



HAL
open science

SNN-Based Online Learning of Concepts and Action Laws in an Open World

Christel Grimaud, Dominique Longin, Andreas Herzig

► **To cite this version:**

Christel Grimaud, Dominique Longin, Andreas Herzig. SNN-Based Online Learning of Concepts and Action Laws in an Open World. 2025. hal-04779695v2

HAL Id: hal-04779695

<https://hal.science/hal-04779695v2>

Preprint submitted on 24 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

SNN-Based Online Learning of Concepts and Action Laws in an Open World

Christel Grimaud, Dominique Longin, Andreas Herzig
Université de Toulouse, CNRS, Toulouse INP, UT3, IRIT, France

Abstract

We present the architecture of a fully autonomous, bio-inspired cognitive agent built around a spiking neural network (SNN) implementing the agent’s semantic memory. This agent explores its universe and learns concepts of objects/situations and of its own actions in a one-shot manner. While object/situation concepts are unary, action concepts are triples made up of an initial situation, a motor activity, and an outcome. They embody the agent’s knowledge of its universe’s actions laws. Both kinds of concepts have different degrees of generality. To make decisions the agent queries its semantic memory for the expected outcomes of envisaged actions and chooses the action to take on the basis of these predictions. Our experiments show that the agent handles new situations by appealing to previously learned general concepts and rapidly modifies its concepts to adapt to environment changes.

1. Introduction

The ability of a cognitive agent to act adequately in a given environment depends on its ability to predict how performing a given activity might affect its current situation, that is, it depends on its knowledge of its environment’s *action laws*. How artificial agents can acquire these laws and how these should be updated if their environment changes has proven a difficult question. In the case where the intended environment is *open*—that is, where the agent’s designer cannot foresee all the situations the agent might encounter in the future—, providing a suitable set of action laws to the agent “by hand” is unfeasible. The only viable solution is that the agent continuously learns the relevant laws from experience, just as natural agents (humans and animals) do. Crucially, this learning process should allow for generalization over disparate experiences, so that the agent is able to behave appropriately in new situations. It should also accommodate environment changes.

The present paper intends to show how this could be done. Its main thrust is that natural agents’ ability to perform well in our open and changing world relies on the fact that they store their knowledge in the form of rapidly updatable concepts with various degrees of generality. Presumably, they first form concepts about the encountered

objects/situations, and then use these as elements for composing concepts of actions supporting their knowledge of their environment’s action laws. We suggest that artificial agents could do just the same, relying on some artificial neural network to learn and store concepts, and then querying it to make predictions about the outcome of envisaged actions.

To test this idea, we here build an artificial hybrid agent with a SNN at its core. We make this agent live in a very simple virtual world, composed of rooms which may be, or not, accessible (hence, knowable) to it. At first, the agent is confined to one single room and learns by itself how to act in it. Then at some point a door opens to a new room containing some never encountered before objects and situations. Yet, although these are new to the agent, some general laws are preserved from one room to the other. We show that having learned these laws in the first room allows the agent to act by and large properly in the second one, as soon as it enters it. We also show that relying on neurally implemented concepts allows the agent to rapidly update its knowledge and adapt to environment changes.

The paper is organized as follows. In Section 2 we discuss related work, and in Section 3 we present the agent and its universe. Section 4 clarifies the notions of concepts and actions laws we use, while Section 5 describes the neural network and its functioning. Section 6 describes the agent’s general functioning and Section 7 presents the results. Finally, Section 8 concludes and discusses future possible developments of the framework.

2. Related Works

Our research problem is autonomous online learning, generalization and updating of concepts and actions laws in an open universe. The intended application is reasoning and planning for autonomous robots. To our knowledge, no existing approach addresses this problem in all its dimensions, but these are investigated in separate research fields.

Continual Learning tackles the problem of lifelong knowledge acquisition Wang et al. (2024); Lesort et al. (2020). Its main challenge is to avoid catastrophic loss of previous knowledge when acquiring new knowledge; a secondary research axis is one-shot/few-shots learning, i.e., the ability to learn online from one or few examples Wang et al. (2020). However, current approaches mostly consider the learning of *tasks* (mostly image classifica-

tion/recognition tasks, but also some more complex tasks such as playing games Kirkpatrick et al. (2017)), not of concepts nor action laws. Furthermore, most of them rely on supervised learning and/or labelled training data, which is unsuitable for open world autonomous agents.

Concept Learning has mainly been studied in view of explainability Gupta and Narayanan (2024), mostly of classification models (e.g., Koh et al. (2020)) but also of decision making in the context of reinforcement learning Das, Chernova, and Kim (2023); Zabounidis et al. (2023). For this reason, many proposals are dedicated to learning a human-predefined set of concepts using some annotated data. In the field of Image Classification some approaches deal with the extraction of concepts from data Wang, Lee, and Qi (2022); Ghorbani et al. (2019); Hase et al. (2019), but then these are extracted from labelled classes of images, which is unsuitable for open world autonomous agents. The vast majority of these methods also disregard the hierarchical organization of concepts from particular to more general, and they generally do not address one-shot learning, online revision or updating of concepts.

Action Learning has been studied from various perspectives. In Dynamic Epistemic Logic, Bolander and Gierasimczuk (2018) proposed a method to learn an action model through successive observations of transitions between states. But this method does not achieve generalization nor accommodate environment changes and only considers the *universally applicable* actions (i.e., actions that can be executed in every logically possible state), a condition real-world actions rarely satisfy. In the field of Planning, Bonet, Frances, and Geffner (2019) showed how to learn abstract actions from a few carefully chosen instances of some general planning problem. But said instances come with their own set of ground actions which must be known beforehand, so this approach cannot be used in open worlds, where an agent needs to incrementally learn from experience.

Our work could be related to model-based approaches of **Reinforcement Learning** (RL) Moerland et al. (2023), but it differs from them in two important aspects. First, we are only interested in learning a model of the environment, not in learning policies. Despite their undeniable successes, RL approaches struggle to adapt to environment changes and to revise learned policies Kirk et al. (2023); Farebrother, Machado, and Bowling (2018). We believe that an agent learning and updating in real time a model of the environment and using it to make its decisions would be able to quickly adapt its behavior. Second, the learning process we propose does not depend on the existence of rewards, which makes it usable in contexts where rewards are scarce.

