

Geometry of Friston’s active inference

Martin Biehl
Araya Inc., Tokyo

MARTIN@ARAYA.ORG

Abstract

We reconstruct Karl Friston’s active inference and give a geometrical interpretation of it.

Keywords: active inference, variational inference, reinforcement learning, pomdp

1. Introduction

In Biehl et al. (2018) we have reconstructed the active inference approach as used in Friston et al. (2015). Here we present a radically shortened and more general account of active inference. We also present for the first time a geometrical interpretation of active inference.

We use the same notation and model as in Biehl et al. (2018). There readers can also find a translation table to the notations used in Friston et al. (2015, 2016).




Figure 1: Bayesian networks of (a) PA-loop and (b) generative model.

The active inference agent we are describing in the following interacts with an environment according to the perception-action loop defined by the Bayesian network in Figure 1(a). There, we write E_t for the environment state, S_t for the sensor value, A_t for the action, and M_t for the memory state of the agent at time step t . We assume that for all time steps t the transition dynamics of the environment $p(e_{t+1}|a_{t+1}, e_t)$ and the dynamics of the sensors $p(s_t|e_t)$ are time-homogenous and fixed. We further assume that the agent has perfect memory, i.e. at all times $m_t := sa_{\prec t}$.

The difference to the setup of reinforcement learning for partially observable Markov decision problems is that there is no explicit reward signal. Instead the agent uses a motivation functional \mathfrak{M} which evaluates the agent’s current beliefs about the consequences of

actions (see Section 3). A standard reward signal R_t can easily be added as an additional sensor value by letting $S_t \rightarrow (S_t, R_t)$.

With the environment, sensor, and memory dynamics fixed it remains to specify the action generation mechanism $p(a_t|m_t)$. We first describe this in two separate steps, *inference* and *action selection*, that can be performed one after the other. Then we show how active inference combines the two steps in one optimization procedure. On the way we also define motivation functionals.

Action generation in more detail: The agent employs a parameterized model in order to predict the consequences of its actions. At each time step it receives a new sensor value in response to an action and updates its model by conditioning on the new memory state. Additionally conditioning on the various possible future actions (or policies) results in a conditional probability distribution which we call the *active posterior*. The active posterior represents the agent’s beliefs about the consequences (for future sensors, latent variables, and internal parameters) of its actions. Obtaining the active posteriors is referred to as the *inference* step. Subsequently, the agent constructs (using softmax) a probability distribution over future actions by assigning high probability to those actions whose entries in the active posterior achieve high values when plugged into the motivation functional (e.g. expected future reward if there is an explicit one). This results in the what we call the *induced policy*. Obtaining the latter is referred to as the *action selection* step. Afterwards the agent can simply sample from this induced policy to generate actions.

In active inference, both inference and action selection are turned into optimization problems and then combined in a multiobjective optimization. The inference step can be turned into an optimization using variational inference (see e.g. Bishop, 2011; Blei et al., 2017). Variational inference introduces a variational version of the active posterior. Since the variational active posterior generally differs from the original/true active posterior it leads to a, generally different, *variational induced policy*. Action selection is then formulated as an optimization by introducing an additional (third) policy whose divergence from the variational induced policy is to be minimized. Active inference then optimizes a sum of the respective objective functions and afterwards the agent can sample from the third policy. With a bit of notational trickery this can be written in the form Friston uses, which looks similar to a variational free energy (or evidence lower bound).

2. Inference

The agent performs inference on the generative model given by the Bayesian network in Figure 1(b). The variables that model variables occurring outside of $p(a|m)$ in the perception-action loop (Figure 1(a)), are denoted as hatted versions of their counterparts. To clearly distinguish the probabilities defined by the generative model from the true dynamics, we use the symbol q instead of p . Here, $\theta^1, \theta^2, \theta^3$ are the parameters of the model. To save space, write $\theta := (\theta^1, \theta^2, \theta^3)$.

The last modelled time step \hat{T} can be chosen as $\hat{T} = T$ (T is the final step of the PA-loop), but it is also possible to always set it to $\hat{T} = t + n$, in which case n specifies a future time horizon from current time step t .

