

Active Inference: Computational Models of Motor Control without Efference Copy

Manuel Baltieri (manuel.baltieri@gmail.com)

EASY group – Sussex Neuroscience, Department of Informatics
University of Sussex, Brighton, UK

Christopher L. Buckley (c.l.buckley@sussex.ac.uk)

EASY group – Sussex Neuroscience, Department of Informatics
University of Sussex, Brighton, UK

Abstract

Computational frameworks for the study of motor systems in neuroscience often rely on a mathematical formulation based on optimal control theory, e.g., forward and inverse models and linear quadratic Gaussian (LQG) control architectures. A forward model maps actions to (predicted) consequences, while an inverse model is thought to define how motor commands are generated from observations. One of the central tenets of the forward/inverse architecture is the presence of a copy of motor commands produced by an inverse model and provided to the forward counterpart. Such copy, usually referred to as “efference copy”, is assumed to be necessary to model, and ultimately explain, motor control and behaviour. Over the years, different results have challenged the idea of an efference copy, suggesting that it may not be physiologically plausible, especially in humans. In this work we focus on a process theory that combines the mathematical richness of LQG models with efference-copy-free architectures, active inference. We provide a minimal computational model discussing and comparing the forward/inverse and the active inference architectures on an idealised model of a single-joint control system.

Keywords: efference copy; LQG; Kalman filter; forward/inverse models; active inference; free energy principle

Introduction

In the last few decades, computational approaches of motor control have emerged as a prominent tool for the study of behaviour in the biological sciences (Jordan, 1996; Kawato, 1999; Wolpert & Ghahramani, 2000; Todorov, 2004; McNamee & Wolpert, 2019). These frameworks are often based on mathematical formulations of optimal control (Stengel, 1994), making use of forward/inverse and linear quadratic Gaussian (LQG) control architectures popular in engineering and robotics (Kawato, 1999; Todorov & Jordan, 2002; McNamee & Wolpert, 2019). In these architectures, perceptual processes can be seen as the transduction of sensory input to some internal (e.g., neural) representation and are often depicted as estimators (Todorov, 2004) or forward models-estimators (McNamee & Wolpert, 2019). These representations produce then (motor) actions via an inverse model (Kawato, 1999) or controller (Todorov, 2004). Crucial to these architectures is the presence of an “efference copy” (Kawato,

1999; Todorov, 2004; McNamee & Wolpert, 2019), representing information regarding an agent’s own actions to be discounted from one’s estimations of sensory inputs. The notion of an efference copy has, however, been comprehensively challenged (Feldman, 2009; Friston, 2011; Adams, Shipp, & Friston, 2013; Feldman, 2016). At the same time, the lack of alternative, powerful mathematical frameworks, made of LQG the dominant model for the study of behaviour and motor control (see Buhrmann and Di Paolo (2014) for a counterexample). An alternative approach maintaining strong connections to Bayesian inference and optimal control theory while disposing with the need for a copy of motor signals is proposed with active inference. In active inference the necessity for an efference copy is bypassed using a more powerful forward, or generative model, and trivial sensorimotor mappings in the form of reflex arcs replacing complex inverse models/controllers, similar to ideas of threshold or referent control (Feldman, 2009). To investigate the potential of this proposal, here we provide a simple mechanistic comparison between forward/inverse-LQG architectures and active inference.

Linear Quadratic Gaussian (LQG) control

In LQG schemes, an estimator (usually a Kalman or Kalman-Bucy filter) and a controller (usually a linear quadratic regulator, LQR) are coupled in a feedback loop and exchange information in two ways, see for instance Wolpert and Ghahramani (2000); Todorov (2004). The estimator generates accurate estimates of latent variables from observations and relays them to the controller, which in turn produces a motor command and sends a copy of it back to the estimator. This copy is crucial to allow the estimator, used as a metaphor for perceptual systems, to discount sensory consequences of an agent’s own motor actions. In the absence of this information, estimates of world variables quickly become imprecise and subsequently motor actions become unstable (Friston, 2011). The notion of a copy of motor signals resonates with the classical idea of efference copy in neuroscience (Crapse & Sommer, 2008; Schwartz, 2016; Straka, Simmers, & Chagnaud, 2018). Efference copy is thought to represent a copy of signals from low-level motor areas that is sent to perceptual processing areas in order to disambiguate movements performed by an agent from environmental stimuli, although its definition is often vague and mixed with the idea of corollary discharge (Crapse & Sommer, 2008; Schwartz, 2016; Straka et al., 2018). In the most prominent examples of LQG-based



architectures in the cognitive sciences, efference copy is necessary for appropriate estimations of hidden variables in the world (Kawato, 1999; Wolpert & Ghahramani, 2000; Todorov, 2004; McNamee & Wolpert, 2019). The neurophysiological evidence supporting efference copy is, however, conflicting (Feldman, 2009; Adams et al., 2013; Feldman, 2016), with alternative models that eschew this idea proposed by, for instance, Friston (2011); Adams et al. (2013).