3. The Agent and its Universe

The agent’s universe is built over a grid of boxes, which we (not the agent) identify using an orthonormal coordinate system (see Figure 1). Each box represents a particular location in the agent’s universe and possesses a particular set of features the agent is able to perceive, drawn from the set $LF = \{OK, KO, NorthWall, EastWall, SouthWall, West-$

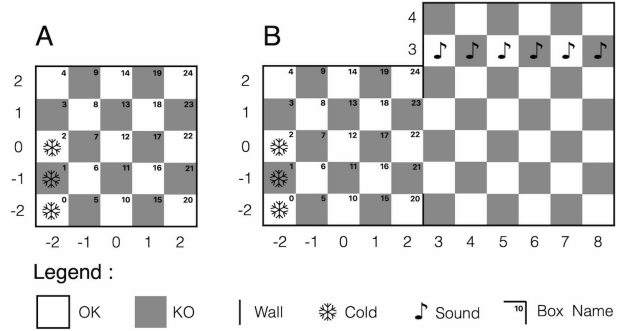


Figure 1. The agent’s accessible world. A: in the first phase, Room 1 only; B: after opening the door, Rooms 1 and 2.

Wall, Cold, Sound, #0, #1, ..., #24}. Although the agent’s universe is not finite, at any time point we only consider the boxes to which it has access, the set of which is always finite. For example, the box with coordinates $(-2, -2)$ has the feature set $LF_{(-2,-2)} = \{OK, SouthWall, WestWall, Cold, \#0\}$, while the box with coordinates $(5, 0)$ has the feature set $LF_{(5,0)} = \{OK\}$ (“#n” is to be taken as a particular name for a box hence a feature, not all boxes need to have one). Two boxes with the same feature set are indistinguishable for the agent. Boxes’ feature sets may change over time, reflecting environment changes. The rooms are made out of boxes, and delimited with impassable walls. Opening a door amounts to removing the wall features from the concerned boxes’ feature sets (as in Figure 1.B). The simplicity of this environment is meant to allow the precise monitoring and testing of the agent’s learning and decision making abilities (section 7 below), so to provide a proof of concept of our approach. Scalability and considered developments of the approach are discussed in our last section.

The agent is composed of a set of sensors, a perceptual system, a semantic memory, a decision system, a motor system and a set of actuators (see Figure 2). Sensors collect data from the external world and feed it to the perceptual system, which performs feature/object recognition. Neural networks doing this while relying on unsupervised learning from unlabelled data already exist (e.g., Thiele, Bichler, and Dupret (2018); Kheradpisheh et al. (2017)), so we simply suppose that the agent’s perceptual system operates as intended and provides the semantic memory with the appropriate inputs, namely, the correct and complete set of features of the agent’s current location. Semantic memory forms concepts by binding together sets of features, and stores them for further retrieval. Its modeling is the main




Figure 2. Schema of the agent

focus of this paper. The decision system is the other important part: it queries the semantic memory to predict the outcome of possible actions, and decides which one to take on the basis of these predictions. This decision is then sent to the motor system, which activates the actuators to perform the corresponding motor activity. Information from the actuators is sent back to semantic memory through proprioception, allowing the agent to memorize the motor-related features of the realized actions.

The agent’s possible actions consist in steps from one box to another adjacent box, in any of the eight directions. Formally, an action is a triple made up of a depart location, a motor activity, and an outcome. By “motor activity” we mean the fact that the agent’s actuators are activated so as to make it move to the immediate next box in the selected direction. The set of motor activities features the agent is able to perceive by proprioception is the set $MF = \{N, NE, E, SE, S, SW, W, NW, Diag., Orth.\}$, where the first eight are specific to each particular direction, while *Diag.* and *Orth.* are shared by all motor activities yielding diagonal/orthogonal moves. In cases where there is a wall at the edge of the depart box in the selected direction, the agent bumps into it and remains at the same place. We then say that the action’s outcome is a failure. Otherwise, the action’s outcome is the agent’s new location.

As for the locations’ features, we suppose that *KO* corresponds to some unpleasant stimulus the agent spontaneously wants to avoid, and *OK* to the absence of such a stimulus, while the others convey some indifferent information.

4. Concepts and Action Laws

The agent can form two kinds of concepts. First, concepts of “things”, in the broad sense. These bind together co-occurrent features, and can be seen as some sort of conjunction in which conjuncts have different “weights”, reflecting the fact that some features are more important than others in a concept’s definition Freund (2008). They are used to store knowledge about locations and more generally any object, so we call them *object concepts*. The second kind is relational concepts. These take other concepts as elements, and bind them together into tuples. Concepts of actions are of this kind: they bind together the agent’s concepts of a depart location, a performed motor activity, and a subsequent outcome, in the order in which they were experienced.

We say that an object concept *X* is *general*, as opposed to *particular*, if there is another concept *Y* such that the set of features composing *X* is a strict subset of the set of features composing *Y*. *Y* is then said to be *more particular* than *X*. We say that an action concept is *general* if the object concept of its initial situation is general or its motor activity component only contains *Diag.* or *Orth.*. We understand the generality of concepts relative to the set of concepts the agent possesses at some point, so no concept is general or particular in itself.

For example, when visiting the box (0, 0) the agent may form the particular object concept $[OK, \#12]$, which is a memory of an OK place with name #12, and only applies to this particular box in its accessible universe. If it then

moves North-East and arrives at box (1, 1), it can form the particular object concept $[OK, \#18]$, and also the particular action concept $[[OK, \#12], [NE, Diag], [OK, \#18]]$ which corresponds to the memory of being in an OK place with name #12 and then moving North-East to arrive at another OK place with name #18. Yet, after visiting a number of locations having the feature *OK* in common, the agent may also form the general object concept $[OK]$. Furthermore, it is a general rule in its accessible universe that moving North-East from an OK location always leads to another OK location, except for when there is a wall at the North or East edge of the depart box. Therefore, after having experienced a number of North-East moves from various OK locations, the agent may form general action concepts such as $[[OK], [NE, Diag], [OK]]$ and $[[OK, NorthWall], [NE, Diag], [Failure]]$. Such general concepts capture the general (non-monotonic) action laws of the agent’s universe, and are the ones it shall rely on to behave in never encountered situations.

5. Implementing the Agent’s Semantic Memory in the Neural Network

Spiking Neural Networks (SNNs) are well suited for autonomous learning in open universes, as they allow for Spike Time Dependent Plasticity (STDP), a family of biologically plausible learning rules which can achieve unsupervised online learning from unlabelled data Thiele, Bichler, and Dupret (2018). They are also known for being energy-efficient, which is interesting for autonomous robots.

We take inspiration in the JAST learning rule Thorpe et al. (2019); Thorpe (2023), which is a simplified version of STDP where the sum of the afferent connections weights on any given neuron remains constant through learning. However, contrary to JAST we do not use binary weights but natural numbers, and we do not freeze neurons after learning, so as to allow updating.