Active posterior At time t the agent plugs its experience $sa_{<t}$ into its generative model by setting $\hat{A}_\tau = a_\tau$, for $\tau < t$ and $\hat{S}_\tau = s_\tau$, for $\tau < t$ and conditioning on these. The

consequences of future actions can be obtained by additionally conditioning on each possible future action sequence $\hat{a}_{t:\hat{T}}$. This leads to the conditional probability distribution that we call the *active posterior* (the experience $sa_{\prec t}$ is considered fixed):

$$q(\hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}, \theta | \hat{a}_{t:\hat{T}}, sa_{\prec t}) \quad (1)$$

Variational active posterior In active inference the active posterior is obtained via variational inference (see e.g. [Blei et al., 2017](#)).

We write r (instead of p or q) to indicate variational probability distributions and ϕ for the entire set of variational parameters. Note, the parameter ϕ contains parameters for each of the future action sequences $\hat{a}_{t:\hat{T}}$ i.e. $\phi = (\phi_{\hat{a}_{t:\hat{T}}})_{\hat{a}_{t:\hat{T}} \in \mathcal{A}^{\hat{T}-t+1}}$. The variational active posterior is of the form $r(\hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}, \theta | \hat{a}_{t:\hat{T}}, \phi)$. To construct the variational active posterior we cycle through the possible future action sequences $\hat{a}_{t:\hat{T}}$ and compute each entry. For a fixed $\hat{a}_{t:\hat{T}}$ the variational free energy, also known as the (negative) evidence lower bound (ELBO) in variational inference literature, is defined as:

$$\mathcal{F}[\hat{a}_{t:\hat{T}}, \phi, sa_{\prec t}] := \sum_{\hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}} \int r(\hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}, \theta | \hat{a}_{t:\hat{T}}, \phi) \log \frac{r(\hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}, \theta | \hat{a}_{t:\hat{T}}, \phi)}{q(s_{\prec t}, \hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}, \theta | \hat{a}_{t:\hat{T}}, a_{\prec t})} d\theta \quad (2)$$

Then variational inference amounts to solving for each $\hat{a}_{t:\hat{T}}$ the optimisation problem:

$$\phi_{\hat{a}_{t:\hat{T}}, sa_{\prec t}}^* := \arg \min_{\phi} \mathcal{F}[\hat{a}_{t:\hat{T}}, \phi, sa_{\prec t}]. \quad (3)$$

3. Action selection and induced policy

Let Δ_{AP} be the space of active posteriors. Then a *motivation functional* is a map $\mathfrak{M} : \Delta_{AP} \times \mathcal{A}^{\hat{T}-t+1} \rightarrow \mathbb{R}$ taking active posteriors $d(\cdot, \cdot, \cdot) \in \Delta_{AP}$ and a sequences of future actions $\hat{a}_{t:\hat{T}}$ to a real value $\mathfrak{M}(d(\cdot, \cdot, \cdot), \hat{a}_{t:\hat{T}}) \in \mathbb{R}$. An example of a motivation functional is the expected value of the sum over future rewards (if one of the sensor values is defined as the reward). Other possibilities can be found in [Biehl et al. \(2018\)](#). Now define for some $\gamma \in \mathbb{R}_0^+$ and some motivation functional \mathfrak{M} a softmax operator $\sigma_{\gamma}^{\mathfrak{M}}$ mapping active posteriors $d(\cdot, \cdot, \cdot)$ to probability distributions over future action sequences:

$$\sigma_{\gamma}^{\mathfrak{M}}[d(\cdot, \cdot, \cdot)](\hat{a}_{t:\hat{T}}) := \frac{1}{Z(\gamma)} e^{\gamma \mathfrak{M}(d(\cdot, \cdot, \cdot), \hat{a}_{t:\hat{T}})}. \quad (4)$$

Then we call

$$q(\hat{a}_{t:\hat{T}} | sa_{\prec t}) := \sigma_{\gamma}^{\mathfrak{M}}[q(\cdot, \cdot, \cdot, sa_{\prec t})](\hat{a}_{t:\hat{T}}) \quad (5)$$

the *induced policy* of active posterior $q(\cdot, \cdot, \cdot, sa_{\prec t})$ and \mathfrak{M} . And we call

$$r(\hat{a}_{t:\hat{T}} | \phi) := \sigma_{\gamma}^{\mathfrak{M}}[r(\cdot, \cdot, \cdot, \phi)](\hat{a}_{t:\hat{T}}) \quad (6)$$

the *induced variational policy* of variational active posterior $r(\cdot, \cdot, \cdot, \phi)$ and \mathfrak{M} .

