Few-shot Learning by Statistical Methods (Part 2)

Dr. Yanwei Fu

School of Data Science, Fudan University

Homepage: http://yanweifu.github.io

CVPR 2023 Tutorial: Few-shot Learning from Meta-Learning, Statistical Understanding to Applications

https://fsl-fudan.github.io

大数据学院
School of Data Science

Few-Shot Learning
Revisited

Machine Learning by Cost

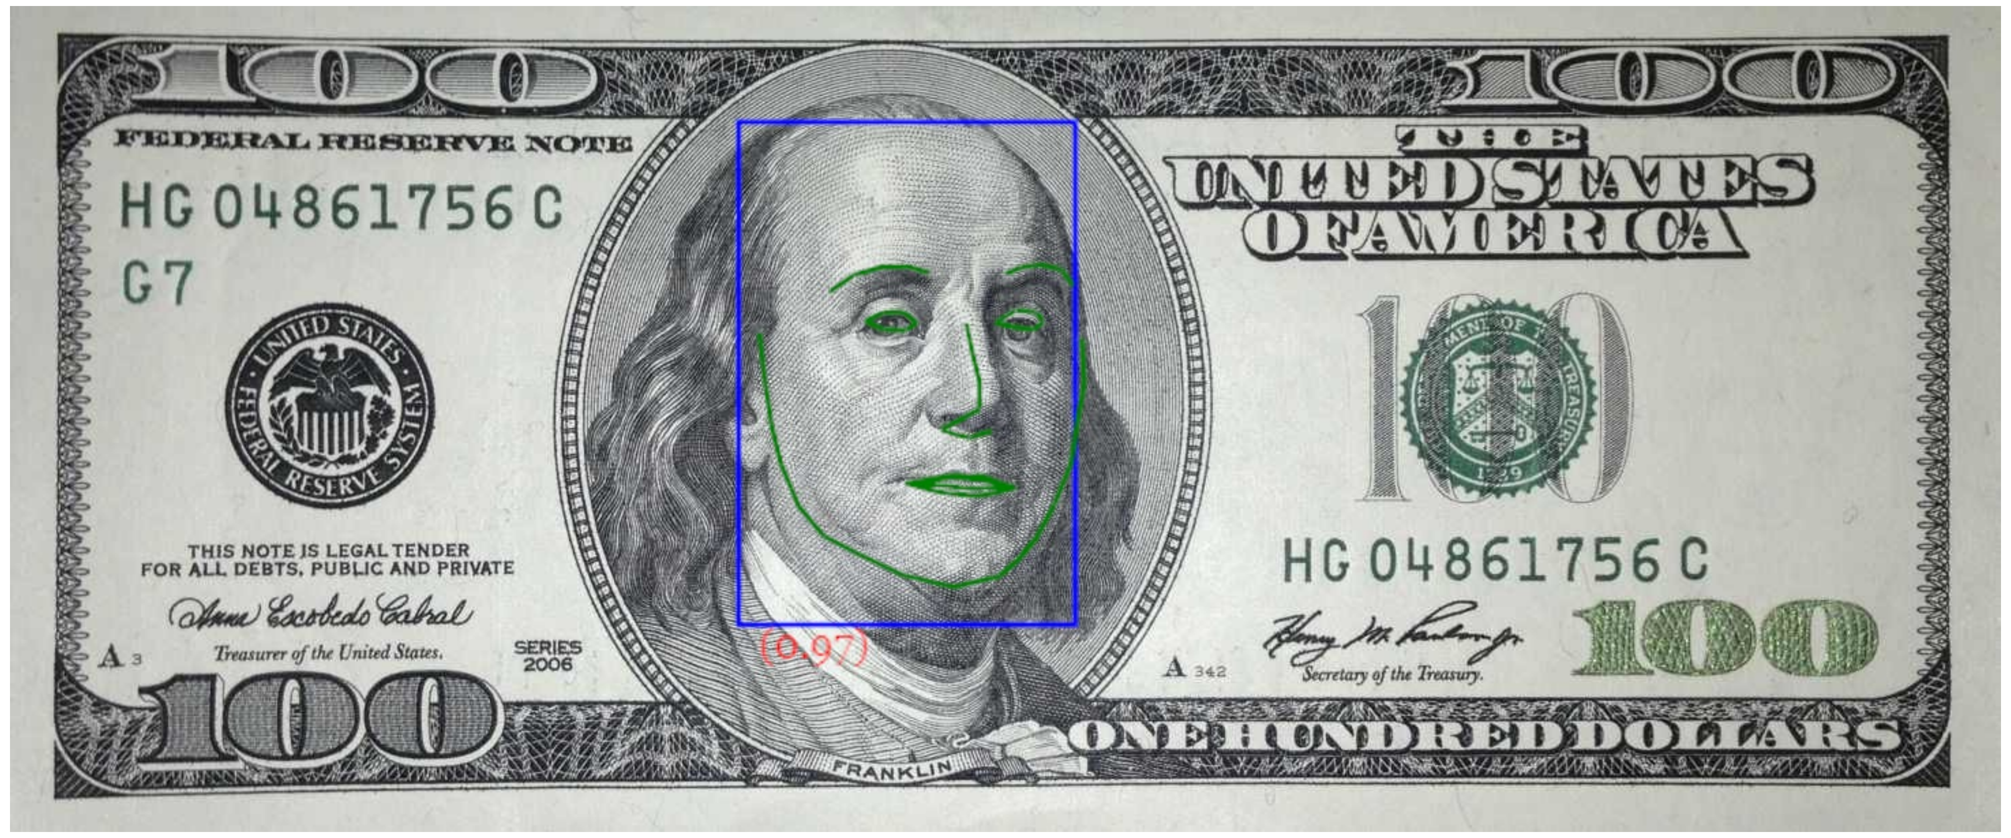

Figure credit: Raul Incze , and Cognifeed

Low cost — Cost: data collection, preprocessing, labeling cost, etc. — **High Cost**

Machine Learning by Cost

Large amount of Unlabeled data	Few Labeled Data,	Some Labeled Data, Many unlabeld data	Large amount of labeled data
⬇	⬇	⬇	⬇
Unsupervised Learning, Self-Supervised Learning	Few-shot Learning	Semi-supervised Learning	(Fully) Supervised Learning

Quality of Labels: Noisy/weakly labeled data ➡ Learning with noisy labels

Fu, Yanwei, et al. "Recent advances in zero-shot recognition: Toward data-efficient understanding of visual content." IEEE SPM, 2018.

From Few-shot Learning to Many-shot Learning

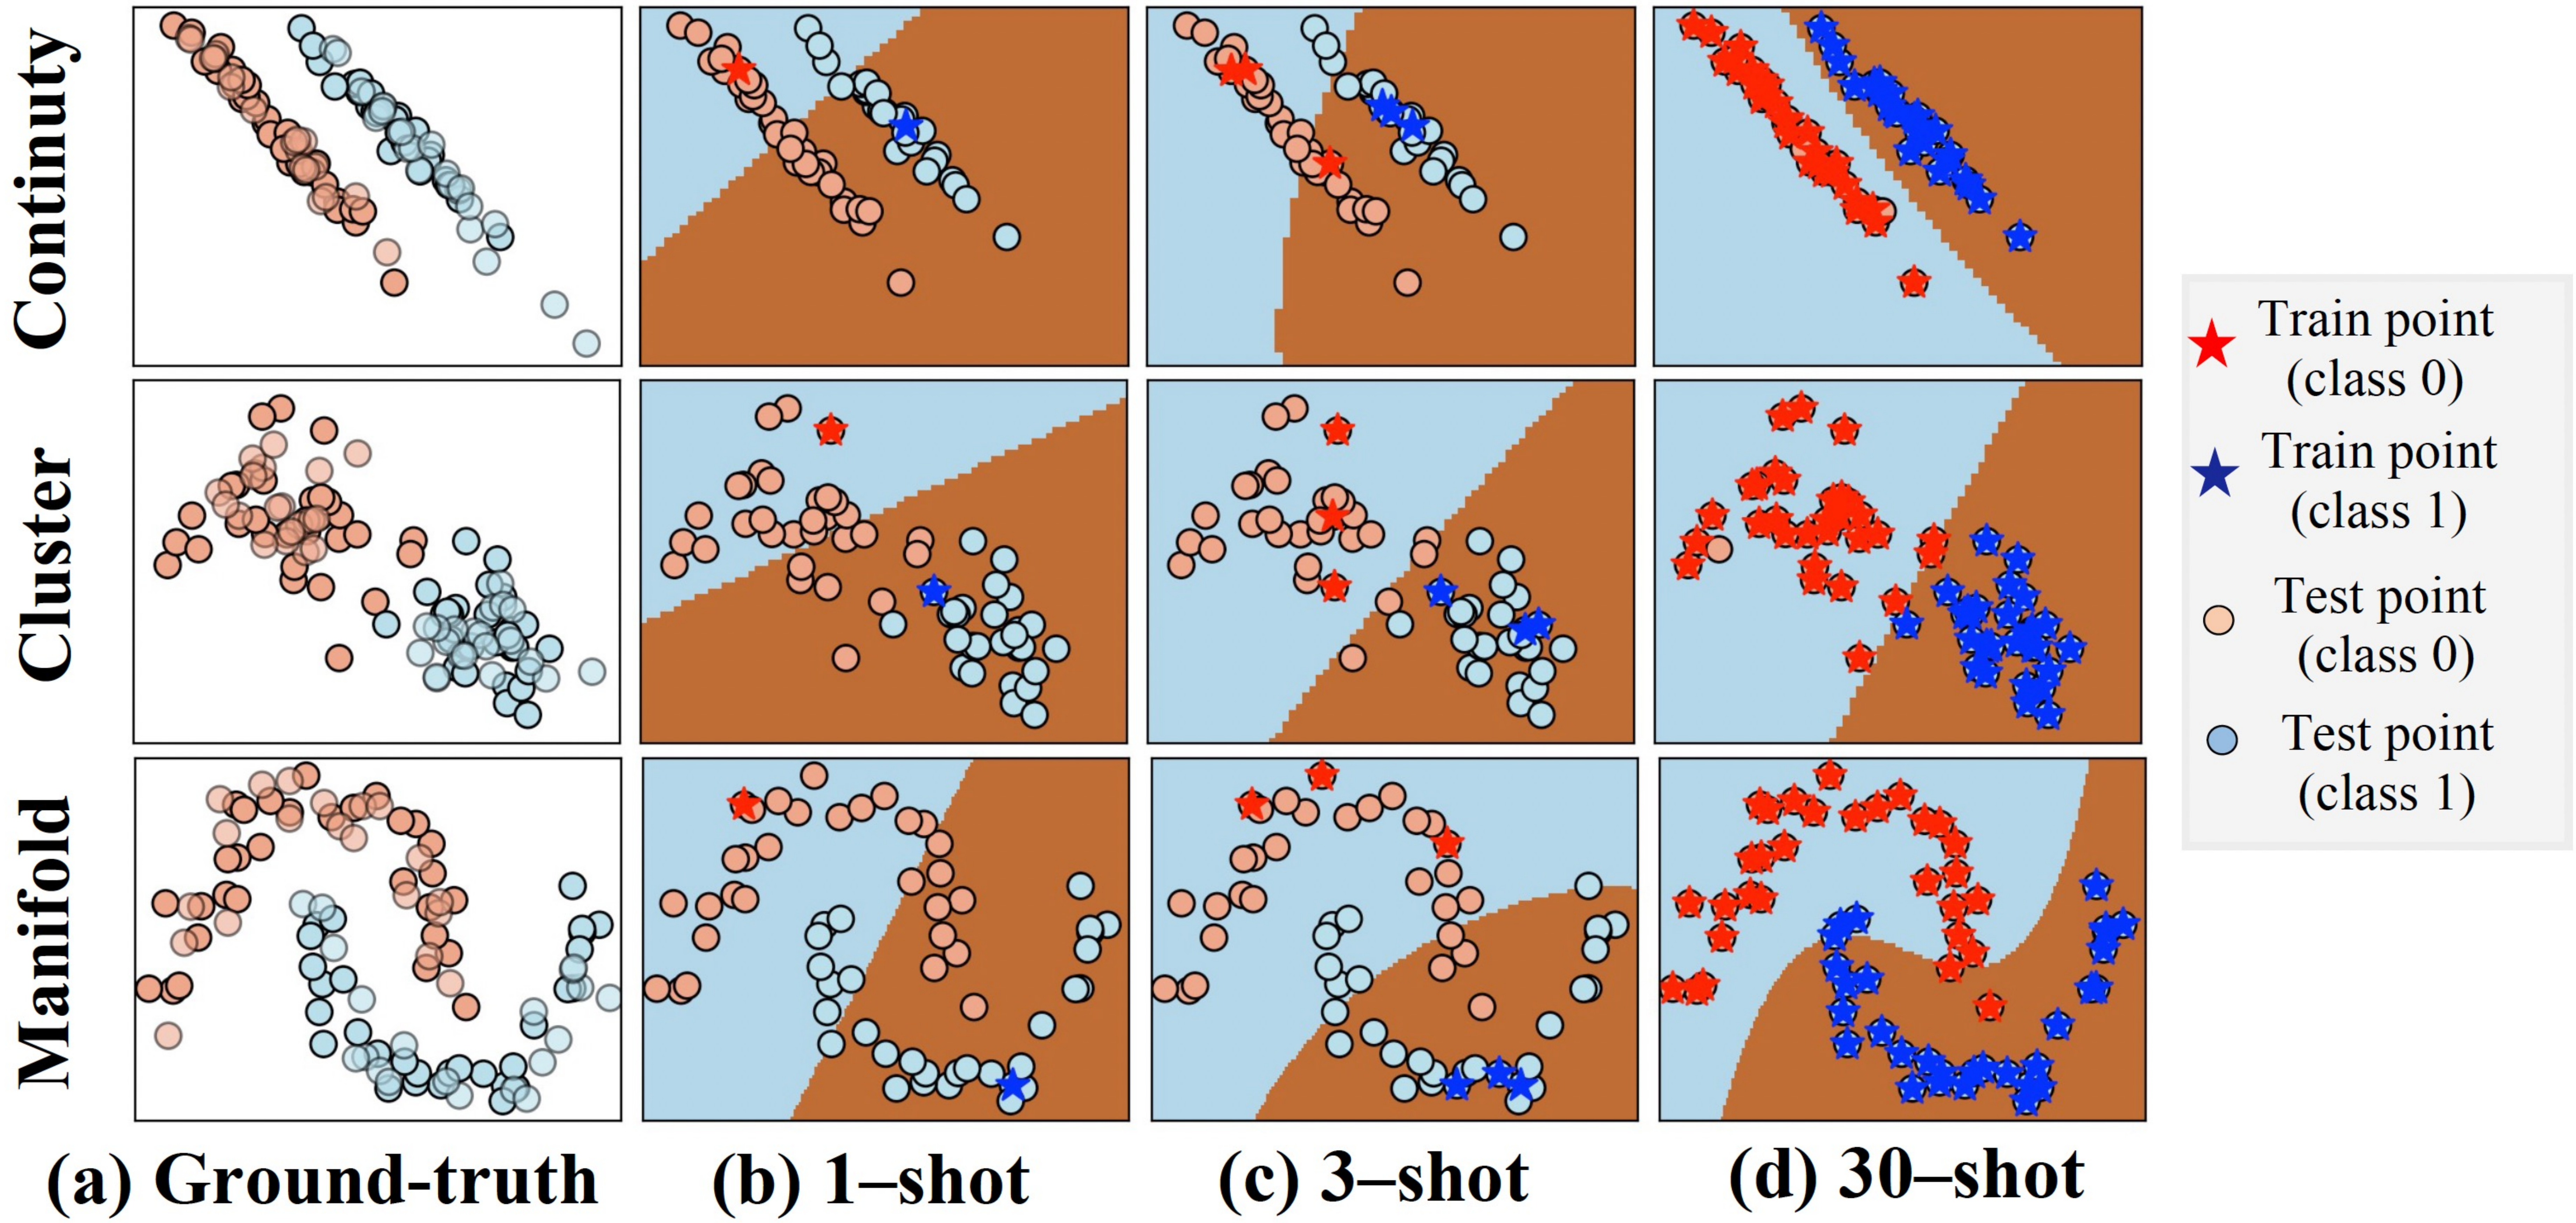

The continuity, cluster, and manifold assumptions of underlying data distribution.

Li, Pan, et al. "Regularising Knowledge Transfer by Meta Functional Learning ." IJCAI, 2021.

Few-Shot Learning Setup

Task formulation

credit: ImageNet

Classes with many samples

credit: CUB-200

Classes with few samples

Base Data

Novel Data

- N-way K-shot meta-learning setting (N random categories)
  - K samples for each category in support set $\mathcal{S} = \{(I_i^{supp}, y_i^{supp})\}$
  - Q samples for each category in query set $\mathcal{Q} = \{I_i^q, y_i^q\}$
- Goal: transfer knowledge from $\mathcal{D}_S = \{(I_i, y_i), y_i \in \mathcal{C}_s\}$ to $\mathcal{D}_t = \{(I_i, y_i), y_i \in \mathcal{C}_t\}$ $(\mathcal{C}_s \cap \mathcal{C}_t = \emptyset)$

Few-Shot Learning by Unlabeled Data

Task formulation

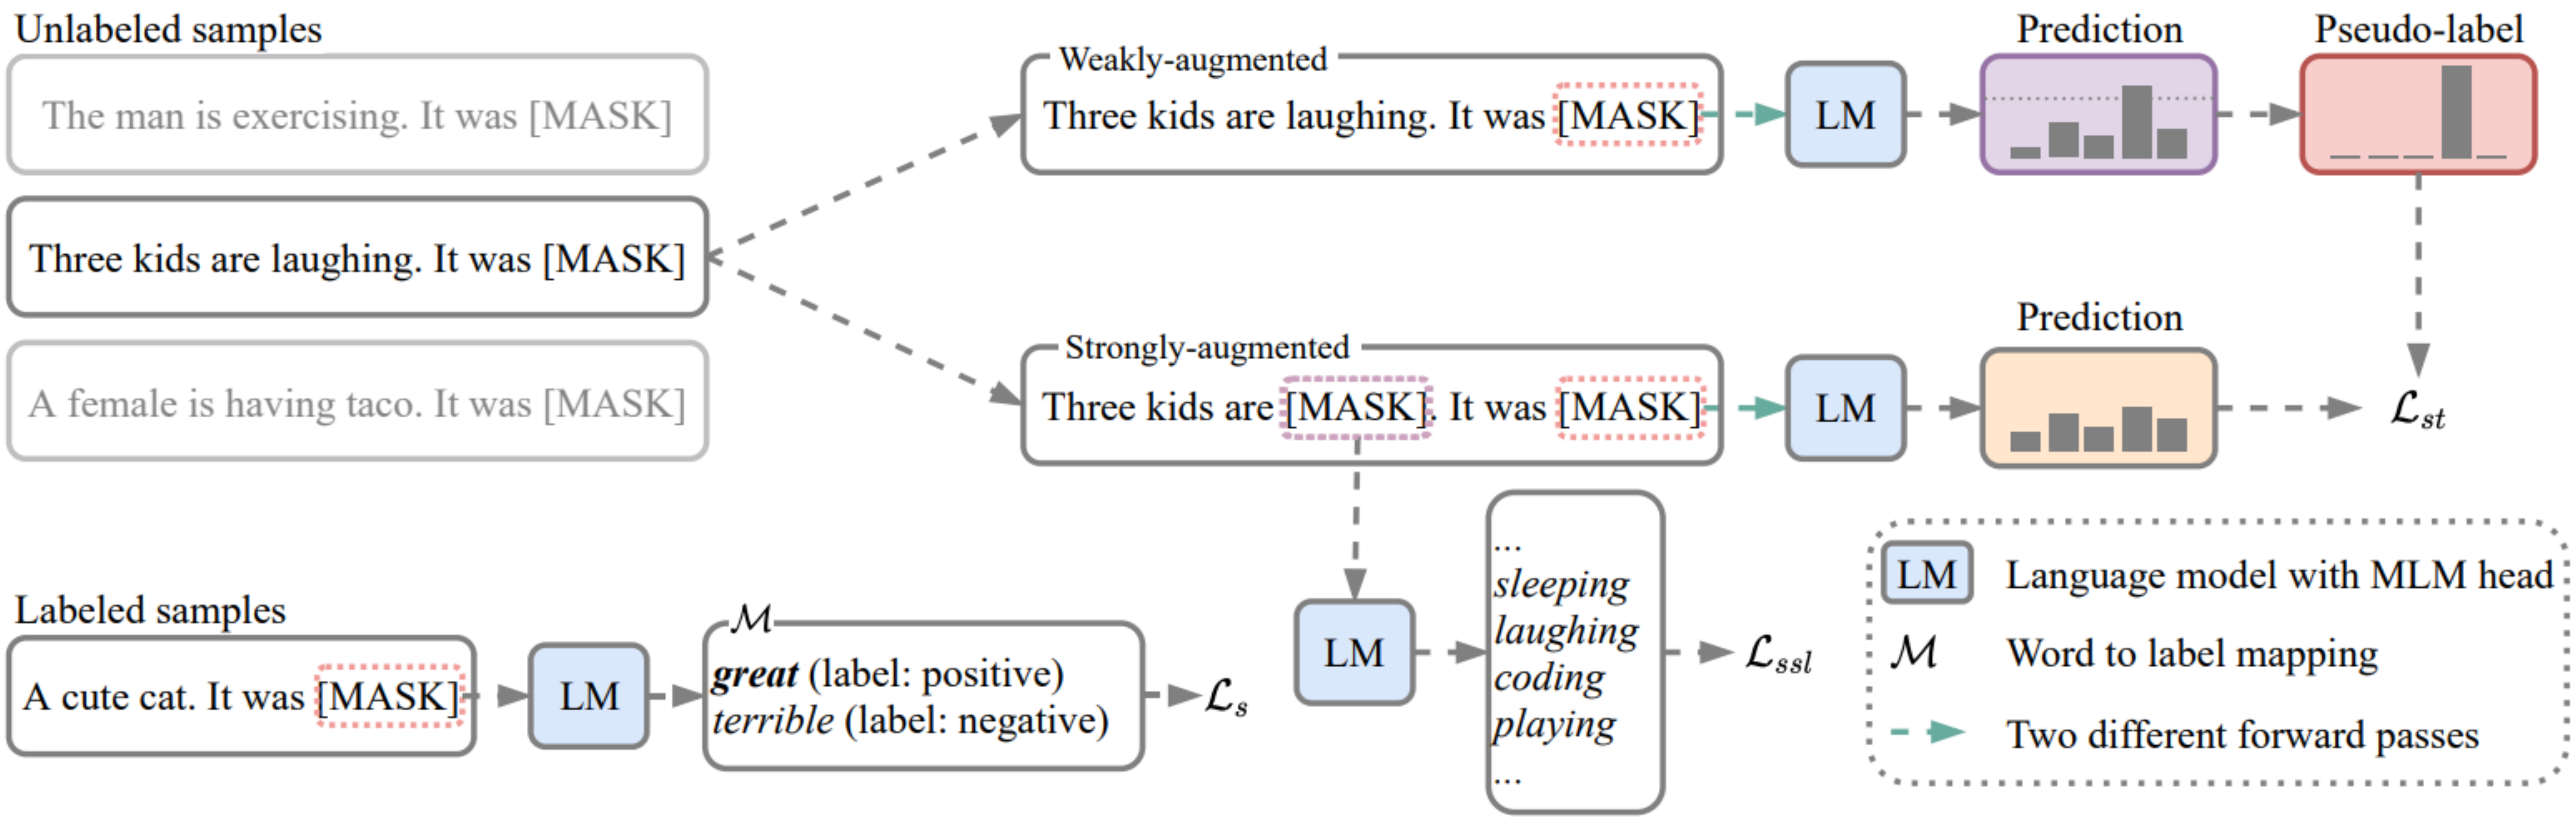

Few-Shot Learning with unlabeled data
as Semi-supervised Learning?

