



Identifying and localizing dynamic affordances to improve interactions with other agents

Simon Gay, Jean-Paul Jamont, Olivier L. Georgeon

► To cite this version:

Simon Gay, Jean-Paul Jamont, Olivier L. Georgeon. Identifying and localizing dynamic affordances to improve interactions with other agents. IEEE International Conference on Development and Learning (ICDL 2022), IEEE, Sep 2022, Londres, United Kingdom. 10.1109/ICDL53763.2022.9962231 . hal-03828062

HAL Id: hal-03828062

<https://hal.science/hal-03828062>

Submitted on 25 Oct 2022

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.

Identifying and localizing dynamic affordances to improve interactions with other agents

Simon L. Gay
LCIS, Univ. Grenoble Alpes
Valence, France
simon.gay@lcis.grenoble-inp.fr

Jean-Paul Jamont
LCIS, Univ. Grenoble Alpes
Valence, France
Jean-Paul.Jamont@lcis.grenoble-inp.fr

Olivier L. Georgeon
EA1598, Lyon Catholic University
LIRIS, Claude Bernard University
Lyon, France
ogergeon@univ-catholyon.fr

Abstract—Allowing robots to learn by themselves to coordinate their actions and cooperate requires that they be able to recognize each other and be capable of intersubjectivity. To comply with artificial developmental learning and self motivation, we follow the *radical interactionism* hypothesis, in which an agent has no *a priori* knowledge on its environment (not even that the environment is 2D), and does not receive rewards defined as a direct function of the environment’s state. We aim at designing agents that learn to efficiently interact with other entities that may be static or may make irregular moves following their own motivation. This paper presents new mechanisms to identify and localize such mobile entities. The agent has to learn the relation between its perception of mobile entities and the interactions that they afford. These relations are recorded under the form of data structures, called *signatures of interaction*, that characterize entities in the agent’s point of view, and whose properties are exploited to interact with distant entities. These mechanisms were tested in a simulated prey-predator environment. The obtained signatures showed that the predator successfully learned to identify mobile preys and their probabilistic moves, and to efficiently target distant preys in the 2D environment.

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

I. INTRODUCTION

We address the problem of how an artificial agent that learns an emergent model of its environment through interaction can acquire knowledge about mobile entities that move freely in the environment (e.g., other agents).

This study is situated within the framework of artificial constructivist learning [e.g., 1] and enactive learning [e.g., 2]. In this framework, the learning occurs through the enaction of control loops that implement Piagetian *sensorimotor schemes* [3], which we call *interactions*. This framework also relates to the notion of *intrinsic motivation* of artificial agents for *developmental learning* [e.g., 4].

More precisely, we investigate a modeling hypothesis called *Radical Interactionism* (RI) [5] and *artificial interactionism* [6]. We implement a kind of self motivation called *interactional motivation* [7]. The agent starts with a predefined set of uninterpreted interactions associated with predefined numerical *valences*, and seeks to enact interactions of positive valence and to avoid interactions of negative valence. Overall, the learning is *unsupervised*. There are no human-defined

labels attached with perceptions, actions, or categories of entities, not even spatial localization of percepts.

Through the learning of interaction with other mobile entities, we pursue the long-term goal of allowing the emergence of social behaviors within groups of artificial agents [e.g., 8], for example collective hunting. Generating such behaviors requires overcoming two main problems:

- 1) learning to define, recognize and localize other agents that make self-motivated movements in the environment,
- 2) inferring the intentions of other agents based on their own environmental contexts.

This paper focuses on the resolution of the first problem. It is subdivided as follows: Section II summarizes and formalizes the Radical Interactionism model, Section III presents a model for defining and recognizing probabilistic affordances, and Section IV presents a mechanism to recognize and localize distant probabilistic affordances in space. Finally, Section V encompasses some conclusive remarks and future development of the intersubjectivity problem.

II. THE RADICAL INTERACTIONIST (RI) HYPOTHESIS

In contrast with most machine learning approaches, an RI agent cannot directly access the state of its environment: its input data is outcome of control loops rather than percepts of the environment’s state. The agent learns and exploits regularities in sequences of control loops offered by its coupling with its environment. The learning mechanism differs from reinforcement learning (e.g., as it is typically implemented in a Partially Observable Markov Decision Process) by the fact that RI agents have no reward defined as a function of the system’s state. Our goal is not to design agents that reach predefined goals or maximize a reward value, but to study the open-ended learning of emergent models of the environment, and to generate social behaviors.

