

What Can You Do with a Rock?

Affordance Extraction via Word Embeddings

Nancy Fulda, Daniel Ricks, Ben Murdoch and David Wingate
 Brigham Young University

nfulda@byu.edu, daniel_ricks@byu.edu, murdoch@byu.edu, wingated@byu.edu

Abstract

Autonomous agents must often detect *affordances*: the set of behaviors enabled by a situation. Affordance detection is particularly helpful in domains with large action spaces, allowing the agent to prune its search space by avoiding futile behaviors. This paper presents a method for affordance extraction via word embeddings trained on a tagged Wikipedia corpus. The resulting word vectors are treated as a common knowledge database which can be queried using linear algebra. We apply this method to a reinforcement learning agent in a text-only environment and show that affordance-based action selection improves performance in most cases. Our method increases the computational complexity of each learning step but significantly reduces the total number of steps needed. In addition, the agent’s action selections begin to resemble those a human would choose.

1 Introduction

The physical world is filled with constraints. You can open a door, but only if it isn’t locked. You can douse a fire, but only if a fire is present. You can throw a rock or drop a rock or even, under certain circumstances, converse with a rock, but you cannot traverse it, enumerate it, or impeach it. The term *affordances* [Gibson, 1977] refers to the subset of possible actions which are feasible in a given situation. Human beings detect these affordances automatically, often subconsciously, but it is not uncommon for autonomous learning agents to attempt impossible or even ridiculous actions, thus wasting effort on futile behaviors.

This paper presents a method for affordance extraction based on the copiously available linguistic information in online corpora. Word embeddings trained using Wikipedia articles are treated as a common sense knowledge base that encodes (among other things) object-specific affordances. Because knowledge is represented as vectors, the knowledge base can be queried using linear algebra. This somewhat counterintuitive notion - the idea that words can be manipulated mathematically - creates a theoretical bridge between the frustrating realities of real-world systems and the im-

mense wealth of common sense knowledge implicitly encoded in online corpora.

We apply our technique to a text-based environment and show that *a priori* knowledge provided by affordance extraction greatly speeds learning. Specifically, we reduce the agent’s search space by (a) identifying actions afforded by a given object; and (b) discriminating objects that can be grasped, lifted and manipulated from objects which can merely be observed. Because the agent explores only those actions which ‘make sense’, it is able to discover valuable behaviors more quickly than a comparable agent using a brute force approach. Critically, the affordance agent is demonstrably able to eliminate extraneous actions without (in most cases) discarding beneficial ones.

All the experiments in this paper were performed using a Wikipedia corpus that had been preprocessed to append part-of-speech tagging to each word.

2 Related Work

Our research relies heavily on word2vec [Mikolov *et al.*, 2013a], an algorithm that encodes individual words based on the contexts in which they tend to appear. Earlier work has shown that word vectors trained using this method contain intriguing semantic properties, including structured representations of gender and geography [Mikolov *et al.*, 2013b; 2013c]. The (by now) archetypal example of such properties is represented by the algebraic expression $vector[‘king’] - vector[‘man’] + vector[‘woman’] = vector[‘queen’]$.

Researchers have leveraged these properties for diverse applications including sentence- and paragraph-level encoding [Kiros *et al.*, 2015; Le and Mikolov, 2014], image categorization [Frome *et al.*, 2013], bidirectional retrieval [Karpathy *et al.*, 2014], semantic segmentation [Socher *et al.*, 2011], biomedical document retrieval [Brokos *et al.*, 2016], and the alignment of movie scripts to their corresponding source texts [Zhu *et al.*, 2015]. Our work is most similar to [Zhu *et al.*, 2014]; however, rather than using a Markov Logic Network to build an explicit knowledge base, we instead rely on the semantic structure implicitly encoded in skip-grams.

Affordance detection, a topic of rising importance in our increasingly technological society, has been attempted and/or accomplished using visual characteristics [Song *et al.*, 2011; 2015], haptic data [Navarro *et al.*, 2012], visuomotor simulation [Schenck *et al.*, 2012; 2016], repeated real-world ex-

perimentation [Montesano *et al.*, 2007; Stoytchev, 2008], and knowledge base representations [Zhu *et al.*, 2014].

