

Supplementary material for *Concept Generalization in Visual Representation Learning*

Mert Bulent Sariyildiz^{1,2}, Yannis Kalantidis¹ Diane Larlus¹ Karteek Alahari²
¹ NAVER LABS Europe ² Inria*

Contents

A Details of ImageNet-CoG levels	1
A.1 Selection criteria for unseen concepts	1
A.2 Implementation details for Lin similarity	2
A.3 Potential label noise in ImageNet-CoG	2
A.4 Statistics for ImageNet-CoG	2
B Evaluation protocol of ImageNet-CoG	2
B.1 Feature extraction and preprocessing	2
B.2 Training classifiers	3
C Extended results	3
C.1 Full set of results	3
C.2 What if we fine-tuned the backbones?	4
D An alternative semantic similarity	4
D.1 ImageNet-CoG with word2vecD	4
E The ImageNet 2021 release	6

This supplementary material is structured as follows. Sec. **A** details the design choices we made in creating the concept generalization levels of ImageNet-CoG (and is an extension of Sections 3.2 and 3.3 of the main paper). Sec. **B** describes the preprocessing pipeline for the models we benchmark in Sec. 4 of the main paper and also provides implementation details of our evaluation protocol (extending Sec. 3.4 of the main paper). Sec. **C** presents the complete set of results for the 31 models we evaluate in Sec. 4 of the main paper. Sec. **D** discusses how creating concept generalization levels with WordNet ontology and Lin similarity compares to creating them using the textual descriptions of concepts and a pretrained language model (and is an extension of Sec. 3.3 of the main paper). Finally, Sec. **E** discusses the impact of a recent update of ImageNet (impacting both ImageNet-21K and ImageNet-1K) on the results and future evaluation of our benchmark.

*UGA, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

A. Details of ImageNet-CoG levels

We begin by detailing the steps to create the concept generalization levels of ImageNet-CoG. They include the selection of eligible unseen concepts in ImageNet-21K [11] (IN-21K) and the implementation details for Lin similarity [34]. We then briefly discuss the potential noise from missing labels in ImageNet-CoG. At the end of the section, we provide basic statistics about the ImageNet-CoG levels, i.e., the exact number of images per concept in each ImageNet-CoG level.

A.1. Selection criteria for unseen concepts

As described in Sec. 3.2 of the main paper, prior to creating the concept generalization levels, we determine a set of *eligible* unseen concepts in IN-21K. To determine these concepts, we implemented the following steps.

- We started with the whole set of IN-21K concepts (21,841) of the Fall 2011 release and excluded the ones from IN-1K, as they are the seen concepts.
- In order to create levels whose size is comparable to IN-1K, following the design choices made for IN-1K, we removed concepts with fewer than 782 images (note that any concept in IN-1K contains at least 782 images and 50 of those are used within the test set).
- It was shown that some of the concepts under the “person” sub-tree in IN-21K can be offensive or visually inappropriate, which may lead to undesirable behavior in downstream applications [66]. We thus excluded the entire “person” sub-tree.
- We also excluded all concepts that are parents in the ontology of eligible concepts, as this could lead to issues with the labeling strategy. Concretely, for any c_1 and c_2 in IN-21K, we exclude c_1 if c_1 is a parent of c_2 .
- Finally, we manually inspected the remaining unseen concepts and found 70 potentially problematic concepts, which may be considered to be offensive, or too ambiguous to distinguish. Examples of such concepts include the very generic “People” (any group of human beings, men or women or children, collectively) or “Orphan” (a young animal without a mother) concepts. The list of such

manually discarded concepts is given in Tab. A. After sequentially applying these steps, we are left with 5146 eligible unseen concepts. The complete list of eligible unseen concepts, along with the concepts in each level $L_{1/2/3/4/5}$ can be found on our project website.

A.2. Implementation details for Lin similarity

As described in Sec. 3.3 of the main paper, to produce the concept generalization levels, we sort the eligible unseen concepts by decreasing semantic similarity to the seen concepts in IN-1K. We used Lin similarity [34] as the semantic measure, which computes the relatedness of two concepts defined in a taxonomy.

Computing this similarity between two concepts requires their *information content* in the taxonomy. Following [44] and [46], we define the information content of a concept as $-\log p(c)$, where $p(c)$ is the probability of encountering concept c in the taxonomy. In our study, the taxonomy is a fragment of WordNet including all the concepts in IN-21K and their parents till the root node of WordNet: “Entity”. Probability of a concept ranges between $[0, 1]$ such that if c_2 is a parent of c_1 then $p(c_1) < p(c_2)$, and the probability of “Entity” becomes 1.

In order to get superior-subordinate relationships between the concepts, we use WordNet-3.0 (the version ImageNet [11] is built on) implementation in the NLTK library [78].

A.3. Potential label noise in ImageNet-CoG

It has been shown recently [70] that ImageNet-1K (IN-1K) has missing-label noise. We can assume this extends to ImageNet-21K. Unfortunately, this type of noise is really difficult to correct and beyond the scope of our benchmark. However, we devise an experiment to get a sense of how much this noise could be. We take ResNet-50 classifiers for L_K and apply them to all the images of the IN-1K val set and vice versa (IN-1K classifiers on L_5 val). After inspecting samples that are predicted with very high confidence (> 0.99 , about 2.7% of the images), we observe *several* cases where an unseen concept has (arguably) been seen during training without its label. Some examples are shown in Fig. A. Given the low percentage of very confident matches and the fact that [70] does not show a big change in performance after re-training with the noise corrected, we believe that this type of labeling noise does not significantly affect our findings.

