
The Natural Language of Actions

Guy Tennenholtz¹ Shie Mannor¹

Abstract

We introduce Act2Vec, a general framework for learning context-based action representation for Reinforcement Learning. Representing actions in a vector space help reinforcement learning algorithms achieve better performance by grouping similar actions and utilizing relations between different actions. We show how prior knowledge of an environment can be extracted from demonstrations and injected into action vector representations that encode natural compatible behavior. We then use these for augmenting state representations as well as improving function approximation of Q-values. We visualize and test action embeddings in three domains including a drawing task, a high dimensional navigation task, and the large action space domain of StarCraft II.




Figure 1. A schematic visualization of action pair embeddings in a navigation domain using a distributional representation method. Each circle represents an action of two consecutive movements in the world. Actions that are close have similar contexts. Relations between actions in the vector space have interpretations in the physical world.

1. Introduction

The question “What is language” has had implications in the fields of neuropsychology, linguistics, and philosophy. One definition tells us that language is “a purely human and non-instinctive method of communicating ideas, emotions, and desires by means of voluntarily produced symbols” (Sapir, 1921). Much like humans adopt languages to communicate, their interaction with an environment uses sophisticated languages to convey information. Inspired by this conceptual analogy, we adopt existing methods in natural language processing (NLP) to gain a deeper understanding of the “natural language of actions” with the ultimate goal of solving reinforcement learning (RL) tasks.

In recent years, many advances were made in the field of distributed representations of words (Mikolov et al., 2013b; Pennington et al., 2014; Zhao et al., 2017). Distributional methods make use of the hypothesis that words which occur in a similar context tend to have similar meaning (Firth, 1957), i.e., the meaning of a word can be inferred from

the distribution of words around it. For this reason, these methods are called “distributional” methods. Similarly, the context in which an action is executed holds vital information (e.g., prior knowledge) of the environment. This information can be transferred to a learning agent through distributed action representations in a manifold of diverse environments.

In this paper, actions are represented and characterized by the company they keep (i.e., their context). We assume the context in which actions reside is induced by a demonstrator policy (or set of policies). We use the celebrated continuous skip-gram model (Mikolov et al., 2013a) to learn high-quality vector representations of actions from large amounts of demonstrated trajectories. In this approach, each action or sequence of actions is embedded in a d -dimensional vector that characterizes knowledge of acceptable behavior in the environment.

As motivation, consider the problem of navigating a robot. In its basic form, the robot must select a series of primitive actions of going straight, backwards, left, or right. By reading the words “straight”, “backwards”, “left”, and “right” the reader already has a clear understanding of their implication in the physical world. Particularly, it is presumed that moving straight and then left should have higher correlation to moving left and then straight, but for the most part contrast to moving right and then backwards. Moreover, it

¹Faculty of Electrical Engineering, Technion Institute of Technology, Israel. Correspondence to: Guy Tennenholtz <guytenn@gmail.com>, Shie Mannor <shie@technion.ac.il>.

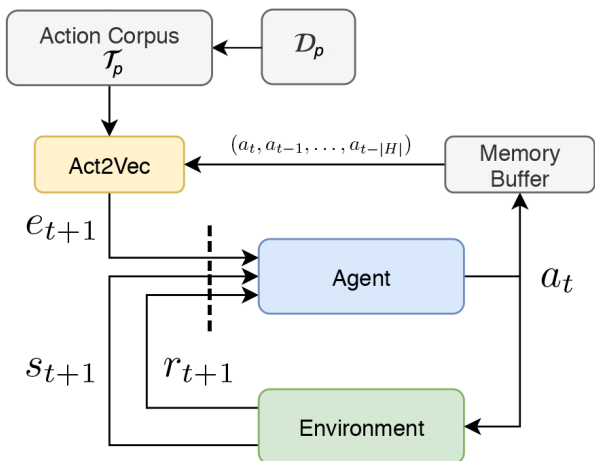


Figure 2. A schematic use of action embedding for state augmentation. Act2Vec is trained over sequences of length $|H|$ actions taken from trajectories in the action corpus \mathcal{T}_p , sampled from \mathcal{D}_p . Representation of action histories (e_{t+1}) are joined with the state (s_{t+1}). The agent π thus maps state and action histories ($s_t, a_{t-1}, \dots, a_{t-|H|}$) to an action a_t .

is well understood that a navigation solution should rarely take the actions straight and backwards successively, as it would postpone the arrival to the desired location. Figure 1 shows a 2-dimensional schematic of these action pairs in an illustrated vector space. We later show (Section 4.2) that similarities, relations, and symmetries relating to such actions are present in context-based action embeddings.

Action embeddings are used to mitigate an agent’s learning process on several fronts. First, action representations can be used to improve expressiveness of state representations, sometimes even replacing them altogether (Section 4.1). In this case, the policy space is augmented to include action-histories. These policies map current state and action histories to actions, i.e., $\pi : \mathcal{S} \times \mathcal{A}^{|H|} \rightarrow \mathcal{A}$. Through vector action representations, these policies can be efficiently learned, improving an agent’s overall performance. A conceptual diagram presenting this approach is depicted in Figure 2. Second, similarity between actions can be leveraged to decrease redundant exploration through grouping of actions. In this paper, we show how similarity between actions can improve approximation of Q -values, as well as devise a cluster-based exploration strategy for efficient exploration in large action space domains (Sections 3, 4.2).

Our main contributions in this paper are as follows. (1) We generalize the use of context-based embedding for actions. We show that novel insights can be acquired, portraying our knowledge of the world via similarity, symmetry, and relations in the action embedding space. (2) We offer uses of action representations for state representation, function approximation, and exploration. We demonstrate the advantage in performance on drawing and navigation domains.

This paper is organized as follows. Section 2 describes the general setting of action embeddings. We describe the Skip Gram with Negative Sampling (SGNS) model used for representing actions and its relation to the pointwise mutual information (PMI) of actions with their context. In Section 3 we show how embeddings can be used for approximating Q -values. We introduce a cluster based exploration procedure for large action spaces. Section 4 includes empirical uses of action representations for solving reinforcement learning tasks, including drawing a square and a navigation task. We visualize semantic relations of actions in the said domains as well as the large action space domain of Starcraft II. We conclude the paper with related work (Section 5) and a short discussion on future directions (Section 6).

2. General Setting

In this section we describe the general framework of Act2Vec. We begin by defining the context assumptions from which we learn action vectors. We then illustrate that embedding actions based on their point-wise mutual information (PMI) with their contexts contributes favorable characteristics. One beneficial property of PMI-based embedding is that close actions issue similar outcomes. In Section 4 we visualize Act2Vec embeddings of several domains, showing their inherent structure is aligned with our prior understanding of the tested domains.

