



# Perspective Taking in Deep Reinforcement Learning Agents

Aqeel Labash<sup>1</sup>, Jaan Aru<sup>1,2\*</sup>, Tambet Matiisen<sup>1</sup>, Ardi Tampuu<sup>1</sup> and Raul Vicente<sup>1\*</sup>

<sup>1</sup> Computational Neuroscience Lab, Institute of Computer Science, University of Tartu, Tartu, Estonia, <sup>2</sup> Institute of Biology, Humboldt University of Berlin, Berlin, Germany

Perspective taking is the ability to take into account what the other agent knows. This skill is not unique to humans as it is also displayed by other animals like chimpanzees. It is an essential ability for social interactions, including efficient cooperation, competition, and communication. Here we present our progress toward building artificial agents with such abilities. We implemented a perspective taking task inspired by experiments done with chimpanzees. We show that agents controlled by artificial neural networks can learn via reinforcement learning to pass simple tests that require some aspects of perspective taking capabilities. We studied whether this ability is more readily learned by agents with information encoded in allocentric or egocentric form for both their visual perception and motor actions. We believe that, in the long run, building artificial agents with perspective taking ability can help us develop artificial intelligence that is more human-like and easier to communicate with.

**Keywords:** deep reinforcement learning, theory of mind, perspective taking, multi-agent, artificial intelligence

## OPEN ACCESS

### Edited by:

Adam Henry Marblestone,  
Harvard University, United States

### Reviewed by:

Joscha Bach,  
Independent Researcher,  
San Francisco, United States  
Jan Lauwereyns,  
Kyushu University, Japan

### \*Correspondence:

Jaan Aru  
jaan.aru@gmail.com  
Raul Vicente  
raulvicente@gmail.com

**Received:** 08 April 2020

**Accepted:** 16 June 2020

**Published:** 23 July 2020

### Citation:

Labash A, Aru J, Matiisen T,  
Tampuu A and Vicente R (2020)  
Perspective Taking in Deep  
Reinforcement Learning Agents.  
*Front. Comput. Neurosci.* 14:69.  
doi: 10.3389/fncom.2020.00069

## 1. INTRODUCTION

Many decisions we take depend on others, what they think, what they believe, and what we know about what they know. This ability to understand and infer the mental states of others is called Theory of Mind (Premack and Woodruff, 1978) or mindreading (Apperly, 2011). Not only humans have the ability to take into consideration what others think and believe. In controlled experiments it has been shown that chimpanzees can know what other conspecifics see and know (Hare et al., 2000). Here we ask whether artificial intelligence (AI) agents controlled by neural networks (Goodfellow et al., 2016) could also learn to infer what other agents perceive and know.

Theory of Mind is an important topic to study in AI mainly because successful human-machine interaction might critically depend on it. In particular, according to prominent theories, understanding the intentions of others is necessary for the emergence of meaningful communication and language (Tomasello, 2010, 2019; Scott-Phillips, 2014; Mercier and Sperber, 2017). According to these views, the basis of communication is not the ability to decode the message from the other, but rather the ability to understand that the other is *trying to communicate* in the first place. Humans excel at language and communication because we have a bias to assume that others are trying to “get something across” to us (Tomasello, 2010, 2019; Scott-Phillips, 2014). Unfortunately, we do not know how this bias is implemented in the human brain and hence we do not know how to implement it in AI. Furthermore, developing agents that are at the level of humans in Theory of Mind will be much more complicated than in computer or board games, as there are no established tasks or benchmarks to train the agents on.

Under these circumstances one way to proceed is to study how much of the Theory of Mind abilities could be learned through reinforcement learning (RL) in simple tasks. This approach

would demonstrate the limits of simple model-free algorithms for understanding Theory of Mind and thus build the basis for more comprehensive approaches. RL is a branch of AI that allows an agent to learn by trial and error while interacting with the environment. In particular, the agent must learn to select the best action in each specific state to maximize its cumulative future reward (Sutton and Barto, 2018). The agent could be for example an autonomous robot (Levine et al., 2016) or a character in a video game (Mnih et al., 2015). The idea behind learning by interacting with an environment is inspired from how human and animal infants learn from the rich cause-effect or action-consequence structure of the world (Thorndike, 1911; Schultz et al., 1997; Sutton and Barto, 2018). Therefore, RL is a biologically plausible mechanism for learning certain associations and behaviors and it can be used to study Theory of Mind.

Theory of Mind was studied by Rabinowitz and colleagues who modeled agents' behavior in a grid world (Rabinowitz et al., 2018). The proposed neural network was trained using meta-learning; in a first stage the network was trained to learn some priors for the behavior of different types of agents to subsequently speed up the learning of a model of a specific agent from a small number of behavioral observations. Their approach was a first step to induce theory of mind faculties in AI agents that were indeed able to pass some relevant tests for Theory of Mind skills. However, as the authors themselves note, their approach was limited in several important aspects that require future work. To name a few, the observer agent learning to model the behavior of others was trained in a supervised manner, it had full observability of the environment and of the other agents, and it was not itself a behaving agent.

In the current work we are interested in the emergence of certain aspects of Theory of Mind in behaving agents trained via RL with partial observability. We believe that these are more plausible conditions to model how humans and other animals might develop these abilities. In particular, we test here the ability of agents trained via RL to acquire one essential part of Theory of Mind: *perspective taking* (Apperly, 2011).

## 1.1. Perspective Taking

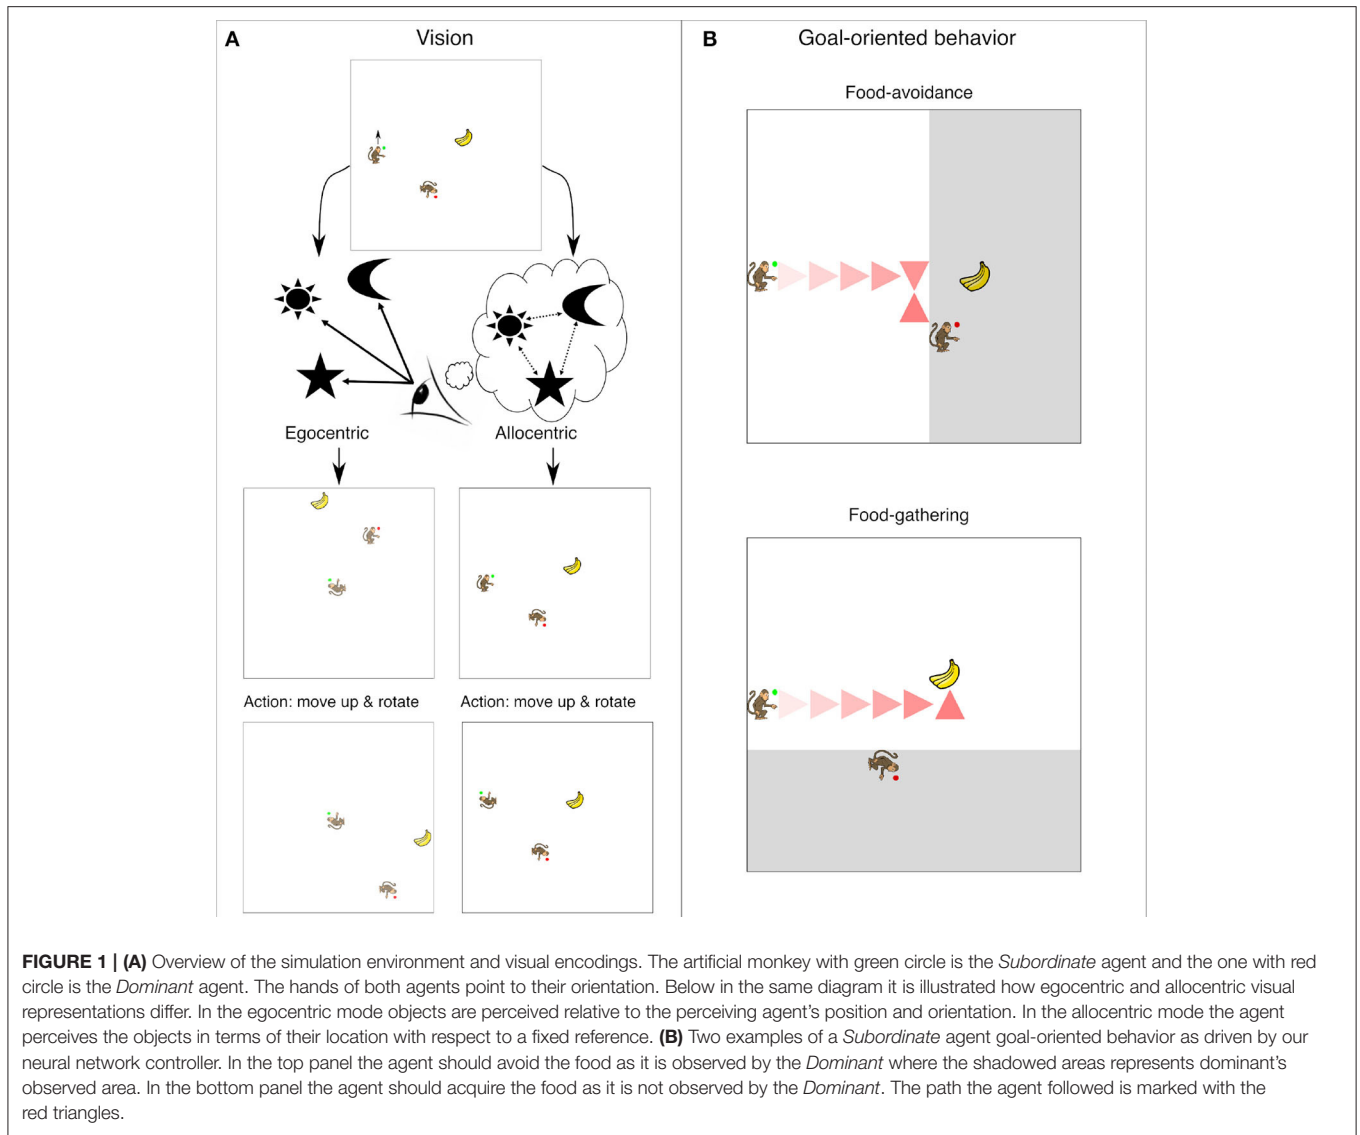
Perspective taking is the ability to consider the circumstances from a perspective that differs from our own. It is a very general ability that is relevant in many different tasks involving communication, memory, language and perception (Ryskin et al., 2015). It could be defined as “the cognitive capacity to consider the world from another individual's viewpoint” (Davis, 1983). It is one of the social competencies that underlies social understanding in many contexts (Galinsky et al., 2008; Apperly, 2011). While perspective taking is a very general ability, in the basic form it boils down to knowing what the other agent does or does not know (Cole and Millett, 2019). The existence of such simple forms of perspective taking is supported by the fact that this ability is not unique to humans and has been observed in other animals like chimpanzees (Hare et al., 2000; Tomasello et al., 2003).