Let I be the set of predefined primitive interactions (control loops). At the beginning of step t , the agent selects an *intended interaction* $i_t \in I$. An example consists in moving forward for a predefined duration. At the end of step t , the agent receives the *enacted interaction* $e_t \in I$ that was actually enacted. If $i_t = e_t$ then the enaction is a success. The agent *did* move forward. Otherwise, the enaction of i_t is a failure. For example, the agent actually enacted another interaction consisting in bumping into an obstacle, which may have a negative valence.

An RI agent learns to anticipate the results of its interactions, and tries to enact interactions of high valence.

To help the agent discover that its environment has a spatial structure (2D in our experiments), we designed the Parallel-RI (PRI) model [9] which allows the simultaneous enaction of multiple interactions. The PRI considers additional stimuli that cannot be separated from the movement that produced them. The optical flow is an example of such stimuli, that must be associated with a movement to characterize a position in space. Thus, the PRI model considers *primary interactions* as control loops (action,result), and *secondary interactions* as a couple (interaction, stimulus). At the end of step t , the agent receives, not one, but a set of k enacted interactions $E_t = \{e_1, \dots, e_k\} \subset I$, containing a primary interaction and a set of secondary interactions associated with this primary interaction.

Previous PRI experiments showed that the agent was able to identify and localize static affordances (possibilities of enacting an interaction [10]), and to store and keep track of them in an emergent structure, called Space Memory [9]. This memory generates an implicit context of affordances that the agent can exploit to generate behaviors in accordance with its interactional motivation. A subsequent model also showed the possibility to identify objects that moved in a straight line, by considering sequences of interactions [11]. These models, however, could not integrate entities moving irregularly (other agents). The present study addresses this limitation.

III. INTEGRATING MOBILE AFFORDANCES

This section explains the *signature mechanism* [9] by which the agent estimates the possibility of enacting interactions in a given context. This mechanism is based on the assumption that the enaction result of an interaction i depends on a limited context of elements in the environment, defining the *affordance* of i . As a PRI agent can only perceive its environment through enacted interactions, we define the signature S_i of an interaction i as an emerging structure characterizing one or several sets of interactions (i.e. $\{j_k\} \subset E_t$) whose enaction can characterize the presence of an element affording i for next step $t + 1$.

Defining objects by learning the affordances they provide is abundant in literature, e.g., [12]–[14]. Most of these approaches define affordances from perception, which limits the detection to next action, or requires prior knowledge on environment and space [e.g., 15] to detect distant affordances. Signatures cope with this limitation by using interactions instead of perception, allowing to exploit spatial properties implicitly encoded in interactions to detect distant affordances.

Formally, a *signature of interaction* is a function $S_i : \mathcal{P}(I) \rightarrow [-1; 1]$, where $\mathcal{P}(I)$ is 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 interaction i at step $t + 1$: 1 for certainty of success, -1 for certainty of failure. The agent adjusts the parameters of S_i each time i succeeds or fails.

Signatures must be *reversible* by defining a pseudo-reverse function $\hat{S}_i : \{-1; 1\} \rightarrow \mathcal{P}(\mathcal{P}(I))$ such that $\hat{S}_i(1)$ gives the set of minimal contexts C_i^1 in which i is possible, i.e. contexts

that *afford* i , and $\hat{S}_i(-1)$ gives the set of minimal contexts in which the enaction of i is impossible.

However, when interacting with mobile entities, the presence of an affordance at the end of step t does not guaranty that the interaction can be enacted during step $t+1$ because the entity may move in the meantime. We thus must separate the estimation of the presence of an affordance from the prediction of success of enacting the interaction.

A. Separating affordances from prediction of success

From an observer's perspective, three types of situations can happen when the agent tries interacting with a mobile entity:

- 1) The affordance is present at the right place, and the agent enacts the interaction successfully (e.g., a prey is in front of the agent, and the agent catches the prey).
- 2) The affordance is present, but the interaction fails (e.g., a prey is in front of the agent, but the prey moves and the agent fails to catch it).
- 3) The affordance is absent, leading to a failure of the interaction (e.g., there is not prey in front of the agent, but the agent tries to catch one and fails).

