Identifying and Localizing Dynamic Affordances to Improve Interactions with Other Agents in Continuous Environments

Simon L. Gay[1], Olivier L. Georgeon[2] and Jean-Paul Jamont[1]

Abstract— **Allowing autonomous agents to learn by themselves to interact with other agents requires that they be able to recognize each other and be capable of inferring their behaviors. To comply with artificial developmental learning, we follow the radical interactionism hypothesis, in which an agent has no *a priori* knowledge on its environment. A previous work has shown that the agent can learn to identify, localize, and predict movements of mobile elements, but was only tested in discrete environments, limiting their applicability on real-world systems. This paper presents new mechanisms for the identification and localization of mobile entities in a continuous environment. These mechanisms learn the relations between the agent's sensorimotor patterns and the entities, static or mobile, affording them, and store discovered properties in data structures called Signatures. The properties of signatures are then exploited to detect distant entities in surrounding environment without relying on a geometrical notion of space. These mechanisms were tested in a simple environment, the results showed how signatures integrate the complexity of a continuous environment through the limited sensory system of the agent, and localize distant entities through data structures that are compatible with previously developed behavior inference and prediction mechanisms.**

Index Terms—developmental learning, interactionism, affordance, autonomous mental development, spatial awareness.

I. Introduction

We address the problem of how an artificial agent with no prior ontological knowledge of its environment can generate an emergent model of its environment, and acquire knowledge about autonomous mobile entities (e.g. other agents), particularly in a continuous environment.

This study relates to the domains of artificial constructivist learning [1] and enactive learning [2], in which learning occurs through the enaction of control loop implementing Piagetian sensorimotor shemes [3], which we call *interactions*. More precisely, this study uses a modeling hypothesis called Radical Interactionism (RI) [4], in which an agent starts with a predefined set of uninterpreted *interactions*. These interactions are associated with in-born numerical values, called *valences*, that define the agent's *drives* without relying on an external interpreter, states or predefined goal. By experiencing its environment through interactions, an RI agent constructs a model that it can exploit to generate behaviors allowing to enact interactions with high valences. Thus, unlike other approaches seeking optimal policies to accomplish specific tasks (e.g. Reinforcement Learning [5], Deep RL [6], MARL [7]), RI focuses on the development of agents capable of constructing a reliable model of their environment from sensorimotor experience. Valences steer decisions as an operating drive, but are not taken into account in learning as a parameter to be optimized, thereby dissociating the generated sensorimotor model from the decision-making model. Thus, the study focuses more on the reliability of the constructed model and predictions than on efficiency in solving a predefined task.

As we study more complex environments, we face the problem of dealing with other agents. In a multiagent perspective, the RI approach must be extended to collective, adversarial or collaborative behaviors. Previous studies [8], [9] demonstrated that an RI agent can integrate mobile entities (reactive agents such as preys) into its model, detect them, infer their behavioral preferences according to surrounding static entities and finally predict their future moves. However, these studies were conducted in discrete environments, limiting their applicability to real-world systems like robots.

Here we extend the agent's abilities to integrate other mobile entities in its internal model; specifically focusing on the integration of such entities and their detection and localization in a continuous space. The paper is organized as follows: the rest of the introduction summarizes the RI formalism devised from the literature on sensorimotor learning [10], affordances [11], [12] and schema mechanisms [13]. Section II presents the model for defining and recognizing mobile entities and Section III presents the model for localizing entities in space. Finally, Section IV encompasses conclusive remarks and future developments.

Formally, an RI agent starts with a predefined set I of *interactions* (control loops), each associated with an inborn valence $\nu_i \in \mathbf{R}$. At the beginning of Step t, the agent selects an *intended interaction* $i_t \in I$ to try to enact. At the end of Step t, it receives the *actually enacted interaction* e_t. The enaction is a *success* if $i_t = e_t$ and a *failure* otherwise. An example failure may be when an agent intends to move forward ($i_t = \texttt{move forward}$), but actually collides with an obstacle ($e_t = \texttt{collide}$). Enacted interactions are the only means for the agent to perceive its environment. It learns to predict the result of future intended interactions in the context of previously enacted interactions to select behaviors enabling the enaction of interactions of high valence.

The *Parallel RI* (PRI) model [14] allows the simultaneous enaction of multiple interactions. It distinguishes between *primary interactions* defined by couples (action, outcome), and *secondary interactions* defined by couples (interaction,

*This work is supported by the French National Research Agency in the framework of the "Investissements d'avenir" program (ANR-15-IDEX-02).

