

Learning Intermediate Features of Object Affordances with a Convolutional Neural Network

Aria Y. Wang (yuanw3@andrew.cmu.edu)

Center for Neural Basis of Cognition (CNBC), Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213 United States

Michael J. Tarr (michaeltarr@cmu.edu)

Department of Psychology, Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213 United States

Abstract

Our ability to interact with the world around us relies on being able to infer what actions objects afford – often referred to as affordances. The neural mechanisms of object-action associations are realized in the visuomotor pathway where information about both visual properties and actions is integrated into common representations. However, explicating these mechanisms is particularly challenging in the case of affordances because there is hardly any one-to-one mapping between visual features and inferred actions. To better understand the nature of affordances, we trained a deep convolutional neural network (CNN) to recognize affordances from images and to learn the underlying features or the dimensionality of affordances. Such features form an underlying compositional structure for the general representation of affordances which can then be tested against human neural data. We view this representational analysis as the first step towards a more formal account of how humans perceive and interact with the environment.

Keywords: affordance; dataset; convolutional neural network;

While interacting with our environment, we naturally infer the functional properties of the objects around us. These properties, typically referred to as affordances, are defined by Gibson (1979), as all of the actions that an object in the environment offers to an observer. For example, kick for a ball and drink for water. Understanding affordances is critical for understanding how humans are able to interact with objects in the world.

In recent years, convolutional neural networks have been successful in performing object recognition in large-scale image datasets (Krizhevsky, Sutskever, & Hinton, 2012). At the same time, convolutional networks trained to recognize objects have been used as feature extractors and can successfully model neural responses as measured by fMRI in human visual cortex (Agrawal, Stansbury, Malik, & Gallant, 2014) or by electrodes in monkey IT cortex (Yamins & DiCarlo, 2016). To understand the relevant visual features in an object that are indicative of affordances, we trained a CNN to recognize affordable actions of objects in images.

Dataset Collection

Training deep CNNs is known to require large amounts of data. Available affordance datasets with images and semantic labels are largely limited at this moment. The only relevant dataset currently available to the public was created by Chao, Wang, Mihalcea, and Deng (2015), and only includes affordance labels for 20 objects from the PASCAL dataset and 90 objects from the COCO dataset. Here we built a large scale affordance dataset with affordances labels attached to all images in the ImageNet dataset (Deng et al., 2009). This dataset forms a more general representation of the affordance space and allows large scale end-to-end training from the image space and to this affordance space. The dataset collection process is shown in Figure 1. Human labelers were presented with object labels from ImageNet object categories and answered the question “What can you do with that object?”. All answers were then co-registered with WordNet (Miller, 1995) action labels so that our labels could be extended to other datasets. The top five responses from labelers were used as canonical affordance labels for each object. 334 categories of actions were labeled for around 500 objects categories. When combined with image to object label mappings from ImageNet, these affordance labels provided us with the image to affordance label mappings that were used to train our CNN.

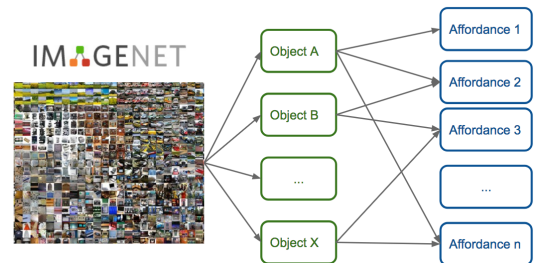


Figure 1: Dataset Collection. The labelers are given object labels, indicated in the green boxes here and assign to them affordances labels, indicated with blue boxes.

Visualization of Affordance Space

In our affordance dataset, each object was represented by a binary vector indicating whether each of the possible actions was available for this object or not. Each object can then be

represented as a point in the affordance space. We used PCA to project these affordance vectors into a 3D space and plotted the object classes as illustrated in Figure 2. In the 3D space created for visualization, the objects appear to be well separated. More specifically, the majority of living things were organized along the top axis; the majority of small household items were organized along the left axis; and transportation tools and machines were organized along the right axis. Human-related categories such as dancer and queen do not belong to any axis and appear as floating points in the space.



Figure 2: ImageNet images in the affordance space.

Results

Network Training

A CNN was trained to predict affordance categories from images. A total of 55 affordances were selected as potential actions after ensuring that each affordance label had at least 8 object categories associated with it (by removing affordances that were associated with too few object categories). Each object category was placed in the training, validation or testing sets. These sets were exclusive, such that, if one object category appeared in one set, it would not appear in the other two sets. Such separation ensures that the learning of affordances was not based on recognizing the same objects and learning linear mappings between objects and affordances.

We used the ResNet18 model (He, Zhang, Ren, & Sun, 2016) (other models such as VGG produced similar results), and trained it using the Adam optimizer (Kingma & Ba, 2014) by minimizing binary cross-entropy loss. Approximately 630,000 images from ImageNet were used in training, and approximately 71,000 images each were used for validation and testing. The trained CNN was evaluated by computing the average percentage of correctly predicted affordance labels, and the results are reported in Table 1. The trained networks showed significantly better performance compared to the baseline.

