A Worked Example of Fokker-Planck-based Active Inference

Magnus T. Koudahl and Bert de Vries

Eindhoven University of Technology, Electrical Eng. Dept.
Groene Loper 19, 5612 AP Eindhoven, NL
corr. m.t.koudahl@tue.nl

Abstract. The Free Energy Principle (FEP) and its corollary active inference describe a complex theoretical framework with a substantial statistical mechanics foundation that is often expressed in terms of the Fokker-Planck equation. Easy-to-follow examples of this formalism are scarce, leaving a high barrier of entry to the field. In this paper we provide a worked example of an active inference agent as a hierarchical Gaussian generative model. We proceed to write its equations of motion explicitly as a Fokker-Planck equation, providing a clear mapping between theoretical accounts of FEP and practical implementation. Code is available at github.com/biaslab/ai_workshop_2020.

Keywords: active inference · free energy · Fokker-Planck equation

1 Introduction

Theoretical treatments of the free energy Principle (FEP) and active inference are often framed in terms of the Fokker-Planck equation [5,6,10,12] and related flows. In this paper we aim to bridge a gap between theory and simulation by providing a worked example of an active inference agent written directly in terms of its Fokker-Planck equation. We provide a brief introduction to the Fokker-Planck description of dynamical systems and implement an agent based on a generative model structure common across the active inference literature. We then successfully apply the agent to a context switching task where it learns to track a harmonic oscillator.

2 The Fokker-Planck Equation for Dynamical Systems

The Fokker-Planck description of dynamical systems [5,6] starts by assuming that the system dynamics can be described by stochastic Langevin equations [5,8] of the form

$$\dot{\mathbf{x}} = \frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}) + \sqrt{2\Gamma(\mathbf{x})}\mathbf{W}(t)\,, \tag{1}$$

Corresponding author is Magnus T. Koudahl. This work is part of the research programme Efficient Deep Learning with project number P16-25 project 5, which is (partly) financed by the Netherlands Organisation for Scientific Research (NWO).

where $\mathbf{x}$ denotes the N-dimensional state of the system, $\Gamma(\mathbf{x})$ an $N \times M$ positive semi-definite diffusion matrix and $\mathbf{W}(t)$ is a standard M-dimensional Wiener process. Eq. 1 describes the evolution of a system under deterministic state-dependent dynamics $\mathbf{f}(\mathbf{x})$ and a stochastic fluctuation (diffusion) term $\mathbf{W}(t)$. Equivalently, we can consider the time derivative of the distribution generated by Eq. 1 in terms of the Fokker-Planck equation

$$\frac{\partial p(\mathbf{x},t)}{\partial t} = -\sum_{i=1}^{I} \frac{\partial}{\partial x_i} f_i(\mathbf{x}) p(\mathbf{x},t) + \sum_{i=1,j=1}^{I,J} \frac{\partial^2}{\partial x_i \partial x_j} \Gamma_{i,j}(\mathbf{x}) p(\mathbf{x},t) \tag{2}$$

where both i and j index over dimensions in $\mathbf{x}$. The Fokker-Planck equation describes the time derivative of the distribution $p(\mathbf{x},t)$ generated by Eq. 1 by a deterministic drift component or drag force (the first term) and a random diffusion process (the second term). The core move here is the switch from stochastic realisations of the SDE in Eq. 1 to the deterministic dynamics of the *distribution* over realisations of the same process in Eq. 2. A steady-state solution to the dynamics of Eq. 2 constitutes a vector field and can be written in potential form [1,5,6,8,10] as

$$\mathbf{f}(\mathbf{x}) = (Q(\mathbf{x}) - \Gamma(\mathbf{x}))\nabla J(\mathbf{x}), \tag{3}$$

where $Q(\mathbf{x})$ denotes an anti-symmetric ($Q = -Q^T$) curl matrix and $\Gamma(\mathbf{x})$ a positive semi-definite diffusion matrix. We use ∇ to denote the gradient and $J(\mathbf{x})$ is a potential function. For a proof of this relation, see [1,8]. Writing $\mathbf{f}(\mathbf{x})$ in this form, the anti-symmetric structure of $Q(\mathbf{x})$ describes a solenoidal flow that is *orthogonal* to gradients of $J(\mathbf{x})$. The positive semi-definiteness of $\Gamma(\mathbf{x})$ on the other hand leads to dissipative flow *along* gradients of $J(\mathbf{x})$.