[1] LCIS, Université Grenoble Alpes, Valence, France simon.gay/jean-paul.jamont@lcis.grenoble-inp.fr

[2] UR Confluence: Sciences et Humanites (EA 1598), UCLy, France ogeorgeon@univ-catholyon.fr

additional outcome). Additional outcomes result from movements produced by an interaction. The optical flow is an example of such outcome that must be associated with a movement to characterize a position in space. At the end of step t, the agent receives the set of enacted interactions $E_t \subset I$ containing the primary enacted interaction and several secondary enacted interactions resulting from the primary.

The PRI model uses *signatures of interactions* to evaluate the possibility to enact an interaction based on the previous enacted context E_{t-1}. The signature S_i of interaction i thus evaluates the presence of the affordance of i by defining one or more sets of interactions $\{j_k\}$ whose enaction (i.e. $\{j_k\} \subset E_{t-1}$) characterizes the possibilty of successfully enact i at Step t. Formally, a signature of interaction is a function $S_i : \mathcal{P}(I) \mapsto [-1; 1]$, (with $\mathcal{P}(I)$ the partition of I, i.e. the set of all possible contexts). $S_i(E_t) \in [-1; 1]$ gives the prediction of successfully enacting i (1 for certainty of success and -1 for certainty of failure). The signature's pseudo-reverse function $\hat{S}_i : \{1, -1\} \mapsto \mathcal{P}(\mathcal{P}(I))$ provides either the minimal context(s) $C_i^k \in I$ affording ($\hat{S}_i(1)$) or preventing ($\hat{S}_i(-1)$) interaction i.

The signature of an interaction i designates sets of interactions $\{j_k\}$ that can have their own signatures S_{jk}. By using signatures S_{jk} of interactions j_k that are related to the same primary interaction j, it is possible to define a context S_i^j that, after enacting j, will affords i. This process is applied recursively through a sequence $\sigma = \langle j^1, ..., j^n \rangle$ of primary interactions to detect *distant affordances* through "projected" signatures S_i^σ. This allows the emergence of an implicit notion of space based on sequences of interactions that extends in the extrapersonal space of the agent. Defining affordances through interactions also overcomes limitations of sensori-defined affordances (e.g. [15], [16]) that are limited to defining the next action or require additional spatial information to make longer-term predictions [17].

II. Integration of Mobile Affordances

The main difficulty in integrating a mobile affordance is that, even if the affordance is present, it can move during the interaction's enaction, leading to the failure of the interaction. In [8], we proposed a model based on the observation that a signature related to a mobile element defines prediction certainties that are greater when the affordance is present than when it is absent. The model thus prevents the signature from updating after a failure when the predicted certainty is above the average prediction value of failure (i.e. lower absolute value), enabling the emergence of the signature.

Also, a mobile entity can afford an interaction from multiple positions, depending on its movement's direction, and each movement may have a different probability. The model defined a signature architecture (Fig. 1) based on a layer of formal neurons, connected with a winner-takes-all rule and a competition mechanism. In the case of a success of i, only the neuron with the strongest output is reinforced as a success, whereas all neurons are reinforced in case the of failure. The competition mechanism forces each neuron N_i^k to specialize on a unique context, noted C_i^k. To integrate

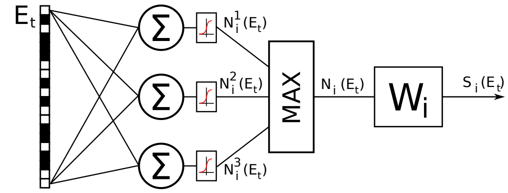

Fig. 1. Implementation of signatures for the integration of mobile affordances from [8]. Neurons are in competition, enabling the integration of multiple positions from which a mobile object can afford the interaction.

"negative affordances" (i.e. the enaction is afforded by the absence of elements), a second layer that has a neuron with a single weight W_i allows to invert the first layer's result, and defines the signature output as $S_i(E_t) = N_i(E_t) \times W_i$.

The competition principle is that we take the strongest weight $w_{k_{max}}^{n_{max}}$ of the winner neuron that is connected to an enacted interaction $i_{k_{max}} \in E_t$, and reset the weight correspnding to $i_{k_{max}}$ of other neurons. This principle enables the identification of independent contexts C_i^k and defines their respective probability. However, in a continuous environment, a mobile affordance can afford an interaction from a continuous set of positions, potentially overlapping, making a single neuron progressively takes all possible contexts.