Skewed Distribution and Oversampling

Since actions such as “hold” and “grab” would be used on objects much more often than actions such as “thrust”, we

obtained an uneven distribution of affordance labels across image categories, as shown in Figure 3. In computer vision, oversampling is a commonly used solution for this problem. However, because of the multi-label nature of the affordance recognition problem, proper oversampling is challenging. Less frequently appearing classes need to be oversampled without over representing the more frequently appearing classes. We used Multi-label Best First Over-sampling (ML-BFO) (Ai et al., 2015), and re-trained the CNN with the resampled data. This produced a considerable increase in prediction performance, as seen in Table 1.

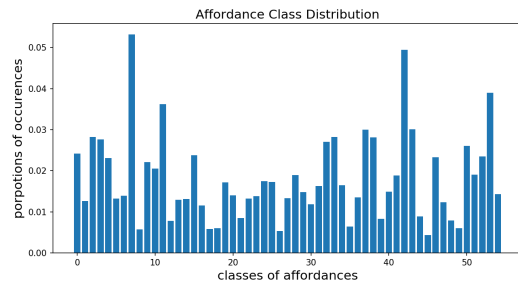


Figure 3: Percentages of objects classes assigned to each affordances categories.

Sample Predictions

Figures 4(a)–(d) demonstrate images where the network was able to predict correctly. However, the presence of distinct features can mislead the network. For example, in Figure 4(e), where white bars stand out in the image, the network predicted “grab” and “drive”, potentially mistaking the image as a bar or a road. On the other hand, human labelers, knowing that it is a image of a wall, provided labels such as “walk” and “enter”. Since ImageNet contains natural scene images, multiple objects are likely to appear in one image, even though each image is assigned only one object label. Such images confuse both the labelers and the network, and therefore can lead to incorrect affordance recognition as shown in Figures 4(f) and (g).

Visualizing the Learned Representation Space

RDM across Layers

To visualize the representations learned by the network, we randomly sampled 10 images from each of 30 objects classes, and extracted activations from the network layers. Pairwise correlation distance between network activation across layers was computed for each pair of images, and is shown in Figure 5. Pairwise distance between affordance labels is shown in the bottom-right matrix. This matrix denotes the ground truth distance in affordance space. Similar patterns begin to emerge in Layer 4 for both the fine-tuned network and the network trained from scratch. Critically, this pattern is not seen for the off-the-shelf network that was not trained on affordances.

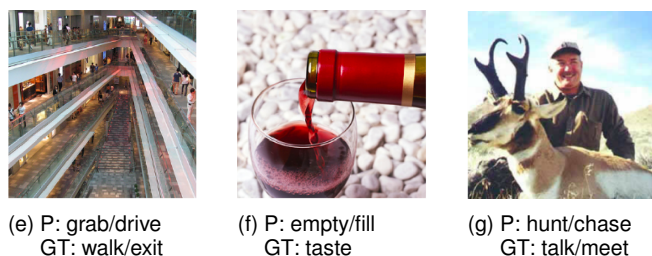


Figure 4: Sample predictions. (a)-(d): Examples of images with correct affordance predictions (correct label below each image). (e)-(g): Examples of images with incorrect affordance predictions (P: indicates the CNN prediction, while GT: indicates the ground truth based on human labeling).

This demonstrates that our network learns representations that effectively separate different affordance categories.

t-SNE

Activations from the second to last layer in the network trained from scratch were visualized using t-SNE (Maaten & Hinton, 2008), as shown in Figure 6. Images are coarsely split into four groups based on their distinct affordances: living things, vehicles, physical spaces and small items. In the 2D t-SNE visualization, the representation of living things (in green), vehicles (in red) and physical spaces (in blue) are visibly separable. Small items (in yellow), in contrast, span the entire

Table 1: Training Results. “Fine-tuning” indicates that the network was pre-trained to predict image categories, while “Training from Scratch” indicates that the network was initialized with random weights. Baseline accuracy was calculated by estimating the most frequent categories.

	Baseline	Fine-tuning	Training from Scratch	Fine-tuning w/ oversampling	Training from scratch w/ oversampling
Training Accuracy (%)	7.61	80.39	71.42	87.60	85.05
Testing Accuracy (%)	6.86	44.62	37.47	55.42	53.43

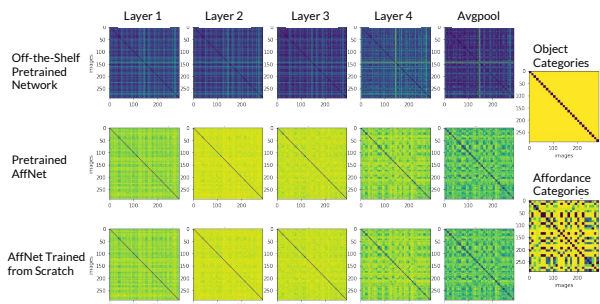


Figure 5: RDM matrix of layers from CNN from off-the-shelf pre-trained network, fine-tuned network and network trained from scratch for affordances.

space. The category of small items does not appear well separated, which is likely due to the visualization being limited to 2 dimensions.