Yiming Chen, et al. "Revisiting Self-Training for Few-Shot Learning of Language Model", EMNLP 2021

大数据学院 School of Data Science

Few-Shot Learning Tasks

Task formulation

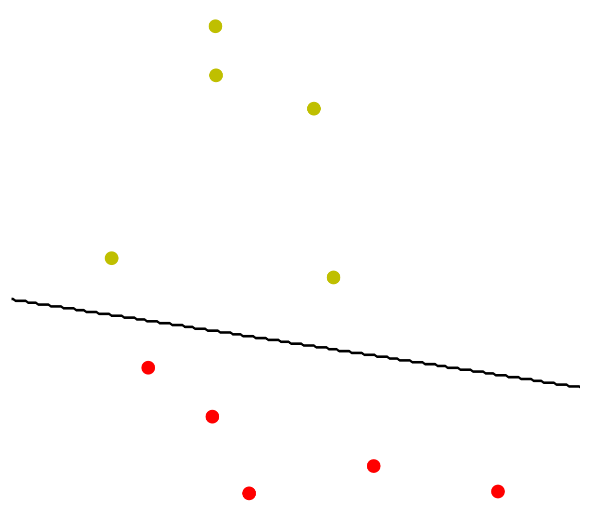

Recognition

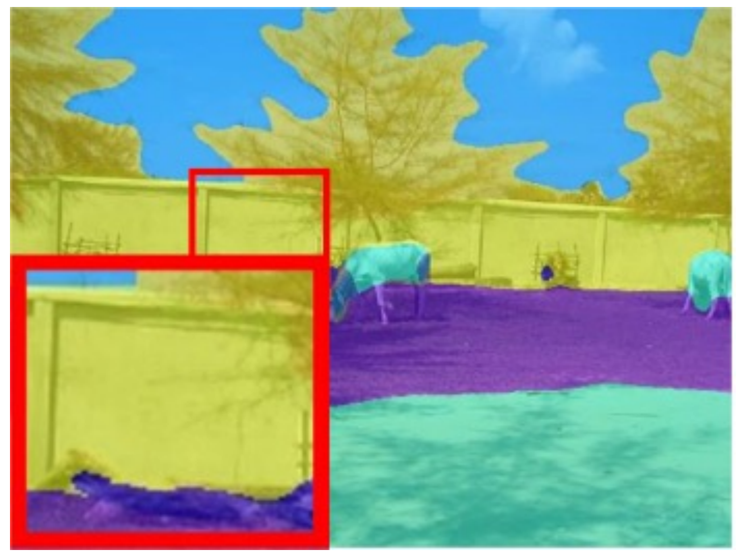

Segmentation

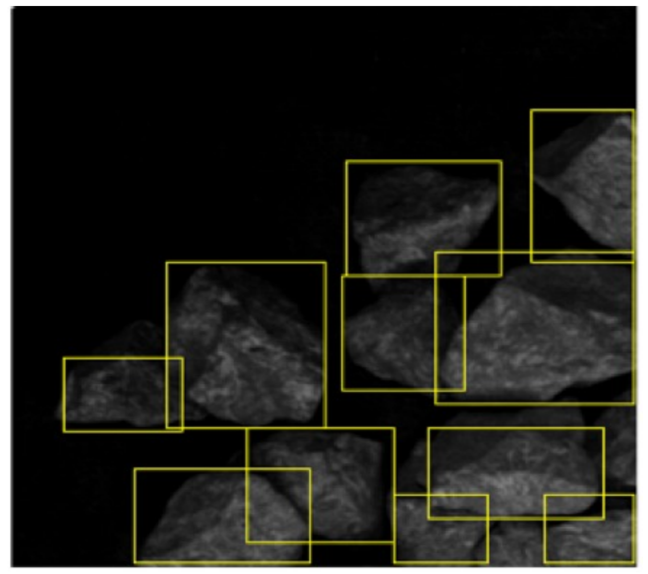

Object Detection

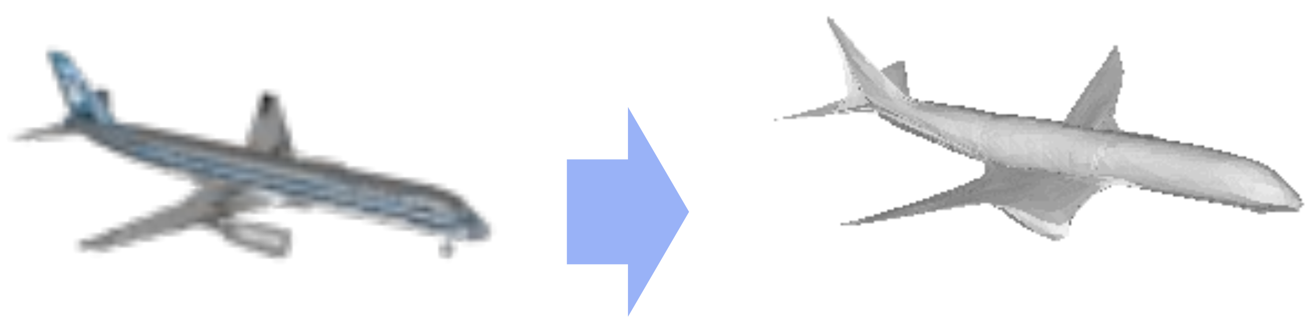

Reconstruction

Robotic Grapsing

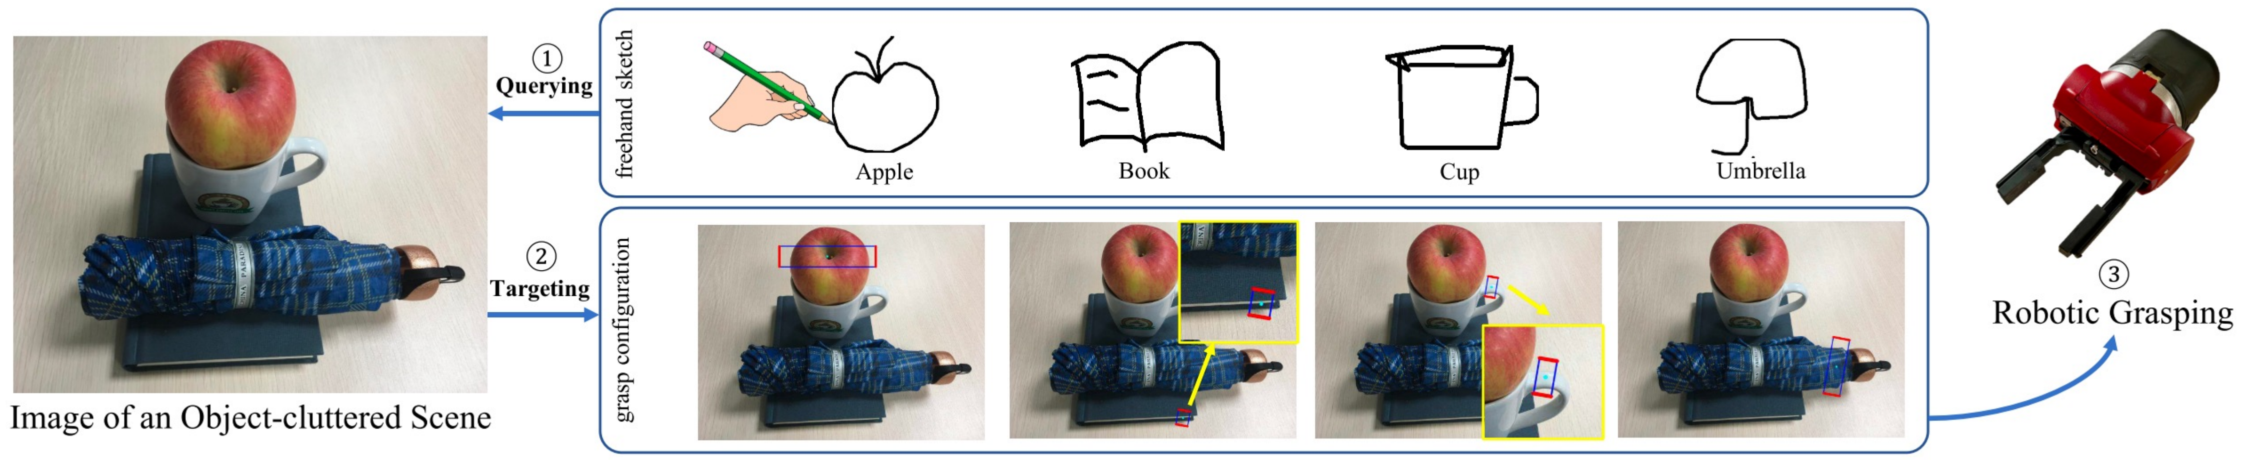

In-context Learning

@credit: [Zheng-CVPR2021][Pixel2mesh] [Lin-ICRA2022] [In-context]

Few-shot Learning: Current State

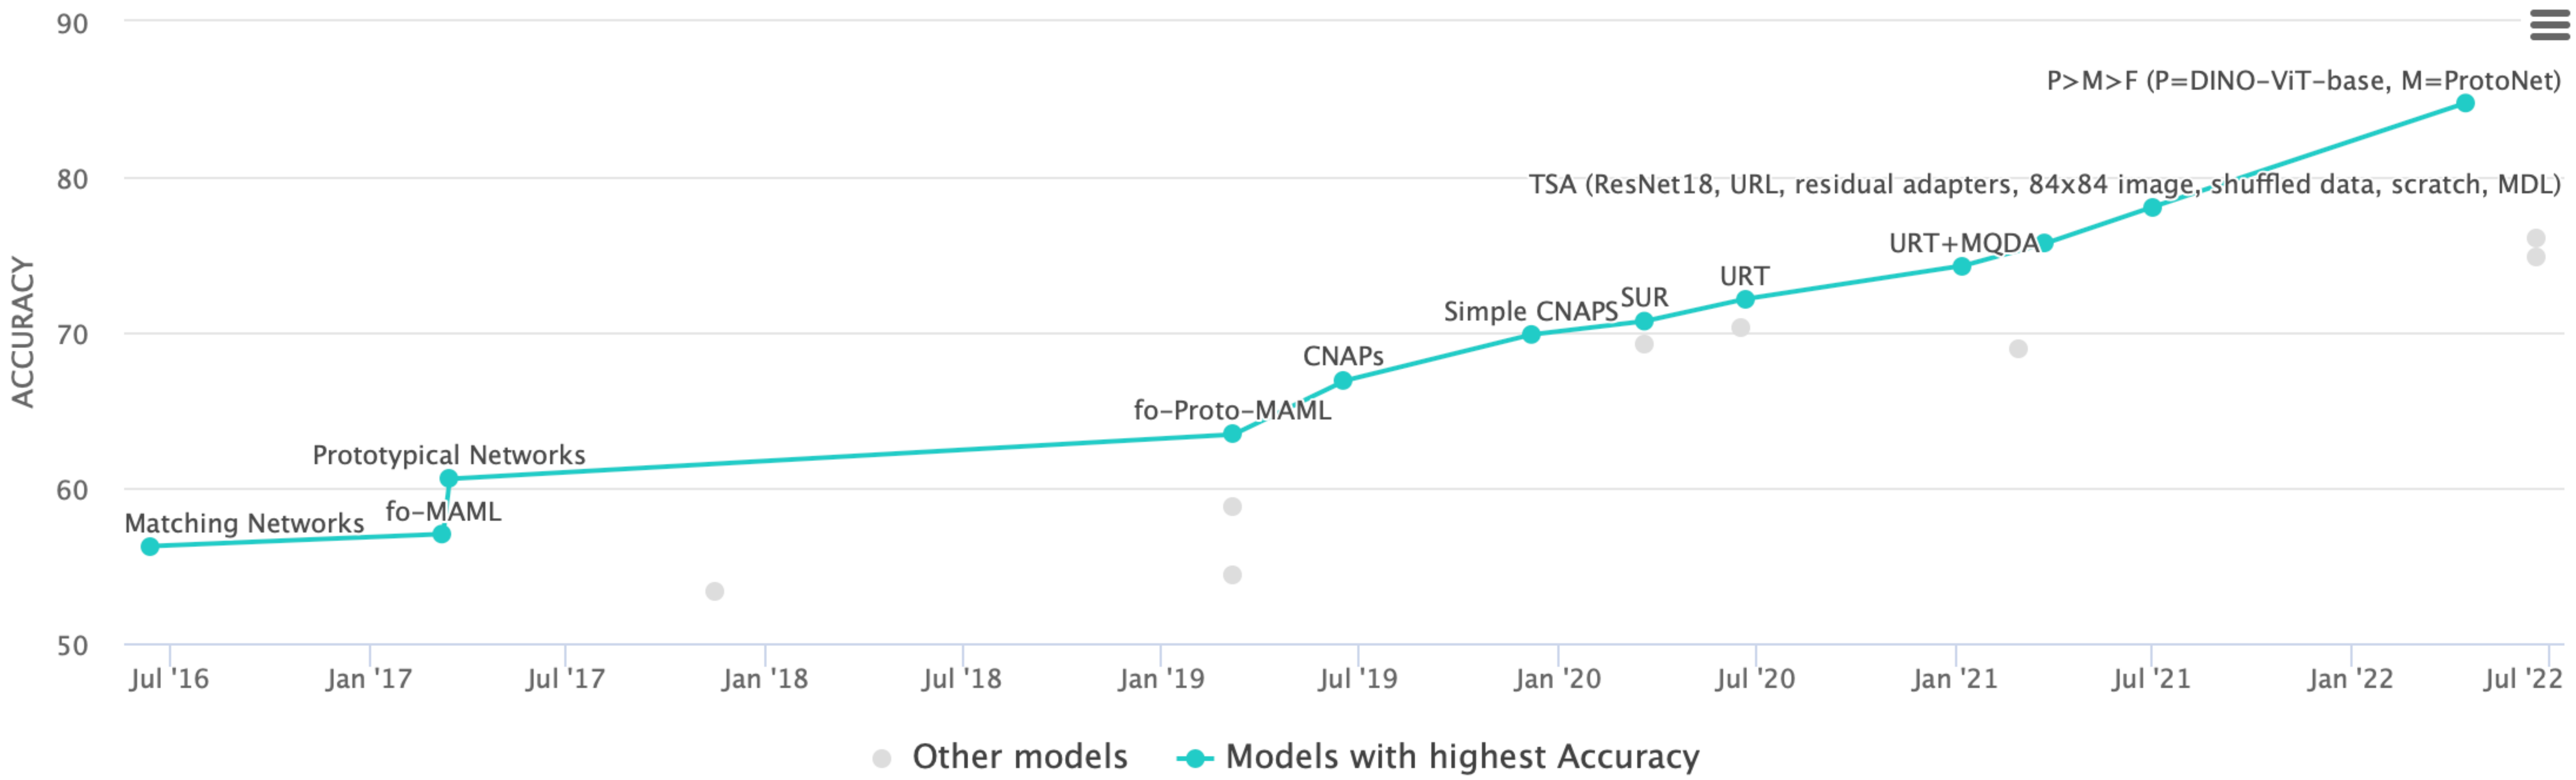

Leaderboard on Meta-Dataset. Image credit: paperswithcode.com

Triantafillou, et al. "Meta-Dataset: A Dataset of Datasets for Learning to Learn from Few Examples." ICLR2020.
https://paperswithcode.com/sota/few-shot-image-classification-on-meta-dataset

What is Required for Few-Shot Learning?

- Prior knowledge
  - A Large and relevant dataset

credit: GluonCV

  - Adaptable Knowledge

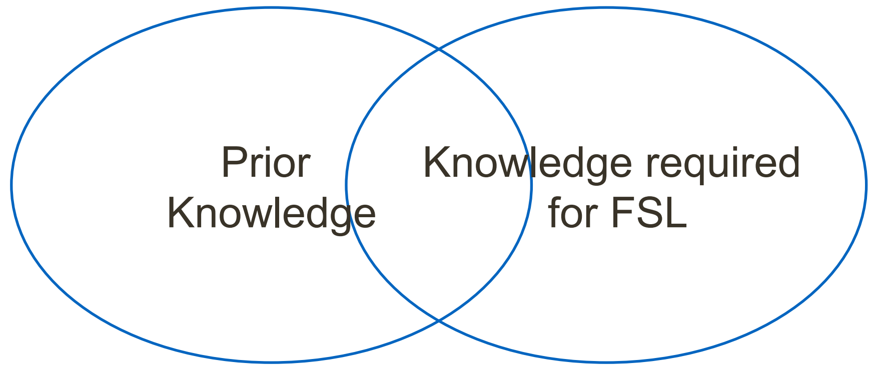

- Fast Adaptation Capacity
  - Pre-train and fine-tune:

Reducing learning difficulty

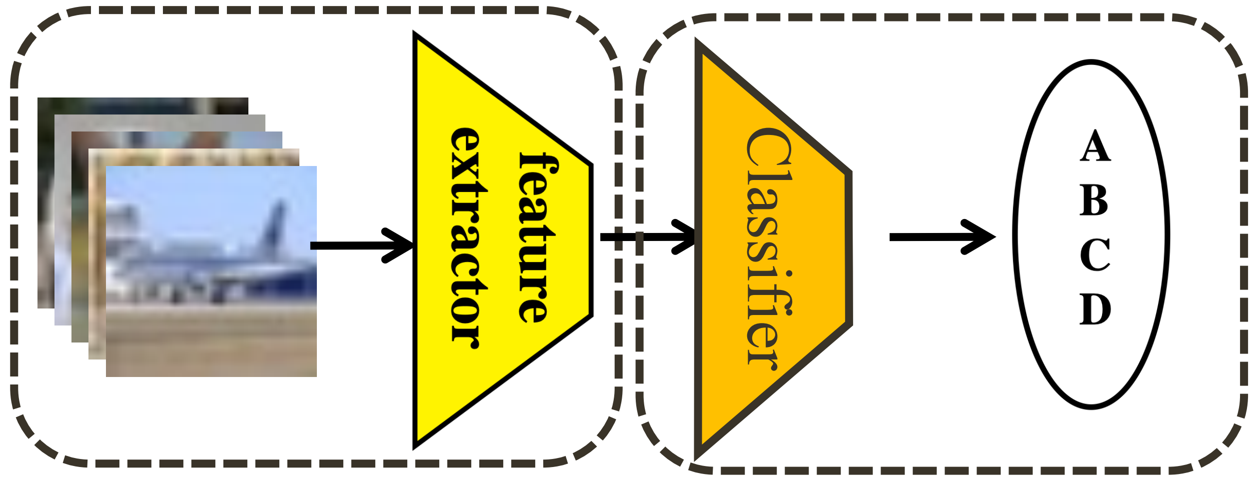

  - Meta-learning:

Learning to learn from few examples

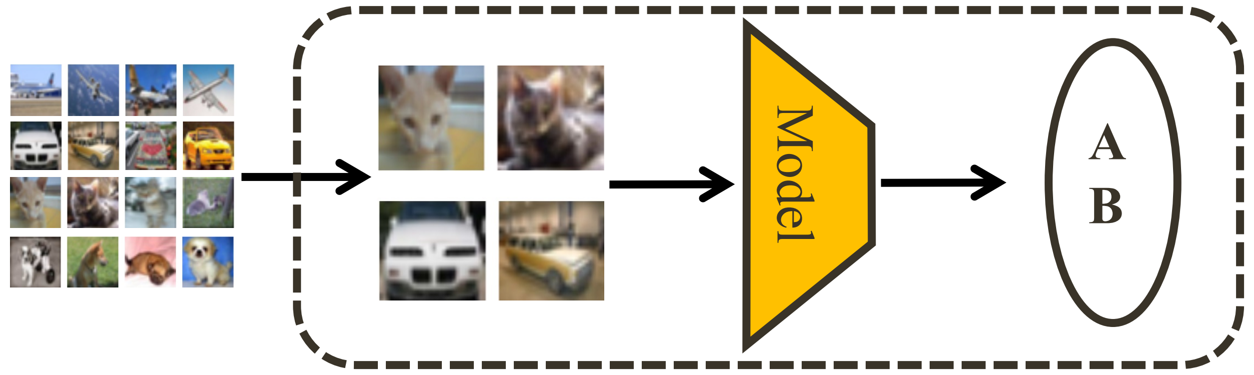

When Statistics Meets Few-Shot Learning

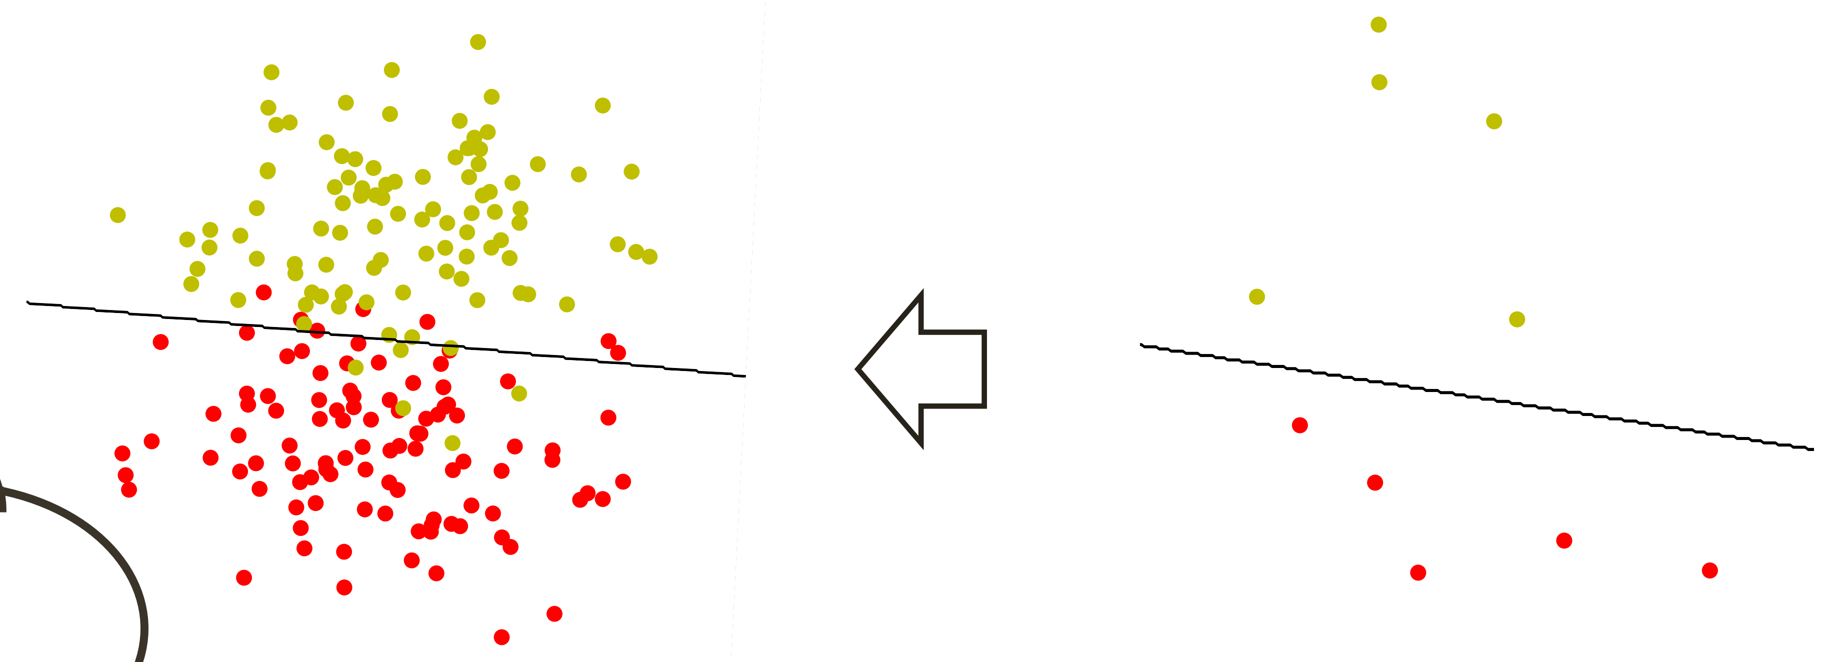

Posterior estimation

What is FSL in statistics?

similarity measure

Calibration

Causality

What can be used to benefit FSL?

Outlier detection

Alignment

Bayesian inference

What can be used to solve FSL?

Gaussian process

Small sample learning in Statistics

Sun, Jun, et al. "Small Sample Learning in Big Data Era", arxiv 2018

Few-Shot Learning

- Learning from base data
  - Causality
  - Similarity Measurement
  - Neural Collapse
- Adaptation on novel data
- FSL in 2020s

The Key Idea of “Causality”

Causality v.s. Correlation:
Find features that *determines* the class instead of *correlated* with the class.