A. A model of signature for continuous environment

We adapted the learning process and competition mechanism for continuous environments. The learning process starts with a deactivated competition mechanism: all neurons learn simultaneously. Failures of i with a prediction above the average prevent weights' update, allowing all neurons to integrate all contexts affording i. While the signature emerges, the certitude values of success increase. When the signature S_i predicts a success with a certitude above a threshold, the competition mechanism is used, forcing the distribution of contexts on neurons.

From our preliminary tests, the following competition principle was retained for its stability. This principle is applied after the neurons' update. The winner neuron $N_{i,max}$ reduces the strength of weights related to non-enacted interactions ($i_k \notin E_t$), making it specialized in a context $\hat{C}_i$ of co-occurring interactions, and leaving other possible contexts to other neurons. The reduction of weights of the winner neuron is of a constant δ_1 weighted by the inverse of the measured probability of success $p_{i,s}$ of its context (1) :

$$w_k^{t+1} \leftarrow w_k^t \quad - sgn(w_k^t).\delta_1.(1 - p_{i,s}) \quad if\ i_k \notin E_t \quad (1)$$

where w_k is a weight of the winner neuron associated with an interaction $i_k \in I$, and sgn a function giving the sign. The winner neuron also slightly removes its context from other neurons. For each neuron $n \neq n_{max}$, positive weights $w_{n,k}$ related to an enacted interaction ($i_k \in E_t$) that is related to a positive weight of neuron n_{max} are reduced as (2):

$$w_{n,k}^{t+1} \leftarrow w_{n,k}^t - \delta_2.w_{n_{max}}^k \quad if\ i_k \in E_t, w_{n,k} > 0, w_{n_{max}}^k > 0 \quad (2)$$

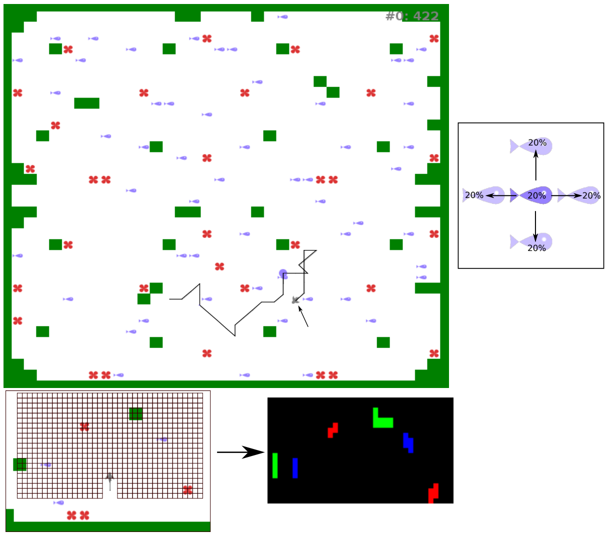

Fig. 2. The test environment. The grey shark (bottom left) is our agent. At each simulation step, the fish (preys) can move randomly up, down, left, and right or remain immobile. The bottom-left frame shows the visual system of our agent. The pixel scale does not match the size of objects.

This second competition principle forces other neurons to progressively abandon the observed context $\hat{C}_{i,n_{max}}$.

B. Test environment

This signature mechanism was tested on an artificial agent moving in the 2-dimensional environment shown in Fig. 2. Although the environment's content is define as a matrix of entities, the agent can move freely in it and quickly misaligns with other elements.

The sensorimotor possibilities of the agent define the following eight primary interactions: *move forward* ▷, *bump* in an obstacle ▶, *eat* ▶, *slide* on a soft object ▶, turn right ◡ and left ◠ by 90°, turn right ▿ and left ▹ by 45°.

This visual system discretizes space in front of the agent as a regular grid of 37x21 positions (unknown to the agent) and perceives 3 colors. The agent can see over algae and fish, but not through walls. The distance between two adjacent positions is 0.4 grid unit. This scale was selected so that wall blocks are never aligned with the visual field's grid. The visual system can generate 37x21x3=2331 additional stimuli, produced by the movement of a primary interaction (except for *bump* that does not generate movement), for a total of 16317 secondary interactions.

We used signatures composed of 9 formal neurons. The competition coefficients used in the presented experiments are $\delta_1 = 0.02$ and $\delta_2 = 0.1$. The signature learning process is driven by a learning mechanism that foster interactions with low certainty of success or failure $|S_i(E_t)|$. Note that testing an interactions leads to a simultaneous update of signatures of interactions associated to the same primary interaction.

