Visual Categorization with Random Projection

Rosa Arriaga[1], **David Rutter**[1], **Maya Cakmak**[2], **Santosh Vempala**[1]

[1]Georgia Tech.

[2]University of Washington.

Keywords: Random projection, Visual Categorization

Abstract

Humans learn categories of complex objects quickly and from a few examples. Random projection has been suggested as a means to learn and categorize efficiently. We investigate how random projection affects categorization by humans and by very simple neural networks on the same stimuli and categorization tasks, and how this relates to the robustness of categories. We find that (i) drastic reduction in stimulus complexity via random projection does not degrade performance in categorization tasks by either humans or simple neural networks, (ii) human accuracy and neural network accuracy are remarkably correlated, even at the level of individual stimuli and (iii) the performance

of both is strongly indicated by a natural notion of category robustness.

1 Introduction

Humans learn to categorize stimuli into two or more sets and are able to do so accurately after being presented with a few training instances (Booth, 2006; Makino and Jitsumori, 2007; Marcus et al., 1999; Mandler, 2003; Smith and Minda, 2000). This finding suggests that humans can identify perceptual information in the stimuli that is relevant to representing a category and discriminating it from others. The vast amount of perceptual information that is continuously sensed by humans suggests that several filtering mechanisms operate on sensory information before it is used in higher-level cognitive processes such as planning, reasoning or categorization. Arriaga and Vempala (1999, 2006) suggested *random projection* as a means by which the human brain could reduce the amount of information that is processed when presented with stimuli. The present paper is motivated by the following question: Does random projection maintain essential properties of visual stimuli to allow for accurate categorization by humans? In other words, does summarizing sensory data in this manner hinder human performance? So far this question has only been considered theoretically.

Human ability to categorize high-dimensional stimuli has been an inspiration for the categorization problem in machine learning. However, computational learning theory has frequently found learning in high dimension to be an intractable problem (Valiant, 1984; Vapnik and Chervonenkis, 1971; Vapnik, 1995). To make the categorization problem more tractable, *dimensionality reduction* is used to focus on relevant information in high-dimensional data. *Attribute selection* methods try to assess the relevance of dimensions in the high-dimensional space, so that only the most relevant subset of features are used for categorization. *Projection* methods, on the other hand, try to find lower-dimensional spaces to represent data so that the categorization problem is easier in the projection space.

Random projection (RP) is the method of mapping sample points in a high-dimensional space into a low-dimensional space whose coordinates are random linear combinations of coordinates in the high-dimensional space. An important property of RP, usually called the Johnson-Lindenstrauss Lemma (Johnson and Lindenstrauss, 1984), is that distances between pairs of points are approximately preserved by random projection, provided the target dimension is not too small (roughly the logarithm of the number of points being projected); see Vempala (2004) for recent proofs. Arriaga and Vempala (1999, 2006) used this property to present a theoretical model of efficient cognitive categorization. In their model, projected samples from well-separated categories remain distinguishable after projection, and categorization can be achieved efficiently using the lower-dimensional projections of stimuli. In particular, they give a simple and efficient algorithm with near-optimal sample complexity for learning a large-margin linear threshold function in a high-dimensional space: project data to a random low-dimensional space, then use a margin-based algorithm such as perceptron; to label a new example, use the same projection and label with the learned threshold function. Their key insight is that (a) margins to well-defined class boundaries are preserved with high probability by random projection (for any single data point), (b) only a small number of points are needed for training in the low-dimensional space, (c) therefore with large probability a sufficiently large sample maintains a large margin in the projected space. They show that the margin of separation of several mathematically well-defined concept classes is preserved by random projection, so that learning the concept is possible and efficient in the projected subspace. Moreover, random projection is easily realized by a simple two-layer neural network with edge weights set independently and randomly. In fact, setting each weight randomly to -1 or 1 suffices, as shown by Arriaga and Vempala (1999, 2006), who called it *neuron-friendly* random projection. Recent work (Allen-Zhu et al., 2014) shows that the weights out of each node can all have the same sign, and still enjoy distance (and margin) preservation. The output of a single neuron is believed to have the same sign on all synapses (excitatory or inhibitory), thereby making random projection even more plausible.

Random projection has been widely applied in conjunction with other methods in machine learning, such as manifold learning (Hegde et al., 2007a,b; Freund et al., 2007), face recognition (Goel et al., 2005), mixture learning (Dasgupta, 1999, 2000) and concept learning (Arriaga and Vempala, 2006; Garg and Roth, 2003). Bingham and Mannila (2001) and Goel et al. (2005) show that random projection has comparable performance with conventional dimensionality reduction methods such as principal component analysis while significantly reducing the computational cost and being data-independent. RP has been proposed as an alternative to kernel methods (Balcan et al., 2006). It has also been found to be useful in training deep neural networks (Saxe et al., 2011) either as a separate layer or as a good initialization.