The Network’s Architecture

The network is composed of an interface, which communicates with the agent’s other components, and a body of hidden neurons which is itself divided into two layers (see Figure 3). The first layer learns object concepts and the second learns action concepts. For this reason we call their neurons, respectively, *object concept neurons* (O-neurons for short) and *action concept neurons* (A-neurons). This architecture draws on neuroanatomical studies according to which concepts are represented in the brain by hierarchically organized *concept neurons*, each receiving information from some lower neurons and sending reciprocal connections to these same neurons so that it can reactivate them for information retrieval Quiroga (2012); Bausch et al. (2021); Shimamura (2010). For simplicity we do not model these reciprocal connections as such, but instead we allow for information to flow in both directions along the same connections: from interface neurons to O-neurons and then to A-neurons for learning and querying, and the other way round for retrieving information. A key point is




Figure 3. Schema of the SNN. O-neuron #1 encodes the concept [#0, Cold, SouthWall, WestWall, OK], and A-neuron #1 encodes the concept [[#0, Cold, SouthWall, WestWall, OK], [E, Orth], [#5, SouthWall, KO]].

that interface neurons are both input and output neurons, depending on the computational phase.

Interface neurons (I-neurons for short) mainly support the representation of features, be it of the visited locations or of the agent’s own motor activities. An additional neuron acts as a failure detector, specifically firing when the agent bumps into a wall and remains at the same place. All of them have their labels fixed from the start.

The first layer of hidden neurons is composed of 100 integrate and fire neurons, with a differentiated dynamics depending on whether their input source is I-neurons or A-neurons. O-neurons learn co-occurrences of perceived features.

The second layer is composed of 400 compartment neurons with three separate input compartments. The first one receives connections from O-neurons, the second one from motor activities I-neurons, and the third one from O-neurons and Failure neuron. Inputs received at each compartment are unable to trigger a spike by themselves, but the first and second compartments make their next compartment ready to receive and transmit inputs for a certain amount of time. In this manner, inputs can only be efficient if they occur in the correct order, so that A-neurons encode sequences of inputs. The use of compartment neurons to learn sequences was suggested in Cui, Ahmad, and Hawkins (2016); Hawkins and Subutai (2016).

The Neural Network’s functioning

We present here the main ideas governing the network’s functioning. Details of their neural implementation are given in Supplementary Material. The network’s digital implementation can be found in the provided code.

O-neurons’ Learning. Each time the agent observes its current location, information from the perceptual system is sent to the network’s interface, inducing the firing of the I-neurons encoding the location’s features. This in turn triggers the firing of a number of O-neurons. If this number reaches some fixed target number, the network directly proceeds to make them learn. Otherwise it looks for additional O-neurons by “boosting” their input. In practice, boosting consists in multiplying the input received by each non-firing O-neuron o by a factor b_o , which is an increasing function of the number of steps performed since o ’s last spike. This favors the firing of neurons that have been inactive for a long time, which can then be re-allocated to a new concept through learning.

The learning process depends on the accuracy of the agent’s knowledge about its current location. To assess it, the (pre-boosting) firing O-neurons send a backward input to I-neurons, and the resulting set of firing I-neurons (the retrieved information) is compared with the initial input (the current observation). If all the observed features can be retrieved from active O-neurons, then all inactive I-synapses (if any) on the learning O-neurons are deleted and replaced with synapses from input neurons. This procedure tends to reinforce O-neurons’ connections with I-neurons encoding well shared features at the expense of connections with those encoding more specific features, fostering the learning of general concepts. If, on the contrary, not all the observed features can be retrieved from the firing O-neurons, then we pick one learning O-neuron with maximal number of steps since its last learning and replace all its synapses, active or inactive, with synapses from input neurons. This neuron thus learns the particular situation with all its features. For the other learning neurons, the learning process is similar as in first case.

Querying the network. To query the neural network, the decision system first sends an input to the I-neurons that encode the features of the initial situation. Their firing brings a number of O-neurons to fire, sending an input to A-neurons’ first compartment. Given an envisaged motor activity m , the decision system then sends an input to the I-neurons that encode m ’s features. These fire, and send an input to A-neurons’ second compartment. A-neurons’ spiking threshold is then gradually lowered until a target number of them fire, sending a backward input to O-neurons through their third compartment’s connections. The O-neurons that fire in response to that input in turn send a backward input to I-neurons, the firing of which is the network’s response to the query. The value of A-neurons’ spiking threshold at the moment a given I-neuron spikes determines the agent’s confidence in the feature’s prediction: the higher its value, the higher the confidence. So, formally, the querying process returns a set $P_m = \{(f_1, c_1), \dots, (f_n, c_n)\}$, where f_i is a feature and c_i the degree of confidence the agent has in its prediction.

Now, among the A-neurons that receive a significant amount of input, some encode more particular action con-

cepts than others. Obviously, we want the neurons supporting more particular concepts to take precedence over those supporting more general concepts, as they convey more adequate information. To achieve this, we first assess the generality of the object concepts (/motor activity features) encoded by O- (/ I-) neurons and then we modulate the input they send to A-neurons by applying to it a multiplying factor that depends on that assessed generality (the more particular the input concept/feature, the greater the multiplying factor). Generality is assessed by checking the weight sum of O- (/ I-) neurons' connections to A-neurons, as neurons with fewer outward connections tend to encode more particular concepts.

A-neurons' Learning. At each step, O-neurons responding to the depart location and I-neurons responding to the performed motor activity send inputs to, respectively, A-neurons' first and second compartments. A-neurons reaching a certain threshold are selected for learning. If their number reaches some fixed target number, the network directly proceeds to make them learn, otherwise it looks for additional A-neurons by boosting their input in a way similar to the one used for O-neurons. The learning process depends on the accuracy of the agent's predictions relative to the action's outcome. If these are correct (i.e., all the expected features are actually present) and complete (i.e., all the actually present features were expected), then the learning process promotes generalization by erasing and replacing inactive synapses on all three compartments, otherwise it allows a few neurons to specialize on the input by also erasing and replacing some active synapses (see Supplementary Material for details).

6. Functioning of the Agent

Suppose the agent is at some depart location and observes it: information from its perceptual system triggers the firing of the I-neurons encoding the location's features, which brings a number of O-neurons to fire and send inputs to A-neurons' first compartments. Concomitantly, the information is transmitted to the decision system, which decides to make a step.

The agent's choice of a motor activity depends on whether it wants to exploit its current knowledge about its environment, or to explore it to improve its knowledge. The exploration/exploitation dilemma is a well-known problem in online learning Watkins (1989); Sutton and Barto (2018), and changing environments make it even more difficult. We therefore do not try to reach an optimal solution here, but instead we simply make the agent's decision system choose at random with equal probability between an *Exploration* and an *Exploitation* mode.

