

A Bayesian perspective on classical control

Manuel Baltieri

Laboratory for Neural Computation and Adaptation, RIKEN Centre for Brain Science, Wako, Saitama, Japan
manuel.baltieri@riken.jp

Abstract—The connections between optimal control and Bayesian inference have long been recognised, with the field of stochastic (optimal) control combining these frameworks for the solution of partially observable control problems. In particular, for the linear case with quadratic functions and Gaussian noise, stochastic control has shown remarkable results in different fields, including robotics, reinforcement learning and neuroscience, especially thanks to the established *duality* of estimation and control processes. Following this idea we recently introduced a formulation of PID control, one of the most popular methods from classical control, based on active inference, a theory with roots in variational Bayesian methods, and applications in the biological and neural sciences. In this work, we highlight the advantages of our previous formulation and introduce new and more general ways to tackle some existing problems in current controller design procedures. In particular, we consider 1) a gradient-based tuning rule for the parameters (or gains) of a PID controller, 2) an implementation of multiple degrees of freedom for independent responses to different types of signals (e.g., two-degree-of-freedom PID), and 3) a novel time-domain formalisation of the performance-robustness trade-off in terms of tunable constraints (i.e., priors in a Bayesian model) of a single cost functional, variational free energy.

Index Terms—PID control, active inference, Bayesian inference, optimal control, optimal tuning, performance-robustness trade-off

I. INTRODUCTION

In the last few decades, the importance of probabilistic approaches to optimal control theory has been highlighted by different applications of Bayesian methods to problems of control. In his pioneering work, Bellman introduced Markov decision processes [1] as part of what is now known as stochastic optimal control [2]–[6]. This formulation captured the intrinsic probabilistic nature of problems of optimal control and decision making, with state transitions, outcomes and actions/decisions that cannot always be easily described in purely deterministic terms. Bellman’s approach extended on his own work on the dynamic programming method for (deterministic) optimal control, defining the Bellman equation, its uses and limitations, including the idea of the curse of dimensionality [7]. Shortly after, Kalman introduced the notions of observability and controllability of a system [8], with the former expressing the degree to which states can be estimated from noisy observations, and the latter representing the degree of control over a system when different manipulations are applied. Kalman also noticed that his filter was *dual* to the linear quadratic regulator (LQR), a now well known method in

(optimal) control theory that he also established [9], showing how both solutions require solving Riccati equations, forward in time for filtering (error covariance matrix) and backward in time for control (Hessian of the cost-to-go function). In the following years, several results improved the treatment of stochastic optimal control problems, including for instance the separation principle [4], [10]¹ and its applications to the treatment of regulation in the presence of uncertainty, i.e., for (linear) partially observable control problems. Due to its analytical tractability and its combination of estimation and control algorithms, the linear quadratic framework (i.e., stochastic optimal control for linear state-space models with Gaussian white noise and quadratic cost functions) has since then become a standard approach in different fields, including not only control theory and engineering [3], [4], but also robotics [13] and neuroscience [14], [15].

In recent years, the results based on the notion of duality in the linear case have been extended to (some classes of) nonlinear systems [16]–[18], highlighting further connections between control and estimation. Notably, these extended dualities often rely on more efficient variational approximations commonly used in problems of inference. For instance, relevant advances in stochastic optimal control and reinforcement learning have been driven by the use of methods commonly adopted to approximate intractable problems of Bayesian inference, e.g., variational Bayes. These methods have been shown to outperform standard dynamic programming and reinforcement learning algorithms for the control of different classes of problems [6], [16], [19], [20]. Building on these ideas, a similar approach has been proposed and adopted in neuroscience in an attempt to characterise brain function and sensorimotor control under a unifying probabilistic framework: active inference. While a full treatment of this framework is beyond the scope of the present work (for some technical reviews see, for instance, [21], [22]), we highlight how active inference combines methods from machine learning (variational Bayes), control theory (stochastic control) and statistical inference (hierarchical and empirical Bayes) to form a theory that includes several existing results from different fields as special cases, from predictive coding, to the infomax principle, to statistical models of learning, to risk-sensitive and KL-control [22]–[27].

Most of these results rely, at the moment, on the application of (approximate) Bayesian approaches to *optimal* control, with

MB is a JSPS International Research Fellow supported by a JSPS Grant-in-Aid for Scientific Research (No. 19F19809).

¹Also known in econometrics as certainty equivalence property [11], [12], but see [4] for a possible distinction.

almost no mention of *classical* control methods. While classical methods can be seen as a special case of optimal control, the possible advantages specific to Bayesian formulations of classical algorithms such as Proportional-Integral-Derivative (PID), remain largely unexplored. In this work we look at classical controllers from the perspective of approximate Bayesian inference, discussing the implications of variational Bayesian methods for the future of, in particular, PID control [28]. This perspective has previously been adopted in, for instance, [29] where a new gradient-based gain tuning rule was derived in closed-form for optimal regulation near the set-point/reference goal. In the next sections we present three cases in support of a new (Bayesian) framework to design and study classical control methods that ought to be seen as complementary to existing ones, e.g., optimal control and frequency-domain analysis. We will first, briefly, 1) recapitulate the previously derived results for gain tuning introducing new connections to path integral control and estimation in the presence of biases, then 2) present a more formal and in depth treatment of the connections between PID controllers with two degrees of freedom and Bayesian inference schemes, and finally 3) focus on the open challenge of framing different competing constraints of the performance-robustness trade-off in PID control, here defined in terms of priors and hyperpriors on a probabilistic generative model.