In this paper, we study random projection in the context of human cognition. We build upon the idea that visual perception may involve a random projection stage. Random projection complements existing mathematical frameworks for cognition (Marcus et al., 1999; Tenenbaum et al., 2011; Xu and Kushnir, 2013), and can be viewed as a general-purpose preprocessing step for human learning. We make predictions based on this hypothesis and test these predictions using behavioral experiments with human subjects as well as simple neural networks designed to capture predicted human behavior. The main idea is that human accuracy at categorization tasks should not degrade if stimuli are randomly projected in advance. Our first hypothesis is the following (in Section 4, we place it in the context of visual psychophysics):

H1. Humans should be able to learn to categorize randomly projected stimuli as easily and accurately as they categorize unprojected stimuli.

Our second hypothesis is based on the idea that predicted human behavior on perceptual categorization tasks will be mimicked by very simple neural networks under the same stimulus conditions. The use of neural networks in machine learning is a sophisticated and high successful technique today, but our goal is not to find a neural network with best possible performance on the categorization tasks at hand. Rather, we use the simplest neural networks in order to draw robust and meaningful conclusions. Moreover, all of our categorization tasks are based on *one-shot* training, i.e., with only a single presentation of one example from each of two categories. While there has been work on Bayesian learning with one-shot training (Li et al., 2006; Lake et al., 2013), multi-layer neural networks typically need a large number of examples to generalize well and achieve low classification error (LeCun and Bengio, 1998; Hinton et al., 2012). In light of this, we hypothesize that:

H2. Very simple neural networks modeling categorization with random projection should achieve similar performance as humans.

Our experimental design to test these predictions is presented in the next section. One challenge in the design is to create stimuli and categories that are both *natural* and *novel* (so as to avoid the bias of prior knowledge of categories, while not ignoring natural features). We developed three different types of stimuli and used two different types of random projection.

In Section 3, we present the results of our human categorization experiments on original and projected visual stimuli. They provide strong support for the above hypotheses, in each of the stimuli types and each of the projection types. We also found similar accuracy in experiments with humans and simple neural networks and observed improved efficiency during categorization when done with random projection. These results are discussed in detail in Section 3. We then turn to a more detailed analysis, comparing the performance of humans and neural networks *on individual stimuli*. We found that these are highly correlated, with high agreement in the subsets of stimuli which humans and neural networks categorized incorrectly. To understand this, we introduce a natural notion of *robustness* of stimuli. Roughly speaking, it measures how clearly a stimulus belongs to one category vs the other, with the robustness value being closer to zero as a stimulus becomes more ambiguous. Placing the images in the order of robustness, we find both that the performance of humans and NN improves with stimulus robustness and that the few stimuli where humans and neural nets disagree are almost all in the region of low robustness. These results are presented in Section 3.1.

Although we cannot definitively claim that human brain actually engages in random projections, our results support the notion that random projection is a plausible explanation for dimensionality reduction within human perception via random summarization. To the best of our knowledge, this is the first study of random projection based on human subjects.

2 Methods

To address the question of how random projections affects categorization ability in humans, we conducted experiments in which human subjects were tested on two-class visual categorization tasks with both unprojected and projected visual stimuli. We designed three types of visual images and two types of random projection methods. We also designed two versions of a simple neural network to categorize the same visual stimuli. In this section, we provide details about the visual stimuli, the projection methods, the experimental design and the neural network design.

2.1 Visual stimulus design

Our choices in creating these stimuli stem from two criteria, *naturalness* and *novelty*. With naturalness we tried to ensure that the stimuli reflect properties of the real world. In particular, they should have a clear structure, with different connected regions distinguishable by differences in color, an attribute that the visual system pays distinct attention to (Read, 2014); they should not be collections of arbitrarily colored pixels. On the other hand, we did not want the images to appear similar to categories that subjects might already have knowledge of or have real world experience with. Using images of dogs and cats would have inevitably caused the subjects to call on prior knowledge of these categories.

We created three types of images to use as visual stimuli: (i) *Geon arrays* which are arrays of simple geometric shapes, (ii) *Dartboards* which are colored Voronoi diagrams, and (iii) *Substrates* which are simple colored layers rotated by some amount. Samples of the two categories from each type of visual stimuli are given in Figure 1. The images for each category were generated by the same algorithm using a different set of parameters for the two classes. In geon arrays, we used four different shapes (geons): balls, cones, cubes and cylinders. Images were generated by placing 9 random geons in a 3×3 array, with different distributions on geons for the two categories (A: 0.15, 0.2, 0.05, 0.6; B: 0.2, 0.6, 0.1, 0.1). Dartboard images were generated by picking 4 points at random in a unit square, computing their Voronoi diagram and coloring each region; again different distributions on the locations of the points were used for the two categories (for category A, a point was chosen in the upper half of the main diagonal of the containing square, giving a partition into 4 rectangles via the horizontal and vertical lines through the chosen point; then four points were picked randomly, one from each rectangle; for B, the point used for subdivision was chosen from the lower half of the diagonal). Substrates were generated using 1-dimensional Brownian random walks from left to right to partition into layers, then the layers were colored and rotated randomly; the angle of rotation was from a different distribution for the two categories (a uniformly chosen anti-clockwise rotation for A, and clockwise for B).