This choice being made, for each motor activity m out of the eight possible the decision system queries the semantic memory for the outcome of the action having the current location for initial situation and m for motor activity. It then rates each of them for its suitability, by building the sets S (for "*Suitable*"), US ("*Unsuitable*") and UD ("*Undecided*"):

- $S = \{(m, c) \mid (OK, c) \in P_m\}$
- $US = \{(m, c) \mid (KO, c) \in P_m \text{ or } (Failure, c) \in P_m\}$

- $UD = \{m \mid \nexists c \text{ s.t. } (m, c) \in S \cup US\}$

The decision system then chooses a motor activity depending on the selected mode. In *Exploration* mode, the agent is willing to take risks and chooses an action with the most uncertain outcome possible: if $UD \neq \emptyset$ it picks one from it, otherwise it goes for one with the least c in $S \cup US$. In *Exploitation* mode by contrast, the agent just wants to avoid KO boxes and failure as much as possible. So, if $S \neq \emptyset$ it chooses one with the greatest c , otherwise if $UD \neq \emptyset$ it picks one from it, and if both S and UD are empty it chooses one with the least c in US .

The decision system then transmits its decision to the motor system to perform the selected motor activity. The I-neurons that encode its features are activated by proprioception and send an input to A-neurons' second compartment.

We simulate the agent's move by computing its arrival location and its features. If the agent bumps into a wall, the *Failure* neuron fires and sends an input to A-neurons' third compartment, and the agent directly learns the action (see A-neurons' learning above). Otherwise, it first learns relevant object concepts about its new location (see O-neurons' learning above), and then the action. After that, the agent is ready for the next step.

7. Results

To test the agent's learning abilities, we placed it at location (0, 0) and prompted it to perform a succession of series of steps, each complete sequence of series of steps being called a *trial*. The results we present here are averaged over 50 trials.

In a first group of tests, trials consisted in 65536 steps in total with the door kept closed all along, so that the agent had no access to the second room. First, we tested the agent's ability to learn an action over one single experience (one-shot learning). To do so, after each step we asked the agent to redo the prediction that led to the just realized action, and compared this new prediction with the action's actual outcome. 96% of these post-learning predictions were both correct and complete (i.e., the predicted features were exactly those of the arrival location; see Table 5 in Supplementary Material).

To test whether the acquired knowledge was retained in the long run, after each series of steps we froze the simulation, deactivated learning and placed the agent successively in each location of each room. There, we asked it for its predictions for each of the eight possible motor activities and compared its answers with the actions' actual outcomes. Tables 1 and 2 show each feature's mean *Hit Rate* (that is, its chances of being predicted when effectively present), and *Correctness* (its chances of being effectively present when predicted)¹ for each room. For lack of space we only show the results for some series. Values for the first room (white lines) show that learned actions are indeed recalled long after having been performed. Values

¹*Hit Rate* is also known as *True Positive Rate*, *Recall* or *Sensitivity*, while *Correctness* is also known as *Precision* or *Positive Predictive Value* – see for example Kohavi and Provost (1998) for definitions. Here we multiplied the obtained figures by 100 to get percentages.

Nb of Steps	Room	OK	KO	Fail.	Wall	Cold	Sound	Box Name
1	1	8.6	18.1	0.0	0.0	0.0	0.0	1.4
	2	8.7	17.6	0.0	0.0	0.0	0.0	0.0
8	1	35.1	44.4	26.0	13.9	8.9	0.0	4.2
	2	30.3	42.0	25.5	11.1	0.0	0.0	0.0
64	1	78.3	77.8	45.4	35.0	31.2	0.0	22.0
	2	64.7	69.6	39.3	21.5	0.0	0.0	0.0
512	1	94.9	93.2	69.8	64.2	57.7	0.0	50.1
	2	86.9	88.2	40.0	24.7	0.0	0.0	0.0
8192	1	96.9	95.1	84.3	80.2	75.8	0.0	73.5
	2	91.6	86.3	56.2	26.8	0.0	0.0	0.0
32768	1	97.3	93.8	85.0	79.7	79.5	0.0	74.7
	2	92.7	85.1	56.0	26.4	0.0	0.0	0.0

Table 1. Predictions’ *Hit Rates* after n steps.

Nb of Steps	Room	OK	KO	Fail.	Wall	Cold	Sound	Box Name
1	1	13.3	23.7	0.0	0.0	0.0	0.0	4.0
	2	14.5	25.3	0.0	0.0	0.0	0.0	0.0
8	1	49.3	42.4	28.8	19.6	6.2	0.0	5.0
	2	49.0	42.3	21.6	13.5	0.0	0.0	0.0
64	1	75.2	74.2	55.9	35.2	33.9	0.0	22.6
	2	72.1	72.5	36.5	19.1	0.0	0.0	0.0
512	1	90.3	86.6	87.0	70.9	78.8	0.0	66.6
	2	83.9	83.3	55.3	28.9	0.0	0.0	0.0
8192	1	93.4	92.6	95.2	88.7	96.1	0.0	84.2
	2	86.3	87.6	66.6	37.1	0.0	0.0	0.0
32768	1	93.1	93.8	94.0	89.9	92.3	0.0	85.4
	2	86.5	87.4	69.8	37.3	0.0	0.0	0.0

Table 2. Predictions’ *Correctness* after n steps.

for the second room (grey lines) show that despite never having been in this room (since we kept the door closed) the agent was able to correctly predict *OK* and *KO* features and to a lesser extent *Failure* —and this, even though locations from the second room have different sets of features, including for some of them a new feature, *Sound*. The poor performance at wall prediction is due to the lack of general rules of the universe observable in the first room regarding the presence of walls in adjacent boxes. The agent thus relies on particular concepts of individual boxes to predict them in the first room, which cannot be used for the second room. The mixed result at failure prediction comes from a competition between general action concepts, the control of which needs to be improved.

To test the agent’s ability to use its knowledge to make appropriate decisions, we recorded the outcome of each action taken in *Exploitation* mode during the trials. Figure 4.A shows the percentages of *OK*, *KO* and *Failure* outcomes thus obtained for each series of steps. We kept track of the visited locations to check that the agent was not looping indefinitely on the same boxes: all boxes kept being visited, be it very rarely, at any point of the trials, due to the *Exploration* mode.

Finally, to test whether the agent would be able to use the knowledge acquired in the first room to act judiciously in the second room, at the end of each series of steps we asked it to chose a move for each of the second room’s type of locations (we say that two locations are of the same type if they have exactly the same features). Figure 4.B shows the percentages of *OK*, *KO* and *Failure* obtained in this manner. These results reflect the agent’s performance

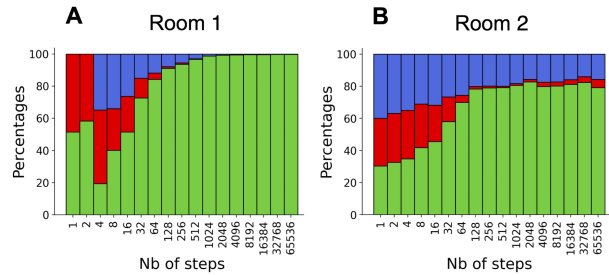


Figure 4. Actions’ mean outcomes with door closed; Green: OK, Red: KO, Blue: Failure, Grey: no data

at making predictions about the second room’s locations: it successfully predicts *OK* and *KO* boxes, but has more difficulties predicting failure.

