Maximizing entropy of the Predictive Processing framework

Maria Otworowska, Iris van Rooij, Johan Kwisthout

Radboud University Nijmegen, Donders Institute for Brain, Cognition, and Behaviour

January 2018

Abstract

The Predictive Processing (PP) framework offers a unifying view on the existence and working of all living systems. The core premise of PP states that as long as agents minimize prediction error, and consequently entropy, they are successful. Current developments and advances in PP indicate that the interaction between agents and their environments is an important component of entropy minimization. In this paper, we explore by means of computer simulations, the interaction between PP-agents and their environments under different conditions. We argue the need to redefine the notion of success in PP in terms of entropy, behavioral and cognitive success, as we show that the environmental conditions that lead to entropy success, are different from conditions that lead to behavioral or cognitive success. Furthermore, we show that being equipped in and applying the mechanisms to minimize prediction error, do not in practice guarantee that the agents will be successful in any sense (entropy, cognitive or behavioral).

Conceiving of the brain as a prediction machine (Clark, 2013) has become a widely spread and influential view in contemporary cognitive neuroscience (Montague and Sejnowski, 1994; Grasso et al., 1996; Poppenk et al., 2010; Friston and Kiebel, 2009). A rapid development of the Predictive Processing (PP) framework has contributed to establishing this perspective, by offering a unified formal guide to explaining mechanisms underlying cognitive processes (Friston, 2009). According to PP, a brain (or an agent), makes predictions about its (sensory¹) inputs and minimizes the difference (prediction error) between these predictions and actual observations (Friston and Kiebel, 2009). This minimization happens either internally (agents update their current beliefs about the world) or externally (agents make a change in the environment to make the observations match

---

¹Here, sensory inputs include both exteroceptive and interoceptive inputs.

their predictions). Minimizing prediction error is a way by which an agent can minimize entropy²³, which is conjectured to be a unifying principle of how the brain works (Friston, 2009). Recent developments in PP framework indicate a crucial role of the environment in entropy minimization (Friston, 2011; Clark, 2015). The environment is assumed to be structured such that PP-agents make accurate and informative predictions, and select actions that influence the environment to minimize entropy (Friston et al., 2015). In this paper, we take the first attempt to investigate this assumption, research the interaction between specific environments and PP-agents, and analyze which conditions make the agents successful.

Traditionally, the role of the environment in achieving success by PP-agents was either ignored or, in the best case, implicitly taken for granted. An example of the first scenario is visible in case of comparator-based models (Allen and Friston, 2016) where actions are chosen based on the internal model, not actual observation (Clark, 2015). These accounts “[...] at best ignore embodied and enactive cognition and are at worst irreconcilable.” (Allen and Friston, 2016, p. 4). On the other side of the spectrum, environments are simply taken for granted to match with agents’ prior expectations and generative models via adaptation processes (Calvo and Friston, 2017). This lack of explicit and formal embedding of environments in the PP framework has led to, what Friston calls ‘revolutionary views’ (Allen and Friston, 2016, p. 5). Clark (2015) proposes to distinguish Conservative Predictive Processing (CPP) from Radical Predictive Processing (RPP). The main distinction between these two accounts lies in the different conceptualization of generative models, and what follows, the involvement of the environments in agents’ success. In CPP, generative models mirror the environment by building its detailed representation, with the assumption that the more detailed the ‘mirror’ is, the better the predictions are and the more successful the PP-agents are (Clark, 2015). The environment, here, is treated like a passive cue-sender or feedback provider to the agents about the quality of their predictions. RPP, on the other hand, postulates that generative models do not mirror the environment, but they are structural complements of the environment (Calvo and Friston, 2017; Clark, 2015), focusing on the dynamic nature of the interactions between agents and their environments. The RPP-models build up fast and frugal strategies that can quickly exploit and interact with environmental structures. Recently, Calvo and Friston (2017) presented a very illustrative example of applying RPP to explain interactions between plants (an example of the PP-agents⁴) and their environments:

---

²³Formally, entropy, here, refers to *conditional* entropy, a quantity conditioned on an agent’s model. This distinction is important conceptually — for different agents, different states are surprising. For example, for a fish, a surprising state will be to be out of the water, while for a cat the surprising state would be to be in the water. More on conditional entropy can be found in Supplementary materials

³For more information about how minimizing prediction error contributes to minimizing entropy, see Box 1.

⁴PP-agents are defined as any self-organizing systems, such as brains, humans, plants or cells (Calvo and Friston, 2017)

Vines provide the sharpest contrast with conservative PP scenarios. Grasping, say, a phone, according to Hohwy (2013) can only be explained via the exploitation of an internal representation of the phone itself. But once PP is fully developed with fast and frugal, morphological, and non-neurally decentralized resources, we can see that detailed internal representations may not be necessary: a climbing plant does not grasp a support because it has an internal representation of it, but rather because of the soft-assembling of the plant body-environment in the context of the approaching manoeuvres. [...] Plants are a good example of radical predictive processing because they are not in the business of enriching a neurocentric model of their surroundings (this is effectively outsourced to natural selection); rather, they are in the business of appropriately coupling with it. (Calvo and Friston, 2017, p. 9)

The emphasis on exploitability of environmental structures directs a new path of research within the PP framework. With our simulation studies, we aimed to contribute to this line of research, by analyzing the role of different environmental properties⁵ in achieving success by PP-agents.

In the PP-framework, the predominant view is that entropy minimization is identified with agents’ success; as long as agents minimize entropy, they, by definition, are successful on this view. From achieving *entropy success*, it follows that (a) the agents are able to survive because they minimize entropy, and (b) they are able to minimize entropy, because they have sufficient knowledge about the environment and themselves, to do so. (We analyze these implications in the next paragraph). The notion of entropy success comes from a proposed solution to the classical *dark room* problem (Friston et al., 2012). The dark room problem challenges entropy minimization assumption by describing a hypothetical situation in which an agent is sitting in a dark room, where it can perfectly predict all its sensory inputs. Why would the agent ever move from this spot where it can fully minimize the prediction error? The classical reply is that the agent’s prior expectation is *not* to be in a dark room. Sitting in a dark room would drastically violate these priors (Friston et al., 2012). This assumption instantly raises a concern, that has already been pointed out by Allen and Friston (2016, p. 15), that “any behaviour can be described as optimal, because one can always write down a prior that prescribes any behaviour in a ‘just so’ fashion”. Allen and Friston (2016), however, further explain and specify that the PP framework is free from this ‘just so’ circularity, because the framework is built on a “provable definition of living systems” (Allen and Friston, 2016, p. 15), that “living beings are those that maintain an upper bound on the entropy of their possible states” and “it follows that the only possible prior belief an agent

---

⁵In our simulations, we analyze more general environmental structures, not the structures specific to an individual or a species, like, e.g., in the case of vines and plants. The environmental structures exploited by vines are different from the structures exploited by humans.

can entertain is that it will behave so as to minimize surprise” (Allen and Friston, 2016, p. 16). The agents are equipped (by evolution, learning or neurodevelopment (Friston, 2012)) with prior expectations of how to act to minimize surprise, and consequently entropy. As long as the agents perform actions that fulfill these expectations, they are, by definition, successful.

Success defined in such a way, however, is not unproblematic. It relies on two implicit assumptions about what is needed for acting to minimize entropy or the entropy success. First, it requires that the agents need to have (in some sense) a veridical or otherwise ‘good’ model of their environments (however represented). Only then, the agents are able to reliably predict the consequences of their actions, and make the decision based on these predictions. Second, the outcomes of the actions need to be not only believed to be true by agents, but actually true in the environment. Otherwise, the agents would not even be able to survive. On the other hand, the agents can take actions based on a wrong or untrue model and still survive, without minimizing surprise.