In 2001 [Laird and van Lent, 2001] identified text-based adventure games as a step toward general problem solving. The same year at AAAI, Mark DePristo and Robert Zubek unveiled a hybrid system for text-based game play [Arkin, 1998], which operated on hand-crafted logic trees combined with a secondary sensory system used for goal selection. The handcrafted logic worked well, but goal selection broke down and became cluttered due to the scale of the environment. Perhaps most notably, in 2015 [Narasimhan *et al.*, 2015] designed an agent which passed the text output of the game through an LSTM [Hochreiter and Schmidhuber, 1997] to find a state representation, then used a DQN [Mnih *et al.*, 2015] to select a Q-valued action. This approach appeared to work well within a small discrete environment with reliable state action pairs, but as the complexity and alphabet of the environment grew, the clarity of Q-values broke down and left them with a negative overall reward. Our work, in contrast, is able to find meaningful state action pairs even in complex environments with many possible actions.

3 Wikipedia as a Common Sense Knowledge Base

Google ‘knowledge base’, and you’ll get a list of hand-crafted systems, both commercial and academic, with strict constraints on encoding methods. These highly-structured, often node-based solutions are successful at a wide variety of tasks including topic gisting [Liu and Singh, 2004], affordance detection [Zhu *et al.*, 2014] and general reasoning [Russ *et al.*, 2011]. Traditional knowledge bases are human-interpretable, closely tied to high-level human cognitive functions, and able to encode complex relationships compactly and effectively.

It may seem strange, then, to treat Wikipedia as a knowledge base. When compared with curated solutions like ConceptNet [Liu and Singh, 2004], Cyc [Matuszek *et al.*, 2006], and WordNet [Miller, 1995], its contents are largely unstructured, polluted by irrelevant data, and prone to user error. When used as a training corpus for the word2vec algorithm, however, Wikipedia becomes more tractable. The word vectors create a compact representation of the knowledge base and, as observed by [Bolukbasi *et al.*, 2016a] and [Bolukbasi *et al.*, 2016b], can even encode relationships about which a human author is not consciously cognizant. Perhaps most notably, Wikipedia and other online corpora are constantly updated in response to new developments and new human insight; hence, they do not require explicit maintenance.

However: in order to leverage the semantic structure implicitly encoded within Wikipedia, we must be able to interpret the resulting word vectors. Significant semantic relationships are not readily apparent from the raw word vectors or from their PCA reduction. In order to extract useful information, the database must be queried through a mathematical process. For example, in Figure 1 a dot product is used to project gendered terms onto the space defined by $vector[‘king’] - vector[‘queen’]$ and $vector[‘woman’] - vector[‘man’]$. In such a projection, the mathematical relationship between the words is readily apparent. Masculine

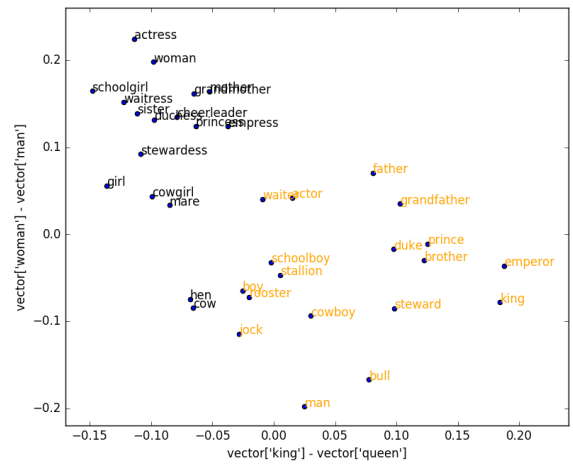


Figure 1: Word vectors projected into the space defined by $vector[‘king’] - vector[‘queen’]$ and $vector[‘woman’] - vector[‘man’]$. In this projection, masculine and feminine terms are linearly separable.

and feminine terms become linearly separable, making it easy to distinguish instances of each group.

These relationships can be leveraged to detect affordances, and thus reduce the agent’s search space. In its most general interpretation, the adjective *affordant* describes the set of actions which are **physically possible** under given conditions. In the following subsections, however, we use it in the more restricted sense of actions which **seem reasonable**. For example, it is physically possible to eat a pencil, but it does not ‘make sense’ to do so.

3.1 Verb/Noun Affordances

So how do you teach an algorithm what ‘makes sense’? We address this challenge through an example-based query. First we provide a canonical set of verb/noun pairs which illustrate the relationship we desire to extract from the knowledge base. Then we query the database using the analogy format presented by [Mikolov *et al.*, 2013a]. Using their terminology, the analogy $sing:song::[?]:[x]$ encodes the following question: If the affordant verb for ‘song’ is ‘sing’, then what is the affordant verb for [x]?