A Markov Decision Processes (MDP) is defined by a 5-tuple $\mathcal{M}_R = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where \mathcal{S} is a set of states, \mathcal{A} is a discrete set of actions, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is a set of transition probabilities, $R : \mathcal{S} \rightarrow [0, 1]$ is a scalar reward, and $\gamma \in (0, 1)$ is a discount factor. We consider the general framework of multi-task RL in which R is sampled from a task distribution \mathcal{D}_R . In addition, we consider a corpus \mathcal{T}_p of trajectories, (s_0, a_0, s_1, \dots) , relating to demonstrated policies. We assume the demonstrated policies are optimal w.r.t. to task MDPs sampled from \mathcal{D}_R . More specifically, trajectories in \mathcal{T}_p are sampled from a distribution of permissible policies \mathcal{D}_p , defined by

$$P_{\pi \sim \mathcal{D}_p}(\pi) = P(\pi \in \Pi_{\mathcal{M}_R}^*) \forall \pi,$$

where $\Pi_{\mathcal{M}_R}^*$ is the set of optimal policies, which maximize the discounted return of \mathcal{M}_R .

We consider a state-action tuple (s_t, a_t) at time t , and define its context of width w as the sequence $c_w(s_t, a_t) = (s_{t-w}, a_{t-w}, \dots, s_{t-1}, a_{t-1}, s_{t+1}, a_{t+1}, \dots, s_{t+w}, a_{t+w})$. We will sometimes refer to state-only contexts and action-only contexts as contexts containing only states or only actions, respectively. We denote the set of all possible contexts by \mathcal{C} . With abuse of notation, we write (a, c) to denote the pair $\langle (s, a), c_w(s, a) \rangle$.

In our work we will focus on the pointwise mutual information (PMI) of (s, a) and its context c . Pointwise mutual in-

formation is an information-theoretic association measure between a pair of discrete outcomes x and y , defined as $PMI(x, y) = \log \frac{P(x, y)}{P(x)P(y)}$. We will also consider the conditional PMI denoted by $PMI(\cdot, \cdot | \cdot)$.

2.1. Act2Vec

We summarize the Skip-Gram neural embedding model introduced in (Mikolov et al., 2013a) and trained using the Negative-Sampling procedure (SGNS) (Mikolov et al., 2013b). We use this procedure for representing actions (in contrast to words) using their contexts, and refer to this procedure as Act2Vec.

Every action $a \in \mathcal{A}$ is associated with a vector $\vec{a} \in \mathbb{R}^d$. In the same manner, every context $c \in \mathcal{C}$ is associated with a vector $\vec{c} \in \mathbb{R}^d$. In SGNS, we ask, does the pair (a, c) come from \mathcal{D}_p ? More specifically, we ask, what is the probability that (a, c) came from \mathcal{D}_p ? This probability, denoted by $P(\mathcal{D}_p = \text{true}; a, c)$, is modeled as

$$P(\mathcal{D}_p = \text{true}; a, c) = \sigma(\vec{a}^T \vec{c}) = \frac{1}{1 + e^{-\vec{a}^T \vec{c}}}.$$

Here, \vec{a} and \vec{c} are the model parameters to be learned.

Negative Sampling (Mikolov et al., 2013b) tries to minimize $P(\mathcal{D}_p = \text{false}; a, c)$ for randomly sampled ‘‘negative’’ examples, under the assumption that randomly selecting a context for a given action is likely to result in an unobserved (a, c) pair. The local objective of every (a, c) pair is thus given by

$$\ell(a, c) = \log \sigma(\vec{a}^T \vec{c}) + k \mathbb{E}_{c_N \sim P_C} \log \sigma(-\vec{a}^T \vec{c}_N),$$

where $P_C(c)$ is the empirical unigram distribution of contexts in \mathcal{T}_p , and k is the number of negative samples.

The global objective can be written as a sum over losses of (a, c) pairs in the corpus

$$\ell = \sum_{a \in \mathcal{A}} \sum_{c \in \mathcal{C}} \#(a, c) \ell(a, c),$$

where $\#(a, c)$ denotes the number of times the pair (a, c) appears in \mathcal{T}_p .

Relation to PMI: Optimizing the global objective makes observed action-context pairs have similar embeddings, while scattering unobserved pairs. Intuitively, actions that appear in similar contexts should have similar embeddings. In fact, it was recently shown that SGNS implicitly factorizes the action-context matrix whose cells are the point-wise mutual information of the respective action and context pairs, shifted by a global constant (Levy & Goldberg, 2014). In what follows, we show that $PMI(a, c)$ is a useful measure for action representation in reinforcement learning.

2.2. State-only context

State-only contexts provide us with information about the environment as well as predictions of optimal trajectories. Let us consider the next-state context $c = s'$. More specifically, given a state action pair (s, a) we are interested in the measure $PMI(a, s' | s)$, where here s' is the random variable depicting the next state. The following lemma shows that when two actions a_1, a_2 have similar PMI w.r.t. (with respect to) their next state context, they can be joined into a single action with a small change in value.

Lemma 1. *Let $K \subseteq \mathcal{A}$ s.t.*

$$|PMI(a_1, s' | s) - PMI(a_2, s' | s)| < \epsilon, \forall a_1, a_2 \in K,$$

where $(a, s') \subset \tau \sim \pi \sim \mathcal{D}_p, \forall a \in K$.

Let $\pi \sim \mathcal{D}_p$ and denote

$$\pi_K(a | s) = \begin{cases} \frac{1}{|K|} \sum_{j=1}^{|K|} \pi(a_j | s) & , a \in K \\ \pi(a | s) & , o.w. \end{cases}$$

Then,

$$\|V^\pi - V^{\pi_K}\|_\infty \leq \frac{6\gamma}{(1-\gamma)^4} \epsilon$$

A proof to the lemma can be found in the supplementary material. The lemma illustrates that neglecting differences in actions of high proximity in embedding space has little effect on a policy’s performance. While state contexts provide fruitful information, their PMI may be difficult to approximate in large state domains. For this reason, we turn to action-only contexts, as described next.

