Perfectly Balanced: Improving Transfer and Robustness of Supervised Contrastive Learning

Dan Fu*, Mayee Chen*, Avanika Narayan, Michael Zhang, Zhao Song, Kayvon Fatahalian, Christopher Ré.

ICML 2022

Mayee F. Chen*, Daniel Y. Fu*, Avanika Narayan, Michael Zhang, Zhao Song, Kayvon Fatahalian, Christopher Ré.
Perfectly Balanced: Improving Transfer and Robustness of Supervised Contrastive Learning. ICML 2022.

Setting the Stage…

Setting the Stage…

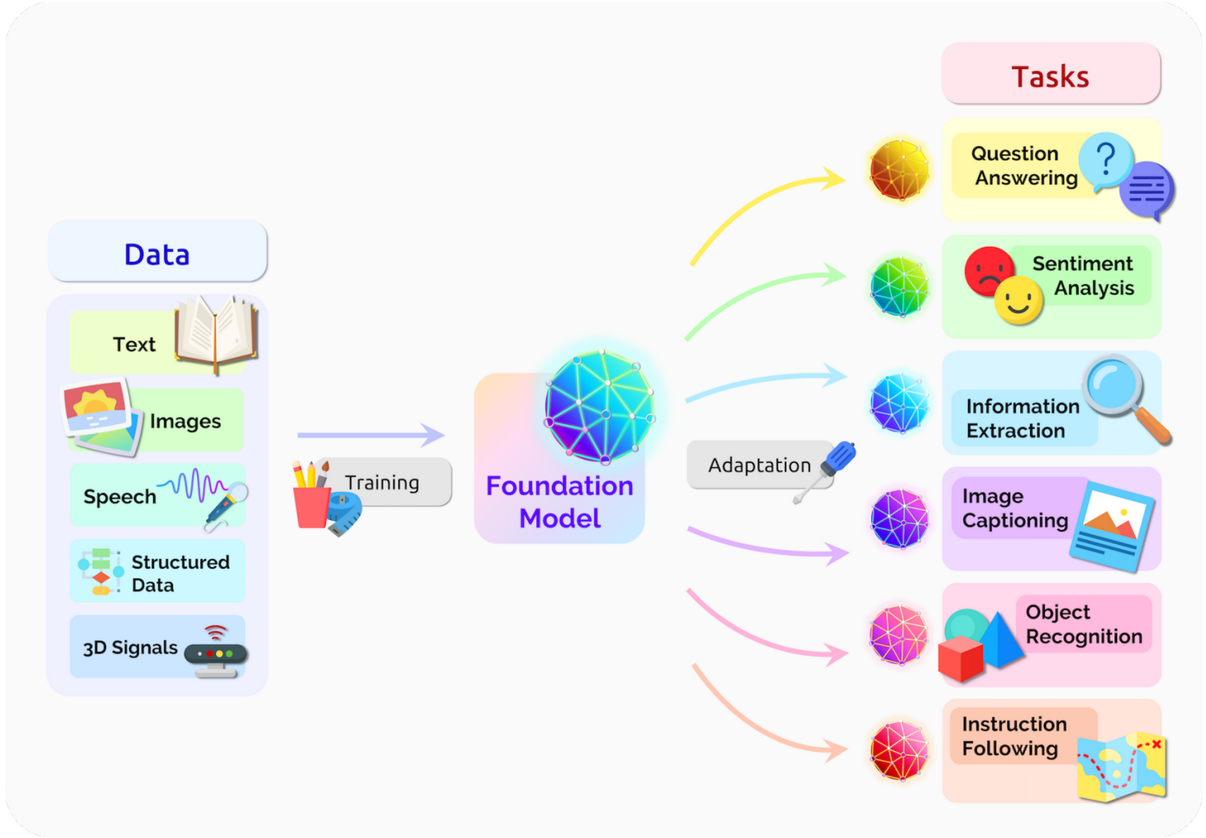

Transfer

Image Source: Bommasani et al, 2021.

Setting the Stage…

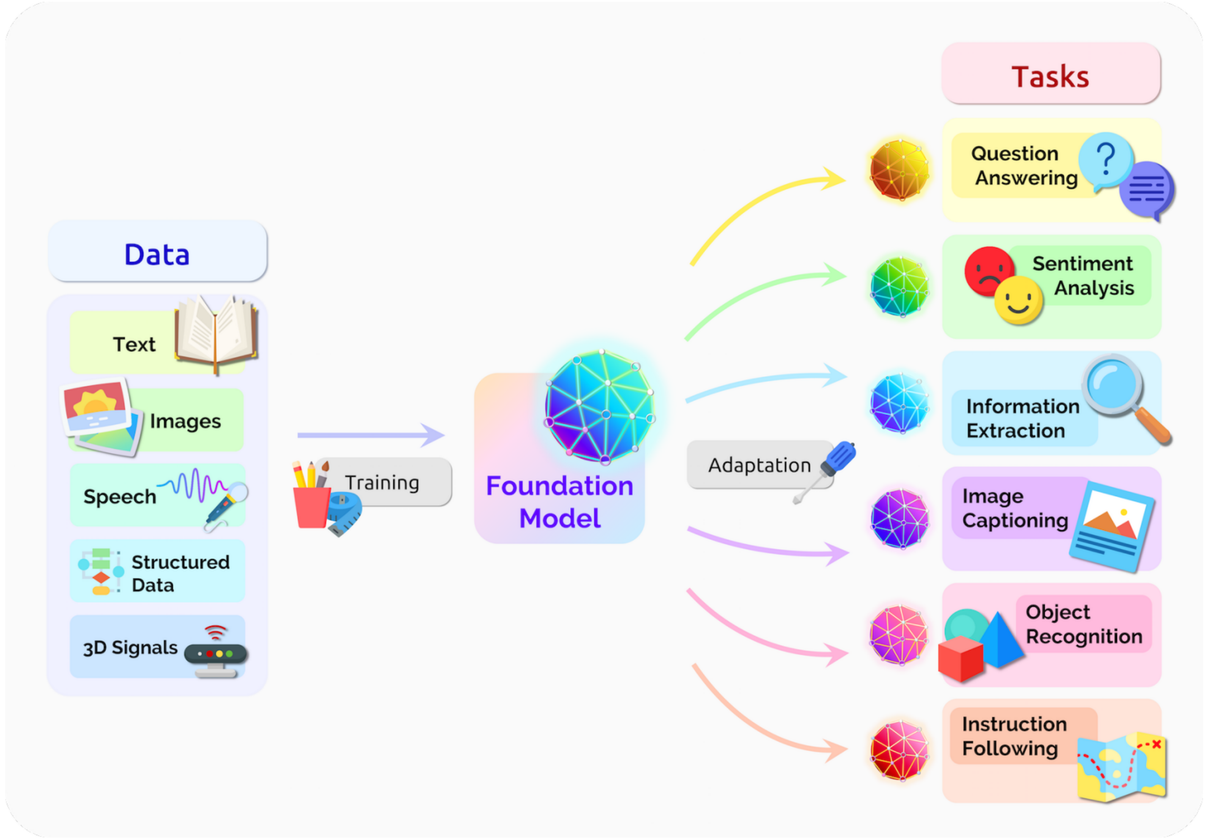

Transfer

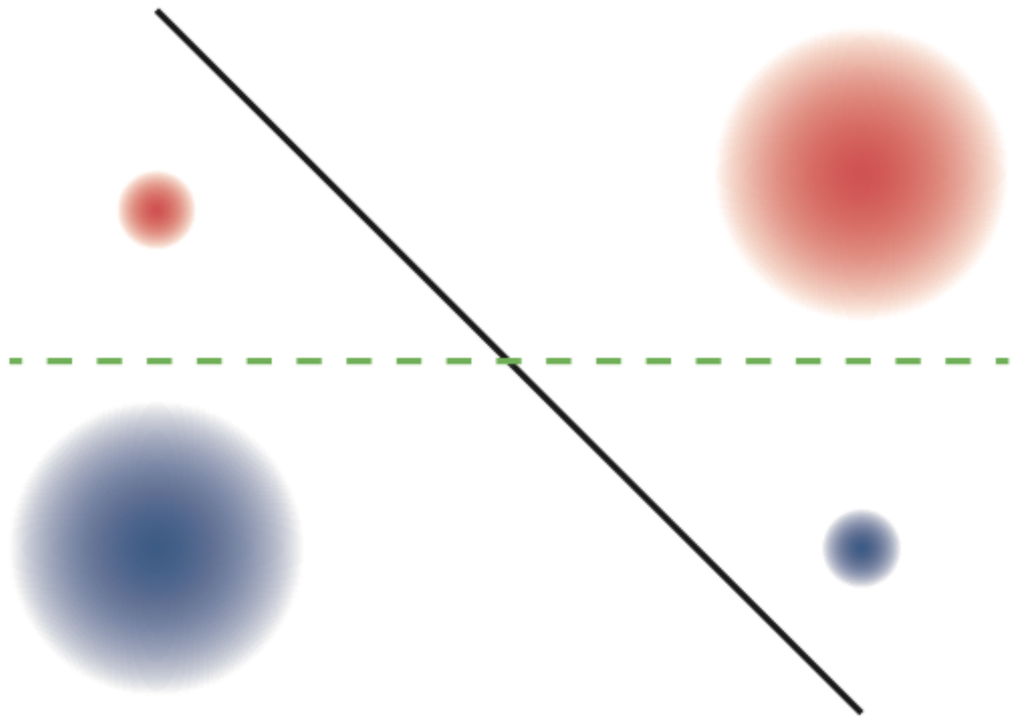

Robustness

Image Source: Bommasani et al, 2021.

