

A probabilistic interpretation of PID controllers using active inference

Manuel Baltieri^{1,2} and Christopher L. Buckley^{1,2}
m.baltieri@sussex.ac.uk

¹Evolutionary and Adaptive Systems Group, Department of Informatics,
University of Sussex, Brighton, UK

²Sussex Neuroscience, University of Sussex, Brighton, UK

Abstract. In the past few decades, probabilistic interpretations of brain functions have become widespread in cognitive science and neuroscience. The Bayesian brain hypothesis, predictive coding, the free energy principle and active inference are increasingly popular theories of cognitive functions that claim to unify understandings of life and cognition within general mathematical frameworks derived from information and control theory, statistical physics and machine learning. The connections between information and control theory have been discussed since the 1950's by scientists like Shannon and Kalman and have recently risen to prominence in modern stochastic optimal control theory. However, the implications of the confluence of these two theoretical frameworks for the biological sciences have been slow to emerge. Here we argue that if the active inference proposal is to be taken as a general process theory for biological systems, we need to consider how existing control theoretical approaches to biological systems relate to it. In this work we will focus on PID (Proportional-Integral-Derivative) controllers, one of the most common types of regulators employed in engineering and more recently used to explain behaviour in biological systems, e.g. chemotaxis in bacteria and amoebae or robust adaptation in biochemical networks. Using active inference, we derive a probabilistic interpretation of PID controllers, showing how they can fit a more general theory of life and cognition under the principle of (variational) free energy minimisation under simple linear generative models.

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, theories inspired the Bayesian brain hypothesis such as predictive coding, the free energy principle and active inference have been used to explain brain processes including perception, action and higher order functions [15, 33, 31, 22, 12, 28, 9].

According to these theories, brains, and biological systems in general, should be thought of as Bayesian inference machines, gathering and representing information from the environment into generative models [22, 12, 28]. Such systems

in fact appear to estimate the latent causes of their sensory input in a process consistent with a Bayesian inference scheme. In particular it has been suggested that perceptual process can be accounted for in terms of predictive coding models whereby feedforward prediction errors and feedback predictions are combined under a generative model to infer the hidden causes of sensory data [33]. More recent theories have extended this proposal to account also for motor control and behaviour [26, 19]. On this view, behaviour is cast as a process of acting on the world to make sensory data better fit existing predictions. This latter process usually falls under the name of *active inference*.

Modelling approaches inspired by control theory are nowadays established methodologies for instance in psychology [32, 10], and are increasingly popular in fields such as biology [41, 40, 1]. Typically inspired by classical control theory and dynamical system theory, they emphasise regulation and concepts such as set-point control and negative feedback for the study of different aspects of living systems, an approach first introduced with cybernetics [4, 38]. In particular, methods such as PID (Proportional-Integral-Derivative) control have been widely used as they represent a very simple methodology with properties that guarantee robustness to perturbations and noise [41, 40, 1, 34].

The relationship between information/probability theory and control theory has long been recognised, with the first intuitions (as far as the authors are aware) emerging from work by Shannon [36] and Kalman [30]. A unifying view of these two theoretical frameworks is nowadays proposed for instance in stochastic optimal control [37] and active inference [20, 35]. How these ideas can be used to combine traditional concepts of control more commonly applied in biology with frameworks like active inference is still however unclear. Here, to address this, we develop an information theoretic interpretation of PID controllers, a very popular control strategy that works with little prior knowledge of the process to regulate. Starting from ideas proposed by the free energy principle, we will show that simple linear generative models approximating the true dynamics of the environment implicitly implement PID control as a process of active inference.

2 The free energy principle

The free energy principle (FEP) was initially introduced by Karl Friston [22] and later elaborated in a series of papers [18, 19, 24]. The FEP is proposed as a unifying theory for biological sciences with roots in information theory, thermodynamics and statistical mechanics. Work on the FEP has so far covered computational neuroscience [16, 17], and behavioural/cognitive neuroscience studies [26, 23]. Furthermore, connections have been implied with theories of biological self-organisation, information theory (e.g. infomax principle), optimal control, cybernetics and economics (e.g. utility theory) among others [19, 20, 21, 35]. According to the FEP, a living system exists only in a limited set of states over time, e.g. a fish can't survive for long out of water. Biological creatures can thus be seen as systems that minimise the surprisal (or surprise/self-information) of their sensory observations to maintain their existence, e.g. a fish' observations

should be limited to states in the water. Since this surprisal is not directly accessible by an agent [22], variational free energy is proposed as a proxy that can be minimised in its place, acting as an upper bound for such quantity.