For the purposes of our simulations, we separate these two implicitly assumed components of entropy success, and we propose to distinguish three definitions of success: behavioral, cognitive and entropy. (1) **Behavioral success**. For agents to be successful, minimally, they need to be able to survive. If the agents display behaviors that lead to their survival, we say they are behaviorally successful. (2) **Cognitive success**. The agents are successful in this sense if they learn causes of their sensory inputs and learn which action to perform in a given environmental situations. Agents’ knowledge is encoded in their generative models and expressed in terms of degrees of beliefs. (3) **Entropy success**. This notion of success is captured by conditional entropy conditioned on agents’ generative models. The lower the conditional entropy, the more successful agents’ models are in terms of predicting the sensory inputs. If the premise of the PP-framework holds true, our simulations should show that the entropy success entails cognitive and behavioral success.

The results from our simulations show that on the contrary, entropy success does not imply behavioral or cognitive success. This means that the distinction of the three notions of success is not only conceptual, but also factual. Our simulation design (for more detailed description, see the Supplementary materials) is relatively simple, but complex enough to show how different environmental conditions contribute to agents’ (behavioral, cognitive and entropy) success. We implemented PP-agents equipped with basic mechanisms to minimize prediction error. These mechanisms include the ability to change predictions to match observations, and the ability to act in the world to bring about the predicted observations⁶. The agents have a basic repertoire of actions and mental states;

---

⁶A wider variety of the mechanisms to minimize prediction error has been proposed by Kwisthout et al. (2017), but they are not part of the original, core assumption proposed by e.g. (Friston, 2009, 2011), and hence, we decided to leave them out in this version of the implementation.

and can perform actions, that in consequence of interacting with the environment, can change agents' mental states and the environment itself. The environment is populated by three types of objects: *edible*, *changeable* and *neutral*. Agents' task was to learn which objects to eat, and how to make an inedible, but changeable, object edible. This motivates to further distinguish three sub-definitions of success within cognitive and behavioral success (see Fig. 1): (1) successful eating; when agents learn which object to eat. (2) successful changing; when agents learn which object to change. (3) temporal success; when agents learn to combine (2) and (1): changing an object in time t , and eating the next object in time t_{+1} . For entropy success, we distinguish entropy related to an edible object and to a changeable object, separately. Furthermore, we implemented three types of environments: 'learnable and changeable', 'learnable only', and neither, to see whether the 'learnable and changeable' environment is always the environment that ensures agents' success. Only in 'learnable and changeable' environments are agents able to learn the generative models and change the environment to bring it closer to the predictions. Technically, only these environments are realistic. These three types of environments vary in terms of the distribution of edible objects, and the environments' exploitability⁷ (see Fig. 2). We analyze how the interaction between different environments and PP-agents determines agents' behavior and success on all dimensions: eating, changing, temporal across the three notions: behavioral, cognitive, and entropy⁸.

Our simulation results point to four important issues:

1. Being equipped with the mechanisms to minimize prediction error does not guarantee success in any sense.
2. Entropy success does not entail cognitive or behavioral success.
3. Different environmental conditions lead to different levels of success both across and within the same definition of success.
4. The agents don't know or learn which action to take to minimize the prediction error

The full analysis and description of the results can be found in the Supplementary materials. Here we give only a preview of the main points of the results and focus specifically on the 'learnable and

---

⁷Traditionally, the exploration-exploitation dilemma is considered an internal problem that agents need to face. Notably, this dilemma is claimed to be resolved by PP theory, as when to explore and when to exploit is implicitly built in agents' prior expectations: being in a surprising state enforces state-exploration until an unsurprising state is reached (Friston, 2010). In our simulation, the exploitability of an environment, or exploitation-exploration rate, is a (general) property of the environment. It means that there are some environmental structures (not directly identified, as they will differ from agent to agent) such that by interacting with these structures, an agent, in turn, will explore or exploit more.

⁸Given that entropy is measured over a longer period of time, we do not distinguish a temporal success here. As for the notions of eating and changing success, they map to entropy of an action given edible or changeable objects, respectively. For more details, see the Supplementary materials.

Figure 1: The figure represents the full overview of the notions and dimensions of success we distinguish. The first level: behavioral, cognitive and entropy success distinction is independent of our simulations. The eating, changing and temporal success level is specific to our simulation setup. It can be thought of as a descriptive measure to compare behavioral, cognitive and entropy success.

changeable’ environments only. The results show that from simply *trying*⁹ to minimize entropy, does not follow that agents will achieve entropy success, nor that they will be successful behaviorally or cognitively. Moreover, the results indicate that entropy success does not entail cognitive nor behavioral success. In our simulations, different conditions led to different levels of behavioral, cognitive and entropy success¹⁰. On average, for behavioral success, the most successful condition was **high exploitation** rate and **high food** distribution condition. For cognitive success, on average, the most successful condition was **high food** distribution and **high exploration** rate condition. For entropy success it was both the condition with **high exploitation** rate and **high food** distribution, as well as the condition with **high exploitation** rate and **low food** distribution. On a single dimensions of success (eating, changing, temporal), different conditions resulted in different level of success (cognitive, behavioral, entropy). For example, the most successful condition for entropy success related to the edible object, is the condition with **high exploitation** rate and **low food** distribution. For behavioral eating success, it is the condition with **high exploitation** rate and **high food** distribution; and for cognitive eating success, it is the condition with **high exploration** rate and **high food** distribution. Furthermore, we observe that environmental conditions that led to

⁹ *Trying* in a sense of being equipped in the mechanisms for minimizing entropy, and applying them.

¹⁰ When we refer to a ‘successful condition’, we mean a condition that has led to a greater success on a given dimension than any other condition on that dimension. There is no normative threshold for success. For example, under the best condition for behavioral eating success, the agents performed the successful action — eating — only in 30% of all situations. Intuitively, one could not consider it a success, but since there was no other condition that would lead to a greater success than this, we call this condition a successful condition.

CONDITIONS

uniform exploitation- exploration	uniform exploitation- exploration	uniform exploitation- exploration
uniform food distribution	high food distribution	low food distribution
high exploration	high exploration	high exploration
uniform food distribution	high food distribution	low food distribution
high exploitation	high exploitation	high exploitation
uniform food distribution	high food distribution	low food distribution

Figure 2: The figure represents all 9 conditions we implemented in our simulation. In total, we proposed two parameters: exploration-exploitation rate (see footnote 8) and food distribution. These two parameters can take three values: uniform, low and high.

agents’ success in terms of eating a correct object are different from conditions under which agents were successful in terms of changing a correct object, both across and within the same definitions of success. For example, the most successful condition for cognitive eating success was the condition with **high exploration** rate and **high food** distribution. While for cognitive changing success, the most successful was condition with **high exploitation** rate and **uniform food** distribution.

Moreover, we discovered that it is relatively easy for an agent to learn sensory effects of their causes; but it is not trivial to choose which sensory inputs to bring about next to best reveal their causes: which action to take to minimize the prediction error. In our simulations, the former refers to the conditional probability distribution of a sensory state (shape and color) given previously taken action (action_{t-1}); the latter refers to the conditional probability distribution of an action given sensory state (see Fig. 3). Notably, this problem has been recently acknowledged by Friston (2017)¹¹. The fact that PP-agents do not know which action to take to minimize prediction error, seems like a problem for the PP-framework; and especially for the assumption that entropy is minimized through action (Calvo and Friston, 2017).

¹¹Friston (2017, p. 11) points out that “The real problem [for an agent] is [...] to select the data that best discloses its causes”. The results of our simulations revealed this problem independently and before Friston’s claim was published.