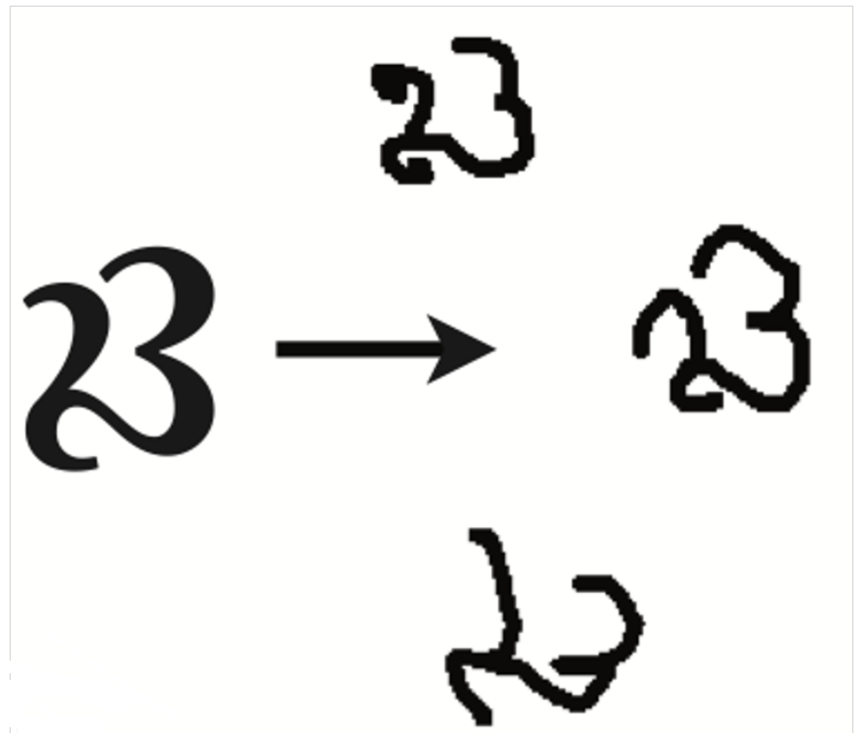

credit: [Brenden-CS2015]

Causality as part of probabilistic program induction for FSL

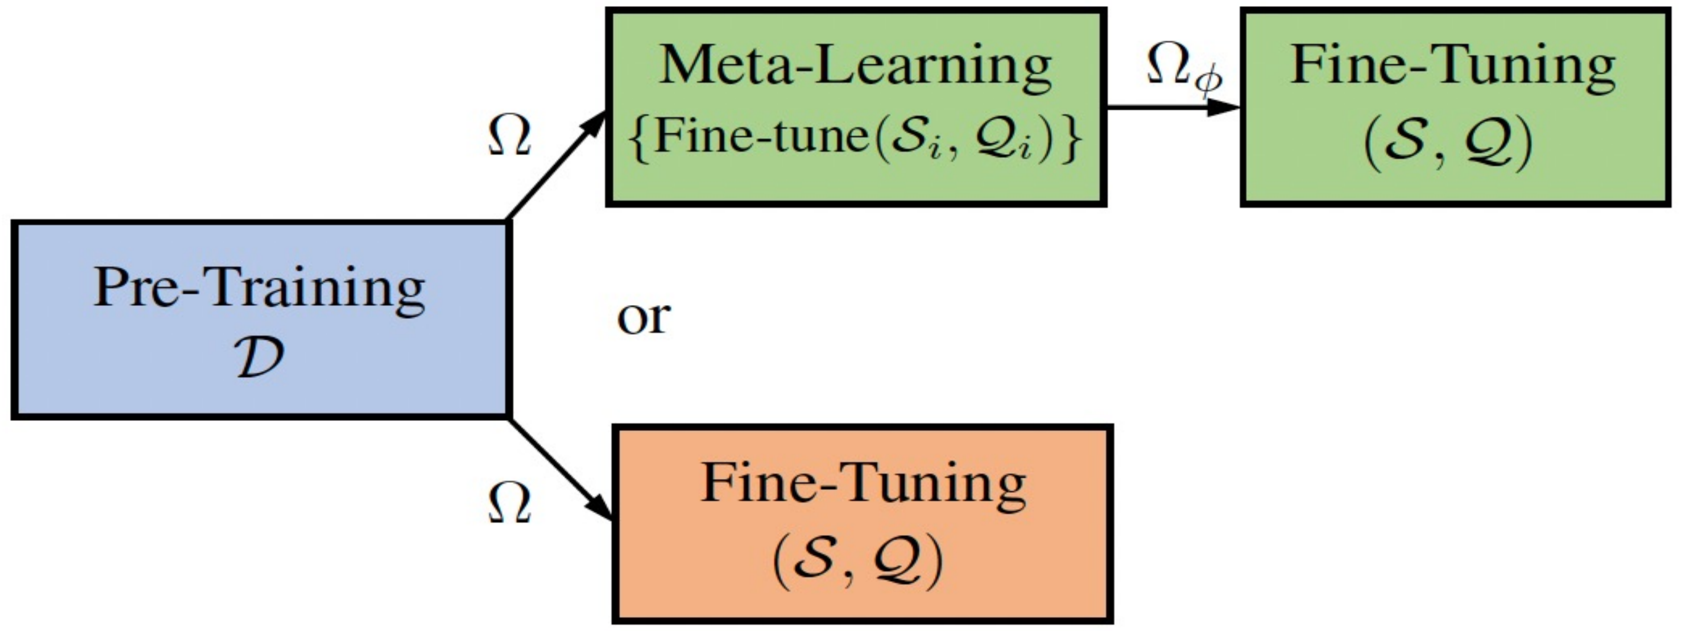

credit: [Yue-NeurIPS2020]

Causality to remove the dependency of pre-knowledge from pre-trained dataset

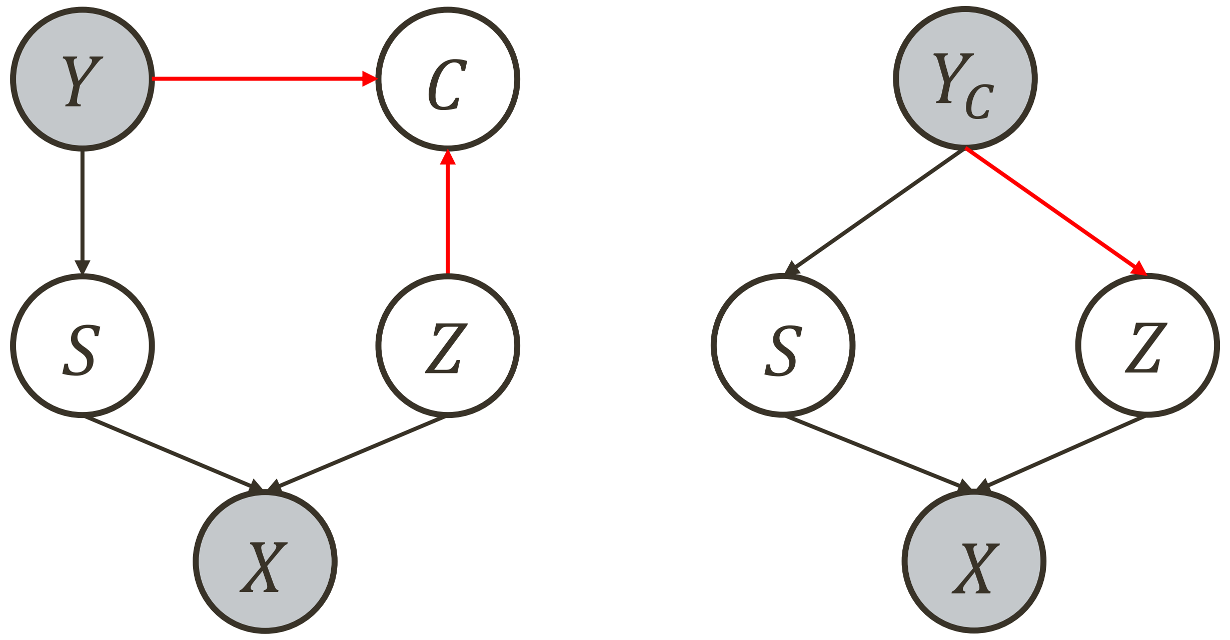

Causality to remove spurious feature

Causality as part of probabilistic program induction

Three key ideas: *compositionality*, *causality*, *learning to learn, in learning new concepts from few examples.*

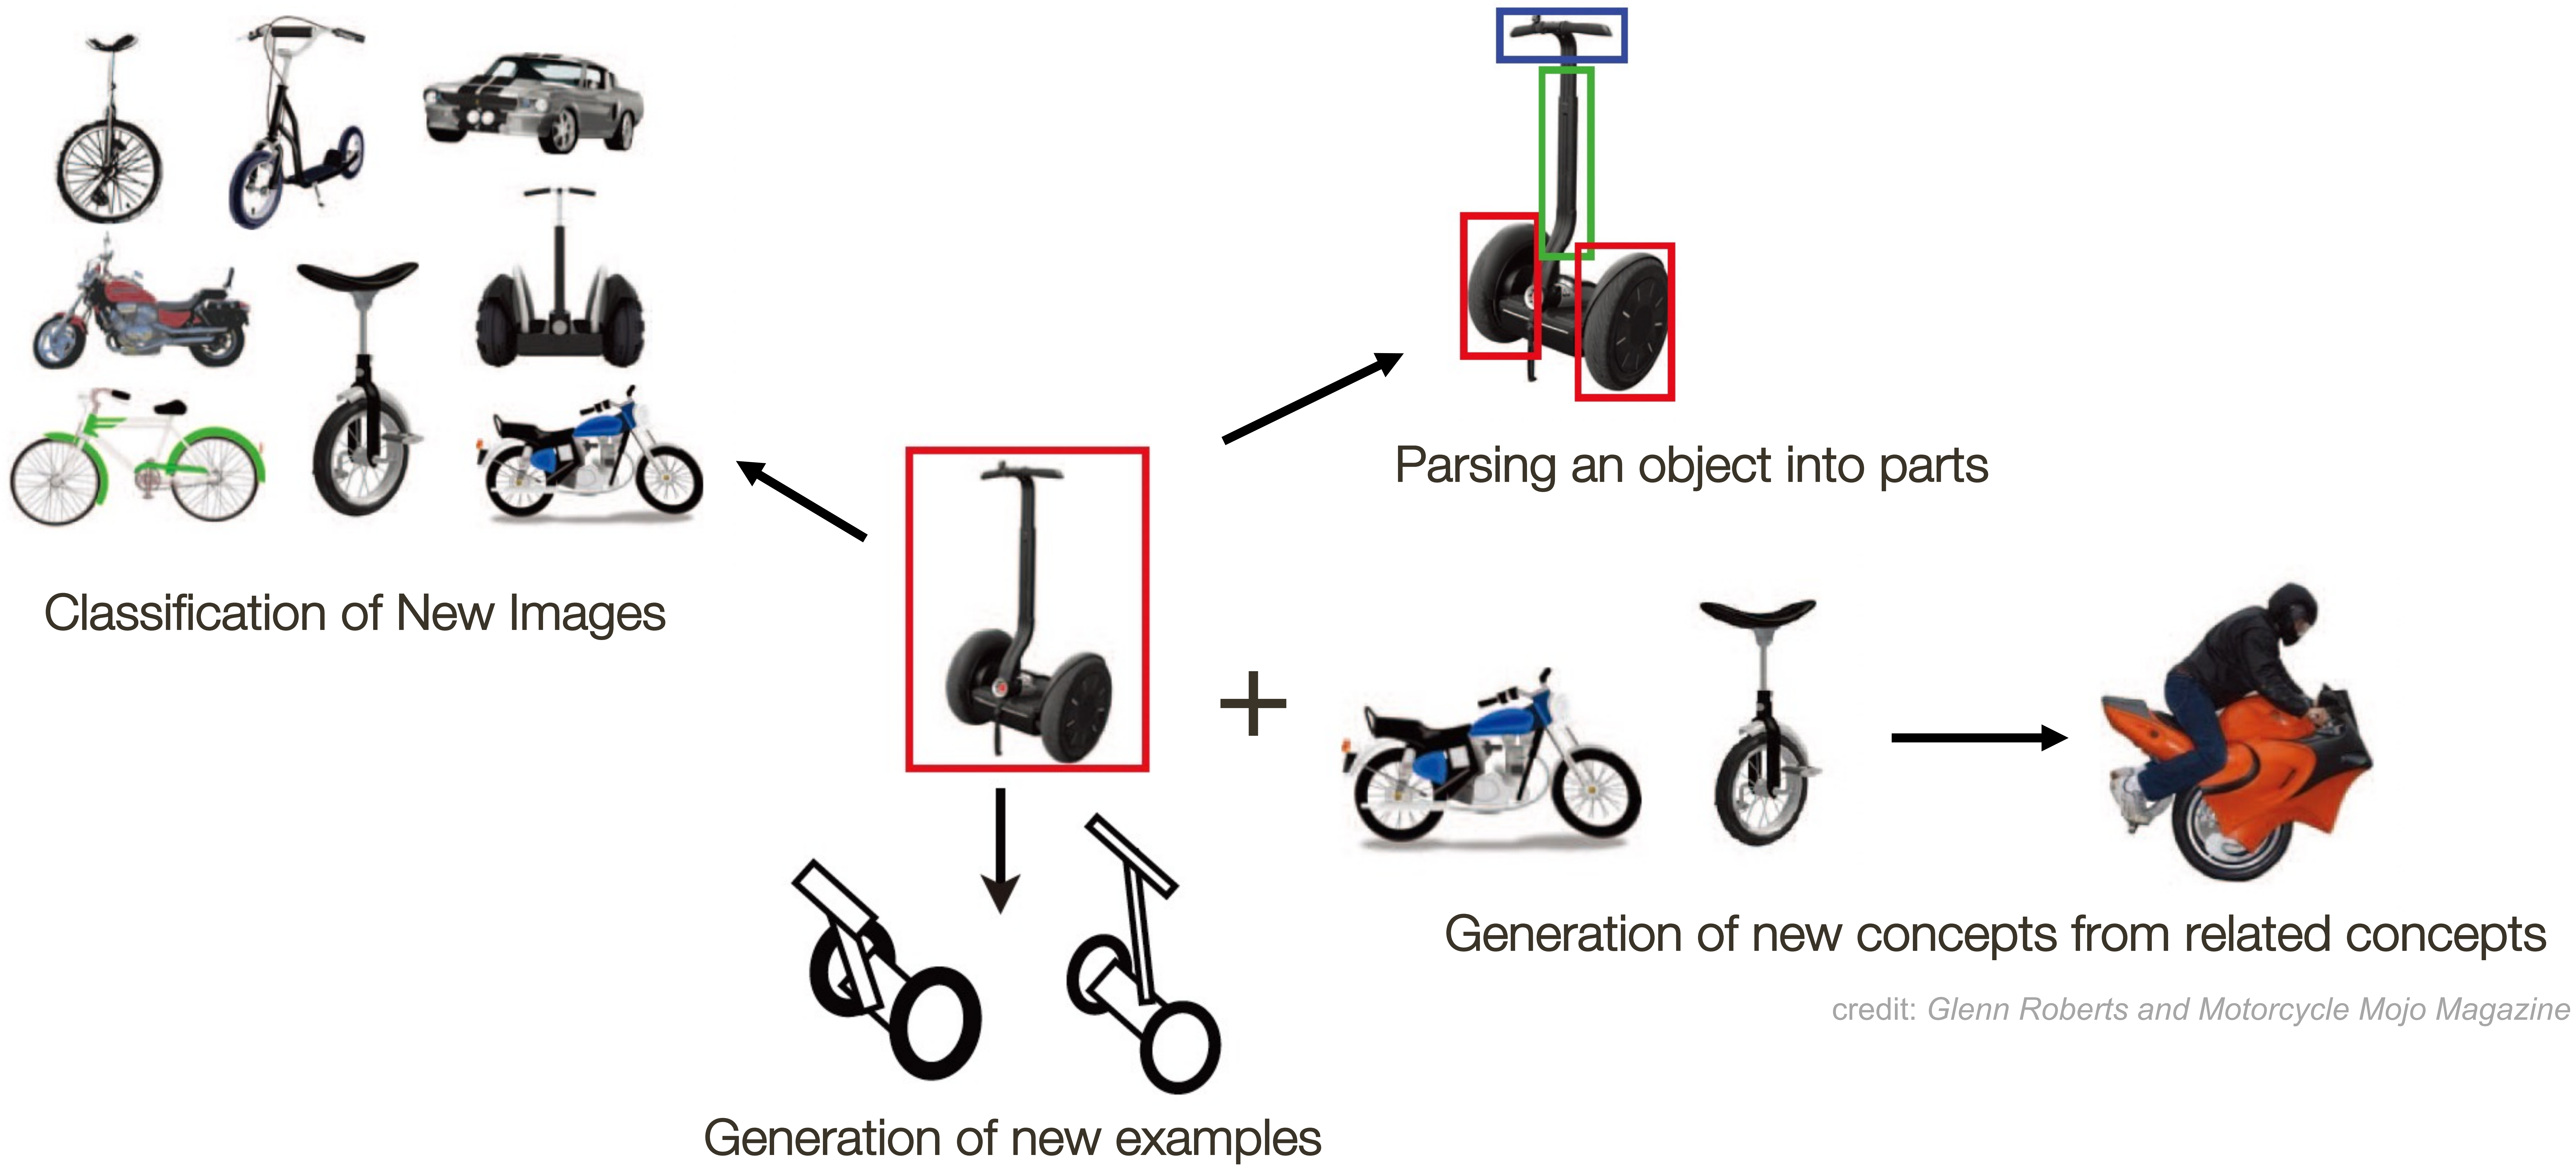

Lake, Brenden, et al. "Human-level concept learning through probabilistic program induction." COGNITIVE SCIENCE. 2015
Lake, Brenden, et al. “One-shot learning by inverting a compositional causal process”, NIPS 2013

Influence of Pre-trained Knowledge

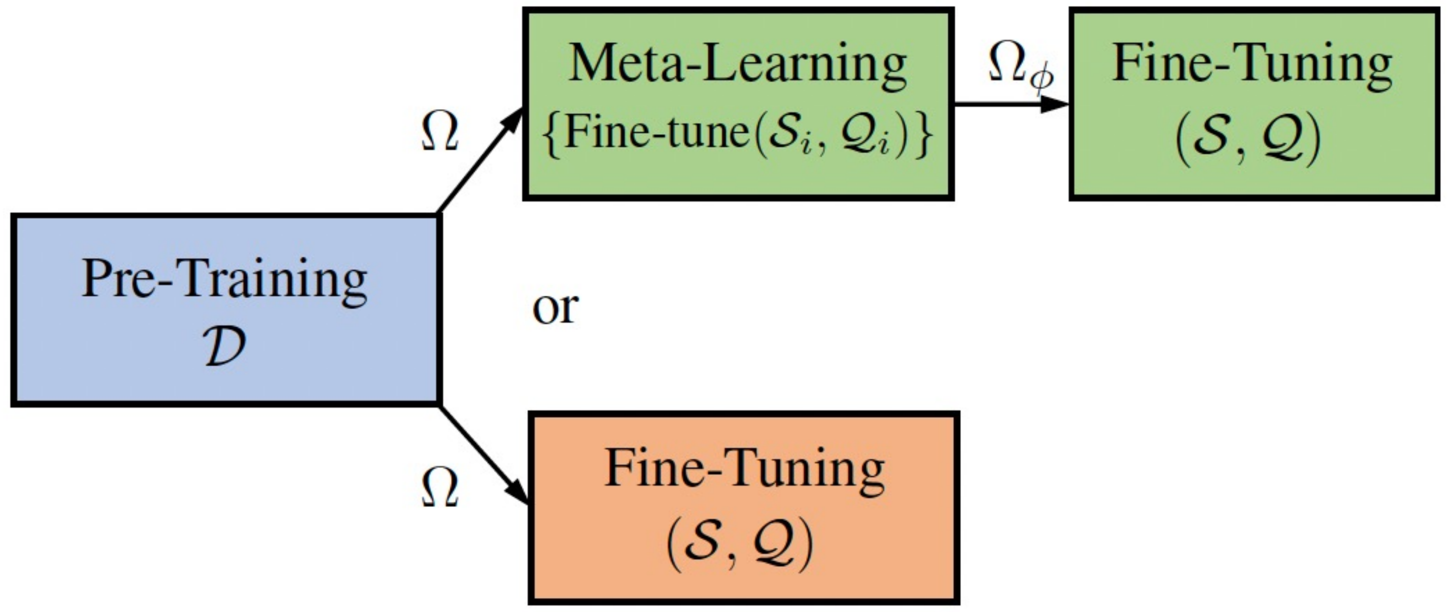

$\mathcal{D}$ = ImageNet, Ω = ResNet

$\mathcal{D}$ = Wikipedia, Ω = BERT in natural language processing

credit: [Yue-NeurIPS2020]

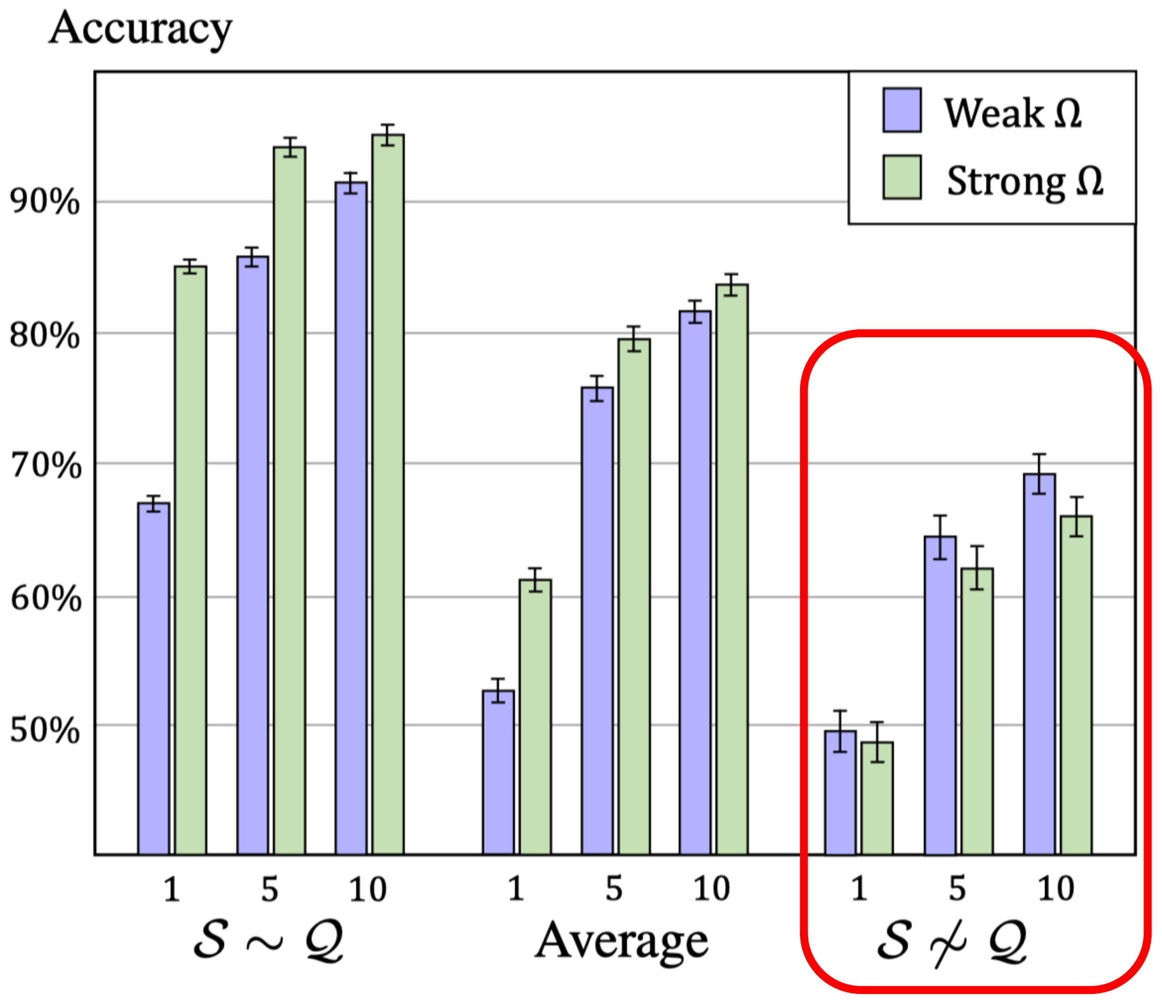

credit: [Yue-NeurIPS2020]

