POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Active Inference: Modelling Motor Control under Multisensory Conflict

Tesi di Laurea Magistrale in
Biomedical Engineering - Ingegneria Biomedica

Author: **Alessandro Arrigoni**

Student ID: 227569
Advisor: Prof. Marta Gandolla
Co-advisors: Pietro Donelli
Academic Year: 2024-25

Abstract

The Free Energy Principle (FEP) and Active Inference represent a promising framework to model motor control. The FEP stems from Helmholtz's proposal that the brain's cognitive processes are not the result of direct elaboration of sensory data; rather, they are the outcome of probabilistic inference about the hidden causes causing the perceived sensations. Starting from classical Bayesian inference, Variational Free Energy (VFE) can be derived as a tractable quantity to be minimized in order to maintain homeostasis. The Active Inference framework suggests that such minimization can occur both through perceptual adaptation and action. Such framework is especially promising since it is able to model specific behaviors that remain widely unaddressed by other motor control theories; specifically, subtle movements are observed in the presence of multisensory conflict, a condition in which different sensory channels provide conflicting information about the environment. An example is represented by the Rubber Hand Illusion, where subjects experience a sense of body ownership over a rubber arm that is placed in a plausible configuration near their real arm (which is concealed from view). During such experiment, subjects perform subtle movements of the real arm towards the fake limb; these movements have the underlying goal of resolving the multisensory conflict arising from their proprioceptive and visual channels, which provide contrasting sensorial information about the arm configuration. The aims of this work are to provide a clear theoretical and analytical description of the Free Energy Principle; moreover, a simplified simulation environment is executed to demonstrate the feasibility of the active inference paradigm both during standard goal-oriented task and while experiencing multisensory conflicts.

Keywords: Free Energy Principle, Active Inference, multisensory conflict, Rubber Hand Illusion

Abstract in lingua italiana

Il Principio di Energia Libera (FEP) e l'Active Inference costituiscono un framework promettente per la modellazione del controllo motorio. Basandosi sulla prospettiva di Helmholtz, il FEP suggerisce che i processi cognitivi non derivino da una semplice elaborazione di dati sensoriali, ma siano il risultato di un'inferenza probabilistica sulle cause nascoste che generano le percezioni. Partendo dall'inferenza bayesiana classica, l'Energia Libera Variazionale (VFE) viene definita come una quantità trattabile che l'organismo minimizza per mantenere l'omeostasi. L'Active Inference propone che tale minimizzazione avvenga sia attraverso l'adattamento percettivo che tramite l'azione. Questo approccio è particolarmente rilevante poiché spiega comportamenti spesso ignorati da altre teorie del controllo motorio, come i micro-movimenti che emergono in presenza di conflitti multisensoriali. Un esempio emblematico è la Rubber Hand Illusion (RHI), in cui i soggetti percepiscono come propria una mano artificiale posta vicino al braccio reale (nascosto alla vista). In questo scenario, si osservano piccoli spostamenti del braccio reale verso quello finto: movimenti che mirano intrinsecamente a risolvere il conflitto tra il canale propriocettivo e quello visivo. L'obiettivo di questa tesi è fornire una rigorosa descrizione teorica e analitica del FEP. Inoltre, viene presentato e implementato un modello di simulazione semplificato per dimostrare l'efficacia del paradigma dell'Active Inference, sia nell'esecuzione di compiti orientati a un obiettivo (goal-oriented), sia durante l'esperienza di conflitti multisensoriali. Parole chiave: Principio di Energia Libera, Active Inference, conflitto multisensoriale, Rubber Hand Illusion.

Parole chiave: Principio di Energia Libera, Inferenza Attiva, conflitto multisensoriale, illusione della mano di gomma

Contents

Introduction

Neuroscience research is actively searching for a unified brain theory capable of integrating the explanation of the brain's various cognitive processes and its ability to interact with the external environment. A promising candidate is the Free Energy Principle (FEP) first introduced by Friston as a general theoretical framework for brain and behaviour [12?]. The FEP aims at explaining perception, action and learning through the minimisation of a quantity called Free Energy.
The origin of FEP stems from Helmholtz's proposal that perception is a cognitive process which is equivalent to a probabilistic inference of the hidden causes of sensory data [46]. In the Bayesian Inference formulation this consists of computing $p(x|s)$, the posterior probability of the causes x given the available sensory signals s. Consult the List of Symbols for a complete description of the notation adopted throughout this work. Helmholtz is also the father of the notion of thermodynamical free energy, a quantity that is minimized in systems at thermal equilibrium; this may be another inspiration in the formulation of Free Energy in the context of neuroscience.
Indeed, the FEP stems from the simple observation that all living organisms must resist the inevitable natural tendency to disorder imposed by the second principle of thermodynamics in order to assure survival. In other words, living organisms must minimize the occurrence of atypical events in their natural habitat in order to maintain their homeostatic properties. For instance, high temperature is a surprising event for a polar bear (that is well adapted to cold environments) which may disrupt its homeostasic processes and lower its chances of survival.
Since the distribution of 'surprising' events is generally unknown and unknowable, organisms need a tractable quantity to be evaluated instead. According to the FEP, this quantity is Free Energy which must be minimized to improve chances of survival.

Content summary

The present work aims at exploring the FEP and providing a proof of concept simulation in a motor control context. Specifically, in Chapter 1 the Free Energy Principle is presented as a theoretical framework explaining the regulatory mechanism behind survival tactics exploited by living organisms to maintain their homeostasis. Active inference is then presented as the practical implementation of the FEP which minimizes Free Energy through perceptual adaptation and action selection. Furthermore, the chapter focuses on specific scenarios which are well addressed by FEP and active inference.

In Chapter 2 the analytical formulation of Free Energy is derived starting from classical Bayesian Inference; moreover, some approximation are introduced to provide a tractable formulation of Free Energy, that can be employed in practical implementations.

In Chapter 3 the context of the simulations implemented is described. The objective of such simulations is to provide a proof of concept implementation of an Active Inference agent performing motor control under different bodily illusions.

In Chapter 4 the results of the simulation presented in the previous chapter are presented and discussed. Finally, future developments and final considerations are drawn in Chapter 5.

The appendices are organized as follows: Appendix A contains a brief overview of the main mathematical tools necessary for the analytical derivation of Free Energy; in Appendix B alternative formulations of Free Energy are presented and the mathematical motifs behind its definition are discussed; finally, Appendix C provides the code utilized for the implementation of the simulations presented in Chapter 3.

1 | The Free Energy Principle

In this chapter the theoretical framework of the Free Energy Principle (FEP) is discussed.

At its core, the FEP interprets cognition processes as means by which biological organisms resist the natural tendency towards thermodynamic equilibrium with their surroundings, i.e. maintain homeostasis. This interpretation stems from the simple observation that a living organism is compelled by its survival instinct to minimize the dispersion of its constituent states; failure to do so would compromise their homeostatic integrity. For instance, a human being must maintain its body temperature at approximately 37 °C; all the organic processes in their body are designed to operate around this temperature. Therefore, fluctuations around this value must be minimal; a broader dispersion would disrupt all the physiological processes, eventually jeopardizing the organism's chances of survival.

Equivalently, it can be stated that a living organisms must avoid atypical events to increase its chances of survival. In the example above, a human being should avoid a body temperature of 40 °C to maintain the organism's homeostatic properties. The atypicality of an event can be measured through a quantity called surprise:

$$I(s) = -\ln p(s), \tag{1.1}$$

where $p(s)$ denotes the probability of observing some particular sensorial pattern s. Consult Section A.1 for a thorough definition of surprise. This quantity assumes large values when the chances of observing a particular sensorial pattern are low; while it decreases the more likely a certain sensation is. In the cited example, surprise associated to the sensations derived from a body temperature of 37 °C will be low; while it will be higher when the sensations the brain experiences result from a body temperature of 40 °C. Note that surprise is evaluated on the sensorial pattern of information that is available to the brain and not on the environment's state variables which are hidden and inaccessible. In the example: the brain doesn't know the exact value of the body's temperature (the state variable); instead it is able to recognize the pattern of sensations that a certain value of body temperature would provoke.

Indeed, the FEP stems from the Helmholtz's perspective according to which perception is the result of probabilistic inference: the brain infers the configuration of the hidden causes of the sensory data it receives. Within the FEP, the brain can be modeled as an Active Inference agent that interacts with its surrounding environment (which could be its own body).

In this framework, a living organism has the objective of minimizing surprise to avoid atypical events that would endanger its chances of survival. However, surprise is an intractable quantity since the probabilistic distribution of all possible sensorial patterns is generally inaccessible. Instead, it must minimize an upper bound of it called Free Energy. Minimizing this quantity would, in turn, lead to a reduction in surprise.

1.1. Active Inference

A clear understanding of the framework within which the FEP operates is required before proceeding with its analytical derivation.

Firstly, a distinction between the Free Energy Principle and Active Inference must be made: while FEP is the general theoretical principle explaining how living organisms should act to ensure their survival, Active Inference is the practical framework employed by self-organizing agents (such as living beings) to minimize Free Energy. Indeed, the FEP postulates the existence of a quantity called Free Energy to be minimized; Active Inference is the framework defining the modalities that could be adopted by an agent to minimize Free Energy: action and perceptual adaptation.

Any implementation of the Active Inference framework requires the definition of two interacting systems: the agent and the environment. Note that in the case the agent is the brain, the environment consists also of the organism's body that must be controlled.

Specifically, the environment is referred to as the 'generative process': it describes how the physical system evolves in time (typically through the laws of physics) and specifies the mapping into the sensory inputs observed by the agent. On the other hand, the agent is referred to as the 'generative model': the agent holds an internal model of the generative process' dynamics along with their expected mappings into sensorial input.

Although the generative process corresponds to the true external dynamics that give rise to sensory input (i.e., the environment and its causal structure), the generative model refers to the agent's internal probabilistic representation of these dynamics. Crucially, these two do not need to share the same form or level of complexity. In practice, the agent's internal model often abstracts or compresses the causal structure of the environment into latent variables — hidden states that capture only the information most relevant for prediction and action. This abstraction allows the agent to operate efficiently: rather than simulating the full external world in detail, it maintains an approximate internal model sufficient to minimize surprise and maintain homeostasis. From the perspective of Active Inference, this mismatch is not a flaw but a feature since it enables generalization and flexibility. The agent does not aim to reproduce the environment exactly but to construct an internal model that is functionally adequate for predicting sensory outcomes and guiding adaptive behavior.

1.2. Free Energy Minimization

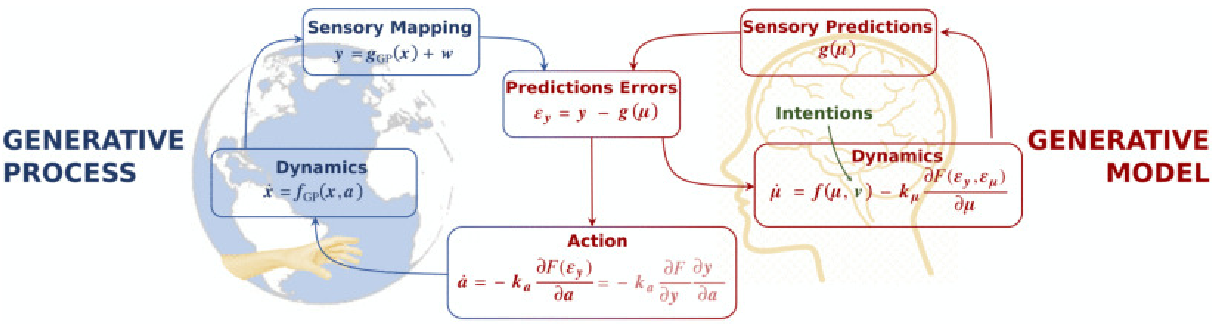

Figure 1.1: Active Inference Framework [38]

The estimation and prediction of the generative process' dynamics allow the agent to compute prediction errors that are exploited to minimize Free Energy – hence surprise. The agent can minimize Free Energy through two parallel and complementary modalities: action and perceptual adaptation.

The former implies that the agents exerts a degree of control over the dynamics of its environment. By altering them, the agent is able to modify the actual sensorial feedback so that it conforms to the predicted outcomes encoded by the generative model. This modality is clearly exemplified in the context of goal-directed behavior: for instance, during a reaching task the generative model predicts future sensorial inputs expected when the movement is executed correctly. The agent then performs the reaching movement to minimize prediction errors between the expected and the actual sensory feedback; in other words, the agent actively infers the optimal action to be exerted on the environment in order to manipulate its dynamics so that the actual sensorial input received follows the trajectory of the internal estimation.

The latter entails the continuous updating of the agent's beliefs about the dynamics of the system it interacts with. Upon receiving novel sensory input, the agent revises the generative model to better reflect observed data. Through this process, the generative model progressively improves its ability to predict future sensory states, thereby reducing the discrepancy between expected and actual sensory outcomes.

A full visual description of the Active Inference framework is well represented in Figure 1.1.

1.3. Comparison with optimal motor control

It is worthwhile to provide a synthesis of the biological foundations of Active Inference and conduct a comparative analysis with Optimal Control theory, the main framework classically adopted to model motor control. Optimal Control is built upon the premise that the brain acts as an efficiency agent: it must minimize a cost function that typically takes into account the accuracy of movements and the effort needed to perform them. Thus the brain acts as a controller of the body inserted in a feedback loop between sensory signals coming from the ascending or afferent pathways and motor commands sent through descending or efferent pathways towards the end-effectors (typically muscles) [44].

At the cognitive level, the brain performs a comparison between desired and actual states of the end-effectors in order to generate a motor command to minimize this discrepancy. This task is performed by the inverse model the brain is equipped with: thus, the inverse model maps the desired spatial trajectory into the specific pattern of neuromuscular activity required to achieve it. Notably such task is non-trivial since multiple muscle activation patterns could result in the desired outcome. The resulting signal is a motor command sent through the spinal cord towards the ultimate end-effector. To overcome sensory and motor delays, the brain is also equipped with a forward model, which receives an efferent copy of the motor command and infers the resulting future states and sensations. Therefore, the inverse model actually performs a comparison between the desired and inferred states; the difference between inferred and actual states could be exploited as an error signal to train the forward model. The left panel in Figure 1.2 provide a schematic diagram of optimal motor control theory.

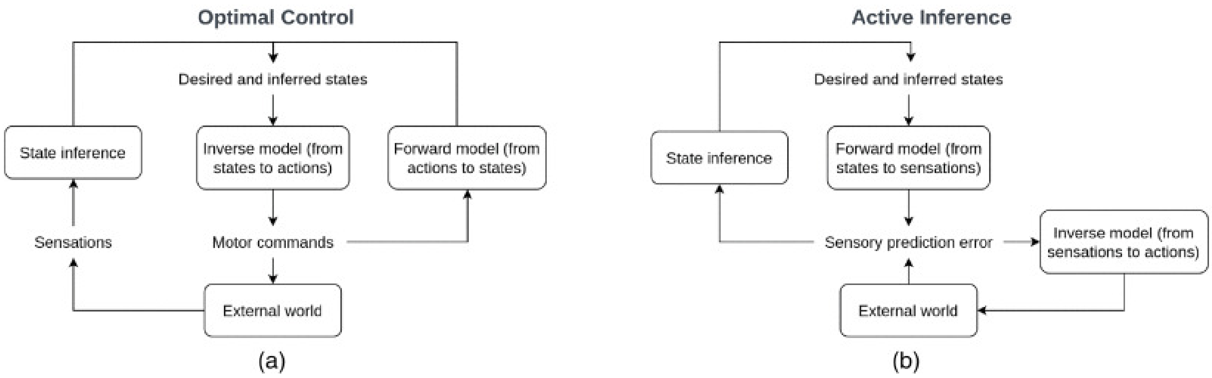

Figure 1.2: Comparison between optimal motor control and active inference [38]

Instead, the right panel of Figure 1.2 provides the schematic diagram of Active Inference. In this framework the brain exploits a forward model (i.e. the generative model) to perform a prediction of sensory patterns from inferred states (i.e. the internal states in the generative model). This prediction is supposed to be the signal sent through the spinal cord, where a low-level comparison between predictions and actual sensory patterns is performed; afferent pathways are then exploited to back-propagate the errors towards cortical areas to train the forward model. In Active Inference, the inverse model is implemented in the arc reflexes at the level of the anterior horns of the spinal cord. Thus, arc reflexes are supposed to intrinsically implement the mapping between prediction errors and motor commands that must be sent to the end-effectors [1, 13].

To summarize, Active Inference and Optimal Control main differences consist in the interpretation of the brain's role and activity and the placement of the inverse model in the central nervous system. Indeed, the Active Inference brain performs prediction over sensory patterns that are compared and mapped into motor commands by low-level arc reflexes in the spinal cord; otherwise, Optimal Control places the inverse model in cortical areas of the motor cortex and suppose the brain to directly produce motor commands.

1.4. Applications

The Active Inference framework described so far has been implemented in a vast variety of scenarios to model problems regarding goal-directed behavior, active sensing, movement planning, pathological behavior. This section provides a brief overview of the main applications of such paradigm.

As an emerging motor control theory, the first benchmark test requires Active Inference to properly model the execution of reaching tasks. In this context, Active Inference serves a dual-purpose: it provides a framework that is able both to model the behavior observed in living organisms and to act as a control paradigm in artificial implementations in the robotics field. For instance, multiple studies [32, 35] propose Active Inference as plausible paradigm implemented by the brain to control eye movements. Particularly, both target fixation and smooth pursuit are modeled. Target fixation consists of maintaining the gaze on a stationary target (a reaching task performed by the eyes); while smooth tracking is the eye movement that allows a moving target to remain stabilized on the fovea by matching the eyes movements with the relative motion between the subject and the target. Notably, in [32] Active Inference was able to model saccades, the sudden and rapid movements produced by the eyes when the target location is far enough from the actual fixation point. Otherwise, [35] investigated the accuracy of an Active Inference model in the presence of sensory and/or motor delays during target pursuit.

The same framework was also introduced in the context of robotics to appropriately control robots during the execution of reaching tasks. For instance, in [37] a robotic arm endowed with 7 degrees of freedom was required to reach a stationary target. The generative model predicted both proprioceptive and visual signals which were then compared to actual signals received by the generative process to compute prediction errors. As will be later discussed in Chapter 2, minimizing Free Energy is equivalent to minimizing prediction errors under plausible approximations. The control paradigm was tested under different noise conditions for the proprioceptive and visual sensory channels, demonstrating a better performance with noisy vision rather than with noisy proprioception. In [31] Active inference was deployed in a humanoid robot to have adaptive body perception and to perform reaching tasks with its upper body and object tracking with its head. Notably, in all the cited studies the presence of a target was modeled in the generative model dynamics as an attractive force that is proportional to the distance between the target and the end-effector; the same strategy will be adopted in the simulations presented in Chapter 3.

Active Inference was also employed to model motor behaviors that remain widely unaddressed by other theories. For instance, other studies were able to implement an active inference framework to model both physical and psychological pathological behavior [22, 26, 33]. Moreover, an Active Inference agent can model active sensing behavior: the process by which an agent actively moves its own sensors in order to gather the most useful information, rather than receiving passively stimuli [38]. An example is represented by the whicker movements performed by rodents in order to scan the surrounding environment: particularly, whiskers display wide anticipatory movements towards expected object location and reduced scanning movements when in touch with an object [45]. In

[28], such behavior was successfully modeled through Active Inference: without any contact the agent progressively enlarged the amplitude of the whiskers' movements to explore the surrounding space; when a contact occurred the movement amplitude was rapidly reduced to scan the object.

Notably, other motor control theories fail to adequately account for multisensory conflicts, as discussed in the next section.

1.5. Multisensory conflicts

The human body must continuously reconstruct a coherent internal representation of both the environment and the self. This objective is achieved by integrating information originating from different sensory channels, comprising both exteroceptive and interoceptive biological sensors that it is equipped with [24]. For instance, the complete representation of an object is constructed by integrating information coming from its sight with tactile information when the subject interacts with it. The richness and completeness of such representation is dependent on the integration of the information of the two channels: exploiting only one channel may be insufficient to explore the variety of features of said object. For example, in the absence of tactile information, a rubber ball and a metal ball of the same size and color may appear indistinguishable. The process of constructing a unified internal representation of the environment by combining information coming from different sensory channels is referred to as multisensory integration [7].

While multisensory integration generally enhances perceptual accuracy and robustness, it also introduces the possibility that information from different sensory channels may provide inconsistent data. Such situation is referred to as multisensory conflict, where two or more sensory channels provide discrepant information about the same object, event or state of the body. The discrepancy may concern any level of description of the environment, ranging from low-level sensory dynamics to high-level abstract representations. Moreover, multisensory conflicts may arise from a variety of sources, including bodily illusions, sensorimotor perturbations or artificial manipulation of the sensory feedback (such as the embodiment of a virtual avatar in VR environment) [26].

When confronted with multisensory conflicts, the brain cannot simply integrate such incongruent information. First, it must resolve the conflict in order to preserve a unified and coherent internal representation of the environment. To this aim, the brain can adopt two main strategies: either adapt its internal representation of the environment (perceptual adaptation) or trigger specific movements to reduce the disparity between conflicting data.

Multisensory conflicts remain widely unaddressed in literature; this is rooted both in limitations of motor control theories and in the difficulty of disentangling such aspect of motor control from others (e.g. goal-directed behavior). Indeed, resolving multisensory conflicts through action requires a null space in the motor mapping between the controlled environment and at least one sensory channel [30]. Equivalently, the effect of the action must have no discernible impact on the dynamics of one of the conflicting sensory channels. A representative example is provided in the next section, which examines one of the most prominent multisensory illusions, namely the Rubber Hand Illusion (RHI).

1.5.1. The Rubber Hand Illusion

The Rubber Hand Illusion (RHI) is a well-known experimental paradigm that was first reported by Botvinick and Cohen in 1998 [27]. In the classical experiment setup, the participant's hand is hidden from view while a realistic rubber hand is placed in plausible anatomical configuration. To set the illusion, both the real and fake hand are stroked simultaneously, causing in the participant a sensation of embodiment over the rubber hand. This sensation is a perceptual adaptation strategy adopted by the brain to resolve the multisensory conflict arising from the illusion [21]. Indeed, the proprioceptive signal received by the agent is provided by the real arm, while the visual information is provided by the sight of the fake hand, which is placed in a different configuration with respect to the real arm. Obviously, the discrepancy between the two must be minimal for the illusion to be maintained: larger discrepancies would break the illusion. More specifically, this adaptive process is termed 'perceptual drift'

Following experimental studies reported another strategy employed by the brain to resolve the multisensory conflict through unintentional action [2]. Specifically, in experimental settings where the concealed real arm was left free to move, subtle movements towards the real arm were observed; equivalently, in settings where the arm was constrained the real limb exerted a steady force towards the rubber hand. The purpose of these type of movements is to reduce the discrepancy between the proprioceptive and visual signals provided to the brain. As before mentioned this is possible under the condition that the action provides minimal changes in one of the two channels: since the real arm is concealed, moving it produces no changes in the visual layout in front of the participants; therefore, the visual signal remains unchanged. This represents a necessary condition to reduce the multisensory signal: indeed, such movements align the proprioceptive information with the visual data received by the rubber hand and consequently reduce the conflict between the two channels. If this requirement is not met, it is not guaranteed that the multisensory conflict could be solved; this case is addressed in the third simulation in Chapter 3, where the agent embodies and controls a virtual arm with dynamics irreconcilable with the real ones.

The Rubber Hand Illusion and multisensory conflicts received minimal attention from a computational modeling perspective. Active Inference has demonstrated to be a suitable paradigm to account for this category of observed behaviors [23, 24, 26]. Specifically, in [30] a model that accounts for both reaching tasks and multisensory conflict resolving is presented: a key aim of the present work is to replicate the simulations and the results of such study, to propose Active Inference as a unified framework for intentional and conflict resolution purposes in motor control.

2 | Mathematical Formulation

In this chapter, the analytical formulation of the Free Energy Principle is presented, followed by a review of the approximations usually adopted in the context of practical implementations.