2.3. Action-only context

Action-only contexts provide us with meaningful information when they are sampled from \mathcal{D}_p . To see this, consider the following property of action-only contexts. If action a_1 is more likely to be optimal than a_2 , then any context that has a larger PMI with a_1 than a_2 will also be more likely to be optimal when chosen with a_1 . Formally, let $\pi^* \sim \mathcal{D}_p$, $\tau^* \sim \pi^*$, and $a_1, a_2 \in \mathcal{A}, s \in \mathcal{S}$ such that

$$P(\pi^*(a_1 | s) \geq \pi^*(a_2 | s)) \geq \frac{1}{2}. \quad (1)$$

Let $c \in \mathcal{C}$. If

$$PMI(a_1, c | s) \geq PMI(a_2, c | s), \quad (2)$$

then

$$P((s, a_1, c) \subseteq \tau^*) \geq P((s, a_2, c) \subseteq \tau^*). \quad (3)$$

To show (3), we write the assumption in Equation (2) explicitly:

$$\frac{P(a_1, c | s)}{P(a_1 | s)P(c | s)} \geq \frac{P(a_2, c | s)}{P(a_2 | s)P(c | s)}.$$

Next, due to the definition of \mathcal{D}_p , we have that

$$\frac{P((s, a_1, c) \subseteq \tau^*)}{P(a_1|s)P(c|s)} \geq \frac{P((s, a_2, c) \subseteq \tau^*)}{P(a_2|s)P(c|s)}.$$

Giving us

$$\begin{aligned} P((s, a_1, c) \subseteq \tau^*) &\geq \frac{P(a_1|s)}{P(a_2|s)} P((s, a_2, c) \subseteq \tau^*) \\ &\geq P((s, a_2, c) \subseteq \tau^*), \end{aligned}$$

where the last step is due to our assumption in Equation 1.

In most practical settings, contexts based embeddings that use state contexts are difficult to train using SGNS. In contrast, action-only contexts usually consist of orders of magnitude less elements. For this reason, in Section 4 we experiment with action-only contexts, showing that semantics can be learned even when states are ignored.

3. Act2Vec for function approximation

Lemma 1 gives us intuition as to why actions with similar contexts are in essence of similar importance to overall performance. We use this insight to construct algorithms that depend on similarity between actions in the latent space. Here, we consider the set of discrete actions \mathcal{A} . Equivalently, as we will see in Section 4.2, this set of actions can be augmented to the set of all action sequences $\{(a_1, \dots, a_k)\}$. We are thus concerned with approximating the Q -value of state-action pairs $Q(s, a)$.

Q -Embedding: When implemented using neural networks, Q -Learning with function approximation consists of approximating the Q -value of state-action pairs by

$$\hat{Q}(s, a) = w_a^T \phi(s), \quad (4)$$

where $\phi(s)$ are the features learned by the network, and w_a are the linear weights learned in the final layer of the network. When the number of actions is large, this process becomes impractical. In NLP domains, it was recently suggested to use word embeddings to approximate the Q -function (He et al., 2016) as

$$\hat{Q}(s, a) = \psi(a)^T \phi(s), \quad (5)$$

where $\psi(a)$ are the learned features extracted from the embedding of a word a . Similar words become close in embedding space, thereby outputting similar Q -Values. This approach can also be applied on action embeddings trained using Act2Vec. These action representations adhere inherent similarities, allowing one to approximate their Q -values, while effectively obtaining complexity of smaller dimension. We will refer to the approximation in Equation 5 as Q -Embedding.

k -Exp: In Q -Learning, the most fundamental exploration strategy consists of uniformly sampling an action. When

the space of actions is large, this process becomes infeasible. In these cases, action representations can be leveraged to construct improved exploration strategies. We introduce a new method of exploration based on action embeddings, which we call k -Exp. k -Exp is a straightforward extension of uniform sampling. First, the action embedding space is divided into k clusters using a clustering algorithm (e.g., k -means). The exploration process then follows two steps: (1) Sample a cluster uniformly, and (2) given a cluster, uniformly sample an action within it. k -Exp ensures actions that have semantically different meanings are sampled uniformly, thereby improving approximation of Q -values.

In Section 4.2 we compare Q -Embedding with k -Exp to Q -learning with uniform exploration, demonstrating the advantage of using action representations.

4. The Semantics of Actions

Word embeddings have shown to capture large numbers of syntactic and semantic word relationships (Mikolov et al., 2013b). Motivated by this, as well as their relation to PMI , we demonstrate similar interpretation on several reinforcement learning environments. This section is divided into three parts. The first and second parts of this section consider the tasks of drawing a square and navigating in 3d space. In both parts, we demonstrate the semantics captured by actions in their respective domains. We demonstrate the effectiveness of Act2Vec in representing a state using the sequence of previously taken actions (see Figure 2). We then demonstrate the use of Act2Vec with Q -Embedding and k -Exp. Finally, in the third part of this section we demonstrate the semantic nature of actions learned in the complex strategy game of StarCraft II.

4.1. Drawing

We undertook the task of teaching an agent to draw a square given a sparse reward signal. The action space consisted of 12 types of strokes: Left, Right, Up, Down, and all combinations of corners (e.g., Left+Up). The sparse reward provided the agent with feedback only once she had completed her drawing, with positive feedback only when the drawn shape was rectangular. Technical details of the environment can be found in the supplementary material. We trained Act2Vec with action-only context over a corpus of 70,000 human-made drawings in the “square category” of the *Quick, Draw!* (Cheema et al., 2012) dataset. Projections of these embeddings are depicted in Figure 3(a). The embedding space projection reflects our knowledge of these action strokes. The space is divided into 4 general regions, consisting of strokes in each of the main axis directions. Strokes relating to corners of the squares are centered in distinct clusters, each in proximity to an appropriate direction. The embedding space presents evident symmetry w.r.t. clockwise vs. counterclockwise drawings of squares.

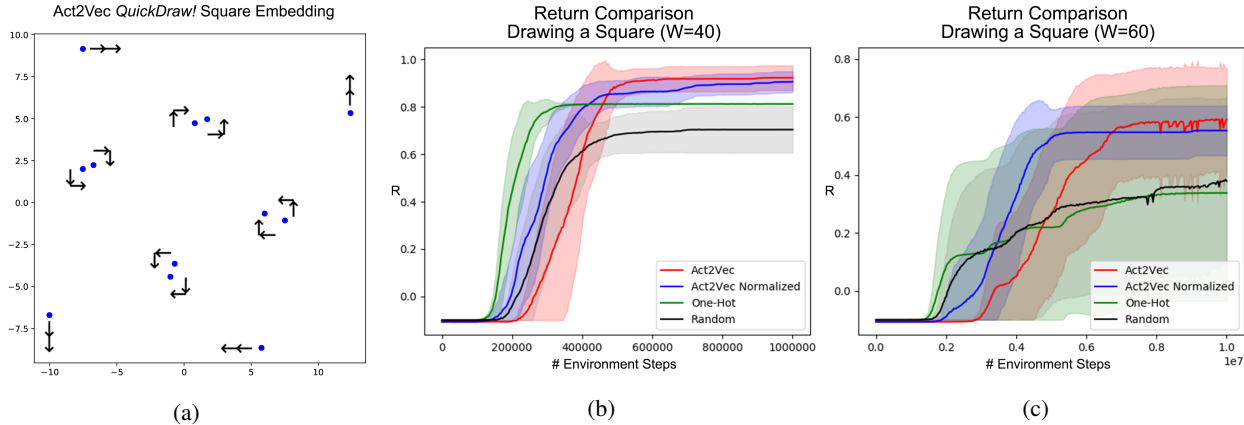


Figure 3. (a): Plot shows Act2Vec embedding of *QuickDraw!*'s square category strokes. Actions are positioned according to direction. Corner strokes are organized in relation to specific directions, with symmetry w.r.t. direction in which squares were drawn. (b-c): Comparison of state representations for drawing squares with different edge lengths ($W = 40, 60$). The state was represented as the sum of previous action embeddings. Results show superiority of Act2Vec embedding as opposed to one-hot and random embeddings.