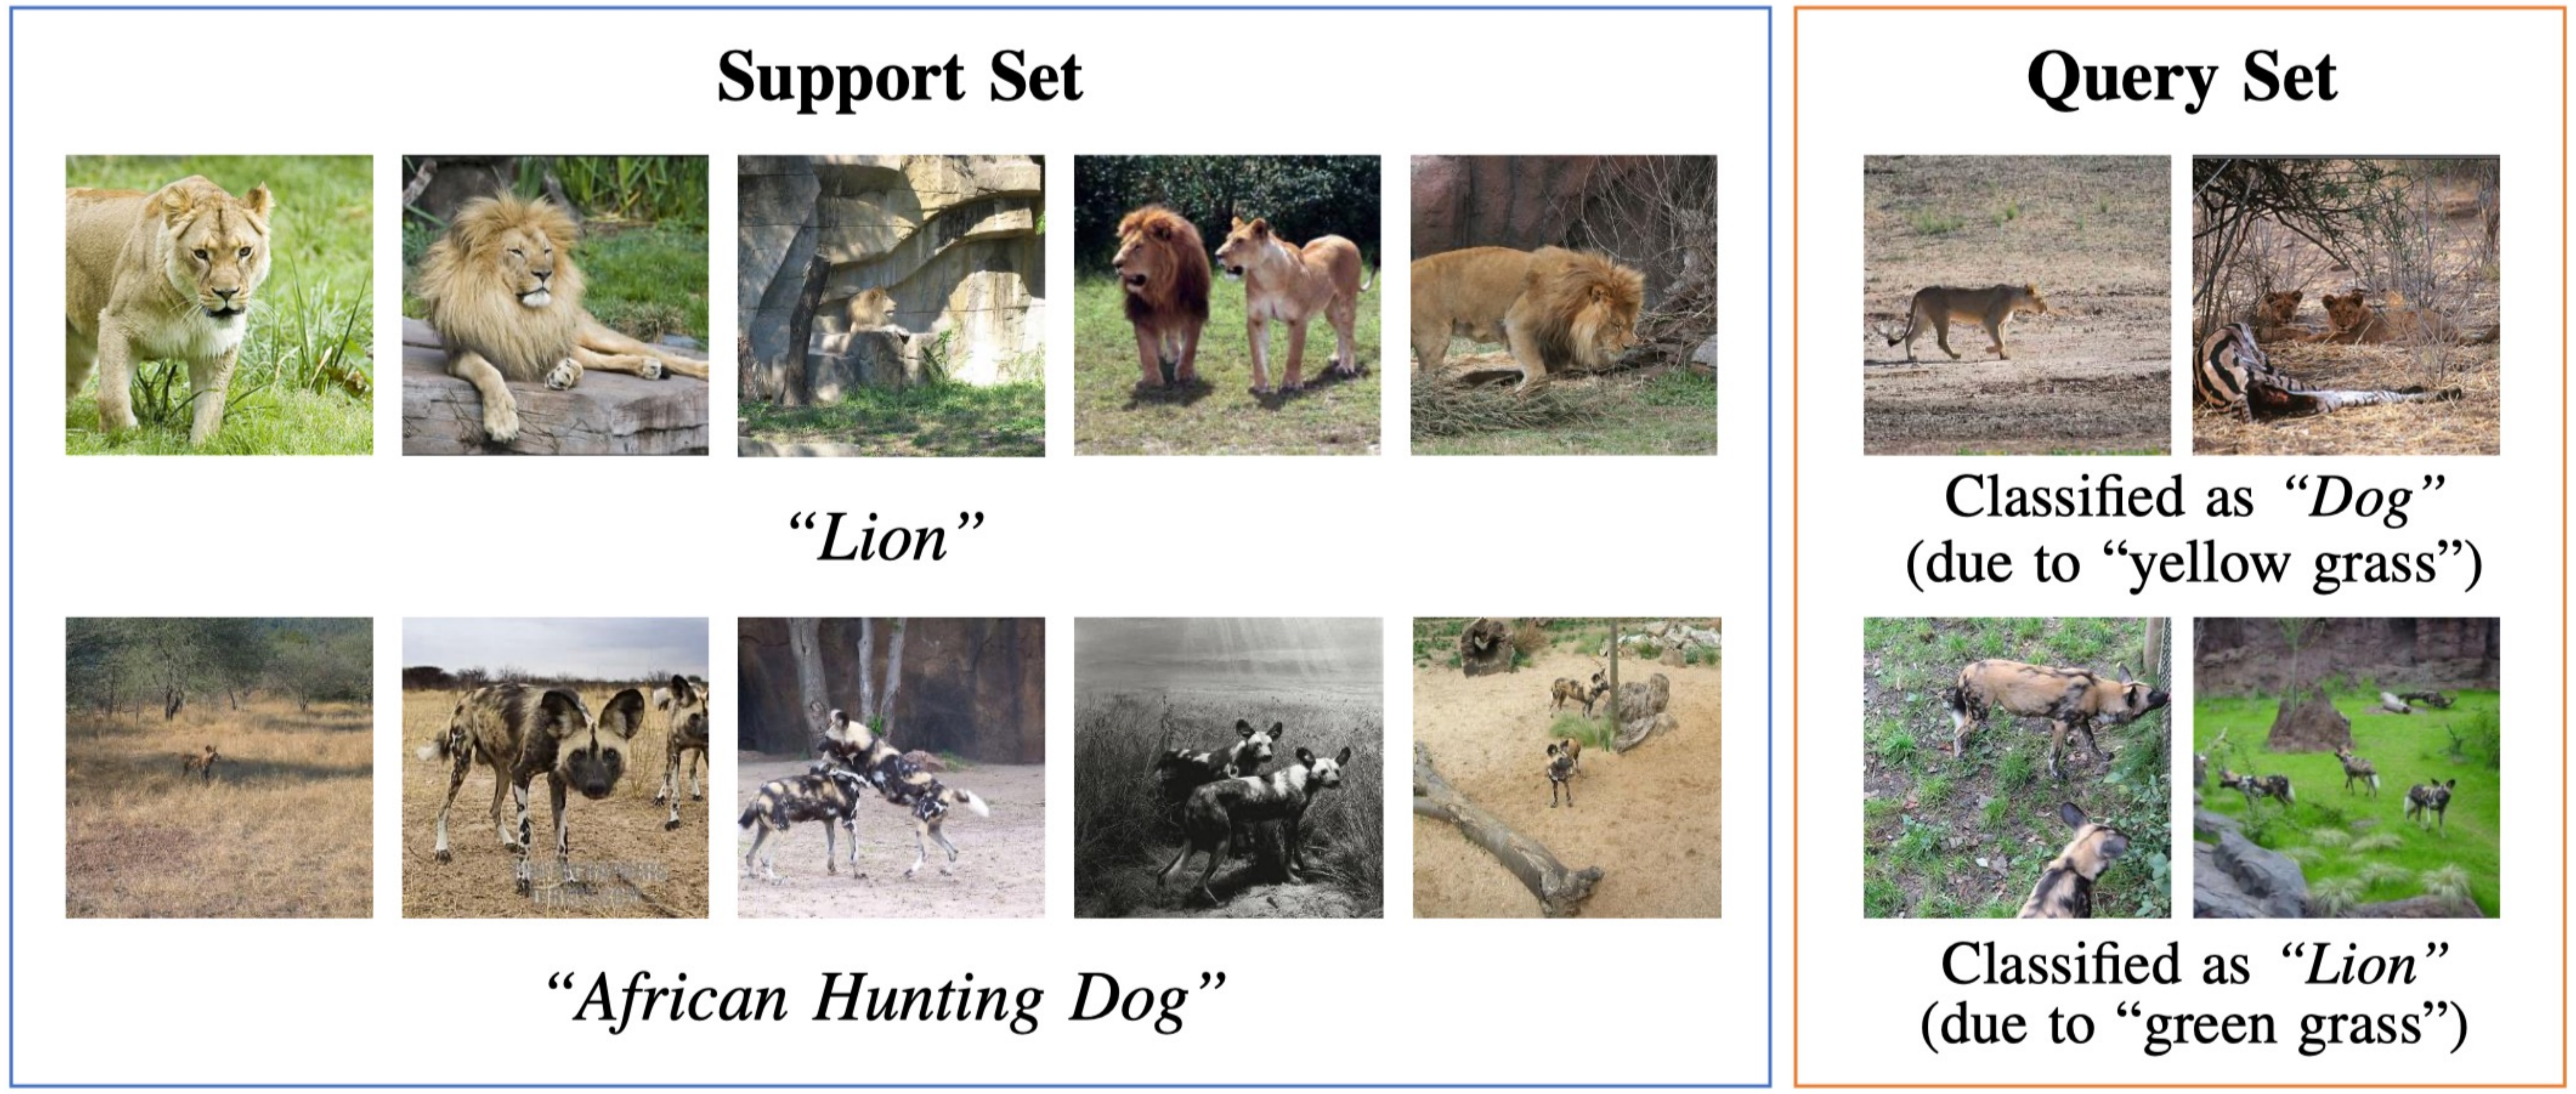

credit: [Yue-NeurIPS2020]

the pre-trained knowledge can do evil in FSL.

Yue, Zhongqi, et al. "Interventional few-shot learning." NeurIPS2020.

Causality to remove the dependency of pre-knowledge

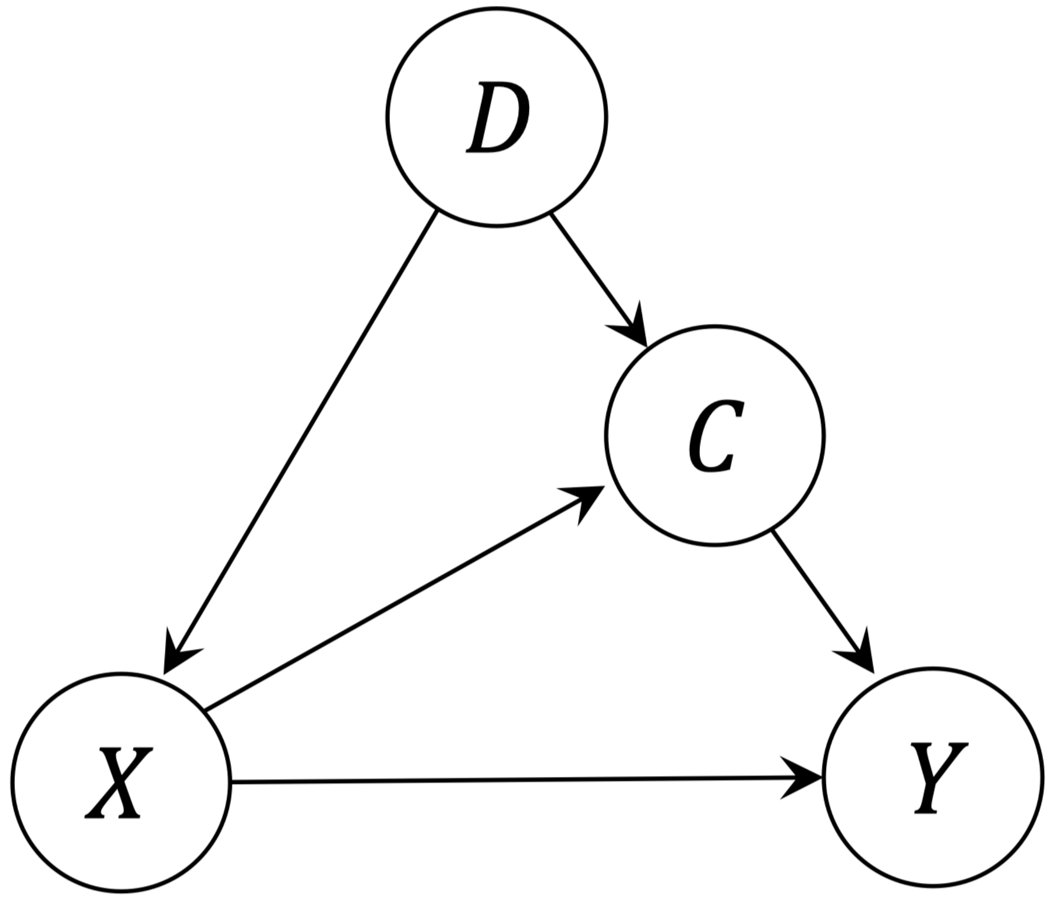

credit: [Yue-NeurIPS2020]

$$D \rightarrow X$$

$$D \rightarrow C \leftarrow X$$

$$X \rightarrow Y \leftarrow C$$

Learn $P(Y|do(X))$ instead of $P(Y|X)$.

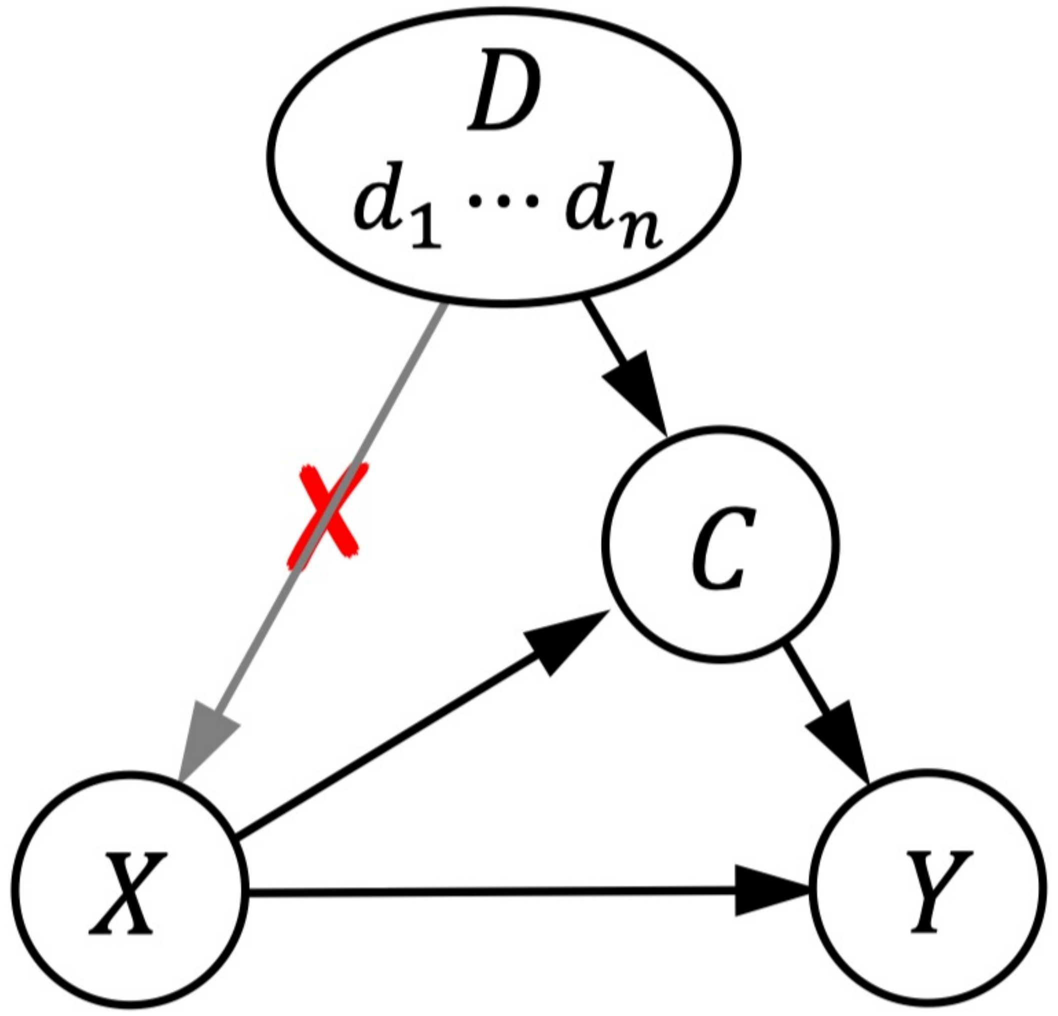

credit: [Yue-NeurIPS2020]

Causal Intervention by backdoor adjustment

$$P\left(Y|do(X=\boldsymbol{x})\right) = \sum_d P\left(Y|X=\boldsymbol{x}, D=d, C=g(\mathbf{x}, d)\right) P(D=d)$$

Yue, Zhongqi, et al. "Interventional few-shot learning." NeurIPS2020.

Structural Causal Model in FSL: data selection perspective

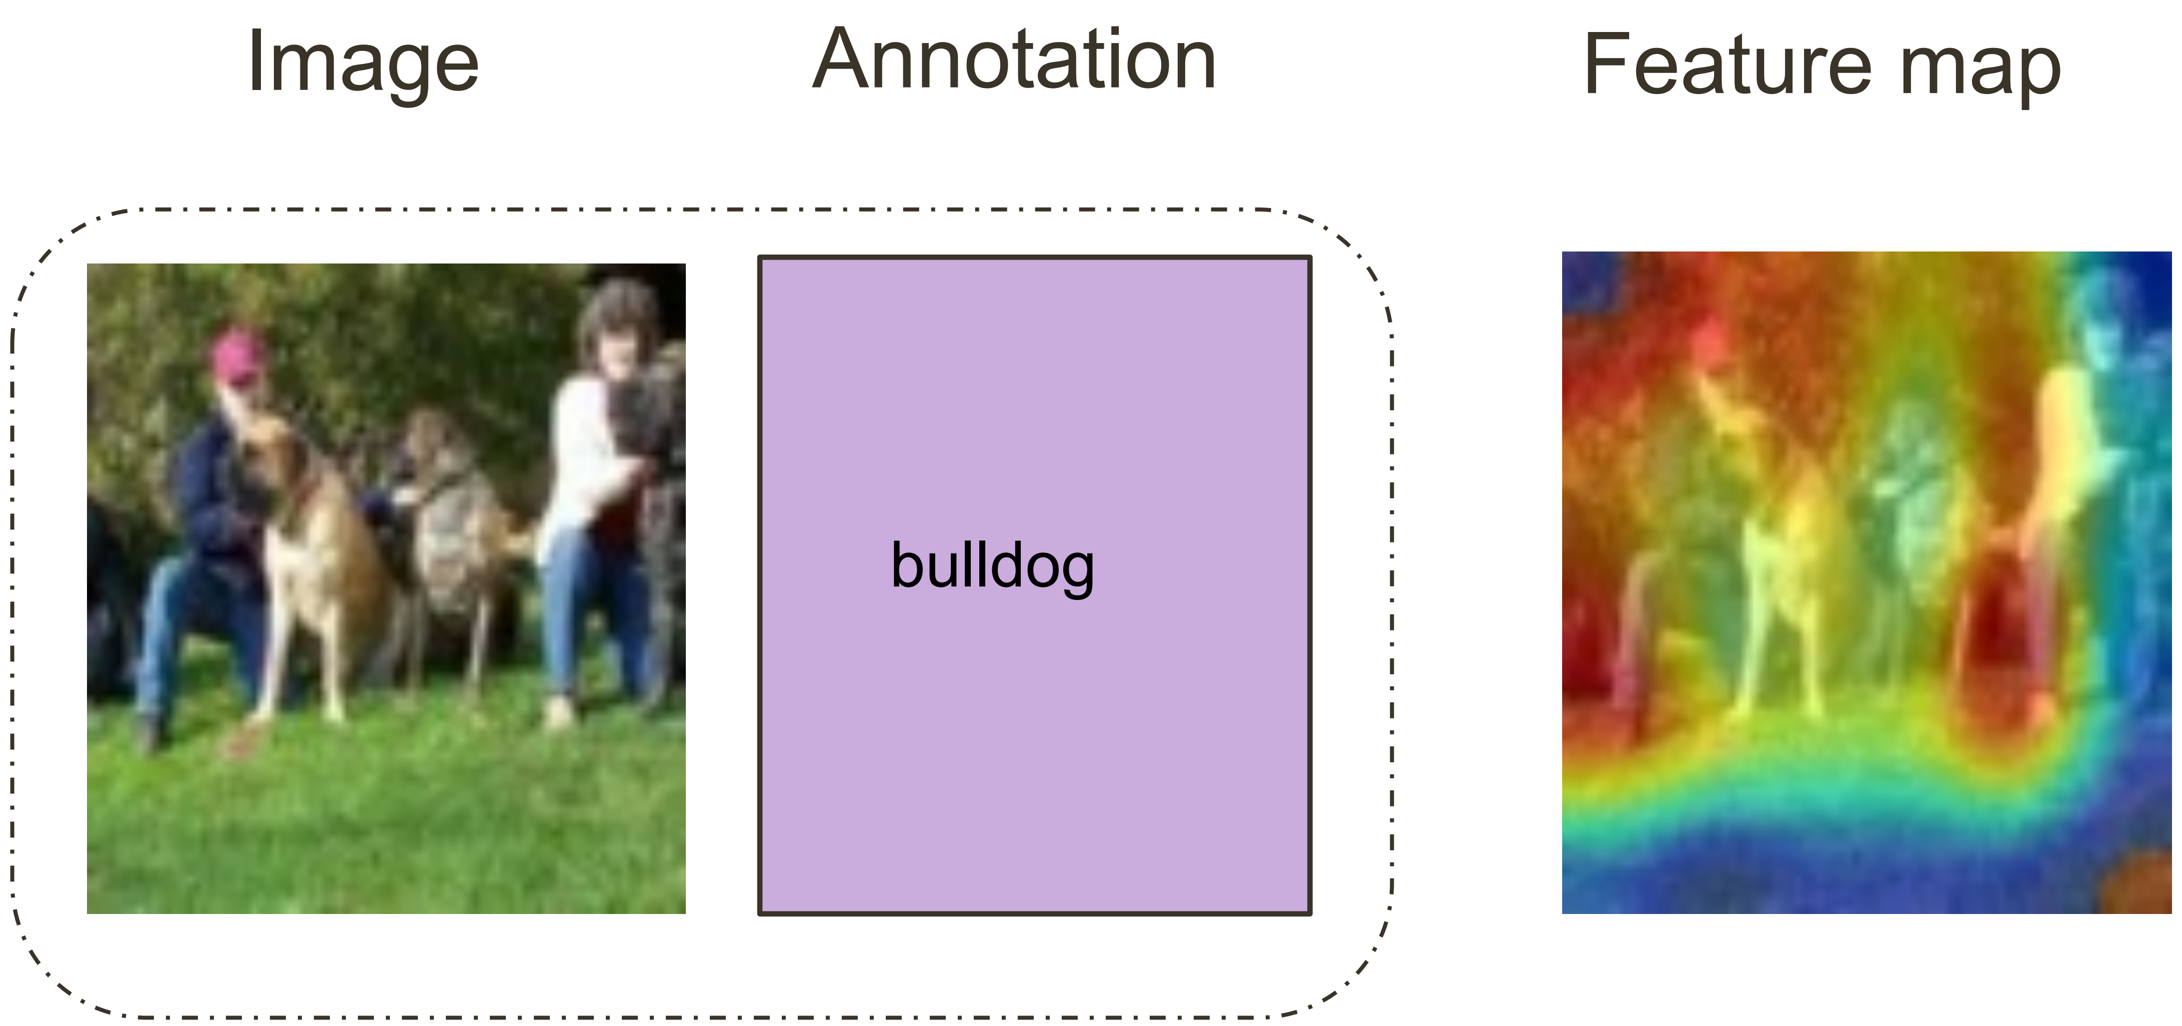

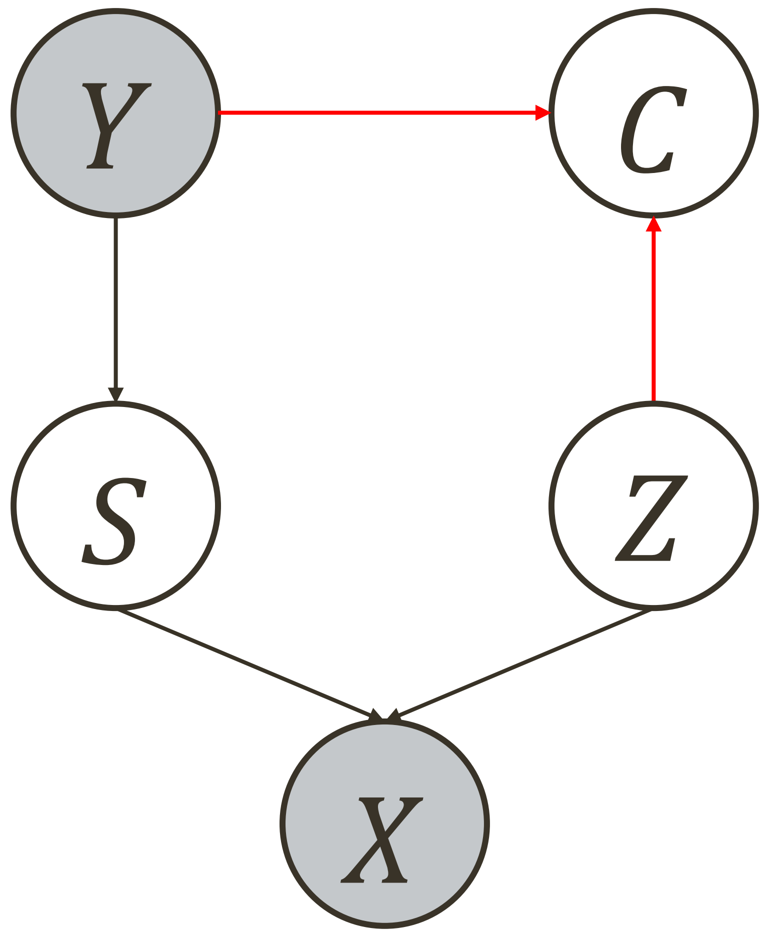

Class Y, Causal feature S and non-causal feature Z,
C=1 (training set)/0 (testing set)

The data selection bias can induce spurious correlation between causal and non-causal features.

Xu, et al. "PatchMix Augmentation to Identify Causal Features in Few-shot Learning." TPAMI2022.