In this study we focus on hypotheses and theories linked to the FEP regarding perception and action in agents, in particular predictive coding and active inference. Predictive coding [33] models of information processing in the brain prescribe a way in which top-down and bottom-up information flows could be combined in the cortex under deep generative models and are, on this view, often described a special case of the FEP [19]. Top-down processes provide predictions about sensory input while bottom-up activity carries prediction errors representing the difference between real and predicted sensations. These errors are then used to train a generative model to produce better predictions. The minimisation of prediction errors achieved by updating the predictions of this model to better represent an agent’s sensations corresponds, on this view, to perception. However, one of the main contributions of the FEP is the extension of predictive coding models to include an account of action, known as *active inference* [25, 26], and thus unify perception and action in a single cohesive mathematical framework where differences between the two processes almost vanish.

2.1 Active inference and control

Active inference provides a second way in which prediction errors, or free energy, can be minimised. While perceptual inference suppresses prediction errors only by updating predictions of a generative model of the incoming sensations [18, 26], active inference minimises errors also by directly acting on the environment to change sensory input to better accord with existing predictions. If a generative model encodes information about favourable states for an agent, then this process constitutes a way by which an agent can change its environment to better meet its needs. Thus, under the FEP, these two processes of error suppression allow an agent to both perceive and control the surrounding environment.

Most agent-based 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 [25, 26] 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 we have argued that this needs not be the case [7], 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* [11]). This idea is however common in other work [25, 26, 23, 28], where cognition and perception are presented as processes of inference to the best explanation, encoding an accurate set of parameters and variables of the environment with agents seen as mainly building sophisticated models of their worlds then used for action and behaviour. One often implicit assumption is that *all* the information available to an agent should be encoded (e.g. an ant modelling the entire environment). This is however not reasonable in dynamic and complex environments: when variables

and parameters in the world change too rapidly, accurate online inference and learning are implausible [3]. A possible alternative introduces action-oriented models entailing a more parsimonious approach where only task-relevant information is encoded [13, 14]. On this view, agents only model environmental properties that are necessary for their behaviour and their goals.

In this work we present an example of such parsimonious, action-oriented model described in [13, 14], connecting them to methods from classic control theory. We focus in particular on Proportional-Integral-Derivative (PID) control, both extensively used in industry [6] and more recently emerging as a model of robust feedback mechanisms in biology, implemented for instance by bacteria [41], amoeba [40] and gene networks [1], and in psychology [34]. PID controllers are ubiquitous in engineering mostly due to the fact that one needs only little knowledge of the process to regulate. In active inference terms, we will show that this corresponds to linear generative models that only approximate properties of the world dynamics. Specifically, our model will describe linear dynamics for a single hidden 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 causes.

3 PID control

Proportional-Integral-Derivative (PID) control is one of the simplest set-point controllers, whereby a desired state (i.e. set-point, reference) 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. The difference between such state and the target value (e.g. 23° C temperature) produces a prediction error whose minimisation drives the controller, e.g. if the temperature is too high, it is decreased and if too low, it is raised. In mathematical terms:

$$e(t) = y_r - y(t) \tag{1}$$

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 repressed (e.g. too many oscillations in the temperature on the trajectory to the desired state may be undesirable). PID controllers 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 steady-state input, while minimising the derivative of the prediction error leads to a decrease in the amplitude of fluctuations. 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.




Fig. 1. A PID controller [2]. 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 $u(t)$.

The popularity of PID controllers is largely due to: 1) their robustness in the presence of uncertainty, i.e. step disturbances and more in general noise, given by the filtering properties of the I term, and 2) an only approximate model of the dynamics of the process to regulate, based on a linearisation around the target state. This might look incompatible with standard work in active inference formulations suggesting a link to optimal control strategies with perfect models of process/environment, however, we argued previously that this needs not be the case [7]. Indeed one of the main strengths of active inference lies, according to us, in its general formulation and in generative models that do not have to mirror the dynamics of the entire environment.