The environment is populated with three types of elements, characterized by different colors to enable their identification through the limited sensori system of the agent. Walls are solid green squares of one unit affording *bump*. Preys and algae are respectively blue elements affording *eat* and red elements affording *slide*. Preys and algae have a round 'hitbox' of radius 0.35 unit. the agent, represented as a gray shark, has a round 'hitbox' of radius 0.4 unit. The agent can move through algae (enacting *slide*). When moving over a fish (enacting *eat*), it is removed and another one is randomly set in an empty position.

The fish move randomly to simulate agents with unknown behavior: at each simulation step, they can stay immobile, or move left, right, up, or down, with a probability of 20% each. If the fish cannot move in the selected direction because of the presence of another object, it remains immobile, making the immobile situation slightly more probable than others.

C. Properties of Signatures of Interactions

During the learning process, we can observe all neurons integrating the contexts affording their associated interaction. Once a signature begins to provide high certainties more frequently, the competition between neurons starts.

In the case of static affordances (i.e. *bump* and *slide*), one of the neurons associates with the context corresponding to the presence of a green (respectively red) element in front of the agent, and inhibits other neurons. This neuron stabilizes after less than 5000 simulation steps and define a high probability ($> 80\%$), in the same way than in discrete environment [8]. However, the evolution of the signature starts to differ after 20000 simulation steps: we can observe the emergence of other contexts with lower probabilities, designating the presence of an entity in a position just next to the position in front of the agent. Indeed, as the visual system of the agent has a low spatial resolution, these contexts refers to a range of positions at the limit from where the element can afford or not the interaction. Thus, in this case, the probability reflects the proportion of positions affording the interaction within the interval separating two adjacent visual positions. We will refer to these contexts as 'limit contexts' in the following descriptions. Fig. 3 and 4 show respectively the signature of interactions *bump* and *slide*.

There are also major differences with signatures of mobile affordances. As observed in discrete environment, the context related to the presence of a fish in front of the agent is the first to emerge and stabilize (after 8000 simulation steps) in the signature of eat, as this context is the most probable. The neuron integrating this context designates a set of visual interactions related to seeing a blue element, in positions that are in front of the agent, thus characterizing the position from which the agent can enact *eat* when the fish remains immobile. We can notice that the size of the 'blue blob' (from an external point of view) is slightly larger than the size of a fish, and the probability is higher than in the discrete environment (around 35% instead of around 25%). The reason is that there are positions partially overlapping the position in front of the agent from where a movement still leads to a success of eat. These positions are thus captured as part of the 'front' context, which also increases its probability. Then, after 10000 simulation steps, other neurons start to integrate contexts related to surrounding positions. As the continuous environment offers a continuous