Structural Causal Model in FSL: data selection perspective

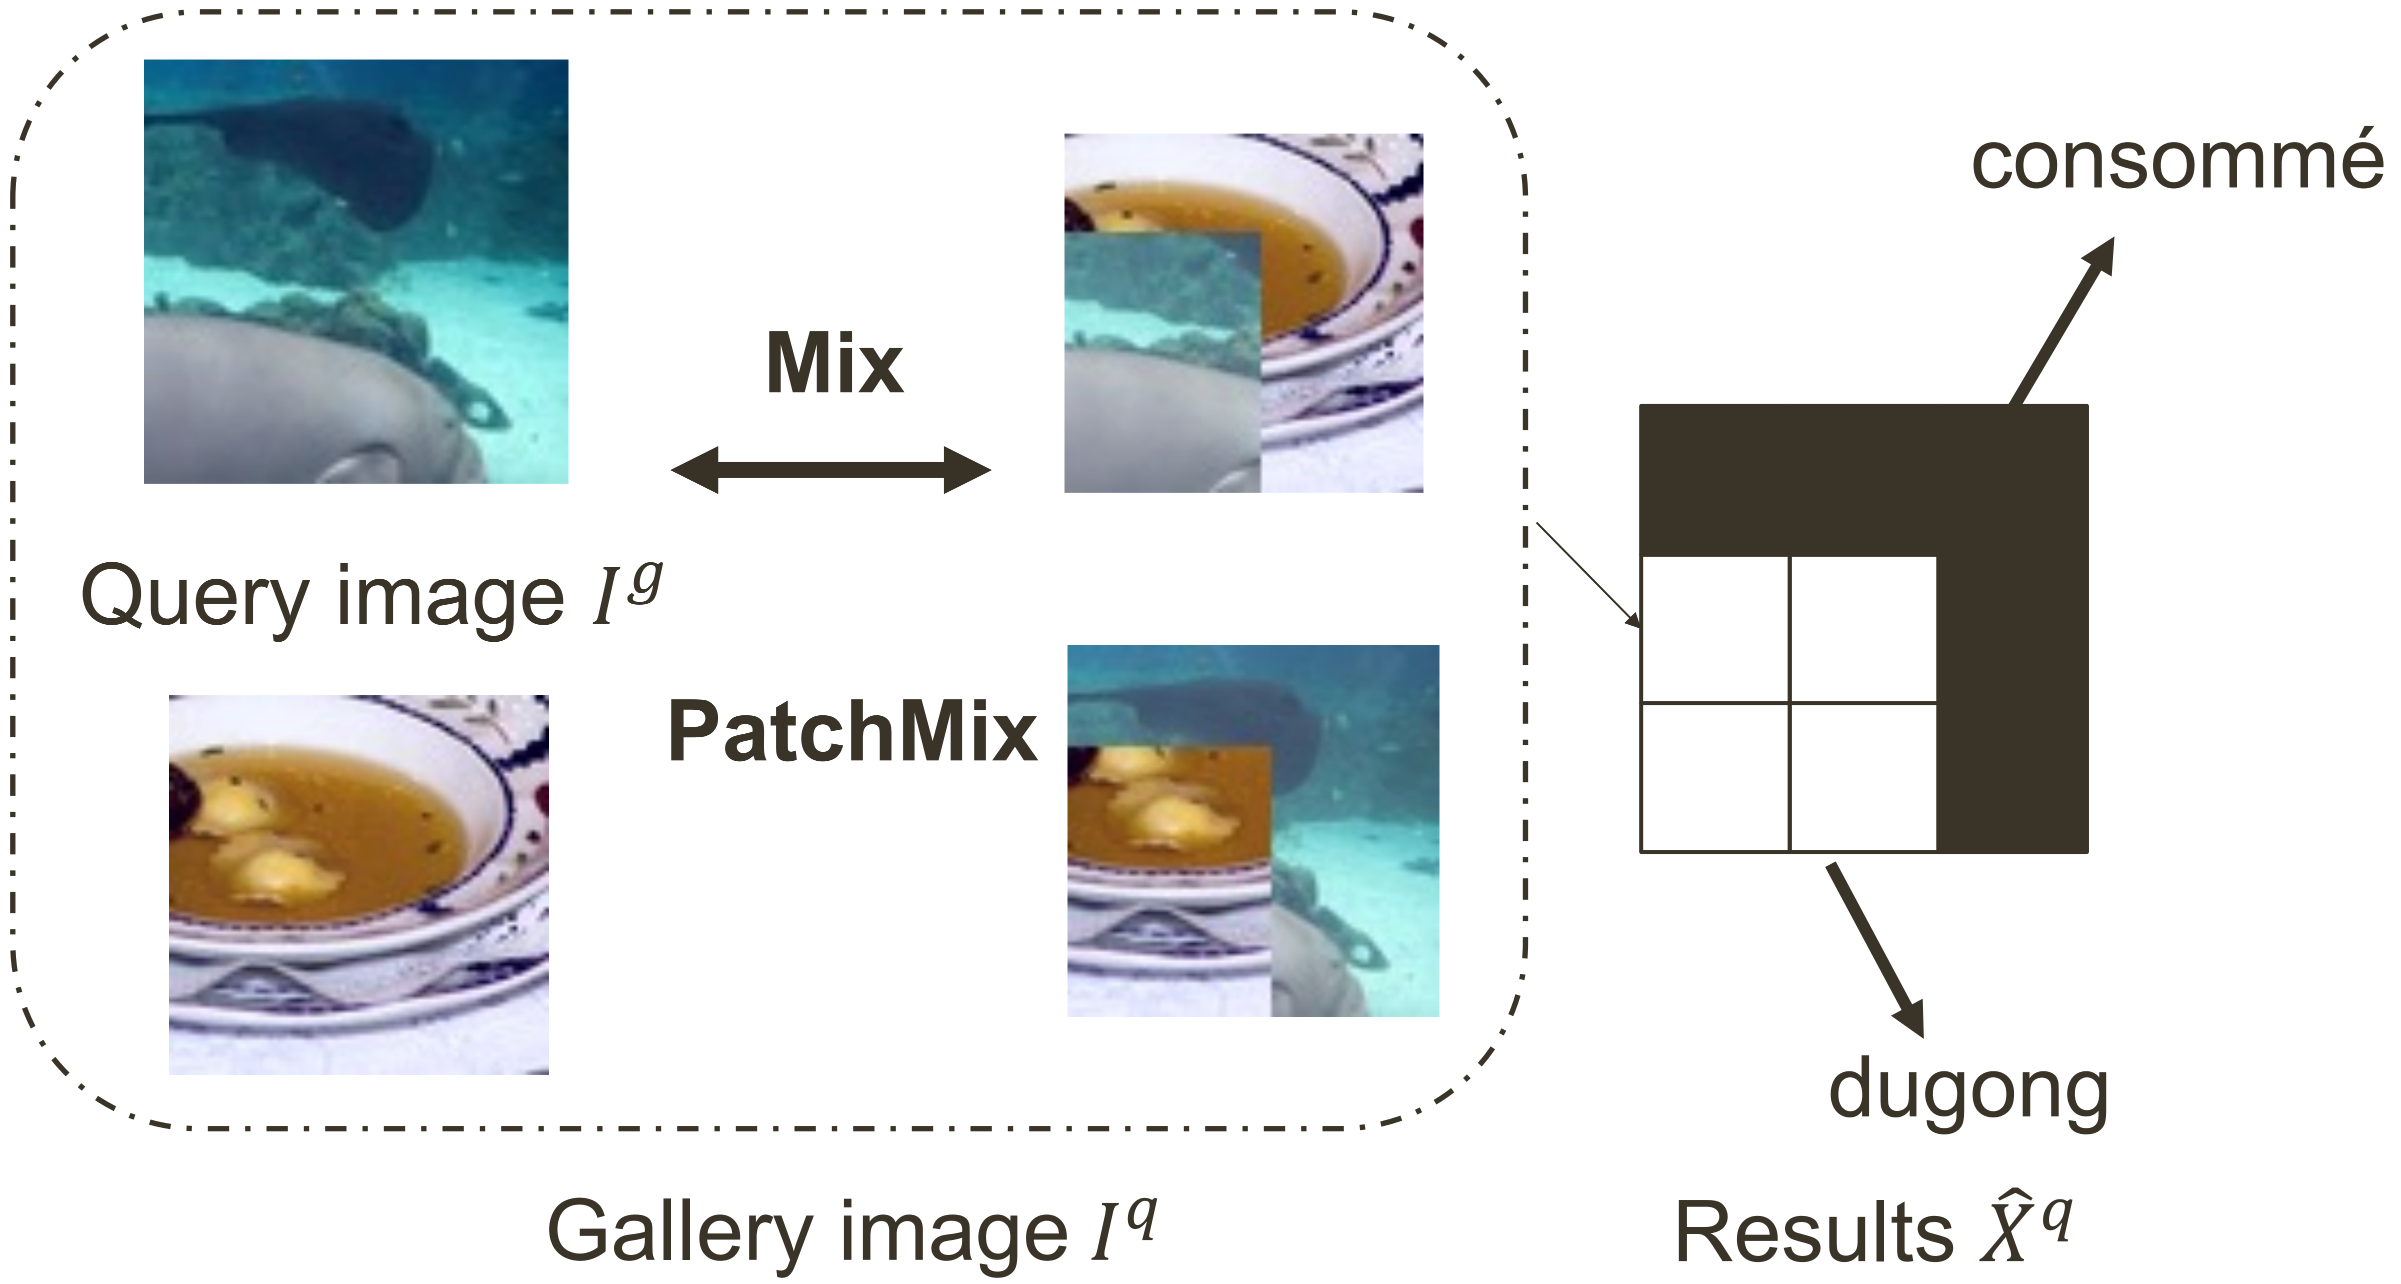

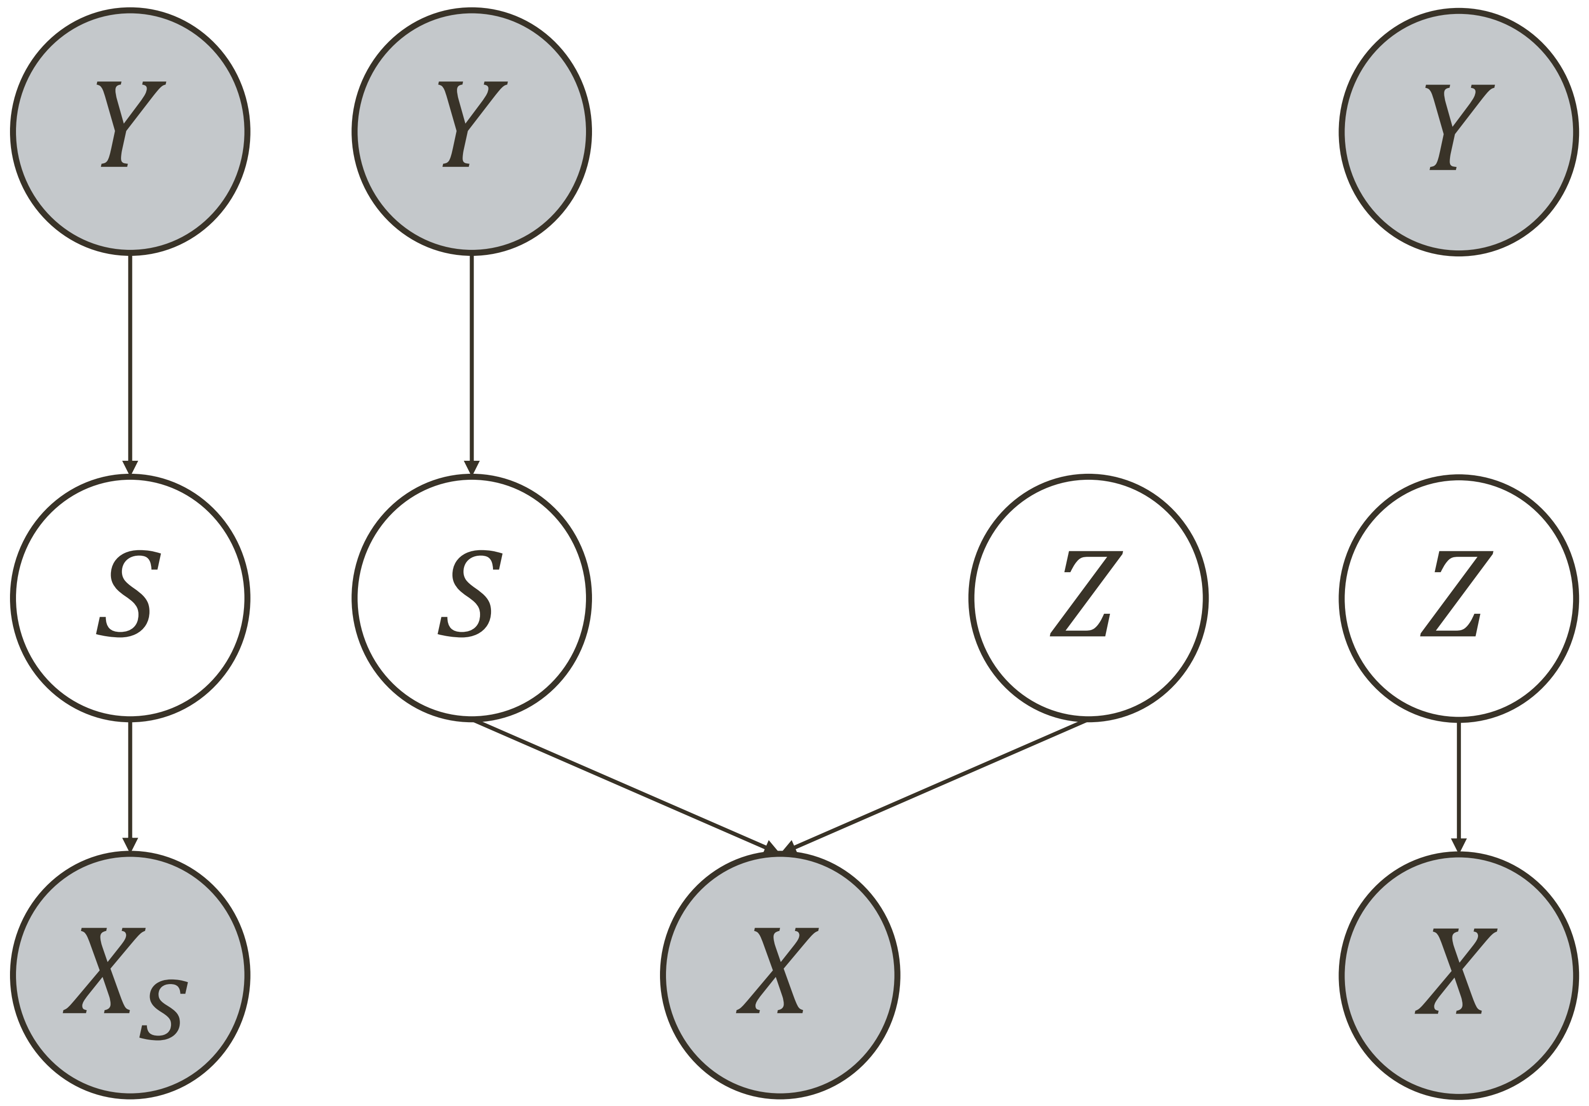

Solving spurious correlation problem by exchanging visual & label information between images.

Xu et al. "PatchMix Augmentation to Identify Causal Features in Few-shot Learning." TPAMI2022.

The Key Idea of “Learning better metrics”

