2021 ICCV OCTOBER 11-17 VIRTUAL

Inria

NAVER LABS Europe


M. Bulent Sariyildiz


Yannis Kalantidis


Diane Larlus


Karteek Alahari

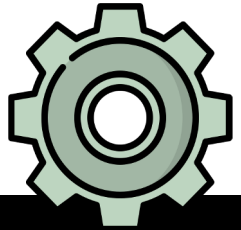

Concept generalization in visual representation learning

International Conference on Computer Vision (ICCV) 2021

Project page: https://europe.naverlabs.com/cog-benchmark

Learning general-purpose visual representations

*Train a model on **ImageNet-1K*** [1]

Model

*Transfer the learned representations to **downstream datasets*** [2,3]


...

DTD


CIFAR10/100


FGVC Aircraft


Places205


MS-COCO


Pascal VOC 2007


SUN397


Oxford-IIIT Pets


Birdsnap


Food-101


Stanford Cars

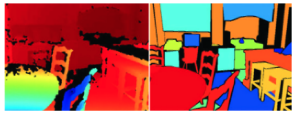
NYU-v2


Oxford 102 Flowers

...

[1] Russakovsky et al. "**ImageNet Large Scale Visual Recognition Challenge**", IJCV 2015
[2] Goyal et al. "**Scaling and Benchmarking Self-Supervised Visual Representation Learning**", ICCV 2019
[3] Kornblith et al. "**Do Better ImageNet Models Transfer Better?**", CVPR 2019

Concept generalization: *Transferring representations across concepts*

cats

*Train a model on **seen** concepts*

Model

*Transfer the learned representations to **unseen** concepts*

dogs

takins

cabbage worms

Sea squirts

Images are from Deng et al. "**Imagenet: A large-scale hierarchical image database.**" CVPR, 2009