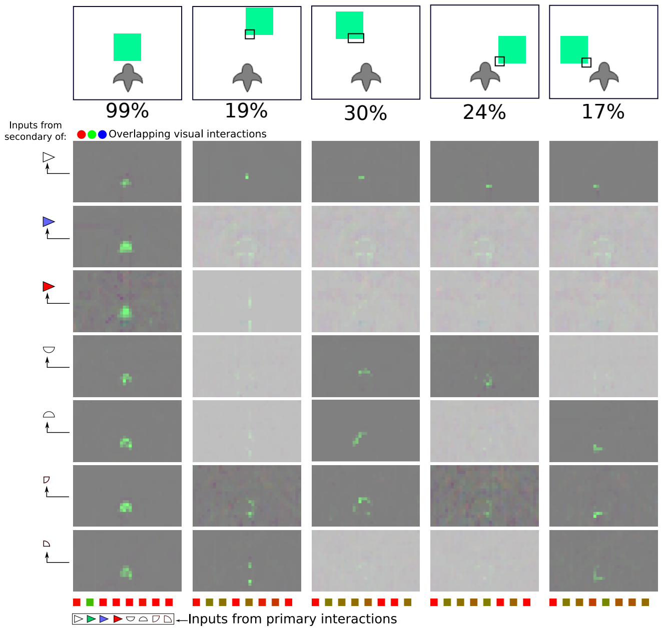

Fig. 3. Signatures of interaction *bump*, recorded after 300 000 simulation steps. A signature is characterized by the weights of 9 formal neurons. The signature only identified five contexts, therefore we only represent 5 out of the 9 neurons, each neuron being represented by a column. As external observers, we can organize weights to make signatures more readable: first, weights related to primary interactions are represented with eight squares below (green for a positive weight, red for a negative weight). Weights associated with secondary interaction are grouped according to their primary interaction, forming the seven groups (from top to bottom: forward, eat, slide, turn left 90°, right 90°, left 45°, right 45°; *bump* does not produce visual interactions). Each group is organized to place visual interaction with their associated position in the visual field. Colors are overlapped to generate signatures under the form of an RGB image. On the first row (weights related to secondary interactions associated with *move forward*), the Signature identified a context that consists of seeing a green element in front of the agent, with a high probability implying this affordance is immobile, and a set of four other contexts that correspond to ambiguous perceptions due to the low resolution of the visual system, leading to a lower probability. The context from external point of view and defined probabilities are represented on top. Other rows show similar structures, although they take more time to emerge as the associated primary interactions are less often enacted (especially *eat* and *slide*). Unreliable contexts are grayed. *bump* is also related to the success of *bump* (green square on the bottom line), since this interaction can be enacted repeatedly.

set of positions, the competition between neurons makes each neuron trying to integrate a particular range of contexts. The signature's neurons stabilize on a stable distribution, with a large set of positions in front of the agent, with a probability around 35%, and a set of nearly equidistant sets of positions around with probabilities depending of their distance from the front position. Fig. 5 shows the signature of eat and the probabilities of the discovered contexts.

Interaction *move forward* is afforded by the absence of element in front of the agent. The signature's weight W (second layer's neuron) quickly converges to -1 (less than 1000 simulation steps), which means that the neurons integrate contexts preventing this interaction. The signature first integrates, in the first 5000 simulation steps, contexts associated with the presence of a green, blue or red element in front of the agent. Then, the contexts starts to split

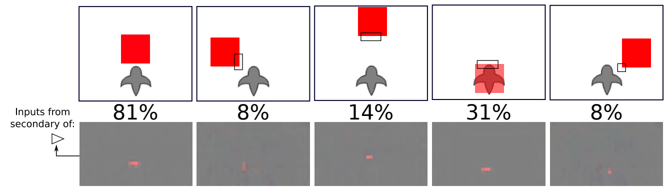

Fig. 4. Signatures of interaction *slide*, recorded after 300 000 simulation steps. The signature only identified five contexts, therefore we only represent 5 out of the 9 neurons, each neuron being represented by a column. We only represent the row of weights associated with secondary interactions related to *move forward*. The signature identified a highly probable context that consist of seeing a red element in front of the agent, and a set of four less probable other contexts that correspond to ambiguous positions.

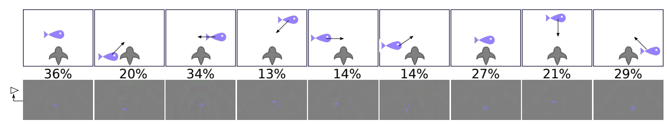

Fig. 5. Signatures of interaction *eat*, recorded after 300 000 simulation steps. All the 9 neurons are used by the signature. We only represent the row of weights associated with secondary interactions related to *move forward*. The signature identified multiple contexts, segmenting the possible positions from which it is possible to interact with the mobile affordance, depending on its movements, with the most probable contexts in front of the agent.

and spread on other neurons. It appears that 9 neurons are not sufficient to integrate all possible contexts, including the multiple contexts related to a fish and 'limit' contexts related to walls and algae: a same neuron can integrate more than one context. We can however observe that a neuron integrates contexts with similar probabilities, as an example, two contexts related to the presence of a fish and a wall in a peripheral position, which, if compared with signature of *eat* and *bump*, both have a probability close to 15%. Due to the competition mechanism, the signature needs more than 100000 simulation steps until finding a stable configuration.

The signature of visual interactions are also analyzed. The signatures related to seeing a green or red element at a specific position of the visual field correspond to seeing a green or red object at a position that is consistent with the movement produced by the associated primary interaction (one step ahead for *move forward*, rotated of 90° left for a *left turn*...). In the case of a green element, we also observe contexts designating positions next to the main one with low probabilities, that we do not observe for red-related interactions. The reason is that walls often form lines in the environment, making it possible to infer the presence of a green element through its neighborhood. The signatures of interactions related to seeing a blue element designate multiple contexts with a most probable context at a position consistent with the movement of the primary interaction, and a set of other less probable positions surrounding it that correspond to different possible movements of fish.

