

---

# Scaling may be all you need for achieving human-level object recognition with human-like visual experience

---

**Emin Orhan**  
New York University  
eo41@nyu.edu

## Abstract

This paper asks whether current self-supervised learning methods, if sufficiently scaled up, would be able to reach human-level visual object recognition capabilities with the same type and amount of visual experience humans learn from. Previous work on this question only considered the scaling of data size. Here, we consider the simultaneous scaling of data size, model size, and image resolution. We perform a scaling experiment with vision transformers up to 633M parameters in size (ViT-H/14) trained with up to 5K hours of human-like video data (long, continuous, mostly egocentric videos) with image resolutions of up to  $476 \times 476$  pixels. The efficiency of masked autoencoders (MAEs) as a self-supervised learning algorithm makes it possible to run this scaling experiment on an unassuming academic budget. We find that it is feasible to reach human-level object recognition capacity at sub-human scales of model size, data size, and image size, if these factors are scaled up simultaneously. To give a concrete example, we estimate that a 2.5B parameter ViT model trained with 20K hours (2.3 years) of human-like video data with a spatial resolution of  $952 \times 952$  pixels should be able to reach roughly human-level accuracy on ImageNet. Human-level competence is thus achievable for a fundamental perceptual capability from human-like perceptual experience (human-like in both amount and type) with extremely generic learning algorithms and architectures and without any substantive inductive biases.

Are modern self-supervised learning (SSL) algorithms as data efficient as humans when it comes to learning powerful internal models of the world? Here, we address this question with respect to a fundamental visual capability, namely real-world visual object recognition. A direct comparison of the data efficiency of deep learning models trained with SSL algorithms vs. humans is challenging for a number of reasons:

- **Mismatch in training data:** These models are typically trained with a very different kind of visual data from humans, both in terms of amount and type of data.
- **Mismatch in model size:** They are typically much smaller in size than the human brain, or even just the visual areas in the human brain, *e.g.* comparing the number of parameters in a model with the number of synapses in a human brain.
- **Mismatch in input size:** They typically operate on much smaller (lower dimensional) inputs than the human brain works with, *e.g.* comparing typical image sizes used in computer vision with the number of photoreceptors in the human retina.

We perform a scaling experiment to address these mismatches and ask if current SSL algorithms can reach human-level visual object recognition capabilities at sub-human scales of (human-like) data size, model size, and input size. An earlier work (Orhan, 2021) considered only data scaling (fixing the model size and input size to relatively small values) and concluded that current SSL algorithms may be orders of magnitude less data efficient than humans. Here, we consider the simultaneous scaling of all three factors and reach a different conclusion (we also train our models on a roughly

4× larger set of human-like videos in the current experiments). Code and models are available from the following public repository: <https://github.com/eminorhan/humanlike-vits>.

**Training data:** The full training data consists of a combination of six video datasets totaling 4971 hours of human-like video. The videos are “human-like” in two prominent aspects: (i) most of the videos are naturalistic, egocentric headcam videos recorded from the perspective of adult or child camera wearers during the course of their daily lives; (ii) they are temporally extended, continuous videos typically lasting tens of minutes to hours long in duration (unlike common video datasets in computer vision, which typically contain much shorter videos). Further details about the individual datasets can be found in the Appendix. The videos are temporally subsampled at the rate of 1 fps, yielding over 17M frames in total for the full dataset.

Our training data is a combination of multiple video datasets collected from multiple individuals, therefore it does not represent the visual experiences of a single individual, which creates a mismatch between the statistics of our training data and the typical visual experiences of a single person. However, a recent work (Orhan and Lake, 2023) found that models trained on length-matched subsets of different natural video datasets with similar temporal characteristics (long, continuous videos) perform remarkably similarly in downstream evaluation tasks, including in ImageNet and out-of-distribution versions of ImageNet (two benchmarks for visual object recognition we focus on in this work), despite notable qualitative differences between these video datasets in terms of both content and style.<sup>1</sup> This result suggests that natural video datasets with similar temporal characteristics may be, by and large, interchangeable with each other insofar as their effect on *overall performance* in downstream evaluation tasks is concerned (this does not necessarily preclude more fine-grained differences between models trained on different datasets or different subsets of the same dataset).