Active inference

Active inference is a process theory proposed to explain brain functioning and other functions of living systems based on Bayesian inference and optimal control theory (Friston, 2010; Buckley, Kim, McGregor, & Seth, 2017). It is an algorithmic implementation of the free energy principle (Friston, 2010), proposing the minimisation of variational free energy, or under simplifying assumptions prediction error, as a driving principle for the study of cognitive processes. The main difference with respect to LQG architectures is that LQG-based models explicitly mirror (by construction, in the linear case) the dynamics of the observed system, thus including knowledge of one's motor actions (i.e., an efference copy), \mathbf{a} from now on, while in active inference this vector is not explicitly modelled by an agent, assuming that no copy of motor signals is available. It is in fact proposed that a deeper duality of estimation and control exists whereby, at the lowest level (i.e., a purely reflexive account of simple motor tasks), actions are just responses to the presence of prediction errors at the proprioceptive level, irrespectively of the cause of sensations, self-generated or external (Friston, 2011; Adams et al., 2013). In recent accounts of more complex behaviour under active inference, action is cast as a problem of estimating (fictitious) control states \mathbf{u} or rather time-dependent policies $\pi_{\mathbf{u}}$ that are inferred via the minimisation of *expected* free energy \mathbf{y} (Friston et al., 2015). Both these proposals support theories in motor neuroscience suggesting that knowledge of such self-produced controls is not available, and not necessary in fact, for motor control in biological systems (Feldman, 2009; Friston, 2011; Adams et al., 2013; Feldman, 2016). In active inference, to replace actions \mathbf{a} in the generative model, a vector \mathbf{v} is introduced that encodes prior beliefs (i.e., desired outcomes) about movement trajectories in the external world (Friston, 2011). See Baltieri and Buckley (2018, 2019a) for a discussion on some other differences implied by the lack of efference copy.

The model

The double integrator is a canonical example used in control theory (Åström & Murray, 2010), modelling single degree of freedom systems, representing a block on frictionless surface or, equivalently, the simplest model of single-joint movement in motor neuroscience (Gottlieb, 1993). In this set up, a limb segment (or a block) moves to reach a new position (for simplicity, $x = 0$) and stop (velocity $x' = 0$). Unlike the more traditional deterministic set up, we will introduce process and measurement noise into the system, making the estimation of hidden states necessary. The equations and parameters for this

simulation follow a standard LQG implementation and can be found in Baltieri and Buckley (2019a)¹.

The LQG solution to the double integrator

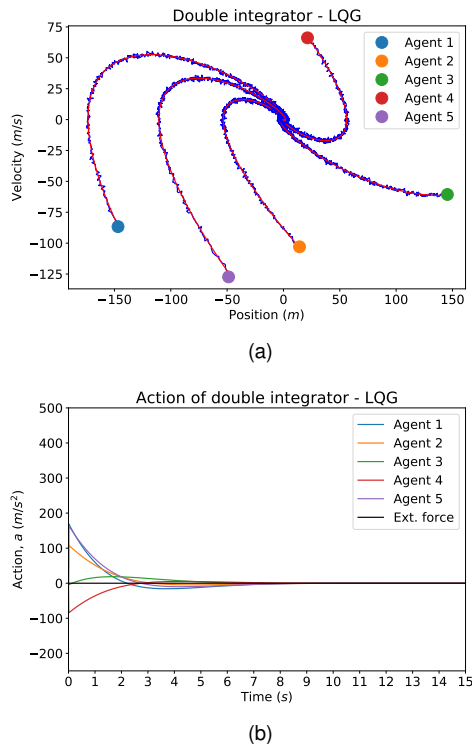


Figure 1: **The double integrator solved using LQG.** (a) Five limb segments implementing LQG with different initial conditions showing the observed trajectories of the blocks in the phase-space (in blue) and the estimates of the trajectories (in red). (b) Actions for the five limbs.