2.1. Analytical Formulation

A complete description of the notation adopted throughout this work is available at List of Symbols.

2.1.1. Bayesian Inference

According to Helmholtz's principle, the brain's cognitive processes are not the result of direct elaboration of sensory data; rather, they are the result of probabilistic inference on the causes of the sensory signals received [18]. In classic Bayesian Inference, this consists of computing the posterior probability of a certain configuration of the environment's hidden states x, given the sensory input s provided to the agent:

$$p(x|s) = \frac{p(s|x)p(x)}{p(s)}. \tag{2.1}$$

Here, $p(s|x)$ represents the likelihood of the sensory signal s given the environment's state configuration x and $p(x)$ represents the probability distribution of the state of the environment (i.e. the probability of the environment holding a certain configuration x) [43].

Although analytically correct, classical Bayesian Inference is practically untractable in this setting [4]. Indeed, the term $p(s)$ at the denominator represents the probability of receiving a given sensory input and should be computed marginalizing it over all possible state configuration of the environment:

$$p(s) = \int p(s|x)p(x)dx. \tag{2.2}$$

However, this operation remains intractable in most cases. For instance, when continuous functions are adopted to approximate the likelihood $p(s|x)$ and the prior $p(x)$, the integral may be analytically untractable. Moreover, in the case of a discrete distribution, the computational load associated to the sum may be impracticable [5]. In general, these probability distributions are unavailable to the brain.

2.1.2. Variational Free Energy

The Free Energy Principle offers a suitable alternative to Bayesian Inference. Its formulation is based on the assumption that the agent holds an auxiliary probability distribution $q(x)$, which is called recognition density. Such distribution represents the approximation of the posterior probability $p(x|s)$ held by the agent.

A measure of discrepancy between these distributions is provided by the Kullback-Leibler divergence between the two:

$$\mathcal{D}_{KL}\big[q(x) \,||\, p(x|s)\big] = \int q(x) \ln \frac{q(x)}{p(x|s)} dx. \tag{2.3}$$

As the recognition density $q(x)$ approximates the posterior probability $p(x|s)$, the value of their Kullback-Leibler divergence decreases, ultimately reaching zero when the distributions coincide exactly (see Section A.4 for a thorough definition of the Kullback-Leibler divergence). Therefore, the Kullback-Leibler divergence holds a measure of the agent's accuracy in approximating the true posterior probability $p(x|s)$. However, this quantity still cannot be evaluated since the true posterior is not knowable.

Therefore, Variational Free Energy is defined as an upper bound of this quantity by evaluating the Kullback-Leibler divergence between the recognition density $q(x)$ and the joint probability $p(x, s)$:

$$\mathcal{F} := \mathcal{D}_{KL}\big[q(x) \,||\, p(x, s)\big] = \int q(x) \ln \frac{q(x)}{p(x, s)} dx. \tag{2.4}$$

VFE's definition can be further arranged by applying the chain rule on the joint probability $p(x, s)$:

$$
\begin{aligned}
\mathcal{F} &= \mathcal{D}_{KL}\big[q(x) \mid\mid p(x,s)\big] \\
&= \int q(x) \ln \frac{q(x)}{p(x,s)} dx \\
&= \int q(x) \ln \frac{q(x)}{p(x|s)p(s)} dx \\
&= \int q(x) \ln \frac{q(x)}{p(x|s)} dx - \int q(x) \ln p(s) dx.
\end{aligned}
$$

Last line's first term can be recognized as the Kullback-Leibler divergence between the recognition density $q(x)$ and the posterior probability $p(x|s)$ as defined in Equation (2.3). Moreover, the second term can be rearranged by noting that $p(s)$ does not depend on x and, therefore, can be taken out of the integral; additionally:

$$\int q(x) dx = 1, \tag{2.5}$$

since the recognition density is a probability distribution.
Therefore, the following formulation of Variational Free Energy is derived:

$$\mathcal{F} = \mathcal{D}_{KL}\big[q(x) \mid\mid p(x|s)\big] - \ln p(s). \tag{2.6}$$

2.1.3. Considerations

According to Equation (2.6) Variational Free Energy is defined as the sum between the Kullback-Leibler divergence between the recognition density $q(x)$ and the posterior probability $p(x|s)$ and the surprise associated to the received sensory input. Since the Kullback-Leibler divergence is a non-negative quantity, VFE also represents an upper bound on surprise. Therefore, minimizing VFE will indirectly minimize both surprise and the divergence between $q(x)$ and $p(x|s)$; this will produce two major consequences: on one hand it assures that the recognition density $q(x)$ progressively provides a more accurate estimation of the posterior probability $p(x|s)$; on the other, the sensory surprise will be minimized accordingly with the Free Energy Principle.

Nonetheless, Variational Free Energy remains analytically untractable since it depends on the hidden states x; however, unlike its constituent single terms, it admits a simplified formulation which can be derived by introducing plausible approximations, most importantly the Laplace approximation.

Through different arrangements of Equation (2.4) different formulations of Variational Free Energy can be derived. Each of them is presented in Section B.1 along with the consequent considerations; also, the analytical equivalence between them is demonstrated.

2.2. The Laplace approximation

Although correct, the theoretical formulation described so far cannot be exploited in practice. In the context of the Free Energy Principle, the agent is endowed with a recognition density $q(x)$ approximating the posterior probability $p(x|s)$. As already noted, minimizing Variational Free Energy ensures a more accurate estimation of this probability distribution. However, a key issue concerns the specification of the distribution class that will be used to parameterize the recognition density $q(x)$.
Under the Laplace approximation, the recognition density can be approximated by a Gaussian explained by sufficient statistics:

$$q(x) = \frac{1}{\sqrt{2\pi\Sigma}} e^{-\frac{(x-\mu)^2}{2\Sigma}} \sim \mathcal{N}(x; \mu, \Sigma), \tag{2.7}$$

where μ and Σ are respectively the mean and variance values [15].

2.2.1. VFE's approximation

For formulation clarity, the univariate case will be presented, without loss of generality for the multivariate case. Moreover, the following parameters are adopted for notational ease:

$$Z = \sqrt{2\pi\Sigma}, \tag{2.8}$$

$$\varepsilon(x) = \frac{(x-\mu)^2}{2\Sigma}, \tag{2.9}$$

so that Equation (2.7) can be rewritten as:

$$q(x) = \frac{1}{Z} e^{-\varepsilon(x)}. \tag{2.10}$$

Firstly we expand Equation (2.4) as follows:

$$\begin{aligned}\mathcal{F} &= \int q(x) \ln \frac{q(x)}{p(x,s)} dx \\ &= \int q(x) \ln q(x) dx - \int q(x) \ln p(x,s) dx \\ &= \int q(x) \ln q(x) dx + \int q(x) I(x,s) dx, \end{aligned} \tag{2.11}$$

where in the last line the definition of surprise was recognized.

By introducing Equation (2.10) in the first term of Equation (2.11):

$$\begin{aligned}\int q(x) \ln q(x) dx &= \int q(x) \ln \frac{1}{Z} e^{-\varepsilon(x)} dx \\ &= \int q(x)\big(-\ln Z - \varepsilon(x)\big) dx \\ &= -\ln\sqrt{2\pi\Sigma} - \frac{1}{2\Sigma}\int q(x)\big(x-\mu\big)^2 dx \\ &= -\frac{1}{2}\ln 2\pi\Sigma - \frac{1}{2\Sigma}\mathbb{E}_q[(x-\mu)^2] \\ &= -\frac{1}{2}\ln 2\pi\Sigma - \frac{1}{2\Sigma}\Sigma \\ &= -\frac{1}{2}\ln 2\pi\Sigma - \frac{1}{2}, \end{aligned} \tag{2.12}$$

where in the third line Equation (2.8) and Equation (2.9) were reintroduced; while in the second to last line the variance definition of the Gaussian recognition density $q(x)$ was recognized.

Moreover, the second term of Equation (2.11) can be elaborated through a second-order Taylor expansion of $I(x,s)$ around μ under the assumptions that $q(x)$ is sharply peaked around its mean μ and that $p(x,s)$ is a smooth function of x:

$$\begin{aligned}\mathcal{I}(x,s) &\approx \mathcal{I}(\mu,s) + \left[\frac{\delta\mathcal{I}}{\delta x}\right]_{x=\mu}(x-\mu) + \frac{1}{2}\left[\frac{\delta^2\mathcal{I}}{\delta x^2}\right]_{x=\mu}(x-\mu)^2 \\ &= \mathcal{I}(\mu,s) + \frac{1}{2}\left[\frac{\delta^2\mathcal{I}}{\delta x^2}\right]_{x=\mu}(x-\mu)^2. \end{aligned} \tag{2.13}$$

Here, the first order term of the Taylor expansion is recognized to be null since the first derivative is evaluated around the mode of the distribution. By substituting into the second term of Equation (2.11):

$$
\begin{aligned}
\int q(x)\mathcal{I}(x,s)dx &\approx \mathcal{I}(\mu,s)\int q(x)dx + \frac{1}{2}\left[\frac{\delta^2\mathcal{I}}{\delta x^2}\right]_{x=\mu}\int q(x)\big(x-\mu\big)^2 dx \\
&= \mathcal{I}(\mu,s) + \frac{1}{2}\left[\frac{\delta^2\mathcal{I}}{\delta x^2}\right]_{x=\mu}\Sigma,
\end{aligned}
\tag{2.14}
$$

where in the last step Equation (2.5) and the definition of variance were adopted.

Finally, substituting Equation (2.12) and Equation (2.14) in Equation (2.11):

$$
\mathcal{F} = \mathcal{I}(\mu,s) + \frac{1}{2}\left(\left[\frac{\delta^2\mathcal{I}}{\delta x^2}\right]_{x=\mu}\Sigma - \ln 2\pi\Sigma - 1\right). \tag{2.15}
$$

Equation (2.15) can be further simplified by computing the optimal variance Σ^* that minimizes VFE:

$$
\begin{aligned}
\frac{\delta\mathcal{F}}{\delta\Sigma} = \frac{1}{2}\left(\left[\frac{\delta^2\mathcal{I}}{\delta x^2}\right]_{x=\mu} - \frac{1}{\Sigma}\right) &\stackrel{!}{=} 0 \\
\implies \Sigma^* = \left[\frac{\delta^2\mathcal{I}}{\delta x^2}\right]_{x=\mu}^{-1}&.
\end{aligned}
\tag{2.16}
$$

By substituting in Equation (2.15):

$$
\mathcal{F} = \mathcal{I}(\mu,s) - \frac{1}{2}\ln 2\pi\Sigma^*. \tag{2.17}
$$

2.2.2. Considerations

Note that by introducing the Laplace approximation, Variational Free Energy assumes the form of a function of the Gaussian recognition density $q(x)$'s statistics (μ and Σ) which is independent of the hidden state x. Therefore, Variational Free Energy is a function of the sensory input received and the belief held by the agent about the environment: remember that μ is the mean of $q(x)$, the agent internal representation of the posterior probability $p(x|s)$.

The second term depends solely on the optimal variance defined in Equation (2.16), which is the variance the agent should assign to the recognition density in order to minimize

Variational Free Energy under the Laplace assumption.
In the Active Inference framework that is later described, this term will have no influence since it is a constant. Indeed, minimizing the VFE entails computing its partial derivatives, which are unaffected by constant terms.

In literature, the first term of Equation (2.17) is usually referred to as the Laplace Energy associated to the joint probability $p(\mu, s)$ under the following notation:

$$\mathcal{L}(\mu, s) = -\ln p(\mu, s). \tag{2.18}$$

This notation stems from the consideration that Variational Free Energy is mainly composed by this term under the Laplace approximation. Indeed, the constant term is irrelevant in the context of minimization of VFE. Thus, semantically speaking, this term also represents 'energy' under the Laplace approximation.
The notation employed throughout this study is derived from the information-theoretic definition of surprise. This choice in notation is motivated by the consistency of terminology in this work. Note that although semantically richer, $\mathcal{L}(\mu, s)$ is mathematically identical to $I(\mu, s)$, since both operators denote the same quantity. Thus, this discrepancy in notation is irrelevant to the objectives of the present work.

2.3. The Generative Density

In Equation (2.17) an estimation of Variational Free Energy was derived under the Laplace approximation. Such formulation comprises two terms: the surprise associated to sensory input evaluated at the mode of the recognition density and a constant term. The latter constant term is disregarded, since it does not influence the minimization of the VFE. Further development of the former term can be derived by introducing a formulation of the approximated generative density (or G-density) $p(\mu, s)$. More specifically, the expressions of the terms constituting the G-density using the chain rule are next presented:

$$p(\mu, s) = p(s|\mu)p(\mu). \tag{2.19}$$

In Equation (2.19) the first term represents the likelihood of observing the sensory input received under the assumption that the hidden state x assumes the configuration described by the sufficient statistics of the recognition density $q(x)$. The second term represents the prior belief about the environmental state μ.

2.3.1. The univariate static case

Firstly, the simple case of a generative model believing in an explanation of the environment dynamics through a single hidden state feeding into a single sensory channel is presented. Specifically, the agent believes that sensory data is mapped through:

$$s = g(\mu, \theta) + n_s, \tag{2.20}$$

where g is a mapping function which in principle could assume any formulation parameterized by θ, while n_s represents a noise term.
Moreover, the agent beliefs about the environmental states assumes the following formulation:

$$\mu = \overline{\mu} + n_\mu, \tag{2.21}$$

where $\overline{\mu}$ is the static prior held by the agent about the hidden states of the environment while n_μ represents the noise term. Note that in this simple case we are assuming static behavior of the agent's prior beliefs suggesting that the agent believes the future hidden state configuration to be history independent. The noise terms in Equation (2.20) and Equation (2.21) are assumed to be Gaussian random variables with null mean and respective variances σ_s and σ_μ:

$$n_s \sim \mathcal{N}(0, \sigma_s^2),$$

$$n_\mu \sim \mathcal{N}(0, \sigma_\mu^2).$$

Therefore, an analytical description of the terms in Equation (2.19) can be derived:

$$p(s|\mu) = \frac{1}{\sqrt{2\pi}\sigma_s} \exp -\frac{(s - g(\mu, \theta))^2}{2\sigma_s^2}, \tag{2.22}$$

$$p(\mu) = \frac{1}{\sqrt{2\pi}\sigma_\mu} \exp -\frac{(\mu - \overline{\mu})^2}{2\sigma_\mu^2}. \tag{2.23}$$

By substituting Equation (2.19), Equation (2.22) and Equation (2.23) in the surprise term of Equation (2.17):

$$\begin{aligned}\mathcal{I}(\mu, s) &= -\ln p(\mu, s) \\ &= -\ln p(s|\mu) - \ln p(\mu) \\ &= -\ln\left[\frac{1}{\sqrt{2\pi}\sigma_s}\exp -\frac{(s-g(\mu,\theta))^2}{2\sigma_s^2}\right] - \ln\left[\frac{1}{\sqrt{2\pi}\sigma_\mu}\exp -\frac{(\mu-\overline{\mu})^2}{2\sigma_\mu^2}\right] \\ &= \frac{(s-g(\mu,\theta))^2}{2\sigma_s^2} + \frac{(\mu-\overline{\mu})^2}{2\sigma_\mu^2} + \ln 2\pi\sigma_s\sigma_\mu \\ &= \frac{1}{2\sigma_s^2}\epsilon_s^2 + \frac{1}{2\sigma_\mu^2}\epsilon_\mu^2 + \ln\sigma_s\sigma_\mu + \ln 2\pi, \end{aligned} \tag{2.24}$$

where in the last line the following auxiliary definitions were introduced:

$$\epsilon_s = s - g(\mu, \theta), \tag{2.25}$$

$$\epsilon_\mu = \mu - \overline{\mu}. \tag{2.26}$$

According to Equation (2.24), the surprise term constituting VFE consists of the sum of prediction errors and a constant depending on the variances σ_s and σ_μ. Indeed, from the definitions of ϵ_s and ϵ_μ in Equation (2.25) and Equation (2.26), it is clear that the former represents a prediction error between the actual sensory input received s and its inference performed by the agent $g(\mu, \theta)$, while the latter represents the deviation of μ from its prior expectation $\overline{\mu}$. Again, the constant terms will not influence VFE's minimization since their derivatives will be null. Thus, in this formulation minimizing VFE is equivalent to minimize the prediction errors ϵ_s and ϵ_μ.

2.3.2. The dynamic case

Here we assume a dynamical generative density on the agent's belief:

$$\frac{d\mu}{dt} = \mu' = f(\mu) + n_\mu, \tag{2.27}$$

where f is a function expressing the dynamics of the agent's beliefs μ. Traditionally, the FEP formulation in the dynamic case exploits generalized coordinates; this consists of expressing the dynamics of higher order derivatives of the hidden states and the consequent sensory input in order to have a complete temporal description of the trajectory followed by such quantities. Particularly, the dynamics of higher order derivatives can be obtained by computing the time derivative of Equation (2.20) and Equation (2.27):

$$s' = \frac{ds}{dt} = \frac{d}{dt}[g(\mu) + n_s] = \frac{\delta g}{\delta \mu}\mu' + n_s', \tag{2.28}$$

$$\mu'' = \frac{d\mu'}{dt} = \frac{d}{dt}[f(\mu) + n_\mu] = \frac{\delta f}{\delta \mu}\mu' + n_\mu'. \tag{2.29}$$

The same computation can be recursively performed on the last equations to derive the following general formulation for the derivative of n-th order:

$$s^{[n]} = \frac{d^n s}{dt^n} = \frac{\delta g}{\delta \mu}\mu^{[n]} + n_s^{[n]}, \tag{2.30}$$

$$\mu^{[n]} = \frac{d^n \mu}{dt^n} = \frac{\delta f}{\delta \mu}\mu^{[n-1]} + n_\mu^{[n-1]}. \tag{2.31}$$

Here nonlinear derivative terms like $\mu^{[n]}\mu^{[n-1]}$ or ${\mu^{[n]}}^2$ are neglected under the assumption of local linearity. Note that the definition in Equation (2.30) implies that the derivative of n-th order of s interacts with the same order derivative of the brain state μ, while Equation (2.31) defines the coupling between consecutive derivative orders of the brain state μ. Conventionally, all higher-order derivatives are stored in a vector denoting the generalized coordinates of a quantity:

$$\tilde{\boldsymbol{s}} = [s^{[0]}, s^{[1]}, s^{[2]}, ...], \tag{2.32}$$

$$\tilde{\boldsymbol{\mu}} = [\mu^{[0]}, \mu^{[1]}, \mu^{[2]}, ...]. \tag{2.33}$$

The generalized coordinates are thus an infinite-dimensional vector; in such space, a point encodes the instantaneous trajectory of a brain state described by infinite time derivatives. The same notation is exploited also for the noise terms and the functions g and f:

$$\tilde{\boldsymbol{n}}_{\boldsymbol{s}} = [n_s^{[0]}, n_s^{[1]}, n_s^{[2]}, ...], \tag{2.34}$$

$$\tilde{\boldsymbol{n}}_{\boldsymbol{\mu}} = [n_\mu^{[0]}, n_\mu^{[1]}, n_\mu^{[2]}, ...], \tag{2.35}$$

$$\tilde{\boldsymbol{g}} = [g^{[0]}, g^{[1]}, g^{[2]}, ...], \tag{2.36}$$

$$\tilde{\boldsymbol{f}} = [f^{[0]}, f^{[1]}, f^{[2]}, ...], \tag{2.37}$$

where

$$g^{[n]} = \frac{\delta g}{\delta \mu}\mu^{[n]}, \tag{2.38}$$

$$f^{[n]} = \frac{\delta f}{\delta \mu}\mu^{[n]}. \tag{2.39}$$

Next, the joint probability $p(\tilde{\boldsymbol{s}}, \tilde{\boldsymbol{\mu}})$ is specified as the product of the likelihood of the sensory data in generalized coordinates $p(\tilde{\boldsymbol{s}}|\tilde{\boldsymbol{\mu}})$ and the agent's prior $p(\tilde{\boldsymbol{\mu}})$. Here we assume statistical independence of noise terms at each dynamical order; therefore, $p(\tilde{\boldsymbol{s}}|\tilde{\boldsymbol{\mu}})$ and $p(\tilde{\boldsymbol{\mu}})$ computation can be marginalized over all dynamical orders in the following manner:

$$p(\tilde{\boldsymbol{s}}|\tilde{\boldsymbol{\mu}}) = p(s^{[0]}, s^{[1]}, s^{[2]}, ...|\mu^{[0]}, \mu^{[1]}, \mu^{[2]}, ...) = \prod_{n=0}^{\infty} p(s^{[n]}|\mu^{[n]}), \tag{2.40}$$

$$p(\tilde{\boldsymbol{\mu}}) = p(\mu^{[0]}, \mu^{[1]}, \mu^{[2]}, ...) \approx \prod_{n=0}^{\infty} p(\mu^{[n+1]}|\mu^{[n]}). \tag{2.41}$$

Note that in Equation (2.41) the factor $p(\mu^{[0]})$ was omitted; the reasoning behind this is the assumption of a uniform probability distribution of the prior belief over the state configuration [30]. As in the static case, the noise fluctuations are believed to be of Gaussian distribution, so that the likelihood of n-th dynamical order of the sensory pattern s and the n-th order conditional prior of μ can be computed through:

$$p(s^{[n]}|\mu^{[n]}) = \frac{1}{\sqrt{2\pi}\sigma_{s[n]}} \exp -\frac{(s^{[n]} - g^{[n]})^2}{2\sigma^2_{s[n]}}, \tag{2.42}$$

$$p(\mu^{[n+1]}|\mu^{[n]}) = \frac{1}{\sqrt{2\pi}\sigma_{\mu[n]}} \exp -\frac{(\mu^{[n+1]} - f^{[n]})^2}{2\sigma^2_{\mu[n]}}. \tag{2.43}$$

By substituting Equation (2.42) and Equation (2.43) in Equation (2.40) and Equation (2.41) and finally evaluating the surprise associated to the joint probability $p(\tilde{\boldsymbol{s}}, \tilde{\boldsymbol{\mu}})$:

$$\begin{aligned}
\mathcal{I}(\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{s}}) &= -\ln p(\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{s}}) \\
&= -\ln p(\tilde{\boldsymbol{s}}|\tilde{\boldsymbol{\mu}}) - \ln p(\tilde{\boldsymbol{\mu}}) \\
&\approx -\ln \prod_{n=0}^{\infty} p(s^{[n]}|\mu^{[n]}) - \ln \prod_{n=0}^{\infty} p(\mu^{[n+1]}|\mu^{[n]}) \\
&= -\sum_{n=0}^{\infty} \ln p(s^{[n]}|\mu^{[n]}) - \sum_{n=0}^{\infty} \ln p(\mu^{[n+1]}|\mu^{[n]}) \\
&= -\sum_{n=0}^{\infty} \ln \frac{1}{\sqrt{2\pi}\sigma_{s[n]}} \exp - \frac{(s^{[n]} - g^{[n]})^2}{2\sigma^2_{s[n]}} - \sum_{n=0}^{\infty} \ln \frac{1}{\sqrt{2\pi}\sigma_{\mu[n]}} \exp - \frac{(\mu^{[n+1]} - f^{[n]})^2}{2\sigma^2_{\mu[n]}}.
\end{aligned}$$

As in the static case the following auxiliary notation are introduced denoting the prediction errors on the n-th dynamical order:

$$\epsilon_s^{[n]} = s^{[n]} - g^{[n]}, \tag{2.44}$$

$$\epsilon_\mu^{[n]} = \mu^{[n+1]} - f^{[n]}. \tag{2.45}$$

Finally, the following formulation for $I(\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{s}})$ is derived:

$$\mathcal{I}(\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{s}}) = \sum_{n=0}^{\infty} \left[\ln 2\pi + \ln \sigma_{s[n]}\sigma_{\mu[n]} + \frac{1}{2\sigma^2_{s[n]}} \epsilon_s^{[n]2} + \frac{1}{2\sigma^2_{\mu[n]}} \epsilon_\mu^{[n]2} \right]. \tag{2.46}$$

According to Equation (2.46) the surprise term to be evaluated to compute VFE is the sum over all dynamical orders of a constant depending on the variances of the statistically independent noise terms and the squared prediction errors. Note that this sum is composed of infinite terms and therefore is intractable; indeed, in practical applications only derivatives up to a finite order are evaluated. This can be done by setting the term corresponding to the highest order considered to random fluctuations with large variance:

$$\mu^{[n_{max}]} = n_\mu^{[n_{max}]}. \tag{2.47}$$

Notably, he large variance $\sigma^2_{\mu[n_{max}]}$ of the noise term $n_\mu^{[n_{max}]}$ assures that the corresponding prediction error term in Equation (2.46) is negligible.