In theory, a single canonical example is sufficient to perform a query. However, experience has shown that results are better when multiple canonical values are averaged.

More formally, let W be the set of all English-language word vectors in our agent’s vocabulary. Further, let $N = \{\vec{n}_1, \dots, \vec{n}_j\}, N \subset W$ be the set of all nouns in W and let $V = \{\vec{v}_1, \dots, \vec{v}_k\}, V \subset W$ be the set of all verbs in W .

Let $C = \{(\vec{v}_1, \vec{n}_1), \dots, (\vec{v}_m, \vec{n}_m)\}$ represent a set of canonical verb/noun pairs used by our algorithm. We use C to define an affordance vector $\vec{a} = 1/m \sum_i (\vec{v}_i - \vec{n}_i)$, which can be thought of as the distance and direction within the embedding space which encodes affordant behavior.

In our experiments we used the following verb/noun pairs as our canonical set:

[‘sing song’, ‘drink water’, ‘read book’, ‘eat food’,

Algorithm 1 Noun Selection With Affordance Detection

```

1: state = game response to last command
2: manipulable_nouns ← {}
3: for each word w ∈ state do
4:   if w is a noun then
5:     if w is manipulable then
6:       add w to manipulable_nouns
7:     end if
8:   end if
9: end for
10: noun = a randomly selected noun from manipulable_nouns
    
```

Algorithm 2 Verb Selection With Analogy Reduction

```

1: navigation_verbs = ['north', 'south', 'east', 'west', 'northeast', 'southeast',
'southwest', 'northwest', 'up', 'down', 'enter']
2: manipulation_verbs = a list of 1000 most common verbs
3: essential_manipulation_verbs = ['get', 'drop', 'push', 'pull', 'open',
'close']
4: affordant_verbs = verbs returned by Word2vec that match noun
5: affordant_verbs = affordant_verbs ∩
manipulation_verbs
6: final_verbs = navigation_verbs ∪ affordant_verbs ∪
essential_manipulation_verbs
7: verb = a randomly selected verb from final_verbs
    
```

respond with a written phrase. Typical actions include commands such as: ‘examine wallet’, ‘eat apple’, or ‘light campfire with matches’. The game engine parses this response and produces a new block of text. The resulting interactions, although syntactically simple, provide a fertile research environment for natural language processing and human/computer interaction. Game players must identify objects that are manipulable and apply appropriate actions to those objects in order to make progress.

In these games, the learning agent faces a frustrating dichotomy: its action set must be large enough to accommodate any situation it encounters, and yet each additional action increases the size of its search space. A brute force approach to such scenarios is frequently futile, and yet factorization, function approximation, and other search space reduction techniques bring the risk of data loss. We desire an agent that is able to clearly perceive all its options, and yet applies only that subset which is likely to produce results.

In other words, we want an agent that explores the game world the same way a human does: by trying only those actions that ‘make sense’. In the following sections, we show that affordance-based action selection provides a meaningful first step towards this goal.



Figure 4: Sample text from the adventure game Zork. Player responses follow a single angle bracket.

4.1 Learning Algorithm

Our agent utilizes Q-learning [Watkins and Dayan, 1992], a reinforcement learning algorithm which attempts to maximize expected discounted reward. Q-values are updated according to the equation

$$\Delta Q_{\pi}(s, a) = \alpha(R(s, a) + \gamma \max_a Q_{\pi}(s', a) - Q_{\pi}(s, a)) \quad (1)$$

where $Q_{\pi}(s, a)$ is the expected cumulative reward for performing action a in observed state s and following the policy π thereafter, α is the learning rate, γ is the discount factor, and s' is the new state observation after performing action a . All experiments in this paper used $\alpha=0.1$ and $\gamma=0.95$

The agent’s state representation is encoded as a hash of the text provided by the game engine. Actions are comprised of verb/object pairs:

$$a = v + ' ' + o, v \in V, o \in O \quad (2)$$

where V is the set of all English-language verbs and O is the set of all English-language nouns. To enable the agent to distinguish between state transitions and merely informational feedback, the agent executes a ‘look’ command every second iteration and assumes that the resulting game text represents its new state. Some games append a summary of actions taken and points earned in response to each ‘look’ command. To prevent this from obfuscating the state space, we stripped all numerals from the game text prior to hashing.