4 PID control as active inference

In this work we will not provide a complete derivation of the active inference scheme, referring to previous treatments [27, 17, 9] for more details. We will begin from the Laplace encoded variational free energy for a univariate case:

$$F = -\ln P(\rho, x) + \text{constants} \quad (3)$$

where ρ represents the observed sensory input of an agent and x encodes the expectation of hidden states in the environment. The remaining constants will not be discussed since they play no role in the minimisation scheme we present.

As previously shown [17, 9], to minimise equation (3) we need to specify the agent’s generative density $P(\rho, x) = P(\rho|x)P(x)$ introducing a likelihood $P(\rho|x)$ and a prior $P(x)$ in terms of an agent’s expectations x . These probabilities can be specified by a generative model in the form of a state space model. In particular, to get the integral, proportional and derivative terms of a PID controller, we will use a *generalised* linear state space model of order 2 [17, 9]:

$$\begin{aligned} \rho &= x + z & \dot{x} &= x' = -\alpha(x + \eta + w) \\ \rho' &= x' + z' & \dot{x}' &= x'' = -\alpha(x' + \eta' + w') \\ \rho'' &= x'' + z'' & \dot{x}'' &= x''' = -\alpha(x'' + \eta'' + w'') \end{aligned} \quad (4)$$

where ρ is the observation of the proprioceptive signal and x is the estimated hidden state, α encodes the decay of x (α is here assumed to be very large, meaning that the generative model represents a belief in an environment that quickly relaxes to equilibrium), η encodes a desired state (e.g. desired temperature) represented mathematically as an exogenous input (or as a prior from higher layers in hierarchical models [17]) and z, w are zero-mean Gaussian random variables. As we shall see later, η is equivalent to y_r in equation (1). However, unlike standard PID schemes, η is here specified as a function of time using generalised coordinates of motion (explained below). Ultimately, this derivation will collapse to a more standard set-point scheme when $\eta = y_r$ and $\eta' = \eta'' = 0$. The prime (e.g. ρ', ρ'') indicates the order of generalised coordinate of motion [17, 9], which are introduced to represent non-Markovian continuous stochastic processes [29], in our case z, w . One could think of them as quantities conveying information about “velocity” (e.g. ρ'), “acceleration” (e.g. ρ''), etc. for each variable. Following [17, 9], we then define:

$$\tilde{\rho} = \tilde{x} + \tilde{z} \quad \dot{\tilde{x}} = -\alpha(\tilde{x} + \tilde{\eta} + \tilde{w}) \quad (5)$$

where the tilde sign (e.g. $\tilde{\rho}$) summarises the generalised state, a variable and its higher orders of motion, into a more compact description (e.g. $\tilde{\rho} = \{\rho, \rho', \rho'', \dots\}$).

With the assumption that random variables \tilde{z}, \tilde{w} are normally distributed (making the likelihood $P(\tilde{\rho}|\tilde{x})$ and the prior $P(\tilde{x})$ of Gaussian form), the variational free energy reduces to:

$$\begin{aligned} F &= \frac{1}{2} \left[\pi_z (\rho - \mu_x)^2 + \pi_{z'} (\rho' - \mu'_x)^2 + \pi_{z''} (\rho'' - \mu''_x)^2 + \pi_w (\mu'_x + \alpha(\mu_x - \eta))^2 \right. \\ &\quad \left. + \pi_{w'} (\mu''_x + \alpha(\mu'_x - \eta'))^2 + \pi_{w''} (\mu'''_x + \alpha(\mu''_x - \eta''))^2 - \ln(\pi_z \pi_w \pi_{z'} \pi_{z''} \pi_{w'} \pi_{w''}) \right] \quad (6) \end{aligned}$$

where we used the means $\tilde{\mu}_x$ of the estimated hidden states rather than the states \tilde{x} themselves since $\tilde{\mu}_x$ are the only sufficient statistics required for the minimisation of free energy under the assumption of optimal (co)variances of the recognition density (see [17, 9]). $\pi_z, \pi_{\tilde{w}}$ are the precision parameters (inverse variances) of \tilde{z}, \tilde{w} respectively. Following [27, 17], the optimisation of the Laplace encoded free energy can be performed via a standard gradient descent procedure:

