plos.org/10.1371/journal.pone.0111967>.
- [105] K. Friston, “Prediction, perception and agency”, *International Journal of Psychophysiology*, vol. 83, no. 2, pp. 248–252, Feb. 2012, ISSN: 01678760. DOI: [10.1016/j.ijpsycho.2011.11.014](https://doi.org/10.1016/j.ijpsycho.2011.11.014). [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0167876011003618>.
- [106] C. Pezzato, R. Ferrari, and C. Hernandez, “A Novel Adaptive Controller for Robot Manipulators based on Active Inference”, pp. 1–27, 2019. [Online]. Available: <http://arxiv.org/abs/1909.12768>.
- [107] G. Oliver, P. Lanillos, and G. Cheng, “Active inference body perception and action for humanoid robots”, 2019. [Online]. Available: <http://arxiv.org/abs/1906.03022>.
- [108] S.-J. Blakemore, C. D. Frith, and D. M. Wolpert, “Spatio-Temporal Prediction Modulates the Perception of Self-Produced Stimuli”, *Journal of Cognitive Neuroscience*, vol. 11, no. 5, pp. 551–559, Sep. 1999, ISSN: 0898-929X. DOI: [10.1162/089892999563607](https://doi.org/10.1162/089892999563607). [Online]. Available: <http://www.mitpressjournals.org/doi/10.1162/089892999563607>.



MEASURES OF SENSE OF AGENCY IN EXPERIMENTAL INVESTIGATION

SoA is typically an unconscious experience, making it difficult to measure in experiments [74]. Unlike for conscious experiences, such as vision, experimenters have to develop paradigms through which to detect SoA. Understanding the measures of SoA in experiments is relevant because the measures (i) help study and validate theories on SoA and (ii) highlight that the quantification of SoA in an artificial agent to further research on the artificial self is challenging.

SoA paradigms can be grouped into *explicit* and *implicit* measures. In *explicit measures*, the experimenters directly ask the participants whether they perceive a SoA. In the mRHI, researchers determine the magnitude of SoA that participants experience with a questionnaire [31][32]. Besides the fact that explicit measures are vulnerable to issues like demand effects, these measures cannot be used in artificial agents. However, artificial agents are a valuable tool in investigating subjective experiences because, unlike humans, the internal states and sensory perception of the robot can be recorded and analyzed [15]. To illustrate, two previous studies replicate the (static) RHI in artificial agents [13][64] which helps uncover the principles of the sense of ownership. Like SoA, the sense of ownership is a subjective and unconscious process that in human experiments is generally determined by a questionnaire. Therefore, rather than directly researching the sense of ownership, the artificial agents perform body estimation using the free-energy principle [13] and Active Inference [64]. Hinz et al.[13] find that the robot's body estimation drifting patterns are similar to drifting patterns from human experiments. In the human experiment, participants indicated an ownership illusion in every experimental condition meaning body-ownership illusions and the drift are two related processes.

To use *implicit measures*, experimenters have to be more inventive. The most widely applied implicit measure of SoA is intentional binding. Intentional binding is the "perceived compression of the time interval between voluntary action and its outcome" [19]. In an extension of the original mRHI by Kalckert and Ehrsson [31], Braun et al. find

A

strong intentional binding effects indicating that intentional binding can also be induced over the rubber hand. Another implicit measure of SoA is sensory attenuation. Sensory consequences by voluntary action have a lower intensity than consequences from passive movement [108].

B

GENERATING DISCRETE-TIME STATES WITH THE PLANT

Since the point mass mRHI is implemented in a simulation, I discretize the continuous state-space equations for the plant (eq. 3.4 and eq. 3.5) using forward Euler to generate the states of each mass. The code that generates the states of the three point masses for the point mass mRHI, i.e. the plant, is available on GitHub at <https://github.com/mtan1503/point-mass-mrhi/blob/main/plant.py>. Each i time step is of length h ,

$$\mathbf{x}[i + 1] = \mathbf{A}_d \mathbf{x}[i] + \mathbf{B}_d \mathbf{u}_p[i] \quad (\text{B.1})$$

where

$$\mathbf{A}_d = I_k + h \cdot \mathbf{A}$$

$$\mathbf{B}_d = h \cdot \mathbf{B}_p.$$

Then to generate the discrete-time derivatives of the states, $\dot{\mathbf{x}}[i]$, and the output, $\mathbf{y}[i]$ with noise, the generated states are used in the state-space equations for the plant. In discrete-time this is as follows,