Setting the Stage…

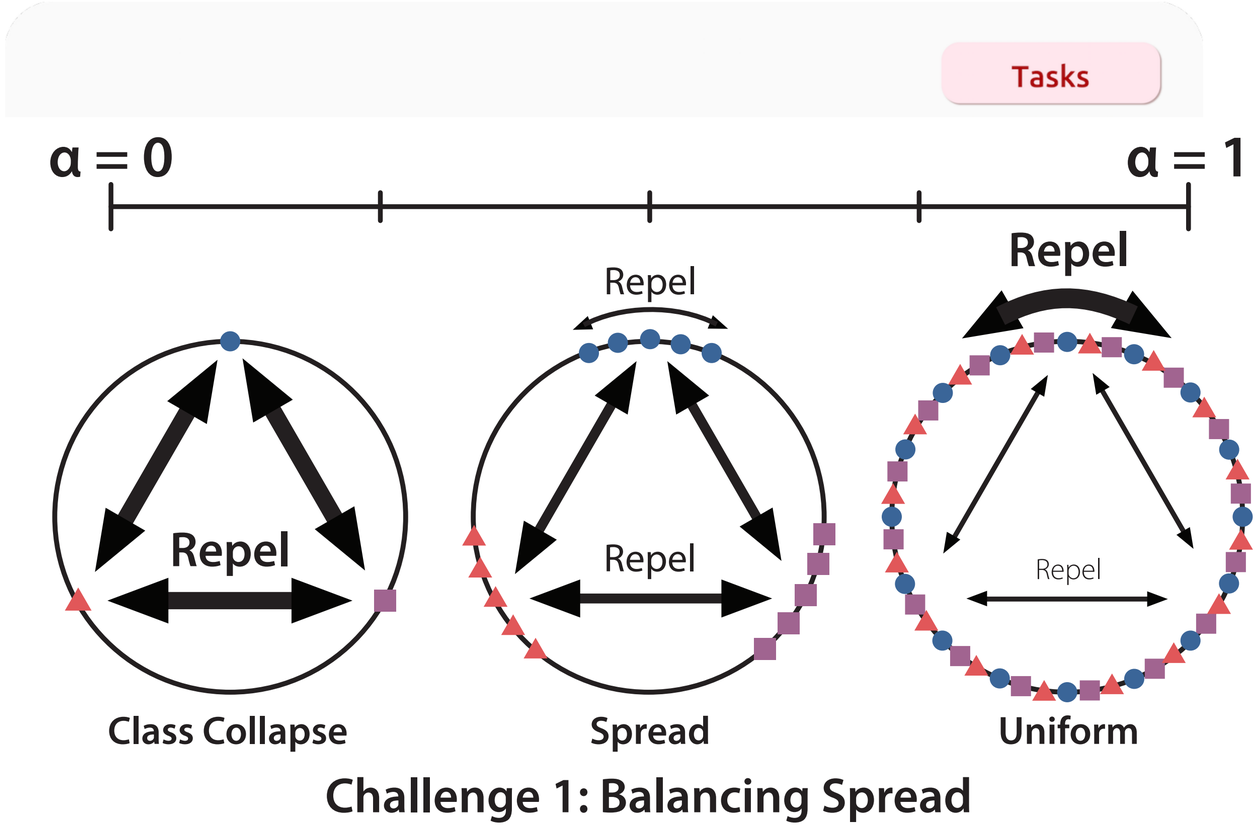

Challenge 1: Balancing Spread

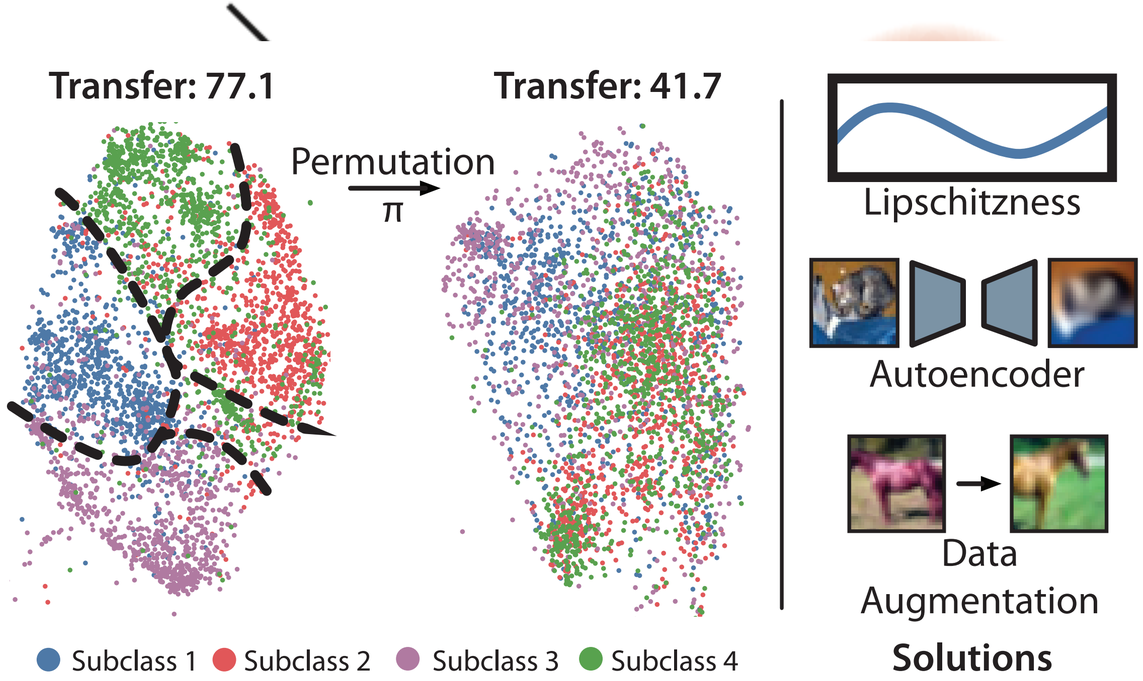

Challenge 2: Class-Fixing Permutation Invariance

Transfer

Robustness

Image Source: Bommasani et al, 2021.

