

Data Imbalance and Fine-Grained Classification

6.S954 Computer Vision and Planetary Health

birds

PASCAL



bird

cats




cat

dogs



dog


ILSVRC




flamingo



cock



ruffed grouse




quail




partridge

...



Egyptian cat



Persian cat



Siamese cat




tabby



lynx

...



dalmatian




keeshond



miniature schnauzer




standard schnauzer



giant schnauzer

...





Two-spotted ladybug
Adalia bipunctata



Seven-spotted ladybug
Coccinella septempunctata

Figure 1. Two visually similar species from the iNat2017 dataset. Through close inspection, we can see that the ladybug on the left has *two* spots while the one on the right has *seven*.



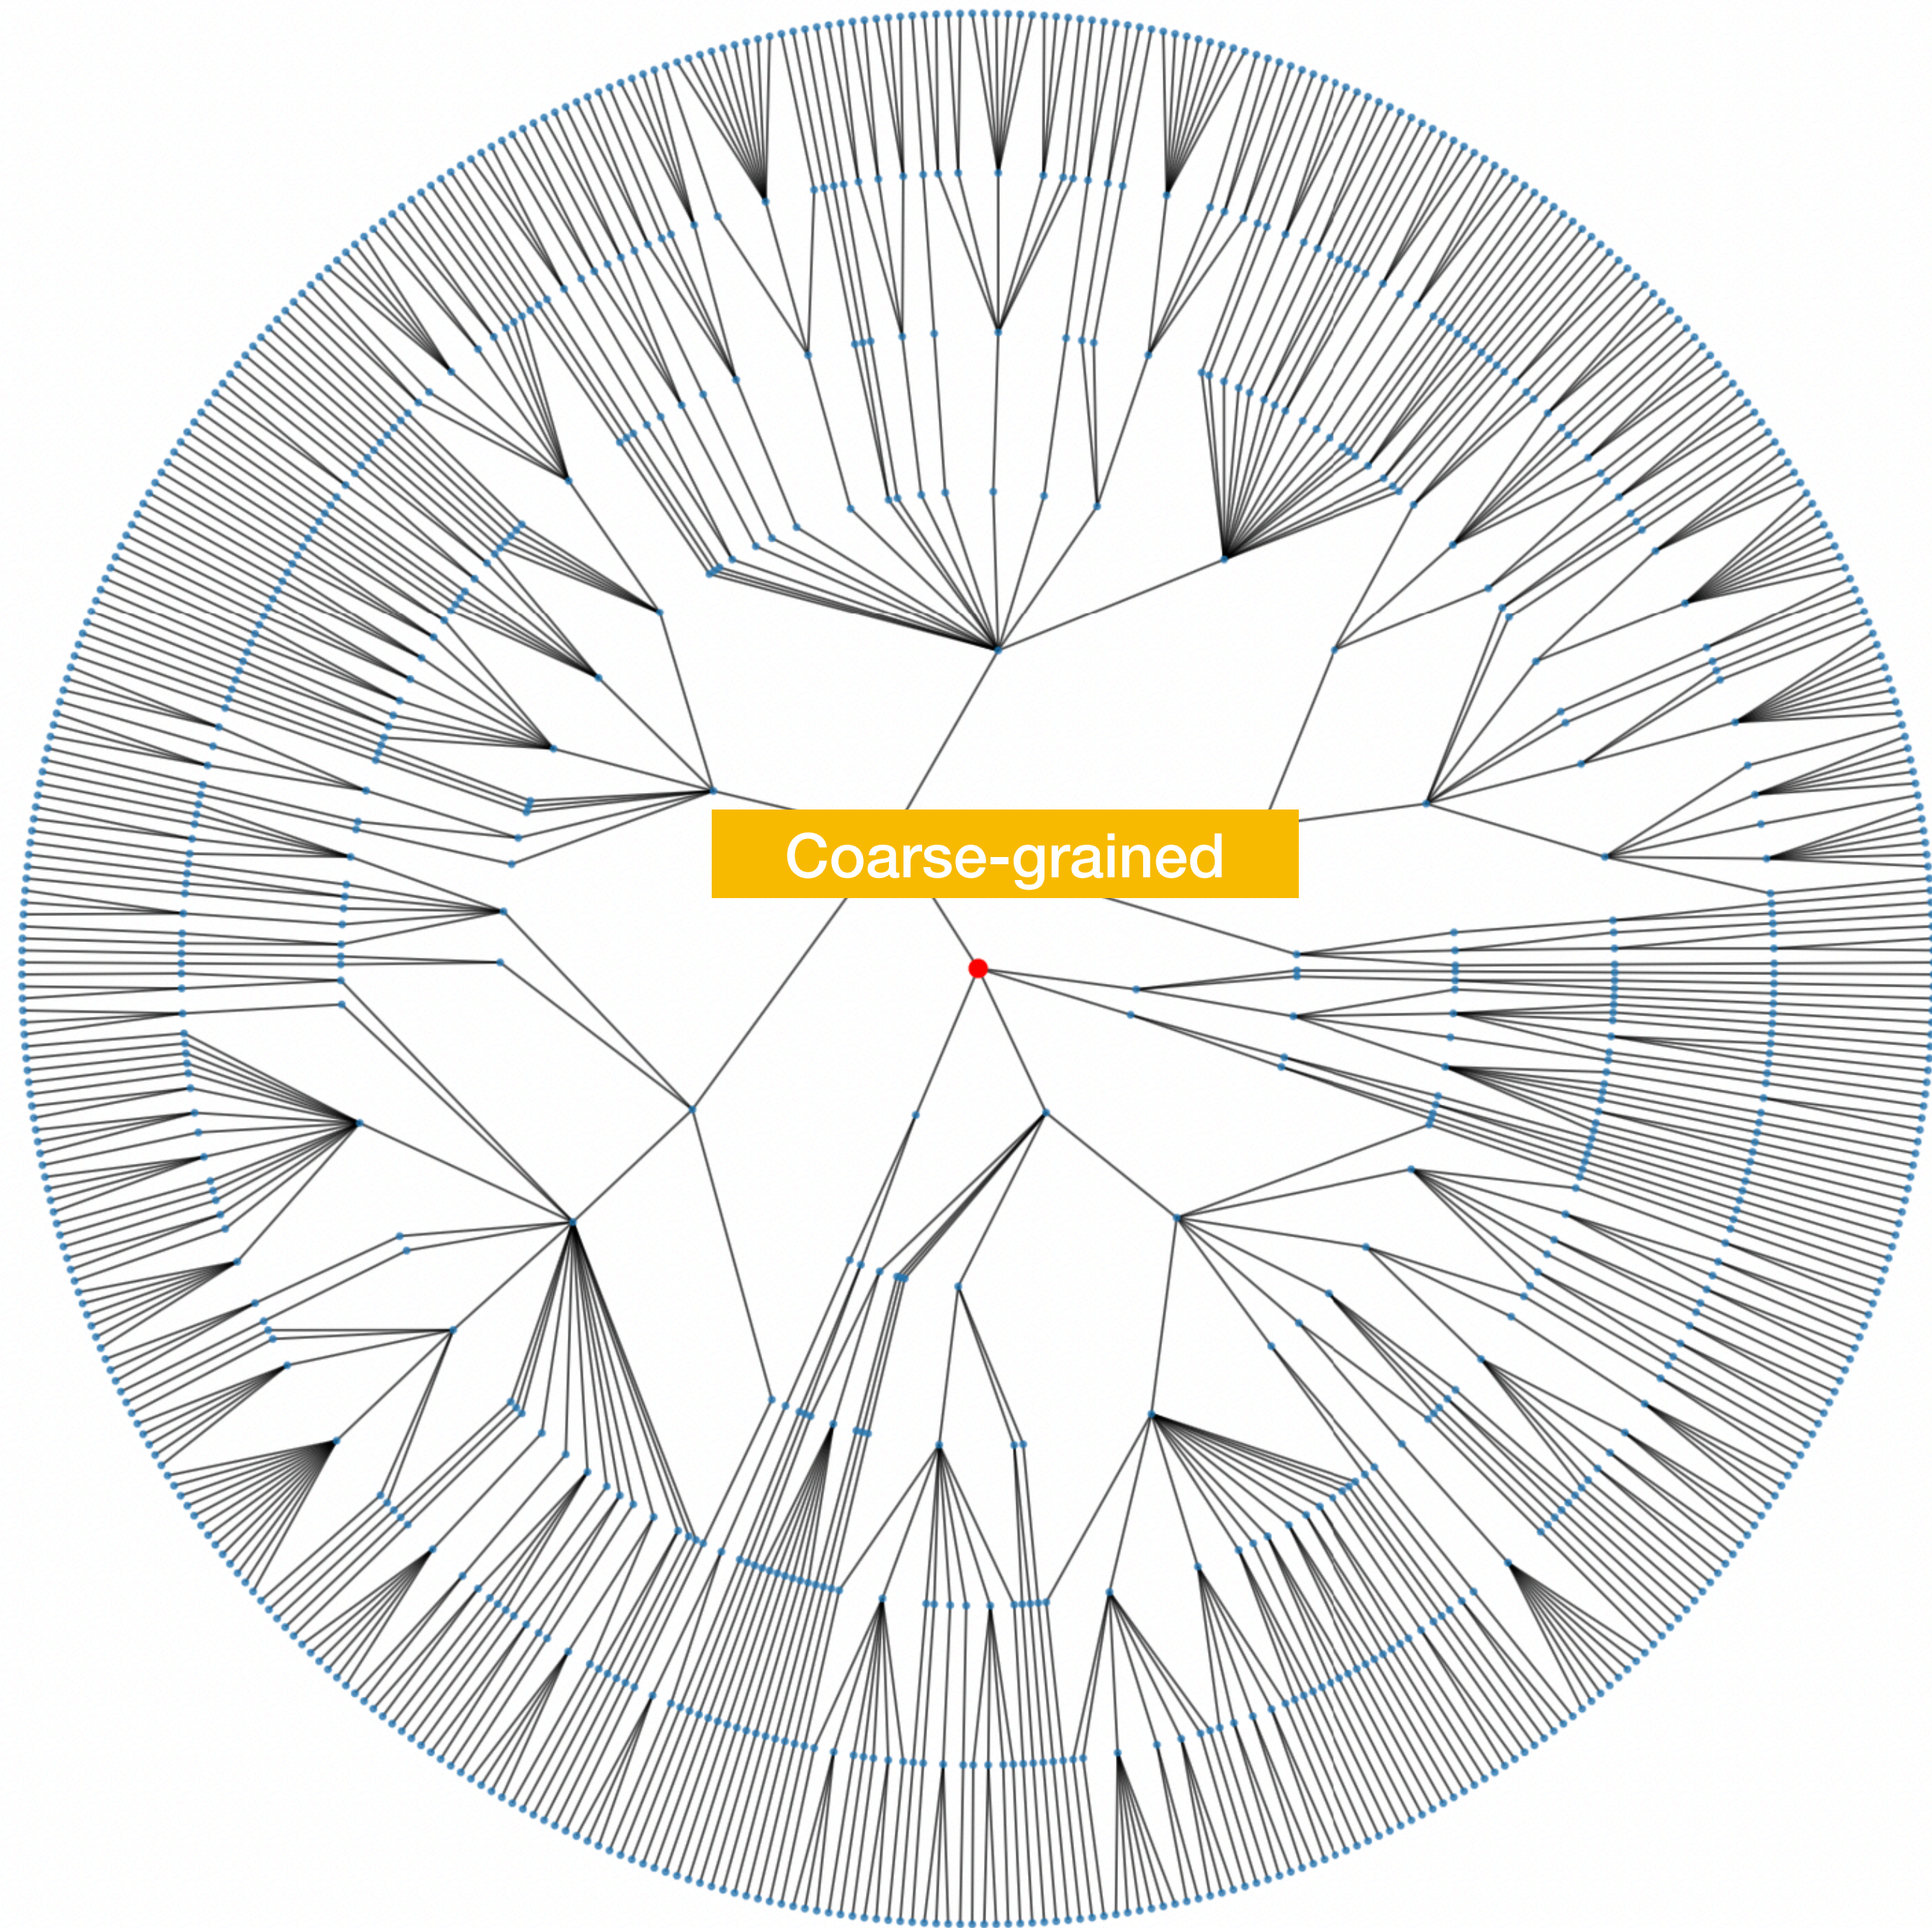
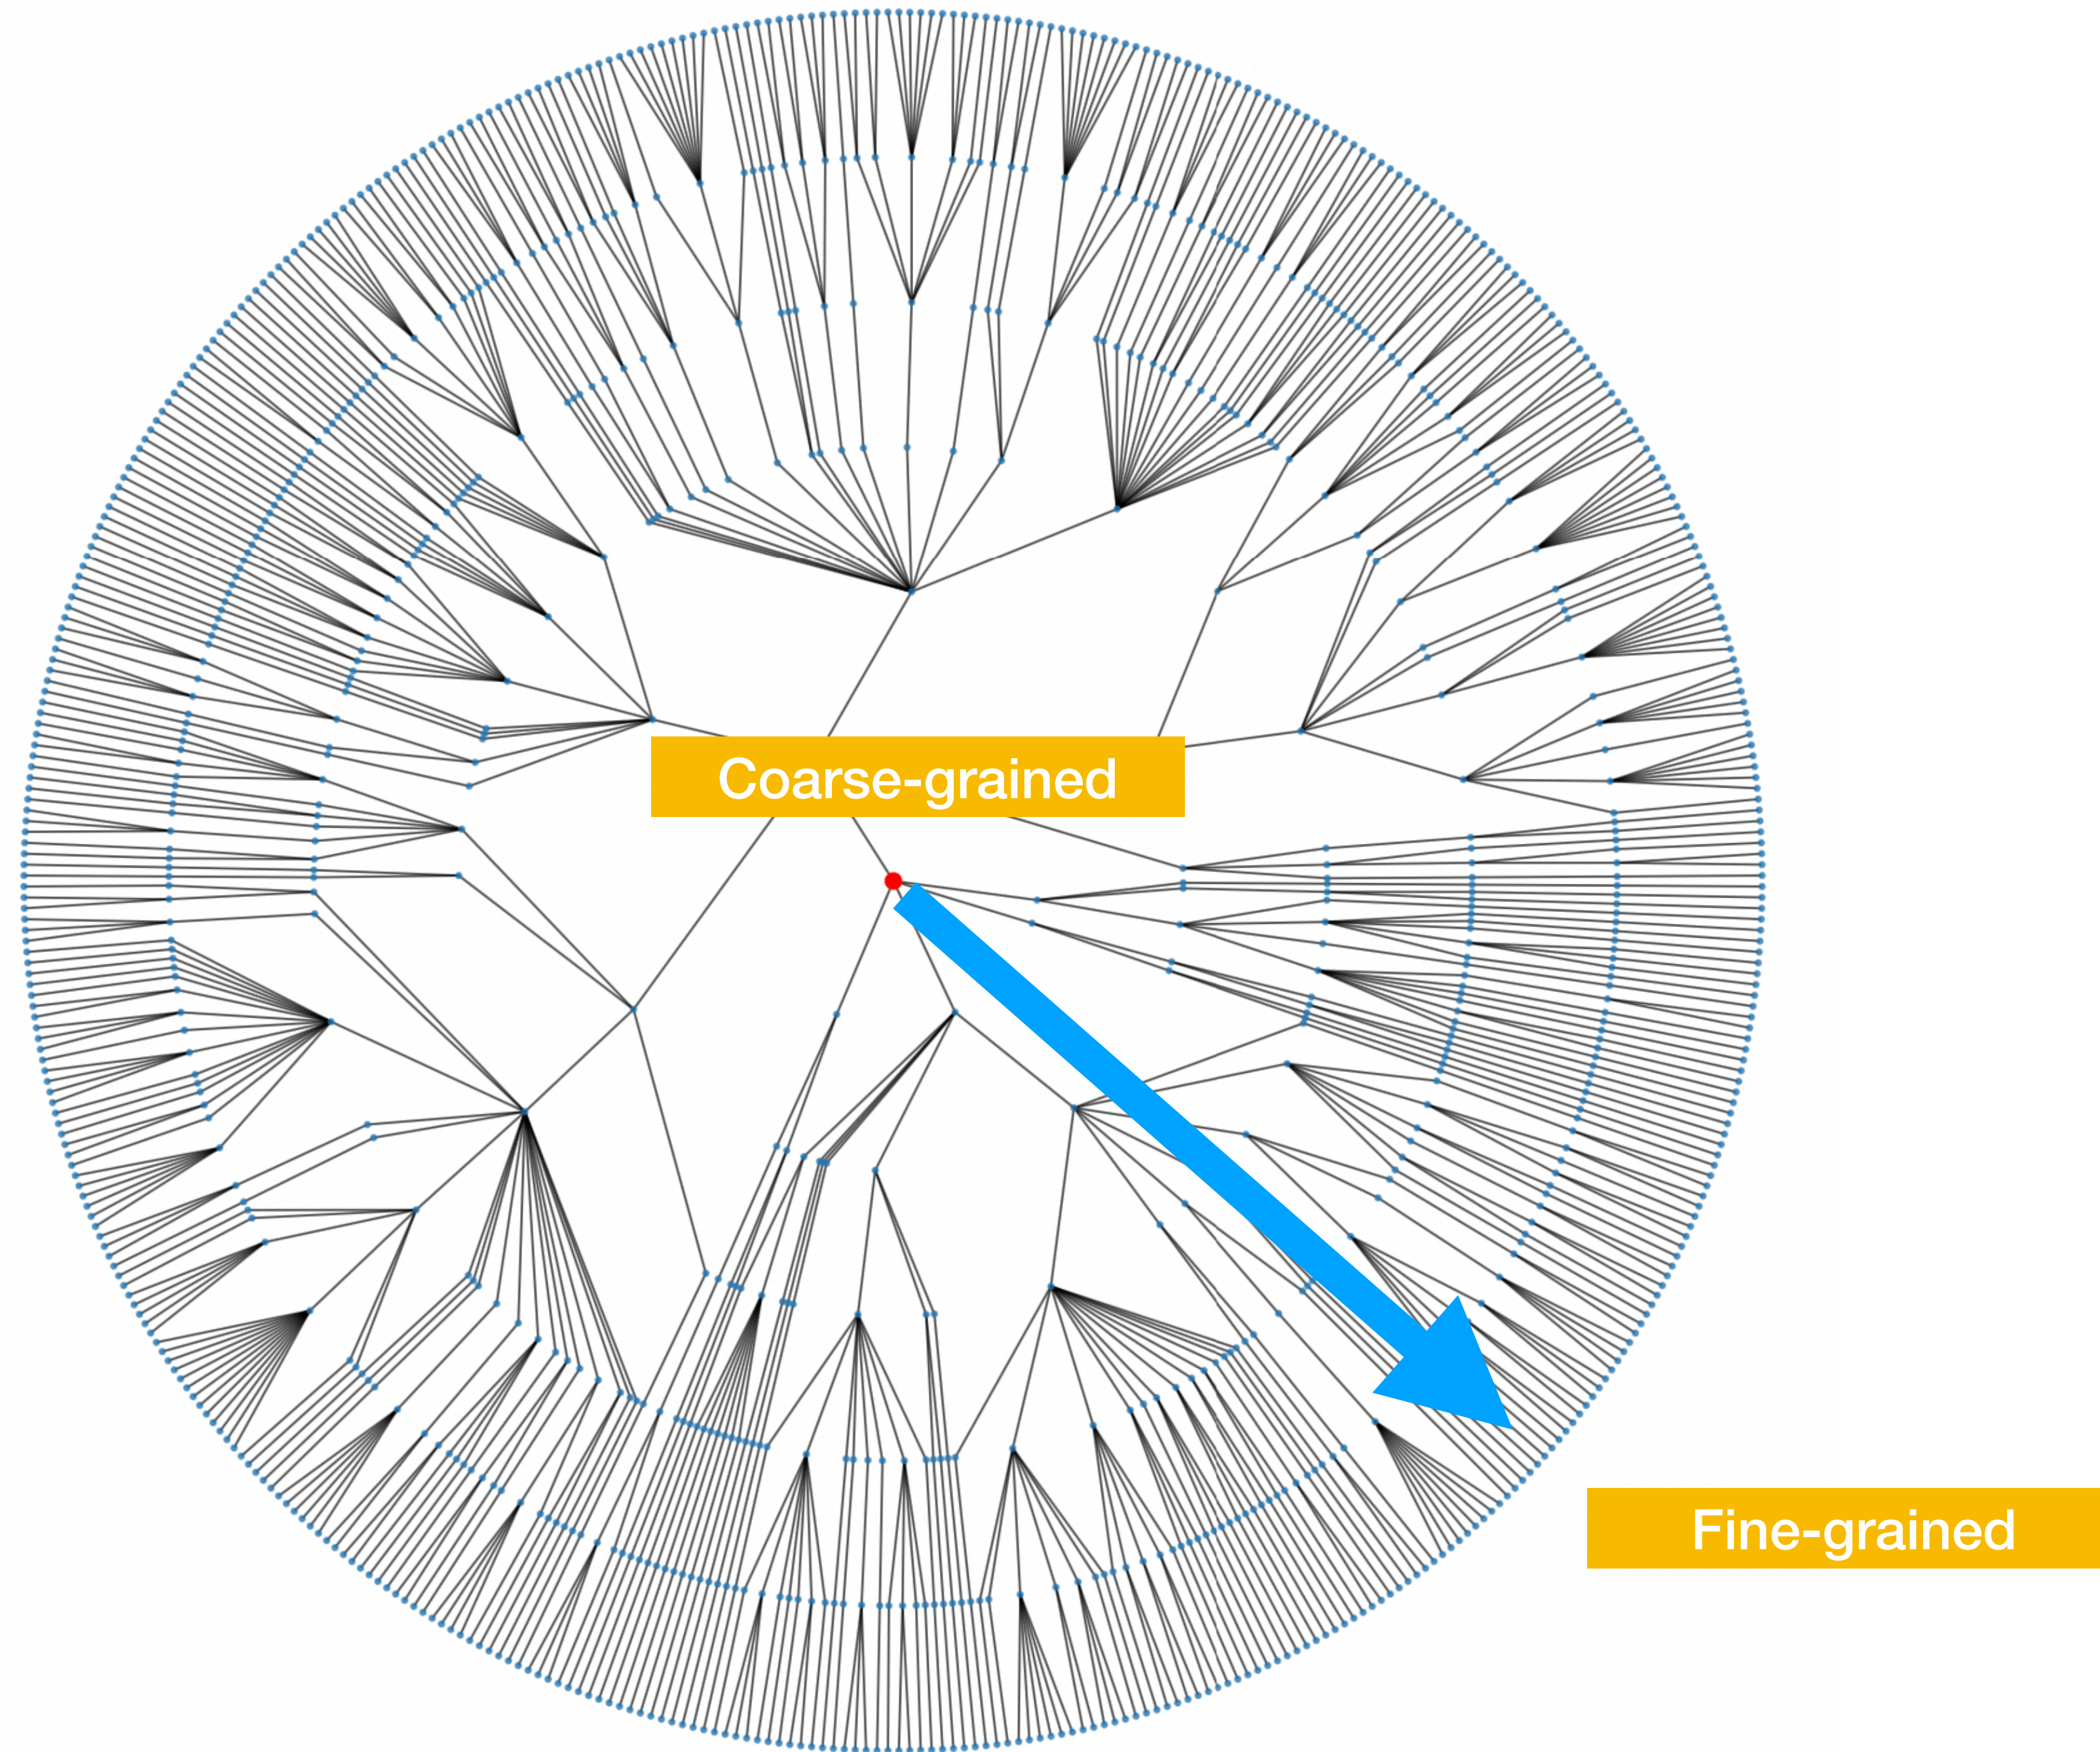
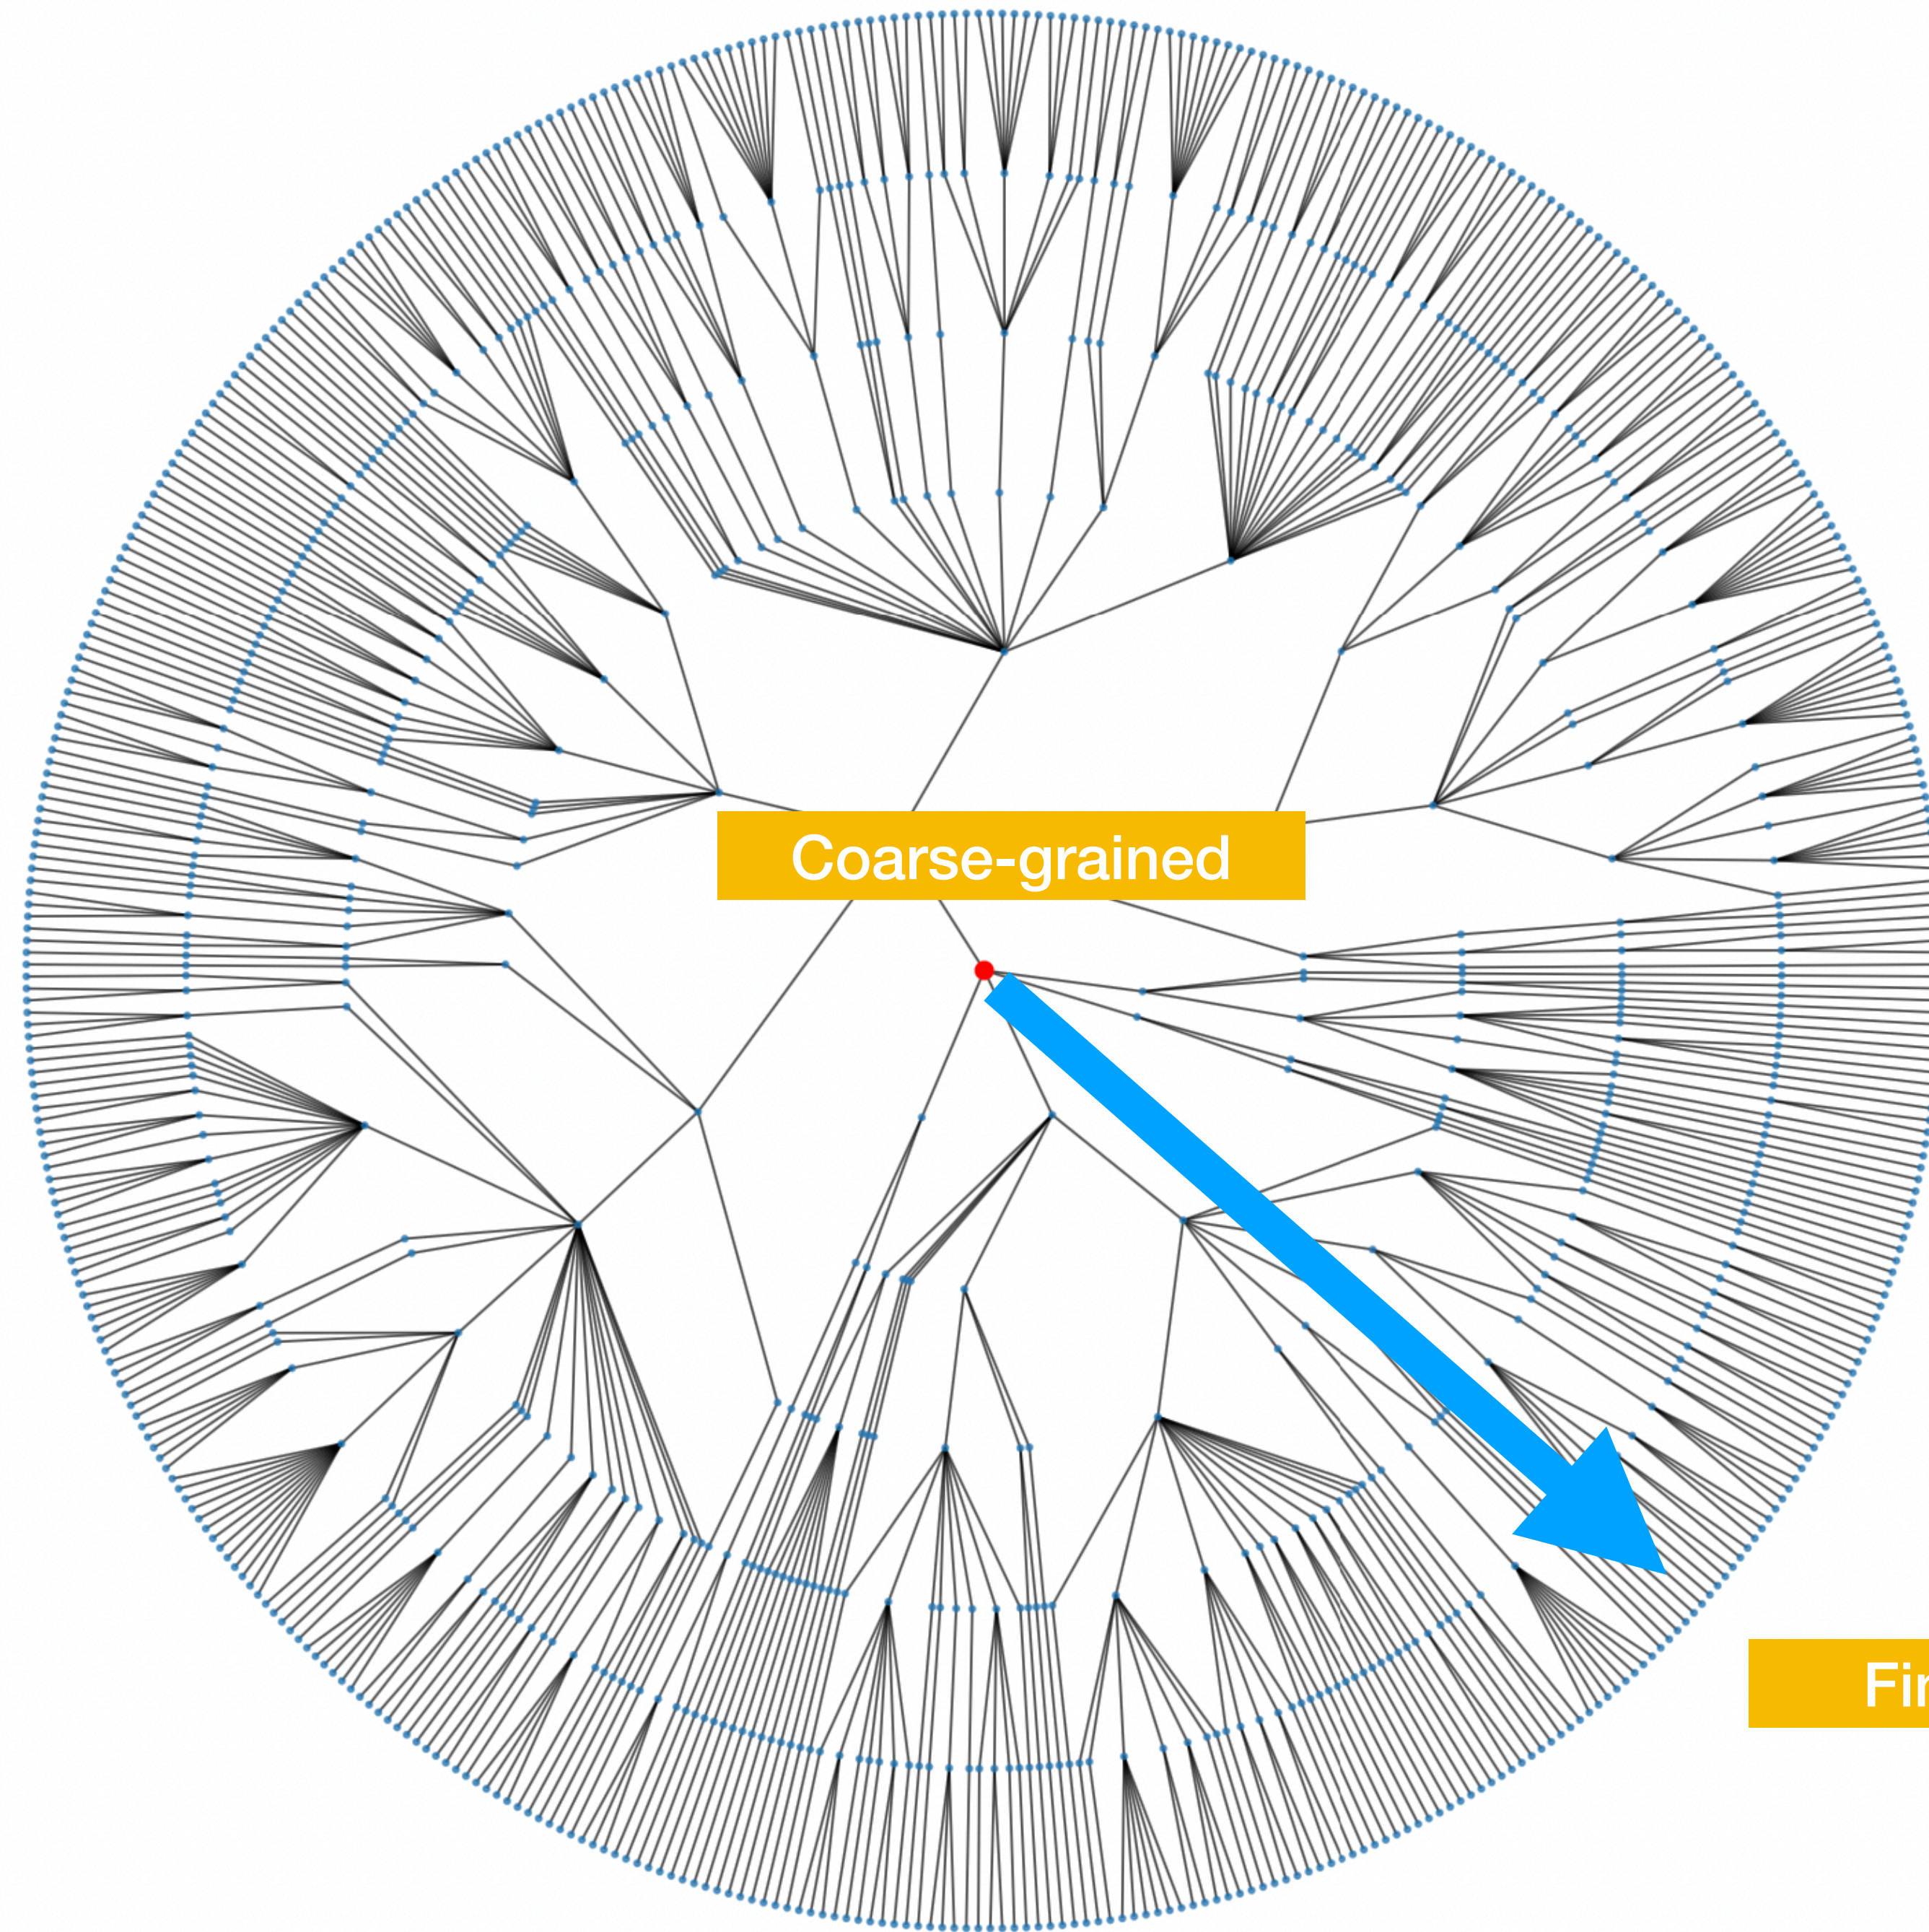


