Cognitive theories and decision-making: Introduction to predictive processing

ISyE 516 Introduction to decision making

Tony McDonald

December 13, 2023

Asst. Professor Department of Industrial and Systems Engineering

Dimensions of cognition

Centralized vs. Embodied (Situated)

Symbol storage and manipulation vs. empirical interactions

Cognitivist (Traditional)

- The brain and body are separate entities
- The brain is a symbol system
- Cognition is just symbol manipulation

Connectionist systems

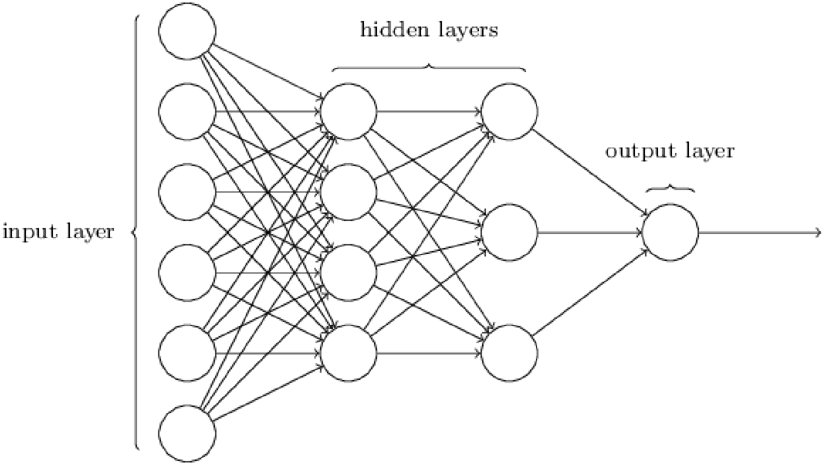

Embodied cognition

Tenants of embodied cognition (Wilson)

1. Cognition is situated in the real-world
2. Cognition is time pressured
3. We use the environment to aid our cognitive process
4. The environment is part of the cognitive system
5. The purpose of cognition is action
6. Off-line cognition uses sensory and motor control

Group exercise 10 min

In your groups, discuss your reflection response about the Hutchins case study.

Today's lecture: What do cognitive theories say about decisions

1. Cognitivist decision-making
2. Embodied decision-making
3. Predictive processing (new for today)

Cognitivist decision making

Recall that cognitivists view cognition as symbol manipulation. The same view applies to decisions.

Consider our car purchasing example from earlier. We can use expected utility and risk to formulate this as a math problem (i.e. symbol manipulation) and make a decision.

This works well for off-line cognition!

Cognitivist decision making

What about online decisions?

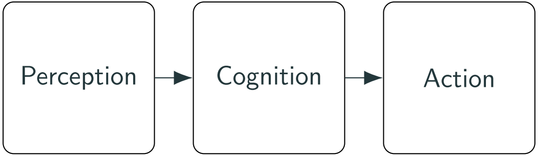

Perception builds a representation of the real world relevant features.
Cognition translates these representations into a plan according to beliefs and preferences, then activates the **Action**.

Decision making resides completely in the cognitive node.

Cognitivist decision making limitations

1. Forcing a separation between perception, cognition, and action and requiring the systems to communicate in a unified way (e.g., through utility representations) is awkward.

2. This division does not reconcile with neurological data, especially for online decisions.

3. There is no provision for epistemic behavior.

Embodied decision making

Many theories, but one plausible theory is Cisek's **affordance competition hypothesis**.

"Decisions emerge from a distributed, probabilistic competition between multiple representations of possible actions which overlap with sensorimotor circuits." (Burr, 2017 p. 4)

"The brain simultaneously is specifying and selecting among representations of multiple "action opportunities"—affordances—which compete with the sensorimotor system itself." (Burr, 2017, p. 5)

Embodied decision making

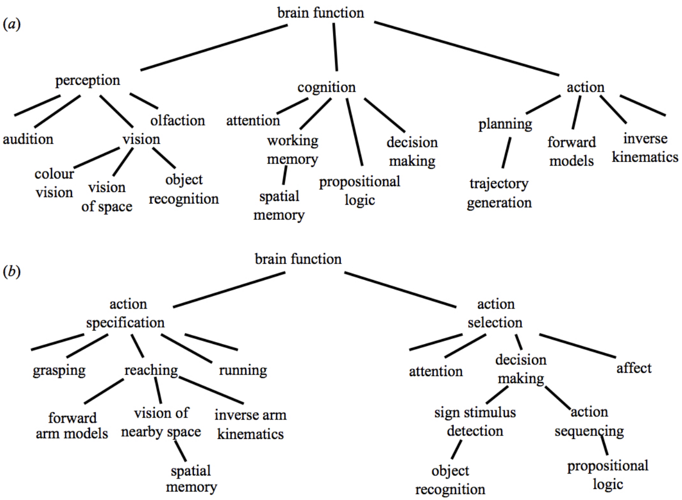

(*a*) Cognitivist organization; (*b*) ACH organization

Embodied decision making

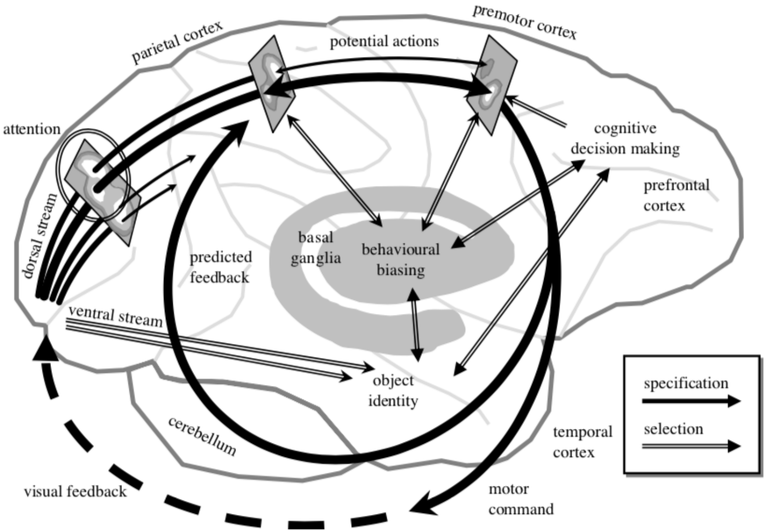

Embodied decision making

The ACH also supports an argument for dynamic decision making

Decisions are continuously analyzed following action initiation through sensory feedback.