II. CASE 1: A BAYESIAN DERIVATION OF PID GAINS AND THEIR OPTIMISATION

PID controllers are the most popular choice for SISO systems regulation in different areas of engineering [30]. Their popularity is mainly due to their simplicity and relatively low number of tunable parameters. However, despite only including a few key parameters, or gains, their tuning (or optimisation) remains largely an open challenge [28], [31]. Existing tuning methods are often limited to specific cases or applications, relying on (ad-hoc) analytical rules, simple heuristics, frequency domain analysis, optimisation (including the use of artificial neural networks) or a combination of the above (for a survey, see [31]), that hardly generalise across different classes of problems. Here we report a more general method than can be explicitly derived by taking a different, Bayesian perspective on control problems.

Previous work relating classical control to optimal observers, and thus indirectly to Bayesian methods, showed that the integral component of PID controllers corresponds to a process of estimation of unknown (but linear/constant or step) perturbations, equivalent to a Kalman-Bucy filter with augmented state for the inference of unknown inputs (or biases) [32], [33]. Using the same approach, this connection was then generalised to higher order polynomial disturbances, equivalent to controllers including further integration terms [34], i.e., corresponding to PIID, PIIID, etc. controllers. In [29], we derived a fully probabilistic version of PID control, highlighting in particular some of the relationships between integral control and an emerging framework in computational and cognitive neuroscience, active inference [25], [26]. Using

active inference, we thus defined a more explicit generative model to describe an underlying stochastic process producing PID control as a gradient descent of a cost functional, variational free energy. While these two approaches, [34] and [29], share a number of features, they also present some core technical differences. Our proposal in fact includes:

- a more direct interpretation of the control matrix R , commonly used as a weight for the cost of control in the value (or cost-to-go) function [3], [4],
- a gradient-based algorithm to optimise R , and
- a generalisation to (some classes of) nonlinear problems.

As shown in [32], [34], the control matrix R is particularly relevant for the computation of the gains of PID controllers, here treated as part of the feedback matrix of a linear quadratic controller. In active inference, such gains correspond to specific hyperparametrisations of the linear state space (generative) model used to approximate the dynamics of the system to control, i.e., the (expected) precision, or inverse covariance, of the observation noise [29]. This result is closely related to Kalman's duality of inference and control [8], [18], [35], highlighting the mathematical correspondence between the processes of stochastic estimation and deterministic control. At the same time, the active inference formulation extends this duality beyond simply noting mathematical similarities, in order to include an account of the *dual role* of action in the context of exploration/exploitation problems [36], [37]. Furthermore, given the role of R as (expected) precision in the generative model, the gain parameters of PID controllers can be optimised via a gradient descent on the same cost functional, i.e., variational free energy [29], following a second order scheme introduced in [38] that, under some assumptions, holds also for some classes of nonlinear problems [24], [38].

III. CASE 2: PID CONTROL WITH 2DOF IN ACTIVE INFERENCE

In many applications of PID control, it is often desirable to build regulators that can respond to external disturbances while avoiding large fluctuations (e.g., overshooting) due to changes of the target of the regulation process. In standard PID control, these requirements are shown to be conflicting [39], [40] thus leading, in the most general case, to a multi-objective optimisation problem whose solutions lie on a Pareto front defined by

- changes in the control target (i.e., set-point response), and
- changes in the amplitude of a step disturbance (i.e., disturbance response).

To overcome the limitations induced by this trade-off, previous work (see [28], [40] and references therein) introduced the idea of controllers with two degrees of freedom, or 2DOF, PID. Multiple degrees of freedom obtained by augmenting controllers with multiple internal loops of PI or PID control (see also equivalent examples such as PI-PD control [28]), then ensure that different constraints can be treated independently, using parameters from different sub-loops to encode separate desired behaviours [28], [40].

Our probabilistic derivation of PID control via a variational approximation of Bayesian inference showcases a clear and direct interpretation of the presence of two degrees of freedom, here derived using rather general arguments. Unlike previous proposals, one need not augment a controller with an extra feedforward component that can separate the effects of a compensator for disturbances or set-point changes [40]. In active inference, the existence of two degrees of freedom is a simple consequence of the probabilistic (Bayesian) description of the generative model used to derive a controller [29]. This becomes more obvious after looking at the variational free energy (see equation (13) in [29]) here reported as²

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

where y, μ_x, μ_v are observations (or measurements), and expected hidden states (the estimate of the state of the system to regulate) and inputs (the set-point) respectively. Hyperparameters $\mu_{\pi_{\tilde{z}}}, \mu_{\pi_{\tilde{w}}}$ are the expected precisions on observation and system noise, and $f(), g()$ are state transition and observation functions. The tilde simply highlights a notation used to group different derivatives, or rather embeddings orders, of a variable, e.g., $\tilde{y} = \{y, y', y''\}$, see [29] for more details. The simplified (i.e., under Gaussian assumptions) variational free energy functional in (1) contains two sets of prediction errors, essentially instantiating two degrees of freedom for the controller. Notice that unlike equation (13) in [29], here we explicitly replaced $\pi_{\tilde{z}}, \pi_{\tilde{w}}$ with $\mu_{\pi_{\tilde{z}}}, \mu_{\pi_{\tilde{w}}}$ from the beginning, to highlight the fact that hyperparameters $\mu_{\pi_{\tilde{z}}}, \mu_{\pi_{\tilde{w}}}$ are only *estimates* of some “true” hyperparameters $\pi_{\tilde{z}}, \pi_{\tilde{w}}$. This follows from a full Bayesian treatment of the control problem, considering all variables to be *random* variables [41] (cf. traditional point-estimates in frequentist frameworks for statistical learning). In our case, to simplify the mathematical treatment, we treat them as Gaussian random variables with means $\mu_{\pi_{\tilde{z}}}, \mu_{\pi_{\tilde{w}}}$ (and covariances to be discussed in the next section). Importantly, these expectations are updated on a slower time scale [29], following schemes found in [23] and in particular [38], emphasising how parameters and hyperparameters of a generative model ought to be considered as fixed quantities over a certain (i.e., long) time scale.