To determine the scaling of object recognition performance with the amount of human-like visual data used during self-supervised pretraining, we train models on the entire dataset and on continuously sampled random subsets of it (Orhan, 2021). Specifically, we train models on 100%, 10%, 1%, 0.1%, and 0.01% of the entire dataset (all sampled as continuous chunks or segments of video). These subsets contain from roughly 5000 to 0.5 hours of human-like video respectively, thus covering a 10000-fold range in data size. Since the subset selection is stochastic, we repeat it three times for each data size. This gives us a total of 13 different datasets (1 entire data + 4 proper subsets × 3 repeats), on which we train all of our models.

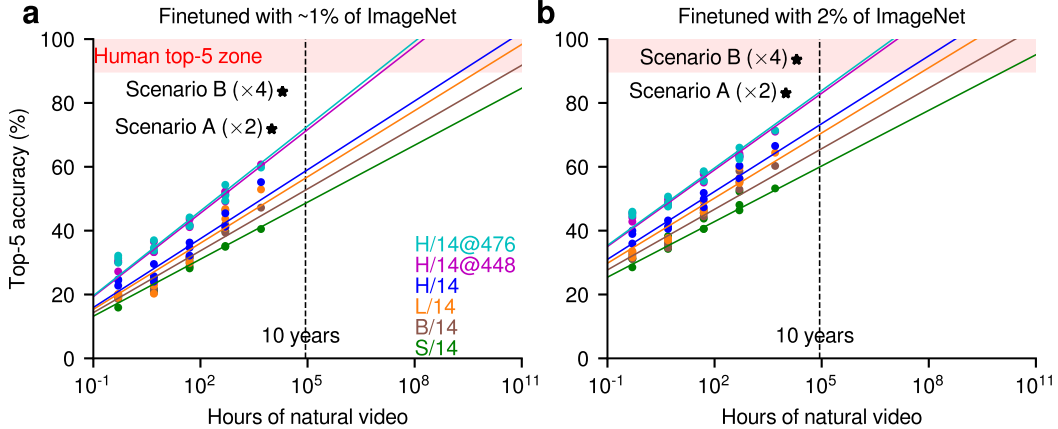
**Models:** We exclusively use vision transformer (ViT) models in our experiments (Dosovitskiy et al., 2020). We consider four standard sizes for our models: ViT-S, ViT-B, ViT-L, and ViT-H, all with  $14 \times 14$  patches. These models respectively have 22M, 87M, 304M, and 633M parameters, covering a 29-fold range in model size from the smallest to the largest model. For the image resolutions, we consider three different spatial resolutions:  $224 \times 224$  (0.05 MP),  $448 \times 448$  (0.2 MP), and  $476 \times 476$  pixels (0.23 MP), covering a roughly 4.5-fold range in spatial resolution (number of pixels). Due to computational costs, we use the larger two resolutions with the ViT-H model only. This gives us a total of 6 different model architectures (ViT-S/14, ViT-B/14, ViT-L/14, ViT-H/14, ViT-H/14@448, and ViT-H/14@476), which we train on each of the 13 datasets described in the previous paragraph for a grand total of  $6 \times 13 = 78$  pretrained models.

**SSL algorithm:** We use masked autoencoders (MAEs) as our SSL algorithm of choice (He et al., 2022). MAEs have a number of advantages over alternative self-supervised visual representation learning algorithms that are relevant for our purposes in this work. First, unlike most other visual SSL algorithms, they require very minimal data augmentation. He et al. (2022) show that using only random crops works well in MAEs. This is advantageous for our purposes, because heavy data augmentations used in other SSL algorithms make the training data less human-like. Second, MAEs work well with very high masking ratios (we use a masking ratio of 80%). Because only visible (unmasked) patches are passed through the encoder in MAEs, we can train much bigger models with larger image sizes than we would be able to do with other algorithms on a modest academic compute budget. We use a standard training configuration adapted from the [original MAE codebase](#).