A.4. Statistics for ImageNet-CoG

Number of images in each level. After selecting 1000 concepts for each level, we ensured that the image statistics are similar to those of IN-1K [47], i.e., we cap the number of images for each concept to a maximum of 1350 (1300 training + 50 testing). Note that we kept the same set of selected

images per concept for all the experiments. We provide the complete list of image filenames in our code repository for reproducibility. In Fig. E, we plot the number of images per concept for each of the five levels and for IN-1K. We note a minor class imbalance in all the generalization levels from these plots. To investigate if this imbalance had any effect on the observations of our benchmark, we further evaluated a subset of the models analyzed in Sec. 4 of the main paper on a variant of the benchmark, where we randomly sub-sampled images from all the selected concepts to result in the same number of 732 training images, i.e., on class-balanced levels. Apart from the overall reduced accuracy as a result of smaller datasets, this experiment produced similar results to the ones shown in the main text, and all our observations continue to hold. We attribute this to the fact that imbalance is minimal.

B. Evaluation protocol of ImageNet-CoG

In this section, we provide additional implementation details of our evaluation protocol, thus extending Sec. 3.4 of the main paper.

B.1. Feature extraction and preprocessing

We establish evaluation protocols for ImageNet-CoG with image features extracted from pretrained visual backbones. To extract these features, we first resize an image such that its shortest side becomes S pixels, then take a center crop of size $S \times S$ pixels. To comply with the testing schemes of the models, for all the backbones we set $S = 224$, except *a*-Inception-v3 [53] ($S = 299$), *a*-DeiT-B-distilled [58] ($S = 384$), *a*-EfficientNet-B1 [54] ($S = 240$) and *a*-EfficientNet-B4 [54] ($S = 380$).

We also adapt their normalization schemes to be compatible with the data augmentation pipeline of the pretrained models. Concretely, we normalize each image by first dividing them by 255 (so that each pixel value is in $[0, 1]$), then applying mean and standard deviation normalization to the pixels, i.e., subtracting $[0.485, 0.456, 0.406]$ from the RGB channels and dividing them by $[0.229, 0.224, 0.225]$, respectively. Note that for *d*-CLIP [43] we use mean $[0.481, 0.457, 0.408]$ and std $[0.268, 0.261, 0.275]$, and do not apply normalization for *s*-SimCLR-v2 [7].

Tab. B lists the set of unique backbone architectures considered in our study, and the dimensionality of the produced feature representations. For all the architectures trained in a supervised way, we extract features from the penultimate layers, i.e., before the last fully-connected layers making class predictions. For self-supervised learning methods, we follow the respective papers and extract features from the layer learned for transfer learning.

n00005787	n00288384	n00466377	n00466524	n00466630	n00474568	n00475014	n00475273	n00475403	n00483313
n00483409	n00483508	n00483848	n01314388	n01314663	n01314781	n01317294	n01317813	n01317916	n01318381
n01318894	n01321770	n01322221	n01323355	n01323493	n01323599	n01324431	n01324610	n01515303	n01517966
n01526521	n01862399	n01887474	n01888181	n02075612	n02152881	n02153109	n02156871	n02157206	n02236355
n02377063	n02377291	n02472987	n02475078	n02475669	n02759257	n02767665	n02771004	n03198500	n03300216
n03349771	n03393017	n03443005	n03680512	n04164406	n04193377	n04224543	n04425804	n04516354	n04979002
n06255081	n06272612	n06274760	n07942152	n08182379	n08242223	n08578517	n09828216	n10300303	n13918274

Table A: WordNet IDs of the 70 concepts considered problematic, therefore removed from the eligible list of unseen concepts.





	Concepts in L_5		Concepts in IN-1K	
				
Ground Truth	Rock climbing	Handball	Ptarmigan	Dalmatian
Predicted labels	Predictions by IN-1K classifier Cliff Soccer ball		Predictions by L_5 classifier Peak Snow	

Figure A: Illustration of the label noise in ImageNet-CoG.

Model	Feature Dim.
All models with ResNet [23] backbone	2048
<i>a</i> -T2T-ViT-t-14 [68]	384
<i>a</i> -DeiT-S [58]	384
<i>a</i> -DeiT-B-distilled [58]	768
<i>a</i> -NAT-M4 [36]	1536
<i>a</i> -EfficientNet-B1 [54]	1280
<i>a</i> -EfficientNet-B4 [54]	1792
<i>a</i> -VGG19 [49]	4096

Table B: Unique architectures used by the models we evaluate in Sec. 4 of the main paper and the dimensionality of the feature vectors we extract from these architectures.

B.2. Training classifiers

In ImageNet-CoG, we perform two types of transfer learning experiments on each set of concepts, i.e., IN-1K or our concept generalization levels $L_{1/2/3/4/5}$ (see Sec. 4.2 of the main paper): (i) linear classification with all the available data, (ii) linear classification with a few randomly selected training samples. Both sets of experiments use the same test set, i.e., all the test samples.

In each of these experiments, we train a classifier with the features extracted using a given model. In order to evaluate each model in a fair manner in each setting, it is important to train each classifier in the best possible way.

