Active inference through whiskers

Francesco Mannella^a, Federico Maggiore^a, Manuel Baltieri^b, Giovanni Pezzulo^{a,*}

^a*Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy*

^b*Laboratory for Neural Computation and Adaptation, RIKEN Centre for Brain Science, Wako-shi, Japan*

Abstract

Rodents use whisking to probe actively their environment and to locate objects in space, hence providing a paradigmatic biological example of active sensing. Numerous studies show that the control of whisking has anticipatory aspects. For example, rodents target their whisker protraction to the distance at which they expect objects, rather than just reacting fast to contacts with unexpected objects. Here we characterize the anticipatory control of whisking in rodents as an active inference process. In this perspective, the rodent is endowed with a prior belief that it will touch something at the end of the whisker protraction, and it continuously modulates its whisking amplitude to minimize (proprioceptive and somatosensory) prediction errors arising from an unexpected whisker-object contact, or from a lack of an expected contact. We will use the model to qualitatively reproduce key empirical findings about the ways rodents modulate their whisker amplitude during exploration and the scanning of (expected or unexpected) objects. Furthermore, we will discuss how the components of active inference model can in principle map to the neurobiological circuits of rodent whisking.

Keywords: Active inference; active sensing; whiskers

1. Introduction

Rodents use whisking to probe actively their environment and to locate objects in space – hence providing a paradigmatic biological example of *active sensing* [1, 2, 3]. Active whisking has been studied in many contexts, such as during the perception of object location [4, 5, 6, 7] and the discrimination of shapes [8].

Studies of whisking behavior in freely behaving animals reveal two main whisking modalities. Away from an object, rodents perform “exploratory” behavior characterized by extended whisking protractions, allowing the animal to scout larger areas. On the other hand, when they are in contact with an object, they adopt a “scanning” strategy, reducing the amplitude of whisking protractions to match their distance to objects.

Matching whisking protraction to object distance is an effective strategy to ensure that contacts happen at the end of the protraction, with a light touch and minimal bending – a strategy called “minimal impingement”. After a contact with an unexpected object, whiskers cease to protract very rapidly (e.g., within 15ms). This rapid cessation of protraction (RCP) suggests a rapid feedback-control mechanism [9, 10, 11].

However, feedback-control alone cannot account for the fact that when objects are removed, the target protraction remains stable for at least one whisking cycle, increasing again only afterwards [7]. This finding suggests that rodents target their whisker protractions to where they *expect* objects to be, rather than just react to unexpected contacts. In other words, the modulation of whiskers’ amplitude appears to be an anticipatory strategy that depends on sensory prediction errors – or the mismatch between expected and sensed inputs – rather than just on current sensory inputs.

Importantly, the anticipatory modulation of whiskers allows the animal to actively perceive – and maintain an expectation about – its distance from an object. Here, perception is “active” in the sense that the alignment

*Corresponding author

Email address: giovanni.pezzulo@istc.cnr.it (Giovanni Pezzulo)

between the animal’s expectations about its distance from an object and the effective distance is achieved by acting (i.e., by changing the amplitude of whiskers’ oscillations), not by merely updating internal beliefs, as normally assumed in inferential theories of perception [12, 13, 14, 15]. Unlike classical theories of perception, which assume that the brain recognizes objects via representations of (hierarchies of) action-independent object features [16, 17], this form of active perception relies on a *generative* model that describes how (touch) sensations change when an object is in contact, i.e., models of sensorimotor contingencies [18] (see [19, 20] for examples of such generative models).

In this paper, we formally characterize rodents’ anticipatory and error-correction whisking strategies in terms of active inference: a framework developed in computational neuroscience to explain animal behavior and neural activity as resulting from the minimization of variational free energy, or under simplifying assumptions, prediction errors [21]. We develop an active inference model of whisking dynamics and use it to simulate an object localization task that exposes anticipatory aspects of whisking dynamics [7]. Furthermore, we discuss how our simulated agent performs active perception, displaying a dynamical alignment of internal expectations and animal-object distance without explicitly encoding such distance. Finally, we illustrate the putative neurobiological substrate of the active inference model of rodent whisking.

1.1. Summary of the active inference perspective on active whisking

The core idea of the active inference approach proposed here is to specify a generative model of sensorimotor contingencies: contingencies between whiskers’ protractions and the resulting proprioceptive and somatosensory (touch) sensations.

The agent’s model continuously generates proprioceptive and touch-related predictions during the action-perception loop. Crucially, the model includes a prior belief that objects will only be touched at the end of a whisker protraction. Hence, the agent will continuously adjust (increase or decrease) whisker protractions to match the prior prediction – or in other words, to minimize prediction errors resulting from either touching unexpected objects, or failing to touch expected objects that are instead missing.

The dynamical adjustment of whiskers’ protractions automatically produces a transition from exploration to scanning with a decrease of whisking amplitude, and from scanning to exploration with an increase of whisking amplitude. Importantly, both the decrease and the increase of whiskers’ amplitude result from the same (error minimization) mechanism, in two different conditions.

During the exploratory phase (characterized by larger oscillations of the whiskers), touching an object with the whiskers before the end of the protraction generates a *somatosensory (touch) prediction error*. This is because such input is unexpected given the prior (that an object will be touched only when the whiskers are fully extended). To minimize this prediction error, the animal can *reduce* the amplitude of the whisking protractions, so that it matches more closely the distance from the object. At the same time, the animal can also update its prior belief about the whisking amplitude, so to select (or predict) a smaller amplitude for the next step. This dynamical adjustment of whisking amplitude entails a shift from an “exploratory” to a “scanning” phase.

During the object scanning phase, when somatosensory or “exploratory” prediction errors have been minimized and hence the amplitude of whisker protractions closely matches animal-object distance, the sudden removal of the object causes a new sensory prediction error to arise – because the animal expected a touch sensation in correspondence of the missing object. To minimize this new prediction error, the animal will thus *increase* the amplitude of the whisking protraction, until another object is located, or the maximum span of the protractions is reached. In parallel, it will also update its prior beliefs about these reaching movements, leading to a larger protraction amplitude for the ensuing cycles. This dynamical adjustment of whisking amplitude characterizes a transition from a “scanning” to an “exploratory” phase.