**Results:** We use validation accuracy on ImageNet (Russakovsky et al., 2015) and OOD accuracy on out-of-distribution (OOD) versions of ImageNet (Geirhos et al., 2021) as our proxies for real-world visual object recognition capacity. These benchmarks emphasize both accuracy and robustness of

---

<sup>1</sup>This result is somewhat reminiscent of the robust emergence of general linguistic capabilities in children despite large differences in the amount and type of linguistic input they receive during their development both across and within cultures (Cristia et al., 2019; Bergelson et al., 2019).



**Figure 1:** Top-5 validation accuracy on ImageNet as a function of the amount of human-like video data used for self-supervised pretraining. Human-level performance is indicated by the red zone at the top ( $>90\%$ ). For reference, the developmentally relevant time scale of 10 years is indicated by the vertical dashed line. Different colors correspond to six different models identified in the legend. The solid lines represent the fits to Equation 1. Based on these fits, projected accuracies under two hypothetical scenarios are shown by the  $\star$  symbols (see Table 1 for a specification of these scenarios). **a** shows the results for the stringent finetuning condition (models finetuned with  $\sim 1\%$  of ImageNet) and **b** shows the results for the more permissive finetuning condition (models finetuned with  $2\%$  of ImageNet).

Scenario	$n$ (years)	$d$ (params)	$r$ (pixels)	Acc ( $\sim 1\%$ )	Acc ( $2\%$ )
Reference (actual)	0.6	0.6B	$476 \times 476$	59.8	71.4
Scenario A ( $\times 2$ proj.)	1.2	1.3B	$672 \times 672$	71.8	83.1
Scenario B ( $\times 4$ proj.)	2.3	2.5B	$952 \times 952$	83.5	93.6

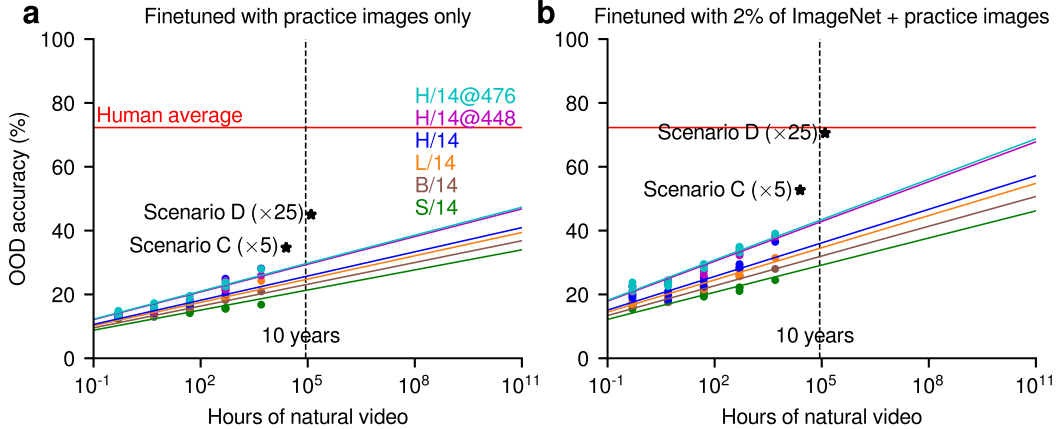
**Table 1:** Validation accuracy on ImageNet under an actual reference scenario and two projections based on the fits to Equation 1. The last two columns show the actual and projected accuracy under the stringent ( $\sim 1\%$ ) and permissive ( $2\%$ ) finetuning conditions respectively.

visual object recognition for real-world objects. Importantly, both benchmarks come with human performance estimates: for ImageNet, we consider  $90\%$  top-5 accuracy as a reasonable lower-bound on human performance based on the human experiments in [Russakovsky et al. \(2015\)](#); for the OOD versions of ImageNet, the average top-1 accuracy for humans is reported to be  $72.3\%$  in [Geirhos et al. \(2021\)](#).

