Endowing the Machine with Active Inference : Application to a P300 Speller BCI

Jelena Mladenovic
Fabien Lotte
Jérémie Mattout

Brain Computer Interface

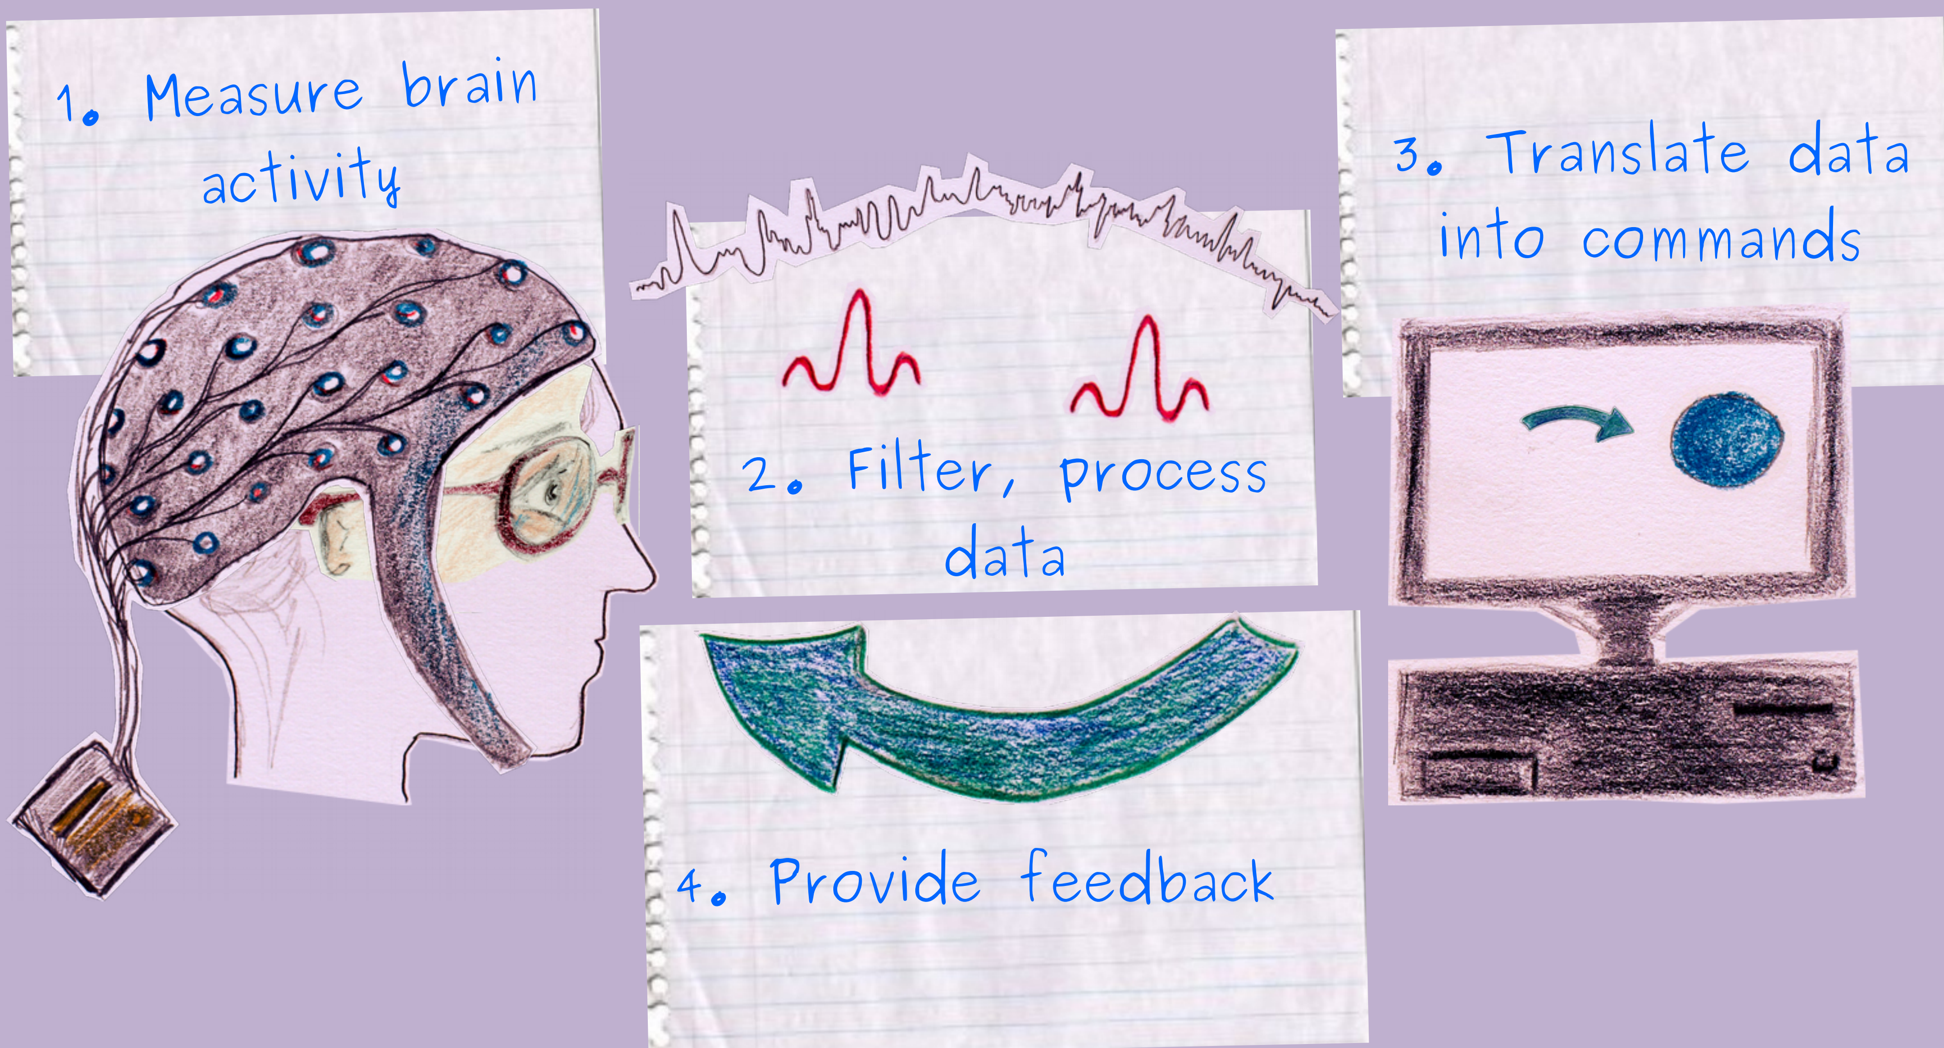

Brain Computer Interface

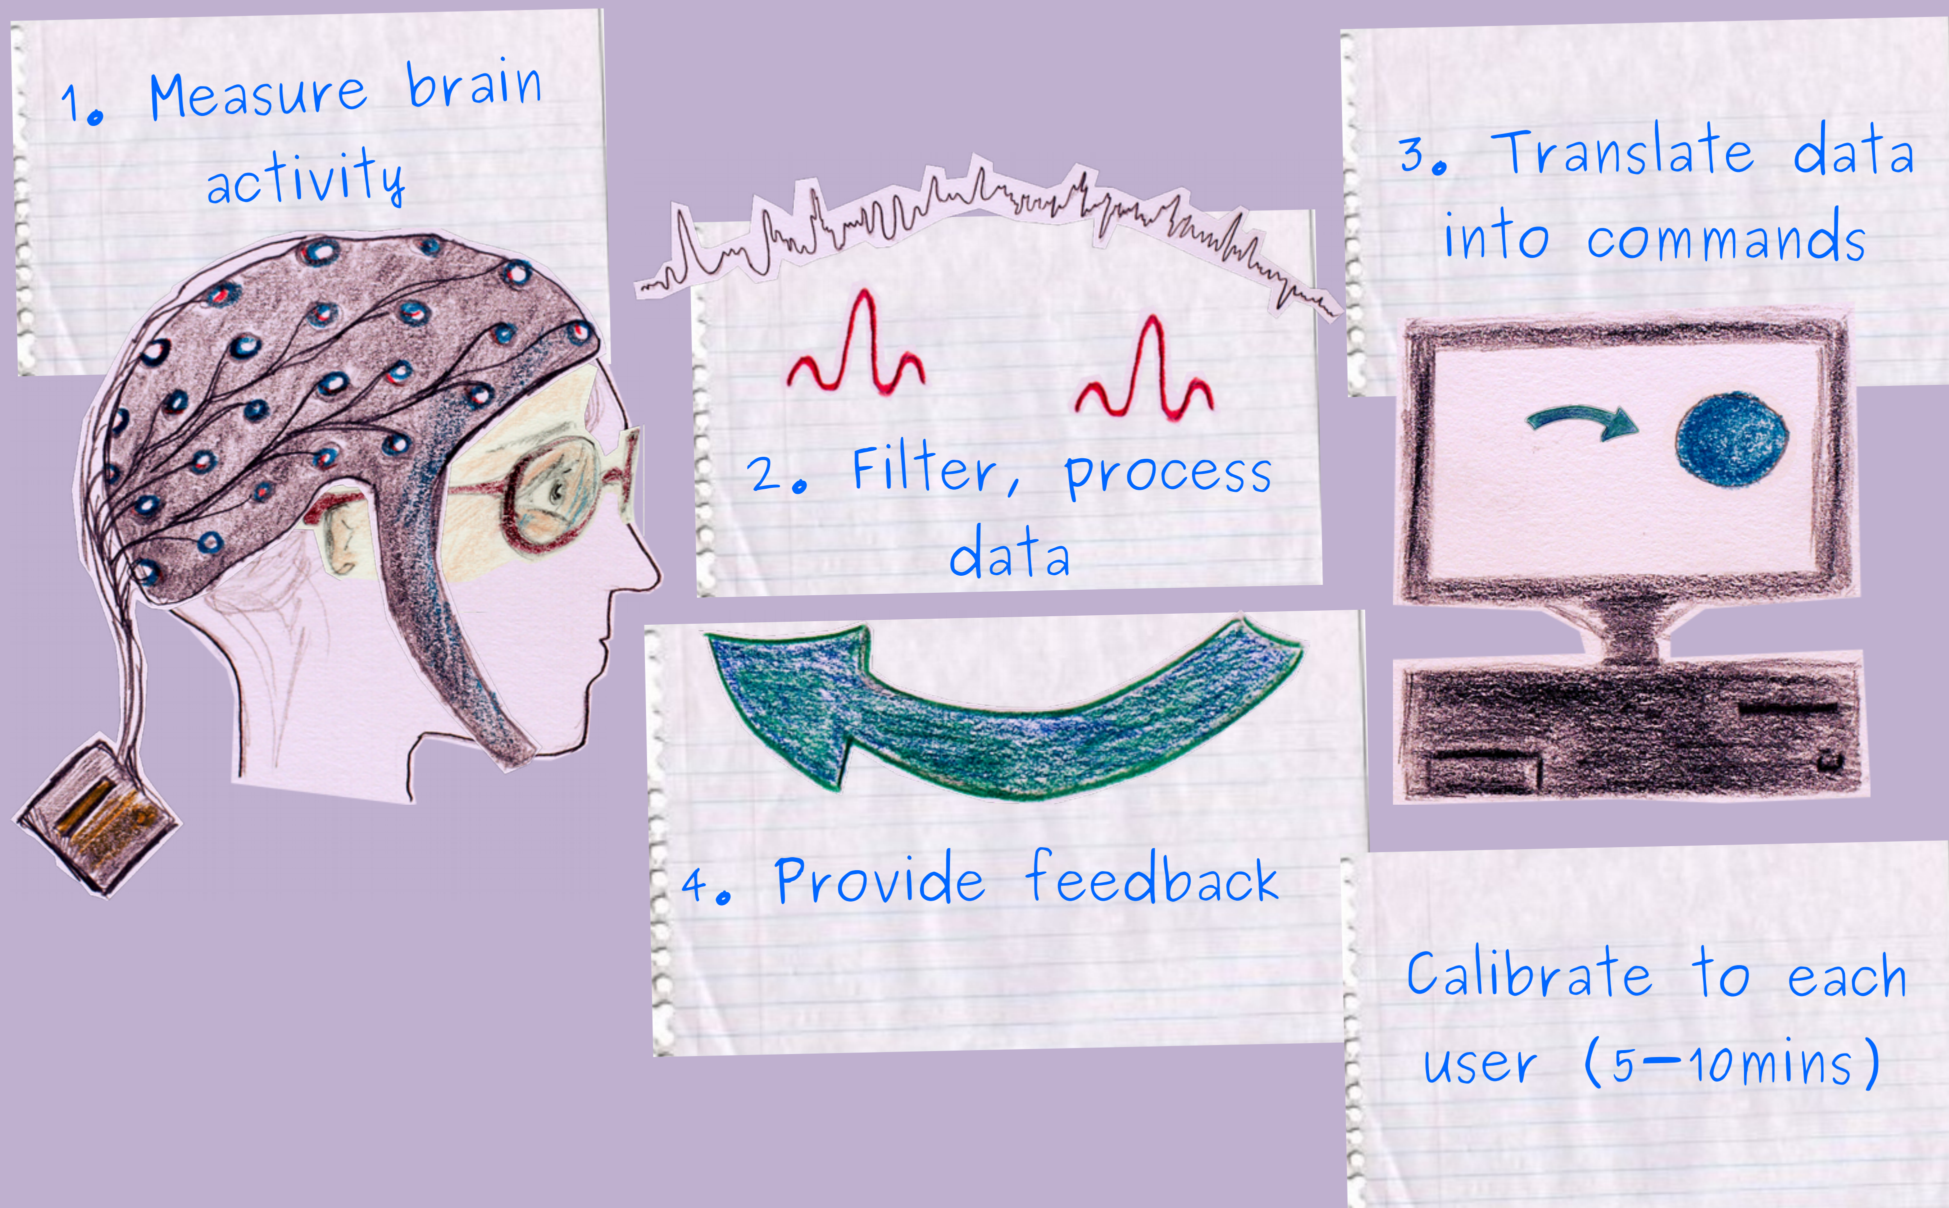

Brain Computer Interface

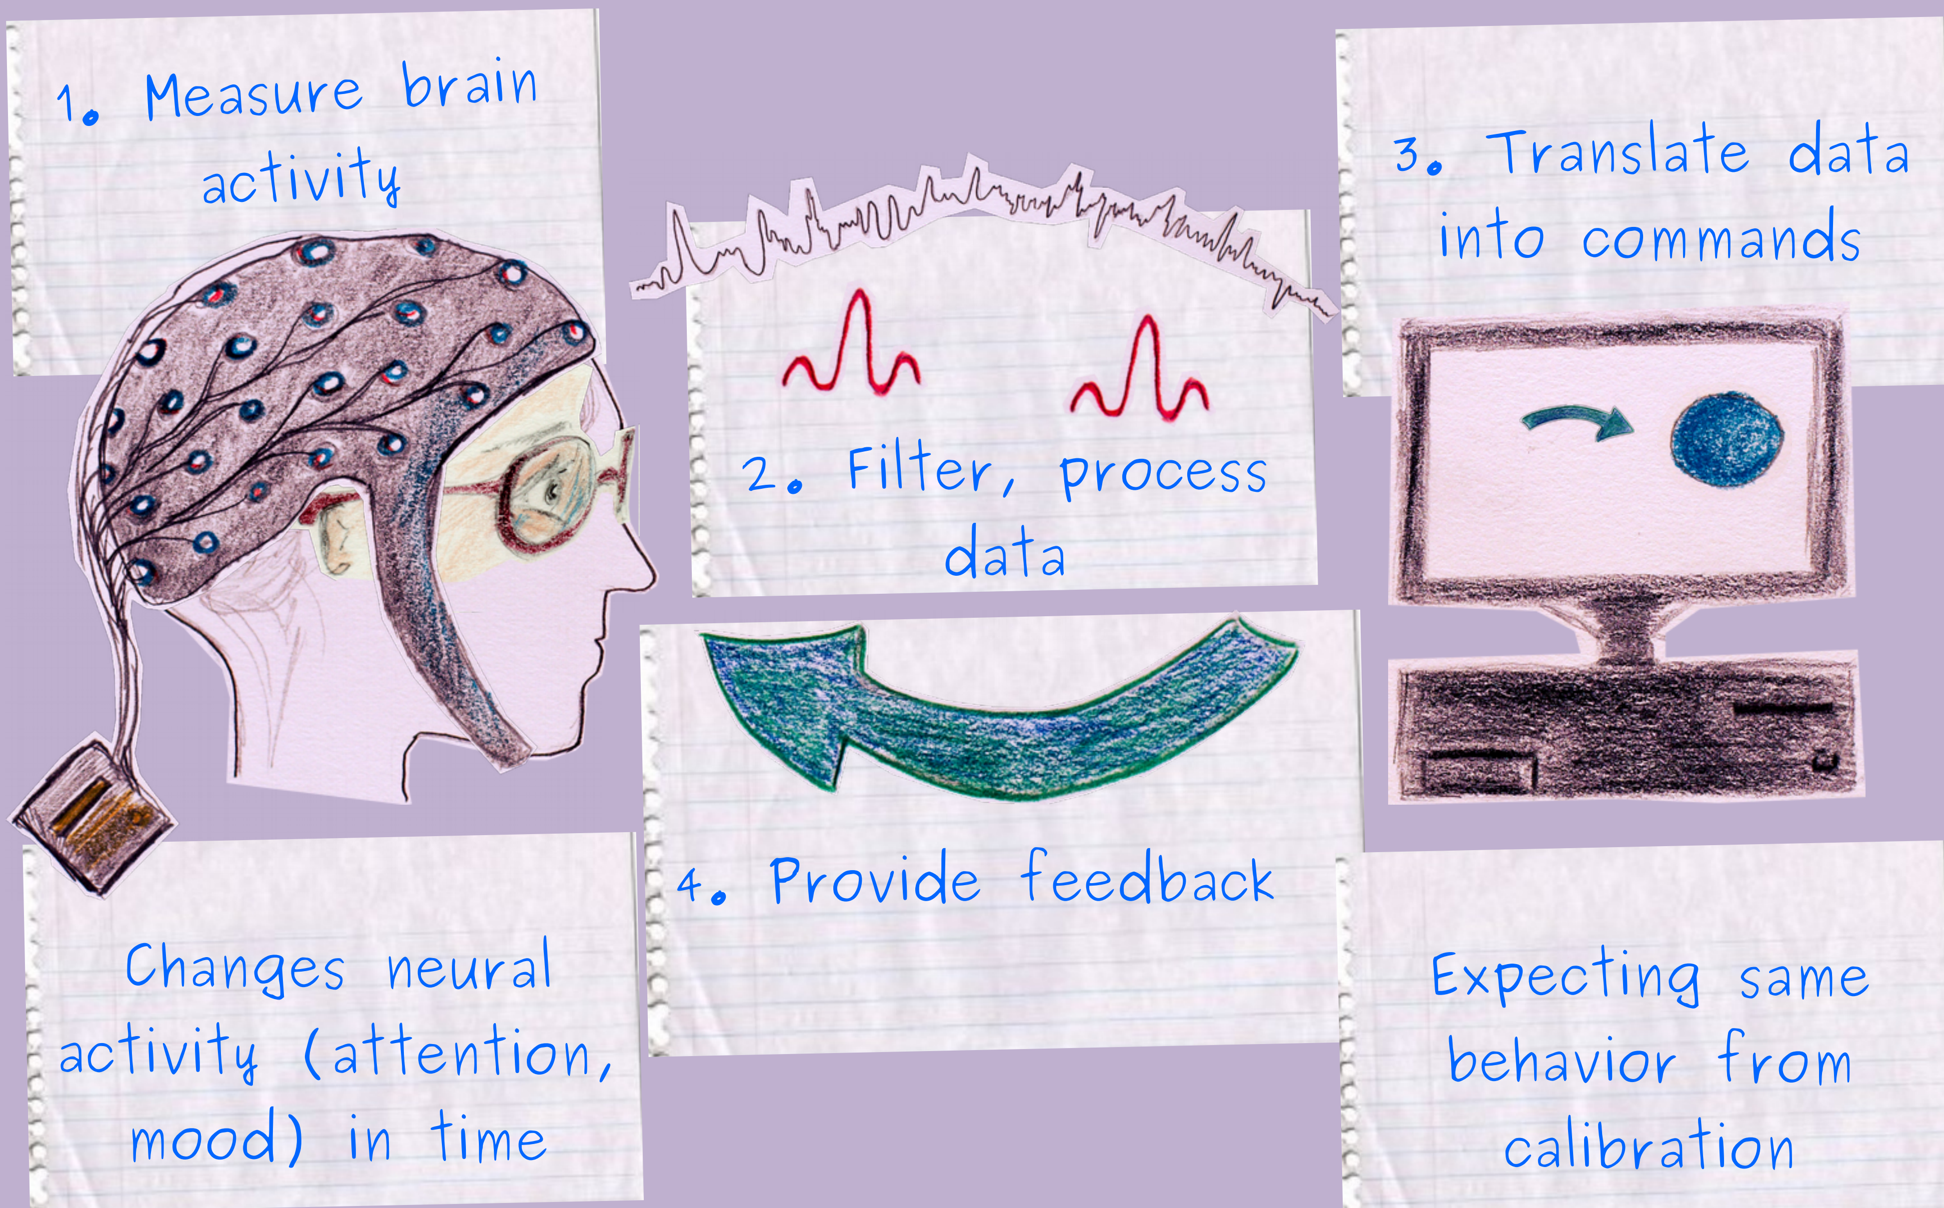

Brain Computer Interface
ADAPTIVE

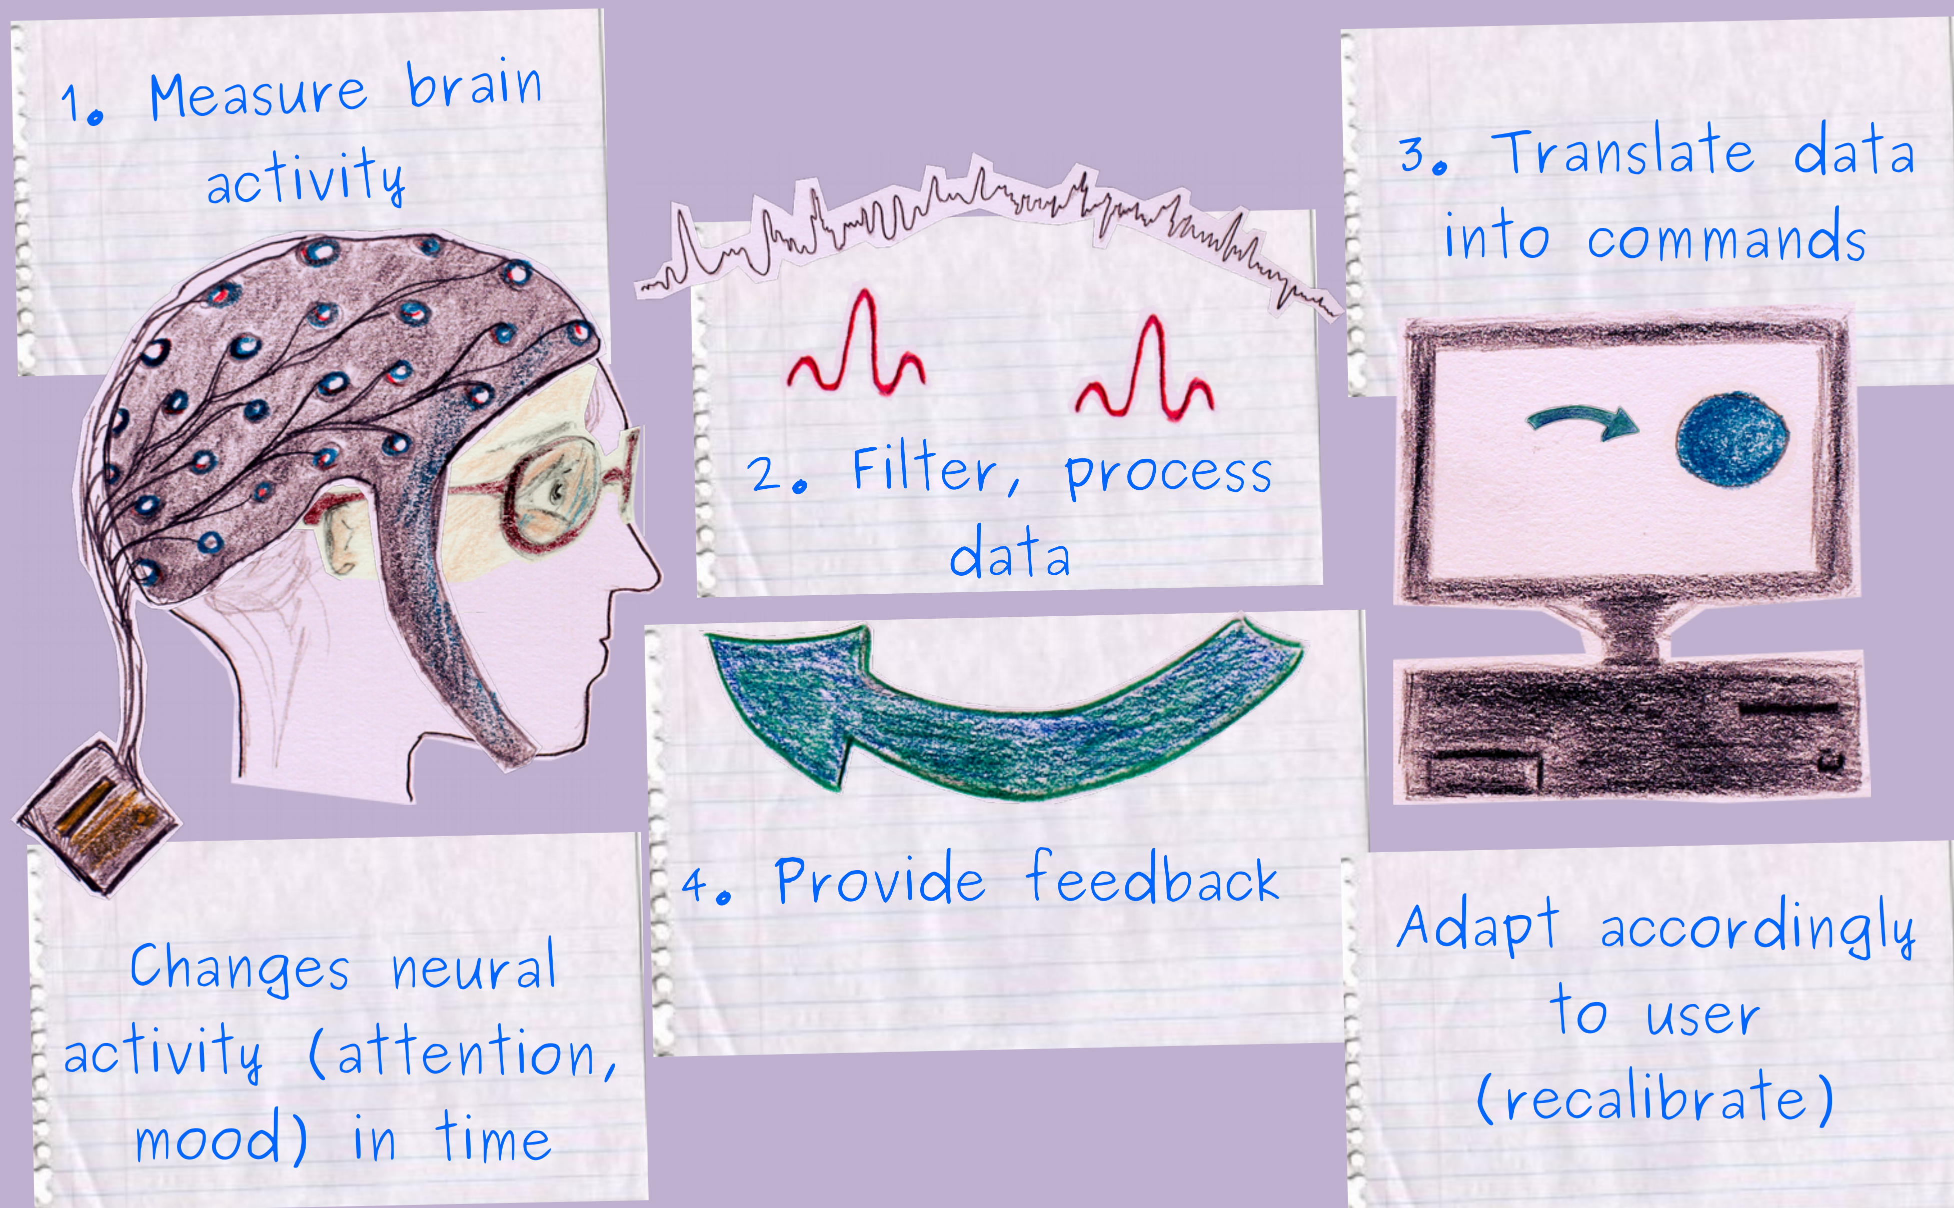

Brain Computer Interface
ADAPTIVE