In a second group of tests, we kept the same setup but opened the door at the 2048th step. The agent spontaneously went in the second room, and spent a variable but significant amount of time in it (36.0% of steps on average, standard deviation = 9.3). Tables 3 and 4 show the features’ *Hit Rates* and *Correctness* from the moment the door was opened. These results show that the agent was able to learn new concepts involving the *Sound* feature. The rather low *Hit Rates* for the feature are due to the lack of observable cues in boxes at the direct south of boxes with sound, which prevents the agent from being able to predict it in 40% of the cases.

It should be noted that this learning did not lead to significant loss of previous knowledge regarding the first room: a moderate drop in *Hit Rates* can be observed for boxes’ names, *Cold* and walls, but not for *OK*, *KO* and *Failure*. Accordingly, the agent’s ability to make appropriate decisions in the first room was not impacted by the door being opened. In fact, the bar chart of *OK*, *KO* and *Failure* outcomes obtained in this second run of tests showed no visible difference with the one obtained with the door closed shown Figure 4.A. This is due to the use of selective forgetting in the learning process. Forgetting is inevitable for an agent with finite memory living in an unbounded world, but the boosting of inputs used in O- and A- neurons’ learning processes (see 5 above) allows the agent to preferentially forget its less used knowledge such as specific details and/or memories that have not been used for long, preserving general concepts and often used memories. How much is forgotten depends on how much needs to be learned (the amount of novelty), and on the initial amount of redundancy in the SNN (how many neurons encode a same concept).

To better assess the agent’s ability to update its knowledge to accommodate world changes, we conducted a third experiment. As before, we kept the agent in the first room for 2048 steps before opening the door, but then we had it run 50 series of 100 steps. Half way through these, we changed sound’s location by removing the *Sound* feature from the boxes (3, 3) - (8, 3) while adding it to the boxes (3, -1) - (8, -1). At the end of each series of steps, we asked the agent to make predictions about the outcome of two actions having *Sound* as the only available information about the depart location: the first one when considering a north move, and the second, when considering a south

Nb of Steps	Room	Hit Rates							Box Name
		OK	KO	Fail.	Wall	Cold	Sound		
2048	1	95.7	94.6	78.3	75.4	69.5	0.0	64.8	
	2	88.5	87.3	46.8	25.5	0.0	0.0	0.0	
4096	1	96.6	96.2	78.2	63.8	56.2	0.0	53.4	
	2	91.3	93.7	56.3	34.4	0.0	33.9	4.0	
8192	1	96.9	95.6	81.0	67.0	60.9	0.0	58.4	
	2	93.5	91.6	58.4	32.0	0.0	28.4	4.7	
16384	1	97.9	96.0	84.3	67.5	63.5	0.0	59.6	
	2	93.6	90.4	70.0	35.5	0.0	32.3	5.1	
32768	1	97.8	97.3	88.0	70.4	69.4	0.0	60.3	
	2	95.2	92.1	74.4	37.3	0.0	33.3	4.0	
65536	1	97.8	96.5	88.9	71.1	68.0	0.0	61.7	
	2	94.3	92.4	77.0	37.0	0.0	31.0	5.8	

Table 3. Predictions’ Hit Rates after n steps.

Nb of Steps	Room	Correctness							Box Name
		OK	KO	Fail.	Wall	Cold	Sound		
2048	1	92.7	89.6	91.8	82.4	90.5	0.0	79.0	
	2	85.0	84.6	58.7	36.1	0.0	0.0	0.0	
4096	1	94.9	91.8	90.7	74.6	94.1	0.0	88.7	
	2	90.7	87.6	72.2	45.2	0.0	61.5	0.9	
8192	1	95.9	92.5	90.3	80.4	95.9	0.0	90.9	
	2	91.1	87.9	73.5	45.5	0.0	66.5	1.0	
16384	1	95.9	94.1	92.7	80.1	96.4	0.0	92.0	
	2	92.1	90.3	76.1	47.1	0.0	60.5	1.9	
32768	1	96.5	95.7	94.9	82.4	99.5	0.0	94.9	
	2	92.6	91.8	81.4	51.0	0.0	57.7	1.3	
65536	1	96.9	95.5	94.5	83.8	98.2	0.0	93.4	
	2	93.7	92.2	82.3	50.3	0.0	71.4	1.7	

Table 4. Predictions’ Correctness after n steps.

move. Once the door was opened, the agent first learned to predict a north wall in the arrival box when considering a north move, but when the *Sound* feature was moved it stopped to do so and learned to predict a south wall in the arrival box when considering a south move instead. Figure 5 shows that learning and updating were both fast and robust, considering that learning occasions are very scarce (less than 0.8% of taken steps on average for each rule).

As regards computing time, it takes about 12 seconds on a conventional computer for the agent to run 2048 steps while keeping track of all the test data. We made no attempt to optimize the computing time, as it seems less critical in the case of online learning of autonomous agents, which can learn while physically performing their actions.

8. Conclusion and Future Developments

In this paper we have given a proof of concept of the architecture of a fully autonomous agent that learns action laws online and accommodates environment changes. This agent relies on general concepts to handle new situations and dynamically adjusts its concepts to its current environment. This makes it well suited for open worlds: if a new door were to open to a third room with new objects and laws, it would learn them just as it did in the second one. Of course, this would come at the cost of the forgetting of its least used concepts, but these are precisely the ones it needs the less. In fact, the agent’s ability to selectively forget its less used knowledge ensures that it will always be able to adapt to new environments by replacing old unused




Figure 5. Prediction of walls in arrival boxes given *Sound* as information about the depart location. Black dotted line: door opens; Red dotted line: sound changes location; Blue: *North-Wall* predictions considering a north move; Orange: *SouthWall* predictions considering a south move (average over 50 trials).

concepts by new useful ones.

As regard scalability to larger environments, the number of neurons needed in each layer of the network (interface, O-neurons, A-neurons) grows linearly with the number of items (features, object concepts, action concepts) to encode; the number of synapses needed on each O-neuron grows linearly with the maximum number of features to be remembered for a given object/situation. However the agent’s ability to form general concepts out of individual experiences and to use them to make decisions makes it unnecessary to encode each individual object with all its features and each performed action, so these are upper bounds.

Further work remains to be done to allow the agent to live in more complex and realistic environments. A first direction of work would be to endow the agent with planning abilities. To do so, one would need a notion of *applicable* action law. It seems that an action law represented by an action concept $[\mathbf{x}, \mathbf{y}, \mathbf{z}]$ could be deemed applicable in a situation s if s satisfies all the features in \mathbf{x} and $\mathbf{z} \neq \text{Failure}$. To live in open worlds the agent would need to build its own set of possible situations (states) online. It appears that the set of its object concepts could be used for this. The decision system should be augmented so as to handle goals and a cost function.