For both benchmarks, we measure the performance of the pretrained models after few-shot supervised finetuning. We evaluate the models under two different finetuning conditions: (i) a stringent setting that exactly replicates the human experiments in [Russakovsky et al. \(2015\)](#) and [Geirhos et al. \(2021\)](#); and (ii) a more permissive setting that uses a small amount of additional labeled examples for supervised finetuning. The more permissive setting is intended to partially account for any prior (supervised) exposure humans might have had to ImageNet categories before the experiment. In the ImageNet benchmark, the strict setting uses exactly 13000 labeled examples from the ImageNet training set ( $\sim 1\%$  of the training set) for finetuning; the more permissive setting uses  $2\%$  of the training set for finetuning. In the OOD ImageNet benchmark, the strict setting only uses the 321 labeled practice images used by [Geirhos et al. \(2021\)](#) to familiarize the human participants in their experiments with the stimuli and the task during an initial practice session; the more permissive setting additionally uses  $2\%$  of the ImageNet training set for supervised finetuning.

Figure 1 shows the results of our scaling experiment on ImageNet. We model the effects of data size, model size, and image resolution on object recognition accuracy as a simple polynomial function in log space:

$$\text{accuracy} = \underbrace{(\alpha_n \log n + \beta_n)}_{\text{data scaling}} \underbrace{(\alpha_d \log d + \beta_d)}_{\text{model size scaling}} \underbrace{(\alpha_r \log r + \beta_r)}_{\text{image res. scaling}} \quad (1)$$



**Figure 2:** Results for the OOD ImageNet benchmark (figure description similar to Figure 1 above).

Scenario	$n$ (years)	$d$ (params)	$r$ (pixels)	Acc (str.)	Acc (perm.)
Reference (actual)	0.6	0.6B	$476 \times 476$	28.0	39.0
Scenario C ( $\times 5$ proj.)	2.8	3.2B	$1064 \times 1064$	34.7	52.7
Scenario D ( $\times 25$ proj.)	14.2	15.8B	$2380 \times 2380$	45.0	70.6

**Table 2:** OOD accuracy on the OOD ImageNet benchmark under an actual reference scenario and two projections based on the fits to Equation 1. The last two columns show the actual and projected accuracy under the stringent (finetuning with practice images only) and permissive (finetuning with 2% of ImageNet + practice images) conditions respectively.

where  $n$  is the amount of human-like video data used for self-supervised pretraining,  $d$  is the model size (number of parameters), and  $r$  is the image resolution (number of pixels). We thus fit 78 data points in each condition with the 6 parameters of this log-polynomial model. The Appendix contains results with two alternative parametric scaling functions. The results with these alternative scaling functions are broadly consistent with the results presented here in the main text using the scaling function in Equation 1.

Based on these fits, Table 1 describes three scenarios: (i) an actual reference scenario that corresponds to the largest and most performant model we have trained, *i.e.* a ViT-H/14 model trained with all  $\sim 5K$  hours of our human-like video dataset with a spatial resolution of  $476 \times 476$  pixels; (ii) a hypothetical scenario A where we double ( $\times 2$ ) each of  $n$ ,  $d$ ,  $r$  with respect to the reference scenario; (iii) a hypothetical scenario B where we quadruple ( $\times 4$ ) each of  $n$ ,  $d$ ,  $r$  with respect to the reference scenario. Table 1 shows the actual and projected object recognition accuracy for each of these scenarios (actual accuracy for the reference scenario and projected accuracy based on the model fits for the hypothetical scenarios). In the hypothetical scenario B ( $\times 4$  projection) and under the more permissive finetuning condition (finetuning with 2% of ImageNet), the projected accuracy already exceeds the 90% top-5 accuracy we have set as the lower bound on human-level accuracy on ImageNet. This result is quite encouraging, since the increases in data size, model size, and image resolution required under this scenario are relatively modest: *i.e.* this scenario only requires a 2.5B parameter ViT model trained with  $\sim 20K$  hours (2.3 years) of human-like video with a spatial resolution of  $952 \times 952$  pixels.