$$\dot{\mathbf{x}}[i] = \mathbf{A} \mathbf{x}[i] + \mathbf{B}_p \mathbf{u}_p[i] + \boldsymbol{\omega}[i] \quad (\text{B.2})$$

$$\mathbf{y}[i] = \mathbf{C} \mathbf{x}[i] + \mathbf{z}[i]. \quad (\text{B.3})$$

C

RESULTS FROM POINT MASS mRHI WITH OLS

This appendix shows the numerical results from the simplest point mass mRHI from Chapter 4, where the agent estimates the agency matrix $\hat{\mathbf{B}}$ with OLS.

C.1. EXPERIMENT 1: UNAMBIGUOUS ENVIRONMENTS

For the unambiguous environments of experiment 1, Table C.1 compares the results of the point mass mRHI of the four performance measures for an agent with and without a prior described in Sec. 3.3.2. Overall the prior improves the estimation accuracy and precision of the agency matrix $\hat{\mathbf{B}}$.

The third performance metric is the accuracy of the prediction of the motion states of mass m_1 for both the agent with and without prior. Table C.1 shows this accuracy is the same because the estimated agency parameter \hat{b}_2 is unchanged. The equal accuracy is due to the agent's prior belief only adjusting the agency parameters that the agent does not have agency over. Applying this prior does not improve the agent's SoA for the mass which it does have agency over (i.e. mass m_1).

The fourth performance metric in Table C.1 is the perceived error. The perceived error decreases as a result of the prior. This is expected because the agent without a prior has a SoA over all three masses. Therefore, the agent accumulates prediction error, $\epsilon_{\dot{x}}$, for all the masses that it tries to control. Mathematically, the error accumulates the corresponding j states where $\hat{b}_j > 0$. The agent with the prior tries to minimize its error accumulation by having a SoA over the mass that the agent believes it can best predict.

Table C.1. This table compares the performance of the agent in the point mass mRHI using the OLS algorithm with and without a “prior” per experiment number of experiment 1. Both agents estimate the agency matrix, $\hat{\mathbf{B}}$, with OLS. However, the prior outlines that the agent can control *one* mass. Figure 4.1 describes this algorithm. The columns quantify the agent's performance according to four performance measures (see Sec. 3.3.2). The green cell color indicates where the agent's performance is better between the prior and no prior conditions.

Exp. nr.	Agency accuracy, $\sum_j^k \mu_{\hat{b}_j} - b_j $		Perceived precision, $\sum_j^k \sigma_{\hat{b}_j}^2$		State accuracy, $\sum_j^2 \mu_{\epsilon_{\dot{x}_j}}$		Perceived error, $\mu_{\epsilon_{\dot{x}_a}}$	
	Prior	No prior	Prior	No prior	Prior	No prior	Prior	No prior
1A	3.61×10 ⁻⁶	8.56×10 ⁻⁶	3.20×10 ⁻⁸	1.49×10 ⁻⁷	2.42×10 ⁻²	2.42×10 ⁻²	1.21×10 ⁻²	1.21×10 ⁻²
1B	4.42×10 ⁻⁵	2.62×10 ⁻¹	5.86×10 ⁻⁷	4.96×10 ⁻⁴	2.42×10 ⁻²	2.42×10 ⁻²	1.21×10 ⁻²	1.13×10 ⁰
1C	5.31×10 ⁻⁵	2.62×10 ⁻¹	3.17×10 ⁻⁶	5.01×10 ⁻⁴	2.42×10 ⁻¹	2.42×10 ⁻¹	1.21×10 ⁻¹	1.26×10 ⁰

C.2. EXPERIMENT 2: AMBIGUOUS ENVIRONMENTS

For the ambiguous environments of experiment 2, Table C.2 compares the results of the point mass mRHI of the four performance measures for an agent with and without a prior described in Sec. 3.3.2. Overall the prior deteriorates the estimation accuracy and precision of the agency matrix $\hat{\mathbf{B}}$. Furthermore, the agency accuracy is also around 5-6 orders of magnitude lower than for experiment 1 (Table C.1). This shows that the prior can improve the SoA in certain environments while promoting incorrect SoA in ambiguous environments.