Chimpanzee social status is organized hierarchically (dominant, subordinate) (Goldberg and Wrangham, 1997),

which is at full display during food gathering: when there is food available that both can reach, the dominant animal almost always obtains it. But what happens if the dominant could potentially reach the food placed behind an obstacle, but does not know that food is there? Can the subordinate take advantage of this? In a series of experiments (Hare et al., 2000) two chimpanzees were set into two separate cages facing each other with food positioned between them. The researchers manipulated what the dominant and the subordinate apes could see. For example in one condition, one piece of food could not be seen by the dominant chimpanzee. The results demonstrated that the subordinate animal exploited this favorable situation and indeed obtained more food in this condition. Hence, it was able to consider what the dominant chimpanzee could and could not see, i.e., take the perspective of the dominant chimpanzee into account (Hare et al., 2000; Tomasello et al., 2003). This work done with chimpanzees was the inspiration for our study.

The aim of the present work is to study whether an AI agent controlled by a neural network can learn to solve a similar perspective taking task using RL. We chose this task because it is relatively simple, while allowing us to study perspective taking with RL. We do not claim that RL captures all aspects of perspective taking or is the exact model of how perspective taking is learned in biological organisms (Aru and Vicente, 2018). We even do not claim that RL allows us to understand or model how chimpanzees solve this particular task. We simply use this task as to probe whether some simple aspects of perspective taking can be learned by RL agents. Understanding the capabilities and limitations of RL in acquiring perspective taking skills will lead to a better algorithmic understanding of the computational steps required for perspective taking in biological organisms.

In this study we are not so much interested in whether the perspective taking task is learnable *per se*. Rather, we seek to compare the speed with which different representations allow the learning of this simple perspective taking task. We are interested in a specific question about perspective taking: is it simpler to learn perspective taking with allocentric or egocentric representations of the environment? With allocentric input the position of other objects and agents is presented in relation to each other independently of the position of the perceiving agent. With egocentric input the position of all objects and other agents is given with respect to the position of the perceiving agent. This means that for example when the agent changes its orientation the whole world will rotate. See **Figure 1** for an illustration of the two encodings of visual input. From neuroscience and behavioral experiments it is known that although animals perceive the world from the egocentric viewpoint, this information is transformed to allocentric code in structures like the hippocampus (Burgess et al., 2001; Wilber et al., 2014; Chersi and Burgess, 2015; Wang et al., 2020). Presumably the fact that this transformation is computed in the brain hints that the allocentric view enables some functions that cannot be achieved through egocentric representation alone (Burgess et al., 2001; Chersi and Burgess, 2015). It is possible that perspective taking is one of these functions. Intuitively it seems that taking the perspective of the other agent demands ignoring own sensory input and taking into account the relations between the other agent and the objects



in the environment, which could be supported by allocentric representation. Hence, another goal of our study is to test the generality of these assumptions and intuitions using minimal models of computational learning agents.

## 2. METHODS

### 2.1. Task and Environment

For the perspective taking task, we generated a grid world environment where each element can spawn randomly within specific regions. The elements considered included two agents (a *Dominant* and a *Subordinate*), and a single food item (reward). In the present experiments only the *Subordinate* agent is controlled by a RL algorithm and can execute actions in the environment by moving around and changing its orientation. The *Dominant* agent is not controlled by any learning algorithm but its role is critical. The value of the reward obtained by the *Subordinate* at reaching the food depends on whether the food is visible

**TABLE 1 |** List of the events rewarded in the experiments and their respective values.

Event	Reward value
Eating food observed by <i>Dominant</i>	-1000
Eating food not observed by <i>Dominant</i>	+1000
Every time step	-0.1

from the *Dominant's* point of view. If food is retrieved by the *Subordinate* when observed by *Dominant* the value of the food item becomes negative (to mimic the likely punishment received from the dominant in the nature). If the food is obtained while not observed by the *Dominant* the value of the reward is positive. See **Table 1** for the list the events rewarded and its correspondent values.

Experiments were conducted using environments created with Python toolbox Artificial Primate Environment Simulator

(APES) (Labash et al., 2018). The toolbox allows to simulate a 2D grid world in which multiple agents can move and interact with different items. Agents obtain information from their environment according to a visual perception model. Importantly, APES includes different visual perception algorithms that allows to calculate visual input based on agents' location, orientation, visual range, visual field angle, and obstacles. In particular, for this work, we simulate two modes for the agents' vision: egocentric and allocentric (for detailed descriptions see the subsection on visual encoding). For further specifics on the toolbox the reader can access the associated GitHub repository<sup>1</sup>.

In all the experiments we considered that both agents have a long range of vision but a limited visual field angle of 180 degrees. In our main scenario for our perspective taking task, the coverage of the food and *Dominant*'s location is distributed as shown in **Figure 2**. Both the *Dominant* and food item can spawn anywhere inside a  $5 \times 5$  area (see **Figure 2**). To successfully solve the task the *Subordinate* agent must learn to navigate to reach the food's location only when the food item is not within the field of vision of the *Dominant* agent. This implies that the *Subordinate* needs to simultaneously integrate three pieces of information to successfully determine whether the food item is observed by the *Dominant* or not: (1) the orientation of *Dominant*, (2) the position of the *Dominant*, and (3) the position of the food. Note that since the subordinate agent moves and rotates around the environment a direct perception of the *Dominant* agent and food is not always present. However, as explained below the agent is equipped with a short-term memory in the form of a LSTM layer that allows it to integrate temporal information (Hochreiter and Schmidhuber, 1997).

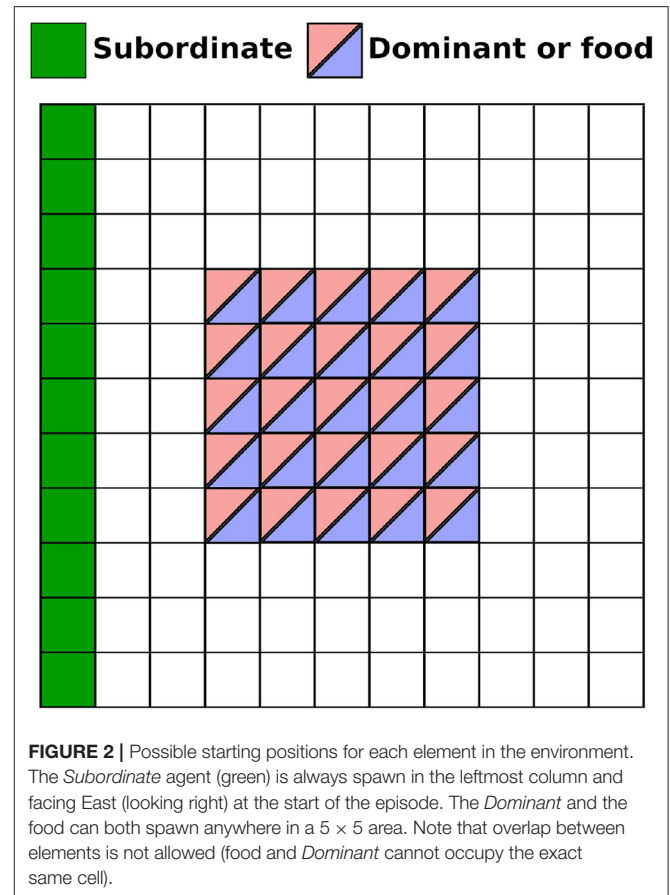
The dimensions of the grid world amounts to  $13 \times 13$  when using allocentric encoding of visual information, and  $11 \times 11$  when using the egocentric encoding. This compensation is needed to balance the fact that egocentric encoding needs a larger input space (since positions are relative to the agent's location and orientation, a  $n \times n$  grid world actually needs a  $n \times (2n - 1)$  input layer for an egocentric agent with 180 degrees vision). The dimensions of the grid worlds were chosen so that the neural network controllers for egocentric and allocentric encodings match in their number of parameters (weights). We also matched the average distance between the initial location of the *Subordinate* agent and the food item.