2.3.3. The multivariate case

Here, the formulation for the multivariate case is derived. Most of the necessary steps will be omitted since they are analogous to the one presented in the univariate case; only the principal differences from the univariate cases and the additional assumptions needed to derive a formulation of the surprise term in Equation (2.17) will be presented.
In the multivariate case such term assumes the following notation:

$$\mathcal{I}(\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\}, \{\tilde{\boldsymbol{s}}_{\boldsymbol{\alpha}}\}) = -\ln p(\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\}, \{\tilde{\boldsymbol{s}}_{\boldsymbol{\alpha}}\}), \tag{2.48}$$

where $\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\}$ and $\{\tilde{\boldsymbol{s}}_{\boldsymbol{\alpha}}\}$ denote two vectors respectively containing the brain states' and the sensory data's generalized coordinates. As before, we define the generative density by specifying the n-th order of the agent's prior belief and of sensory data:

$$s_\alpha^{[n]} = g_\alpha^{[n]} + n_{s,\alpha}^{[n]}, \tag{2.49}$$

$$\mu_\alpha^{[n+1]} = f_\alpha^{[n]} + n_{\mu,\alpha}^{[n]}. \tag{2.50}$$

Similarly to the univariate case we assume a normal distribution for the noise terms in Equation (2.49) and Equation (2.50). Note that in general the noise terms of each variable may be correlated. Assuming statistical independence of noise terms at each dynamical order and a uniform distribution for $p(\{\mu_\alpha^{[0]}\})$, the following formulations for the sensory likelihood and the agent's prior:

$$\begin{aligned} p(\{\tilde{\boldsymbol{s}}_{\boldsymbol{\alpha}}\}|\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\}) &= \prod_{n=0}^{\infty} p(\{s_\alpha^{[n]}\}|\{\mu_\alpha^{[n]}\}) \\ &= \prod_{n=0}^{\infty} \frac{1}{\sqrt{2\pi|\Sigma_{s[n],\alpha}|}} \exp\left[-\frac{1}{2}(s_\alpha^{[n]} - g_\alpha^{[n]})\Sigma_{s[n],\alpha}^{-1}(s_\alpha^{[n]} - g_\alpha^{[n]})^T\right], \end{aligned} \tag{2.51}$$

$$\begin{aligned} p(\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\}) &\approx \prod_{n=0}^{\infty} p(\{\mu_\alpha^{[n+1]}\}|\{\mu_\alpha^{[n]}\}) \\ &= \prod_{n=0}^{\infty} \frac{1}{\sqrt{2\pi|\Sigma_{\mu[n],\alpha}|}} \exp\left[-\frac{1}{2}(\mu_\alpha^{[n+1]} - f_\alpha^{[n]})\Sigma_{\mu[n],\alpha}^{-1}(\mu_\alpha^{[n+1]} - f_\alpha^{[n]})^T\right]. \end{aligned} \tag{2.52}$$

An ulterior simplification of Equation (2.53) and Equation (2.54) can be performed by assuming statistical independence between different variables:

$$
\begin{aligned}
p(\{\tilde{\boldsymbol{s}}_{\boldsymbol{\alpha}}\}|\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\}) &= \prod_{n=0}^{\infty}\prod_{\alpha=1}^{N} p(s_\alpha^{[n]}|\mu_\alpha^{[n]}) \\
&= \prod_{n=0}^{\infty}\prod_{\alpha=1}^{N} \frac{1}{\sqrt{2\pi}\sigma_{s[n],\alpha}} \exp -\frac{(s_\alpha^{[n]} - g_\alpha^{[n]})^2}{2\sigma^2_{s[n],\alpha}}, \qquad (2.53) \\
p(\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\}) &= \prod_{n=0}^{\infty}\prod_{\alpha=1}^{N} p(\mu_\alpha^{[n+1]}|\mu_\alpha^{[n]}) \\
&= \prod_{n=0}^{\infty}\prod_{\alpha=1}^{N} \frac{1}{\sqrt{2\pi}\sigma_{\mu[n],\alpha}} \exp -\frac{(\mu_\alpha^{[n+1]} - f_\alpha^{[n]})^2}{2\sigma^2_{\mu[n],\alpha}}. \qquad (2.54)
\end{aligned}
$$

A similar derivation as the one presented in Section 2.3.2 can be performed also in the multivariate case. The following formulation for Equation (2.48) is finally obtained:

$$
\mathcal{I}(\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\},\{\tilde{\boldsymbol{s}}_{\boldsymbol{\alpha}}\}) = \sum_{n=0}^{\infty}\sum_{\alpha=1}^{N}\left[\ln 2\pi + \ln \sigma_{s[n],\alpha}\sigma_{\mu[n],\alpha} + \frac{1}{2\sigma^2_{s[n],\alpha}}{\epsilon_{s_\alpha}^{[n]}}^2 + \frac{1}{2\sigma^2_{\mu[n],\alpha}}{\epsilon_{\mu_\alpha}^{[n]}}^2\right], \qquad (2.55)
$$

where the following auxiliary notation was exploited to define the prediction errors on sensory data and hidden states' estimates:

$$
\begin{aligned}
\epsilon_{s_\alpha}^{[n]} &= s_\alpha^{[n]} - g_\alpha^{[n]}, \qquad (2.56) \\
\epsilon_{\mu_\alpha}^{[n]} &= \mu_\alpha^{[n+1]} - f_\alpha^{[n]}. \qquad (2.57)
\end{aligned}
$$

2.4. Active Inference: Free Energy minimization

In the previous sections, a tractable version of Variational Free Energy was derived. The analytical formulation described so far stems from the Free Energy Principle and the relative plausible approximations. As before mentioned, Active Inference is the practical framework adopted by self-organizing agents to minimize Free Energy through perceptual adaptation and action. In analytical terms, the action selected by the agent and the internal state of the generative model are updated through a gradient descent method on VFE:

$$\tilde{\boldsymbol{\mu}}(t_{i+1}) = \tilde{\boldsymbol{\mu}}(t_i) + dt\left[\dot{\tilde{\boldsymbol{\mu}}} - k_\mu \frac{\delta \mathcal{F}}{\delta \tilde{\boldsymbol{\mu}}}\right], \tag{2.58}$$

$$A(t_{i+1}) = A(t_i) - dt \cdot k_A \frac{\delta \mathcal{F}}{\delta A}. \tag{2.59}$$

Please note that the update of said variables is performed in discrete time (which will be later used for the simulations), where dt is the time interval between consecutive instants. Notably, the internal state $\tilde{\boldsymbol{\mu}}$ is updated with the contributions of both its expected dynamics $\dot{\tilde{\boldsymbol{\mu}}}$ (which need to be specified in the generative model) and the gradient descent on VFE; otherwise, action is only updated in order to minimize VFE.

In Equation (2.58) and Equation (2.59), k_μ and k_A are usually referred to as learning rates [5]; their purpose is to modulate the gradient descent on VFE to ensure that it unfolds at the same time scales of the expected dynamics. In their absence, expected dynamics and gradient descent could potentially unfold at mismatched time scales making one of the two contributions negligible with respect to the other.

Moreover, a considerable issue arises in the computation of the gradient descents. According to Equation (2.46), the surprise term of VFE is a weighted sum of squared errors at each dynamical order. As before mentioned the constant terms are irrelevant to the computation of the gradients. The computation of $\delta\mathcal{F}/\delta\tilde{\boldsymbol{\mu}}$ is straightforward, since $\tilde{\boldsymbol{\mu}}$ is directly represented in VFE's formulation. Otherwise, the action A is not analytically present in Equation (2.46). The role of action is to change the dynamics of the generative process in order to minimize sensory prediction errors. As it will be later discussed in more detail, the agent is assumed to have learned the causal relation between action and sensory pattern changes $\delta\tilde{\boldsymbol{s}}/\delta A$ so that the computation of the relative gradient could be performed through the chain rule:

$$\frac{\delta \mathcal{F}}{\delta A} = \frac{\delta \mathcal{F}}{\delta \tilde{\boldsymbol{s}}}\frac{\delta \tilde{\boldsymbol{s}}}{\delta A}. \tag{2.60}$$

Analogous to the gradient with respect to $\tilde{\boldsymbol{\mu}}$, the computation of $\delta\mathcal{F}/\delta\tilde{\boldsymbol{s}}$ is trivial since $\tilde{\boldsymbol{s}}$ is directly present in the formulation of VFE. The analytical computation of such gradients is addressed later in the context of the presented simulations.

3 | Motor control simulations

In this chapter a simulation framework is introduced to investigate how active inference can account for movement generation under multisensory conflict and goal-directed tasks. The chosen framework serves as a proof of concept implementation of active inference and does not intend to be a realistic representation of the dynamics occurring in a real organism's nervous system. The simulations presented in this chapter were conducted within the MATLAB environment.

3.1. Simulation setup

The simulation setup chosen includes a sitting agent with their arm distended over a table. The only movement permitted is a external/internal rotation of the shoulder joint which results in a rotation of the agent's forearm around the axis perpendicular to the table's plane. The ordinate axis is positioned perpendicular to the agent's frontal plane, while the abscissa axis is perpendicular to both the previous axes with the positive semiaxis pointing to the agent's right. The origin of the reference system is placed on the elbow joint. The system configuration is fully described by a single degree of freedom, namely the elbow joint angle, which is defined starting from the ordinate positive semiaxis in the counterclockwise direction. A schematic representation of the reference frame is provided in Figure 3.2.

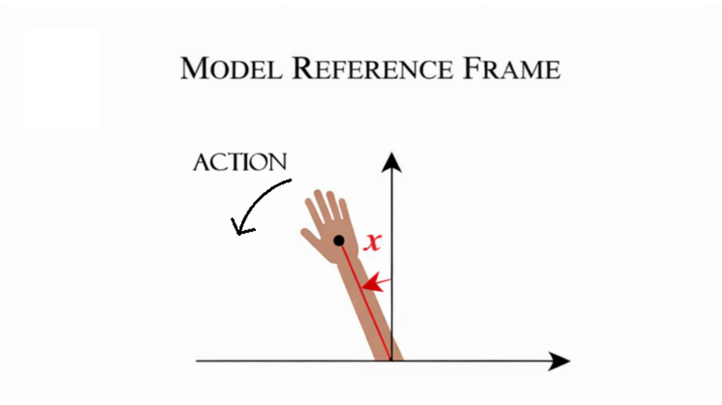

Figure 3.1: Simulation setup: reference system [30]

The agent is assumed to receive sensory data from the environment consisting in proprioceptive information about the joint configuration and visual information about the hand position in the reference system above defined.

In the first simulation, the agent performs a reaching task in the absence of bodily illusions. The target is to be reached with the arm's end-effector (i.e., the hand); therefore, it must be positioned such that it lies within the subject's reachable workspace. In a more complex and realistic simulation, the agent would need to infer the target position from sensory data. Thus, the agent would receive visual data in the form of retinal position of the target image and would reconstruct the spatial coordinates of the target. To ensure the reachability requirement within a simplified simulation framework, the target is defined as a predetermined elbow joint angle configuration rather than as spatial coordinates of a target point (which in principle could result to be unreachable in the setup herein described); moreover, the target is assumed to be known by the agent, who doesn't need to perform inference about its position.

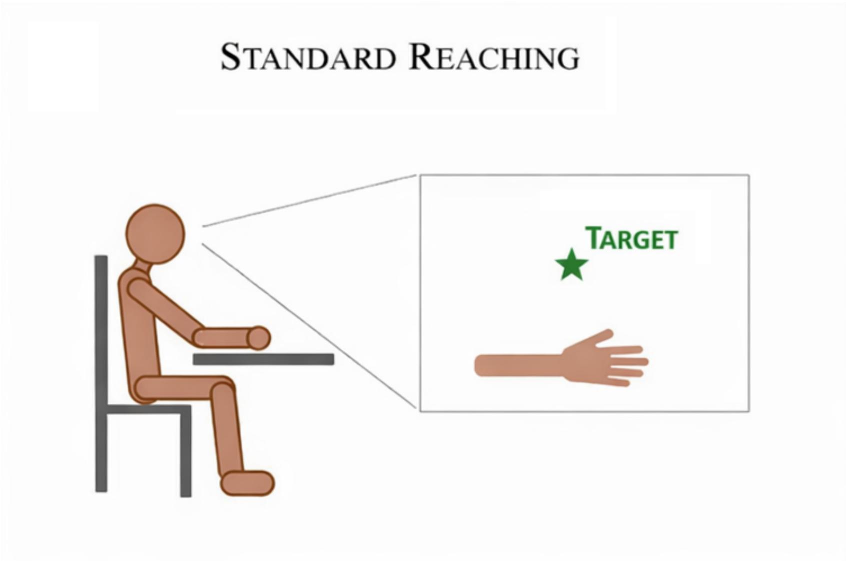

Figure 3.2: First simulation setup: reaching task [30]

In the second simulation, the agent has no direct goal assigned, but is subdue to the Rubber Hand Illusion. Thus, their real arm is concealed from view and the visual information is received from a fake rubber arm placed in a different configuration with respect to the real one. Proprioceptive sensory information is still provided by the real arm. The resulting movement will only have the underlying motive of resolving the multisensory conflict arising from the illusion.

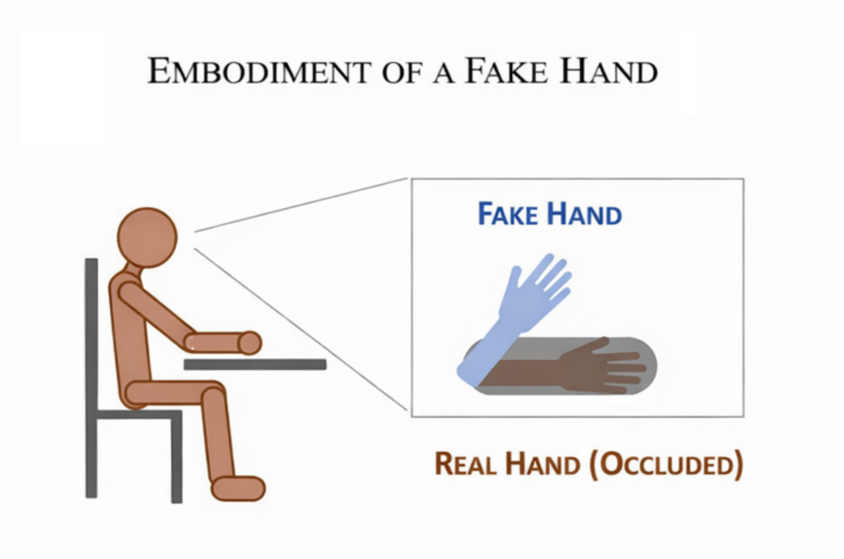

Figure 3.3: Second simulation setup: Rubber Hand Illusion [30]

The third simulation consists of a reaching task performed under multisensory conflict: the agent is assumed to wear a visor displaying a virtual arm that the agent can control by moving their real arm. Therefore, the agent receives visual information about the system configuration from the virtual hand, which has an incorrect visuo-motor mapping producing a spatial misalignment with respect to the real limb: more specifically, the virtual arm moves faster than the real one. Yet again, proprioceptive sensory information is provided by the real arm.

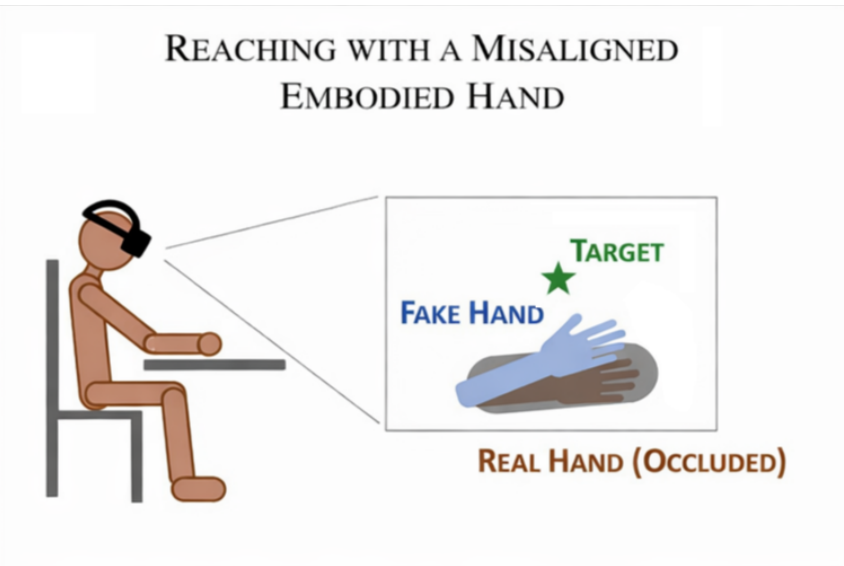

Figure 3.4: Third simulation setup: virtual hand embodiment [30]

In the next section, the generative process (i.e. the environment dynamics) and the generative model (i.e. the agent internal dynamics) will be analytically defined.

3.1.1. Active inference framework

In order to be implement the Free Energy principle and active inference for the experimental setup, the operational framework must be defined. This includes defining the generative process and the generative model along with their dynamics.

The generative process is the mathematical description of the system the active inference agent interacts with, namely the agent's own arm. The generative process is fully described by four elements: a state vector, its dynamics, a sensory vector and the mapping between the state vector and the sensory vector.

The state vector describes the system configuration in generalized coordinates. In the simple setup chosen the arm configuration is fully explained by the elbow joint angle, which represents the hidden state $\tilde{\boldsymbol{x}}$ in generalized coordinates. For the implementation, components up to the first order derivative (the angular velocity) will be considered. Thus, the state vector is:

$$\tilde{\boldsymbol{x}} = \begin{bmatrix} x \\ x' \end{bmatrix}. \tag{3.1}$$

While the state vector contains the instantaneous value of the state variable (the elbow joint) and its considered derivatives, the generative process' full description requires the definition of a set of differential equations defining the evolution in time of said variables. The dynamics of the state vector describes how the each derivative of the state evolves in time under the effect of its own expected dynamics and the action selected by the agent. The generative process dynamics can be represented as a vector of the the same dimension of the state vector:

$$\dot{\tilde{\boldsymbol{x}}} = \begin{bmatrix} \dot{x} \\ \dot{x}' \end{bmatrix}. \tag{3.2}$$

Note that in the adopted notation, the time derivative of the state vector in generalized coordinates is denoted with the prime symbol $'$, while the dot notation i exploited for the differential equation describing the dynamics of a component of the state vector. For instance, x' denotes the value of the angular velocity of the arm, while $\dot{x}$ denotes the differential equation describing the temporal evolution of the elbow joint angle x. Moreover, $\dot{x}'$ denotes the differential equation defining the temporal evolution of the state's first order derivative x'.

The following set of equations is adopted:

$$\dot{\tilde{\boldsymbol{x}}} = \begin{bmatrix} \dot{x} \\ \dot{x}' \end{bmatrix} = \begin{bmatrix} x' \\ A - \phi x'^{\beta}/m_{arm} \end{bmatrix} = \begin{bmatrix} f_{x,1} \\ f_{x,2} \end{bmatrix} = \boldsymbol{f_x}. \tag{3.3}$$

The dynamics of the state variable x is straightforwardly defined by the value of its first order derivative. Otherwise, the angular velocity x' dynamics are modeled as a damped system with a power dependence on angular velocity expressed through the exponent β. ϕ represents the viscosity of the system, while m_{arm} is the arm mass. Finally, the angular velocity's dynamics is affected by the action A selected by the agent, which has the same dimensions of an angular acceleration.

The sensory vector s comprises the sensory input the agent receives from the environment through different sensory channels. In the setup chosen such channels are proprioception and vision. Specifically, the proprioceptive information regarding arm configuration is assumed to be a mere measure of the elbow joint angle x affected by random Guassian noise n_p (with standard deviation σ_{np}). For simplicity the visual information gathered by the agent is assumed to be the coordinates of the hand in the two-dimensional reference system adopted; more precisely, the coordinates of the center of the hand's palm. A more accurate and complex modeling would have provided the agent with the retinal position of the hand's palm and a mapping between the two reference (which must be learned by the agent).
Likewise proprioception, vision is also affected by random Gaussian noise n_v (with standard deviation σ_{nv}), which is assumed to be the same along the two directions of the reference system [30]. In the case of proprioception, the mapping between the state vector and the sensory vector is straightforward; while in the case of visual information simple trigonometry rules are adopted. Thus, the sensory vector assumes the following form:

$$\boldsymbol{s} = \begin{bmatrix} s_p \\ s_{vx} \\ s_{vy} \end{bmatrix} = \begin{bmatrix} x + n_p \\ L_{arm} \cdot cos(x + \pi/2) + n_v \\ L_{arm} \cdot sin(x + \pi/2) + n_v \end{bmatrix} = \boldsymbol{g_x} + \boldsymbol{n_x}, \tag{3.4}$$

where L_{arm} is the length of the arm.

On the other hand, the generative model is a mathematical description of the active inference agent. It represents the agent's internal representation of the environment and of its dynamics. Its full description is similar to that one of the generative process: the agent has its own internal state vector $\tilde{\boldsymbol{\mu}}$, with its own intrinsic dynamics $\dot{\tilde{\boldsymbol{\mu}}}$, an expected sensory vector $\boldsymbol{\mu_s}$ and the mapping between $\tilde{\boldsymbol{\mu}}$ and $\boldsymbol{\mu_s}$. Although in general the internal state domain and dimensions can be uncorrelated to those of the generative process, for simplicity it is assumed that the agent's internal representation mirrors that one of the generative process. Thus, the internal state vector assumes the following form:

$$\tilde{\boldsymbol{\mu}} = \begin{bmatrix} \mu \\ \mu' \\ \mu'' \end{bmatrix}, \tag{3.5}$$

where μ and μ' represent the internal belief about the configuration of the arm (i.e. the elbow joint angle and its angular velocity). Note that the internal state vector comprises also the second order of time derivative accordingly with the theoretical formulation of active inference presented in Chapter 2. The generative model also includes the dynamics of the internal state vector:

$$\dot{\tilde{\boldsymbol{\mu}}} = \begin{bmatrix} \dot{\mu} \\ \dot{\mu}' \end{bmatrix} = \begin{bmatrix} \mu' \\ (\overline{K}(\mu_T - \mu) - \overline{\phi}\mu')/\overline{m}_{arm} \end{bmatrix} = \begin{bmatrix} f_{\mu,1} \\ f_{\mu,2} \end{bmatrix} = \boldsymbol{f_\mu}. \tag{3.6}$$

Similar to the generative process' state dynamics, the dynamics of the internal belief μ is straightforwardly defined by the value of its first order derivative. In the expected angular velocity's dynamics an attractor point μ_T has been added which represents the target configuration to be reached. When $\mu = \mu_T$ the agent does not intend to trigger any movement: indeed the term $\overline{K}(\mu_T - \mu)$ represents a force that is proportional to the distance between the hand and the target and always directed towards the target. Performing a visually-guided reaching task requires also the estimation of the target location (i.e. the target coordinates in the reference system) through the elaboration of the visual signal coming from the target. Once the target location is estimated, the agent needs to estimate the state configuration that allows the completion of the reaching task, which is exactly μ_T. Thus, conceptually μ_T is a parameter to be inferred exploiting internal models that are supposed to have been trained across the agent's lifespan. For simplicity, the target configuration μ_T is assumed to be known apriori by the agent [31]. As in the generative process, the internal belief's dynamics is modeled with a damping coefficient $\overline{\phi}$, while $\overline{m}_{arm}$ represents the agent's belief about the arm's mass. In accordance with previous implementations of the the Free Energy Principle and Active Inference the highest order dynamics are set to random gaussian fluctuations [5, 30].