The two sets of prediction errors, $\mu_{\pi_{\tilde{z}}} \left(\tilde{y} - g(\tilde{\mu}_x, \tilde{\mu}_v) \right)$ and $\mu_{\pi_{\tilde{w}}} \left(\tilde{\mu}'_x - f(\tilde{\mu}_x, \tilde{\mu}_v) \right)$ weighted by hyperparameters $\mu_{\pi_{\tilde{z}}}$ and $\mu_{\pi_{\tilde{w}}}$, represent likelihood and prior of a Bayesian update scheme formulated using generative models under Gaussian assumptions, the Laplace [42] and the variational Gaussian [43] approximations (to clarify their role see discussion in Chapter 3 of [27]). The update equations minimising these prediction errors [29] (also called recognition dynamics [44]), are similar to the update and prediction steps of standard

²Some terms in the free energy functional are hereby dropped for clarity. For the treatment of an extra set of terms important for the optimisation of PID gains, see [29]. For a more complete discussion of other terms which are constant during the optimisation phase, see [21], [23], [38].

algorithms from estimation theory such as Kalman-Bucy filters [45], and equivalent to feedback and feedforward loops in 2DOF PID controllers [40]. In this set up, PID control with a single degree of freedom can be derived as the limit case for fully observable states, i.e., $\tilde{y} = \tilde{x}$ (cf. state feedback methods in [30]). The independence of set-point and disturbance responses crucial for 2DOF PID controllers then corresponds, in this framework, to a generative model having system and measurement noise independent of one another, a standard assumption for linear state space models.

IV. CASE 3: THE PERFORMANCE-ROBUSTNESS TRADE-OFF FOR PID CONTROLLERS UNDER ACTIVE INFERENCE

The presence of conflicting criteria for the design of PID controller is a well known issue in the control theory literature [46], as partially highlighted in the previous section. This conflict is often referred to as the performance-robustness trade-off [28], [47], [48]. Controllers are usually designed to optimise some given performance criteria while, at the same time, attempting to maintain a certain level of robustness in face of uncertainty and unexpected conditions during the regulation process. The performance of a controller is normally assessed using one or more of the following criteria [28], [47]:

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

while robustness is mainly evaluated based on:

- *robustness to model uncertainty*, or how uncertainty on 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, unified and tractable framework that can be applied to a large class of compensation problems. An example in this direction is presented in [49] (see also [50]–[52] for other partial attempts). This methodology is based on the maximisation of the integral gain (equivalent, near the reference point, to the minimisation of the integral of the error from the set-point [39]), subject to constraints derived from a frequency domain analysis related to the Nyquist stability criterion applied to the controlled system [49]. Here, we propose our Bayesian formulation as an alternative (and in many cases complementary) framework for the design of PID controllers that leverages the straightforward interpretation of the performance-robustness trade-off for PID controllers in terms of uncertainty parameters (i.e., hyperparameters, precisions or inverse covariances) in the variational free energy [29]. To highlight its potential, we discuss the four standard criteria listed above as part of performance-robustness trade-off to address what can be gained using a Bayesian perspective.

A. Load disturbance response

A classic design principle for PID controllers is based on the response of a controller to perturbations that drive

a process away from the target value [39]. Random, zero-mean disturbances are commonly modelled as white Gaussian variables, and the parameters of the controller are simply tuned to reject such noise. Integral control then guarantees an appropriate response to step disturbances, equivalent to *non-zero*-mean noise (or to a bias term [33]), by accumulating and compensating for the ensuing steady-state error [32], [39], [53], [54]. The load disturbance response is usually expressed in terms of a minimisation of the Integral Absolute Error (IAE) between the state of the system to regulate and its target:

$$IAE = \int_0^{\infty} |e(t)| dt \quad (2)$$

or approximated by the Integral Error (IE) for non-oscillating or oscillating but well-damped systems [39]:

$$IE = \int_0^{\infty} e(t) dt \quad (3)$$

The IE criterion is especially relevant because it gives a straightforward intuition of the role of integral gain since, under a few simplifying assumptions (including a system's initial state close to the target value), the IE is equal to the inverse of k_i as $t \rightarrow \infty$ [39]. This implies that for large (theoretically, infinite) integral gains, the IE is minimised. While useful for its straightforward interpretation of this free parameter, practical and physical limitations often restrict the maximisation of the integral gain.

Our formulation builds on previous work showing how the use of integral control is optimal for unknown step perturbations applied to a system [32], [54]. In statistical terms, the presence of such disturbances can be formally seen as a *bias* term in an estimation process [55], showing how rejecting (step) perturbations is equivalent to estimating biases [33]. In active inference we can extend this (exact) result for linear systems and disturbances to nonlinear cases (not limited to polynomial perturbations as in [34]), where a more general (but often only approximate or limited to special classes of nonlinearities) duality of estimation and control is obtained using variational and path integral formulations [16], [18], or via probability integral transforms in the form of hierarchical generative models [24].