The number of possible combinations for initial configurations of the environment and agents exceeds 20,000. Upon movement of the *Subordinate* agent along the grid the number of possible states becomes  $> 1,000,000$ .

## 2.2. Model

### 2.2.1. Input

The input to the network controlling the *Subordinate* actions is a set of binary maps. They encode the different agents and other elements properties in the environment. The list of inputs to the network include:



**FIGURE 2** | Possible starting positions for each element in the environment. The *Subordinate* agent (green) is always spawned in the leftmost column and facing East (looking right) at the start of the episode. The *Dominant* and the food can both spawn anywhere in a  $5 \times 5$  area. Note that overlap between elements is not allowed (food and *Dominant* cannot occupy the exact same cell).

- **Spatial location of elements:**  $13 \times 13$  or  $11 \times 21$  binary one-hot map for each element represented by a 1 at the corresponding element position. In egocentric vision, the agent own location is not required. This is because the focal agent location does not change (in an egocentric framework the agent is always at the center of its own visual field).
- **Observability mask:**  $13 \times 13$  or  $11 \times 21$  binary mask which indicates the field of vision of the *Subordinate*. It helps to distinguish whether a cell in the grid world is empty or out of the field of vision.
- **Orientations:**  $1 \times 4$  binary one-hot vector for each agent with 1 at the corresponding agent orientation. In egocentric vision, only the orientation for *Dominant* agent is required since in egocentric the *Subordinate* is always looking forward. An important difference between allocentric and egocentric vision is orientation encoding. In allocentric vision the orientation is encoded as (North, South, East, West) in comparison to egocentric vision where orientation is in relation to the focal agent (toward the agent, same direction as the agent, to its left, to its right).

### 2.2.2. Visual Encoding: Allocentric vs. Egocentric

In this work we compare two types of visual perception. With allocentric input the locations and orientations of items in the environment are encoded in reference to a fixed system of

<sup>1</sup><https://github.com/aeel13932/APES>



coordinates (as if the vision is provided by a fixed camera with a top-down view). With egocentric input, the items are perceived from the eyes of the *Subordinate* agent, and hence they change in relation to the agent movements and rotations. In both allocentric and egocentric there is an observability binary array which represents the observed and non-observed areas.

In allocentric encoding, we feed to the network  $13 \times 13$  arrays that represent positional information of items in the environment in addition to  $1 \times 4$  array for each agent in order to encode their orientation. In this mode, when the *Subordinate* changes its orientation and moves, four bits will change corresponding to its previous location, current location, previous orientation, and current orientation. We note that the observability and the *Dominant* agent orientation array might also change in case the subordinate movement leads to some item switch from observed to not observed (or viceversa).

In egocentric encoding, *Subordinate's* position and orientation remain fixed despite the agent's movements or rotations. We humans, similarly to other animals, when we turn left or right we still look forward and in the same position from our perspective. Hence, in egocentric encoding the network is not fed *Subordinate's* orientation, but still it is fed the relative *Dominant's* orientation. Thus, *Dominant's* orientation input is based on where it looks in relation to the *Subordinate's* (toward the *Subordinate*, same direction as the *Subordinate*, looking to its left or right). In the egocentric condition the input arrays that represent the environment have dimension  $11 \times 21$ . Although the base environment is  $11 \times 11$ , the input layer is augmented in the egocentric encoding to accommodate that the agent (with 180 degrees of vision) should always perceive the environment from in reference to its own centered view and of the same size regardless of its position. For example, when the agent is located at bottom right corner of the environment and looking North (forward from its perspective) it should have the left of its visual field encoding the  $11 \times 11$  environment. However, when sitting on the bottom left corner and facing the same orientation, now the environment should be displayed as its right visual field. Hence, the input layer is augmented to  $11 \times 21$  to accommodate a common range of centered vision regardless of the agent's location.

### 2.2.3. Action Encoding: Allocentric vs. Egocentric

The action space of the network controller (which architecture we explain next) depends on the framework for motor output simulated for the agent. In the allocentric encoding of motor output the action space is composed by (moving North, moving South, moving East, moving West, and no move). In the egocentric encoding of motor output the action space is (move forward, move backward, move right, move left, and no move). Note that each moving action is accompanied by a rotation so that the agent is always looking at the direction is heading. For example, if the agent moves North, it will also rotate to face North. This conforms to the fact that humans and most animals advance in the same direction they are facing. We also note that in most classical video games the combination of allocentric actions and allocentric vision is used.

### 2.2.4. Architecture

In our model we used a neural network to control the actions of the *Subordinate* agent. The architecture and hyperparameters are the same as in Jaques et al. (2019) with two important exceptions. First, additional inputs are fed to the network. Orientation of both agents are fed after the convolutional layer as shown in **Figure 3**. Note that in egocentric encoding only the *Dominant's* orientation was fed. Second, we used a dueling Q-network (Wang et al., 2016) instead of the advantage actor-critic model.

In a dueling Q-network the state-action value (Q-value) calculation is based on two separate estimates: the state value (how good the current state is) and advantages (which benefit is obtained from each action) as described by

$$Q(s, a; \theta) = V(s; \theta_V) + \left[ A(s, a; \theta_A) - \max_{a'} A(s, a'; \theta_A) \right], \quad (1)$$

where  $s$  is a state,  $a$  is an action,  $V$  is the state value,  $A$  is the advantage, and  $a'$  is the next action.  $\theta$  represents the network parameters, while  $\theta_V$  is the subset of parameters used in the value network and  $\theta_A$  is the subset of parameters used in the advantage network.

Using a dueling network architecture involves updating two network models: a training model (parameterized by  $\theta$ ) which weights are updated using gradient descent, and a target model (parameterized by  $\theta^-$ ) which weights are periodically  $\tau$ -averaged with training model's weights as described by

$$\theta^- = \theta\tau + \theta^-(1 - \tau). \quad (2)$$

The  $\epsilon$ -greedy policy  $\pi(s; \theta)$  chooses a random action with probability  $\epsilon$  and an action with maximum Q-value otherwise as in Equation (3):

$$\pi(s; \theta) = \begin{cases} \text{random action} & u < \epsilon, u \sim U(0, 1) \\ \operatorname{argmax}_a Q(s, a; \theta) & \text{otherwise} \end{cases} \quad (3)$$

To summarize, **Figure 3** illustrates the network architecture, its input layers and the output to control the actions of the *Subordinate* agent.