The expected sensory vector μ_s comprises the expected proprioceptive and visual sensory data. It mirrors the generative process sensory vector $\boldsymbol{s}$: indeed, the proprioceptive data is represented by expected elbow joint angle mu, while the visual information is represented by the hand's expected coordinates in the two-dimensional reference system:

$$\boldsymbol{\mu_s} = \begin{bmatrix} \mu_p \\ \mu_{vx} \\ \mu_{vy} \end{bmatrix} = \begin{bmatrix} \mu \\ \overline{L}_{arm} \cdot cos(\mu + \pi/2) \\ \overline{L}_{arm} \cdot sin(\mu + \pi/2) \end{bmatrix} = \begin{bmatrix} g_{\mu,1} \\ g_{\mu,2} \\ g_{\mu,3} \end{bmatrix} = \boldsymbol{g_\mu}, \tag{3.7}$$

where $\overline{L}_{arm}$ is internal belief about the arm's length.

3.1.2. Free Energy

In Chapter 2 Variational Free Energy was defined as the Kullback-Leibler divergence between the joint probability $p(x, s)$ and the recognition density $q(x)$. Under the Laplace approximation a tractable Variational Free Energy was derived as in Equation (2.46) where, essentially, VFE is a weighted sum of squared prediction errors over all considered dynamical orders plus a constant. By applying such equation in the context of the simulation herein described the following definition of VFE is derived:

$$\begin{aligned} \mathcal{F} = \frac{1}{2\sigma_{\mu'}}(\mu' - f_{\mu,1})^2 + \frac{1}{2\sigma_{\mu''}}(\mu'' - f_{\mu,2})^2 + \frac{1}{2\sigma_p}(s_p - g_{\mu,1})^2 + ... \\ ... + \frac{1}{2\sigma_v}\left((s_{vx} - g_{\mu,2})^2 + (s_{vy} - g_{\mu,3})^2\right) + C. \end{aligned} \tag{3.8}$$

The analytical derivation of the constant C is not necessary since the perceptual and active adaptation of the agent is performed exploiting the gradients of VFE in order to minimize it. Moreover, given the definition of $f_{\mu,1}$ in Equation (3.6), it is evident that the first error term $\mu' - f_{\mu,1}$ is identically zero. In light of such considerations the actual VFE formulation that is actually adopted in the simulation is:

$$\mathcal{F} = \frac{1}{2\sigma_{\mu''}}(\mu'' - f_{\mu,2})^2 + \frac{1}{2\sigma_p}(s_p - g_{\mu,1})^2 + \frac{1}{2\sigma_v}\left((s_{vx} - g_{\mu,2})^2 + (s_{vy} - g_{\mu,3})^2\right). \tag{3.9}$$

For simulation purposes a discrete time interval can be defined:

$$\boldsymbol{t} = [t_1 : dt : t_N], \tag{3.10}$$

where t_1 and t_N are respectively the simulation starting and ending instants, while dt is the discrete period separating consecutive time instants (i.e. the temporal resolution); N is the number of elements in the vector $\boldsymbol{t}$.

To carry out the simulation the initial condition of both the generative process and the generative model must be established:

$$\begin{aligned}
&\tilde{\boldsymbol{x}}(t_1) = \boldsymbol{x}_0, \\
&\tilde{\boldsymbol{\mu}}(t_1) = \boldsymbol{\mu}_0, \\
&\boldsymbol{s}(t_1) = \boldsymbol{g}_{\boldsymbol{x}}(t_1) + \boldsymbol{n}_{\boldsymbol{x}}(t_1), \\
&\boldsymbol{\mu}_{\boldsymbol{s}}(t_1) = \boldsymbol{g}_{\boldsymbol{\mu}}(t_1).
\end{aligned}$$

Subsequently, an iterative process performs the computation of the dynamics of both interacting systems over the discrete time interval described by $\boldsymbol{t}$. Firstly, the inferred state and the action chosen by the agent are updated using the following equations:

$$\tilde{\boldsymbol{\mu}}(t_{i+1}) = \tilde{\boldsymbol{\mu}}(t_i) + dt \left[\dot{\tilde{\boldsymbol{\mu}}}(t_i) - \boldsymbol{k}_{\mu} \frac{\delta VFE}{\delta \tilde{\boldsymbol{\mu}}} \right], \tag{3.11}$$

$$A(t_{i+1}) = A(t_i) - dt \left[\boldsymbol{k}_{\boldsymbol{A}} \frac{\boldsymbol{\delta VFE}}{\boldsymbol{\delta A}} \right]. \tag{3.12}$$

Critical considerations must be examined regarding Equation (3.11) and Equation (3.12). The internal state vector is updated by considering both its expected dynamics $\dot{\tilde{\boldsymbol{\mu}}}(t_i)$ and a second term minimizing VFE through a gradient descent; otherwise, the selected action A is updated only via VFE's minimization. The simulation needs to take into account that the intrinsic dynamics of the system (i.e. $\dot{\tilde{\boldsymbol{\mu}}}$) and the gradient descent of VFE may unfold at different time scales [5, 30]. For instance, if the minimization of VFE was much faster, then the the generative model dynamics' contribution to the internal state vector $\tilde{\boldsymbol{\mu}}$ update would be negligible. Hence, the introduction of additional weights $\boldsymbol{k}_{\mu}$ and $\boldsymbol{k}_{\boldsymbol{A}}$ is required. The vector $\boldsymbol{k}_{\mu}$ contains the specific weights over the corresponding dynamical order of the internal state:

$$\boldsymbol{k}_{\mu} = \begin{bmatrix} k_{\mu} \\ k_{\mu'} \\ k_{\mu''} \end{bmatrix}. \tag{3.13}$$

Otherwise, the action weight $\boldsymbol{k}_{\boldsymbol{A}}$ is composed of two components, corresponding to the weighting of proprioceptive and visual contributions to action. As will be discussed in more detail later, although action is represented as a scalar quantity, it is conceptually treated as comprising two distinct components, associated with proprioceptive and visual modalities, respectively:

$$\boldsymbol{k_A} = \begin{bmatrix} k_{A,p} \\ k_{A,v} \end{bmatrix}. \tag{3.14}$$

Once the internal state and the selected action are updated, the real system state vector $\tilde{\boldsymbol{x}}$ and the consequential sensory input $\boldsymbol{s}$ can be updated through the following equations:

$$\tilde{\boldsymbol{x}}(t_{i+1}) = \tilde{\boldsymbol{x}}(t_i) + dt\left[\dot{\tilde{\boldsymbol{x}}}\big(\tilde{\boldsymbol{x}}(t_i), A(t_{i+1})\big)\right], \tag{3.15}$$

$$\boldsymbol{s}(t_{i+1}) = \boldsymbol{g_x}\big(\tilde{\boldsymbol{x}}(t_{i+1})\big) + \boldsymbol{n_x}(t_{i+1}). \tag{3.16}$$

3.1.3. Gradient descents computation

To complete the mathematical framework of the simulations described thus far, the analytical formulations for the gradients involved in the agent's belief and action updates (Equation (3.11), Equation (3.12)) must be derived.

VFE's gradient with respect to the internal state vector $\tilde{\boldsymbol{\mu}}$ is a vector comprising the partial derivatives of VFE (as defined by Equation (3.9)) with respect to all considered dynamical orders contained in $\tilde{\boldsymbol{\mu}}$:

$$-\frac{\delta\mathcal{F}}{\delta\tilde{\boldsymbol{\mu}}} = \begin{bmatrix} -\frac{\delta\mathcal{F}}{\delta\mu} \\ -\frac{\delta\mathcal{F}}{\delta\mu'} \\ -\frac{\delta\mathcal{F}}{\delta\mu''} \end{bmatrix}. \tag{3.17}$$

By exploiting the definition of VFE given in Equation (3.9), such components are easily derived. The first component is the partial derivative of VFE with respect to the internal state variable μ:

$$\begin{aligned} -\frac{\delta\mathcal{F}}{\delta\mu} = \frac{1}{\sigma_{\mu''}}(\mu'' - f_{\mu,2})\frac{\delta f_{\mu,2}}{\delta\mu} + \frac{1}{\sigma_p}(s_p - g_{\mu,1})\frac{\delta g_{\mu,1}}{\delta\mu} + \dots \\ \dots + \frac{1}{\sigma_v}\left((s_{vx} - g_{\mu,2})\frac{\delta g_{\mu,2}}{\delta\mu} + (s_{vy} - g_{\mu,3})\frac{\delta g_{\mu,3}}{\delta\mu}\right). \end{aligned} \tag{3.18}$$

The remaining partial derivatives are obtained by differentiating the relative components in Equation (3.6) and Equation (3.7):

$$\frac{\delta f_{\mu,2}}{\delta\mu} = -\frac{\overline{K}}{\overline{m}_{arm}}, \tag{3.19}$$

$$\frac{\delta g_{\mu,1}}{\delta\mu} = 1, \tag{3.20}$$

$$\frac{\delta g_{\mu,2}}{\delta\mu} = -\overline{L}_{arm} \cdot sin(\mu + \pi/2), \tag{3.21}$$

$$\frac{\delta g_{\mu,3}}{\delta\mu} = \overline{L}_{arm} \cdot cos(\mu + \pi/2). \tag{3.22}$$

The second component is the partial derivative of VFE with respect to the internal angular velocity estimate μ':

$$-\frac{\delta\mathcal{F}}{\delta\mu'} = \frac{1}{\sigma_{\mu''}}(\mu'' - f_{\mu,2})\frac{\delta f_{\mu,2}}{\delta\mu'}. \tag{3.23}$$

Once again the remaining partial derivative can be computed by differentiating $f_{\mu,2}$ definition given in Equation (3.6):

$$\frac{\delta f_{\mu,2}}{\delta\mu'} = -\frac{\overline{\phi}}{\overline{m}_{arm}}. \tag{3.24}$$

Finally, the last component is the partial derivative of VFE with respect to the last considered order derivative of the internal state vector μ'':

$$-\frac{\delta\mathcal{F}}{\delta\mu''} = -\frac{1}{\sigma_{\mu''}}(\mu'' - f_{\mu,2}). \tag{3.25}$$

Unlike the case of the internal state vector $\tilde{\boldsymbol{\mu}}$, the derivation of the VFE gradient with respect to the action A is non-trivial: indeed, the action selected by the agent is not explicitly represented into VFE's definition in Equation (3.9). To address this apparent impasse, active inference posits that the agent possesses an internal model describing how actions translate into sensory state changes (i.e. $\delta\boldsymbol{s}/\delta A$). This model is assumed to be hardwired into the agent and trained across his lifespan through experience [1, 30]. This allows to compute the gradient through the chain rule of differentiation. From a mathematical standpoint, the result is that thee gradient with respect to action will be treated as a vector with two components: one associated to the proprioceptive stimulus and one associated to the visual information gathered by the agent:

$$\begin{aligned}\frac{\boldsymbol{\delta\mathcal{F}}}{\boldsymbol{\delta A}} = \frac{\delta\mathcal{F}}{\delta \boldsymbol{s}}\frac{\delta \boldsymbol{s}}{\delta A} &= \left[\frac{\delta\mathcal{F}}{\delta s_p}\frac{\delta s_p}{\delta A}, \frac{\delta\mathcal{F}}{\delta \boldsymbol{s_v}}\frac{\delta \boldsymbol{s_v}}{\delta A}\right] = \\ &= \left[\frac{\delta\mathcal{F}}{\delta s_p}\frac{\delta s_p}{\delta A}, \frac{\delta\mathcal{F}}{\delta \boldsymbol{s_v}}\frac{\delta \boldsymbol{s_v}}{\delta s_p}\frac{\delta s_p}{\delta A}\right] = \\ &= \frac{\delta s_p}{\delta A}\left[\frac{\delta\mathcal{F}}{\delta s_p}, \frac{\delta\mathcal{F}}{\delta \boldsymbol{s_v}}\frac{\delta \boldsymbol{s_v}}{\delta s_p}\right].\end{aligned} \tag{3.26}$$

Indeed, the action selected by the agent can be though of as the compromise between two components: the former which tries to minimize the proprioceptive error, while the latter aims at minimizing visual errors [31]. Once again, VFE's definition in Equation (3.9) can be straightforwardly differentiated with respect to the components of the sensory vector $\boldsymbol{s}$:

$$\frac{\delta\mathcal{F}}{\delta s_p} = \frac{1}{\sigma_p}(s_p - g_{\mu,1}), \tag{3.27}$$

$$\frac{\delta\mathcal{F}}{\delta \boldsymbol{s_v}} = \begin{bmatrix}\delta\mathcal{F}/\delta s_{vx} \\ \delta\mathcal{F}/\delta s_{vy}\end{bmatrix} = \frac{1}{\sigma_v}\begin{bmatrix}s_{vx} - g_{\mu,2} \\ s_{vy} - g_{\mu,3}\end{bmatrix}. \tag{3.28}$$

Conversely, some approximations and considerations are necessary to derive the remaining partial derivatives. The principal approximation adopted consists of substituting the sensory vector components with their estimated counterparts in the generative model. Thus:

$$\frac{\delta \boldsymbol{s_v}}{\delta s_p} \approx \frac{\delta \boldsymbol{\mu_v}}{\delta \mu_p} = \frac{\delta \boldsymbol{g_{\mu,v}}}{\delta \mu} = \begin{bmatrix}-\overline{L}_{arm} \cdot sin(\mu + \pi/2) \\ \overline{L}_{arm} \cdot cos(\mu + \pi/2)\end{bmatrix}, \tag{3.29}$$

where $\boldsymbol{g_{\mu,v}}$ denotes the vector comprising the second and third components of $\boldsymbol{g_\mu}$, which represent the estimations of the visual sensations the agent receives.

Finally, $\delta s_p/\delta A$ represents the mapping between the action performed and the resulting proprioceptive sensation. In biological organisms this mapping is thought to be hard-wired in a series of arc reflexes that allow to change the proprioceptive signals through action without requiring any motor planning [1]. Nonetheless, in a simulation context an analytical formulation of such term must be derived. The same approximation as in the case of Equation (3.29) can be exploited:

$$\frac{\delta s_p}{\delta A} \approx \frac{\delta \mu_p}{\delta A} = \frac{\delta \mu}{\delta A}. \tag{3.30}$$

An additional consideration is necessary since no relationship between the belief μ and the action A is available: given Equation (3.3) it is clear that action has thee same dimensions of an angular acceleration; hence, we can impose the action to be equal to the dynamics of the belief's first order derivative $\dot{\mu}'$. By rearranging its definition given in Equation (3.6) we can express the belief μ in function of A and finally compute its derivative:

$$\begin{aligned} \dot{\mu}' &= \frac{\overline{K}(\mu_T - \mu) - \overline{\phi}\mu'}{\overline{m}_{arm}} \stackrel{!}{=} A, \\ \implies \mu &= \mu_T - \frac{\overline{\phi}\mu' + \overline{m}_{arm}A}{\overline{K}}. \end{aligned} \tag{3.31}$$

Therefore:

$$\frac{\delta s_p}{\delta A} \approx \frac{\delta \mu}{\delta A} = -\frac{\overline{m}_{arm}}{\overline{K}}. \tag{3.32}$$

The gradients herein derived can be plugged into Equation (3.11) and Equation (3.12) to update the internal state vector and the action selected by the agent with the aim of minimizing VFE during the simulations. As provisioned by the theoretical framework description, VFE is the sum of squared errors of two kinds: model prediction errors and sensory prediction errors. It is worth noticing that the action update is driven solely by the minimization of the latter type, while the perceptual inference is driven by both. Indeed, the agent's action only plays a role in the dynamics of the generative process and affects the subsequent sensory data $\boldsymbol{s}$; whereas, the perceptual adaptation of the agent internal model can minimize prediction errors on its own dynamics and prediction errors on the sensory information. It is noteworthy that whereas action modifies sensory input to align with internal expectations, the update of internal beliefs minimizes prediction errors by adjusting the state vector $\tilde{\boldsymbol{\mu}}$ to enhance the accuracy of the sensory prediction $\boldsymbol{\mu_s}$.

3.1.4. Parameters

This section provides an overview of the parameters adopted in the three simulations. The generative process' and generative model's characterizations were kept mostly equal across the three simulations.

The three simulations were run over the same time interval ranging between $t_1 = 0\,s$ and $t_N = 30\,s$ with a temporal resolutions of $dt = 10\,ms$. For all simulations $t_{on} = 3\,s$ denotes the onset time instant at which the reaching tasks and illusions are initiated.

As for the generative process, the real arm's length and mass in Equation (3.4) were set to the average values for an adult human being: $L_{arm} = 45\,cm$ and $m_{arm} = 1.5\,kg$. The noise standard deviations for the sensory acquisitions are set to plausible values: $\sigma_{np} = 0.1\,rad$ and $\sigma_{nv} = 1\,cm$. The parameters for the damped dynamics in Equation (3.3) were set to $\phi = 4$ and $\beta = 0.5$. A power dependence on x' was adopted to avoid critical damping, i.e. to avoid oscillations around the target [30]. Nonetheless, some precautions must be adopted in order to avoid imaginary parts when x' assumes negative values; Equation (3.3) was implemented in the following form in the code:

$$\dot{x}' = A - \frac{\phi}{m_{arm}} sign[x'] \cdot |x'|^{\beta}. \tag{3.33}$$

This adaptation ensures that the argument of the square root (i.e. $|x'$) remains non-negative, thereby preventing imaginary parts in the result; nonetheless, its sign is preserved and contributes to define the direction in which x' is updated.

The parameters in the equations of the generative model (Equation (3.6) and Equation (3.7)) were kept consistent with their counterparts in the generative process: therefore $\overline{m}_{arm} = m_{arm} = 1.5\,kg$, $\overline{L}_{arm} = L_{arm} = 45\,cm$ and $\overline{\phi} = \phi = 4$. This suggests that the agent maintains an unbiased representation of the system's characteristics, presumably developed through cumulative experience across their lifetime. The proportionality constant for the attractive force toward the target was set to $\overline{K} = 1.3$. It is worth noticing that such force represents the elastic component in Equation (3.6) that is absent in the generative process counterpart (Equation (3.3)). This implies a strict interplay between the action A in the generative process dynamics and the elastic force attracting the belief μ towards the target configuration μ_T. Such target was set to $\mu_T = \pi/3$ for the first and third simulations where the agent is performing a reaching task, while it was set to a null value for the second simulation. In the second simulation the rubber hand is placed in a configuration $x_{RH} = \pi/3$. In the third simulation a velocity gain for the virtual hand dynamics is defined as $v_{gain} = x'_{VH}/x' = 1.3$.

For all the simulations the initial condition were set with $\tilde{\boldsymbol{x}}(t_1) = \mathbf{0}$ and $\tilde{\boldsymbol{\mu}}(t_1) = \mathbf{0}$.

The uncertainties relative to the various prediction errors in Equation (3.9) were set to $\sigma_{\mu''} = 0.1$ and $\sigma_p = 0.1$. The uncertainty relative to the visual prediction error was set to $\sigma_v = 0.01$ for the first and third simulations; while it was set to $\sigma_v = 0.07$ for the second simulation, implying an higher uncertainty about the visual prediction during the Rubber Hand Illusion. Indeed, these parameters modulate the prediction errors in VFE evaluation and the subsequent gradient descents employed to update both the internal state vector and the selected action.

The learning rates for the action and internal state vector updates were set to: $k_{A,p} = 0.3$, $k_\mu = 0.1$, $k_{\mu'} = 0.01$, $k_{\mu''} = 0.001$. The weight for the visual contribution to the action's update was set to $k_{A,v} = 0.3$ for the fist and third simulations; for the second simulation this parameter was set to a null value, i.e. during the Rubber Hand Illusion the action update is solely driven by the minimization of the proprioceptive prediction error. In the first and third simulations, the equal setting of action gains ($k_{A,p}$ and $k_{A,v}$) implies that any difference in the magnitude of action updates is solely attributable to the variance in uncertainty values σ_p and σ_v.

Table 3.1 provides a summary of the parameters employed in the three simulations.

Parameter	**Description**	**Value**		
		I simulation	II simulation	III simulation
t_1, t_N	Time interval		$[0, 30]\,$s	
dt	Temporal resolution		10 ms	
t_{on}	Onset time		3 s	
L_{arm}, $\overline{L}_{arm}$	Arm length		45 cm	
m_{arm}, $\overline{m}_{arm}$	Arm mass		1.5 kg	
ϕ, $\overline{\phi}$	Damping coefficient		4	
β	Damping exponent		0.5	
$\overline{K}$	Attractive force constant		1.3	
σ_{np}	Proprioceptive noise		0.1 rad	
σ_{nv}	Velocity noise		1 cm	
$\sigma_{\mu''}$	Internal state uncertainty		0.1	
σ_p	Proprioceptive uncertainty		0.1	
σ_v	Visual uncertainty	0.01	0.07	0.01
$\tilde{\boldsymbol{x}}(t_1)$, $\tilde{\boldsymbol{\mu}}(t_1)$	Initial conditions		$\mathbf{0}$	
k_μ, $k_{\mu'}$, $k_{\mu''}$	State learning rates		0.1, 0.01, 0.001	
$k_{A,p}$	Action learning rate (prop.)		0.3	
$k_{A,v}$	Action learning rate (vis.)	0.3	0	0.3
μ_T	Target configuration	$\pi/3$	0	$\pi/3$
x_{RH}	Rubber Hand configuration	-	$\pi/3$	-
v_{gain}	Velocity gain	-	-	1.3

Table 3.1: Summary of the parameters employed across the three simulations

4 | Results discussion

In this chapter the results stemming from the simulations presented in Chapter 3 are presented and discussed.

4.1. First simulation

In the first simulation, the active inference agent must perform a standard reaching task of a fixed target placed within its reachable space. No multisensory conflict is present: therefore, proprioceptive and visual sensory channels provide consistent information about the arm's configuration which can be correctly interpreted to reach the target configuration.

