Evaluating the scalability of deep active inference

Tomasz Korbak

University of Warsaw

Institue of Philosophy and Sociology, Polish Academy of Sciences

tomasz.korbak@gmail.com

Abstract

Active inference is a modeling framework in computational neuroscience build upon the assumption that cognitive systems encode a hierarchical generative model of the world and act as to minimize their prediction errors. Existing implementations of active inference, however, are poorly comparable to alternative models and fall short of meeting reproducibility standards. Moreover, these simulations are usually limited to toy environments, which casts a shadow of doubt on whether active inference scales up to complex cognitive systems in the wild. The goal of the project is to build a universal active inference agent based on recent advances in scaling up variational inference by coupling it with deep neural networks. The active inference agent will be evaluated against state-of-the-art reinforcement-learning algorithms in several Atari environments (a received benchmark in machine learning).

Goals of the project

- To implement a framework for conducting experiments with active inference agents,
- To investigate whether active inference agents can efficiently operate in complex environments,
- To investigate whether employing a neural network makes the agent scale better with increasing complexity of the environment,
- To investigate whether active inference agent can efficiently explore complex environments.

Prior work

Active inference and the free energy principle

Active inference is a modeling framework in computational neuroscience build upon the assumption of the free energy principle (FEP). According to FEP, cognitive systems encode a hierarchical generative model of the world and act as to minimize their prediction errors. Perception, learning and action can be cast as Bayesian variational inference, i.e. gradient descent on free energy, which imposes an upper bound on prediction error [2].

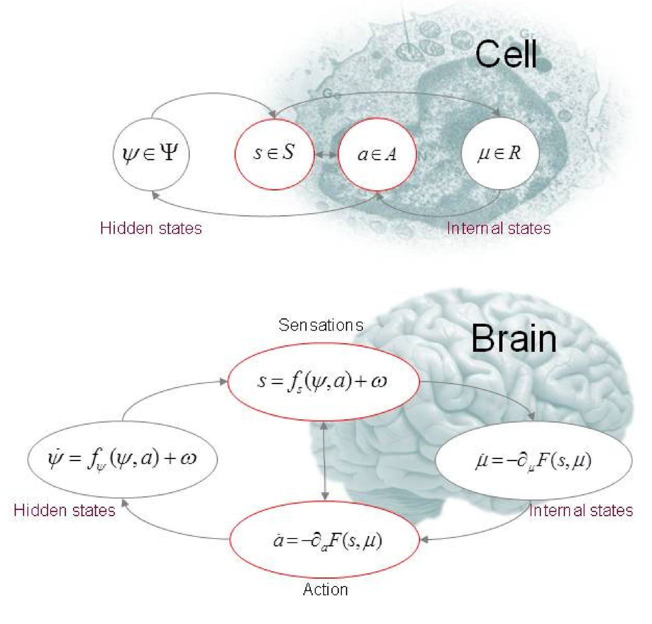

Figure 1: A Bayesian network formulation of of active inference. Nodes represent four sets of states (action, sensory, external, internal), while arrows represent the dependencies between them. Source: Wikipedia, CC BY-SA 3.0.

Problems with existing implementations of active inference

- They fall short of meeting reproducibility standards
- They fail to generate comparable, quantitative predictions
- Active inference was only evaluated on toy problems such as visual perception and saccadic eye movements. On the other hand, FEP claims to pertain cognition in general
- They were only used to model biological cognition, leaving unaddressed the issue of applying active inference to solving real-world engineering problems given industrial constraints

From reinforcement learning to active inference

- Reinforcement learning is the standard approach to sensorimotor control problems.
- In reinforcement learning we search for a policy (describing what to do in a given state) that maximizes expected reward.
- Equivalently, a reinforcement learning problem can be posed as an inference problem, i.e. computing the posterior over policies given a prior and received rewards.
- Therefore, active inference can be evaluated on reinforcement learning benchmarks and compared against a growing body of deep reinforcement learning algorithms.

Disclaimer: active inference is an unsupervised learning method, i.e. it does optimize directly for an external reward, but for surprise minimization.

From deep generative modeling to active inference

- Active inference framework is based on variational approximation of Bayesian inference, i.e. posing it as tractable optimization of the parameters of a given probability distribution.
- Variational methods were recently employed and greatly advanced in the context of deep learning.
- Most of these advancements come down to optimizing the parameters of a (deep) neural network parametrizing a distribution rather than optimizing the parameters of a distribution directly.

Key questions

Can active inference scale up to solve non-trivial problems in complex, real-world environments?

Active inference employs a variational approximation scheme to compute an intractable posterior. While the problem under variational assumptions is definitely simpler, is is still an open question whether active inference converges in real environments in a reasonable time.

Does employing a deep neural network influence its performance?

[4] combines recent advances in variational inference with active inference framework into what he calls deep active inference. A proof-of-concept implementation of deep active inference on mountain car problem is available. No attempts were made at solving more complex problems.

Is it comparable to state-of-the-art (deep) reinforcement learning methods?

Reinforcement learning is a fast-moving field [3]. It is also frequently criticizes for decoupling from constraints that biological agents face: intrinsic motivation, curiosity and lack of single well-defined scalar reward. Gym environments, established approach to comparing different agents, can be also employed for evaluating active inference agents.

Does it outperform reinforcement learning when it comes to exploration (maximizing epistemic value)?

Reinforcement learning approaches have well-known difficulties when it comes to taking novel paths and dealing with uncertainty. The free energy functional, on the other hand, naturally admits decomposition into extrinsic and epistemic value. Thus an active inference agent should be able to (Bayes-)optimally trade off exploration and exploitation.

Methods

Gym is a framework for evaluating reinforcement learning agents. It provides a diverse set of environments, each with a well-defined API [1].

We will implement a general-purpose (deep) active inference agent in PyTorch framework and provide it with a Gym-compatible API.

Then we will train and evaluate the agent on a set of environments, such as Mountain Car or Space Invaders. Various hyperparameters (including the use of neural network to parametrize the distribution) will be explored. Performance of shallow active inference deep active inference, and deep reinforcement learning will be compared on exploration-hungry problems.

Figure 2: Atari games, such as Space Invaders pictured above, are the E. coli of reinforcement learning.

Conclusions

There are two possible outcomes:

- (Deep) active inference will prove scalable, or
- It won't.

In the first case, the profit will be twofold: an unsupervised alternative to reinforcement learning will be delineated and a case will be made in favor of FEP as a tractable formal model of cognition.

If active inference fails to scale up, the endeavor will still be of value for the cognitive science community. A case will be made for narrowing the explanatory scope of FEP.

References

[1] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *CoRR*, abs/1606.01540, 2016.

[2] Karl Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, and G. Pezzulo. Active inference: A process theory. *Neural computation*, 2017.

[3] Yuxi Li. Deep reinforcement learning: An overview. *CoRR*, abs/1701.07274, 2017.

[4] Kai Ueltzhoffer. Deep active inference. *Biological Cybernetics*, 2018.