We perform SGD to train classifiers, with momentum=0.9

updates, using batches of size 1024, and apply weight decay regularization to parameters. We choose learning rate and weight decay hyper-parameters on a validation set randomly sampled from the training set of each concept domain (20% of the training set is randomly sampled as a validation set for each concept domain). We sample 30 (learning rate, weight decay) pairs using Optuna [1] with a parzen estimator [77]. We then train the final classifier (with the hyper-parameters chosen from the previous validation step) on the full training set and report performance on the test set. We repeat this process 5 times with different seeds. This means that, in each repetition, we take a different random subset of the training set as a validation set and start hyper-parameter tuning with different random pairs of hyper-parameters. Despite this stochasticity, the overall pipeline is quite robust, with standard deviation in most cases less than 0.2, therefore, not clearly visible in figures. We will release training configurations along with our benchmark on our project website.

C. Extended results

C.1. Full set of results

In Sec. 4 of the main paper, we evaluate concept generalization performance for 31 models (listed in Tab. 1 of the main paper) on ImageNet-CoG. Figures 3 and 4 of the main paper report the results of training logistic regression classifiers with all the available training data for each concept (discussed in Sec. 4.2.1), and training it with a few samples per concept (discussed in Sec. 4.2.2), respectively. Due to

space constraints, although Fig. 3 includes the results for all the models on all concept generalization levels, Fig. 4 provides only a selection of the few-shot results. In this section, we present the full set of results for all the methods when training with few and all data samples in table form. We also present the full set of figures for all the methods and levels when training with a few training samples per concept.

How fast can models adapt to unseen concepts. For completeness, we present the scores of all the models for $N = \{1, 2, 4, 8, 16, 32, 64, 128, \text{All}\}$ on IN-1K and $L_{1/2/3/4/5}$ in Fig. F (raw scores) and Fig. G (relative scores). These results, grouped by levels (i.e., for IN-1K and for $L_{1/2/3/4/5}$ separately) are also presented in Tables D, E, F, G, H, I respectively. These additional results complement Sec. 4.2.2 of the main paper.

Generalization to unseen concepts. To access the raw numbers of the results discussed in Sec. 4.2.1 of the main paper, we refer the reader to Tables D, E, F, G, H, I and the $N = \text{All}$ columns, which correspond to the scores shown in Fig. 3(a) of the main paper.

C.2. What if we fine-tuned the backbones?

Our benchmark and evaluation protocol are based on the assumption that good visual representations should generalize to different tasks with minimal effort. In fact, we explicitly chose to decouple representation learning and only consider frozen/pretrained backbones as feature extractors. We then evaluate how well pretrained representations transfer to concepts *unseen* during representation learning. Fine-tuning the models would therefore go against the main premise of our benchmark: after fine-tuning all concepts are “seen” during representation learning, i.e., the feature spaces can now be adapted. It would then be unclear: are we measuring the generalization capabilities of the pretraining strategy or of the fine-tuning process? How much does the latter affect generalization? We consider such questions out of the scope of our study. In fact, learning linear classifiers on top of pre-extracted features additionally allows us to exhaustively optimize hyper-parameters for all the methods and levels (see Section B.2), making sure that comparisons are fair across all models.

Measuring performance *relative to fine-tuning*, would however verify that the observed performance drops are due to increasing semantic distance and not variabilities across the levels. To this end, we fine-tune ResNet50 (pretrained on IN-1K) on IN-1K and on levels $L_{1/2/3/4/5}$ separately. Then we compare their performance with the protocol we chose for our benchmark, i.e. the case where we learn linear classifiers on top of pre-extracted features. In Fig. B, we show the *relative* scores of the linear classifiers on top of pre-extracted (labeled as “Pre-extracted”) against fine-tuned ResNet50s (labeled as “Fine-tuned”).

We observe that pre-extracted features become less and

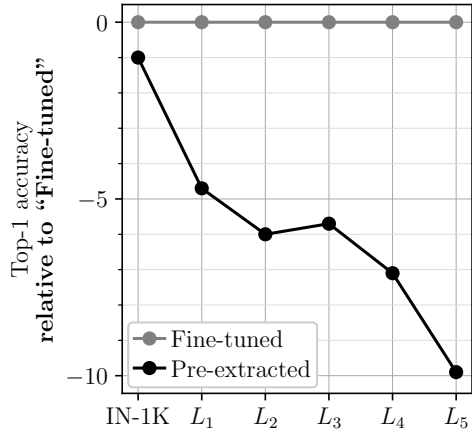


Figure B: Comparison of training linear classifiers on **pre-extracted features** vs. **fine-tuning** backbones on each level. Y-axis shows the top-1 accuracies obtained **relative** to the accuracy of the fine-tuned models.

less informative for unseen concepts as we move from IN-1K to L_5 , supporting our main assumption that semantically less similar concepts are harder to classify.

D. An alternative semantic similarity

D.1. ImageNet-CoG with word2vecD

One of the requirements for studying concept generalization in a controlled manner is a knowledge base that provides the semantic relatedness of any two concepts. As IN-21K is built on the concept ontology of WordNet [41], in Sec. 3.3 of the main paper we leverage its graph structure, and propose a benchmark where semantic relationships are computed with the Lin measure [34].

As mentioned in Sec. 3 of the main paper, the WordNet ontology is hand-crafted, requiring expert knowledge. Therefore similarity measures that exploit this ontology (such as Lin) are arguably reliable in capturing the semantic similarity of concepts. However, it could also be desirable to learn semantic similarities automatically, for instance, using other knowledge bases available online such as Wikipedia. In this section, we investigate if such knowledge bases could be used in building our ImageNet-CoG.

With this motivation, we turn our attention to semantic similarity measures that can be learned over textual data describing the IN-21K concepts. Note that each IN-21K concept is provided with a name¹ and a short description.² The idea is to use this information to determine the semantic relatedness of any two concepts.