Figure 3: The figure shows agents’ degrees of beliefs under the **most successful** environmental conditions per query. The left plot captures degrees of beliefs in the successful action given the observation. $P(\text{eat}|\text{green triangle})$ represents agents’ degrees of beliefs that the edible object (green triangle) is to be eaten. $P(\text{change}|\text{green square})$ represents agents’ degrees of beliefs that the changeable object (green square) is to be changed. $P(\text{eat}|\text{change}_{t-1})$ represent temporal knowledge of the agents, that after changing a changeable object, the agent should eat the edible object that appears next. We can see that the agents did not learn to eat the edible object. The degrees of beliefs, even under the best conditions, are still rather low (less than 0.3). The most successful condition for learning to change the changeable object led to only a relatively higher degrees of beliefs (less than 0.7). The right plot represents agents’ degrees of beliefs in ‘green’ color and ‘triangle’ shape in time t , given that the previous action in $t-1$ was ‘change’. We see that the agents expect with a very high certainty (0.9) that they will encounter green triangle in time t .

Given our results, it becomes apparent that the many implicit assumptions that the Predictive Processing framework is built on, upon a closer investigation, do not hold as promised. We can see that the conditions under which the agents achieve entropy success are different from conditions that lead to behavioral or cognitive success. Moreover, we can conclude that being equipped with mechanisms to minimize prediction error seems to be not enough in making decisions which action to take to actually minimize the prediction error. In terms of environmental conditions, we learned that conditions with high exploration rate, on average, lead to a greater cognitive success. And the conditions with high exploitation rate lead to a greater behavioral success. The food availability parameter is more important for behavioral than cognitive or entropy success. For example, for entropy success, the agents are equally successful both in conditions with low food distribution as well as in conditions with high food distribution. These results show that, contrary to the PP-assumptions, achieving success is not guaranteed by simply applying mechanisms to minimize prediction error, and that entropy success does not entail behavioral or cognitive success.

All in all, it is clear that the theory needs further specification, and revising implicit, unreliable assumptions is a starting point. We can speculate that (1) there are some underspecified parts of the theory missing, that once formally explicated will account for the problems, (2) there is *something* special about the environments, that by interaction with the special structures, it's *somehow* making agents successful. It is crucial for the future research to focus on exploring what this special *something* could be, and analyzing how the *somehow* could work.

To conclude, once PP mechanisms were implemented in a relatively simple setup, the premise that minimizing entropy will lead to agents success, does not hold anymore. It still remains an open question of how, by interacting with the environment, PP-agents are able to learn successful, or otherwise veridical, generative models.

References

- Micah Allen and Karl J Friston. From cognitivism to autopoiesis: towards a computational framework for the embodied mind. *Synthese*, pages 1–24, 2016.
- Paco Calvo and Karl Friston. Predicting green: really radical (plant) predictive processing. *Journal of The Royal Society Interface*, 14(131):20170096, 2017.

- Andy Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3):181–204, 2013.
- Andy Clark. Radical predictive processing. *The Southern Journal of Philosophy*, 53(S1):3–27, 2015.
- Karl Friston. The free-energy principle: a rough guide to the brain? *Trends in cognitive sciences*, 13(7):293–301, 2009.
- Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- Karl Friston. Embodied inference: Or i think therefore i am, if i am what i think. *The implications of embodiment (Cognition and Communication)*, pages 89–125, 2011.
- Karl Friston and Stefan Kiebel. Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1521):1211–1221, 2009.
- Karl Friston, Christopher Thornton, and Andy Clark. Free-energy minimization and the dark-room problem. *Frontiers in psychology*, 3, 2012.
- 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.
- Karl J Friston. Free energy and global dynamics. In Mikhail I Rabinovich, Karl J Friston, and Pablo Varona, editors, *Principles of brain dynamics: global state interactions*, pages 261–292. MIT Press, 2012.
- Karl J Friston. Predictive coding, active inference and sentient robots. In Pierre-Yves Oudeyer, editor, *The Newsletter of the Technical Committee on Cognitive and Developmental Systems*, volume 14, page 12. IEEE CIC, 2017.
- Renato Grasso, Stefan Glasauer, Yasuhiko Takei, and Alain Berthoz. The predictive brain: anticipatory control of head direction for the steering of locomotion. *Neuroreport*, 7(6):1170–1174, 1996.
- Jakob Hohwy. *The predictive mind*. Oxford University Press, 2013.
- Johan Kwisthout, Harold Bekkering, and Iris van Rooij. To be precise, the details don’t matter: on predictive processing, precision, and level of detail of predictions. *Brain and cognition*, 112: 84–91, 2017.

- P Read Montague and Terrence J Sejnowski. The predictive brain: temporal coincidence and temporal order in synaptic learning mechanisms. *Learning & Memory*, 1(1):1–33, 1994.
- J Poppenk, Morris Moscovitch, Anthony Randal McIntosh, E Ozcelik, and FIM Craik. Encoding the future: successful processing of intentions engages predictive brain networks. *Neuroimage*, 49(1):905–913, 2010.

Supplementary materials for:

Maximizing entropy of the Predictive Processing framework

In these supplementary materials we describe the details of the simulations. In Section 1, we present the setup of the simulations including formalization of PP-agents (Section 1.1), environments (Section 1.2) and conditions (Section 1.3). In Section 2, we explain the simulation procedure. Finally, we present results of the behavioral success (Section 3), cognitive success (Section 4), and entropy success (Section 5).

1 Setup

In this simulation study, we implemented a simple PP-agent embedded in a simple environment, and manipulate different environmental parameters to investigate which conditions make an agent successful. We define three notions of success: (a) behavioral success; (b) cognitive success; (c) entropy success. The agents have a basic repertoire of actions and mental states; and can perform actions, that in consequence of interacting with the environment, change agents' mental states. The environment is populated by three types of objects: *edible*, *changeable* and *neutral*. Edible objects can be eaten by an agent (that would lead to a change in a mental state of the agent), changeable objects can be changed by an agent to ensure the appearance of an edible object in next time frame. This means, if an agent changes a changeable object in time t , then an edible object appears in time $t+1$. Furthermore, there are three different types of environments: (i) causal (which is learnable and changeable), (ii) correlations (which is only learnable but not changeable), and (iii) random (which is neither learnable nor changeable), and two types of environmental parameters: (i) exploitation rate and (ii) accessibility to edible objects. In this section, we will first explain the reasons behind the choices we made in our implementation, and further, we will discuss the individual components of the simulation.

1.1 Agents

Each agent is formalized as a dynamic causal Bayesian network (CBN) representing the set of beliefs the agent has about the world. A CBN is a probabilistic graph model which represents causal relations between variables as a directed acyclic graph (DAG), where nodes represent variables and directed edges represent the causal relations between them (Pearl, 2009). The network consists of five variables: (1) Actions = ['change', 'eat', 'nothing'], (2) Mental state = ['full', 'hungry']; (3) Edible = ['yes', 'no']; (4) Color = ['green', 'red']; (5) Shape = ['triangle', 'square']; (6) Prev Action

$= \text{Action}_{t-1}$ (see Figure 1a). The network can further be divided into two levels of a hierarchy. At each level of the hierarchy, we can distinguish hypothesis nodes and prediction nodes (Kwisthout et al., 2017). The value of prediction nodes given hypothesis nodes is a *prediction* at that level. Predictions at a higher level are hypothesis nodes at the level below.

The PP-agents try to minimize entropy by minimizing prediction errors between predictions and observations. The prediction error is used to update and revise agents’ beliefs and to update agents’ hyperpriors (see Section 2.3). These hyperpriors represent the information the beliefs are based on; that is, they describe probability densities over the probabilities that indicate how confident the agents are that the probability of the belief is correct. Hyperpriors are described by hyperparameters, which can be based on frequency information as well as so-called pseudo-counts or ‘virtual observations’ that represent agents’ background knowledge (Koller and Friedman, 2009).