In Fig. 1a we can see how, in a standard implementation of LQG, our limb (or block) is driven to the desired position $x = 0$ and velocity $x' = 0$ from a set of 5 randomly initialised conditions (zero-mean Gaussian distributed, $sd=100$). In Fig. 1b we then simply show the actions over time of the same 5 example limbs, all converging to zero since the systems effectively reach their desired target. The main feature of LQG, and from which active inference departs, is the reliability of estimates of both position and velocity (the red line in the phase space). In LQG, accurate estimates are necessary to then enact the LQR component implementing a negative feedback mechanism based on estimates $\hat{\mathbf{x}}$ rather than true hidden states \mathbf{x} . When knowledge of the motor signals \mathbf{a} is removed from a Kalman-Bucy filter in an LQG set up, estimates of the hidden properties of the world become inaccurate and unstable, as shown in Fig. 2 for the double integrator. In this example, rather than converging towards the desired state, our simu-

¹Code available at <https://github.com/mbaltieri/doubleIntegrator>.

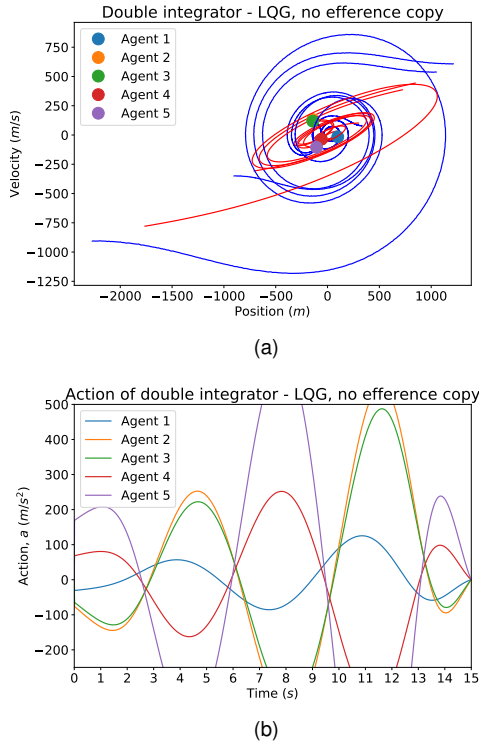


Figure 2: **The double integrator solved using LQG with no efference copy.** (a) Five limb segments with different initial conditions showing the observed trajectories in the phase-space (in blue) and the estimates of the same trajectories (in red) when an (efference) copy motor signals a is removed. (b) Actions for the five limbs.

lated limbs get away from it (Fig. 2a) since the new observations are too inaccurate given the lack of mechanisms to discount the effects of a . In Fig. 2b we can then see that actions a begin to exponentially oscillate rather than converging to zero, as in Fig. 1. This is due to one of the assumptions for observability in Kalman(-Bucy) filters (Stengel, 1994), explicitly requiring knowledge of *all* inputs, including one’s motor commands, and outputs of a system in order to determine its latent states.

The double integrator with active inference

The generative model for the double integrator introduces priors representing an imaginary spring that pulls the limb back to the origin ($x = 0$) and an imaginary damper that slows it down ($x' = 0$), see Baltieri and Buckley (2019a) where the equations and parameters are introduced. In the simplest case, the active inference solution is equivalent to a PID controller (Baltieri & Buckley, 2019b), the “optimal” linear solution when knowledge of one’s own actions a is not available in the generative model. In Fig. 3 we can see an example implementation of the double integrator using active inference. Five simulated limb segments are initialised at random posi-




Figure 3: **The double integrator solved using active inference.** (a) Five limb segments implementing active inference agents with different initial conditions showing the observed trajectories in the phase-space (in blue) and the estimates of the trajectories (in red). (b) Actions for the five limbs.

tion and velocity (zero-mean Gaussian distributed, $sd=100$) and Fig. 3a, and converge to the target solution where the output actions are essentially zero (excluding some noise), as expected Fig. 3b. The most striking feature is that estimates of both position and velocity of the block are very inaccurate but the limb nonetheless reaches the target in the phase space. These differences are given by the generative model implemented by the limb, encoding an imaginary spring-damper system that pulls it towards its “desired” state.