To this end, we leverage language models to map the

¹<http://www.image-net.org/archive/words.txt>

²<http://www.image-net.org/archive/gloss.txt>




Figure C: **Semantic similarities** of the concepts captured by (i) Lin similarity [34] on WordNet graph [41] and (ii) cosine similarity of word2vec embeddings [79] extracted from textual descriptions of concepts, *vs.* **visual similarities** encoded by ResNet50, on IN-1K and generalization levels $L_{1/2/3/4/5}$ of ImageNet-CoG. We report the performance of linear logistic regression classifiers trained on features extracted from the global average pooling layer of ResNet50. The **orange line** shows results obtained on 1000 *random* unseen concepts (line represents the mean accuracy obtained over 15 random splits).

textual description of any concept into an embedding vector, such that the semantic similarity between two concepts can be measured as the similarity between their representations in that embedding space. We achieve this through the skip-gram language model [40], which has been extensively used in many natural language processing tasks, to extract “word2vec” representations of all concepts. However, we note that the name of many IN-21K concepts are *named entities* composed of multiple words, yet the vanilla skip-gram model tokenizes a textual sequence into words. We address this issue following [80] that learns a skip-gram model by taking into account such named entities. Specifically, we use the skip-gram model trained on Wikipedia³ by the Wikipedia2Vec software [79].

We compute the word2vec embeddings of IN-21K concepts as follows. Firstly, we combine the names and descriptions of all concepts and learn tf-idf weights for each unique word. Secondly, for each concept, we compute two word2vec representations: one for the concept name, and one for the concept description, by averaging the word2vec representations of the words that compose them. These two average vectors are added and used as the final word2vec representation of the concept. Finally, as the semantic similarity measure, we simply use the cosine similarity between

Dataset	Split	# Images
IN-1K [47]	train	1281167
IN-1K [47]	val	50000
IN-1K-blurred [81]	train	1281066
IN-1K-blurred [81]	val	49997

Table C: **Comparison of the number of images in IN-1K and IN-1K-blurred.**

the word2vec representations of two concepts:

$$\text{sim}_{\text{w2v}}(c_1, c_2) = \frac{\omega_{c_1} \top \omega_{c_2}}{\|\omega_{c_1}\| \cdot \|\omega_{c_2}\|}, \quad (1)$$

where ω_c denotes the word2vec representation of concept c .

Recall that in Sec. 3.3 of the main paper, first we rank the 5146 eligible unseen concepts in IN-21K (which remain after our filtering, as explained in Sec. 3.3 of the main paper and Sec. A.1), w.r.t. their Lin similarity to the concepts in IN-1K. Then, we sub-sample 5000 concepts to construct concept generalization levels. To create another benchmark based on the textual information of the concepts as described above, we could repeat this procedure by replacing Lin similarity with the cosine similarity we defined in Eq. (1). However, this could select a different sub-set of 5000 concepts, which, in turn, would produce two benchmarks with different sets of unseen concepts. To prevent this, we re-rank the 5000 concepts selected by the Lin similarity, based on their text-based cosine similarity to IN-1K concepts. Then we simply divide the re-ordered concepts into 5 disjoint sequential sets.⁴

We compare the two benchmarks constructed with different knowledge bases (i.e., using the WordNet graph *vs.* textual descriptions) for our baseline model ResNet50 [23] that is pretrained on the seen concepts (IN-1K) for image classification, following our standard protocol. Concretely, first, we extract image features from the penultimate layer of the ResNet50, then we train linear classifiers on each concept domain separately.

We report results in Fig. C for the two benchmarks as well as randomly selected subsets of 1000 concept each. We see that the benchmark constructed using the WordNet ontology [41] and the Lin similarity [34] yield much more challenging concept generalization levels than the one obtained using textual data and a skip-gram language model [79] pretrained on Wikipedia. This is especially visible when comparing classification performance on the levels $L_{3/4/5}$ produced by each technique. We argue that this could be due to the fact that WordNet is an ontology hand-crafted by experts and is able to better approximate the semantic similarity of two concepts compared to the learned skip-gram model.

⁴Note that, given that the percentage of discarded concepts is very small (less than 3%, as 146 concepts are discarded from the 5146 eligible ones), this choice has minimal impact anyway.

³April 2018 version of the English Wikipedia dump.

ResNet	Transformer	NAS & Other	Self-Supervision	Web data	Regularization
● ResNet50 (23.5M)	▲ <i>a</i> -T2T-ViT-t-14 (21.1M)	★ <i>a</i> -Inception-v3 (25.1M)	■ <i>s</i> -DINO	★ <i>s</i> -SimCLR-v2	■ <i>d</i> -Semi-Sup
■ <i>a</i> -ResNet152 (58.1M)	▶ <i>a</i> -DeiT-S (21.7M)	✦ <i>a</i> -EfficientNet-B1 (6.5M)	▲ <i>s</i> -SwAV	★ <i>s</i> -MoCo-v2	▲ <i>d</i> -Semi-Weakly-Sup
	▶ <i>a</i> -DeiT-S-distilled (21.7M)	✦ <i>a</i> -EfficientNet-B4 (17.5M)	▶ <i>s</i> -BarlowTwins	★ <i>s</i> -MoChi	▲ <i>d</i> -MoPro
	▶ <i>a</i> -DeiT-B-distilled (86.1M)	✦ <i>a</i> -NAT-M4 (7.6M)	▶ <i>s</i> -OBoW	◆ <i>s</i> -CompReSS	▲ <i>d</i> -CLIP
		✦ <i>a</i> -VGG19 (139.6M)	▶ <i>s</i> -BYOL	◆ <i>s</i> -InfoMin	
					■ <i>r</i> -ReLabel
					▲ <i>r</i> -CutMix
					▶ <i>r</i> -MixUp
					▶ <i>r</i> -Manifold-MixUp
					▼ <i>r</i> -Adv-Robust
					★ <i>r</i> -MEAL-v2