$$\dot{\tilde{\mu}}_x = D\tilde{\mu}_x - \frac{\partial F}{\partial \tilde{\mu}_x} \quad \dot{a} = -\frac{\partial F}{\partial a} = -\frac{\partial F}{\partial \tilde{\rho}} \frac{\partial \tilde{\rho}}{\partial a} \quad (7)$$

where the two equations prescribe perception and action processes respectively. The first equation includes an extra term $D\tilde{\mu}_x$ that represents the “mean of the motion” in the minimisation of variables in generalised coordinates of motion [17, 9], with D as a differential operator with respect to time, i.e. $D\tilde{\mu}_x = \dot{\tilde{\mu}}_x$. More intuitively, since we are now minimising the components of a generalised state representing a trajectory rather than a static variable, variables are in a moving framework of reference where the minimisation is achieved for $\dot{\tilde{\mu}}_x = \tilde{\mu}'_x$ rather than $\dot{\tilde{\mu}}_x = 0$. At this point, the mean of the motion becomes the motion of the mean, thereby satisfying Hamilton’s principle of least action [17]. In the second equation, an assumption of active inference is that actions a only affect sensory input $\tilde{\rho}$ and furthermore that this mapping is known to the agent and enacted as a reflex mechanism, see [26] for discussion. By applying the gradient descent described in equation (7) to our free energy function in equation (6), we then obtain the following update equations for perception:

$$\begin{aligned}\dot{\mu}_x &= \mu'_x - \left[-\pi_z(\rho - \mu_x) + \pi_w\alpha(\mu'_x + \alpha(\mu_x - \eta)) \right] \\ \dot{\mu}'_x &= \mu''_x - \left[-\pi_{z'}(\rho' - \mu'_x) + \pi_{w'}\alpha(\mu''_x + \alpha(\mu'_x - \eta')) + \pi_w(\mu'_x + \alpha(\mu_x - \eta)) \right] \\ \dot{\mu}''_x &= \mu'''_x - \left[-\pi_{z''}(\rho'' - \mu''_x) + \pi_{w''}\alpha(\mu'''_x + \alpha(\mu''_x - \eta'')) + \pi_{w'}(\mu''_x + \alpha(\mu'_x - \eta')) \right]\end{aligned}\quad (8)$$

and for action:

$$\dot{a} = - \left[\pi_z(\rho - \mu_x) \frac{\partial \rho}{\partial a} + \pi_{z'}(\rho' - \mu'_x) \frac{\partial \rho'}{\partial a} + \pi_{z''}(\rho'' - \mu''_x) \frac{\partial \rho''}{\partial a} \right] \quad (9)$$

The mapping of these equations to a PID control scheme becomes more clear under a few simplifying assumptions, starting from an agent that will have strong priors (desires) on the causes of its proprioceptive observations. Intuitively, these priors will be used to define actions that change the observations to better fit the agent’s desires. This is implemented in the weighting mechanism of prediction errors by precision parameters in equation (6); see also [26, 8, 7] for similar discussions on precisions and behaviour. Here we want to weight prediction errors on the expected causes, $\pi_w(\tilde{\mu}'_x + \tilde{\mu}_x - \tilde{\eta})$, more than the ones on observations, $\pi_z(\tilde{\rho} - \tilde{\mu}_x)$. To achieve this, we decrease precisions π_z on proprioceptive observations, effectively biasing the gradient descent procedure towards minimising errors on the priors [8]. We then set the decay parameter α to a large value, meaning that the agent encodes beliefs in a world that quickly settles to an equilibrium state, with higher orders of generalised motion in each line of equation (8) not considered during the minimisation. Perception is then approximated as:

$$\begin{aligned}\mu_x &\approx \mu'_x - \pi_w\alpha(\mu'_x + \alpha(\mu_x - \eta)) \\ \mu'_x &\approx \mu''_x - \pi_{w'}\alpha(\mu''_x + \alpha(\mu'_x - \eta')) \\ \mu''_x &\approx \mu'''_x - \pi_{w''}\alpha(\mu'''_x + \alpha(\mu''_x - \eta''))\end{aligned}\quad (10)$$

where $\mu_x''' = 0$ since we truncated our generalised state-space model to order 2 (i.e. anything beyond that is zero-mean Gaussian noise). This system of equations sets, at steady state, the expected hidden states $\tilde{\mu}_x$ to the priors $\tilde{\eta}$, $\tilde{\mu}_x = \tilde{\eta}$.