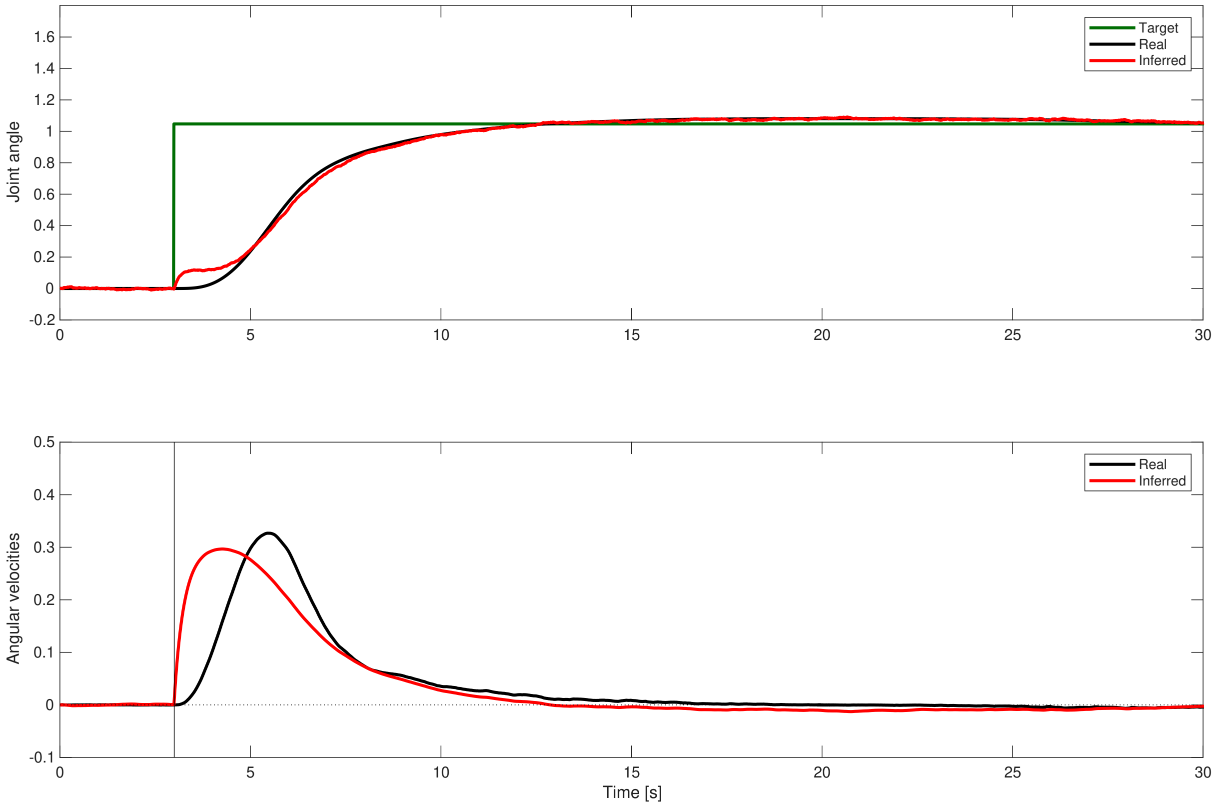

Figure 4.1: First simulation: real and inferred kinematics

The real and inferred arm's kinematics are observable in Figure 4.1. Prior to the onset time t_{on}, the target configuration is set to coincide with the initial state of the system. Subsequently, the target is presented to the agent, triggering the reaching task. Delaying the target presentation by a predefined interval is a deliberate choice: it facilitates the observation of the agent's baseline behavior and ensures that the internal state reaches a stable equilibrium before any goals are introduced. Specifically, before triggering the reaching task, no considerable movement od both the inferred and real kinematics are appreciable. When the target is provided to the agent the internal state vector $\tilde{\boldsymbol{\mu}}$ starts to move towards the target; subsequently the state vector $\tilde{\boldsymbol{x}}$ follows the trajectory of the internal one under the effect of the action selected by the agent.

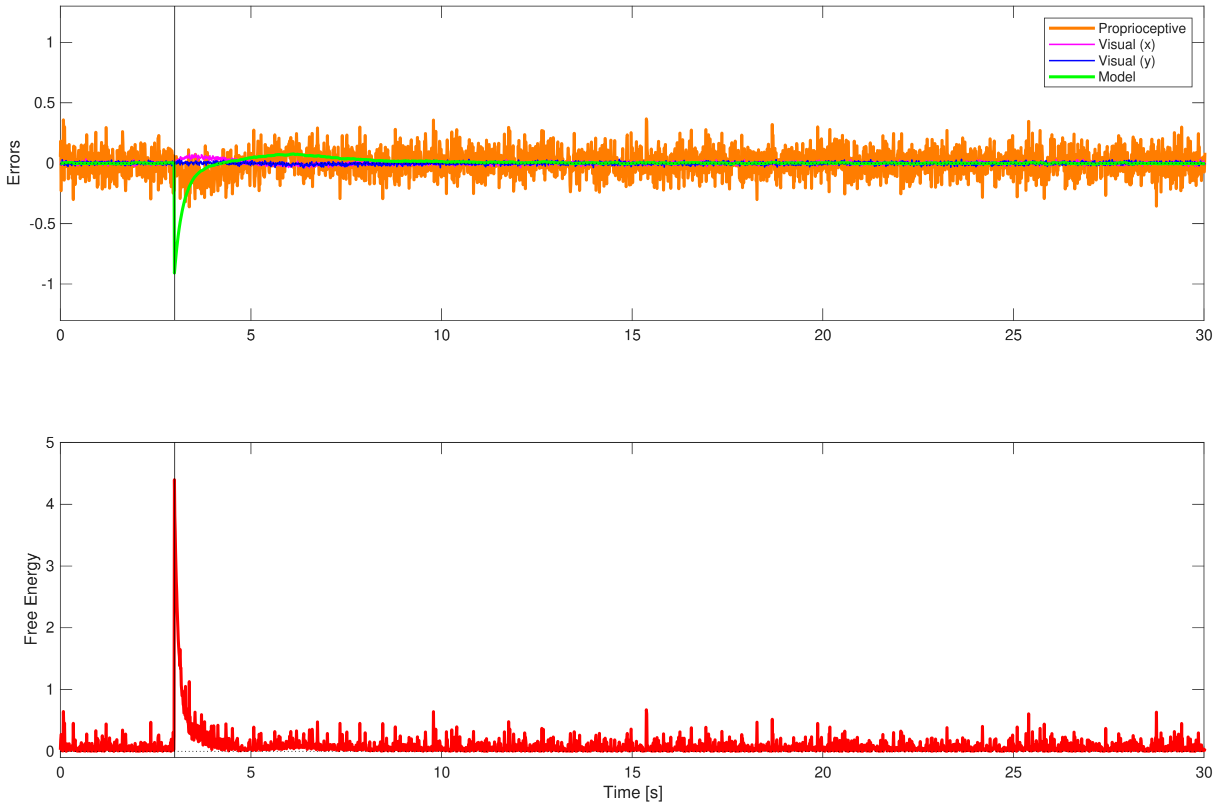

Figure 4.2: First simulation: prediction errors and Variational Free Energy

As visible in Figure 4.2, the anticipatory movement of the internal state vector is driven by a sudden spike in the model prediction error. This spike is also observable in the evaluation of VFE. Subsequently, the difference in the kinematics of the state vector $\tilde{\boldsymbol{x}}$ and the internal state vector $\tilde{\boldsymbol{\mu}}$ causes an increase in the sensory prediction errors. Over time, the minimization of VFE successfully drives the errors towards null values: the residual oscillations in the sensory errors are attributable to the presence of stochastic noise perturbations, which are also reflected in VFE's profile.

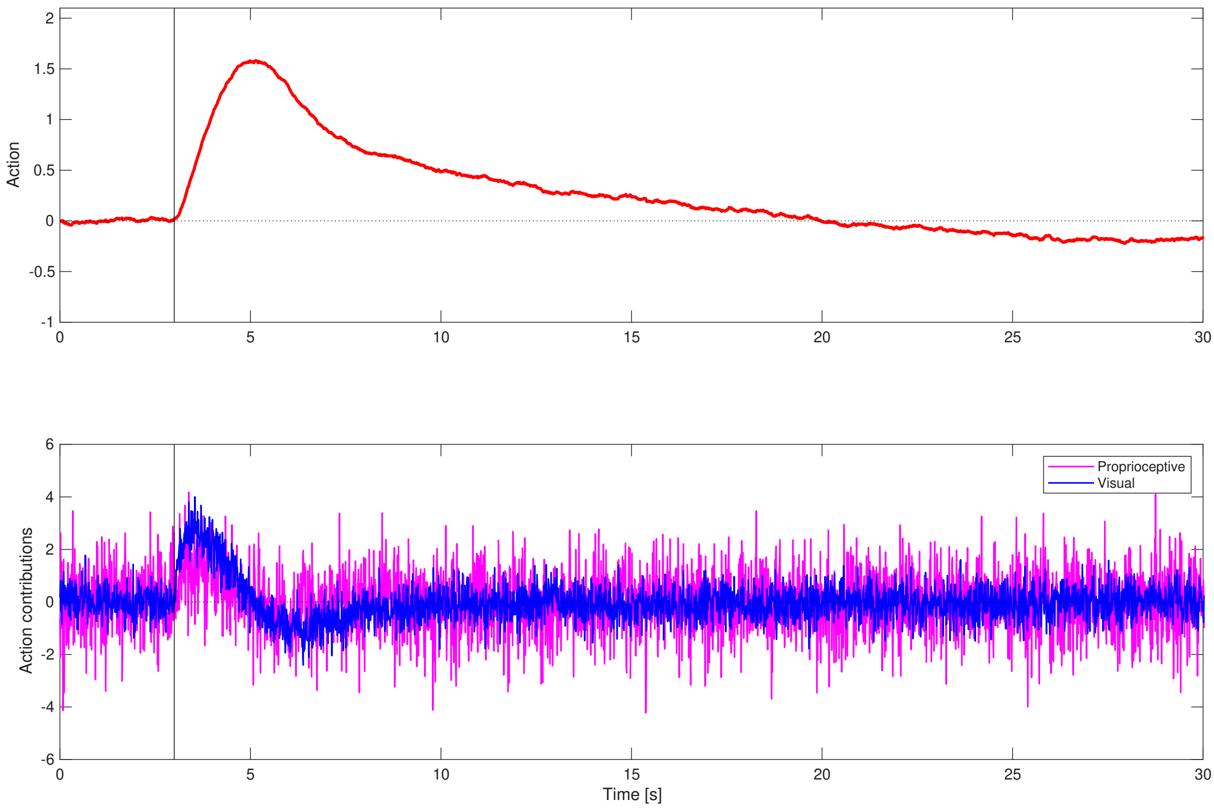

Figure 4.3: First simulation: action and sensory contributions

Figure 4.3 displays the action selected by the agent alongside the different contributions from proprioceptive and visual prediction errors. Prior to the onset time, the action is negligible: the minimal observable fluctuations serve exclusively to counteract sensorial perturbations. It is worth emphasizing that action can only be triggered and modulated by sensory prediction errors. Hence, the initial spike of the model prediction errors after the onset time does not elicit an immediate increase in action. The modulation of action is contingent upon the increase of sensory prediction errors, functioning to drive their minimization. Notably, action assumes negative values towards the end of the simulation to counter a slight overshoot of the target.

Sequentially, the onset of a different target configuration cause a spike in the model prediction error. Consequently, the internal state vector's trajectory is modified in an anticipatory fashion with the purpose of increasing sensory prediction errors. Only at this stage, action is triggered and modulated to modify the trajectory of the state vector in order to keep sensory errors as small as possible: to this end, the trajectory of the state vector follows the internal state vector's one. Finally, the model prediction error is reduced as the internal state approaches the target configuration.

4.2. Second simulation

In the second simulation, the active inference must reconcile the multisensory conflict arising from the Rubber Hand Illusion in the absence of a target configuration. Not having a target ensures a complete decoupling between the reaching task and the multisensory conflict context: thus, any triggered movement will have the only purpose of resolving the multisensory conflict.

Mirroring the first simulation, the illusion is triggered at the onset instant t_{on}. In the illusory state, the agent is visually informed by the rubber hand ($x_{RH} = \pi/3$), while it still receives proprioceptive signals from the real arm; hence establishing multisensory conflict. Therefore, the agent's visual input in Equation (3.4) is modified in:

$$\boldsymbol{s} = \begin{bmatrix} s_p \\ s_{vx} \\ s_{vy} \end{bmatrix} = \begin{bmatrix} x + n_p \\ L_{arm} \cdot cos(x_{RH} + \pi/2) + n_v \\ L_{arm} \cdot sin(x_{RH} + \pi/2) + n_v \end{bmatrix}. \tag{4.1}$$

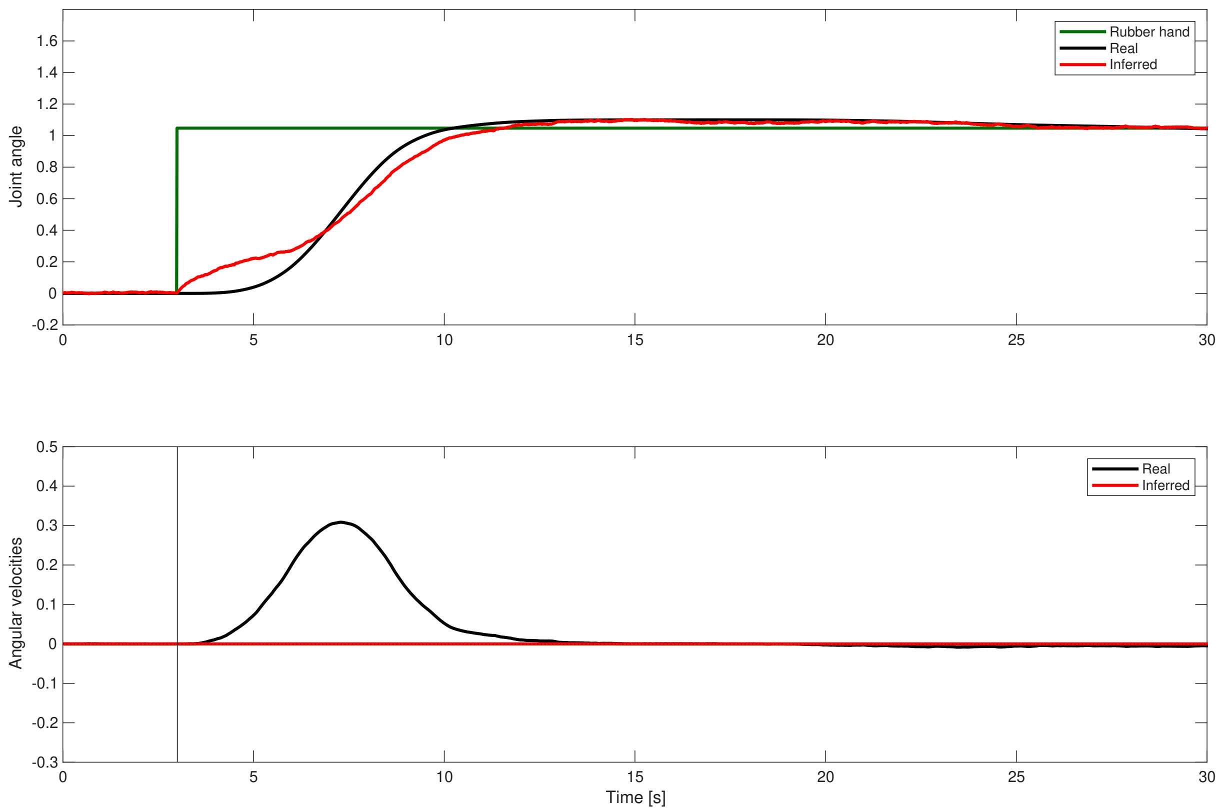

Figure 4.4: Second simulation: real and inferred kinematics

Figure 4.4 displays the kinematics of both the real and inferred state alongside the static rubber hand configuration. As illustrated in the first panel, the agent converges towards the rubber hand configuration despite the absence of a defined target; however, the underlying mechanism driving this convergence is clearly different, as evident in the second panel. Similarly to the first simulation, the inferred state μ displays an anticipatory trajectory relative to the physical state x; In turn, the physical state is driven to track the inferred state. Otherwise, the inferred angular velocity μ' is identically null, while the real one has its own free dynamics. Such a discrepancy is dictated by the agent not being subject to a reaching task. To ensure this condition, the elastic component in Equation (3.6) is removed: effectively, without an explicit goal, the internal model does not favor any specific state, rendering all configurations equally permissible; equivalently, the target configuration is always equal to the inferred state ($\mu_T = \mu$). Consequently, the dynamics of μ' in Equation (3.6) turns out to be:

$$\dot{\mu}' = f_{\mu,2} == \frac{\overline{K}(\mu - \mu) - \overline{\phi}\mu'}{\overline{m}_{arm}} = -\frac{\overline{\phi}}{\overline{m}_{arm}}\mu'. \tag{4.2}$$

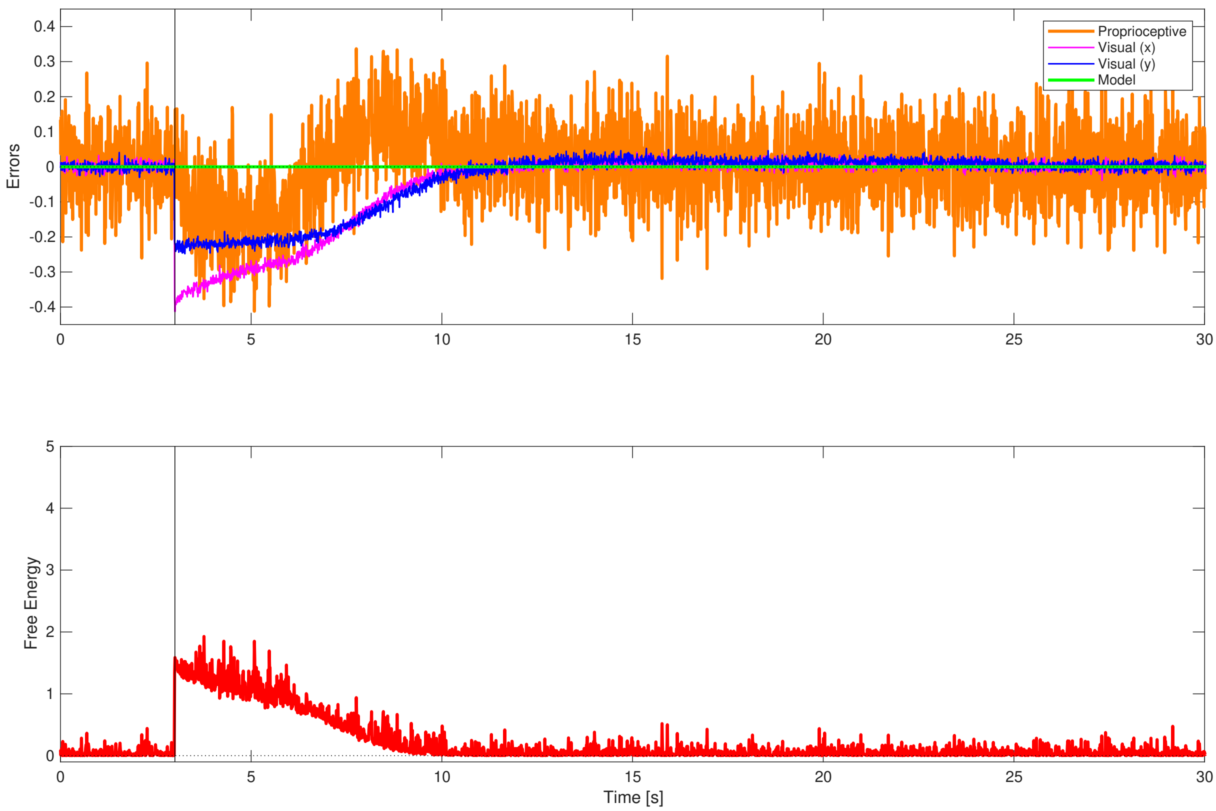

Figure 4.5: Second simulation: prediction errors and Variational Free Energy

Moreover, as confirmed by the first panel in Figure 4.5, the model prediction error is identically equal to zero as well. Indeed, expanding Equation (3.11) with Equation (3.23), Equation (3.25) and Equation (4.2):

$$\begin{aligned}\mu'(t_{i+1}) &= \mu'(t_i) + dt\left[\dot{\mu}'(t_i) - k_{\mu'}\frac{\delta VFE}{\delta\mu'}\right] = \\ &= \mu'(t_i) + dt\left[-\frac{\overline{\phi}}{\overline{m}_{arm}}\mu'(t_i) - k_{\mu'}\frac{1}{\sigma_{\mu''}}(\mu''(t_i) + \frac{\overline{\phi}}{\overline{m}_{arm}}\mu'(t_i))\right] = \\ &= c_1\,\mu'(t_i) + c_2\,\mu''(t_i).\end{aligned} \tag{4.3}$$

Analogously, it can be demonstrated that:

$$\mu''(t_{i+1}) = c_3\,\mu'(t_i) + c_4\,\mu''(t_i). \tag{4.4}$$

Consequently, the updates of μ' and μ'' are defined as linear combinations of their respective values at the preceding time step. Given that their initial conditions are null, both variables remain identically zero throughout the second simulation. Importantly, this result reflects the fact that the agent does not have any explicit intention to move. Consequently, the agent incorrectly infers that its velocity remains zero, as if it executed no movement. This finding provides a conjectural explanation for the lack of movement awareness experienced during the Rubber Hand Illusion [20].

As evident in the first panel of Figure 4.5, the onset of the illusion triggers a pronounced spike in visual prediction errors, which is responsible for the correspondent spike in VFE. Consequentially, the internal state μ is induced to shift its trajectory towards the rubber arm, causing an increase in the proprioceptive error. Finally, such error solicits the agent to select an action so that the physical state x follows μ toward the rubber hand configuration.

It is worth noticing that in this simulation action can be triggered solely by proprioceptive prediction errors ($k_{A,v} = 0$). Indeed, as observable in the second panel of Figure 4.6, the visual contribution to action would be negative, directed away from the rubber hand. This stems from the fact that visual signals originate from the rubber hand itself; hence, the contribution is negative, as the gradient descent effectively attempts to shift the rubber hand position towards the physical one to minimize the spatial discrepancy. Such consideration provides the rationale behind selecting $k_{A,v} = 0$ for this simulation, since the agent has no control over the rubber arm. Nonetheless, visual errors still provide a contribution to the dynamics of the internal state μ, allowing it to trigger a movement by modulating the proprioceptive error. It is worth remarking that in this simulation an higher visual uncertainty σ_v was employed, reflecting the reduced reliability of the visual input from the rubber hand, as perceived by the agent.

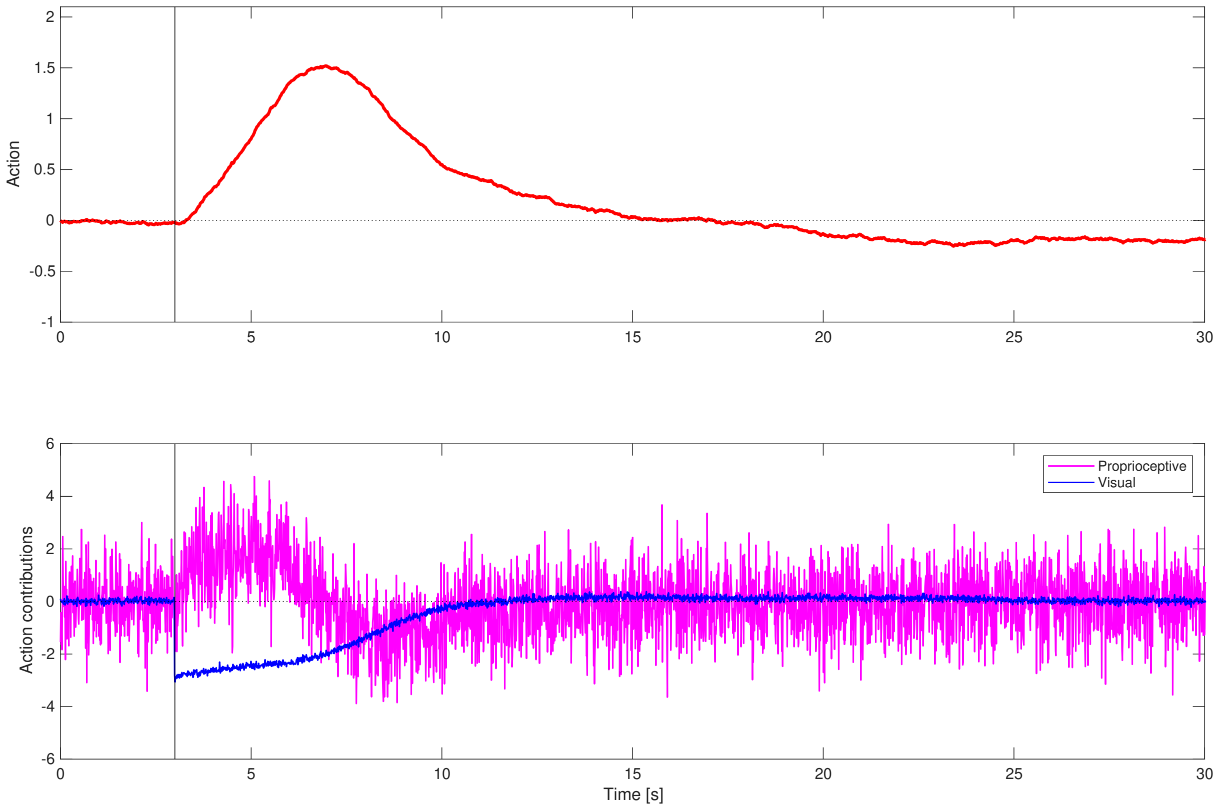

Figure 4.6: Second simulation: action and sensory contributions

The rationale for omitting the elastic force in the internal state dynamics (achieved by setting $\mu_T = \mu$) requires further elaboration. Setting the target to the initial value of the internal state might seem equivalent to removing the reaching objective, reflecting the agent's will to remain in its initial position. When the illusion is triggered, the agent would still display the behavior described so far. However, this configuration would exert a restorative force toward the starting position when μ shifts towards the rubber hand. Consequently, this condition would not disentangle the reaching task from the multisensory conflict. The outcome would be an equilibrium between the rubber hand configuration and the initial condition, as verified by the control simulation's results in Figure 4.7. In the fifth panel it is evident that the visual errors remain stable and do not decrease, resulting in a failure in minimizing VFE (as observable in the sixth panel). This implies that performing a reaching task under a non-compatible multisensory conflict is essentially impossible: the target configuration could be achieved only if it coincided with the rubber arm position.