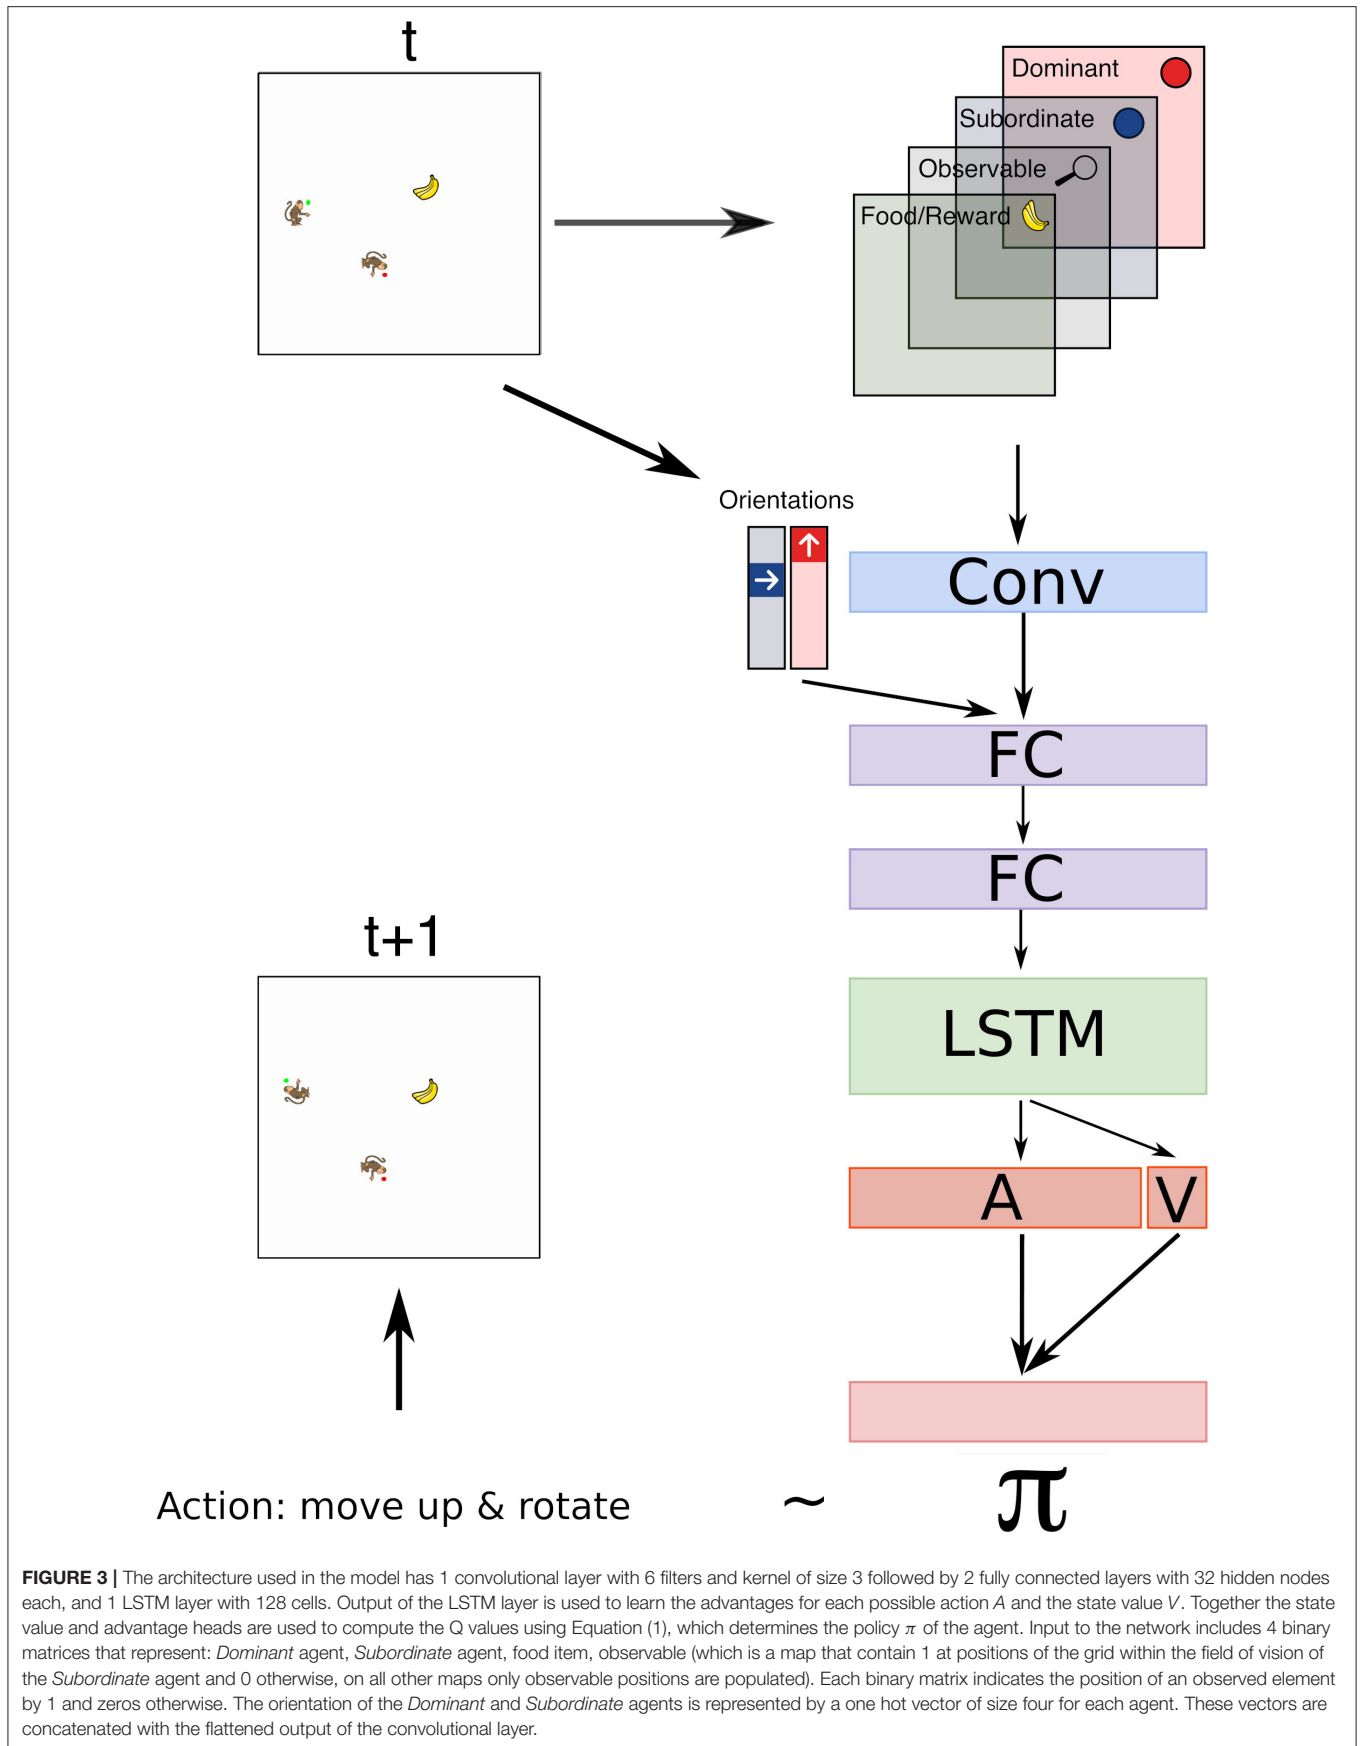
## 2.3. Training

### 2.3.1. Reinforcement Models

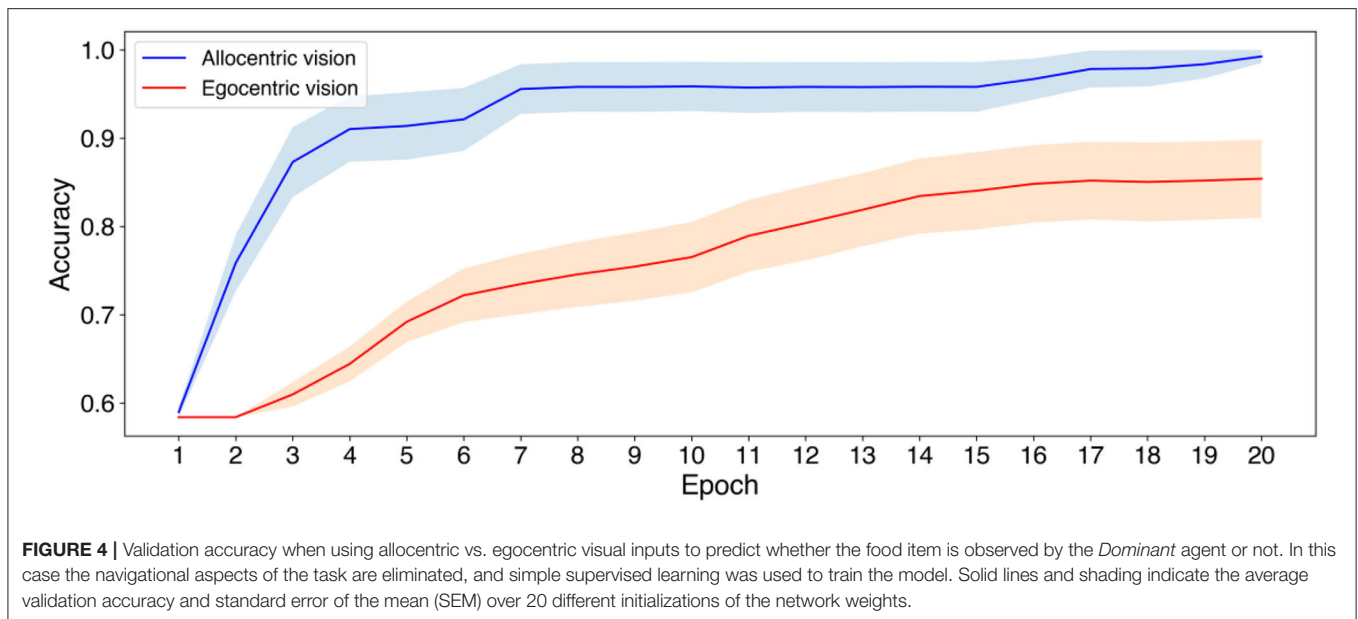
All RL models were trained for 20 million steps. An episode is terminated when the food is eaten by the *Subordinate* agent or the food is not eaten after 100 time steps.

We used replay memory to remove sequential correlations and smooth distribution changes during training as it is usually done in other studies (Hausknecht and Stone, 2015). The replay buffer size used is  $10^3$  trajectories. Maximum length of each trajectory is 100 time steps. Shorter trajectories were padded with zeros.

For the neural network implementation we used the Keras library (Chollet, 2015). We used Adam optimizer (batch size set to 16) with a fixed learning rate at 0.001 and annealed the exploration probability with a schedule from an initial value of 1 until reaching 0.1 at the 75% of the total number of steps. We clipped the gradient at 2 to prevent the gradient explosion problem.



**FIGURE 3 |** The architecture used in the model has 1 convolutional layer with 6 filters and kernel of size 3 followed by 2 fully connected layers with 32 hidden nodes each, and 1 LSTM layer with 128 cells. Output of the LSTM layer is used to learn the advantages for each possible action  $A$  and the state value  $V$ . Together the state value and advantage heads are used to compute the Q values using Equation (1), which determines the policy  $\pi$  of the agent. Input to the network includes 4 binary matrices that represent: *Dominant* agent, *Subordinate* agent, food item, observable (which is a map that contain 1 at positions of the grid within the field of vision of the *Subordinate* agent and 0 otherwise, on all other maps only observable positions are populated). Each binary matrix indicates the position of an observed element by 1 and zeros otherwise. The orientation of the *Dominant* and *Subordinate* agents is represented by a one hot vector of size four for each agent. These vectors are concatenated with the flattened output of the convolutional layer.



