



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. 2024. hal-04779695

HAL Id: hal-04779695

<https://hal.science/hal-04779695v1>

Preprint submitted on 18 Nov 2024

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. The 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 *concepts*, which can have various degrees of generality. Presumably, they first form concepts about the encountered objects and situations, and then use these as building blocks for relational concepts, among which are 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 according to its own interests. 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, even though these are investigated in separate research fields.

Continual Learning tackles the problem of life-long 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 classification/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 in these approaches concepts are extracted from labelled classes of images, which is, again, unsuitable for open world autonomous agents. Furthermore, the vast majority of proposed methods 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. However 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. However said instances come with their own set of ground actions which must be known beforehand, hence this approach cannot be used in open worlds, where an agent needs to incrementally learn from sequential observations.

Reinforcement Learning (RL): Our work could be related to model-based approaches of RL Moerland et al. (2023), but 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 that would be able to learn online a model of the environment and to dynamically use it to make 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 suitable for 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).




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

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 $LF = \{OK, KO, NorthWall, EastWall, SouthWall, WestWall, Cold, Sound, \#0, \#1, \dots, \#24\}$. 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 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




Figure 2. Schema of the agent

feature/object recognition. Since we want the agent to learn by itself, we need this process to rely on unsupervised learning with unlabelled data. Neural networks doing this already exist 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 features of the agent’s current location. Furthermore, we suppose that the agent’s observations are always correct and complete.

Semantic memory forms concepts by binding together the perceived features, and stores them for further retrieval. Its modeling is the main 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 composed 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 common features 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 indifferently information.

4. Concepts and Action Laws

The agent is able to 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. Moreover, 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 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




Figure 3. Schema of the SNN. O-neuron #1 supports the concept [#0, Cold, SouthWall, WestWall, OK], and O-neuron #2 the concept [#5, SouthWall, KO]. A-neuron #1 supports the concept [[#0, Cold, SouthWall, WestWall, OK], [E, Orth], [#5, SouthWall, KO]], and A-neuron #3 the concept [[OK], [Diag], [OK]].

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 compartment 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 Subtai (2016).

The Network's functioning

We briefly describe the network's functioning. A more detailed account is provided in Supplementary Material.

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 that encode 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 procedure favors the firing of O-neurons that have been inactive for a long time, which are then re-used for 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 I-neurons 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. The learning process is similar to the first case for the other learning neurons.

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. Then, given an envisaged motor activity m , the decision system 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.

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.

For each learning A-neuron a we compute a learning rate LR_a , which is an increasing function of the number of steps performed since the neuron's last learning. The idea is that seldom used neurons tend to encode more particular concepts than often used neurons, and should thus be able to learn more rapidly to retain more features from a given situation.

The learning process depends on the accuracy of the agent's predictions relative to the action's outcome (see Section 6). 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 replaces all inactive synapses in a 's first two compartments with synapses from their respective input neurons, but only $\min(LR_a, i)$ in the third one, where i is the number of inactive synapses. If they are not correct, then all inactive synapses plus some active ones are replaced in a 's first two compartments, with $\max(LR_a, i)$ the number of synapses replaced in each one. The learning process for the third compartment is the same as in the first case. If the agent's predictions are not complete, then the number of replaced synapses is $\max(LR_a, i)$ in all three compartments. 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.

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 and is transmitted from there 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 the 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 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 UD , otherwise it goes for one with the least c in $S \cup US$. In *Exploitation* mode, by contrast, the agent just wants to land on an OK box and 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 UD . 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* and consisting in 65536 steps in total. The results we present here are averaged over 50 trials.

A first group of tests was carried out with the door kept closed all along, so the agent had no access to

the second room. First, we tested its ability to learn an action over one single experience (one-shot learning). To do so, after each step we asked it to redo the prediction that led to the just realized action, and compared this new prediction with the action’s actual outcome (see Table 1). 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. These results show a good performance at immediate recall after learning.

Table 1. Mean post-learning predictions percentages over 50 trials. CC: Correct and Complete, MF: Missing Features, PE: Prediction Errors.