Figure 1: Figure 1a represents an agent’s model about the world in a form of a causal Bayesian network. There are two kinds of nodes: regular (Color, Shape, Mental State, Edible, Action) and temporal (Previous Action). Temporal node takes a value of a node Action from a previous time frame. Figure 1b represents the model in a form of a two-level hierarchy, where the predictions on the level 1 are hypothesis on the level 2.

1.2 Environments

There are three types of environments: causal, correlations and random. Each environment consists of a set of objects $\{ \blacktriangle \blacksquare \blacktriangle \blacksquare \}$, some of the objects are edible $= \blacktriangle$, and some changeable $= \blacksquare$.

1.2.1 Causal Environment

There is a temporal- and action- dependency between changeable and edible objects. If an agent performs action ‘change’ on a ‘changeable’ object in time t , then in time t_{+1} , the edible object will appear. In case on non-edible and non-changeable objects, i.e., $\blacksquare \blacktriangle$, the action performed by an agent does not influence the next observation. Below we present a list of action-observation dependencies:

- $\blacktriangle$ AND ACTION=“EAT” $\rightarrow$ MENTAL STATE = “FULL”, NEXT OBSERVATION: RANDOM
- $\blacktriangle$ AND ACTION!=“EAT” $\rightarrow$ MENTAL STATE = “HUNGRY”, NEXT OBSERVATION: EDIBLE=“NO”
- $\blacksquare$ AND ACTION=“CHANGE” $\rightarrow$ MENTAL STATE = “HUNGRY”, NEXT OBSERVATION: EDIBLE=“YES”
- $\blacksquare$ AND ACTION!=“CHANGE” $\rightarrow$ MENTAL STATE = “HUNGRY”, NEXT OBSERVATION: EDIBLE=“NO”
- $\blacksquare$ or $\blacktriangle$ $\rightarrow$ MENTAL STATE = “HUNGRY”, NEXT OBSERVATION: EDIBLE=“NO”

1.2.2 Correlations Environment

There is only a temporal dependency between a changeable and an edible object. The edible object always appears after the changeable one, regardless of the action taken by the agent.

- $\blacktriangle$ $\rightarrow$ MENTAL STATE = “FULL”, NEXT OBSERVATION: RANDOM
- $\blacksquare$ $\rightarrow$ MENTAL STATE = “HUNGRY”, NEXT OBSERVATION: EDIBLE=“YES”
- $\blacksquare$ or $\blacktriangle$ $\rightarrow$ MENTAL STATE = “HUNGRY”, NEXT OBSERVATION: RANDOM

1.2.3 Random Environment

There is no temporal dependency between changeable and edible objects. They appear at random, and the actions taken by the agents cannot influence the environment.

1.3 Conditions

We define two environmental parameters: the exploitation rate and food accessibility. In Figure 2b, the x-axis represents a dimension of food accessibility, or a ratio of not edible to edible objects. This parameter can take three values: (i) uniform distribution of food and no-food; (ii) more food (5 times

more food than no-food); (iii) less food (5 times less food than no-food). The y-axis represents the exploitation rate. This dimension can be interpreted in two ways. First, it can be seen as an agent’s internal bias to exploit or to explore more. Second, it can be interpreted as a result of an agent’s interaction with the environment; meaning there is something in the environment that makes the agent exploit or explore more. The nature of this ‘something’ is left without further specification. The simulation results for exploitation rate are interpreted in the second sense (see the main text). There are some structures in the environment that would make agents explore or exploit more. For the simplicity of the simulations, these structures are not defined. The exploitation rate can, similarly to food availability, be (i) uniformly distributed; (ii) low: more exploration (0.8 chance); (iii) high: more exploitation (0.8 chance). In total, there are nine conditions $\times$ three environments. Figure 2 represents the interaction of these two dimensions resulting in the definition of individual conditions (1 – 9).

Figure 2: The figures show (a) verbally described environmental conditions (the left figure), and (b) a numerical representation of (a). For example, the environment with high exploration rate and high food distribution is referred to as condition 5.

2 The simulations

In this section, the basic simulation procedure is explained. Each of the agents is given the network structure and needs to learn the conditional probabilities. The agents are granted unlimited memory – they remember all the observations and actions they take. However, to decide on the action, an

agent has access only to its previous action and a current observation. The agents start off with uniform hyperpriors and uniform probability distributions. The lifespan of each agent takes 300 time units (300 runs). In every t , an agent, based on its generative model tries to predict the next observation and choose the best action given the prediction. Then, given the observation, an agent (1) performs an action; and (2) calculates prediction error (the divergence between prediction and observation). In case of causal environment (CE) agents, (1) results in (a) determining the observation in t_{+1} , (b) changing agent’s current mental state, and (c) updating agents memory. (2) results in (a) hierarchical updating of the agent’s beliefs, and (b) updating the generative model by adjusting hyperpriors. At the end of each time t , the agent updates its probabilities distributions.

2.1 Action Selection

An agent performs the action based on the current observation. There are two ways of deciding on the action. One of these two ways is chosen depending on the exploitation rate parameter. If the exploitation parameter is high, and the agent will exploit, rather than explore, then the agent chooses its actions that have the highest probability. If the exploitation parameter is low, and the agent will explore rather than exploit, then the action is chosen based on the whole probability distribution. For example, if the distribution of actions is [‘eat’ = 0.2, ‘change’ = 0.7, ‘nothing’ = 0.1], then the next action would be a weighted random choice with [0.2, 0.7, 0.1] weights.

The observation in time t_{+1} depends on the action taken by an agent in time t . The observation is defined in terms of whether or not the next object will be edible. If the observation in time t was a green square, and the action taken by the agent was ‘change’, then (in case of CE agents), the next observation will be an edible object. The only edible object is a green triangle, so the next observation will be a green triangle. If the next observation is determined to be not edible, then the actual shape and color of the next object are assigned randomly.

Moreover, with every performed action, there is a change (or not) in the mental state of the agent. For example, if the agent eats a green triangle, then it instantaneously changes its mental state to ‘full’. It means that change still happens in time t . In time t_{+1} , the agent will be ‘hungry’ again.

2.2 Hierarchy

As it’s shown in Figure 1b, the network is divided into two levels of the hierarchy. At each level, there are a hypothesis node and a prediction node. Formally, a hypothesis node is a prior joint probability distribution, and a prediction node is a posterior probability distribution (Kwisthout

et al., 2017). The prediction nodes at the higher level, are the hypothesis nodes at the lower level.

When observing an object, an agent calculates a difference, so-called prediction error, between the observed (or inferred) distribution (P_{Obs}) and predicted distribution (P_{Pred}) at each level of the hierarchy. Formally, the distance is defined as Kullback-Leibler divergence $D_{KL} = (P_{Obs}||P_{Pred})$. The result is the prediction error, i.e., the net difference between prediction and observation, and it is used to update the agent’s generative model (Kwisthout et al., 2017).

After the prediction error is computed, that amount is used to update agent’s beliefs about the world. This means to find such a value assignment for hypothesis nodes that ensures that the prediction error is minimized. To calculate this, we use a local search algorithm for minimizing a scalar function. The output is a new value assignment for a hypothesis node and the new divergence measure. Since there are two levels of the hierarchy, some of the nodes (Color, Shape and Mental State) are both the hypothesis and the prediction node. To update the beliefs on both levels, we need to recalculate probabilities distributions of Color, Shape and Mental State. We do it as follows:

1. Calculate marginal probability of a variable V , $P(V)$ from the original probability distribution

$$P(V) = \sum_p P(V|p) \times P(p)$$

2. Calculate marginal probability of a variable V , $P'(V)$ after revision, here $P'(V)$ doesn’t have any parents, since it’s a hypothesis node.

$$P'(V) = val$$

3. Calculate the ratio of $P(V)$ to $P'(V)$

$$\frac{P(V)}{P'(V)} = ratio$$

4. Calculate a new $P(V)$:

$$\frac{P(v|p)}{ratio} \times P(p) + \frac{P(v|\neg p)}{ratio} \times P(\neg p)$$

After belief revision, we perform a basic form of model revision (Kwisthout et al., 2017), that involves recalculating of conditional probabilities given the new observation. Lastly, the hyperpriors are updated.