These categories can be assigned semantic labels, at least before projection e.g., "more cylinders" vs "more cones", "more green" vs "more white" and "sloping up" vs "sloping down". However, these are not familiar categories, rather familiar features

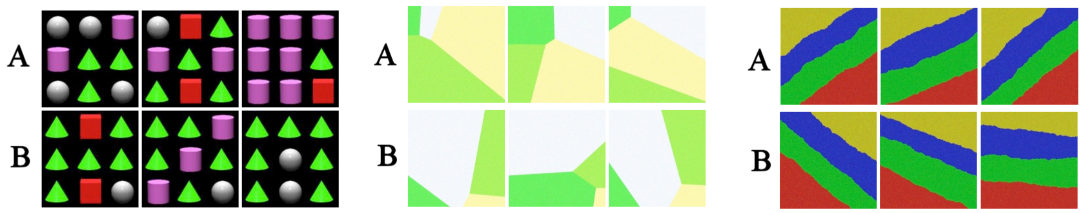

(a) Geon (b) Dartboard (c) Substrate

Figure 1: Samples from the two categories A and B for the three types of stimuli.

(number, color, angle) in an unfamiliar context. In our experiments, subjects were presented with a single pair of images from each stimuli type, one from each category, for a few seconds. Immediately following this, they had to classify other images, one at a time; the labels of images used for testing were not revealed to the subject. Thus, while subjects must have used some cognitive rule to categorize, it is unlikely that they had the time to consciously define the right semantic categories and then use them for classification; even more so in the case of projected stimuli, which we describe next.

2.2 Random Projection of Images

Random projection maps points in n-dimensional space to k-dimensional space, where k is smaller than n. The process of projection can be achieved by multiplying a given point (taken as an n-vector) by a random $n \times k$ matrix to produce a k-vector. Typical choices for the random matrix with provable guarantees are picking each entry independently from a standard Normal distribution (Johnson and Lindenstrauss, 1984; Indyk and Motwani, 1998; Arriaga and Vempala, 1999; Dasgupta and Gupta, 2003), or uniformly from two discrete values $\{-1, 1\}$ (Arriaga and Vempala, 1999, 2006). In designing our projections, we tried to ensure that the projected stimuli have a visually salient representation so that comparison with the original stimuli is meaningful. This suggests that some spatial properties of the image should be preserved by projection and that projections have meaningful color values. We used two projection methods (Figure 2), as well as unprojected images. Nevertheless the projected images were only of size 6×6, and the projection methods were generic and applied to all three stimulus types without any modification.

In the *sliding-window* random projection, a single color is generated from a random 0-1 combination of the colors in a window that slides over the image, and that color is used to fill a corresponding location in the projected image. The window size and the increment with which the window is moved determines the reduced image dimension. This is a random convolution of a particularly simple type, with random weights chosen from $\{0, 1\}$ and a very small number of units. Neural networks based on convolution have been extremely successful in large-scale machine learning (LeCun and Bengio, 1998; Sermanet et al., 2014). They are usually initialized with random weights and then trained with a large number of examples. Biological support for such mechanisms indicates that subsets of sensory inputs are aggregated at lower levels (e.g., line detectors for various parts of the visual field). Our choice of simplified convolution is considerably more restricted than general random projection, and thus any support for our hypothesis (of no degradation in performance) would arguably be stronger. In our experiments, we used a window to reduce the resolution of original images from 150×150 to 6×6 (for the geon arrays) and 500×500 to 6×6 (for larboards and substrates). In the former case, the window size was 25×25 and in the latter it was 100×100.

Another aspect of the human visual system is the existence of specialized feature detectors, for colors, lines, corners etc. We designed a random projection that takes this into account by first identifying corners in an image. In *corner* random projection, corners in the original image are detected using the FAST corner detector (Rosten and Drummond, 2006) and windows of pixels are extracted from a small region, 31×31 around each corner. The projections are computed as the weighted average of these windows using a set of randomly chosen weights (Figure 2). In this case, the resolution reduces from 150×150 and 500×500 to 20×20.

Samples of projected images obtained with both methods are given in Figure 3. These families of random projections do not have distance (or margin) preservation guarantees for arbitrary inputs. Our prediction is that human performance on categorizing these visual stimuli will not be affected in any significant way by these projections.