This allows for action adjustments that consider evolving biomechanical costs.

Embodied decision making

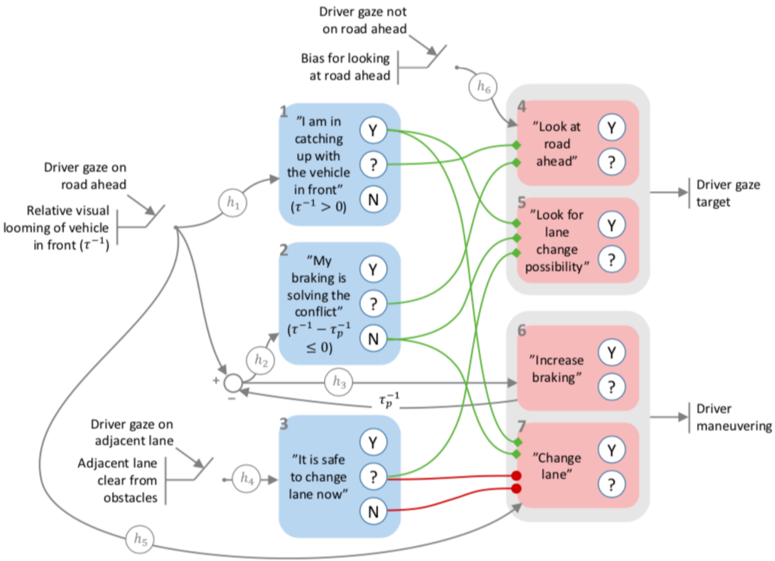

Embodied decision making limitations

1. The account of offline decision making is incomplete

2. There is no notion of the impact of affect on decision making

Predictive processing

Predictive processing is the proposal that the principle purpose of the brain is to reduce errors identified by the cognitive system. The cognitive system is a hierarchical generative model that integrates signals from the perceptual systems of the body with higher level goals and representations.

The *hierarchical generative model* is developed over time based on how states and events in the world and the body generate sensory input. The brain attempts to minimize errors between predicted sensory input and actual sensory input through perception and action.

Perception is updating the sensory prediction based on sensory input.
Action is behavior that changes the sensory input to match predictions.

Predictive processing

There are four key elements of predictive processing:

1. Hierarchical generative model

2. Active inference

3. Precision

4. Model tuning

Hierarchical generative model

We can think of the cognitive system as a hierarchy. The lowest levels represent basic sensory input signals. Higher levels represent increased abstraction and aggregation.

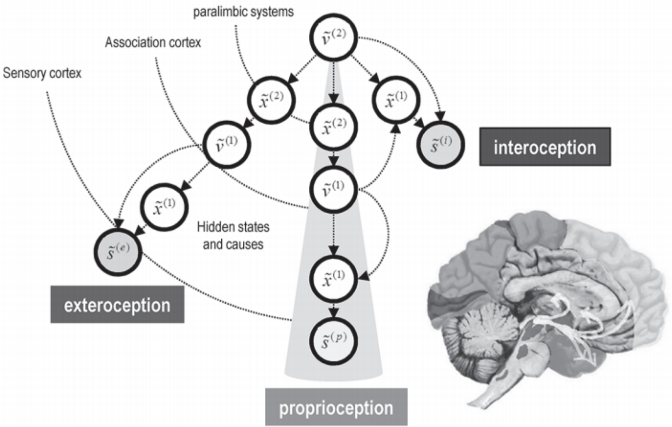

Active inference

Active Inference is the process in which humans infer environmental and bodily states from sensory input.

Sensory input can be derived from one of three sources:

1. *exteroceptive* - generated by the environment.
2. *interoceptive* - generated by the internal organs (i.e. emotions).
3. *proprioceptive* - generated by the states of the muscles and joints.

Active inference

Hierarchical Generative Model

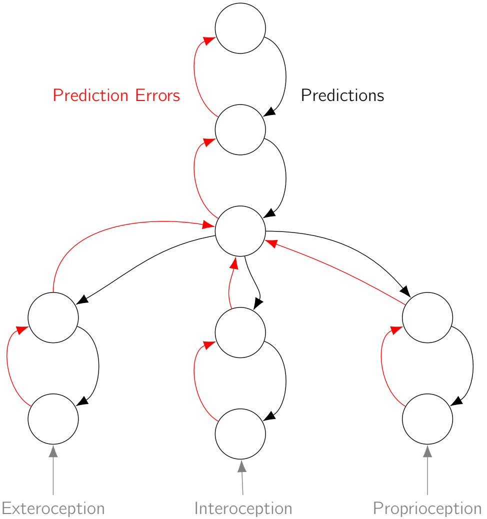

Active inference

Example: Consider that you are driving a vehicle and are approaching an intersection. You believe that you are in a lane that permits proceeding straight through the intersection or turning. You decide to go straight. However you are in a turn-only lane. As you enter the intersection other drivers begin to honk their horns.