Another beneficial improvement would be to make the agent able to use incomplete information as input for learning and querying, as real world agents’ observations are rarely complete. We could then make the agent able to query its semantic memory for object properties given some partial input. It would also be useful to implement negation in the network, to allow the agent to represent the fact that a given object does not have a given feature. Neural inhibition could probably be used for this, but the appropriate learning rules remain to be found. It seems to us that taken together these two improvements would bring the agent to draw non-monotonic inferences in the spirit of Grimaud (2016).

It would also be useful to allow the agent to distinguish between objects and their locations, as actions can modify one, the other or both. Biological brains process the “what” and the “where” components of observations in two separate pathways before reunifying them, and this could be an inspiration source.

A last line of research would be to investigate how an agent should decide between *Exploration* and *Exploitation* modes in an open world.

References

- Bausch, M.; Niediek, J.; Reber, T. P.; Mackay, S.; Boström, J.; Elger, C. E.; and Mormann, F. 2021. Concept neurons in the human medial temporal lobe flexibly represent abstract relations between concepts. *Nature communications*, 12(1): 6164.
- Bolander, T.; and Gierasimczuk, N. 2018. Learning to act: qualitative learning of deterministic action models. *Journal of Logic and Computation*, 28(2): 337–365.
- Bonet, B.; Frances, G.; and Geffner, H. 2019. Learning features and abstract actions for computing generalized plans. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2703–2710.
- Cui, Y.; Ahmad, S.; and Hawkins, J. 2016. Continuous online sequence learning with an unsupervised neural network model. *Neural computation*, 28(11): 2474–2504.
- Das, D.; Chernova, S.; and Kim, B. 2023. State2explanation: Concept-based explanations to benefit agent learning and user understanding. *Advances in Neural Information Processing Systems*, 36: 67156–67182.
- Farebrother, J.; Machado, M. C.; and Bowling, M. 2018. Generalization and regularization in DQN. *arXiv preprint arXiv:1810.00123*.
- Freund, M. 2008. On the notion of concept I. *Artificial Intelligence*, 152(1): 105–137.
- Ghorbani, A.; Wexler, J.; Zou, J. Y.; and Kim, B. 2019. Towards automatic concept-based explanations. *Advances in neural information processing systems*, 32.
- Grimaud, C. 2016. Modelling reasoning processes in natural agents: a partial-worlds-based logical framework for elemental non-monotonic inferences and learning. *Journal of Applied Non-Classical Logics*, 26(4): 251–285.
- Gupta, A.; and Narayanan, P. 2024. A survey on Concept-based Approaches For Model Improvement. *arXiv preprint arXiv:2403.14566*.
- Hase, P.; Chen, C.; Li, O.; and Rudin, C. 2019. Interpretable image recognition with hierarchical prototypes. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, 32–40.
- Hawkins, J.; and Subutai, A. 2016. Why Neurons Have Thousands of Synapses, a Theory of Sequence Memory in Neocortex. *Frontiers in Neural Circuits*, 10.
- Kheradpisheh, S. R.; Ganjtabesh, M.; Thorpe, S. J.; and Masquelier, T. 2017. STDP-based spiking deep convolutional neural networks for object recognition. *Neural Networks*, 99: 56–67.
- Kirk, R.; Zhang, A.; Grefenstette, E.; and Rocktäschel, T. 2023. A survey of zero-shot generalisation in deep reinforcement learning. *Journal of Artificial Intelligence Research*, 76: 201–264.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Koh, P. W.; Nguyen, T.; Tang, Y. S.; Mussmann, S.; Pierson, E.; Kim, B.; and Liang, P. 2020. Concept bottleneck models. In *International conference on machine learning*, 5338–5348. PMLR.
- Kohavi, R.; and Provost, F. 1998. Glossary of Terms.
- Lesort, T.; Lomonaco, V.; Stoian, A.; Maltoni, D.; Filliat, D.; and Díaz-Rodríguez, N. 2020. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information fusion*, 58: 52–68.
- Moerland, T. M.; Broekens, J.; Plaat, A.; Jonker, C. M.; et al. 2023. Model-based reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 16(1): 1–118.
- Quiroga, R. Q. 2012. Concept cells: the building blocks of declarative memory functions. *Nature Reviews Neuroscience*, 13: 587–597.
- Shimamura, A. P. 2010. Hierarchical relational binding in the medial temporal lobe: the strong get stronger. *Hippocampus*, 20(11): 1206–1216.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Thiele, J.; Bichler, O.; and Dupret, A. 2018. Event-based, timescale invariant unsupervised online deep learning with STDP. *Front. Comput. Neurosci.* 12, 46 (2018).
- Thorpe, S. 2023. Timing, Spikes, and the Brain. In *Time and Science: Volume 2: Life Sciences*, 207–236. World Scientific Publishing Europe Ltd.
- Thorpe, S.; Masquelier, T.; Martin, J.; Yousefzadeh, A. R.; and Linares-Barranco, B. 2019. Method, digital electronic circuit and system for unsupervised detection of repeating patterns in a series of events. Patent US20190286944A1.
- Wang, A.; Lee, W.-N.; and Qi, X. 2022. Hint: Hierarchical neuron concept explainer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10254–10264.
- Wang, L.; Zhang, X.; Su, H.; and Zhu, J. 2024. A comprehensive survey of continual learning: theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wang, Y.; Yao, Q.; Kwok, J. T.; and Ni, L. M. 2020. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3): 1–34.
- Watkins, C. J. C. H. 1989. Learning from delayed rewards.

Zabounidis, R.; Campbell, J.; Stepputtis, S.; Hughes, D.; and Sycara, K. P. 2023. Concept learning for interpretable multi-agent reinforcement learning. In *Conference on Robot Learning*, 1828–1837. PMLR.

Supplementary Material

Post Learning Predictions

The table below is referred to in section 7, first group of tests. Results are averaged over 50 trials.

We call a prediction *Correct and Complete* (CC) if the predicted features are exactly those of the arrival location. The table’s first line shows the mean percentage of steps leading to a CC post-learning prediction, for each series of steps. The second line (MF for “*Missed Features*”) shows the average percentage of features occurrences that the agent failed to predict after learning. The third line (PE for “*Predictions Errors*”) shows the average percentage of wrongly predicted features occurrences.