2.3 Hyperpriors

Hyperpriors represent second-order information about the beliefs and are represented in the form of (conditional) hyperparameters of a Dirichlet distribution that describes the confidence the agents have that the probability assessment is correct. These hyperparameters are based on both sequences of observations (frequency information or *counts*) as well as background or innate knowledge (*pseudocounts*) that alter the frequency information. For example, for a variable V with possible states $V = [v_1, v_2]$, and a sequence of observations, $O = [v_1, v_1, v_1, v_2]$, the probability distribution that can be derived purely from the frequency information without hyperpriors is $P(V = v_1) = \frac{\#v_1}{\#total} = \frac{3}{4}$ and $P(V = v_2) = \frac{\#v_2}{\#total} = \frac{1}{4}$. In this case, $P(V = v_2) < P(V = v_1)$. Pseudocounts for variable V , PC_V , are defined as a tuple representing a number of additional “virtual” observations for each variable state: $PC_V = [(v_1, 1), (v_2, 5)]$. The probability distribution, when taken these pseudocounts into account changes as follows: $P(V = v_1) = \frac{3+1}{4+6} = 0.4$ and $P(V = v_2) = \frac{1+5}{4+6} = 0.6$. In this case, $P(V = v_2) > P(V = v_1)$. For PP-agents, if an agent performs a successful action (e.g., eats a green triangle), then this will be added both to counts ($\#(eat \wedge green\ triangle) + 1$ for eating a green triangle), and to pseudocounts ($PC(eat \wedge green\ triangle, +1)$ for eating a green triangle). In case of an unsuccessful action, only counts will change.

2.4 The notions of success

Given the setup and that the agent’s task is to learn when and what to eat, there are three ways in which an agent’s behavior can be successful: (1) successfully eating the edible object, (2) successfully changing the changeable object, (3) performing a sequence of (2) and (1). Furthermore (1)–(3) can be assessed on the (a) ‘behavioral’ or (b) ‘cognitive’ scale, expressed in terms (a) taking successful actions, (b) degrees of beliefs about the world. Finally, an agent can be successful in terms of (4) minimizing conditional entropy. We are interested whether the same environmental conditions will result in agents’ success on all of the dimensions.

3 Results: Behavioral Success

As explained in previous sections, there are three ways in which an agent can be successful: (1) behavioral success, (2) cognitive success, and (3) entropy success. In what follows, we analyze (1)–(3), present the results and provide a summary and conclusions.

3.1 ‘Behavioral Success’: General overview

The basic measure of the behavioral success – *absolute* success – can be expressed in terms of simple counting: how many times an agent performed a successful action given an observation. Here, a successful action is the action that is considered correct given environmental conditions. Henceforth, a successful action will be referred to as a correct action. Each agent always performs 300 actions. The total number of actions per condition is 7800, and per environment 70200. The second measure of success is a relative measure – *ratio* – of success: how many times an agent performed the correct action given an observation to the times it had had the chance to perform it. We distinguish:

1. **Successful eating:** performing action ‘eat’ given ‘green triangle’ (see Table 1)
2. **Successful changing:** performing action ‘change’ actions given ‘green square’ (see Table 2)
3. **Temporal success:** Performing the sequence of actions (see Table 3):
  - (a) Action ‘change’ actions given ‘green square’ in time t_{-1}
  - (b) Action ‘eat’ given ‘green triangle’ in time t

Tables 1 to 3 show the absolute and relative scores, together with a mean and a standard deviation measure. We can see that in case of **successful eating**, correlations environment (CorrE) agents are the most successful. However, the average ratio of success is rather low: only in 30% of the time the agents used the opportunity to eat, when observing green triangle. A similar trend is visible for random environment (RE) agents. For causal environment (CE) agents, the success is low, even comparing to the other environments agents. The agents performed the correct action only in 13% of the time.

In case of **successful changing**, the causal environment (CE) agents improved. Here, all three environments (CE, CorrE, RE) agents performed similarly, performing the successful action almost 50% of the time. In case of **temporal success**, again all three environments agents performed the correct sequence of actions around 11% (causal)–15% (correlations) of the time.

	Absolute	Mean	SD	Ratio
causal	6296	26.906	57.764	0.135
correlations	13385	57.201	69.217	0.309
random	11255	48.098	70.288	0.302

	Absolute	Mean	SD	Ratio
causal	4590	19.615	20.037	0.482
correlations	4331	18.509	17.509	0.486
random	5204	22.239	25.277	0.460

	Absolute	Mean	SD	Ratio
causal	447	1.91	5.444	0.112
correlations	1304	5.573	10.006	0.151
random	583	2.491	5.055	0.133

Table 1: **Successful eating.** The absolute number of successes (action ‘eat’ when ‘green triangle’ was observed), mean and standard deviation, and the ratio of performed successful action to the number of potential successful actions (the cases when an agent observed ‘green triangle’, but it didn’t ‘eat’ it).

Table 2: **Successful changing.** The absolute number of successes (action ‘change’ when ‘green square’ was observed), mean and standard deviation, and the ratio of performed successful action to the number of potential successful actions (the cases when an agent observed ‘green square’, but it didn’t ‘change’ it).

Table 3: **Temporal Success.** The absolute number of temporal successes (action ‘change’ when ‘green square’ was observed in time t_{-1} and action ‘eat’ when ‘green triangle’ was observed in time t), mean and standard deviation, and the ratio of performed successful sequence of actions to the number of potential successful sequence of actions (the cases when an agent ‘changed’ ‘green square’, but didn’t ‘eat’ ‘green triangle’).

3.2 Scores per conditions

In the previous subsection, we compared the ratio scores between different environments. We observed that for successful eating and temporal success, the ratio scores were much lower in the causal environment compared to the correlations and random environments. In this section, we analyze how the different conditions contributed to the success of agents. We look at (a) the differences between ratio scores (Tables 4 to 6) per condition per environment, (b) whether the difference between different environments in (a) is significant (Tables 7 to 9). To analyze (b), we applied two-sided t-test for the null hypothesis that two independent samples have identical average values. We ran a pair-wise comparison between: (i) causal-correlations, (ii) causal-random, and (iii) correlations-random environments. Assuming p-value threshold of 0.05, the null hypothesis can be rejected for most of the conditions for successful eating ratio scores and for some of the conditions for temporal success ratio scores, but only between causal and correlations and causal and random environments. These conditions mostly include high exploration rate, or uniform exploitation-exploration rate. If we take p-value threshold to be 0.01, then the significant differences appear mainly for successful eating scores. This includes conditions poor in food (condition 3, 6, 9) or with uniform food distribution (condition 1, 4). It seems that for conditions with high exploitation rate, there are no significant differences across the different environments.

3.2.1 Successful eating

In case of causal environment (CE) agents, the most successful conditions are fully determined by accessibility to food (conditions 2, 5 and 8), it means that the exploitation rate didn't make a difference here. A high level of food, combined with a high exploitation rate (condition 8), creates the best condition for causal environment agents to be successful (almost 30% success score). Condition 8 is also a condition with the highest standard deviation (0.43). CorrE and RE agents are also successful in this condition (31%- and 42%). There is no significant difference between these scores.

There is no significant difference between correlations environment (CorrE) agents and random environment (RE) agents, across all conditions. For CorrE agents, the most successful condition is rich in food and has a higher exploration rate (condition 5), the standard deviation for this condition is also relatively high (0.27). For RE agents, many conditions seem to be (relatively) successful, with the most successful being condition 8. The standard deviation in case of RE agents was always high, across all conditions.