The action space in this environment is relatively small. One way of representing the state is through the set of all previous actions, since in this case $s_t = \cup_{i=0}^{t-1} a_i$. The state was therefore represented as the vector equal to the sum of previous action vectors. We compared three types of action embeddings for representing states: Act2Vec, normalized Act2Vec (using the l_2 norm), one-hot, and randomized embeddings. Figure 3(b,c) shows results of these representations for different square sizes. Act2Vec proved to be superior on all tasks, especially with increased horizon - where the sparse reward signal drastically affected performance. We also note that normalized Act2Vec achieved similar results with higher efficiency. In addition, all methods but Act2Vec had high variance in their performance over trials, implying they were dependent on the network's initialization. A detailed overview of the training process can be found in the supplementary material.

4.2. Navigation

In this section we demonstrate how sequences of actions can be embedded using Act2Vec. We then show how embeddings based on trajectories captured in a simple navigation domain can *transfer knowledge* to a more complex navigation domain, thus improving its learning efficiency.

In physical domains, acceptable movements of objects in space are frequently characterized by smoothness of motion. As such, when we open a door, we move our hand smoothly through space until reaching the knob, after which we complete a smooth rotation of the doorknob. An agent learning in such an environment may tend to explore in a manner that does not adhere to patterns of such permissible behavior (e.g., by uniformly choosing an arbitrary action). Moreover, when inspecting individual tasks, actions incorporate various properties that are particular to the task at hand. Looking left and then right may be a use-

less sequence of actions when the objective is to reach a goal in space, while essential for tasks where information gathered by looking left contributes to the overall knowledge of an objective on the right (e.g., looking around the walls of a room). In both cases, looking left and immediately right is without question distinct to looking left and then looking left again. These semantics, when captured properly, can assist in solving any navigation task.

When studying the task of navigation, one is free to determine an action space of choice. In most applications, the primitive space of actions is either defined by fixed increments of movement and rotation or by physical forces. Let us consider the former case and more specifically examine the action space consisting of moving forward (marked by \uparrow), and rotating our view to the left or to the right (marked by \leftarrow and \rightarrow , respectively). These three actions are in essence sufficient for solving most navigation tasks. Nonetheless, semantic relations w.r.t. these actions become particularly evident when action sequences are used. For this case, we study the augmented action space consisting of all action sequences of length k , i.e., $(a_1, a_2, \dots, a_k) \in \mathcal{A}^k$. For example, for the case of $k = 2$, the action sequence (\uparrow, \uparrow) would relate to taking the action \uparrow twice, thereby moving two units forward in the world, whereas (\uparrow, \leftarrow) would relate to moving one unit forward in the world and then rotating our view one unit to the left. This augmented action space holds interesting features, as we see next.

We trained Act2Vec with action-only context on a corpus of 3000 actions taken from a 2d navigation domain consisting of randomly generated walls. Given a random goal location, we captured actions played by a human player in reaching the goal. Figure 4 (a,b) shows Principal Component Analysis (PCA) projections of the resulting embeddings for action sequences of length $k = 2, 3$. Examining




Figure 4. (a-b): The action corpus of size 3000 actions was generated by actions executed in a 2d navigation domain consisting of three actions: move forward (\uparrow), rotate view left (\leftarrow), and rotate view right (\rightarrow). Plots show PCA projection of Act2Vec embedding for sequences of length 2 and 3. (c): Comparison of techniques on the Seek-Avoid environment. Plots show results for different sequence lengths, with and without Q -Embedding. Sequences of length 3 only showed improvement when cluster based exploration was used.

the resulting space, we find two interesting phenomena. First, the embedding space is divided into several logical clusters, each relating to an aspect of movement. In these, sequences are divided according to their forward momentum as well as direction. Second, we observe symmetry w.r.t. the vertical axis, relating to looking left and right. These symmetrical relations capture our understanding of the consequences of executing these action sequences. In the next part of this section we use these learned embeddings in a *different* navigation domain, teaching an agent how to navigate while avoiding unwanted objects.

Knowledge Transfer: We tested the Act2Vec vector embedding trained on the 2d navigation domain on the DeepMind Lab (Beattie et al., 2016) Seek-Avoid environment¹. Here, an agent must navigate in 3d space, learning to collect good apples while avoiding bad ones. The agent was trained using Q-learning with function approximation. We tested sequences of length $k = 1, 2, 3$ using both methods of function approximation (Equations 4 and 5).

Results, as depicted in Figure 4(c), show superiority of using embeddings to approximate Q-values. Action sequences of length $k = 2$ showed superiority over $k = 1$ with ϵ -uniform exploration, with 20 percent increase in total reward. Sequences of length 3 did not exhibit an increase in performance. We speculate this is due to the uniform exploration process. To overcome this matter, we used k-means in order to find clusters in the action embedding space. We then evaluated k -Exp on the resulting clusters. While results were dependent on the initialization of clusters (due to randomization of k -means), sequences of

¹While the tasks of navigation in the 2d domain and the Seek-Avoid domain are different, the usage of their actions presents similar semantic knowledge, thereby incorporating transfer of knowledge between these domains.

length 3 showed increase in total reward. Additional technical details can be found in the supplementary material.

Regularization: Learning to navigate using sequence of actions led to smoothness of motion in the resulting tests. This property arose due to the unique clustering of sequences. Particularly, the sequences (\leftarrow, \rightarrow) and (\rightarrow, \leftarrow) were discarded by the agent, allowing for a smooth flow in navigation. Consequently, this property indicates that action sequences act as a regularizer to the navigation task.

4.3. StarCraft II

StarCraft II is a popular video game that presents a very hard challenge for reinforcement learning. Its main difficulties include a huge state and action space as well as a long-time horizon of thousands of states. The consequences of any single action (in particular, early decisions) are typically only observed many frames later, posing difficulties in temporal credit assignment and exploration. Act2Vec offers an opportunity to mitigate some of these challenges, by finding reliable similarities and relations in the action space.