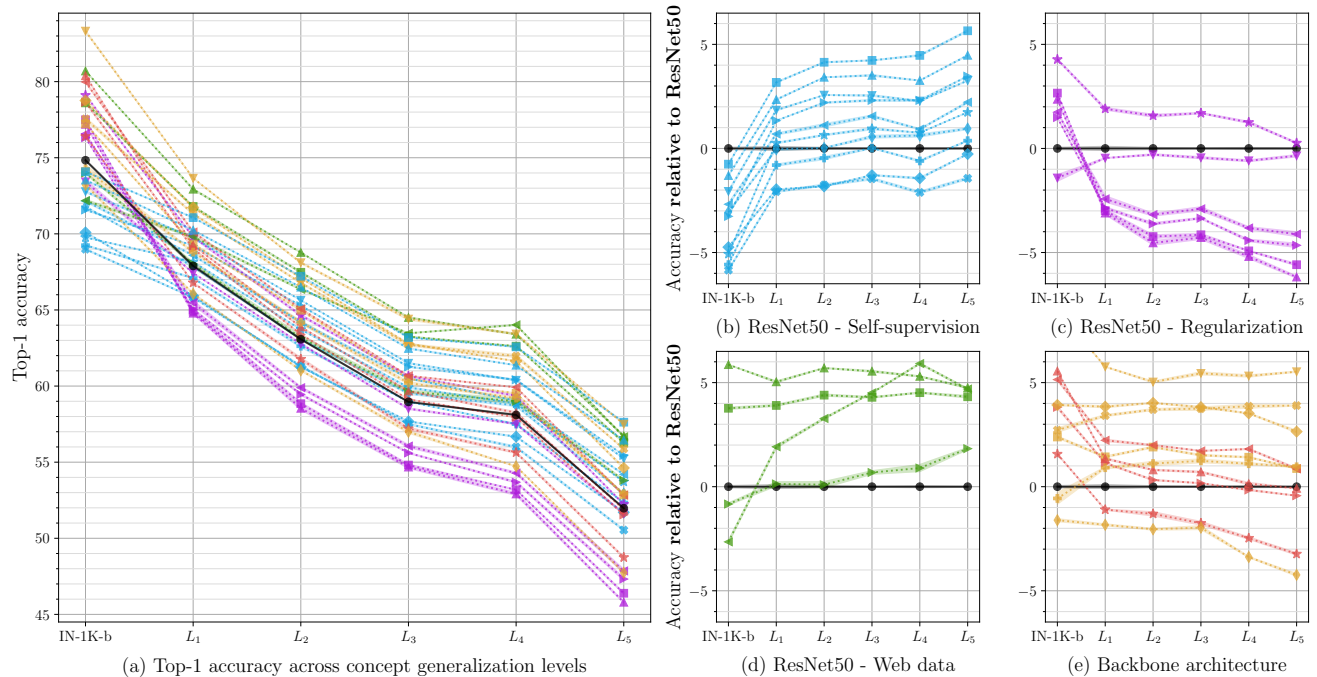


Figure D: **Linear classification on ImageNet-CoG using blurred images for IN-1K.** Top-1 accuracies for all the 31 models listed in Tab. 1 of the main paper, after training logistic regression classifiers on the **blurred version** of IN-1K (IN-1K-b in the plots) and each level $L_{1/2/3/4/5}$. (a) Absolute top-1 accuracy on all levels. (b)-(e) accuracy relative to the baseline ResNet50 for all the models, split across the four model categories presented in Sec. 4.1 of the main paper.

We see that, for a given level L_i , WordNet combined with Lin similarity manages to gather concepts that are harder to discriminate and that the resulting classification performance is lower. This experiment, however, shows that it is possible to create a similar benchmark using automatically produced semantic similarity scores, the main alternative in the absence of any reliable hand-crafted ontology.

E. The ImageNet 2021 release

The ImageNet team recently released a new version of the full ImageNet dataset (IN-21K) as well as the ILSVRC-2012 dataset (IN-1K).⁵ with this release, both the datasets are now available for download directly from the official website.⁶

The 2021 Winter version of IN-21K. We built ImageNet-CoG on the 2011 Fall release of IN-21K, which was the only version available in 2020, when we started constructing our benchmark. The 2011 Fall version contained 21841 concepts, while the new release has only 19167 concepts—a subset of

the concepts from the Fall 2011 release. This follows recent studies from the ImageNet team, which identify potentially problematic concepts [66]. They were removed from the latest ImageNet version, including all the concepts under the “Person” sub-tree in WordNet.

With this modified version we successfully verified that: i) all the concepts of ImageNet-CoG are available in the new release, and ii) the images for all the 5000 concepts of ImageNet-CoG are identical in both releases. Consequently, all the results in our work *can also be reproduced* using the Winter 2021 version of IN-21K.

Blurred version of IN-1K. To protect the privacy of people present in some of the IN-1K images, the ImageNet team released a new version of this dataset, which we refer to as IN-1K-blurred [81]. In this version, the faces of people are blurred in the images. The statistics of these two versions are compared in Tab. C.