Note that, if ϕ is not the optimized parameter $\phi_{sa_{\prec t}}^*$ then the induced variational policy cannot be expected to lead to actions that actually reflect the preferences encoded in \mathfrak{M} . On the other hand, if the true active posterior, or the optimized variational active posterior are used and $\gamma \rightarrow \infty$ the induced policy should correspond to the best guess for an agent with the given generative model, variational distributions, and motivation \mathfrak{M} .




Figure 2: On top the space of probability distributions over future action sequences $\Delta_{\hat{A}^{\hat{T}-t+1}}$ on the bottom the space of active posteriors Δ_{AP} .

4. Active inference

Now introduce an additional variational policy $s(\hat{a}_{t:\hat{T}}|\rho)$ parameterized by ρ in order to approximate the induced variational policy $r(\hat{a}_{t:\hat{T}}|\phi)$ of the variational posterior (action selection). To make this happen we can minimize the Kullback-Leibler divergence between the two.

To do inference and action selection at once we can then minimize the sum of the variational free energy (Equation (2)) and the Kullback-Leibler divergence w.r.t. ρ, ϕ (for an illustration of the situation see Figure 2):

$$\underbrace{\sum_{\hat{a}_{t:\hat{T}}} s(\hat{a}_{t:\hat{T}}|\rho) \mathcal{F}[\hat{a}_{t:\hat{T}}, \phi, sa_{\prec t}]}_{D_1} + \underbrace{\text{KL}[s(\hat{A}_{t:\hat{T}}|\rho) || r(\hat{A}_{t:\hat{T}}|\phi)]}_{D_2} \quad (7)$$

If ϕ and ρ are optimized the agent can then sample actions from $s(\hat{a}_{t:\hat{T}}|\rho)$.

Now if we change notation and let $s(\hat{a}_{t:\hat{T}}|\rho) \rightarrow r(\hat{a}_{t:\hat{T}}|\rho)$ and $r(\hat{a}_{t:\hat{T}}|\phi) \rightarrow q(\hat{a}_{t:\hat{T}}|\phi)$ then the above can be rewritten as:

$$\sum_{\hat{a}_{t:\hat{T}}, \hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}} \int r(\hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}, \theta, \hat{a}_{t:\hat{T}}|\rho, \phi) \log \frac{r(\hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}, \theta, \hat{a}_{t:\hat{T}}|\rho, \phi)}{q(s_{\prec t}, \hat{s}_{t:\hat{T}}, \hat{e}_{0:\hat{T}}, \theta, \hat{a}_{t:\hat{T}}|\phi, a_{\prec t})} d\theta. \quad (8)$$

This looks similar to a variational free energy or evidence lower bound and is the form of the free energy found in [Friston et al. \(2016\)](#). What distinguishes it from a true variational free energy is the occurrence of the parameter ϕ in numerator and denominator.

Acknowledgments

I want to thank Manuel Baltieri and Thomas Parr for helpful discussions as well as Christian Guckelsberger, Daniel Polani, Christoph Salge, and Simón C. Smith, for feedback on this article.

References

- Martin Biehl, Christian Guckelsberger, Christoph Salge, Simón C. Smith, and Daniel Polani. Expanding the Active Inference Landscape: More Intrinsic Motivations in the Perception-Action Loop. *Frontiers in Neurorobotics*, 12, 2018. ISSN 1662-5218. doi: 10.3389/fnbot.2018.00045. URL <https://www.frontiersin.org/articles/10.3389/fnbot.2018.00045/abstract>.
- Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. Springer, New York, April 2011. ISBN 978-0-387-31073-2.
- David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877, April 2017. ISSN 0162-1459, 1537-274X.
- Karl Friston, Francesco Rigoli, Dimitri Ognibene, Christoph Mathys, Thomas Fitzgerald, and Giovanni Pezzulo. Active Inference and Epistemic Value. *Cognitive Neuroscience*, 6(4):187–214, 2015. ISSN 1758-8936.
- Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo. Active Inference: A Process Theory. *Neural Computation*, 29(1):1–49, November 2016. ISSN 0899-7667.