

Spotlight

Beyond learnability:
understanding human
visual development with
DNNsLei Yuan ^{1,2,*}

Recently, Orhan and Lake demonstrated the computational plausibility that children can acquire sophisticated visual representations from natural input data without inherent biases, challenging the need for innate constraints in human learning. The findings may also reveal crucial properties of early visual learning and inform theories of human visual development.

How to build intelligence? Building intelligence involves three crucial components: the input data, the learning machinery, and the learning outcome. Advances in wearable sensors that can capture learners' input data on the scale of daily life [1,2], in machine learning algorithms that can simulate human learning [3], and in measured performance across cultures and contexts [4] contribute new insights into the interplay of input learning data, learning machinery, and learning outcomes.

The foundations for building human intelligence begin in infancy, and it is the properties of that foundation that motivated a recent paper by Orhan and Lake [5]. The richness, complexity, and noise of real-life input of young children challenge explanations of how children rapidly learn all that they know. Traditional theory proposed a solution in terms of 'inductive biases'; that is, the children's learning machinery was biased to make certain conclusions from the data, driven by innate constraints [6]. Orhan and Lake demonstrated the

computational plausibility that children can build sophisticated representations of visual objects from naturalist input data without built-in inductive biases. The authors used an input data set collected from three infants who wore head cameras capturing their point-of-view scenes (or egocentric views) for brief durations (1–2 h) weekly between the ages of 6 and 31 months. They then used the collected video frames to train deep neural networks (DNNs) in a self-supervised manner with no names or category labels for objects.

One such model used by Orhan and Lake was a Vision Transformer (ViT) embedding model trained via various self-supervised learning (SSL) algorithms (e.g., DINO). Rather than learning labels for instances (as in supervised learning), embedding models are optimized to represent perceptually similar images closer in the representation space and more different images farther apart. A method known as data augmentation has a critical role in this [7]. In this method, the training data are augmented by creating a series of variations of the original input data through augmentation methods, such as rotation or cropping. In the current study, the models were optimized by training with variations (i.e., augmented views) as belonging to the same object, and then tested in a suite of downstream benchmark tasks (e.g., image classification and semantic segmentation). The overall results demonstrated that a relatively small subset of augmented children's input data yielded a strong performance, achieving ~70% of the performance of models trained through the same unsupervised learning algorithms with the full data set of clear, photographed images from ImageNet.

Data augmentation has been argued to mimic the properties of natural human experiences [7]. As an active perceiver moves, uses, and looks at objects in the world,

they create many varied ego-centric images of the same objects. As such, natural human experiences are fundamentally constrained by time and space. Images in the same physical space (e.g., sitting at the kitchen table) are more similar compared with images captured from different contexts (e.g., sitting at the kitchen table versus playing at the park). Images sampled from the same context are likely to be about the same objects or people (e.g., different views of the same spoon pushing around food on a plate). The computational efficiency of the children's view images over that of ImageNet may well stem from the proportion of objects with these 'natural' augmentations. Classic theories of human concept learning have posited that similarity has an important role in concept learning and categorization [8]. These traditional theories contrast with more recent ones, which place more emphasis on how learning the names of objects (in a supervised way) teaches children object categories [9]. Orhan and Lake's results demonstrate that similarity itself may be a powerful pathway to categorization.

The second model used by Orhan and Lake is generative, with the ability to generate new images or fill in missing portions of an existing image. It works by first learning a discrete, compressed representation of input images using a vector quantized generative adversarial network (VQ-GAN). A transformer is then trained to generate images sequentially from top to bottom within this compressed representation. Models trained with children's data generated completions that exhibited broad consistency with the original scene, regarding the color, texture, orientation, and outline of the objects or people depicted. Interestingly, compared with models trained with clear photograph images from ImageNet, models trained with children's egocentric view data were less successful at generating finer details. A similar result was also obtained in the embedding models trained

with children's views, which were less object-centric and showed more sensitivity to global context and background scene information. Do these performance differences between models trained with children's egocentric images and those with photographs unveil relevant principles of visual category learning?

From certain perspectives, both infants' early visual experiences and their developing visual systems can be considered immature; however, these immaturities may be uniquely adaptative for developing a robust visual representation system [3]. Infants are born with low visual acuity and show preferences for simple images with fewer edges and high spatial contrast [10]. These properties of the infants' visual system may be beneficial for similarity-based category learning because they may promote the discovery and learning of similar global structures across images with different views. DNN models combined with realistic ego-centric view data

may be useful for simulating individual learning trajectories and testing theoretical hypotheses about the development of visual representations. Conversely, data augmentation may approximate, but not yield, optimal within-object variation. Incorporating properties of human visual development (e.g., context-dependent, multiple views of the same object or low acuity) may inspire a new generation of machine learning algorithms that mimic human-level visual intelligence.

Acknowledgments

This work was supported by the National Science Foundation DRK12 Grant 2200781 to L.Y. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author (s) and do not necessarily reflect the views of the National Science Foundation.

Declaration of interests

None declared by author.

¹Department of Psychology and Neuroscience, University of Colorado Boulder, Boulder, CO, USA

²Website: www.colorado.edu/lab/del/

*Correspondence:

lei.yuan@colorado.edu (L. Yuan).

<https://doi.org/10.1016/j.tics.2024.05.002>

© 2024 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

References

1. Smith, L.B. *et al.* (2015) Contributions of Head-mounted cameras to studying the visual environments of infants and young children. *J. Cogn. Dev.* 16, 407–419
2. Bergelson, E. *et al.* (2019) Day by day, hour by hour: naturalistic language input to infants. *Dev. Sci.* 22, e12715
3. Sheybani, S. *et al.* (2024) Curriculum learning with infant egocentric videos. *Adv. Neural Inf. Process. Syst.* 36, 1–14
4. Frank, M.C. *et al.* (2021) *Variability and Consistency in Early Language Learning: The Wordbank Project*, MIT Press
5. Orhan, A.E. and Lake, B.M. (2024) Learning high-level visual representations from a child's perspective without strong inductive biases. *Nat. Mach. Intell.* 6, 271–283
6. Tenenbaum, J.B. *et al.* (2011) How to grow a mind: statistics, structure, and abstraction. *Science* 331, 1279–1285
7. Shorten, C. and Khoshgoftaar, T.M. (2019) A survey on image data augmentation for deep learning. *J. Big Data* 6, 60
8. Rosch, E. and Mervis, C.B. (1975) Family resemblances: studies in the internal structure of categories. *Cogn. Psychol.* 7, 573–605
9. Yu, C. and Smith, L. (2007) Rapid word learning under uncertainty via cross-situational statistics. *Psychol. Sci.* 18, 414–420
10. Ayzenberg, V. *et al.* (2023) Development of visual object recognition. *Nat. Rev. Psychol.* 3, 73–90