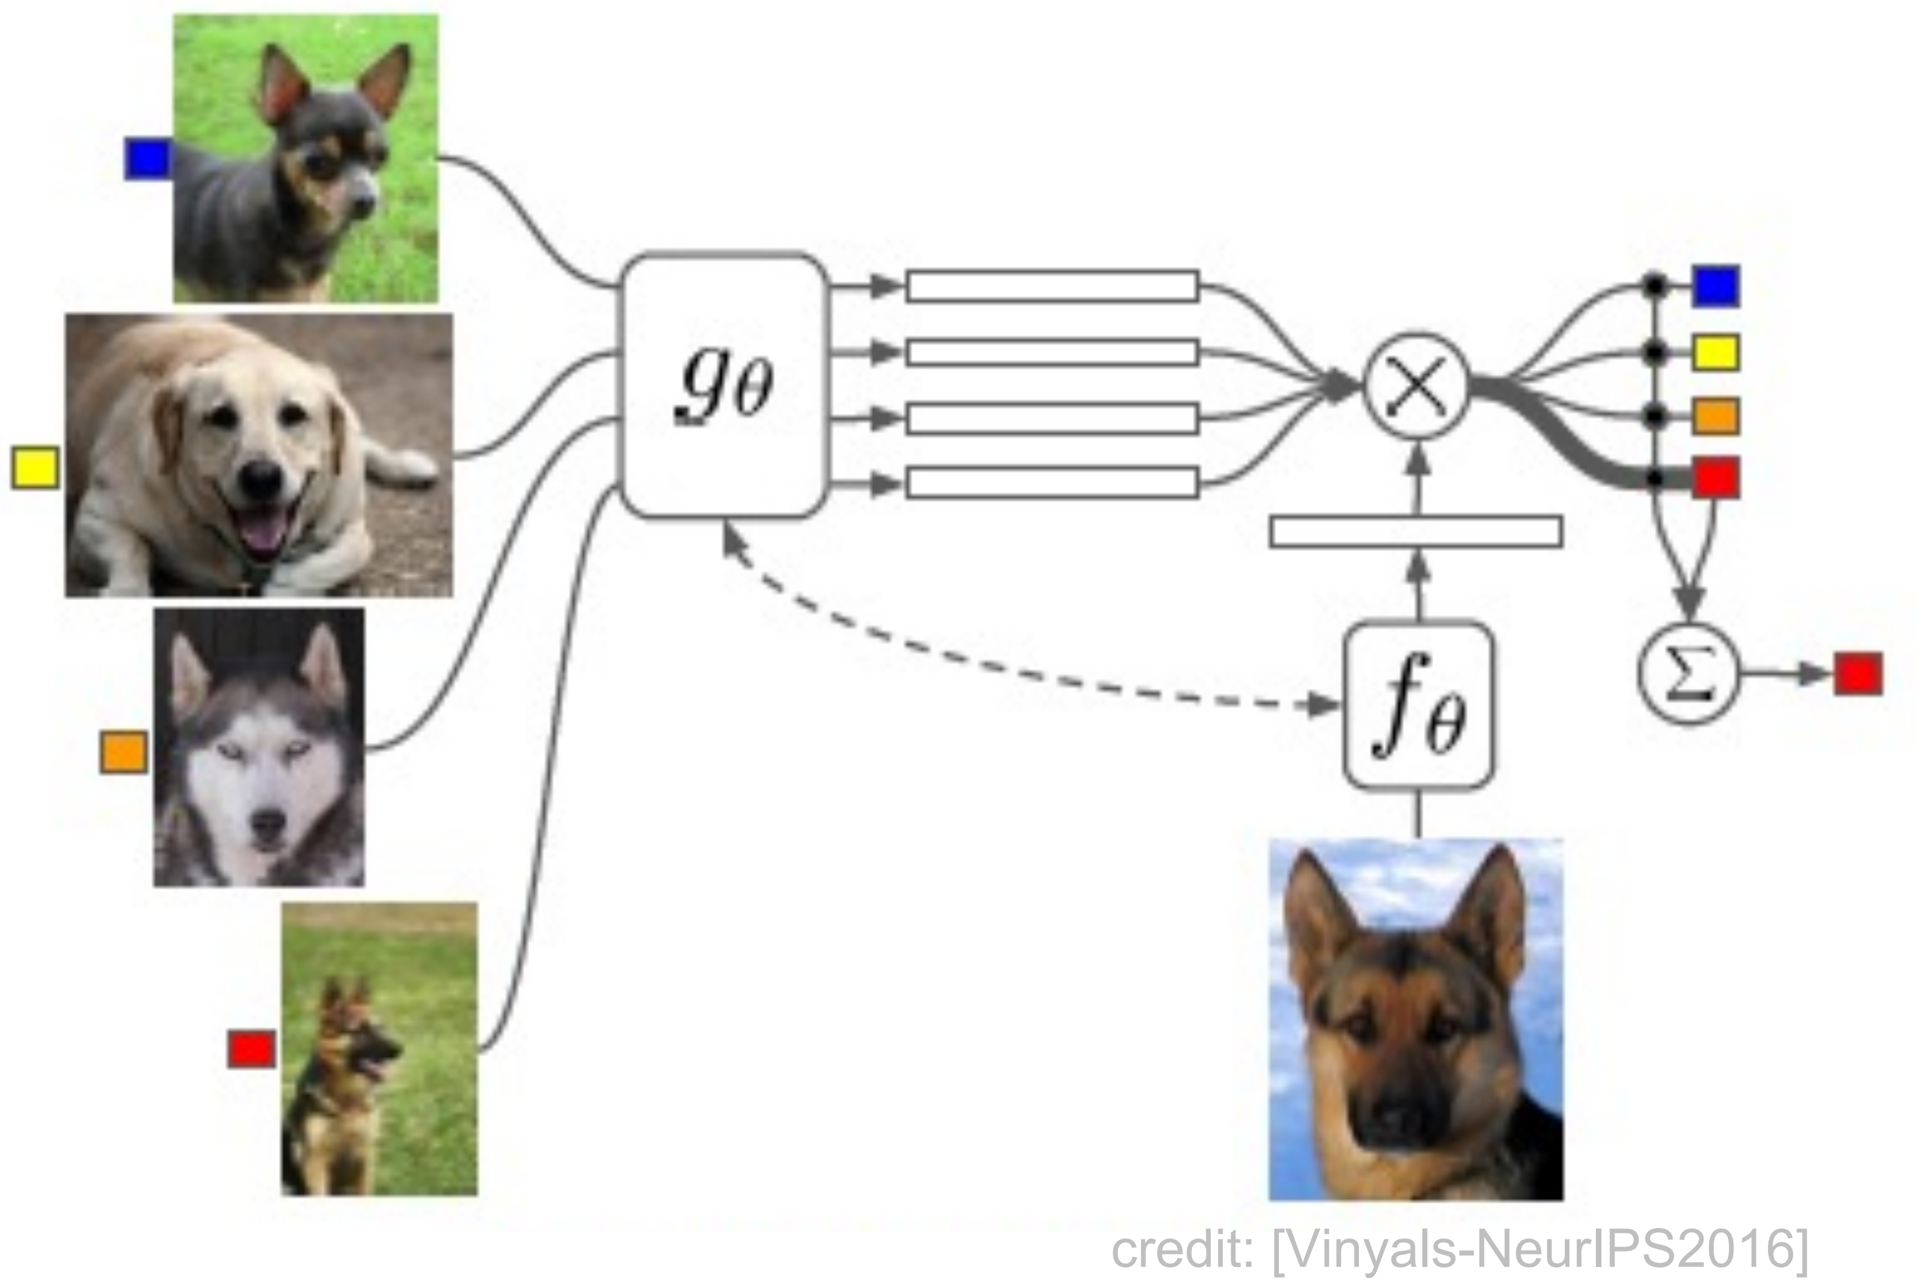

Matching Network

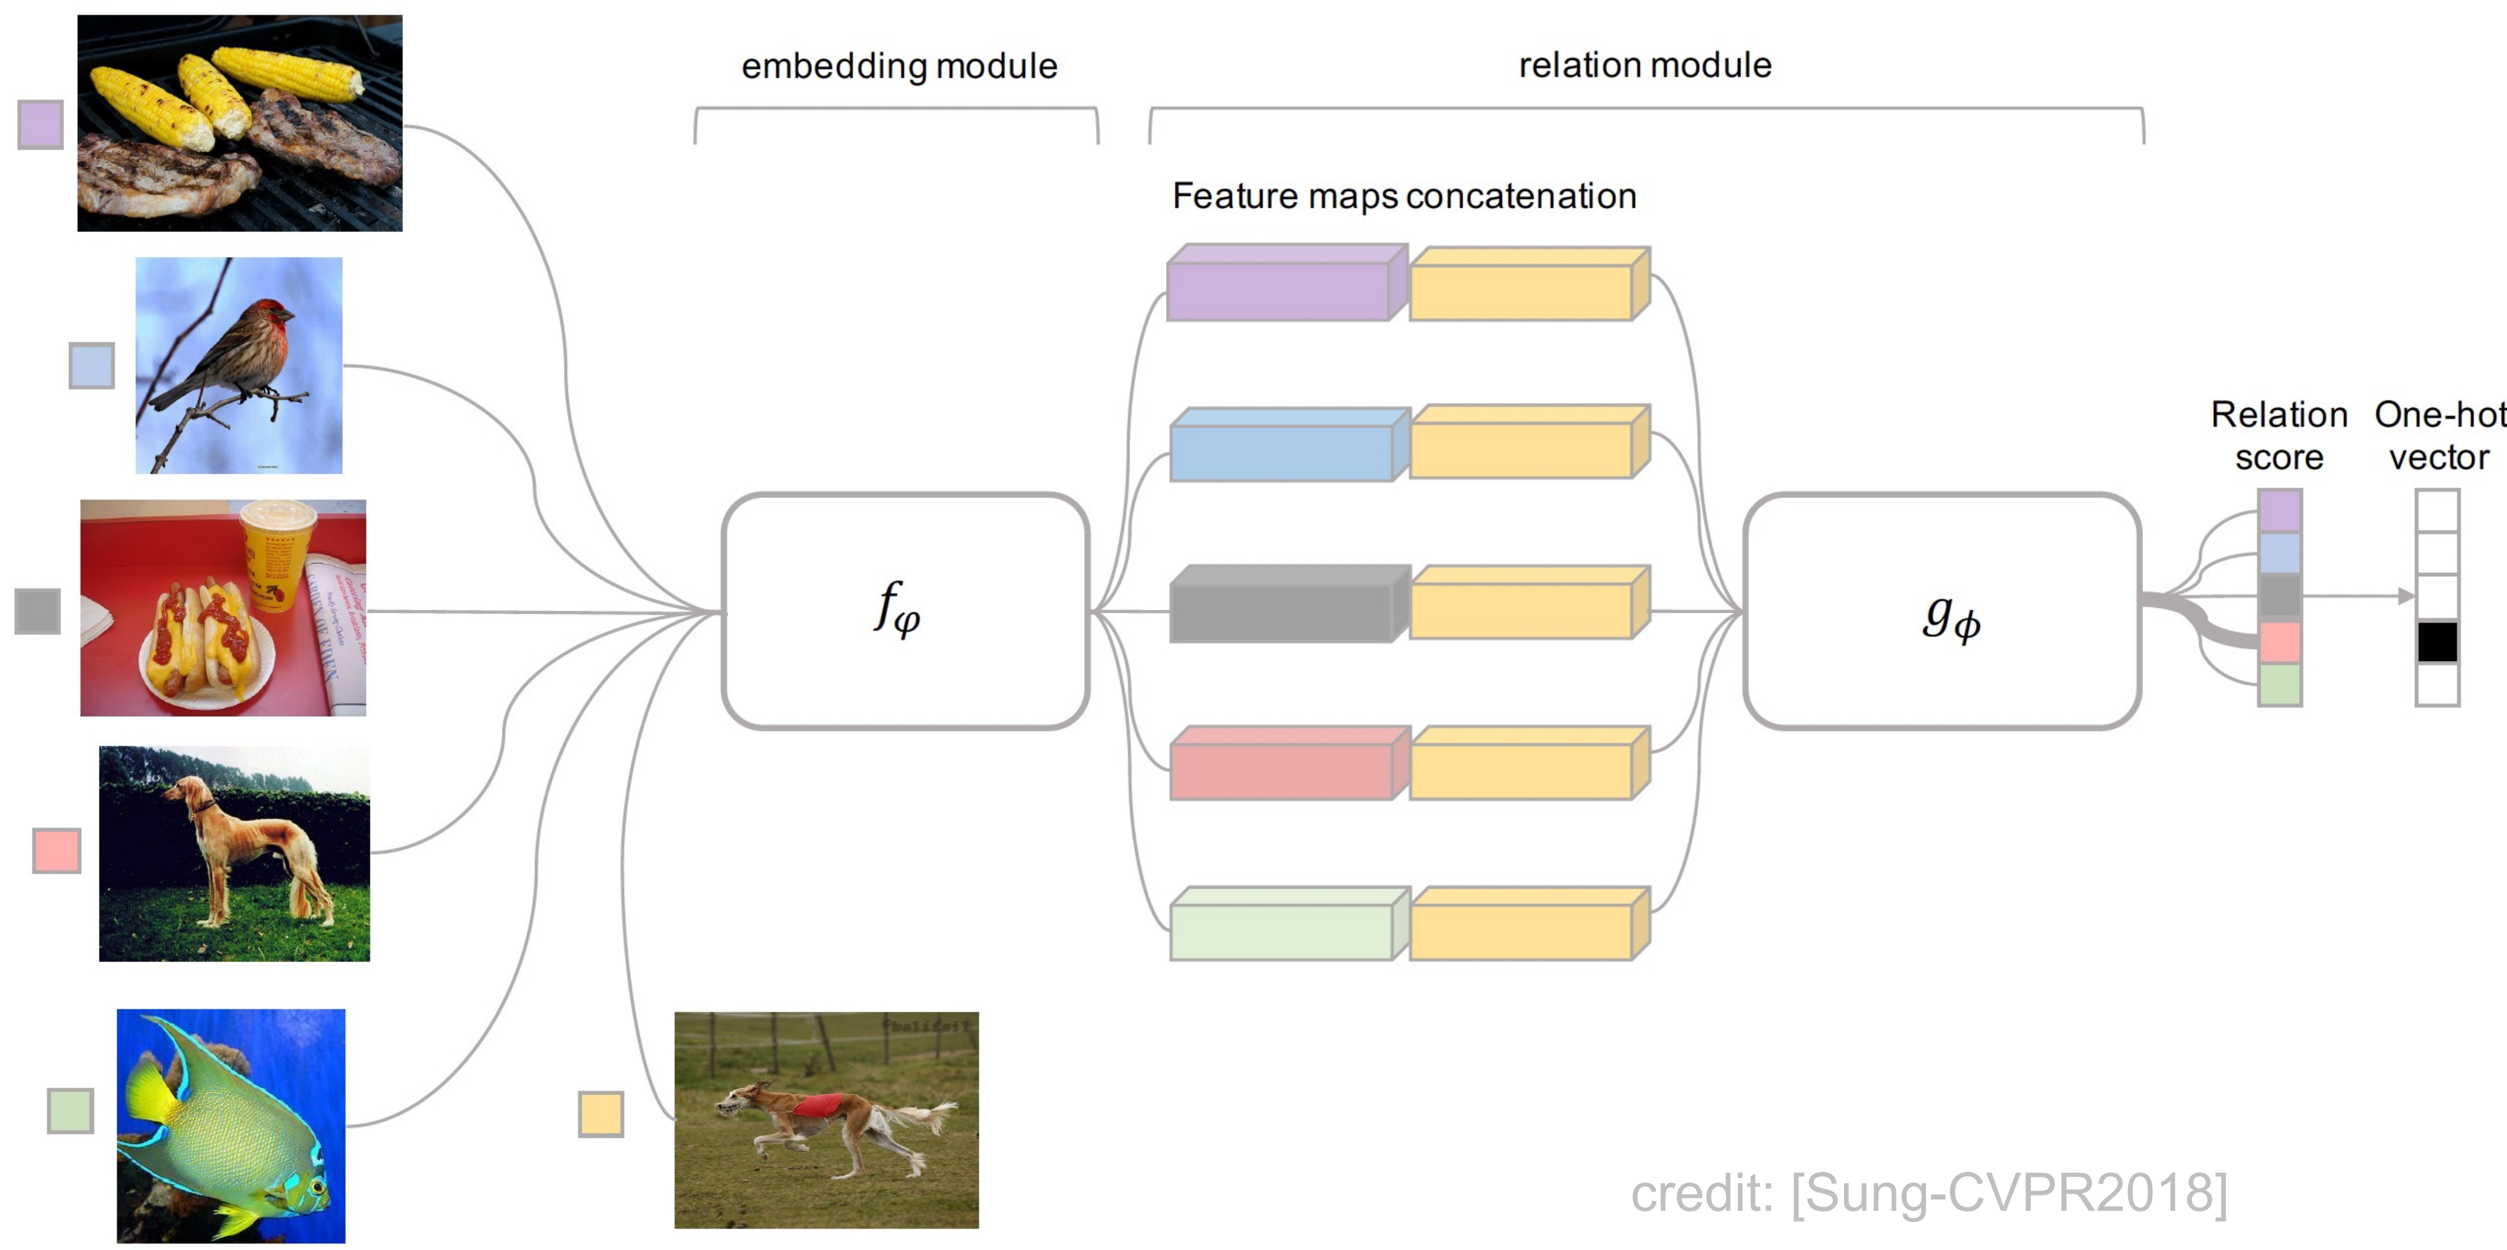

Relation Network

Can we learn the better metrics by probability distribution?

Vinyals et al. “Matching Networks for One Shot Learning.” NeurIPS 2016.
Sung et al. ” Learning to Compare: Relation Network for Few-Shot Learning”, CVPR 2018.

Earth Mover's Distance

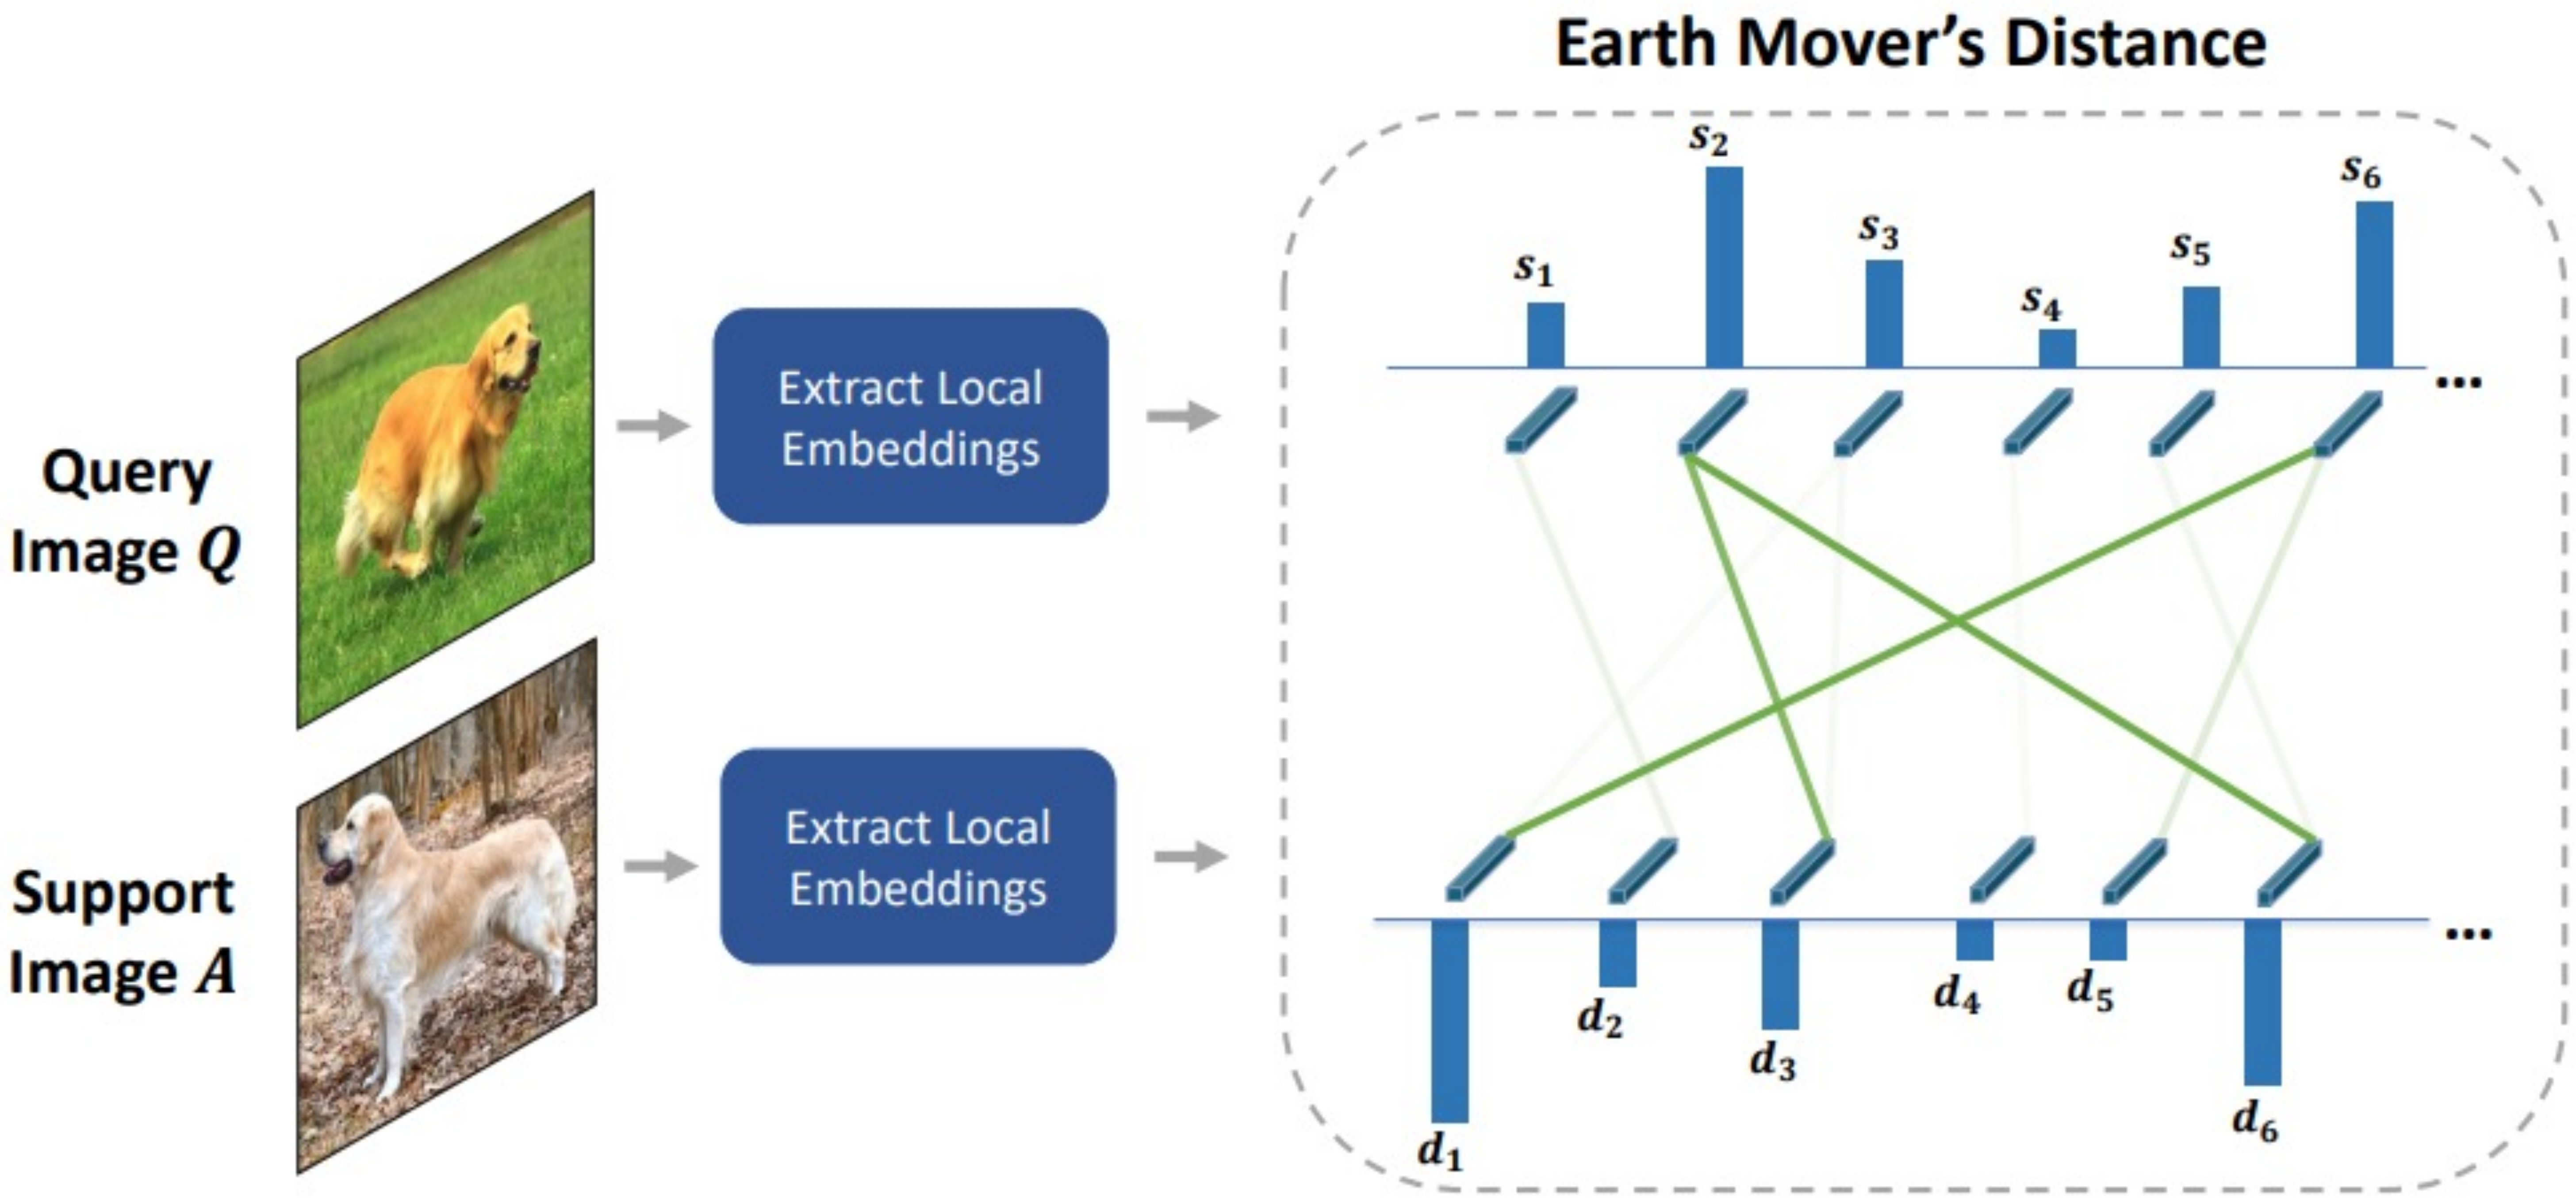

$$distance(\boldsymbol{Q}, \boldsymbol{A}) = \sum_{ij} \tilde{x}_{ij} c_{ij}$$

$\widetilde{\boldsymbol{X}}$: values of optimal matching flows
c_{ij}: matching cost between vector $\boldsymbol{u}_i$ and $\boldsymbol{v}_j$
s_i : weight of vector $\boldsymbol{u}_i$
d_j : weight of vector $\boldsymbol{v}_j$

credit: [Zhang-CVPR2020]

Zhang et al. "Deepemd: Few-shot image classification with differentiable earth mover's distance and structured classifiers." CVPR2020.
Zhang, et al. "Deepemd: Differentiable earth mover's distance for few-shot learning." TPAMI2022.

Brownian Distance Covariance in FSL

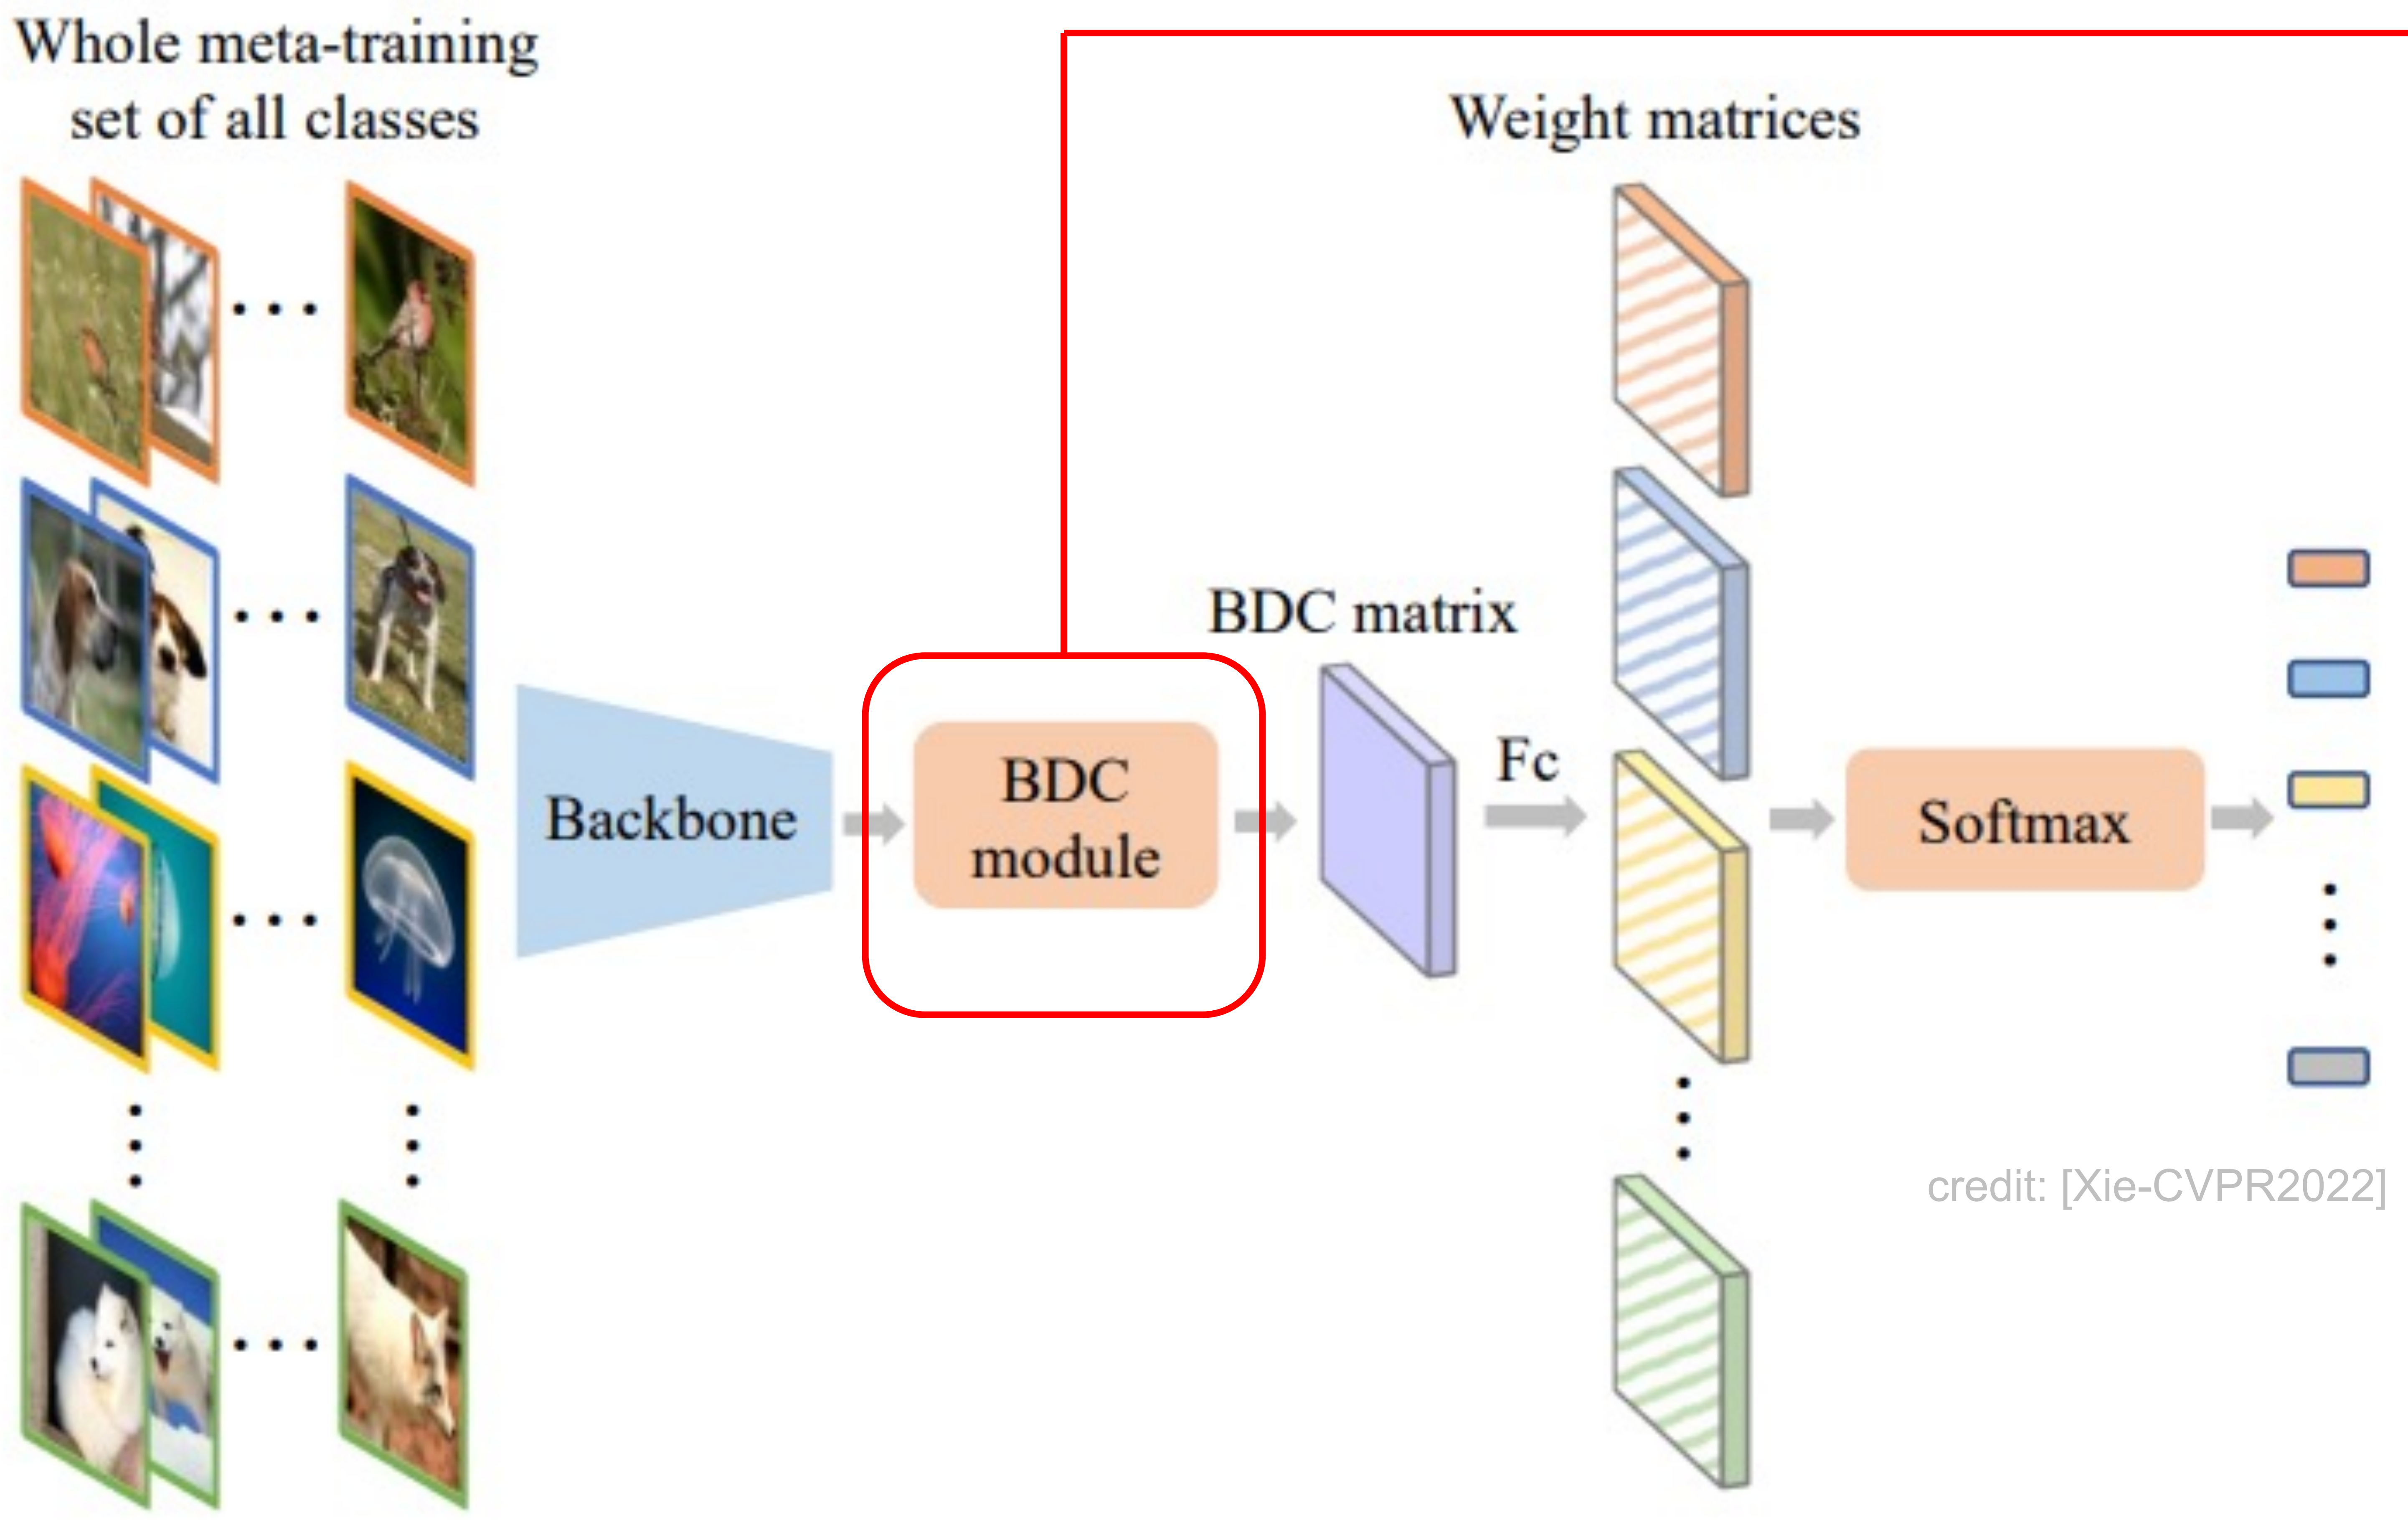

BDC metric ρ between two sets $\{(x_1, y_1), \dots, (x_m, y_m)\}$:

$$\hat{A} = (\hat{a}_{kl}), \hat{a}_{kl} = ||x_k - x_l||$$
$$\hat{B} = (\hat{b}_{kl}), \hat{b}_{kl} = ||y_k - y_l||$$

$\rho(X, Y) = tr(A^T B)$ after normalization

- Non-negative, 0 iff X, Y are independent
- Characterize linear and non-linear dependency
- Invariant to individual translations and orthonormal transformations

Xie et al. "Joint distribution matters: Deep brownian distance covariance for few-shot classification." CVPR2022.