Nb of Steps	1	2	4	8	16	32	64	128	256
CC	100.0	100.0	100.0	98.5	98.0	97.0	96.2	93.4	88.9
MF	0.0	0.0	0.0	0.6	0.8	1.2	2.0	3.0	4.8
PE	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3
Nb of Steps	512	1024	2048	4096	8192	16384	32768	65536	
CC	91.4	93.5	95.1	95.2	95.9	96.5	96.8	96.7	
MF	3.5	2.6	1.9	1.9	1.6	1.4	1.3	1.3	
PE	0.3	0.2	0.1	0.1	0.1	0.1	0.1	0.1	

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 2 and 3 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 for the second room (grey lines) show that despite never having been in this room (since we kept the door closed) the agent is able to correctly predict *OK* and *KO* features and to a lesser extent *Failure* —and this, even though locations

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.

Table 2. Predictions’ *Hit Rates* after n steps, door closed.

Nb of Steps	Room	OK	KO	Fail.	Wall	Cold	Sound	Box Name
1	1	13.4	11.9	0.0	0.0	0.0	0.0	1.4
	2	14.0	11.6	0.0	0.0	0.0	0.0	0.0
8	1	32.6	39.1	28.6	14.0	9.1	0.0	4.1
	2	28.3	35.7	28.2	11.1	0.0	0.0	0.0
64	1	77.5	77.7	45.3	35.0	26.0	0.0	22.4
	2	66.2	69.3	38.8	22.8	0.0	0.0	0.0
512	1	94.3	92.2	70.5	64.5	57.2	0.0	50.9
	2	84.5	84.6	40.2	23.4	0.0	0.0	0.0
8192	1	97.3	94.7	83.0	79.6	72.5	0.0	72.6
	2	91.7	87.2	52.0	26.0	0.0	0.0	0.0
32768	1	96.5	94.1	85.7	80.7	81.5	0.0	75.0
	2	89.5	84.6	52.2	23.8	0.0	0.0	0.0

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 regarding the presence of walls in adjacent boxes: the agent uses particular concepts to predict them in the first room hence it is helpless in the second room. The mixed result for failure prediction comes from a competition between general action concepts, the control of which needs to be improved.

Table 3. Predictions’ *Correctness* after n steps, door closed.

Nb of Steps	Room	OK	KO	Fail.	Wall	Cold	Sound	Box Name
1	1	18.0	17.5	0.0	0.0	0.0	0.0	4.0
	2	20.0	18.4	0.0	0.0	0.0	0.0	0.0
8	1	45.2	42.1	26.6	18.8	5.1	0.0	5.2
	2	43.4	40.6	20.4	11.5	0.0	0.0	0.0
64	1	74.8	74.5	58.7	34.0	28.6	0.0	22.7
	2	72.0	72.8	41.8	18.9	0.0	0.0	0.0
512	1	90.5	87.2	83.9	69.5	72.8	0.0	68.6
	2	83.4	83.1	47.4	25.9	0.0	0.0	0.0
8192	1	93.0	92.2	94.4	87.7	95.2	0.0	84.3
	2	85.5	86.5	66.2	36.3	0.0	0.0	0.0
32768	1	94.2	93.0	93.4	89.9	94.0	0.0	87.1
	2	85.3	86.2	60.7	38.1	0.0	0.0	0.0

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

Finally, we tested whether the agent would be able to use the knowledge acquired in the first room to act judiciously in the second room. To do so, 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 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 (36.6% of steps on average, standard deviation = 10.7).

Tables 4 and 5 show the features’s *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 seemingly low *Hit Rates* for the feature are explained by 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. These results also show that changing the agent’s environment does not lead to significant loss of previously acquired knowledge. What happens is that the network accommodates the new concepts by recruiting neurons that have not been used for a long time to encode them (remember the boosting of inputs sums in the learning processes of O- and A-neurons). As a result, the agent tends to forget the details of particular objects and actions such as boxes’ names or the presence of *Cold* and walls (see the moderate drop of these features’ *Hit Rates*), and more generally to forget old and unused concepts. But general concepts, which are used on a regular basis, are preserved, at least as long as they keep being used.

Accordingly, the agent’s ability to use its knowledge to make appropriate decisions in the first room is 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 and shown in Figure 4.A.