Active inference

As the car enters the intersection:

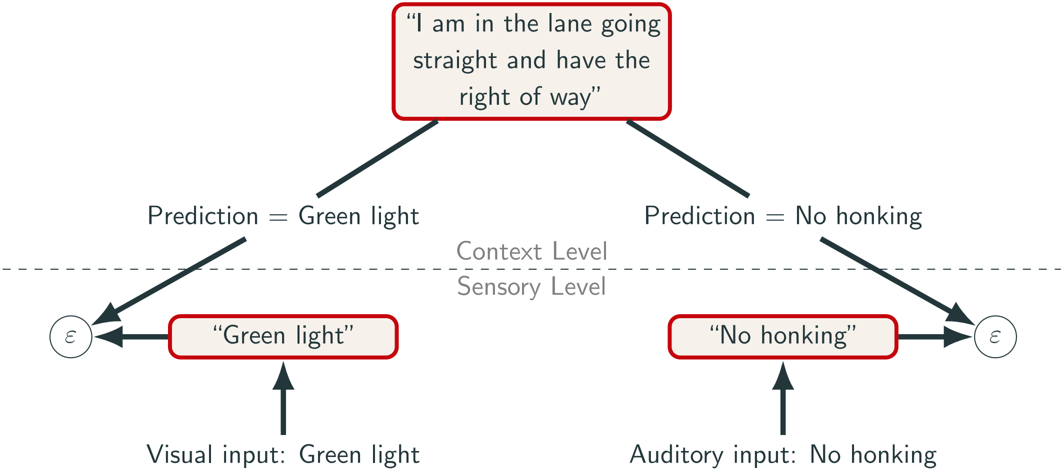

Active inference

As other vehicles begin to honk at the driver:

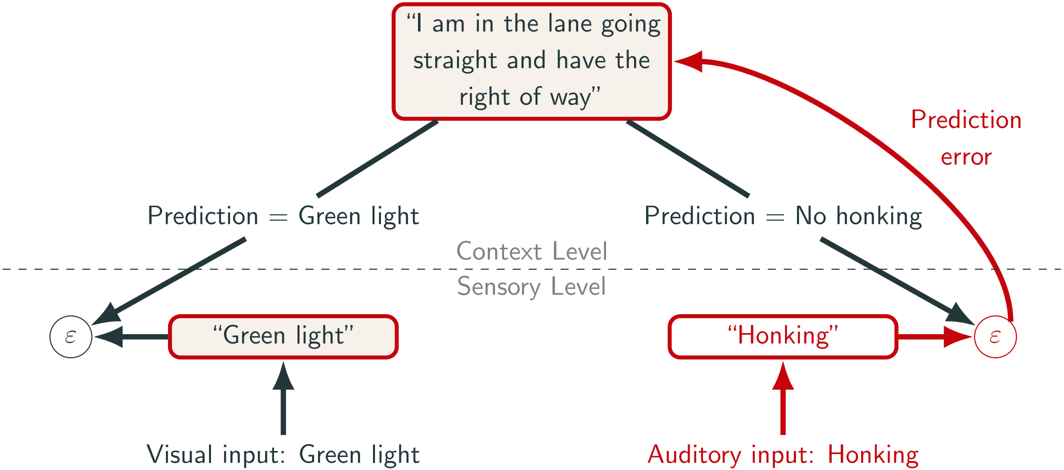

Active inference

As the driver begins to understand their error:

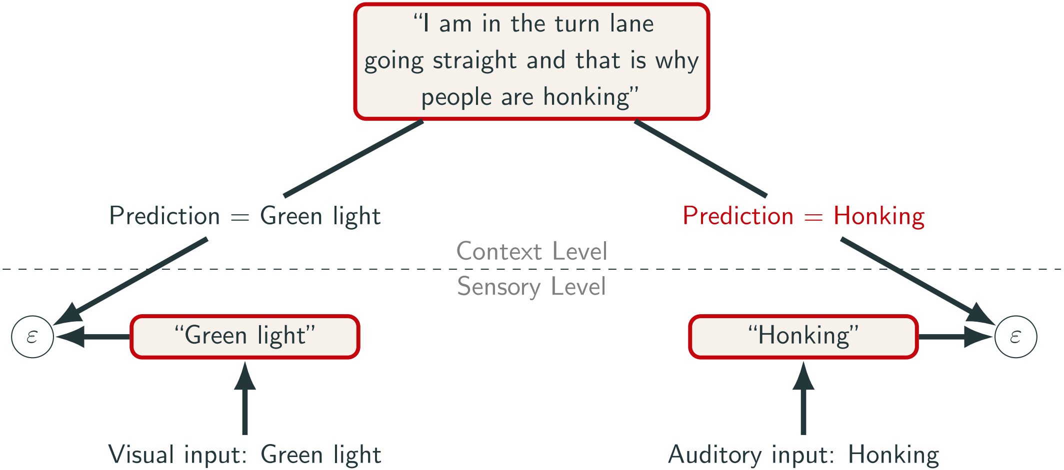

Active inference

A few key additional points:

1. Only unexplained sensory input is passed up the hierarchy
2. This mechanism allows us to explain decision updates (similar to the ACH)
3. When all prediction errors are resolved, the system reaches a stable state.

Precision

Predictions from the hierarchical generative model are probabilistic. These predictions include an expected value and an expected **precision** *(i.e. inverse variance). Prediction errors are* **precision weighted** *as they are passed up the hierarchy.*

Precision plays a key role in decisions between pragmatic and epistemic actions. Epistemic actions are used to increase the precision of the current prediction. Pragmatic actions are used to correct explicit errors in expectation.