### 2.3.2. Supervised Models

For the supervised models we used the same network as described in **Figure 3** with the exception that we removed the LSTM layer since we trained the model to predict from the initial time step and hence no memory effects were necessary.

The learning signal was the ground truth label of whether the *Dominant* agent observes the food item within its field of vision for each environment initialization. Thus, the datasets consisted of all possible variations for the initial configurations of the environments (amounting to 26,400 samples for egocentric vision and 32,100 for allocentric vision). We used a 80/20 split for training and validation data. The training proceeded by minimizing the cross-entropy loss function using Adam optimizer (batch size was set to 64 while the learning rate was set to 0.001). To ensure the representativity of the models, all accuracies were averaged over 20 random initializations of the model weights.

## 3. RESULTS

Here we present the main experiments to test the ability of the agent to solve the present perspective taking task. In particular, we are interested in comparing how the visual and action encoding affects the readiness to learn the task by trial and error.

As described in section 2, both the *Dominant* agent and food item can spawn anywhere inside a  $5 \times 5$  area (see **Figure 2**), giving rise to a large number of possible combinations for their relative position and orientation. To solve the task the *Subordinate* agent must learn to navigate to reach the food's location only when the food item is not within the field of vision of the *Dominant* agent. This implies that the *Subordinate* needs to integrate three pieces of information to successfully determine whether the food item is observed by the *Dominant* or not: (1) the orientation of *Dominant*, (2) the position of the *Dominant*, and (3) the position

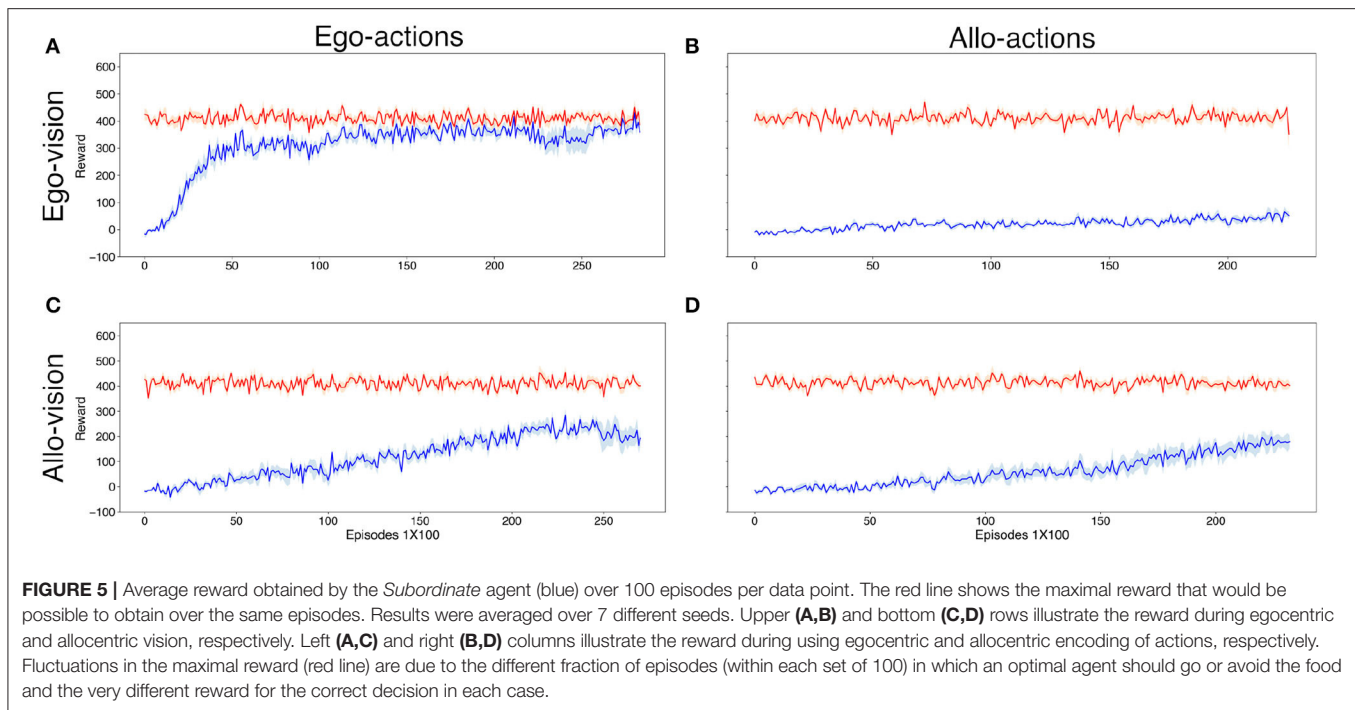
of the food. In addition, it needs to maintain in memory the result of that integration of information, and navigate successfully toward the food item.

Thus, solving the perspective taking task presumably involves both an estimation of whether the food item is being observed by the *Dominant* as well as memory and navigational aspects to reach the food item (or avoid it).

To compare both types of visual processing (egocentric vs. allocentric encoding) independently of memory and navigational aspects, first we trained the model to output a binary decision about whether the food item is visible to the *Dominant* in the first time step (initial configuration of each episode). The model architecture is the same as depicted in **Figure 3** (used also for the RL models) with the exception of the removal of the LSTM layer which is not needed since we aim to decode from the initial configuration of each episode. The model was trained by supervised learning with a cross-entropy loss function to predict whether food is visible from the *Dominant* agent's point of view.

**Figure 4** shows the accuracy at validation samples as a function of the number of training epochs. As observed in **Figure 4** when visual input is fed in allocentric coordinates the model exhibits a quicker learning of the decision of whether the food item is observed by the *Dominant* agent. In particular, with the allocentric mode of visual processing the decision is learned with high accuracy ( $> 90\%$ ) from 4 epochs of training. Similar level cannot be achieved by egocentric viewpoint in 20 epochs of training (reaching  $\sim 83\%$ ), an amount of training by which allocentric view already provides close to perfect decoding.

We then proceed to compare how an egocentric and allocentric encoding of visual input affects the performance of an agent in solving the perspective taking task when learning from trial and error. That is we trained our agent anew using only RL. During our RL experiments the *Subordinate* receives a reward of  $-0.1$  per time step regardless of the status of the food (seen or not seen by the dominant). The other reward is received only at



the end of the episode when the agent either reaches the food or the episode terminates after a given amount of time. We note that the full task now involves the need of not only integrating the necessary information to determine whether to approach the food, but also to maintain and convert this information into a motor plan. For this study we also considered the egocentric and allocentric representations for moving actions in addition to the egocentric and allocentric encodings for visual perception. Thus, we explore 4 combinations of visual and action representations ( $\text{ego}_{\text{vis}}\text{-ego}_{\text{act}}$ ,  $\text{allo}_{\text{vis}}\text{-ego}_{\text{act}}$ ,  $\text{ego}_{\text{vis}}\text{-allo}_{\text{act}}$ , and  $\text{allo}_{\text{vis}}\text{-allo}_{\text{act}}$ ).

In this case agents who had egocentric visual input and action encoding quickly learned the task, whereas agents with allocentric input were not able to efficiently learn the perspective taking task irrespective of the action encoding (see **Figure 5**).

Indeed, we can see from **Figure 5** that out of the four combinations, only the “ $\text{ego}_{\text{vis}}\text{-ego}_{\text{act}}$ ” condition performed well in the task. This is at first surprising, given that: (i) the previous result that without navigation (in the supervised learning setting) the allocentric agents are better in deciding whether the *Dominant* can see the food, and (ii) the agents with allocentric vision underperform the agents with egocentric vision even when using the same action encoding. Also, agents with egocentric vision cannot learn the task when using allocentric actions compared to the efficient and almost perfect score when using also egocentric actions.

Hence, part of the difficulty seems to come from the coupling of visual to navigational aspects, i.e., not only extracting the relevant information from the visual input but also its conversion into the appropriate actions. No single factor (visual encoding or action encoding) seems to individually explain the success at efficiently learning the task, rather what matters is their specific

combination. Therefore, we analyzed in more detail the required computations for the different cases to succeed in solving the task.

In the case of agents with allocentric actions (right column in **Figure 5**) the movements are in reference to fixed directions in the space, namely North, South, West, East, and standstill. This implies that when the *Subordinate* agent with egocentric vision but allocentric actions ( $\text{ego}_{\text{vis}}\text{-allo}_{\text{act}}$ ) sees food in front of it, it will not automatically know which action it needs to take to get closer to the food. In comparison to that, with egocentric actions ( $\text{ego}_{\text{vis}}\text{-ego}_{\text{act}}$ ) it is enough to move forward to approach the food. Similarly, simple heuristics also exist whenever the food item is within the field of vision of the agent. Thus, when working with egocentric actions, for agents with egocentric vision the selection of the optimal action is solely a function its location relative to the goal, an information which is directly accessible from its visual encoding. However, when working with allocentric actions, the same agent would need to know their own orientation which is not directly accessible with an egocentric visual encoding. Hence, the difference of performance between the panels **Figures 5A,B** when learning by trial and error.