Squared root of the Euclidean distance and the Norm distance for dissimilarity measurement

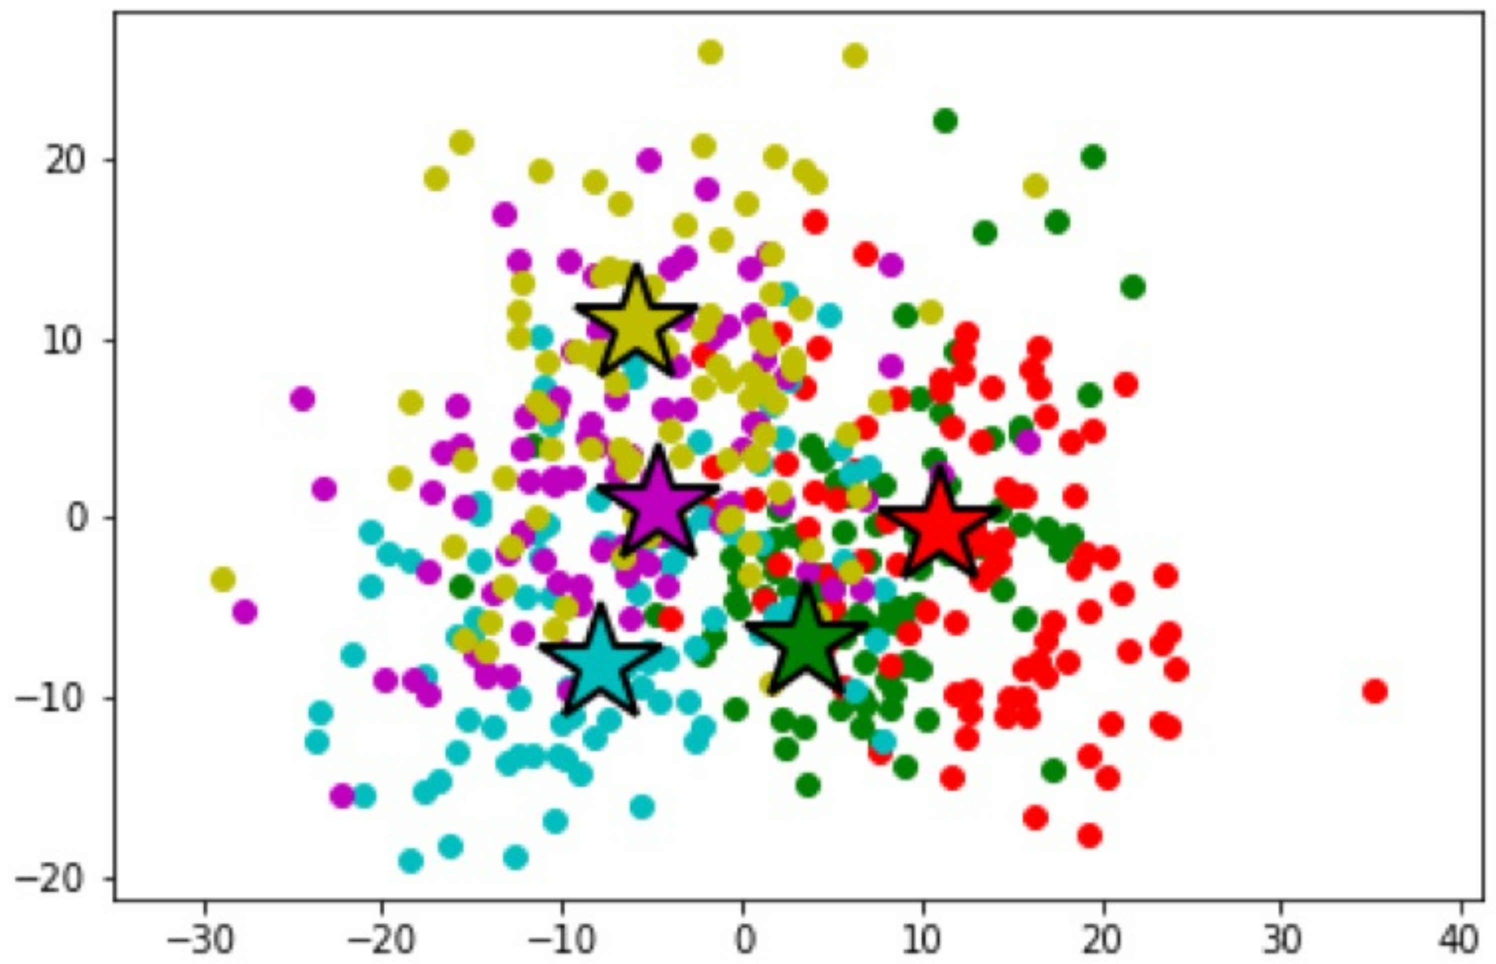

Prototypical Network

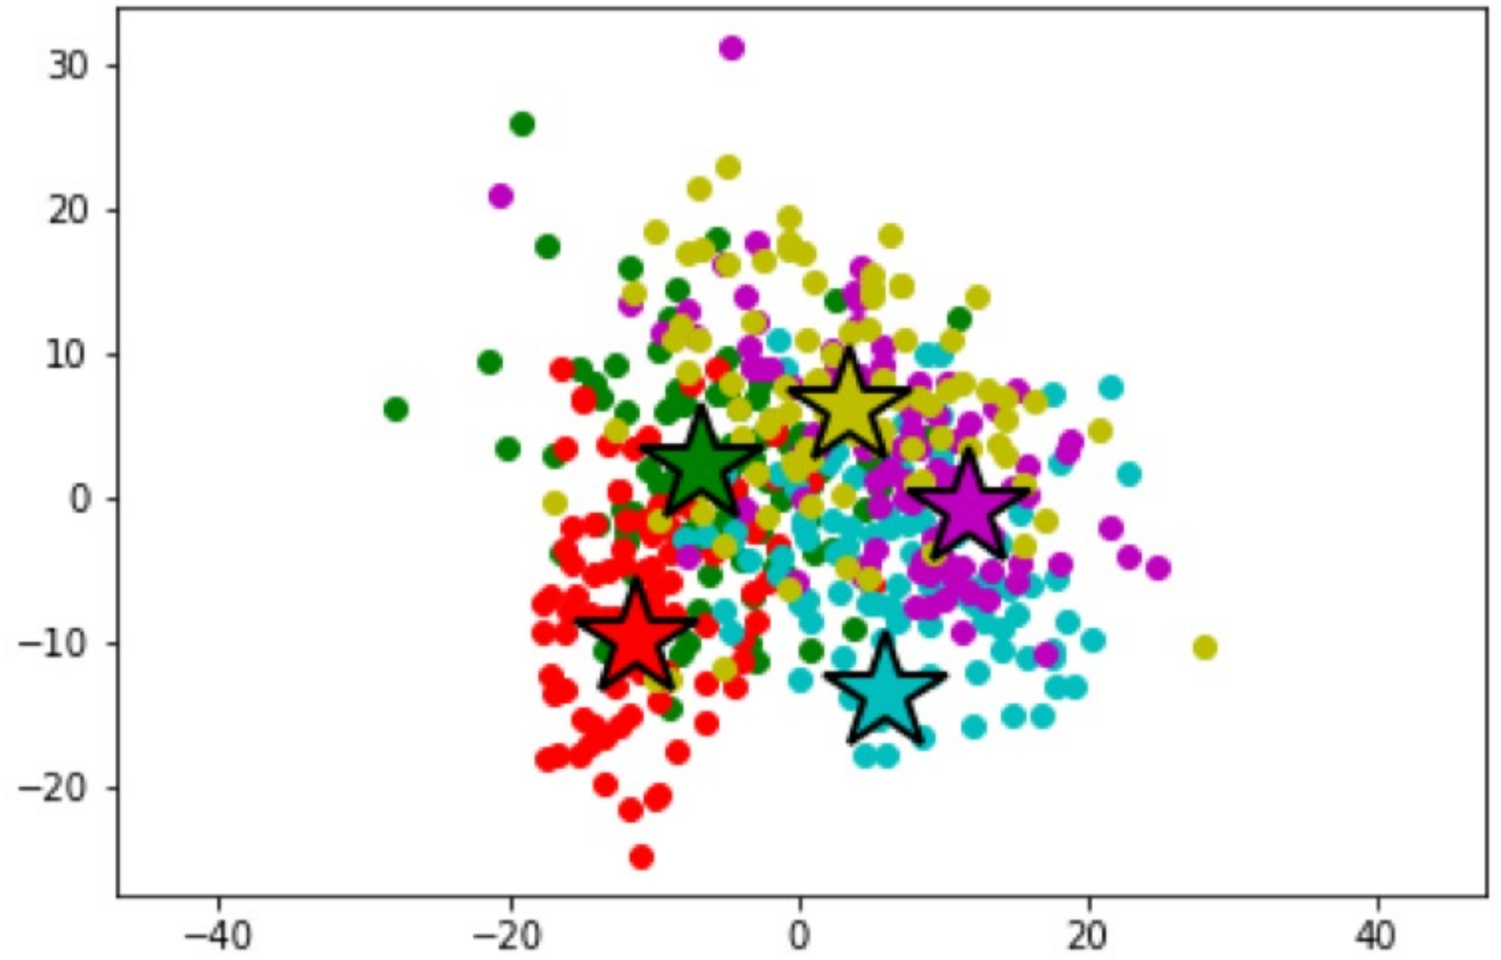

Prototypical Network with ring loss

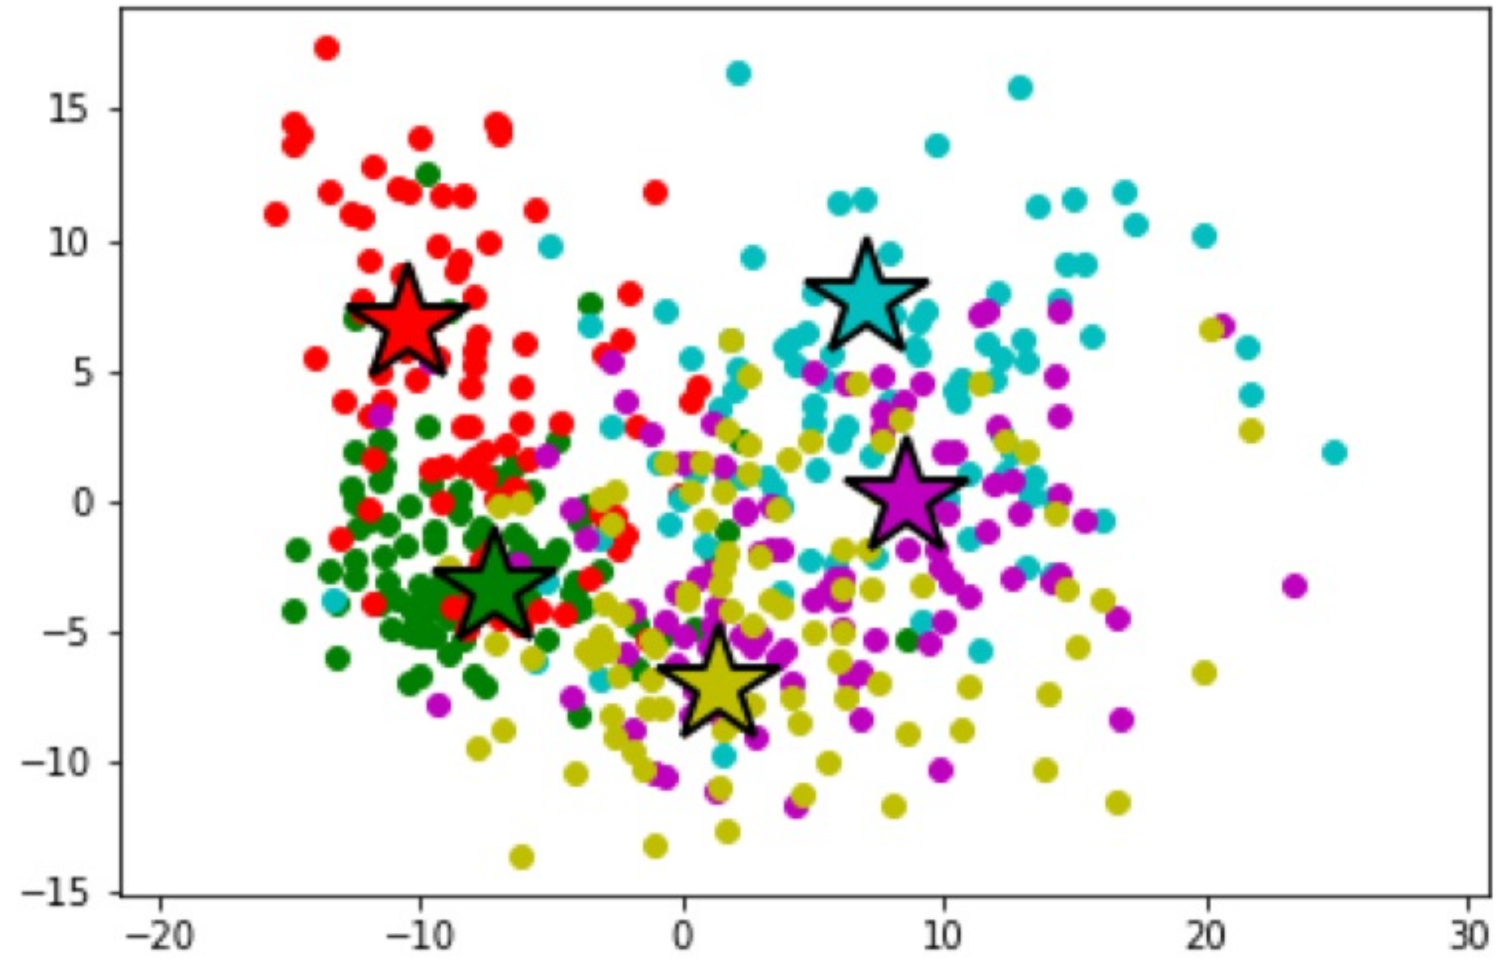

SEN

credit: [Nguyen-ECCV2020]

SEN dissimilarity between query feature z and prototype c:

$$d_s(z,c) = \sqrt{\mathrm{d}_e(z,c) + \epsilon d_n(z,c)}\,, \mathrm{d}_e(z,c) = ||z-c||^2, d_n(z,c) = \left(||z|| - ||c||\right)^2$$

Nguyen et al. "Sen: A novel feature normalization dissimilarity measure for prototypical few-shot learning networks." ECCV 2020.

The Key Idea of “Neural Collapse”

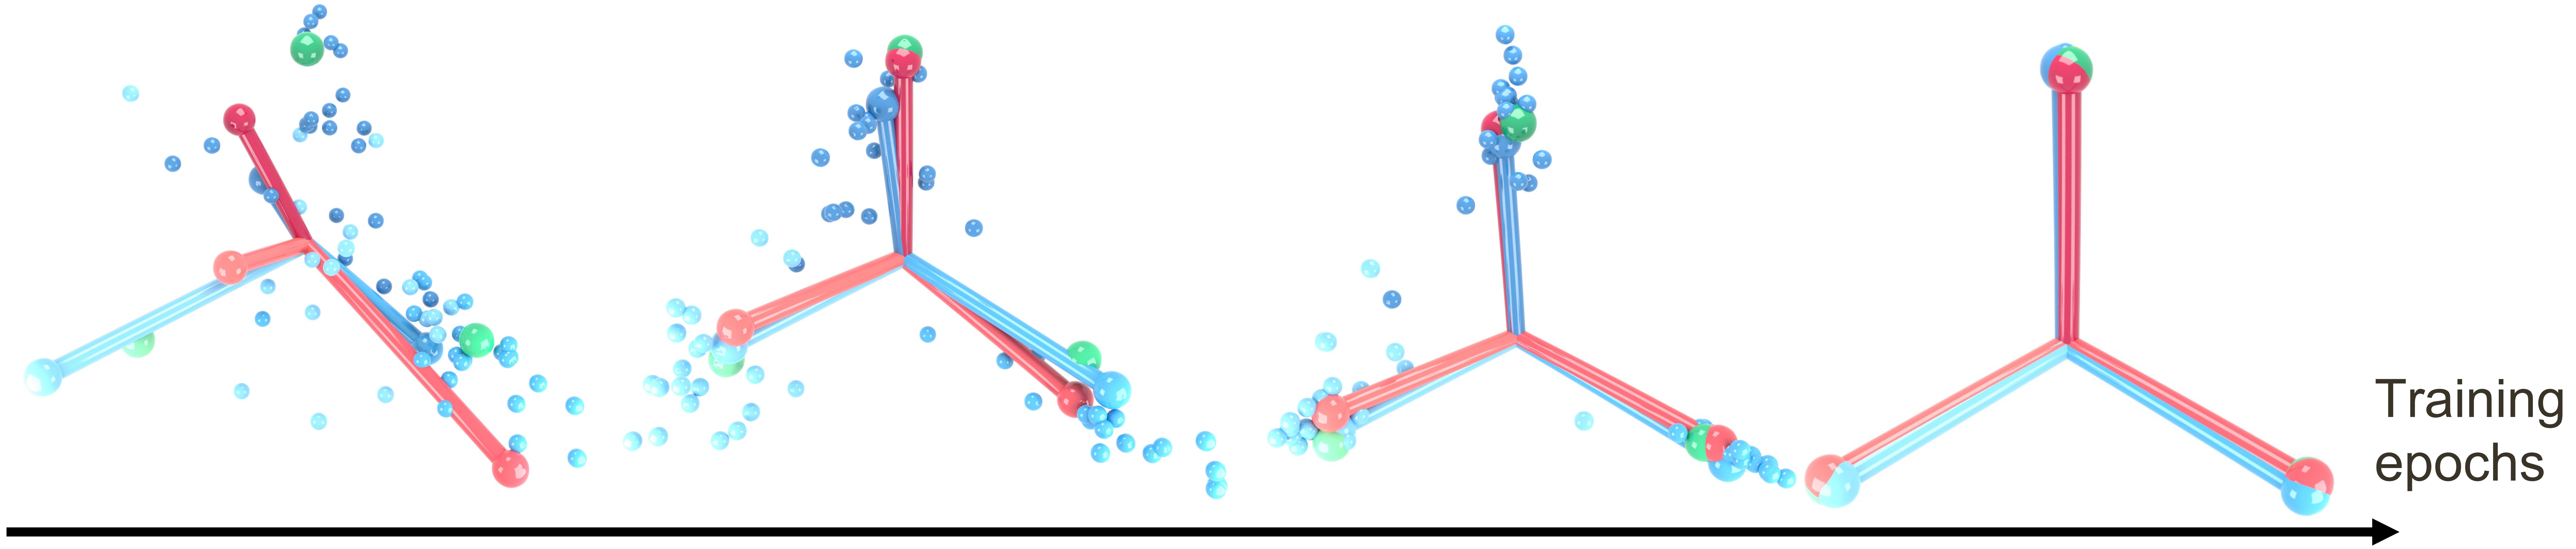

credit: [Han-PNAS2020]
Green spheres: vertices of Simplex ETF
Red ball-and-sticks: linear classifiers
Blue ball-and-sticks: class-means
Small blue spheres: last-layer features.

Neural Collapse is characterized by four manifestations in the classifier and last-layer activations

- NC1: Intra-class variation collapse to 0
- NC2: Class centers converge to simplex ETF
- NC3: Linear classifiers converge to class centers
- NC4: Classifier acts like nearest class center

Papyan, Han, and Donoho. "Prevalence of neural collapse during the terminal phase of deep learning training." PNAS2020.