Although the models we evaluated in the paper were pretrained on IN-1K, with non-blurred images, for future reference, we performed our evaluation also on the blurred


⁵<https://image-net.org/update-mar-11-2021.php>

⁶<https://image-net.org/download-images.php>


version of IN-1K (IN-1K-blurred) for all the models. Concretely, for each model, we follow our evaluation protocol on IN-1K-blurred by extracting features of the blurred images and training logistic regression classifiers on them. We report these results in Fig. D. Note that Fig. D is the new version of Fig.3 in the main paper, with results obtained on IN-1K-blurred instead of IN-1K. We observe that the scores drop on average 0.91%, which is comparable to the 0.68% drop observed on popular models [81].

References


- [77] James S Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Proc. NeurIPS*, 2011. 3
- [78] Steven Bird, Ewan Klein, and Edward Loper. *Natural Language Processing with Python*. 2009. 2
- [79] Ikuya Yamada, Akari Asai, Jin Sakuma, Hiroyuki Shindo, Hideaki Takeda, Yoshiyasu Takefuji, and Yuji Matsumoto. Wikipedia2vec: An efficient toolkit for learning and visualizing the embeddings of words and entities from wikipedia. In *Proc. EMNLP*, 2020. 5
- [80] Ikuya Yamada, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. Joint learning of the embedding of words and entities for named entity disambiguation. In *Proc. CONLL*, 2016. 5
- [81] Kaiyu Yang, Jacqueline Yau, Li Fei-Fei, Jia Deng, and Olga Russakovsky. A study of face obfuscation in ImageNet. *arXiv preprint arXiv:2103.06191*, 2021. 5, 6, 7




(a) IN-1K [47]




(b) Lin L_1




(c) Lin L_2



(d) Lin L_3



(e) Lin L_4



(f) Lin L_5

Figure E: The number of images per concept for IN-1K [47] and each of the concept generalization levels obtained by Lin similarity. We end up with 1.17M, 1.17M, 1.15M, 1.16M, 1.14M images in total for levels $L_{1/2/3/4/5}$ respectively. Note that IN-1K has 1.33M images in total.

ResNet	Transformer	NAS & Other	Self-Supervision	Web data	Regularization	
<ul style="list-style-type: none"> ● ResNet50 (23.5M) ■ a-ResNet152 (58.1M) 	<ul style="list-style-type: none"> ▲ a-T2T-ViT-t-14 (21.1M) ▲ a-DeiT-S (21.7M) ▲ a-DeiT-S-distilled (21.7M) ▲ a-DeiT-B-distilled (86.1M) 	<ul style="list-style-type: none"> ★ a-Inception-v3 (25.1M) ▲ a-EfficientNet-B1 (6.5M) ▲ a-EfficientNet-B4 (17.5M) ▲ a-NAT-M4 (7.6M) ▲ a-VGG19 (139.6M) 	<ul style="list-style-type: none"> ■ s-DINO ▲ s-SwAV ▲ s-BarlowTwins ▲ s-OBOW ▲ s-BYOL 	<ul style="list-style-type: none"> ★ s-SimCLR-v2 ▲ s-MoCo-v2 ▲ s-MoChi ▲ s-CompReSS ▲ s-InfoMin 	<ul style="list-style-type: none"> ■ d-Semi-Sup ▲ d-Semi-Weakly-Sup ▲ d-MoPro ▲ d-CLIP 	<ul style="list-style-type: none"> ■ r-ReLabel ▲ r-CutMix ▲ r-MixUp ▲ r-Manifold-MixUp ▲ r-Adv-Robust ▲ r-MEAL-v2