With allocentric vision (for which the learning of whether the food is visible by the *Dominant* was best in the supervised setting), the performance and speed of learning by reinforcement are significantly inferior to the case with egocentric vision. In this case, agents with allocentric visual input also need to take into account their orientation in order to choose the appropriate action toward the goal. For example, if the agent uses an allocentric viewpoint but an egocentric action space, some of the simple heuristics available for egocentric representations such as “move forward when food is in front of you” seem harder to



discover. This might be due to the fact that additional processing is required to extract the relevant variables (e.g., is the food item in front of the agent?) to implement certain heuristics using allocentric visual encoding.

So far, we have quantified the success of the different types of *Subordinate* agents by the amount of reward obtained. Next, we refine the analysis of the success by evaluating which behavior and types of errors the agents committed. In particular, we tested all 4 models by initializing the agents and environment in all possible allowed combinations. Then we counted in how many episodes the agent performed a correct or incorrect behavior keeping track on whether in the given trial the food item should have been eaten or not (food item is within the visual field of the *Dominant* agent).

**Figure 6A** shows the results of the analysis. The agent equipped with egocentric viewpoint and action space ( $\text{ego}_{\text{vis}}\text{-ego}_{\text{act}}$ ), obtained the food item in more than 93.38% of the cases when it should eat the food (non-observed by the dominant) and avoided it 99.3% of the cases when it should avoid it (observed by the dominant). Compared to the 78.49 and 95.89% for the  $\text{allo}_{\text{vis}}\text{-allo}_{\text{act}}$  case, we see that this type of agent had a main issue in obtaining the food when this was edible (not seen by the dominant). Similar performance in both types of trials is realized by the agent endowed with egocentric actions ( $\text{allo}_{\text{vis}}\text{-ego}_{\text{act}}$ ). Most dramatically, agents with egocentric visual processing and allocentric action space obtain the food item in only 32.59% of the cases when the food is rewarded. **Figures 6B,C** panels show an example of incorrect behavior while **Figures 6D,E** panels show an example of correct behavior.

Finally, we also analyzed the information that each layer of the architecture contains about whether the food is observed by the *Dominant* agent or not. This is the essential bit of information that needs to be extracted from the visual input to guide the decision of the agent of whether to approach the food or not. To this end we added a linear decoder in each layer of the neural architecture for each of the 4 RL models. That is, we trained a linear decoder (linear discriminant analysis) from different layers of the architecture to predict whether the food is visible by the *Dominant* agent. **Figure 7** shows that indeed it is possible to reliably decode this information from 3 out of the 4 types of trained agents. Also we observed that the information about whether the food is visible from the *Dominant* agent increases with further layers of processing until peaking around the last convolutional layer and the LSTM layer to slightly decrease near the output layer.

## 4. DISCUSSION

In this work we aimed to develop RL agents that could solve a basic perspective taking task. For that we devised a perspective taking task inspired by work done with chimpanzees. Our goal was not to design a comprehensive model of perspective taking behavior in humans or other animals. Rather, we explored whether relatively simple deep reinforcement learning algorithms can capture basic computational aspects of perspective taking. The behavior of the agents showed evidence for basic perspective

taking skills, which shows that at least a part of perspective taking skills might indeed be learned through RL.

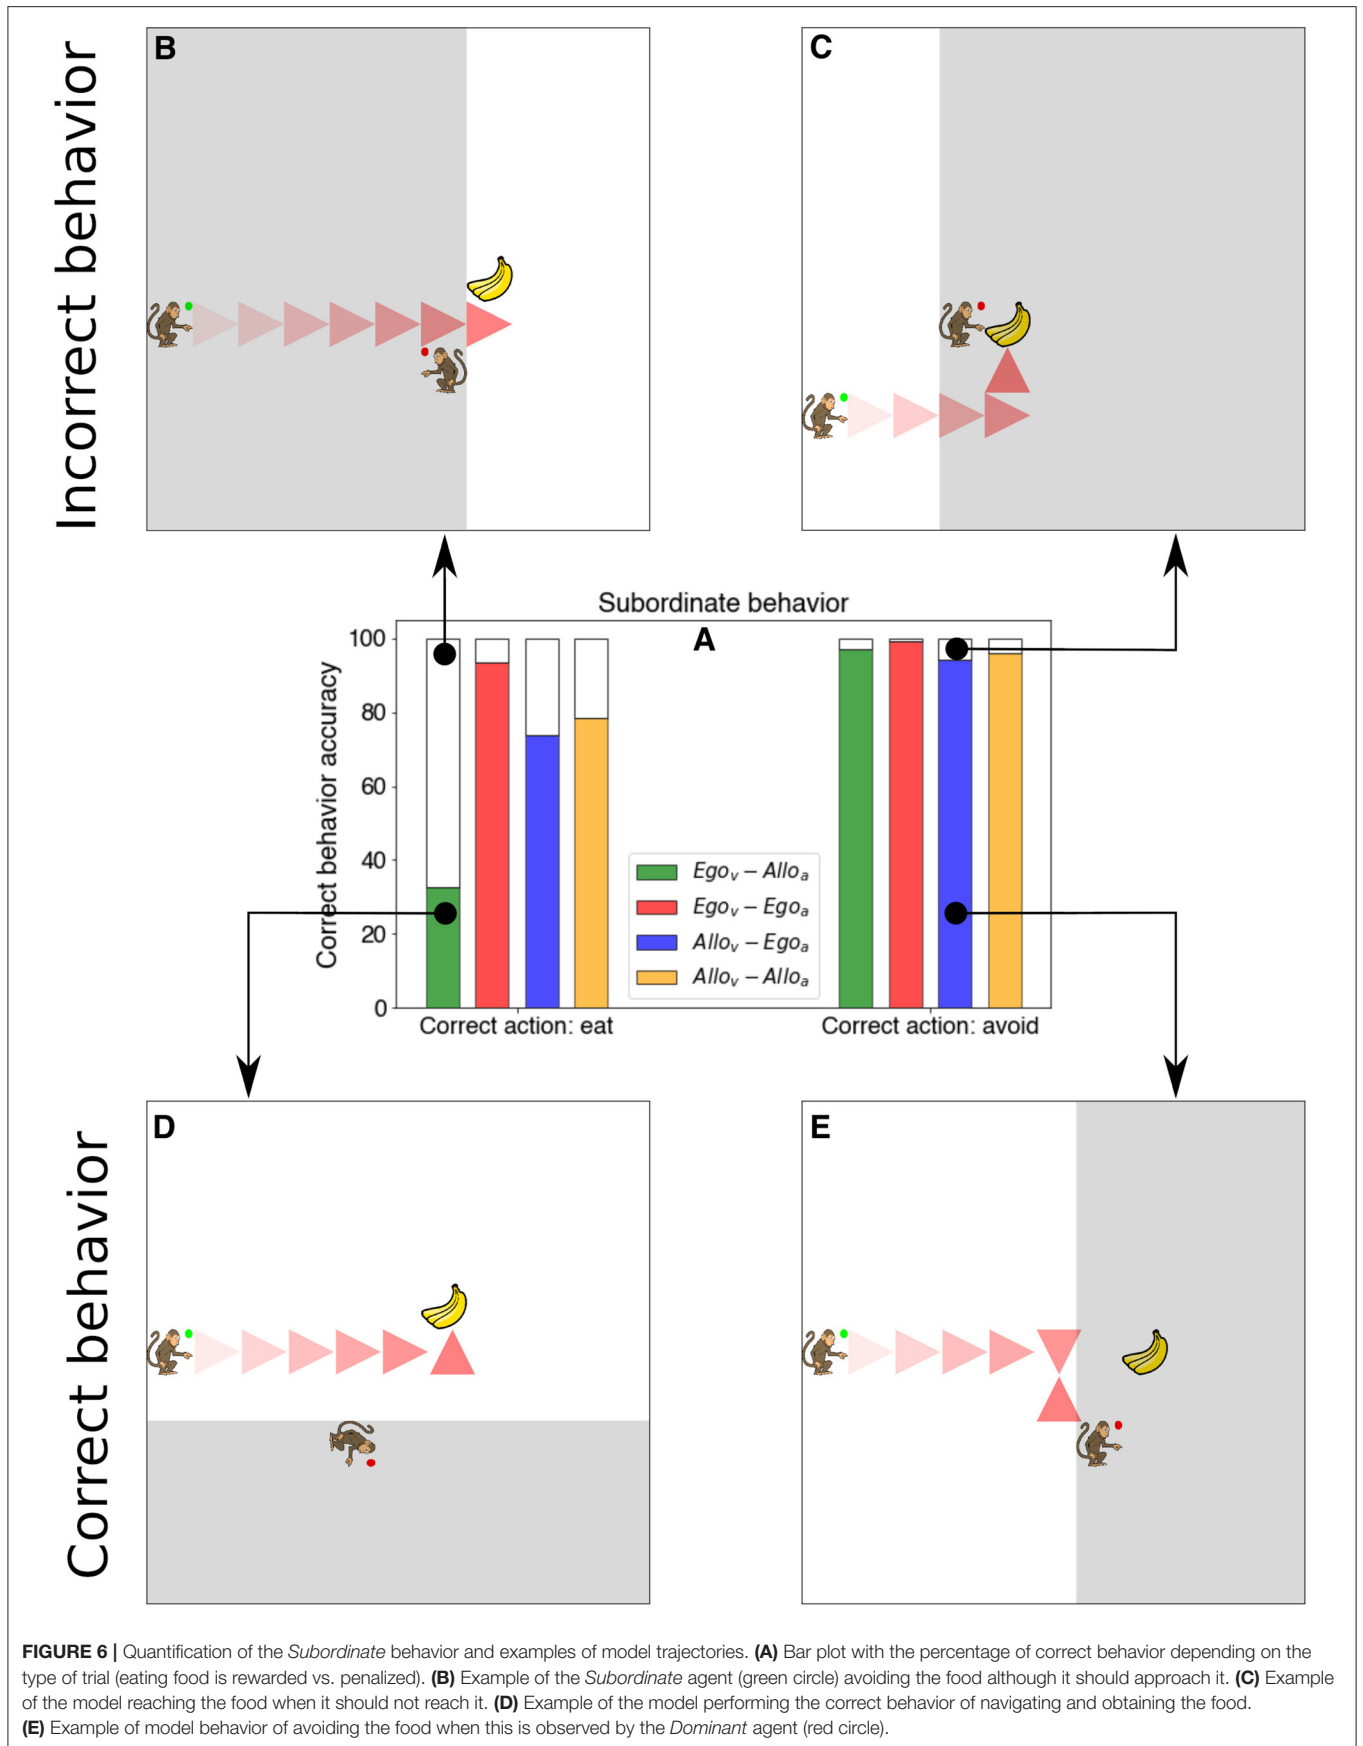
The real advantage of AI algorithms is that they allow us to deconstruct the studied process. In this case we deconstructed the perspective taking task by (a) separating the decision and navigational aspects of the task and (b) using different combinations of visual input (ego-vision vs. allo-vision) and action (ego-action vs. allo-action).