From the agent's perspective, situations 2 and 3 cannot be distinguished, as they have the same result. Situations 2 distorts the learning of signatures, as situations 1 and 2 can occur in the same context E , causing the prediction $S_i(E)$ to remain negative even though the affordance is present.

However, our preliminary tests showed that, despite remaining negative, signature predictions are slightly higher in case of situations 1-2 than in situations 3. Indeed, situations 1 allow contexts of interactions designating the affordance to emerge, while remaining insufficient to predict with a positive value.

We thus use the average prediction in case of failure \bar{S}_i^f as a threshold to distinguish between situations 2 and 3. When the interaction fails in an assumed situation 2 (i.e., $S_i(E_t) > \bar{S}_i^f$), the signature is not reinforced, which limits the influence of situations 2 in signature construction.

Symmetrically, some interactions may fail due to the presence an entity in a specific location (e.g., trying to move forward will succeeds unless an obstacle appears). As success of such interactions are expected to be more frequent than failures, the average of predictions will converge to a positive value. In this case, the average of positive predictions \bar{S}_i^s is used as a threshold to prevent the signature reinforcement in case of success in an assumed situation 2 (i.e. $S_i(E_t) < \bar{S}_i^s$).

It is then possible to define the ratio of success when the prediction $S_i(E_t) > \bar{S}_i^f$ (or failure when $S_i(E_t) < \bar{S}_i^s$), implying that the agent is in a situation of type 1 or 2. A ratio $p_{C_i^1}^i$ is thus defined for each context $C_i^1 \in \hat{S}_i(1)$ measuring the probability that the interaction will succeed in the presence of an affordance containing a mobile entity.

B. Implementation of Signatures

We extend Gay *et al*'s signature architecture [9] using multiple neurons as illustrated in Fig. 1. Each signature S_i consists of m neurons N_i^1 to N_i^m ; the neuron with the

strongest output defines the prediction of S_i . In case of success, the neuron with the strongest output is reinforced, while a failure reinforces all neurons. This competition leads to a specialization of each neuron for a specific context, while they are desensitized from other contexts. Thus, with a sufficient number of neurons, a signature can identify contexts affording its interaction independently.

Formally, a neuron N_i^n is defined as a set of weights $\{w_k^n\}$, with $\text{Card}(\{w_k^n\}) = \text{Card}(I)$, and an output defined as:

$$N_i^n(E_t) = f\left(\sum_k E_t[k] \times w_k^n\right), f(x) = \frac{1}{1 + e^{-x}} \quad (1)$$

where $E_t[k] = 1$ when $i_k \in E_t$ and $E_t[k] = 0$ otherwise.

Then, the response of the group is defined as the maximum output, and remapped to a range in $[-1; 1]$:

$$N_i(E_t) = \max_n(N_i^n(E_t)) \times 2 + 1 \quad (2)$$

In order to consider interactions that are afforded by the absence of an entity instead of its presence, we added an output weight W_i defining the output of the signature:

$$S_i(E_t) = N_i(E_t) \times W_i \quad (3)$$

The weight W_i is restrained in the interval $[-1, 1]$, allowing to inverse the result of the prediction, which makes neurons able to integrate contexts preventing the enaction of i .

The learning process uses a classical gradient descent¹ and prediction values obtained with E_{t-1} .

A consequence of this implementation is that high weights of neurons characterize contexts affording i . Weights of neurons can be grouped by primary interaction, each group containing a weight related to a primary interaction and weights related to its associated secondary interaction. Thus, a signature S_i can be subdivided into minimal contexts $C_{j,n}^i$, associated with a primary interaction j and a neuron n . It is then possible to define the ratio of success $p_{j,n}^i$ of each context $C_{j,n}^i$ by updating it when $j \in E_{t-1}$ and neuron n has the highest activity, and $S_i(E_{t-1}) < \overline{S}_i^s$.




Fig. 1. Signature architecture based on multiple neurons. As the signature output relies on the neuron with the greatest output, neurons are in competitions, leading to a specialization of each neuron for a specific context. The last function remaps the output to range $[-1; 1]$ and uses a global weight W_i , that can inverse the signature result, allowing the representation of contexts preventing the enaction of an interaction.