Figure 2 shows the results of our scaling experiment on the OOD ImageNet benchmark. Similar to the case of ImageNet above, Table 2 describes an actual reference scenario and two hypothetical projections (a  $\times 5$  projection and a  $\times 25$  projection with respect to the reference scenario) based on the log-polynomial fits to the scaling data and shows the actual and projected OOD accuracies under these scenarios. In this case, we estimate that  $n$ ,  $d$ , and  $r$  need to be scaled up more compared to the case of ImageNet in order to reach human-level performance, but the  $\times 25$  scenario (scenario D) is projected to achieve human-level OOD accuracy under the permissive finetuning condition and, with its requirements of 14.2 years of data, a 15.8B parameter model, and a spatial resolution of  $2380 \times 2380$  pixels (5.7 MP), we believe this scenario is still broadly within the human bounds in all of  $n$ ,  $d$ , and  $r$ . The Appendix contains further discussion of the results.

## References

- Bergelson, E., Casillas, M., Soderstrom, M., Seidl, A., Warlaumont, A. S., and Amatuni, A. (2019). What do North American babies hear? A large-scale cross-corpus analysis. *Developmental Science*, 22(1):e12724.
- Cristia, A., Dupoux, E., Gurven, M., and Stieglitz, J. (2019). Child-directed speech is infrequent in a forager-farmer population: A time allocation study. *Child Development*, 90(3):759–773.
- Damen, D., Doughty, H., Farinella, G. M., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., and Wray, M. (2018). Scaling Egocentric Vision: The EPIC-KITCHENS Dataset. In *European Conference on Computer Vision (ECCV)*.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Geirhos, R., Narayanappa, K., Mitzkus, B., Thieringer, T., Bethge, M., Wichmann, F. A., and Brendel, W. (2021). Partial success in closing the gap between human and machine vision. *Advances in Neural Information Processing Systems*, 34:23885–23899.
- Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al. (2022). Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012.
- Gu, C., Sun, C., Ross, D. A., Vondrick, C., Pantofaru, C., Li, Y., Vijayanarasimhan, S., Toderici, G., Ricco, S., Sukthankar, R., et al. (2018). AVA: A video dataset of spatio-temporally localized atomic visual actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6047–6056.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. (2022). Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009.
- Lee, Y. J., Ghosh, J., and Grauman, K. (2012). Discovering important people and objects for egocentric video summarization. In *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1346–1353. IEEE.
- Orhan, A. E. (2021). How much human-like visual experience do current self-supervised learning algorithms need in order to achieve human-level object recognition? *arXiv preprint arXiv:2109.11523*.
- Orhan, A. E. and Lake, B. M. (2023). What can generic neural networks learn from a child’s visual experience? *arXiv preprint arXiv:2305.15372*.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252.
- Singh, K. K., Fatahalian, K., and Efros, A. A. (2016). KrishnaCam: Using a longitudinal, single-person, egocentric dataset for scene understanding tasks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*.
- Sullivan, J., Mei, M., Perfors, A., Wojcik, E., and Frank, M. C. (2021). SAYCam: A large, longitudinal audiovisual dataset recorded from the infant’s perspective. *Open Mind*, 5:20–29.
- Williamson, S. J. and Cummins, H. Z. (1983). *Light and Color in Nature and Art*, volume 1. Wiley New York.

## Appendix

### Comparing deep learning models vs. humans

**Model size:** A conservative estimate of 10% of the brain dedicated to visual processing and 10K synapses per neuron suggests an estimate of 100T synapses or connections dedicated to visual processing in the human brain. This number is several orders of magnitude larger than the number of parameters used in current deep learning models.

**Input size:** The human retina contains roughly 6M color sensitive cone receptors, very tightly concentrated within a few degrees of visual angle around the fovea (Williamson and Cummins, 1983). By comparison, the average size of an ImageNet image (after resizing) is  $310 \times 256$  pixels, or 0.08 megapixels (MP), in the most commonly used preprocessing pipeline today, which is about two orders of magnitude smaller in spatial resolution.