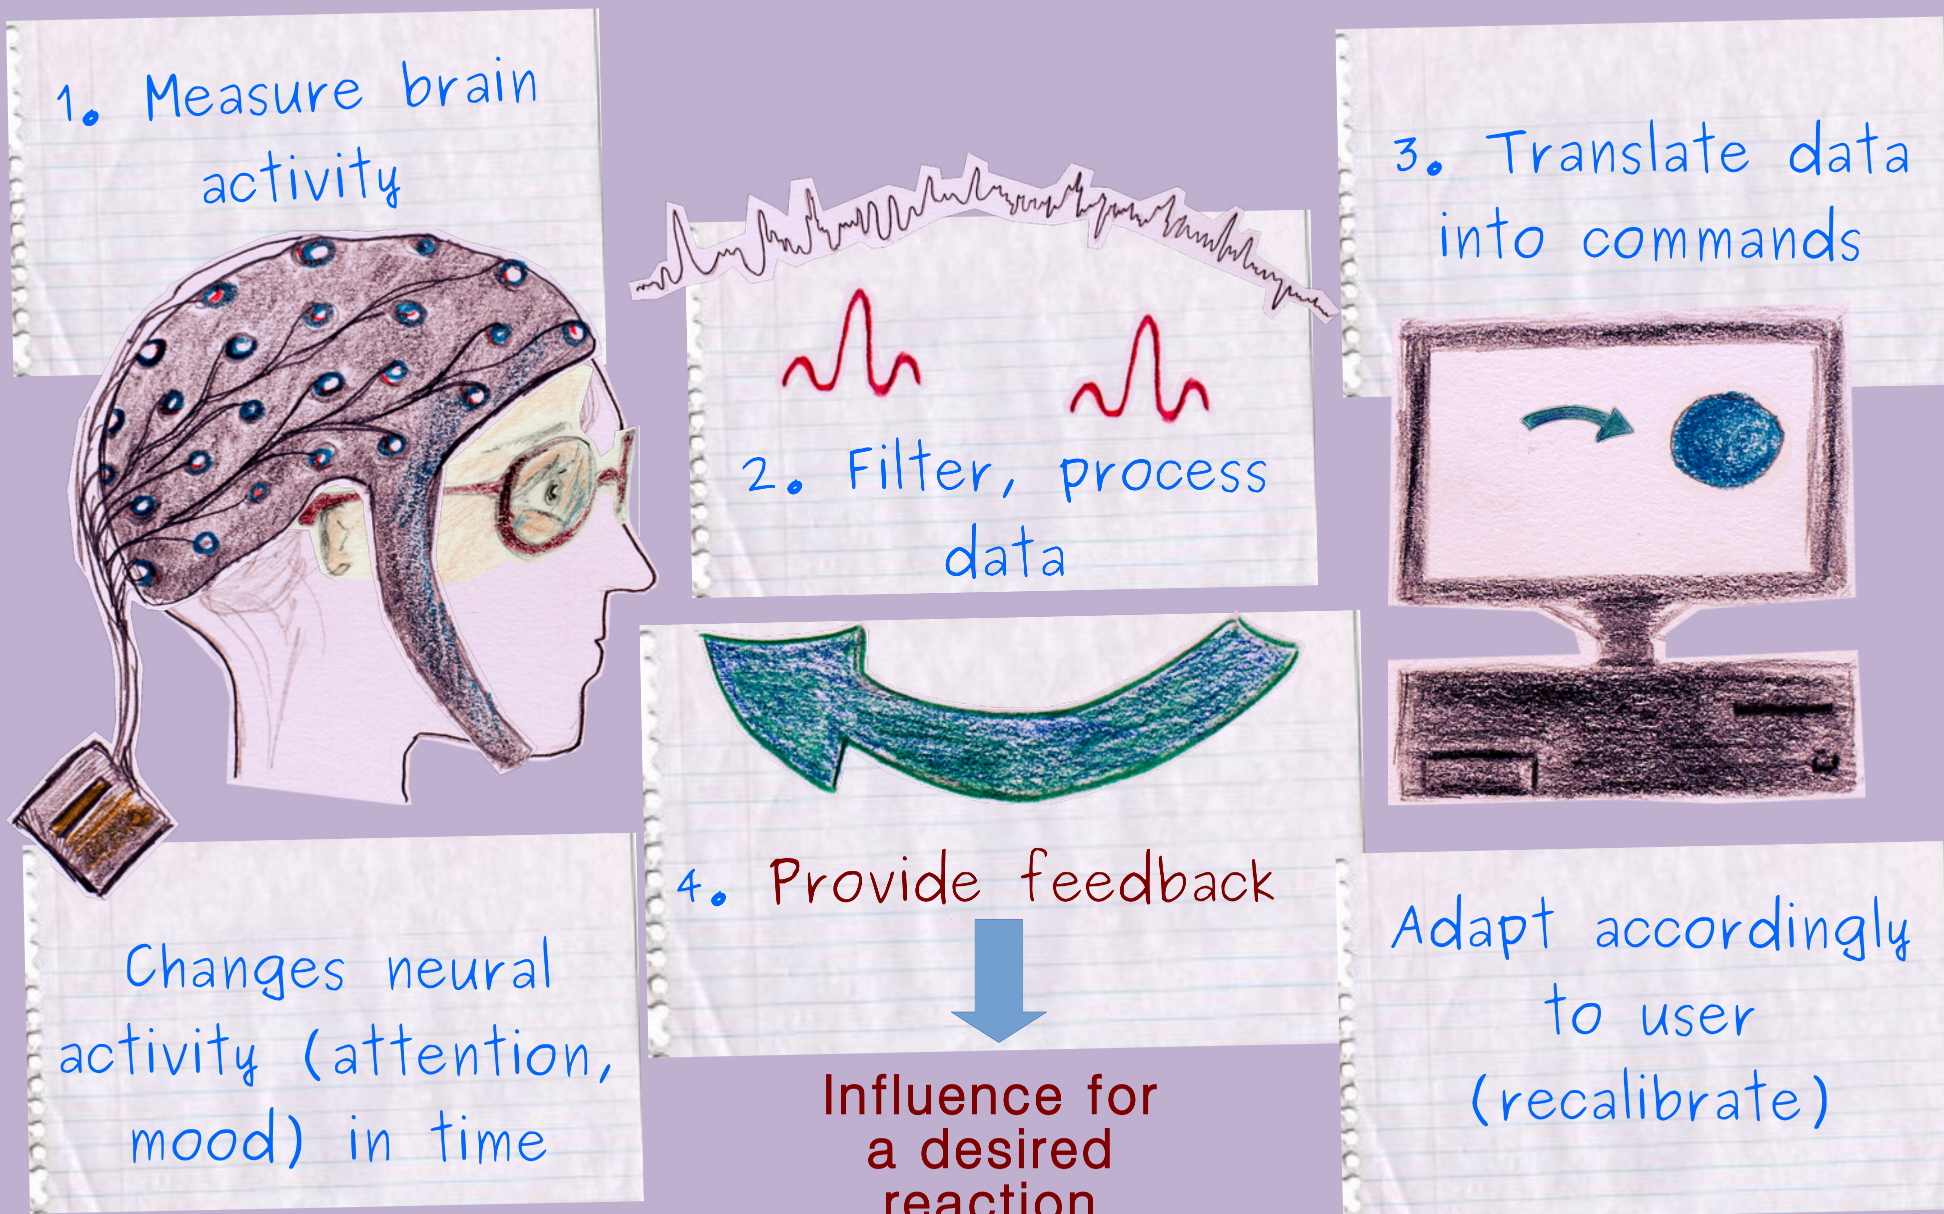

Unifying, Generic Framework (Conceptual)

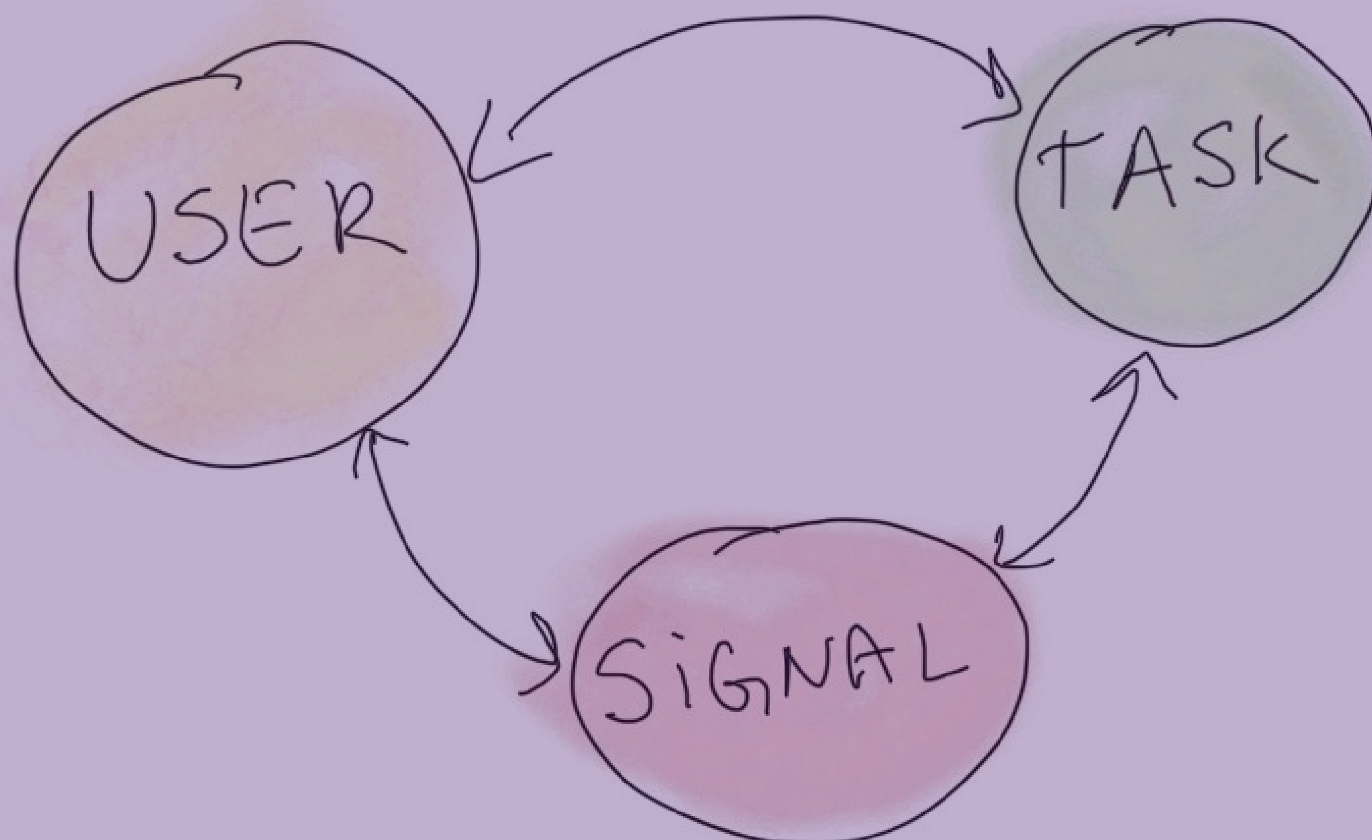

Unifying, Generic Framework (Conceptual)

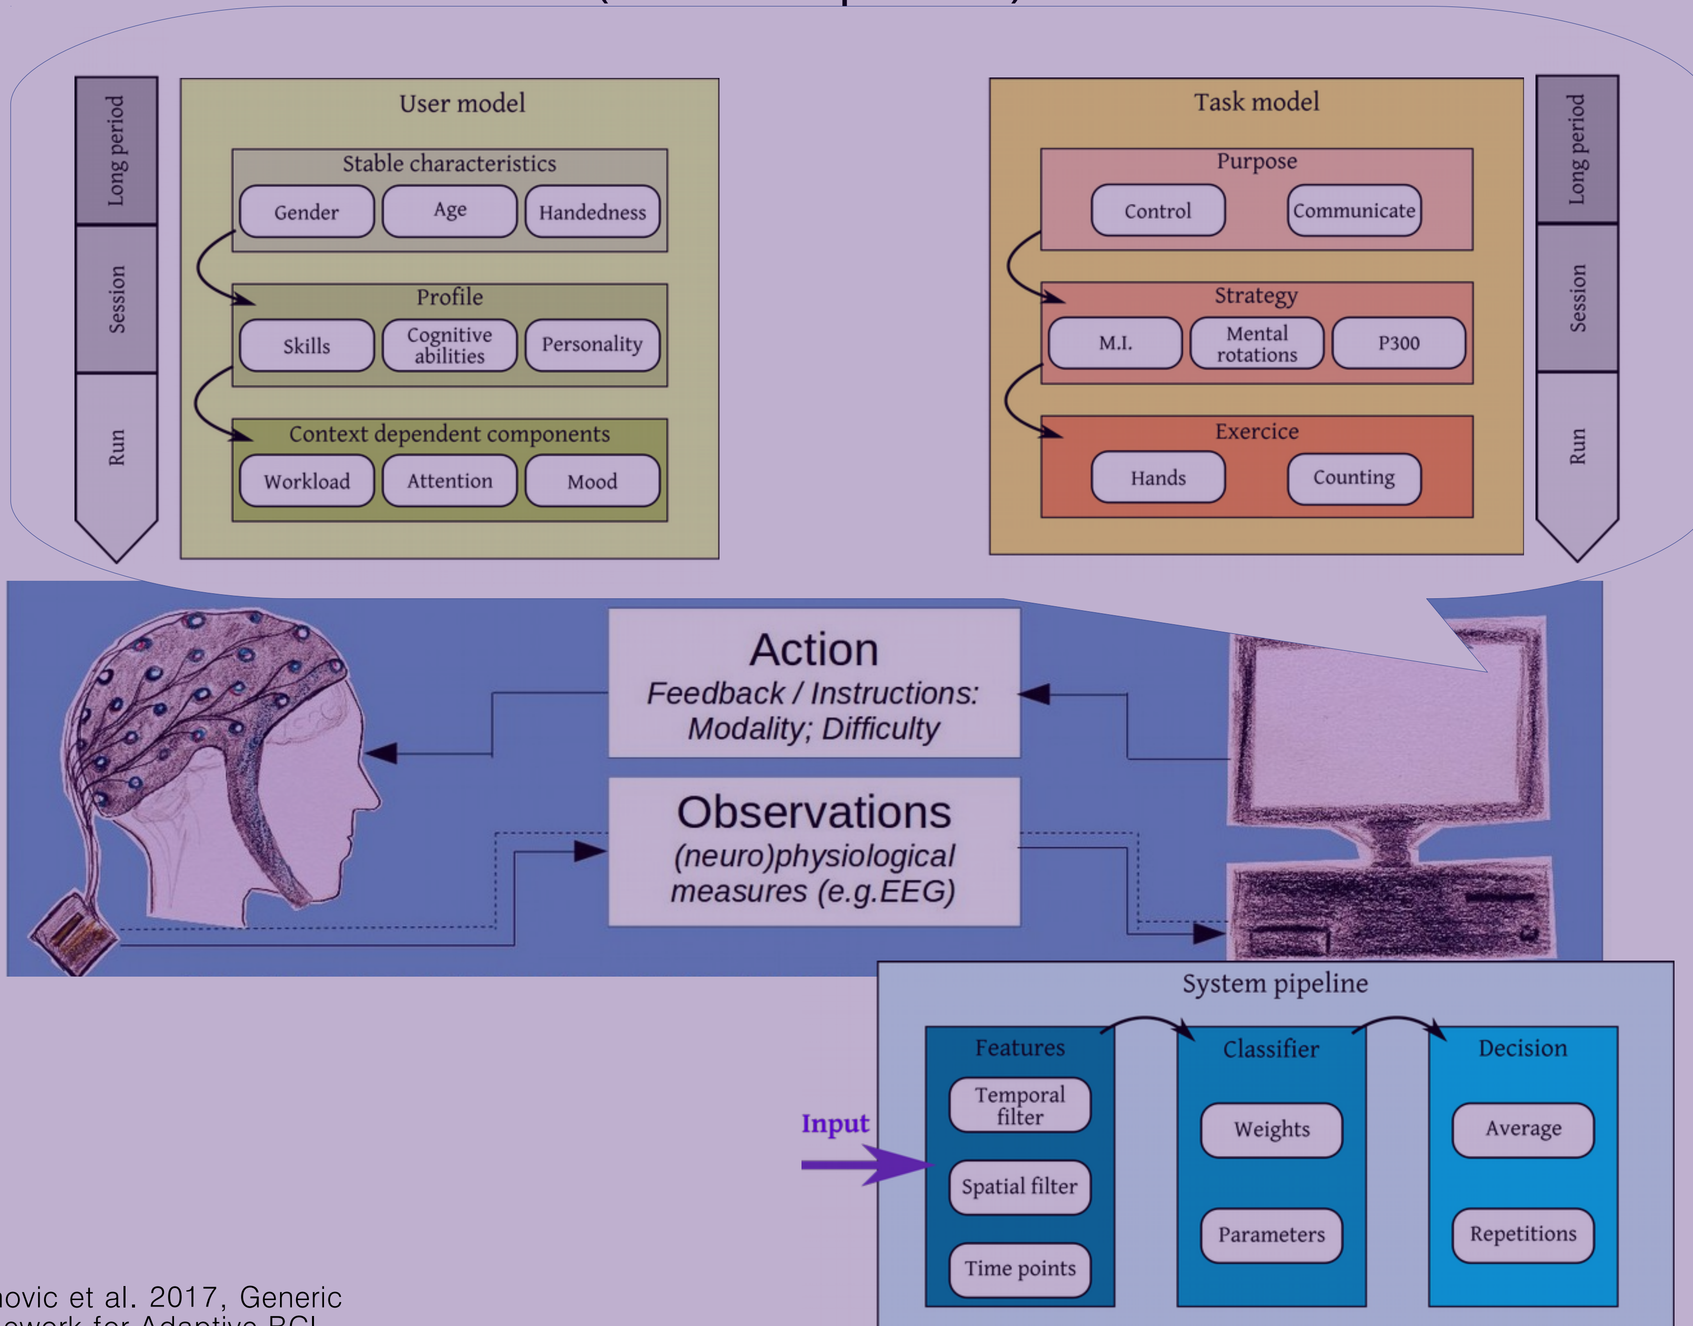

Mladenovic et al. 2017, Generic Framework for Adaptive BCI

Unifying, Generic Framework (Computational)

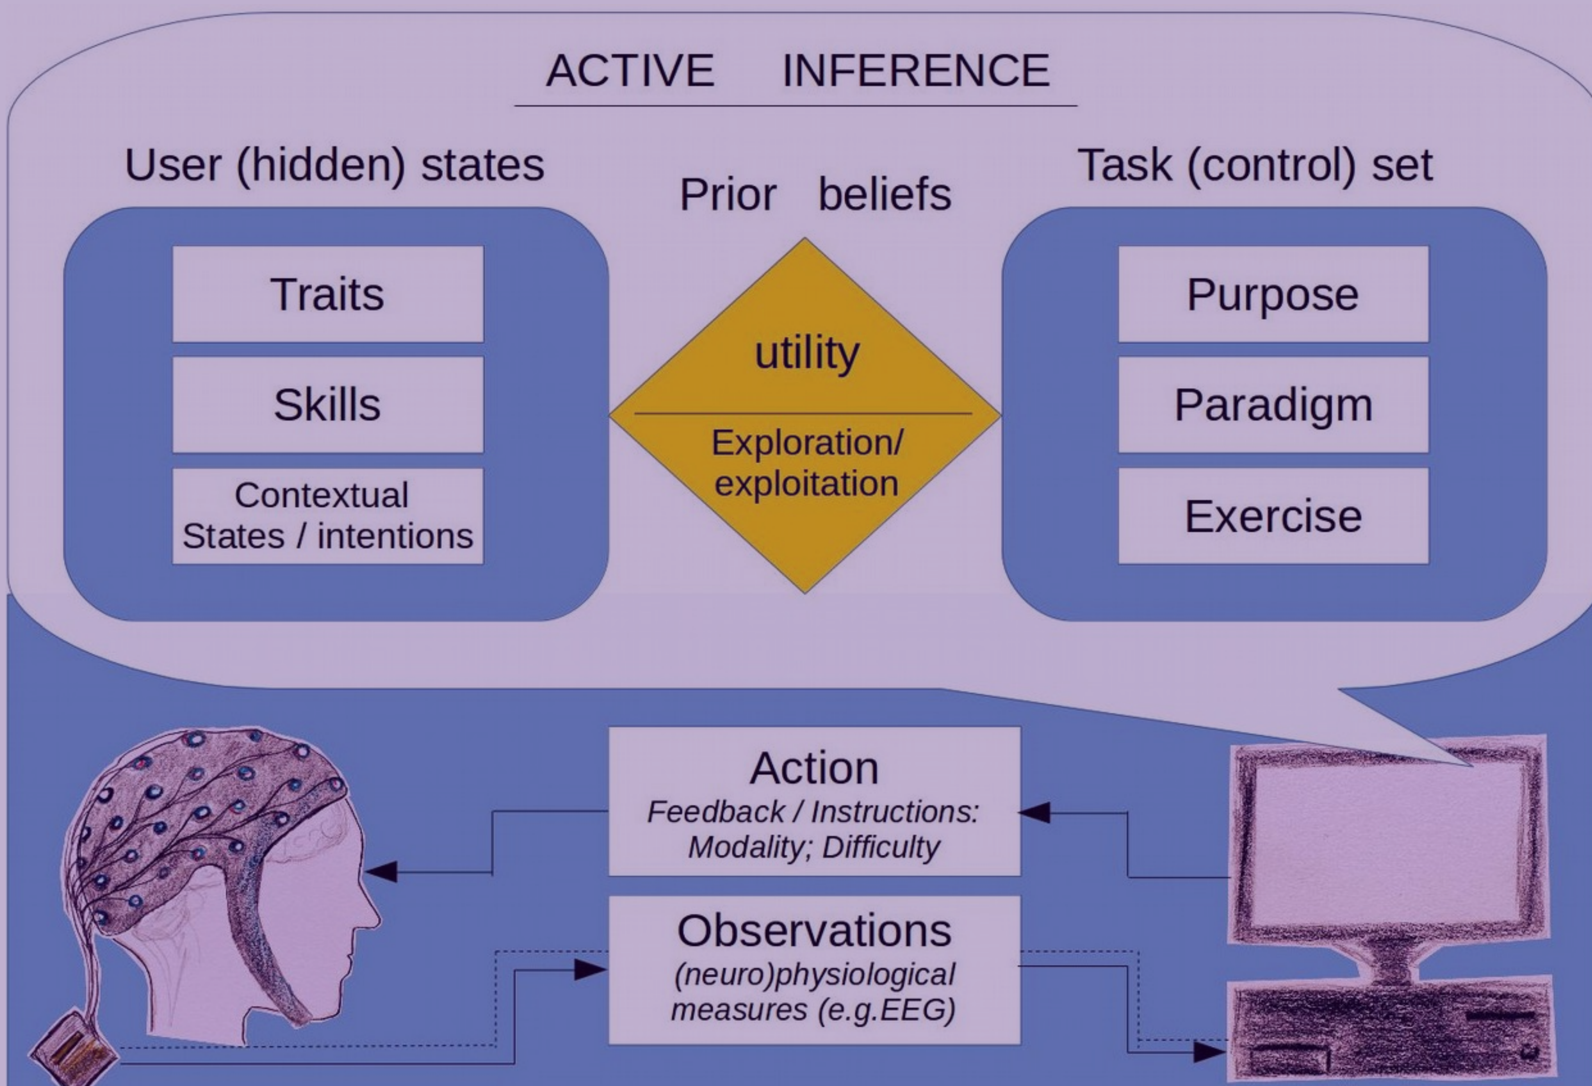

Mladenovic et al. 2018 Active Inference for Adaptive BCI: Application to the P300 Speller