3 Laplace-encoded free energy and generative models

To apply the Fokker-Planck equation to active inference, we follow [5,6] and let $J(\mathbf{x})$ denote a variational free energy functional. We now need to specify a generative model. A common choice in active inference literature [3,6] is a hierarchical generative model of the form

$$\begin{aligned} \mu_1 &= h_1(\mu_0) + \omega_1 & \phi_0 &= g_1(\mu_0) + w_0 \\ \mu_2 &= h_2(\mu_1) + \omega_2 & \phi_1 &= g_2(\mu_1) + w_1 \\ &\vdots & &\vdots \end{aligned} \tag{4}$$

We let μ_n denote internal states of the agent, ϕ_n sensory states and let $h_n(\cdot)$ and $g_n(\cdot)$ denote arbitrary link functions. Further assuming all noise terms w_n, ω_n are iid Gaussian, we can rewrite each layer of the hierarchical generative model

$p(\mu_{0:n}, \phi_{0:n})$ as

$$p(\mu_{n+1}|\mu_n) = \mathcal{N}(\mu_{n+1}|\mu_n, \sigma^2_{\mu_{n+1}}) = \frac{1}{\sqrt{2\pi\sigma^2_{\mu_{n+1}}}} \exp\left(-\frac{(\mu_{n+1} - h_{n+1}(\mu_n))^2}{2\sigma^2_{\mu_{n+1}}}\right)$$
$$p(\phi_n|\mu_n) = \mathcal{N}(\phi_n|\mu_n, \sigma^2_{\phi_n}) = \frac{1}{\sqrt{2\pi\sigma^2_{\phi_n}}} \exp\left(-\frac{(\phi_n - g_n(\mu_n))^2}{2\sigma^2_{\phi_n}}\right) \quad (5)$$

where $\sigma^2_{\bullet_n}$ denotes prior variance at the n-th level of $\bullet \in \{\mu, \phi\}$. Once the generative model has been specified, we need to constrain the recognition factors in order to compute the required gradients. Following [3,6] we assume a fully factorised Gaussian recognition density, also known as the mean-field variational Laplace approximation. Under these assumptions the free energy reduces to a sum of precision-weighted prediction errors between internal states at each level μ_n and the level above μ_{n+1}, and internal μ_n and sensory states ϕ_n. This chain can theoretically continue forever. To terminate the chain, we can assume excessive variance (i.e., negligible precision) at the highest level under consideration which renders higher order contributions to the free energy negligible. For a thorough derivation we refer to [2,3]. Ignoring constant terms and summing over levels, the free energy thus takes the form

$$J(\mu_{0:n}, \phi_{0:n}, a) = \sum_n \left(\frac{1}{2\sigma^2_{\mu_{n+1}}}(\mu_{n+1} - h_{n+1}(\mu_n))^2 + \frac{1}{2\sigma^2_{\phi_n}}(\phi_n(a) - g_n(\mu_n))^2\right). \quad (6)$$

Note that we additionally assume that ϕ_n depends on active states a. This is the inverse model assumption that augments the generative model [3,4]. The purpose of the inverse model is to update active states by allowing the derivative $\frac{\partial J}{\partial a}$ through

$$\dot{a} = \frac{\partial a}{\partial t} = -\frac{\partial J(\mu_{0:n}, \phi_{0:n})}{\partial a} = -\frac{\partial J(\mu_{0:n}, \phi_{0:n})}{\partial \phi_{0:n}} \frac{\partial \phi_{0:n}}{\partial a} \quad (7)$$

where we explicitly mediate the effects of action on free energy through the agents sensory states ϕ_n [3,4,6]. This move is usually justified by an appeal to reflex arcs in a neuroscience context [3,4,9] and has successfully been applied in simulation [3,6] as well as robotics [9]. Note that Eq. 7 is effectively a gradient flow on free energy, following the functional form of Eq. 3.