III. Detection of Distant Affordances

The detection of distant affordances exploits two properties of signatures: first, a signature designates an affordance as sets of interactions $\{jk\} \subset \hat{S}_i(1)$, and each interaction j_k may have its own signature. Then, signatures of secondary

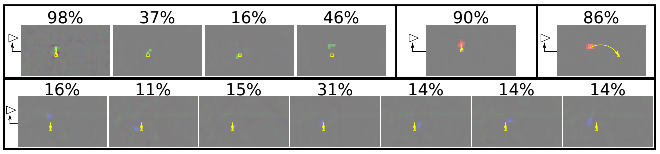

Fig. 6. Signatures of a sample of visual interactions, recorded after 300 000 simulation steps. We only display weights related to secondary interactions associated with *move forward* of contexts considered as reliable. The visual interactions are associated with a position shown with a yellow square in the visual field. Top left: *seeing green while moving forward*. The signature identifies a highly reliable other position matching the movement produced by *move forward* (yellow arrow). Three other contexts designate green elements around this position, due to the walls that often form lines. Top-middle and top-right: interactions *seeing red while moving forward* and *seeing red while turning left by* $90°$. A unique context is discovered at a position matching the movement of *move forward* and *turn left by* $90°$. Bottom: *seeing blue while moving forward*. A set of seven contexts is discovered. The most probable one matches the movement of *move forward* (the blue object does not move), while others, with lower probabilities, are shifted positions (the blue object moves in the right direction to afford the visual interaction). These examples show how signatures encode movements of primary interactions.

interactions encode the movement produced by primary interactions. These properties enable the recursive projection of a signature S_i through a sequence σ of primary interactions. The projected signature, noted S_i^σ allows the detection of a distant affordance instance of i at a *position* characterized by the movement produced by the enaction of sequence σ.

In [8], we extended the projection mechanism to mobile affordances by integrating the probability of individual contexts. The interactions designated by a signature are projected individually to define *projection sequences*, that are triplets (σ, λ, p), where σ is the path sequence, λ is the sequence of projected interactions, with $\lambda[0]$ being the final interaction of the projection sequence, and p the probability of the sequence, obtained by recursively multiplying with the probability of contexts. When two projection sequences have the same final interaction $\lambda[0]$, they are merged if they have the same path σ and start from the same context C_i^k, otherwise the sequence with the lowest probability p is removed. This filter ensures that there is only one projection sequence for each path σ and context C_i^k leading to an interaction j ($\lambda[0]$). The consequence is that the presence of a mobile affordance in the surrounding space will trigger a set of projection sequences with the same $\lambda[0]$, each starting from a different context C_i^k and whose path sequences indicate the possible next positions of the affordance [9].

In a continuous environment, a signature S_i designates contexts composed of multiple interactions j_k, some of them characterizing the 'edge' of the element affording i. Some projection sequences thus project through a sequence λ of interactions characterizing such 'edges', producing a projected signature designating an affordance with an overestimated size. If these projected sequences start from a context C_i^k with a high probability, they can replace more pertinent sequences (i.e. characterizing a position centered on the affordance) but with lower probabilities. As observed in signatures, the center of an affordance recognized by a context C_i^k is designated by interactions with the highest weights. We thus propose to integrate the weights' strengths into the projections process. The strength of a weight w_m^k in a context C_i^k indicates how much the presence or absence of the associated interaction j_m in E_t affects the result of $C_i^k(E_t)$. We propose to add to projection sequences a normalized value $\omega_m^k = w_m^k/|w_{max}^k|$, where w_{max}^k is the weight with the greatest absolute value in C_i^k.

The modified process to generate projection sequences is defined as follows: a projection sequence is a tuple $(\sigma, \lambda, p, \omega)$. The projection starts with a context C_i^k of a signature S_i that is decomposed into projection sequences $(\sigma_0, \lambda_0, p_0, \omega_0)$, where $\sigma_0 = \langle \, \rangle$ is an empty sequence, $\lambda_0 = \langle j_m \rangle$, with j_m an interaction of $\hat{C}_i^k(1)$, p_0 the probability of context C_i^k and $\omega_0 = w_k/|w_{max}|$, with w_k the weight connected to j_k. The recursive projection through a primary interaction j allows the generation, from a projection sequence $(\sigma, \lambda, p, \omega)$, a set of sequences $(\langle j, \sigma \rangle, \langle j_m, \lambda \rangle, p \times p_m, \omega \times \omega_m)$, where j_m is an interaction of a context $\hat{C}_{\lambda[0]}^l(1)$ associated with primary interaction j, p_m the probability of this context and ω_m the normalized weight value of j_m in this context. Thus, the normalized value ω of a projection value characterizes how much the projected interaction $\lambda[0]$ characterizes a position that is "closed" from the distant affordance's center.