¹The enaction result of i_t is defined as $r_t = 1$ in case of success ($i_t \in E_t$) and $r_t = -1$ in case of failure ($i_t \notin E_t$). The weight W_i is updated as follows: $W_i \leftarrow W_i + \Delta_i \cdot (\alpha \cdot N_i - \alpha/2)$, with $\Delta_i = r_t - S_i(E_{t-1})$ and α the learning rate. In case of success, only the neuron with the highest output is reinforced. In case of failure, all neurons are reinforced: $w_k^n \leftarrow w_k^n + \alpha \cdot \Delta_i \cdot N_i^n$, with $\Delta_i = (r_t \cdot W_i + 1)/2 - N_i^n$.

C. Test Environment

This signature mechanism was tested on an artificial agent moving in the 2-dimensional discrete environment shown in Fig. 2. The sensorimotor possibilities of the agent define the following five primary interactions:

- \triangleright move forward by one step,
- \blacktriangleright bump in a solid element,
- \blacktriangleleft eat something edible,
- \curvearrowright turn left by 90°
- \curvearrowleft turn right by 90°

Interactions *move forward*, *bump* and *eat* are considered as mutually alternative: intending one of these interactions may lead to the enaction of one of the two others instead.

We add a set of secondary interactions provided by the agent's visual system, that can detect colors and measure distances, with a field of view of 180° . Secondary interactions consist in *seeing* the displacement of a red, green or blue entity at a certain (but unknown) position in egocentric space, while enacting a primary interaction. Interaction *bump* does not generate visual interactions (no movement). We discretize the visual field as a grid of 15×8 positions in front of the agents that matches the environment's grid. We thus define $4 \times 3 \times 15 \times 8 = 1440$ secondary interactions. Signatures are implemented using sets of $m = 6$ neurons. The signature learning process is driven by a learning mechanism that foster interactions with low certainty of success or failure (low $|S_i(E_t)|$).

The environment contains three types of objects offering spatial regularities that the agent can discover by interacting with them, and characterized by a color that makes them recognizable through its sensorimotor system: 1) wall (green), affording bump, 2) algae (red), that are walkthroughable (and thus useless in the agent's perspective), and 3) fish (blue), affording eat. The fish move randomly: 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 an obstacle (wall, alga or other fish), it remains immobile, making the immobile situation slightly more probable than other directions. This random movement simulates agents with unknown behavior.




Fig. 2. Test environment. The agent is represented as a grey shark (bottom left), wall as green blocks, algae as red leaves and mobile preys as blue fishes. At each simulation step, the fish has 20% of remaining at current position or to move up, down, left, or right.

We then let the agent behave in its environment, driven by the signature learning mechanism. Signature of *bump* (Fig. 3) emerges and stabilizes within 5000 simulation steps.

The signature is similar to signatures obtained in previous environments [9][11]. It associates the success of *bump* with the presence of *seeing a green element moving right in front of the agent*, and of a previously enacted *bump*. The signature thus gathered every interaction allowing to detect the presence of a wall in front of the agent, even through they come from different sensory modalities.




Fig. 3. Signatures of interaction *bump* after 100 000 steps. It is characterized by the weights of 6 formal neurons, each neuron being represented by a column. As the signature identified a unique context, we only represent weights of one neuron. As an external observer, knowing that the environment is 2D, we organize weights of a neuron to make signatures more readable: first, weights associated with primary interactions are represented with five squares below (green for a positive weight, red for a negative weight). Weights associated with secondary interactions are grouped according to their primary interactions, forming the four groups (from top to bottom: forward, eat, turn left, turn right ; bump does not produce visual interactions). Each group is organized to place visual interactions with their associated position in space, relative to the agent (orange triangle). Colors associated with visual interactions are overlapped to generate signatures under the form of a RGB image. Signature of bump identified a context that consist of seeing a green element in front of the agent, which corresponds to the presence of a wall in front of the agent. Bump is also related to the success of previous bumps, since the agent can bump repeatedly.

Signatures of secondary interactions related to static elements (*seeing red* and *seeing green*) progressively stabilize, depending on their frequency of occurrence. After 50 000 simulation steps, most of these signatures stabilized. These signatures are also similar to signatures obtained with static entities [9][11]. They designate elements of the same color but on a different position in space. From an external point of view, the spatial offset between the visual interaction and the entity designated by its signature matches the movement performed by the enaction of its associated primary interaction. This property is used for distant affordance detection [9] (details in Section IV).