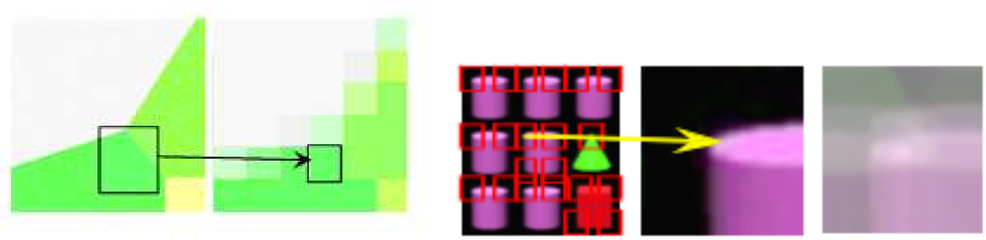

(a) (b)

Figure 2: Description of the random projection methods for visual stimuli: (a) sliding-window random projection and (b) corner random projection.

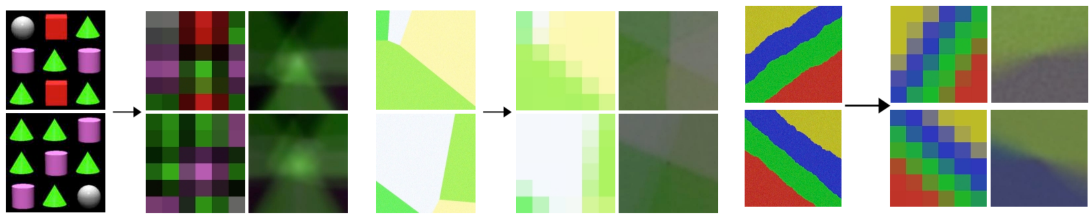

(a) Geon (b) Dartboard (c) Substrate

Figure 3: Samples of images projected using sliding window and corner projection methods.

2.3 Procedure

There were 9 experimental conditions resulting from all combinations of the three stimuli types (geon arrays, dartboards, substrates) and the three projections (unprojected, sliding-window projected, corner projection). Each of the 9 conditions was tested on 16 subjects. Subjects were randomly assigned to 3 conditions each, chosen to range across the three stimuli types and the three projections and balanced across them. For example, a subject might be tested on unprojected geons, sliding-window projected dartboards and corner projected substrates. The total number of human subjects was thus 48 (= 9 x 16 / 3). All the subjects were college students ages 18-24 with an equal number of males and females. They signed an IRB-approved consent form before parsticiipating in the study

In each categorization task, subjects were first shown two sample images, one from each category, side by side for up to 10 seconds (the training phase). Then each subject was shown a sequence of 16 single images of the same type of stimulus and asked to label each one as A or B within 5 seconds. Images from the two categories were presented in random order. The number of correct responses and the reaction time were measured.

2.4 Neural network categorization task

For the categorization of visual stimuli into two classes we used two simple feed-forward neural network structures. The first one assumes that all images (original or projected) are scaled to a 150×150 resolution. The input to the neural network is the visual stimuli with RGB values normalized and the sign of the output is the category label. The neural networks contain $150 \times 150 \times 3 = 77500$ input units and one output unit. The input units are directly connected to the output unit and all transfer functions used are linear. Essentially, the category is decided by the sign of the inner product of the input vector and the weights, i.e., a linear threshold function with zero as the threshold. The weights are initialized randomly in the range $[-1, 1]$ and the vector of weights is normalized. The weights are updated with the delta rule and the neural network is trained until the total weight update within an epoch falls below a certain threshold.

The second neural network differs from the first in that projected images are used as inputs without being rescaled. Therefore, the neural networks are much smaller for projected stimuli. The complete neural network is the combination of a projection layer, in which no learning occurs (the weights are chosen randomly and fixed), and a categorization layer. The input to these networks is therefore always the unprojected stimuli. The two neural networks are illustrated in Figure 4.

To mirror the task performed by humans, we trained the neural networks on only one image pair (one from each category) and tested them on the remaining 16 images for that task, identical to the human experiment. First, we trained using the exact same exemplars that the humans were presented. Later, to obtain an average performance of the neural network, we did a nine-fold cross-validation, whereby the neural network is trained with a single stimulus pair and tested on the remaining images. Every stimulus in the set is used once as a training example, for a total of nine iterations.

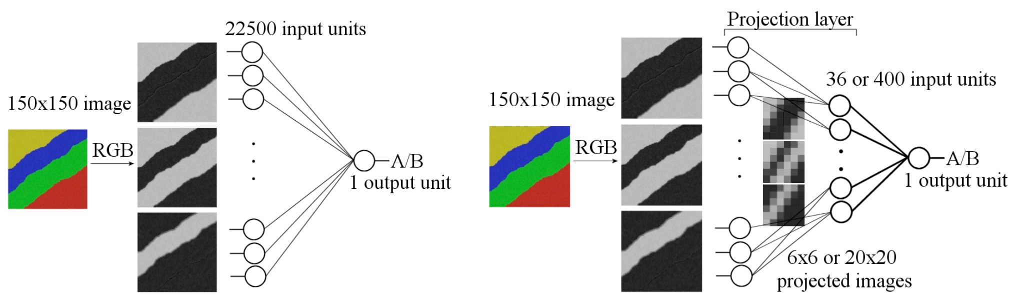

(a) Standard neural network

(b) Neural network with a projection layer

Figure 4: The simple neural network used for visual stimuli categorization.