The sequence filter integrates this value: when two projection sequences lead to the same final interaction $\lambda[0]$ through the same path σ, but come from different contexts C_i^k, we only keep the sequence with the greatest term $p \times \omega$. The selection is thus a compromise between the probability and the pertinence, i.e. the proximity to the affordance center.

Then, when the final interaction $\lambda[0]$ of a projection sequence is in E_t, a candidate affordance instance of i is detected at position σ. In a continuous environment, the same entity can be detected through multiple enacted interactions, as shown in Fig. 2. We thus propose to detect and gather sequences locating the same affordance instance of an interaction i. First, the detection mechanism gathers detected projection sequences that have the same final interaction $\lambda[0]$. An affordance instance A_i is thus characterized by a set Ξ of sequences σ_k and a final interaction $j = \lambda[0]$. Then, when two instances have a sufficient number of path sequences in common, i.e. $A_{i,1}, A_{i,2} \setminus Card(\Xi_1 \cap \Xi_2) \neq \emptyset$, these instances are merged. At this point, an affordance instance is defined as a data structure containing a set Ξ of sequences of interactions leading to it and a set $\{j_k\}$ of interactions characterizing it from the agent's point of view.

It is then possible to extract information from these data structures and filter sets Ξ: if the set contains at least one sequence projected from a context C_i^k with a high probability (we use a threshold of 80%), then the affordance instance can be considered as a *static* element. Therefore, sequences projected from less probable contexts, such as "limit" contexts observed in signatures of *bump* and *slide*, can be removed, leading to a more accurate localization of the affordance instance. Otherwise, the affordance instance is considered as a *mobile* affordance. The position of static affordances can be deduced by the length and the first interaction of the shortest

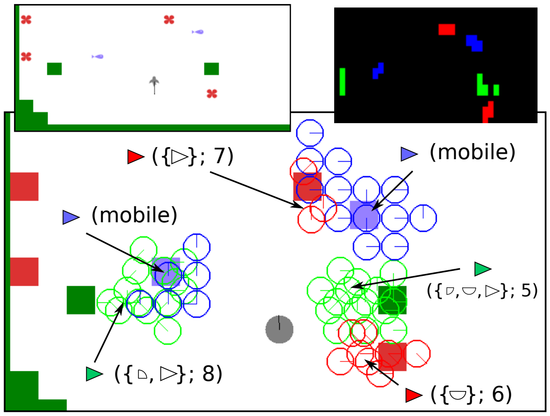

Fig. 7. Detection of distant affordance instances in the configuration shown in the top-left insert. The top right insert shows the visual outcome after enacting *move forward* (the vector of visual outcomes is organized to match positions and colors in the visual field for better readability). Bottom: the affordance view (observer's point of view). Each circle represents the position and orientation that can be reached by enacting a sequence σ of a detected affordance instance. Circles thus show the position from which the interaction should be enacted (i.e. in front of the entity). The detection mechanism identified and segmented 6 affordance instances. The two instances of *bump* and the two instances of *slide* are identified by sequences projected from high probability contexts, allowing to define them as *static* affordances. They are localized with sets of sequences defining a dense set of positions in front of the entities that form the affordances (walls and algae). From these sequences, the detection mechanism can extract the direction and distance from shortest sequences: the first interactions of shortest sequences define the 'orientation', and the length of shortest sequences, the 'distance' (including *turns* interactions). These data structures can be stored in subsequent mechanism [9] to define a allocentric reference. The two instances of *eat* are recognized as mobile affordances, and localized with a set of sequences leading to a sparse set of positions, which indicate possible next positions of the affordance. Integrating the affordance context as a set of sequences of interactions makes it possible to use subsequent mechanisms developed in discrete environment for behavioral preference inference and movement predictions.

sequences of the set Ξ, characterizing the distance and direction in egocentric reference, while sequences leading to a mobile affordance characterize the set of possible next positions. Fig. 7 shows a distant affordance detection in an example of environment configuration.

The fact that affordance instances are characterized by data structures consisting of sequences of interactions produces a form of discretization of the continuous environment context. Moreover, these generated data structures are the same as those used and exploited in the discrete environment mechanisms [9] to store static affordances, detect and measure the movements of mobile affordances through the allocentric context of static affordances, infer their behavioral preferences and predict their future movements.