Unifying, Generic Framework

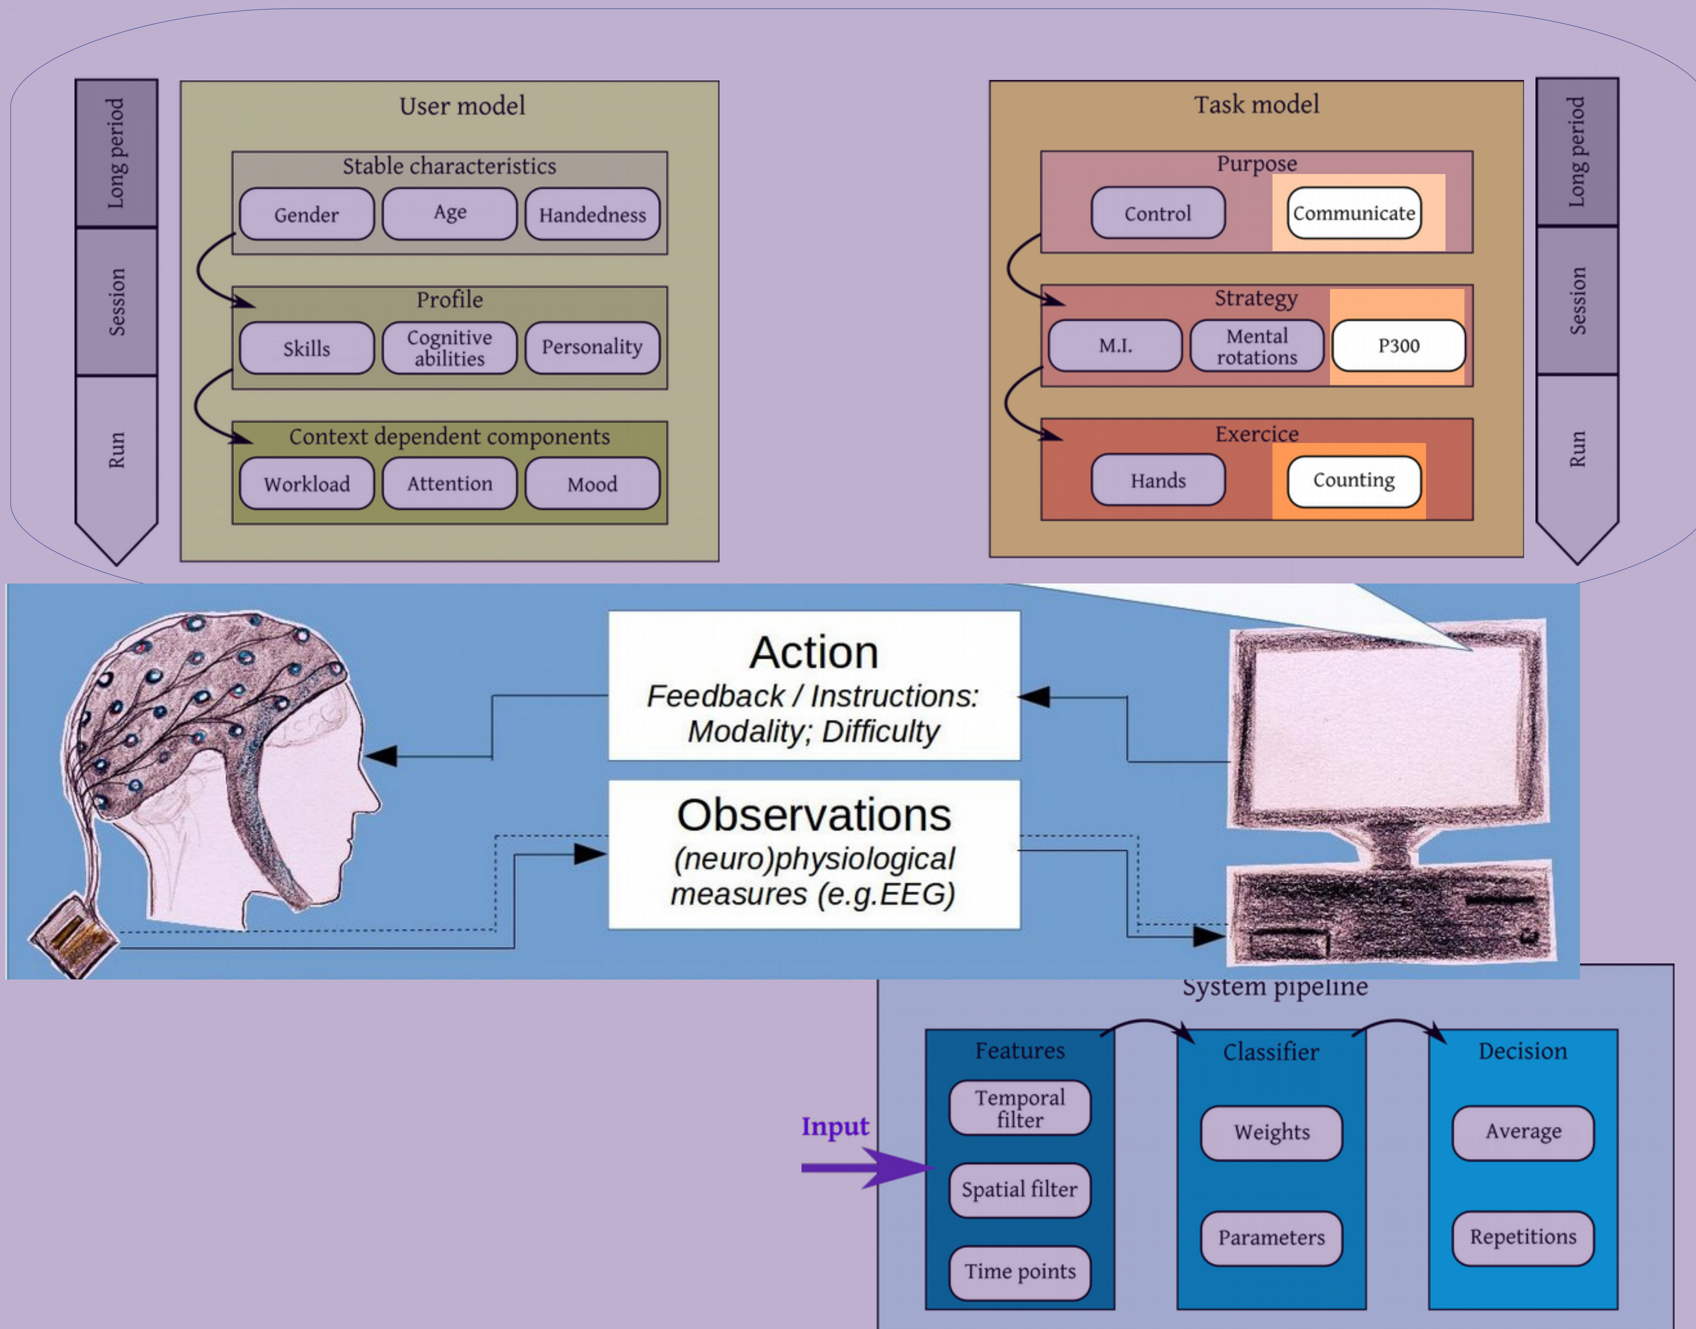

Principle
(P300 Speller)

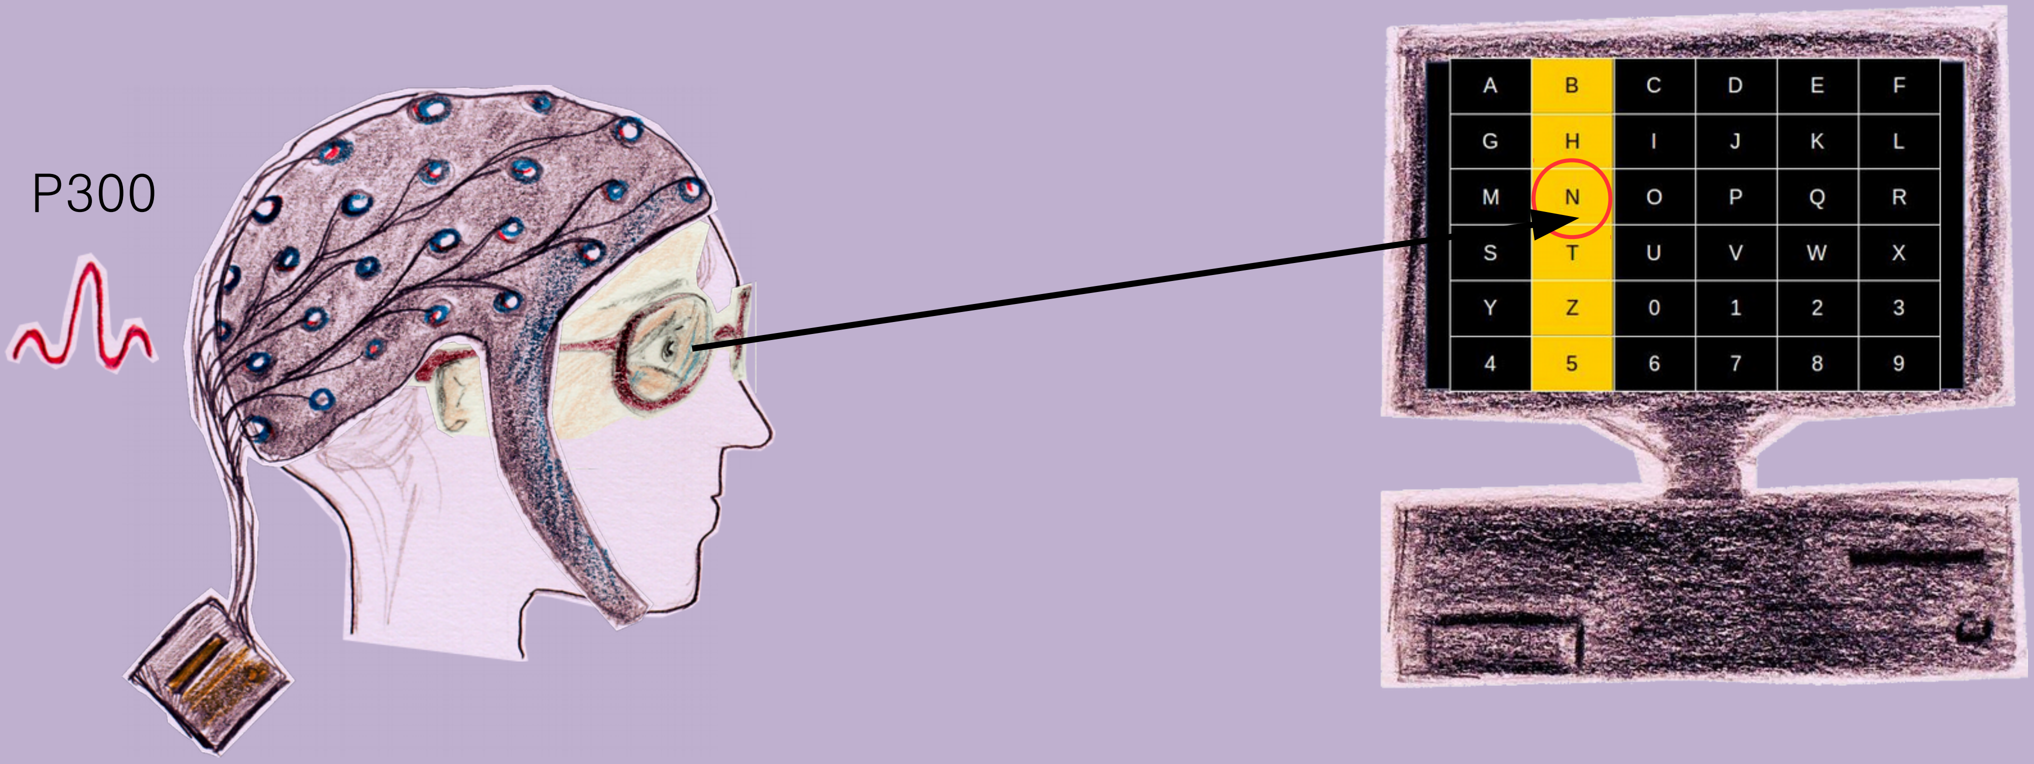

Principle
(P300 Speller)

P300

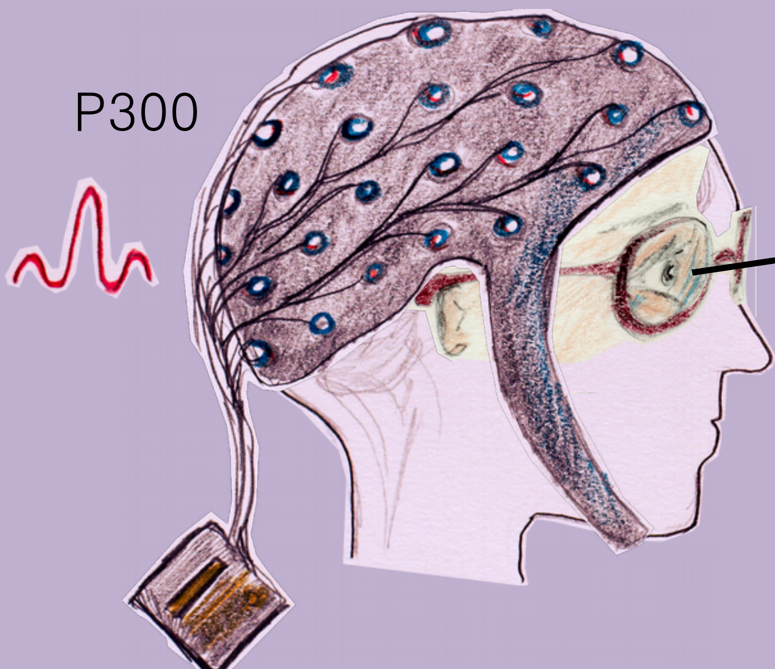

1. Spell after n number of flashes

2. spell when enough evidence has been accumulated

optimal stopping

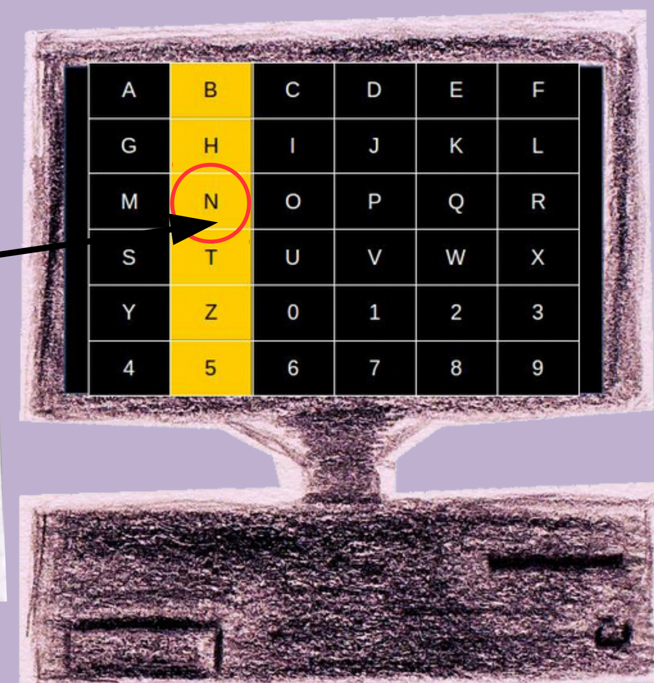

Principle
(P300 Speller)

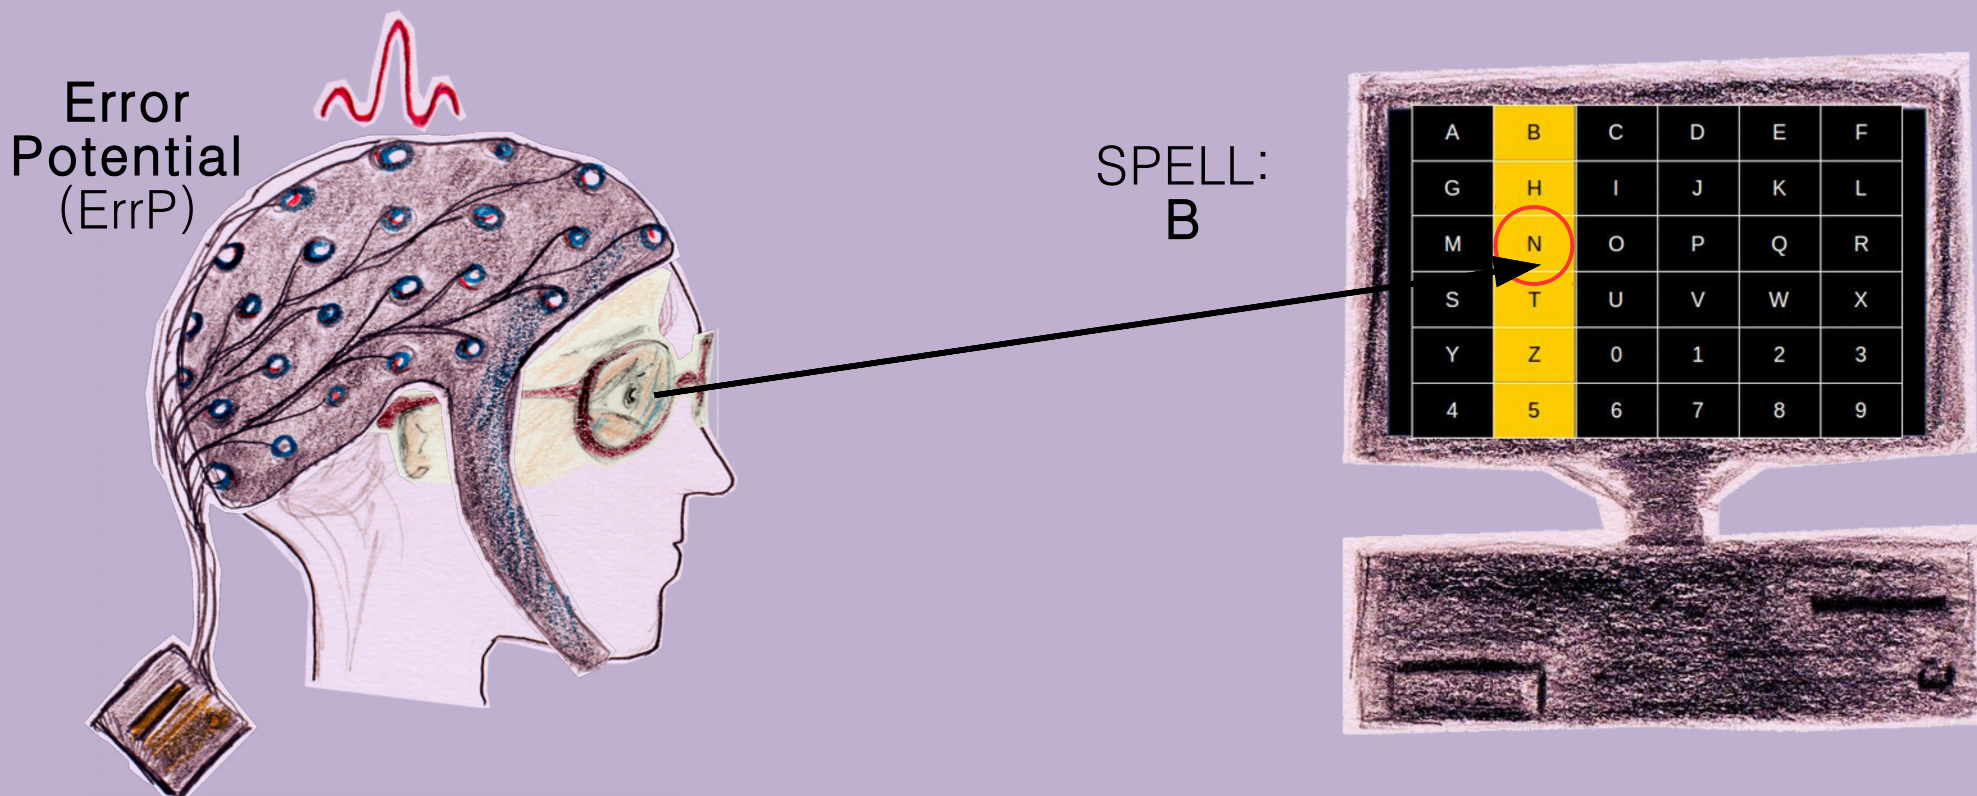

Unifying, Generic Framework
(for P300)

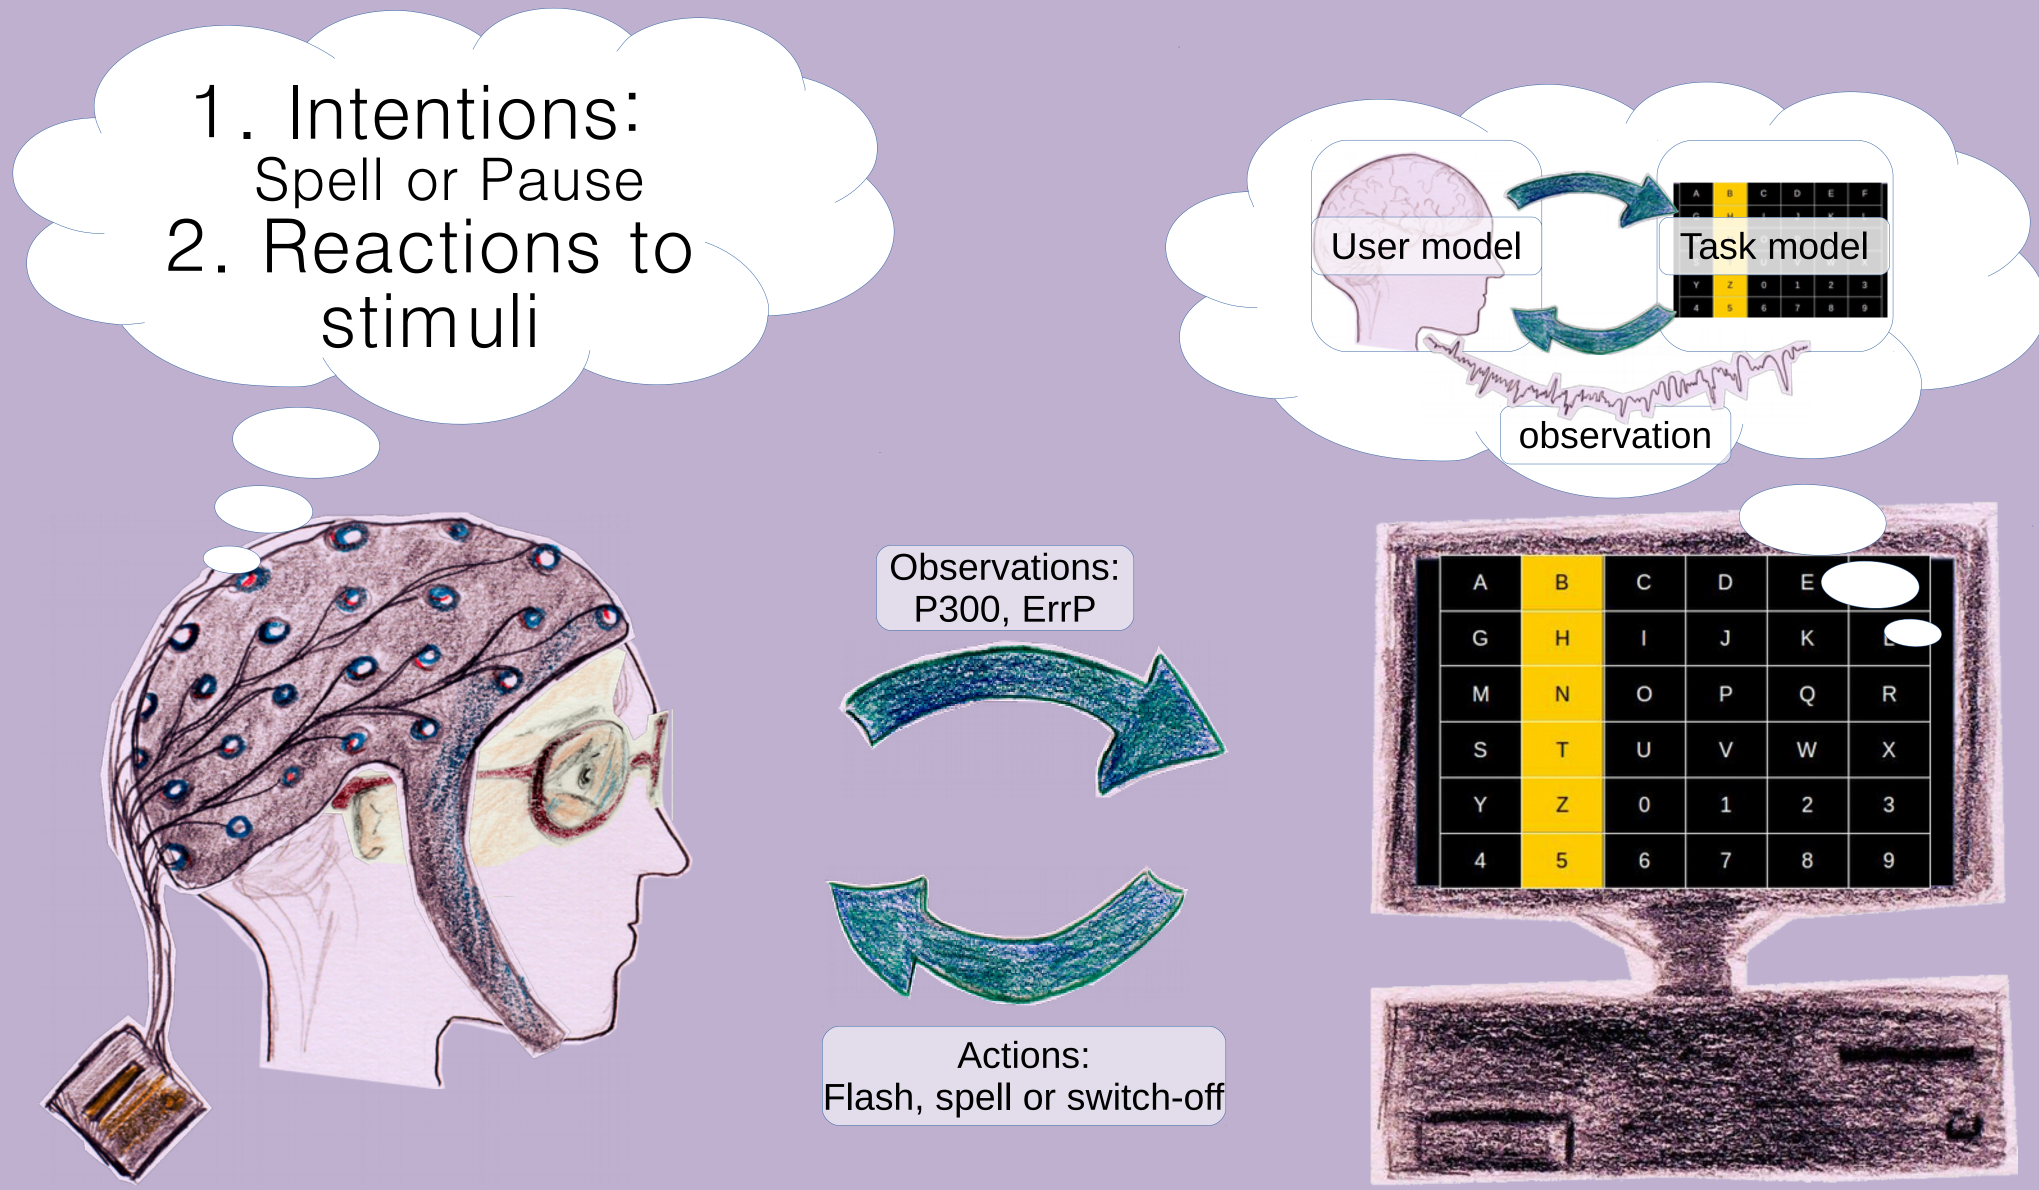

Unifying, Generic Framework
(Computational, P300)