Figure 2. Distribution of training images per species. iNat2017 contains a large imbalance between classes, where the top 1% most populated classes contain over 16% of training images.







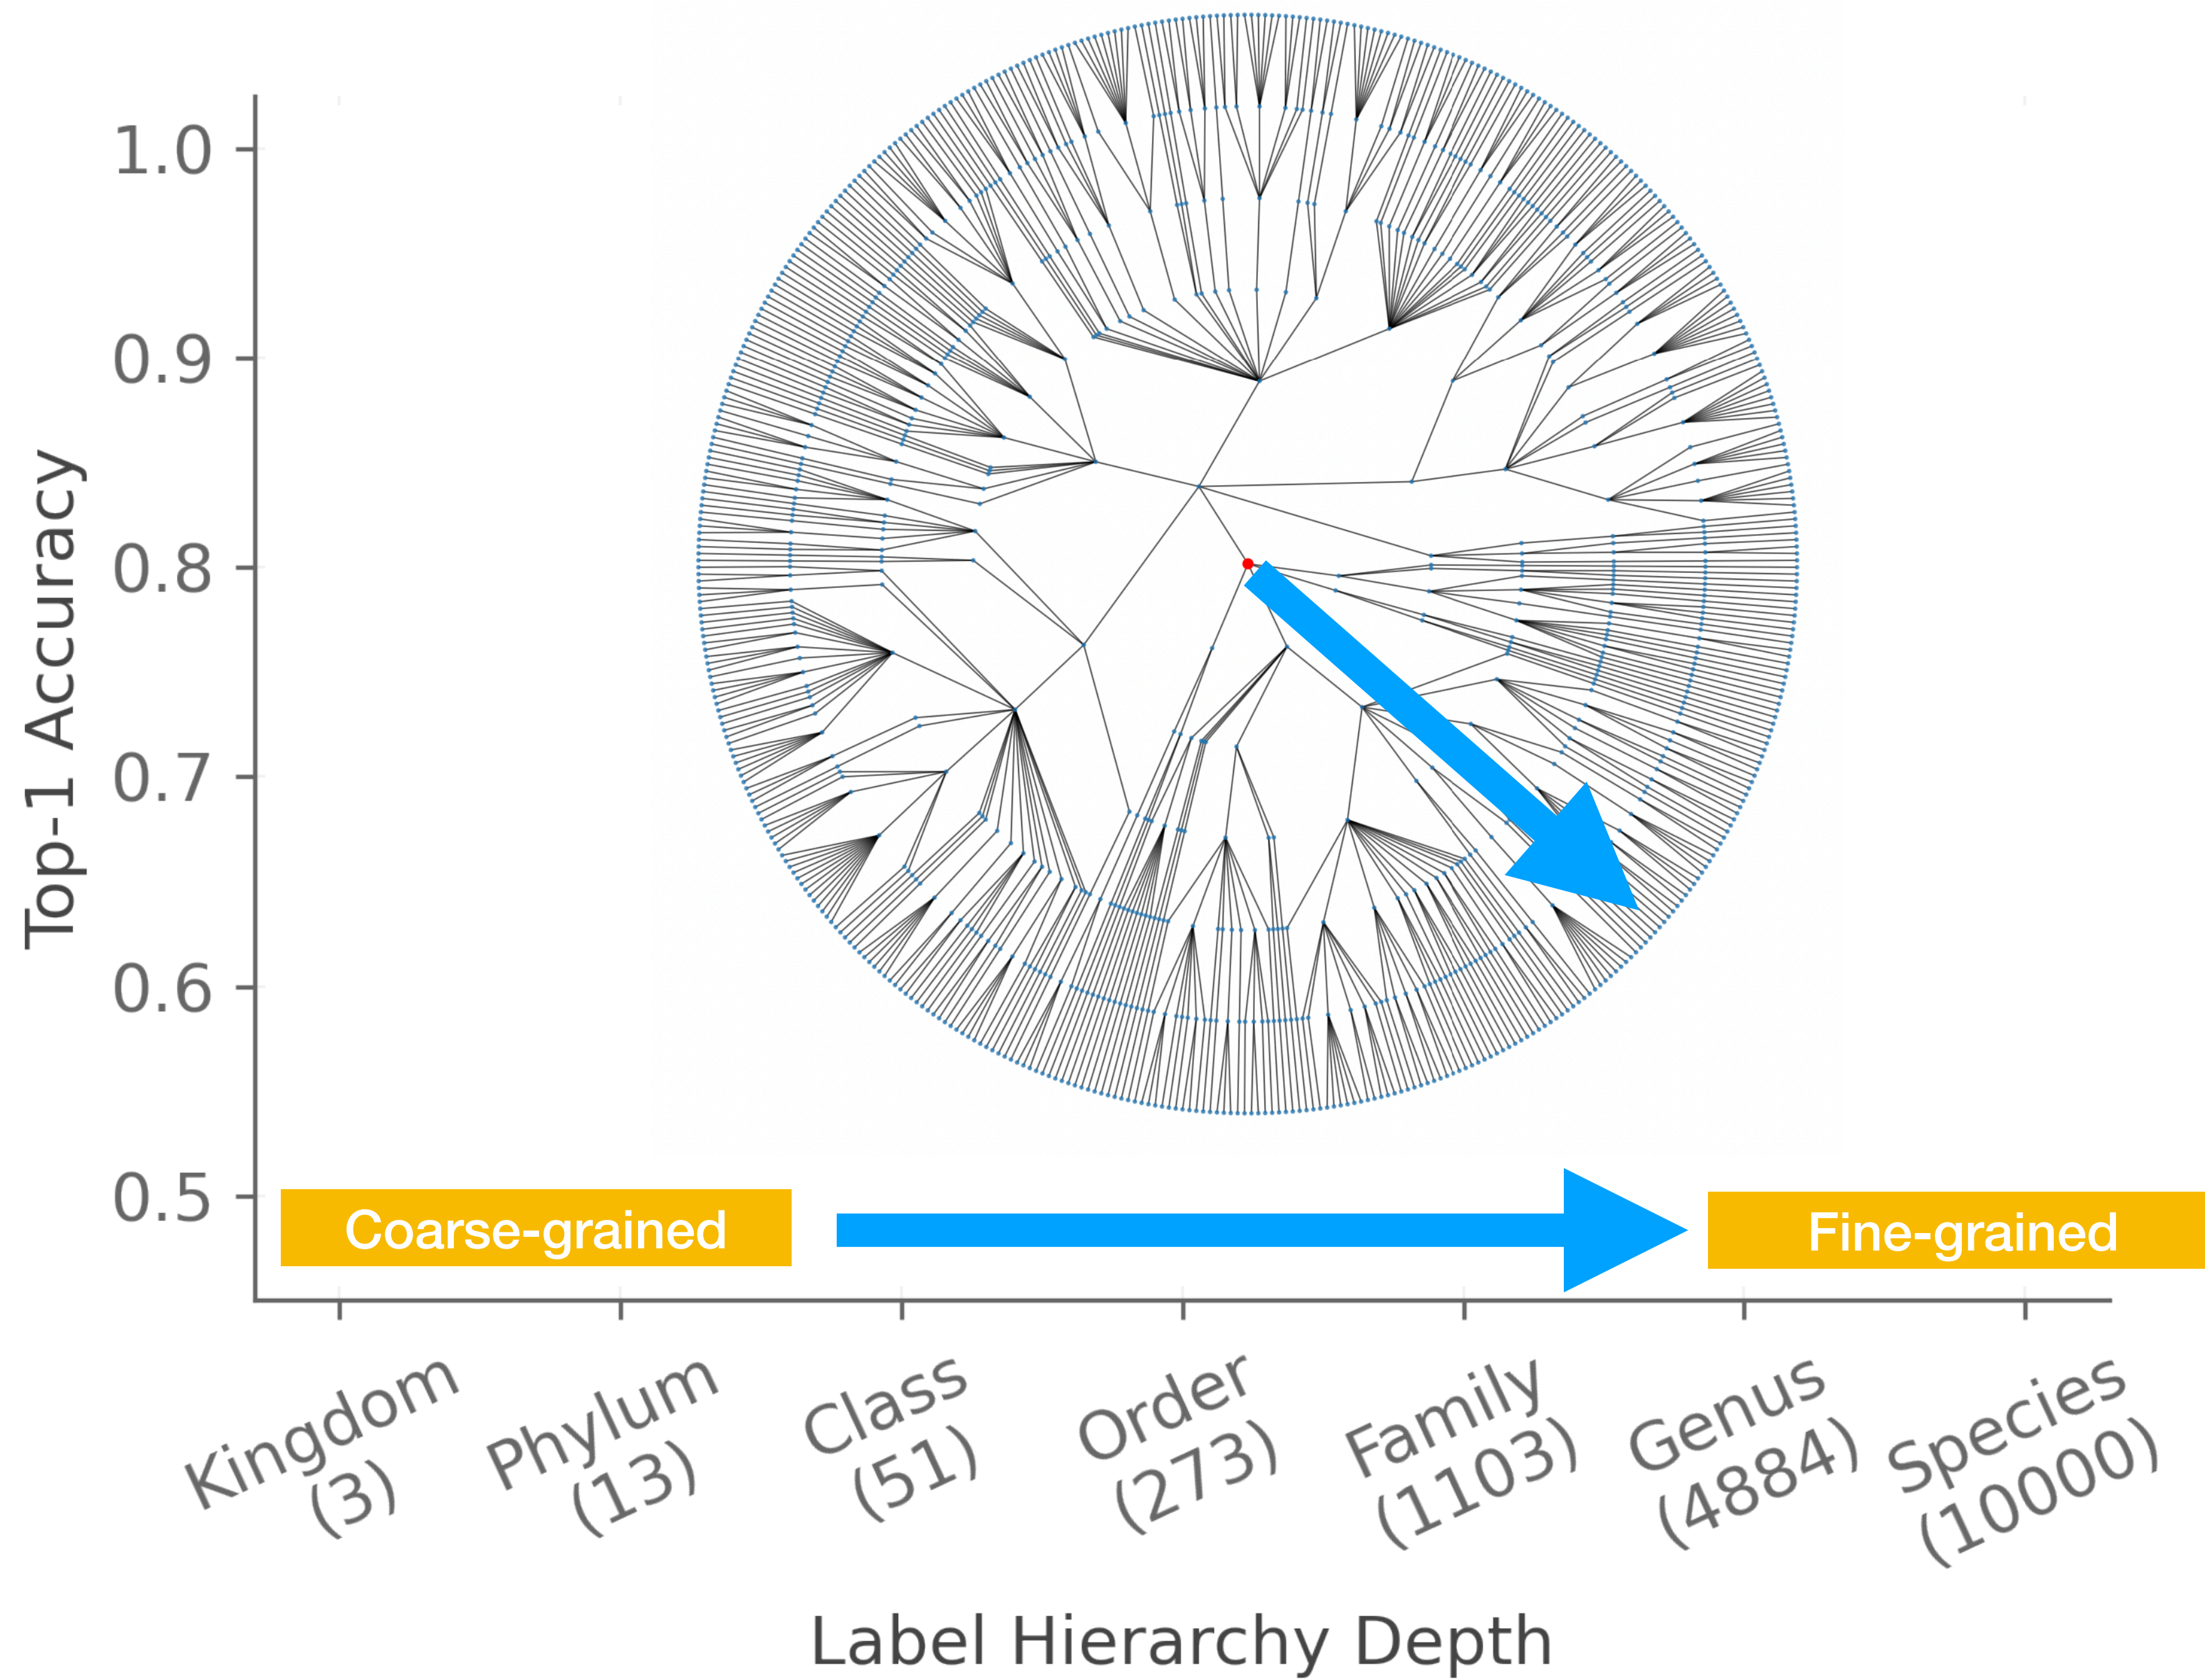


S. umblicata




S. ornata








Two-spotted ladybug
Adalia bipunctata




Seven-spotted ladybug
Coccinella septempunctata


Figure 1. Two visually similar species from the iNat2017 dataset. Through close inspection, we can see that the ladybug on the left has *two* spots while the one on the right has *seven*.



Granularity spectrum




(a)



(b)




(c)



LoveDA-rural

- Agricultural
- Forest
- Barren
- Water
- Road
- Building
- Background



Other sources of imbalance

Not just class distribution!



Figure 2. **Common data challenges:** (1) **Illumination:** Animals are not always salient. (2) **Motion blur:** common with poor illumination at night. (3) **Size of the region of interest (ROI):** Animals can be small or far from the camera. (4) **Occlusion:** e.g. by bushes or rocks. (5) **Camouflage:** decreases saliency in animals' natural habitat. (6) **Perspective:** Animals can be close to the camera, resulting in partial views of the body.



Male and Female Blue Grosbeak (*Passerina caerulea*)
 Image credits: Ed Schneider/Shutterstock.com (left), Steve Byland/Shutterstock.com (right)






Figure 3. Number of annotations for each **location**, over the two classes. The distribution images per location is long-tailed, and each location has a different and peculiar class distribution.

Why are these hard problems in computer vision?