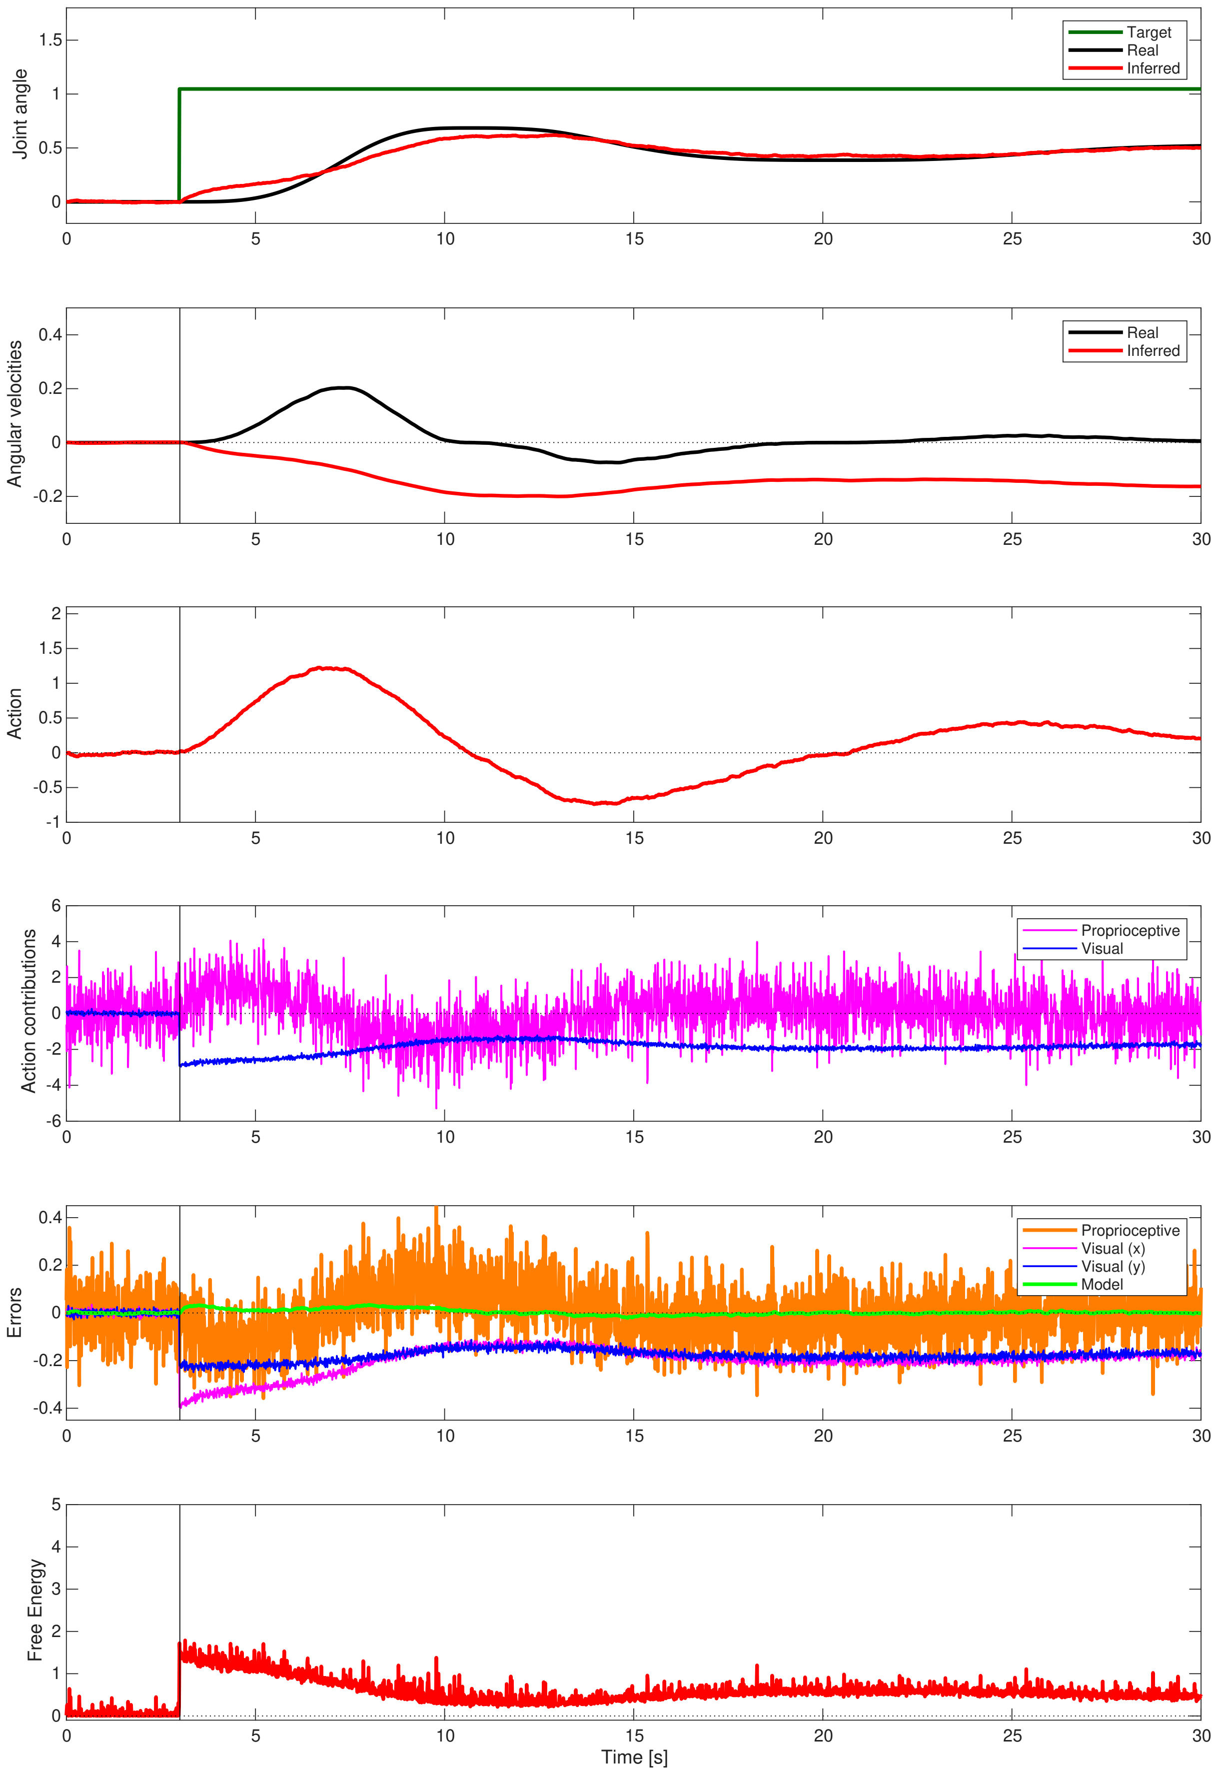

Figure 4.7: RHI: control simulation

4.3. Third simulation

The third simulation consists of a reaching task superimposed to a multisensory conflict which is solicited by a sense of body ownership over a virtual hand displayed by a visor. The agent receives visual information from the virtual hand while receiving proprioceptive information from the real arm. The generative process is extended by the state vector of the virtual limb $\tilde{\boldsymbol{x}}_{\boldsymbol{VH}}$ with its own dynamics $\dot{x}_{VH}$:

$$\tilde{\boldsymbol{x}}_{\boldsymbol{VH}} = \begin{bmatrix} x_{VH} \\ x'_{VH} \end{bmatrix}, \tag{4.5}$$

$$\dot{x}_{VH} = x'_{VH} = v_{gain} \cdot x', \tag{4.6}$$

where $v_{gain} = 1.3$ is the gain adopted to map the real limb's angular velocity onto the virtual arm's one. The simulation is started with the virtual arm and the real one aligned in the configuration $x = x_{VH} = 0$; however, the different angular velocities will produce a spatial misalignment. The consequent multisensory conflict remains unsolvable due to the structural constraints of the generative process.

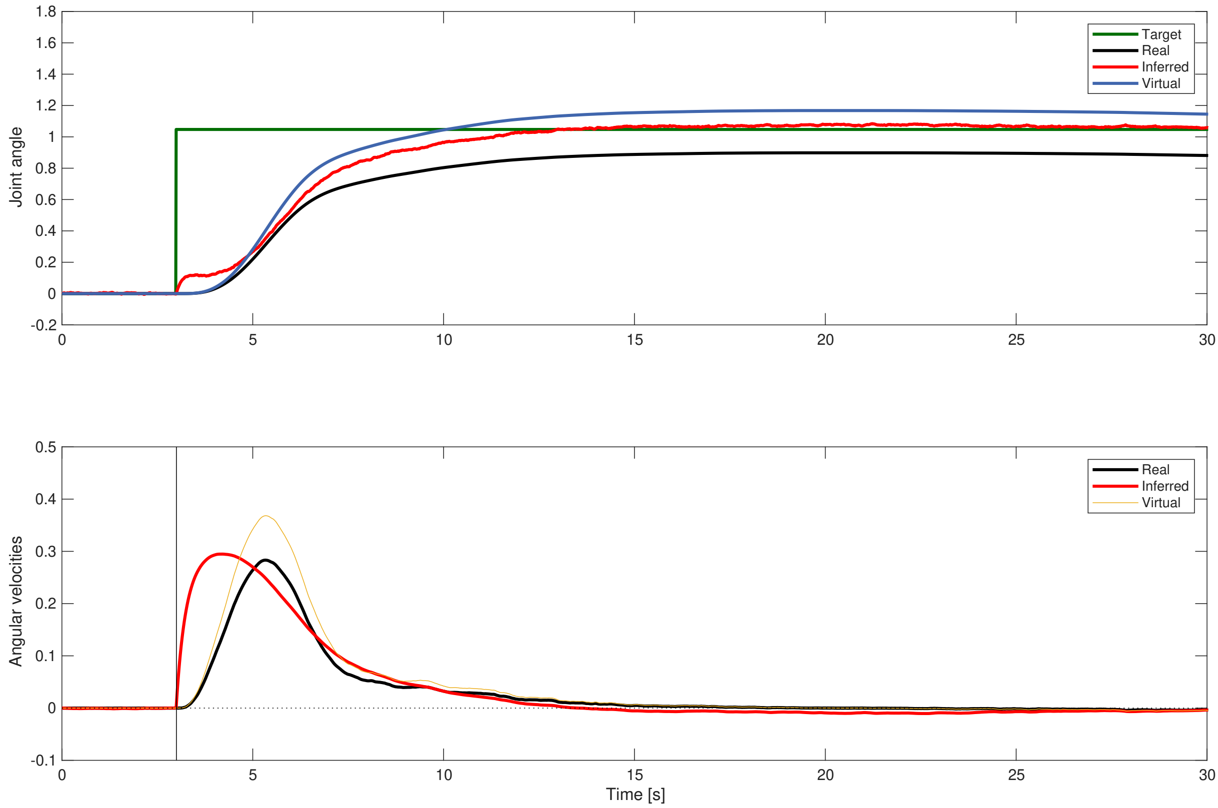

Figure 4.8: Third simulation: real, inferred and virtual kinematics

As anticipated in Table 3.1, the target of the reaching task is the same of the first simulation ($\mu_T = \pi/3$), along with all the other parameters. The resulting kinematics are illustrated in Figure 4.8. As evident in the first panel, the agent is unable to successfully perform the reaching task under multisensory conflict: indeed, the virtual hand overshoots the target, while the real one undershoots it. The inferred state trajectory lies in between the virtual and the real ones: this is the result of an equilibrium condition between visual and proprioceptive prediction errors. Consistent with the other simulations, the internal state vector $\tilde{\boldsymbol{\mu}}$ displays an anticipatory behavior that drives the movement towards the target. Notably, only the inferred state reaches the target, as its convergence is driven by the elastic force in the generative model dynamics. Moreover, the inferred state is closer to the virtual hand since vision is characterized by a lower sensory uncertainty compared to proprioception.

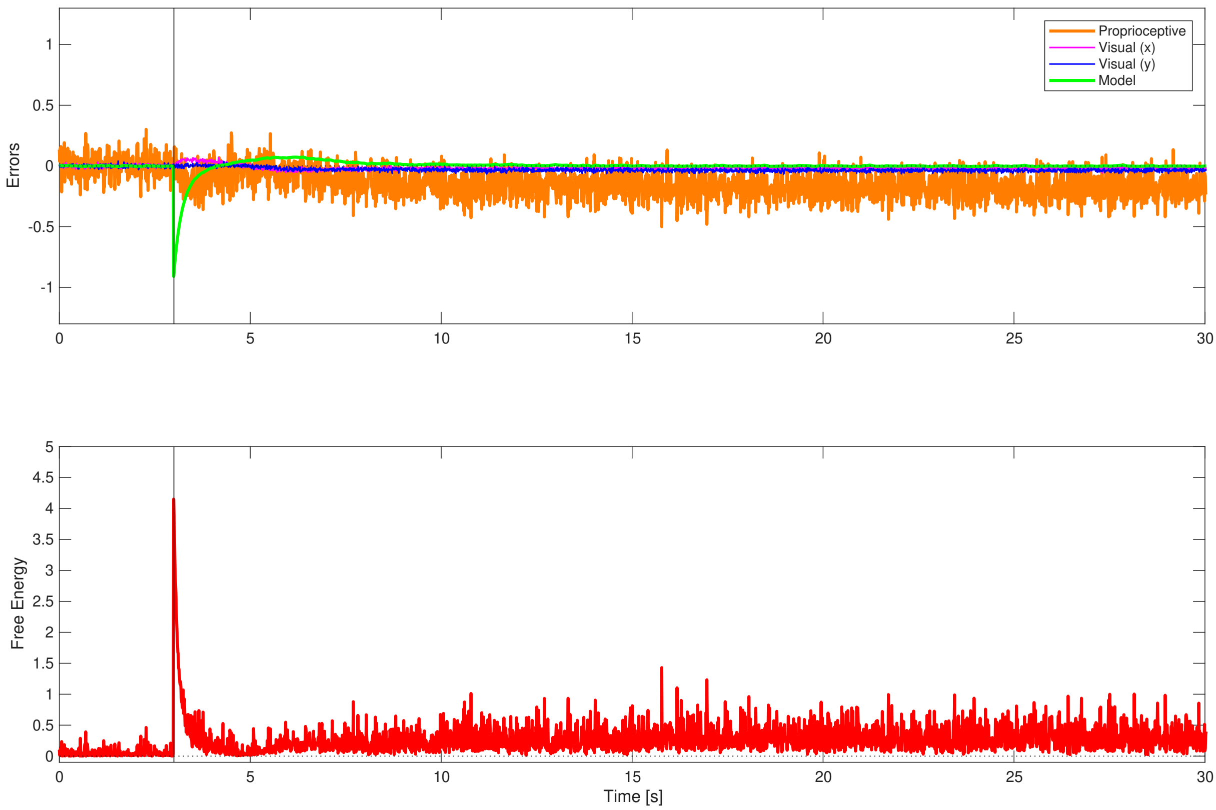

Figure 4.9: Third simulation: prediction errors and Variational Free Energy

The results depicted in Figure 4.9 indicate a behavior similar to the one observed in the first simulation: first, the target onset causes a spike in the model prediction error (and consequently in VFE); then, this error triggers the anticipatory movements of the internal state vector, increasing the proprioceptive error; finally, such sensory error elicits an action that causes a movement of the real arm towards the target. Importantly, the proprioceptive error is affected by oscillations caused by stochastic noise, as in the other simulations; however, unlike previous cases, its mean value is not null, resulting in a higher value of VFE at steady state. This is attributable to the physical discrepancy between the real and virtual hand positions.

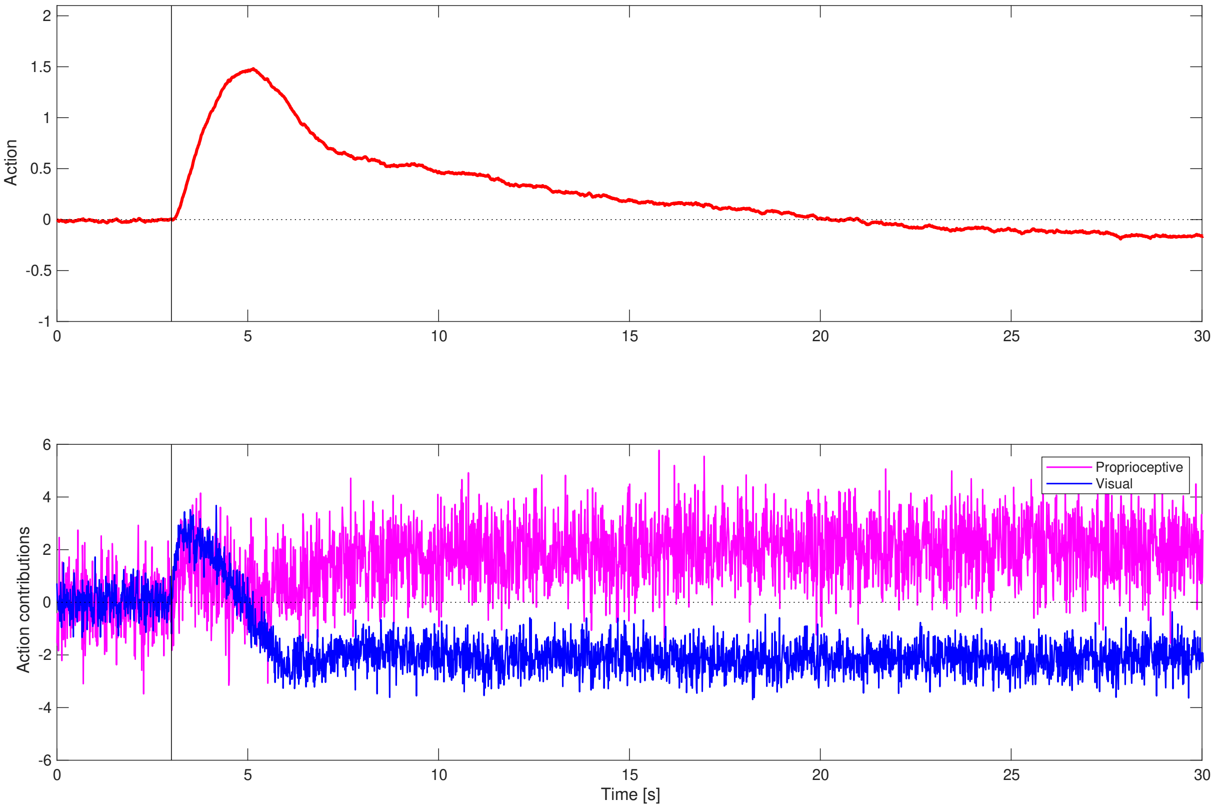

Figure 4.10: Third simulation: action and sensory contributions

Finally, Figure 4.10 depicts the temporal evolution of the action selected by the agent and the different contributions elicited by proprioceptive and visual errors. At steady state, only the model prediction error converges towards zero, while the sensory prediction errors exhibit non-null residuals (at least in mean values); these, in turn, provide a persistent stimulus for the agent's action update (as illustrated in the second panel). In contrast with the previous simulations, the visual and proprioceptive errors elicit opposing contributions: indeed, the visual prediction error (driven by the virtual hand's position) generates a negative update to counteract the observed overshoot of the target; conversely, the proprioceptive error produces a positive contribution attempting to move the real hand towards the target. At steady state, this opposing forces reach an equilibrium condition, failing, however, to reach the target with either of the arms. Indeed, due to the simulation setup, there is no possible way of suppressing both errors: any attempt at minimizing one of the two would increase the other; the achieved steady state represents the optimal compromise between the two. Different combinations of action gains ($k_{A,p}$, $k_{A,v}$) and sensory uncertainties (σ_p, σ_v) would cause a shift of this equilibrium in favor of either the real or the virtual limb, i.e. allowing a better performance in the reaching task for one of the two at the expense of the other's performance; however, there is no combination of parameters that would allow both limbs to reach the target.

This simulation illustrates that during a reaching task under multisensory conflict which cannot be resolved via action, the intentional component largely dominates over the compensatory movements aiming at resolving (or minimizing) the conflict. Hence, demonstrating the difficulty of disentangling the two components in most scenarios.

5 | Conclusions

Building upon Helmholtz's proposal of the brain as an inference machine, this work has successfully modeled motor control as a process of minimizing Variational Free Energy (VFE). The analytical derivation presented in Chapter 2 coupled with the simulations presented in Chapter 3, reinforces the idea that perception and action are not separate functions but are intertwined modalities cooperating during motor control tasks.

The present work's core accomplishments are highlighted below. The first real contribution lies in the rigorous definition and analytical derivation of Variational Free Energy. Indeed, most existing literature do not provide a clear derivation of the Free Energy Principle and focus their work on applying it to specific use cases. Moreover, the lack of a consistent notation and nomenclature hinders the interested reader from comprehending such framework.

Furthermore, a simulation setup was presented applying active inference to specific scenarios. This simulation was deliberately kept simple, since its real objective was not to provide a detailed and realistic model of the brain's behavior during motor control; instead, it served as a proof of concept application demonstrating the feasibility of applying the active inference framework to motor control scenarios. Indeed, the simulation demonstrated the versatility of active inference in different contexts, namely task-oriented movements and multisensory conflict during bodily illusion. Unlike other prominent theories of motor control, active inference demonstrated its ability to handle very different scenarios without having to modify the model; only slight modifications of the parameters were needed for the same agent to model all the three presented scenarios. Specifically, the first scenario illustrated the agent's ability to perform reaching task, an established benchmark performance for motor control theories. The real contribution stems from the second simulation, in which the agent was subject to the Rubber Hand Illusion; through active inference, the same agent was able to model the slight movements observed in experimental settings [2, 20]. Additionally, the agent incorrectly infers its velocity to remain null, therefore providing a possible explanation behind the lack of awareness exhibited by subjects during bodily illusions. However, it is worth noticing that in the described model the agent doesn't receive any proprioceptive information about its velocity: such information would provide a new sensory proprioceptive error that could result in the agent correctly inferring the angular velocity of the arm. Finally, in the third simulation the agent performed a reaching task under an artificially induced multisensory conflict that could not be resolved via action. In this context, the reaching task and the multisensory conflict both play a role in the action selection process, however the intentional contribution prevails resulting in a clear attempt at reaching the target, despite receiving conflicting information about the position of the arm.