As we will see in our results, in addition to producing switches from exploration to scanning and vice versa, the dynamic adjustment of whisker movements allows an animal to perceive its distance from an object. Indeed, in a regime of dynamical convergence (i.e., after some cycles of adjustments) the whiskers’ amplitude can be used as a proxy for animal-object distance – even if the generative model has no explicit notion of either objects or distance, but simply embodies sensorimotor contingencies, or relations, about how whisker protractions generate touch sensations. This idea is in accordance with theoretical views of

perception as a closed-loop convergence process, during which motor variables (e.g., whisking velocity or amplitude) are dynamically controlled over time, until they converge to steady state [22, 6, 23].

In the next section, we formally introduce the active inference model and illustrate simulations of whisking behavior during animal-object interactions [7]. Our simulations will show that minimizing somatosensory prediction errors guides both active exploration and the scanning of objects by whiskers.

2. Methods

2.1. Active inference

In theories of agency and perception inspired by Bayesian principles, the goal of an inferential agent is to generate a purposeful engagement with the world via the estimation of the probability of hidden variables $\mathbf{u} = \{\mathbf{x}, \mathbf{v}\}^1$ given by hidden states $\mathbf{x} = \{x_i\}_{i=1}^D$ and hidden causes $\mathbf{v} = \{v_i\}_{i=1}^F$ for some sensory inputs $\mathbf{s} = \{s_i\}_{i=1}^S$

$$P(\mathbf{u}|\mathbf{s}) = \frac{P(\mathbf{u}, \mathbf{s})}{P(\mathbf{s})} = \frac{P(\mathbf{s}|\mathbf{u})P(\mathbf{u})}{P(\mathbf{s})}. \quad (1)$$

A common issue with exact Bayesian schemes is that the marginal likelihood or model evidence $P(\mathbf{s})$ is often analytically intractable or computationally difficult to calculate. Moreover, the posterior $P(\mathbf{u}|\mathbf{s})$ may not follow a standard distribution and thus have no tractable summary statistics.

Active inference [25, 21, 26, 27, 28] brings forward a biologically plausible (variational) approximation to this problem, where an auxiliary distribution $Q(\mathbf{u})$ called *recognition density* has to be optimized to become a good approximation of the posterior. To implement this, the Kullback-Leibler divergence (a measure of the dissimilarity between probability distributions with roots in information theory) is minimized:

$$D_{KL}(Q(\mathbf{u}) || P(\mathbf{u}|\mathbf{s})) = \int Q(\mathbf{u}) \ln \frac{Q(\mathbf{u})}{P(\mathbf{u}|\mathbf{s})} d\mathbf{u} = \mathcal{F} + \ln P(\mathbf{s}) \quad (2)$$

where $\mathcal{F} \equiv -\langle \ln P(\mathbf{s}, \mathbf{u}) \rangle_Q + \langle \ln Q(\mathbf{u}) \rangle_Q$ is the *variational free energy* (VFE). This quantity depends on the recognition density and the agent's knowledge about the environment recapitulated by the joint density $P(\mathbf{s}, \mathbf{u}) = P(\mathbf{s}|\mathbf{u})P(\mathbf{u})$, while $\ln P(\mathbf{s})$ does not depend on the recognition density $Q(\mathbf{u})$. Hence, minimizing $\mathcal{F}$ with respect to $Q(\mathbf{u})$ will minimize the whole D_{KL} .

Optimizing $\mathcal{F}$ for arbitrary $Q(\mathbf{u})$ can however be a complex task. A common choice is thus to introduce a Laplace approximation [29] assuming that the joint density is a smooth function of $\mathbf{u}$, and that its logarithm can be approximated with a quadratic function near the mode (i.e. the logarithm of a Gaussian distribution). This is equivalent to a variational Gaussian approximation of the recognition density as a multivariate Gaussian form over the D-dimensional space $\mathbf{u}$ centered around the estimated mode $\boldsymbol{\mu}$ [30]. The Laplace approximation is therefore here used to evaluate, up to quantities that act as constants during the minimization process, $\mathcal{F}$ as

$$\mathcal{F} \approx -\ln P(\mathbf{s}, \mathbf{u})|_{\mathbf{u}=\boldsymbol{\mu}} \quad (3)$$

where $\boldsymbol{\mu}$ has the role of the agent's expectations about the hidden states $\mathbf{x}$ and hidden causes $\mathbf{v}$ (see [31] for a pedagogical treatment).

To evaluate the VFE, we then specify the joint density $P(\mathbf{s}, \mathbf{u})$ expressing hidden states and sensory inputs. In particular, we will assume that

$$\begin{aligned} P(\mathbf{s}, \mathbf{u}) &= P(\mathbf{s}, \mathbf{x}, \mathbf{v}) \\ &= P(\mathbf{s}|\mathbf{x}, \mathbf{v})P(\mathbf{x}, \mathbf{v}) \end{aligned} \quad (4)$$

thus specifying dependencies typical of state space model in continuous time, with the likelihood $P(\mathbf{s}|\mathbf{x}, \mathbf{v})$ corresponding to observation law and the prior $P(\mathbf{x}, \mathbf{v})$ encoding dynamics following the treatment found,

¹For a treatment including other variables, e.g., parameters evolving on different time scales, see for instance [24].

for instance, in [24]. At the same time, we also assume that the probability of hidden causes $\mathbf{v}$, $P(\mathbf{v})$, is Gaussian with precision approaching zero (i.e., whose covariance tends to infinity), meaning that it will not affect the remainder of the prior $P(\mathbf{x}|\mathbf{v})$. In the case of a system with one hidden state, one hidden cause and one sensory input this is equivalent to

$$s = g(x) + z_s \quad x' = f(x, v) + z_x \quad (5)$$

where $g(\cdot)$ maps hidden states to observations, $f(\cdot)$ encodes the dynamics of the hidden state and z_s, z_x represent (white) noise terms. Under zero-mean Gaussian assumptions about random variables z_s and z_x , the VFE is then

$$\begin{aligned} \mathcal{F} &\approx \frac{1}{2} \left[\frac{(s - g(x))^2}{\Sigma_s} + \frac{(x' - f(x, v))^2}{\Sigma_x} - \ln(\Sigma_s \Sigma_x) \right] \Big|_{x=\mu_x, v=\mu_v} \\ &= \frac{1}{2} \left[\frac{(s - g(\mu_x))^2}{\Sigma_s} + \frac{(\mu'_x - f(\mu_x, \mu_v))^2}{\Sigma_{\mu_x}} - \ln(\Sigma_s \Sigma_x) \right] \end{aligned} \quad (6)$$

where Σ_s and Σ_x are the variances of the random variables z_s and z_x , respectively.

