

Article

PID control as a process of active inference with linear generative models

Manuel Baltieri ^{†,‡}  and Christopher L. Buckley [‡] 

EASy group - Sussex Neuroscience, Department of Informatics, University of Sussex, Brighton, UK;
m.baltieri@sussex.ac.uk, c.l.buckley@sussex.ac.uk

* Correspondence: m.baltieri@sussex.ac.uk

† This paper is an extended version of our paper published in From Animals to Animats 15: 15th International Conference on Simulation of Adaptive Behavior, SAB 2018.

Version February 27, 2019 submitted to Preprints

Abstract: In the past few decades, probabilistic interpretations of brain functions have become widespread in cognitive science and neuroscience. In particular, the free energy principle and active inference are increasingly popular theories of cognitive functions that claim to offer a unified understanding of life and cognition within a general mathematical framework derived from information and control theory, and statistical mechanics. However, we argue that if the active inference proposal is to be taken as a general process theory for biological systems, it is necessary to understand how it relates to existing control theoretical approaches routinely used to study and explain biological systems. For example, recently, PID control has been shown to be implemented in simple molecular systems and is becoming a popular mechanistic explanation of behaviours such as chemotaxis in bacteria and amoebae, and robust adaptation in biochemical networks. In this work, we will show how PID controllers can fit a more general theory of life and cognition under the principle of (variational) free energy minimisation when using approximate linear generative models of the world. This more general interpretation provides also a new perspective on traditional problems of PID controllers such as parameter tuning as well as the need to balance performances and robustness conditions of a controller. Specifically, we then show how these problems can be understood in terms of the optimisation of the precisions (inverse variances) modulating different prediction errors in the free energy functional.

Keywords: approximate Bayesian inference, active inference, PID control, generalised state-space models, sensorimotor loops, information theory, control theory

1. Introduction

Probabilistic approaches to the study of living systems and cognition are becoming increasingly popular in the natural sciences. In particular for the brain sciences, the Bayesian brain hypothesis, predictive coding, the free energy principle and active inference have been proposed to explain brain processes including perception, action and higher order cognitive functions [1–8]. According to these theories, brains, and biological systems more generally, should be thought of as Bayesian inference machines, since they appear to estimate the latent states of their sensory input in a process consistent with a Bayesian inference scheme. Given the complexity of exact Bayesian inference, however, approximated schemes are believed to provide a more concrete hypothesis on the underlying mechanisms. One candidate scheme is the free energy principle (FEP), which was introduced in [4] and later elaborated in a series of papers, e.g. [9–11], and has its roots in information theory, control theory, thermodynamics and statistical mechanics. While initially the theory emerged in the computational [12] and behavioural/cognitive neurosciences [13,14], over time, further connections with the fields of biological self-organisation, information theory, optimal control, cybernetics and economics among others, have also been suggested [10,15–17]. According to the FEP, living systems exist in a limited set of physical states and thus must minimise the entropy of those physical states (see fluctuation

theorems for non-equilibrium thermodynamics, e.g. [18]). To achieve this, organisms can minimise the informational entropy of their sensory states, which, under ergodic assumptions, is equivalent to the time average of *surprisal* (or self-information) [9]. Surprisal quantifies how improbable an outcome is for a system, i.e. a fish out of water is in a surprising state. Biological creatures can thus be seen as minimising the surprisal of their sensations to maintain their existence, e.g. a fish' observations should be limited to states in water. Since this surprisal itself is not directly accessible by an agent, variational free energy is proposed as an upper bound on this quantity which can be minimised in its stead [4,19]. It has also been suggested that cognitive functions such as perception, action, learning and attention can be accounted for in terms approximate Bayesian inference schemes such as the FEP. In particular, according to this hypothesis, perception can be described using predictive coding models of the cortex. These models describe perception as a combination of feedforward prediction errors and feedback predictions combined under a generative model to infer the hidden causes and states of sensory data [2]. More recent work has connected these ideas to control theory and cybernetics [15,17,20], extending existing accounts of (optimal) motor control and behaviour [10,13,21,22]. On this view, behaviour is cast as a process of acting on the world to make sensory data better fit existing predictions, with (optimal) motor control cast as a Bayesian inference problem. The most recent attempt to unify predictive coding and optimal control theory, usually falls under the name of *active inference* [10,13].

While in standard accounts of perceptual inference prediction errors can be suppressed only by updating predictions of the incoming sensations, in active inference errors can also be minimised by directly acting on the environment to change sensory input to better accord with existing predictions [9,13]. If a generative model encodes information about favourable states for an agent, then this process constitutes a way by which the agent can change its environment to better meet its needs. Thus, under the FEP, these two processes of error suppression allow a system to both infer the properties of, and control, the surrounding environment. Most models implementing the FEP and active inference assume that agents have a deep understanding of their environment and its dynamics in the form of an accurate and detailed generative model. For instance, in [13,23] the generative model of the agent explicitly mirrors the *generative process* of the environment, i.e. the dynamics of the world the agent interacts with. In recent work, it has been argued that this need not be the case [24–27], especially if we consider simple living systems with limited resources. We intuitively don't expect an ant to model the entire environment where it forages, performing complex simulations of the world in its brain (cf. the concept of *Umwelt* [28]). When states and parameters in the world change too rapidly, accurate online inference and learning are implausible [29]. This idea is however common in the literature, e.g. [6,13,14,23], where cognition and perception are presented as processes of inference to the best explanation, and agents are primarily thought to build sophisticated models of their worlds with only a secondary role for action and behaviour. A possible alternative introduces action-oriented models entailing a more parsimonious approach where only task-relevant information is encoded [24,25]. On this normative view, agents only model a minimal set of environmental properties, perhaps in the form of sensorimotor contingencies [26], that are necessary to achieve their goals.

The relationship between information/probability theory and control theory has long been recognised, with the first intuitions emerging from work by Shannon [30] and Kalman [31]. A unifying view of these two theoretical frameworks is nowadays proposed for instance in stochastic optimal control [32,33] and extended in active inference [15], with connections to ideas of sensorimotor loops in biological systems [11,13]. These connections emphasise homeostasis, regulation and concepts such as set-point control and negative feedback for the study of different aspects of living systems, with roots in the cybernetics movement [34,35]. It remains, however, unclear how the active inference formulation directly relates to more traditional concepts of control theory. PID control, a popular control strategy working with little prior knowledge of the process to regulate, is commonly applied in engineering [36–38] and more recently used in biology and neuroscience modelling [39–43]. In this work, we develop an information theoretic interpretation of PID control, showing how it can be derived in a

more general Bayesian (active) inference framework. We will show that approximate models of the world are often enough for regulation, and in particular that simple linear generative models that only approximate the true dynamics of the environment implement PID control as a process of inference. Using this formulation we also propose a new method for the optimisation of the gains of PID controllers based on the same principle of variational free energy minimisation, and implemented as a second order optimisation process. Finally, we will show that our implementation of PID controllers as approximate Bayesian inference lends itself to a general framework for the formalisation of different (conflicting) criteria in the design of a controller, the so-called performance-robustness trade-off [38,44], as a cohesive set of constraints implemented in a free energy functional. In active inference, these criteria will be mapped to precisions, or inverse variances, of observations and dynamics of a state-space model with a straightforward interpretation in terms of uncertainty on different variables of a system.

In section 2 we will introduce PID control and give a brief overview of the recent literature highlighting the most common design principles used nowadays for PID controllers. The free energy principle will be presented in section 3, followed by a complete derivation of PID control as a form of active inference. In this section we will also propose that the parameters of a PID controller, its gains, can be optimised following the active inference formulation, which also captures modern design constraints and desiderata of PID controllers.

2. PID control

Proportional-Integral-Derivative (PID) control is one of the most popular types of controllers used in industrial applications, with more than 90% of total controllers implementing PID or PI (no derivative) regulation [38,45]. It is one of the simplest set-point regulators, whereby a desired state (i.e. set-point, reference, target) represents the final goal of the regulation process, e.g. to maintain a room temperature of 23° C. PID controllers are based on closed-loop strategies with a negative feedback mechanism that tracks the real state of the environment. In the most traditional implementation of negative feedback methods, the difference between the measured state of the variable to regulate (e.g. the real temperature in a room) and the target value (e.g. 23° C) produces a prediction error whose minimisation drives the controller's output, e.g. if the temperature is too high, it is decreased and if too low, it is raised. In mathematical terms:

$$e(t) = y_r - y(t) \quad (1)$$

where $e(t)$ is the error, y_r is the *reference* or set-point (e.g. desired temperature) and $y(t)$ is the observed variable (e.g. the actual room temperature).

This mechanism is, however, unstable in very common conditions, in particular when a steady-state offset is added (e.g. a sudden and unpredictable change in external conditions affecting the room temperature which are not under our control), or when fluctuations need to be suppressed (e.g. too many oscillations while regulating the temperature may be undesirable). PID controllers elegantly deal with both of these problems by augmenting the standard negative feedback architecture, here called *proportional* or *P term*, with an *integral* or *I* and a *derivative* or *D term*, see Fig. 1. The integral term accumulates the prediction error over time in order to cancel out errors due to unaccounted steady-state input, while minimising the derivative of the prediction error leads to a decrease in the amplitude of fluctuations of the controlled signal. The general form of the control signal $u(t)$ generated by a PID controller is usually described by:

$$u(t) = k_p e(t) + k_i \int_0^t e(\tau) d\tau + k_d \frac{de(t)}{dt} \quad (2)$$

where $e(t)$ is, again, the prediction error and k_p, k_i, k_d are the so called proportional, integral and derivative gains respectively, a set of parameters used to tune the relative strength of the P, I and D

terms of the controller. The popularity of PID controllers is largely due to their simple formulation and




Figure 1. A PID controller [46]. The prediction error $e(t)$ is given by the difference between a reference signal $r(t)$, y_r in our formulation, and the output $y(t)$ of a process. The different terms, one proportional to the error (P term), one integrating the error over time (I term) and one differentiating it (D term), drive the control signal $u(t)$.

implementation. One of the major challenges on the other hand, lies with the tuning of parameters k_p, k_i, k_d , that have to be adapted to deal with different (often conflicting) constraints on the regulation process [36,44].