Precision

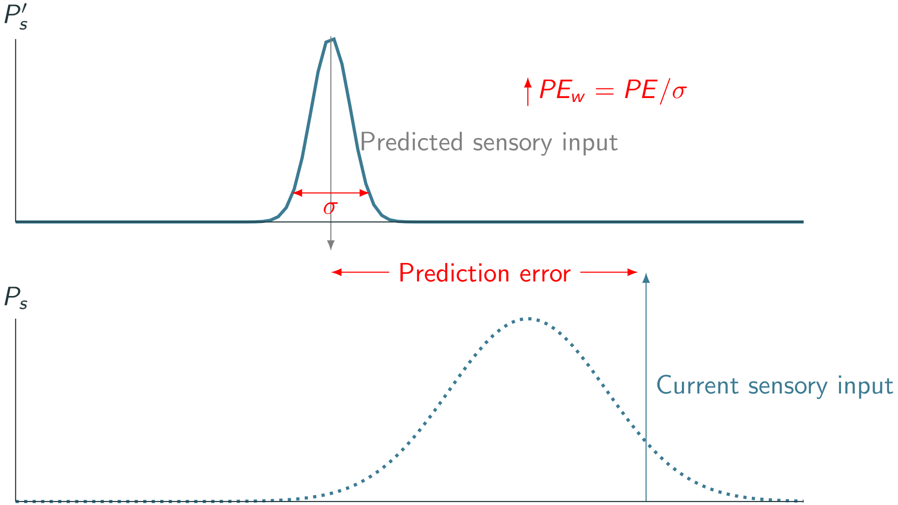
P'_s
$PE_w = PE/\sigma$
Predicted sensory input
σ
Prediction error
P_s
Current sensory input

Precision: Inaccurate model and probable event

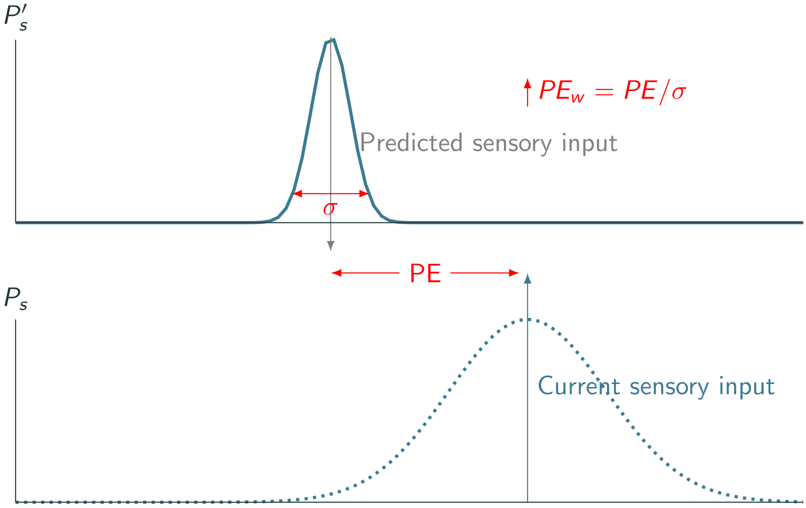

Example: Volleyball vs. the lab compared to a professional.

Precision: Accurate model but improbable event

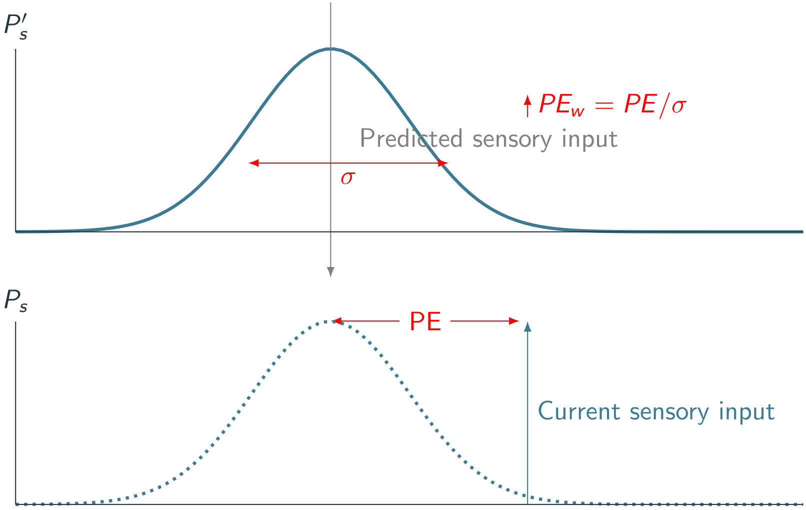

Example: "Bad luck" events, such as a hard hit baseball that is blown back into the stadium with a gust of wind.

Precision: False certainty

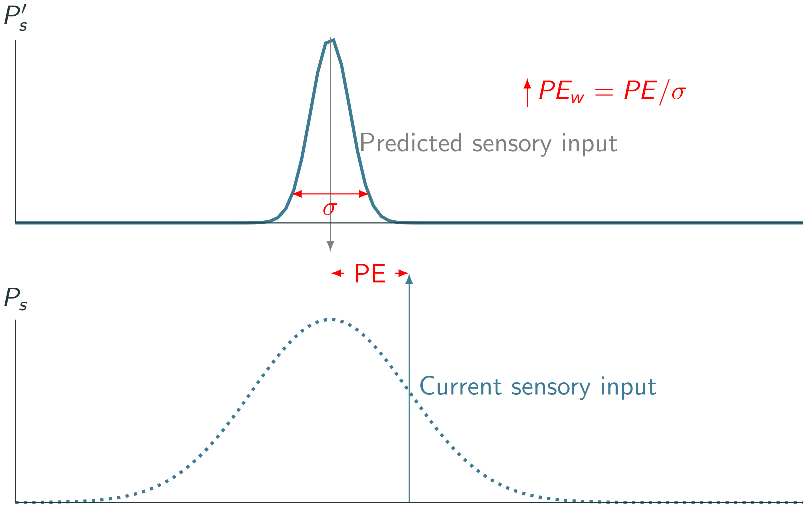

Example: Teenage driver looking away from the road to write a text message with the belief that the vehicle in front of them will not brake.