Furthermore, in our (Bayesian) formulation we gain a second and arguably deeper intuition on the role of the integral gain, which is now explicitly represented as one of the expected precisions (or inverse covariance) of observations \tilde{y} , i.e., μ_{π_z} , see [29]. This prescribes a simple and alternative way of understanding why the maximisation of k_i is usually a good heuristic for regulation problems where PID control is used [28]: maximising k_i is in fact equivalent to minimising uncertainty on measurements y , by maximising (minimising) the expected precision μ_{π_z} (variance μ_{σ_z}) of the measurements of the system to regulate. At the same time, this can also explain some of the limitations of this heuristic, discussed in the frequency domain for instance in [49]. The maximisation

of k_i , without any constraints, corresponds to the minimisation of the expected measurement variance μ_{σ_z} , such that $t \rightarrow \infty$,

$$\mu_{\sigma_z} \rightarrow 0. \quad (4)$$

In practice, however, one should always consider a certain level of intrinsic, i.e., aleatoric, uncertainty whose variance is fundamentally irreducible. Even an optimal controller can't overcome the limited sensitivity of a sensor (here represented by the "real" σ_z , as opposed to its estimate μ_{σ_z}), bringing μ_{σ_z} down to 0 is thus not possible if $\sigma_z > 0$. In other proposals [49], the same aleatoric uncertainty σ_z is effectively approximated with a measure that captures the levels of controllability of a system through the definition of appropriate sensitivity functions in the frequency domain.

Our Bayesian implementation also extends the intuition behind the integral gain as a precision of observations to the other two gains, k_p and k_d . In our formulation, these gains become in fact estimated precisions of higher embedding orders of the observations, y' , y'' , often also called *generalised coordinates of motion* [23], [38]. These embedding orders essentially represent a Taylor expansion (in time) of continuous random variables defined according to a Stratonovich (rather than Ito) interpretation, equivalent to non-Markovian (semi-Markovian, of Markovian of order n) stochastic processes [24], [29], [45]. In practice, for measurements taken at a high enough frequency, and with controllers having a short enough intrinsic time scale to regulate such high frequency measurements (i.e., a time scale approaching the underlying continuous models of the systems to regulate), observation noise should be treated as coloured, rather than white as in standard delta-autocorrelated noise. Under these assumptions, the implementation of PID control and its extensions (e.g., multiple I and D terms) becomes simply a linear approximation of a measured non-Markovian trajectory. Perhaps in a more intuitive way, we can see expected precisions $\mu_{\pi_z} = \{\mu_{\pi_z}, \mu_{\pi_{z'}}, \mu_{\pi_{z''}}\}$ as simultaneously 1) representing the precision of a trajectory in the state-space (rather than the precision on a point) and 2) regulating the convergence rate of measurements to a *set-trajectory* (rather than point), specifying how quickly a controller ought to respond to a sudden change in a set of observations and their higher orders of motion.

B. Set-point response

Following the load disturbance rejection property, a second performance criterion used for the design of PID controllers is their set-point response, i.e., how controllers respond to variations in the set value used as a target to regulate a system. Naively, this could be seen as closely related to load disturbances: rather than changes in the measurement, we now have changes in the target value, both of them used to define some error term, e . In practice however, it is desirable to decouple these two problems, creating a controller with different sensitivities to load disturbances or set-point updates whenever necessary [39]. This requires a controller with two degrees of freedom, as discussed in more detail in section III, which is an inherent feature of the active inference formulation

where expectations of hidden states $\tilde{\mu}_x$ are updated using a (Bayesian) scheme that balances (via a set of independent expected precisions, or weights) prediction errors on both

- observations, $\mu_{\pi_{\tilde{y}}}(\tilde{y} - g(\tilde{\mu}_x, \tilde{\mu}_v))$, where load disturbances can appear as part of the measurements \tilde{y} , and
- system dynamics, $\mu_{\pi_{\tilde{w}}}(\tilde{\mu}'_x - f(\tilde{\mu}_x, \tilde{\mu}_v))$, where set-trajectories can be updated as inputs/priors $\tilde{\mu}_v$.

Mirroring the role of $\mu_{\pi_{\tilde{z}}}$ for load disturbances, expected precisions $\mu_{\pi_{\tilde{w}}}$ on dynamic prediction errors effectively implement a response mechanism to set-trajectories updates, with high expected precisions implying a fast response, and low precisions entailing a slow one. Equivalently, from a probabilistic perspective this can be explained with the idea that the former describes a model with low uncertainty on dynamics (high precision = low covariance) meaning that any variation from such dynamics should quickly be dealt with; the latter encodes, on the other hand, the fact that high expected covariance allows for changes in set-trajectories, i.e., sudden updates are not surprising, therefore changes can be slow (and in the limit for very large covariances, almost absent).

C. Measurement noise response

A third common requirement for PID controllers is related to their performance in face of noisy or uncertain measurements. These may be due, for instance, to physical constraints/sensitivities of the sensors. In the literature, high frequency measurement noise [30] is usually tackled via a careful and ad-hoc controller design, including for example pre-filtering of the observed data [47]. In the Bayesian formulation of PID controllers that we introduced, we have a direct measure of (the best estimate of) the measurement noise: the expected precision or inverse covariance $\mu_{\pi_{\tilde{z}}}$ of the random variable z and its higher orders of motion. Measurement noise is thus related to the same set of hyperparameters used to explain load disturbances rejection which, on the other hand, can be seen as low frequency noise. This shows another trade-off between design criteria, in this case related to the high frequency properties of measurement noise and the (usually) low frequency of external disturbances.

The previously identified maximisation of expected precision $\mu_{\pi_{\tilde{z}}}$ (integral gain k_i) implies an increased cutoff frequency of the low-pass filter implemented by linear generative models of the kind we introduced to approximate PID control [56]. This suggests that, while low-frequency disturbances can be suppressed more quickly (even if at the cost of possibly overshooting), this comes at the expenses of a “hypersensitivity” to high-frequency noise, i.e., not rejecting as much noise as otherwise possible with slower load disturbance responses (as shown in a simple model, for instance, in [56]). In our framework, this can easily be noticed by looking at the role played by expected precision $\mu_{\pi_{\tilde{z}}}$ in the time domain, encoding expected variability in observed data without a clear distinction between rare perturbations and persistent noise.