We used a corpus of over a million game replays played by professional and amateur players. The corpus contained over 2 billions played actions, which on average are equivalent to 100 years of consecutive gameplay. The action space was represented by over 500 action functions, each with 13 types of possible arguments, as described in (Vinyals et al., 2017). We trained Act2Vec with action-only context to embed the action functions into action vectors of dimension d , ignoring any action arguments. T-SNE (Maaten & Hinton, 2008) projections of the resulting action embeddings are depicted in Figure 5.

In StarCraft II players choose to play one of three species:




Figure 5. Plots show t-SNE embedding of Starcraft II action functions for all races (a) as well as the Terran race (b). Action representation distinguish between all three races through separate clusters. Plot (b) depicts clusters based on categorial types: building, training, researching, and effects. Clusters based on common player strategies appear in various parts of the embedding space.

Terran, Protoss, or Zerg. Once a player chooses her race, she must defeat her opponent through strategic construction of buildings, training of units, movement of units, research of abilities, and more. While a myriad strategies exist, expert players operate in conventional forms. Each race admits to different strategies due to its unique buildings, units, and abilities. Embeddings depicted in Figure 5(a) show distinct clusters of the three different races. Moreover, actions that are commonly used by all three races are projected to a central position with equal distance to all race clusters. Figure 5(b) details a separate t-SNE projection of the Terran race action space. Embeddings are clustered into regions with definite, distinct types, including: training of units, construction of buildings, research, and effects. While these actions seem arbitrary in their raw form, Act2Vec, through context-based embedding, captures their relations through meaningful clusters.

Careful scrutiny of Figure 5(b), shows interesting captured semantics, which reveal common game strategies. As an example, let us analyze the cluster containing actions relating to training Marines, Marauders, Medivacs, and Widow Mines. Marines are an all-purpose infantry unit, while Marauders being almost opposite to the Marine units, are effective at the front of an engagement to take damage for Marines. Medivacs are dual purpose dropships and healers. They are common in combination with Marines and Marauders, as they can drop numbers of Marines and Marauders and then support them with healing. A Widow Mine is a light mechanical mine that deals damage to ground or air units. Widow Mines are used, as a standard opening move, in conjunction with Marines and Medivacs, to zone out opposing mineral fields. Other examples of strategic clusters

include the Ghost and the Nuke, which are commonly used together, as well as the Stimpack, Combat Shield, and Concussive Shells abilities, which are Marine and Marauder upgrades, all researched at the Tech Lab attached to a Barracks.

The semantic representation of actions in StarCraft II illustrates how low dimensional information can be extracted from high dimensional data through an elementary process. It further emphasizes that knowledge implicitly incorporated in Act2Vec embeddings can be compactly represented without the need to solve the (many times challenging) task at hand.

5. Related Work

Action Embedding: (Dulac-Arnold et al., 2015) proposed to embed discrete actions into a continuous space. They then find optimal actions using a nearest neighbor approach. They do not, however, offer a method in which such action representations can be found. Most related to our work is that of (Chandak et al., 2019) in which an embedding is used *as part of the policy's structure* in order to train an agent. Our work provides a complementary aspect, with an approach to directly inject prior knowledge from expert data. In addition, we are able to capture semantics without the need to solve the task at hand.

Representation Learning: Representation learning is concerned with finding an appropriate representation of data in order to perform a machine learning task (Goodfellow et al.). In particular, deep learning exploits this concept by its very nature (Mnih et al., 2015). Other work related

to representation in RL include Predictive State Representations (PSR) (Littman & Sutton, 2002), which capture the state as a vector of predictions of future outcomes, and a Heuristic Embedding of Markov Processes (HEMP) (Engel & Mannor, 2001), which learns to embed transition probabilities using an energy-based optimization problem. In contrast to these, actions are less likely to be affected by the “curse of dimensionality” that is inherent in states.

One of the most fundamental work in the field of NLP is word embedding (Mikolov et al., 2013b; Pennington et al., 2014; Zhao et al., 2017), where low-dimensional word representations are learned from unlabeled corpora. Among most word embedding models, Word2Vec (Mikolov et al., 2013b) (trained using SGNS) gains its popularity due to its effectiveness and efficiency. It achieves state-of-the-art performance on a range of linguistic tasks within a fraction of the time needed by previous techniques. In a similar fashion, Act2Vec represents actions by their context. It is able to capture meaningful relations between actions - used to improve RL agents in a variety of tasks.

Learning from Demonstration (LfD): Imitation learning is primarily concerned with matching the performance of a demonstrator (Schaal, 1997; 1999; Argall et al., 2009). Demonstrations typically consist of sequences of state-action pairs $\{(s_0, a_0), \dots, (s_n, a_n)\}$, from which an agent must derive a policy that reproduces and generalizes the demonstrations. While we train Act2Vec using a similar corpus, we do not attempt to generalize the demonstrator’s mapping $\mathcal{S} \rightarrow \mathcal{A}$. Our key claim is that vital information is present in the order in which actions are taken. More specifically, the context of an action masks acceptable forms and manners of usage. These natural semantics cannot be generalized from state-to-action mappings, and can be difficult for reinforcement learning agents to capture. By using finite action contexts we are able to create meaningful representations that capture relations and similarities between actions.

Multi-Task RL: Multitask learning learns related tasks with a shared representation in parallel, leveraging information in related tasks as an inductive bias, to improve generalization, and to help improve learning for all tasks (Caruana, 1997; Ruder, 2017; Taylor & Stone, 2009a;b; Gupta et al., 2017). In our setting, actions are represented using trajectories sampled from permissible policies². These representations advise on correct operations for learning new tasks, though they only incorporate local, relational information. They provide an approach to implicitly incorporate prior knowledge through representation. Act2Vec can thus be used to improve efficiency of multi-task RL methods.

²In our setting, policies are given as optimal solutions to tasks. In practical settings, due to finite context widths, policies need not be optimal in order to capture relevant, meaningful semantics.

Skill Embedding: Concisely representing skills allow for efficient reuse when learning complex tasks (Pastor et al., 2012; Hausman et al.; Kroemer & Sukhatme, 2016). Many methods use latent variables and entropy constraints to decrease the uncertainty of identifying an option, allowing for more versatile solutions (Daniel et al., 2012; End et al., 2017; Gabriel et al., 2017). While these latent representations enhance efficiency, their creation process is dependent on the agent’s ability to solve the task at hand. The benefit of using data generated by human demonstrations is that it lets one learn expressive representations without the need to solve any task. Moreover, much of the knowledge that is implicitly acquired from human trajectories may be unattainable by an RL agent. As an example of such a scenario we depict action embeddings learned from human replays in StarCraft II (see Section 4.3, Figure 5). While up-to-date RL algorithms have yet to overcome the obstacles and challenges in such problems, Act2Vec efficiently captures evident, valuable relations between actions.