Another indication of an incorrect SoA is that the agent with prior knowledge has a lower accuracy for the predicted states ("state accuracy") while still perceiving a lower error than the agent without a prior. This means that the choice of agency lowers the agent's perceived error, despite being the incorrect choice.

Table C.2. This table compares the performance of the agent in the point mass mRHI using the OLS algorithm with and without a "prior" per experiment number of experiment 2. This table is a continuation of Table C.1 that describes experiment 1.

Exp. nr.	Agency accuracy, $\sum_j^k \mu_{\hat{b}_j} - b_j $		Perceived precision, $\sum_j^k \sigma_{\hat{b}_j}^2$		State accuracy, $\sum_j^2 \mu_{\epsilon_{x_j}}$		Perceived error, $\mu_{\epsilon_{x_a}}$	
	Prior	No prior	Prior	No prior	Prior	No prior	Prior	No prior
2A	7.57×10^{-1}	4.04×10^{-1}	2.69×10^{-3}	1.67×10^{-4}	1.26×10^0	2.42×10^{-1}	1.79×10^{-2}	6.11×10^{-1}
2B	7.87×10^{-1}	4.49×10^{-1}	3.79×10^{-3}	1.67×10^{-4}	1.24×10^0	2.43×10^{-1}	1.95×10^{-2}	6.12×10^{-1}
2C	6.21×10^{-1}	2.97×10^{-1}	4.12×10^{-3}	5.70×10^{-4}	1.21×10^0	2.42×10^{-1}	6.38×10^{-1}	1.30×10^0

D

RESULTS FORM POINT MASS MRHI WITH EM

This appendix shows the numerical results from the point mass mRHI from Chapter 5, where the agent estimates the agency matrix $\hat{\mathbf{B}}$ with EM.

D.1. EXPERIMENT 1: UNAMBIGUOUS ENVIRONMENTS

For the unambiguous environments of experiment 1, Table D.1 compares the results of the point mass mRHI with and without prior knowledge to see if the prior improves the agent's SoA. The table compares the results according to four performance measures for an agent with and without a prior described in Sec. 3.3.2. The algorithm with prior is the EM algorithm and the algorithm without prior is the OLS algorithm (see Figure 4.2 in Sec. 4.1).

With the prior the agent is more accurate (experiment 1) and more precise (experiment 1A and 1B) in estimating the agency matrix $\hat{\mathbf{B}}$ improves. The increase in noise between experiment 1B and 1C, increases the variance of the agency matrix $\hat{\mathbf{B}}$ by approximately a factor 10 for the agent with a prior, but barely increases the variance for an agent without a prior. However, without the prior, the agent's variance barely increases (4.96×10^{-4} vs. 5.01×10^{-4}) despite the inaccurate estimation of the matrix (1.51×10^0). Thus, overall the prior moves the agency estimation to the correct matrix \mathbf{B} .

However, the prior does not improve the accuracy of the estimation of the motion states of the mass. This indicates that, while the agent's overall estimation of the agency matrix $\hat{\mathbf{B}}$ improves, the agent's estimation of the agency parameter \hat{b}_{i1} deteriorates.

Table D.1. This table shows for unambiguous environments the agent's agency improves with a prior w.r.t. without a prior. The agent with prior knowledge uses the EM algorithm (see Figure 5.6). The agent without prior knowledge, only performs OLS to estimate the agency parameter, $\hat{\mathbf{B}}$. The columns show the agent's performance according to the different performance measures described in Sec. 3.3.2. The green background color indicates where the agent's performance is better between the prior and no prior conditions.

Exp. nr.	Agency accuracy, $\sum_j^k \mu_{\hat{b}_j} - b_j $		Agency variance, $\sum_j^n \sigma_{\hat{b}_j}^2$		Accuracy mass $m_1, \sum_j^2 \mu_{\varepsilon_{\hat{x}_j}}$		Perceived error, $\mu_{\varepsilon_{\hat{x}_a}}$	
	Prior	No prior	Prior	No prior	Prior	No prior	Prior	No prior
1A	1.42×10^{-6}	8.56×10^{-6}	5.88×10^{-8}	1.49×10^{-7}	2.42×10^{-2}	2.42×10^{-2}	1.20×10^{-2}	7.27×10^{-2}
1B	8.93×10^{-4}	2.62×10^{-1}	9.59×10^{-5}	4.96×10^{-4}	2.50×10^{-2}	2.42×10^{-2}	1.28×10^{-2}	1.20×10^0
1C	4.33×10^{-2}	2.62×10^{-1}	1.42×10^{-3}	5.01×10^{-4}	3.20×10^{-1}	2.42×10^{-1}	1.26×10^{-1}	1.41×10^0

D.2. EXPERIMENT 2: AMBIGUOUS ENVIRONMENTS

This section provides some extra results for the ambiguous environments defined by experiment 2.