To minimise free energy in presence of strong priors, the agent will necessarily have to modify its sensory input $\tilde{\rho}$ to better match expectations $\tilde{\mu}_x$ which in turn will be shaped by the priors (i.e. desires) $\tilde{\eta}$. Effectively, the agent ‘‘imposes’’ its desires on the world, driving actions that will minimise the prediction errors arising at the proprioceptive sensory layers. In essence, an active inference agent implements set-point regulation by acting to make its sensations accord with its strong priors/desires. After these assumptions, action can be written as:

$$\dot{a} \approx - \left[\pi_z (\rho - \eta) \frac{\partial \rho}{\partial a} + \pi_{z'} (\rho' - \eta') \frac{\partial \rho'}{\partial a} + \pi_{z''} (\rho'' - \eta'') \frac{\partial \rho''}{\partial a} \right] \quad (11)$$

where we assumed $\tilde{\mu}_x = \tilde{\eta}$, but still need to specify partial derivatives $\partial \tilde{\rho} / \partial a$. As discussed previously [26], this step also highlights the fundamental differences between the FEP and the more traditional forward/inverse models formulation of control problems [39]. While these derivatives define a form of inverse model, unlike more traditional approaches this does not involve a mapping between actions and hidden states \tilde{x} but is cast in terms of sensory data $\tilde{\rho}$ directly. It is claimed that this provides an easier implementation for such an inverse model [20], one that is grounded in an extrinsic frame of reference (i.e. the real world, $\tilde{\rho}$) rather than in an intrinsic one in terms of hidden states \tilde{x} . To achieve PID-like control, we finally assume that the agent adopts the simplest (i.e. linear) relationship between its actions (controls) and their effects on sensory input across all generalised coordinates of motion:

$$\frac{\partial \rho}{\partial a} = \frac{\partial \rho'}{\partial a} = \frac{\partial \rho''}{\partial a} = 1 \quad (12)$$

This reflects a very simple reflex-arc-like mechanism that is triggered any time a proprioceptive prediction is made. Intuitively, positive actions increase the values of the sensed variables $\tilde{\rho}$, while negative actions decrease them. There is however an apparent inconsistency here that we need to dissolve: the proprioceptive input ρ and its higher order states ρ', ρ'' are *all* linearly dependent on actions a as represented in equation (12). While an action may not change position, velocity and acceleration of a variable in the same way, the goal of an agent is not to perfectly represent its physical reality but just to encode sensorimotor properties that allow it to achieve its goals. In the same way, PID controllers are, in most cases, effective but only approximate solutions for control [5]. This allows us to understand the encoding of an inverse model from the perspective of an agent rather than assuming a perfect, objective mapping from sensations to actions that reflects exactly how actions affect sensory input [26]. This also points at possible investigations of generative/inverse models in simpler living systems where accurate internal models are not needed, and where strategies like PID control are implemented [41, 40, 1]. By combining equation (11) and equation (12), action can then be simplified to:

$$\dot{a} \approx \pi_z (\eta - \rho) + \pi_{z'} (\eta' - \rho') + \pi_{z''} (\eta'' - \rho'') \quad (13)$$

consistent with the “velocity form” or algorithm of a PID controller [5]:

$$\dot{u} = k_i(y_r - y) + k_p \frac{d}{dt}(y_r - y) + k_d \frac{d^2}{dt^2}(y_r - y) \quad (14)$$

where we removed the explicit dependence on time t . Velocity forms are used in control problems where, for instance, integration is provided by an external mechanism outside the controller [5, 6]. This algorithm is often described using discrete systems to avoid the definition of the derivative of random variables, often assumed to be white noise (i.e. Markov 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 process. This effectively drops the Markov assumption of white noise and implements the same definition of analytic (i.e. differentiable) noise and related 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 non-Markov processes. To ensure that the active inference implementation approximates the velocity form of PID control we then need to clarify the relationship between generalised coordinates of motion in equation (13) and the differential operators $d/dt, d^2/dt^2$ in equation (14). As pointed out in previous work [27, 9], the two of them are equal at the minimum of the free energy landscape, when the gradient descent has reached its steady state. To simplify our formulation and show this more directly, we could consider the case for $\eta' = \eta'' = 0$, defining the more standard set-point control where the desired trajectory collapses to a single point, equivalent in the velocity form to the case where y_r is a constant and $dy_r/dt = d^2y_r/dt^2 = 0$.