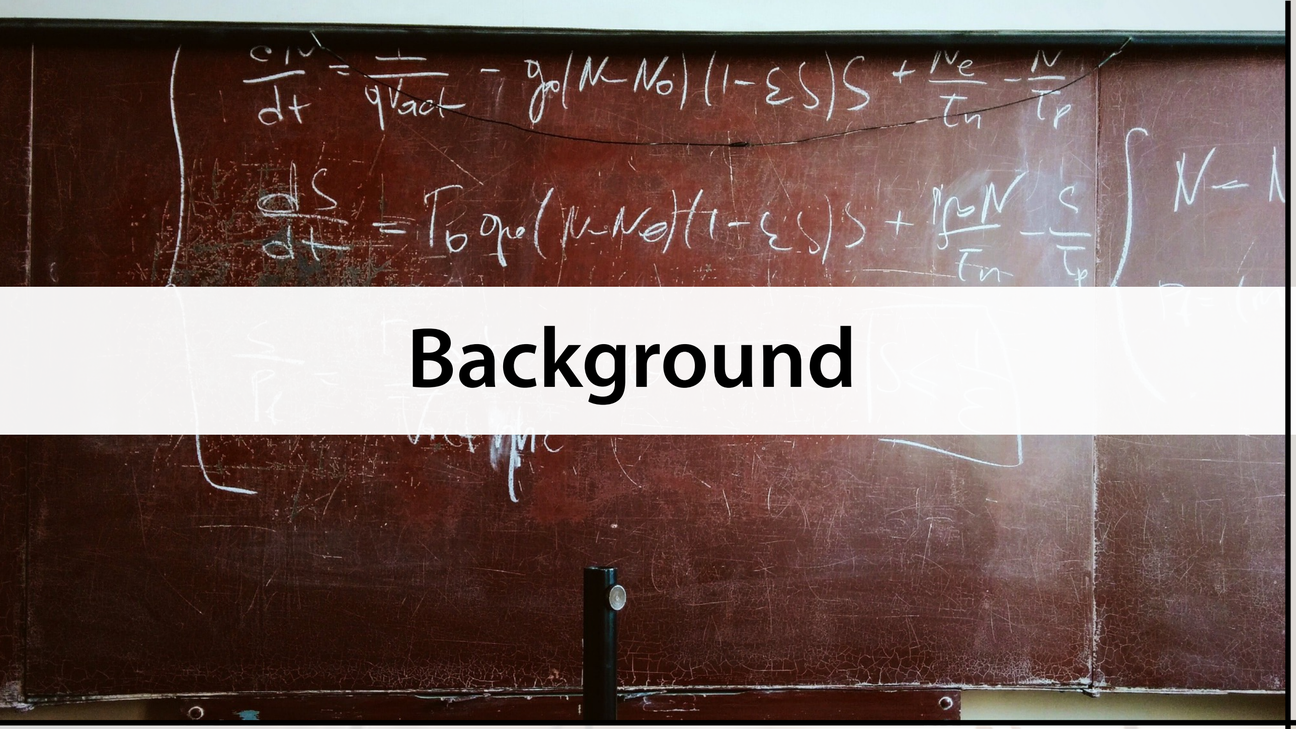

Background

Challenge 1: Balancing Spread

Challenge 2: Subclass Clustering

Experiments

Supervised Contrastive Learning

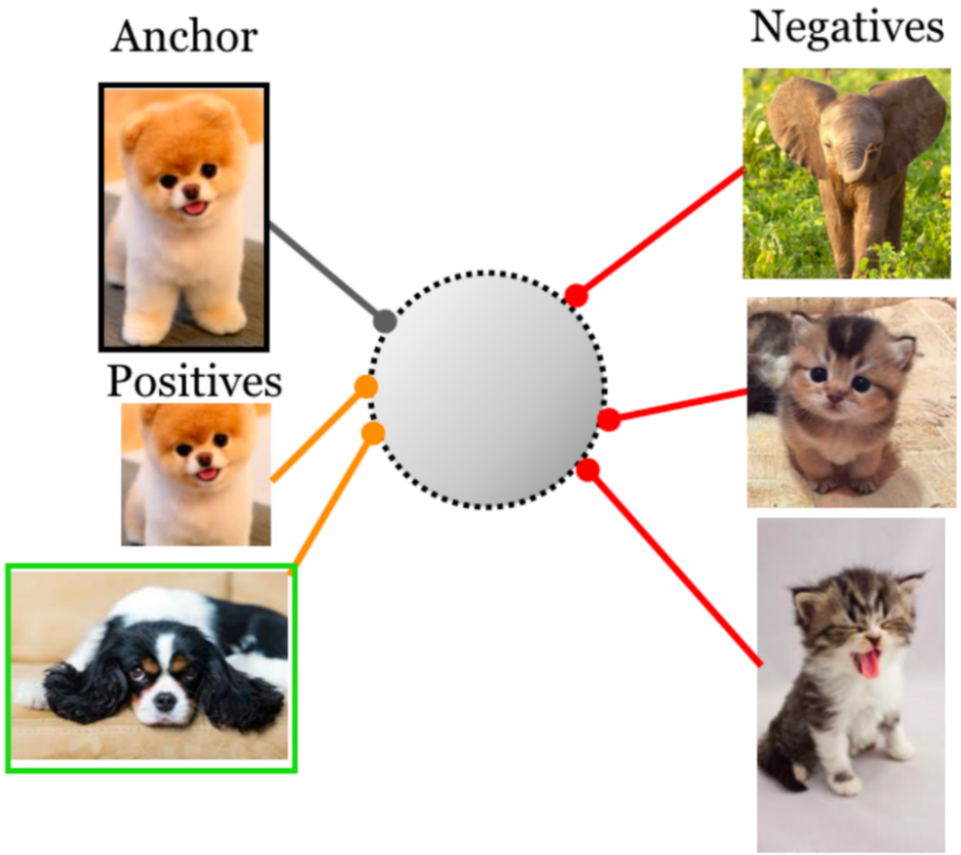

Supervised Contrastive (SupCon)

Image Source: https://ai.googleblog.com/2021/06/extending-contrastive-learning-to.html

Supervised Contrastive Learning

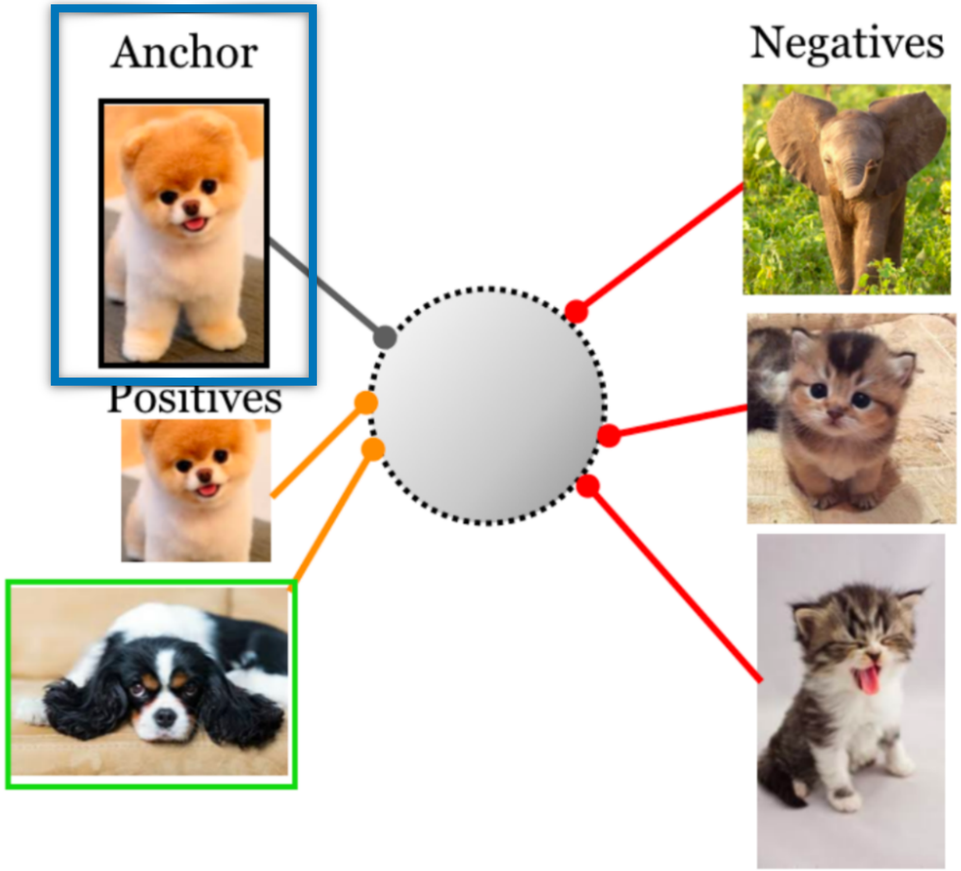

Supervised Contrastive (SupCon)

Image Source: https://ai.googleblog.com/2021/06/extending-contrastive-learning-to.html