At the same time, however, the active inference formulation can be used to treat this problem in a more principled way,

introducing informative priors on expected precisions $\mu_{\pi_{\tilde{z}}}$, i.e., hyperpriors $\eta_{\pi_{\tilde{z}}}$ ³ (or more complicated functions $h(\eta_{\pi_{\tilde{z}}})$), see [24] for a formal treatment. The variational free energy functional then includes another set of prediction errors, (cf. (1)),

$$F \approx \frac{1}{2} \left[\mu_{\pi_{\tilde{z}}} (\tilde{y} - g(\tilde{\mu}_x, \tilde{\mu}_v))^2 + \mu_{\pi_{\tilde{w}}} (\tilde{\mu}'_x - f(\tilde{\mu}_x, \tilde{\mu}_v))^2 + \mu_{p_{\tilde{z}}} (\mu_{\pi_{\tilde{z}}} - h(\eta_{\pi_{\tilde{z}}}))^2 \right] \quad (5)$$

with $\mu_{p_{\tilde{z}}} (\mu_{\pi_{\tilde{z}}} - h(\eta_{\pi_{\tilde{z}}}))$ playing the role of L2 (or Tikhonov) regularisation terms for the ensuing recognition dynamics derived as a gradient descent on (5). Using these prediction errors one can effectively encode, for instance, constraints that reject strong high frequency noise by specifically targeting frequent large instantaneous fluctuations of expected precision $\mu_{\pi_{\tilde{z}}}$, penalising them with a Gaussian hyperprior (an L2 regularisation term that affects “outliers”) centred at $\eta_{\pi_{\tilde{z}}}$. While such hyperprior would certainly then also influence the response to sparse step changes, expected precisions $\mu_{\pi_{\tilde{z}}}$ could be updated by slowly shifting hyperpriors $\eta_{\pi_{\tilde{z}}}$ to reflect biases in measurements \tilde{y} that persist over a long period of time. Importantly, while the cost functional presents in this case some new terms, the underlying minimisation scheme remains the same: the recognition dynamics will simply include extra regularisation terms while still following a gradient descent on the (augmented) variational free energy.

In the same way expected precisions $\mu_{\pi_{\tilde{z}}}$ regulate the response to changes in the observations due to load disturbances, expected precisions on higher order stochastic properties (e.g., expected precisions on expected precisions, $\mu_{p_{\tilde{z}}}$) can then be seen as regulating how a controller adapts to varying levels of measurement noise covariance given some (informative) priors $h(\eta_{\pi_{\tilde{z}}})$. For example, in cases where the variance of measurement noise changes over time, e.g., due to the natural degradation of sensors, our formulation can include mechanisms that take into account existing prior knowledge and that can be used by a controller to dynamically adapt to new levels of noise. More in general, in the presence of *stochastic volatility* (i.e., models where the covariances of different random variables are themselves random variables [41]), one can easily encode prior knowledge of higher order properties of random variables by including extra hierarchical layers on the generative model we introduced for PID control.

D. Robustness to model uncertainty

PID controllers are usually designed to withstand some level of model uncertainty, inherent in any system we observe, interact with and try to regulate. In control theory, this problem affects compensators attempting to regulate a system while having access only to a limited amount of information regarding the dynamics of the system itself. PID controllers are especially popular as a “model free” strategy, or rather, for the

³To maintain a notation similar to the one used in [24], [29].

small number of tunable parameters that are necessary to afford robust, although often suboptimal, control [46]. In control problems, this robustness is sometimes captured by sensitivity functions [47], [49], providing a proxy for, among other things, the sensitivity of a feedback system to variations in models of process dynamics. In our derivation of PID control as a process of Bayesian (active) inference, the uncertainty of the dynamics is represented by the expected precisions of system dynamics, $\mu_{\pi_{\tilde{w}}}$, in the linear generative model defined in [29]. For instance, low expected precisions $\mu_{\pi_{\tilde{w}}}$, expressing high uncertainty/covariance, encode the (Bayesian) belief that large fluctuations in the dynamics can be expected, while high expected precisions express the fact that dynamics should show only small fluctuations. Moreover, using our formulation we can describe the behaviour of a PID controller such that under controllability assumptions [4], [8], it effectively “imposes” its own (linear) dynamics/priors on a system through larger weighted prediction errors $\mu_{\pi_{\tilde{w}}}(\tilde{\mu}'_x - f(\tilde{\mu}_x, \tilde{\mu}_v))$, by forcing it into an attractor encoded in the set-trajectory represented by the controller’s priors. The state of affairs of the world is only partially relevant to a PID controller, since as long as conditions of reachability and controllability [4] are met, all it does is try to drive a (controllable) system towards the desired equilibrium encoded by its priors on a set-trajectory.