6. Discussion and Future Work

If we recognize actions as symbols of a natural language, and regard this language as “expressive”, we imply that it is an instrument of inner mental states. Even by careful introspection, we know little about these hidden mental states, but by regarding actions as thoughts, beliefs, strategies, we limit our inquiry to what is objective. We therefore describe actions as modes of behavior in relation to the other elements in the context of situation.

When provided with a structured embedding space, one can efficiently eliminate actions. When the number of actions is large, a substantial portion of actions can be eliminated due to their proximity to known sub-optimal actions. When the number of actions is substantially small, action sequences can be encoded instead, establishing an extended action space in which similar elimination techniques can be applied.

Recognizing elements of the input as segments of an expressive language allow us to create representations that adhere unique structural characteristics. While the scope of this paper focused on action representation, distributional embedding techniques may also be used to efficiently represent states, policies, or rewards through appropriate contexts. Interpreting these elements relative to their context withholds an array of possibilities for future research.

Lastly, we note that distributional representations can be useful for debugging RL algorithms and their learning process. By visualizing trajectories in the action embedding space, an overseer is able to supervise over an agents progress, finding flaws and improving learning efficiency.

References

- Abbeel, P. and Ng, A. Y. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 1. ACM, 2004.
- Argall, B. D., Chernova, S., Veloso, M., and Browning, B. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- Beattie, C., Leibo, J. Z., Teplyaev, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.
- Caruana, R. Multitask learning. *Machine learning*, 28(1): 41–75, 1997.
- Chandak, Y., Theodorou, G., Kostas, J., Jordan, S., and Thomas, P. S. Learning action representations for reinforcement learning. *arXiv preprint arXiv:1902.00183*, 2019.
- Cheema, S., Gulwani, S., and LaViola, J. Quickdraw: improving drawing experience for geometric diagrams. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1037–1064. ACM, 2012.
- Daniel, C., Neumann, G., and Peters, J. Hierarchical relative entropy policy search. In *Artificial Intelligence and Statistics*, pp. 273–281, 2012.
- Dulac-Arnold, G., Evans, R., van Hasselt, H., Sunehag, P., Lillicrap, T., Hunt, J., Mann, T., Weber, T., Degris, T., and Coppin, B. Deep reinforcement learning in large discrete action spaces. *arXiv preprint arXiv:1512.07679*, 2015.
- End, F., Akrou, R., Peters, J., and Neumann, G. Layered direct policy search for learning hierarchical skills. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pp. 6442–6448. IEEE, 2017.
- Engel, Y. and Mannor, S. Learning embedded maps of markov processes. In *in Proceedings of ICML 2001*. Citeseer, 2001.
- Firth, J. R. A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis*, 1957.
- Gabriel, A., Akrou, R., Peters, J., and Neumann, G. Empowered skills. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pp. 6435–6441. IEEE, 2017.
- Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. *Deep learning*, volume 1.
- Gupta, A., Devin, C., Liu, Y., Abbeel, P., and Levine, S. Learning invariant feature spaces to transfer skills with reinforcement learning. *ICLR*, 2017.
- Hausman, K., Springenberg, J. T., Wang, Z., Heess, N., and Riedmiller, M. Learning an embedding space for transferable robot skills. *ICLR*.
- He, J., Chen, J., He, X., Gao, J., Li, L., Deng, L., and Ostendorf, M. Deep reinforcement learning with a natural language action space. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pp. 1621–1630, 2016.
- Kroemer, O. and Sukhatme, G. S. Learning relevant features for manipulation skills using meta-level priors. *arXiv preprint arXiv:1605.04439*, 2016.
- Levy, O. and Goldberg, Y. Neural word embedding as implicit matrix factorization. In *Advances in neural information processing systems*, pp. 2177–2185, 2014.
- Littman, M. L. and Sutton, R. S. Predictive representations of state. In *Advances in neural information processing systems*, pp. 1555–1561, 2002.
- Maaten, L. v. d. and Hinton, G. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov): 2579–2605, 2008.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013a.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013b.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518 (7540):529, 2015.
- Pastor, P., Kalakrishnan, M., Righetti, L., and Schaal, S. Towards associative skill memories. In *Humanoid Robots (Humanoids), 2012 12th IEEE-RAS International Conference on*, pp. 309–315. IEEE, 2012.
- Pennington, J., Socher, R., and Manning, C. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Ruder, S. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.

- Sapir, E. *An introduction to the study of speech*. Citeseer, 1921.
- Schaal, S. Learning from demonstration. In *Advances in neural information processing systems*, pp. 1040–1046, 1997.
- Schaal, S. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242, 1999.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Taylor, M. E. and Stone, P. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685, 2009a.
- Taylor, M. E. and Stone, P. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685, 2009b.
- Vinyals, O., Ewals, T., Bartunov, S., Georgiev, P., Vezhn-evets, A. S., Yeo, M., Makhzani, A., Küttler, H., Agapiou, J., Schrittwieser, J., et al. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*, 2017.
- Zhao, Z., Liu, T., Li, S., Li, B., and Du, X. Ngram2vec: Learning improved word representations from ngram co-occurrence statistics. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 244–253, 2017.

7. Supplementary Material

7.1. Proof of Lemma 1

We show a more general formulation of the lemma. More specifically, we show that it holds for

$$\pi_K(a|s) = \begin{cases} \mu(a|s) \frac{\sum_{a' \in K} \pi(a'|s)}{\sum_{a' \in K} \mu(a'|s)} & , a \in K \\ \pi(a|s) & , o.w. \end{cases},$$

where μ is any (possibly adversarial) distribution over the set of actions in K . In the case of the lemma, μ is the uniform distribution.