4 A worked example

We proceed by defining an environmental process as a harmonic oscillator with an added friction term. The environmental system dynamics are described by a Hamiltonian (a potential function) that decomposes into potential and kinetic energy terms, plus added friction administered by the agent through action

$$H(x, \dot{x}, a) = \underbrace{\frac{1}{2m}x^2}_{\text{potential}} + \underbrace{\frac{1}{2}k\dot{x}^2}_{\text{kinetic}} - \underbrace{\dot{x}u\tanh(a)}_{\text{friction}} . \tag{8}$$

Here x denotes the position of the system, $\dot{x}$ the velocity, m the mass, k is a constant, u is a force term that bounds the amount of friction the agent can administer and a still represents action. The system obeys standard Hamiltonian dynamics

$$\frac{dx}{dt} = -\frac{\partial H}{\partial \dot{x}}, \frac{d\dot{x}}{dt} = \frac{\partial H}{\partial x} . \tag{9}$$

Hamiltonian dynamics are commonly applied to the description of conservative systems [5]. However in the present example the additional friction term in Eq. 8 means the system no longer conserves energy. In other words, introducing action dependent friction allows the agent to systematically add or subtract energy from the system. If no action is taken ($a = 0$) the third term vanishes and the environmental process describes a standard conservative simple harmonic oscillator.

We let $J(\mu_{0:1}, \phi_0, a)$ denote the free energy of a two-layer model that receives observations only at the first level. The agent thus only observes position and not velocity. Formally this means setting $\phi_0 = x$ and omitting higher orders of ϕ. We then define a new potential vector J' as the concatenation of the Hamiltonian of the environmental process $H(x, \dot{x}, a)$ and the free energy functional $J(\mu_{0:1}, \phi_0, a)$ of the agent

$$J' = \begin{bmatrix} H(x, \dot{x}, a) \\ J(\mu_{0:1}, \phi_0, a) \end{bmatrix} \Rightarrow \nabla J' = \begin{bmatrix} \frac{x}{m} \\ k\dot{x} - u\tanh(a) \\ \frac{1}{\sigma^2_{\mu_1}}\left(\mu_1 - h_1(\mu_0)\right) + \frac{1}{\sigma^2_{\phi_0}}\left(x - g_0(\mu_0)\right) \\ \frac{1}{\sigma^2_{\mu_1}}\left(h_1(\mu_0) - \mu_1\right) \\ -\frac{1}{\sigma^2_{\phi_0}}\left(x - g_0(\mu_0)\right)u\,\mathrm{sech}^2(a) \end{bmatrix} . \tag{10}$$

We further assume an accurate inverse model for the effect of action a on observations ϕ_0. Concretely this means the agent is able to accurately calculate the gradient flow described in Eq.7 which is given by

$$\dot{a} = -\frac{\partial J(\mu_{0:1}, \phi_0)}{\partial \phi_0}\frac{\partial \phi_0}{\partial a} = \frac{1}{\sigma^2_{\phi_0}}\left(x - g_0(\mu_0)\right)u\,\mathrm{sech}^2(a) \tag{11}$$

where $\mathrm{sech}(\cdot)$ is the hyperbolic secant. This derivative can be found in 5th row of $\nabla J'$. Note that the sign is opposite before multiplication by $-\Gamma$. Since the agent does not observe velocity $\dot{x}$, the corresponding sensory prediction error involving ϕ_1 is absent at the second level (4th row of $\nabla J'$).

Concatenating the vector J' allows for simultaneous integration of the agent and the environment by using block matrices for Q and Γ. Assuming Hamiltonian dynamics for the agent as well, we can write the block system matrices as

$$Q = \begin{bmatrix} 0 & -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \gamma_1 & 0 & 0 \\ 0 & 0 & 0 & \gamma_2 & 0 \\ 0 & 0 & 0 & 0 & \gamma_3 \end{bmatrix}, \mathbf{x} = \begin{bmatrix} x \\ \dot{x} \\ \mu_0 \\ \mu_1 \\ a \end{bmatrix}. \tag{12}$$

Here $\mathbf{x}$ denotes the state vector of the combined system similarly to Eq. 3 and Q encodes two blocks of Hamiltonian dynamics. Internal states of the agent additionally perform gradient descent on $J(\mu_{0:1}, \phi_0, a)$ with learning rates γ_1 and γ_2. Action is updated by gradient descent on $J(\mu_{0:1}, \phi_0, a)$ with learning rate γ_3. Note that by virtue of the Fokker-Planck formalism, the learning rates acquire an interpretation in terms of the amplitude of random fluctuations. In other words, maintaining nonequilibrium steady state in a noisy environment mandates high learning rates. Substituting Eq. 10 and Eq. 12 into Eq. 3 now finishes the dynamics that underwrite active inference for a partition of states (evolving under the dynamics of Eq. 3) into external states, internal states, sensory states and action.