Interestingly, for CE agents, almost all (except condition 7), conditions that were not (relatively) successful, were especially unsuccessful. Conditions 1, 3, 4, 6 and 9 yielded poor results. These

conditions are low in food and/or have low exploitation rate.

3.2.2 Successful changing

For CE agents, condition 7 is the most successful, reaching the score of almost 70% of correct actions. The success under this condition is determined by the high exploitation score (the food distribution is uniform here). On average, CE agents scored similarly. There is no significant difference between the different conditions. For CorrE agents and RE agents, the best condition was condition 8 (high exploitation, rich in food). None of the conditions is particularly unfavorable for performing successful actions. Condition 3 that was the worst for CE agents for eating, now it's (48%) successful. The standard deviation is high across all conditions and environments.

3.2.3 Temporal success

For CE agents, condition 8 was the most successful condition in terms of performing the right sequence of actions: (1) 'change' green square and 'eat' green triangle. The standard deviation for this condition is high (almost 0.4). For CorrE agents, condition 8 was also the most successful condition, with an equally high standard deviation. For RE agents, conditions 1 (uniform for both food and exploitation rate) was the most successful (with a high standard deviation). Significantly worse for CE agents, compared to CorreE and RE agents, were conditions 1, 3, and 6.

	condition 1	condition 2	condition 3	condition 4	condition 5	condition 6	condition 7	condition 8	condition 9
causal	0.046	0.239	0.028	0.094	0.215	0.085	0.143	0.295	0.059
correlations	0.337	0.302	0.285	0.29	0.357	0.312	0.314	0.31	0.28
random	0.347	0.188	0.363	0.369	0.315	0.328	0.198	0.42	0.154

Table 4: The ratio score for **successful eating**. The times when agents performed the successful action ‘eat’ when observing ‘green triangle’, to the times when the agents observed ‘green triangle’, but took different action. The ratio score is an average ratio per condition. For causal and random environments agents, the most successful condition was condition 8 (high exploitation rate, rich in food). The least successful condition for causal environments agents was condition 3 (uniform exploration-exploitation rate, poor in food). The least successful condition for random and correlations environments agents was condition 9 (high exploitation rate, poor in food). The most successful condition for correlations environment agents was condition 5 (high exploration rate, rich in food)

	condition 1	condition 2	condition 3	condition 4	condition 5	condition 6	condition 7	condition 8	condition 9
causal	0.406	0.338	0.483	0.34	0.382	0.311	0.668	0.496	0.406
correlations	0.521	0.457	0.491	0.338	0.389	0.319	0.547	0.733	0.57
random	0.521	0.402	0.537	0.357	0.354	0.423	0.487	0.603	0.5

Table 5: The ratio score for **successful changing**. The times when agents performed the successful action ‘change’ when observing ‘green square’, to the times when the agents observed ‘green square’, but took different action. The ratio score is an average ratio per condition. For causal environments agents, the most successful condition was condition 7 (high exploitation rate, uniform food distribution). The most successful condition for correlations and random environment agents was condition 8 (high exploitation rate, rich in food). The least successful condition for causal and correlations environments agents was condition 6 (high exploration, poor in food). The least successful condition for random environments agents was condition 5 (high exploration rate, rich in food).

	condition 1	condition 2	condition 3	condition 4	condition 5	condition 6	condition 7	condition 8	condition 9
causal	0.056	0.146	0.021	0.086	0.143	0.069	0.163	0.215	0.05
correlations	0.129	0.142	0.142	0.088	0.122	0.122	0.194	0.258	0.16
random	0.207	0.069	0.204	0.124	0.11	0.125	0.118	0.171	0.087

Table 6: The ratio score for **temporal success**. The times when agents performed the successful sequence of actions: ‘change’ when ‘green square’ in time t_{-1} and ‘eat’ when ‘green triangle’ in time t . The ratio score is an average ratio per condition. For causal and correlations environments agents, the most successful condition was condition 8 (high exploitation rate, rich in food). The most successful condition random environment agents was condition 1 (uniform exploitation rate, uniform food distribution). The least successful condition for causal environments agents was condition 3 (uniform exploitation-exploration rate, poor in food), for correlations environments agents was condition 4 (high exploration, uniform food distribution), and for random environments agents was condition 2 (uniform exploitation-exploration rate, rich in food).

	SD_{CS}	SD_{CR}	SD_R	pval(CS-CR)	pval(CS-R)	pval(CR-R)
condition 1	0.059	0.33	0.401	6.874e-05	0.0004271	0.891
condition 2	0.353	0.367	0.307	0.5915	0.5843	0.271
condition 3	0.035	0.277	0.389	4.425e-05	6.915e-05	0.3563
condition 4	0.077	0.173	0.299	2.776e-06	3.337e-05	0.2437
condition 5	0.218	0.267	0.295	0.04852	0.181	0.6183
condition 6	0.061	0.16	0.269	1.416e-08	3.717e-05	0.776
condition 7	0.325	0.411	0.383	0.1216	0.6192	0.3119
condition 8	0.433	0.433	0.476	0.9345	0.3524	0.3939
condition 9	0.122	0.392	0.33	0.009146	0.1396	0.2582

Table 7: The standard deviation and the significance scores for successful eating. There is a significant difference between causal (CS) and correlations (CR) environments ratio scores for conditions (in red) 1, 3–6, and 9. And between causal and random (R) environments ratio scores for condition 1, 3, 4, and 6. There is no significant difference between ratio scores of correlations and random environments agents.

	SD_{CS}	SD_{CR}	SD_R	pval(CS-CR)	pval(CS-R)	pval(CR-R)
condition 1	0.4	0.297	0.394	0.4801	0.6906	0.7942
condition 2	0.369	0.29	0.365	0.5594	0.8941	0.4624
condition 3	0.401	0.356	0.412	0.682	0.9604	0.6486
condition 4	0.313	0.159	0.253	0.6692	0.925	0.7067
condition 5	0.265	0.231	0.235	0.768	0.5843	0.7844
condition 6	0.303	0.19	0.254	0.7706	0.288	0.1013
condition 7	0.395	0.434	0.472	0.1217	0.05834	0.6771
condition 8	0.416	0.289	0.45	0.2069	0.7961	0.138
condition 9	0.482	0.427	0.475	0.5045	0.9822	0.4863

Table 8: Standard deviation and significance scores for successful changing. There is no significant difference among ratio scores of causal (CS), correlations(CR) and random (R) environments agents.

3.3 Summary of the results

For CE agents, the most successful conditions are 8 and 7 (both high exploitation rate). The least successful are 3 and 9 (low in food). For CorrE agents, best conditions are also 5 (high exploration) and 8 (high exploitation), the least 9, 6 and 4 (low or uniform in food). For RE agents, best are 8 (high exploitation) and 3 (low food). The least successful is 9 (low in food, high exploitation), 5 (high food, exploration) and 2 (high in food).

On average, for all three types of agents, condition 8 was the most successful. This condition is rich in food and has a high exploitation rate. As for the least successful condition, there is no clear

	SD_{CS}	SD_{CR}	SD_R	pval(CS-CR)	pval(CS-R)	pval(CR-R)
condition 1	0.083	0.182	0.307	0.03933	0.02394	0.3661
condition 2	0.329	0.194	0.16	0.3286	0.031	0.0903
condition 3	0.039	0.195	0.318	0.003127	0.007252	0.4567
condition 4	0.084	0.093	0.166	0.3878	0.08965	0.2699
condition 5	0.208	0.163	0.109	0.3074	0.09414	0.5147
condition 6	0.057	0.103	0.153	0.02684	0.03064	0.6087
condition 7	0.306	0.337	0.308	0.5294	0.8477	0.4193
condition 8	0.384	0.372	0.335	0.9408	0.6489	0.5842
condition 9	0.12	0.334	0.213	0.06284	0.5822	0.1792

Table 9: Standard deviation and significance scores for temporal success score. There is a significant difference between causal (CS) and correlations (CR) environments ratio scores for conditions (in red) 1, 3, and 6. And between causal and random (R) environments ratio scores for condition 1,2, 3, and 6. There is no significant difference between ratio scores of correlations and random environments agents.