3 Results

In this section we present comparative results of categorization performance for unprojected and projected images for each type of stimulus. We compare human subjects and the simple neural networks used and we analyze both average performance and performance on individual images.

In Figure 5(a), average successful categorization rate among human subjects is compared for unprojected and projected images. Humans were able to categorize successfully in all three types of images (t-test with $p < 0.001$). Our results indicate that in spite of the reduction in the feature complexity, there was no statistically significant loss in subjects' ability to correctly categorize stimuli. Thus, the random projections used in our experiments preserved the essential similarities and differences between stimuli from different categories as perceived by humans. Also, after projection, the accuracy of classification was similar across all three stimuli types and both projection methods. Moreover, post-test responses to the question, "what did you use to categorize?" were highly varied, except for unprojected substrates, suggesting that human subjects were not consciously creating the same semantic labels for categories.

Similar results were obtained in the neural network categorization task. Here, for the projected images, we used two types of networks as described in the previous section (Figure 4; the first one uses images rescaled to 150×150, regardless of projection, while the second uses a fixed projection layer and does not rescale). As expected, there was little difference in the performance of the two types of neural networks. The networks with and without the rescaling had identical performance in 6 of the 9 conditions, and the difference in the classification rate was just under 7% in the remaining 3 conditions (unprojected and sliding-window projected Geon arrays and corner-projected substrates). We remark that projections improved the efficiency in categorization by reducing the size of the neural networkIn. Figure 5 includes NN results without the projection layer. Figure 5(a) and (b) compare performance on original and projected stimuli for humans and neural networks. The neural network results are the average performance in the 9-fold tests. Figure 5 (c) and (d) compare performance across the different types of stimuli. Our findings were similar across all stimulus types, indicating that the general-purpose random projection methods used did not affect human or neural network performance.

We extended our analysis by investigating variations in successful categorization on

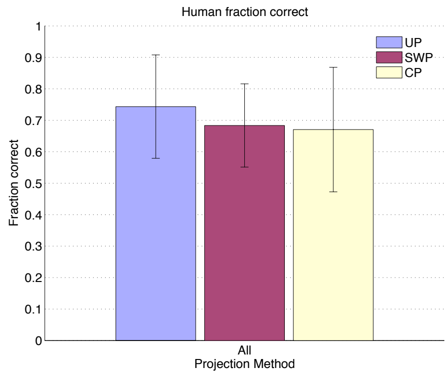

(a) Human

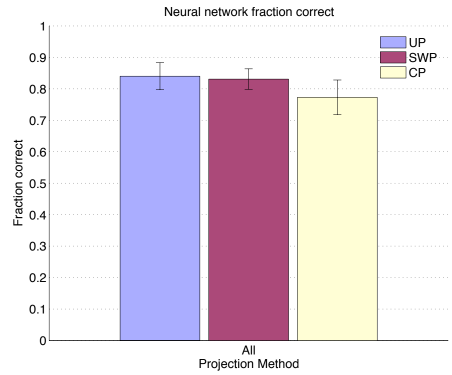

(b) Neural network

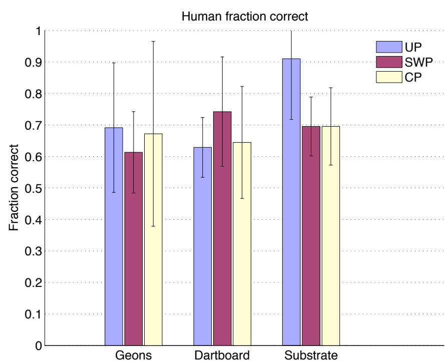

(c) Human

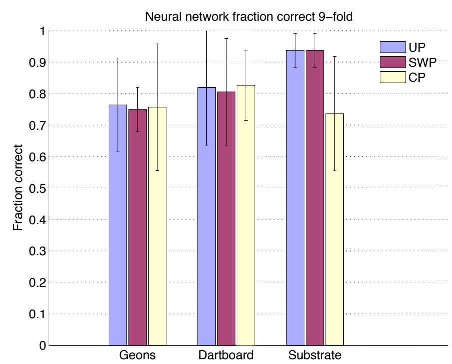

(d) Neural network

Figure 5: Experimental Results. The fraction of correctly categorized stimuli across all stimulus types without projection (blue), after sliding-window random projection (red), and after corner projection (yellow) by (a) humans and (b) neural networks through the 9-fold cross validation are shown. The graphs (c) and (d) give the performance for each stimulus type. The standard deviations in are given by the error bars.

individual images.