Concept generalization: *Transferring representations across concepts*

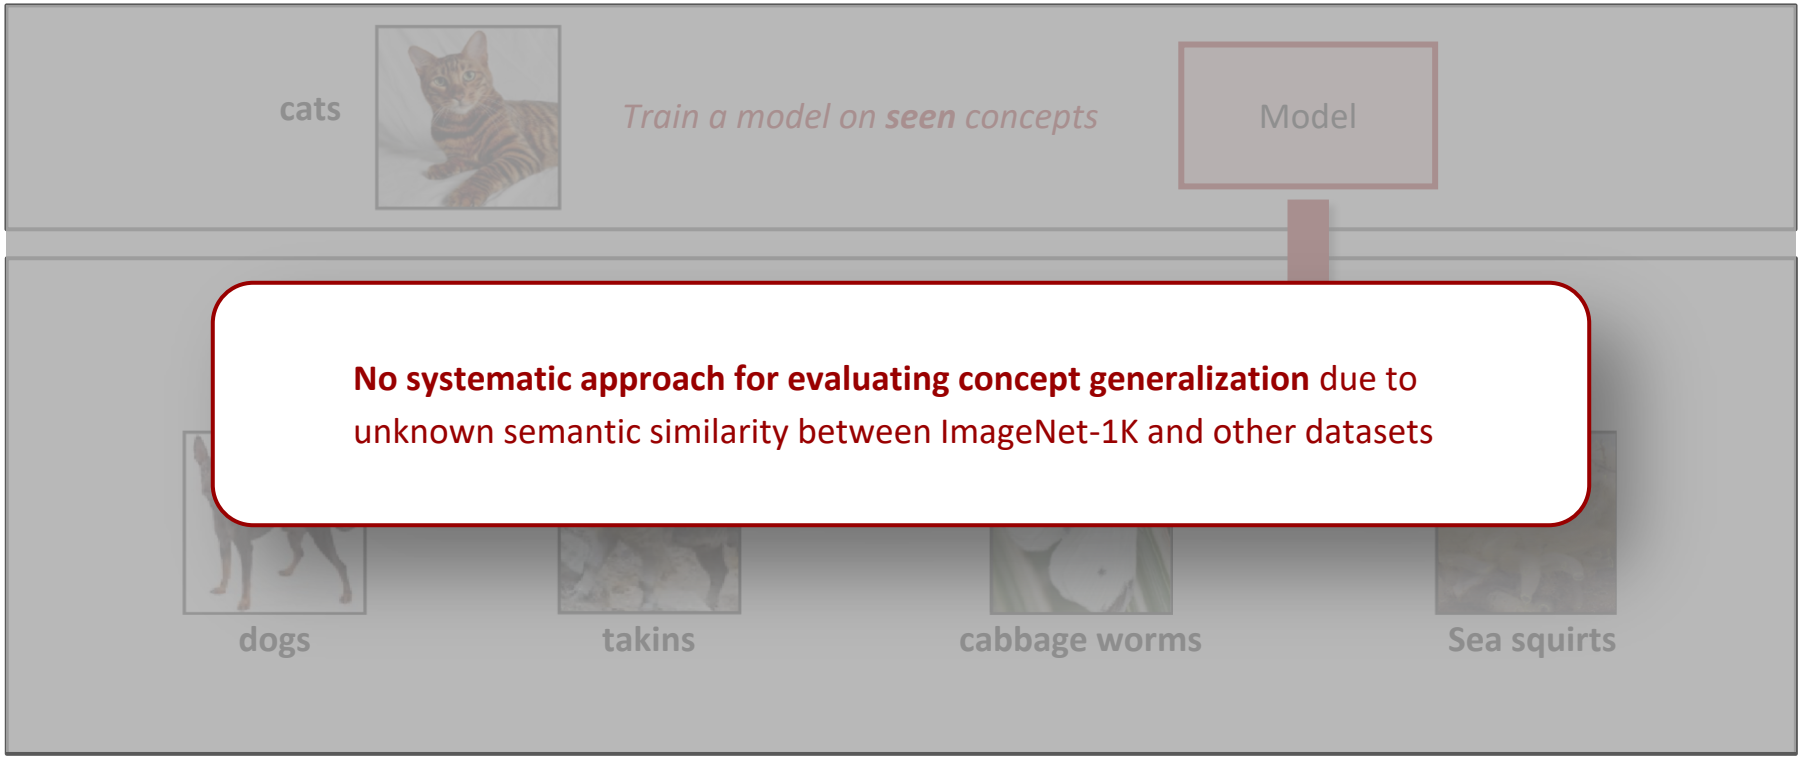

No systematic approach for evaluating concept generalization due to unknown semantic similarity between ImageNet-1K and other datasets

Images are from Deng et al. "**Imagenet: A large-scale hierarchical image database.**" CVPR, 2009

Our ImageNet-CoG Benchmark

A benchmark tailored for concept generalization, built on full ImageNet [2]

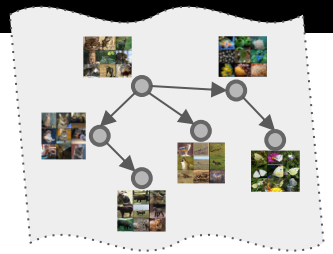

WordNet[1]

Seen and unseen concepts are in the same concept ontology[1] where semantic similarity between concepts is defined by linguists

[1] Miller. et al. "**WordNet: A Lexical Database for English**". ACM-Comm 1995
[2] Deng et al. "**Imagenet: A large-scale hierarchical image database.**" CVPR, 2009

Creating ImageNet-CoG

WordNet[1]

1) Measure the **semantic distance between concepts** using Lin similarity [2]

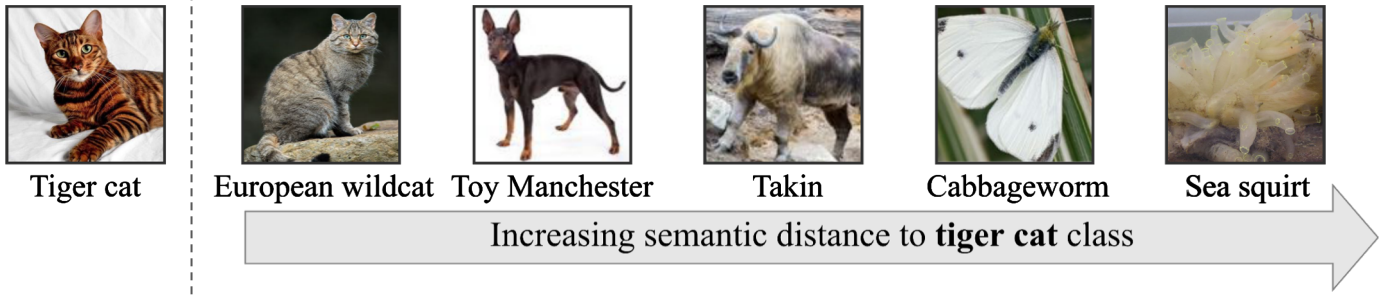

[1] Miller. et al. "**WordNet: A Lexical Database for English**". ACM-Comm 1995
[2] Lin et al. "**An information-theoretic definition of similarity.**" ICML, 1998
Images are from Deng et al. "**Imagenet: A large-scale hierarchical image database.**" CVPR, 2009