Supervised Contrastive Learning

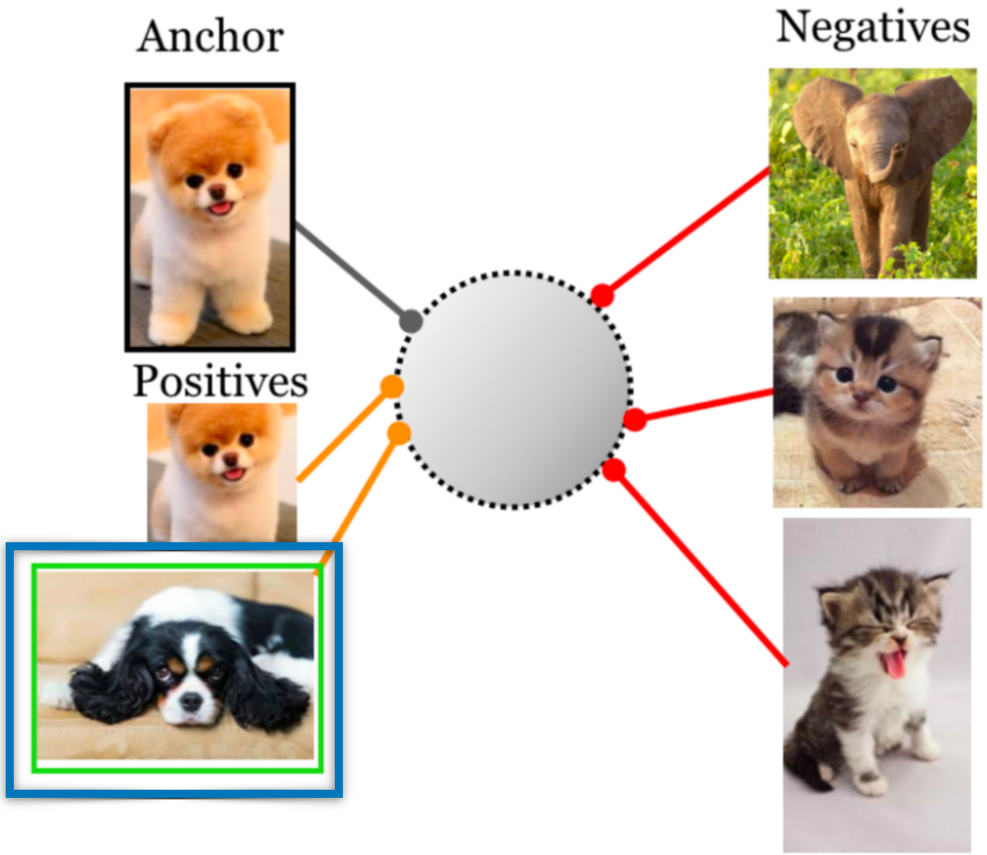

Supervised Contrastive (SupCon)

Image Source: https://ai.googleblog.com/2021/06/extending-contrastive-learning-to.html

Supervised Contrastive Learning

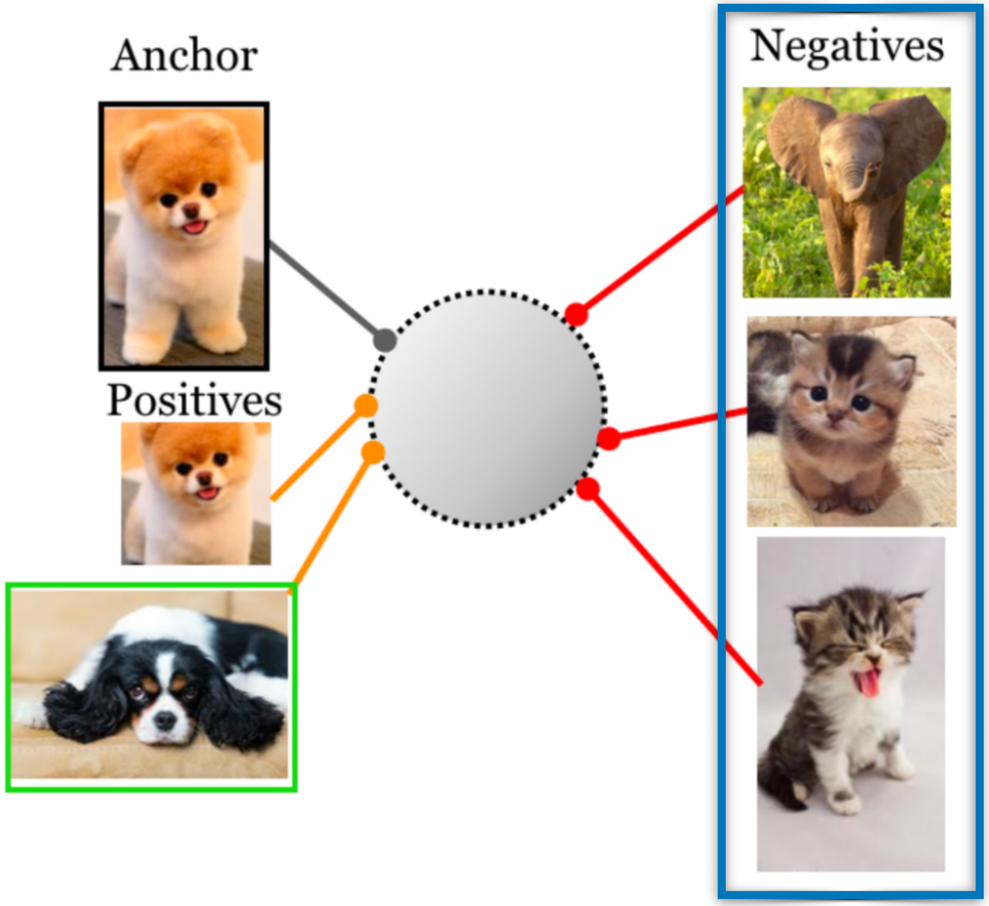

Supervised Contrastive (SupCon)

Image Source: https://ai.googleblog.com/2021/06/extending-contrastive-learning-to.html