At this point, the optimization of the VFE can be achieved either by changing the expected states μ, μ' through a (modified) gradient descent (see [24] for details)

$$\begin{aligned} \dot{\mu}_x &= \mu'_x - \frac{\partial \mathcal{F}}{\partial \mu_x} \\ \dot{\mu}_v &= -\frac{\partial \mathcal{F}}{\partial \mu_v}, \end{aligned} \quad (7)$$

or by selecting an action through a free energy gradient descent where

$$\dot{\alpha} = -\frac{\partial \mathcal{F}}{\partial \alpha} = -\frac{\partial \mathcal{F}}{\partial s} \frac{\partial s}{\partial \alpha} \quad (8)$$

The gradient with respect to actions highlights a key assumption of active inference treatments, whereby the action variable a is not itself part of the generative model (to be specified in the next section) but instead appears only in the recognition dynamics (i.e., the minimisation of variational free energy for action and perception embodied by an agent) via the relation $\frac{\partial s}{\partial \alpha}$, capturing how actions a affect sensory input s given an agent's specific "implementation", e.g., bodily constraints. In the active inference literature this is usually associated with the idea of reflex arcs, proposed to be the central construct for motor execution (here thought of essentially as a chain of reflexes) via the minimization of proprioceptive prediction errors induced by top-down modulatory signals from motor areas [32]. Such signals (i.e., predictions about proprioceptive sensory input) are thought to generate mismatches between *expected* proprioceptive signals and effective proprioceptive sensations, based on an equilibrium-trajectory (referent) control [33, 34] framework encoding desired movements in the (low-dimensional) space of observables, as opposed to classical treatments of motor control relying on inverse models that specify motor commands as (high-dimensional) signals in an intrinsic (bodily) frame of reference describing movements explicitly (e.g., motor signals describing stretching and compressing of muscle fibers).

2.2. Active whisking model

Whisker control in rodents is simulated in a rat-like agent that controls a single whisker, attached near its nose, in a 2D kinematic environment (see Figure 1A). The agent moves forward with a constant speed until it reaches a specified final position (e.g., the end of a platform). At that point, an object can appear within reach of the whisker, and later disappear. Our simulations (discussed in Section 3) will show how the agent controls its whiskers' protraction when objects appear or disappear.

Importantly, in active inference, there is a fundamental distinction between the *generative process*, corresponding to the actual dynamics that govern the whisker's movements in the environment (Figure 1) and the

agent's (whisker) controller, implicitly specified as a *generative model* of the whisker's dynamics (Figure 2) giving rise to the agent's recognition dynamics (see previous section). The recognition dynamics and generative process are thus engaged in a closed action-perception loop: the recognition dynamics instantiate actions for the generative process and the generative process provides sensory evidence to the recognition dynamics. To set up such action-perception loop, we will describe next the generative process and the generative model used to derive the recognition dynamics of the agent, respectively.

Figure 1: The simulated environment, or generative process. (A) The 2D environment: the simulation includes a rat-like agent with a single whisker. (B) Generative process and whisker dynamics. The generative process is specified by the dynamics of the state variable $\mathbf{x}^{\text{GP}}$. The first two components of $\mathbf{x}^{\text{GP}}$ (x_o^{GP} and x_w^{GP}) represent a central pattern generator described by an harmonic oscillator, while the last component x_w^{GP} describes the motor command for the angular position of the whisker's joint. The parameter ω determines the frequency of the central pattern generator oscillation and is kept constant during all simulations. The value of α specifies a scaling factor for the oscillation in the x_w^{GP} component. The generated sensory states are defined by s_p^{GP} and s_s^{GP} , with k representing the current upper limit of x_w^{GP} (the position of the external object in terms of whisker base angle), and $\beta = 500$ being the smoothness of the pressure response of the whisker.

Generative process. The whisker's dynamics are simulated within the generative process (Figure 1B). The whiskers' oscillatory movements are given by a dynamical system which comprises of three variables. The first two (x_o^{GP} and x_w^{GP}) define a harmonic oscillator of unitary amplitude as a first order system, which (for simplicity) will remain fixed in our simulations. The third one (x_w^{GP}) defines the actual movement of the joint angle of the whisker. This auxiliary variable rescales (with a leak term) the amplitude of the oscillator response in x_o^{GP} through a parameter α . Importantly, in this model, the agent's actions correspond to updates of α , the parameter that regulates the amplitude of the oscillation of x_w^{GP} . In other words, in this model, the agent can only change the way it samples its environment, by modulating the amplitude of the protraction of its whiskers.

The term x_w^{GP} is not only controlled by the agent, but also constrained by an external factor (k) that corresponds to the current limit of the whisker angle, due for example to the presence of an object that limits its oscillation. When the whisker angle (provided by x_w^{GP}) approaches the limit k , a somatosensory touch event is produced, which corresponds to a peak in the sensory state s_s^{GP} of the agent's generative model. Proprioceptive sensory states are simulated as linear and noisy mappings of $\dot{x}_w^{\text{GP}}$.