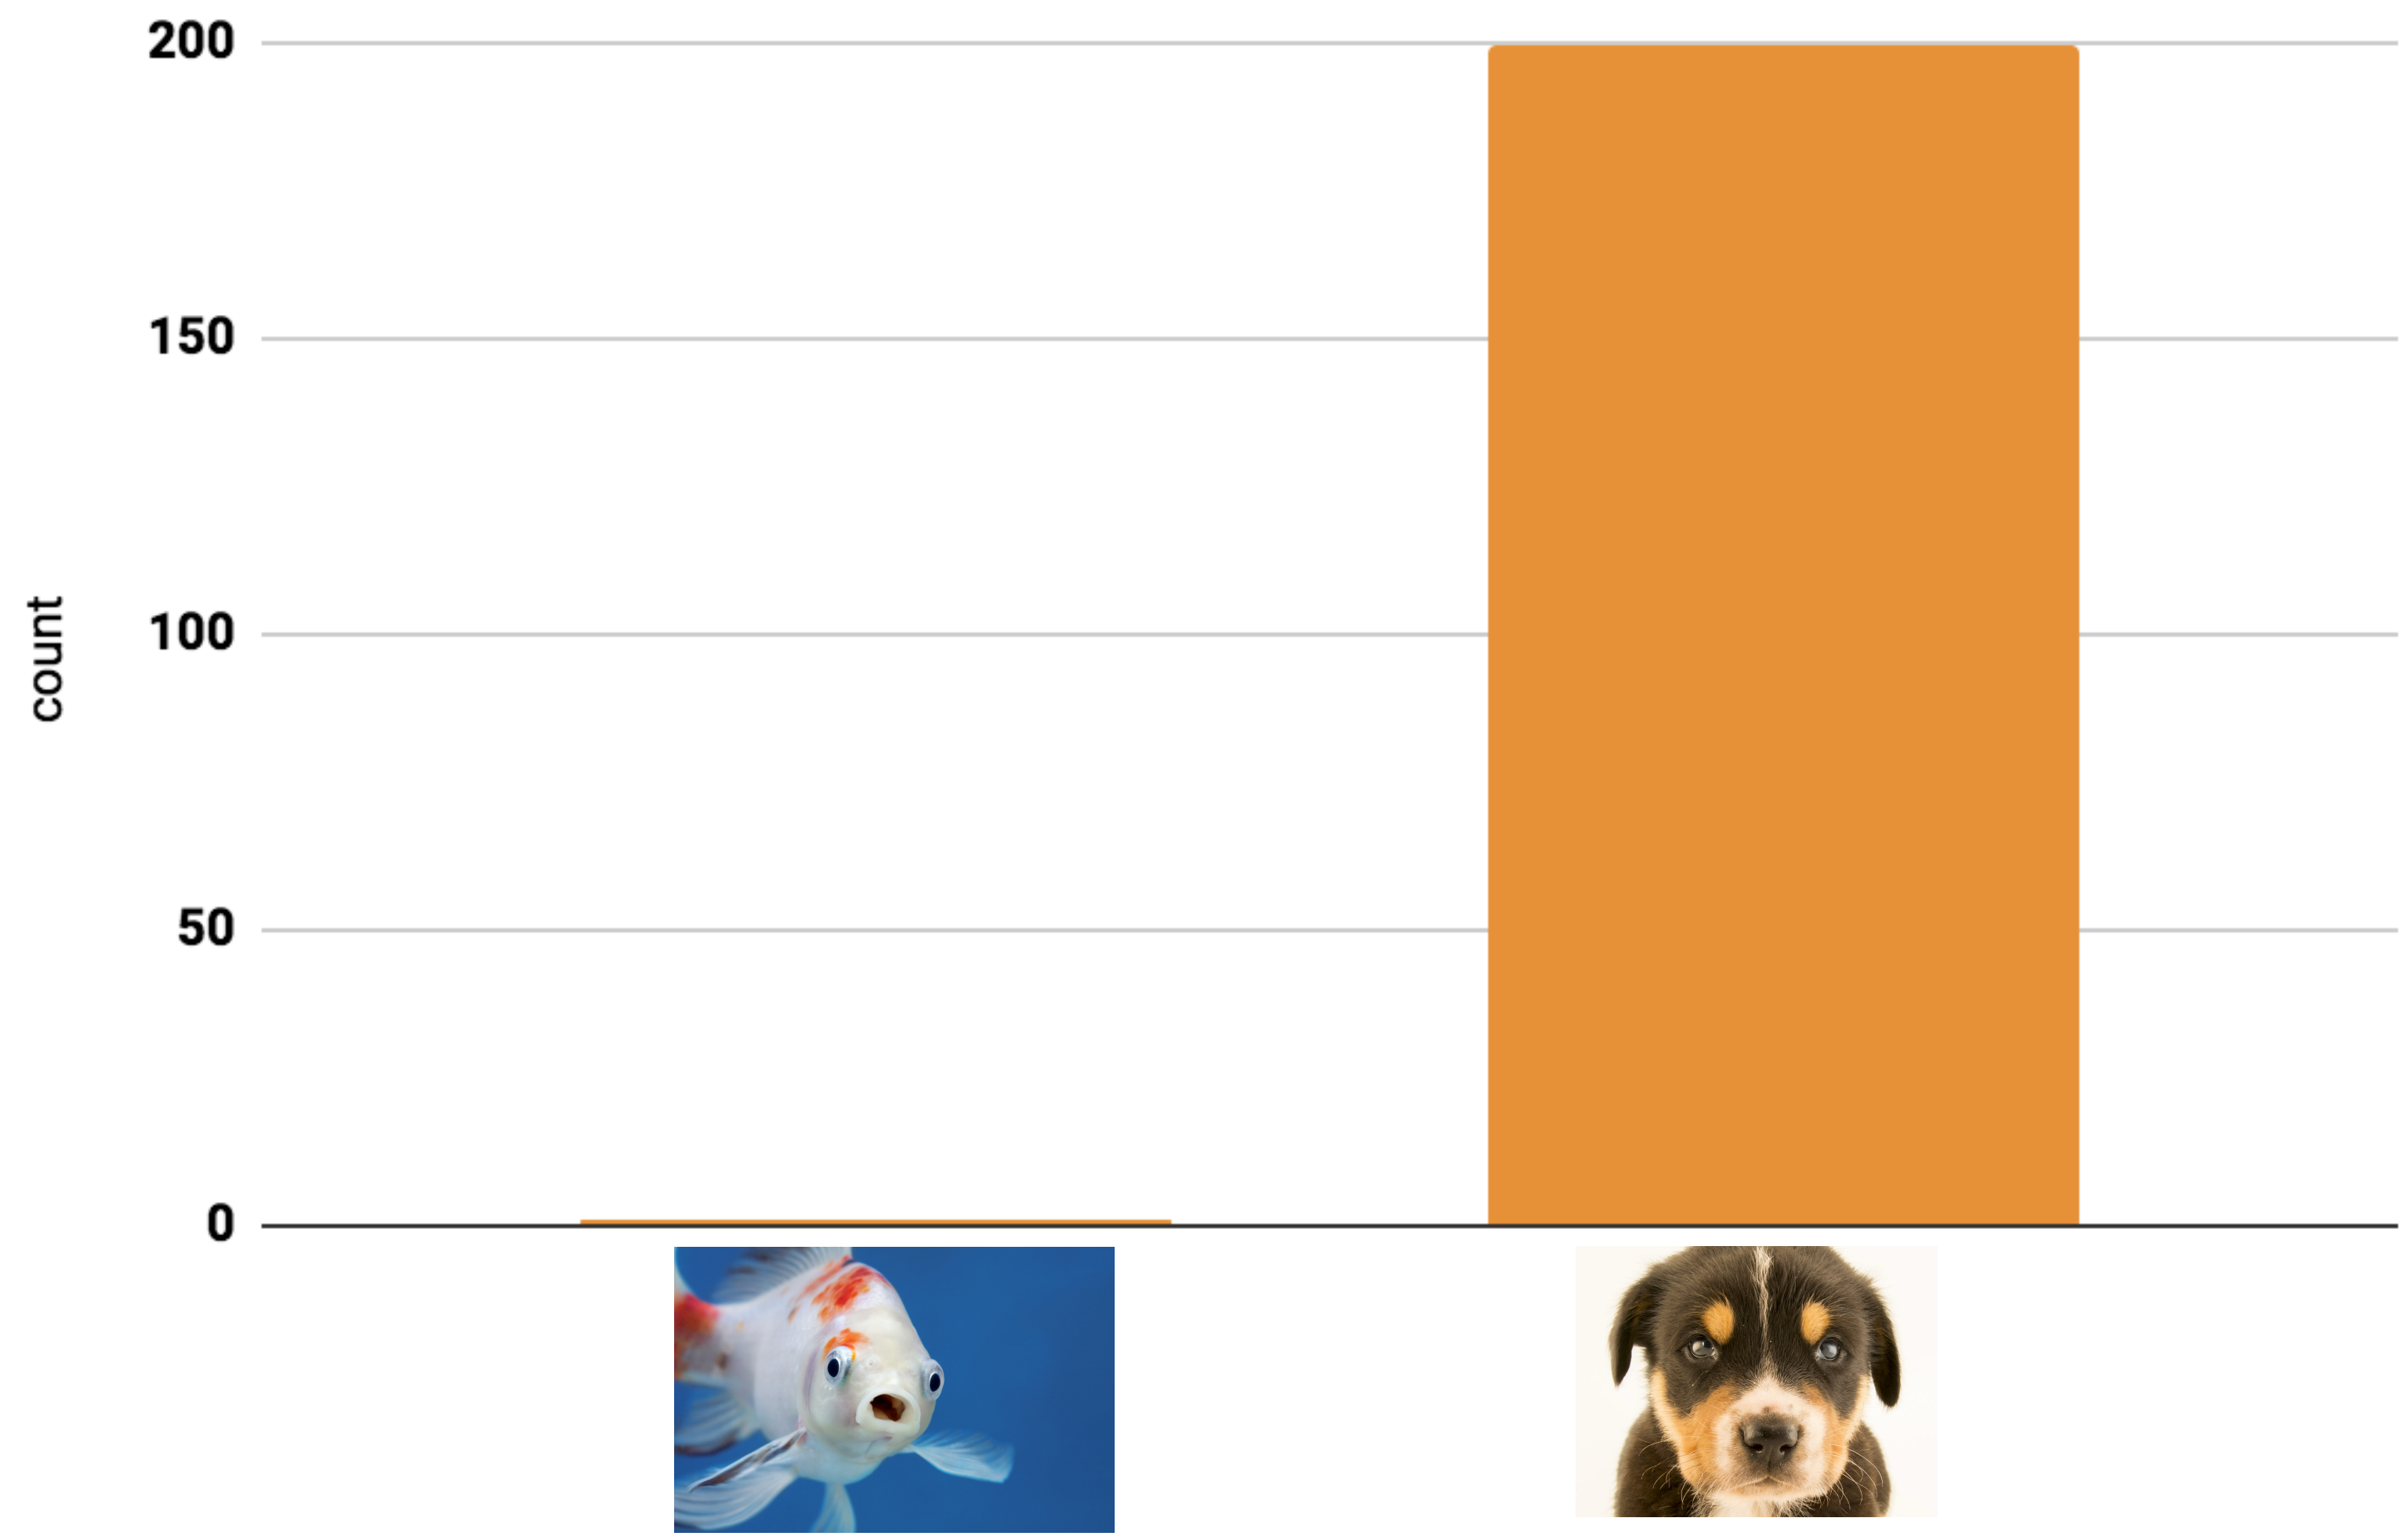


Optimization




Representation


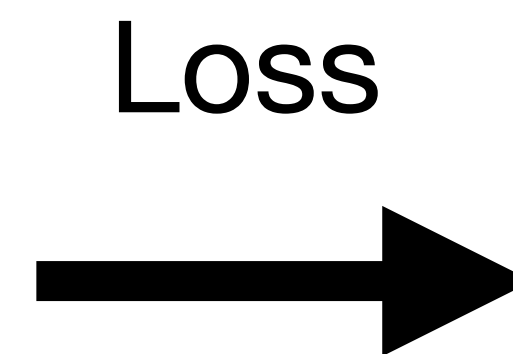
Data Imbalance




Optimization challenges





Mini batch



Optimization challenges




Loss







Mini batch


Loss








Representation challenges






Representation challenges








example



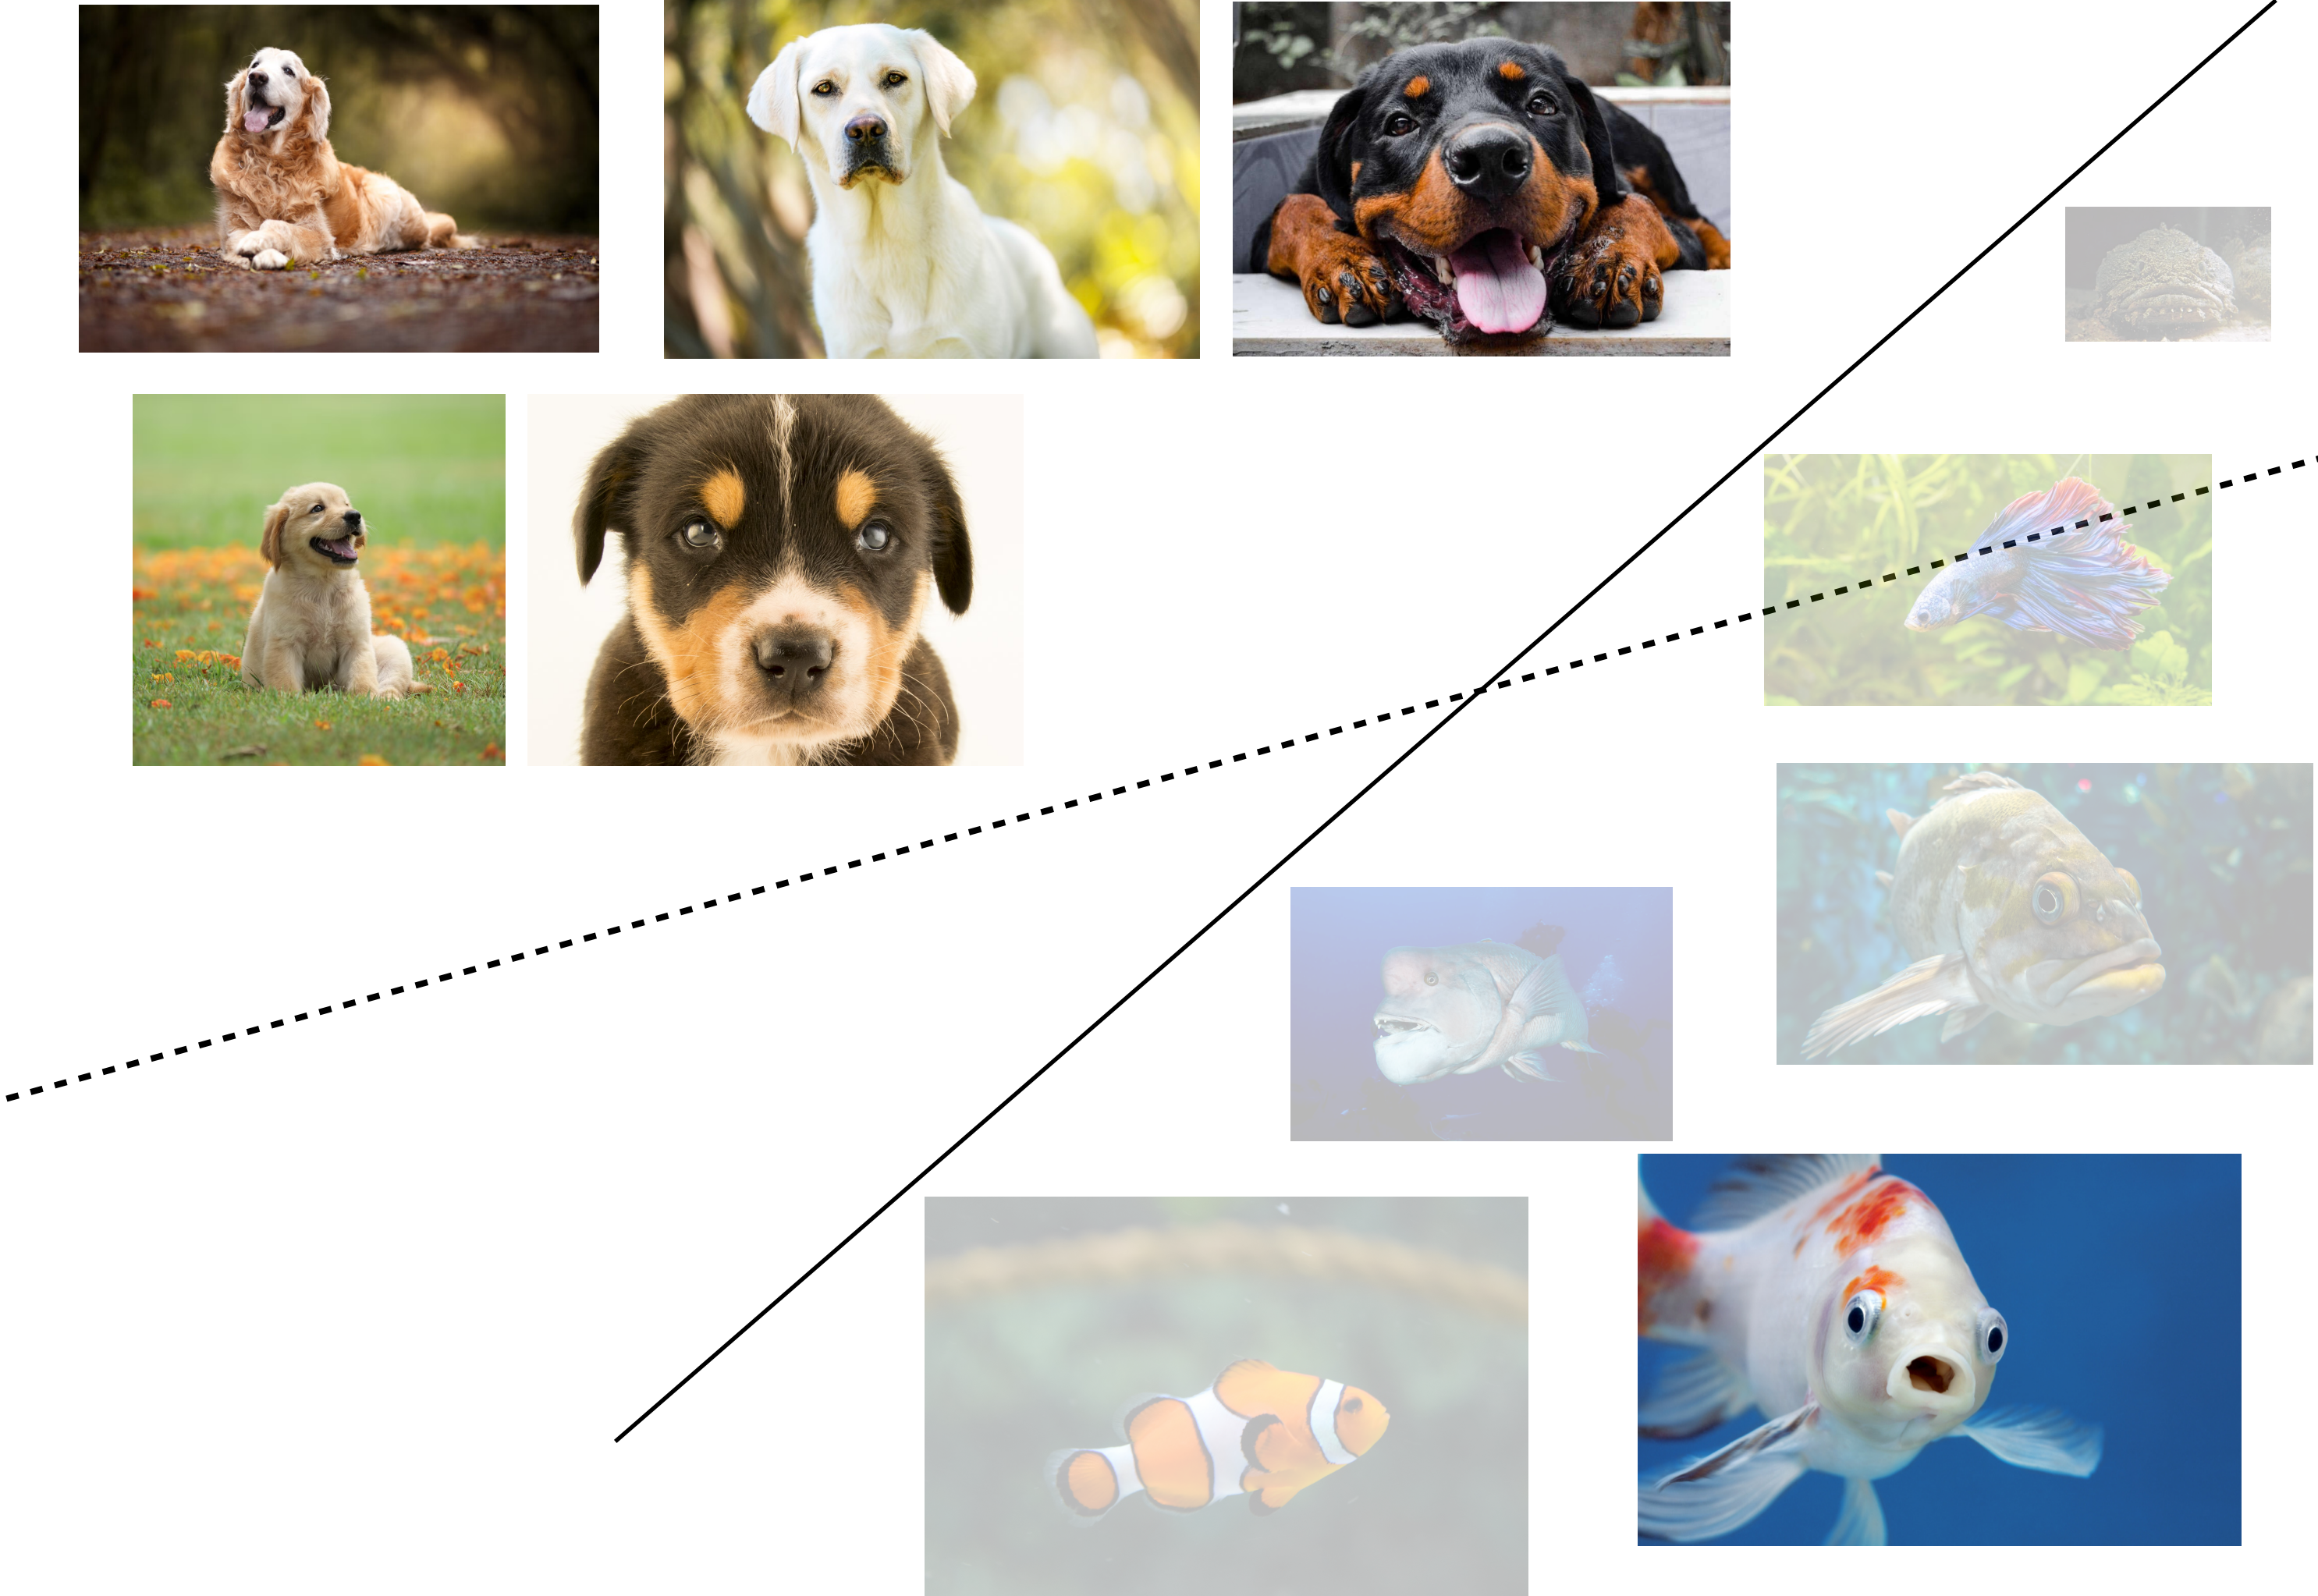


Representation challenges



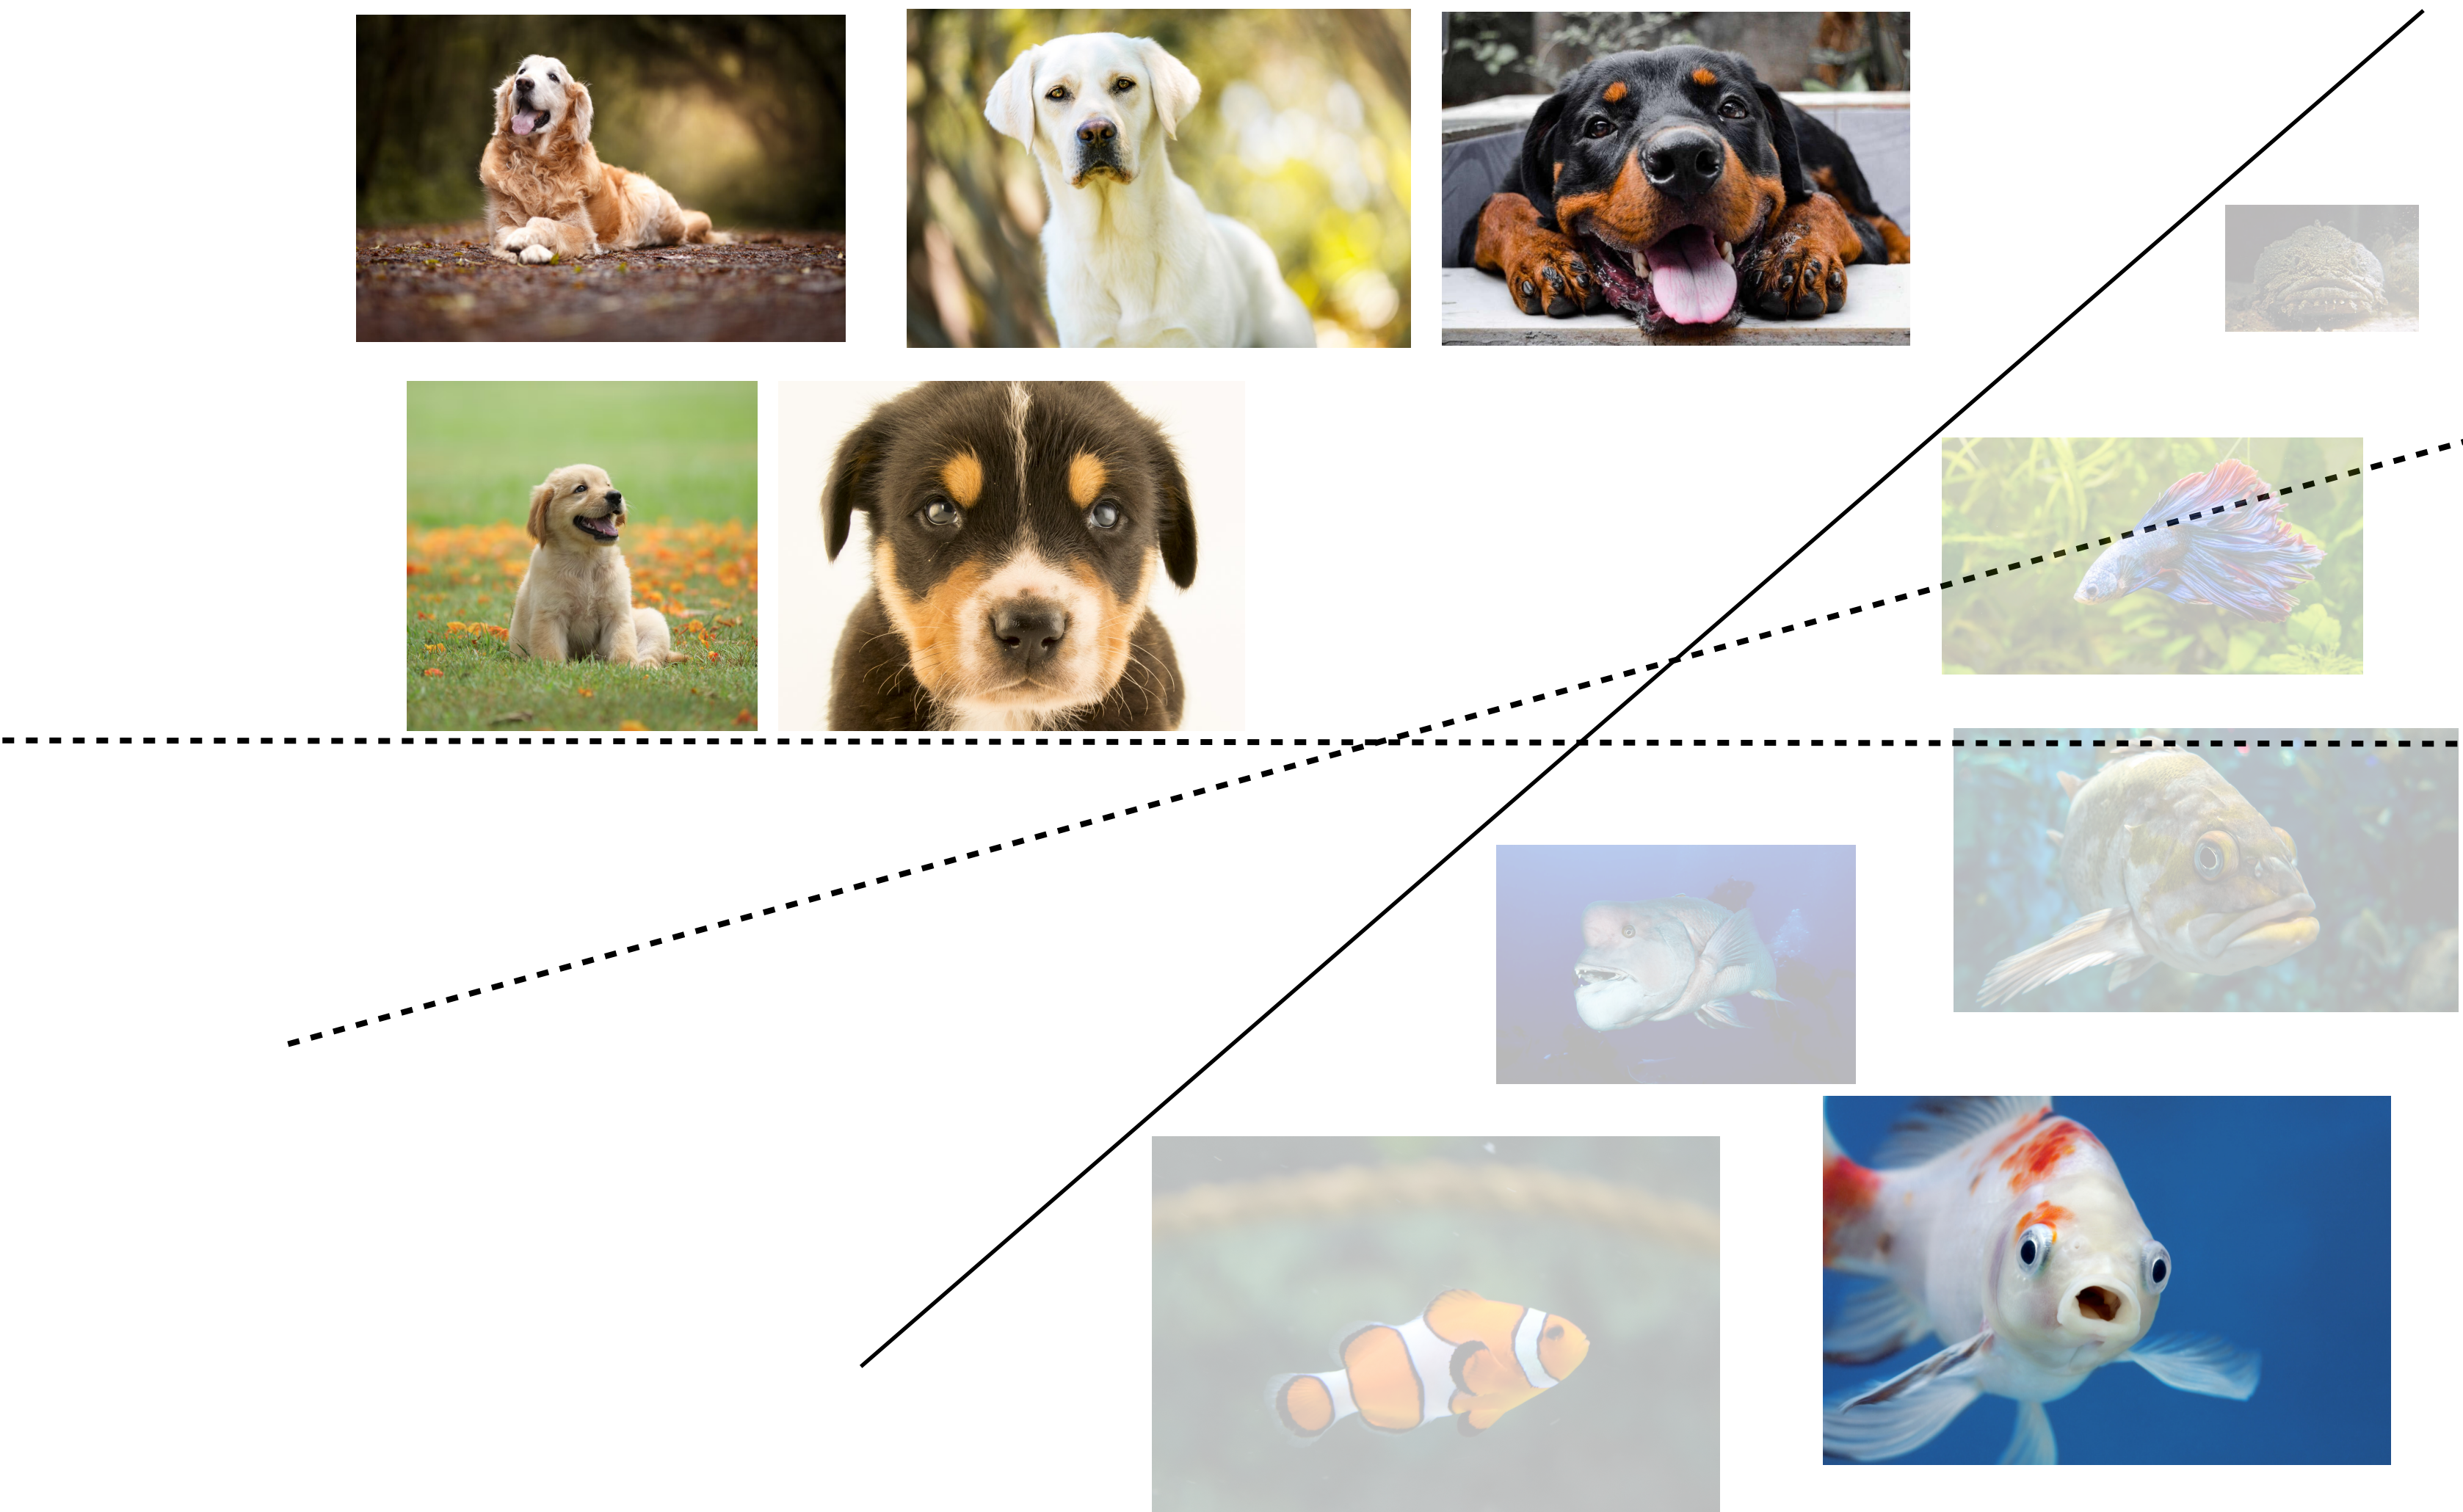


Representation challenges



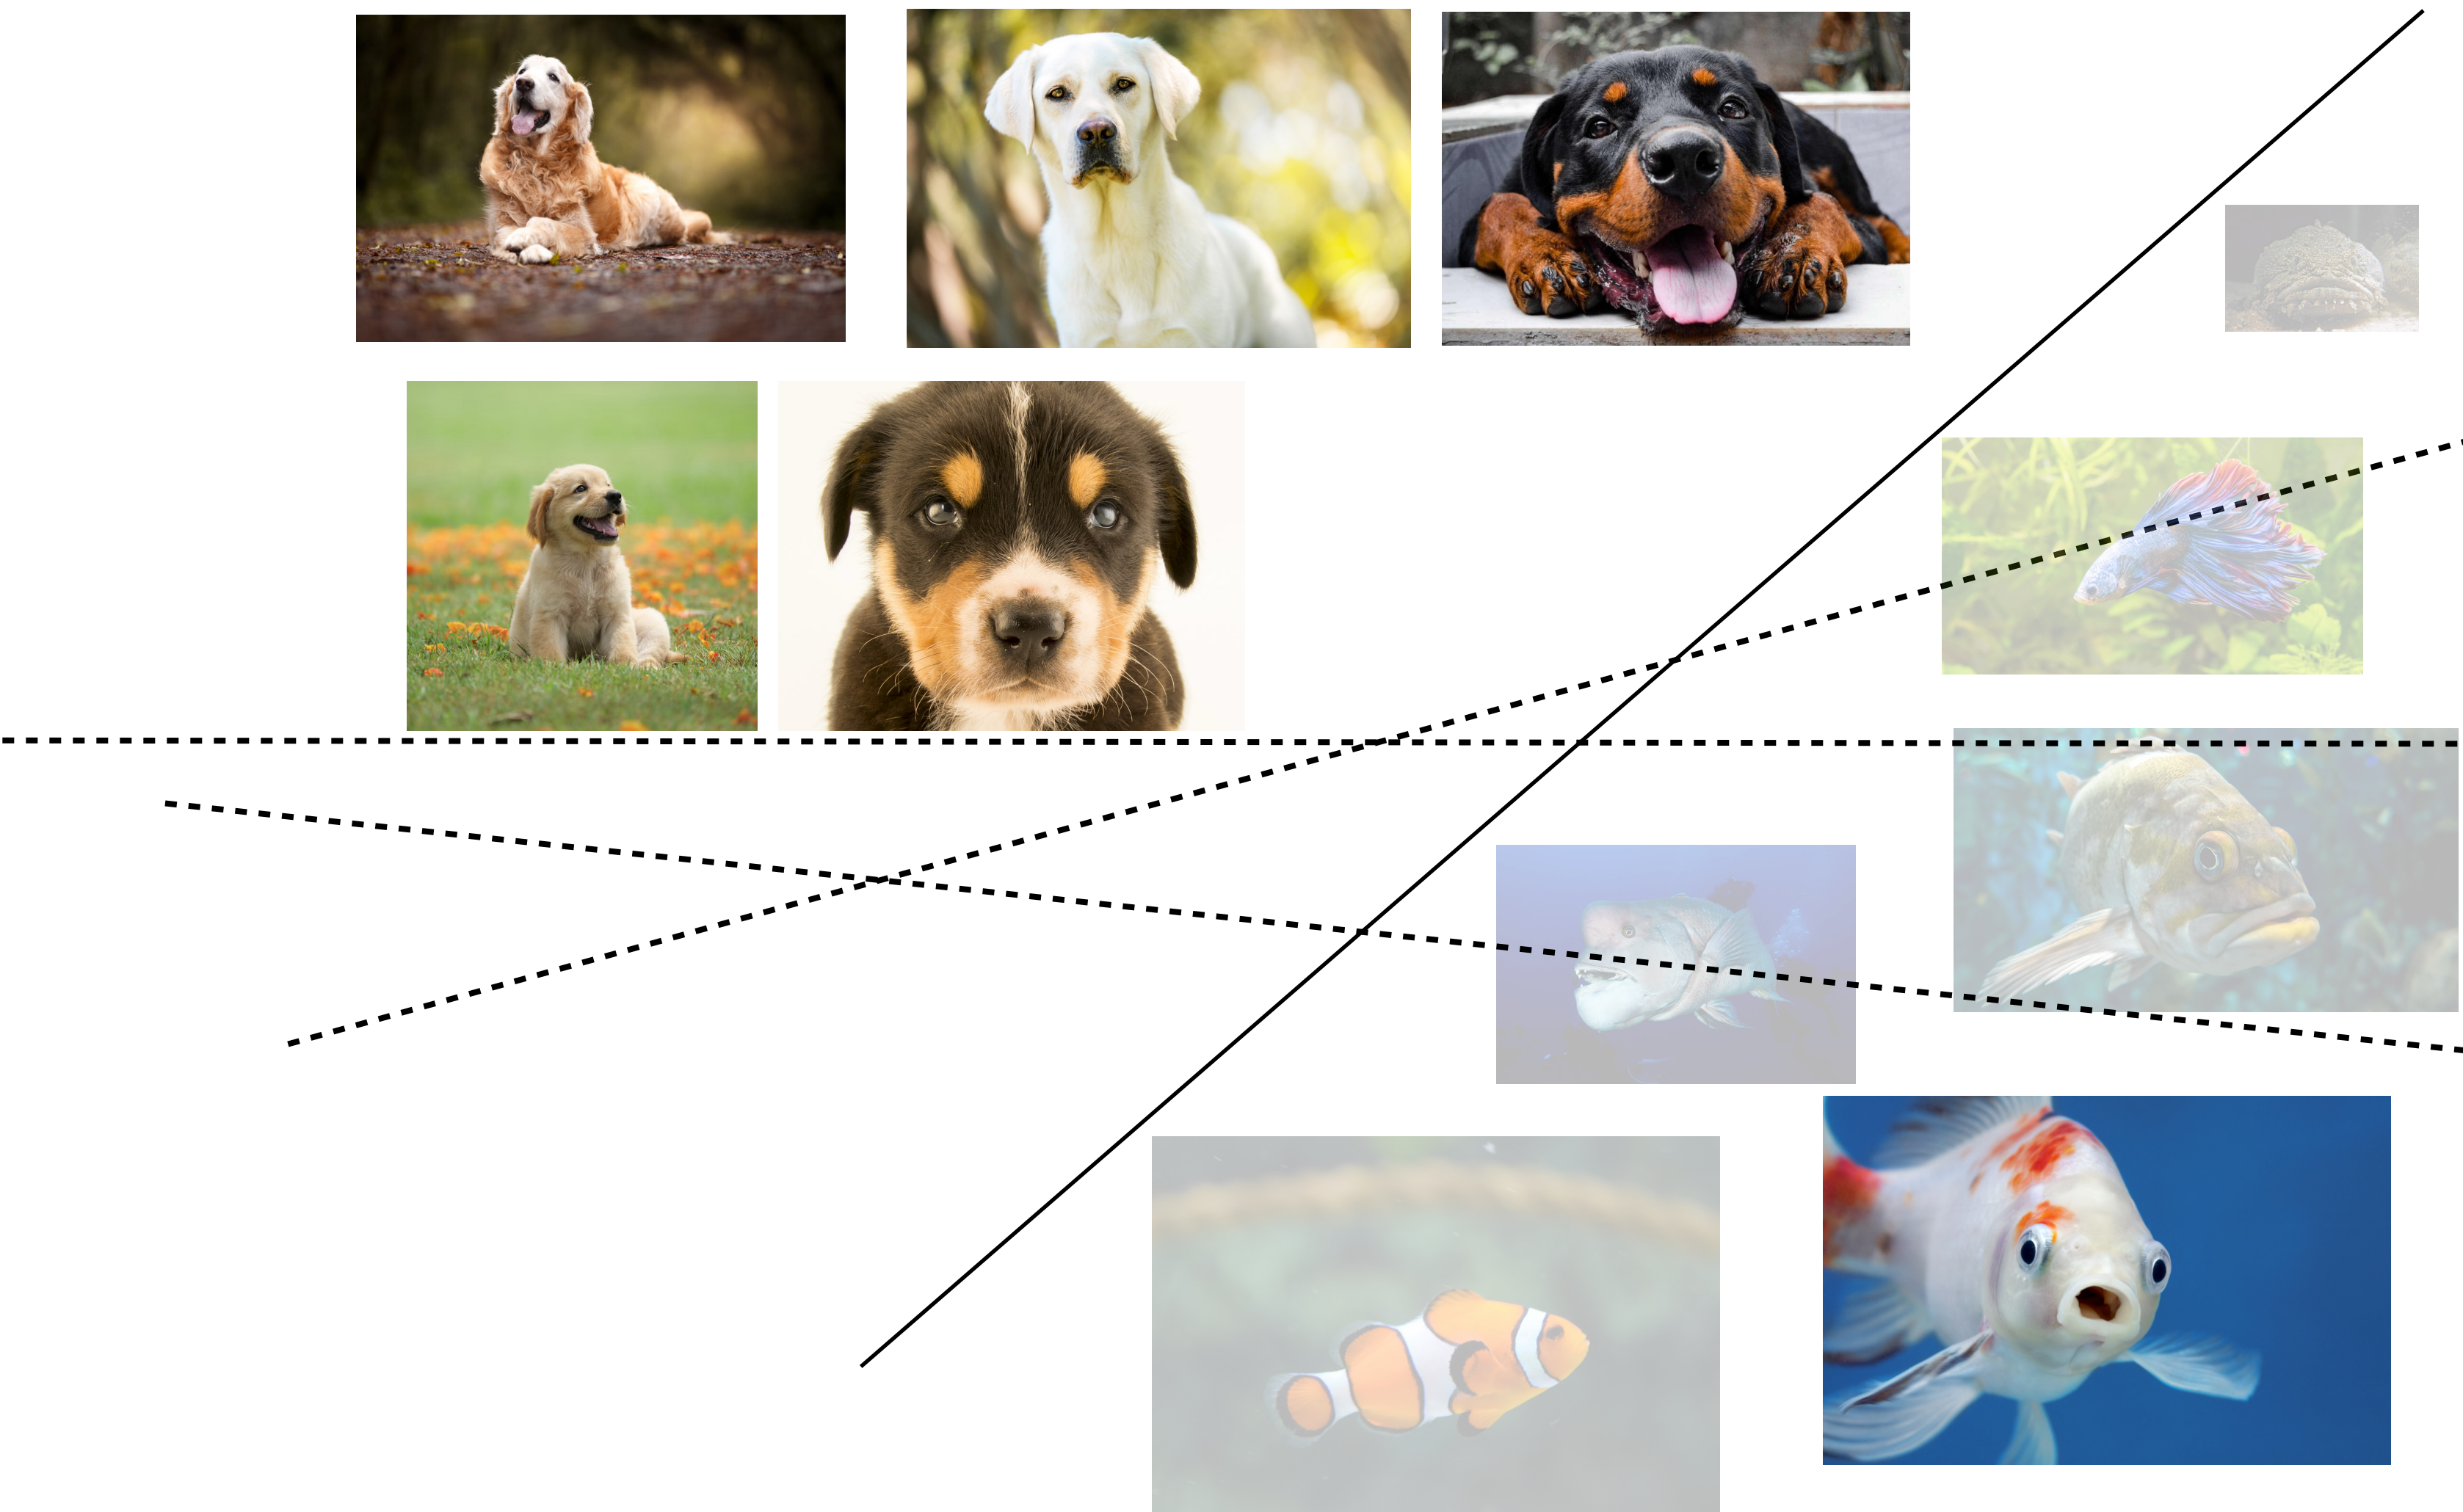