Precision: False uncertainty

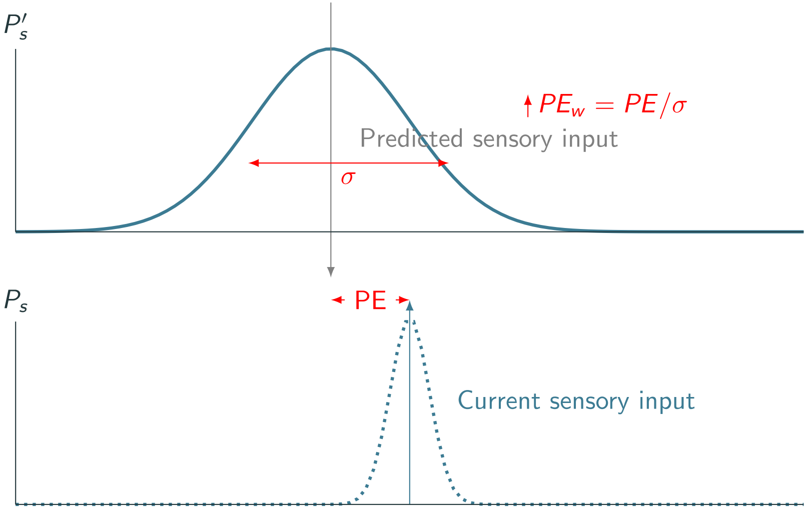

Example: Driving in Wisconsin after driving in Texas for 5 years.

Group Exercise

In your groups (15 min):

1. Develop your own example for 2 of the four cases of precision/prediction errors.
2. For each example explain how active inference would function in the scenario.
3. Document your response in a paragraph.

Model tuning

Model tuning is the predictive processing framework's method of explaining learning.

As statistical regularities emerge in the environment the hierarchical generative model is updated to align. This process is different from Active inference which deals with moment-to-moment model updates.

Model tuning

Predictive processing model tuning follows a reinforcement learning pattern. Reinforcement learning in the predictive processing involves a gradual modification of the generative model that improves predictions and precision, thus tuning the system to relevant statistical regularities in the world.

In this context, learning involves both error reduction and model simplification. Learning will result in minimally complex models that have reasonable accuracy (i.e. *Satisficing* behavior).

Making Predictive Processing Quantitative

Two perspectives:

1. Bayesian inference

2. Markov Decision Process (Friston et al. 2012)

Bayesian predictive processing

Bayes rule:

$$P(B|E) = \frac{P(E|B) * P(B)}{P(E)}$$

Where B is our belief and E is evidence.

In predictive processing higher levels of the generative model make predictions about our beliefs ($P(B)$ or our prior). The lower levels of the hierarchy collect data, or Evidence (E) from a hypothesis space ($P(E)$). Our brain learns the likelihood of the evidence given the beliefs ($P(E|B)$). We can use Bayes rule to update $P(B|E)$, which is the prediction error.

Markov Decision Process (MDP)

Definition A Markov decision process is the tuple (S, A, T, r), where;

S is a finite set of states (s).

A is a finite set of actions (a).

$T(s'|s, a) = Pr(\{s_{t+1} = s'|s_t = s, a_t = a\})$ is the transition probability that the state $s' \in S$ at time $t + 1$ follows action $a \in A$ in state $s \in S$ at time t

$r(s)$ is some reward received at state $s \in S$.

Problem The goal is to find a *policy* $\pi : S \rightarrow A$ that maximizes the cumulative rewards. We can express this in terms of the sequence of actions that maximizes the value:

$$V(s) = \max_{a_{0:T}} \left\{ r(s) + \sum_{i=1}^{T} \sum_{s'} Pr(\{s_i = s; |s_0 = s, a_0, ..., a_i\}) r(s') \right\}$$

POMDPs and predictive processing

MDPs assume that the human knows what state they are in. This is often not reasonable. POMDP's relax this assumption to accomodate partially observed states.

Definition A Partially Observable Markov Decision Process is the tuple (S, A, T, r, Ω, O), where;

S, A, T, r are the same as the MDP formulation.

Ω is the finite set of observations or outcomes.

$O(o|s) = Pr(\{o_t = o|s_t = s\})$ is the observation probability of $o \in \Omega$ given the agent is in state $s \in S$ at time t.

POMDPs can be converted to an MDP using beliefs about the current state and Bayes rule: $b'(s') = P(s'|o, a, b) = \frac{P(o|s',a,b)P(s'|a,b)}{P(o|a,b)} \propto O(o|s', a) \sum s \in S T(s'|s, a)b(s)$.
We can now treat the beliefs as states and create a "Belief Markov Decision Process"

Definition A Belief Markov Decision Process is the tuple (B, A, T, r), where

B is a finite set of beliefs (b).

A is a finite set of actions (a).

$T(b'|b, a) = Pr(\{b_{t+1} = b'|b_t = b, a_t = a\})$ is the probability that the belief state $b' \in B$ at time $t + 1$ follows action $a \in A$ in belief state $b \in B$ at time t

$r(b) = \sum s \in S b(s) r(s)$ is the reward expected in belief state $b \in B$.

POMDPs and predictive processing

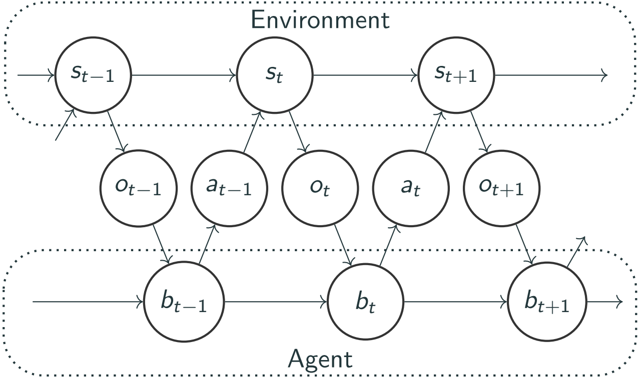

A pause...