Generative model. The generative model is schematically illustrated in Figure 2. For simplicity, this model only describes only the dynamics of a single whisker; the forward movement of the simulated agent is hardwired in our simulations and hence not included in the generative model. Here, the variable x_o^{GP} encodes the central pattern generator of the generative process. It is assumed to be known (analogously to standard problems of filtering with known inputs [35], and in line with dynamical accounts of sensorimotor

Figure 2: The generative model. (A) Schematic illustration of the model as a Bayesian network. The model includes x , a variable that filters the amplitude of the whisker oscillation given by the central pattern generator x_o^{GP} . The variable x depends also on an input (v), encoding a prior over the desired amplitude. x and x' are then mapped to predictions about somatosensory (s_s) and proprioceptive (s_p) input. (B) Generative model in analytical form. The somatosensory mapping function g_{s_s} between μ_x and s_s has $\theta = 50$ that sets the smoothness of the pressure sensory prediction. The yellow arrow from x_o^{GP} describes a known input to the internal variable x' , under the assumption of an additional observation with an identity mapping to the pattern generator x_o^{GP} , omitted from the figure for simplicity (i.e., identity mappings imply prediction errors with infinite precision, or equivalently, zero variance, and thus variational updates instantaneously reaching steady-state).

interactions with no fixed set-point control, cf. the role of inputs φ_0 in the centrifugal governor found in [36]) and thus equivalent to a sharply peaked Gaussian distribution with (in the limit) zero variance.

The variable x represents the (stochastic) dynamics of the whisker and is described in Langevin form, with a deterministic term rescaling (with a leak) the amplitude of the oscillations of x_o^{GP} by a factor v . Here we also implicitly assumed the presence of an additional observation variable with an identity mapping (i.e., infinite precision) to the pattern generator x_o^{GP} . Due to the mapping appearing as a simple identity, we however omitted this from the generative model and thus from the associated free energy functional.

On the other hand, the hidden cause v models an exogenous input (from a dynamical systems perspective) in the dynamics of x , scaling the amplitude of whisker oscillations while influenced by bottom-up prediction errors, see below. Crucially, $P(v)$ is modelled as a Gaussian distribution with arbitrarily high variance (zero precision), meaning that its role as prior on the desired amplitude is entirely driven by external stimuli, i.e., v , or rather its mode μ_v , will act a proxy that registers the presence of unknown objects in the environment.

The variable $\mathbf{s}$ includes two sensory predictions: a proprioceptive prediction (s_p), proportional to the current angle velocity of the whisker, and a somatosensory (touch) prediction (s_s), corresponding to the actual pressure over the whisker during the contact with an external object. Both predictions are assumed to be generated from Gaussian distributions, with means given by the two mapping functions g_{s_s} and g_{s_p} that link the dynamics of the variable x . The prediction s_s is modelled as a radial function centered on the local maximum of the hidden states encoding the whisker's angle (i.e. the component x of the internal state variable and its velocity x'). It is then interesting to note that while s_s and s_s^{GP} refer to the same physical quantity, there is a fundamental difference between the radial function used to model s_s and the equation used in the generative process for s_s^{GP} (see Figure 1), see also [19, 20].

Variational updates, or recognition dynamics. As explained in Section 2.1, once the joint density of the generative model (Equation 4) is specified, it is possible to compute the variational free energy using the current means (modes) of the model's hidden variables μ_x, μ'_x, μ_v and the current sensory states (see equation

6). With the VFE specific to the generative model, we can then find the gradients of VFE, generating updates for perception and action of our simulated agent (see equations 7 and 8). For perception, the model is updated by changing the posterior beliefs about both the expected hidden states μ_x and μ'_x and the expected causal variable v . For action, the model is updated by descending the free energy gradient $\frac{\partial \mathcal{F}}{\partial \alpha}$ with respect to the parameter α . This gradient can be factored into two components $\frac{\partial \mathcal{F}}{\partial s} \frac{\partial s}{\partial \alpha}$. The former ($\frac{\partial \mathcal{F}}{\partial s}$) is the gradient of the free energy with respect to the sensory states. The latter ($\frac{\partial s}{\partial \alpha}$) is the gradient of sensory states with respect to the action. It corresponds to a learned (or hardwired) mapping between action and sensory states at the level of the generative process and it is usually identified with the spinal neural circuits which manage reflex arcs [37]. This gradient $\frac{\partial s}{\partial \alpha}$ is implemented here as a composition of $\frac{\partial s_p}{\partial \alpha} = x_o^{\text{GP}}$ and $\frac{\partial s_s}{\partial \alpha} \approx 1$, where x_o^{GP} is the exact solution of the proprioceptive part of the gradient ($\frac{\partial s_p}{\partial \alpha}$), while the somatosensory gradient ($\frac{\partial s_s}{\partial \alpha}$) is approximated to 1, which is sufficient to inform the process about the direction of the update of α based on somatosensory states.

The update of v showcases how the agent adapts to external constraints, such as the presence of an object that blocks whisker movements, updating its beliefs about the causes of its sensory input and preventing whisker oscillations from reaching the desired, but unattainable, amplitude specified in the current prior expectation μ_v . In other words, when the agent cannot change the world (e.g., cannot use the whisker to push the object) via its own actions, it can nonetheless adapt its internal prior beliefs μ_v to account for the actual constraints of the generative process.

3. Results

We used the active inference model of whisking to simulate the experimental setup of Figure 1a, and the control of a single whisker in the presence or the absence of an object. The simulation can be conceptually divided into three phases. (1) Initially there is no object and the animal's whiskers operate in "exploration mode", a phase characterized by large whisker oscillations. (2) An object then appears and the animal progressively reduces whisking amplitude, shifting from "exploration" to the "scanning" of the object. (3) Finally, the object is removed and the amplitude progressively increases once again, marking a transition back from "scanning" to "exploration".

Figure 3 shows the results of the simulations. At first, the recognition dynamics and the generative process are in the same initial conditions. Figure 3A shows the dynamics of the variable x_w^{GP} that controls the amplitude of the whisker oscillation in the generative process, as a function of time. Without objects, the animal is in "exploration" mode, characterized by large whisker oscillations and extended protractions, which slightly increase over time. In the second phase, an object (black line) appears, limiting whiskers' protractions. Finally, in the third phase, the object disappears and whisker protractions increase once again. Crucially however, they initially remain significantly smaller than in the first phase. This is especially evident in the first cycle after the disappearance of the object, when they almost match (expected) animal-object distance, even if the object is not there to stop whiskers' protractions. This is the key empirical observation that has lead to the proposal that whisker movements are guided by anticipatory mechanisms that consider expected animal-object distance, rather than feedback-control mechanisms [7].