2.1. The performance-robustness trade-off

The presence of conflicting criteria for the design of PID controller is a well known issue in the control theory literature, often referred to as the performance-robustness trade-off [38,44,47–49]. A controller needs to optimise pre-specified performance criteria while, at the same time, preserving some level of robustness in face of uncertainty and unexpected conditions during the regulation process. In recent attempts to formalise and standardise these general principles [38,44], the performance of a controller has been proposed to be evaluated through:

- load disturbance response, how a controller reacts to changes in external inputs, e.g. a step input,
- set-point response, how a controller responds to different set-points over time,
- measurement noise response, how noise on the observations impacts the regulation process,

while robustness to be assessed on:

- robustness to model uncertainty, how uncertainty on the plant/environment dynamics affects the controller.

The goal of a general methodology for the design and tuning of PID controllers is to bring together these (and possibly more) criteria into a formal and tractable framework that can be used for a large class of problems. One possible example is presented in [48] (see also [50,51] for other attempts). This methodology is based on the maximisation of the integral gain (equivalent to the minimisation of the integral of the error from the set-point, see [36]), subject to constraints derived from a frequency domain analysis related to the Nyquist stability criterion applied to the controlled system [48]. In this work, we propose our formulation also as a general framework for the design and tuning of PID controllers leveraging the straightforward interpretation of the performance-robustness trade-off for PID controllers in terms of uncertainty parameters (i.e. precisions or inverse variances) in the variational free energy.

3. The active inference framework

According to the free energy principle, living system must minimise the surprisal, or self-information, of their observations [4,9,10,19], defined as:

$$-\ln p(\psi|m) \quad (3)$$

where ψ is a set of sensory inputs conditioned on an agent m . Surprisal, in general, can in fact differ from agent to agent, with states that are favourable for a fish (in water), different from those favourable for a bird (out of water) (see [52] for a review on the value of information). According to the FEP, agents that minimise the surprisal of their sensory states over time will also minimise the entropy of their sensations, thus limiting the number of states they can physically occupy [4,19]. This minimisation is, however, intractable in any practical scenario since surprisal can be seen as the negative log-model evidence or negative log-marginal likelihood of observations ψ , with (omitting m for simplicity from now on) the marginal likelihood or model evidence expressed as:

$$p(\psi) = \int_{\vartheta} p(\psi, \vartheta) d\vartheta. \quad (4)$$

This integral is defined over all possible hidden variables, ϑ , of observations ψ . In many cases, the marginalisation is intractable since the latent space of ϑ may be high dimensional or the distribution may have a complex (analytical) form. In statistical mechanics, an approximation under variational formulations transforms this into an optimisation problem. The approximation goes by several names, including variational Bayes and ensemble learning [53,54], and constitutes the mathematical basis of the free energy principle. Using variational Bayes, surprisal can then be decomposed into [54]:

$$-\ln p(\psi) = F - D_{\text{KL}}(q(\vartheta) \parallel p(\vartheta|\psi)), \quad (5)$$

where

$$D_{\text{KL}}(q(\vartheta) \parallel p(\vartheta|\psi)) = \int q(\vartheta) \ln \frac{q(\vartheta)}{p(\vartheta|\psi)} d\vartheta, \quad (6)$$

is the Kullback-Leibler (KL) divergence [55], or relative entropy [54], an asymmetrical non-negative measure of the difference between two probability distributions. The first one, $p(\vartheta|\psi)$, represents the posterior distribution specifying the probability of hidden states, causes and parameters (ϑ) given observations ψ , while the second one $q(\vartheta)$, is the variational or recognition density which encodes current beliefs over hidden variables ϑ . The latter is introduced with the idea of approximating the (also) intractable posterior $p(\vartheta|\psi)$ with a simpler distribution, $q(\vartheta)$, and then minimising their difference through the KL divergence: when the difference is zero (following Jensen's inequality the divergence is always non-negative [54]), $q(\vartheta)$ is a perfect description of $p(\vartheta|\psi)$. Analogously, from the point of view of an agent, its goal is to explain the hidden states, causes and parameters ϑ of sensations ψ by approximating the posterior $p(\vartheta|\psi)$ with a known distribution, $q(\vartheta)$. The first term in equation 5 can be written as

$$F = \int q(\vartheta) \ln \frac{q(\vartheta)}{p(\vartheta, \psi)} d\vartheta \quad (7)$$

and is defined as (variational) free energy [8,12,56,57] for its mathematical analogies with free energy in thermodynamics, or [54] (negative) evidence lower bound in machine learning. Since the KL divergence is always non-negative we arrive at

$$D_{\text{KL}}(q(\vartheta) \parallel p(\vartheta|\psi)) \geq 0 \Rightarrow F \geq -\ln p(\psi) \quad (8)$$

which demonstrates that variational free energy is an upper bound to surprisal, since by minimising F we are guaranteed to minimise $-\ln p(\psi)$. To evaluate the variational free energy F , we must formalise

a recognition density $q(\theta)$ and a generative density $p(\theta, \psi)$ specific to an agent. Starting from the latter, we define a generative model formulated as a one dimensional *generalised* state-space model [12]:

$$\begin{aligned} \psi &= g(x, v) + z & \dot{x} &= x' = f(x, v) + w \\ \psi' &= g_x(x, v)x' + g_v(x, v)v' + z' & \dot{x}' &= x'' = f_x(x, v)x' + f_v(x, v)v' + w' \\ \psi'' &= g_x(x, v)x'' + g_v(x, v)v'' + z'' & \dot{x}'' &= x''' = f_x(x, v)x'' + f_v(x, v)v'' + w'' \\ &\vdots & &\vdots \end{aligned} \quad (9)$$

where ψ are the observations and $\theta = \{x, v, \theta, \gamma\}$, with x as the hidden states and v as the exogenous inputs, while θ and γ follow a partition in terms of parameters and hyperparameters defined in [12] and are specified later to simplify the notation now. Functions $g(\cdot)$ and $f(\cdot)$ map hidden states/inputs to observations and the dynamics of hidden states/inputs respectively. The prime symbols, e.g. x', x'', x''' are used to define higher orders of motion of a variable. Generalised coordinates of motion are introduced to represent non-Markovian continuous stochastic processes based on Stratonovich calculus, with strictly non-zero autocorrelation functions [12,58,59]. Ito's formulation of stochastic processes, on the other hand, is based on Wiener noise, where the autocorrelation can be seen as strictly equal to a delta function [59,60]. In general, the Stratonovich formulation is preferred in physics, where it is assumed that perfect white noise does not exist in the real world [61], while Ito's calculus is extensively used in mathematics/economics for its definition preserving the Martingale property [62]. It is proposed that models of biological systems should be based on the Stratonovich derivation [12], to accommodate more realistic properties of the physical world (i.e. non-Markovian processes). Using the Stratonovich interpretation, random processes can be described as analytic (i.e. differentiable) and become better approximations of real-world (weakly) coloured noise [60,63,64]. In this formulation, standard state-space models are extended, describing dynamics and observations for higher "orders of motion" encoding, altogether, a trajectory for each variable. The more traditional state-space description is based on Markovian processes (i.e. white noise) and can be seen as a special case of *generalised* state-space models defined here and in, for instance, [8,12]. When coloured noise is introduced, one should either define a high order autoregressive process expressed in terms of white noise [65] or embrace the Stratonovich formulation defining all the necessary equations in a state-space form [12]. The higher "orders of motion" introduced here can be thought of as quantities specifying "velocity" (e.g. $(\psi)'$), "acceleration" (e.g. $(\psi)''$), etc. for each variable, which is neglected in more standard formulations. For practical purposes, in equation (9) we also made a local linearity approximation on higher orders of motion suppressing nonlinear terms [8,12]. We introduce then a more compact form:

$$\tilde{\psi} = g(\tilde{x}, \tilde{v}) + \tilde{z} \quad \tilde{x}' = f(\tilde{x}, \tilde{v}) + \tilde{w} \quad (10)$$

where the tilde sign (e.g. $\tilde{\psi}$) summarises a variable and its higher orders of motion (e.g. $\tilde{\psi} = \{\psi, \psi', \psi'', \dots\}$). The stochastic model in equation (9) can then be described in terms of a generative density:

$$P(\tilde{\psi}, \tilde{x}, \tilde{v}; \theta, \gamma) = P(\tilde{\psi}|\tilde{x}, \tilde{v}; \theta, \gamma)P(\tilde{x}, \tilde{v}; \theta, \gamma) \quad (11)$$