5 Results

We simulated the system for 50 timesteps with $\gamma_1 = \gamma_2 = 0.1$, $\gamma_3 = 1$, $\sigma^2_{\bullet n} = 0.1$, $h_n(\mu_n) = g_n(\mu_n) = \mu_n$, $m = k = 1$,$u = 0.5$ and initial state $x = 2$, $\dot{x} = 2$, $\mu_0 = 0$, $\mu_1 = 0$, $a = 0$. Note that the initial states of the agent and the environment are different. At $t = 25$, we change the parameters of the environmental process, setting $m = 10$, $k = 0.1$ and resetting the states of the environment to $x = 10, \dot{x} = 2$. This results in an abrupt change in the environmental process. The task of the agent is then twofold: (1) it needs to learn environmental dynamics to accurately predict incoming observations and (2) it needs to flexibly adapt to a change in previously learnt dynamics. Results are shown in Fig. 5. After an initial learning period we observe that the agent accurately learns to track the environmental process. The agents active states settle into an oscillatory pattern to smooth out the trajectory and dampen noise. When the environmental process changes, we observe a new learning period as the agent adapts to the context switch. Prediction errors are quickly attenuated and the agent resumes accurately tracking the environmental process.

6 Discussion

In this paper we showed a worked example of an agent in the form of a common model structure and specified its equations of motion directly in terms of a Fokker-Planck equation. Writing the agent as a Fokker-Planck equation renders the coupling between theory such as [5,6] immediate, with the goal of providing

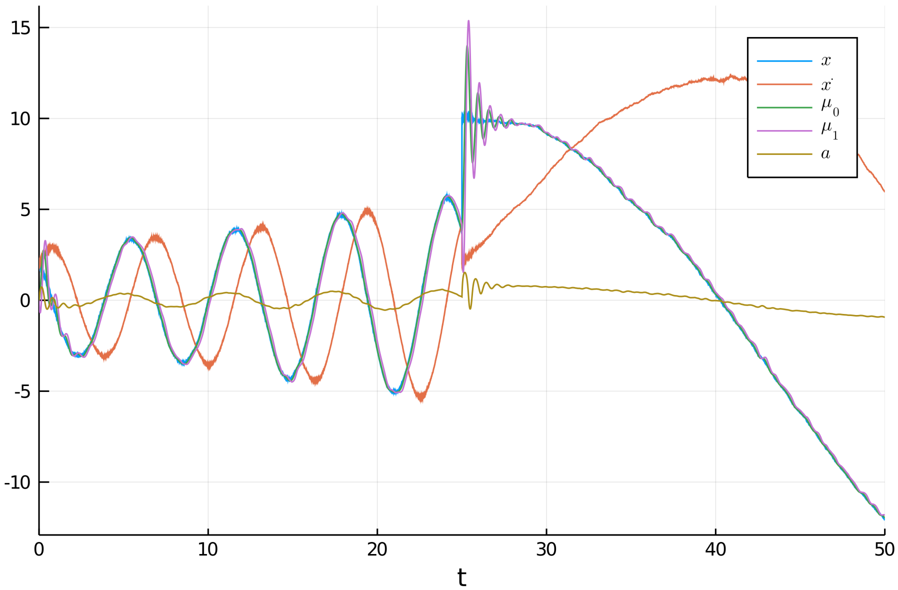

Fig. 1. Trajectory of agent and environmental process. Note the close correspondence between the blue and green lines (x and μ_0), showing that the agent succesfully learns to predict the environmental process.