In the first phase, when there is no object, the variables x_w^{GP} and x that control the amplitude of whisker oscillation in the generative process (Figure 3A) and model (Figure 3B) slightly increase over time. This is because the agent expects to make contact with something at the end of the whisker protraction, but does not receive any touch input (Figure 3C) – and hence generates a negative somatosensory prediction error (Figure 3D). The action increment that cancels out this prediction error (Figure 3F) increases the value of α (Figure 3E), which in turn amplifies the whisker oscillation / protraction (x_w^{GP} in Figure 3A).

In the second phase, when the object (black line) appears, the agent's expected amplitude is significantly greater than the actual amplitude of the generative process; this is because x_w^{GP} is physically blocked by the presence of an unexpected object, but μ_x is not (note also that there is no explicit representation of the object in the generative model). However, the variable μ_x rapidly re-aligns to the true dynamics of the generative process. This re-alignment process is guided by an error-correction mechanism, as we can see with sensory signals and prediction errors shown in Figure 3C and Figure 3D, respectively. As shown in

the third panel, the agent receives somatosensory (touch) and proprioceptive sensations from the generative process, especially evident with the first prominent somatosensory input (first orange peak). This input is completely unexpected and hence generates a large positive somatosensory prediction error, which is shown in Figure 3D. The (negative) action increment that cancels out this prediction error (Figure 3F) suddenly decreases the value of the α parameter (Figure 3E), which in turn decreases the amplitude of the whisker oscillation / protraction (x_w^{GP} in Figure 3A). This decrease stops when the whisker touches the object gently and the α parameter finds a dynamic equilibrium around the correct value.

Finally, in the third phase, when the object (black line) disappears, the μ_x variable (Figure 3B) continues to guide whisker protraction at the expected distance from the object, at least in the first cycle. The sudden disappearance of the object implies that the agent does not receive its expected touch sensation (i.e., there is no orange peak in Figure 3C) and hence generates a negative somatosensory prediction error (Figure 3D), which is compensated by progressively increasing the amplitude of whisker oscillations (Figure 3E and Figure 3F). This marks thus a transition back to the exploratory mode characterising phase one, bringing the whisker oscillations back to baseline, i.e., extended protractions to explore the environment.

The simulation of the three phases closely resembles the empirical results reported in [7], where upon appearance of an object (in [7], a platform), whiskers' protractions match precisely the distance between the animal and the object – as seen in the second phase of our simulation. Crucially, when the object is suddenly removed, whiskers' protractions do not adapt immediately, but remain proportional to the last (expected) animal-object distance – which is exactly what happens in the third phase of our simulation. This suggests that the control of whiskers may be anticipatory at its core (based on expectations about animal-object distance) [7]. Furthermore, our simulations show that error-correction (via a free energy minimization) mechanism can explain both the shift from “exploration” to “scanning” when the object is introduced (second phase) and the reverse one from “scanning” to “exploration” when the object is removed (third phase). In our model, what guides the changes in whisker protraction is the somatosensory (touch) prediction error. The proprioceptive prediction error instead plays a stabilizing role on whisking oscillation, allowing our agent to compensate for small and unexpected changes caused by external disturbances of noise.

3.1. Control of multiple whiskers

To showcase further applications of our model, we finally consider an agent with multiple (here, two) whiskers. Given the structure of our framework, the formulation is rather straightforward, showing thus its potential for the development of more complex models of active perception. In this extended generative model (Figure 4A), we simply add a copy of the variables specified in the previous generative model, except for the oscillator variable x_o^{GP} , which is shared across all the whiskers. Given this duplication, the parameters of different whiskers can be tuned independently, hence we here simulate two whiskers with slightly different lengths and centers of oscillations (Figure 4). However, since the model relies on a shared oscillator variable x_o^{GP} , the resulting whisker dynamics remain in phase during scanning behaviour (Figure 4B) – as observed empirically [6].

4. Discussion

We proposed an active inference model of active sensing through whiskers. We started from the premise that rodents use their whiskers' dynamics to make (perceptual) inferences about the world; and described the active control of whiskers as an error-correction mechanism ensuring that a prior prediction (that something will be touched at the end of the protraction) can be realized in the environment.

Our simulations illustrate that the active inference model reproduces key anticipatory aspects of whisker movement control, such as the fact that mice target the amplitude of their whisking movements to the expected distance from an object; and (contrary to purely reactive models) that such expectations are robust and stable for some time even when objects disappear [7]. Furthermore, the model characterizes the alternation of exploration and scanning phases, exploration to scanning when an object is sensed and scanning to exploration when such object disappears, in terms of a unique prediction error minimization mechanism. Finally, this model explains how the control of whisker movements supports active perception

by dynamically aligning expectations (about whisker protraction) in the agent’s generative model to effective object distance. This is because the same error (or more precisely, free energy) minimization process has the dual role of selecting the motor actions to execute and updating the internal variables of the generative model.

The proposed model builds on a tradition of early cybernetic models and in particular perceptual control theory [38], which stresses the importance of continuously controlling motor variables (e.g. modulating the force impressed on the acceleration pedal while driving) to keep a preferred (prior) perceptual variable constant (e.g., to ensure that I always see 120Km/h on my speedometer). Furthermore, our model is also in line with the (more recent) theoretical view of active perception as a closed-loop convergence process, where motor variables (here, whisker amplitude) are continuously modulated for perceptual purposes [22, 23].

The active view of perception that emerges from the model differs in at least two ways from views of perception as the recognition of objective features of the environment [16] or the inference of external variables [12, 13, 14, 15]. First, traditional theories of perception assume that in order to recognize external variables (e.g., objects, faces or places), an agent has to be endowed with a model of these variables and their features. Rather, the proposed model stresses the perception of external variables (here, animal-object distance) without an explicit, action-independent model of objects and their features. The generative model only includes knowledge about touch sensations changing as an effect of whisking actions (i.e., sensorimotor contingencies [18, 39, 19, 40]), without explicit action-independent perceptual features or external variables like objects or distances, as commonly assumed in other theories of perception.