In this case, we also make the conditional dependence on θ, γ explicit, defining θ as slowly changing parameters coupling hidden states and causes to observations, and hyperparameters γ as encoding properties of random fluctuations/noise $\tilde{w}$ and $\tilde{z}$. $P(\psi|x, v; \theta, \gamma)$ is a likelihood function describing the measurement law in equation (10), while the prior $P(\tilde{x}, \tilde{v}; \theta, \gamma)$ describes the system's dynamics. Under the Laplace approximation [66,67], the form of the recognition density $q(\theta)$ is specified in terms of a Gaussian distribution centred around the estimated mode (i.e. the mean for a Gaussian distributions) which can be evaluated using an extension of the EM algorithm [56,57]. Furthermore, (co)variances

can be solved analytically in terms of the Hessian of the free energy evaluated at the mode [8,67,68]. The variational free energy in equation (7) can then be simplified, without constants, to [8]:

$$F \approx -\ln P(\tilde{\psi}, \tilde{x}, \tilde{v}; \theta, \gamma) \Big|_{\tilde{\theta}=\tilde{\mu}_{\theta}} \quad (12)$$

where the condition $\tilde{\theta} = \tilde{\mu}_{\theta}$ represents the fact that the generative density $P(\tilde{\psi}, \tilde{x}, \tilde{v}; \theta, \gamma)$ will be approximated by a Gaussian distribution centred around the best estimates $\tilde{\mu}_{\theta}$ of the unknown $\tilde{\theta}$, following the Laplace method implemented in a variational context [66]. With Gaussian assumptions on random variables $\tilde{z}$ and $\tilde{w}$ in equation (10), the likelihood and prior in equation (11) are also Gaussian, and the variational free energy can be expressed as:

$$F \approx \frac{1}{2} \left[\pi_{\tilde{z}} \left(\tilde{\psi} - g(\tilde{\mu}_x, \tilde{\mu}_v) \right)^2 + \pi_{\tilde{w}} \left(\tilde{\mu}'_x - f(\tilde{\mu}_x, \tilde{\mu}_v) \right)^2 - \ln(\pi_{\tilde{z}} \pi_{\tilde{w}}) \right] \quad (13)$$

where $\tilde{x}$ and $\tilde{v}$ are replaced by their sufficient statistics, means/modes $\tilde{\mu}_x, \tilde{\mu}_v$, and sensory and dynamics/process precisions $\pi_{\tilde{z}}, \pi_{\tilde{w}}$, or inverse variances, of random variables $\tilde{z}$ and $\tilde{w}$. Following [12,56], the optimisation of the (Laplace-encoded) free energy with respect to expected hidden states $\tilde{\mu}_x$, equivalent to estimation or perception, can be implemented via a gradient descent:

$$\dot{\tilde{\mu}}_x = D\tilde{\mu}_x - \frac{\partial F}{\partial \tilde{\mu}_x} \quad (14)$$

while, considering how, from the perspective of agent, only observations ψ are affected by actions a (i.e. $\psi(a)$), control or action can be cast as:

$$\dot{a} = -\frac{\partial F}{\partial a} = -\frac{\partial F}{\partial \tilde{\psi}} \frac{\partial \tilde{\psi}}{\partial a} \quad (15)$$

representing a set of coupled differential equations describing a closed sensorimotor loop in terms of a physically plausible minimisation scheme [12]. The first equation includes a term $D\tilde{\mu}_x$ that represents the “mode of the motion” (also the mean for Gaussian variables) in the minimisation of states in generalised coordinates of motion [8,12,69], with D as a differential operator “shifting” the order of motion of $\tilde{\mu}_x$ such that $D\tilde{\mu}_x = \tilde{\mu}'_x$. More intuitively, since we are now minimising the components of a generalised state representing a trajectory rather than a static state, variables are in a moving frame of reference in the phase-space, and the minimisation is achieved when the temporal dynamics of the gradient descent match the ensemble dynamics of the estimates of hidden states, so for $\dot{\tilde{\mu}}_x = \tilde{\mu}'_x$ rather than for $\dot{\tilde{\mu}}_x = 0$ (which assumes that the mode of the motion is zero, as in standard state-space formulations with Markov assumptions). In the second equation, active inference makes the assumption that agents have innate knowledge of the mapping between actions a and observations $\tilde{\psi}$ (i.e. $\partial \tilde{\psi} / \partial a$) as reflex arcs, acquired on an evolutionary time scale, see [13,15] for discussion.

4. Results

4.1. PID control as active inference

To implement PID control as a process of active inference, we will first describe an agent’s generative model as a generalised linear state-space model of second order (i.e. only two higher orders of motion, anything beyond that is zero-mean Gaussian noise):

$$\begin{aligned} \psi &= x + z & \dot{x} &= x' = -\alpha(x + v) + w \\ \psi' &= x' + z' & \dot{x}' &= x'' = -\alpha(x' + v') + w' \\ \psi'' &= x'' + z'' & \dot{x}'' &= x''' = -\alpha(x'' + v'') + w'' \end{aligned}$$

where $\alpha \in \theta$ is a parameter. As previously suggested, with a Gaussian assumption on $\tilde{z}, \tilde{w}$, the likelihood is reduced to:

$$P(\tilde{\psi}|\tilde{x}, \tilde{v}; \theta, \gamma) = P(\tilde{\psi}|\tilde{x}; \theta, \gamma) = N(\tilde{\mu}_x, \sigma_{\tilde{z}}^2) \quad (16)$$

where we assume no direct dependence of observations $\tilde{\psi}$ on external inputs $\tilde{v}$, while the prior is described by:

$$P(\tilde{x}, \tilde{v}; \theta, \gamma) = P(\tilde{x}|\tilde{v}; \theta, \gamma)P(\tilde{v}; \theta, \gamma) \quad (17)$$

with

$$\begin{aligned} P(\tilde{x}|\tilde{v}; \theta, \gamma) &= N(-\alpha(\tilde{\mu}_x + \tilde{\mu}_v), \sigma_{\tilde{w}}^2) \\ P(\tilde{v}; \theta, \gamma) &= N(\tilde{\eta}_x, \sigma_{\tilde{v}}^2) \end{aligned} \quad (18)$$

The Laplace-encoded variational free energy in equation (13) then becomes:

$$\begin{aligned} F \approx \frac{1}{2} \left[\pi_z (\psi - \mu_x)^2 + \pi_{z'} (\psi' - \mu'_x)^2 + \pi_{z''} (\psi'' - \mu''_x)^2 + \pi_w (\mu'_x + \alpha(\mu_x - \eta_x))^2 + \right. \\ \left. + \pi_{w'} (\mu''_x + \alpha(\mu'_x - \eta'_x))^2 + \pi_{w''} (\mu'''_x + \alpha(\mu''_x - \eta''_x))^2 - \ln(\pi_z \pi_w \pi_{z'} \pi_{w'} \pi_{z''} \pi_{w''}) \right] \end{aligned} \quad (19)$$

To simplify our formulation, we assume that precisions $\pi_{\tilde{v}}$ tend to infinity (i.e. no uncertainty on the priors for $\tilde{v}$), so that $P(\tilde{v}; \theta, \gamma)$ in equation (18) becomes a delta function and inputs $\tilde{v}$ reduce to their prior expectations $\tilde{\eta}_x$, i.e. $\tilde{\mu}_v = \tilde{\eta}_x$. With this simplification, prior precisions $\pi_{\tilde{v}}$ and respective predictions errors ($\tilde{\mu}_v - \tilde{\eta}_x$) are not included in our formulation (see [56,57] for more general treatments). By applying the gradient descent described in equation (14) and equation (15) to our free energy functional, we then get the following update equations for perception (estimation):

$$\begin{aligned} \dot{\mu}_x &= \mu'_x - \left[-\pi_z (\psi - \mu_x) + \pi_w \alpha (\mu'_x + \alpha(\mu_x - \eta_x)) \right] \\ \dot{\mu}'_x &= \mu''_x - \left[-\pi_{z'} (\psi' - \mu'_x) + \pi_{w'} \alpha (\mu''_x + \alpha(\mu'_x - \eta'_x)) + \pi_w (\mu'_x + \alpha(\mu_x - \eta_x)) \right] \\ \dot{\mu}''_x &= \mu'''_x - \left[-\pi_{z''} (\psi'' - \mu''_x) + \pi_{w''} \alpha (\mu'''_x + \alpha(\mu''_x - \eta''_x)) + \pi_{w'} (\mu''_x + \alpha(\mu'_x - \eta'_x)) \right] \end{aligned} \quad (20)$$

and for action (control):

$$\dot{a} = - \left[\pi_z (\psi - \mu_x) \frac{\partial \psi}{\partial a} + \pi_{z'} (\psi' - \mu'_x) \frac{\partial \psi'}{\partial a} + \pi_{z''} (\psi'' - \mu''_x) \frac{\partial \psi''}{\partial a} \right]. \quad (21)$$

The mapping of these equations to a PID control scheme becomes more clear under a few simplifying assumptions. First, we assume strong priors on the causes of proprioceptive observations ψ ¹. Intuitively, these priors are used to define actions that change the observations to better fit the agent's desires, i.e. the target of the PID controller. This is implemented in the weighting mechanism of prediction errors by precisions in equation (19); see also [13,26,70] for similar discussions on the role of precisions for behaviour. In our derivation, weighted prediction errors on system dynamics, $\pi_{\tilde{w}}(\tilde{\mu}'_x + \tilde{\mu}_x - \tilde{\eta}_x)$, will be weighted more than weighted errors on observations, $\pi_{\tilde{z}}(\tilde{\psi} - \tilde{\mu}_x)$. To achieve

¹ For consistency with previous formulations, e.g. [8,13,15], we will define ψ as *proprioceptive* observations. Proprioception is the sense of position and movement of different parts of our body. For the car model we introduce later, this is equivalent for instance to readings of the velocity of the car.

this, we decrease sensory precisions $\pi_{\tilde{z}}$ on proprioceptive observations, effectively biasing the gradient descent procedure towards minimising errors on the prior dynamics [70]. Secondly, we set the decay parameter α to a large value (theoretically $\alpha \rightarrow \infty$, in practice $\alpha = 10^5$ in our simulations), obtaining a set of differential equations including only terms of order α^2 for perception:

$$\begin{aligned}\dot{\mu}_x &\approx -\pi_w \alpha \left(\alpha (\mu_x - \eta_x) \right) \\ \dot{\mu}'_x &\approx -\pi_{w'} \alpha \left(\alpha (\mu'_x - \eta'_x) \right) \\ \dot{\mu}''_x &\approx -\pi_{w''} \alpha \left(\alpha (\mu''_x - \eta''_x) \right)\end{aligned}\quad (22)$$