3.1 Comparison based on robustness

Here we investigate how categorization of images relates to their *robustness*, which we will define presently. Although category robustness has been discussed in the cognitive science literature in relation to categorization, a concise mathematical definition of robustness is not available. Arriaga and Vempala (1999) define the robustness of a category as the minimum, over all examples, of the distance of an example to the category boundary. This assumes that the categories are defined by concepts with well-defined boundaries, while in our case no assumption about the category representation is made. The linear discriminant analysis method, also known as Fisher's method (Fisher, 1938), estimates the hyperplane which best separates two categories by maximizing the ratio of the distance between categories to the scatter within categories, after approximating categories by Normal distributions. We adopt the ratio between the distance of the sample to the mean of its category to the distance of the sample to the mean of the other category. In order to unify the robustness calculation for the two categories, we take the difference between the ratio and its reciprocal. The precise definition of our robustness measure is as follows. Here $\mathbf{x}$ is the sample stimuli, $\mu_{\mathbf{A}}$ and $\mu_{\mathbf{B}}$ are the means of the two categories:

$$R_{AB}(\mathbf{x}) = \frac{\|\mathbf{x} - \mu_{\mathbf{A}}\|}{\|\mathbf{x} - \mu_{\mathbf{B}}\|} - \frac{\|\mathbf{x} - \mu_{\mathbf{B}}\|}{\|\mathbf{x} - \mu_{\mathbf{A}}\|}$$

Note that this provides a measure of robustness for individual samples rather than for a whole category. For an example whose category is ambiguous, i.e., its distances to the two sample means are nearly equal, the robustness measure will be close to zero; if the distance to one of the categories is smaller by a factor of r, the robustness measure will go up as $r - 1/r$ in magnitude (positive for examples from one category and negative for examples from the other category).

We investigated the role of robustness in categorization by looking at the categorization performance of humans and NNs on individual images, taking into account their robustness. We first observed that samples on which human subjects made the most mistakes are the ones that are frequently misclassified by the neural network. In Figure 6 performance on individual images of human subjects and neural networks (in the 9-fold cross-validation) are compared. Although there exists outliers, performance in general is very similar, *even on individual stimuli!*

We argue that the performance on individual images is strongly related with the robustness of the image as defined above. In Figure 6, stimuli are sorted according to their robustness, i.e., the two extremes being the most robust samples of the two categories and the middle ones being the least robust ones. It is again clear from this figure that images with low robustness are the ones on which both humans and neural networks perform poorly. Another interesting observation is that most samples on which human and NN performance are not correlated (outliers) are ones with low robustness.

The distribution of image samples on a robustness scale in Figure 6 shows that the trend of degrading performance towards the origin can be observed in all categories. The figure also provides an idea of how projections affect robustness. While sliding window projection does not seem to change the robustness of samples, corner projection stretches the robustness distribution, making low-robustness images worse, and highly robust images better.

4 Discussion

The main findings of this study can be summarized as follows. (i) The categorization ability of human subjects does not degrade with random projection applied as a task-