Supervised Contrastive Learning

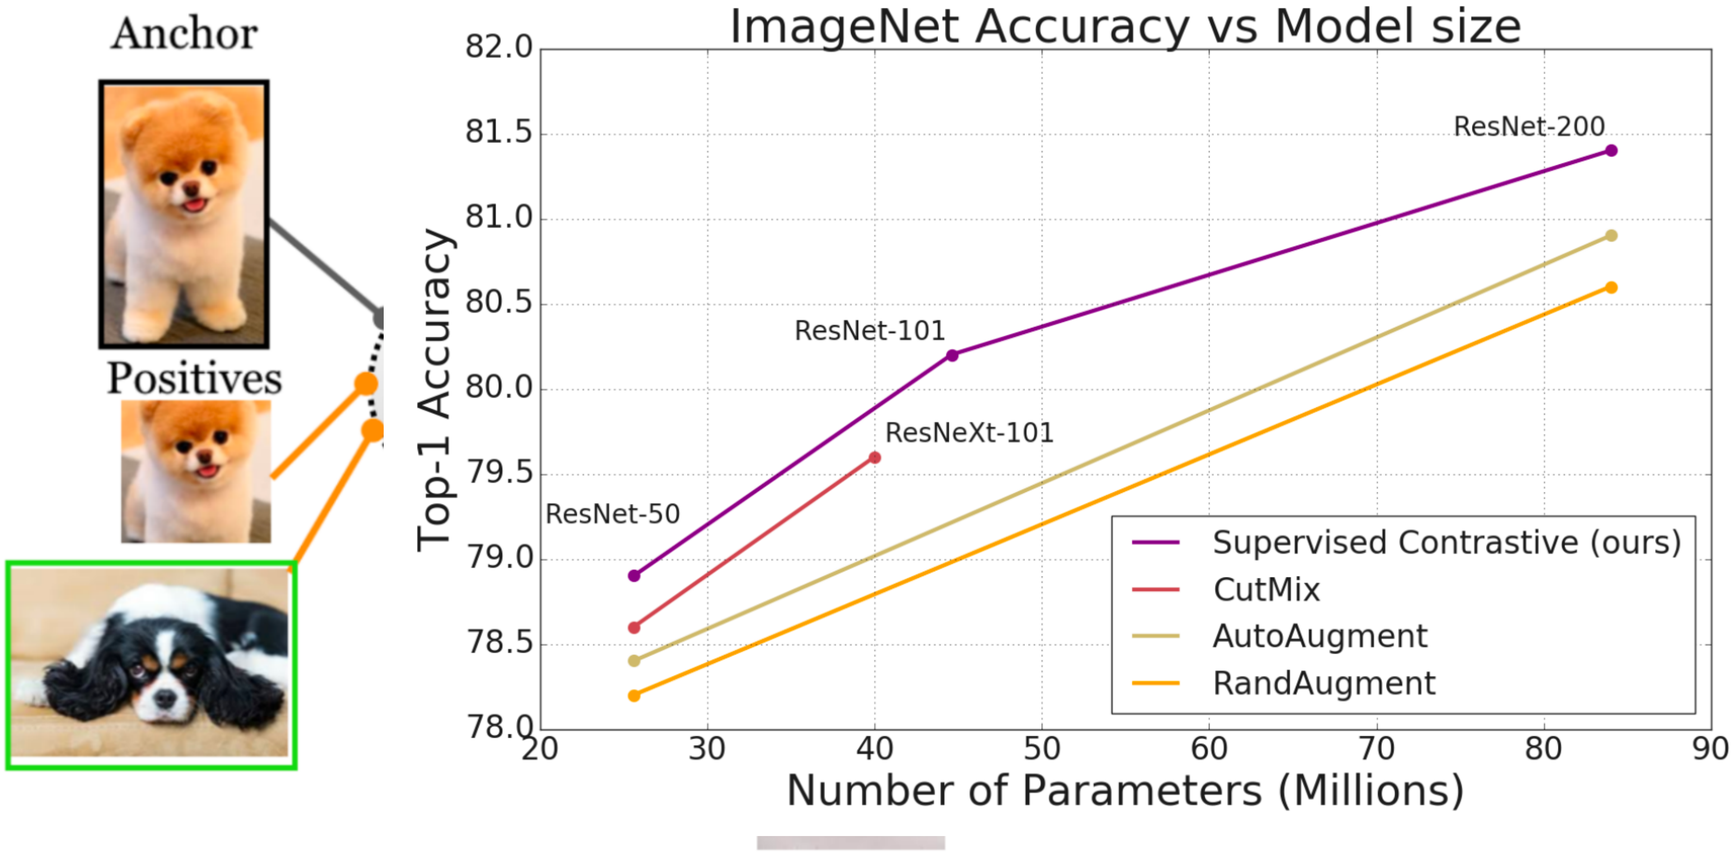

Supervised Contrastive (SupCon)

Image Source: https://ai.googleblog.com/2021/06/extending-contrastive-learning-to.html

Supervised Contrastive Learning

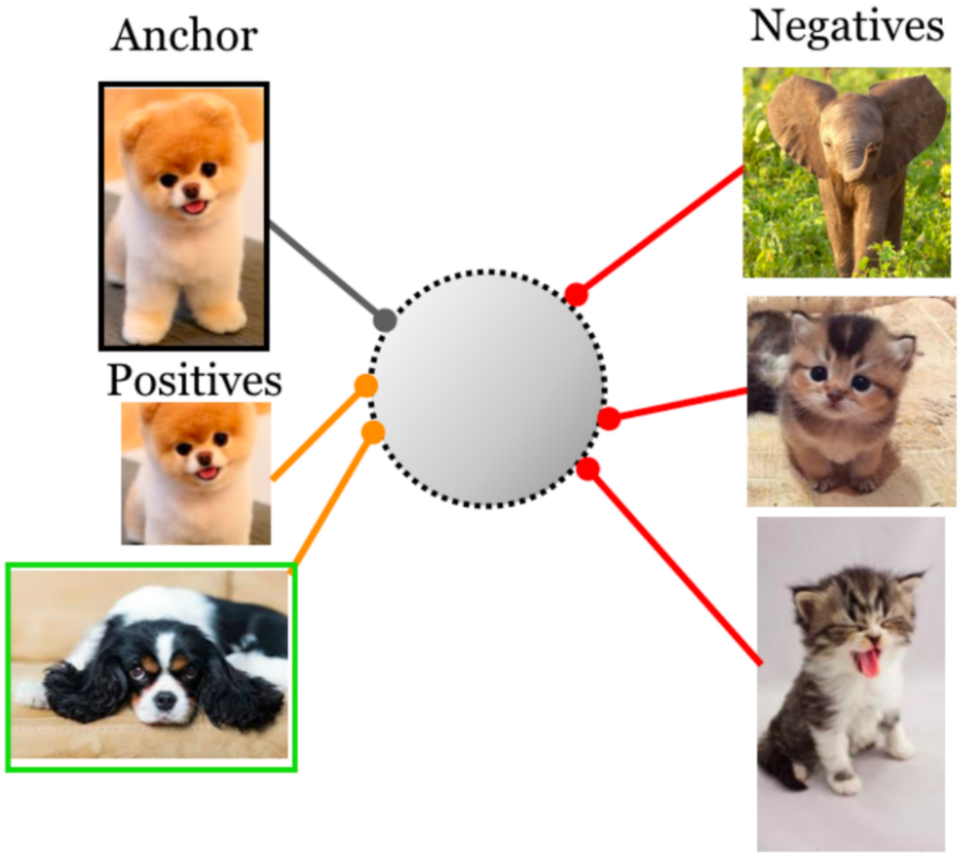

Supervised Contrastive (SupCon)

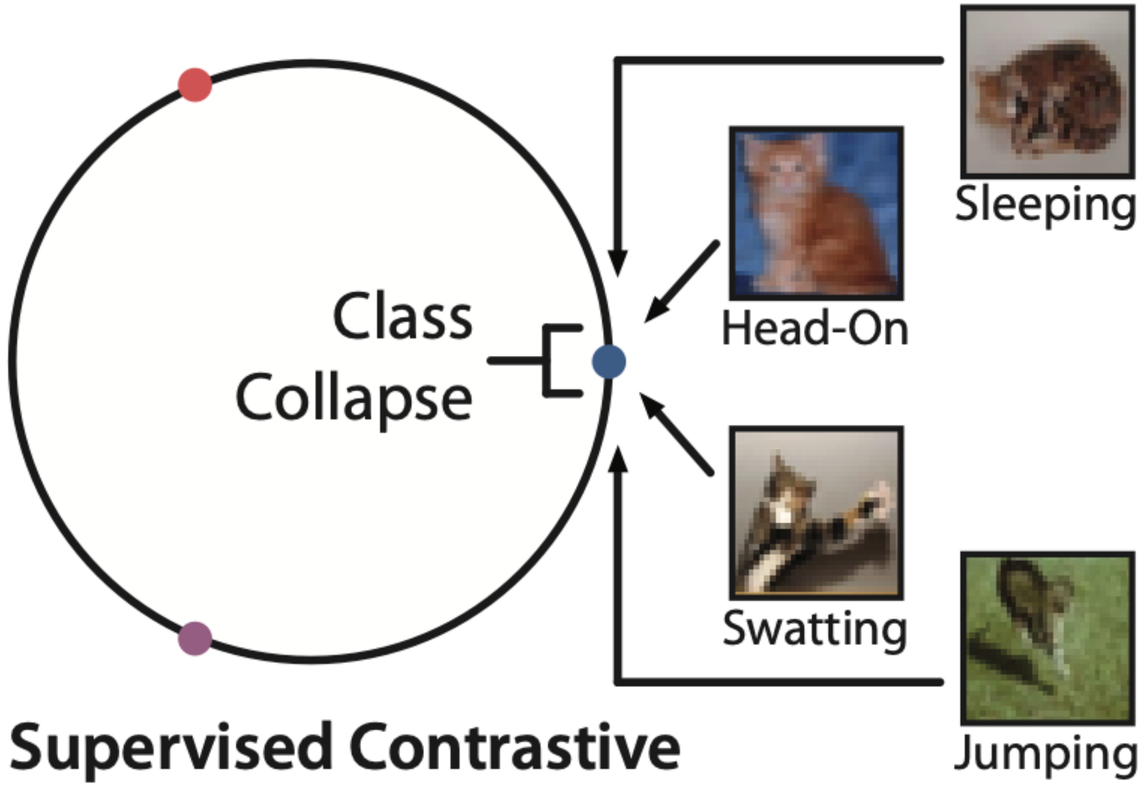

…at the cost of class collapse

Image Source: https://ai.googleblog.com/2021/06/extending-contrastive-learning-to.html