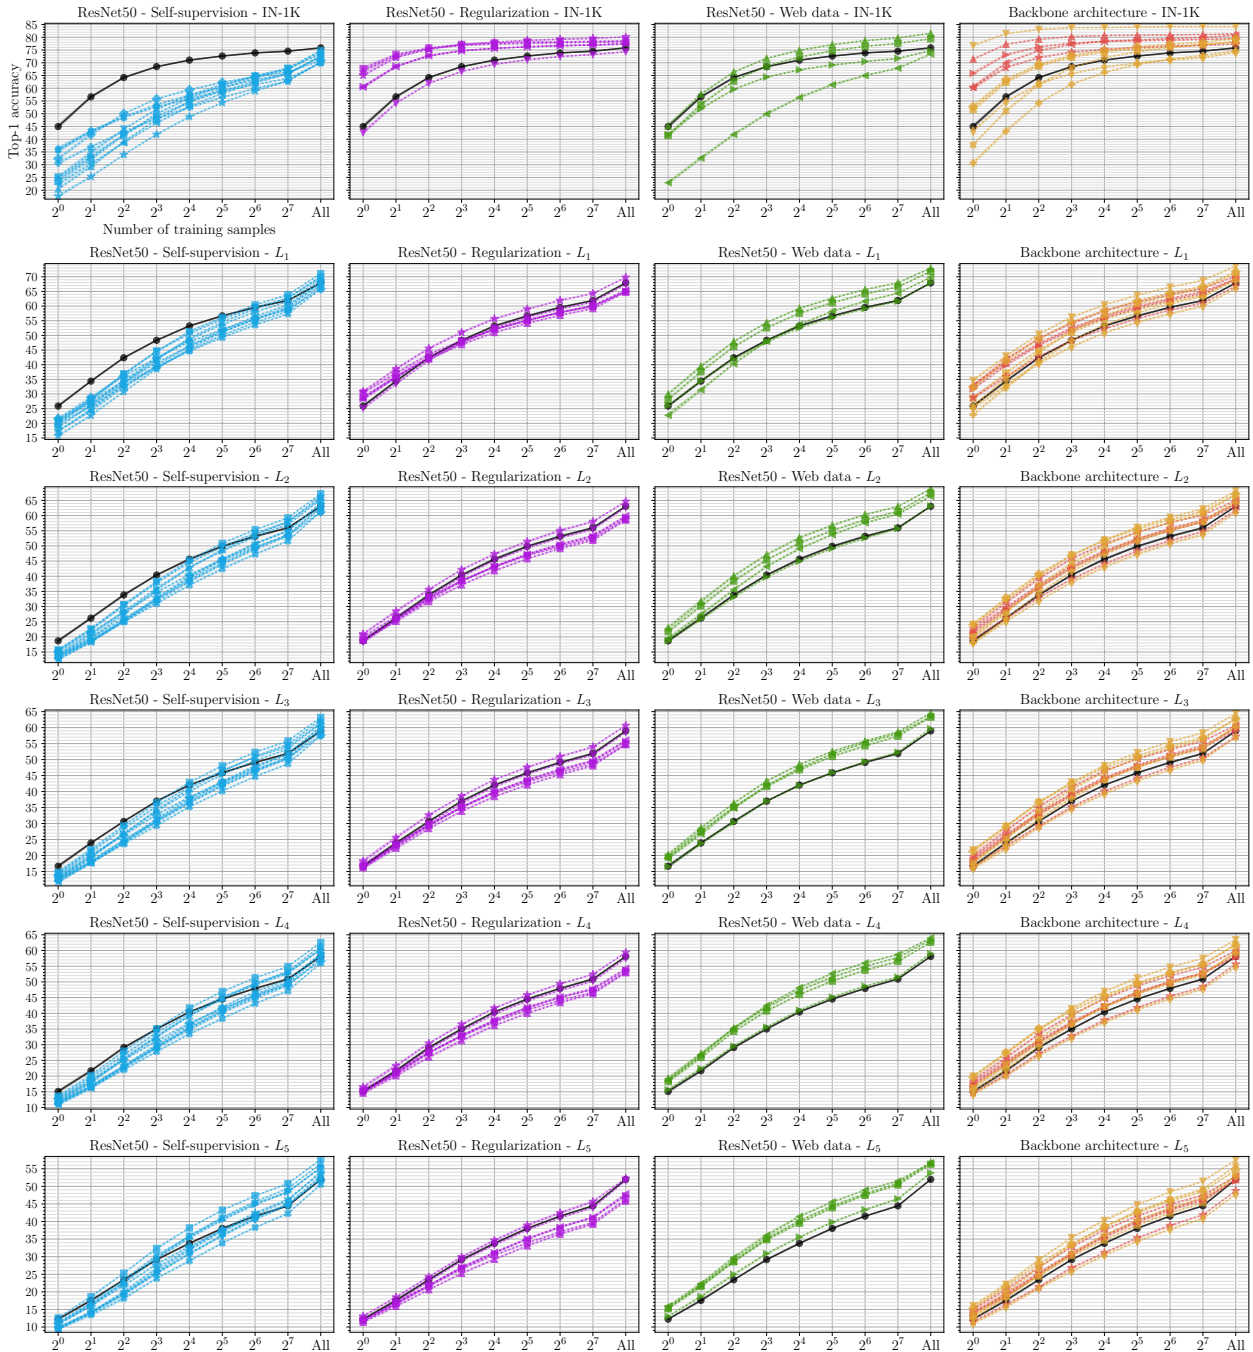


Figure F: **Few-shot linear classification on ImageNet-CoG.** Top-1 accuracy for each method using logistic regression classifiers. We train them on pre-extracted features for the concepts in IN-1K and our generalization levels ($L_{1/2/3/4/5}$), with a few training samples per concept, i.e., $N = \{1, 2, 4, 8, 16, 32, 64, 128\}$. “All”, the performance when all the samples are used, is also shown for reference.

ResNet	Transformer	NAS & Other	Self-Supervision	Web data	Regularization	
<ul style="list-style-type: none"> ● ResNet50 (23.5M) ■ a-ResNet152 (58.1M) 	<ul style="list-style-type: none"> ▲ a-T2T-ViT-t-14 (21.1M) ▶ a-DeiT-S (21.7M) ◆ a-DeiT-S-distilled (21.7M) ▼ a-DeiT-B-distilled (86.1M) 	<ul style="list-style-type: none"> ★ a-Inception-v3 (25.1M) ✦ a-EfficientNet-B1 (6.5M) ✦ a-EfficientNet-B4 (17.5M) ◆ a-NAT-M4 (7.6M) ◆ a-VGG19 (139.6M) 	<ul style="list-style-type: none"> ■ s-DINO ▲ s-SwAV ✦ s-BarlowTwins ◆ s-OBOW ▼ s-BYOL 	<ul style="list-style-type: none"> ★ s-SimCLR-v2 ✦ s-MoCo-v2 ✦ s-MoChi ◆ s-CompReSS ▼ s-InfoMin 	<ul style="list-style-type: none"> ■ d-Semi-Sup ▲ d-Semi-Weakly-Sup ✦ d-MoPro ▼ d-CLIP 	<ul style="list-style-type: none"> ■ r-ReLabel ▲ r-CutMix ✦ r-MixUp ▼ r-Manifold-MixUp ★ r-Adv-Robust ★ r-MEAL-v2