Writing the assumption explicitly we have that for all $a_1, a_2 \in K$

$$\left| \frac{P_{(a_1, s') \sim \mathcal{D}_p}(a_1, s'|s)}{P_{a_1 \sim \mathcal{D}_p}(a_1|s)P_{s' \sim \mathcal{D}_p}(s'|s)} - \frac{P_{(a_2, s') \sim \mathcal{D}_p}(a_2, s'|s)}{P_{a_2 \sim \mathcal{D}_p}(a_2|s)P_{s' \sim \mathcal{D}_p}(s'|s)} \right| < \epsilon,$$

which can be rewritten

$$\left| \frac{P(s'|s, a_1)}{P_{s' \sim \mathcal{D}_p}(s'|s)} - \frac{P(s'|s, a_2)}{P_{s' \sim \mathcal{D}_p}(s'|s)} \right| < \epsilon,$$

giving us

$$|P(s'|s, a_1) - P(s'|s, a_2)| < \epsilon. \quad (6)$$

Next, let \mathcal{M} be the MDP defined by $(\mathcal{S}, \mathcal{A}, P, R, \gamma)$, and let $\hat{\mathcal{M}}$ be the MDP defined by $(\mathcal{S}, \mathcal{A}, \hat{P}, R, \gamma)$, where

$$\hat{P}(s'|s, a) = \begin{cases} \sum_{a' \in K} P(s'|s, a')\mu(a'|s) & , a \in K \\ P(s'|s, a) & , o.w. \end{cases}$$

For all $s \in \mathcal{S}, a \in \mathcal{A}$ we have that

$$\begin{aligned} & \left| P(s'|s, a) - \hat{P}(s'|s, a) \right| \\ & \leq \sup_{a_1 \in K} \left| P(s'|s, a_1) - \sum_{a' \in K} P(s'|s, a')\mu(a'|s) \right| \\ & = \sup_{a_1 \in K} \left| \sum_{a' \in K} (P(s'|s, a_1) - P(s'|s, a'))\mu(a'|s) \right| \\ & \leq \sup_{a_1 \in K} \sum_{a' \in K} |P(s'|s, a_1) - P(s'|s, a')|\mu(a'|s) < \epsilon, \end{aligned}$$

where in the last step we used Equation 6 and the fact that $\sum_{a' \in K} \mu(a'|s) = 1$. For the remainder of the proof we will use the following result proven in (Abbeel & Ng, 2004):

Lemma 2. *Let $\mathcal{M}, \hat{\mathcal{M}}$ be MDPs as defined above. If*

$$\sum_{s'} \left| P(s'|s, a) - \hat{P}(s'|s, a) \right| < \epsilon, \forall s, a$$

then

$$\sum_{s_t} \left| P(s_t|s_0) - \hat{P}(s_t|s_0) \right| < t\epsilon.$$

By definition of $\hat{\mathcal{M}}$ and π_K we have that $\hat{P}^\pi = P^{\pi_K}$. Then, for all $s \in \mathcal{S}$

$$\begin{aligned} & |V^\pi(s) - V^{\pi_K}(s)| \\ &= \left| \mathbb{E}_{P^\pi} \left(\sum_{t=0}^{\infty} \gamma^t r(s_t) \middle| s_0 = s \right) - \mathbb{E}_{\hat{P}^\pi} \left(\sum_{t=0}^{\infty} \gamma^t r(s_t) \middle| s_0 = s \right) \right| \\ &= \left| \sum_{t=0}^{\infty} \gamma^t [\mathbb{E}_{P^\pi}(r(s_t)|s_0 = s) - \mathbb{E}_{\hat{P}^\pi}(r(s_t)|s_0 = s)] \right| \\ &\leq \sum_{t=0}^{\infty} \gamma^t |\mathbb{E}_{P^\pi}(r(s_t)|s_0 = s) - \mathbb{E}_{\hat{P}^\pi}(r(s_t)|s_0 = s)|. \end{aligned}$$

Writing the above explicitly we get

$$|V^\pi(s) - V^{\pi_K}(s)| \leq \sum_{t=0}^{\infty} \gamma^t \left| \sum_s r(s) (P(s_t|s_0 = s) - \hat{P}(s_t|s_0 = s)) \right|. \quad (7)$$

Next denote the sets

$$\begin{aligned} Q_t &= \{(s_t) : P(s_t|s_0 = s) > \hat{P}(s_t|s_0 = s)\}, \text{ and} \\ Q_t^c &= \{(s_t) : P(s_t|s_0 = s) \leq \hat{P}(s_t|s_0 = s)\}. \end{aligned}$$

Then,

$$\begin{aligned} & \sum_{s_t} r(s_t) (P(s_t|s_0 = s) - \hat{P}(s_t, a)) \\ &= \sum_{s_t \in Q_t} r(s_t) (P(s_t|s_0 = s) - \hat{P}(s_t|s_0 = s)) + \sum_{s_t \in Q_t^c} r(s_t) (P(s_t|s_0 = s) - \hat{P}(s_t|s_0 = s)) \\ &\leq \sum_{s_t \in Q_t} r(s_t) (P(s_t|s_0 = s) - \hat{P}(s_t|s_0 = s)) \\ &= \sum_{s_t \in Q_t} r(s_t) |P(s_t|s_0 = s) - \hat{P}(s_t|s_0 = s)| \\ &\leq \sum_{s_t} r(s_t) |P(s_t|s_0 = s) - \hat{P}(s_t|s_0 = s)| \\ &\leq \sum_{s_t} |P(s_t|s_0 = s) - \hat{P}(s_t|s_0 = s)|. \end{aligned}$$

By switching the roles P_t and \hat{P}_t we get the opposite inequality:

$$\sum_{s_t} r(s_t) (\hat{P}(s_t|s_0 = s) - P(s_t|s_0 = s)) \leq \sum_{s_t} |\hat{P}(s_t|s_0 = s) - P(s_t|s_0 = s)|.$$

Overall,

$$\sum_{s_t} r(s_t) |\hat{P}(s_t|s_0 = s) - P(s_t|s_0 = s)| \leq \sum_{s_t} |\hat{P}(s_t|s_0 = s) - P(s_t|s_0 = s)|. \quad (8)$$

Plugging (8) in (7) and using Lemma 2 we get

$$|V^\pi(s) - V^{\pi_K}(s)| \leq \sum_{t=0}^{\infty} \gamma^t \sum_{s_t} |\hat{P}(s_t|s_0 = s) - P(s_t|s_0 = s)| \leq \sum_{t=0}^{\infty} \gamma^t t \epsilon.$$

The proof follows immediately due to

$$\sum_{t=0}^{\infty} \gamma^t t = \gamma \frac{1 + 4\gamma + \gamma^2}{(1 - \gamma)^4} \leq \frac{6\gamma}{(1 - \gamma)^4}$$

for $|\gamma| < 1$.