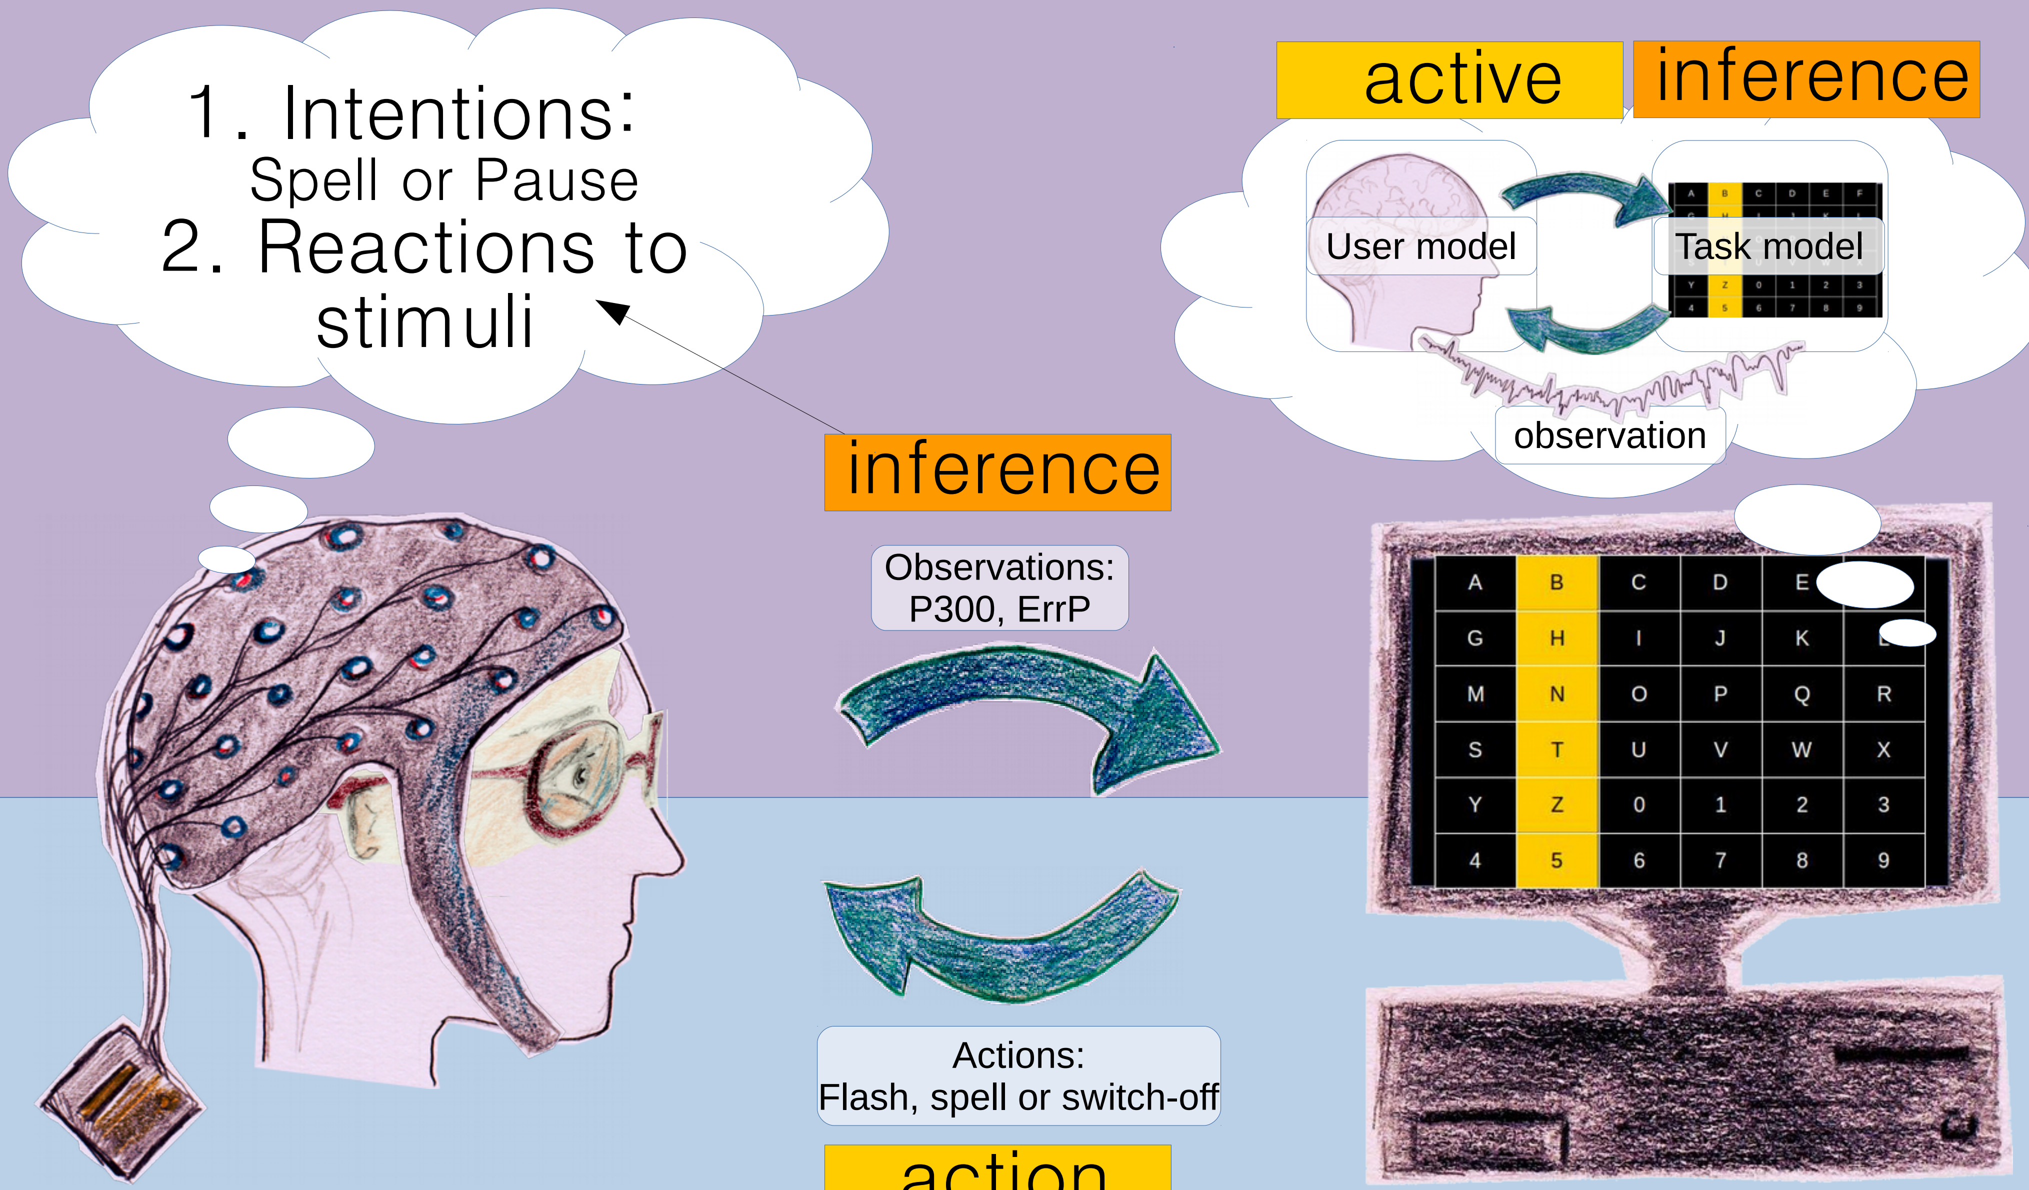

Active Inference
(Optimal stopping & flashing)

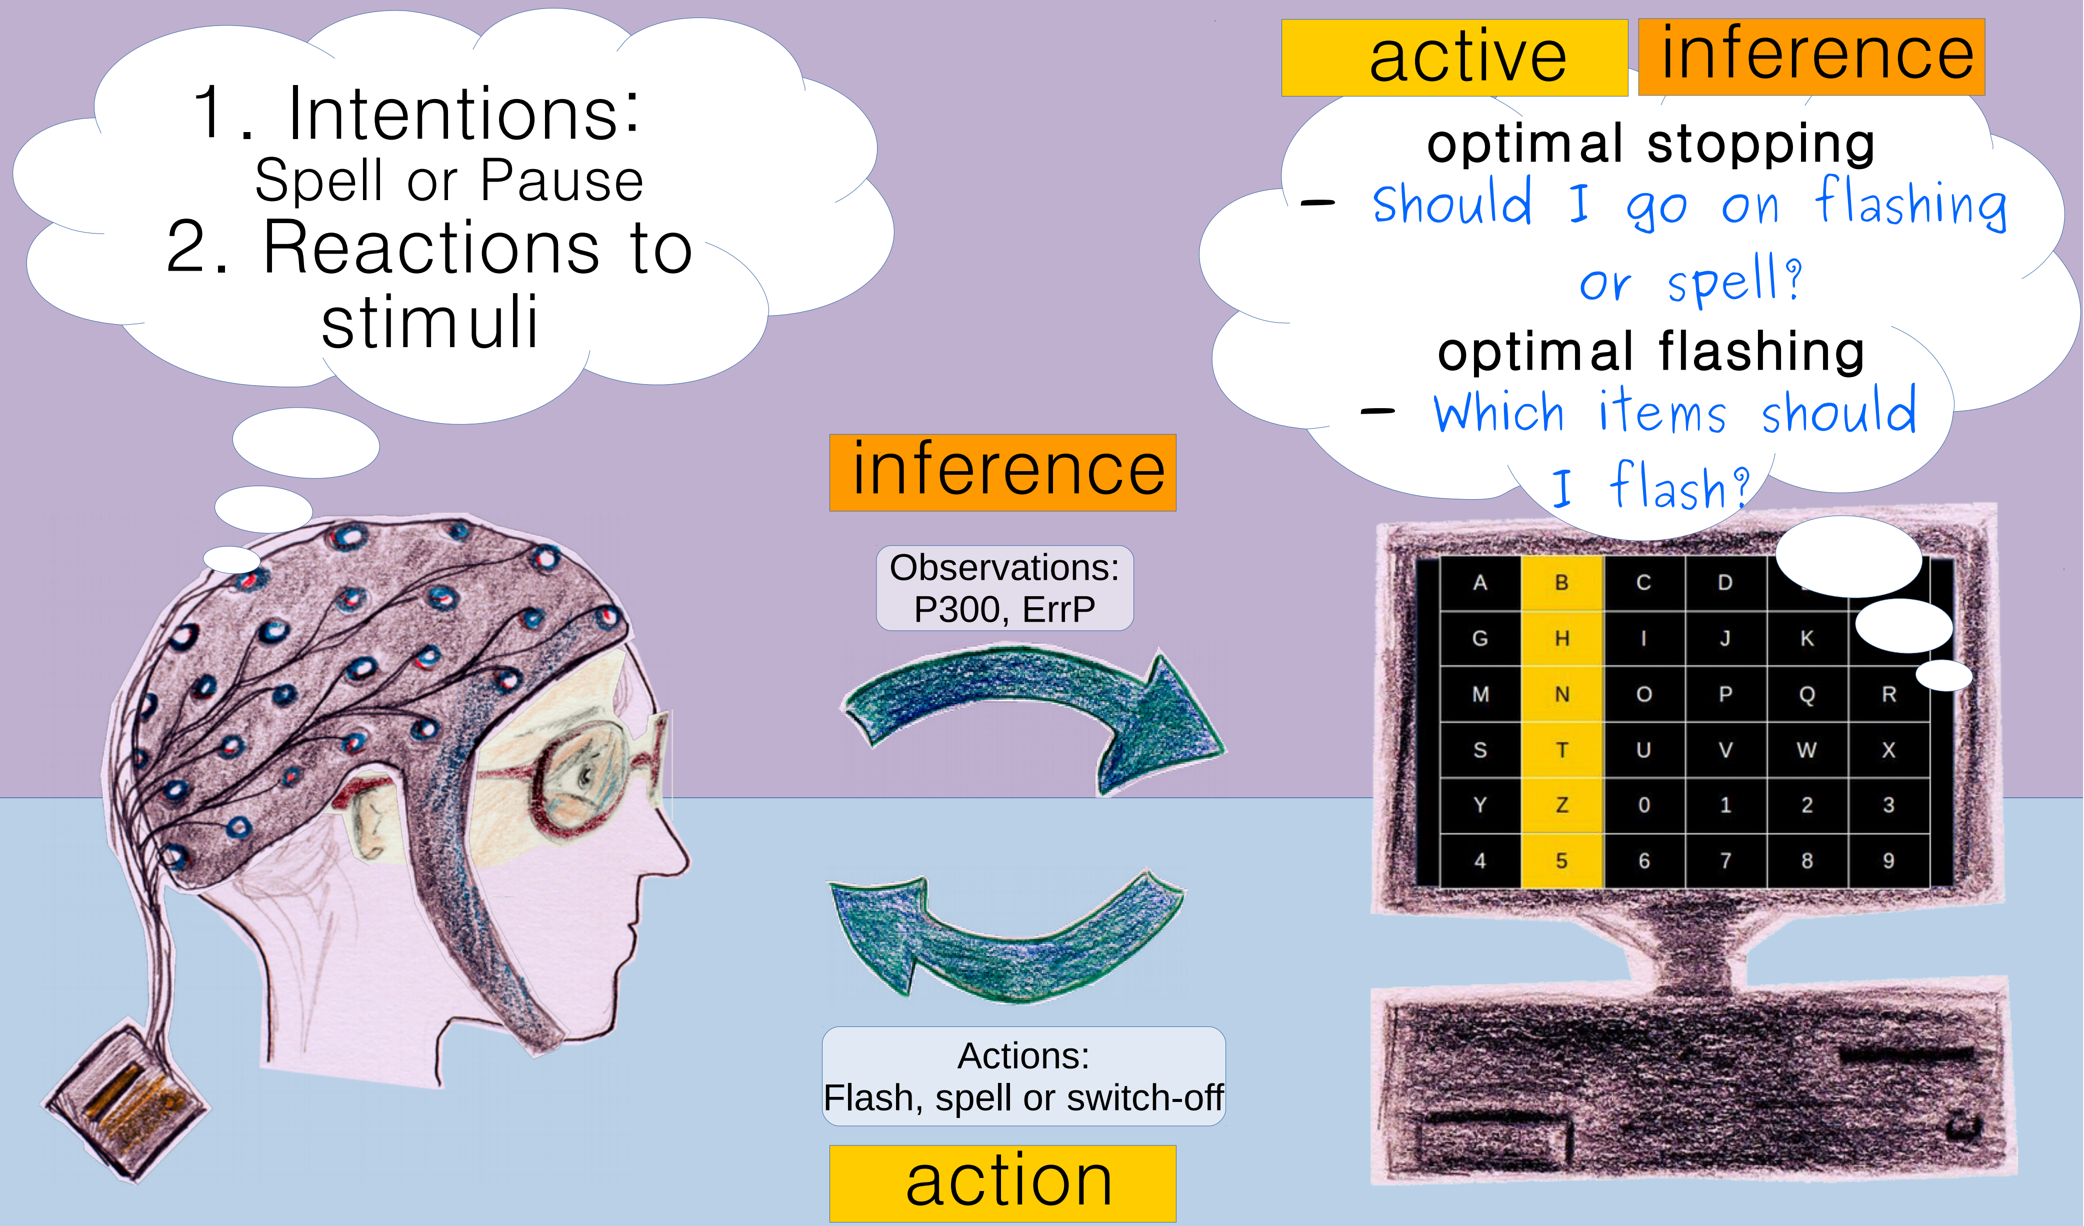

Active Inference:
The Origins
Karl Friston

Nature's tendency to disorder – increase of entropy.

Active Inference:
The Origins

Nature's tendency to disorder –
increase of entropy.

Living beings resist dispersion –
minimize entropy,
and limit themselves into a
finite number of states!

Active Inference:
The Origins

Markov model

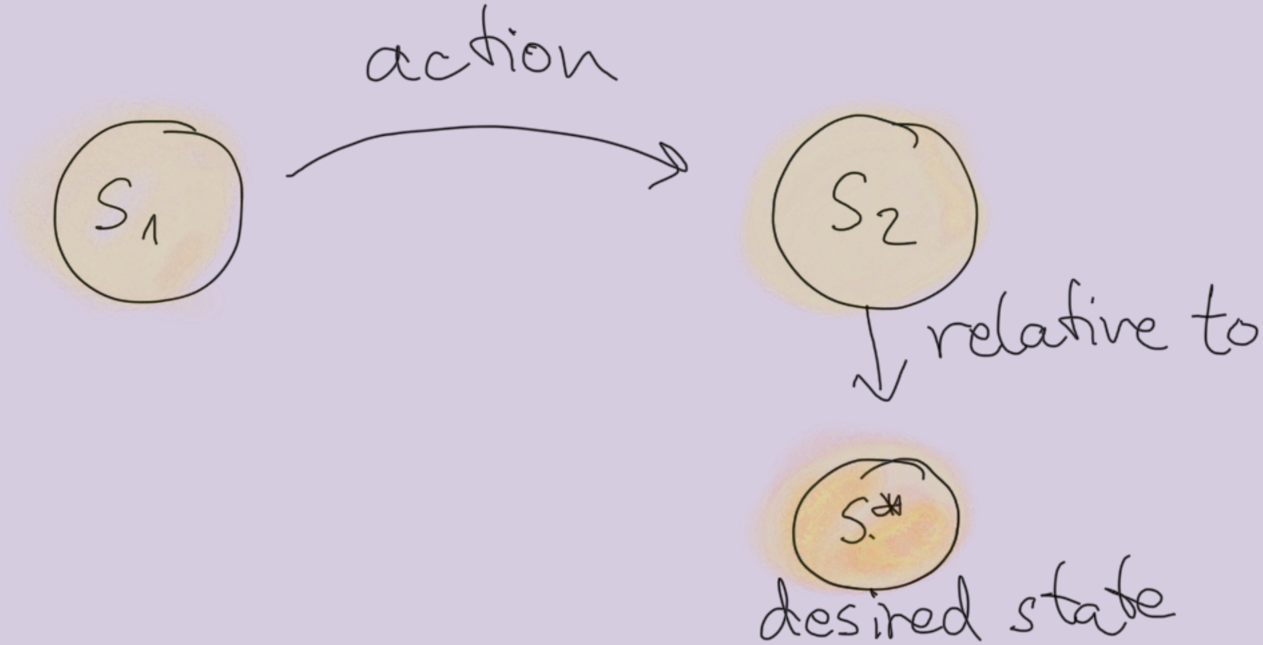

Passing from one state to another is possible through performing an action. The choice of future state depends on min of the relative entropy – relative to a desired state!

Active Inference: The Origins

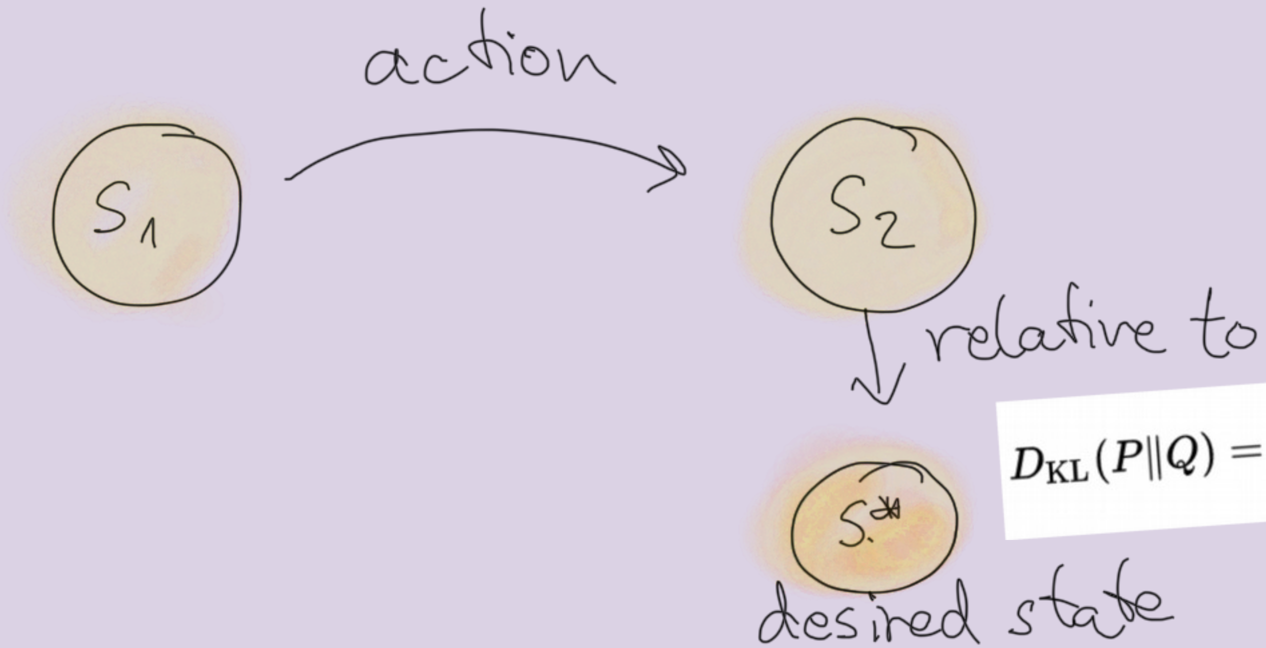

$$D_{\mathrm{KL}}(P\|Q) = \sum_i P(i)\,\log\frac{P(i)}{Q(i)}.$$

Relative entropy
Or
KL Divergence

Active Inference:
The Origins

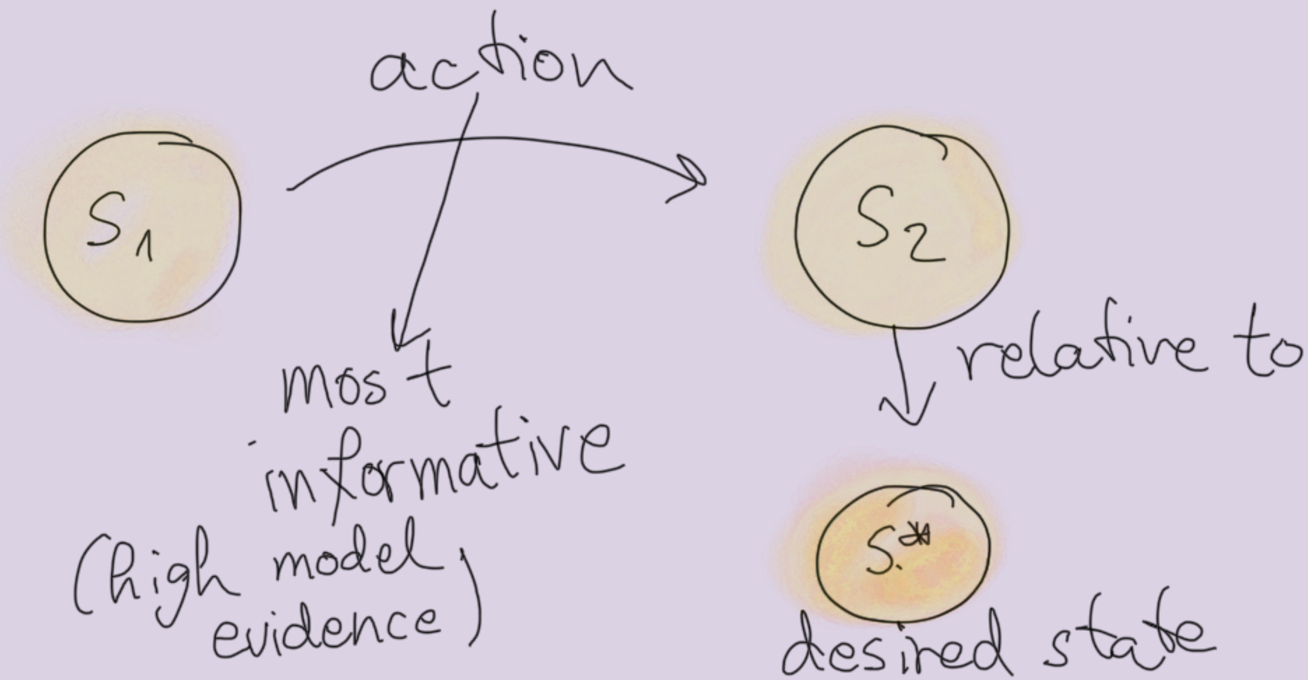

min surprise!

Active Inference:
The Origins

Bayesian brain!
Internal model of the environment.
A Bayesian inference (updating certainty of the environment with new observations)

Bayes Theorem (reminder)

$$P(A|B) = \frac{P(A)P(B|A)}{P(B)}$$

is it from a fire?

you see smoke.