A few important points:

1. Classical approaches to MPDs involve the optimization of future rewards to specify a policy in terms of an action from any given state. The human tries to maximize future rewards against a world that is governed by laws the agent can infer.
2. POMDPs explicitly model inference through a probabilistic mapping between hidden states of the world and observations. Thus beliefs can be exploited to optimize behavior.

POMDPs and predictive procesing

Definition The free-energy formulation refers to the tuple $(\Omega, S, A, \nu, P, Q, R)$ comprising:

A finite set of observations Ω.

A finite set of hidden states S.

Real valued parameters $\nu \ \in \ \mathbb{R}^d$.

A sampling probability $R(o'|o, a) = Pr(\{o_{t+1} = o'|o_t = o, a_t = a\})$ that observation $o' \in \Omega$ at time $t + 1$ follows action $a \in \ A$, given observation $o' \in \Omega$ at time t.

A generative probability $P(o, s, \theta|m) = Pr(\{o_0, ..., o_t\} = o, \{s_0, ..., s_T\} = s, \nu = \theta)$ over observations to time t, states at all times and parameters

A recognition probability $Q(s, \theta|\mu) = Pr(\{s_0, ..., s_T\} = s, \nu = \theta)$ over states at all times and parameters with sufficient statistics $\mu \in \mathbb{R}^d$.

In this presentation m is the form of a generative model or probability $P_m(o, s, \theta) := P(o, s, \theta|m)$. The sufficient statistics of the recognition probability $Q_\mu(s, \theta) := Q(s, \theta|\mu)$ encode a probability distribution over a sequence of hidden states and the parameters of the model $\theta \in \nu$. This recognition probability encode hidden states in the future and past, which themselves can change with time.

Distinctions between free-energy and MDP

There are several important distinctions between the free energy formulation and MDP:

1. The transition probability over states (from POMDP) is replaced with a sampling probability over observations. i.e. the agent does not need to know the actual result of their action on the world, just the coupling between actions and sensory consequences.
2. The free-energy formulation introduces generative and recognition probabilities used to infer hidden states, both past and future. i.e. the agent represents a sequence over states rather than just the current state.
3. There are no reward or cost functions in the free-energy formulation. Optimal behavior does not maximize rewards, rather it minimizes free-energy.

Translating active inference principles to decision-making models

1. We assume humans consider the task environment as a set of states, s. (e.g., "a crash is imminent", "a crash is not imminent")
2. Humans have probabilistic preferences, $P(s)$, about being in each state and will take actions, a, to remain in preferred states.
3. Humans have an internal model of the relationship between states and actions modeled as a predictive distribution, $Q(s|a)$.

Translating active inference principles to decision-making models

4. The Kullback-Leibler divergence (i.e., relative entropy) between the preferences, $P(s)$, and predicted states, $Q(s|a)$, is the Expected Free Energy (EFE or $\mathcal{G}(a)$) of an action.[19]
$$\mathcal{G}(a) = D_{KL}(\, Q(s|a) \| \, P(s)) := \textstyle\sum_{s \in \mathcal{S}} Q(s|a) \log \frac{Q(s|a)}{P(s)}$$
5. Humans make decisions/take actions to minimize the EFE. We can model this based on the relative log-likelihood between actions controlled by the precision, γ, where $\pi(a)$ is the probability of action, a.

$$\log \frac{\pi(a)}{\pi(a')} \propto -\gamma(\mathcal{G}(a) - \mathcal{G}(a'))$$

Translating active inference principles to decisions over time

6. We can extend the model to decision making/action in dynamic settings by assuming a Markovian predictive distribution, $Q(s_{t+1}|s_t, a_t)$. In this case EFE can be extended to:[20]

$$\mathcal{G}(a_{1:T}) = D_{KL}(\,Q(s_{1:T}|a_{1:T})\|\,P(s_{1:T}))$$

and action selection can be modeled as:

$$\pi(\,a_{1:T}|\,s_1) \propto \exp(-\gamma\mathcal{G}(a_{1:T}))$$

Translating active inference principles to partial observability

7. We can also extend the model to the case where the state is not observable but the human can collect observations, o, related to the state.[20] e.g., in driving a driver can observe looming but not the state of the forward vehicle (slowing to turn, stopping for a family of ducks, etc.). These observations can be exteroceptive, interoceptive, or proprioceptive.

Translating active inference principles to partial observability

8. In this case, the human interacts with the environment by performing an action, a_t, at each time step. The action influences a transition from the current state, s_t, to the next state, s_{t+1}, with dynamics governed by, $P(s_{t+1}|s_t, a_t)$. Since the human cannot observe the state, they infer it based on their observation, o_t, generated through the probability distribution, $P(o_t|s_t)$ (estimated with variational inference).
9. EFE can be modeled as:

$$\mathcal{G}(a_{1:T}) = D_{KL}(\, Q(s_{1:T}|a_{1:T}) \|\, P(s_{1:T})) + \mathbf{E}_{Q(s_{1:T}|a_{1:T})}[\mathbf{H}(P(o_{1:T}|s_{1:T}))]$$

Translating active inference principles to driving

1. We used looming as our observations.
2. We analyzed two states, wait (0) or brake (1).
3. We assume that drivers plan over a fixed time horizon.
4. We assume that the precision in the driver's beliefs changes with time.[21]

Active inference driver decision making model

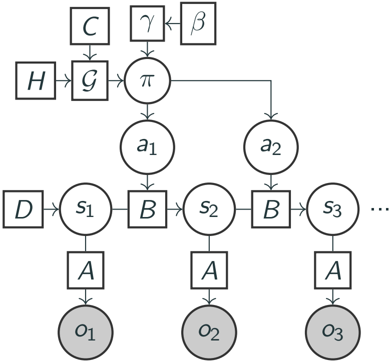

Parameter	Meaning	Context
A	Looming observation distributions $P(o\|s)$	The driver's learned association between the need to brake and looming.
B	State transition dynamics $P(s'\|s, a)$	The driver's understanding of the environment dynamics.
C	Expected state distribution $P(s)$	The drivers belief about the state of the environment.
D	Initial state belief $P(s_0)$	The drivers initial belief about the state of the environment.
H	Planning horizon H	The amount of time into the future that the driver considers.
β	Precision rate	The change in the driver's precision over time.