Signatures of interactions related to mobile elements require more steps, as they relate to a larger variety of contexts to identify: at least 45 000 steps are required to identify contexts affording *move forward* and *eat*. Signature of interaction *eat* (Fig. 5) characterizes five contexts corresponding to the five positions of fish that can lead to a success of eat. Note that the position under the agent does not appear in contexts of interactions associated to move forward, as this situation is not possible. The signature of *move forward* (Fig. 4) has a negative weight W . The signature thus shows the affordance that prevents this interaction. The signature designates five contexts associated with the presence of a fish, and one context

associated with the presence of a wall in front of the agent.

Signatures of secondary interactions consisting in seeing blue (Fig. 6) elements designate five contexts, corresponding to the five positions leading to a success of these interactions, with an offset (from an external point of view) corresponding to the movement of the associated primary interaction.




Fig. 4. Signature of *move forward*, recorded after 100 000 simulation steps. Each column represents a neuron of the signature. The weight W is negative: the signature thus represents contexts *preventing* moving forward. The signature identifies six contexts, represented (from an external point of view) above. As a fish cannot be below the agent after forward or eat, only 5 contexts are related to forward primary interaction (greyed context has low weights and is thus unused by the signature). As eat interaction is rarely enacted, contexts related to this primary interaction (second line) are still constructing.

We also analyze the ratio of successful enaction after a prediction of success. The ratios obtained in contexts implying static objects (such as walls) are close to 1 indicating that the presence of this type of affordance implies the success of the interaction. Ratios obtained with mobile fish are close to 20%, which correspond to the probability that the fish moves in the right direction when the agent tries to eat it. The contexts with a fish in front of the agent is however slightly greater. This can be explained by the fact a prey cannot move to a different position when blocked by a wall or an alga, increasing the probability of eating the fish when in front of the agent. These ratios, summarized in table I, show that signatures can integrate and encode stochastic properties of the environment.

TABLE I
AVERAGE RATIOS OF SUCCESSFUL PREDICTION IN CONTEXTS: WALL, FISH IN FRONT, AND FISH IN OTHER POSITIONS (SURROUNDING)

interaction	wall	front	surround.
forward	0.96	0.25	0.19
bump	0.97	/	/
eat	/	0.24	0.19
seeing blue (Fig. 6)	/	0.24	0.18

IV. LOCALIZING DISTANT AFFORDANCES

The detection of distant affordances relies on a property of signatures: a signature of an interaction designates an affordance as sets of interactions $\{j_k\} \in \hat{S}_i(1)$ allowing to detect the presence of this affordance. However, each interaction j_k can have its own signature. It is thus possible to define, from

is completed (or until $\Theta_{\xi_k^i}$ is empty). Interaction i is then predicted using $S_i(E^l)$. If the signature predicts a success, the affordance of i is confirmed at position $\sigma_{\xi_k^i}$.

C. Test Environments

The affordance detection mechanism was tested with signatures recorded after 200 000 simulation steps. The projection mechanism generates projection sequences with a maximum length up to 7 interactions.

The projection sequence construction mechanism was adapted for a signature implementation based on neurons. First, we only project interactions designated by a signature with a weight with an absolute value that is greater than a threshold, eliminating non-significant weights. Then, we added a new property to projection sequences, the global weight, characterizing the pertinence of the sequence to represent the affordance. This global weight is computed as follows: first sequences have a global weight defined as $w^{global} = W_S \times w_k$. Then each backmove through a weight w_k of a signature S , the global weight of the new sequence is updated as $w^{global} = w^{global} \times W_S \times w_k$. The filter mechanism then compare values $p \times w^{global}$ instead of p alone, offering a good compromise between probability and pertinence of sequences.

The agent is presented to different environment configurations. An enaction cycle is performed to let the agent perceive its environment, and sequences of detected affordances are analyzed. Fig. 7 shows the detection in a context containing two wall blocks, two fish and an alga. Sequences localizing static objects (walls) allows moving toward them. Sequences localizing fish do not reach the position of the prey, but a position just next to it. Indeed, as the agent can eat a fish on its side, the resulting sequence is a compromise between probability and length of the sequence. The alga is ignored since it has the same property as an empty space. Thus, the affordance detection mechanism can still detect and localize distant affordances under the form of sequences of interactions, which can be stored and exploited by the space memory similar as in a static environment.