D.2.1. TABLE OF THE PERFORMANCE MEASURES

For the unambiguous environments of experiment 2, Table D.2 compares the results of the point mass mRHI with and without prior knowledge to see if the prior induces an incorrect SoA. This table is a continuation of Table D.1.

Table D.2 shows that the agent's estimation of the agency matrix $\hat{\mathbf{B}}$ is generally less accurate and precise than the OLS algorithm without prior for the synchronous ambiguous environments (experiment 2A and 2B). The decrease in accuracy and increase in variance with respect to both experiment 1 (in Table D.1) and the no prior OLS algorithm (in Table D.2) indicates that the agent incorrectly chooses to gain SoA over mass m_2 . The agent's perceived error with the EM algorithm is around the same magnitude as in experiment 1 and smaller than the perceived error with the OLS algorithm. This means that while the agent's SoA is incorrect, the agent does have a SoA over mass m_2 . One exception to the incorrect SoA in experiment 2 is the asynchronous environment (experiment 2C). In experiment 2C, the agent seems to improve its SoA when compared to the no prior algorithm.

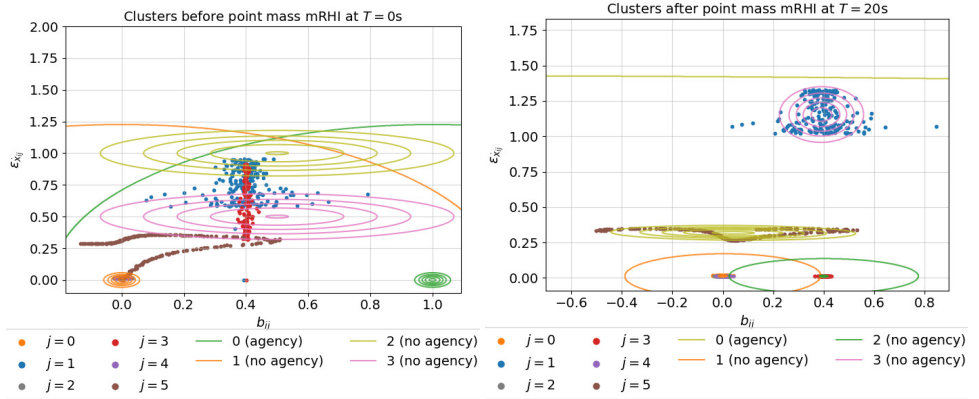
Table D.2. This table shows for ambiguous environments the agent's agency deteriorates with a prior w.r.t. without a prior. The calculation are performed in the same way as Table D.1.

Exp. nr.	Agency accuracy, $\sum_j^k \mu_{\hat{b}_j} - b_j $		Agency variance, $\sum_j^n \sigma_{\hat{b}_j}^2$		Accuracy mass $m_1, \sum_j^2 \mu_{\varepsilon_{\hat{x}_j}}$		Perceived error, $\mu_{\varepsilon_{\hat{x}_a}}$	
	Prior	No prior	Prior	No prior	Prior	No prior	Prior	No prior
2A	7.22×10^{-1}	4.04×10^{-1}	3.27×10^{-3}	1.67×10^{-4}	1.18×10^0	2.42×10^{-1}	6.99×10^{-2}	6.11×10^{-1}
2B	7.39×10^{-1}	4.49×10^{-1}	5.19×10^{-3}	1.67×10^{-4}	1.15×10^0	2.43×10^{-1}	8.85×10^{-2}	6.12×10^{-1}
2C	2.62×10^{-1}	2.97×10^{-1}	1.95×10^{-2}	5.70×10^{-4}	5.95×10^{-1}	2.42×10^{-1}	1.73×10^{-1}	1.30×10^0

D.2.2. ESTIMATION OF THE PARAMETERS THE CLUSTERS

The agent in the ambiguous environments (experiment 2) starts with the same initial conditions as in the unambiguous environments (experiment 1). Despite that the initial conditions are the same, the agent's SoA is different. I will explain how the interaction of the prior and the sensory evidence results in a different duration of incorrect SoA over mass m_2 (see Table 5.4).