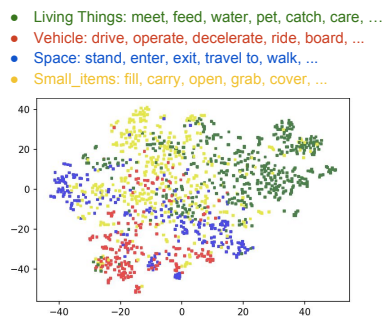


Figure 6: t-SNE visualization of the second to last layer in the CNN trained from scratch. Representations of images are coarsely split into four groups based on the distinct affordances of the images: living things (green), vehicles (red), physical spaces (blue) and small items (yellow).

Unit Visualization

We were able to visualize the output layer units of the CNN by optimizing in pixel space to determine which images maximally activated a specific unit. Figure 7 shows such visualization of 6 units from the output layer. The “ride” unit, for example, shows human- and horse-like structures; the “wear” unit

shows a coarse clothing pattern and details of common textures often associated with clothing. Similarly, units “climb”, “sit”, and “fill” show stairs-like, chair-like, and container-like structures respectively. Interestingly, the “watch” unit shows preference for dense textures in the center of the image space, which may correlate with image characteristics from objects that are related to watching (e.g., TV). It should be noted that unit visualization is very limited for capturing the learned intermediate features. Interpreting features in a limited 2D space is inherently biased and subjective.

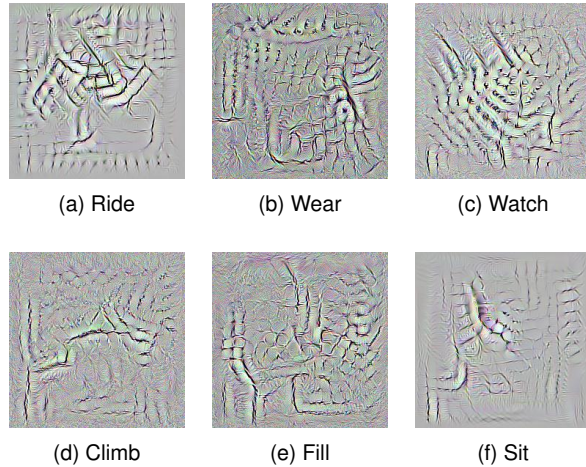


Figure 7: Visualization of 6 last layer units in the CNN.

Discussion

We successfully trained a CNN to predict affordances from images, as a means for learning the underlying dimensionality of object affordances. The intermediate features in the CNN constitute an underlying compositional structure for the representation of affordances.

Future Work

To ensure the objectivity of the affordance labeling, affordance labels for images – as opposed to just object categories labels – are being collected currently using Amazon Mechanical Turk. This dataset will be made publicly available after verification.

With a CNN trained for affordance recognition, weights from the intermediate layers can be extracted and used to featurize each image. A model can then be trained to predict the BOLD responses to each image. Correlations between the predicted responses and the true responses can be used to measure model performance. If a linear model is built to perform this task, the model weights could then be used as a proxy to localize where information about affordances is represented in the human brain.

Finally, affordance categories can be split into two large groups: semantically relevant ones, such as “eat”, which requires past experience with the objects in question; and non-semantically relevant ones, such as “sit”, which may be in-

ferred directly from the shapes of the objects. If semantic affordances are being processed in the brain, top-down information about the objects is potentially necessary in order to inform an observer about affordable actions, while the non-semantic ones would not require top-down information. Given such differences we may be able to differentiate between top-down and bottom-up visual processing in the human brain using our model; in particular, by distinguishing the different brain regions that are engaged in either or both of these two processes.

Acknowledgments

This project is funded by MH R90DA023426-12 Interdisciplinary Training in Computational Neuroscience (NIH/NIDA)

References

- Agrawal, P., Stansbury, D., Malik, J., & Gallant, J. L. (2014). Pixels to voxels: modeling visual representation in the human brain. *arXiv preprint arXiv:1407.5104*.
- Ai, X., Wu, J., Sheng, V. S., Yao, Y., Zhao, P., & Cui, Z. (2015). Best first over-sampling for multilabel classification. In *Proceedings of the 24th acm international on conference on information and knowledge management* (pp. 1803–1806).
- Chao, Y.-W., Wang, Z., Mihaleca, R., & Deng, J. (2015). Mining semantic affordances of visual object categories. In *Computer vision and pattern recognition (cvpr), 2015 ieee conference on* (pp. 4259–4267).
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Computer vision and pattern recognition, 2009. cvpr 2009. ieee conference on* (pp. 248–255).
- Gibson, J. (1979). *The ecological approach to visual perception* (pp. 127-143). Boston: Houghton Mifflin.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the ieee conference on computer vision and pattern recognition* (pp. 770–778).
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097–1105).
- Maaten, L. v. d., & Hinton, G. (2008). Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov), 2579–2605.
- Miller, G. A. (1995). Wordnet: a lexical database for english. *Communications of the ACM*, 38(11), 39–41.
- Yamins, D. L., & DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3), 356.