Second, traditional theories assume that perception is only achieved by changing beliefs internally. Here, instead, the perception of animal-object distance is achieved in an active manner: by dynamically adjusting a motor variable (whisker amplitude) based on a prior encoding the sensorimotor coupling between agent and environment, rather than a fixed set-point [36]. This ensures that touch sensations only happen at the end of the protraction – hence minimizing any discrepancy between prior predictions about touch and current sensations. An interesting “side-effect” of this adjustment is that it provides an implicit estimate of animal-object distance. Given this lawful correlation between whisker protraction (in the stationary regime) and animal-object distance, the animal can use the former as a proxy for the latter, so to, for example, approach or avoid the object. Our model therefore shows how to endow active inference models with minimalistic (sensorimotor or action-oriented) models or action-world interactions, in place of more complete models that describe and predict external dynamics independent of an agent’s actions [41, 42, 43, 44, 45, 46, 47] (see also e.g., [19, 48, 36, 20] for example implementations). An objective for future research is testing to what extent this approach scales up to robotic settings [49, 50].

4.1. The neurobiology of whisking from an active inference perspective

The neurobiology of active whisking involves a widely distributed brain network, which spans sensory, motor, premotor/prefrontal areas and the brain stem [1]. Figure 5 shows part of this network and highlights the putative neuronal underpinnings of the main variables of the generative model used in this article. We map sensory inputs and prediction errors in the two sensory modalities (somatosensory information about touch, such as the mechanical pressure over whiskers s_s and proprioceptive information s_p) to the barrel cortex (vibrissal somatosensory cortex) [51, 52, 53], the two variables of the central pattern generator (x_o^{GP} and $x_o'^{GP}$) to the brain stem, the hidden variable that modulates the central pattern generator μ_x to vibrissal motor cortex and the hidden cause μ_v to the premotor / prefrontal cortex.

The proposed neurobiological scheme follows a dynamical systems perspective on cortical computations and movement control [54, 55]. In this scheme, premotor/prefrontal areas define a prior over the amplitude of the oscillation of whiskers. Vibrissal motor cortex controls whisker movements (until they reach the desired set point) by modulating the amplitude and midpoint of whiskers protractions [56, 57]. The actual control of movement engages synergistically the vibrissal motor cortex and central pattern generators in the brain stem. Evidence indicates that vibrissal motor units indirectly control the activity of vibrissal motoneurons, through the modulation of sub-cortical central pattern generators; and in turn the motoneurons control the facial muscles responsible for whiskers’ movements [58].

Somatosensory prediction errors (discrepancies between somatosensory predictions and sensations) are computed by the barrel cortex and are related to whisking amplitude change rather than whisking frequency

or velocity [59, 7]. These prediction errors are sent to vibrissal motor cortex and premotor/prefrontal cortical areas, to keep the agent’s generative model in register with external constraints, such as the (unpredicted) presence of an object, which may prevent the whisker from reaching the “desired” amplitude specified in the current prior expectation (μ_v). Specifically, prediction errors modify the oscillation amplitude of whiskers at the level of motor cortex; and the prior about oscillation amplitude at the level of premotor/prefrontal cortical areas. Finally, the model includes a second type of (proprioceptive) prediction error, which measures the difference between the expected and currently sensed whisker protraction, and ultimately helps stabilizing whisker movements in face of small external disturbances, such as wind.

5. Conclusions

The model introduced in this article explains the active control of whisker movements as an active inference process which continuously minimizes prediction errors and the discrepancy between expected somatosensory sensations and current observations. This anticipatory, error correction mechanisms allows smooth transitions between phases of exploration and scanning of objects, and vice versa, empowering the animal with the ability to (actively) perceive their distance from objects.

Our model is in agreement with key empirical evidence supporting the anticipatory, expectation-based nature of whisking control [7]. Furthermore, it provides preliminary evidence for a new interpretation of neural signals in the brain circuits for whisking in terms of Bayesian (active) inference constructs, such as priors about amplitude of oscillation in premotor / prefrontal cortices and sensory prediction errors in barrel cortex. While some aspects of the model qualitatively align with the known neurophysiology of whisking (see Section 4.1) the specific mappings between neuronal circuits and active inference signals will be an object of future studies.

Acknowledgments

This research has received funding from the European Union’s Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3) and the European Research Council under the Grant Agreement No. 820213 (ThinkAhead). MB is a JSPS Postdoctoral Research Fellow supported by a KAKENHI Grant-in-Aid for Scientific Research (No. 19F19809).

Data Availability

The python code to reproduce the results of this paper can be found at:

<https://github.com/francesco-mannella/Active-Inference-In-Continuous-Time-With-Whiskers>

References

- [1] M. E. Diamond, M. Von Heimendahl, P. M. Knutsen, D. Kleinfeld, E. Ahissar, ‘where’ and ‘what’ in the whisker sensorimotor system, *Nature Reviews Neuroscience* 9 (8) (2008) 601–612.
- [2] B. Mitchinson, R. A. Grant, K. Arkley, V. Rankov, I. Perkon, T. J. Prescott, Active vibrissal sensing in rodents and marsupials, *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1581) (2011) 3037–3048.
- [3] N. E. Bush, S. A. Solla, M. J. Hartmann, Whisking mechanics and active sensing, *Current Opinion in Neurobiology* 40. doi:10.1016/j.conb.2016.08.001.
- [4] E. Ahissar, P. M. Knutsen, Object localization with whiskers, *Biological cybernetics* 98 (6) (2008) 449–458.