The simulations revealed a consistent underlying mechanism: the internal state displays anticipatory movements that drive physical execution by purposefully increasing the proprioceptive prediction error. However, the causes behind these trajectory modifications differ between goal oriented tasks and multisensory conflict scenarios: in the former case, the assigned target would cause a spike in the model prediction error; while, in the latter case, the onset of a bodily illusion would cause a spike in the sensory prediction errors. Thus providing a possible explanation of the different underlying mechanism behind action selection in these conditions. The presence of abrupt spikes in such errors is attributable to the artificial onset of a predefined target or of a multisensory conflict. A more complex agent would be provided with an hierarchical structure that could model the gradual onset of a desired target or of the illusion. Hierarchical models were purposefully avoided in the present work since they are out of its scope and would require a comprehensive and rigorous treatise; in essence, the hierarchical structure is organized so that higher levels define the priors in the levels below. Therefore, an implementation of such models could explain the parameter setting that was adopted; moreover, these parameters would have their own dynamics that could be adapted depending on the context. For example, in the second simulation the visual uncertainty was artificially modified to model an higher degree of mistrust felt towards the visual information coming from the rubber hand; additionally, the action selection was set to be modulated solely by proprioceptive error. The specific characterization of the second simulation could be achieved by an appropriate hierarchical structure that would independently modify the values of such parameters. See [19] for a nice example of an hierarchical active inference agent emulating the higher-level role of dopaminergic signals in decision making processes. Future research requires validating the results herein obtained through actual experimental settings.

In summary, the simulations demonstrated how active inference represents a central theory in motor control that provides a comprehensive model for different scenarios.

Bibliography

[1] R. A. Adams, S. Shipp, and K. J. Friston. Predictions not commands: Active inference in the motor system, 5 2013. ISSN 18632653.

[2] T. Asai. Illusory body-ownership entails automatic compensative movement: for the unified representation between body and action. *Experimental Brain Research*, 233 (3):777–785, 3 2015. ISSN 14321106. doi: 10.1007/s00221-014-4153-0.

[3] M. J. Beal. Variational algorithms for approximate bayesian inference. Technical report, 1998.

[4] R. Bogacz. A tutorial on the free-energy framework for modelling perception and learning. *Journal of Mathematical Psychology*, 76:198–211, 2 2017. ISSN 10960880. doi: 10.1016/j.jmp.2015.11.003.

[5] C. L. Buckley, C. S. Kim, S. McGregor, and A. K. Seth. The free energy principle for action and perception: A mathematical review, 12 2017. ISSN 10960880.

[6] A. Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36:181–204, 2013. ISSN 14691825. doi: 10.1017/S0140525X12000477.

[7] H. H. Ehrsson. Q Stein-The New Handbook of Multisensory Processes 43 The Concept of Body Ownership and Its Relation to Multisensory Integration. Technical report.

[8] A. Ehud and A. Eldad. Perception as a closed-loop convergence process, 2016. URL `https://elifesciences.org/articles/12830#abstract`.

[9] K. Friston. A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360:815–836, 4 2005. ISSN 09628436. doi: 10.1098/rstb.2005.1622.

[10] K. Friston. Variational filtering. *NeuroImage*, 41:747–766, 7 2008. ISSN 10538119. doi: 10.1016/j.neuroimage.2008.03.017. URL `https://linkinghub.elsevier.com/retrieve/pii/S1053811908002462`.

[11] K. Friston. Hierarchical models in the brain. *PLoS Computational Biology*, 4, 2008. ISSN 15537358. doi: 10.1371/journal.pcbi.1000211.

[12] K. Friston. The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13(7):293–301, 7 2009. ISSN 13646613. doi: 10.1016/j.tics.2009.04.005.

[13] K. Friston. What is optimal about motor control?, 11 2011. ISSN 08966273.

[14] K. Friston, J. Kilner, and L. Harrison. A free energy principle for the brain. *Journal of Physiology Paris*, 100:70–87, 7 2006. ISSN 09284257. doi: 10.1016/j.jphysparis.2006.10.001.

[15] K. Friston, J. Mattout, N. Trujillo-Barreto, J. Ashburner, and W. Penny. Variational free energy and the laplace approximation. *NeuroImage*, 34:220–234, 1 2007. ISSN 10538119. doi: 10.1016/j.neuroimage.2006.08.035.

[16] K. Friston, K. Stephan, B. Li, and J. Daunizeau. Generalised filtering. *Mathematical Problems in Engineering*, 2010, 1 2010. ISSN 1024-123X. doi: 10.1155/2010/621670. URL `https://onlinelibrary.wiley.com/doi/10.1155/2010/621670`.

[17] K. J. Friston. Can a single brain cell be surprised?, 12 2025. ISSN 20411723.

[18] K. J. Friston, J. Daunizeau, J. Kilner, and S. J. Kiebel. Action and behavior: A free-energy formulation. *Biological Cybernetics*, 102:227–260, 3 2010. ISSN 03401200. doi: 10.1007/s00422-010-0364-z.

[19] K. J. Friston, T. Shiner, T. FitzGerald, J. M. Galea, R. Adams, H. Brown, R. J. Dolan, R. Moran, K. E. Stephan, and S. Bestmann. Dopamine, affordance and active inference. *PLoS Computational Biology*, 8, 2012. ISSN 15537358. doi: 10.1371/journal.pcbi.1002327.

[20] C. D. Frith, S. J. Blakemore, and D. M. Wolpert. Abnormalities in the awareness and control of action, 12 2000. ISSN 09628436.

[21] X. Fuchs, M. Riemer, M. Diers, H. Flor, and J. Trojan. Perceptual drifts of real and artificial limbs in the rubber hand illusion. *Scientific Reports*, 6, 4 2016. ISSN 20452322. doi: 10.1038/srep24362.

[22] H. W. Harris. Active inference and psychodynamics: a novel integration with applications to depression and stress disorders. *Frontiers in Psychiatry*, 16, 2025. ISSN 16640640. doi: 10.3389/fpsyt.2025.1630858.

[23] 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. 7 2018. doi: 10.1109/DEVLRN.2018.8761005. URL http://arxiv.org/abs/1806.06809http://dx.doi.org/10.1109/DEVLRN.2018.8761005.

[24] P. Lanillos, S. Franklin, A. Maselli, and D. W. Franklin. Active strategies for multisensory conflict suppression in the virtual hand illusion. *Scientific Reports*, 11, 12 2021. ISSN 20452322. doi: 10.1038/s41598-021-02200-7.

[25] P. Lanillos, C. Meo, C. Pezzato, A. A. Meera, M. Baioumy, W. Ohata, A. Tschantz, B. Millidge, M. Wisse, C. L. Buckley, and J. Tani. Active inference in robotics and artificial agents: Survey and challenges. 12 2021. URL http://arxiv.org/abs/2112.01871.

[26] J. Limanowski and K. Friston. Active inference under visuo-proprioceptive conflict: Simulation and empirical results. *Scientific Reports*, 10, 12 2020. ISSN 20452322. doi: 10.1038/s41598-020-61097-w.

[27] C. J. B. M. Rubber hands 'feel' touch that eyes see. *Nature*, 1998.

[28] F. Mannella, F. Maggiore, M. Baltieri, and G. Pezzulo. Active inference through whiskers. *Neural Networks*, 144:428–437, 12 2021. ISSN 18792782. doi: 10.1016/j.neunet.2021.08.037.

[29] A. Maselli, P. Lanillos, and G. Pezzulo. Supplementary materials active inference unifies intentional and conflict-resolution imperatives of motor control. Technical report, 2022.

[30] A. Maselli, P. Lanillos, and G. Pezzulo. Active inference unifies intentional and conflict-resolution imperatives of motor control. *PLOS Computational Biology*, 18: e1010095, 6 2022. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1010095. URL https://dx.plos.org/10.1371/journal.pcbi.1010095.

[31] G. Oliver, P. Lanillos, and G. Cheng. An empirical study of active inference on a humanoid robot. *IEEE Transactions on Cognitive and Developmental Systems*, 14: 462–471, 6 2022. ISSN 23798939. doi: 10.1109/TCDS.2021.3049907.

[32] T. Parr and K. J. Friston. The anatomy of inference: Generative models and brain structure, 11 2018. ISSN 16625188.

[33] T. Parr, J. Limanowski, V. Rawji, and K. Friston. The computational neurology of movement under active inference. *Brain*, 144:1799–1818, 2021. ISSN 14602156. doi: 10.1093/brain/awab085.

[34] T. Parr, G. Pezzullo, and K. J. Friston. The free energy principle in mind,

brain, and behavior. Technical report, 2022. URL `http://direct.mit.edu/books/oa-monograph/chapter-pdf/2004452/f000200_9780262369978.pdf`.

[35] L. U. Perrinet, R. A. Adams, and K. J. Friston. Active inference, eye movements and oculomotor delays. *Biological Cybernetics*, 108:777–801, 12 2014. ISSN 14320770. doi: 10.1007/s00422-014-0620-8.

[36] G. Pezzulo, T. Parr, and K. Friston. Active inference as a theory of sentient behavior, 2 2024. ISSN 18736246.

[37] L. Pio-Lopez, A. Nizard, K. Friston, and G. Pezzulo. Active inference and robot control: A case study. *Journal of the Royal Society Interface*, 13, 9 2016. ISSN 17425662. doi: 10.1098/rsif.2016.0616.

[38] M. Priorelli, F. Maggiore, A. Maselli, F. Donnarumma, D. Maisto, F. Mannella, I. P. Stoianov, and G. Pezzulo. Modeling motor control in continuous time active inference: A survey. *IEEE Transactions on Cognitive and Developmental Systems*, 16:485–500, 4 2024. ISSN 23798939. doi: 10.1109/TCDS.2023.3338491.

[39] J. Raber, A. Casile, T. Verbelen, I. P. Stoianov, and M. Priorelli. Flexible intentions: An active inference theory. Technical report.

[40] J. Raber, A. Casile, T. Verbelen, I. P. Stoianov, and M. Priorelli. Flexible intentions: An active inference theory. Technical report, 2023.

[41] C. E. Rasmussen and M. Opper. The variational gaussian approximation revisited. Technical report, 2007.

[42] T. R. Ruikes. *Neural dynamics and coding in the cortico-hippocampal hierarchy during multisensory and mnemonic processing*. [Thijs Ruikes], 2025. ISBN 9789465104331.

[43] R. Smith, K. J. Friston, and C. J. Whyte. A step-by-step tutorial on active inference and its application to empirical data. *Journal of Mathematical Psychology*, 107, 4 2022. ISSN 10960880. doi: 10.1016/j.jmp.2021.102632. URL `https://github.com/rssmith33/Active-Inference-Tutorial-Scriptshttps://github.com/infer-actively/pymdp`.

[44] E. Todorov and M. I. Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5:1226–1235, 2002. ISSN 10976256. doi: 10.1038/nn963.

[45] J. Voigts, D. H. Herman, and T. Celikel. Tactile object localization by anticipatory whisker motion. *Knutsen and Ahissar*, 113:620–632, 2015. doi: 10.1152/jn.00241.2014.-Ro. URL `www.jn.org`.

[46] H. von Helmholtz. *Treatise on Physiological Optics*, volume 3. Optical Society of America, 1924.

A | Mathematical tools

This appendix provides an overview over the mathematical tools necessary for a comprehensive mathematical formulation of the Free Energy Principle. This will not be a complete description of such tools; instead, it will provide the minimum knowledge necessary to understand the FEP and its formulation.

A.1. Surprise

In information theory, surprise is a measure of the likelihood of an event occurring. It is also referred to as information content. The latter definition stems from the following principle: the 'information value' of a communicated message depends on how surprising its content is; the more surprising the content will be the more valuable the message will be. For instance, the knowlodge that it won't snow on the 15th of August is an information of little value since it is very likely that such an event won't occur; otherwise, the information that it will indeed snow on the 15th of August is very valuable since the occurence of such an event is highly surprising.

Mathematically, surprise is defined as:

$$\mathcal{I}(x) = -\ln p(x), \tag{A.1}$$

where x is the event and $p(x)$ denotes the probability of the event x occurring. Straightforwardly, when the event's occurrence is highly likely ($p(x)$ approaching 1), surprise will be null; while the surprise associated to an unlikely event ($p(x)$ close to 0), will approach infinity. Note that the logarithm is added to assure that the joint surprise of independent events is additive. Indeed, the joint probability of independent events is the product of the individual probabilities of such events:

$$p(x_1, x_2) = p(x_1)p(x_2). \tag{A.2}$$

Therefore:

$$\begin{aligned}\mathcal{I}(x_1, x_2) &= -\ln p(x_1, x_2) = \\ &= -\ln p(x_1)p(x_2) = \\ &= -\ln p(x_1) - \ln p(x_2) = \\ &= \mathcal{I}(x_1) + \mathcal{I}(x_2).\end{aligned} \tag{A.3}$$

The following example provides a clear and simple explanation of surprise and will later be used to explain entropy, cross-entropy and the Kullback-Leibler divergence.
For instance, let's assume that the weather forecasting for today was 85% probability of having a sunny day and a 15% probability of having a rainy day. Based on this distribution, the expected outcome is that today will be sunny, while the surprising outcome is that today turns out to be rainy. We can quantify the surprise associated to such events by evaluating their information content (or, indeed, surprise):

$$\mathcal{I}(\text{sunny}) = -\ln p(\text{sunny}) = -\ln 0.85 \approx 0.16, \tag{A.4}$$

$$\mathcal{I}(\text{rainy}) = -\ln p(\text{rainy}) = -\ln 0.15 \approx 1.90. \tag{A.5}$$

Notably, surprise is able to quantify the atypical outcome of actually experiencing a rainy day, when a sunny day was forecasted.

A.2. Shannon Entropy

The entropy of a stochastic process quantifies the average surprise associated to the possible states of said variable. In the case of a discrete random variable X, it is computed as the weighted sum of the surprises associated to all its possible states:

$$\mathcal{H}(X) = -\sum_{x \in X} p(x) \ln p(x) = \sum_{x \in X} p(x)\mathcal{I}(x), \tag{A.6}$$

where x is a particular occurrence of the random variable X. Equivalently, it can be evaluated for a continuous distribution with:

$$\mathcal{H}(X) = -\int p(x) \ln p(x) dx = \int p(x)\mathcal{I}(x) dx. \tag{A.7}$$

Entropy measures the unpredictability of a random variable; indeed, entropy approaches a null value (its minimum) when an event of the random variable is certain (i.e. its probability is 1); on the other hand it assumes its maximum value for a uniform distribution (the case of maximum uncertainty on the outcome):

$$\mathcal{H}_{max}(X) = -\sum_{x \in X} \frac{1}{N} \ln \frac{1}{N} = N \frac{1}{N} \ln N = \ln N, \tag{A.8}$$

where N is the number of possible outcomes of an event.

Revisiting the example discussed in Section A.1, an extended version is presented to better understand the concept of entropy. Consider the following probabilities in weather forecasting:

Event	**Probability** $p(x)$	**Surprise** $\mathcal{I}(x)$
Sunny	0.5	0.69
Rainy	0.25	1.39
Cloudy	0.25	1.39

Table A.1: Weather forecasting example: probability distribution and surprise evaluation

Entropy can be computed as follows:

$$\mathcal{H}(X) = p(sunny)\mathcal{I}(sunny) + p(rainy)\mathcal{I}(rainy) + p(cloudy)\mathcal{I}(cloudy) = 1.04. \tag{A.9}$$

In the case of a sure forecasting entropy will tend to a null value. The computation of the limit is omitted, but it is of simple resolving since the surprise associated to events that differ from the one with $p(x) = 1$ will approach infinity: indeed, such events which are impossible ($p(x) = 0$) according to the probability distribution. The nil value assures complete predictability over the weather. Indeed, according to this distribution, the weather will certainly be sunny.

Event	**Probability** $p(x)$	**Surprise** $\mathcal{I}(x)$
Sunny	1	0
Rainy	0	$\to +\infty$
Cloudy	0	$\to +\infty$

Table A.2: Weather forecasting example: sure forecasting

Finally, in the case of a uniform distribution entropy will assume its maximum value, since there is complete uncertainty on the outcome of the event:

Event	**Probability** $p(x)$	**Surprise** $\mathcal{I}(x)$
Sunny	$1/3$	$\ln 3$
Rainy	$1/3$	$\ln 3$
Cloudy	$1/3$	$\ln 3$

Table A.3: Weather forecasting example: uniform distribution

$$\mathcal{H}(X) = 3 \cdot \frac{1}{3} \cdot \ln 3 = \ln 3. \tag{A.10}$$

Such forecasting is essentially meaningless, as it amounts to mere guessing without any informative basis.

A.3. Cross-entropy

Cross-entropy evaluates the average surprise experienced when the real probability distribution $p(x)$ is modeled by an auxiliary distribution $q(x)$. It is calculated as the weighted sum of the surprise experienced by $q(x)$ where the weights are the real probabilities $p(x)$:

$$\mathcal{H}(p, q) = -\sum_{x \in X} p(x) \ln q(x) = \sum_{x \in X} p(x) \mathcal{I}_q(x). \tag{A.11}$$

Equivalently, it can be evaluated for a continuous distribution with:

$$\mathcal{H}(p, q) = -\int p(x) \ln q(x) dx = \int p(x) \mathcal{I}_q(x) dx. \tag{A.12}$$

It can be demonstrated that cross-entropy assumes its minimum value when the distribution $q(x)$ models exactly the true distribution $p(x)$. In such a case the cross-entropy essentially becomes the entropy of the real distribution $p(x)$:

$$\mathcal{H}_{min}(p, q) = -\sum_{x \in X} p(x) \ln p(x) = \mathcal{H}(p) \tag{A.13}$$

Note that cross-entropy is not a commutative operator:

$$\mathcal{H}(p, q) \neq \mathcal{H}(q, p). \tag{A.14}$$

A.4. Kullback-Leibler divergence

The Kullback-Leibler divergence (or relative entropy) is defined as the difference between the cross-entropy $\mathcal{H}(p,\, q)$ and the entropy $\mathcal{H}(p)$:

$$\begin{aligned}\mathcal{D}_{KL}\big[p||q\big] &= \mathcal{H}(p,\, q) - \mathcal{H}(p) = \\ &= -\sum_{x\in X} p(x)\ln q(x) - \left(-\sum_{x\in X} p(x)\ln p(x)\right) = \\ &= \sum_{x\in X} p(x)\ln\frac{p(x)}{q(x)}. \end{aligned} \tag{A.15}$$

Analogously, it can be defined for continuous distributions as:

$$\mathcal{D}_{KL}\big[p(x)||q(x)\big] = \int_{x\in X} p(x)\ln\frac{p(x)}{q(x)}dx. \tag{A.16}$$

It represents the information gain achieved when moving from a prior belief to a posterior belief. In Machine Learning, it is often used as a loss function to minimize the "distance" between a model's prediction and the ground truth: indeed, it is a measure of discrepancy between the two distributions $p(x)$ and $q(x)$. Therefore, it assumes its minimum value when $q(x) = p(x)$; indeed:

$$\mathcal{D}_{KL}\big[p||q\big]_{min} = \mathcal{H}(p,\, q) - \mathcal{H}(p) = \mathcal{H}(p,\, p) - \mathcal{H}(p) = \mathcal{H}(p) - \mathcal{H}(p) = 0. \tag{A.17}$$

Notably, the Kullback Leibler divergence is an always positive quantity: because of this property it is often adopted as a loss function to be minimized (as in the case of Free Energy).

B | VFE - Definition and Forms

This appendix clarifies different formulations of Variational Free Energy (VFE) that can be encountered in literature. Although equivalent, these formulations shed light on different features and properties that VFE possesses. Moreover, the definition of VFE is debated. In particular, the decision to exploit the reverse Kullback-Leibler divergence instead of the forward formulation is discussed.

B.1. VFE alternative formulations

All VFE's formulations derive from different manipulations of its definition that was presented in Equation (2.4):

$$\mathcal{F} = \mathcal{D}_{KL}\big[q(x) \,||\, p(x,s)\big] = \int q(x) \ln \frac{q(x)}{p(x,s)} dx, \tag{B.1}$$

where the Kullback-Leibler definition (Section A.4) was exploited. The following formulation of $\mathcal{F}$ was derived in Equation (2.6):

$$\mathcal{F} = -\ln p(s) + \mathcal{D}_{KL}\big[q(x) \,||\, p(x|s)\big]. \tag{B.2}$$

Elaborating differently VFE's definition another formulation of VFE can be derived:

$$\begin{aligned}\mathcal{F} &= \int q(x) \ln \frac{q(x)}{p(x,s)} dx \\ &= -\int q(x) \ln p(x,s) dx + \int q(x) \ln q(x) dx \\ &= \mathcal{H}\big(q(x), p(x,s)\big) - \mathcal{H}\big(q(x)\big).\end{aligned} \tag{B.3}$$

This corresponds to applying the definition of the Kullback-Leibler divergence given in Appendix A. According to this formulation, VFE is the difference between the cross-entropy between the recognition density $q(x)$ and the joint probability $p(x, s)$ and the entropy of the recognition density. Minimizing VFE is equivalent to finding the most "economical" explanation of sensory data—one that is accurate enough to explain the observations (low cross-entropy) but simple enough enough to avoid overfitting (high entropy) [18].

The third and last formulation can be derived by applying the chain rule to the joint probabilty $p(x, s)$ in the following manner:

$$\begin{aligned}
\mathcal{F} &= \int q(x) \ln \frac{q(x)}{p(x,s)} dx \\
&= \int q(x) \ln \frac{q(x)}{p(s|x)p(x)} dx \\
&= -\int q(x) \ln p(s|x) dx + \int q(x) \ln \frac{q(x)}{p(x)} dx \\
&= \mathcal{H}\big(q(x), p(s|x)\big) + \mathcal{D}_{KL}\big[q(x) \,||\, p(x)\big]
\end{aligned} \tag{B.4}$$

where in the last step the definitions of cross-entropy and Kullback-Leibler divergence (given respectively in Section A.3 and Section A.4) were adopted. Such formulation expresses VFE as a measure of complexity minus accuracy. Indeed, the second term is the Kullback-Leibler divergence between the recognition density $q(x)$ (the agent's prior belief) and the probability distribution of hidden states $p(x)$; such term grows larger with higher discrepancy between the two probability distributions; in other words, it assumes large values when the agent needs to make larger revisions to its belief to explain novel empirical data [43]. thus providing a measure of complexity. A larger complexity results in a limitation of the predictive power of a model. Conversely, the first term is the cross-entropy between the recognition density $q(x)$ and the posterior density $p(s|x)$: a metric that quantifies the amount of sensory detail successfully captured by the model. Consequently, minimizing VFE is equivalent to minimizing the difference between complexity and accuracy of the chosen model; i.e. the agent will try acquire an optimal model which is simple and accurate at the same time.

B.2. VFE definition discussion

Here, the definition and consequent formulations of Variational Free Energy are discussed. Indeed, upon its initial introduction in Section 2.1, alternative definitions of VFE may have been considered.

The definition of VFE in Equation (2.4) was derived to construct an upper bound of the Kullback-Leibler divergence between the recognition density $q(x)$ and the posterior probability $p(x|s)$. By confronting with the definition given in Section A.4, it is evident that during VFE's formulation we are considering the reverse Kullback-Leibler divergence instead of the forward one that is usually adopted in information theory:

$$\mathcal{D}_{KL}\big[p(x|s) \,||\, q(x)\big]. \tag{B.5}$$

This decision is driven by different factors, here explained.

Following the original definition's line of thought, one may define VFE as the forward Kullback-Leibler divergence between the joint probability $p(x, s)$ and $q(x)$:

$$\mathcal{F} := \mathcal{D}_{KL}\big[p(x,s) \,||\, q(x)\big] \tag{B.6}$$

The consequent issues to such a definition arise when expanding Equation (B.6) as follows:

$$\begin{aligned}
\mathcal{F} &= \mathcal{D}_{KL}\big[p(x,s) \,||\, q(x)\big] \\
&= \int p(x,s) \ln \frac{p(x,s)}{q(x)} dx \\
&= \int p(x|s)p(s) \ln \frac{p(x|s)p(s)}{q(x)} dx \\
&= p(s) \left[\int p(x|s) \ln \frac{p(x|s)}{q(x)} dx + \ln p(s) \int p(x|s) dx \right] \\
&= p(s) \left[\int p(x|s) \ln \frac{p(x|s)}{q(x)} dx + \ln p(s) \right] \\
&= p(s) \Big[\mathcal{D}_{KL}\big[p(x|s) \,||\, q(x)\big] - \mathcal{I}(s) \Big]
\end{aligned} \tag{B.7}$$

Equation (B.7) highlights two critical issues regarding this definition.
Firstly, even if we introduced approximations analogous to the ones described in Sec-

tion 2.2, the factored out $p(s)$ (consequently also $\mathcal{I}(s)$) would remain intractable as in the case of classical Bayesian Inference presented in Section 2.1.1.