This can be interpreted as an agent encoding beliefs in a world that quickly settles to a desired equilibrium state. This assumption effectively decouples orders of generalised motion, with higher embedding orders not affecting the minimisation of lower ones in equation (20), since terms from lower orders are modulated by α directly. The remaining terms effectively impose constraints on the generalised motion only close to equilibrium, playing a minor role in the control process away from the target/equilibrium (the more interesting part of regulation). These terms are necessary for the system to settle to a proper steady state when $(\tilde{\mu}_x - \tilde{\eta}_x) \rightarrow 0$ and maintain consistency across generalised orders of motion for small fluctuations at steady state, but have virtually no influence at all in conditions far from equilibrium. Following equation (22), at steady state, expectations on hidden states $\tilde{\mu}_x$ are mainly driven by priors $\tilde{\eta}_x$:

$$\tilde{\mu}_x = \tilde{\eta}_x \quad (23)$$

but are still not met by appropriate changes in observations $\tilde{\psi}$ which effectively implement the regulation around the desired target. To minimise free energy in presence of strong priors, this agent will necessarily have to modify its observations $\tilde{\psi}$ to better match expectations $\tilde{\mu}_x$, which in turn are shaped by priors (i.e. desires) $\tilde{\eta}_x$. Effectively, the agent “imposes” its desires on the world, acting to minimise the prediction errors arising at the proprioceptive sensory layers. In essence, an active inference agent implements set-point regulation by behaving to make its sensations accord with its strong priors/desires. After these assumptions, action can be written as:

$$\dot{a} \approx - \left[\pi_z \left(\psi - \eta_x \right) \frac{\partial \psi}{\partial a} + \pi_{z'} \left(\psi' - \eta'_x \right) \frac{\partial \psi'}{\partial a} + \pi_{z''} \left(\psi'' - \eta''_x \right) \frac{\partial \psi''}{\partial a} \right] \quad (24)$$

where we still need to specify partial derivatives $\partial \tilde{\psi} / \partial a$. As discussed in [13], this step highlights the fundamental differences between the FEP and the more traditional forward/inverse models formulation of control problems in biological systems [71,72]. While these derivatives help in the definition of an inverse model (i.e. finding the correct action for a desired output), unlike more traditional approaches, active inference does not involve a mapping from hidden states $\tilde{x}$ to actions a , but is cast in terms of (proprioceptive) sensory data $\tilde{\psi}$ directly. This is thought to simplify the problem: from a mapping between *unknown* hidden states and actions, to a mapping between *known* proprioceptive observations $\tilde{\psi}$ and actions a . It is claimed that this provides an easier implementation for an inverse model [15], one that is grounded in an extrinsic frame of reference, i.e. the real world ($\tilde{\psi}$), rather than in an intrinsic one in terms of hidden states ($\tilde{x}$) to be inferred first. To achieve PID-like control, we assume that the agent adopts the simplest (i.e. linear) relationship between its actions (controls) and their effects on sensory input across all orders of motion:

$$\frac{\partial \psi}{\partial a} = \frac{\partial \psi'}{\partial a} = \frac{\partial \psi''}{\partial a} = 1. \quad (25)$$

This reflects a very simple reflex-arc-like mechanism that is triggered every time a proprioceptive prediction is generated: positive actions (linearly) increase the values of the sensed variables $\tilde{\psi}$, while

negative actions decrease them. There is, however, an apparent inconsistency here that we need to dissolve: the proprioceptive input ψ and its higher order states ψ', ψ'' are *all* linearly dependent with respect to actions a as represented in equation (25). While an action may not change position, velocity and acceleration of a variable in the same way, a generative model doesn't need to perfectly describe the system to regulate: these derivatives only encode sensorimotor dependencies that allow for, in this case, sub-optimal control. In the same way, PID controllers are, in most cases, effective but only approximate solutions for control [36,73]. This allows us to understand the encoding of an inverse model from the perspective of an agent (i.e. the controller) rather than assuming a perfect, objective mapping from sensations to actions that reflects exactly how actions affect sensory input [13]. This also points at possible investigations of generative/inverse models in simpler living systems where accurate models are not perhaps needed, and where strategies like PID control are implemented [39–41]. By combining equation (24) and equation (25), action can then be simplified to:

$$\dot{a} \approx \pi_z(\eta_x - \psi) + \pi_{z'}(\eta'_x - \psi') + \pi_{z''}(\eta''_x - \psi'') \quad (26)$$

which is consistent with the “velocity form” or algorithm of a PID controller [36]:

$$\dot{u} = k_i(y_r - y) + k_p \frac{d}{dt}(y_r - y) + k_d \frac{d^2}{dt^2}(y_r - y). \quad (27)$$

Velocity forms are used in control problems where, for instance, integration is provided by an external mechanism outside the controller [36,73]. Furthermore, velocity algorithms are the most natural form for the implementation of integral control to avoid windup effects of the integral term, emerging when actuators can't regulate an indiscriminate accumulation of steady-state error in the integral term due to physical limitations [36,74]. This algorithm is usually described using discrete systems to avoid the definition of the derivative of random variables, often assumed to be white noise in the Ito's sense (i.e. Markovian processes). In the continuous case, if the variable y is a Markov process, its time derivative is in fact not well defined. For this form to exist in continuous systems, y must be a smooth (stochastic) process. Effectively, this drops the Markov assumption of white noise and implements the same definition of analytic (i.e. differentiable) noise related to Stratonovich calculus and the generalised coordinates of motion we described earlier. The presence of extra prediction errors beyond the traditional negative feedback (proportional term) can, in this light, be seen as a natural consequence of considering linear non-Markovian processes with simple reflex mechanisms responding to position, velocity and acceleration in the generalised motion phase space (see equation (25)). To ensure that the active inference implementation approximates the velocity form of PID control we still need to clarify the relationship between the generalised coordinates of motion in equation (26) and the differential operators $d/dt, d^2/dt^2$ in equation (27). As pointed out in previous work, when the variational free energy is minimised, the two of them are equal since the motion of the mode becomes the mode of the motion [8,56]. To simplify our formulation and show PID control more directly, we can consider the case for $\eta'_x = \eta''_x = 0$, defining the more standard set-point control where a desired or set-trajectory collapses to a single set-point in the state-space and equivalent, in the velocity form, to the case where y_r is a constant and $dy_r/dt = d^2y_r/dt^2 = 0$.

To show an implementation of PID control through active inference we use a standard model of cruise control, i.e. a car trying to maintain a certain velocity over time². While only a toy model, the intuitions and results we derive can easily be transferred to the regulation of proteins in bacterial chemotaxis [39] or yeast osmoregulation [75], and more generally to any homeostatic mechanism [34], especially when including limited knowledge of external forces [76]. In this setup, a controller receives the speed of the car as an input and adapts the throttle of the vehicle based on a negative feedback

² Our code is available at <https://github.com/mbaltieri/PIDControlActiveInferenceFEP>.

mechanism to achieve the desired, or target, cruising velocity. In real-world scenarios, this mechanism needs to be robust in presence of external disturbances, essentially represented by changes in the slope of the road, wind blowing, etc., see Fig. 2d. For simplicity, we will use the model based on the formulation in [73], see also Appendix A. In this particular instance, we will provide a simple proof of concept, simplifying PID to PI control as in [73], hence implementing only a first order generalised state-space model (see equation (16)). The controller receives noisy readings ψ, ψ' of the true velocity and acceleration of the car, x, x' , following the formulation in equation (16). The controller is provided with a Gaussian prior in generalised coordinates encoding desired velocity and acceleration with means $\eta_x = 10 \text{ km/h}$, $\eta'_x = 0 \text{ km/h}^2$. This prior represents a target trajectory for the agent that, as we saw in equation (26), will be equivalent to integral and proportional terms of a PI controller in velocity form. The recognition dynamics ([69]) are then specified in equation (20) and equation (21).




Figure 2. A cruise controller based on PI control under active inference. (a) The response of the car velocity over time with a target state, or prior in our formulation, $\eta_x = 10 \text{ km/h}$, $\eta'_x = 0 \text{ km/h}^2$. (b) The acceleration of the car over time with a specified prior $\eta'_x = 0 \text{ km/h}^2$. (c) The external force v , introduced at $t = 150\text{s}$, models a sudden change in the environmental conditions, for instance wind or change in slope. Action obtained via the minimisation of variational free energy with respect to a and counteracts the effects of v . The motor action is never zero since we assume a constant slope, $\lambda = 4^\circ$ (see table A1, Appendix A). (d) The model car we implemented, where v could be thought as a sudden wind or a changing slope.

In Fig. 2 we show the behaviour of a standard simulation of active inference implementing PI-like control for the controller of the speed of a car. The sensory and process precisions $\pi_{\tilde{z}}, \pi_{\tilde{w}}$ are fixed, to show here only the basic disturbance rejection property of PID controllers [36,76]. In Fig. 2a, after the car is exposed to some new external condition (e.g. wind) represented in Fig. 2c and not encoded in the controller's generative model, the regulation process brings the velocity of the car back to the desired state after a short transition period. Fig. 2b shows how sudden changes in the acceleration of

the car are quickly cancelled out in accord with the specified prior $\eta'_x = 0 \text{ km/h}^2$. The action of the car is then shown, as one would expect [76], to counteract the external force v , Fig. 2c.

4.2. Responses to external and internal changes

It is often desirable for a PID regulator to provide different responses to external perturbations (e.g. wind), which should be rather rapid, and to internal updates (e.g. a shift in target velocity) which should be relatively smooth [36,45], see also section 2.1. It is not, however, trivial to identify and isolate parameters that contribute to these effects [37,77,78], and thus to tune these properties independently. It has been suggested that in order to achieve such decoupling, a controller with two degrees of freedom is necessary [45,77]. Such controller can be thought to contain a feedforward model of the dynamics of the observed/regulated system [73]. In our implementation, this is elegantly achieved by construction, since active inference is based on generative (forward) models. Specifically, we can fix the response to external forces by setting the expected sensory precisions $\pi_{\tilde{z}}$ (i.e. PI gains) but then independently tune the response to changes in the setpoint by altering the expected process precisions $\pi_{\tilde{w}}$ on the priors, see Fig. 3a and Fig. 3b.