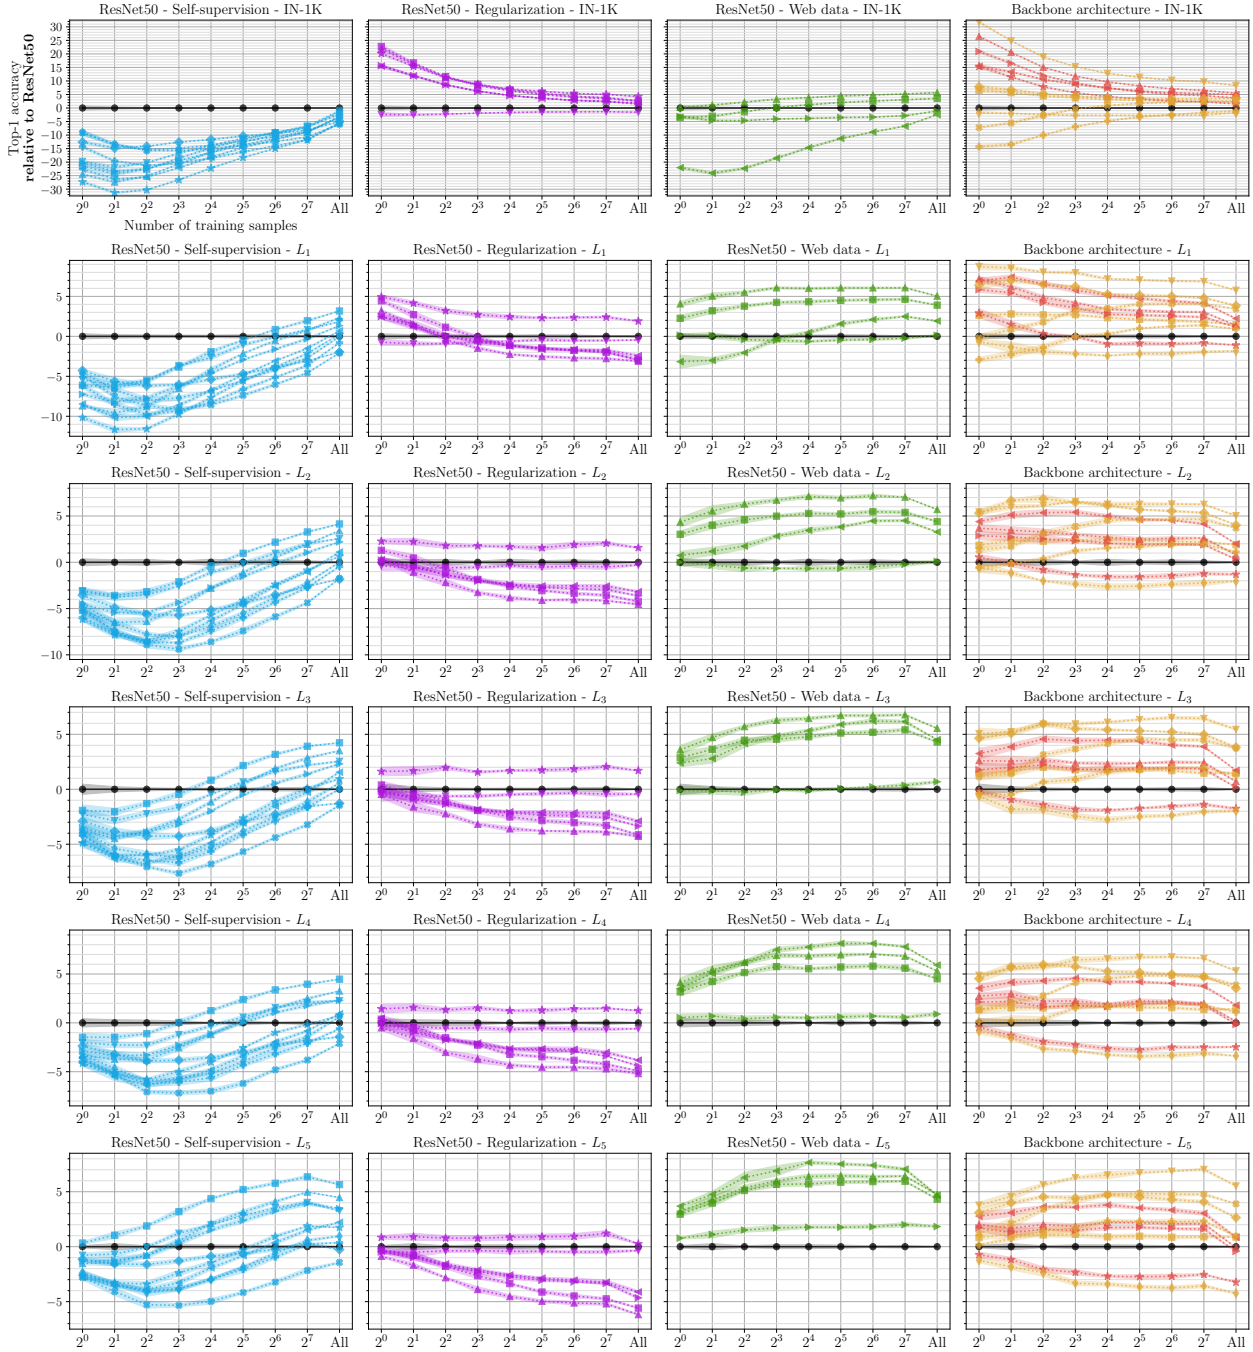


Figure G: **Relative few-shot linear classification on ImageNet-CoG.** The scores shown in Fig. F from a different perspective: all scores are relative to ResNet50.

