Active Inference for Dummies

(translating Friston)

We observe a phenomenon in our environment, and we take action onto it. The explanation how we make a decision to take that exact action, is what active inference, among other things tries to explain.

Principle of Active Inference.

We have an implicit model of the world; the more experiences or observations **o** we had in our past, the more our model **m** is rich. In that way we can recognize the new situation (observation) and apply the actions that worked out best in the past.

These possible actions (linked to a certain observation) that we can have in our “memory”, are called **control states**, **U,** and represent all the possible future actions we can choose to conduct in reality. So, the actions we applied to the world in the past for a certain observation enter our “memory” with a degree of success. We often tend to use the same action or behavior which had success in the past. This success assigns a high value to that action (arranged in a **vector C**, explained later). So using this high-value action also maximizes gain or utility.

Also, we wish to minimize the difference of what we **expect** to happen as a result of our possible action (control state) and the **true** *result of our real action*. Here Active Inference differs from other methods because instead of using the *true result of a real action*, it uses priors set in the goal or utility (set in **C**). We wish to minimize the surprise of an event and maximize our gain. We are beings who want to control our environment, to control the effect of our own behavior (influence) on the world, but also successfully reach our goals. In order to do that, Active Inference proposes that beings can take a finite set of states, making the whole problem computationally plausible. In order to avoid surprise (which is bound) we must develop a way to predict the consequences of our actions.

However if it is a completely new situation (observation), it becomes harder to act (almost random) but we try to infer the possible best actions from our model of the world, and predict the outcome (observations) -- the richer the given model m, the smaller the proba of being surprised. This minimization of surprise uses (also referred to as) a free-energy minimization approach.

This is a closed-loop system, where the action we just took becomes a cause of the following observation, which in turn causes us to make another action (a posterior), while storing the previous action in our “memory”, becoming part of the **empirical priors**.

Model of the world

What is interesting is the model **m** of the world we have, which is a generative model (from our empirical priors, we choose the best action, by referring to the sufficient statistics). I have already mentioned the **control states, U** (possible not yet performed actions, often linked to a certain observation). It is not all, there is something called the **hidden states S** that also take part in the generative model. Why are they hidden? Because they represent our interpretation of the true state of the world, which is in fact hidden to us.

So, the world is in fact in a state, a true state, and we try to infer it.

What we know for sure is what we can observe, so from there we can infer the possible true state of the world, representing it as the hidden states and applying possible control states we have in our " memory storage". Both control and hidden states are called the empirical priors.

We not only have the **empirical priors** (which probabilities to appear are updated) in our model of the world, but also the **full priors** which define our goal and its importance (how far are we willing to go to get it). They assist in choosing a specific action when in a current hidden state. One of these priors (vector C) contains the utility or the desired outcome or observation we wish to get. This simply means, that we will favorize to some degree one action more over the others, in order to get the expected outcome (gain). So, the vector c is a full prior over the desired outcomes. Active Inference also takes time into consideration, represented as the precision which will control the exploration-exploitation trade-off, arranged in ***gamma*** (e.g. favouring precise over rapid actions), it is called gamma because it uses a gamma distribution.

The probability to pass from a state $\mathbf{s_1}$ using control state $\mathbf{u_i}$, to another state $\mathbf{s_2}$, is regulated by the **transition matrix B**, {state1, possible_action1, state2}, which considers all the priors, in order to maximize gain. This is a typical Markov Decision Process example (state, action, state) following a policy (a cost of transition). So, these policies in Active Inference are presented in the expectations to choose an action.

Bayesian model

The Bayesian theory only refers to the posterior probability prediction, i.e. inferring the next future observation by taking knowledge from the past (the empirical priors) or the full priors.

A very simplistic eg. of the classical Bayes is: if you toss two coins, of which one is fake (always showing heads) and the other is true. If you toss them 3 times what is the probability to get heads? The prior probability, given we look for heads, for the fake one is always heads so p = 1, and for the other coin is p = 1/2. After we have tossed them 3 times and got p= 1 for fake and p=⅛ for the true coin, the posterior question would be the inverse from the first question (the likelihood) that is: after tossing 3 times and each time getting heads, what is the probability that the coin is the fake one?

So, by drawing an analogy, the more times we toss, the more certain we will be, and the future observation (that it will be heads or not) is more predictable. And we infer the next observation by using the knowledge from the past (likelihood and priors).

The ABC (matrices) of Active Inference

A represents **the likelihood** that an observation **o** will occur given a state **s**, the higher the likelihood the better estimations one has over the future.

B is the probability to transition from one state to the other.

C will influence the likelihood (transitions between the states), the choice of the control states in order to achieve the desired outcome. In it we assign values for all possible final outcomes, meaning we will penalize undesired outcomes and favour desired ones (often with a softmax function).

Summary, brief:

1. There is an observation **o_1**
2. We interpret **o_1** with our model of the world, **m**
   With the model we infer the true states of the world:
   (**i)** representing them with the hidden states **s** which
   (ii) limiting the set of possible actions, control states linked to this observation or hidden state. Simply put, the hidden states assist in choosing a subset of actions (otherwise it is an enormous amount of possible actions). So, from the current state s_t we can choose from u_t…. u_T possible future actions and lead to s_t+1,
   - We can also choose from true actions we just performed in the past (a_1,a_2...a_t).

   ***(iii) control states have values of success from the past, which gives them a degree of probability of being used again or not (expectation)
3. We wish to choose such action (to pass from one state to another while predicting posterior outcome):
   a. so that the surprise is minimal (free-energy minimization).
   b. so that the utility is maximized (defined in C)
   c. According to likelihood -- $P(o|s_t)$ - likelihood is tied to the past observations, and with them it infers the probability that observation o would occur given the state s.
   d. According to precision (gamma)

Dependencies (closed loop)

Observations → from which we infer true states of the world (**hidden states** to us) → which lead us to select a subset of possible actions (**control states**) → the following action will modify/create the next observation → which we again interpret with our model of the world.

Conclusion

The more observations we have, the richer our model is. In this way we reduce the possibility of being surprised, and maximize the expected gain.
The world is in a state of which we have a limited interpretation -- it is in a hidden state which we infer from what we can observe. With these hidden states we infer the possible actions we can take to maximize our gain, and minimize the surprise of the next future true state.
Active inference uses the Bayes theory to explain the choice of action a bounded, rational being makes to infer the future observations, in order to minimize the surprise.