Background

Challenge 1: Balancing Spread

Challenge 2: Subclass Clustering

Experiments

Challenge 1: Balancing Spread

SupCon → Class-Conditional InfoNCE

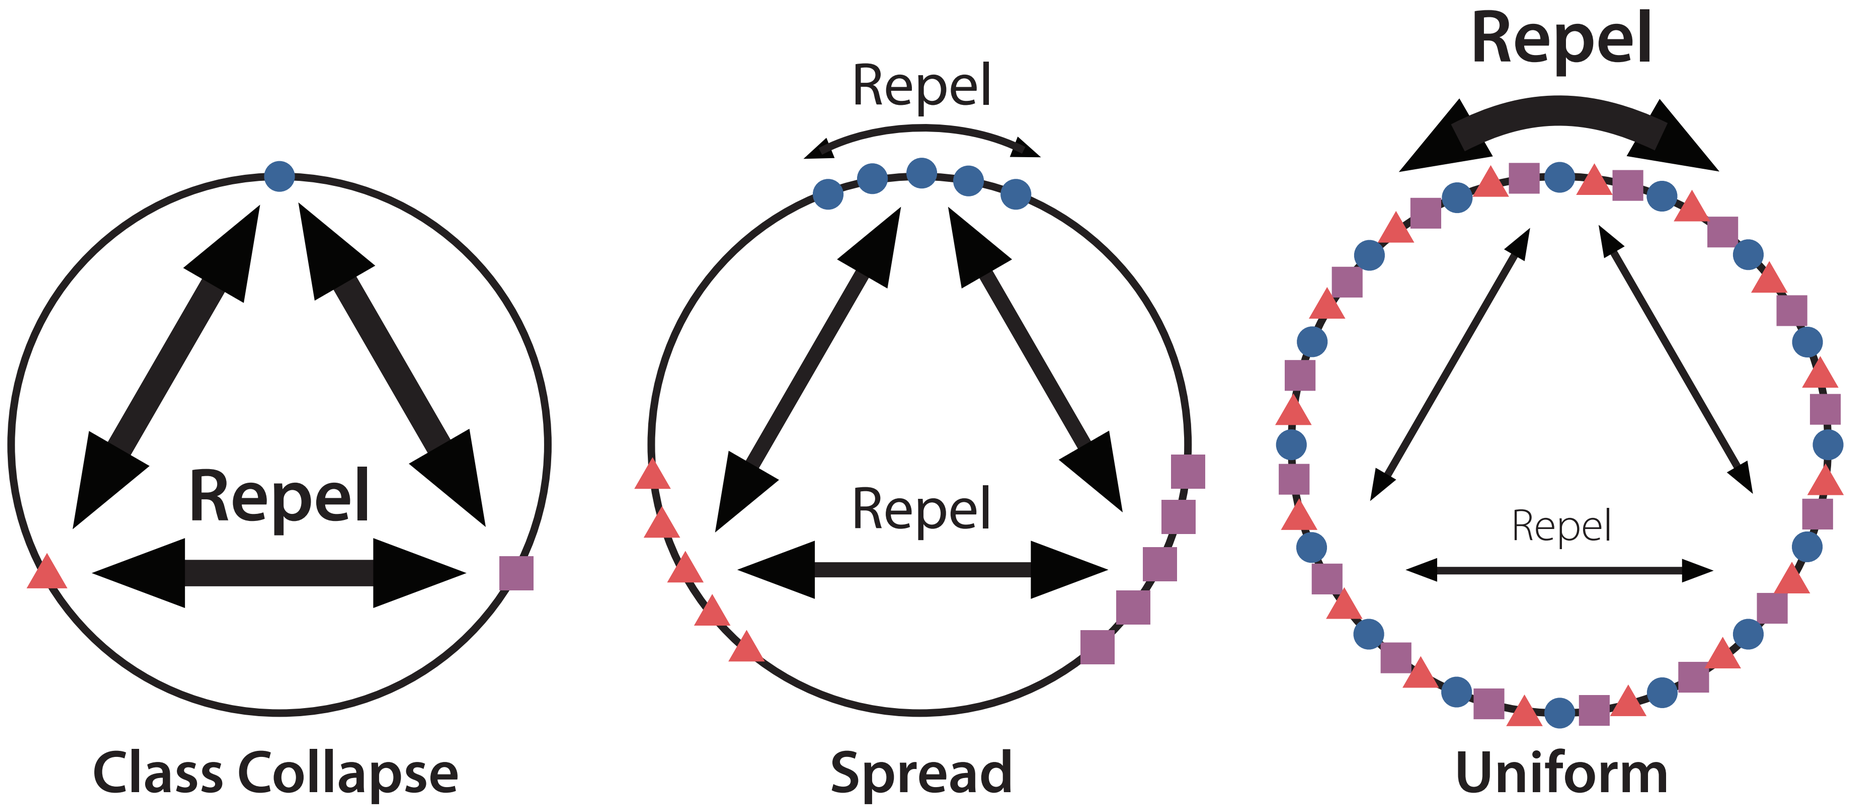

Adding a Weighted Class-Conditional InfoNCE Loss

Lspread = (1-α) Supcon + α Class-Conditional InfoNCE

collapse **uniform per class**

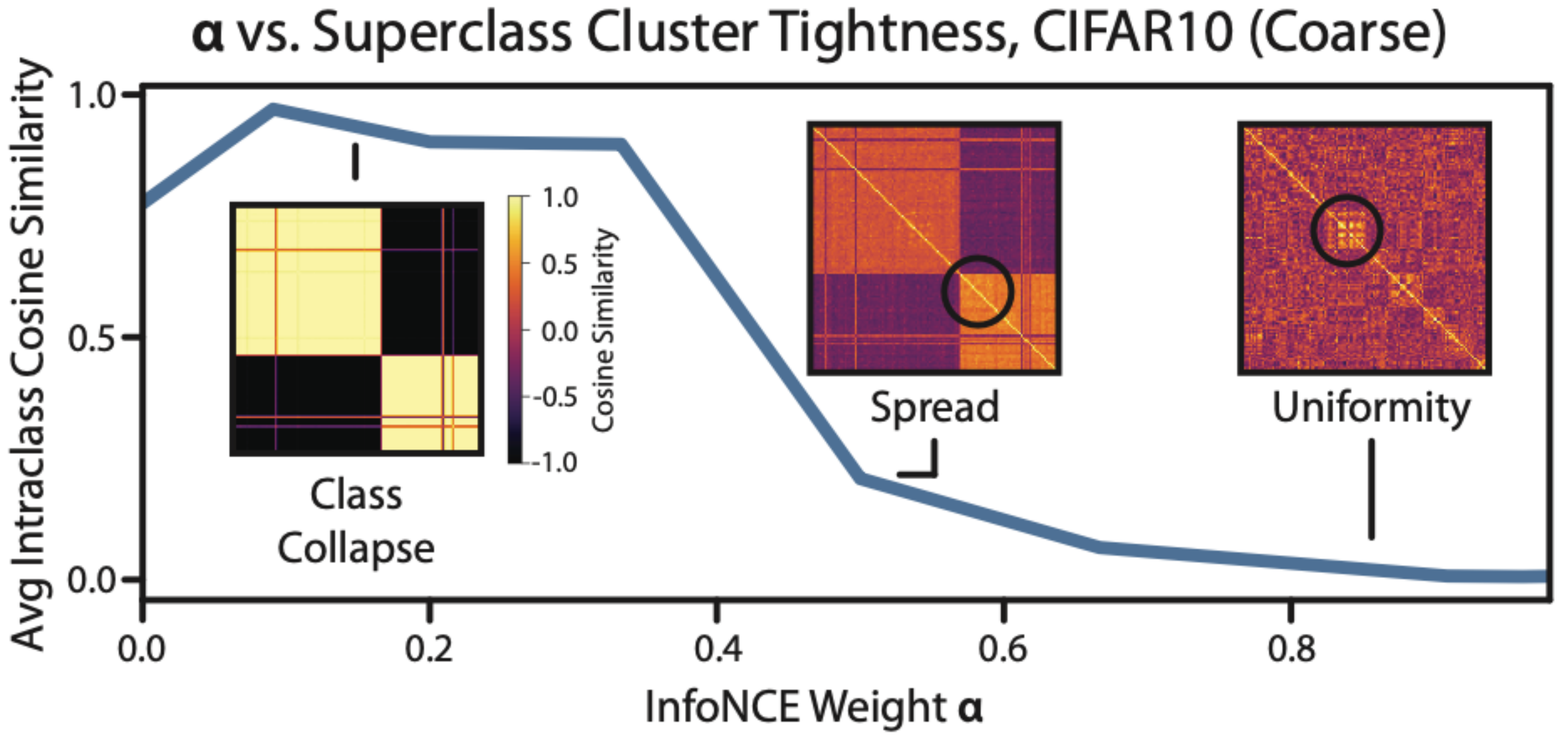

Background

Challenge 1: Balancing Spread

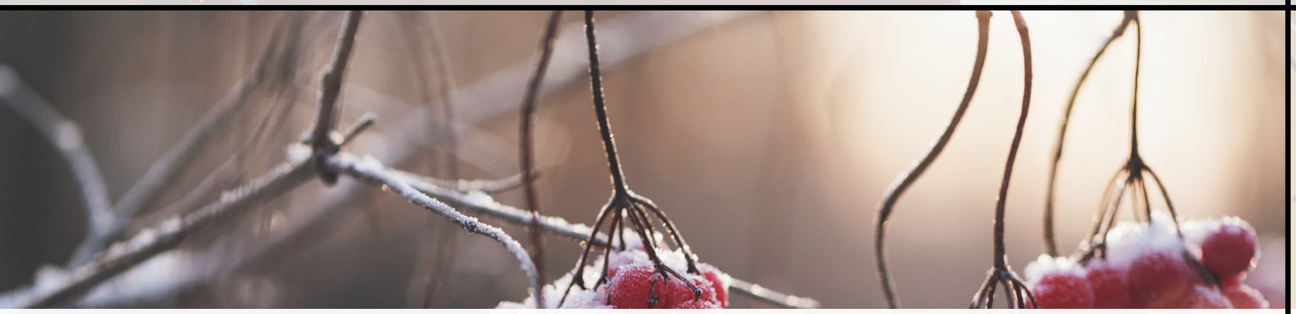

Challenge 2: Subclass Clustering