As in the case of measurement noise, our formulation allows for the construction of an extra layer of hyperpriors to handle model uncertainty: in the active inference formulation we can in fact include priors on expected precisions $\mu_{\pi_{\tilde{w}}}$ to represent existing information on the expected/desired dynamics of a system to regulate

$$F \approx \frac{1}{2} \left[\mu_{\pi_z} \left(\tilde{y} - g(\tilde{\mu}_x, \tilde{\mu}_v) \right)^2 + \mu_{\pi_{\tilde{w}}} \left(\tilde{\mu}'_x - f(\tilde{\mu}_x, \tilde{\mu}_v) \right)^2 + \mu_{p_{\tilde{w}}} \left(\mu_{\pi_{\tilde{w}}} - k(\eta_{\pi_{\tilde{w}}}) \right)^2 \right] \quad (6)$$

For instance, it is not hard to imagine that, following standard hierarchical or empirical Bayes methods in statistical inference [41], information on existing control problems could be used to define classes of systems whose shared statistical properties form generic priors $\eta_{\pi_{\tilde{w}}}$. These priors could then be used to initialise our model in a suitable part of the state space to ensure a desired level of robustness (and if a similar approach were to be adopted for η_{π_z} , to guarantee some desired performances). In such settings, expected precisions $\mu_{\pi_{\tilde{w}}}$ can still be optimised via a simple gradient descent, now L2-regularised with newly introduced priors entering the variational free energy equation in the form of weighted prediction errors, $\mu_{p_{\tilde{w}}}(\mu_{\pi_{\tilde{w}}} - k(\eta_{\pi_{\tilde{w}}}))$. This approach, especially when employing empirical Bayes, is similar in spirit to the clever initialisation achieved in deep learning approaches via “pre-training”, where introducing an unsupervised learning phase before supervised training showed substantial improvements in the performance and generalisation properties of neural networks [57], [58].

V. DISCUSSION

The duality of estimation and control has long been recognised and exploited in problems of regulation under constraints of partial observability, i.e., stochastic control [3], [4], [10], [16], [18], [35]. This property relies on the mathematical equivalence of some classes of estimation and regulation problems, formulated as Bayesian inference and optimal control respectively. The applications of this duality have led to a series of significant new results in different areas, such as reinforcement learning [20], robotics [59] and neuroscience [14] where methods of *approximate* Bayesian inference are now often employed to improve existing solutions. In this work we built on some of these previous ideas, discussing possible applications of Bayesian inference theories and related approximations to methods from classical control. In particular, we focused on PID control and on our previous implementation of this method in terms of Bayesian *active* inference [29], proposing this as a general unifying framework for the design of PID controllers still largely missing to date [28], [31].

In [29] we recently introduced a gradient-based procedure for gain tuning, using an interpretation of these parameters as stochastic properties (i.e., expected precisions, or inverse variances) of the system to regulate. Here we expanded on this formulation by providing direct links to Kalman’s duality [8], [18], [35] and Bayesian estimation in the presence of bias terms, i.e., unknown step inputs [33].

We then discussed standard problems such as the necessity of two degrees of freedom in order to afford independent responses to load and set-point changes [40]. Using the probabilistic interpretation given in [29], we then drew a comparison between a pragmatic introduction of two degrees of freedom [40], represented by feedback and feedforward sub-loops in standard 2DOF PID control, and the more principled formulation of active inference, aligned with update and prediction equations of filtering algorithms (e.g., Kalman-filters [9]), and the use of prior and likelihood densities in recursive Bayesian update schemes [45].

Crucially, we then proposed to frame one of the major open challenges for methods like PID control, the general performance-robustness trade-off due to the presence of conflicting design criteria [28], [31], in terms of variational free energy minimisation [17], [23], [29], [38], [60]⁴

$$F \approx \frac{1}{2} \left[\mu_{\pi_z} \left(\tilde{y} - g(\tilde{\mu}_x, \tilde{\mu}_v) \right)^2 + \mu_{\pi_{\tilde{w}}} \left(\tilde{\mu}'_x - f(\tilde{\mu}_x, \tilde{\mu}_v) \right)^2 + \mu_{p_z} \left(\mu_{\pi_z} - h(\eta_{\pi_z}) \right)^2 + \mu_{p_{\tilde{w}}} \left(\mu_{\pi_{\tilde{w}}} - k(\eta_{\pi_{\tilde{w}}}) \right)^2 \right] \quad (7)$$

In this formulation, simple constraints (load disturbance response and set-point change response) can easily be mapped to first order weighting parameters on the mean estimates of the state of the system to regulate. More complex ones, on the other hand, (measurement noise response and robustness to model uncertainty) can be introduced in terms of stochastic volatility [41], i.e., by treating second moments (expected

⁴The following equation combines (5) and (6).

precisions, or hyperparameters) as random variable having appropriate hyperpriors encoded in the generative model. This mapping provides an immediate understanding of different desired statistical properties of the system to govern (see table I), now summarised in Table I.

TABLE I: Active inference as a general framework for PID controllers design (adapted from [29] and here extended).

Criterion	Mapped to	Interpretation in Active Inference
Load disturbance response	$\mu\pi_z$	Expected inverse covariance of the observations (i.e., precision), with low covariance implying a fast response, and vice versa
Set-point change response	$\mu\pi_{\bar{w}}$	Two degrees of freedom derived from the presence of two sets of prediction errors, sensory and dynamics, mapping to likelihood and priors of a Bayesian inference process
Measurement noise response	(priors on) $\mu\pi_z$	Direct mapping of measurement noise to inverse covariance of the observations (i.e., precision), with hyperpriors (priors on expected precisions) introduced to differentiate high frequency noise from low frequency disturbances
Robustness to model uncertainty	(priors on) $\mu\pi_{\bar{w}}$	Direct mapping of model uncertainty to expected covariances of system fluctuations, representing the hidden dynamics of the system to control, with hyperpriors that can describe initial knowledge of, for example, a class of similar regulation problems to facilitate the optimisation of states/parameters (similar to the role of unsupervised “pre-training” in deep learning [58])