Figure 3. Different responses to load disturbances and set-point changes. The simulations were 300s long, with an external disturbance/different target velocity introduced at $t = 150\text{s}$. Here we report only a 20 seconds time window around the change in conditions. (a) The same load disturbance ($v = 3.0\text{km/h}^2$) is applied with varying expected process precisions $\pi_{\tilde{w}}$ where $\pi_w = \{\exp(-24), \exp(-22), \exp(-20)\}$. Expected sensory log-precisions $\pi_{\tilde{z}}$ are fixed over the duration of the simulations, with $\mu_{\gamma_z} = 1$. (b) A similar example for changes in the target velocity of the car, from $\eta_x = 13\text{km/h}$ to $\eta_x = 10\text{km/h}$, tested on varying expected process precisions $\pi_{\tilde{w}}$ where $\pi_w = \{\exp(-24), \exp(-22), \exp(-20)\}$.

In the limit for process prediction errors $\pi_{\tilde{w}}(\tilde{\mu}'_x + \alpha(\tilde{\mu}_x - \tilde{\eta}_x))$ much larger than the sensory ones $\pi_{\tilde{z}}(\tilde{\psi} - \tilde{\mu}_x)$ and with fixed expected sensory precisions $\pi_{\tilde{z}}$, the response to load disturbances is invariant (Fig. 3a). A new target velocity for the car creates different responses with varying $\pi_w = \{\exp(-24), \exp(-22), \exp(-20)\}$ ³. Larger $\pi_{\tilde{w}}$ values imply an expected low uncertainty on the dynamics (i.e. changes to the set-point are not encoded and therefore not expected) and are met almost instantaneously with an update of expected hidden states $\tilde{\mu}_x$, matched by suitable actions a . On the other hand, smaller $\pi_{\tilde{w}}$ account for higher variance/uncertainty and thus changes in the target velocity are to be expected, making the transitions to new reference values slower, as seen in Fig. 3b.

³ Precisions on higher embedding orders are built, in both cases, using a smoothness (i.e. decay) factor of $1/2$, see [12].

4.3. Optimal tuning of PID gains

One of the main goals of modern design principles for PID controllers is to find appropriate tuning rules for the gains on the prediction errors: proportional, integral and derivative terms. However, existing approaches are often limited [37,38,44,48,78]. In general, the proportional term must bring a system to the target state in the first place, the integral of the error should promptly deal with errors generated by steady state inputs not accounted by a model [76], while the derivative term should reduce the fluctuations by controlling changes in the derivative of a variable [73]. In our car example, this could mean for example controlling the velocity of the vehicle in spite of changes such as the presence of wind or variations in slope of the road (I term) and avoiding unnecessary changes in accelerations close to the target (D term, even if sometimes not used for cruise control problems [73]). In our treatment of PID controllers as approximate Bayesian inference, the controllers' gains k_i, k_p, k_d become equivalent to sensory precisions $\pi_z, \pi_{z'}, \pi_{z''}$, cf. equation (26) and equation (27). Following [12,56,57], we thus propose to optimise these precisions to minimise the path integral of variational free energy (or free action), assuming that parameters and hyperparameters change on a much slower time scale. To do so, we extend our previous formulation and replace fixed sensory precisions $\pi_z, \pi_{z'}, \pi_{z''}$ with *expected* sensory precisions $\mu_{\pi_z}, \mu_{\pi_{z'}}, \mu_{\pi_{z''}}$, derived from a Laplace approximation applied not only to hidden states x but extended also to these hyperparameters, now considered as random variables to be estimated, rather than fixed quantities [56,57].

Active inference provides then an analytical criterion for the tuning of PID gains in the temporal domain, where otherwise mostly empirical methods or complex methods in the frequency domain have insofar been proposed [36,38,47,48]. In frameworks used to implement active inference, such as DEM [12,56], parameters and hyperparameters are usually assumed to be conditionally independent of hidden states based on a strict separation of time scales (i.e. a mean-field approximation). This assumption prescribes a minimisation scheme with respect to the path-integral of free energy, or free action, requiring the explicit integration of this functional over time. In our work, however, for the purposes of building an online self-tuning controller, we will treat expected sensory precisions as conditionally dependent but changing on a much slower time-scale with respect to states x , using a second order online update scheme based on generalised filtering [57]. The controller gains, $\mu_{\pi_z}, \mu_{\pi_{z'}}, \mu_{\pi_{z''}}$, will thus be updated specifying instantaneous changes of the curvature of expected precisions with respect to variational free energy rather than first order updates with respect to free action:

$$\ddot{\mu}_{\pi_z} = -\frac{\partial F}{\partial \mu_{\pi_z}} \quad (28)$$

Expected precisions μ_{π_z} should however be non-negative since variances need to be positive, a fact also consistent with the negative feedback principle behind PID controllers (i.e. negative expected precisions would apply a positive feedback). To include this constraint, following [66] we thus parametrise sensory precisions π_z (and consequently expected sensory precisions μ_{π_z}) in the generative model as:

$$\pi_z = \exp(\gamma_z) \quad (29)$$

creating, effectively, log-normal priors and making them strictly positive thanks to the exponential mapping of hyperparameters γ . The scheme in equation (28) is then replaced by one in terms of expected sensory log-precisions μ_{γ_z} :

$$\ddot{\mu}_{\gamma_z} = -\frac{\partial F}{\partial \mu_{\gamma_z}} \quad (30)$$

For practical purposes, the second order system presented in equation (30) is usually reduced to a simpler set of first order differential equations [8]:

$$\begin{aligned}\dot{\mu}_{\gamma_z} &= \mu'_{\gamma_z} \\ \dot{\mu}'_{\gamma_z} &= -\frac{\partial F}{\partial \mu_{\gamma_z}} - \kappa \mu'_{\gamma_z}\end{aligned}\quad (31)$$

where μ'_{γ_z} is a prior on the motion of hyperparameters γ which encodes a “damping” term for the minimisation of free energy F ⁴. This term enforces hyperparameters to converge to a solution close to the real steady state thanks to a drag term for $\kappa > 0$ ⁵. The parametrisation of expected precisions in terms of log-precisions γ_z , in fact, makes the derivative of the free energy with respect to log-precisions strictly positive ($\partial F / \partial \gamma_z > 0$), not providing a steady-state solution for the gradient descent [57]. This “damping” term stabilises the solution, reducing the inevitable oscillations around the real equilibrium of the system. Given the free energy defined in equation (19), with $\exp(\mu_{\gamma_z})$ replacing π_z , the minimisation of expected sensory log-precisions (or “log- PID gains”) is prescribed by the following equations:

$$\begin{aligned}\dot{\mu}_{\gamma_z} &= \mu'_{\gamma_z} \\ \dot{\mu}'_{\gamma_z} &= -\frac{\partial F}{\partial \mu_{\gamma_z}} - \kappa \mu'_{\gamma_z} = -\frac{1}{2} \left[\exp(\mu_{\gamma_z}) (\psi - \mu_x)^2 - 1 \right] - \kappa \mu'_{\gamma_z} \\ \dot{\mu}_{\gamma_{z'}} &= \mu'_{\gamma_{z'}} \\ \dot{\mu}'_{\gamma_{z'}} &= -\frac{\partial F}{\partial \mu_{\gamma_{z'}}} - \kappa \mu'_{\gamma_{z'}} = -\frac{1}{2} \left[\exp(\mu_{\gamma_{z'}}) (\psi' - \mu'_x)^2 - 1 \right] - \kappa \mu'_{\gamma_{z'}} \\ \dot{\mu}_{\gamma_{z''}} &= \mu'_{\gamma_{z''}} \\ \dot{\mu}'_{\gamma_{z''}} &= -\frac{\partial F}{\partial \mu_{\gamma_{z''}}} - \kappa \mu'_{\gamma_{z''}} = -\frac{1}{2} \left[\exp(\mu_{\gamma_{z''}}) (\psi'' - \mu''_x)^2 - 1 \right] - \kappa \mu'_{\gamma_{z''}}\end{aligned}\quad (32)$$

236 This scheme introduces a new mechanism for the tuning of the gains of a PID controller, allowing
237 the controller to adapt to adverse and unexpected conditions in an optimal way, in order to avoid
238 oscillations around the target state.

In Fig. 4 the controller for the car velocity is initialised with suboptimal sensory log-precisions μ_{γ_z} , i.e. log-PI gains. The parameters were initially not updated (Fig. 4d) to allow the controller to settle around the desired state, see Fig. 4a. The adaptation begins at $t = 30$ s and is stopped at $t = 150$ s, when an external force is introduced, to test the response of the controller after the gains have been optimised. With the adaptation process, the controller becomes more responsive when facing external disturbances (cf. Fig. 2), quickly and effectively counteracted by prompt changes in controls, see Fig. 4c. As a trade-off, the variances of the velocity and the acceleration are however increased, see Fig. 4a and see Fig. 4b. The optimisation of the gains through μ_{γ_z} without extra constraints (if not the stopping condition we imposed at $t = 150$ s, after the adaptation reaches a steady-state) effectively introduces an extremely responsive controller: cancelling out the effects of unwanted external inputs, such as wind in our cruise control example, but also more sensitive to measurement noise. In Fig. 5 we show summary statistics with the results of the adaptation of the gains. Following the examples in Fig. 2 and Fig. 4, we simulated 20 different cars with expected sensory log-precisions μ_{γ_z} sampled uniformly in the interval $[-4, -2]$ and expected process log-precisions μ_{γ_w} in the interval $[-23, -21]$. We initially maintained (i.e. no adaptation) the same hyperparameters and introduced a load disturbance at $t = 150$ s, then

⁴ In [57] we can see that this is equivalent to the introduction of a prior $p(\tilde{\gamma})$ on the motion of $\tilde{\gamma}$ to be zero (i.e. zero mean) with precision 2κ .