Bayes Theorem
(reminder)

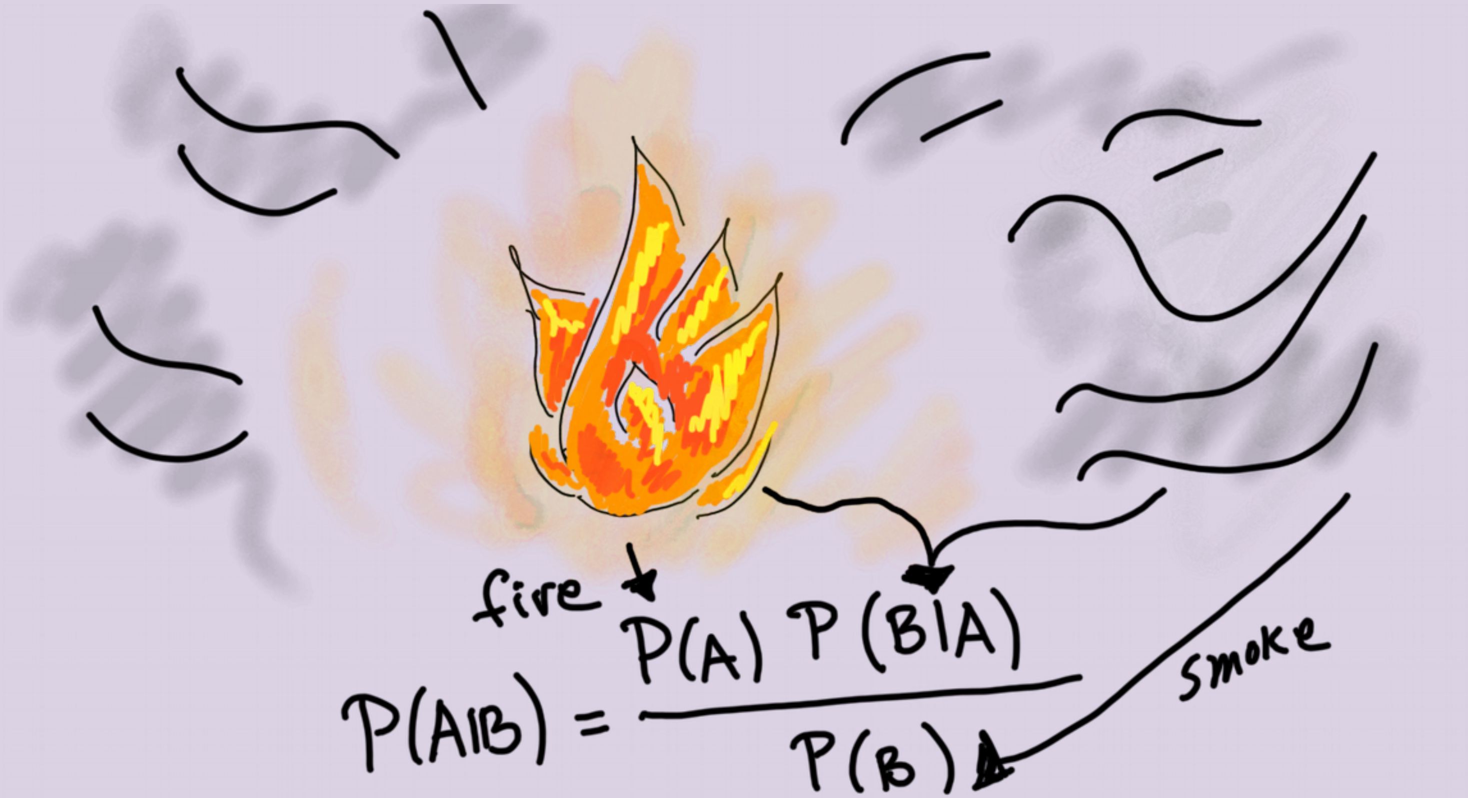

P(Fire|Smoke) = P(Fire) P(Smoke|Fire) / P(Smoke)
= 1% x 90% / 10% = 9%

Bayes Theorem
(reminder)

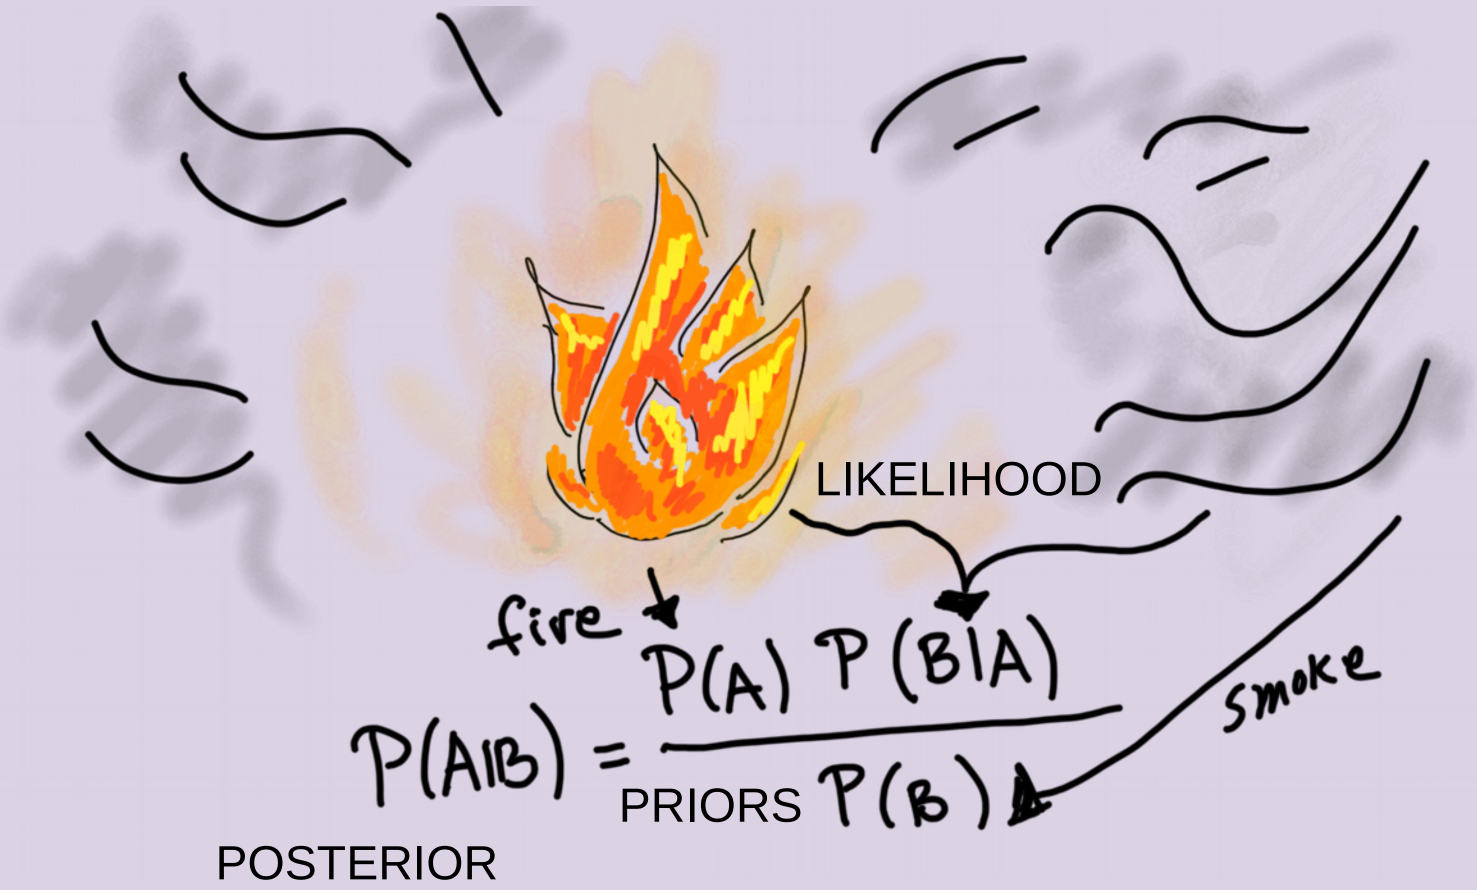

Active Inference:
The Origins

Bayesian brain!
Internal model of the environment.
A Bayesian inference (updating certainty of the environment with new observations)

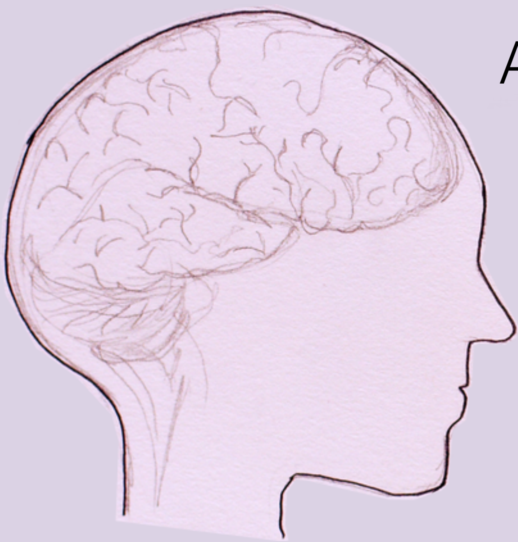

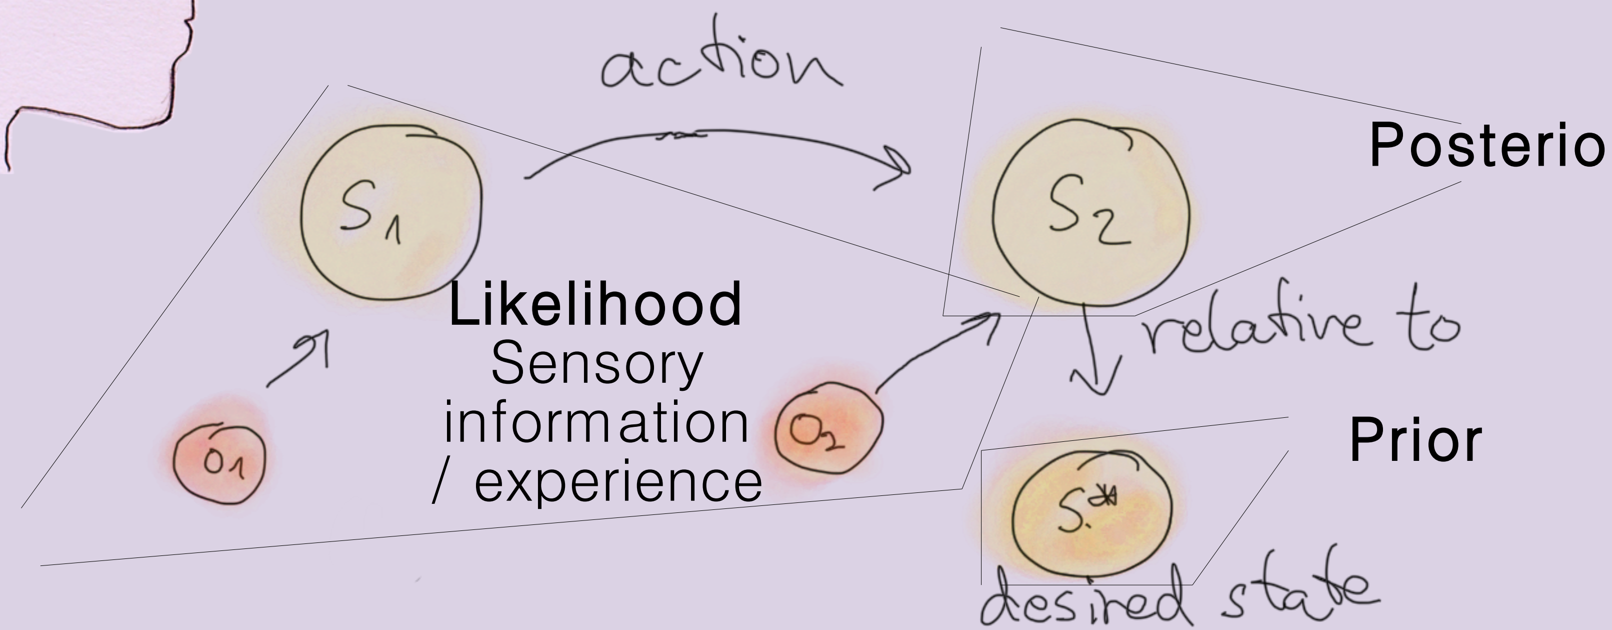

Active Inference:
The Origins

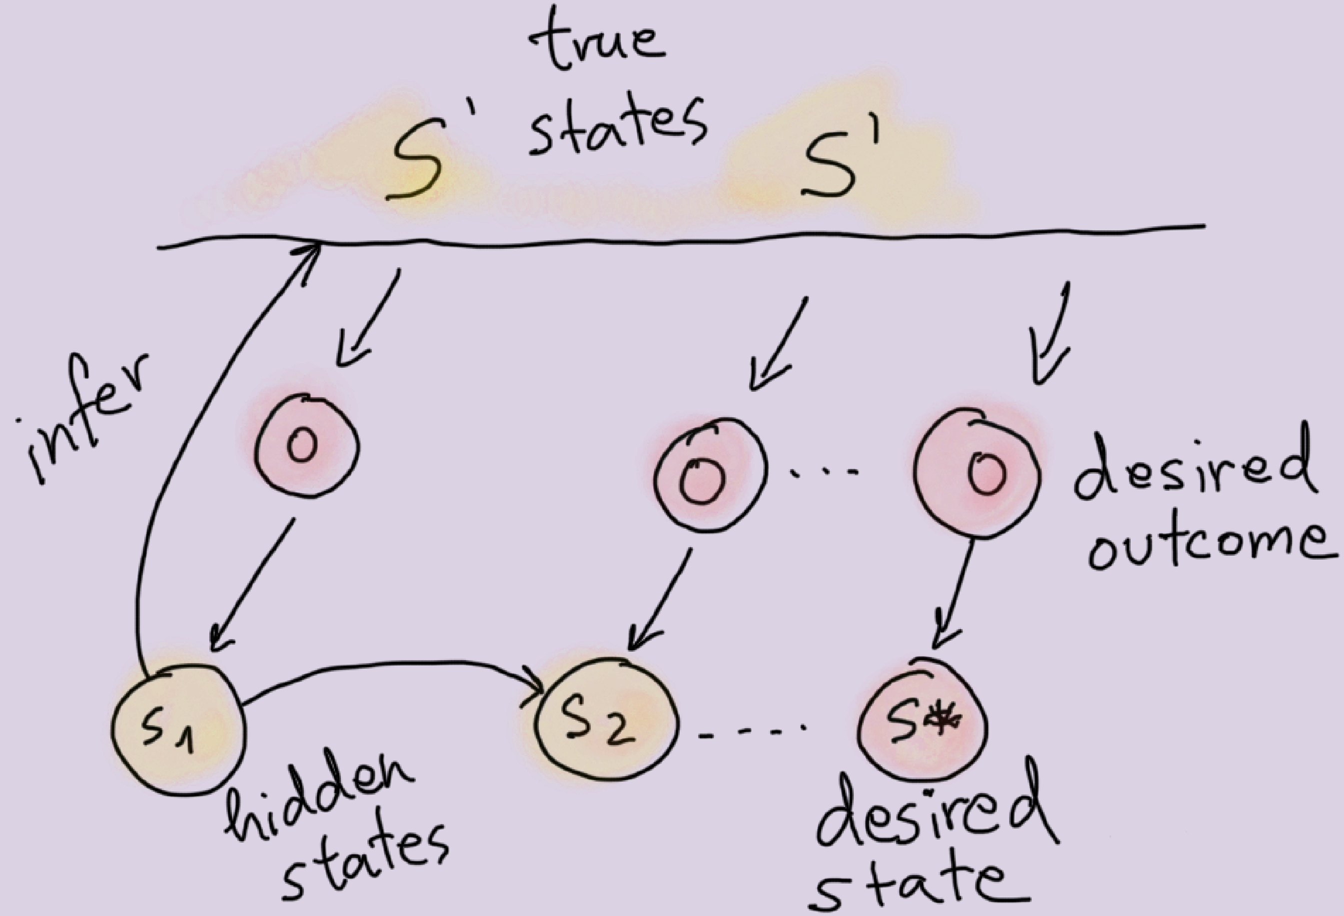

One cannot know its true states or the true states of the world
but only infer it through sensory observations

Active Inference: The Origins

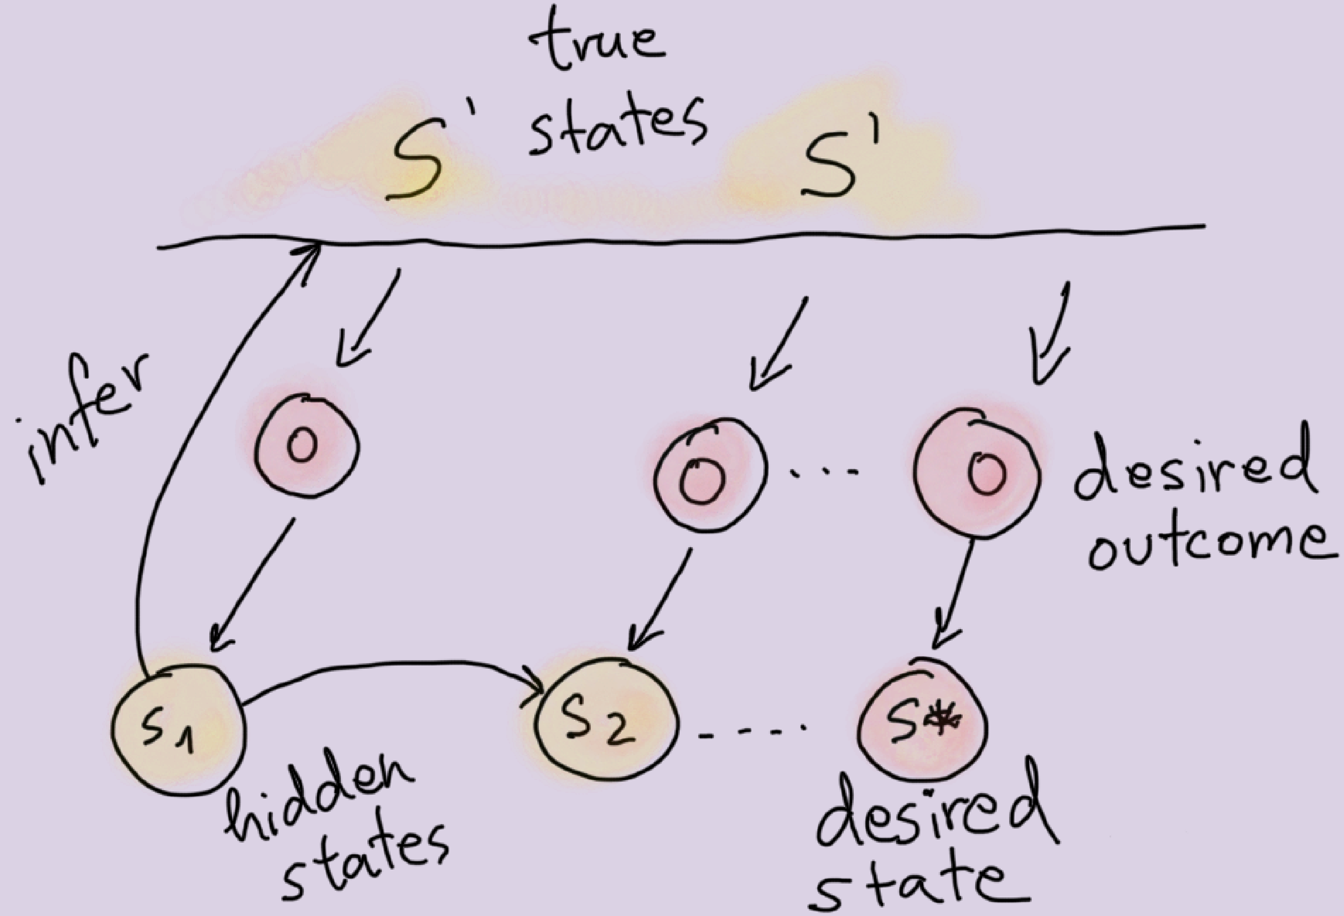

Mapping between states and observations is diffeomorphic.
Entropy over hidden states is bound by entropy over observations.

Active Inference:
The Origins

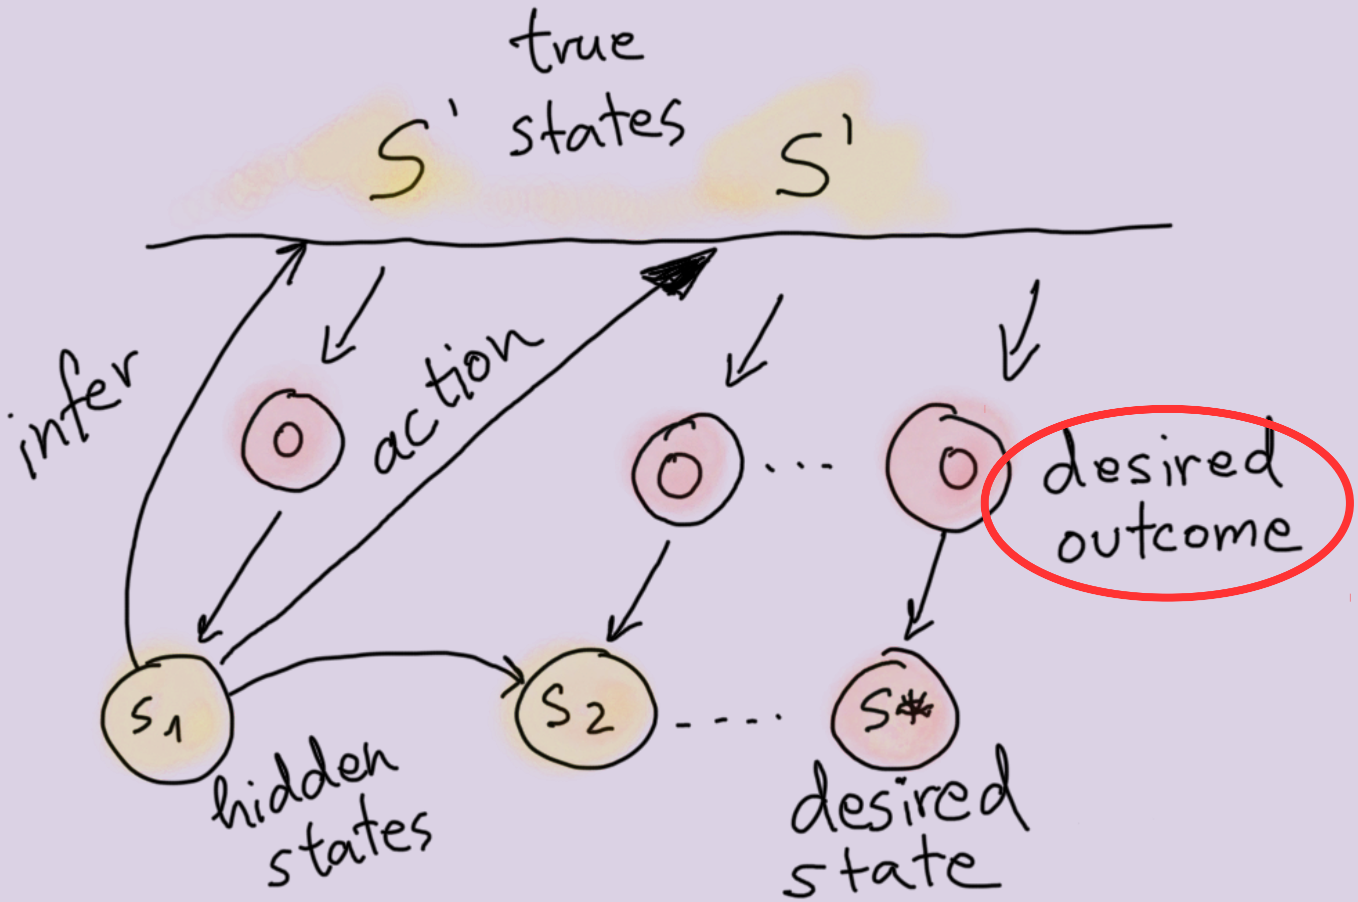

To reach the desired state, one can update their model of the world or perform action and influence the world.

Active Inference:
The Origins

Predict Posterior is difficult
without knowing the Prior.

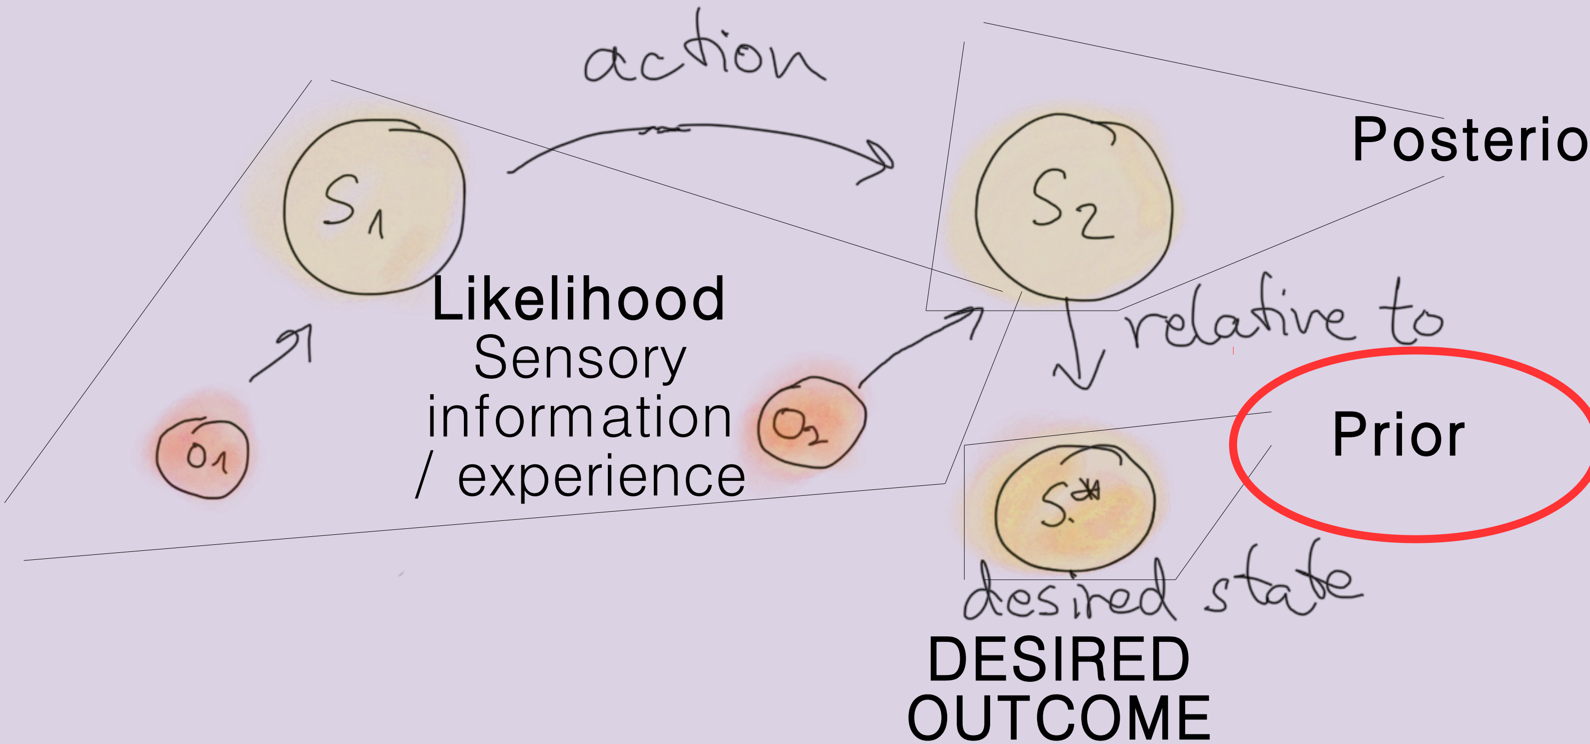

Active Inference:
The Origins

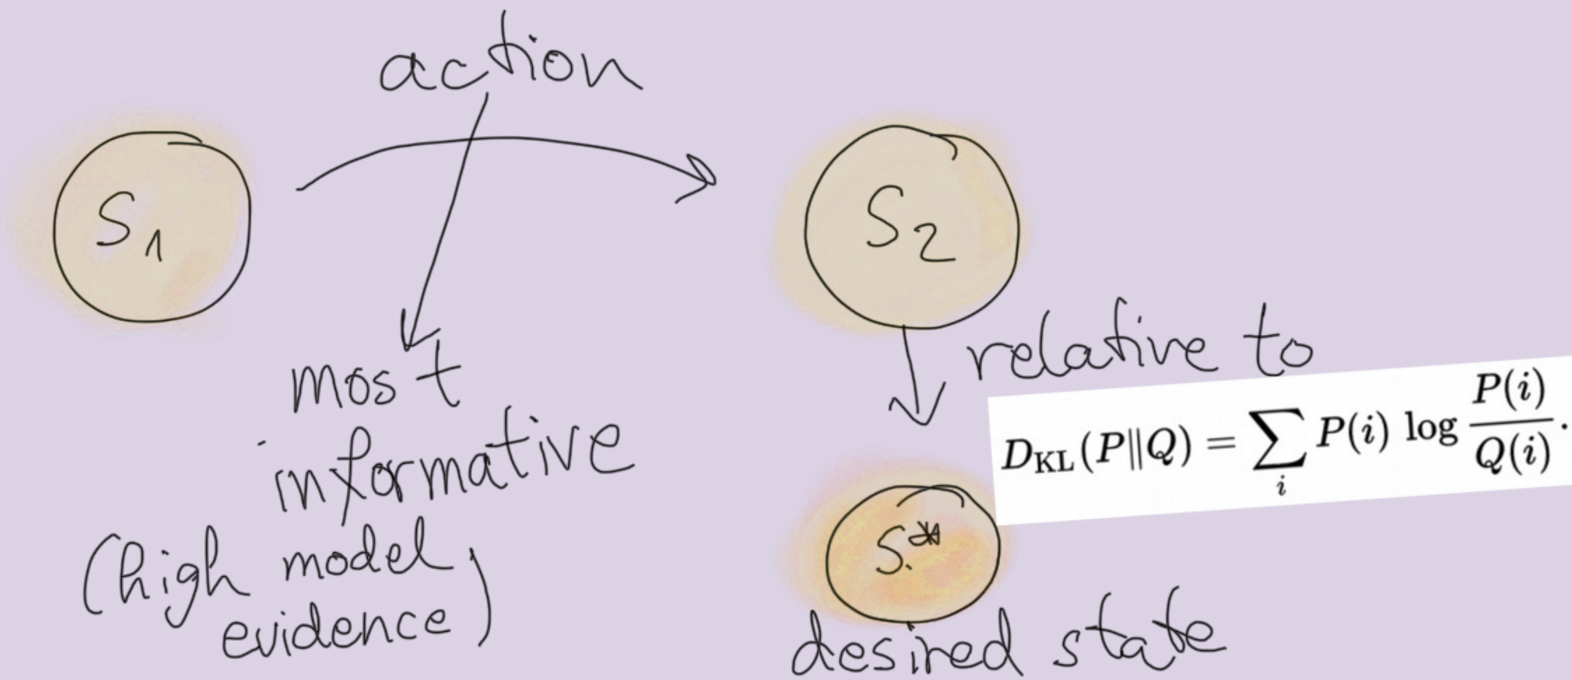

min surprise!