IV. Conclusion

This work presents an adaptation of our model to continuous environment, allowing an artificial agent with no prior knowledge to construct an exploitable model of an environment containing agents with unknown decision mechanisms. We focus here on the mechanisms enabling the integration of mobile entities and their localization in the surrounding environment. This work shows that adapting the model to constraints of a continuous environment does not require major changes to the model. However, we can observe interesting additional data from the generated structures, illustrating how signatures of interaction can integrate properties of the environment that can be exploited by subsequent mechanisms. The generated data structures used to localize distant affordance instances are similar to those generated by the discrete version of the model, suggesting that subsequent mechanisms for inferring and predicting the behavior of mobile entities can be adapted with minor changes.

Our tests also showed the importance of dynamically adapting the number of contexts, i.e. the number of neurons in signatures, a problem that we will investigate. Our future work will also study the complete prediction mechanism, and the emergence of collaborative/competitive behaviors, opening intersubjectivity possibilities between agents in real-world systems.

References

[1] K. R. Thórisson, *A New Constructivist AI: From Manual Methods to Self-Constructive Systems*. Atlantis Press, 2012, pp. 145–171.

[2] T. Froese and T. Ziemke, "Enactive artificial intelligence: Investigating the systemic organization of life and mind," *Artificial Intelligence*, vol. 173, no. 3, pp. 466–500, 2009.

[3] J. Piaget, *The construction of reality in the child.* New York: Basic Books, 1954.

[4] O. L. Georgeon and D. Aha, "The Radical Interactionism Conceptual Commitment," *Journal of Artificial General Intelligence*, vol. 4, no. 2, pp. 31–36, 2013.

[5] A. K. Shakya, G. Pillai, and S. Chakrabarty, "Reinforcement learning algorithms: A brief survey," *Expert Systems with Applications*, vol. 231, p. 120495, 2023.

[6] Q. Li, T. Lin, Q. Yu, H. Du, J. Li, and X. Fu, "Review of deep reinforcement learning and its application in modern renewable power system control," *Energies*, vol. 16, no. 10, 2023.

[7] C. Zhu, M. Dastani, and S. Wang, "A survey of multi-agent deep reinforcement learning with communication," *AAMAS*, vol. 38, no. 1, Jan. 2024.

[8] S. Gay, J.-P. Jamont, and O. L. Georgeon, "Identifying and localizing dynamic affordances to improve interactions with other agents," in *IEEE ICDL*, 2022.

[9] S. Gay, F. Suro, O. Georgeon, and J.-P. Jamont, "An affordance-based intersubjectivity mechanism to infer the behaviour of other agents," in *IEEE ICDL*, 2023, pp. 403–408.

[10] M. Hoffmann, H. Marques, A. Arieta, H. Sumioka, M. Lungarella, and R. Pfeifer, "Body schema in robotics: A review," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 4, pp. 304–324, 2010.

[11] J. J. Gibson, "The theory of affordances. perceiving, acting, and knowing: toward an ecological psychology," *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*, pp. 67–82, 1977.

[12] E. Sahin, M. Cakmak, M. Dogar, E. Ugur, and G. Üçoluk, "To afford or not to afford: A new formalization of affordances toward affordance-based robot control," *Adaptive Behavior*, vol. 15, pp. 447–472, 2007.

[13] O. l. Georgeon, F. S. Perotto, K. R. Thórisson, A. Sheikhlar, and P. Robtertson, "Schema mechanism 2.0 for developmental artificial intelligence," in *Proceedings of the international workshop on situated self-guided learning*, 2025.

[14] S. Gay, A. Mille, O. Georgeon, and A. Dutech, "Autonomous construction and exploitation of a spatial memory by a self-motivated agent," *Cognitive Systems Research*, vol. 41, 08 2016.

[15] J. Baleia, P. Santana, and J. Barata, "On exploiting haptic cues for self-supervised learning of depth-based robot navigation affordances," *Journal of Intelligent & Robotic Systems*, vol. 80, 2015.

[16] K. Qian, X. Jing, Y. Duan, B. Zhou, F. Fang, J. Xia, and X. Ma, "Grasp pose detection with affordance-based task constraint learning in single-view point clouds," *J. of Intelligent and Robotic Systems*, vol. 100, 2020.

[17] C. Pohl, K. Hitzler, R. Grimm, A. Zea, U. D. Hanebeck, and T. Asfour, "Affordance-based grasping and manipulation in real world applications," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2020, pp. 9569–9576.