Creating ImageNet-CoG

1) Measure the **semantic distance between concepts** using Lin similarity
2) Generalize the semantic distance to **sets of concepts**, and create **levels**

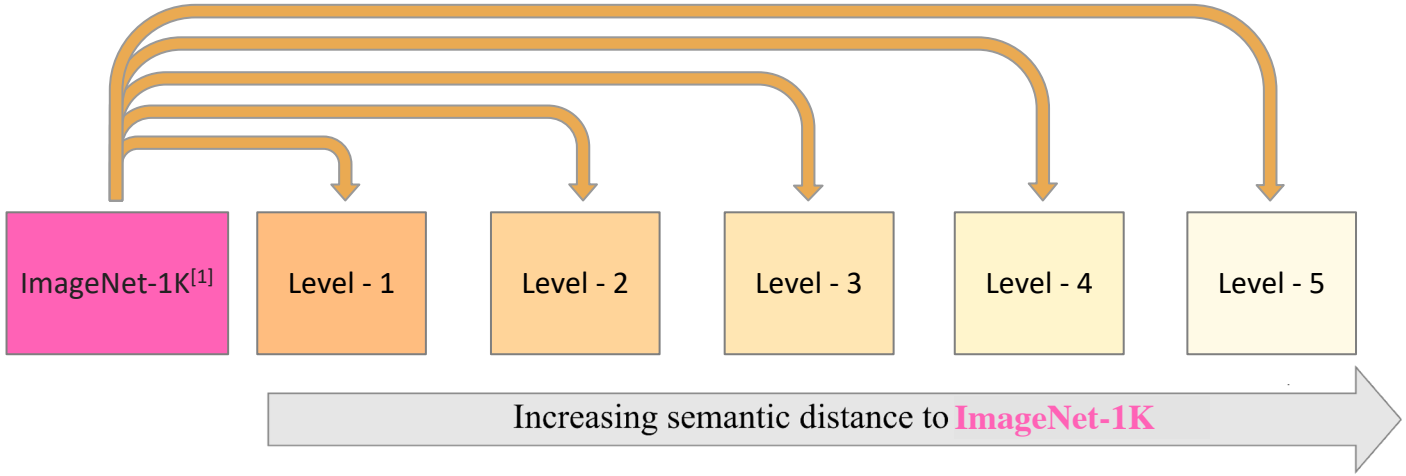

[1] Russakovsky et al. "**ImageNet Large Scale Visual Recognition Challenge**", IJCV 2015

Evaluation protocol for ImageNet-CoG

1) Extract image features for all the concepts
2) Train linear classifiers on each level separately