### 4.1. Allocentric vs. Egocentric Perspective Taking

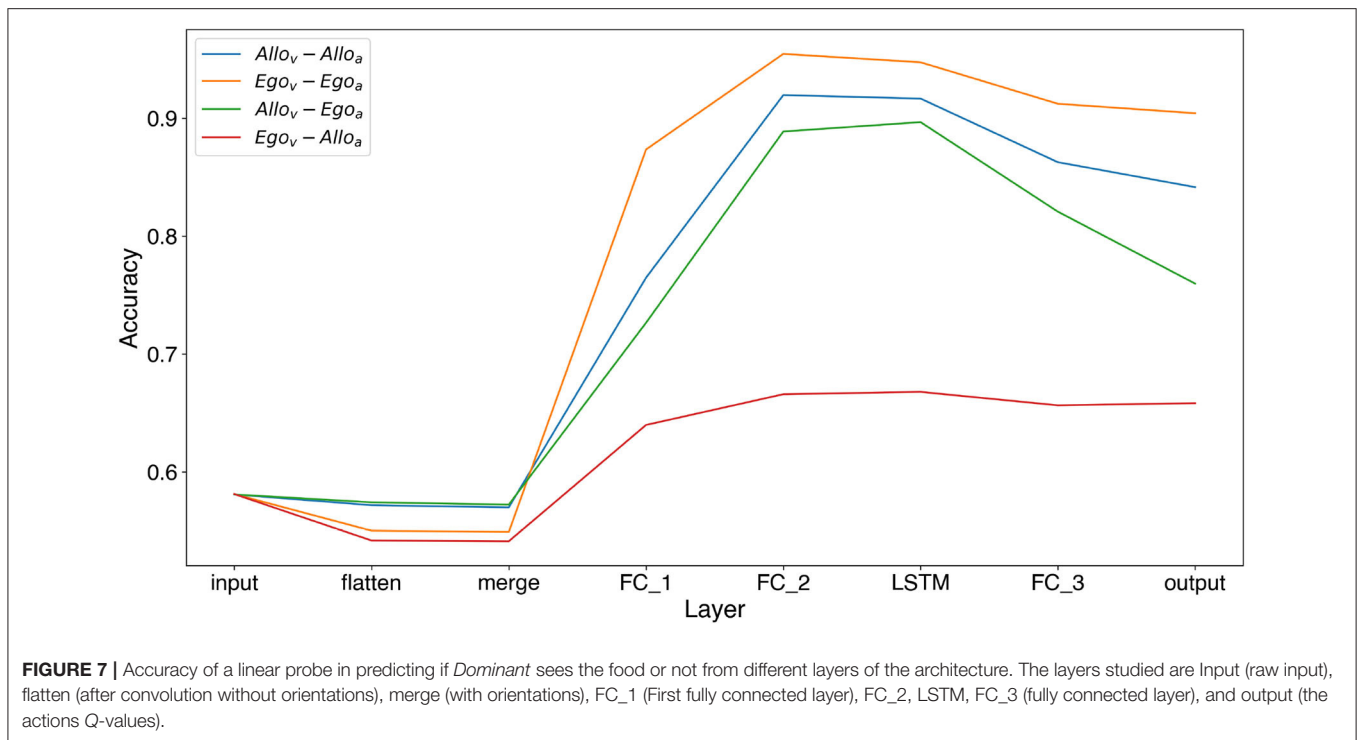
The condition with egocentric vision corresponds to the natural way how animals interact with the world: they perceive objects and other agents from their viewpoint. However, for some reason the brains of animals have also developed specific systems where objects are represented in allocentric fashion - they are represented in relation to each other as in a map (Burgess et al., 2001; Chersi and Burgess, 2015; Wang et al., 2020). This allocentric representation allows animals to compute certain aspects of the world more easily. One of such functions might be perspective taking. Indeed, in our work we found that in the supervised setting without navigation agents can much more readily learn perspective taking skills from allocentric input representations. Had we stopped there we would have concluded that the allocentric input representation is the optimal one for solving tasks involving perspective taking.

However, we then studied agents in the RL setting, where the agents needed to demonstrate their perspective taking skills by navigating toward the food item. Under these settings we found that in the simple environment studied  $\text{ego}_{\text{vis}}\text{-ego}_{\text{act}}$  agents clearly outperformed all other visual-action combinations, including the  $\text{allo}_{\text{vis}}\text{-allo}_{\text{act}}$  case. This seems to show that when navigation is involved the egocentric representation is actually much more efficient for learning. That said, it has to be kept in mind that we studied a very simple setting. In all of our experiments, the goal was in the field of the vision of the *Subordinate*. Hence, this is like navigating toward a building you see in front of you. Using a map probably will make things more complicated.

One could assert that our agents had learned a coordinate transformation, not true perspective taking. However, we believe that coordinate transformations that occur in the parietal cortex and hippocampus in the mammalian brain also support the computations supporting basic levels of perspective taking in biological systems. An important question in neuroscience is how this transformation from egocentric to allocentric coordinates is computed in the brain (Chersi and Burgess, 2015; Bicanski and Burgess, 2018). Also, it is clear that in the animals these two systems interact (Wilber et al., 2014; Bicanski and Burgess, 2018; Wang et al., 2020). In the present work, we did not study the interactions of these two systems. In the future work we seek to study how the allocentric representation is computed from the egocentric input and how these two systems interact during online decision making. Moreover, we aim to design more complex experiments to distinguish between simple coordinate transformation and other aspects of perspective taking. In future



**FIGURE 6 |** Quantification of the *Subordinate* behavior and examples of model trajectories. **(A)** Bar plot with the percentage of correct behavior depending on the type of trial (eating food is rewarded vs. penalized). **(B)** Example of the *Subordinate* agent (green circle) avoiding the food although it should approach it. **(C)** Example of the model reaching the food when it should not reach it. **(D)** Example of the model performing the correct behavior of navigating and obtaining the food. **(E)** Example of model behavior of avoiding the food when this is observed by the *Dominant* agent (red circle).



work we seek to develop tasks such that the agents would have to integrate not only the poses and orientations of objects but rather the contents of the field of view of the other agent, the attention of the other agent, and whether the goal object has been observed by the other agent in the past.

## 4.2. Different Levels of Perspective Taking in Humans and AI

Although some agents in this study could solve the visual perspective taking tasks presented here, we do not claim that the agents learned to reconstruct the perspective of another agent. They used more simple cues to solve the task at hand. The subordinate agent does not have any pressure to reconstruct the view from the dominant, other than the existence of the food item in its view (which is the event associated with punishment or reward). Hence, while the agent likely did not reconstruct the view from another agent point of view, it is able to infer the event of interest which is the existence of a certain item in the dominants point of view. In particular, our agent learned to infer the relationship between reward, the food item and its relation to the position and angle of the dominant agent. While this seems to be a very low level mechanism of perspective taking, it is at least one concrete mechanism: We have shown that deep RL agents can learn this type of knowledge necessary for simple perspective taking tasks.

We acknowledge that the complexity of the task is nowhere near what humans encounter in their everyday perspective taking tasks (Apperly, 2011). Also, although the current task was inspired by work done with chimpanzees, the chimpanzees were not trained on the task, they were just tested at it; they had acquired perspective taking from various encounters

with different chimpanzees under natural conditions (Hare et al., 2000). Furthermore, our RL agents could compute their decision to approach or avoid the food based on the orientation of the *Dominant* and the food position, whereas in the original experiments with chimpanzees (Hare et al., 2000) the different conditions also involved obstacles (to hide the food). Hence, our current setup is a very simplified version of perspective taking, but hopefully it lies the groundwork for more elaborate experiments. In future work, it would be desirable to make direct comparisons between primate behavior and the performance of RL agents under exactly the same task and learning conditions.