Nb of Steps										
	1	2	4	8	16	32	64	128	256	
CC	100.0	100.0	99.0	99.5	99.8	98.0	96.4	92.8	89.2	
MF	0.0	0.0	0.3	0.2	0.1	0.9	2.0	3.9	5.6	
PE	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3	

Nb of Steps									Global
	512	1024	2048	4096	8192	16384	32768	65536	
CC	90.7	92.8	94.7	95.3	95.5	96.1	96.7	96.5	96.3
MF	4.5	3.4	2.5	2.2	2.1	1.8	1.6	1.6	1.7
PE	0.4	0.2	0.2	0.2	0.1	0.1	0.1	0.1	0.1

Table 5. Mean post-learning predictions percentages over n trials. CC: Correct and Complete, MF: Missing Features, PE: Prediction Errors.

These results show a good performance at immediate recall after learning. The lower CC prediction scores around 128-2048 steps can be explained as follows. At start up, the network has a lot of unused neurons at its disposal and simply learns the particular concepts with all their features, so specific features are well learned. But as the network accumulates knowledge, unused neurons get scarce and the network recruits less used neurons for learning. A trade-off is made between learning the current observation’s specific features and not forgetting the common features already encoded by the neurons, which makes specific features more difficult to learn. Yet after repeated encounters the specific features are finally learned, so the score improves again.

The Network’s implementation

We provide here a more detailed account of the network’s implementation. All parameters can be found in the supplied code.

We respectively note L , M , and F the sets of interface neurons that encode location’s features, motor activity features, and failure. $I = L \cup M \cup F$ is the set of interface neurons.

We note O the set of O-neurons, and A the set of A-neurons. For convenience we call “*L-neurons*” the neurons from L , and similarly for the other above sets. For any $l \in L$ and $o \in O$, $w_{l,o}$ denotes the weight of the connection between l and o . Similarly, for $n \in \{1, 2, 3\}$,

$x \in O \cup M \cup F$, and $a \in A$, $w_{x,a}^n$ denotes the weight of x ’s connection on a ’s n^{th} compartment. All weights are positive integers and are initialized at random, in a way such that the sum of weights on any O-neuron is equal to some fixed integer WS_O , and the sum of weights on any A-neuron’s n^{th} compartment is equal to some fixed integer WS_A . We store these connexions in separate arrays $cnxLO$, $cnxOA1$, $cnxMA2$ and $cnxOA3$. $cnxLO$ reads “*connexions from L to O*”, and similarly for the others. The agent’s learning consists in the update of these arrays.

For each O-neuron we set at random a fixed spike threshold ST_o comprised between some minimal and maximal possible values. L-neurons all have the same spike threshold, and the *Failure* neuron has its own spike threshold. We do not specify spike thresholds for M neurons, as we shall not need them. Spike thresholds for A-neurons are dynamically adjusted, in a way we shall see below.

O-Neurons

At any point of time, we note $inputLO$ the set of L-neurons that fire in response to an input from either the perceptual system or the decision system, and send an input to O-neurons. Inputs from the perceptual system correspond to the fact that the agent is observing its current location, while inputs from the decision system correspond to the fact that the agent is querying its semantic memory.

Firing. The set of O-neurons that fire in response to an input $inputLO$ is noted $firedO$. To compute it, for any neuron o in O we first calculate the total input $inputSumO_o = \sum_{l \in inputLO} w_{l,o}$ received by o , and then the difference $overT_o$ with o ’s spike threshold (that is, $overT_o = inputSumO_o - ST_o$). Starting with $firedO = \emptyset$, we then successively add to it the O-neurons with the highest positive $overT_o$ value, until either no O-neuron with positive $overT_o$ remains or the number $|firedO|$ of firing O-neurons reaches some target number $TnbFiredO$ (so that we keep the number of firing neurons under control).

For any neuron o in O , we note $lastSpikedO_o$ the number of steps performed by the agent since the last time o underwent learning. These values are initialized at random using some upper bound, and updated each time the neuron learns or the agent makes a step.

Learning. Learning of object concepts occurs whenever the agent observes its current situation. To compose the set $learningO$ of learning O-neurons, we first set $learningO = firedO$. If $|learningO|$ reaches some target number $TnbQueryO$, we directly proceed to the neurons’ learning ($TnbQueryO < TnbFiredO$). Otherwise, we look for some more neurons to make learn, by boosting the input received by O-neurons. In practice, for any $o \in O - firedO$ we define $boostedInputSumO_o = inputSumO_o * (lastSpikedO_o + b)/b$, with $b \in \mathbb{N}$ a fixed parameter. We then pick the neurons with the highest boosted input sum values and successively add them to $learningO$ until $|learningO| \geq TnbQueryO$. This method allows to recruit O-neurons that did not learn for a long time and ensures that all neurons are used by the network in the long run.

The learning of O-neurons essentially consists in

replacing their afferent connections from non-firing L-neurons with afferent connections from firing L-neurons (namely, those from $inputLO$). To decide which neuron from $inputLO$ will establish new connections with a given O-neuron, we define a probability distribution $probaNewSynapsesLO$ on $inputLO$. For any neuron l in $inputLO$, this probability is a decreasing function of $\sum_{o \in O} w_{l,o}$, accounting for the fact that neurons cannot indefinitely grow new forward connections, so a neuron having already a lot of them should be less prone to establish new ones than a neuron having fewer of them.

The learning process depends on the accuracy of the agent’s knowledge relative to the current situation. To assess it, we use the set $firedO$ as a backward input to L-neurons, and compare the resulting set $firedL$ of firing L-neurons with the initial input $inputLO$. This corresponds to the agent comparing what it sees (represented by $inputLO$) with what it expected based on its current knowledge (represented by $firedL$). If $inputLO \subseteq firedL$, that is, if all the observed features can be retrieved from the firing O-neurons, then for any neurons o in $learningO$ and l in $L - inputLO$ the connection $w_{l,o}$ is reset to 0. In other words, we delete all O-neurons’ inactive synapses. We then replace them with as many new synapses from neurons from $inputLO$, following the above mentioned probability distribution. This procedure tends to reinforce O-neurons’ connections with interface neurons encoding well shared features, at the expense of connexions with interface neurons encoding more specific features, and is the driving force for the formation of general object concepts.

If, on the contrary, $inputLO \not\subseteq firedL$, we first check if $|learningO|$ reaches $TnbFiredO$. If not, we look for one more neuron to add to $learningO$, using the same method as previously except that only one neuron is added. The learning process is then similar to the one above, except that we pick one neuron with maximal $lastSpikedO_o$ value and replace all its synapses, whether active or inactive. This neuron thus evades the general trend towards generalization and learns the particular situation instead, with all this features. Note that in both learning procedures the sum of afferent connections weights $\sum_{l \in L} w_{l,o}$ remains unchanged for all neurons o , in line with the JAST learning rule’s spirit.