Given that the English language contains at least 20,000 verbs and 100,000 nouns in active use, a naive application of Q-learning is intractable. Some form of action-space reduction must be used. For our baseline comparison, we use an agent with a vocabulary consisting of the 1000 most common verbs in Wikipedia plus an 11-word navigation list and a 6-word essential manipulation list as depicted in Algorithm 2. The navigation list contains words which, by convention, are used to navigate through text-based games. The essential manipulation list contains words which, again by convention, are generally applicable to all in-game objects.

The baseline agent does not use a fixed noun vocabulary. Instead, it extracts nouns from the game text using part-of-speech tags. To facilitate game interactions, the baseline agent augments its noun list using adjectives that precede them. For example, if the game text consisted of ‘You see a red pill and a blue pill’, then the agent’s noun list for that state would be [‘pill’, ‘red pill’, ‘blue pill’]. (And its next action is hopefully ‘swallow red pill’).

In Sections 5.1 and 5.2 the baseline agent is contrasted with an agent using affordance extraction to reduce its manipulation list from 1000 verbs to a mere 30 verbs for each state, and to reduce its object list to a maximum of 15 nouns per state. We compare our approach to other search space reduction techniques and show that the *a priori* knowledge provided by affordance extraction enables the agent to achieve results which cannot be paralleled through brute force methods. All agents used epsilon-greedy exploration with a decaying epsilon.

The purpose of our research was to test the value of affordance-based search space reduction. Therefore, we did not add augmentations to address some of the more challenging aspects of text-based adventure games. Specifically, the

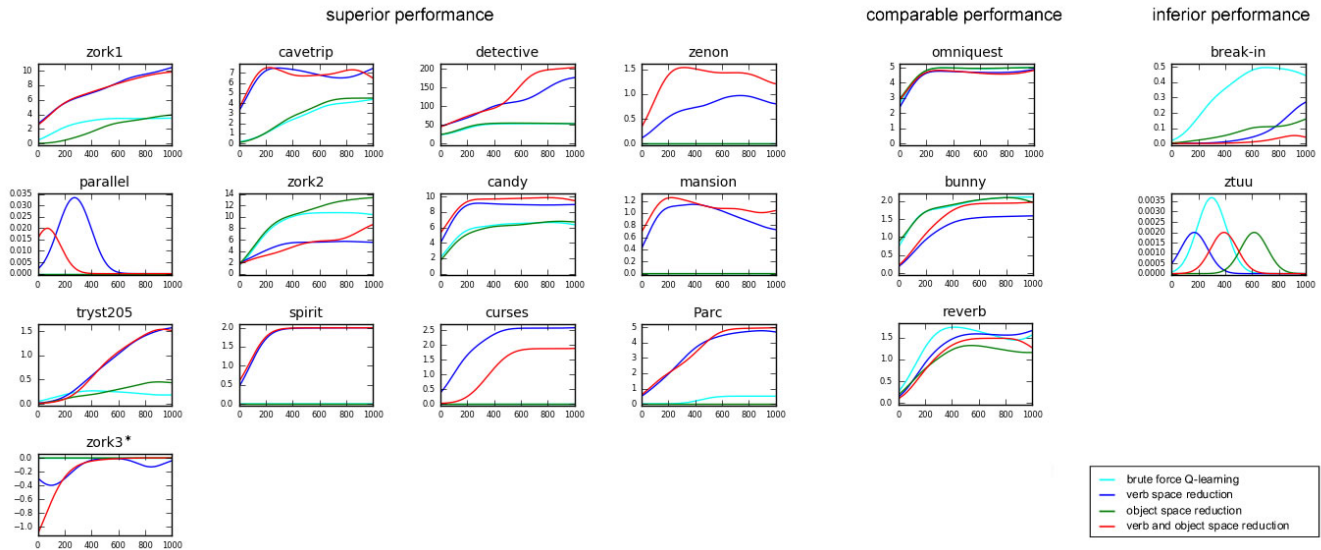


Figure 5: Learning trajectories for eighteen Z-machine games. No agent scored any points on the remaining 31 games. Average of 10 data runs. *In zork3, standard smoothing obscured small, sporadic positive rewards obtained by our reduction algorithms.

agent maintained no representation of items carried in inventory or of the game score achieved thus far. The agent was also not given the ability to construct prepositional commands such as ‘put book on shelf’ or ‘slay dragon with sword’. Our algorithm does not account for physical or temporal dependencies in which the possession of an item or the execution of a specific action in the past affect the outcome of the current behaviors.