⁵ $\kappa = 5$ in our simulations




Figure 4. Optimising PID gains as expected sensory log-precisions μ_{γ_z} . This example shows the control of the car velocity before and after the optimisation of μ_{γ_z} (before and after the vertical dash dot black line) is introduced. (a) The velocity of the car. (b) The acceleration of the car. (c) The action of the car, with an external disturbance introduced at $t = 150$ s. (d) The optimisation of expected sensory precisions μ_{γ_z} and their convergence to an equilibrium state, after which the optimisation is stopped before introducing an external force. The blue line represents the true log-precision of observation noise in the system, $\gamma_z = \gamma_{z'} = 5$.

repeated the simulations (20 cars) with the same initial conditions allowing for the adaptation of expected sensory log-precisions as log-PI gains after $t = 30$ s, as in Fig. 4. Following [79], we measured the performance of the controllers by defining the integral absolute error (IAE):

$$IAE = \int_t^{t+\tau} |e(t)| dt \quad (33)$$

between two zero-crossings: the last time the velocity was at the target value before a disturbance is introduced, assumed to be $t = 150$ in our case, and the first time the velocity goes back to the target after a disturbance is introduced ($t + \tau$). To compute $t + \tau$, we took into account the stochasticity of the system and errors due to numerical approximations, considering the case for the real velocity to be within a ± 0.5 km/h interval away from the target value. The IAE captures the impact of oscillations on the regulation problem by integrating the error over the temporal interval where the car is pushed away from its target due to some disturbance (for more general discussions on its role and uses see [36]). As we can see in Fig. 5, the IAE converges to a single value for all cars (taking into account our approximation of a ± 0.5 km/h interval while measuring it) and is clearly lower when the adaptation mechanism for expected sensory log-precisions is introduced, making the controller very responsive to external forces and thus reducing the time away from the target velocity, see Fig. 4 for an example.




Figure 5. Performance of PID controllers with and without adaptation of the gains based on the minimisation of free energy. The integral absolute error (IAE) is used to measure the effects of the oscillations introduced by a single load disturbance at $t = 150s$ (see text for the exact definition of the IAE).

5. Discussion

In this work we developed a minimal account of regulation and control mechanisms based on active inference, a process theory for perception, action and higher order functions expressed via the minimisation of variational free energy [4,8,10,13]. Our implementation constitutes an example of the parsimonious, action-oriented models described in [24,25], connecting them to methods from classic control theory. We focused in particular on Proportional-Integral-Derivative (PID) control, both extensively used in industry [36–38,78] and more recently emerging as a model of robust feedback mechanisms in biology, implemented for instance by bacteria [39], amoeba [40] and gene networks [41], and in psychology [42]. PID controllers are ubiquitous in engineering mostly due to the fact that one needs only little knowledge of the process to regulate. In the biological sciences, this mechanism is thought to be easily implemented even at a molecular level [43] and to constitute a possible account for limited knowledge of the external world in simple agents [76].

Following our previous work on minimal generative models [26], we showed that this mechanism corresponds, in active inference terms, to linear generative models for agents that only approximate properties of the world dynamics. Specifically, our model describes linear dynamics for a single hidden or latent state and a linear mapping from the hidden state to an observed variable, representing knowledge of the world that is potentially far removed from the real complexity behind observations and their hidden variables. To implement such model, we defined a generative model that only approximates the environment of an agent and showed how under a set of assumptions including analytic (i.e. non-Markovian, differentiable) Gaussian noise and linear dynamics, this recapitulates PID control. A crucial component of our formulation is the presence of low sensory precision parameters on proprioceptive prediction errors of our free energy function or equivalently, high expected variance of proprioceptive signals. These low precisions play two roles during the minimisation of free energy: (1) they implement control signals as predictions of proprioceptive input influenced by strong priors (i.e. desires) rather than by observations, see equation (24) and [13], and (2) they reflect a belief that there are large exogenous fluctuations (low precision = high variance) in the observed proprioceptive input. This last point can be seen as the well known property of the Integral term [73,76] of PID controllers, dealing with unexpected external input (i.e. large exogenous fluctuations). The model represented by derivatives $\partial\tilde{\psi}/\partial a$ encodes then how actions a approximately affect observed proprioceptive sensations $\tilde{\psi}$, with an agent implementing a sensorimotor mapping that does not match the real dynamics of actions applied to the environment. The formulation in equation (20) and equation (21) can in general

be applied to different tasks, in the same way PID control is used in different problems without specific knowledge of the system to regulate.

The generative model we used is expressed in generalised coordinates of motion, a mathematical construct used to build non-Markovian continuous stochastic models based on Stratonovich calculus. Their importance has been expressed before [12,56,57], for the treatment of real world processes best approximated by continuous models and for which Markov assumptions don't really hold (see also [69] for discussion). The definition of a *generalised* state-space model provides then a series of weighted prediction errors and their higher orders of motion from the start, with PID control emerging as the consequence of an agent trying to impose its desired prior dynamics on the world via the approximate control of its observations on different embedding orders (for I, P and D terms). In this light, the ubiquitous efficacy of PID control may thus reflect the fact that the simplest models of controlled dynamics are first-order approximations to generalised motion. This simplicity is mandated because the minimisation of free energy is equivalent to the maximisation of model evidence, which can be expressed as accuracy minus complexity [10,24]. On this view, PID control emerges via the implementation of constrained (parsimonious, minimum complexity) generative models that are, under some constraints, the most effective (maximum accuracy) for a task.

In the control theory literature, many tuning rules for PID gains have been proposed (e.g. Ziegler-Nichols, IMC, etc., see [36,38] for a review) and used in different applications [36–38,48,78], however most of them produce quite different results, highlighting their inherent fit to only one of many different goals of the control problem. With our active inference formulation, we argue that different criteria can and should be expressed within the same set of equations in order to better understand their implications for a system. Modern approaches to the study of PID controllers propose four points as fundamental features to be considered for the design of a controller [44]:

- load disturbance response
- set-point response
- measurement noise response
- robustness to model uncertainty.

In our formulation, these criteria can be interpreted using precision (inverse variance) parameters of different prediction errors in the variational free energy, expressing the the uncertainty associated to observations and priors, as reported in table 1, see also Appendix B for further reference.

Table 1. Active inference as a general framework for PID controllers.

Criterion	Mapped to	Advantages in active inference
Load disturbance response	$\mu\pi_z$	Intuitively expressed via the expected inverse variance of the observations (i.e. precision), with low variance implying a fast response and vice versa (see section 4.2 and section 4.3)
Set-point change response	$\mu\pi_\phi$	Natural formulation of PID controllers with two degrees of freedom derived from sensory and process precisions and expressed as a Bayesian inference process (see section 4.2)
Measurement noise response	$\mu\pi_z$	Straightforward interpretation of PID gains as (expected) inverse variances of different embedding orders of measurement noise (see Appendix B)
Robustness to model uncertainty	$\mu\pi_\phi$	Direct mapping of model uncertainty to expected variances of the fluctuations, representing unknown dynamics, of the system to control (see Appendix B)

After establishing the equivalence between PID control and linear approximations of generalised motion in generative models, we showed that the controllers' gains, k_i, k_p, k_d , are in our formulation equivalent to expected precisions, $\mu\pi_z, \mu\pi_{z'}, \mu\pi_{z''}$, for which a minimisation scheme is provided in [12,56,57]. The basic version of this optimisation produces also promising results in presence of time-varying measurement (white) noise in the simulated car (see Fig. A1 in Appendix B). If the

adaptation is halted on a system with fixed measurement noise, it can be used to effectively deal with load disturbances, external forces acting against a system reaching his target (see Fig. 4), e.g. a change in chemicals concentration for a bacterium.

Future extensions could provide a more principled way of dealing with these two (and possibly other) conflicting cases, an issue that can be solved by introducing suitable hyperpriors (priors on hyperparameters) expressing the confidence of a system regarding changes in measurement noise via the use of precisions on hyperpriors [12]. High confidence (i.e. high precision on hyperpriors) would imply that a system should quickly react to sudden changes, both in measurement noise and other disturbances, since they are unexpected. On the other hand, low confidence (i.e. low precision on hyperpriors) would make a system's reaction to new conditions slower since such changes are expected. A trade-off between these conditions, with appropriate knowledge of a system or a class of systems introduced in the form of hyperpriors, would then make the process completely automatised, taking advantage of, for instance, empirical Bayes for learning such hyperpriors [10]. By extending our proposition with priors on precisions we can also, in principle, cast more criteria for the controller, expressing different requirements for more complex regulation processes. Given the fact that any optimality criterion can be recast as a prior, following the complete class theorem [80,81], as long as we know how to represent these rules as priors for the controller, we can provide any combination of requirements and tune the parameters in a straightforward fashion.

6. Conclusion

PID controllers are robust controllers used as a model of regulation for noisy and non-stationary processes in different engineering fields [38,73]. More recently, they have also been proposed as behavioural models of adaptive learning in humans [42] and as mechanistic explanations of different functions of systems in microbiology [39–41]. Their utmost relevance to the natural sciences is becoming clear, with implementations now proposed at the level of simple biomolecular interactions [43,82]. PID controllers are renowned for their simplicity and straightforward interpretation in control theory, however a general interpretation in probabilistic frameworks (e.g. Bayesian inference) is still missing.