Experiments

Challenge 2: Subclass Clustering

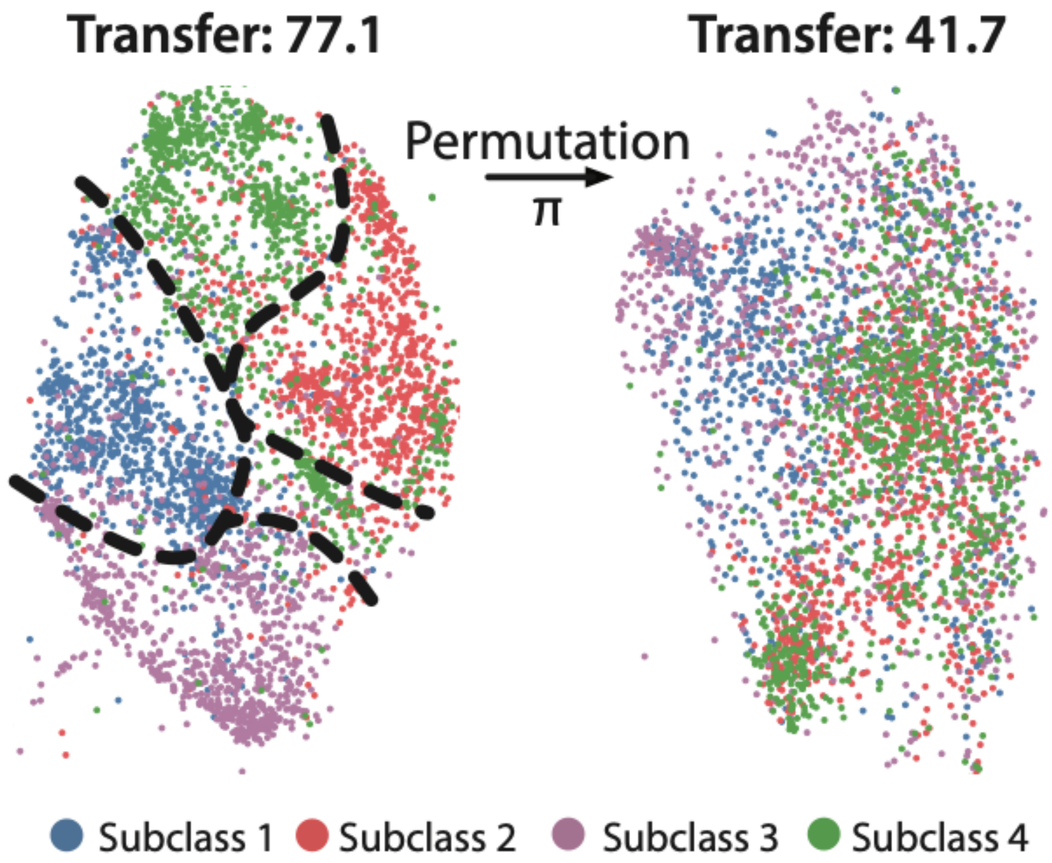

Superclass structure is not enough, we also need subclass clustering.

Breaking Permutation Invariance

How do we break invariance for good downstream transfer?

Breaking Permutation Invariance

How do we break invariance for good downstream transfer?

Unrealistic for modern neural networks

Breaking Permutation Invariance

How do we break invariance for good downstream transfer?