5 Conclusion

PID controllers are robust controllers used as a model of regulation for noisy and non-stationary processes in different disciplines, from engineering to biology and psychology. They however do not guarantee optimality, so a straightforward interpretation of this control strategy in terms of optimal control is missing. Active inference is often described as an extension of optimal control theory with deep connections to Bayesian inference [20]. While active inference has been proposed as a general mathematical theory of life and cognition according to the minimisation of variational free energy [19], methods such as PID control are still widely adopted as models of biological systems [41, 40, 1]. 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 once simplified (i.e. linear) generative models are introduced. 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 [19, 14]. On this view, PID

control emerges via the implementation of parsimonious (minimum complexity) generative models that are the most effective (maximum accuracy) for a task.

Following our previous work [7], we defined a generative model that only approximates the agent’s environment and showed how under a set of assumptions including analytic (i.e. non-Markovian, differentiable) Gaussian noise and linear dynamics, this model recapitulates PID control. A crucial component of our formulation is the presence of low precision parameters on proprioceptive prediction errors of our free energy function or equivalently, beliefs about high 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 (11) and [26, 7], and (2) they reflect a belief into the presence of large exogenous fluctuations (low precision = high variance) as part of the observed proprioceptive input. This last point can be seen as the well known property of the Integral term [6] of PID controllers, dealing with unexpected external input (i.e. large exogenous fluctuations). The model represented by derivatives $\partial\tilde{\rho}/\partial a$ encodes then how actions a approximately affect observed proprioceptive sensations $\tilde{\rho}$, with an agent implementing a sensorimotor mapping that does not match the real dynamics of actions applied to the environment. The generative model we proposed 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.

In future work we will explore the implications of PID control as active inference for the study of biological systems. In particular we suggest that given our formalisation it is trivial to generalise the set-point definition of PID controllers, based on point attractors, to trajectories (e.g. a reference temperature changing during the day with pre-specified properties such as the rate of change, etc.) using generalised coordinates of motion [27, 9]. We may also be able to provide a Bayes-optimal algorithm for the optimisation of the gains of a PID controller, k_i, k_p, k_d (i.e. precision parameters in our free energy formulation, $\pi_w, \pi_{w'}, \pi_{w''}$), for which only heuristic methods exist at the moment [5].