5 Results

We tested our agent on a suite of 50 text-based adventure games compatible with Infocom’s Z-machine. These games represent a wide variety of situations, ranging from business scenarios like ‘Detective’ to complex fictional worlds like ‘Zork: The Underground Empire’. The core objective of the games vary widely. In some games, the player’s goal is simply to survive, while in others specific items must be acquired. Significantly, the games provide little or no information about the agent’s goals or which actions might provide reward.

During training, the agent interacted with the game engine for 1000 epochs with 1000 training steps per epoch. On each training step, the agent received a reward corresponding to the change in game score. After each epoch the game was restarted and the game score reset, but the agent retained its learned Q-values. Some games, notably zork3, allowed the player to ‘die’ as play advanced, triggering a negative reward and prematurely terminating the training epoch.

One game experienced a scoring malfunction and was excluded from subsequent analysis; however, our verb-space reduction method obtained a score identical to the baseline agent’s in the discarded partial results.

Our affordance-based search space reduction algorithms enabled the agent to score points on 18/49 games, with a peak performance (expressed as a percentage of possible game

score) of 57.44% for verb space reduction, 23.90% for object space reduction, and 56.84% when both methods were combined. The baseline agent (see Sec. 4.1) scored points on 12/49 games, with a peak performance of 20.84%. Peak performance is defined as the maximum score achieved over all epochs, a metric that expresses the agent’s ability to comb through the search space and discover areas of high reward.

Figures 5 and 7 show the performance of our reduction techniques when compared to the baseline. Affordance-based search space reduction improved overall performance on 13/18 games, and decreased performance on only 2 games.

Examination of the 31 games in which no agent scored points (and which are correspondingly not depicted in Figures 5 and 7) revealed three prevalent failure modes: (1) The game required prepositional commands such as ‘look at machine’ or ‘give dagger to wizard’, (2) The game provided points only after an unusually complex sequence of events, (3) The game required the user to infer the proper term for manipulable objects. (For example, the game might describe ‘**something** shiny’ at the bottom of a lake, but required the agent to ‘get shiny **object**’.) Our test framework was not designed to address these issues, and hence did not score points on those games. A fourth failure mode (4) might be the absence of a game-critical verb within the 1000-word manipulation list. However, this did not occur in our coarse examination of games that failed.

5.1 Alternate Reduction Methods

We compared our affordance-based reduction technique with four other approaches that seemed intuitively applicable to the test domain. Results are shown in Figure 7.

Intrinsic rewards: This approach guides the agent’s exploration of the search space by allotting a small reward each time a new state is attained. The agent is thus encouraged to traverse multiple game locations, enabling possibilities for

Affordant selection	Random selection
decorate glass	continue quantity
open window	break sack
add table	result window
generate quantity	stay table
ring window	build table
weld glass	end house
travel passage	remain quantity
climb staircase	discuss glass
jump table	passage

Figure 6: Sample exploration actions produced by a Q-learner with and without affordance detection. The random agent used nouns extracted from game text and a verb list comprising the 200 most common verbs in Wikipedia.

higher reward. We call these awards intrinsic because they are tied to the agent’s assessment of its progress rather than to external events.

Random reduction: When applying search space reductions one must always ask: ‘Did improvements result from my specific choice of reduced space, or would *any* reduction be equally effective?’ We address this question by randomly selecting 30 manipulation verbs to use during each epoch.

ConceptNet reduction: In this approach we used ConceptNet’s CapableOf relation to obtain a list of verbs relevant to the current object. We then reduced the agent’s manipulation list to include only words that were also in ConceptNet’s word list (effectively taking the intersection of the two lists).

Co-occurrence reduction: In this method, we populated a co-occurrence dictionary prior to training using the 1000 most common verbs and 30,000 most common nouns in Wikipedia. The dictionary tracked the number of times each verb/noun pair occurred within a 9-word radius of each other. During game play, the agent’s manipulation list was reduced to include only words which exceeded a low threshold (co-occurrences > 3).