Unrealistic for modern neural networks

"Reverse Lipschitz" assumption only on the decoder

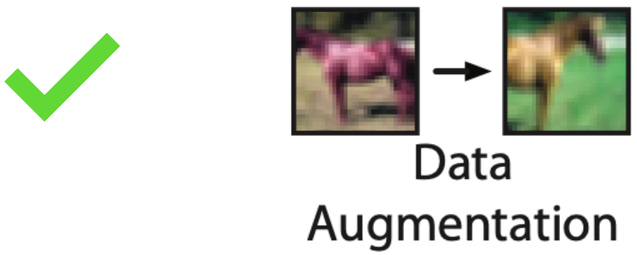

Lipschitz assumption on augmented data, not on all data

Background

Challenge 1: Balancing Spread

Challenge 2: Subclass Clustering

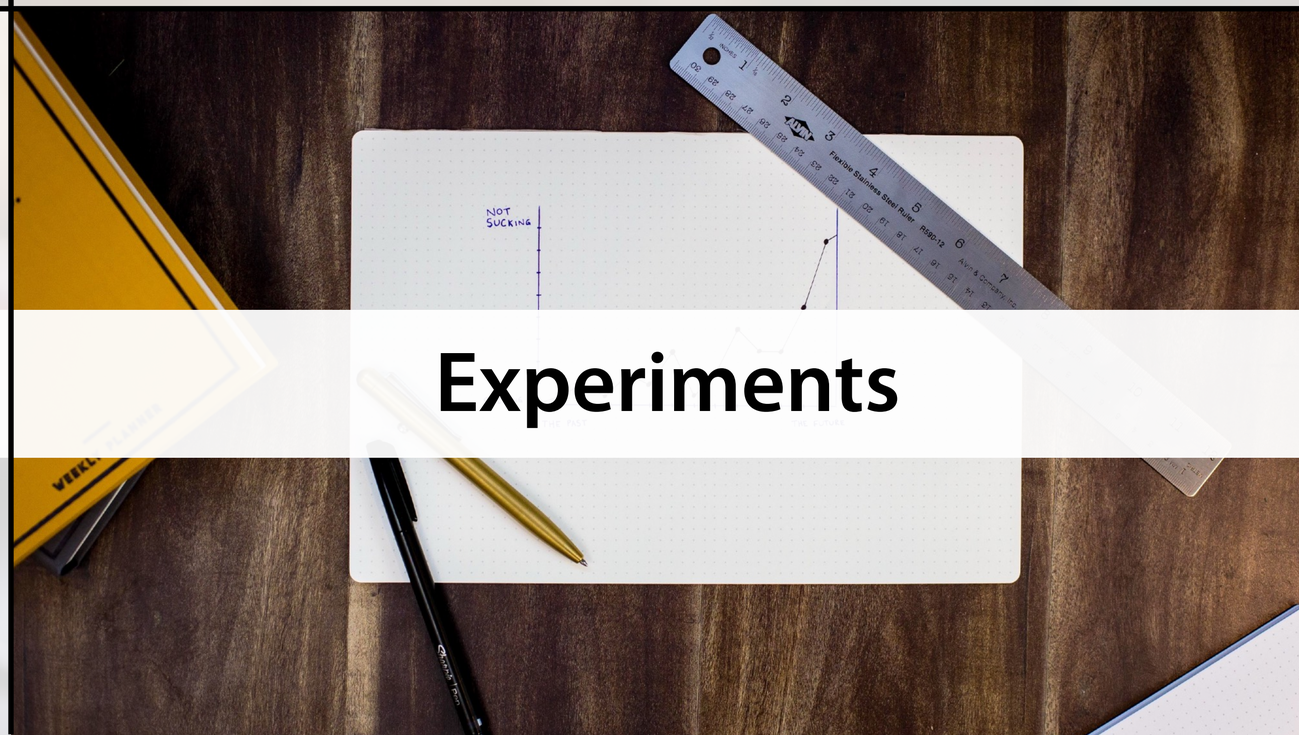

Experiments

Thanos 

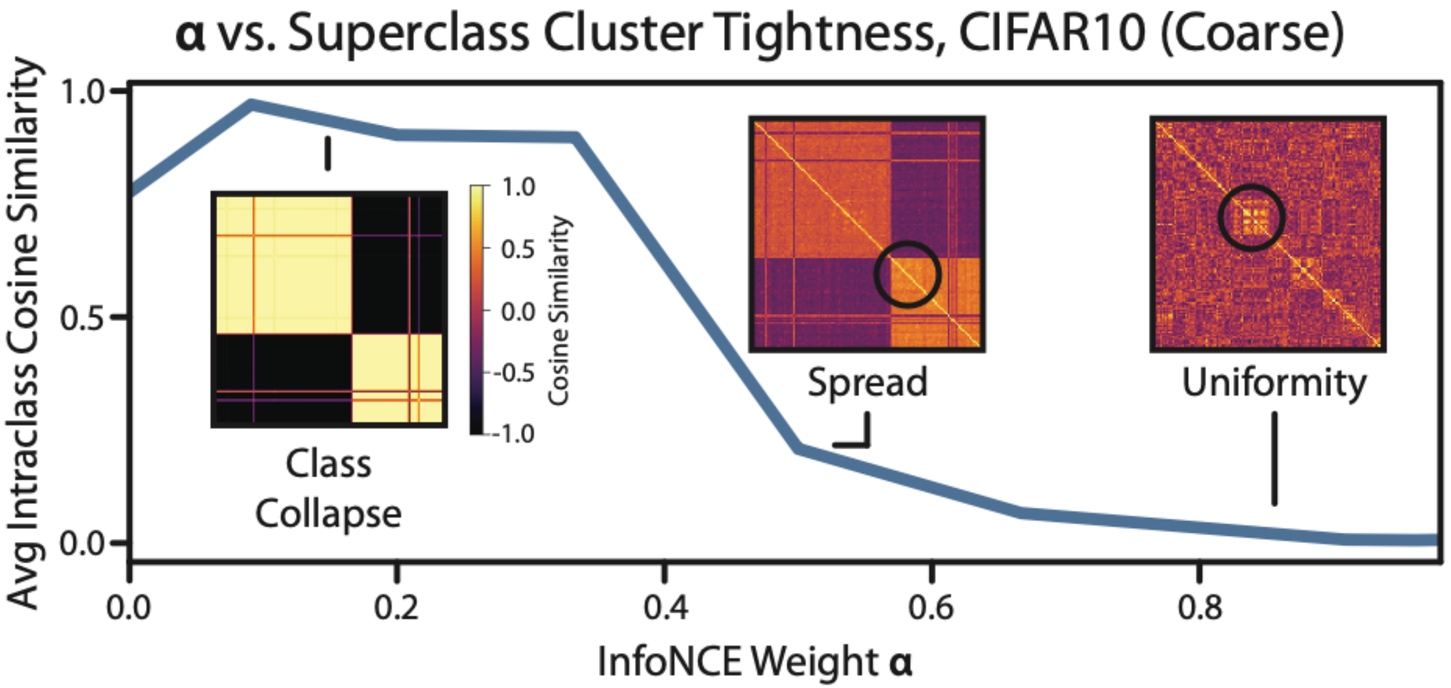

+

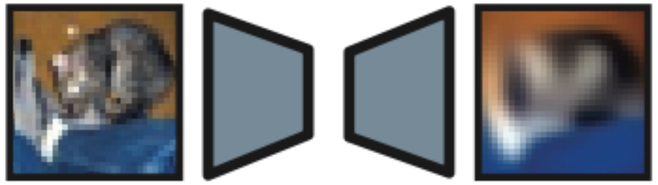

Autoencoder

Data Augmentation

Better Transfer (11.1 points on average)

Table 3. Coarse-to-fine transfer learning performance.

	Method	CIFAR10	CIFAR100	CIFAR100-U	MNIST	TinyImageNet
Baselines	InfoNCE (Chen et al., 2020a)	77.6 ± 0.1	60.5 ± 0.1	56.4 ± 0.3	98.4 ± 0.1	44.9 ± 0.1
	SupCon (Khosla et al., 2020)	51.8 ± 1.2	56.1 ± 0.1	49.8 ± 0.3	95.4 ± 0.1	43.9 ± 0.1
	SupCon + InfoNCE (Islam et al., 2021)	77.6 ± 0.1	55.7 ± 0.1	48.0 ± 0.2	98.6 ± 0.1	46.1 ± 0.1
Ours	cAuto	71.4 ± 0.1	62.9 ± 0.1	58.7 ± 0.5	98.7 ± 0.1	47.1 ± 0.1
	SupCon + cNCE (L_{spread})	77.1 ± 0.1	58.7 ± 0.2	53.5 ± 0.4	98.5 ± 0.1	45.8 ± 0.1
	SupCon + cAuto	71.7 ± 0.1	63.8 ± 0.6	$\mathbf{59.8 \pm 0.3}$	98.7 ± 0.1	49.3 ± 0.1
	SupCon + cNCE + cAuto (THANOS)	$\mathbf{79.1 \pm 0.2}$	$\mathbf{65.0 \pm 0.2}$	59.7 ± 0.3	$\mathbf{99.0 \pm 0.1}$	$\mathbf{49.6 \pm 0.1}$

Better Transfer (11.1 points on average)

Table 3. Coarse-to-fine transfer learning performance.

	Method	CIFAR10	CIFAR100	CIFAR100-U	MNIST	TinyImageNet
Baselines	InfoNCE (Chen et al., 2020a)	77.6 ± 0.1	60.5 ± 0.1	56.4 ± 0.3	98.4 ± 0.1	44.9 ± 0.1
	SupCon (Khosla et al., 2020)	51.8 ± 1.2	56.1 ± 0.1	49.8 ± 0.3	95.4 ± 0.1	43.9 ± 0.1
	SupCon + InfoNCE (Islam et al., 2021)	77.6 ± 0.1	55.7 ± 0.1	48.0 ± 0.2	98.6 ± 0.1	46.1 ± 0.1
Ours	cAuto	71.4 ± 0.1	62.9 ± 0.1	58.7 ± 0.5	98.7 ± 0.1	47.1 ± 0.1
	SupCon + cNCE (L_{spread})	77.1 ± 0.1	58.7 ± 0.2	53.5 ± 0.4	98.5 ± 0.1	45.8 ± 0.1
	SupCon + cAuto	71.7 ± 0.1	63.8 ± 0.6	$\mathbf{59.8 \pm 0.3}$	98.7 ± 0.1	49.3 ± 0.1
	SupCon + cNCE + cAuto (THANOS)	$\mathbf{79.1 \pm 0.2}$	$\mathbf{65.0 \pm 0.2}$	59.7 ± 0.3	$\mathbf{99.0 \pm 0.1}$	$\mathbf{49.6 \pm 0.1}$

…and Robustness! (4.7 points on average, 11.5 on CelebA)

Table 4. Unsupervised subclass recovery (top, F1), and worst-group performance (AUROC for ISIC, Acc for others).

Method	Group Labels	Waterbirds	ISIC	CelebA
		Sub-Group Recovery		
Sohoni et al. (2020)	✗	56.3	74.0	24.2
SupCon	✗	47.1	92.5	19.4
THANOS	✗	**59.0**	**93.8**	**24.8**
		Worst-Group Robustness		
Sohoni et al. (2020)	✗	88.4	92.0	55.0
JTT (Liu et al., 2021)	✗	83.8	91.8	77.9
SupCon	✗	86.8	**93.3**	66.1
THANOS	✗	**88.6**	92.6	**89.4**
GroupDRO	✓	90.7	92.3	88.9

Thank You!

Poster Session: Wednesday, July 20, 6:30-8:30 PM

Contact:
Dan Fu (danfu@cs.stanford.edu, @realDanFu)
Mayee Chen (mfchen@stanford.edu, @MayeeChen)

arXiv: https://arxiv.org/abs/2204.07596
Mayee F. Chen*, Daniel Y. Fu*, Avanika Narayan, Michael Zhang, Zhao Song, Kayvon Fatahalian, Christopher Ré. Perfectly Balanced: Improving Transfer and Robustness of Supervised Contrastive Learning. *ICML 2022*.

AIBSD workshop paper: https://aibsdworkshop.github.io/2022/index.html
Daniel Y. Fu*, Mayee F. Chen*, Michael Zhang, Kayvon Fatahalian, Christopher Ré. The Details Matter: Preventing Class Collapse in Supervised Contrastive Learning. *AIBSD @ AAAI 2022*. **Best Paper.**

Blog: https://hazyresearch.stanford.edu/blog/2022-04-19-contrastive-2
Code: https://github.com/HazyResearch/thanos-code

Dan Fu

Mayee Chen

Avanika Narayan

Michael Zhang