For the synchronous ambiguous environment (experiment 2A), Figure D.1b shows that the agent's incorrect SoA over the observations from $j = 3$ because the green contours overlap the data points from $j = 3$ (see Box 5.1 for an explanation of the plots). Again, the agent has a SoA over the data with a positive b_{ij} and a low prediction error, $\varepsilon_{\hat{x}_{ij}}$. This contrasts with Figure D.1a of the initial time window of data, δ_N at $t = 0s$, where the agent's initial parameters for the agency cluster do not overlap any data points. Thus, with time the agent gains SoA over the incorrect data points.

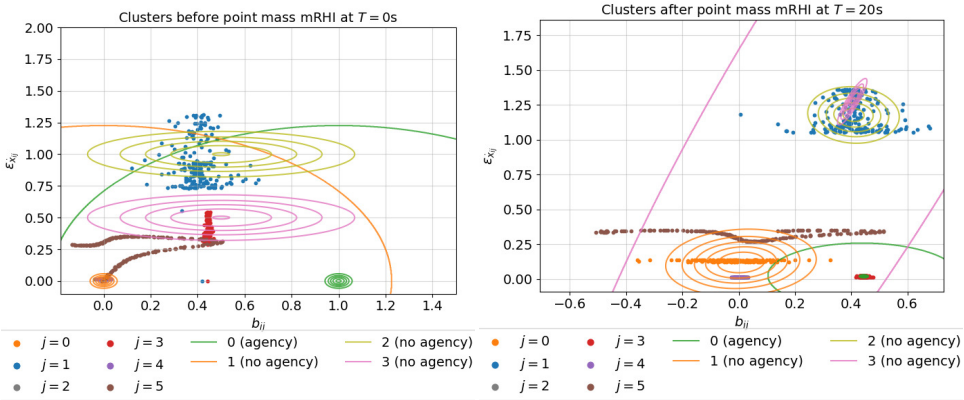


(a) The initial time window of data and clusters before of the simulation ($T = 0s$) (b) The final time window of data and clusters at the end of the simulation ($T = 20s$).

Figure D.1. This figure shows that in a synchronous ambiguous environment (experiment 2A) the agent gains an incorrectly SoA over the data from $j = 3$. The lines represent the probability density distributions of the agency and no-agency clusters. The dots represents the data for the prediction error, ε_{x_j} , and observed b_{ij} for each j . The initial conditions, in Figure a, do not promote the agent to gain SoA over the data points from $j = 1$ because the agent believes that the prediction error of the states \hat{x} should be low. Figure b shows that the agent incorrectly believes that the hidden cause of the red dots ($j = 3$), is its own force input, F_{int} , because the agency probability density distribution (green lines) tries to describe these data points.

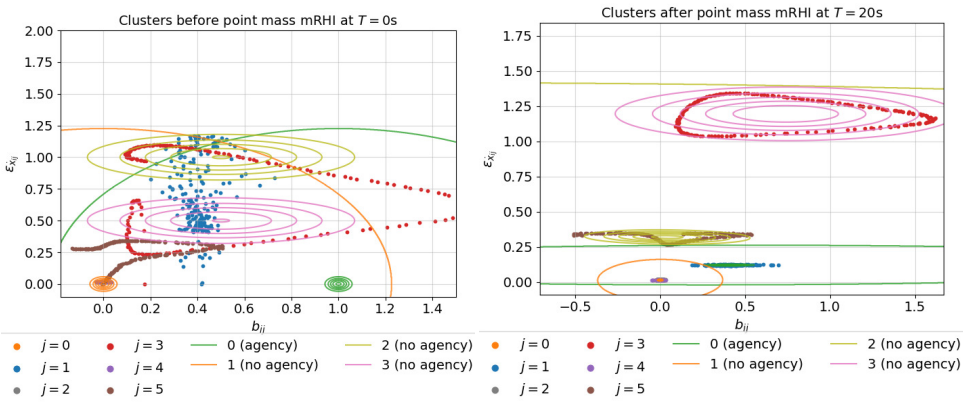
Figure D.2 explains why in the ambiguous environment where mass m_2 is lower (experiment 2B), the agent has a more SoA over mass m_2 than in experiment 2A. The agent compensates for the decrease in the mass of m_2 with respect to mass m_1 , by estimating parameter \hat{b}_{i1} with a larger value (recall, $b_j = 1/m_j$). Therefore, in experiment 2A, the plant's $b_2 = b_4$ whereas in experiment 2B the plant's $b_2 \neq b_4$. In Figure D.1a, we see the observations from $j = 1$ are initially more overlapped with the observations from $j = 3$. To explain this overlap, the agent in the ambiguous environment (2A) initially tries to describe the observations with *one* probability distribution. In other words, in experiment 2A the agent initially believes that the data from mass m_1 and m_2 have the same hidden cause whereas in experiment 2B the agent can more easily distinguish the observations that have a lower prediction error. Thus, the agent has more incorrect SoA over mass m_2 .