an entry point for researchers interested in FEP. A second and more subtle point speaks to the Fokker-Planck equation as a way of writing equations of motion. Writing system dynamics in terms of the Fokker-Planck equation allows for interpreting the equations of motion as a "mechanics". The agent presented here operates under Hamiltonian mechanics but [5] opens up the possibility of investigating quantum- or electro-mechanical agents as well since these can also be written in terms of Q and Γ. Additionally, FEP literature offers a number of alternative free energies that are available as alternatives for the potential function $J(\mathbf{x})$, for example the Expected, Generalised and constrained Bethe free energies [7,11,13]. By combining choices for Q, Γ and $J(\mathbf{x})$ it is immediately clear that Fokker-Planck-based agents represent a sizeable class of agents that are mostly unexplored in practical applications.

References

1. Ao, P.: Potential in stochastic differential equations: novel construction. Journal of Physics A: Mathematical and General **37**(3), L25–L30 (Jan 2004). https://doi.org/10.1088/0305-4470/37/3/L01, `http://stacks.iop.org/0305-4470/37/i=3/a=L01?key=crossref.d94e3c0885d8b4979280533ef70299fd`

2. Baltieri, M., Buckley, C.: PID Control as a Process of Active Inference with Linear Generative Models. Entropy **21**, 257 (Mar 2019). https://doi.org/10.3390/e21030257
3. Buckley, C.L., Kim, C.S., McGregor, S., Seth, A.K.: The free energy principle for action and perception: A mathematical review. Journal of Mathematical Psychology **81**, 55–79 (Dec 2017). https://doi.org/10.1016/j.jmp.2017.09.004, https://linkinghub.elsevier.com/retrieve/pii/S0022249617300962
4. Friston, K.: What Is Optimal about Motor Control? Neuron **72**(3), 488–498 (Nov 2011). https://doi.org/10.1016/j.neuron.2011.10.018, http://www.sciencedirect.com/science/article/pii/S0896627311009305
5. Friston, K.: A free energy principle for a particular physics. arXiv:1906.10184 [q-bio] (Jun 2019), http://arxiv.org/abs/1906.10184, arXiv: 1906.10184
6. Friston, K., Ao, P.: Free Energy, Value, and Attractors. Computational and Mathematical Methods in Medicine **2012**, 1–27 (2012). https://doi.org/10.1155/2012/937860, http://www.hindawi.com/journals/cmmm/2012/937860/
7. Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., Pezzulo, G.: Active inference and epistemic value. Cognitive Neuroscience **6**(4), 187–214 (Oct 2015). https://doi.org/10.1080/17588928.2015.1020053, https://doi.org/10.1080/17588928.2015.1020053
8. Ma, Y.A., Chen, T., Fox, E.B.: A Complete Recipe for Stochastic Gradient MCMC. arXiv:1506.04696 [math, stat] (Oct 2015), http://arxiv.org/abs/1506.04696, arXiv: 1506.04696
9. Oliver, G., Lanillos, P., Cheng, G.: Active inference body perception and action for humanoid robots. arXiv:1906.03022 [cs] (Aug 2019), http://arxiv.org/abs/1906.03022, arXiv: 1906.03022
10. Parr, T., Da Costa, L., Friston, K.: Markov blankets, information geometry and stochastic thermodynamics. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **378**(2164), 20190159 (Feb 2020). https://doi.org/10.1098/rsta.2019.0159, https://royalsocietypublishing.org/doi/10.1098/rsta.2019.0159
11. Parr, T., Friston, K.J.: Generalised free energy and active inference: can the future cause the past? preprint, Neuroscience (Apr 2018). https://doi.org/10.1101/304782, http://biorxiv.org/lookup/doi/10.1101/304782
12. Parr, T., Sajid, N., Friston, K.J.: Modules or Mean-Fields? Entropy **22**(5), 552 (May 2020). https://doi.org/10.3390/e22050552, https://www.mdpi.com/1099-4300/22/5/552, number: 5 Publisher: Multidisciplinary Digital Publishing Institute
13. Zhang, D., Wang, W., Fettweis, G., Gao, X.: Unifying Message Passing Algorithms Under the Framework of Constrained Bethe Free Energy Minimization. arXiv:1703.10932 [cs, math] (Mar 2017), http://arxiv.org/abs/1703.10932, arXiv: 1703.10932