### Training data

The combined training dataset consists of the following video datasets:

**Ego4D:** Ego4D is a large dataset of egocentric headcam videos recorded by 923 unique participants while performing daily activities (Grauman et al., 2022). The total length of the videos in Ego4D is about 3670 hours. The dataset can be accessed from [this address](#) after signing a license agreement.

**AVA (v2.2):** AVA is a dataset of films (Gu et al., 2018). The dataset comes with rich annotations for parts of these films. However, we use the entire films (without annotations) for self-supervised learning in this work. We use all videos in both `trainval` and `test` splits (299 and 131 videos respectively). The total length of the videos from this dataset is about 636 hours. The dataset is publicly accessible from [this address](#).

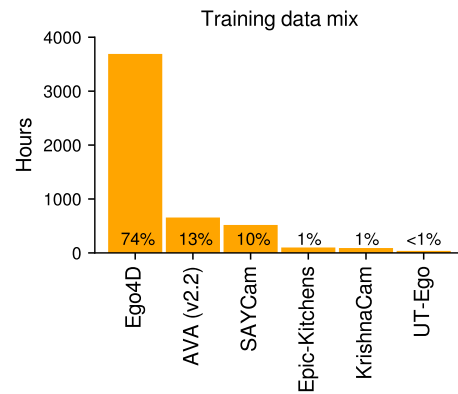
**SAYCam:** SAYCam is a large, longitudinal dataset of headcam videos recorded from the perspective of three young children between the ages of 6 to 32 months (Sullivan et al., 2021). The combined length of the videos from all three children is about 498 hours. Recordings were made approximately once a week for 1-2 hours (usually continuously) over the course of a 2.5 year period during the early development of the children. This dataset is not publicly accessible, but researchers affiliated with an institution can apply for access on the [Databrary](#) repository for behavioral science.

**Epic-Kitchens:** This dataset contains relatively long headcam videos from multiple participants performing daily culinary chores in their kitchens (Damen et al., 2018). The total length of the videos in this dataset is about 80 hours. The dataset is publicly accessible from [this address](#).

**KrishnaCam:** KrishnaCam is a dataset of headcam videos recorded by a graduate student (Singh et al., 2016). It contains long, continuous, egocentric videos of daily episodes in the life of the graduate student. The total length of the videos in the dataset is about 70 hours. The dataset is publicly accessible from [this address](#).

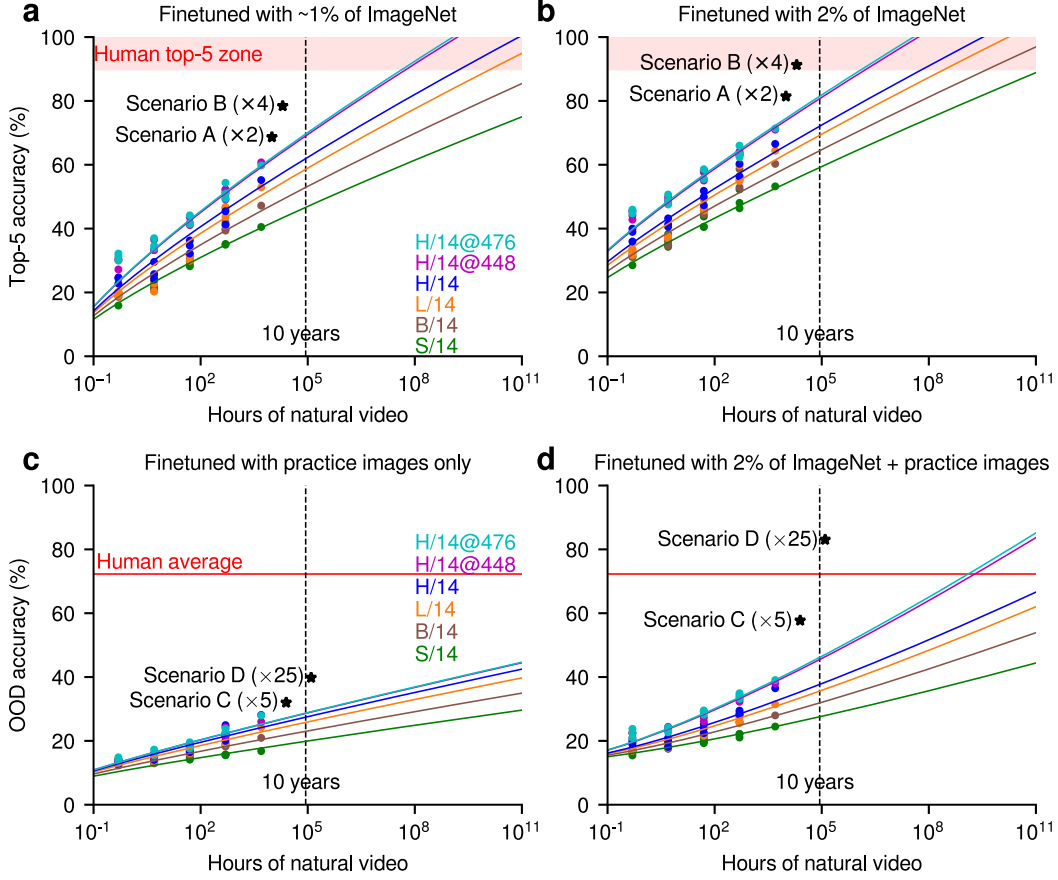
**UT Ego:** UT Ego contains four continuous, egocentric headcam videos, each 3-5 hours long (Lee et al., 2012) and is publicly accessible from [this address](#).