	Eat	Change	Temporal
causal	8	7	8
correlations	5	8	8
random	8	8	1

Table 10: Most successful conditions

	Eat	Change	Temporal
causal	3	9	3
correlations	9	6	4
random	9	5	2

Table 11: Least successful conditions

winner, however, it seems that it’s mostly determined by food accessibility. All the conditions low in food (3, 6, 9) are among the least successful.

4 Results: Cognitive Success

In this section we present a measure of agents’ internal success, cognitive success, expressed in terms of the degrees of beliefs of an agent about the world. In what follows, first, we query agents’ model (an average model per condition) to check how well the agents’ model captures the true world model. For the purpose of analyzing conditional probabilities of ‘Action’, we follow the notion of ‘correct action’ from the previous measure. Further, we present a measure based on the degrees of beliefs – a query score.

4.1 Successful eating

$P(\text{Action}|\text{Color} = \text{green}, \text{Shape} = \text{triangle})$

Figure 3: The plots show probability distributions of action given green triangle under different conditions and in different environments. Each row of the plots represents different environment: causal, correlations and random. Each column represents different values of exploitation-exploration rate parameter: uniform rate, high exploration, high exploitation. Lastly, every single plot represents different values of the second parameter: food distribution. The bars represent degrees of beliefs of agents in which action to take given the observation.

Figure 3 represents the probability distribution $P(\text{action}|\text{color} = \text{green}, \text{shape} = \text{triangle})$ of an average agent per condition. In all of the conditions, the probability of 'eat' (which would be the successful action) is lower than 0.5. Only in conditions 5 (high exploration rate and rich in food)

for CorrE agents, and conditions 4 (uniform food distribution, high exploration) and condition 8 (high exploitation rate, rich in food) for RE agents, the probability of ‘eat’ was higher than any other possible action. The probability distributions for CE agents implies that the agents did not learn to ‘eat’ ‘green triangle’. Below, we analyze conditions for the cases when the agents actually did ‘eat’ ‘green triangle’ which resulted in the change in their mental state.

Additionally, conditioning action on agent’s mental state, allows us to analyze which conditions were the most successful only in the cases when agents did ‘eat’ ‘green triangle’. Now, we can see that condition 5 (high exploration, rich in food) was the most successful, yielding the highest probability of ‘eat’. This is interesting because according to the measure of success from the previous section, condition 5 was not the most successful (avg. of 0.2). The probability distributions for CorrE and RE agents gives similar results as for $P(action|color = green, shape = triangle)$.

In Figure 5 we see that all types of agents learned (with a very high probability, especially for CE agents), that ‘green triangle’ is edible.

$P(\text{Action} | \text{Color} = \text{green}, \text{Shape} = \text{triangle}, \text{Mental State} = \text{full})$

Figure 4: The plots show probability distributions of action given green triangle and mental state under different conditions and in different environments. Each row of the plots represents different environment: causal, correlations and random. Each column represents different values of exploitation-exploration rate parameter: uniform rate, high exploration, high exploitation. Lastly, every single plot represents different values of the second parameter: food distribution. A combination of exploitation rate and food distribution constitutes a condition. Each row contains 1 (left)–9 (right) conditions. For example, condition 4 is a condition with high exploration rate and uniform food distribution. The bars represent degrees of beliefs of agents in which action to take given the observation.

$P(\text{Edible} | \text{Color} = \text{green}, \text{Shape} = \text{triangle})$

Figure 5: The plots show probability distributions of whether the green triangle is edible, under different conditions and in different environments. Each row of the plots represents different environment: causal, correlations and random. Each column represents different values of exploitation-exploration rate parameter: uniform rate, high exploration, high exploitation. Lastly, every single plot represents different values of the second parameter: food distribution. A combination of exploitation rate and food distribution constitutes a condition. Each row contains 1 (left)–9 (right) conditions. For example, condition 4 is a condition with high exploration rate and uniform food distribution. The bars represent degrees of beliefs of agents in whether the object is edible given the observation.

4.2 Successful changing

$P(\text{Action} | \text{Color} = \text{green}, \text{Shape} = \text{square})$

Figure 6: The plots show probability distributions of action given green square under different conditions and in different environments. Each row of the plots represents different environment: causal, correlations and random. Each column represents different values of exploitation-exploration rate parameter: uniform rate, high exploration, high exploitation. Lastly, every single plot represents different values of the second parameter: food distribution. A combination of exploitation rate and food distribution constitutes a condition. Each row contains 1 (left)–9 (right) conditions. For example, condition 4 is a condition with high exploration rate and uniform food distribution. The bars represent degrees of beliefs of agents in which action to take given the observation.

In the case of choosing the successful action for ‘green square’ (Fig. 6), almost all CE agents were successful here, with an exception of agents from condition 6 (high exploration, poor in food) and

9 (high exploitation, poor in food). Similarly, for CorrE agents, almost all conditions led to a success, except from condition 6. For RE agents, only condition 4 (high exploration, uniform food distribution) was unsuccessful. The ‘success’, however, is not really big, as the degrees of belief are not high. The highest probability of the successful action for CE agents is in condition 7 (high exploitation, uniform food distribution).

4.3 Temporal success

Below we present a series of queries that constitutes temporal success. The temporal aspect is captured by ‘previous action’. If the ‘previous action’ was changed, the agents could expect an edible object next.

Almost all agents learned that if the previous action was changed, then the next object most likely will be edible (Fig. 10), with color = green (Fig. 8) and shape = triangle (Fig. 9). In case of action (Fig. 7), almost none (except from CorrE agents in condition 5) of the agents learned the temporal dependency that changing the object in time t_{-1} will more likely require eating the object in time t .

P(Action|Previous Action = change)

Figure 7: The plots show probability distributions of action given that the previous action was 'change' under different conditions and in different environments. Each row of the plots represents different environment: causal, correlations and random. Each column represents different values of exploitation-exploration rate parameter: uniform rate, high exploration, high exploitation. Lastly, every single plot represents different values of the second parameter: food distribution. A combination of exploitation rate and food distribution constitutes a condition. Each row contains 1 (left)–9 (right) conditions. For example, condition 4 is a condition with high exploration rate and uniform food distribution. The bars represent degrees of beliefs of agents in which action to take given the previously taken action.

$P(\text{Color} | \text{Previous Action} = \text{change})$

Figure 8: The plots show probability distributions of a color of an object given that the previous action was 'change' under different conditions and in different environments. Each row of the plots represents different environment: causal, correlations and random. Each column represents different values of exploitation-exploration rate parameter: uniform rate, high exploration, high exploitation. Lastly, every single plot represents different values of the second parameter: food distribution. A combination of exploitation rate and food distribution constitutes a condition. Each row contains 1 (left)–9 (right) conditions. For example, condition 4 is a condition with high exploration rate and uniform food distribution. The bars represent degrees of beliefs of agents in the object's color given the previously taken action.

P(Shape|Previous Action = change)

Figure 9: The plots show probability distributions of a shape of an object given that the previous action was 'change' under different conditions and in different environments. Each row of the plots represents different environment: causal, correlations and random. Each column represents different values of exploitation-exploration rate parameter: uniform rate, high exploration, high exploitation. Lastly, every single plot represents different values of the second parameter: food distribution. A combination of exploitation rate and food distribution constitutes a condition. Each row contains 1 (left)–9 (right) conditions. For example, condition 4 is a condition with high exploration rate and uniform food distribution. The bars represent degrees of beliefs of agents in the object's shape given the previously taken action.