Active Inference:
The Origins

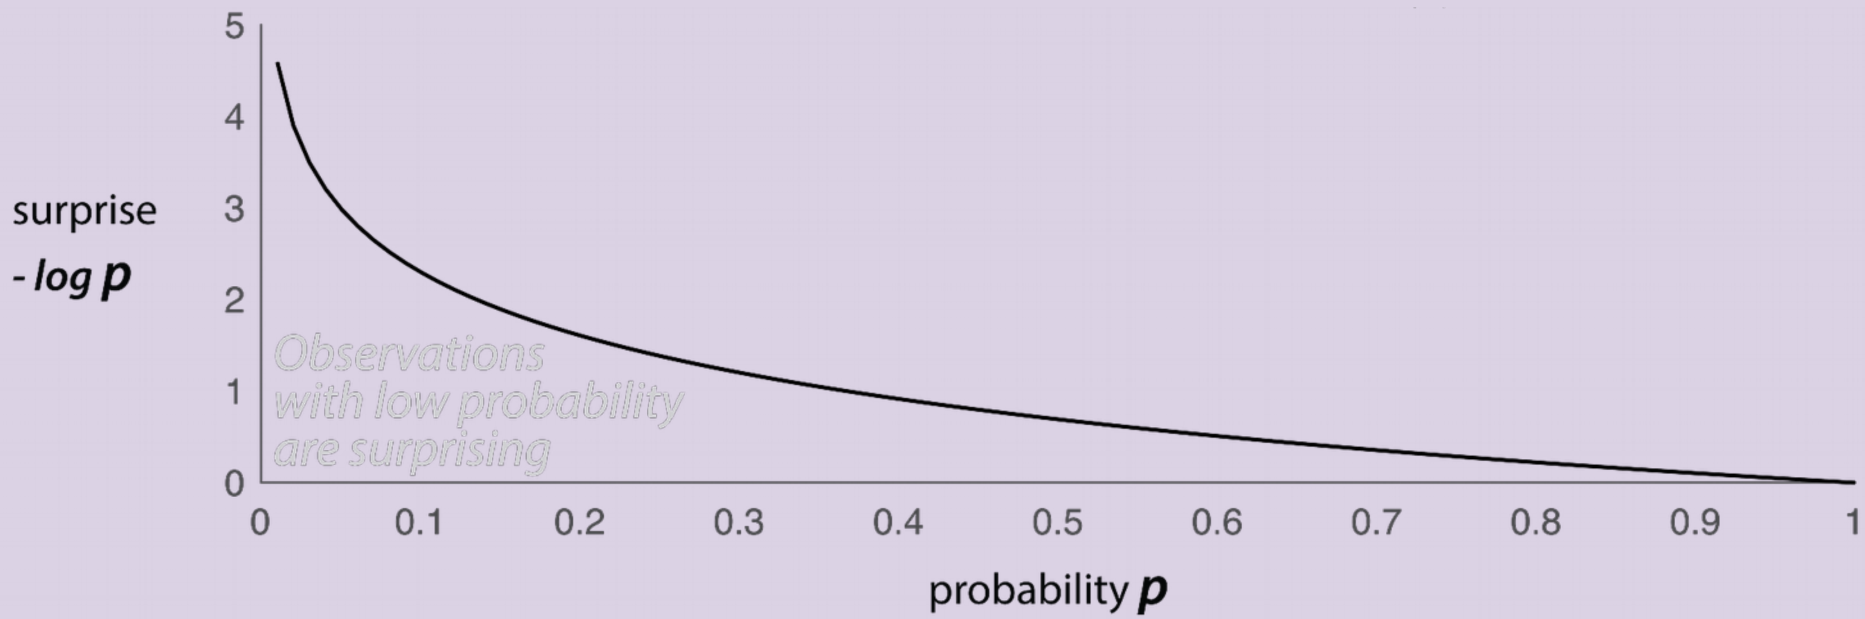

Courtesy of Oleg Solopchuk
(tutorial of Active Inference)

Active Inference:
The Origins

Maximize model evidence
logP(o)
= minimize surprise
−logP(o)

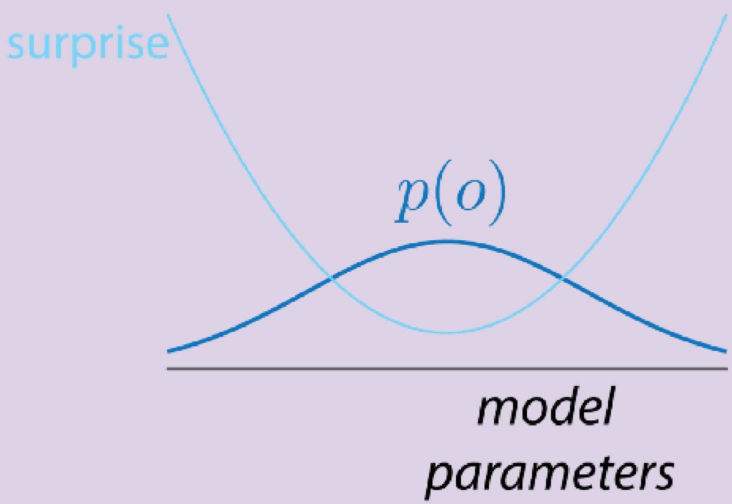

Courtesy of Oleg Solopchuk
(tutorial of Active Inference)

Active Inference: The Origins

Model evidence =
marginal likelihood
"Sum out" s from P(o, s)

$$-\log p(o) = -\log \sum_s p(o,s)$$

$$p(a) = \sum_b p(a,b)$$

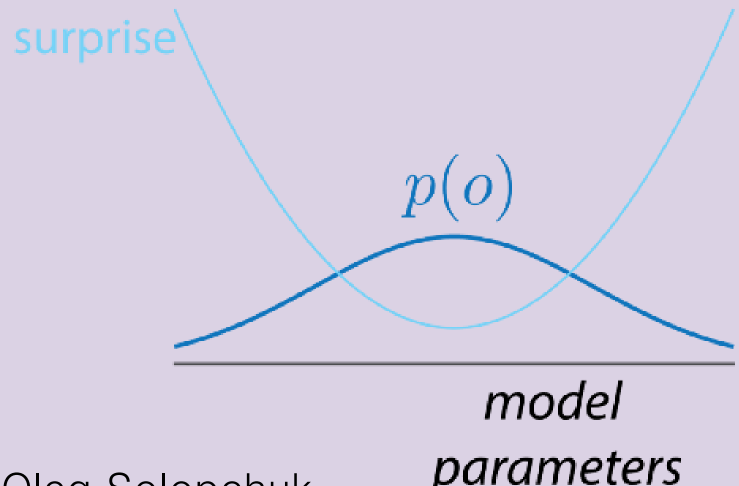

Courtesy of Oleg Solopchuk
(tutorial of Active Inference)

Active Inference:
The Origins

Model evidence =
marginal likelihood
"Sum out" s from P(o, s)

Variational
(approximate)
Bayes

$$-\log \sum_s p(o,s) = -\log \sum_s q(s) \frac{p(o,s)}{q(s)}$$

multiply by 1

Courtesy of Oleg Solopchuk
(tutorial of Active Inference)

Active Inference:
The Origins

$$-\log \sum_s p(o,s) = -\log \sum_s q(s)\frac{p(o,s)}{q(s)}$$

multiply by 1

Courtesy of Oleg Solopchuk
(tutorial of Active Inference)

Active Inference:
The Origins

$$-\log \sum_s p(o,s) = -\log \sum_s q(s)\frac{p(o,s)}{q(s)}$$

multiply by 1

Function of expectation $\leq$ ***Expectation of a function***

$$-log \sum_s q(s)\frac{p(o,s)}{q(s)} \leq -\sum_s q(s)\log\frac{p(o,s)}{q(s)}$$

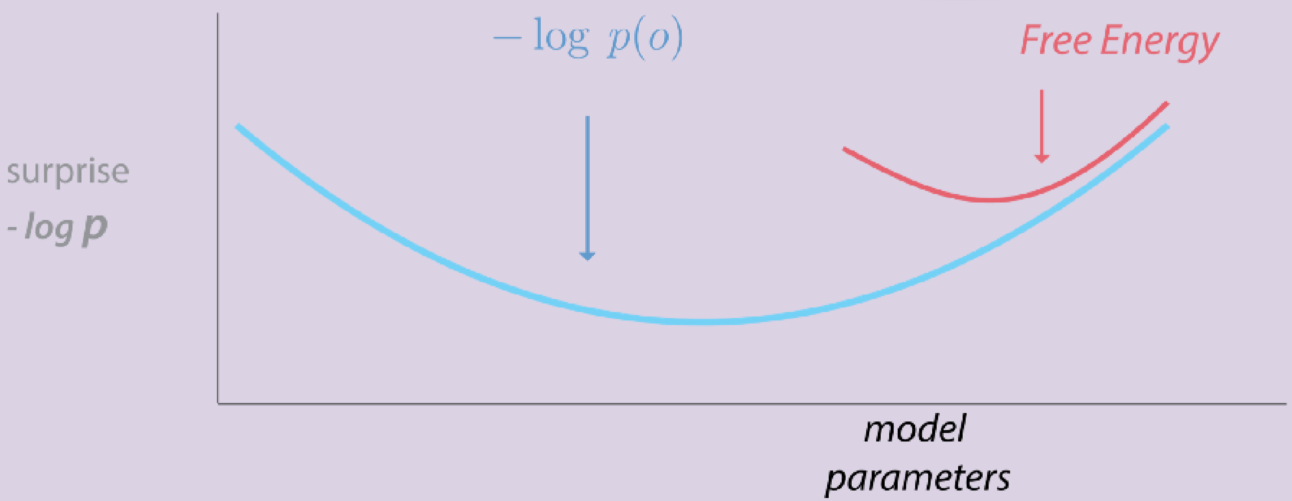

Courtesy of Oleg Solopchuk
(tutorial of Active Inference)

Active Inference: Variational Free Energy

Variational Free Energy F

$$= \sum_s q(s) \log \frac{q(s)}{p(o,s)}$$

$$p(a,b) = p(a|b)\,p(b) = p(b|a)\,p(a)$$

$$= \sum_s q(s) \log \frac{q(s)}{p(s|o)\,p(o)}$$

$$= \sum_s q(s) \log \frac{q(s)}{p(s)\,p(o|s)}$$

$$KL(p(a)||q(a)) = \sum_a p(a) \log \frac{p(a)}{q(a)}$$

$$\log(a/b) = \log(a) - \log(b)$$

$$= KL(q(s)||p(s|o)) - log\, p(o)$$

$$= KL(q(s)||p(s)) - \sum_s q(s)\, log\, p(o|s)$$

Courtesy of Oleg Solopchuk
(tutorial of Active Inference)

Active Inference:
Variational Free Energy

Variational Free Energy F

$$= \sum_s q(s) \log \frac{q(s)}{p(o,s)}$$

$$p(a,b) = p(a|b)\,p(b) = p(b|a)\,p(a)$$

$$= \sum_s q(s) \log \frac{q(s)}{p(s|o)\,p(o)}$$

$$= \sum_s q(s) \log \frac{q(s)}{p(s)\,p(o|s)}$$

$$KL(p(a)||q(a)) = \sum_a p(a) \log \frac{p(a)}{q(a)}$$

$$\log(a/b) = \log(a) - \log(b)$$

$$= KL(q(s)||p(s|o)) - log\, p(o)$$

$$= KL(q(s)||p(s)) - \sum_s q(s)\, log\, p(o|s)$$

Complexity

Accuracy

Courtesy of Oleg Solopchuk
(tutorial of Active Inference)

Active Inference: Policies

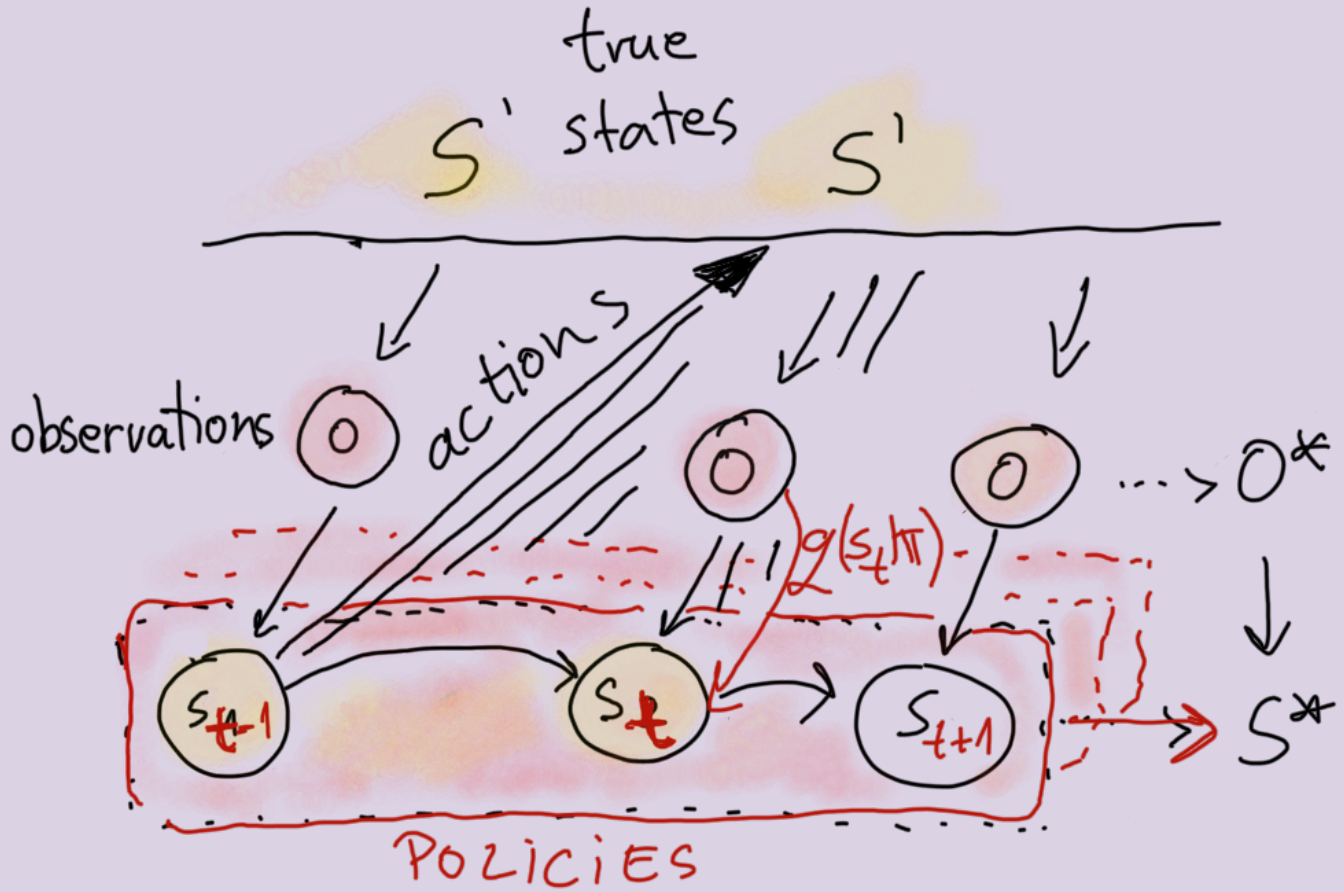

Active Inference:
Free Energy

$$= \sum_{o,s} p(o_t|s_t)\, q(s_t|\pi) \log \frac{q(s_t|\pi)}{p(o_t, s_t|\pi)} \qquad p(a,b) = p(b|a)\, p(a)$$

$$= \sum_{o,s} p(o_t|s_t)\, q(s_t|\pi) \log \frac{q(s_t|\pi)}{p(s_t|o_t,\pi)\, p(\boldsymbol{o}_t)}$$

prior preferences on future outcomes

- Epistemic value

$$= \sum_{o,s} p(o_t|s_t)\, q(s_t|\pi) \log \frac{q(s_t|\pi)}{p(s_t|o_t,\pi)} - \sum_{o,s} p(o_t|s_t)\, q(s_t|\pi) \log p(\boldsymbol{o}_t)$$

Bayes rule

$$\boxed{\frac{q(s_t|\pi)}{p(s_t|o_t,\pi)} = \frac{q(s_t|\pi)\, q(o_t|\pi)}{p(o_t|s_t,\pi)\, q(s_t|\pi)} = \frac{q(o_t|\pi)}{p(o_t|s_t,\pi)}} \qquad p(a|b) = \frac{p(b|a)\, p(a)}{p(b)}$$

$$= \sum_{o,s} p(o_t|s_t)\, q(s_t|\pi) \log \frac{q(o_t|\pi)}{p(o_t|s_t,\pi)} - \sum_{o,s} p(o_t|s_t)\, q(s_t|\pi) \log p(\boldsymbol{o}_t)$$

- Epistemic value

Epistemic value tells us how much we could learn from the environment if we followed this policy.

Active Inference: Free Energy

$$= \sum_{o,s} p(o_t|s_t)\, q(s_t|\pi) \log \frac{q(o_t|\pi)}{p(\boldsymbol{o}_t)\, p(o_t|s_t,\pi)}$$ ← *predicted outcomes*

$$= \sum_{o,s} q(s_t|\pi)\; p(o_t|s_t)\; \log \frac{q(o_t|\pi)}{p(\boldsymbol{o}_t)} - \sum_s q(s_t|\pi) \sum_o p(o_t|s_t) \log p(o_t|s_t)$$

$$KL(p(a)||q(a)) = \sum_a p(a) \log \frac{p(a)}{q(a)} \qquad H[p(a)] = -\sum_a p(a) \log p(a)$$

$$= KL(q(o_t|\pi)||p(\boldsymbol{o}_t)) + \sum_s q(s_t|\pi) H[p(o_t|s_t)]$$

Expected cost *Expected Ambiguity*

Active Inference:

Free Energy F

$$= KL(q(o_t|\pi)||p(o_t)) + \sum_s q(s_t|\pi)H[p(o_t|s_t)]$$

Action choice:
Min(F) x precision

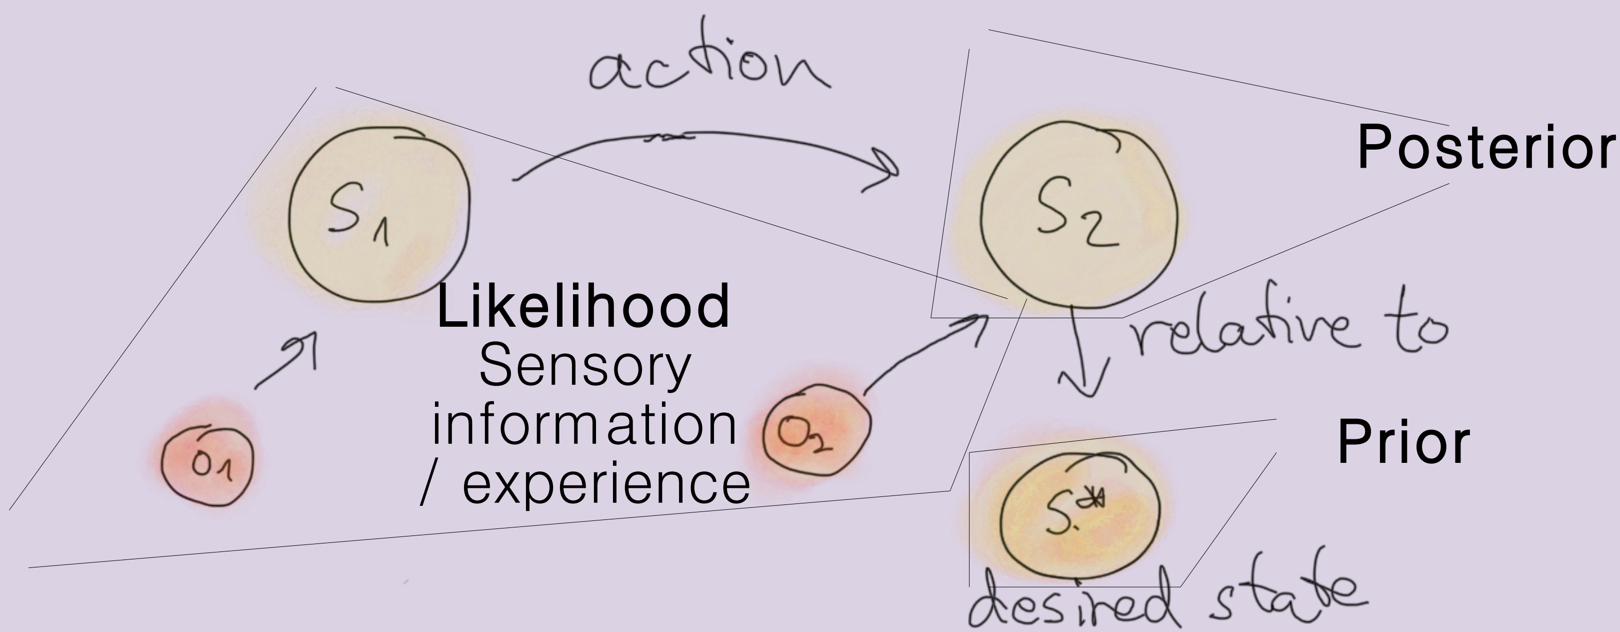

Unifying, Generic Framework (Computational, P300)