Active inference has been proposed as a general mathematical theory of life and cognition according to the minimisation of variational free energy [10]. On this view, biological agents are seen as homeostatic systems maintaining their existence via the the minimisation of free energy. This process is implemented via the estimation and prediction of latent variables in the world (equivalent to perception) and the control of sensory inputs with behaviours accommodating normative constraints of an agent. Active inference is often described as an extension of optimal control theory with deep connections to Bayesian inference [15]. While methods such as PID control are still widely adopted as models of biological systems, it is unclear how general theories such as active inference connect to practical implementation of homeostatic principles such as PID control. In this work we proposed a way to connect these two perspectives showing how PID controllers can be seen as a special case of active inference. This account is based on the definition of a linear generative model for an agent approximating the dynamics of its environment, potentially very different from the information represented by the model. The model is expressed in generalised coordinates of motion [8,12,69] with prediction errors at different embedding orders for integral, proportional and derivative components emerging naturally as a consequence of an agent assuming non-Markovian dynamics on its sensory input. Through the use of active inference we also proposed the implementation of a mechanism for the optimisation of the gains of a PID controller, i.e. the weights of different prediction errors, now interpreted as precision parameters encoding the uncertainty of different variables from the perspective of an agent.

Author Contributions: MB: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing. CLB: Conceptualization, Methodology, Supervision, Writing.

Funding: This work was supported in part by a BBSRC Grant BB/P022197/1.

Acknowledgments: The authors would like to thank Karl Friston for thought-provoking discussions and insightful feedback on a previous version of this manuscript, and Martijn Wisse and Sherin Grimbergen for important comments on the initial mathematical derivation.

Conflicts of Interest: The authors declare no competing interests.

Appendix A The car model

The equation of motion of the car is:

$$m \frac{d^2 s}{dt^2} = F - F_d \quad (\text{A1})$$

where s is the position, F the force generated by the engine and F_d a disturbance force that accounts for a gravitational component F_g , a rolling friction F_r and an aerodynamic drag F_a , such that $F_d = F_g + F_r + F_a$, see again Fig. 2d. The forces will be modelled as following:

$$\begin{aligned} F &= r_g a(t) T_m \left(1 - \beta \left(\frac{\omega}{\omega_m} - 1 \right)^2 \right) \\ F_g &= mg \sin \lambda \\ F_r &= mg C_r \operatorname{sgn} \dot{s} \\ F_a &= \frac{1}{2} \rho C_d A \dot{s}^2 \end{aligned} \quad (\text{A2})$$

with all the constants and variables reported and explained in table A1.

Table A1. Cruise control problem, constants and variables.

	Description	Value
$s(t)$	car position	-
r_g	gear ratio divided by wheel radius	12
$a(t)$	control	-
T_m	maximum torque	190 Nm
β	motor constant	0.4
ω	engine speed	$\alpha_n v$
ω_m	speed that gives maximum torque	420 rad/s
m	car mass	100 kg
g	gravitational acceleration	9.81 m/s ²
λ	slope of the road	4°
C_r	coefficient of rolling friction	0.01
ρ	density of the air	1.3 kg/m ³
C_d	aerodynamic drag coefficient	0.32
A	frontal area of the car	2.4 m ²

Appendix B Measurement noise and model uncertainty in active inference

Nowadays, it is common to include two more desiderata for the design of PID controllers (see section 2 and [44]) in order to characterise and tune their response to (1) different types of measurement noise and (2) their robustness to model uncertainty, inherent in simple approximate controllers [38,44]. In our example, these properties map, respectively, to the response of a car given time-varying noise and to the available knowledge of a system, e.g. the working range of a controller or the type of disturbances affecting the car.

In particular, the former describes the behaviour of a PID controller in presence of noise on the observed variables by modulating the decay of different prediction errors in equation (26). It is known that this response is (in the limit for $t \rightarrow \infty$ and with the assumption of a system at equilibrium) inversely proportional to the integral gain [36,38]. In our case however, we have a more general and trivial relationship where the integral gain k_i is, by construction, equivalent to the inverse variance

(i.e. precision) of the measurement noise π_z , see equation (26) and equation (27). The remaining gains k_p, k_d can then be seen as encoding the uncertainty (i.e. precision) of higher orders of motion when the measurement noise is effectively coloured, otherwise just approximating possible correlations of the observed data over time.

On the other hand, the robustness to model uncertainty is related to expected process log-precisions $\mu_{\gamma_{\bar{w}}}$ encoding (again by construction) the amplitude of fluctuations due to unknown effects on the dynamics [12]. By modulating the prior dynamics of a system, these hyperparameters assume then a double role, they can either: (1) passively describe (estimate) the dynamics of a system (cf. Kalman filters [83]) or (2) actively impose desired trajectories on observations that can be implemented through actions on the world, as explained in section 4.1. With these conditions at the extremes, a spectrum of intermediate behaviours is also possible, with $\mu_{\gamma_{\bar{w}}}$ enacting different sensorimotor couplings by weighting the importance of objective information and desired states/priors of a system.

In the majority of the formulations of control problems, the properties of measurement noise and model uncertainty (especially their (co)variance) are assumed to be constant over time. Often, these parameters need also to be adapted to different systems since their properties are likely to be different. In section 4.3, we proposed an optimisation method for the gains of a PID controller based on active inference that here we exploit for time changing properties of the noise of a system, and that we show in an example when the measurement noise suddenly increases. In our car example, we could think of a drop in performance of the sensors recording velocity.




Figure A1. Performance of PID controllers with a sudden increase in measurement noise. 20 cars simulated in the case where measurement noise is increased at $t = 150s$ during the 300s simulations. We report aggregate results with the variance from the target value measured over the last 25% ($225 < t < 300s$) of a simulation. We show (1) the case for adaptation of the gains of the PI controller (through expected sensory log-precisions, or log-PI gains, μ_{γ_z}) interrupted before the measurement noise drastically changes, and (2) the case where the adaptation process persists for the entire duration of the simulations.

We simulated 20 cars for 300s with adaptation of expected sensory log-precisions (or log-PI gains) μ_{γ_z} , introduced at $t = 30s$ and stopped at $t = 150s$. At $t = 150s$ we then decreased the log-precision of measurement noise (n.b. not the expectation on the log-precision) from $\gamma_z = 5$ to $\gamma_z = 2$ for the rest of the simulations to simulate the partial failure of a sensor, and stopped the adaptation process. We then simulated 20 cars where adaptation was not halted after the increased measurement noise. To represent the difference, we measured the variance of the real velocity of the cars (without measurement noise to avoid biases), from $t = 225s$ to $t = 300s$ to allow the velocity to settle after the transient due to the sudden change. Agent that kept adapting their gains are shown to be more robust to persistent changes in noise, see Fig. A1.

In the case of model uncertainty, given the dual role of $\mu_{\gamma_{\bar{w}}}$ explained above, i.e. encoding prior dynamics reflecting both real properties of the environment and desired trajectories imposed on the system to regulate, it is harder to show the update of expected precisions without compromising the control of the car. The optimisation of variational free energy is, in fact, not intrinsically biased towards the control of a system, i.e. we externally imposed that as a condition for the agent. While having more flexible priors, an agent could potentially begin to account for uncertainty in the world rather than forcibly change its observations to reach its target.

References

- Dayan, P.; Hinton, G.E.; Neal, R.M.; Zemel, R.S. The Helmholtz Machine. *Neural computation* **1995**, *7*, 889–904.
- Rao, R.P.; Ballard, D.H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience* **1999**, *2*, 79–87.
- Knill, D.C.; Pouget, A. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences* **2004**, *27*, 712–719.
- Friston, K.J.; Kilner, J.; Harrison, L. A free energy principle for the brain. *Journal of Physiology-Paris* **2006**, *100*, 70–87.
- Clark, A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences* **2013**, *36*, 181–204.
- Hohwy, J. *The predictive mind*; OUP Oxford, 2013.
- Bogacz, R. A tutorial on the free-energy framework for modelling perception and learning. *Journal of mathematical psychology* **2017**, *76*, 198–211.
- Buckley, C.L.; Kim, C.S.; McGregor, S.; Seth, A.K. The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology* **2017**, *14*, 55–79.
- Friston, K.J. The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences* **2009**, *13*, 293–301.
- Friston, K.J. The free-energy principle: a unified brain theory? *Nature reviews. Neuroscience* **2010**, *11*, 127–138.
- Friston, K.J.; Rigoli, F.; Ognibene, D.; Mathys, C.; Fitzgerald, T.; Pezzulo, G. Active inference and epistemic value. *Cognitive neuroscience* **2015**, pp. 1–28. doi:10.1080/17588928.2015.1020053.
- Friston, K.J. Hierarchical models in the brain. *PLoS Computational Biology* **2008**, *4*.
- Friston, K.J.; Daunizeau, J.; Kilner, J.; Kiebel, S.J. Action and behavior: A free-energy formulation. *Biological Cybernetics* **2010**, *102*, 227–260.
- Friston, K.J.; Mattout, J.; Kilner, J. Action understanding and active inference. *Biological Cybernetics* **2011**, *104*, 137–160.
- Friston, K.J. What is optimal about motor control? *Neuron* **2011**, *72*, 488–498.
- Friston, K.J. Life as we know it. *Journal of the Royal Society Interface* **2013**, *10*, 20130475.
- Seth, A.K. The Cybernetic Bayesian Brain. In *Open MIND*; Open MIND. Frankfurt am Main: MIND Group, 2014.
- Evans, D.J.; Searles, D.J. The fluctuation theorem. *Advances in Physics* **2002**, *51*, 1529–1585.
- Friston, K.J. A free energy principle for biological systems. *Entropy* **2012**, *14*, 2100–2121.
- Pezzulo, G.; Cisek, P. Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends in cognitive sciences* **2016**, *20*, 414–424.
- Körding, K.P.; Wolpert, D.M. Bayesian decision theory in sensorimotor control. *Trends in cognitive sciences* **2006**, *10*, 319–326.
- Franklin, D.W.; Wolpert, D.M. Computational mechanisms of sensorimotor control. *Neuron* **2011**, *72*, 425–442.
- Friston, K.J.; Daunizeau, J.; Kiebel, S.J. Reinforcement learning or active inference? *PLoS One* **2009**, *4*, e421.
- Clark, A. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*; Oxford University Press, 2015.
- Clark, A. Radical predictive processing. *The Southern Journal of Philosophy* **2015**, *53*, 3–27.
- Baltieri, M.; Buckley, C.L. An active inference implementation of phototaxis. *Proc. Eur. Conf. on Artificial Life*, 2017, pp. 36–43.