Secondly, unlike the case of VFE's reverse definition, it is unclear if the forward definition represents an upper bound on the forward Kullback-Leibler divergence between $p(x|s)$ and $q(x)$. Further examination is required:

$$\mathcal{F} = p(s)\Big[\mathcal{D}_{KL}\big[p(x|s) \,||\, q(x)\big] - \mathcal{I}(s)\Big] \overset{?}{>} \mathcal{D}_{KL}\big[p(x|s) \,||\, q(x)\big]$$

$$\big(p(s) - 1\big)\mathcal{D}_{KL}\big[p(x|s) \,||\, q(x)\big] > p(s)\mathcal{I}(s)$$

Remembering that $p(s)$ is a probability distribution and thus assumes values in the interval $[0, 1]$, it can be noticed that $p(s) - 1$ is in general negative; therefore:

$$\mathcal{D}_{KL}\big[p(x|s) \,||\, q(x)\big] < \frac{p(s)\mathcal{I}(s)}{p(s) - 1} < 0 \tag{B.8}$$

Disequation (B.8) is impossible since the Kullback-Leibler divergence is an always positive quantity as discussed in Section A.4. Therefore, for no combination of sensory data and hidden states the forward VFE is an upper bound on the forward divergence between $p(x|s)$ and $q(x)$. Its minimization would not imply a better approximation of the posterior probability adopting the recognition density.

C | Code

This appendix contains both the code that was employed to run the simulations and the code used to print the graphs presented in Chapter 4. All the code is also available in a repository on GitHub (https://github.com/alessandroarrigoni00/Master_thesis-Active_inference/tree/main).

C.1. Simulations code

Listing C.1: Simulation code

```
clc;
clear all;
close all;

%% Simulation parameters

simTime = 30; % seconds
dt = 0.01;
time = 0:dt:simTime;
N = length(time);
onset_time = 3;

%% Generative process

A = zeros (N, 3); % Action
x = zeros (2, N, 3); % Hidden states
s = zeros (3, N, 3); % Sensory input (proprioceptive + visual)
xp = x; % Environment dynamics

% Sensory noise
Sigma_p = .1;
Sigma_v = .01;
```

```
n_p = randn(1,N)*Sigma_p; % Sensory noise on proprioception
n_v = randn(2,N)*Sigma_v; % Sensory noise on visual coordinates

% Parameters
L_arm = .45; % Length of the arm (m)
m_arm = 1.5; % Mass of the arm
visc = 4; % Viscosity
beta = .5; % Power dependence between damping and velocity

%% Generative model

mu_x = zeros (3, N, 3); % Hidden states estimates
mu_s = zeros (3, N, 3); % Sensory input estimates
mup = zeros (3, N, 3); % Internal dynamics

K = 1.3;
Sigma_mup = .01;
Sigma_mupp = .1;

% Gains
k_mu = [.1; .01; .001];

% Gradients
dF_dmu = zeros (3, N, 3);
dF_dA = zeros (2, N, 3);
dgvx_dmu = zeros (N, 3);
dgvy_dmu = zeros (N, 3);

%% VFE and errors

VFE = zeros (N, 3);
epsilon_s = zeros (3, N, 3);
epsilon_mu = zeros (2, N, 3);

% Target

mu_teta_t = zeros (1, N);
mu_teta_t (onset_time/dt:end) = pi/3;

```

```
% Parameters

Sigma_sv = .01;
Sigma_sp = .1;
k_A = [.3, .3]; % k_A = [k_Ap, k_Av]
v_gain = 1.3;

% Initial conditions

x(:,1,:) = zeros(size(x(:,1,:)));
mu_x(:,1,:) = zeros(size(mu_x(:,1,:)));
mup(3,:,:) = zeros(size(mup(3,:,:)));
A(1,:) = zeros(size(A(1,:)));

x_vh = zeros(size(x(:,:,3)));

%% First simulation

for i = 1:N
    mup(1:2, i, 1) = [mu_x(2,i,1), (K*(mu_teta_t(i)-mu_x(1,i,1))-
       visc*mu_x(2,i,1))/m_arm];
    mu_s(:,i,1) = [mu_x(1,i,1), L_arm*cos(mu_x(1,i,1) + pi/2),
       L_arm*sin(mu_x(1,i,1) + pi/2)];
    s(:,i,1) = [x(1,i,1) + n_p(i),      L_arm*cos(x(1,i,1)+pi/2) +
       n_v(1,i),      L_arm*sin(x(1,i,1)+pi/2) + n_v(2,i)];

    epsilon_s(:,i,1) = s(:,i,1) - mu_s(:,i,1);
    epsilon_mu(:,i,1) = mu_x(2:3,i,1) - mup (1:2,i,1);
    VFE(i, 1) = 1/(2*Sigma_mupp)*epsilon_mu(2,i,1)^2 + 1/(2*
       Sigma_sp)*epsilon_s(1,i,1)^2 + 1/(2*Sigma_sv)*(epsilon_s(2,
       i,1)^2 + epsilon_s(3,i,1)^2);

    dgvx_dmu(i,1) = -L_arm * sin(mu_x(1,i,1) + pi/2);
    dgvy_dmu(i,1) = L_arm * cos(mu_x(1,i,1) + pi/2);

    dF_dmu(3,i,1) = -1/Sigma_mupp * epsilon_mu(2,i,1);
    dF_dmu(1,i,1) = (K/m_arm)*dF_dmu(3,i,1) + 1/Sigma_sp*epsilon_s
       (1,i,1) + 1/Sigma_sv*(epsilon_s(2,i,1)*dgvx_dmu(i,1) +
       epsilon_s(3,i,1)*dgvy_dmu(i,1));
```

```
    dF_dmu(2,i,1) = -dF_dmu(3,i,1) * (-visc/m_arm);

    dF_dA(1,i,1) = -(m_arm/K) * (1/Sigma_sp) * epsilon_s(1,i,1);
    dF_dA(2,i,1) = -(m_arm/K) * (1/Sigma_sv) * (epsilon_s(2,i,1)*
       dgvx_dmu(i,1) + epsilon_s(3,i,1)*dgvy_dmu(i,1));

    if i~=N
        mu_x(:,i+1,1) = mu_x(:,i,1) + dt*(mup(:,i,1)+k_mu.*dF_dmu
           (:,i,1));
        A(i+1,1) = A(i,1) + dt*k_A*dF_dA(:,i,1);
        x(:,i+1,1) = x(:,i,1) + dt*[x(2,i,1);     A(i+1,1)-(visc/
           m_arm)*sign(x(2,i,1))*abs(x(2,i,1))^beta];
    end
end

%% Third simulation

for i = 1:N
    mup(1:2, i, 3) = [mu_x(2,i,3), (K*(mu_teta_t(i)-mu_x(1,i,3))-
       visc*mu_x(2,i,3))/m_arm];
    mu_s(:,i,3) = [mu_x(1,i,3), L_arm*cos(mu_x(1,i,3) + pi/2),
       L_arm*sin(mu_x(1,i,3) + pi/2)];
    s(:,i,3) = [x(1,i,3) + n_p(i),     L_arm*cos(x_vh(1,i)+pi/2) +
        n_v(1,i),     L_arm*sin(x_vh(1,i)+pi/2) + n_v(2,i)];

    epsilon_s(:,i,3) = s(:,i,3) - mu_s(:,i,3);
    epsilon_mu(:,i,3) = mu_x(2:3,i,3) - mup (1:2,i,3);
    VFE(i, 3) = 1/(2*Sigma_mupp)*epsilon_mu(2,i,3)^2 + 1/(2*
       Sigma_sp)*epsilon_s(1,i,3)^2 + 1/(2*Sigma_sv)*(epsilon_s(2,
       i,3)^2 + epsilon_s(3,i,3)^2);

    dgvx_dmu(i,3) = -L_arm * sin(mu_x(1,i,3) + pi/2);
    dgvy_dmu(i,3) = L_arm * cos(mu_x(1,i,3) + pi/2);

    dF_dmu(3,i,3) = -1/Sigma_mupp * epsilon_mu(2,i,3);
    dF_dmu(1,i,3) = (K/m_arm)*dF_dmu(3,i,3) + 1/Sigma_sp*epsilon_s
       (1,i,3) + 1/Sigma_sv*(epsilon_s(2,i,3)*dgvx_dmu(i,3) +
       epsilon_s(3,i,3)*dgvy_dmu(i,3));
    dF_dmu(2,i,3) = -dF_dmu(3,i,3) * (-visc/m_arm);
```

```
    dF_dA(1,i,3) = -(m_arm/K) * (1/Sigma_sp) * epsilon_s(1,i,3);
    dF_dA(2,i,3) = -(m_arm/K) * (1/Sigma_sv) * (epsilon_s(2,i,3)*dgvx_dmu(i,3) + epsilon_s(3,i,3)*dgvy_dmu(i,3));

    if i~=N
        mu_x(:,i+1,3) = mu_x(:,i,3) + dt*(mup(:,i,3)+k_mu.*dF_dmu(:,i,3));
        A(i+1,3) = A(i,3) + dt*k_A*dF_dA(:,i,3);
        x(:,i+1,3) = x(:,i,3) + dt*[x(2,i,3);     A(i+1,3)-(visc/m_arm)*sign(x(2,i,3))*abs(x(2,i,3))^beta];
        x_vh(2,i+1) = v_gain*x(2,i,3);
        x_vh(1,i+1) = x_vh(1,i) + dt*x_vh(2,i);
    end
end

%% Second simulation

% Rubber hand
teta_rub = zeros (1,N);
teta_rub (onset_time/dt:end) = pi/3;

s(2,:,2) = L_arm * cos(teta_rub + pi/2) + n_v(1,:);
s(3,:,2) = L_arm * sin(teta_rub + pi/2) + n_v(2,:);

% Parameters

k_A(2) = 0;
Sigma_sv = .07;

for i = 1:N
    mup(1:2, i, 2) = [mu_x(2,i,2), (-visc*mu_x(2,i,2))/m_arm];
    mu_s(:,i,2) = [mu_x(1,i,2), L_arm*cos(mu_x(1,i,2) + pi/2), L_arm*sin(mu_x(1,i,2) + pi/2)];
    s(1,i,2) = x(1,i,2) + n_p(i);

    epsilon_s(:,i,2) = s(:,i,2) - mu_s(:,i,2);
    epsilon_mu(:,i,2) = mu_x(2:3,i,2) - mup (1:2,i,2);
```

```
    VFE(i, 2) = 1/(2*Sigma_mupp)*epsilon_mu(2,i,2)^2 + 1/(2*
       Sigma_sp)*epsilon_s(1,i,2)^2 + 1/(2*Sigma_sv)*(epsilon_s(2,
       i,2)^2 + epsilon_s(3,i,2)^2);

    dgvx_dmu(i,2) = -L_arm * sin(mu_x(1,i,2) + pi/2);
    dgvy_dmu(i,2) = L_arm * cos(mu_x(1,i,2) + pi/2);

    dF_dmu(3,i,2) = -1/Sigma_mupp * epsilon_mu(2,i,2);
    dF_dmu(1,i,2) = (K/m_arm)*dF_dmu(3,i,2) + 1/Sigma_sp*epsilon_s
       (1,i,2) + 1/Sigma_sv*(epsilon_s(2,i,2)*dgvx_dmu(i,2) +
       epsilon_s(3,i,2)*dgvy_dmu(i,2));
    dF_dmu(2,i,2) = -dF_dmu(3,i,2) * (-visc/m_arm);

    dF_dA(1,i,2) = -(m_arm/K) * (1/Sigma_sp) * epsilon_s(1,i,2);
    dF_dA(2,i,2) = -(m_arm/K) * (1/Sigma_sv) * (epsilon_s(2,i,2)*
       dgvx_dmu(i,2) + epsilon_s(3,i,2)*dgvy_dmu(i,2));

    if i~=N
        mu_x(:,i+1,2) = mu_x(:,i,2) + dt*(mup(:,i,2)+k_mu.*dF_dmu
           (:,i,2));
        A(i+1,2) = A(i,2) + dt*k_A*dF_dA(:,i,2);
        x(:,i+1,2) = x(:,i,2) + dt*[x(2,i,2);     A(i+1,2)-(visc/
           m_arm)*sign(x(2,i,2))*abs(x(2,i,2))^beta];
    end
end
```

C.2. Figures code

Listing C.2: Figures code

```
load("results.mat");
x_lim = [0, simTime];

for sim =1:3
    figure
    % Joint angle
    subplot(2,1,1)
    plot(time, mu_teta_t,'color',[0, 109/255,2/255] ,'linewidth',
       2);
    hold on
    plot(time, x(1,:,sim),'k','linewidth', 2)
    plot(time, mu_x(1,:,sim),'r','linewidth', 2)
    if(sim == 3)
        plot(time, x_vh(1,:),'color',[64/255, 103/255,174/255] ,'
           linewidth', 2)
    end
    xlim(x_lim)
    yrange = [-0.2, 1.8];
    ylim(yrange);
    ylabel('Joint angle');
    labels = {'Target', 'Real', 'Inferred', 'Virtual'};
    if (sim ==2)
        labels{1} = 'Rubber hand';
    end
    legend(labels(1:end-1))
    if (sim == 3)
        legend(labels)
    end

    % Joint angle velocity
    subplot(2,1,2)
    plot(time, x(2,:,sim),'k','linewidth', 2);
    hold on;
    plot(time, mu_x(2,:,sim),'r','linewidth', 2)
    if (sim == 3)
        plot(time, x_vh(2,:))
```

```
    end
    xlim(x_lim);
    yrange = [-0.1, 0.5];
    if (sim ==2)
        yrange = [-0.3, 0.5];
    end
    ylim(yrange);
    plot([time_t time_t], yrange,'k')
    plot([0 simTime], [0 0],'k:')
    ylabel('Angular␣velocities')
    labels = {'Real', 'Inferred', 'Virtual'};
    legend(labels(1:end-1))
    if (sim == 3)
        legend(labels)
    end
    xlabel('Time␣[s]')

    figure
    % Action
    subplot(2,1,1)
    plot(time, A(:,sim),'r','linewidth', 2);
    hold on;
    xlim(x_lim);
    yrange = [-1, 2.1];
    ylim(yrange)
    plot([time_t time_t], yrange,'k')
    plot([0 simTime], [0 0],'k:')
    ylabel('Action')

    % Action contribution
    subplot(2,1,2)
    plot(time, dF_dA(1,:,sim),'m','linewidth', 1);
    hold on;
    plot(time, dF_dA(2,:,sim),'b','linewidth', 1)
    xlim(x_lim)
    yrange = [-6, 6];
    ylim(yrange)
    plot([time_t time_t], yrange,'k')
    plot([0 simTime], [0 0],':k')
```

```
    ylabel('Action contributions')
    legend('Proprioceptive', 'Visual')
    xlabel('Time [s]')

    figure
    % Errors
    subplot(2,1,1)
    plot(time, epsilon_s(1,:,sim),'color',[1, 125/255, 0],'linewidth', 2); % proprioceptive sensory error
    hold on;
    plot(time, epsilon_s(2,:,sim),'m','linewidth', 1) % visual sensory error (x axis)
    plot(time, epsilon_s(3,:,sim),'b','linewidth', 1) % visual sensory error (y axis)
    plot(time, epsilon_mu(2,:,sim),'color','g','linewidth', 2) % model error
    xlim(x_lim)
    yrange = [-1.3, 1.3];
    if (sim == 2)
        yrange = [-0.45, 0.45];
    end
    ylim(yrange)
    plot([time_t time_t], yrange,'k')
    plot([0 simTime], [0 0],'k:')
    ylabel('Errors')
    legend ('Proprioceptive', 'Visual (x)', 'Visual (y)', 'Model')

    % Free Energy
    subplot(2,1,2)
    plot(time, VFE(:,sim),'r','linewidth', 2); % VFE
    hold on;
    xlim(x_lim)
    yrange = [-0.1, 5];
    ylim(yrange)
    plot([time_t time_t], yrange,'k')
    plot([0 simTime], [0 0],'k:')
    ylabel('Free Energy')
    xlabel('Time [s]')
end
```

List of Figures

List of Tables

List of Symbols

Table C.1: List of Symbols

Symbol	Description
x	Environment's hidden state
s	Sensory input received
$\mathcal{I}(\,\cdot\,)$	Surprise
$\mathcal{D}_{KL}\big[\,\cdot\,\|\,\cdot\,\big]$	Kullback-Leibler divergence
$q(x)$	Recognition density
$\mathcal{F}$	Variational Free Energy
μ	Mean under the Laplace approximation
Σ	Variance under the Laplace approximation
Z	$q(x)$'s normalization factor under the Laplace approximation
ε	$q(x)$'s exponent under the Laplace approximation
$\mathbb{E}\big[\,\boldsymbol{\cdot}\,\big]$	Expected value
Σ^*	Optimal variance under the Laplace approximation
$\mathcal{N}(\cdot\,;\cdot\,,\cdot)$	Normal distribution
$\mathcal{L}(\cdot)$	Laplacian encoded energy
$g(\mu,\theta)$	Mapping function
$\overline{\mu}$	Constant mean in the static case
n_s	Sensory noise
n_μ	Noise term in the internal dynamics
σ_s^2	Variance of the sensory noise term
σ_μ^2	Variance of the internal noise term
ϵ_s	Sensory prediction error

Continued on next page...

Continued from previous page

Symbol	Description
ϵ_{μ}	Model prediction error
$f(\cdot)$	Dynamics mapping function
$(\cdot)^{[n]}$	N-th order derivative
$\tilde{\boldsymbol{s}}$	Sensory vector in generalized coordinates
$\tilde{\boldsymbol{\mu}}$	Internal state vector in generalized coordinates
$\tilde{\boldsymbol{n}}_{\boldsymbol{s}}$	Sensory noise vector in generalized coordinates
$\tilde{\boldsymbol{n}}_{\boldsymbol{\mu}}$	Internal noise vector in generalized coordinates
$\tilde{\boldsymbol{g}}$	Sensory mapping function in generalized coordinates
$\tilde{\boldsymbol{f}}$	Internal dynamics mapping function in generalized coordinates
$\sigma^2_{s[n]}$	Variance of the sensory noise term at the n-th order derivative
$\sigma^2_{\mu[n]}$	Variance of the sensory noise term at the n-th order derivative
$\epsilon_s^{[n]}$	Sensory prediction error at the n-th order derivative
$\epsilon_\mu^{[n]}$	Model prediction error at the n-th order derivative
$\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{\alpha}}\}$	Multivariate internal state vector
$\{\tilde{\boldsymbol{s}}_{\boldsymbol{\alpha}}\}$	Multivariate sensory vector
$\Sigma_{s[n],\alpha}$	Sensory covariance matrix
$\Sigma_{\mu[n],\alpha}$	Internal state covariance matrix
$\dot{\tilde{\boldsymbol{\mu}}}$	Internal state expected dynamics in generalized coordinates
dt	Discrete time interval
k_μ	Internal state learning rate
A	Action selected by the agent
k_A	Action learning rate
$\tilde{\boldsymbol{x}}$	Generative process hidden state in generalized coordinates
$\tilde{\boldsymbol{s}}$	Generative process sensory vector in generalized coordinates
$\dot{\tilde{\boldsymbol{x}}}$	Generative process hidden state dynamics in generalized coordinates
$\boldsymbol{f}_{\boldsymbol{x}}$	Generative process hidden state dynamics mapping
ϕ	Generative process viscosity coefficient
β	Generative process viscosity power dependence
m_{arm}	Real arm's mass

Continued on next page...

Continued from previous page

Symbol	Description
s_p	Proprioceptive sensory signal
s_{vx}	Visual sensory signal (on the x axis)
s_{vy}	Visual sensory signal (on the x axis)
L_{arm}	Real arm's length
n_p	Proprioceptive noise term
n_v	Visual noise term
$\boldsymbol{g_x}$	Generative process sensory mapping vector
$\boldsymbol{n_x}$	Generative process sensory noise vector
$\overline{K}$	Direct proportionality constant for the attractive force
μ_T	Target configuration
$\overline{m}_{arm}$	Generative model arm's mass estimation
$\overline{\phi}$	Generative model viscosity coefficient estimation
$\boldsymbol{f_\mu}$	Generative model internal state dynamics mapping
$\boldsymbol{\mu_s}$	Generative model sensory prediction vector
$\overline{L}_{arm}$	Generative model arm's length estimation
μ_p	Proprioceptive estimated sensory signal
μ_{vx}	Visual estimated sensory signal (on the x axis)
μ_{vy}	Visual estimated sensory signal (on the x axis)
$\boldsymbol{g_\mu}$	Generative model sensory mapping vector
C	Constant term in VFE formulation
$\boldsymbol{t}$	Simulation time vector
t_1	Simulation starting time
N	Number of considered time instants
t_N	Simulation ending time
$\boldsymbol{x_0}$	Initial hidden state configuration
$\boldsymbol{\mu_0}$	Initial internal state configuration
$\boldsymbol{k_\mu}$	Internal state learning rate vector
$\boldsymbol{k_A}$	Action learning rate vector
x_{RH}	Rubber Hand configuration

Continued on next page...

Continued from previous page

Symbol	**Description**
v_{gain}	Velocity gain
c_i	Generic constant
$\tilde{\boldsymbol{x}}_{\boldsymbol{VH}}$	Virtual hand state in generalized coordinates
$\dot{x}_{VH}$	Virtual hand dynamics
$\mathcal{H}(\cdot)$	Entropy
$\mathcal{H}(\cdot\,,\cdot)$	Cross-Entropy

Ringraziamenti

Desidero esprimere la mia profonda gratitudine alla Professoressa Marta Gandolla e a Pietro Donelli, per avermi guidato con pazienza e competenza durante la stesura di questo lavoro; nonché per essere venuti incontro alle mie esigenze per la progettazione iniziale del progetto.

Un particolare ringraziamento alla mia famiglia, senza la quale non avrei raggiunto questo traguardo. Grazie ai miei genitori, che mi hanno supportato e sopportato in questi anni tra alti e bassi. A mio fratello, per essere stato la sfida silenziosa e il riferimento costante che mi ha spinto a guardare sempre avanti. A tutti i miei parenti che mi hanno supportato anche solo psicologicamente nel raggiungimento di questo obiettivo: senza di voi non avrei mai portato avanti e concluso gli studi.

Menzione d'onore alla mia ragazza Virginia, che si è sforzata di seguire e comprendere il mio percorso accademico quasi quanto si è impegnata a ostacolarlo, riuscendo a trascinarmi lontano dai libri proprio nei momenti meno opportuni. La ringrazio soprattutto perché, con ammirevole coerenza, sta cercando di distrarmi anche in questo preciso istante, proprio mentre scrivo le ultime righe di questa tesi. Seriamente, la ringrazio per essermi stata accanto come nessun altro avrebbe saputo fare e per avermi aspettato - con una pazienza tutta sua - mentre concludevo gli studi.

Ai miei amici, per essere stati la mia valvola di sfogo indispensabile: grazie per avermi regalato quei momenti di svago e spensieratezza necessari a staccare la spina quando ne avevo più bisogno. La mia salute mentale ringrazia.

Infine, un ringraziamento del tutto narcisistico e orgoglioso al sottoscritto: per la disciplina che mi ha portato a non mollare anche quando l'idea di chiudere tutto sembrava la più saggia.