6 Acknowledgements

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

Bibliography

- [1] J. Ang, S. Bagh, B. P. Ingalls, and D. R. McMillen. Considerations for using integral feedback control to construct a perfectly adapting synthetic gene network. *Journal of theoretical biology*, 266(4):723–738, 2010.
- [2] Arturo Urquizo. PID controller – wikipedia, the free encyclopedia, 2011. [Online; accessed March 30, 2018].
- [3] W. Ashby. Requisite variety and its implications for the control of complex systems. *Cybernetica*, 1:83–99, 1958.
- [4] W. R. Ashby. *An introduction to cybernetics*. Chapman & Hall Ltd., 1957.
- [5] K. J. Åström. PID controllers: theory, design and tuning. *Instrument society of America*, 1995.
- [6] K. J. Åström and R. M. Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.
- [7] M. Baltieri and C. L. Buckley. An active inference implementation of phototaxis. In *Proc. Eur. Conf. on Artificial Life*, pages 36–43, 2017.
- [8] H. Brown, R. A. Adams, I. Parees, M. Edwards, and K. Friston. Active inference, sensory attenuation and illusions. *Cognitive processing*, 14(4):411–427, 2013.
- [9] C. L. Buckley, C. S. Kim, S. McGregor, and A. K. Seth. The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 14:55–79, 2017.
- [10] C. S. Carver and M. F. Scheier. *Attention and self-regulation: A control-theory approach to human behavior*. Springer Science & Business Media, 1981.
- [11] A. Clark. *Being there: Putting brain, body, and world together again*. MIT press, 1998.
- [12] A. Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(03):181–204, 2013.
- [13] A. Clark. Radical predictive processing. *The Southern Journal of Philosophy*, 53(S1):3–27, 2015.
- [14] A. Clark. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, 2015.
- [15] P. Dayan, G. E. Hinton, R. M. Neal, and R. S. Zemel. The Helmholtz Machine. *Neural computation*, 7(5):889–904, 1995.
- [16] K. Friston. A theory of cortical responses. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 360(1456):815–836, 2005.
- [17] K. Friston. Hierarchical models in the brain. *PLoS Computational Biology*, 4(11), 2008.
- [18] K. Friston. The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13(7):293–301, 2009.
- [19] K. Friston. The free-energy principle: a unified brain theory? *Nature reviews. Neuroscience*, 11(2):127–138, 2010.
- [20] K. Friston. What is optimal about motor control? *Neuron*, 72(3):488–498, 2011.
- [21] K. Friston. Life as we know it. *Journal of the Royal Society Interface*, 10(86):20130475, 2013.
- [22] K. Friston, J. Kilner, and L. Harrison. A free energy principle for the brain. *Journal of Physiology-Paris*, 100(1):70–87, 2006.
- [23] K. Friston, J. Mattout, and J. Kilner. Action understanding and active inference. *Biological Cybernetics*, 104(1-2):137–160, 2011.

- [24] K. Friston, F. Rigoli, D. Ognibene, C. Mathys, T. Fitzgerald, and G. Pezzulo. Active inference and epistemic value. *Cognitive neuroscience*, pages 1–28, mar 2015.
- [25] K. J. Friston, J. Daunizeau, and S. J. Kiebel. Reinforcement learning or active inference? *PLoS One*, 4(7):e6421, 2009.
- [26] K. J. Friston, J. Daunizeau, J. Kilner, and S. J. Kiebel. Action and behavior: A free-energy formulation. *Biological Cybernetics*, 102(3):227–260, 2010.
- [27] K. J. Friston, N. Trujillo-Barreto, and J. Daunizeau. DEM: A variational treatment of dynamic systems. *NeuroImage*, 41(3):849–885, 2008.
- [28] J. Hohwy. *The predictive mind*. OUP Oxford, 2013.
- [29] A. H. Jazwinski. *Stochastic Processes and Filtering Theory*, volume 64. Academic Press, 1970.
- [30] R. E. Kalman et al. Contributions to the theory of optimal control. *Bol. Soc. Mat. Mexicana*, 5(2):102–119, 1960.
- [31] D. C. Knill and A. Pouget. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719, 2004.
- [32] W. T. Powers. *Behavior: The control of perception*. Aldine Chicago, 1973.
- [33] R. P. Rao and D. H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.
- [34] H. Ritz, M. R. Nassar, M. J. Frank, and A. Shenhav. A control theoretic model of adaptive behavior in dynamic environments. *bioRxiv*, page 204271, 2017.
- [35] A. K. Seth. The Cybernetic Bayesian Brain. In *Open MIND*. Open MIND. Frankfurt am Main: MIND Group, 2014.
- [36] C. E. Shannon. Coding theorems for a discrete source with a fidelity criterion. *IRE Nat. Conv. Rec*, 4(142-163):1, 1959.
- [37] E. Todorov. General duality between optimal control and estimation. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 4286–4292. IEEE, 2008.
- [38] N. Wiener. *Cybernetics or Control and Communication in the Animal and the Machine*, volume 25. MIT press, 1961.
- [39] D. M. Wolpert and M. Kawato. Multiple paired forward and inverse models for motor control. *Neural networks*, 11(7-8):1317–1329, 1998.
- [40] L. Yang and P. A. Iglesias. Positive feedback may cause the biphasic response observed in the chemoattractant-induced response of dictyostelium cells. *Systems & control letters*, 55(4):329–337, 2006.
- [41] T.-M. Yi, Y. Huang, M. I. Simon, and J. Doyle. Robust perfect adaptation in bacterial chemotaxis through integral feedback control. *Proceedings of the National Academy of Sciences*, 97(9):4649–4653, 2000.