Figure 7 shows the performance of these four algorithms, along with a baseline learner using a 1000-word manipulation list, on the sixteen games presented in Figure 5. Affordance-based verb selection improved performance in most games, but the other reduction techniques fell prey to a classic danger: they pruned precisely those actions which were essential to obtain reward.

5.2 Fixed-length Vocabularies vs. Free-form Learning

An interesting question arises from our research. What if, rather than beginning with a 1000-word vocabulary, the agent was free to search the entire English-language verb space?

A traditional learning agent could not do this: the space of possible verbs is too large. However, the Wikipedia knowledge base opens new opportunities. Using the action selection mechanism described in Section 4.1, we allowed the agent to construct its own manipulation list for each state (see Section 3.1). The top 15 responses were unioned with the agent’s navigation and essential manipulation lists, with actions selected randomly from that set.

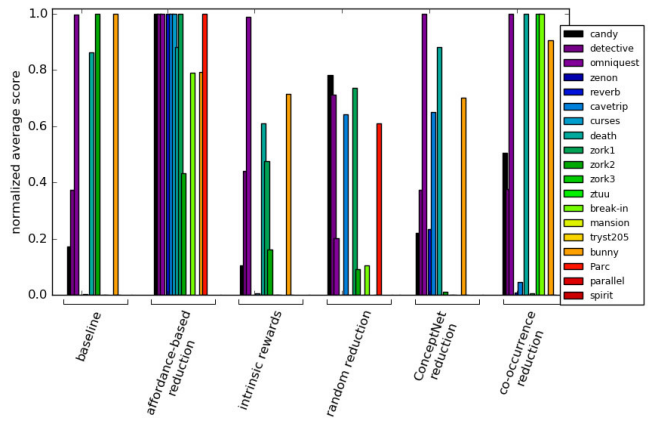


Figure 7: Five verb space reduction techniques compared over 100 exploration epochs. Average of 5 data runs. Results were normalized for each game based on the maximum reward achieved by any agent.

A sampling of the agent’s behavior is displayed in Figure 6, along with comparable action selections from the baseline agent described in Section 4.1. The free-form learner is able to produce actions that seem, not only reasonable, but also rather inventive when considered in the context of the game environment. We believe that further research in this direction may enable the development of one-shot learning for text-based adventure games.

6 Conclusion

The common sense knowledge implicitly encoded within Wikipedia and other online corpora opens new opportunities for autonomous agents. In this paper we have shown that previously intractable search spaces can be efficiently navigated when word embeddings are used to identify context-dependent affordances. We have also shown that, in the domain of text-based adventure games, this approach is superior to several other intuitive methods.

We emphasize that, although our initial experiments have been restricted to text-based environments, the underlying principles apply to any domain in which mappings can be formed between words and objects. Steady advances in the fields of object recognition and semantic segmentation, combined with improved precision in robotic systems, suggests that our methods are readily applicable to systems including self-driving cars, domestic robots, and UAVs.

Acknowledgements

Our experiments were run using Autoplay: a learning environment for interactive fiction (<https://github.com/danielricks/autoplay>). We thank Nvidia, the Center for Unmanned Aircraft Systems, and Analog Devices, Inc. for their generous support.