Algorithm 1 Q -Embedding with k -Exp

Create Act2Vec Embeddings \mathcal{E} for actions in \mathcal{A}
 Create k clusters $\{C_i\}_{i=1}^k$ based on \mathcal{E} using k -means
 Initialize replay memory \mathcal{R}
for episode = 1 **to** m **do**
 Reset simulator
 for $t = 1$ **to** T **do**
 $Q(s_t, a_t^i; \theta) \leftarrow \psi_\theta(\mathcal{E}(a_t^i))^T \phi_\theta(s_t)$
 Sample U uniformly in $[0, 1]$
 if $U < \epsilon$ **then**
 Sample i uniformly in $\{1, \dots, k\}$
 Sample action a uniformly in C_i
 else
 Choose action a which maximizes $Q(s_t, a; \theta)$
 end if
 Execute action a in simulator.
 Observe reward r_t and next state s_{t+1}
 Store transition (s_t, a_t, r_t, s_{t+1}) in \mathcal{R} .
 Sample random minibatch of transitions (s_j, a_j, r_j, s_{j+1}) from \mathcal{R}
 $y_j \leftarrow \begin{cases} r_j & s_{j+1} \text{ is terminal} \\ r_j + \gamma \max_{a' \in \mathcal{A}} Q(s_{j+1}, a'; \theta) & \text{o.w.} \end{cases}$
 Perform gradient descent step on $(y_j - Q(s_j, a_j; \theta))^2$ w.r.t. network parameters θ
 end for
end for

7.2. Experimental Details

7.2.1. DRAWING

We used the *Quick, Draw!* (Cheema et al., 2012) dataset, consisting of 50 million drawings ranging over 345 categories. While the raw data describes strokes as relative positions of a user’s pen on the canvas, we simplified the action space to include four primitive actions: left, right, up, down, i.e., movement of exactly one pixel in each of the directions. This was done by connecting lines between any two consecutive points in a stroke, and ignoring any drawings consisting of multiple strokes. Finally, we defined our action space to be any sequence of length k primitive actions (a_1, \dots, a_k) , which we will hereinafter denote as an action stroke.

For our task, we were concerned with drawing a square. We used the 70,000 human-drawn squares in the “square” category of the *Quick, Draw!* dataset. Scanning the corpus of square drawings we trained Act2Vec on 12 types of actions strokes:

1. Strokes of length 20 pixels in each of the axis directions: Left \times 20, Right \times 20, Up \times 20, Down \times 20.
2. Strokes of length 20 pixels consisting of corners in each of the possible directions: Left \times 10+Up \times 10, Up \times 10+Right \times 10, Right \times 10+Down \times 10, Down \times 10+Left \times 10, Right \times 10+Up \times 10, Up \times 10+Left \times 10, Left \times 10+Down \times 10, Down \times 10+Right \times 10.

Act2Vec was trained with SGNS over 50 epochs using a window size of $w = 2$, embedding dimension $d = 5$, and $k = 5$ negative samples.

The reward function was given by $R = \frac{\sum_{i=1}^4 \min(W, l_i)}{4W}$, where l_i denotes the length of the i^{th} side of the drawn shape, and W is the desired length of each side. When the drawn shape had more or less than four sides, a reward of -0.1 was given. In our case we tested values of $W = 40, 60$.




Figure 6. Images show the 2d navigation (a) and 3d Seek-Avoid (b) environments used in Section 4.2. Action embeddings learned in the 2d environment were used as knowledge transfer to improve the solution of the task in the 3d environment.

The state was represented as the sum of embeddings of previously taken actions. That is,

$$s_t = \sum_{i=0}^{t-1} \mathcal{E}(a_i),$$

where \mathcal{E} are the used embeddings. We tested four types of embeddings: Act2Vec, Normalized Act2Vec (by l_2 norm), one-hot embedding, and random embedding. The one-hot embedding is defined as the unit vector e_i of length $|\mathcal{A}|$, with zeros everywhere except for position i , corresponding to action a_i . Random embeddings were initialized uniformly in $[0, 1]^d$.

The agent was trained using Proximal Policy Optimization (PPO) (Schulman et al., 2017) over 1 million iterations for the case of $W = 40$ and 10 million for $W = 60$ using 256 asynchronous threads. PPO was run using the following parameters: learning rate = 0.0005, $\lambda = 0.95$, $\gamma = 0.99$, 3 experience buffer passes, entropy coefficient = 0.01, clipping range = 0.1. The value and policy networks were both represented as feedforward neural networks with two hidden layer with 64 units each. Each of the state representation methods were tested for a set of 15 trials, except for random embeddings, which were tested for a set of 100 trials. A different random embeddings was sampled for each tested trial.

7.2.2. NAVIGATION

We start by describing the 2d environment in which the action embeddings where learned. An image of the environment can be seen in Figure 6a. The environment was divided to a 25x25 grid, in which 300 walls were randomly placed as well as 5 numbers relating to goals. The player was initialized in the center of the grid. The player’s movement was not constrained to the grid, but rather a more continuous movement - similar to that of the Seek-Avoid environment. The player was given the following three actions: Move forward one unit, rotate left 25 degrees, and rotate right 25 degrees. The player was given the task to move from its initial position and reach each of the goal numbers (in order).

Actions of a demonstrator player were recorded for different initializations of the 2d environment. Act2Vec was trained over a corpus of 3000 actions. Act2Vec was trained with SGNS over 50 epochs using different window sizes of $w = 2, 3$, embedding dimensions $d = 3, 5$, and $k = 5$ negative samples.

The 3d environment used to test the DQN agent was the Seek-Avoid lab (Beattie et al., 2016). An image of the environment can be seen in Figure 6b. The input image was resized to a resolution of 80×80 pixels with 3 channels (RGB). All experiments used the same state representation: a convolutional neural network with 3 layers: (1) 2d convolution with 8 channels, window size of 3×3 and stride of 2, (2) 2d convolution with 16 channels, window size of 3×3 , and stride of 2, and (3) fully connected layer with 128 hidden features, $\phi_\theta(s)$. For the case of the traditional DQN, the network output was $w_a^T \phi_\theta(s)$, where w_a were the network’s parameters learned for each action. For the case of Q -Embedding, we used another network to approximate w_a as $\psi_\theta(\mathcal{E}(a))$, where $\mathcal{E}(a)$ is the embedding of action a , and ψ_θ is a feedforward neural network with one hidden layer of 64 units and \tanh activation function, and 128 outputs. In this case, the network output was given by $\psi_\theta(\mathcal{E}(a))^T \phi_\theta(s)$, for each of the actions.

The action space consisted of: move forward one unit, turn view 25 degrees to the right, and turn view 25 degrees to the left. Every action was repeated for a sequence of 10 frames. In the case of action sequences, every action in the sequence was repeated for 10 frames, i.e., a sequence of k actions was executed for $10k$ frames.

The agent used a replay memory of one million transitions and a training batch of 64 transitions. The agent was trained over 60000 iterations with a learning rate of 0.00025, and discount factor of $\gamma = 0.99$. For uniform and k -Exp exploration decaying ϵ was used, starting at a value of 1 and decaying linearly to a value of 0.1 once reached to iteration 19800.

7.2.3. Q-EMBEDDING ALGORITHM

The Q -Embedding with k -Exp algorithm is presented in Algorithm 1. An NLP version of this algorithm (with uniform exploration) has already been proposed in (He et al., 2016).