Representation challenges



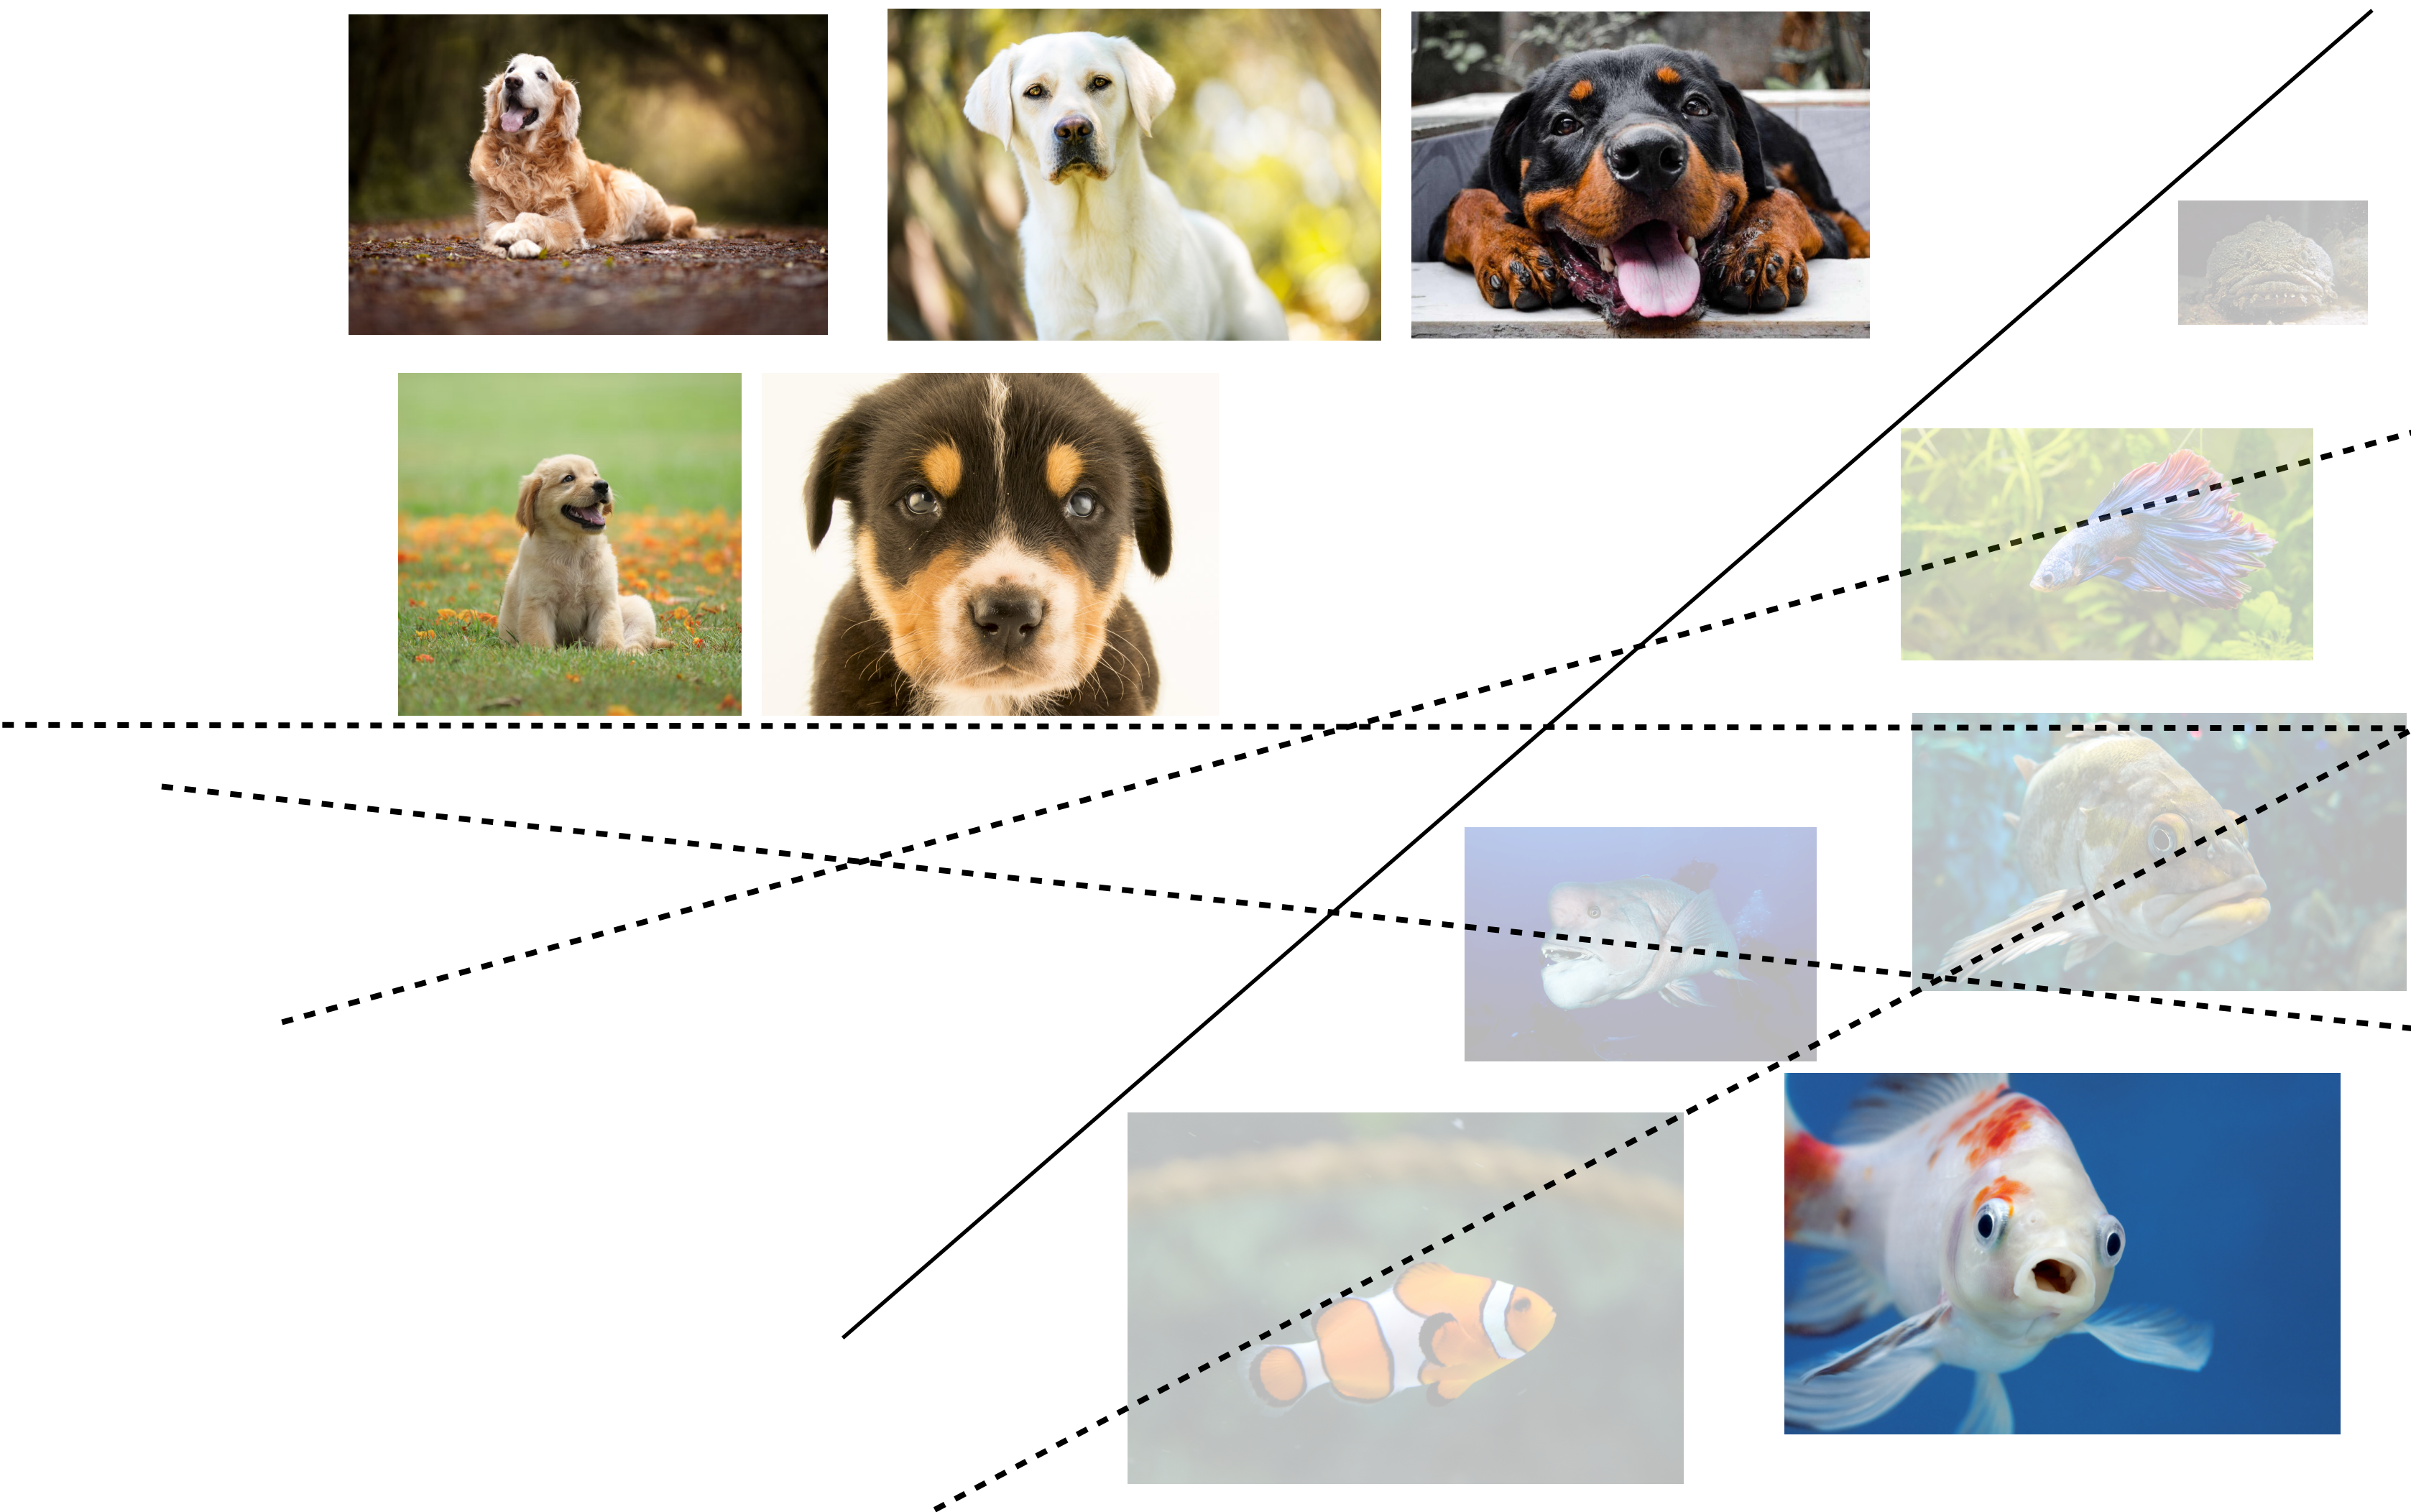






Representation challenges



Representation challenges



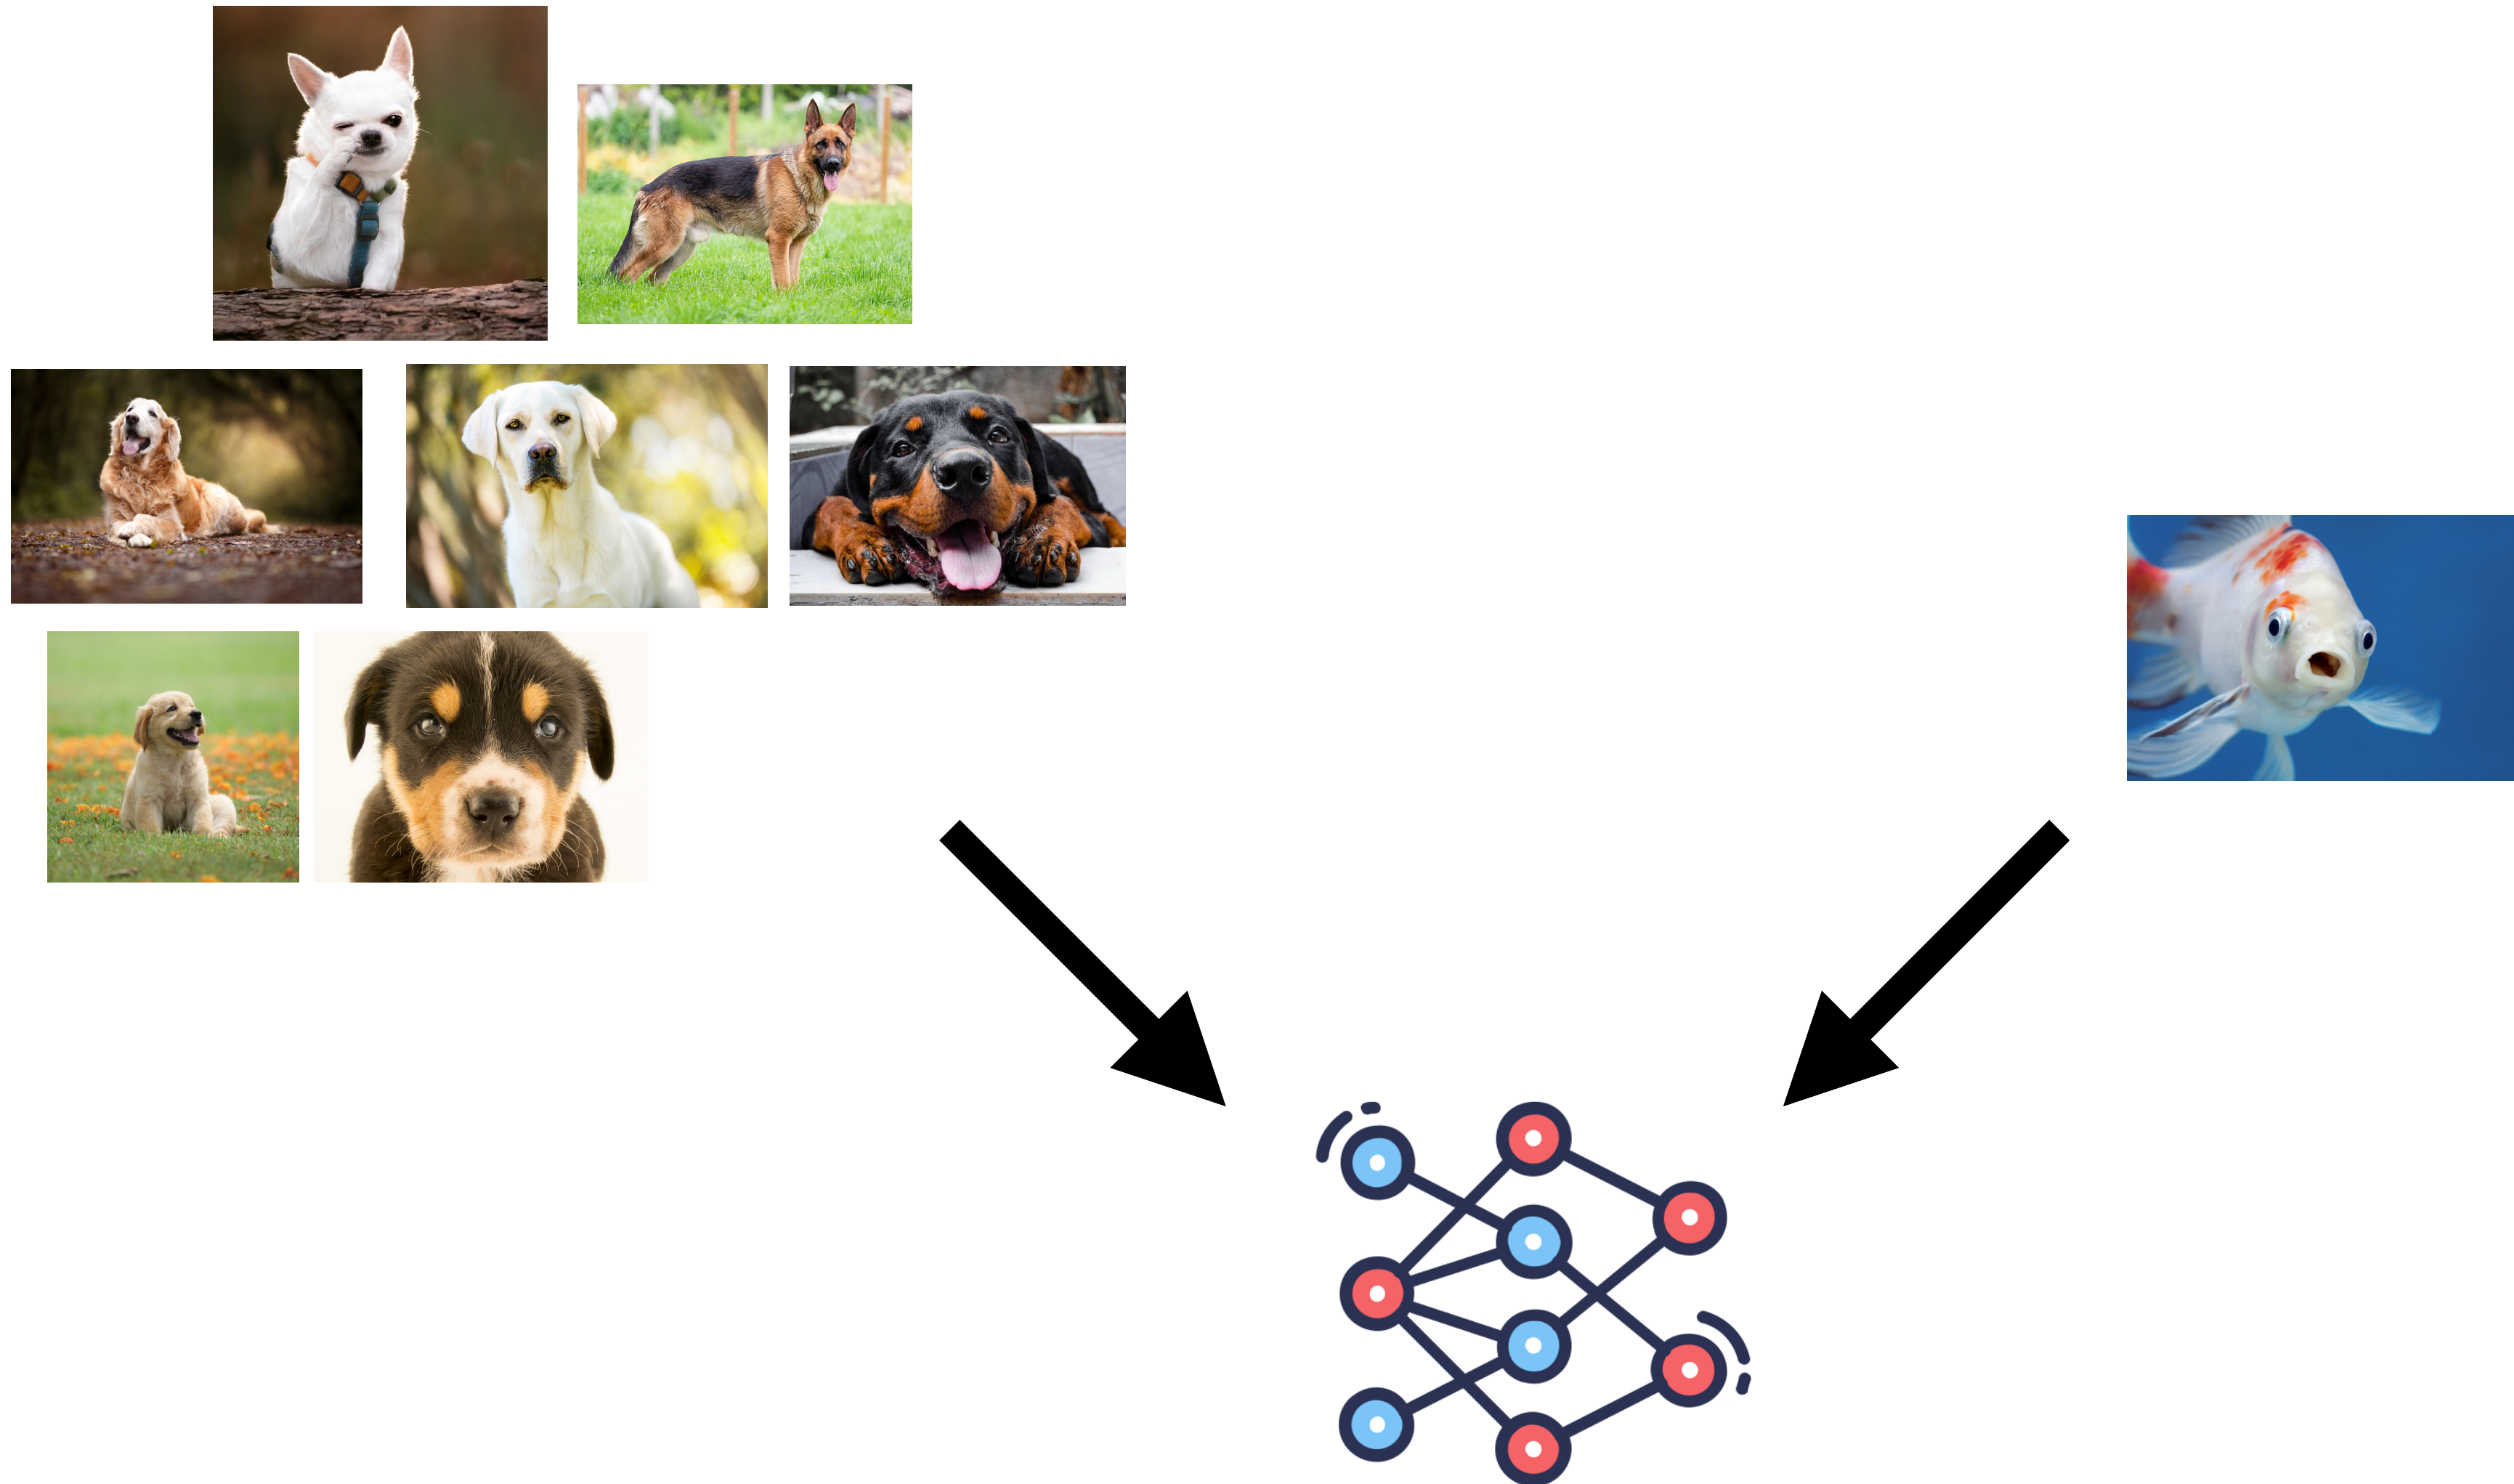
Representation challenges






Imbalance

A few strategies





Over/under-sampling



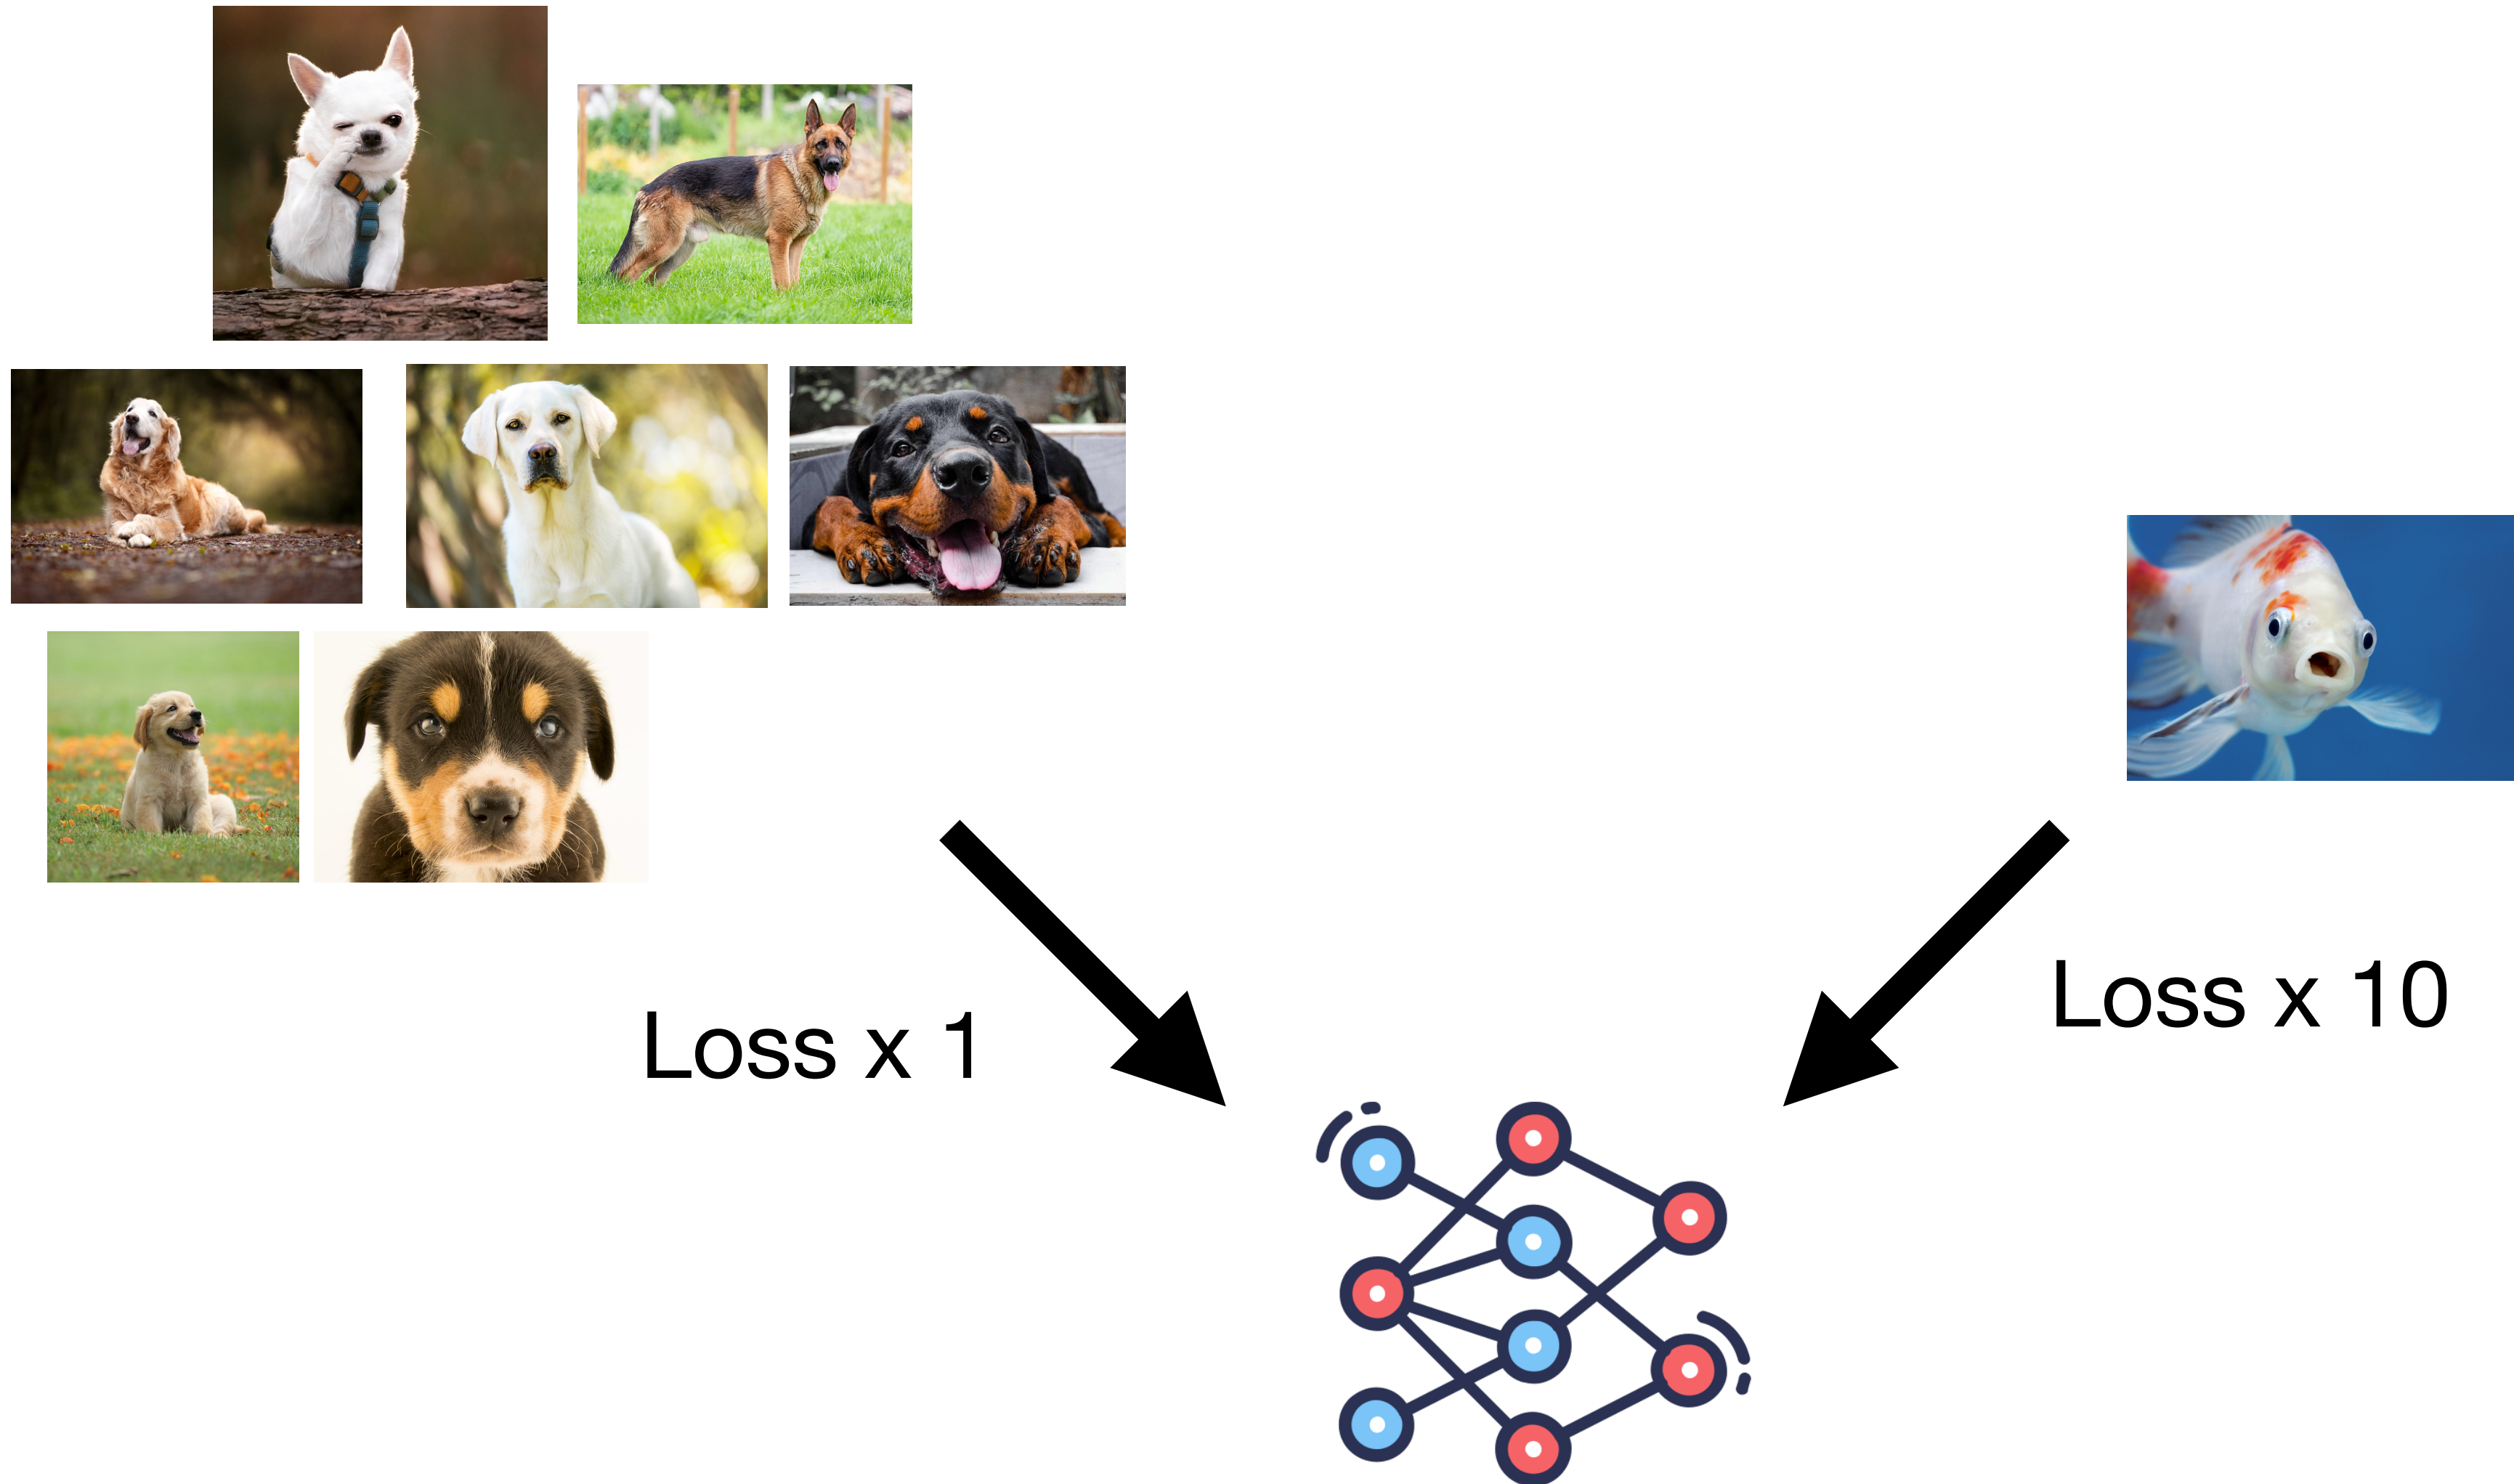
Over/under-sampling





Over/under-sampling




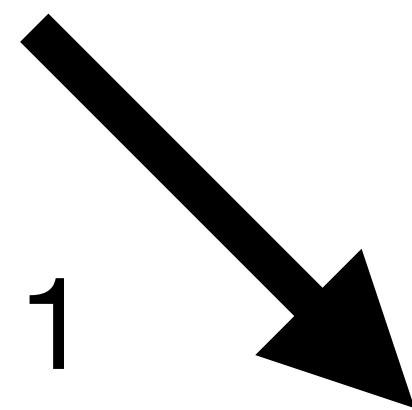
Loss weighting





Loss weighting






Loss x 1





Loss x 10




Data augmentation/generation




Data augmentation/generation



In practice: What is “good”?



In practice: What is “good”?




$$\text{Accuracy} = \frac{\text{Correct predictions}}{\text{All predictions}}$$

```
def my_classifier():  
    return 0
```

99.5% accuracy!

In practice: What is “good”?




$$\text{Macro Accuracy} = \frac{\text{Correct predictions class 1}}{\text{All predictions class 1}} + \frac{\text{Correct predictions class 0}}{\text{All predictions class 0}}$$

$$\text{Macro Accuracy} = \frac{\text{Correct predictions class 1} + \text{Correct predictions class 0}}{\text{Num classes}}$$


```
def my_classifier():  
    return 0
```

50.25% accuracy

Fine-Grained Classification



Two-spotted ladybug
Adalia bipunctata




Seven-spotted ladybug
Coccinella septempunctata

Figure 1. Two visually similar species from the iNat2017 dataset. Through close inspection, we can see that the ladybug on the left has *two* spots while the one on the right has *seven*.