VI. CONCLUSION AND FUTURE WORK

In an influential paper, Åström and Hägglund asked whether PID control can play a role in the future of control theory and engineering [28]. Despite being the most used controller in industry, the emergence of more specialised and better performing methods over the years, such as model predictive control, cast doubts on its long term applications and uses. Åström and Hägglund however argued that due to its combined effectiveness and simplicity, PID is likely to remain relevant for the foreseeable future, perhaps in conjunction with other methods. At the same time, they highlighted a series of existing problems and open challenges faced by PID, including a relatively limited number of theoretical results in areas such as gain tuning and general (PID) controller design. In this work we built on our previous formulation of PID control in terms of active inference, a modern theory combining stochastic control and probabilistic Bayesian inference under the umbrella of variational free energy minimisation, to propose new applications of Bayesian methods to PID controllers in order to establish a more general design framework. After introducing a new practical implementation of optimal gain tuning in [29], here we extended our proposal highlighting the connections between different design principles for PID, from the importance of multiple degrees of freedom to optimal tuning with conflicting performance-robustness criteria. This framework gives an interpretation of a series of different constraints as

first and second order properties of a generative model that generates a PID controller as a gradient-based minimisation of a single cost functional, variational free energy. In the future, we will focus on simulations testing the current proposal using standard control benchmarks and following a vast literature on Bayesian models (see [24], [41] and references therein). We will then also draw more direct connections to modern machine and reinforcement learning, combining the present work with methods from [61], where preliminary results based on these and other ideas are utilised in the field of deep reinforcement learning with large neural networks performing amortised inference.

REFERENCES

- [1] R. E. Bellman, “A markovian decision process,” *Journal of Mathematics and Mechanics*, pp. 679–684, 1957.
- [2] K. J. Åström, *Introduction to stochastic control theory*. Academic Press, 1970.
- [3] B. Anderson and J. B. Moore, *Optimal control: linear quadratic methods*. Prentice-Hall, Inc., 1990.
- [4] R. F. Stengel, *Optimal control and estimation*. Courier Corporation, 1994.
- [5] E. Todorov, “Optimal control theory,” *Bayesian brain: probabilistic approaches to neural coding*, pp. 269–298, 2006.
- [6] H. J. Kappen, *Optimal control theory and the linear Bellman equation*. Cambridge University Press, 2011, p. 363–387.
- [7] R. E. Bellman, *Dynamic Programming*. Courier Dover Publications, 1957.
- [8] R. E. Kalman, “Contributions to the theory of optimal control,” *Bol. Soc. Mat. Mexicana*, vol. 5, no. 2, pp. 102–119, 1960.
- [9] —, “A new approach to linear filtering and prediction problems,” *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [10] W. M. Wonham, “On the separation theorem of stochastic control,” *SIAM Journal on Control*, vol. 6, no. 2, pp. 312–326, 1968.
- [11] H. A. Simon, “Dynamic programming under uncertainty with a quadratic criterion function,” *Econometrica, Journal of the Econometric Society*, pp. 74–81, 1956.
- [12] H. Theil, “A note on certainty equivalence in dynamic planning,” *Econometrica: Journal of the Econometric Society*, pp. 346–349, 1957.
- [13] F. L. Lewis, D. M. Dawson, and C. T. Abdallah, *Robot manipulator control: theory and practice*. CRC Press, 2003.
- [14] E. Todorov and M. I. Jordan, “Optimal feedback control as a theory of motor coordination,” *Nature neuroscience*, vol. 5, no. 11, p. 1226, 2002.
- [15] E. Todorov, “Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system,” *Neural computation*, vol. 17, no. 5, pp. 1084–1108, 2005.
- [16] S. K. Mitter and N. J. Newton, “A variational approach to nonlinear estimation,” *SIAM journal on control and optimization*, vol. 42, no. 5, pp. 1813–1833, 2003.
- [17] H. J. Kappen, “Linear theory for control of nonlinear stochastic systems,” *Physical review letters*, vol. 95, no. 20, p. 200201, 2005.
- [18] E. Todorov, “General duality between optimal control and estimation,” in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*. IEEE, 2008, pp. 4286–4292.
- [19] H. Attias, “Planning by probabilistic inference.” in *AISTATS*. Citeseer, 2003.
- [20] E. Todorov, “Efficient computation of optimal actions,” *Proceedings of the national academy of sciences*, vol. 106, no. 28, pp. 11 478–11 483, 2009.
- [21] 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*, vol. 14, pp. 55–79, 2017.
- [22] L. Da Costa, T. Parr, N. Sajid, S. Veselic, V. Neacsu, and K. J. Friston, “Active inference on discrete state-spaces: a synthesis,” *arXiv preprint arXiv:2001.07203*, 2020.
- [23] K. J. Friston, N. Trujillo-Barreto, and J. Daunizeau, “DEM: A variational treatment of dynamic systems,” *NeuroImage*, vol. 41, no. 3, pp. 849–885, 2008.
- [24] K. J. Friston, “Hierarchical models in the brain,” *PLoS Computational Biology*, vol. 4, no. 11, 2008.