As regards computing time, it takes about 10 sec-

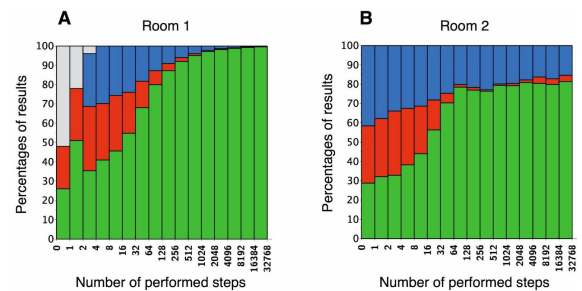


Figure 4. Actions’ mean outcomes with door closed; Green: *OK*, Red: *KO*, Blue: *Failure*, Grey: no data

Table 4. Predictions’ *Hit Rates* after n steps, door opened at 2048th step.

Nb of Steps	Room	OK	KO	Fail.	Wall	Cold	Sound	Box Name
2048	1	95.8	94.2	78.5	74.1	69.2	0.0	64.7
	2	89.3	86.2	45.7	24.2	0.0	0.0	0.0
4096	1	97.9	96.5	75.6	62.7	56.6	0.0	53.1
	2	93.7	92.9	57.4	34.2	0.0	35.0	5.1
8192	1	96.8	96.9	81.1	66.4	58.9	0.0	55.6
	2	93.7	93.6	64.0	36.5	0.0	34.2	4.2
16384	1	96.8	95.4	82.5	68.5	67.4	0.0	60.0
	2	92.4	91.3	66.0	33.9	0.0	32.1	5.3
32768	1	96.9	97.2	85.6	67.8	59.1	0.0	57.0
	2	93.2	92.5	75.5	39.7	0.0	38.4	7.6
65536	1	96.9	94.4	88.8	68.4	59.2	0.0	60.5
	2	93.6	90.7	80.5	40.5	0.0	39.6	7.1

Table 5. Predictions’ *Correctness* after n steps – door opened at 2048th step.

Nb of Steps	Room	OK	KO	Fail.	Wall	Cold	Sound	Box Name
2048	1	92.9	89.5	91.5	82.7	91.0	0.0	78.9
	2	85.2	83.7	58.0	36.6	0.0	0.0	0.0
4096	1	94.6	90.7	93.0	75.1	91.8	0.0	88.1
	2	90.9	87.8	75.9	44.7	0.0	64.0	1.2
8192	1	94.9	93.2	93.3	77.7	94.3	0.0	87.9
	2	91.1	89.2	79.7	44.8	0.0	61.5	1.4
16384	1	95.6	93.0	91.7	81.5	96.0	0.0	90.3
	2	91.3	89.2	75.9	47.3	0.0	66.2	1.1
32768	1	96.2	94.5	94.1	78.3	95.2	0.0	91.0
	2	93.2	91.5	79.9	50.8	0.0	64.0	2.5
65536	1	96.2	95.9	90.6	82.1	94.6	0.0	94.7
	2	93.6	92.8	79.8	51.7	0.0	65.4	2.2

onds 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.

Finally, we highlight that the choice of a localist representation for the neural network makes it easy to read, since the concept encoded by each neuron can be known simply by looking at its connections.

8. Conclusion and Future Developments

In this paper we have designed and implemented 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 for the second room. 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 ensures that it

will always be able to learn about new environments, by replacing old unused concepts by new useful ones.

Further work remains to be done to endow the agent with planning abilities. Notably, a notion of *applicable* action law would be needed. Intuitively, it seems that an action law represented by an action concept $[\mathbf{x}, \mathbf{y}, \mathbf{z}]$ should be deemed applicable in a situation s if s satisfies all the features in \mathbf{x} and $\mathbf{z} \neq \text{Failure}$. Furthermore, to comply with open world requirements the agent would need to be able to build its own set of possible situations (states) online. The set of its object concepts could probably be used to this end. A cost function should also be added, and the decision system should be augmented so as to handle goals.

Additionally, a number of other improvements would be desirable in order to allow the agent to live in more realistic environments. A first one would be to implement negation in the network, so that the agent would be able to represent the fact that a given object does not have a given feature. We believe that this could be done by the means of neural inhibition, but the appropriate learning rules remain to be found. Another essential improvement would be to have the network use incomplete information as input for learning and querying, and to make the agent able to query its semantic memory for object properties given some partial input. It seems to us that this 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, since actions can modify one, the other or both. Biological brains achieve this by using two separate pathways to process the “what” and the “where” components of observations before reunifying them, and this could be an inspiration source. Lastly, a completely different line of research would be to investigate how the 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.