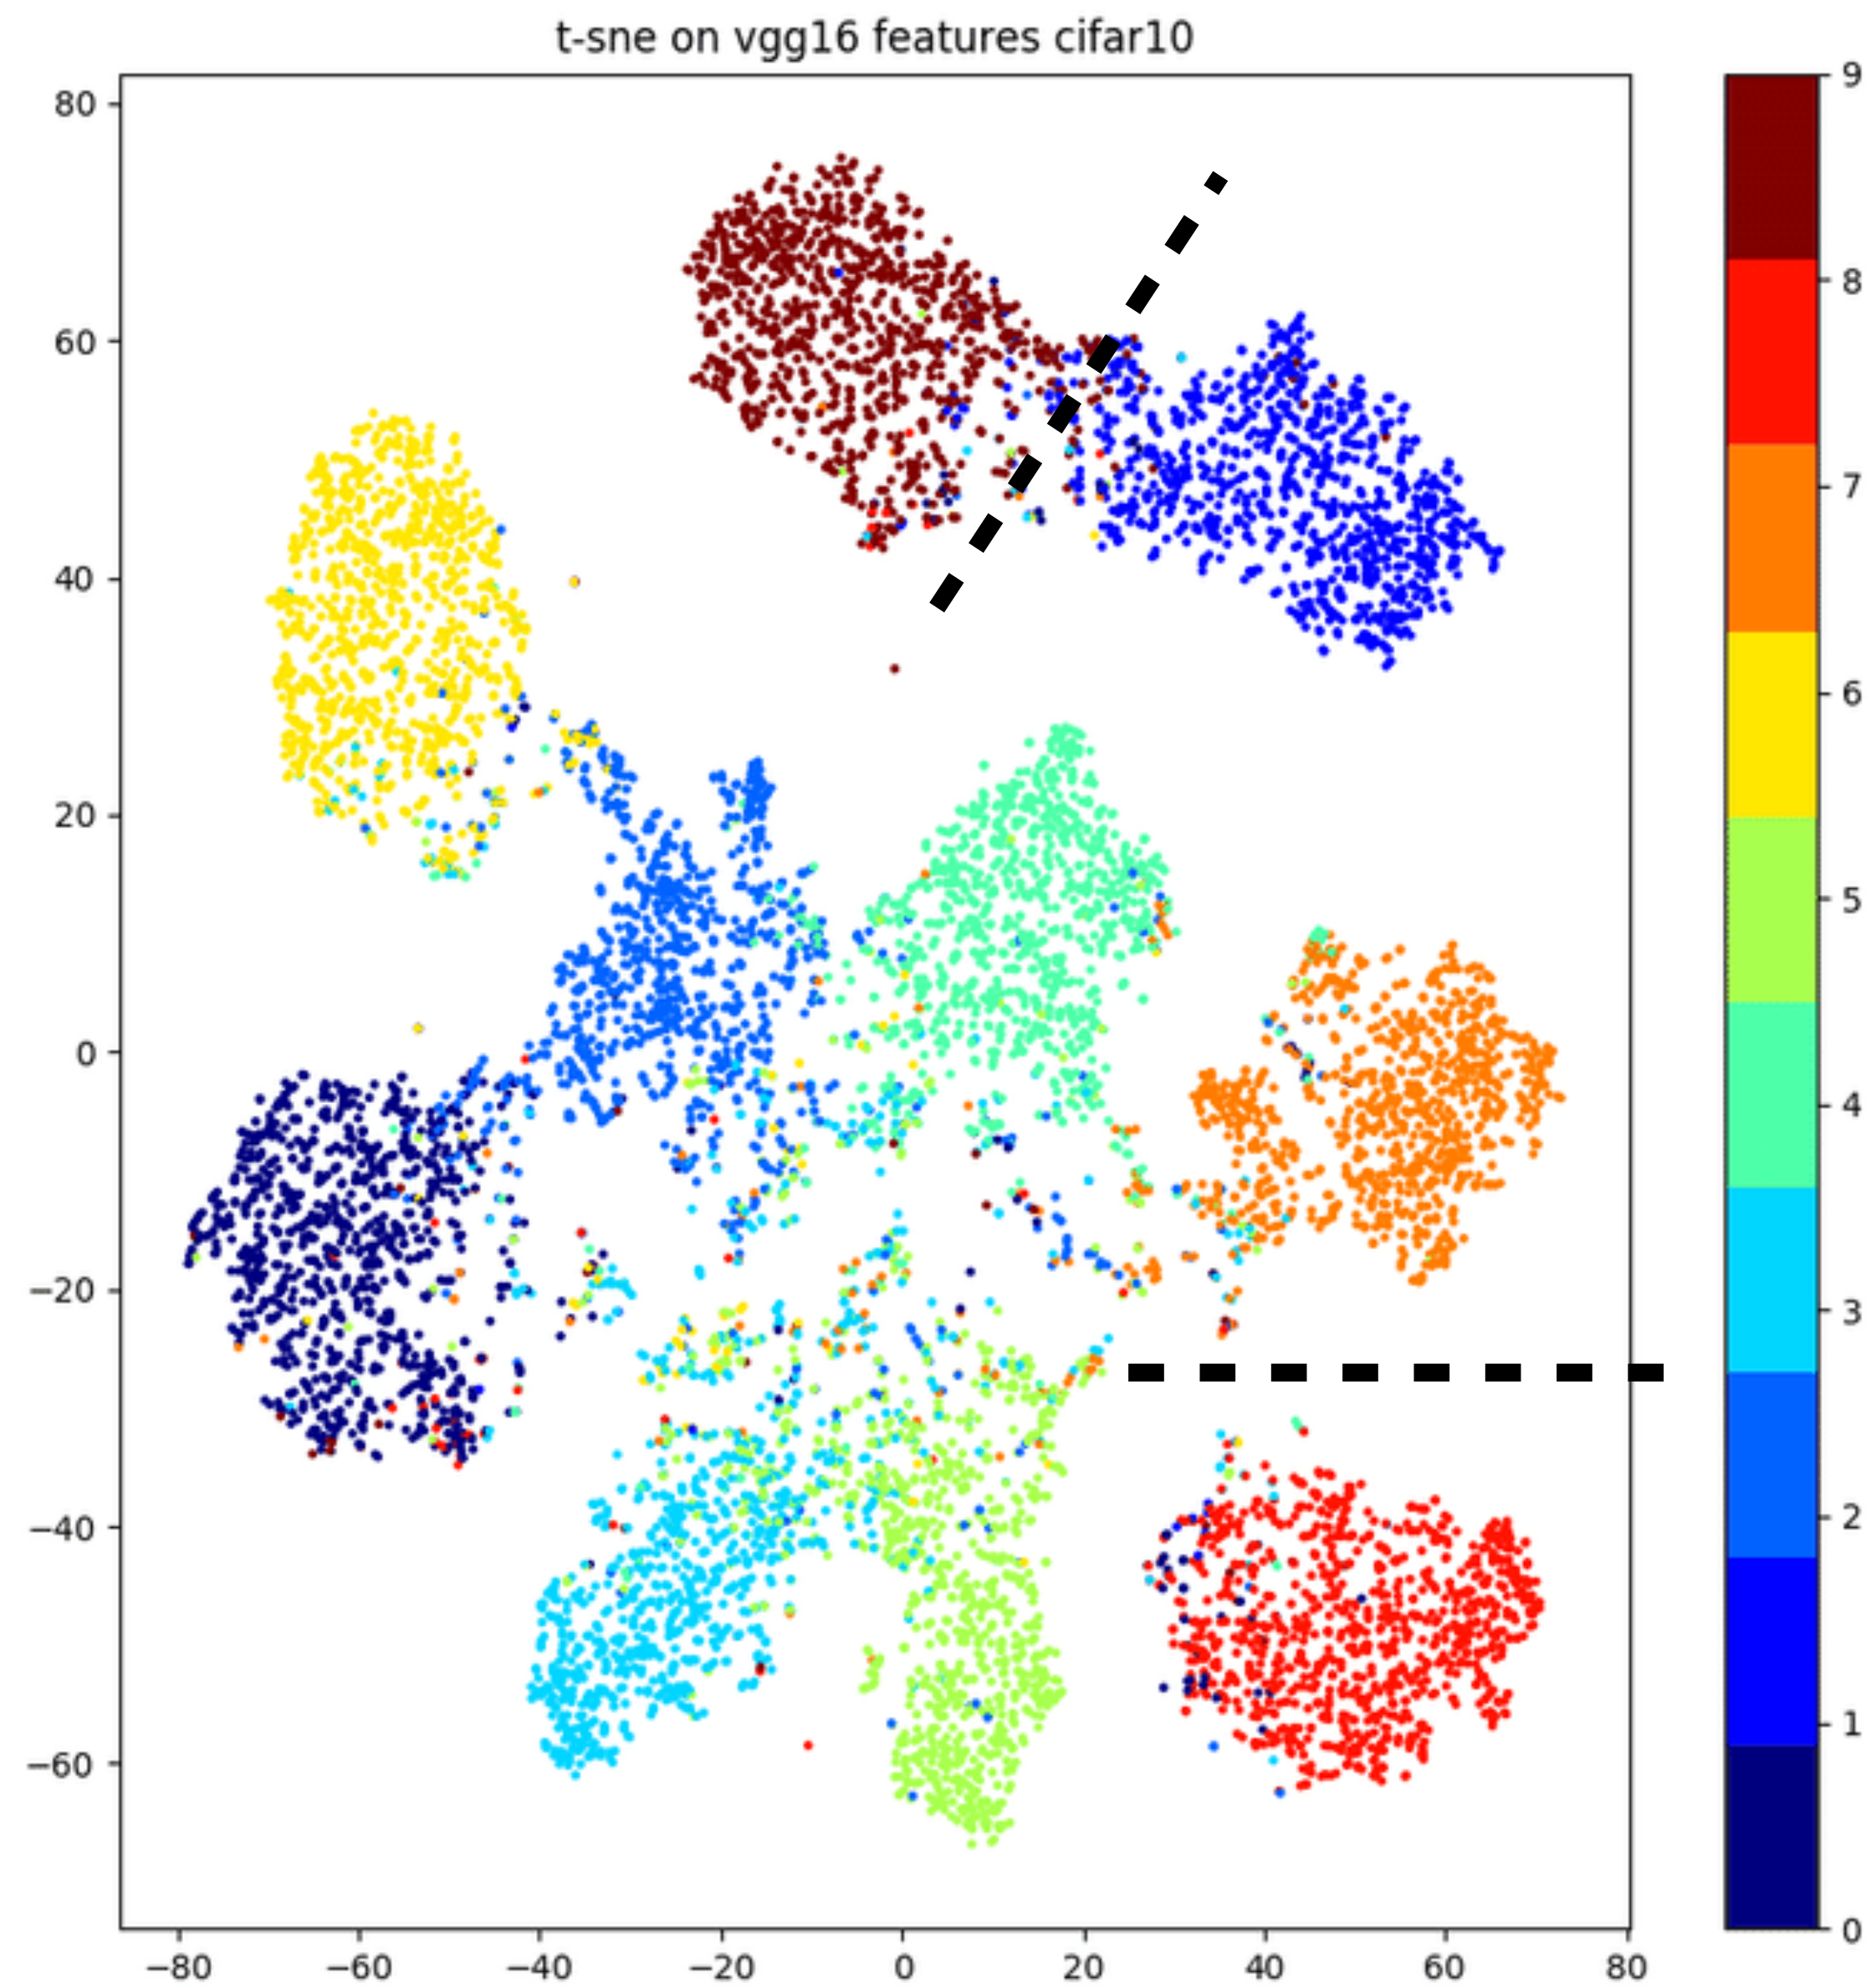


Granularity spectrum

Representation challenges



Representation challenges



Simple solution: More representation capacity

Image Classification on iNaturalist

Leaderboard

Dataset

View


Top 1 Accuracy

by

Date

for

All models



Simple solution: More representation capacity

Caveat: need enough data to separate the classes

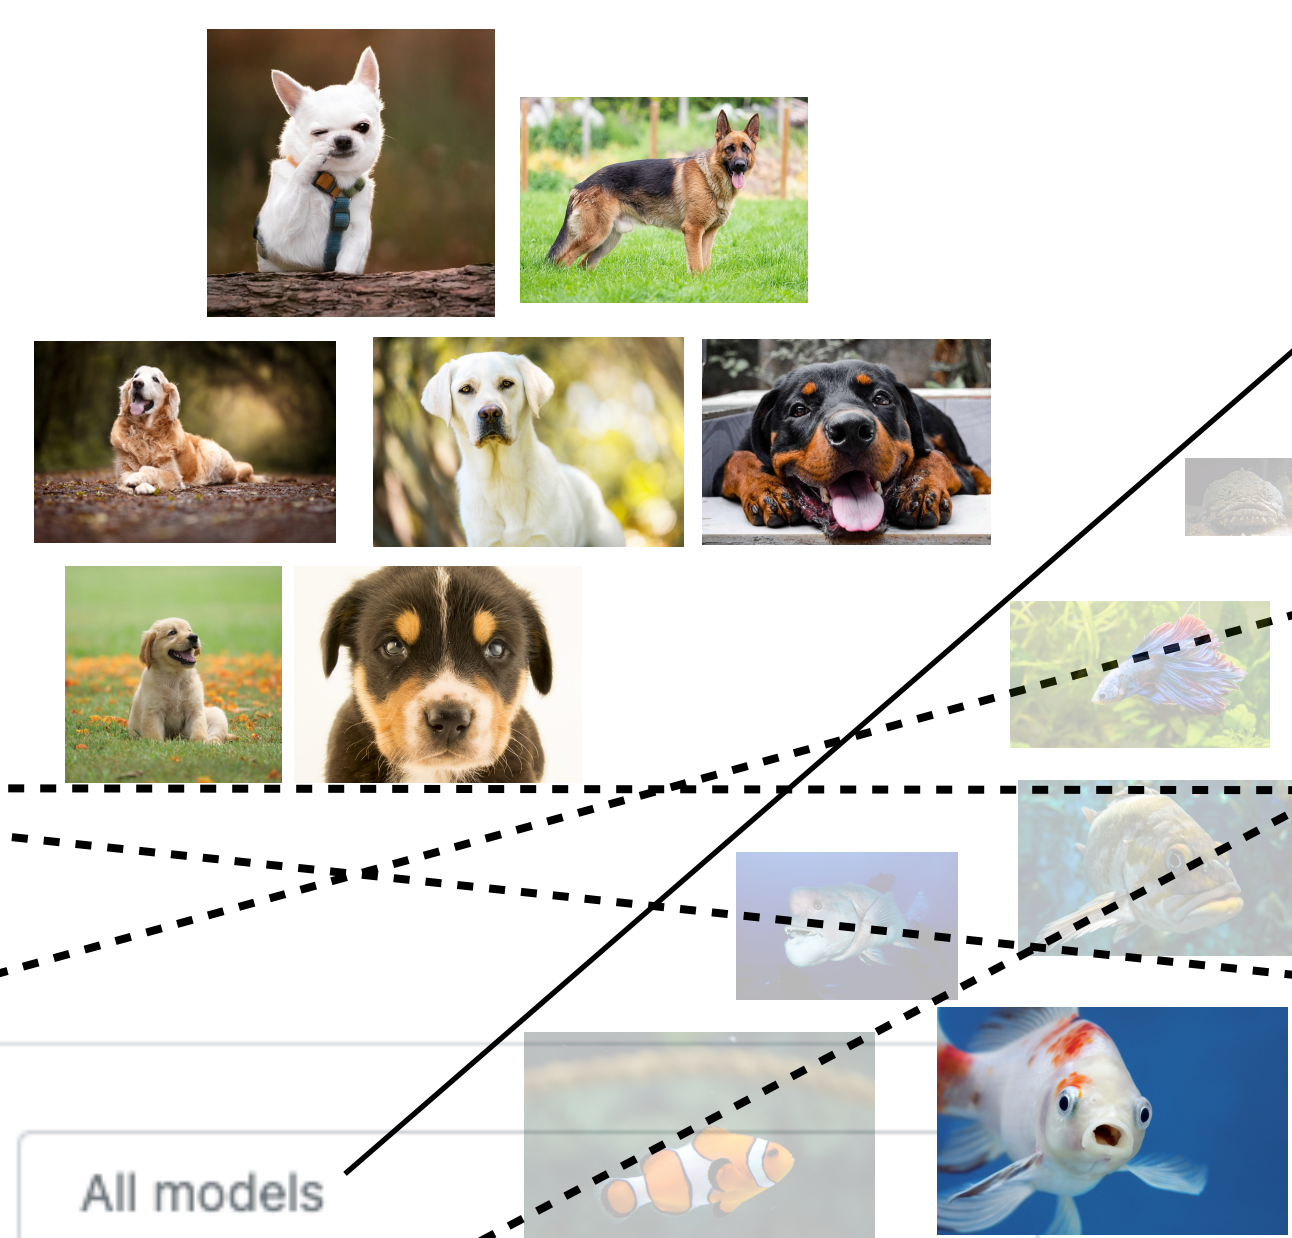


Image Classification on iNaturalist

Leaderboard

Dataset

View


Top 1 Accuracy

by


Date

for


All models




How do we get here?