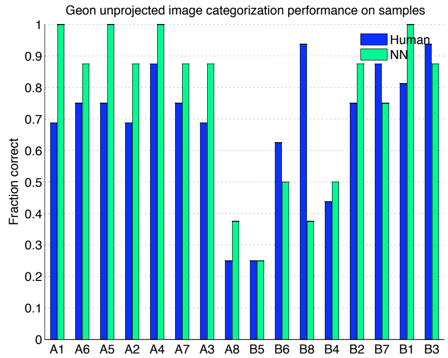

(a) Unprojected

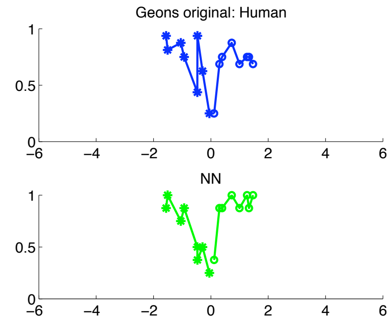

(b) Unprojected

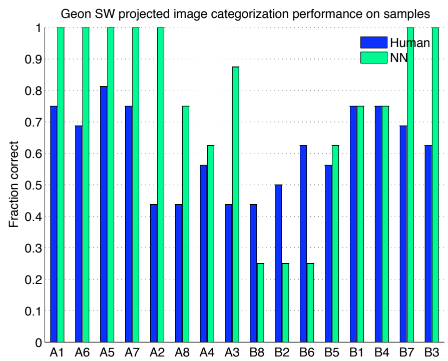

(c) Sliding window projected

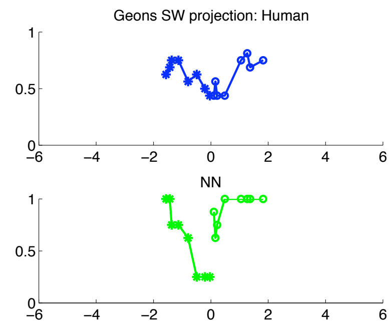

(d) Sliding window projected

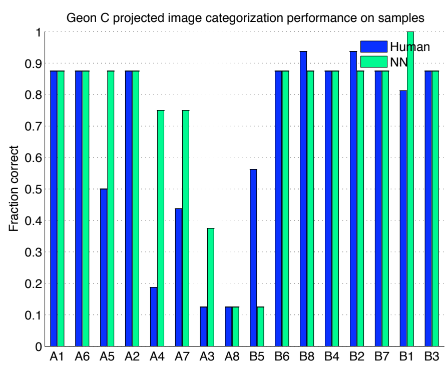

(e) Corner projected

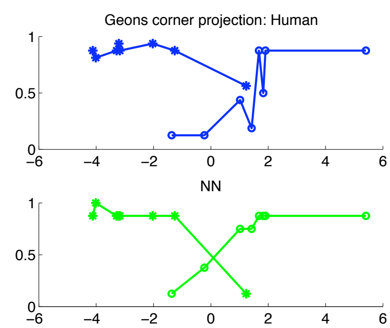

(f) Corner projected

Figure 6: (Left) Performance comparison on individual image samples (Geons) sorted according to robustness. (Right) Distribution of image samples according to their robustness and change of human (top) and neural network (bottom) performance with robustness.

independent initial processing step, (ii) Simple neural networks have remarkably similar performance to humans on categorization tasks for both original and projected stimuli. Despite their simplicity, these networks are able to capture human categorization capabilities at the level of individual samples. (iii) Human and NN performance are highly correlated to the robustness of categories. A measure of robustness that is in agreement with human performance can be obtained by considering both the distance of samples to the other category and the distance to the prototype of their own categories. There exists a three-way correlation between robustness of image samples and the categorization performance by humans and neural networks on those images. As far as we know this is the first study using random projections with human subjects, and comparing their performance on one-shot learning and categorization with neural networks.

Turning to support from psychophysics (the analysis of perceptual process using quantitative measures), we note that it is possible to compare monkey neuronal data to human psychophysics performance because the two are generally understood to have very similar visual systems. Results indicate that humans and non-human primates can process visual stimuli rapidly, at speeds between 14-50 ms (Proctor and Brosnan, 2013). This happens under a variety of presentation paradigms from single stimuli presentation of clip art to a continuous sequence of naturalistic images (Keysers et al., 2001). The pyschophysics literature also shows that there is consistency of encoding among naturalistic and geometric stimuli in monkey primary visual cortex (Wiener et al., 2001). Our results raise a number of questions about how the visual system might be able to make random projections of stimuli. The literature indicates that if RP is a mechanism in the visual system then it 1) it happens quickly, 2) in both human and non-human primates and 3) across varied stimuli that are static and in motion.

To conclude, random projection is computationally efficient and could account for the efficiency of categorization by humans. In neural networks, random projection increases the efficiency of learning by reducing the size of the network. Taken together with our empirical findings that pre-projecting visual stimuli does not significantly affect human ability to categorize, random projection appears to be a highly plausible mechanism for information processing in general, and categorization in particular. Confirming this, by understanding the projections used more precisely, is a compelling research direction for cognition and learning. As one reviewer suggested, it would be interesting to compare human and neural network performance on more complex images, with more training examples and using deeper neural networks. Future studies could (a) develop stimuli by varying the ease with which they can be labeled semantically (easy to hard) and (b) vary the random projection method from what we used and other neurally plausible methods to fully random positive and negative weights. In this regard, we are inspired by the fact that the trichromatic nature of human vision was deduced via behavioral experiments hundreds of years before objective knowledge from physiology confirmed its existence (Read, 2014).

Acknowledgments

We thank David Mann, Dawn Finney, Ya-Lin Huang and Parikshit Ram for their help at various stages of this project. We also thank the anonymous reviewers for many useful comments that improved the paper.

References

Allen-Zhu, Z., Gelashvili, R., Micali, S., and Shavit, N. (2014). Sparse sign-consistent johnsonlindenstrauss matrices: Compression with neuroscience-based constraints. *Proceedings of the National Academy of Sciences*, 111(47):16872–16876.

Arriaga, R. and Vempala, S. (1999). An algorithmic theory of learning: Robust concepts and random projection. In *FOCS*, pages 616–623.

Arriaga, R. and Vempala, S. (2006). An algorithmic theory of learning: Robust concepts and random projection. *Machine Learning*, 63(2):161–182.