Discussion and conclusions

In this work we introduced a minimal model of motor control for a system with a single degree of freedom (Gottlieb, 1993). Following standard computational approaches for the modelling of motor functions relying on LQG architectures (Todorov & Jordan, 2002; McNamee & Wolpert, 2019), we discussed the importance of an “efference copy” (Crapse & Sommer, 2008; Schwartz, 2016; Straka et al., 2018) of motor signals for solving even the simplest of tasks. A growing body of work, in fact, nowadays challenges its definition and proposes new paradigms that can eschew such copy (Adams et al., 2013; Feldman, 2016). Without such a signal, representing a system’s own motor commands to be disambiguated

from external stimuli, standard implementations LQG architectures cannot solve even simple motor control tasks. To investigate alternative approaches, we discussed and compared a proposal based on efference-copy-free architectures, active inference. Active inference is a process theory derived from the free energy principle (Friston, 2010), advocating the minimisation of variational free energy, or under simplifying assumptions prediction error, as a driving principle for describing functions including perception, motor control and learning. To solve the same control problem we introduced, active inference relies on the generation of predictions of proprioceptive sensations (position, velocity and acceleration of the agent in this case), followed by the implementation of actions in the world via (trivial) reflex arcs. The proprioceptive modality is essentially treated as other inputs (vision, audition, etc.) and estimates/predictions are generated using the same generative model taking advantage of incoming proprioceptive sensations. This produces a considerably different control system, with state estimates and actions now created by the same (generative) model, making it hard to clearly separate processes of perception and action, see Fig. 2 and related discussion in Baltieri and Buckley (2018). The copy of motor control signals (i.e., efference copy), necessary in standard LQG settings to meet the observability constraints of Kalman-Bucy filters (Stengel, 1994) is not included in this formulation. Active inference postulates, in fact, that direct representations of the causes of self-generated sensations need not be discounted during the prediction of new incoming sensory inputs. This could be seen as a limitation of active inference accounts but, on the contrary, this may speak to the robustness of this approach in face of unknown inputs, i.e., motor actions produced by an agent as seen here or exogenous forces from the environment as proposed in Baltieri and Buckley (2019b, 2019a).

Acknowledgments

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

References

Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Structure and Function*, *218*(3), 611–643.

Åström, K. J., & Murray, R. M. (2010). *Feedback systems: an introduction for scientists and engineers*. Princeton university press.

Baltieri, M., & Buckley, C. L. (2018). The modularity of action and perception revisited using control theory and active inference. In T. Ikegami, N. Virgo, O. Witkowski, M. Oka, R. Suzuki, & H. Iizuka (Eds.), *The 2018 conference on artificial life* (p. 121-128).

Baltieri, M., & Buckley, C. L. (2019a). Nonmodular architectures of cognitive systems based on active inference. *arXiv preprint arXiv:1903.09542 (accepted at IJCNN 2019)*.

Baltieri, M., & Buckley, C. L. (2019b). PID control as a process of active inference with linear generative models. *Entropy*, *21*(3).

Buckley, C. L., Kim, C. S., McGregor, S., & Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, *14*, 55–79.

Buhrmann, T., & Di Paolo, E. A. (2014). Spinal circuits can accommodate interaction torques during multijoint limb movements. *Frontiers in computational neuroscience*, *8*, 144.

Crapse, T. B., & Sommer, M. A. (2008). Corollary discharge across the animal kingdom. *Nature Reviews Neuroscience*, *9*(8), 587.

Feldman, A. G. (2009). New insights into action–perception coupling. *Experimental Brain Research*, *194*(1), 39–58.

Feldman, A. G. (2016). Active sensing without efference copy: referent control of perception. *Journal of neurophysiology*, *116*(3), 960–976.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews. Neuroscience*, *11*(2), 127–138.

Friston, K. (2011). What is optimal about motor control? *Neuron*, *72*(3), 488–498.

Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive neuroscience*, 1–28. doi: 10.1080/17588928.2015.1020053

Gottlieb, G. L. (1993). A computational model of the simplest motor program. *Journal of Motor behavior*, *25*(3), 153–161.

Jordan, M. I. (1996). Computational aspects of motor control and motor learning. In *Handbook of perception and action* (Vol. 2, pp. 71–120). Elsevier.

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current opinion in neurobiology*, *9*(6), 718–727.

McNamee, D., & Wolpert, D. M. (2019). Internal models in biological control. *Annual review of control, robotics, and autonomous systems*, *2*, 339–364.

Schwartz, A. B. (2016). Movement: how the brain communicates with the world. *Cell*, *164*(6), 1122–1135.

Stengel, R. F. (1994). *Optimal control and estimation*. Courier Corporation.

Straka, H., Simmers, J., & Chagnaud, B. P. (2018). A new perspective on predictive motor signaling. *Current Biology*, *28*(5), R232–R243.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature neuroscience*, *7*(9), 907–915.

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, *5*(11), 1226.

Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature neuroscience*, *3*(11s), 1212.