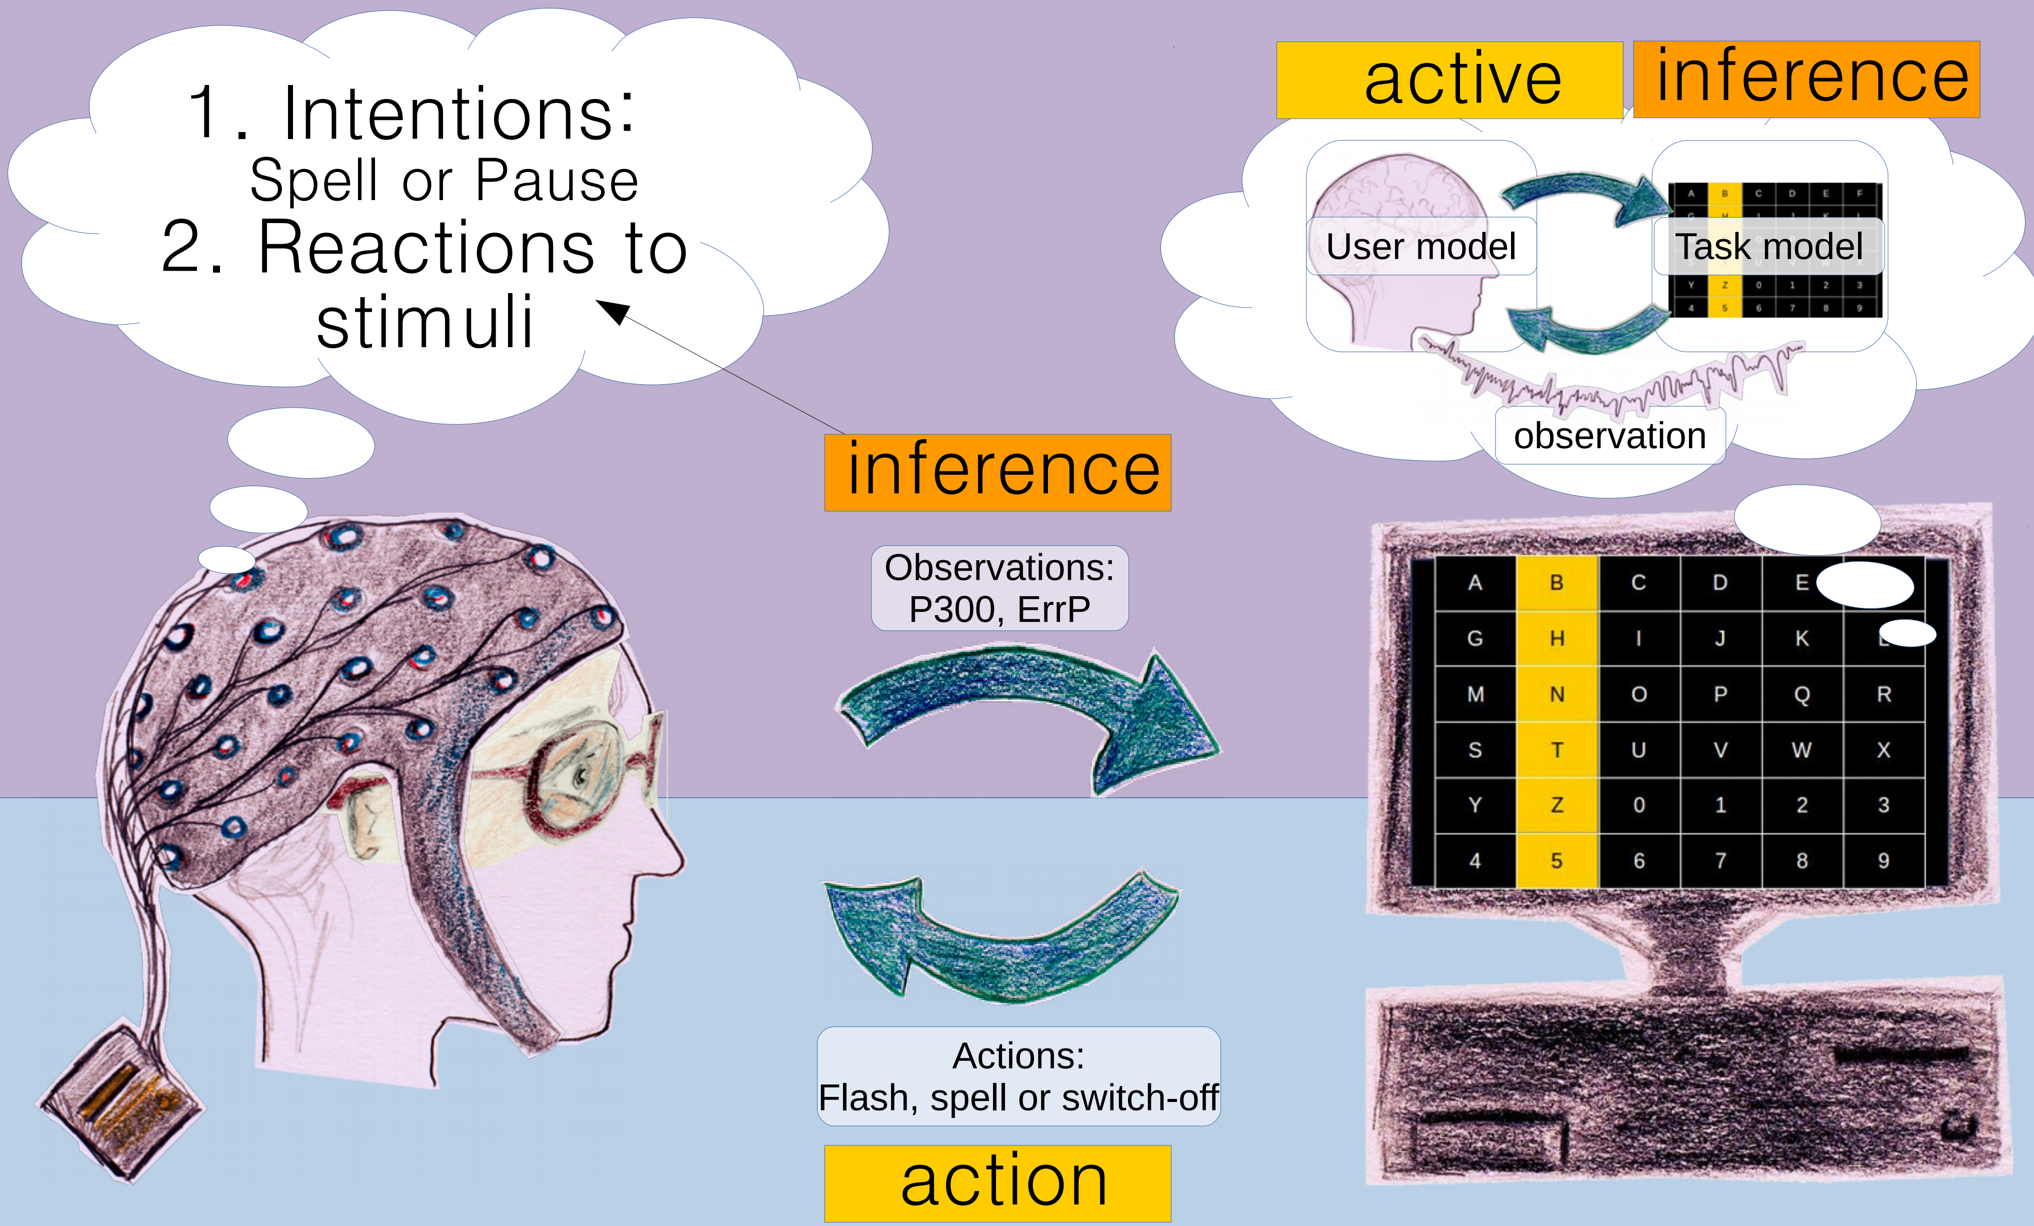

Active Inference:
Parallel brain and machine

The brain as an adaptive system,
described by Active Inference:

accumulates sensory input
to update its internal model of the environment,

The machine incorporates the same behavior :

Active Inference:
Parallel brain and machine

The brain as an adaptive system, described by Active Inference:

The machine incorporates the same behavior :

accumulates sensory input to update its internal model of the environment, → EEG data, to update its internal model of the **user**;

Active Inference:
Parallel brain and machine

The brain as an adaptive system, described by Active Inference:

The machine incorporates the same behavior :

accumulates sensory input to update its internal model of the environment, → EEG data, to update its internal model of the **user**;

optimizes its interactions through making inference about the environment, → with the updated user model, make inference about the user intentions,

Active Inference:
Parallel brain and machine

The brain as an adaptive system, described by Active Inference:

The machine incorporates the same behavior :

accumulates sensory input to update its internal model of the environment, → EEG data, to update its internal model of the **user**;

optimizes its interactions through making **inference** about the environment, → with the updated user model, make predictions about the user state,

and

through **acting** upon the environment; → reinforce predictions or reduce prediction error with optimal action (feedback or stimuli).

Dialogue between the two,
co-adaptation

Active Inference:
Application on P300 Speller

Generic (Bayes) Model
States, Observations, Action

User intentions
(spell, pause)
and reactions to
stimuli

36 items and
37th “look–away” for pause
x 4 reaction types:
“My letter is flashed/spelled”
or not.

Active Inference:
Application on P300 Speller

Generic (Bayes) Model
States, Observations, Action

User intentions
(spell, pause)
and reactions to
stimuli

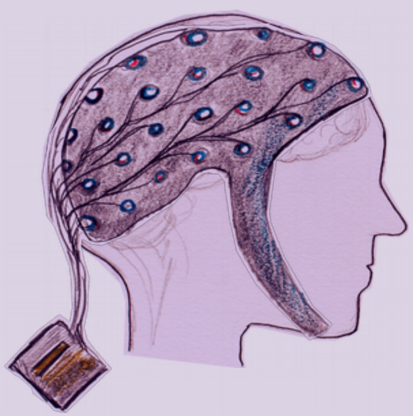

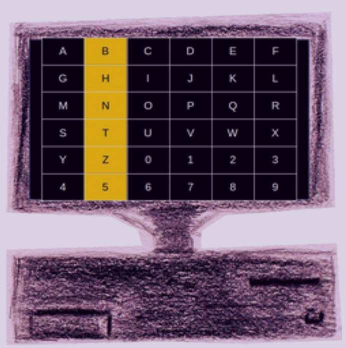

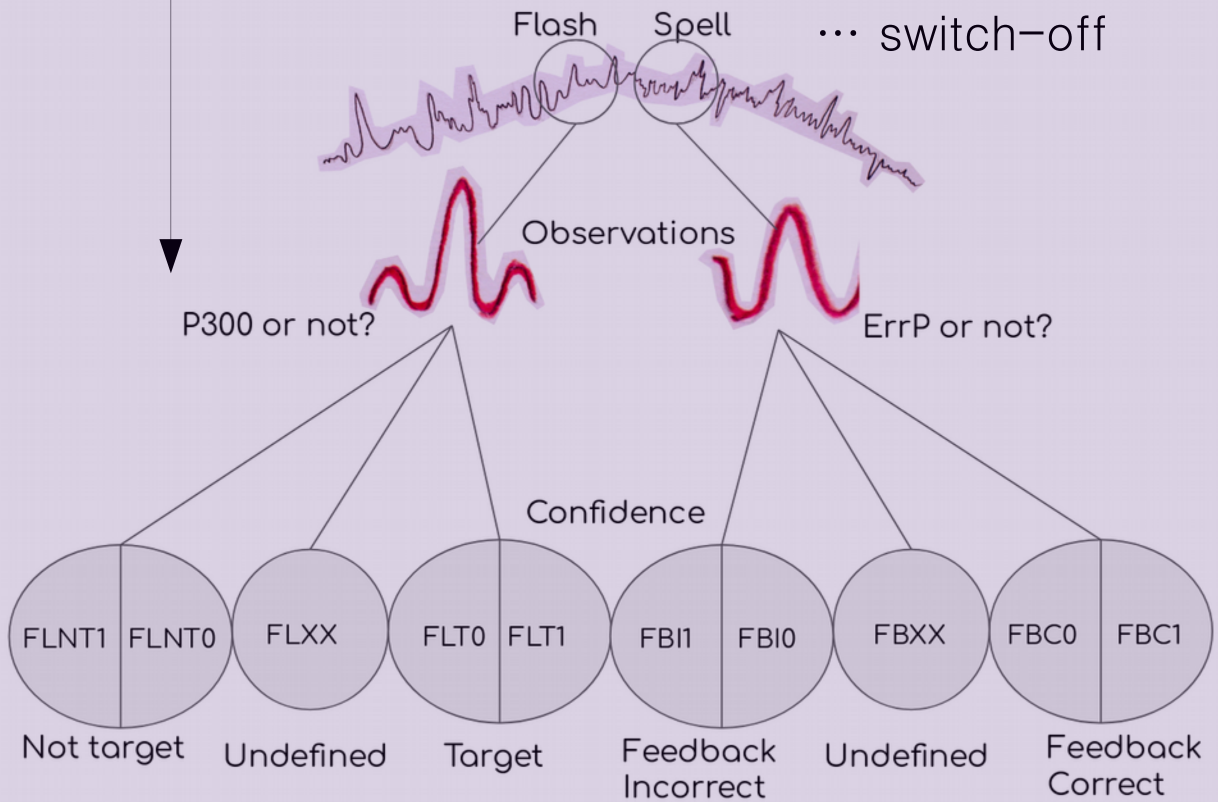

Active Inference:
Application on P300 Speller

Generic (Bayes) Model
States, Observations, Action

User intentions (spell, pause) and reactions to stimuli

Note:
For each subject, the probability to get an observation (likelihood) is learned from training

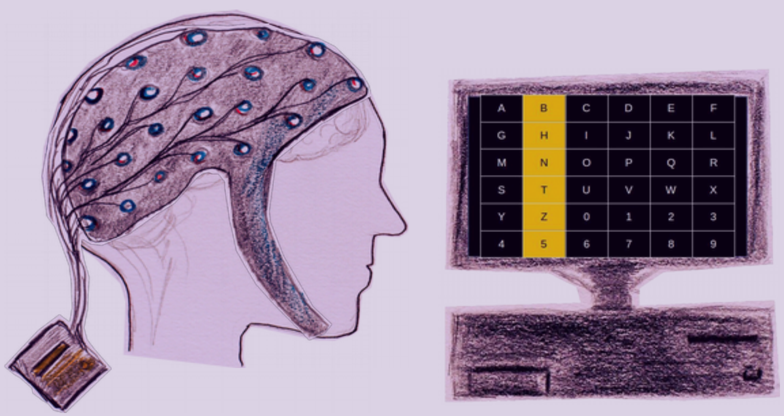

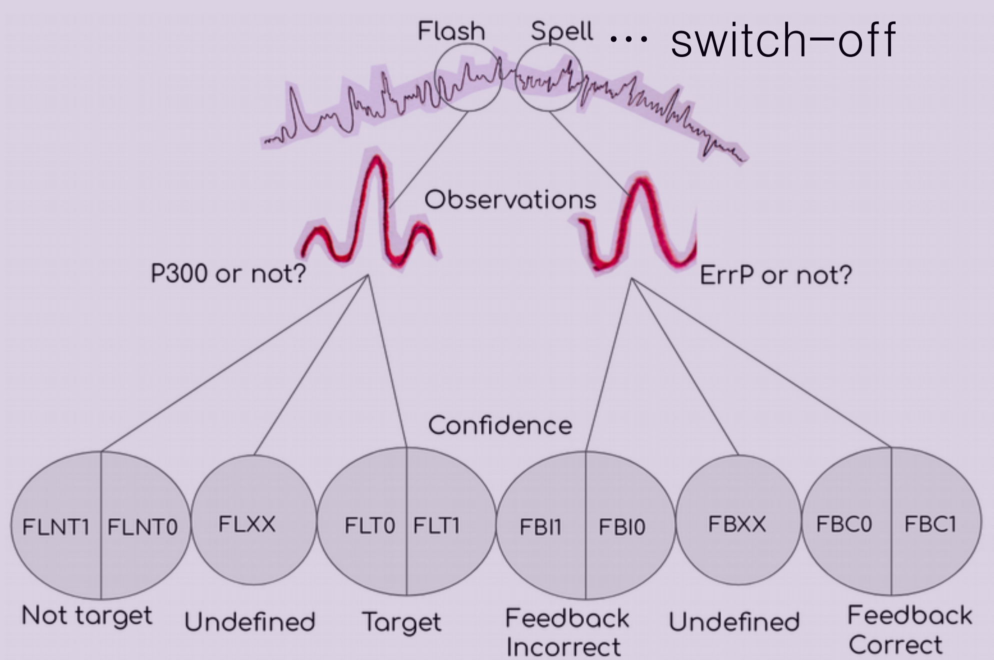

Active Inference:
Application on P300 Speller

Generic (Bayes) Model
States, Observations, Action

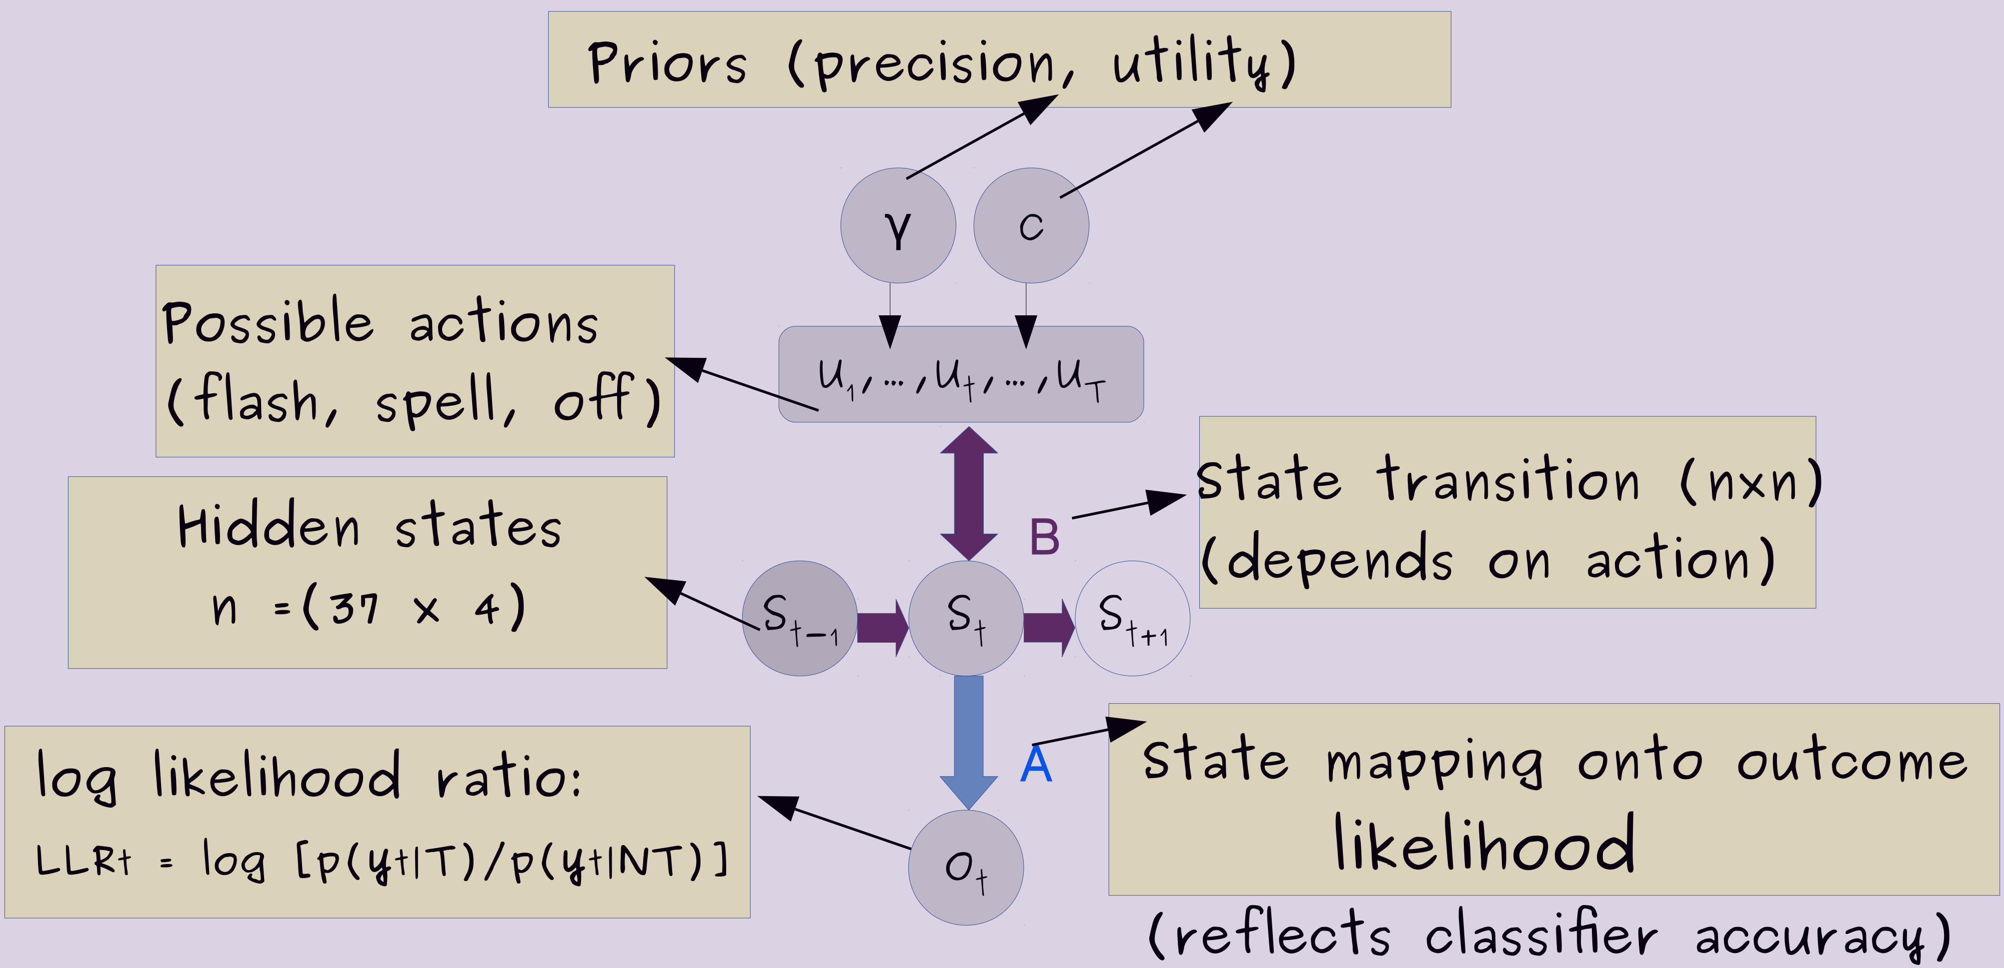

Active Inference:
Application on P300 Speller

Generic (Bayes) Model
States, Observations, Action

Passing from one state to another (through action) using relative entropy and by choosing such action that provides most information (min surprise)
= **Reveal the feature of interest: the target letter!!!**