Balcan, M., Blum, A., and Vempala, S. (2006). Kernels as features: On kernels, margins, and low-dimensional mappings. *Machine Learning*, 65(1):79–94.

Bingham, E. and Mannila, H. (2001). Random projection in dimensionality reduction: applications to image and text data. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '01)*, pages 245–250, New York, NY, USA. ACM.

Booth, A. (2006). Object function and categorization in infancy: Two mechanisms of facilitation. *Infancy*, 10(2):145–169.

Dasgupta, S. (1999). Learning mixtures of gaussians. In *FOCS*, pages 634–644.

Dasgupta, S. (2000). Experiments with random projection. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, UAI'00, pages 143–151, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Dasgupta, S. and Gupta, A. (2003). An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures and Algorithms*, 22(1):60–65.

Fisher, R. (1938). The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7:179–188.

Freund, Y., Dasgupta, S., Kabra, M., and Verma, N. (2007). Learning the structure of manifolds using random projections. In *Neural Information Processing Systems (NIPS)*.

Garg, A. and Roth, D. (2003). Margin distribution and learning. In *ICML*, pages 210–217.

Goel, N., Bebis, G., and Nefian, A. (2005). Face recognition experiments with random projection. In *Proceedings of SPIE, Biometric Technology for Human Identification II*, volume 5779, pages 426–437.

Hegde, C., Davenport, M., Wakin, M., and Baraniuk, R. (2007a). Efficient machine learning using random projections. In *Workshop on Efficient Machine Learning, Neural Information Processing Systems (NIPS)*.

Hegde, C., Wakin, M., and Baraniuk, R. (2007b). Random projections for manifold learning. In *Neural Information Processing Systems (NIPS)*.

Hinton, G. E., Deng, L., Yu, D., Dahl, G. E., Mohamed, A., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T. N., and Kingsbury, B. (2012). Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Process. Mag.*, 29(6):82–97.

Indyk, P. and Motwani, R. (1998). Approximate nearest neighbors: Towards removing the curse of dimensionality. In *STOC*.

Johnson, W. and Lindenstrauss, J. (1984). Extensions of lipshitz mapping into hilbert space. *Contemporary Mathematics*, 26:189–206.

Keysers, C., Xiao, D.-K., Fldik, P., and Perrett, D. I. (2001). The speed of sight. *Journal of Cognitive Neuroscience*, 13:90–101.

Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2013). One-shot learning by inverting a compositional causal process. In *Advances in Neural Information Processing Systems*, pages 2526–2534.

LeCun, Y. and Bengio, Y. (1998). The handbook of brain theory and neural networks. chapter Convolutional Networks for Images, Speech, and Time Series, pages 255–258. MIT Press, Cambridge, MA, USA.

Li, F., Fergus, R., and Perona, P. (2006). One-shot learning of object categories. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(4):594–611.

Makino, H. and Jitsumori, M. (2007). Discrimination of artificial categories structured by family resemblances: A comparative study in people (homo sapiens) and pigeons (columba livia). *Journal of Comparative Psychology*, 121(1):22–33.

Mandler, J. (2003). Early category and concept development: Making sense of the blooming, buzzing confusion. In Rakinson, D. and Oakes, L., editors, *Conceptual Categorization*. Oxford University Press.

Marcus, G., Vijayan, S., Rao, S. B., and Vishton, P. (1999). Rule-learning in seven-month-old infants. *Science*, 283:77–80.

Proctor, D. and Brosnan, S. (2013). Visual processing speed in capuchin monkeys (cebus apella) and rhesus macaques (macaca mulatta). *International Journal of Comparative Psychology*, 26:166–175.

Read, J. (2014). The place of human psychophysics in modern neuroscience. *Neuroscience*.

Rosten, E. and Drummond, T. (2006). Machine learning for high-speed corner detection. In *European Conference on Computer Vision*, volume 1, pages 430–443.

Saxe, A. M., Koh, P. W., Chen, Z., Bhand, M., Suresh, B., and Ng, A. Y. (2011). On random weights and unsupervised feature learning. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 1089–1096.

Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., and LeCun, Y. (2014). Overfeat: Integrated recognition, localization and detection using convolutional networks. In *International Conference on Learning Representations (ICLR2014)*. CBLS.

Smith, D. and Minda, J. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(1):3–27.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., and Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022):1279–1285.

Valiant, L. (1984). A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142.

Vapnik, V. (1995). *The nature of statistical learning theory*. Springer.

Vapnik, V. and Chervonenkis, A. (1971). On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its applications*, 26(2):264–280.

Vempala, S. (2004). *The Random Projection Method*. Dimacs Series in Discrete Mathematics and Theoretical Computer Science.

Wiener, M., Oram, M. W., Liu, Z., and Richmond, B. (2001). Consistency of encoding in monkey visual cortex. *Journal of Neuroscience*, 21:8210–8221.

Xu, F. and Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, 22(1):28–32.