(a) Initial Data



(b) Cross-Entropy Loss



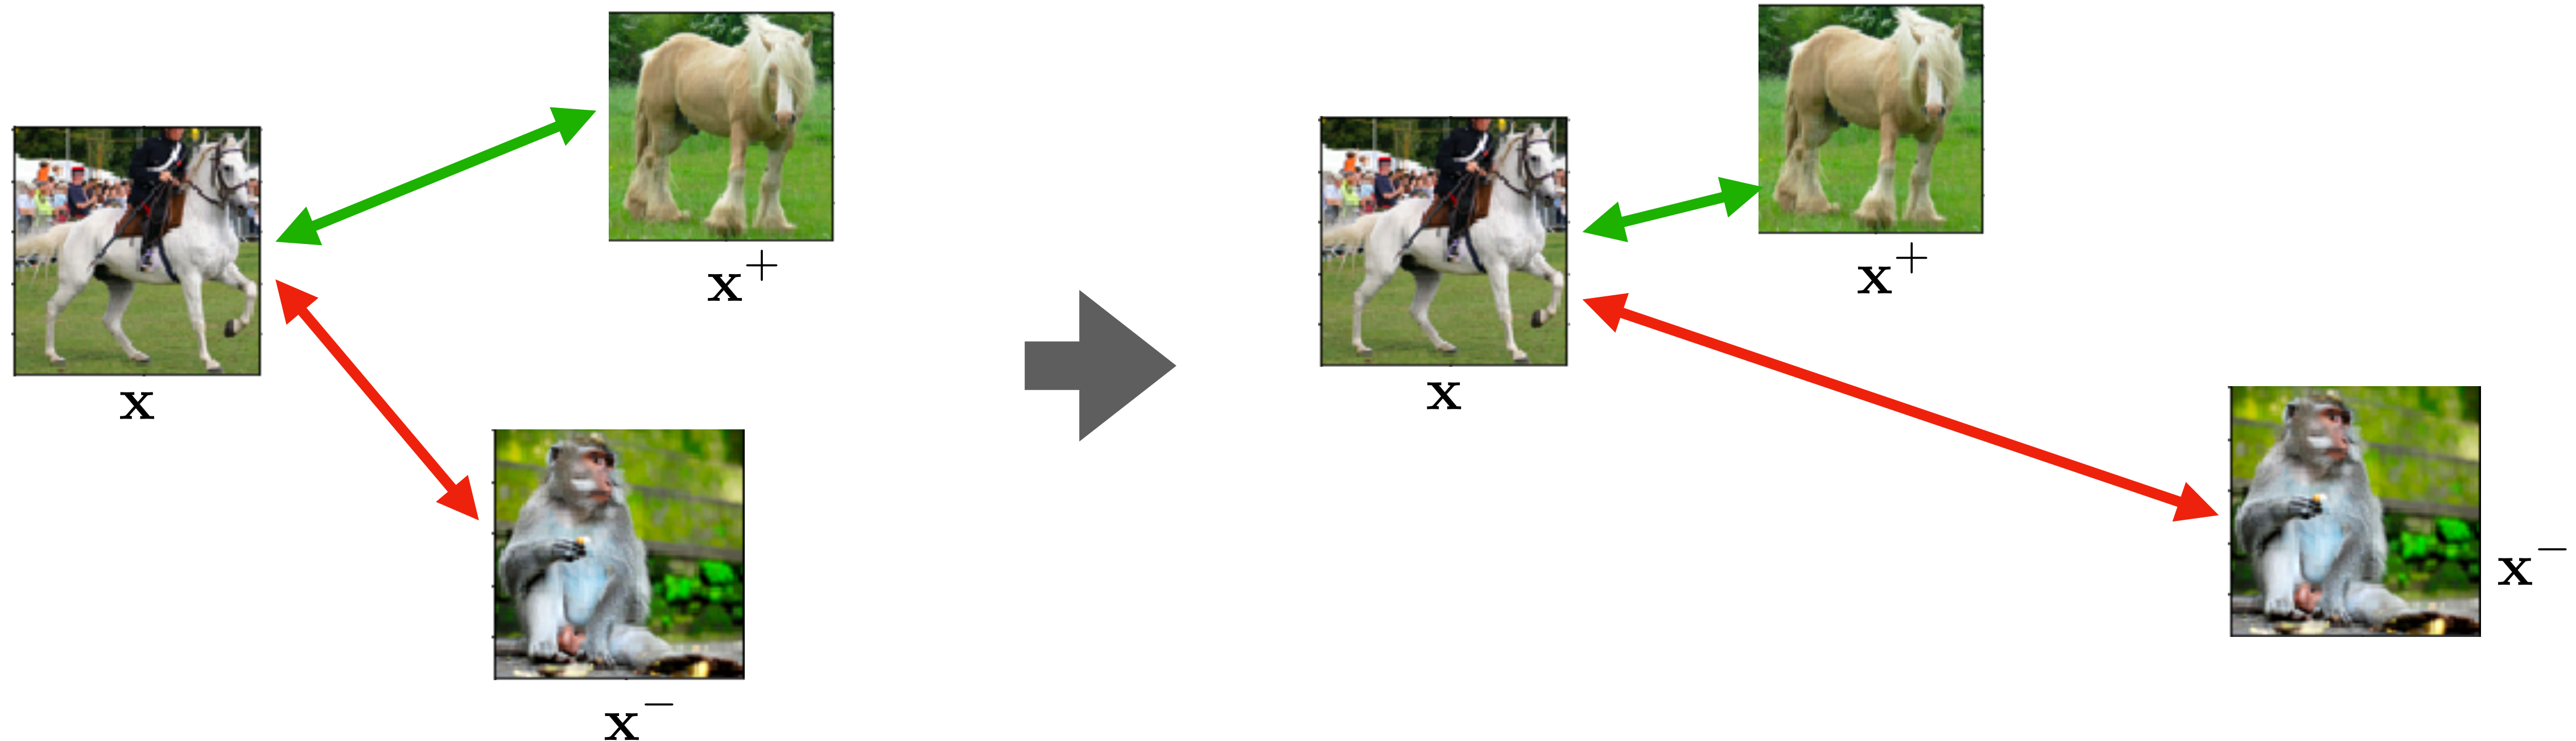
Metric / Contrastive Learning

distance of dissimilar pair(s) \gg **distance of similar pair(s)**


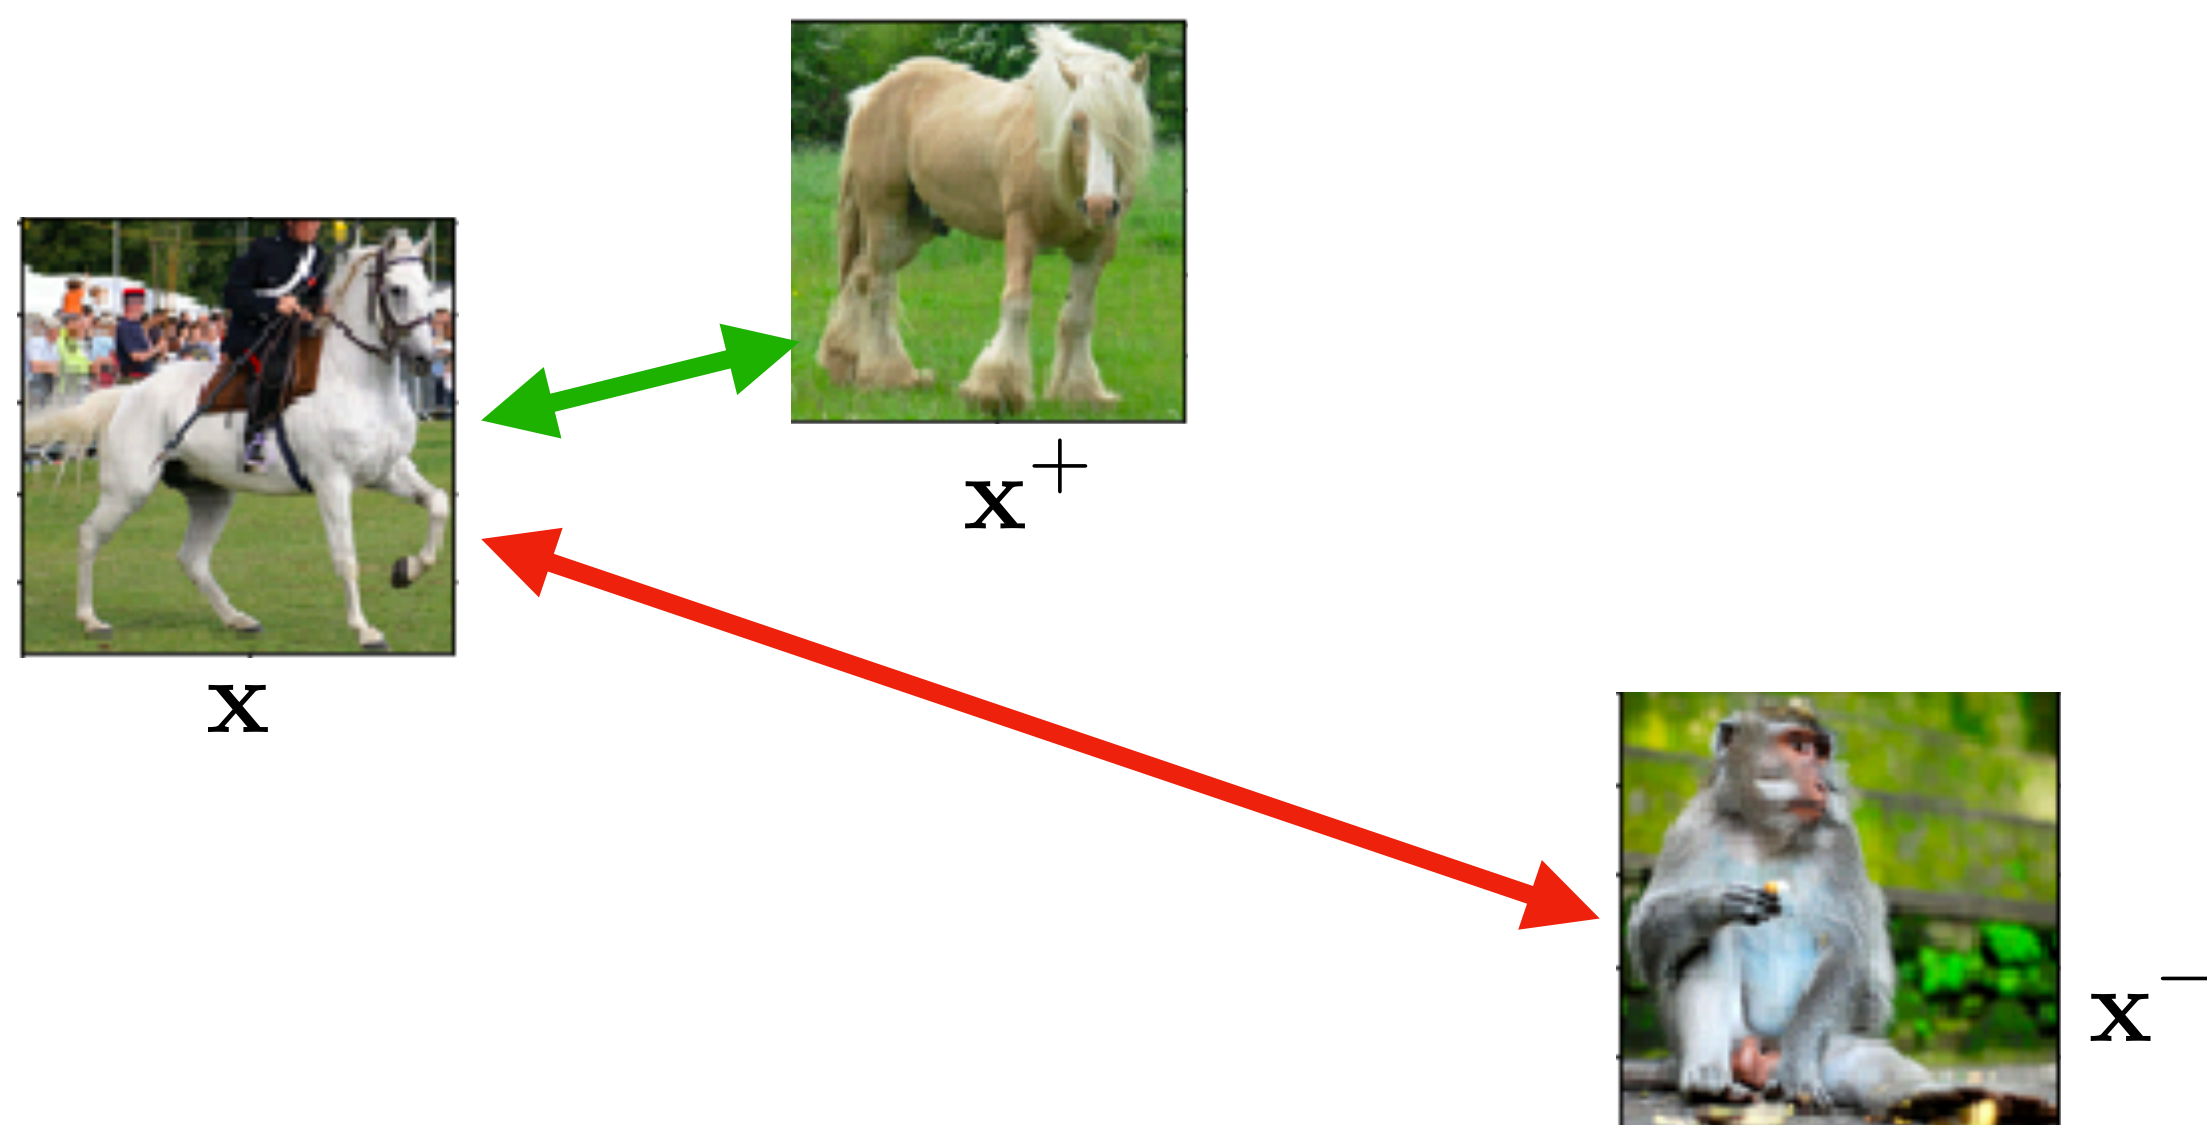
- Triplet loss (Schroff et al 2015):

$$\mathcal{L}_{\text{triplet}}(\mathbf{x}, \mathbf{x}^+, \mathbf{x}^-) = \sum_{\mathbf{x} \in \mathcal{X}} \max(0, \underbrace{\|f(\mathbf{x}) - f(\mathbf{x}^+)\|_2^2}_{\text{green line}} - \underbrace{\|f(\mathbf{x}) - f(\mathbf{x}^-)\|_2^2}_{\text{red line}} + \epsilon)$$

margin \swarrow



Importance of hard negatives




Beware the objective function



Beware the objective function











Two-spotted ladybug
Adalia bipunctata

Seven-spotted ladybug
Coccinella septempunctata

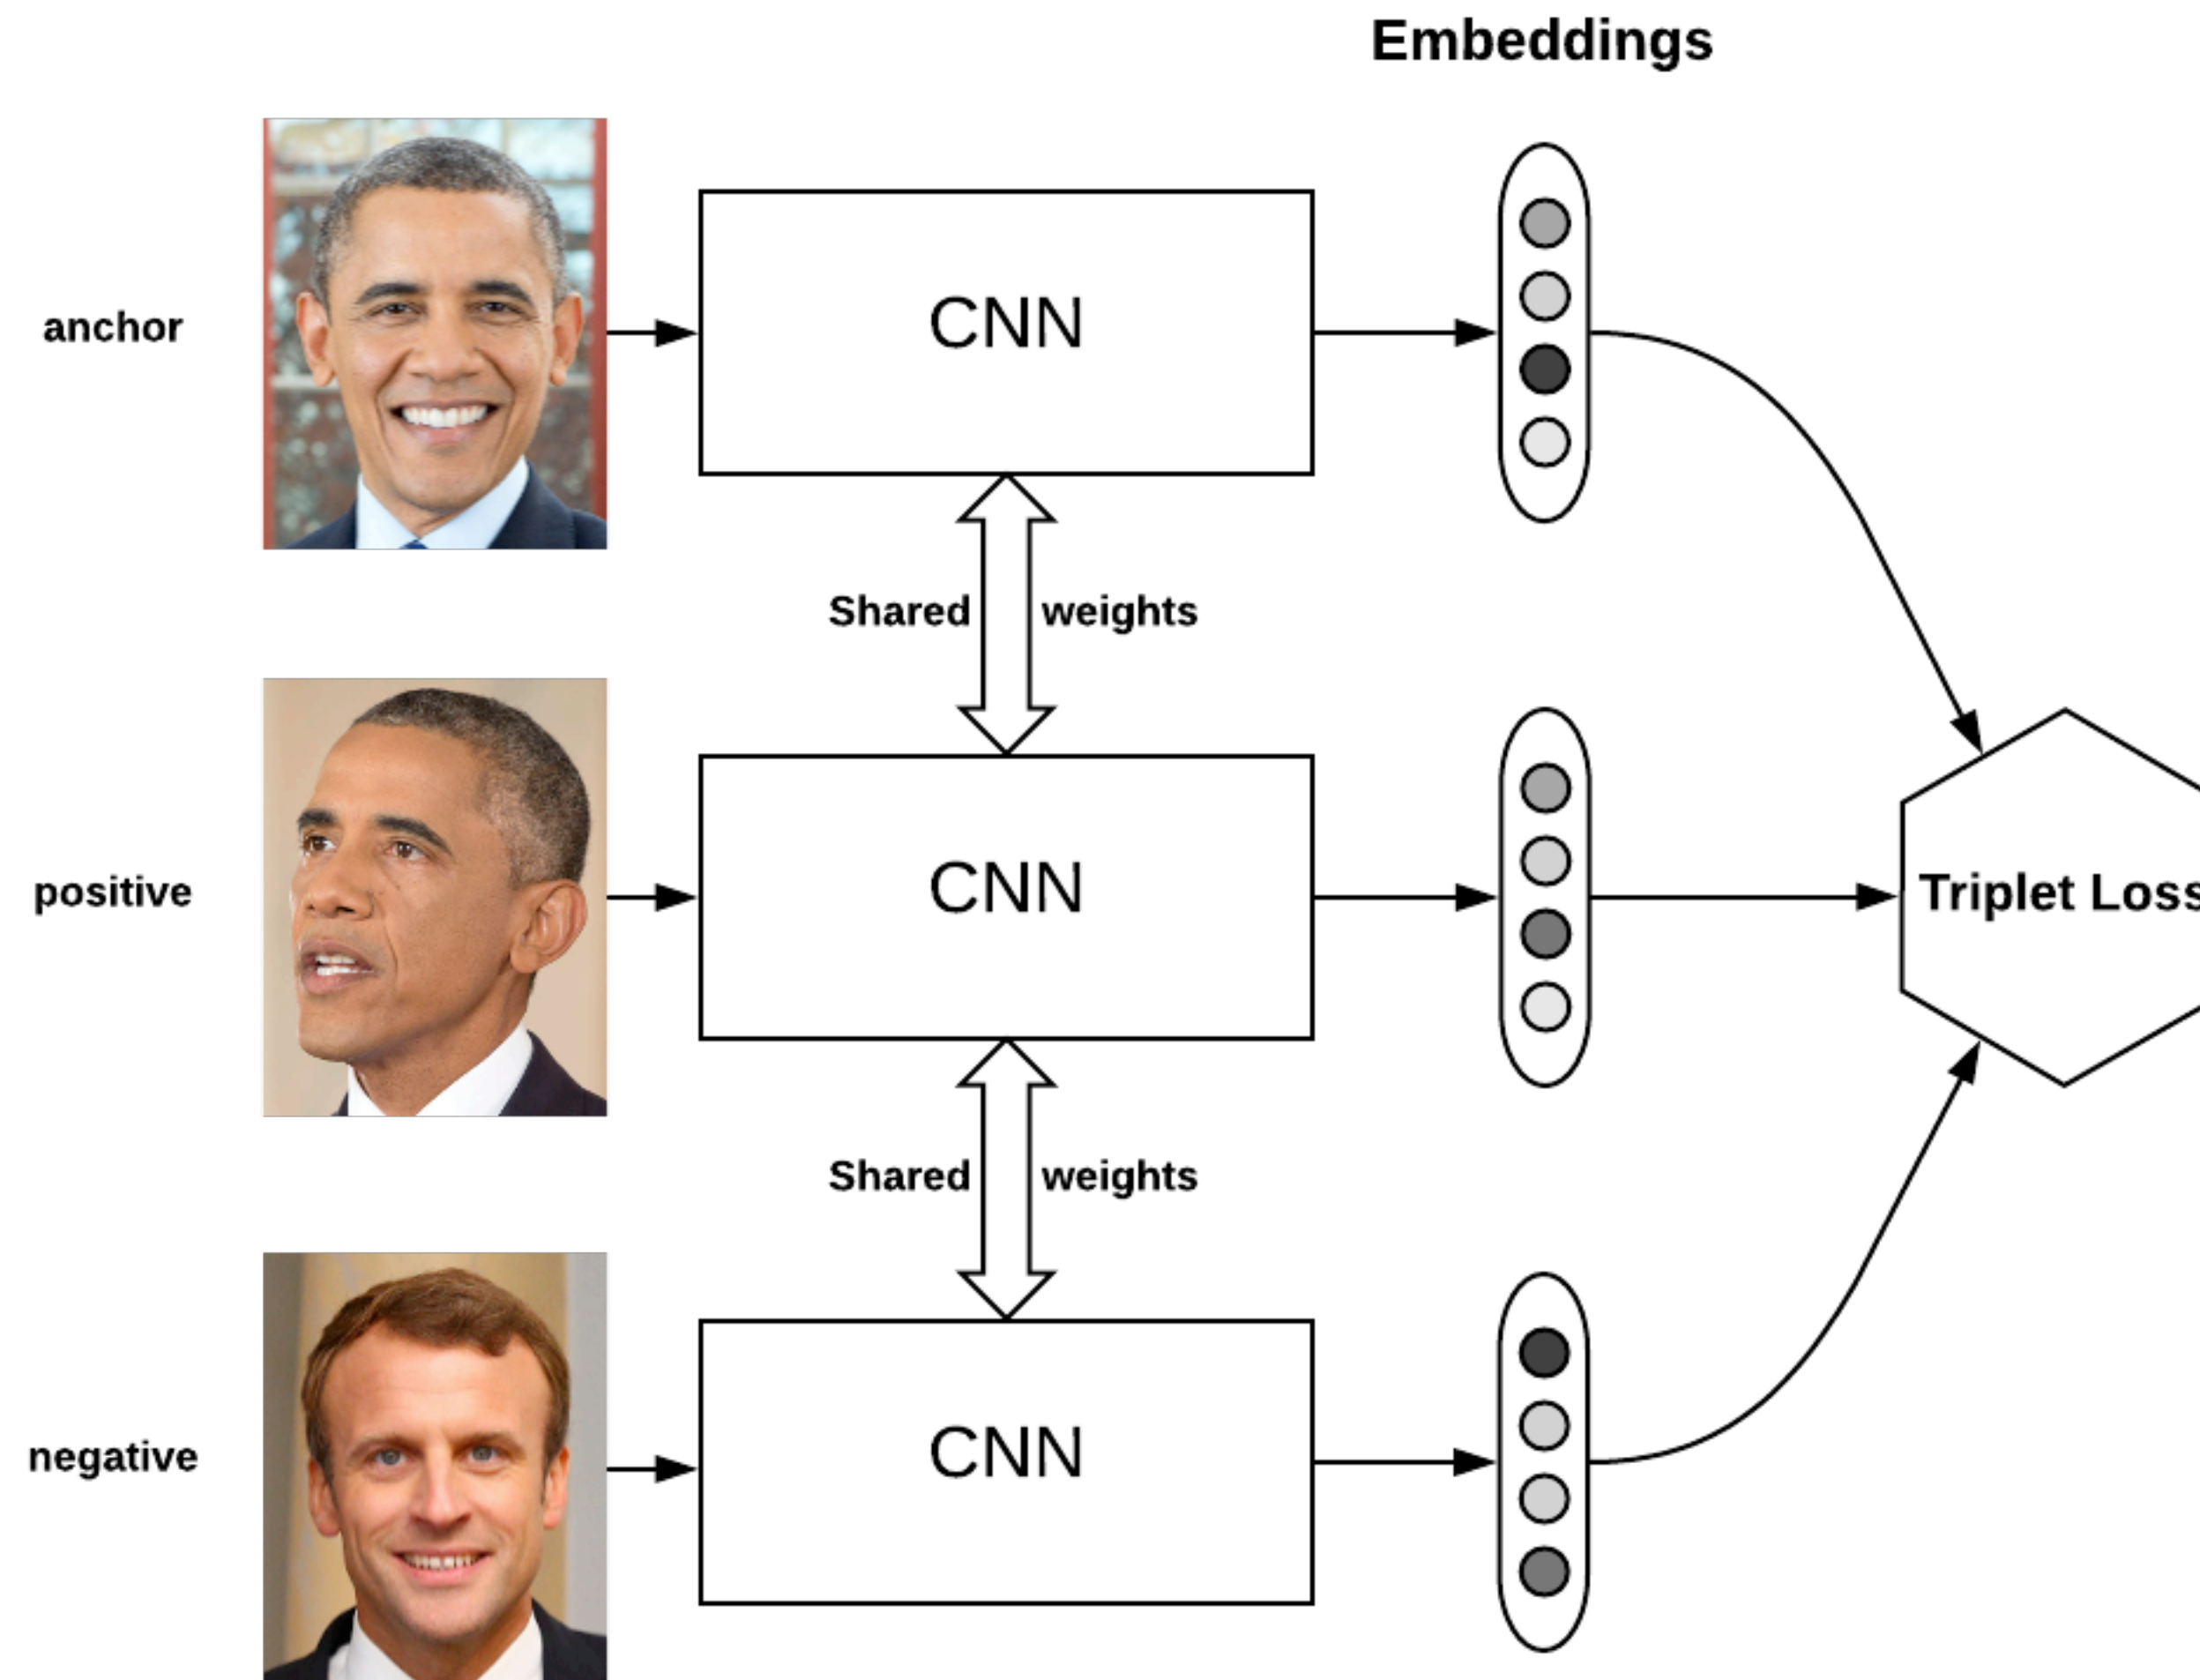
Figure 1. Two visually similar species from the iNat2017 dataset. Through close inspection, we can see that the ladybug on the left has *two* spots while the one on the right has *seven*.



Explicit feature matching



Triplet network



+ more data




Figure 8. **Pre-trained models performance evaluation.** We compare **DINOv2** (ViT-L/p14-518), **CLIP** (ViT-L/p14-336), and **MegaDescriptor-L** (Swin-L/p4-w12-384) on 29 selected datasets.

Data Imbalance and Fine-Grained Classification

6.S954 Computer Vision and Planetary Health