- [25] —, “The free-energy principle: a unified brain theory?” *Nature reviews. Neuroscience*, vol. 11, no. 2, pp. 127–138, 2010.
- [26] K. J. Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, and G. Pezzulo, “Active inference: a process theory,” *Neural Computation*, vol. 29, no. 1, pp. 1–49, 2017.
- [27] M. Baltieri, “Active inference: building a new bridge between control theory and embodied cognitive science,” Ph.D. dissertation, University of Sussex, 2019.
- [28] K. J. Åström and T. Häggglund, “The future of PID control,” *Control engineering practice*, vol. 9, no. 11, pp. 1163–1175, 2001.
- [29] M. Baltieri and C. L. Buckley, “PID control as a process of active inference with linear generative models,” *Entropy*, vol. 21, no. 3, p. 257, 2019.
- [30] K. J. Åström and R. M. Murray, *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.
- [31] K. H. Ang, G. Chong, and Y. Li, “PID control system analysis, design, and technology,” *IEEE transactions on control systems technology*, vol. 13, no. 4, pp. 559–576, 2005.
- [32] C. D. Johnson, “Optimal control of the linear regulator with constant disturbances,” *IEEE Transactions on Automatic Control*, vol. 13, no. 4, pp. 416–421, 1968.
- [33] —, “On observers for systems with unknown and inaccessible inputs,” *International journal of control*, vol. 21, no. 5, pp. 825–831, 1975.
- [34] —, “Further study of the linear regulator with disturbances – The case of vector disturbances satisfying a linear differential equation,” *IEEE Transactions on Automatic Control*, vol. 15, no. 2, pp. 222–228, 1970.
- [35] R. E. Kalman, “On the general theory of control systems,” in *Proceedings First International Conference on Automatic Control, Moscow, USSR*, 1960.
- [36] Y. Bar-Shalom and E. Tse, “Dual effect, certainty equivalence, and separation in stochastic control,” *IEEE Transactions on Automatic Control*, vol. 19, no. 5, pp. 494–500, 1974.
- [37] M. Baltieri and C. L. Buckley, “On Kalman-Bucy filters, linear quadratic control and active inference,” *arXiv preprint arXiv:2005.06269*, 2020.
- [38] K. J. Friston, K. Stephan, B. Li, and J. Daunizeau, “Generalised filtering,” *Mathematical Problems in Engineering*, vol. 2010, 2010.
- [39] K. J. Åström, *PID controllers: theory, design and tuning*. Research Triangle Park, 1995.
- [40] M. Araki and H. Taguchi, “Two-degree-of-freedom PID controllers,” *International Journal of Control, Automation, and Systems*, vol. 1, no. 4, pp. 401–411, 2003.
- [41] C. Robert, *The Bayesian choice: from decision-theoretic foundations to computational implementation*. Springer Science & Business Media, 2007.
- [42] D. J. MacKay, *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- [43] M. Opper and C. Archambeau, “The variational Gaussian approximation revisited,” *Neural computation*, vol. 21, no. 3, pp. 786–792, 2009.
- [44] C. S. Kim, “Recognition dynamics in the brain under the free-energy principle,” *Neural Computation*, vol. 30, no. 10, pp. 2616–2659, 2018.
- [45] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. Academic Press, 1970, vol. 64.
- [46] D. E. Rivera, M. Morari, and S. Skogestad, “Internal model control: PID controller design,” *Industrial & engineering chemistry process design and development*, vol. 25, no. 1, pp. 252–265, 1986.
- [47] K. J. Åström and T. Häggglund, *Advanced PID control*. ISA-The Instrumentation, Systems, and Automation Society Research Triangle Park, NC, 2006.
- [48] O. Garpinger, T. Häggglund, and K. J. Åström, “Performance and robustness trade-offs in PID control,” *Journal of Process Control*, vol. 24, no. 5, pp. 568–577, 2014.
- [49] K. J. Åström, H. Panagopoulos, and T. Häggglund, “Design of PI controllers based on non-convex optimization,” *Automatica*, vol. 34, no. 5, pp. 585–601, 1998.
- [50] M. Zhuang and D. Atherton, “Automatic tuning of optimum PID controllers,” in *IEE Proceedings D (Control Theory and Applications)*, vol. 140, no. 3. IET, 1993, pp. 216–224.
- [51] M. Grimble and M. Johnson, “Algorithm for PID controller tuning using LQG cost minimization,” in *Proceedings of the 1999 American Control Conference*, vol. 6. IEEE, 1999, pp. 4368–4372.
- [52] R. T. O’Brien and J. M. Howe, “Optimal PID controller design using standard optimal control techniques,” in *Proceedings of the 2008 American Control Conference*. IEEE, 2008, pp. 4733–4738.
- [53] B. A. Francis and W. M. Wonham, “The internal model principle of control theory,” *Automatica*, vol. 12, no. 5, pp. 457–465, 1976.
- [54] E. D. Sontag, “Adaptation and regulation with signal detection implies internal model,” *Systems & control letters*, vol. 50, no. 2, pp. 119–126, 2003.
- [55] B. Friedland, “Treatment of bias in recursive filtering,” *IEEE Transactions on Automatic Control*, vol. 14, no. 4, pp. 359–367, 1969.
- [56] B. W. Andrews, T.-M. Yi, and P. A. Iglesias, “Optimal noise filtering in the chemotactic response of *Escherichia coli*,” *PLoS Comput Biol*, vol. 2, no. 11, p. e154, 2006.
- [57] I. Harvey and J. V. Stone, “Unicycling helps your French: Spontaneous recovery of associations by learning unrelated tasks,” *Neural Computation*, vol. 8, no. 4, pp. 697–704, 1996.
- [58] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [59] E. A. Theodorou, J. Buchli, and S. Schaal, “A generalized path integral control approach to reinforcement learning,” *Journal of machine learning research*, vol. 11, no. Nov, pp. 3137–3181, 2010.
- [60] E. A. Theodorou and E. Todorov, “Relative entropy and free energy dualities: Connections to path integral and kl control,” in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*. IEEE, 2012, pp. 1466–1473.
- [61] A. Tschantz, M. Baltieri, A. Seth, C. L. Buckley *et al.*, “Scaling active inference,” *arXiv preprint arXiv:1911.10601*, 2019.