Figure 3 shows a breakdown of the individual video datasets in our combined training data together with their sizes (in hours). By far the biggest contributor to our combined training data is Ego4D (74%). The private SAYCam dataset constitutes only about 10% of the training data. Therefore, we expect that those who wish to replicate the experiments in this work using only publicly accessible data sources should get very similar results to those reported here.



**Figure 3:** Human-like video datasets used for SSL.





**Figure 4:** Results with the alternative scaling function in Equation 2. **a** and **b** show the results for the ImageNet benchmark; **c** and **d** show the results for the OOD ImageNet benchmark.

### Alternative scaling functions

Figure 4 shows the results with a 10-parameter scaling function of the form:

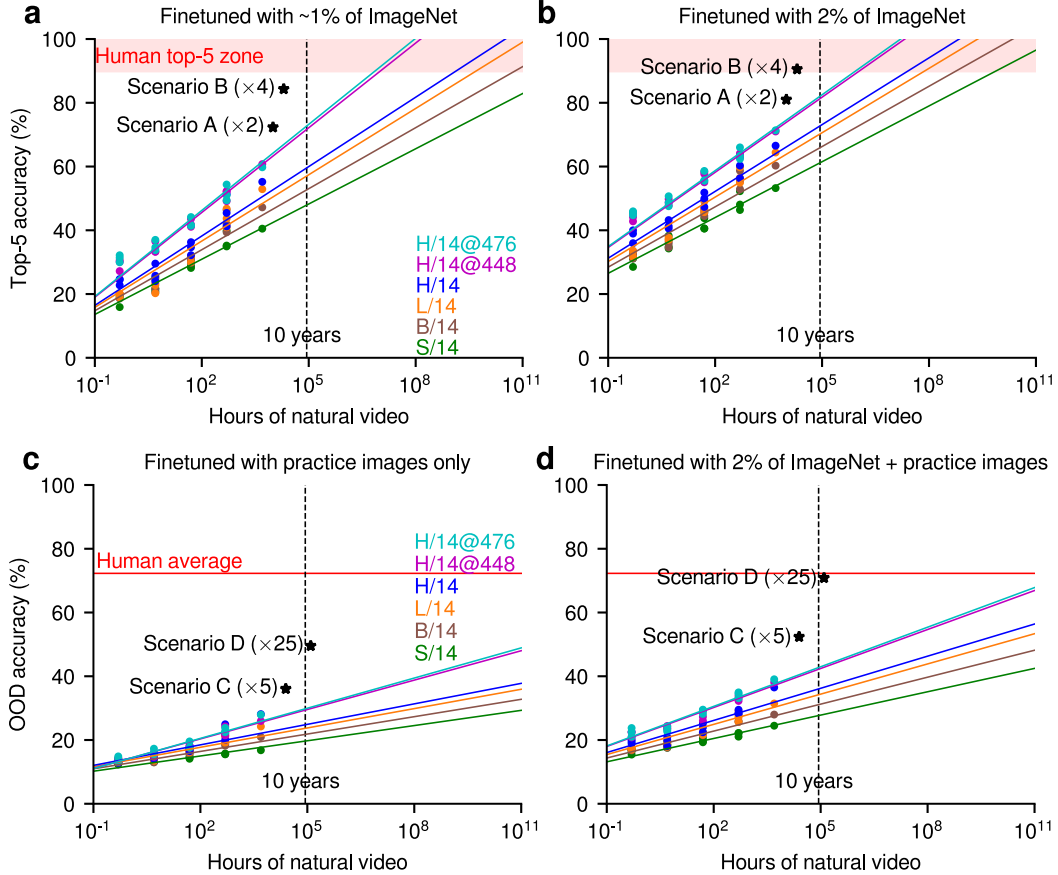
$$\text{accuracy} = (\alpha_n \tilde{n}^{\beta_n} + \gamma_n)(\alpha_d \tilde{d}^{\beta_d} + \gamma_d)(\alpha_r \tilde{r}^{\beta_r} + \gamma_r) + \delta \quad (2)$$