- [5] S. B. Mehta, D. Whitmer, R. Figueroa, B. A. Williams, D. Kleinfeld, Active spatial perception in the vibrissa scanning sensorimotor system, *PLoS Biol* 5 (2) (2007) e15.
- [6] I. Saraf-Sinik, E. Assa, E. Ahissar, Motion makes sense: an adaptive motor-sensory strategy underlies the perception of object location in rats, *Journal of Neuroscience* 35 (23) (2015) 8777–8789.
- 355 [7] J. Voigts, D. H. Herman, T. Celikel, Tactile object localization by anticipatory whisker motion, *Journal of neurophysiology* 113 (2) (2015) 620–632.
- [8] F. Anjum, H. Turni, P. G. Mulder, J. van der Burg, M. Brecht, Tactile guidance of prey capture in etruscan shrews, *Proceedings of the National Academy of Sciences* 103 (44) (2006) 16544–16549.
- 360 [9] D. Deutsch, M. Pietr, P. M. Knutsen, E. Ahissar, E. Schneidman, Fast feedback in active sensing: touch-induced changes to whisker-object interaction, *PloS one* 7 (9) (2012) e44272.
- [10] B. Mitchinson, C. J. Martin, R. A. Grant, T. J. Prescott, Feedback control in active sensing: rat exploratory whisking is modulated by environmental contact, *Proceedings of the Royal Society B: Biological Sciences* 274 (1613) (2007) 1035–1041.
- 365 [11] C. C. Rodgers, R. Nogueira, B. C. Pil, E. A. Greeman, J. M. Park, Y. K. Hong, S. Fusi, R. M. Bruno, Sensorimotor strategies and neuronal representations for shape discrimination, *Neuron* doi:10.1016/j.neuron.2021.05.019.
- [12] K. Doya, S. Ishii, A. Pouget, R. P. N. Rao (Eds.), *Bayesian Brain: Probabilistic Approaches to Neural Coding*, 1st Edition, The MIT Press, 2007.
- [13] H. Von Helmholtz, *Helmholtz’s treatise on physiological optics*, Vol. 3, Optical Society of America, 1925.
- 370 [14] T. S. Lee, D. Mumford, Hierarchical bayesian inference in the visual cortex., *J Opt Soc Am A Opt Image Sci Vis* 20 (7) (2003) 1434–1448.
- [15] R. P. Rao, D. H. Ballard, Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects, *Nat Neurosci* 2 (1) (1999) 79–87. doi:10.1038/4580.
URL <http://dx.doi.org/10.1038/4580>
- 375 [16] D. Marr, *A computational investigation into the human representation and processing of visual information*, WH San Francisco: Freeman and Company, 1982.
- [17] M. Riesenhuber, T. Poggio, Hierarchical models of object recognition in cortex, *Nature neuroscience* 2 (11) (1999) 1019–1025.
- 380 [18] J. K. O’Regan, A. Noë, A sensorimotor account of vision and visual consciousness, *Behavioral and brain sciences* 24 (5) (2001) 939.
- [19] M. Baltieri, C. L. Buckley, An active inference implementation of phototaxis, in: *Artificial Life Conference Proceedings* 14, MIT Press, 2017, pp. 36–43.
- [20] A. Tschantz, A. K. Seth, C. L. Buckley, Learning action-oriented models through active inference, *PLoS computational biology* 16 (4) (2020) e1007805.
- 385 [21] K. Friston, The free-energy principle: a unified brain theory?, *Nature reviews neuroscience* 11 (2) (2010) 127–138.
- [22] E. Ahissar, E. Assa, Perception as a closed-loop convergence process, *Elife* 5 (2016) e12830.
- [23] C. L. Buckley, T. Toyoizumi, A theory of how active behavior stabilises neural activity: Neural gain modulation by closed-loop environmental feedback, *PLoS computational biology* 14 (1) (2018) e1005926.

- [24] K. Friston, Hierarchical models in the brain, *PLOS Computational Biology* 4 (11) (2008) 1–24. doi: 10.1371/journal.pcbi.1000211.
- [25] Friston, Kilner, Harrison, A free energy principle for the brain, *Journal of physiology, Paris* 100 (2006) 70–87. doi:10.1016/j.jphysparis.2006.10.001.
- [26] K. Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, G. Pezzulo, Active inference: A process theory, *Neural Computation* 29 (1) (2016) 1–49, [IF 2015: 1.626]. doi:doi:10.1162/NECO_a_00912.
- [27] K. J. Friston, T. Parr, B. de Vries, The graphical brain: belief propagation and active inference, *Network Neuroscience* 1 (4) (2017) 381–414.
- [28] G. Pezzulo, F. Rigoli, K. Friston, Active inference, homeostatic regulation and adaptive behavioural control., *Progress in neurobiology* 134 (2015) 17–35, [IF 2015: 13.177]. doi:10.1016/j.pneurobio.2015.09.001.
URL <http://dx.doi.org/10.1016/j.pneurobio.2015.09.001>
- [29] Friston, Mattout, Trujillo-Barreto, Ashburner, Penny, Variational free energy and the laplace approximation, *NeuroImage* 34 (2007) 220–34. doi:10.1016/j.neuroimage.2006.08.035.
- [30] M. Opper, C. Archambeau, The variational gaussian approximation revisited, *Neural computation* 21 (2008) 786–92. doi:10.1162/neco.2008.08-07-592.
- [31] C. L. Buckley, C. S. Kim, S. McGregor, A. K. Seth, The free energy principle for action and perception: A mathematical review, *Journal of Mathematical Psychology* 81 (2017) 55 – 79. doi:<https://doi.org/10.1016/j.jmp.2017.09.004>.
- [32] R. A. Adams, S. Shipp, K. J. Friston, Predictions not commands: active inference in the motor system, *Brain Structure and Function* 218 (3) (2013) 611–643.
- [33] T. Buhrmann, E. A. Di Paolo, Spinal circuits can accommodate interaction torques during multijoint limb movements, *Frontiers in computational neuroscience* 8 (2014) 144.
- [34] A. G. Feldman, et al., Referent control of action and perception, *Challenging Conventional Theories in Behavioral Neuroscience*.
- [35] M. Baltieri, C. L. Buckley, On kalman-bucy filters, linear quadratic control and active inference, *arXiv preprint arXiv:2005.06269*.
- [36] M. Baltieri, C. L. Buckley, J. Bruineberg, Predictions in the eye of the beholder: an active inference account of watt governors, in: *Artificial Life Conference Proceedings*, MIT Press, 2020, pp. 121–129.
- [37] K. J. Friston, J. Daunizeau, J. Kilner, S. J. Kiebel, Action and behavior: a free-energy formulation., *Biol Cybern* 102 (3) (2010) 227–260. doi:10.1007/s00422-010-0364-z.
- [38] W. T. Powers, W. T. Powers, *Behavior: The control of perception*, Aldine Chicago, 1973.
- [39] A. K. Engel, A. Maye, M. Kurthen, P. König, Where’s the action? the pragmatic turn in cognitive science, *Trends in cognitive sciences* 17 (5) (2013) 202–209.
- [40] M. Baltieri, C. L. Buckley, Generative models as parsimonious descriptions of sensorimotor loops, *Behavioral and Brain Sciences* 42.
- [41] G. Pezzulo, F. Dommarumma, P. Iodice, D. Maisto, I. Stoianov, Model-based approaches to active perception and control, *Entropy* 19 (6) (2017) 266.
- [42] G. Pezzulo, P. Cisek, Navigating the affordance landscape: Feedback control as a process model of behavior and cognition., *Trends in Cognitive Sciences* 20 (6) (2016) 414–424, [IF 2015: 17.850]. doi: 10.1016/j.tics.2016.03.013.