References

- [Arkin, 1998] Ronald C. Arkin. *Behavior-Based Robotics*. MIT Press, 1998.
- [Bolukbasi *et al.*, 2016a] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *NIPS*, pages 4349–4357. Curran Associates, Inc., 2016.
- [Bolukbasi *et al.*, 2016b] Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam Tauman Kalai. Quantifying and reducing stereotypes in word embeddings. *CoRR*, abs/1606.06121, 2016.
- [Brokos *et al.*, 2016] Georgios-Ioannis Brokos, Prodromos Malakasiotis, and Ion Androutsopoulos. Using centroids of word embeddings and word mover’s distance for biomedical document retrieval in question answering. *CoRR*, abs/1608.03905, 2016.
- [Frome *et al.*, 2013] Andrea Frome, Greg S. Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *In NIPS*, 2013.
- [Gibson, 1977] James J. Gibson. The theory of affordances. In Robert Shaw and John Bransford, editors, *Perceiving, Acting, and Knowing*. 1977.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [Karpathy *et al.*, 2014] Andrej Karpathy, Armand Joulin, and Li Fei-fei. Deep fragment embeddings for bidirectional image sentence mapping. In *In arXiv:1406.5679*, 2014.
- [Kiros *et al.*, 2015] Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. Skip-thought vectors. *CoRR*, abs/1506.06726, 2015.
- [Laird and van Lent, 2001] John E. Laird and Michael van Lent. Human-level AI’s killer application: Interactive computer games. *AI Magazine*, 22(2):15–26, 2001.
- [Le and Mikolov, 2014] Quoc V. Le and Tomas Mikolov. Distributed representations of sentences and documents. *CoRR*, abs/1405.4053, 2014.
- [Liu and Singh, 2004] H. Liu and P. Singh. Conceptnet — a practical commonsense reasoning tool-kit. *BT Technology Journal*, 22(4):211–226, 2004.
- [Matuszek *et al.*, 2006] Cynthia Matuszek, John Cabral, Michael Witbrock, and John Deoliveira. An introduction to the syntax and content of cyc. In *Proceedings of the 2006 AAI Spring Symposium on Formalizing and Compiling Background Knowledge and Its Applications to Knowledge Representation and Question Answering*, pages 44–49, 2006.
- [Mikolov *et al.*, 2013a] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781, 2013.
- [Mikolov *et al.*, 2013b] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *NIPS*, pages 3111–3119. Curran Associates, Inc., 2013.
- [Mikolov *et al.*, 2013c] Tomas Mikolov, Wen tau Yih, and Geoffrey Zweig. Linguistic regularities in continuous space word representations. Association for Computational Linguistics, May 2013.
- [Miller, 1995] George A. Miller. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, November 1995.
- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [Montesano *et al.*, 2007] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor. Modeling affordances using bayesian networks. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4102–4107, Oct 2007.
- [Narasimhan *et al.*, 2015] Karthik Narasimhan, Tejas D. Kulkarni, and Regina Barzilay. Language understanding for text-based games using deep reinforcement learning. *CoRR*, abs/1506.08941, 2015.
- [Navarro *et al.*, 2012] Stefan Escaida Navarro, Nicolas Gorges, Heinz Wörn, Julian Schill, Tamim Asfour, and Rüdiger Dillmann. Haptic object recognition for multi-fingered robot hands. In *2012 IEEE Haptics Symposium (HAPTICS)*, pages 497–502. IEEE, 2012.
- [Russ *et al.*, 2011] Thomas A Russ, Cartic Ramakrishnan, Eduard H Hovy, Mihail Bota, and Gully APC Burns. Knowledge engineering tools for reasoning with scientific observations and interpretations: a neural connectivity use case. *BMC bioinformatics*, 12(1):351, 2011.
- [Schenck *et al.*, 2012] Wolfram Schenck, Hendrik Hasenbein, and Ralf Möller. Detecting affordances by mental imagery. In Alessandro G. Di Nuovo, Vivian M. de la Cruz, and Davide Marocco, editors, *Proceedings of the SAB Workshop on “Artificial Mental Imagery”*, pages 15–18, Odense (Denmark), 2012.
- [Schenck *et al.*, 2016] Wolfram Schenck, Hendrik Hasenbein, and Ralf Möller. Detecting affordances by visuomotor simulation. *arXiv preprint arXiv:1611.00274*, 2016.
- [Socher *et al.*, 2011] Richard Socher, Cliff C. Lin, Chris Manning, and Andrew Y Ng. Parsing natural scenes and natural language with recursive neural networks. *ICML*, pages 129–136, 2011.
- [Song *et al.*, 2011] Hyun Oh Song, Mario Fritz, Chunhui Gu, and Trevor Darrell. Visual grasp affordances from appearance-based cues. In *ICCV Workshops*, pages 998–1005. IEEE, 2011.
- [Song *et al.*, 2015] Hyun Oh Song, Mario Fritz, Daniel Goehring, and Trevor Darrell. Learning to detect visual grasp affordance. In *IEEE Transactions on Automation Science and Engineering (TASE)*, 2015.
- [Stoytchev, 2008] Alexander Stoytchev. *Learning the Affordances of Tools Using a Behavior-Grounded Approach*, pages 140–158. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [Watkins and Dayan, 1992] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- [Zhu *et al.*, 2014] Yuke Zhu, Alireza Fathi, and Li Fei-Fei. Reasoning about object affordances in a knowledge base representation. In *ECCV*, 2014.
- [Zhu *et al.*, 2015] Yukun Zhu, Ryan Kiros, Richard S. Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. *CoRR*, abs/1506.06724, 2015.