Figure 5 shows the results with an 8-parameter scaling function of the form:

$$\text{accuracy} = \alpha_n \tilde{n} + \alpha_d \tilde{d} + \alpha_r \tilde{r} + \alpha_{nd} \tilde{n} \tilde{d} + \alpha_{nr} \tilde{n} \tilde{r} + \alpha_{dr} \tilde{d} \tilde{r} + \alpha_{ndr} \tilde{n} \tilde{d} \tilde{r} + \gamma \quad (3)$$

where, in both equations,  $\tilde{n}$ ,  $\tilde{d}$ ,  $\tilde{r}$  represent  $n$ ,  $d$ ,  $r$  in the log space: *i.e.*  $\tilde{n} = \log n$ ,  $\tilde{d} = \log d$ , and  $\tilde{r} = \log r$ , respectively. We note that purely additive scaling functions with no multiplicative interactions between different factors do a poor job of fitting the data, as the data (especially the OOD accuracy) display a clear multiplicative pattern (see *e.g.* Figure 2 and panels **c-d** in Figures 4-5). This is why we consider more general parametric scaling functions that combine both additive and multiplicative effects in Equations 1-3.

In Figures 4-5, panels **a** and **b** show the results for the ImageNet benchmark, panels **c** and **d** show the results for the OOD ImageNet benchmark under the stringent and permissive finetuning conditions, respectively. The reference scenario and the hypothetical scenarios considered here are identical to those in the main text. The scaling functions here and in the main text were fit with the `scipy.optimize.minimize` function using the adaptive Nelder-Mead optimization method and the parameters were initialized to small values.



**Figure 5:** Results with the alternative scaling function in Equation 3. **a** and **b** show the results for the ImageNet benchmark; **c** and **d** show the results for the OOD ImageNet benchmark.

### Further discussion

Our results suggest that human-level accuracy and robustness in visual object recognition are achievable from human-like visual experience at sub-human scales of data size, model size, and image resolution, using highly generic self-supervised learning algorithms and deep learning architectures without any strong inductive biases. This result is not inconsistent with the earlier work on this question (Orhan, 2021), which only considered the scaling of data size (fixing the model size and image resolution to relatively small values). In fact, in the current set of experiments, when we just consider the data size scaling for the ViT-S/14 model, which is similar in size to the ResNeXt-50 model trained with the DINO algorithm in Orhan (2021), and fix the spatial resolution to  $224 \times 224$  pixels, we get estimates in the range of 1M-100M years of human-like visual experience in order to reach human-level accuracy on ImageNet (see the green lines in Figures 1a and 1b). These estimates are comparable to the estimated range of 1M-1B years reported in Orhan (2021).