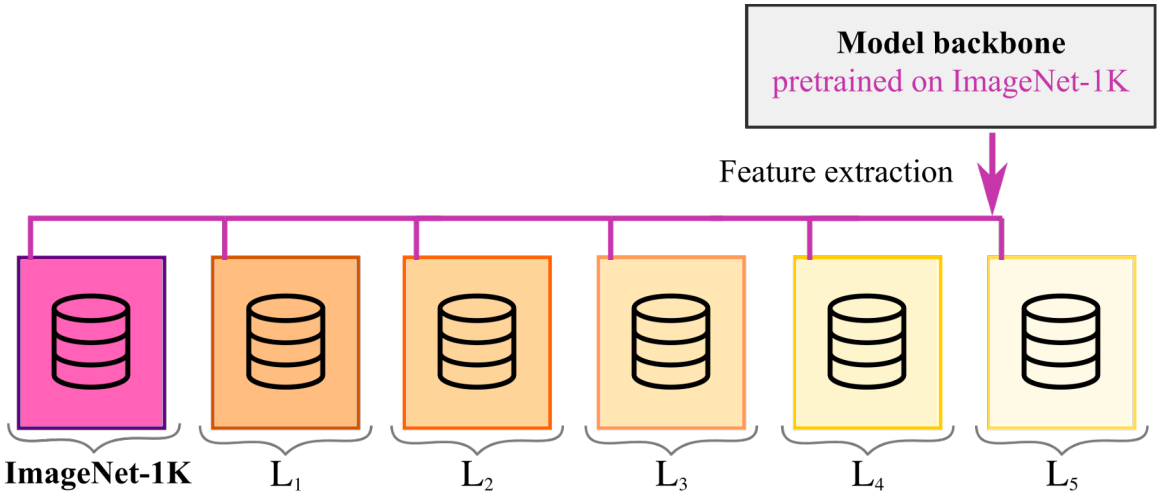

Evaluations with 31 representation learning methods

Baseline:
Supervised ResNet50 in torchvision

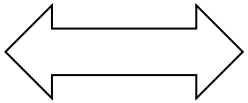

Compare against
4 groups of methods

Architecture:
10 CNN, Transformer or NAS models

Self-supervision:
10 self-supervised models

Regularization:
6 models trained with regularization techniques

Data:
4 models pre-trained with additional data

A subset of our experiments *(come to our poster for more!)*

Baseline:
Supervised ResNet50 in torchvision

2) Self-supervised models generalize well

3) Label-based regularization models might be overfitting

Baseline vs.
Self-supervised models

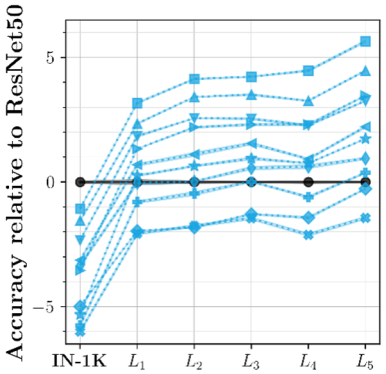

Baseline vs.
Regularization techniques

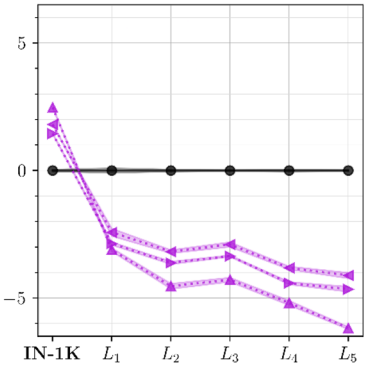

1) Performance of **all 31 models**
monotonically decreases from ImageNet-1K to Level-5

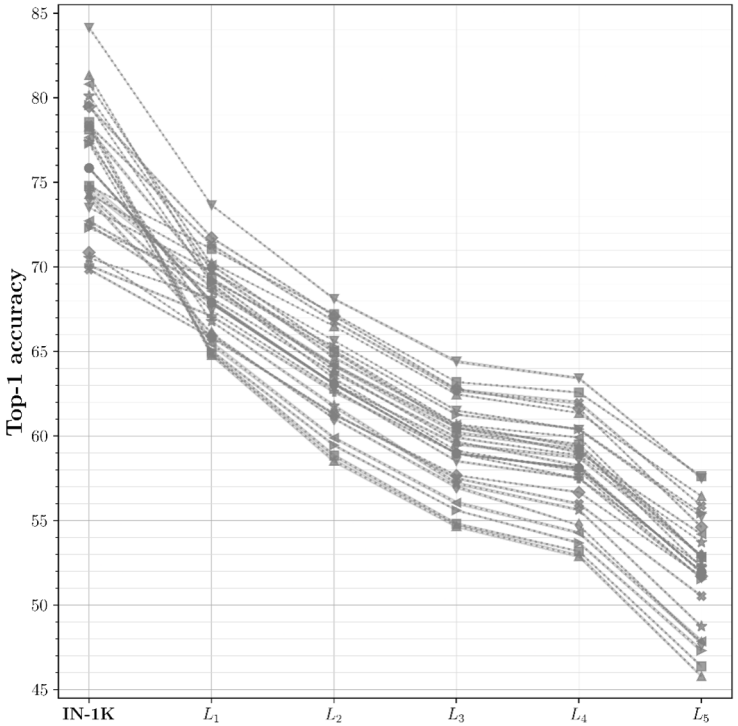

Conclusion

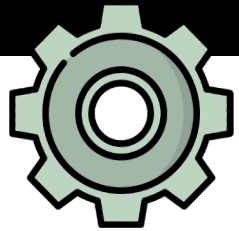

- ➔ **We proposed the ImageNet-CoG benchmark**
  - ◆ Enables measuring concept generalization in a controlled way
  - ◆ **Seen** concepts ⇒ ImageNet-1K concepts
  - ◆ **Unseen** concepts ⇒ Sampled from the full ImageNet-21K dataset
    - ● **5 Levels** ⇒ Increasingly challenging transfer datasets
- ➔ **Try your own method**
  - ◆ Can be used out-of-the-box for ImageNet-1K pretrained models
- ➔ **31 models already evaluated** on ImageNet-CoG
  - ◆ Interesting insights on popular state-of-the-art methods

Please come to our poster to discuss
Session-8 *→ Oct. 13 Wednesday 9:00-10:00 AM EDT*
Oct. 15 Friday 4:00-6:00 PM EDT

Project page: https://europe.naverlabs.com/cog-benchmark