A-Neurons

At any point of time, we respectively note $inputOA1$, $inputMA2$ and $inputOA3$ the sets of neurons from O and M and $O \cup F$ that send inputs to A-neurons’ first, second and third compartments. We note $firedA$ the set of firing A-neurons. Inputs on the first compartment encode an initial situation, inputs on the second compartment encode motor activity features, and inputs on the third compartment encode an action’s outcome. These inputs may occur in two different situations: when the decision system queries the neural network for the expected outcome of an envisaged action, and when the neural network learns about a given action.

Querying the network. To query the neural network, the decision system sends an input to interface neurons so as to trigger the firing of the L-neurons that encode the features of the initial situation. These firing L-neurons are an

input set $inputLO$ for O-neurons, and bring some of them to fire. The set $firedO$ of spiking O-neurons is an input set $inputOA1$ to A-neurons’ first compartment, but it does not provide enough input to bring them to spike. The decision system then sends a second input to interface neurons, so as to provoke the firing of the M-neurons that encode the features of the envisaged motor activity. The set of firing M-neurons is an input set $inputMA2$ to A-neurons’ second compartment. Querying the network amounts to bringing some A-neurons to spike in response to this incomplete input and send backward input to O-neurons through the third compartment’s connexions, so that these spike and send backward input to interface neurons. The resulting firing of interface neurons is the network’s answer to the query.

Now, among the A-neurons that receive a significant amount of input, some support the representation of general action concepts (such as, e.g., $[OK]$, $[NE]$, $[OK]$), while others support the representation of more particular action concepts (such as, say, $[OK, NorthWall]$, $[NE]$, $[Failure]$). Obviously, we need the neurons that support the more particular concepts to take precedence over those supporting more general concepts, as more particular concepts convey more adequate information. To achieve this, we modulate the inputs $inputOA1$ and $inputMA2$ received by A-neurons by applying a multiplying factor to their connections weights. For any neuron o in $inputOA1$, the multiplying factor is a decreasing function of $\sum_{a \in A} w_{o,a}^1$. The idea behind this is that neurons with fewer output connections generally support more particular concepts than neurons with more output connections, and so that the A-neurons to which they connect should also encode more particular concepts. Multiplying factors for $cnxMA2$ are defined in the same manner. The modulated input sums are then computed in the usual way for each compartment of each A-neuron, and summed up to get the total input received by the neuron. The set $firedA$ of firing A-neurons is obtained by first adjusting A-neurons’ spike threshold so that only the A-neurons with the highest input sum fire, and then gradually lowering it until $|firedA|$ reaches a target number $TnbQueryA$. Backward input to O-neurons is computed incrementally as A-neurons fire, bringing O-neurons to fire as soon as their spike threshold is reached, and similarly for interface neurons.

The value of A-neurons’ spike threshold at the time a given A-neuron a fires measures how well a responds to $inputOA1$ and $inputMA2$. It depends on both the accuracy of the action concept encoded by a relative to the terms of the query (i.e., the initial situation and the envisaged motor activity), and the concept’s degree of generality: the more adequate and particular the encoded concept, the higher the received input, hence the spike threshold. Since information conveyed by adequate particular concepts is more reliable than information conveyed by less adequate or more general concepts, the value of A-neurons’ spike threshold at the time an interface neuron spikes under backward input is a good indicator of the prediction’s reliability. This allows the agent to associate a confidence degree to each predicted feature.

Learning. For each A-neuron a , we note $lastSpikedA_a$ the number of steps performed by the agent since the last time a underwent learning. These

values are initialized at random using some upper bound, and updated each time the neuron learns or the agent makes a step. Learning of action concepts takes place after each step made by the agent. In this case, *inputOA1* is the set *firedO* resulting from the agent observing its depart location, *inputMA2* is the set of M-neurons firing by proprioception as the agent performs the motor activity, and *inputOA3* is either the set *firedO* resulting from the agent observing its arrival location or *{Failure}*, depending on the case.

To select the set *learningA* of A-neurons to make learn, we compute, for each *a* in *A*, the total input sum *inputSumA1A2_a* received by *a* on its first and second compartments. If *inputSumA1A2_a* reaches a certain threshold *learnAT*, we put *a* in *learningA*. If the number of A-neurons obtained in this manner reaches the target number *TnbQueryA*, we directly proceed to the learning. Otherwise, we search for some more neurons, by boosting the inputs *inputOA1* and *inputMA2* in a way similar to the one used for O-neurons.

For each input set *i* in *{inputOA1, inputMA2, inputOA3}*, we compute a probability distribution *probaNewSynapsesA_i* on *i* that models the propensity of each neuron *x* in *i* to grow new connections to A-neurons' corresponding compartment. As previously for O-neurons, for each *i* and *x* this probability is a decreasing function of $\sum_{a \in A} w_{x,a}^n$ (with $n \in \{1, 2, 3\}$ depending on the input set).

For each neuron *a* from *learningA*, we compute *a*'s learning rate *LR_a*. The learning rate is an integer comprised between 1 and the sum *WS_A* of the weights of the connections to A-neurons' compartments. It represents the neuron's basic propensity to modify its synapses through learning, and is calculated in a way such that A-neurons with maximal *lastSpikedA* value among those in *learningA* get the maximal learning rate *WS_A*, those with minimal *lastSpikedA* value get the learning rate 1, and the others get learning rates in between.

The learning process depends on the accuracy of the agent's expectations relative to the action's outcome. If these expectations were both correct (i.e., all expected features are actually present) and complete (i.e., all actually present features were expected), then for any A-neuron *a* in *learningA*, the learning process replaces all of *a*'s inactive synapses in the first two compartments, but only some of them in the third one. Here, the exact number of replaced inactive synapses is *min(LR_a, i)*, where *i* is the number of inactive synapses, the synapses to replace being chosen at random among the inactive ones. In all three compartments, the new synapses replacing the deleted ones are chosen at random, following the input set's probability distribution (the process is similar to the one described above for O-neurons). If the agent's expectations were not correct, then for any *a* in *learningA* the learning process not only replaces all inactive synapses on *a*'s first two compartments, but also some active ones. Here the number of synapses to be replaced is *max(LR_a, i)*, where *i* is the number of inactive synapses, and the deleted active synapses are chosen at random. As for *a*'s third compartment, the learning process is the same as in the first case. If the agent's expectations were not complete, then

for any *a* in *learningA* the number of replaced synapses is *max(LR_a, i)* in all three compartments. In the case where the predictions were neither correct nor complete, the learning process successively runs the two above procedures. This differentiated learning process aims at promoting generalization when predictions are correct and complete while allowing a few neurons to specialize when they are not.

Computing Infrastructure

Research was carried out on a MacBook Pro with Apple M1 Max chip (2022). Operating System: macOS-15.2-arm64-arm-64bit. We used Python with Spyder IDE.

- . Spyder version: 6.0.0 (standalone)
- . Python version: 3.11.9 64-bit
- . Qt version: 5.15.8
- . PyQt5 version: 5.15.9