$P(\text{Edible} | \text{Previous Action} = \text{change})$

Figure 10: The plots show probability distributions of a edibility of an object given that the previous action was 'change' under different conditions and in different environments. Each row of the plots represents different environment: causal, correlations and random. Each column represents different values of exploitation-exploration rate parameter: uniform rate, high exploration, high exploitation. Lastly, every single plot represents different values of the second parameter: food distribution. A combination of exploitation rate and food distribution constitutes a condition. Each row contains 1 (left)–9 (right) conditions. For example, condition 4 is a condition with high exploration rate and uniform food distribution. The bars represent degrees of beliefs of agents in edibility of the object given the previously taken action.

4.4 Summary and Total Score

There was only one condition under which agents were successful in all dimensions: condition 5 for CorrE agents. Condition 5 is high in food and features higher exploration rate.

To assess the cognitive success across different conditions, we introduce a query score. Given a query, first, it is checked in which conditions agents made a correct inference, second, a scale of the conditions is created: from the conditions with the highest probability to the lowest. Then, a score is assigned depending on the place on the scale. The highest possible score is 9 (the total number of conditions), the lowest 0. For example, a condition X scored 9 means that (a) the agents learned the correct distribution in all nine conditions; (b) condition X is a condition with the highest probability from the query.

Using the query score measure (Fig. 11), we can see which of the conditions were the most successful (in terms of having the highest, correct probability distribution per query).

Figure 11: Query score per condition per environment. In case of the causal environment, condition 5 has the highest score. For correlations and random environments, condition 8 is the highest.

For CE agents, the most successful, in total, was condition 5 (rich in food, high exploration rate). According to the behavioral success measure, the most successful condition (8) was rich in food, with high exploitation rate. The most successful for ‘eating’ was condition 5, for ‘changing’ was condition 7. For temporal success, the most successful was condition 2. For CorrE agents, the most successful, in total, was condition 8. For ‘eating’, the most successful was condition 5. For ‘changing’, condition 8. For temporal success, condition 5. For RE agents, the most successful condition was also condition 8. For ‘eating’, it was condition 8, for ‘changing’, condition 3, and for temporal success, condition 2.

5 Results: Entropy success

The entropy success measure aims at calculating and comparing conditional entropy of agents' model across different conditions. It is important to distinguish entropy conditioned on an agent's model (conditional entropy) $H(X|m)$ from entropy $H(X)$ (Friston and Ao, 2011; Schwartenbeck et al., 2013). Conditional entropy is equivalent to minimizing the long-term surprise given a specific generative model. What is, and what is not, surprising for an agent is determined by prior beliefs that are built-in by evolution (Schwartenbeck et al., 2013). States that are surprising for one agent (a fish being out of the water), might be not surprising for the other agent (a cat being out of the water). This is why it is important to condition entropy on agents' generative model.

In case of our simulations, we are interested in $H(Action|Color, Shape)$, the conditional entropy of action given color and shape:

$$H(Action|Color, Shape) = - \sum_{c \in C, s \in S, a \in A} p(c, s, a) \log p(c, s, a) + \sum_{c \in C, s \in S} \left(\sum_{a \in A} p(c, s, a) \right) \times \log p(c, s)$$

First, we present the general results (Tab. 12), then we redefine the conditional entropy to make it directly comparable with other success measures by conditioning on a particular value assignment for Color and Shape (Tab. 13).

As we can see in the Table 12, for causal environment (CE) agents, the lowest conditional entropy of action given shape and color is achieved in condition 8. This condition is rich in food and features high exploitation rate. For correlations environment (CorrE) agents, it is also condition 8. For random environment (RE) agents, the condition leading to the lowest conditional entropy measure is condition 5 (rich in food, featuring high exploration rate).

Next, we are interested in calculating a modified version of conditional entropy (H^*) of 'Action', specifically $H^*(Action|green\ triangle)$ and $H^*(Action|green\ square)$. This way, we can directly map the entropy measure of $H^*(Action|green\ triangle)$ to the measures of successful (both behaviorally and cognitively) eating; and $H^*(Action|green\ square)$ to successful changing.

As we can see in the Table 13, the lowest entropy for CE agents for successful eating is in condition 9 (high exploitation, poor in food). Interestingly, this condition was among the least successful for behavioral and cognitive eating success. The conditions that were the most successful for behavioral (condition 8) and cognitive (condition 5) eating success, are among the least successful in terms of entropy eating success. For changing success, the most successful condition was condi-

tion 8 (high exploitation rate, rich in food). This condition is again different from the successful cognitively and behaviorally conditions.

For CorrE agents, the conditions for successful eating and changing are similar to CE agents. Condition 9 and 8 results in the lowest entropy for ‘green triangle’. Condition 8, for ‘green square’. The value is, however, higher for ‘green triangle’ than for ‘green square’ (it’s the highest across all environments, but still the lowest in this case). For RE agents, the most successful condition for successful eating was condition 6, for successful changing, condition 5. Condition 6 is poor in food and has high exploration rate; condition 5 is rich in food, with high exploration rate.

	condition 1	condition 2	condition 3	condition 4	condition 5	condition 6	condition 7	condition 8	condition 9
causal	0.559	0.441	0.605	0.776	0.746	0.836	0.399	0.26	0.424
correlations	0.71	0.505	0.754	0.963	0.837	0.963	0.497	0.276	0.536
random	0.835	0.733	0.881	0.364	0.214	0.502	0.59	0.447	0.628

Table 12: The Conditional entropy results $H(Action|Color, Shape)$. The lowest entropy is marked in red. For causal and correlations environment agents, the condition that yielded the lowest entropy was condition 8. For random environment agents, it was condition 5.

		condition 1	condition 2	condition 3	condition 4	condition 5	condition 6	condition 7	condition 8	condition 9
causal	green triangle	0.183	0.229	0.107	0.255	0.419	0.142	0.138	0.133	0.076
	green square	0.07	0.026	0.106	0.093	0.032	0.144	0.045	0.017	0.072
corrs	green triangle	0.283	0.279	0.212	0.393	0.49	0.257	0.19	0.1435	0.1431
	green square	0.073	0.023	0.103	0.092	0.031	0.137	0.051	0.017	0.077
random	green triangle	0.296	0.423	0.133	0.116	0.111	0.076	0.196	0.248	0.097
	green square	0.093	0.032	0.156	0.042	0.011	0.092	0.068	0.021	0.111

Table 13: The conditional entropy results of $H^*(Action|green\ triangle)$ and $H^*(Action|green\ square)$. $H^*(Action|green\ triangle)$ can be mapped to (behavioral and cognitive) eating success, $H^*(Action|green\ square)$ to changing success. The lowest entropy is marked in red.

References

- Karl Friston and Ping Ao. Free energy, value, and attractors. *Computational and mathematical methods in medicine*, 2012, 2011.
- Daphne Koller and Nir Friedman. Parameter estimation. In Daphne Koller and Nir Friedman, editors, *Probabilistic graphical models: principles and techniques*, pages 717–773. MIT press, 2009.
- Johan Kwisthout, Harold Bekkering, and Iris van Rooij. To be precise, the details don’t matter: on predictive processing, precision, and level of detail of predictions. *Brain and cognition*, 112: 84–91, 2017.
- Judea Pearl. *Causality*. Cambridge University Press, 2009.
- Philipp Schwartenbeck, Thomas FitzGerald, Raymond J Dolan, and Karl Friston. Exploration, novelty, surprise, and free energy minimization. *Frontiers in psychology*, 4, 2013.