Fitting an active inference model to driving data

- 2x2x2 experimental design
- Within Subjects: Scenario (rear-end collision, obstacle avoidance), Kinematic urgency (critical/non-critical)
- Between Subjects: Alert Type (Silent failure, Alerted transition)
- Each driver drove 4 drives in a counterbalanced order

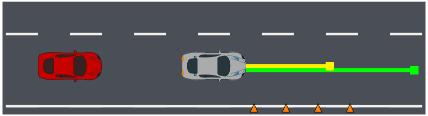

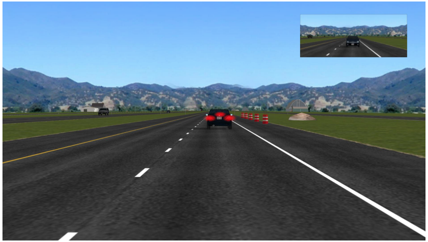

Compare model predictions and experimental results

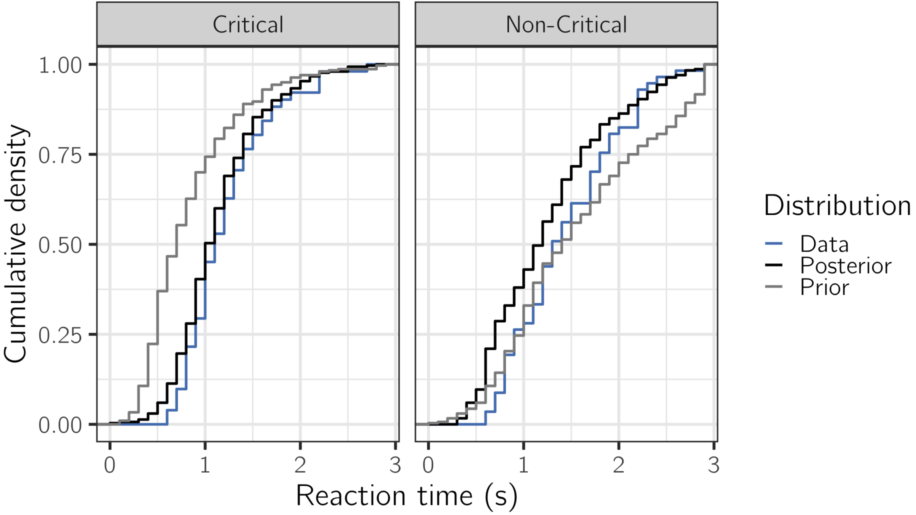

Extending the active inference driver model (AIDM)

Automation failure scenario:

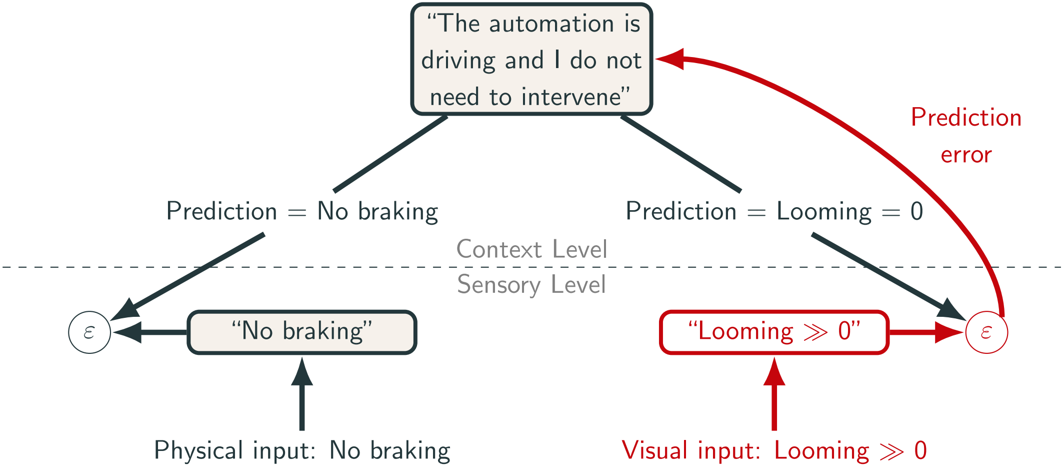

Extending the active inference driver model (AIDM)

As the driver begins to understand the failure:

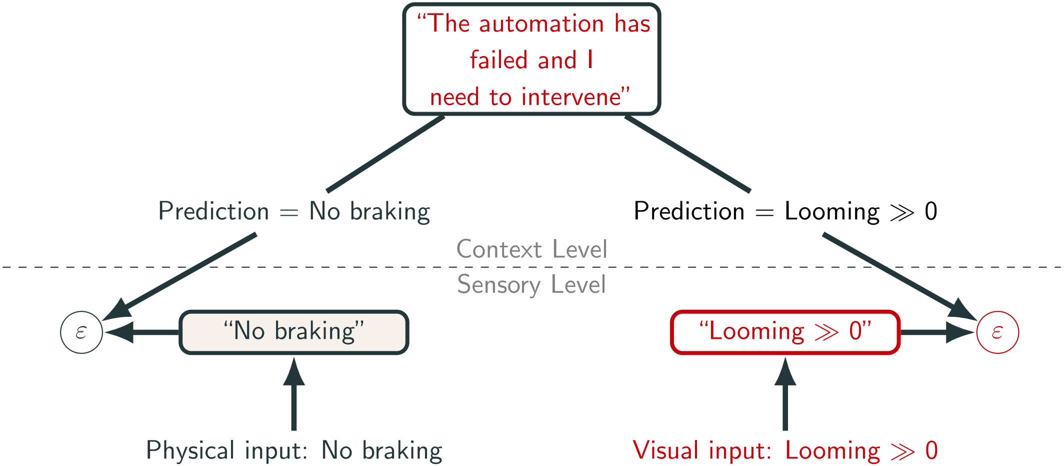

Extending the active inference driver model (AIDM)

What happens next?