- [43] A. Clark, Surfing uncertainty: Prediction, action, and the embodied mind, Oxford University Press, 2015.
- [44] A. K. Seth, A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synesthesia, *Cognitive neuroscience* 5 (2) (2014) 97–118.
- 435 [45] A. K. Seth, The cybernetic bayesian brain, Open MIND. Frankfurt am Main: MIND Group.
- [46] A. W. Corcoran, G. Pezzulo, J. Hohwy, From allostatic agents to counterfactual cognisers: active inference, biological regulation, and the origins of cognition, *Biology & Philosophy* 35 (3) (2020) 1–45.
- [47] M. Sims, G. Pezzulo, Modelling ourselves: what the free energy principle reveals about our implicit notions of representation, *Synthese*.
- 440 [48] M. Baltieri, C. L. Buckley, Pid control as a process of active inference with linear generative models, *Entropy* 21 (3) (2019) 257.
- [49] M. J. Pearson, B. Mitchinson, J. C. Sullivan, A. G. Pipe, T. J. Prescott, Biomimetic vibrissal sensing for robots, *Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1581) (2011) 3085–3096.
- 445 [50] J. C. Sullivan, B. Mitchinson, M. J. Pearson, M. Evans, N. F. Lepora, C. W. Fox, C. Melhuish, T. J. Prescott, Tactile discrimination using active whisker sensors, *IEEE Sensors Journal* 12 (2) (2011) 350–362.
- [51] N. J. Sofroniew, Y. A. Vlasov, S. A. Hires, J. Freeman, K. Svoboda, Neural coding in barrel cortex during whisker-guided locomotion, *eLife* 4 (2015) e12559. doi:10.7554/eLife.12559.
- 450 [52] A. Ayaz, A. Stäuble, M. Hamada, M.-A. Wulf, A. B. Saleem, F. Helmchen, Layer-specific integration of locomotion and sensory information in mouse barrel cortex, *Nature Communications* 10 (2019) 2585. doi:10.1038/s41467-019-10564-8.
- [53] C. C. H. Petersen, Sensorimotor processing in the rodent barrel cortex, *Nature Reviews Neuroscience* 20 (2019) 533–546. doi:10.1038/s41583-019-0200-y.
- 455 [54] F. Mannella, G. Baldassarre, Selection of cortical dynamics for motor behaviour by the basal ganglia, *Biological cybernetics* 109 (6) (2015) 575–595.
- [55] K. V. Shenoy, M. Sahani, M. M. Churchland, Cortical control of arm movements: a dynamical systems perspective, *Annual review of neuroscience* 36 (2013) 337–359.
- 460 [56] D. N. Hill, J. C. Curtis, J. D. Moore, D. Kleinfeld, Primary motor cortex reports efferent control of vibrissa motion on multiple timescales, *Neuron* 72 (2011) 344–356. doi:10.1016/j.neuron.2011.09.020.
- [57] W. A. Friedman, H. P. Zeigler, A. Keller, Vibrissae motor cortex unit activity during whisking, *Journal of Neurophysiology* 107 (2012) 551–563, pMID: 21994257. doi:10.1152/jn.01132.2010.
- 465 [58] N. P. Cramer, A. Keller, Cortical control of a whisking central pattern generator, *Journal of Neurophysiology* 96 (2006) 209–217, pMID: 16641387. doi:10.1152/jn.00071.2006.
- [59] R. N. S. Sachdev, R. W. Berg, G. Champney, D. Kleinfeld, F. F. Ebner, Unilateral vibrissa contact: changes in amplitude but not timing of rhythmic whisking, *Somatosensory & Motor Research* 20 (2003) 163–169. doi:10.1080/08990220311000405208.

Figure 3: **Simulation results.** (A) Dynamics of the variable x_w^{GP} , which corresponds to the whisker's base angle in the generative process. (B) Dynamics of the variable μ_x , which corresponds to the predicted whisker's base angle in the generative model. By comparing A and B, it is possible to notice that the appearance of the object in phase 2 affects immediately the dynamics of the whisker's base angle in the generative process (x_w^{GP}); however, the dynamics of the predicted whisker's base angle μ_x take slightly longer to adapt. (C) Agent's sensory inputs (somatosensory (touch) in orange and proprioceptive in blue). (D) Agent's prediction errors (somatosensory (touch) in orange and proprioceptive in blue). (E-F) Dynamics of α (action) and its derivative $\dot{\alpha}$. Negative peaks of touch prediction error correspond to situations in which the agent would have expected to touch the object but did not – and hence increases the amplitude of the whisker protraction (by increasing α). Rather, negative peaks of touch prediction error correspond to situations in which the agent feels unexpected touch sensations – and hence decreases the amplitude of the whisker protraction (by decreasing α).

Figure 4: Simulation of the control of two whiskers having slightly different dimensions and central points of oscillations. (A) Extended generative model for the control of two whiskers. The model's variables are the same as Figure 2, but there is a separate set of variables for each whisker, except for the oscillator variable x_o^{GP} , which is shared across both whiskers. (B) Simulation results. The notation is the same as in Figure 3. Here, the key thing to notice is that the two whiskers (blue and orange) have oscillate in phase. See the main text for an explanation.

Figure 5: Putative neuronal underpinnings of the generative model for whisker control discussed in this article. This simplified neuronal scheme comprises premotor / prefrontal areas (PMPF), vibrissal motor cortex (VMC), central pattern generators (CPGs) in the brain stem, and Barrel Cortex / vibrissal somatosensory cortex. Colored circles represent variables of the generative model, whereas the ξ symbols represent prediction errors. Colored edges show the (neuronal) message passing between variables of the model. The red edge originating from the v (causal) variable represents top-down signals. Blue edges originating from the ξ symbols represent prediction errors that are propagated bottom-up. See the main text for explanation.