The results from the synchronous ambiguous environments (experiment 2A and 2B) contrast with the asynchronous ambiguous environment (experiment 2C) because the agent SoA becomes more correct with time. The shift of force input to mass m_2 by 500 ms, means the agent can no longer accurately predict the motion states, \hat{x} , of mass m_2 with its own input force F_{int} . Figure D.3a plots the initial higher prediction error of ε_{x_3} relative to ε_{x_1} . Therefore, the agency cluster describes the blue points of $j = 1$ in Figure D.3b. Hence, in experiment 2C, the agents prediction error minimization to gain SoA over *one* mass persuades the agent to eventually gain a correct SoA.



(a) The initial time window of data and clusters before of the simulation ($T = 0s$) (b) The final time window of data and clusters at the end of the simulation ($T = 20s$).

Figure D.2. This figure shows that in the synchronous ambiguous environment (experiment 2B) the agent incorrectly clusters the data according to their hidden causes with the EM algorithm. Despite the 10% decrease in mass m_2 with respect to mass m_1 , from 2.5 kg to 2.25 kg, the agent still has SoA over mass m_2 . The decrease in mass result in smaller amplitudes of the motion states, \hat{x} , which the agent can compensate for by increasing its estimation of \hat{b}_3 . Thus, the prediction error $\varepsilon_{\hat{x}_4}$ remains low.



(a) The initial time window of data and clusters before of the simulation ($T = 0s$) (b) The final time window of data and clusters at the end of the simulation ($T = 20s$).


Figure D.3. This figure shows that in this trial of experiment 2C the agent correctly clusters the data according to their hidden causes with the EM algorithm. This is a result of the time lag of the force input of mass m_2 w.r.t. mass m_1 .

E


THE INFLUENCE OF THE INITIAL COVARIANCE MATRIX ON SOA

The EM algorithm is known to be sensitive to the initial conditions. In the point mass mRHI, the initial covariance matrix of the agency cluster, Σ_k , is important because it quantifies the agent's initial belief of the strength of its agency. Therefore, I predict that the estimation accuracy and variance of parameter $\hat{\mathbf{B}}$ is dependent on the initial covariance matrix. Figure E.1 plots the average error in the estimated \hat{b}_{i1} per trial for different initial covariance matrices. I test 15 different values for the diagonal elements of the initial covariance matrix between 10^{-6} and 10^{-1} that are equally spaced over a logarithmic scale.

In both experiment 1B and 1C, when the initial covariance matrix is larger than around $10^{-2.8}$ the error significantly increases. For the lower initial covariance values, the resulting error varies but remains relatively low. Larger initial values for the elements of the covariance matrix means the agent initially believes that it has a less precise agency parameter, which translates to a lower SoA. This shows that the belief of strength of the causality at the beginning of the point mass mRHI influences the agent's agency accuracy. Moreover, this indicates that the changing the initial elements of the covariance matrix to from the current 10^{-3} to $10^{-3.5}$ will improve the agent's SoA.



(a) Experiment 1B.



(b) Experiment 1C.

Figure E.1. This figure shows that the estimation error of \hat{b}_{i1} is dependent on the initial covariance matrix, Σ , of the agency cluster. The plot shows each of the 10 trial, where the mean of the 10 trials is represented by a dashed line. I also plot the mean estimation error for \hat{b}_1 with OLS.