Fig. 7. Distant affordances are detected and localized through sequences of interactions allowing to reach them. Circles show the position and orientation associated with these sequences (red: not moving forward, blue: eat, green: bump). We can notice that the agent ignores the alga since it has the same interactional property as an empty space.

V. CONCLUSION AND FUTURE WORK

This work proposes a new mechanism to enable an interactionist agent to consider mobile and non-predictable entities in its emergent model of the environment. Results obtained in a simulated environment showed that the agent can still localize

distant affordances without the notion of space, and define their positions in a similar way than in static environment, allowing the use of a Space Memory [9].

Future work will study how the space memory can be used to better interact with other agents, and how the intrinsically motivated decisional mechanism can integrate probabilities of presence of affordance into consideration. In multi-agent contexts, this will allow us to study the mutual integration of agents in each other's environmental model, and how these models can be exploited for generating behaviors involving collaborative tasks, such as coordinated hunt of large preys.

We will also study how an agent can predict another agent's intentions through the observation of the other agent's context, as a previous implementation of the space memory demonstrated the possibility of reference change, opening intersubjectivity possibilities between agents.

REFERENCES

- [1] K. R. Thórisson, "A new constructivist ai: From manual methods to self-constructive systems," in *Theoretical Foundations of Artificial General Intelligence*, P. Wang and B. Goertzel, Eds., vol. 4, pp. 145–171. Atlantis Press, 2012, Series Title: Atlantis Thinking Machines.
- [2] T. Froese and T. Ziemke, "Enactive artificial intelligence: Investigating the systemic organization of life and mind," *Artificial Intelligence*, vol. 173, no. 3–4, pp. 466–500, 2009.
- [3] J. Piaget, *The construction of reality in the child*, New York: Basic Books, 1954.
- [4] P.-Y. Oudeyer, F. Kaplan, and V. Hafner, "Intrinsic motivation systems for autonomous mental development," *IEEE transactions on evolutionary computation*, vol. 11, no. 2, pp. 265–286, 2007.
- [5] O. L. Georgeon and D. Aha, "The radical interactionism conceptual commitment," *Journal of Artificial General Intelligence*, vol. 4, pp. 31–36, 12 2013.
- [6] M. Guillermin and O. L. Georgeon, "Artificial interactionism: Avoiding isolating perception from cognition in ai," *Frontiers in Artificial Intelligence*, vol. 5, 2022.
- [7] O. L. Georgeon, J. B. Marshall, and S. Gay, "Interactional motivation in artificial systems: Between extrinsic and intrinsic motivation," in *IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, 2012, pp. 1–2, IEEE.
- [8] A. Netanyahu, T. Shu, B. Katz, A. Barbu, and J. Tenenbaum, "Phase: Physically-grounded abstract social events for machine social perception," in *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [9] S. L. Gay, A. Mille, O. L. Georgeon, and A. Dutech, "Autonomous construction and exploitation of a spatial memory by a self-motivated agent," *Cognitive Systems Research*, vol. 41, pp. 1–35, 08 2016.
- [10] J. J. Gibson, "The theory of affordances," *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*, Hilldale, USA, vol. 1, no. 2, pp. 67–82, 1977.
- [11] S. L. Gay and S. Hassas, "Autonomous object modeling based on affordances in a dynamic environment," in *Conference on Biologically Inspired Cognitive Architectures*, 2015, vol. 71, pp. 150–156.
- [12] A. Paola, P. Èric, L. S. Katrin, R. Subramanian, and P. P. A. Ronald, "Building affordance relations for robotic agents - a review," in *International Joint Conference on Artificial Intelligence*, Zhi-Hua Zhou, Ed. 8 2021, pp. 4302–4311, International Joint Conferences on Artificial Intelligence Organization, Survey Track.
- [13] J. Baleia, P. Santana, and J. Barata, "On exploiting haptic cues for self-supervised learning of depth-based robot navigation affordances," *Journal of Intelligent and Robotic Systems*, vol. 80, 02 2015.
- [14] E. Uğur and E. Şahin, "Traversability: A case study for learning and perceiving affordances in robots," *Adaptive Behavior*, vol. 18, no. 3–4, pp. 258–284, 2010.
- [15] A. Manoury, S. M. Nguyen, and C. Buche, "Hierarchical affordance discovery using intrinsic motivation," in *International Conference on Human-Agent Interaction*, New York, NY, USA, 2019, HAI '19, pp. 186–193, Association for Computing Machinery.