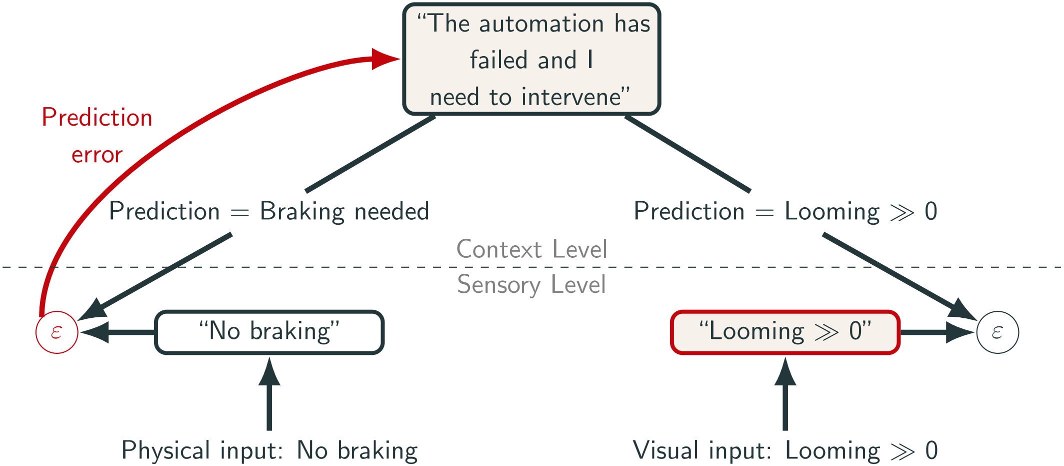

Extending the active inference driver model (AIDM)

To capture decision making and actions over time we need to expand the model.

1. Add an action module to map discrete decisions (accelerate, decelerate, wait) to continuous actions (acceleration).
2. Add additional features to represent the environment (relative speed to a lead vehicle, road geometry).
3. Add regularization to the optimization process to reduce identifiability issues.

Extending the active inference driver model (AIDM)

We examined this extension with the INTERACTION dataset:

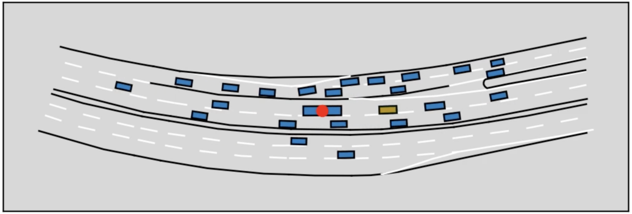

Compared the AIDM with a traditional rule based models (IDM), and 2 deep neural networks (behavior cloning (BC) with recurrence and without).

Trained the models on data from one lane and tested them on held out samples from the same lane and samples from a second lane.

Extending the active inference driver model (AIDM)

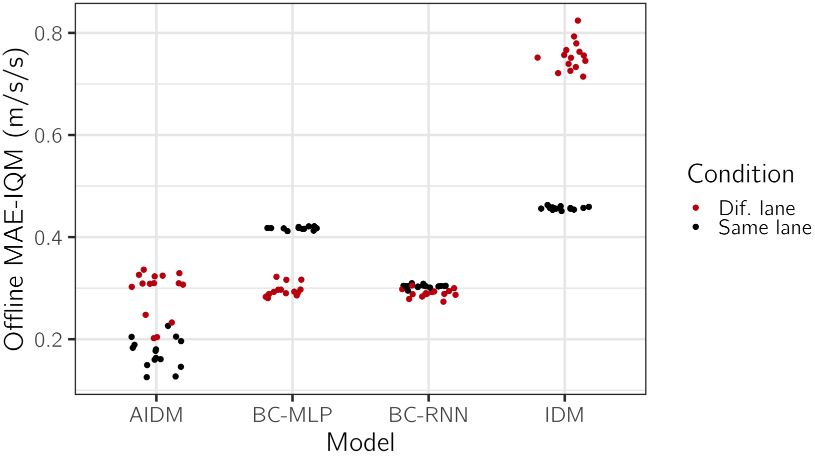

Extending the active inference driver model (AIDM)

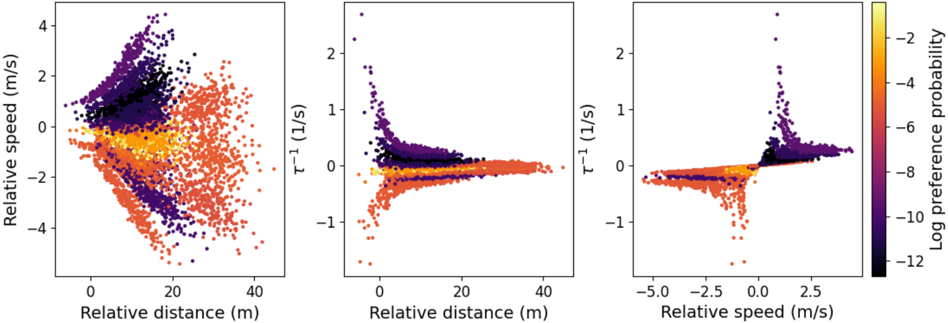

Future plans with Active Inference models

1. Continue to extend the model to capture continuous car following
2. Investigate the role of trust and situation awareness in more detail
3. Extend the model to continuous control scenarios in healthcare (e.g., care decision-making, telerobotic surgery)