Active Inference:
Application on P300 Speller

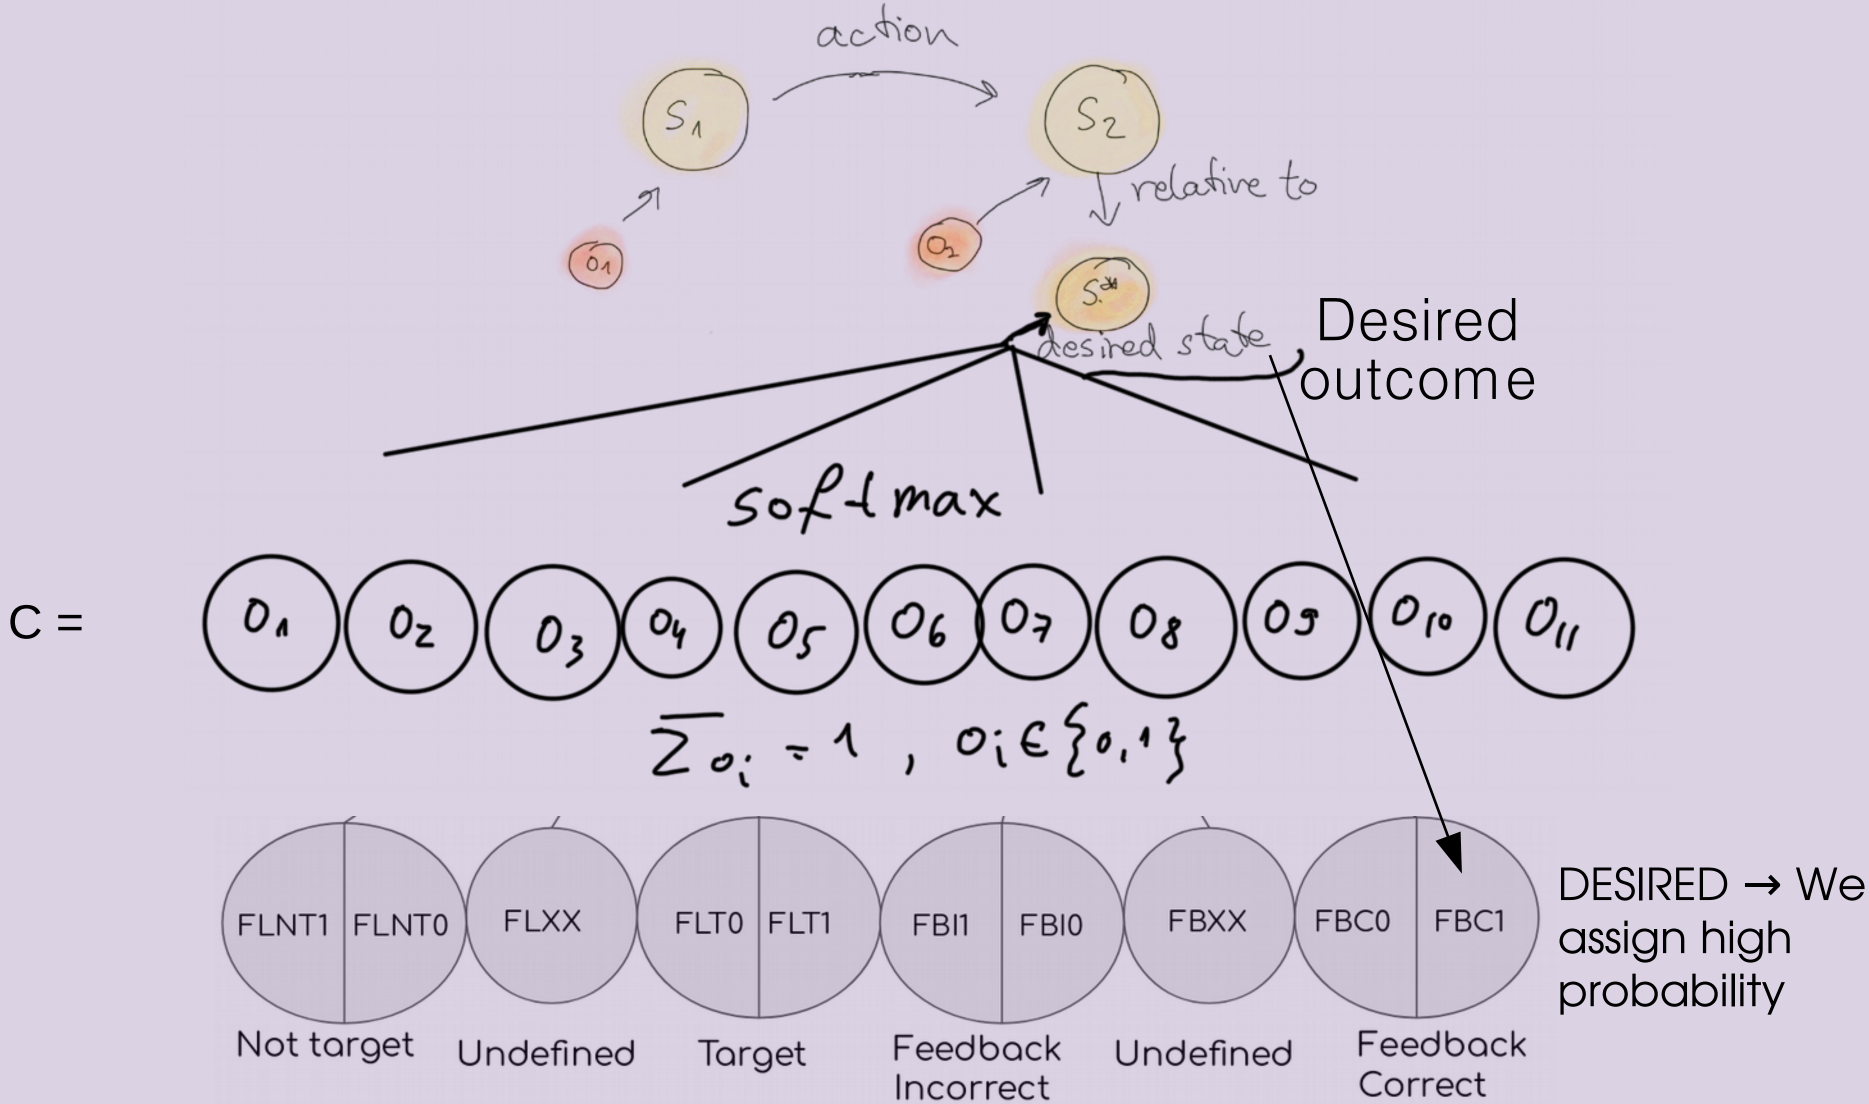

Active Inference:
Application on P300 Speller

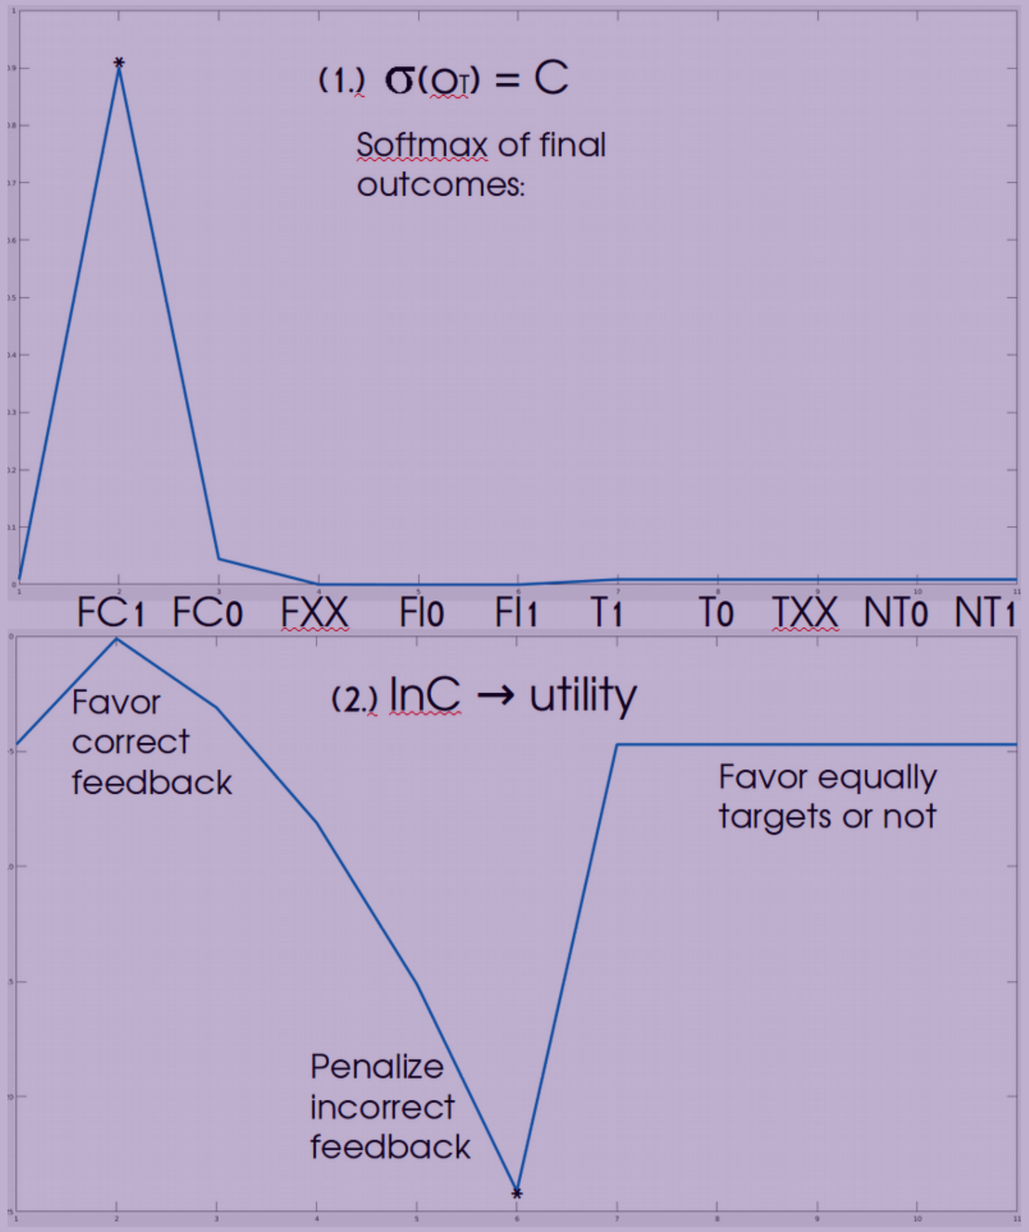

Active Inference:
Application on P300 Speller

All is in one equation, making any available action possible (switch between any action automatically and by rate of precision and speed defined).

Active Inference: Experimental Design

Observations – already classified data (target/not) with Riemann distances transformed into log likelihoods, from 18 real subjects.

Active Inference: Simulation

Observations – already classified data (target/not) with Riemann distances transformed into log likelihoods, from 18 real subjects.

After each flash, we pick randomly from the pool of data (target/not), map it to our observations and get 1 out of 5 possible observations.

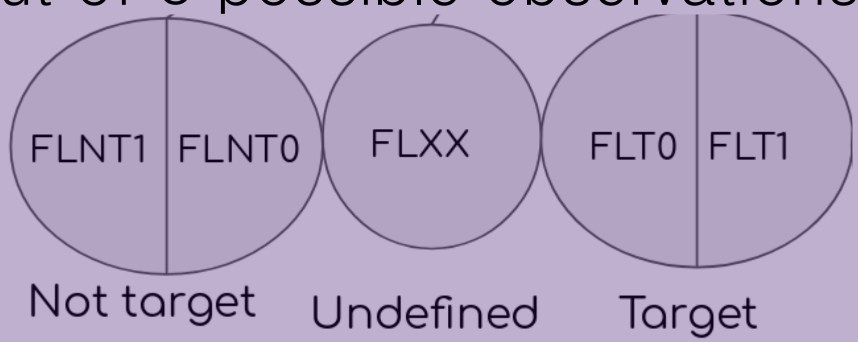

So, even though it was a target in the beginning, it could be a non-target after the mapping – that depends on the likelihood

Active Inference:

Demo:

FeedBack classifier (ErrP) is simulated (specificity: 0.95, sensitivity: 0.75, p(FBXX)= 0)

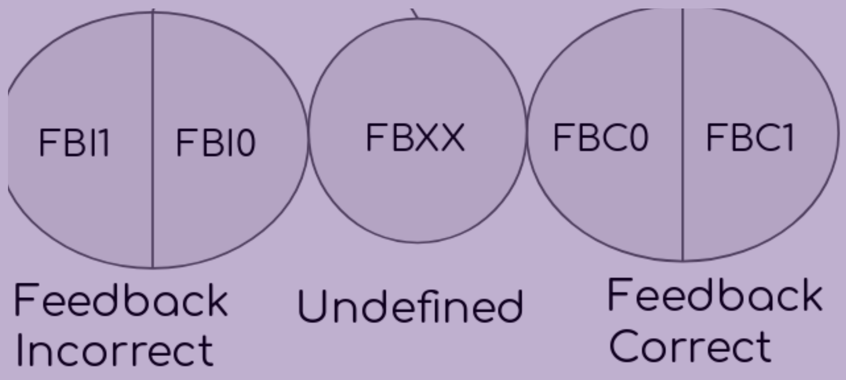

Active Inference: Results

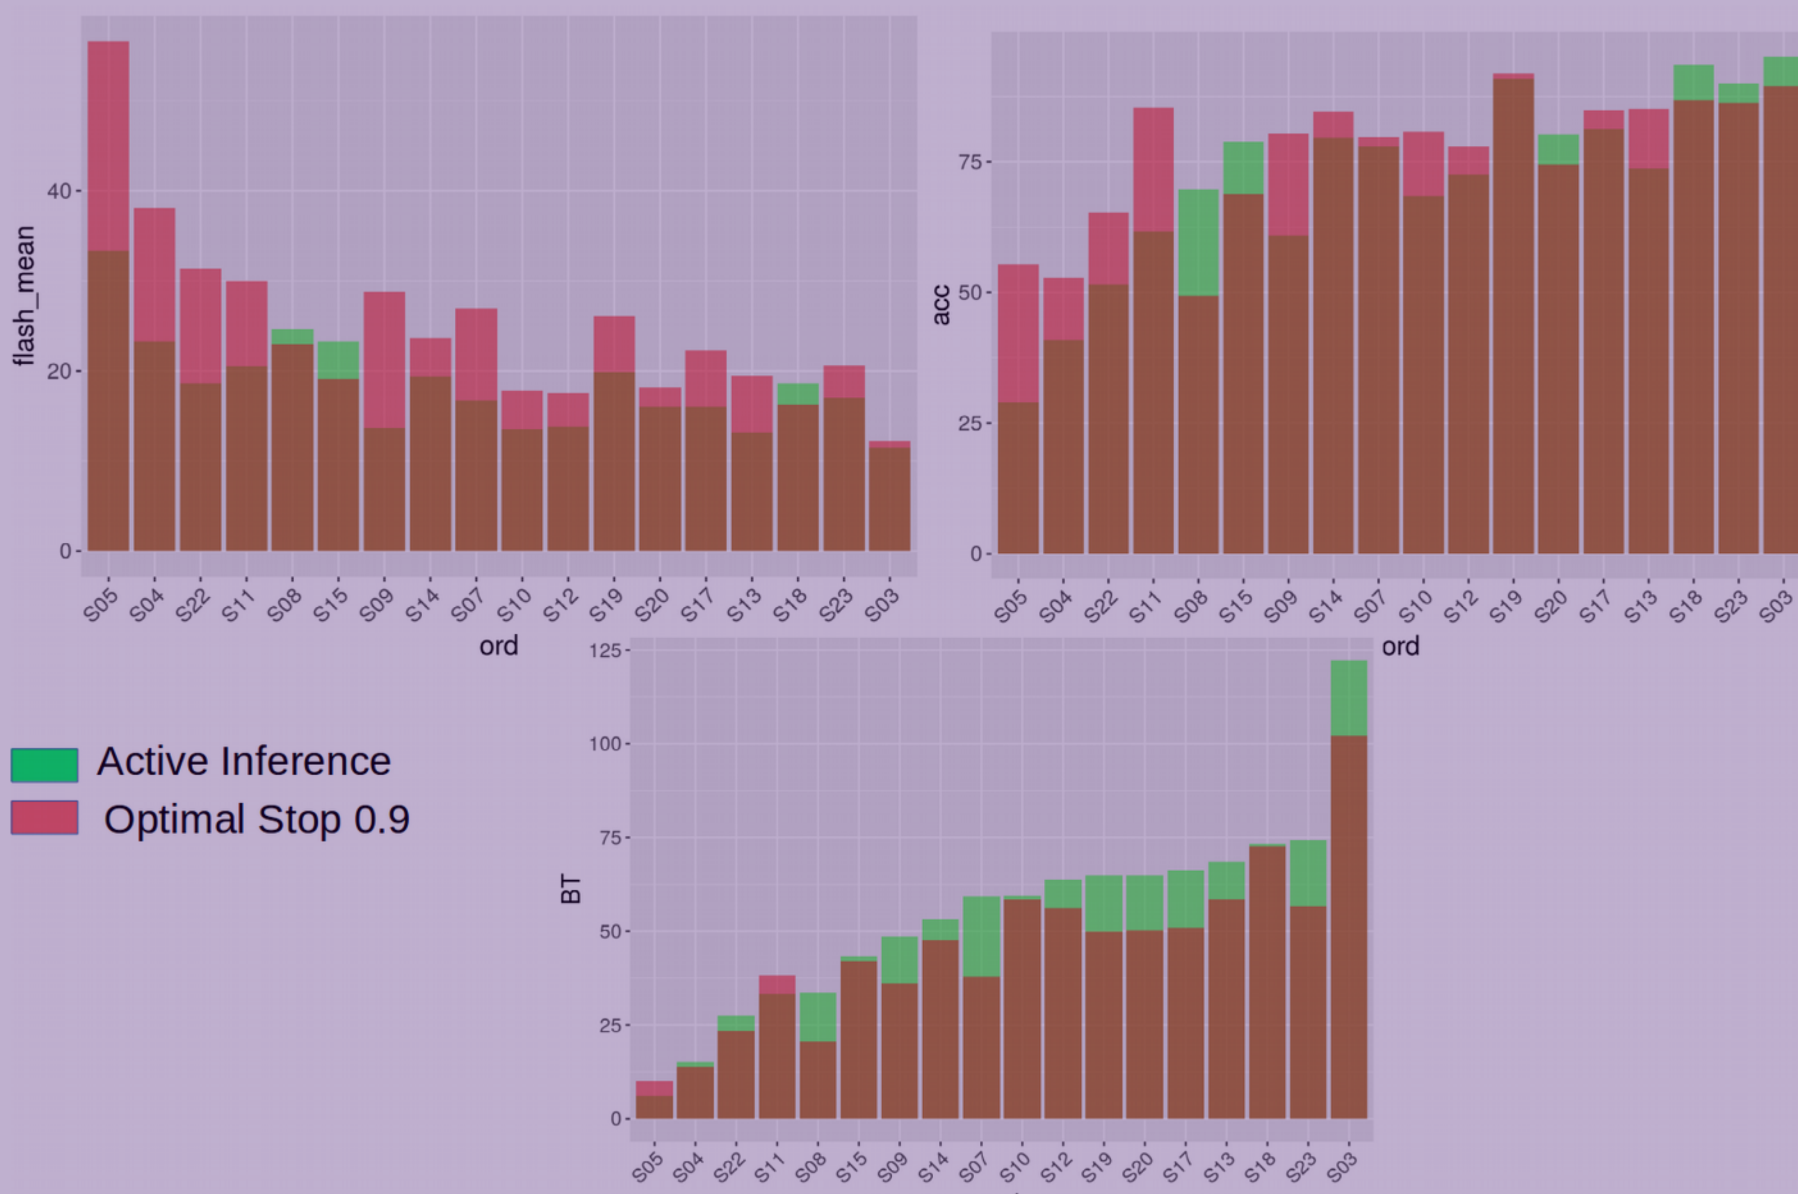

Active Inference: Results

12000 letters, 18 subjects

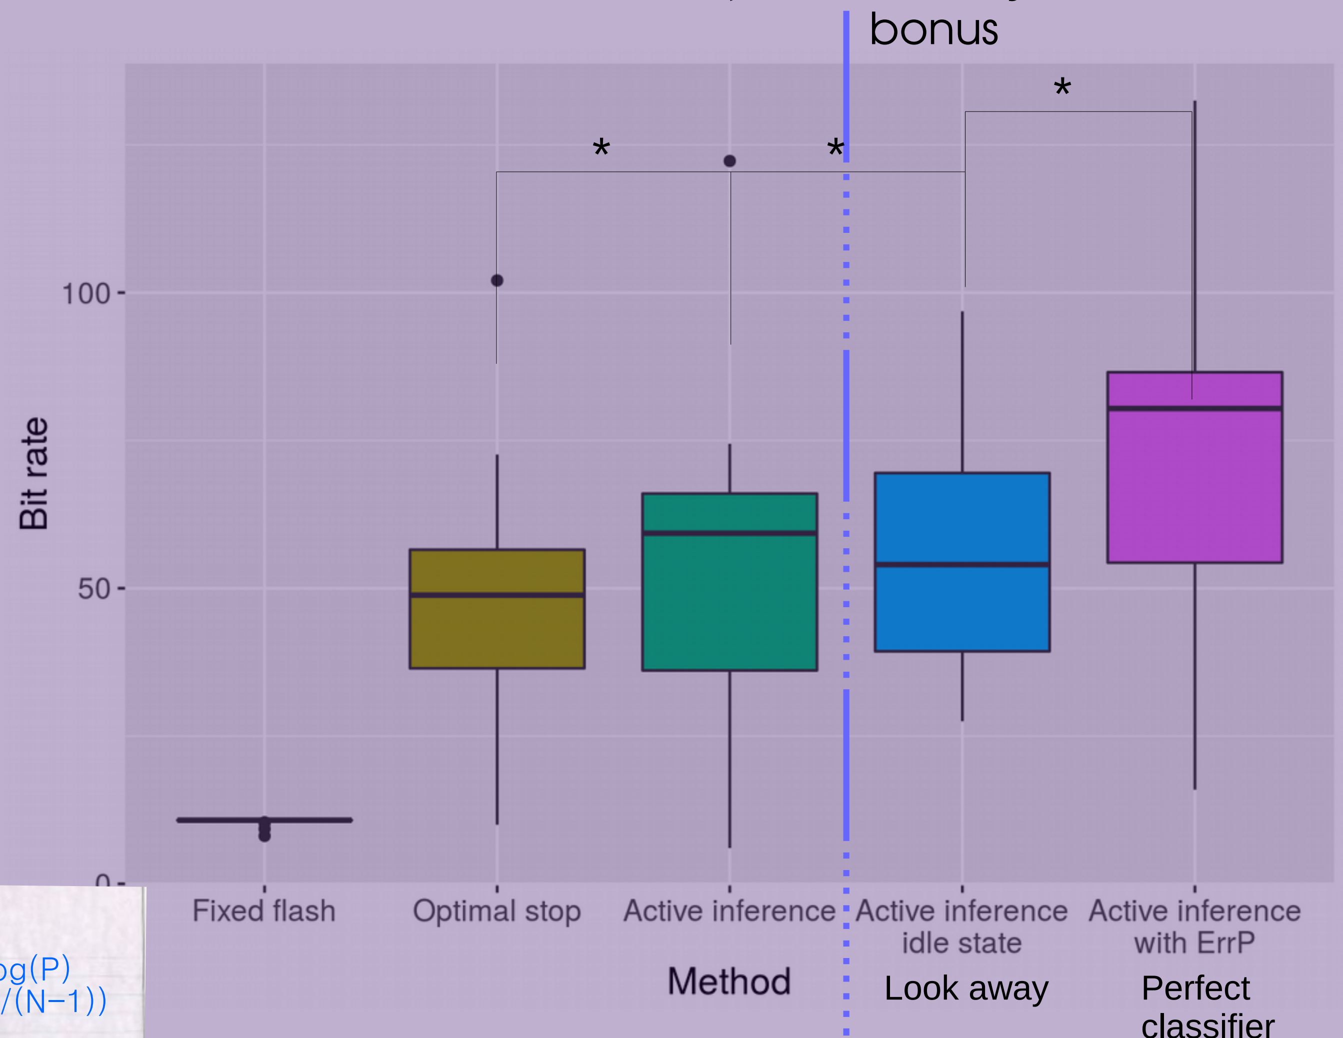

Bit Rate:

```
B <- log(N) + P * log(P)
- (1-P) * log((1-P)/(N-1))

T <- 0.2 * nFlash
BR <- B * (60/T)
```

Active Inference:
Take Away

Active Inference is available in spm toolbox (matlab)
for free, with many examples;

Provide to Active Inference a well defined model:
States, Observations, Action and Desired Outcome!
Likelihood needs to be learned.

All is integrated in
one unified Bayesian framework
minimizing Free Energy;

It is good for BCI: co-adaptation,
action influences user automatically.

CHECK OUT tutorial from Oleg:
https://medium.com/@solopchuk/tutorial-on-active-inference-30edcf50f5dc

Active Inference:
Back to the Future

Tux Flow
with Active Inference!!

Active Inference:
Changing C values

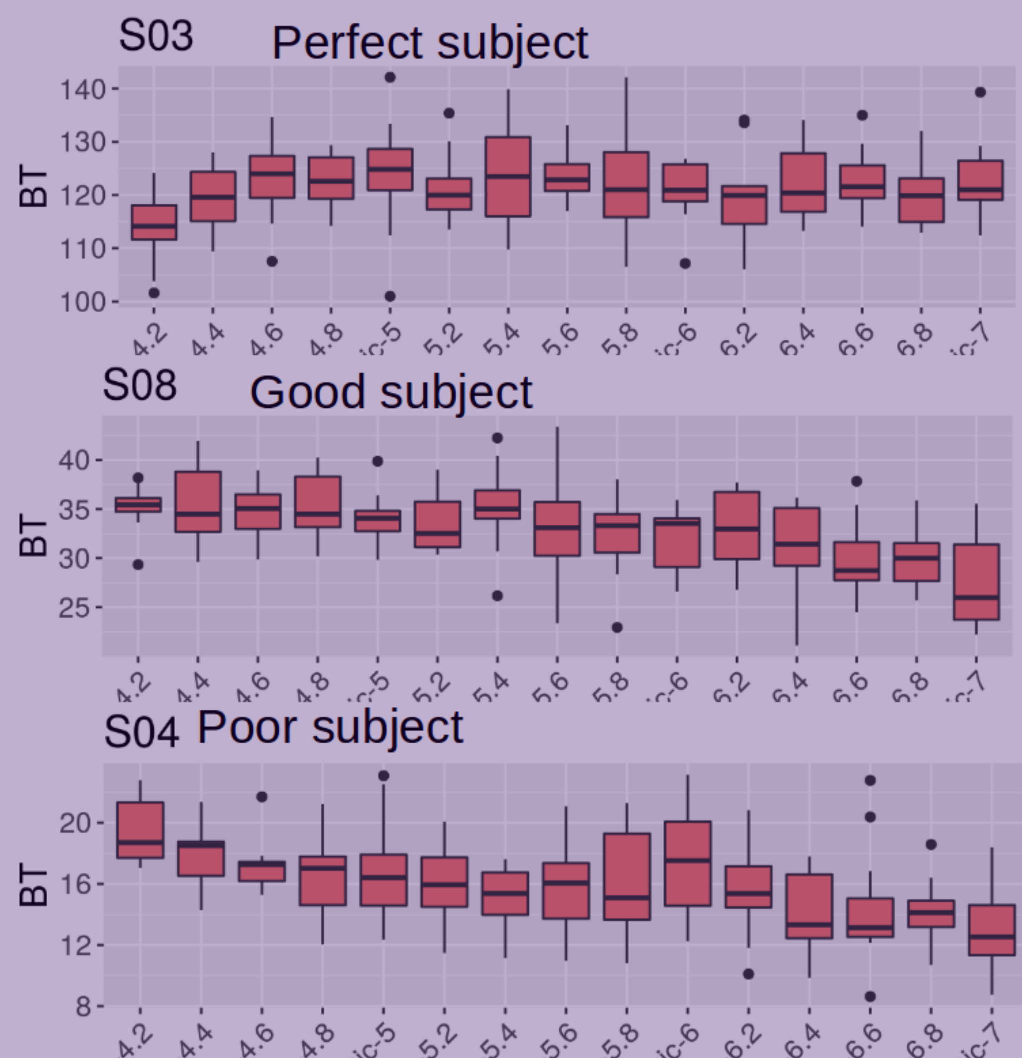

Active Inference:
Experimental Design

After each flash, we pick randomly from the pool of data (target/not), map it to our observations and get 1 out of 5 possible observations.

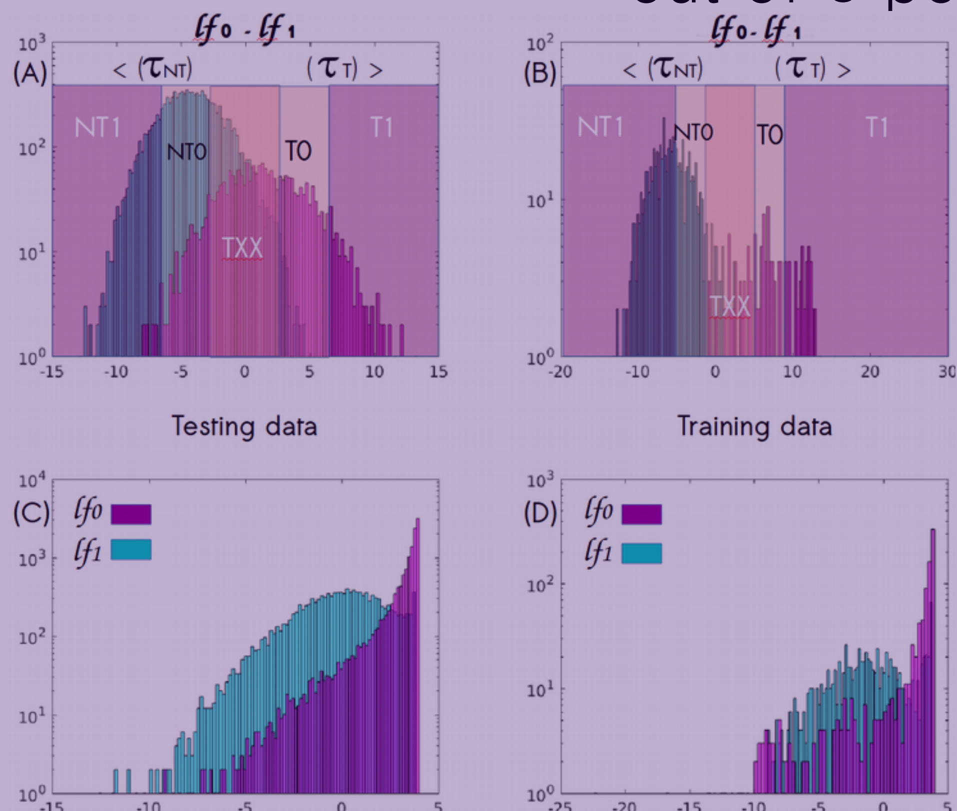

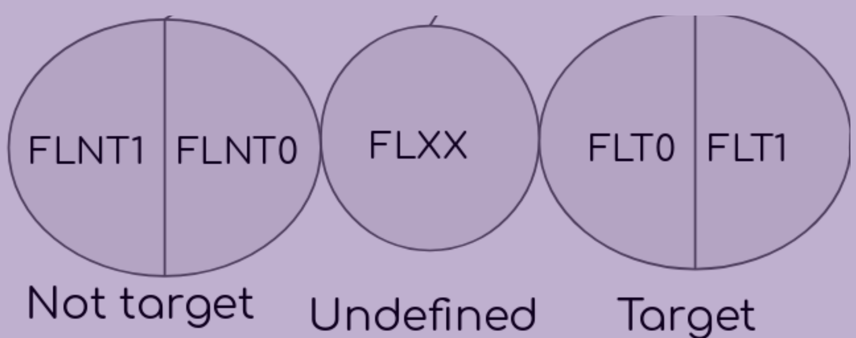

So, even though it was a target in the beginning, it could be a non-target after the mapping – that depends on the likelihood