Based on human studies, perspective taking has been divided into two levels (Apperly, 2011), namely level 1 and 2 perspective taking. Level 1 perspective taking is about the question whether the other agent knows about a particular object or not. This system is thought to be more basic and automatic and is functional also in small children and chimpanzees (Hare et al., 2000; Tomasello et al., 2003; Apperly, 2011). Level 2 perspective taking is more complicated: here it is not only about whether the other agent knows about the object but also about how the other agent sees it (Apperly, 2011). Our agents mastered tasks that have been used to probe level 1 perspective taking in chimpanzees, but even here the caveat is that our RL agents were trained for thousands of episodes in the very same task. When children and chimpanzees are given level 1 perspective taking tasks, they solve it without training (Apperly, 2011). In this sense we claim that the current RL agents solve level 0 perspective taking tasks that we define as “achieving perspective taking level 1 behavior after extensive training.” From this viewpoint it is also clear which tasks should our RL agents try to solve next-level

1 and 2 perspective taking without extensive training on the same task.

Rabinowitz and colleagues used a meta-learning strategy to close the gap between level 0 and level 2 perspective taking (Rabinowitz et al., 2018), however their study had several assumptions (training by supervision, full observability and non-behaving agents) that were addressed in our study. In general, a future avenue of work should address more complex environments and perspective taking tasks involving more agents and stringent generalization tests.

Another avenue for future work is that of opening the networks that successfully implement perspective-taking capabilities during RL. In particular, it will be interesting to search and study the receptive fields of specific neurons in the network whose activity correlates with a decision requiring perspective-taking skills.

## 5. CONCLUSION

Perspective taking, like any other cognitive ability, has multiple facets and for the scientific understanding of such abilities it is necessary to study all of these facets (Apperly, 2011). Here we studied the simplest possible case where agents controlled by artificial neural networks learned with the help of RL in a simple task.

Theory of Mind involves many other processes and assumptions lacking in current AI agents. Here we studied the efficiency of different visuomotor representations in solving a task that requires to take into account the perspective of another agent. Investigating the capabilities and limitations of RL agents in acquiring perspective taking is a first step toward dissecting the algorithmic and representational options underlying perspective taking and more generally, Theory of

Mind. As human communication heavily relies on Theory of Mind (Tomasello, 2010, 2019; Scott-Phillips, 2014), a better understanding of Theory of Mind is a prerequisite for developing AI algorithms that can take the perspective and comprehend the intentions of humans.

## DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found here: [https://github.com/aeel13932/APES\\_PT](https://github.com/aeel13932/APES_PT).

## AUTHOR CONTRIBUTIONS

RV and JA designed research. AL and RV performed research. AT and TM helped with research. RV and JA supervised research. JA, RV and AL wrote the manuscript.

## FUNDING

This work was supported by PUT1476 from the Estonian Ministry of Education and Research. RV also thanks the Estonian Centre of Excellence in IT (EXCITE) funded by the European Regional Development Fund, through the research grant TK148. JA was also supported by the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Grant Agreement No. 799411.

## ACKNOWLEDGMENTS

We would like to thank Daniel Majoral and Michael Wibral for fruitful discussions and comments. This manuscript has been released as a pre-print at [arxiv.org](https://arxiv.org) (Labash et al., 2019).

## REFERENCES

- Apperly, I. (2011). *Mindreaders: The Cognitive Basis of "Theory of Mind"*. Psychology Press. doi: 10.4324/9780203833926
- Aru, J., and Vicente, R. (2018). What deep learning can tell us about higher cognitive functions like mindreading? *arXiv [Preprint]*. arXiv:1803.10470. doi: 10.31234/osf.io/9skek
- Bicanski, A., and Burgess, N. (2018). A neural-level model of spatial memory and imagery. *eLife* 7:e33752. doi: 10.7554/eLife.33752
- Burgess, N., Becker, S., King, J. A., and O'Keefe, J. (2001). Memory for events and their spatial context: models and experiments. *Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci.* 356, 1493–1503. doi: 10.1098/rstb.2001.0948
- Chersi, F., and Burgess, N. (2015). The cognitive architecture of spatial navigation: hippocampal and striatal contributions. *Neuron* 88, 64–77. doi: 10.1016/j.neuron.2015.09.021
- Chollet, F. (2015). *Keras*. Available online at: <https://keras.io/>
- Cole, G. G., and Millett, A. C. (2019). The closing of the theory of mind: a critique of perspective-taking. *Psychonom. Bull. Rev.* 26, 1787–1802. doi: 10.3758/s13423-019-01657-y
- Davis, M. H. (1983). Measuring individual differences in empathy: evidence for a multidimensional approach. *J. Pers. Soc. Psychol.* 44, 113–126. doi: 10.1037/0022-3514.44.1.113
- Galinsky, A. D., Maddux, W. W., Gilin, D., and White, J. B. (2008). Why it pays to get inside the head of your opponent: the differential effects of perspective taking and empathy in negotiations. *Psychol. Sci.* 19, 378–384. doi: 10.1111/j.1467-9280.2008.02096.x
- Goldberg, T. L., and Wrangham, R. W. (1997). Genetic correlates of social behaviour in wild chimpanzees: evidence from mitochondrial DNA. *Anim. Behav.* 54, 559–570. doi: 10.1006/anbe.1996.0450
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. Available online at: <http://www.deeplearningbook.org>
- Hare, B., Call, J., Agnetta, B., and Tomasello, M. (2000). Chimpanzees know what conspecifics do and do not see. *Anim. Behav.* 59, 771–785. doi: 10.1006/anbe.1999.1377
- Hausknecht, M., and Stone, P. (2015). "Deep recurrent q-learning for partially observable MDPs," in *2015 AAAI Fall Symposium Series*.
- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780. doi: 10.1162/neco.1997.9.8.1735
- Jaques, N., Lazaridou, A., Hughes, E., Gülçehre, Ç., Ortega, P. A., Strouse, D., et al. (2019). "Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning," in *Proceedings of the 36th International Conference on Machine Learning*, eds C. Kamalika and S. Ruslan (Long Beach, CA: PMLR), 97, 3040–3049. Available online at: <http://proceedings.mlr.press/v97/jaques19a/jaques19a.pdf>
- Labash, A., Aru, J., Matiisen, T., Tampuu, A., and Vicente, R. (2019). Perspective taking in deep reinforcement learning agents. *arXiv [Preprint]*. arXiv:1907.01851.

- Labash, A., Tampuu, A., Matiisen, T., Aru, J., and Vicente, R. (2018). APES: a python toolbox for simulating reinforcement learning environments. *arXiv [Preprint]*. [arXiv:1808.10692](https://arxiv.org/abs/1808.10692).
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. (2016). End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.* 17, 1334–1373.
- Mercier, H., and Sperber, D. (2017). *The Enigma of Reason*. Harvard University Press.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature* 518, 529–533. doi: 10.1038/nature14236
- Premack, D., and Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behav. Brain Sci.* 1, 515–526. doi: 10.1017/S0140525X00076512
- Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, C., Eslami, S., and Botvinick, M. (2018). “Machine theory of mind,” in *Proceedings of the 35th International Conference on Machine Learning*, eds J. Dy and A. Krause (Stockholmsmässan: PMLR), 80, 4218–4227. Available online at: <http://proceedings.mlr.press/v80/rabinowitz18a/rabinowitz18a.pdf>
- Ryskin, R. A., Benjamin, A. S., Tullis, J., and Brown-Schmidt, S. (2015). Perspective-taking in comprehension, production, and memory: an individual differences approach. *J. Exp. Psychol. Gen.* 144:898. doi: 10.1037/xge0000093
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599. doi: 10.1126/science.275.5306.1593
- Scott-Phillips, T. (2014). *Speaking Our Minds: Why Human Communication is Different, and How Language Evolved to Make It Special*. Macmillan International Higher Education.
- Sutton, R. S., and Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press.
- Thorndike, E. L. (1911). *Animal Intelligence: Experimental Studies*. Macmillan. doi: 10.5962/bhl.title.55072
- Tomasello, M. (2010). *Origins of Human Communication*. MIT Press.
- Tomasello, M. (2019). *Becoming Human: A Theory of Ontogeny*. Belknap Press. doi: 10.4159/9780674988651
- Tomasello, M., Call, J., and Hare, B. (2003). Chimpanzees understand psychological states—the question is which ones and to what extent. *Trends Cogn. Sci.* 7, 153–156. doi: 10.1016/S1364-6613(03)00035-4
- Wang, C., Chen, X., and Knierim, J. J. (2020). Egocentric and allocentric representations of space in the rodent brain. *Curr. Opin. Neurobiol.* 60, 12–20. doi: 10.1016/j.conb.2019.11.005
- Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., and Freitas, N. (2016). “Dueling network architectures for deep reinforcement learning,” in *International Conference on Machine Learning*, 1995–2003.
- Wilber, A. A., Clark, B. J., Forster, T. C., Tatsuno, M., and McNaughton, B. L. (2014). Interaction of egocentric and world-centered reference frames in the rat posterior parietal cortex. *J. Neurosci.* 34, 5431–5446. doi: 10.1523/JNEUROSCI.0511-14.2014

**Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Labash, Aru, Matiisen, Tampuu and Vicente. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.