27. Baltieri, M.; Buckley, C.L. A Probabilistic Interpretation of PID Controllers Using Active Inference. From Animals to Animats 15; Manoonpong, P.; Larsen, J.C.; Xiong, X.; Hallam, J.; Triesch, J., Eds. Springer International Publishing, 2018, pp. 15–26.
28. Clark, A. *Being there: Putting brain, body, and world together again*; MIT press, 1998.
29. Ashby, W.R. Requisite variety and its implications for the control of complex systems. *Cybernetica* **1958**, *1*, 83–99.
30. Shannon, C.E. Coding theorems for a discrete source with a fidelity criterion. *IRE Nat. Conv. Rec* **1959**, *4*, 1.
31. Kalman, R.E. Contributions to the theory of optimal control. *Bol. Soc. Mat. Mexicana* **1960**, *5*, 102–119.
32. Stengel, R.F. *Optimal control and estimation*; Courier Corporation, 1994.
33. Todorov, E. General duality between optimal control and estimation. Decision and Control, 2008. CDC 2008. 47th IEEE Conference on. IEEE, 2008, pp. 4286–4292.
34. Ashby, W.R. *An introduction to cybernetics*; Chapman & Hall Ltd., 1957.
35. Wiener, N. *Cybernetics or Control and Communication in the Animal and the Machine*; Vol. 25, MIT press, 1961.
36. Åström, K.J. *PID controllers: theory, design and tuning*; Research Triangle Park, 1995.
37. Ang, K.H.; Chong, G.; Li, Y. PID control system analysis, design, and technology. *IEEE transactions on control systems technology* **2005**, *13*, 559–576.
38. Åström, K.J.; Hägglund, T. *Advanced PID control*; ISA-The Instrumentation, Systems, and Automation Society Research Triangle Park, NC, 2006.
39. Yi, T.M.; Huang, Y.; Simon, M.I.; Doyle, J. Robust perfect adaptation in bacterial chemotaxis through integral feedback control. *Proceedings of the National Academy of Sciences* **2000**, *97*, 4649–4653.
40. Yang, L.; Iglesias, P.A. Positive feedback may cause the biphasic response observed in the chemoattractant-induced response of Dictyostelium cells. *Systems & control letters* **2006**, *55*, 329–337.
41. Ang, J.; Bagh, S.; Ingalls, B.P.; McMillen, D.R. Considerations for using integral feedback control to construct a perfectly adapting synthetic gene network. *Journal of theoretical biology* **2010**, *266*, 723–738.
42. Ritz, H.; Nassar, M.R.; Frank, M.J.; Shenhav, A. A Control Theoretic Model of Adaptive Learning in Dynamic Environments. *Journal of cognitive neuroscience* **2018**, pp. 1–17.
43. Chevalier, M.; Gomez-Schiavon, M.; Ng, A.; El-Samad, H. Design and analysis of a Proportional-Integral-Derivative controller with biological molecules. *bioRxiv* **2018**. doi:10.1101/303545.
44. Åström, K.J.; Hägglund, T. The future of PID control. *Control engineering practice* **2001**, *9*, 1163–1175.
45. Åström, K.J.; Hägglund, T. Revisiting the Ziegler–Nichols step response method for PID control. *Journal of process control* **2004**, *14*, 635–650.
46. Arturo Urquiza. PID controller – Wikipedia, the free encyclopedia, 2011. [Online; accessed March 30, 2018].
47. Rivera, D.E.; Morari, M.; Skogestad, S. Internal model control: PID controller design. *Industrial & engineering chemistry process design and development* **1986**, *25*, 252–265.
48. Åström, K.J.; Panagopoulos, H.; Hägglund, T. Design of PI controllers based on non-convex optimization. *Automatica* **1998**, *34*, 585–601.
49. Garpinger, O.; Hägglund, T.; Åström, K.J. Performance and robustness trade-offs in PID control. *Journal of Process Control* **2014**, *24*, 568–577.
50. Grimbale, M.; Johnson, M. Algorithm for PID controller tuning using LQG cost minimization. Proceedings of the 1999 American Control Conference (Cat. No. 99CH36251). IEEE, 1999, Vol. 6, pp. 4368–4372.
51. O'Brien, R.T.; Howe, J.M. Optimal PID controller design using standard optimal control techniques. 2008 American Control Conference. IEEE, 2008, pp. 4733–4738.
52. Kolchinsky, A.; Wolpert, D.H. Semantic information, autonomous agency and non-equilibrium statistical physics. *Interface Focus* **2018**, *8*, 20180041.
53. Beal, M.J. *Variational algorithms for approximate Bayesian inference*; University of London London, 2003.
54. Bishop, C.M. *Pattern Recognition and Machine Learning*; Springer-Verlag New York, 2006.
55. Kullback, S.; Leibler, R.A. On information and sufficiency. *The annals of mathematical statistics* **1951**, *22*, 79–86.
56. Friston, K.J.; Trujillo-Barreto, N.; Daunizeau, J. DEM: A variational treatment of dynamic systems. *NeuroImage* **2008**, *41*, 849–885.
57. Friston, K.J.; Stephan, K.; Li, B.; Daunizeau, J. Generalised filtering. *Mathematical Problems in Engineering* **2010**, *2010*.

- 513 58. Stratonovich, R.L. *Topics in the theory of random noise*; Vol. 2, CRC Press, 1967.
- 514 59. Jazwinski, A.H. *Stochastic Processes and Filtering Theory*; Vol. 64, Academic Press, 1970.
- 515 60. Fox, R.F. Stochastic calculus in physics. *Journal of statistical physics* **1987**, *46*, 1145–1157.
- 516 61. Wong, E.; Zakai, M. On the convergence of ordinary integrals to stochastic integrals. *The Annals of*
517 *Mathematical Statistics* **1965**, *36*, 1560–1564.
- 518 62. Moon, W.; Wettlaufer, J. On the interpretation of Stratonovich calculus. *New Journal of Physics* **2014**,
519 *16*, 055017.
- 520 63. Van Kampen, N.G. *Stochastic processes in physics and chemistry*; Vol. 1, Elsevier, 1992.
- 521 64. Klöden, P.E.; Platen, E. *Numerical Solution of Stochastic Differential Equations*; Springer, 1992.
- 522 65. Chui, C.K.; Chen, G. *Kalman filtering with Real-Time Applications*; Springer, 2017.
- 523 66. Friston, K.J.; Mattout, J.; Trujillo-Barreto, N.; Ashburner, J.; Penny, W. Variational free energy and the
524 Laplace approximation. *Neuroimage* **2007**, *34*, 220–234.
- 525 67. MacKay, D.J. *Information theory, inference and learning algorithms*; Cambridge university press, 2003.
- 526 68. Särkkä, S. *Bayesian filtering and smoothing*; Vol. 3, Cambridge University Press, 2013.
- 527 69. Kim, C.S. Recognition Dynamics in the Brain under the Free-Energy Principle. *Neural Computation* **2018**,
528 *30*, 2616–2659.
- 529 70. Brown, H.; Adams, R.A.; Parees, I.; Edwards, M.; Friston, K.J. Active inference, sensory attenuation and
530 illusions. *Cognitive processing* **2013**, *14*, 411–427.
- 531 71. Kawato, M. Internal models for motor control and trajectory planning. *Current opinion in neurobiology* **1999**,
532 *9*, 718–727.
- 533 72. Wolpert, D.M.; Ghahramani, Z. Computational principles of movement neuroscience. *Nature neuroscience*
534 **2000**, *3*, 1212.
- 535 73. Åström, K.J.; Murray, R.M. *Feedback systems: an introduction for scientists and engineers*; Princeton university
536 press, 2010.
- 537 74. Svrcek, W.Y.; Mahoney, D.P.; Young, B.R.; others. *A real-time approach to process control*; Wiley New York,
538 2006.
- 539 75. Muzzey, D.; Gómez-Urbe, C.A.; Mettetal, J.T.; van Oudenaarden, A. A systems-level analysis of perfect
540 adaptation in yeast osmoregulation. *Cell* **2009**, *138*, 160–171.
- 541 76. Sontag, E.D. Adaptation and regulation with signal detection implies internal model. *Systems & control*
542 *letters* **2003**, *50*, 119–126.
- 543 77. Araki, M.; Taguchi, H. Two-degree-of-freedom PID controllers. *International Journal of Control, Automation,*
544 *and Systems* **2003**, *1*, 401–411.
- 545 78. Johnson, M.A.; Moradi, M.H. *PID control*; Springer, 2005.
- 546 79. Hägglund, T. A control-loop performance monitor. *Control Engineering Practice* **1995**, *3*, 1543–1551.
- 547 80. Ferguson, T.S. *Mathematical statistics: A decision theoretic approach*; Academic press, 1967.
- 548 81. Brown, L.D. A complete class theorem for statistical problems with finite sample spaces. *The Annals of*
549 *Statistics* **1981**, pp. 1289–1300.
- 550 82. Briat, C.; Gupta, A.; Khammash, M. Antithetic integral feedback ensures robust perfect adaptation in noisy
551 biomolecular networks. *Cell systems* **2016**, *2*, 15–26.
- 552 83. Kalman, R.E. A new approach to linear filtering and prediction problems. *Journal of basic Engineering* **1960**,
553 *82*, 35–45.