NC in FSL: Within-class Variation Collapse

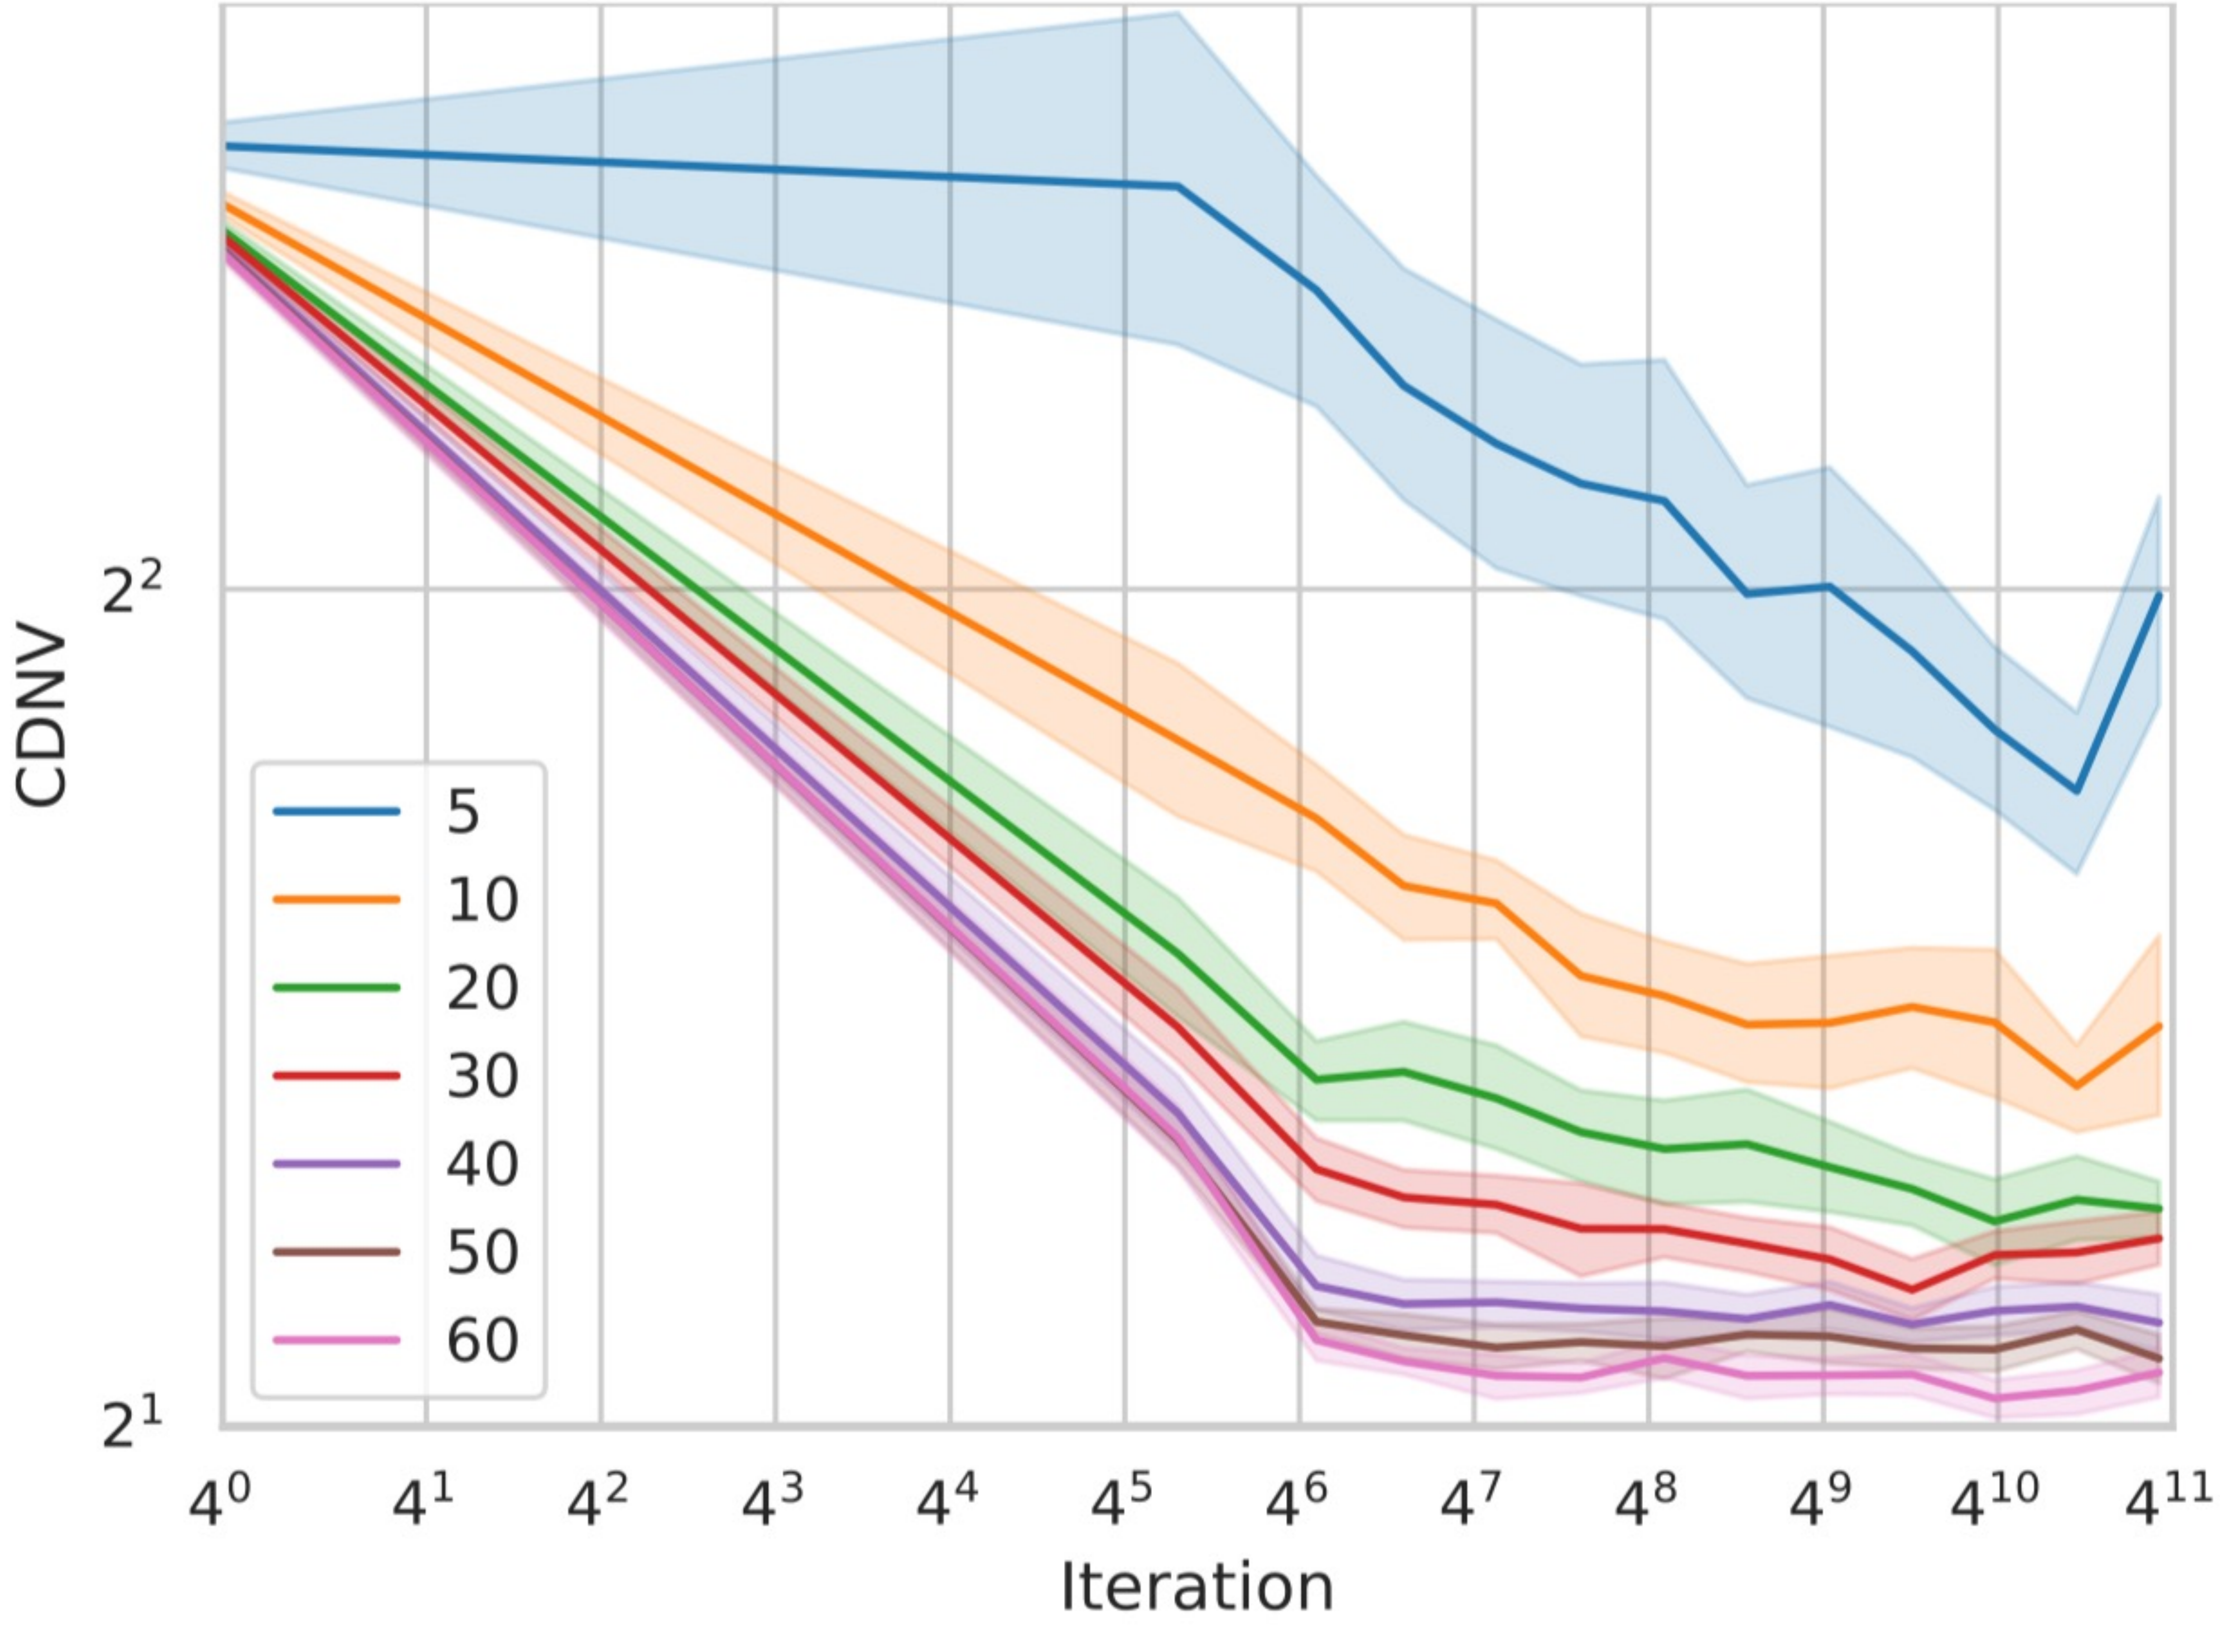

(c) target CDNV

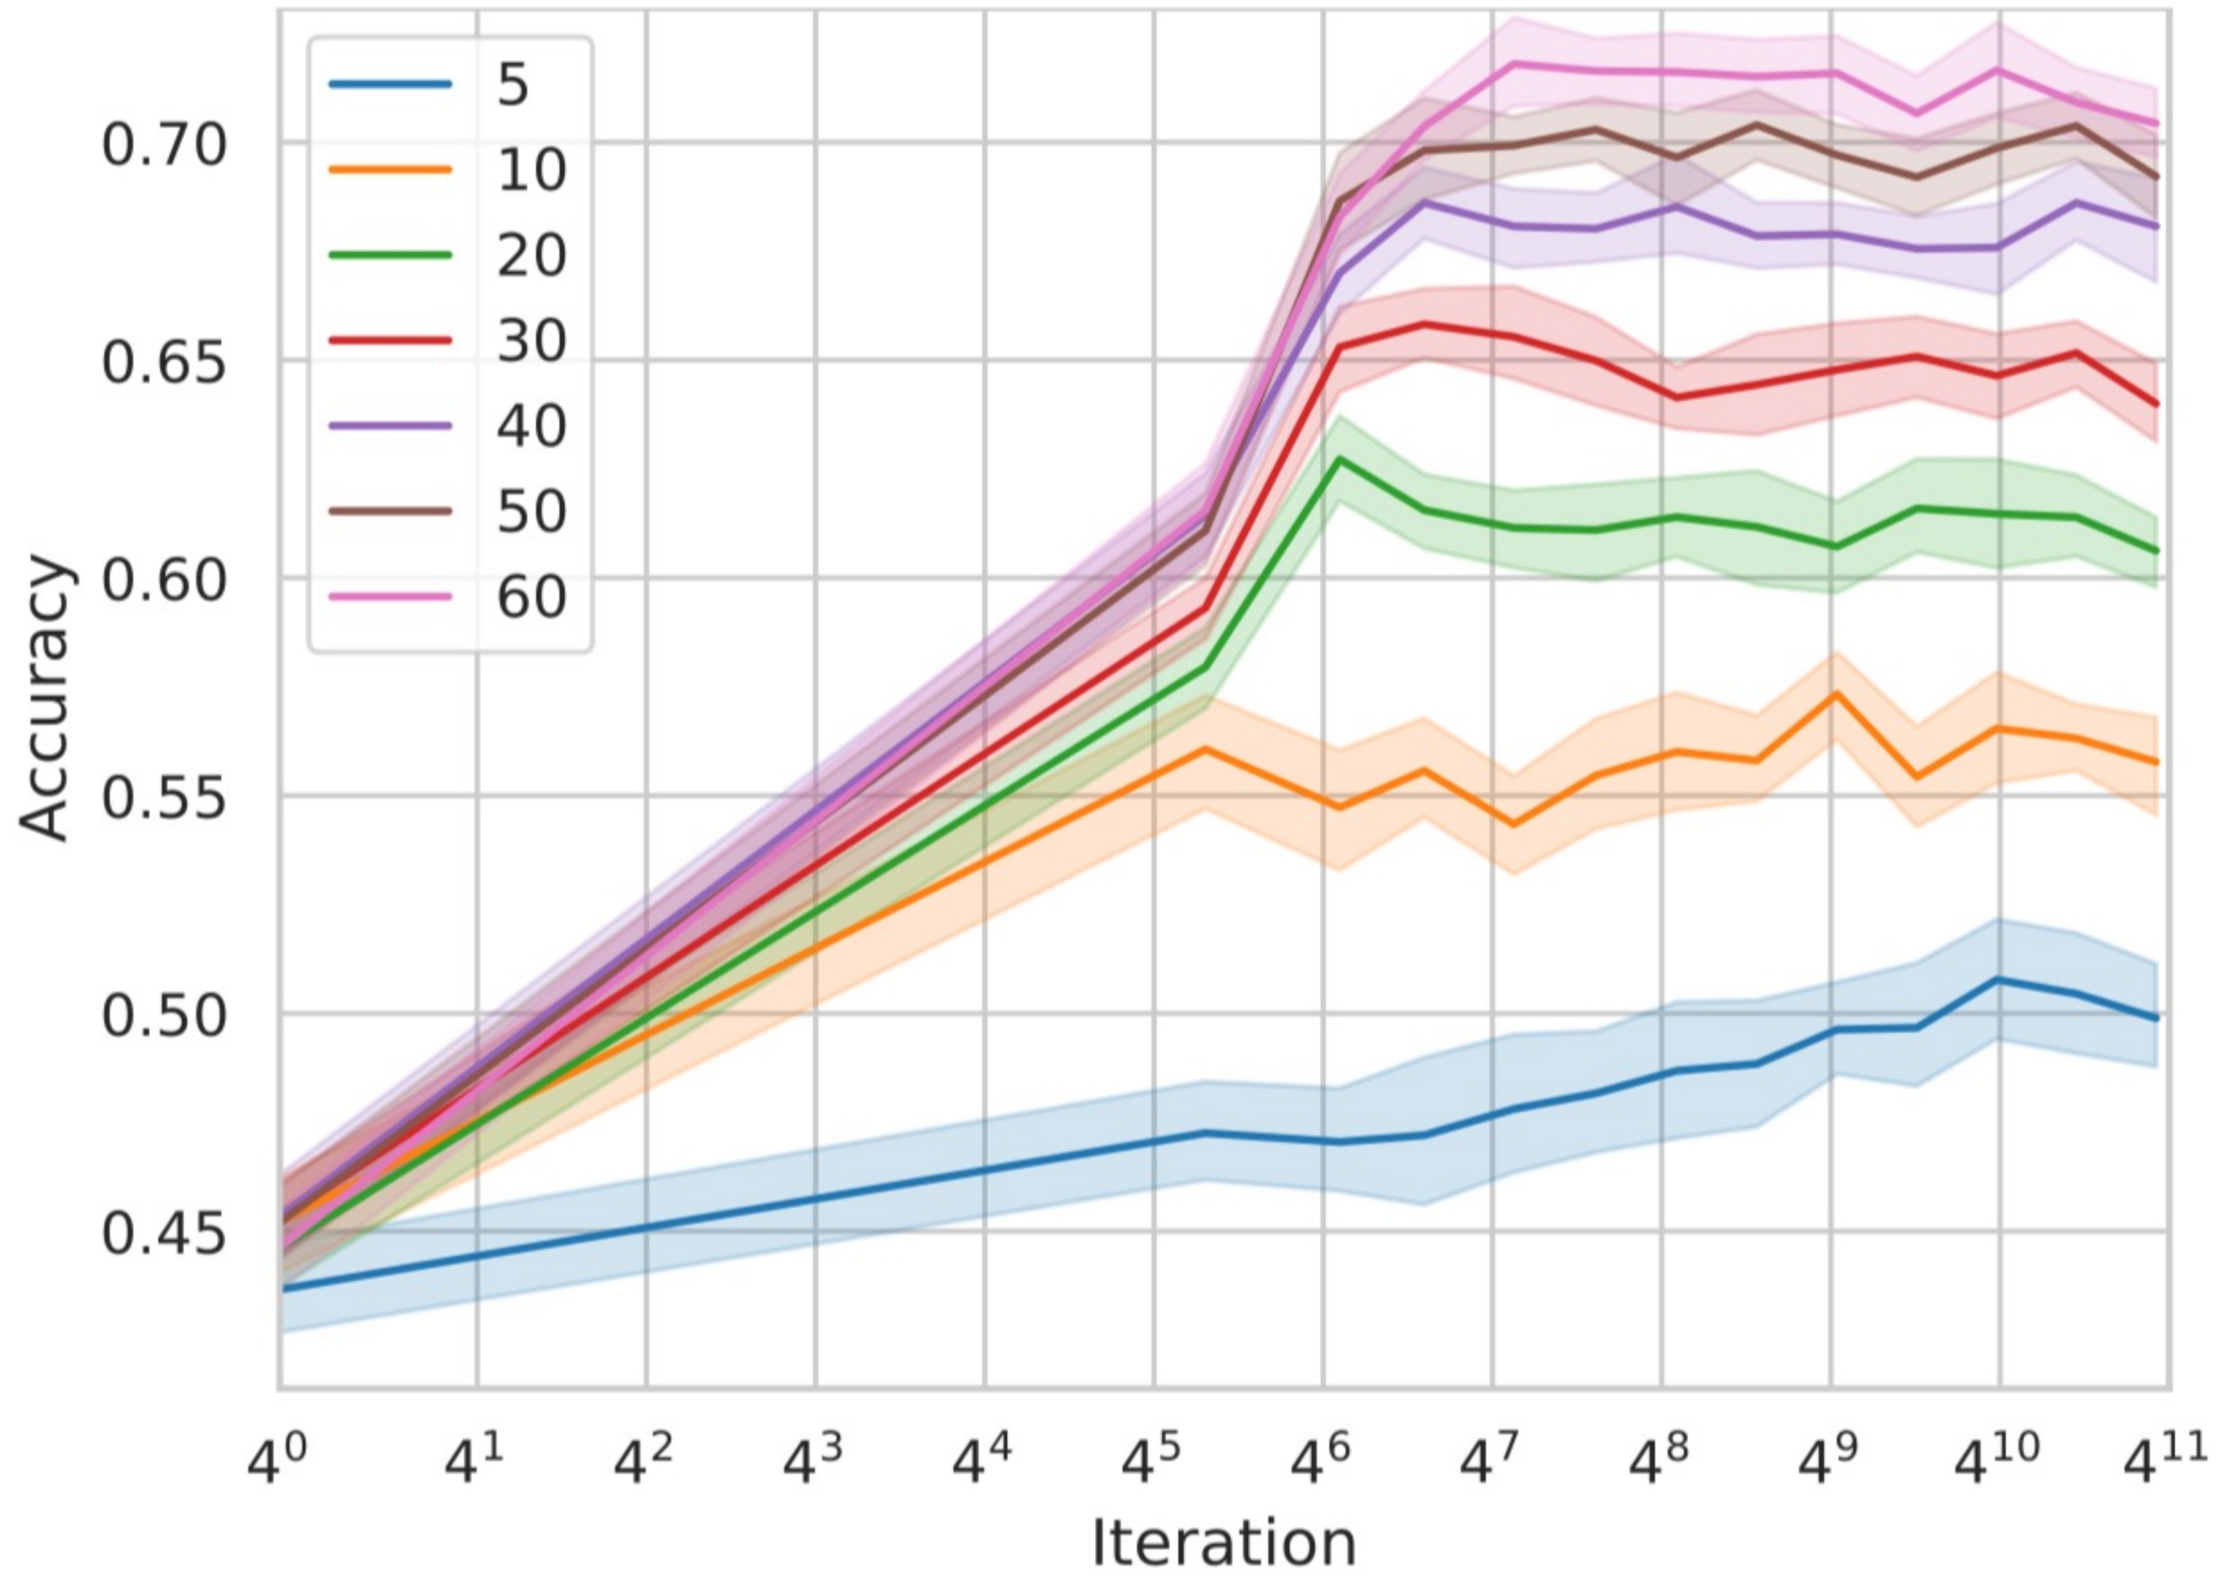

credit: [Galanti-ICLR2021]

(d) target accuracy

CDNV: $V_f(Q_1, Q_2) = \frac{Var_f(Q_1)+Var_f(Q_2)}{2||\mu_f(Q_1)-\mu_f(Q_2)||^2}$

Galanti, Tomer, András György, and Marcus Hutter. "On the role of neural collapse in transfer learning." ICLR2021.

NC in imbalanced learning: stronger regularization

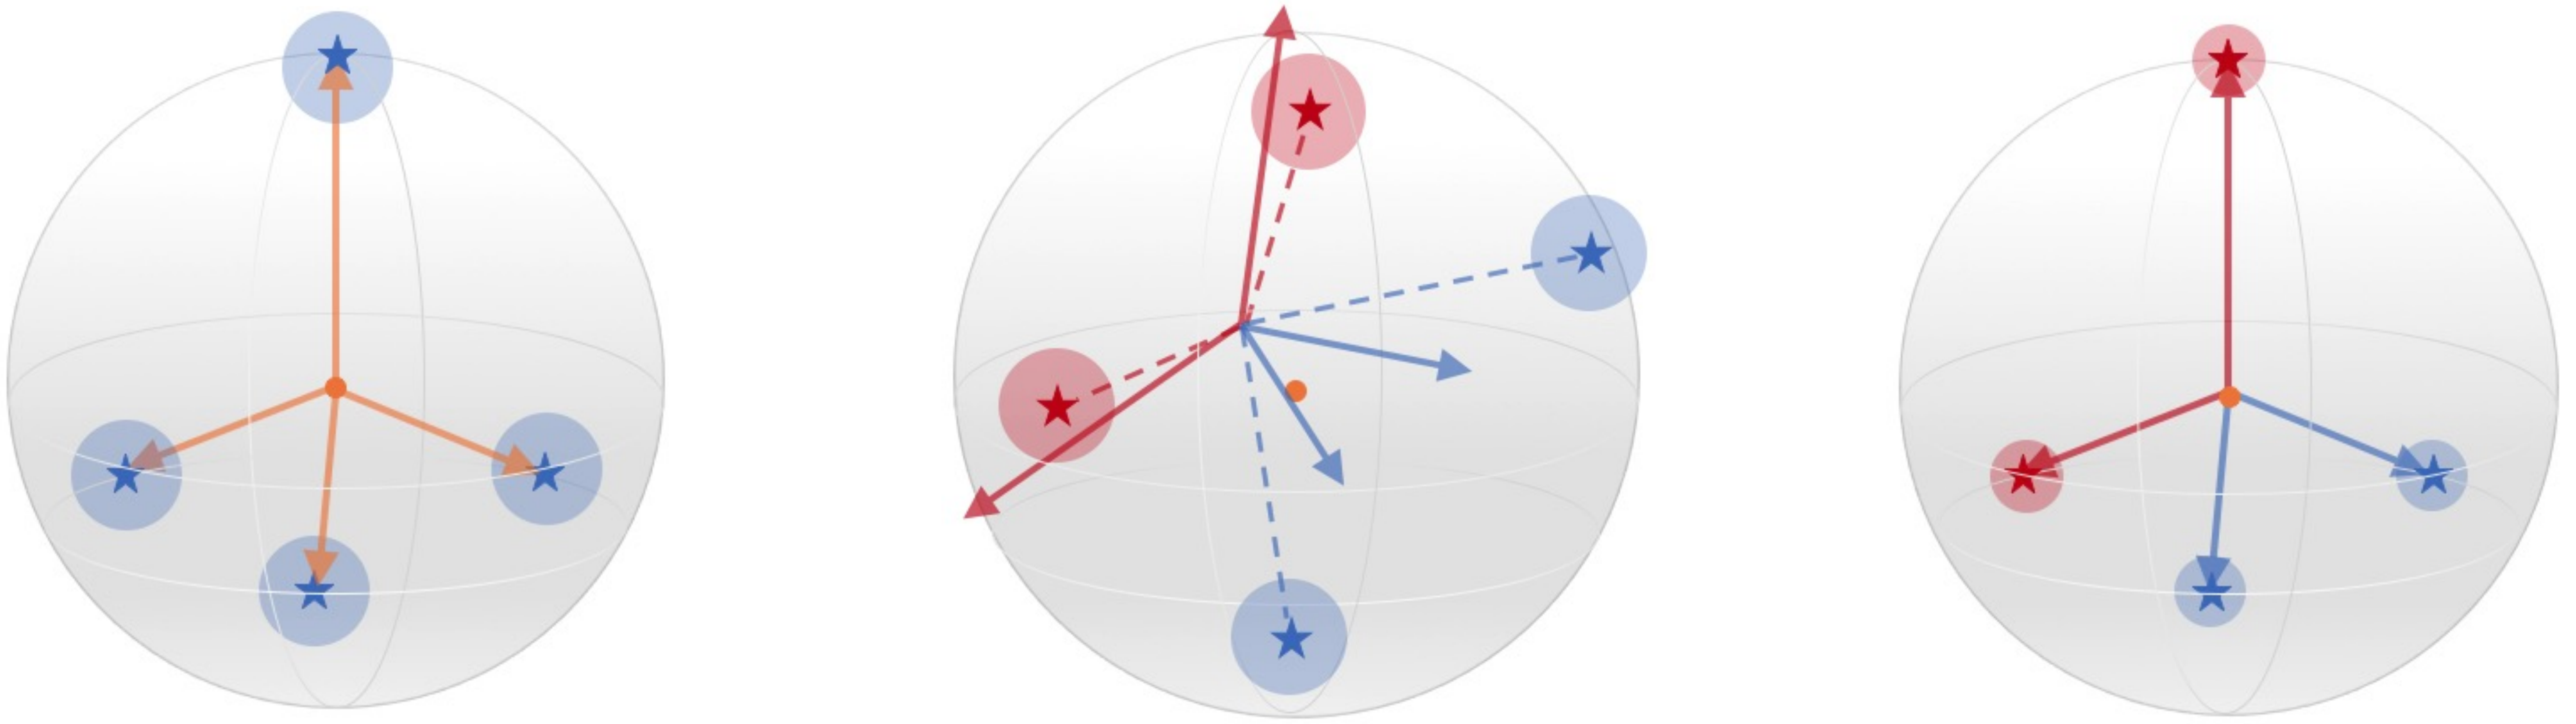

credit: [Liu-AISTATS2023]

Balanced dataset | imalanced dataset | imalanced dataset with NC

- NC1: Intra-class variation collapse to 0
  - Compact within-class features: $L_W = \sum_{k=1}^{K} \sum_{y_k=k} \frac{1}{n_k} ||h_i - \mu_k||_2^2$
- NC2: Class centers converge to simplex ETF
  - Distinct between-class features: $L_B = -\frac{1}{K} \sum_{k=1}^{K} \min_{k',k' \neq k} \arccos \frac{<\dot{\mu}_k, \dot{\mu}_{k'}>}{||\dot{\mu}_k|| \cdot ||\dot{\mu}_{k'}||}$

Liu, et al. "Inducing Neural Collapse in Deep Long-tailed Learning. " AISTATS2023.

NC in imbalanced learning: ETF classifier

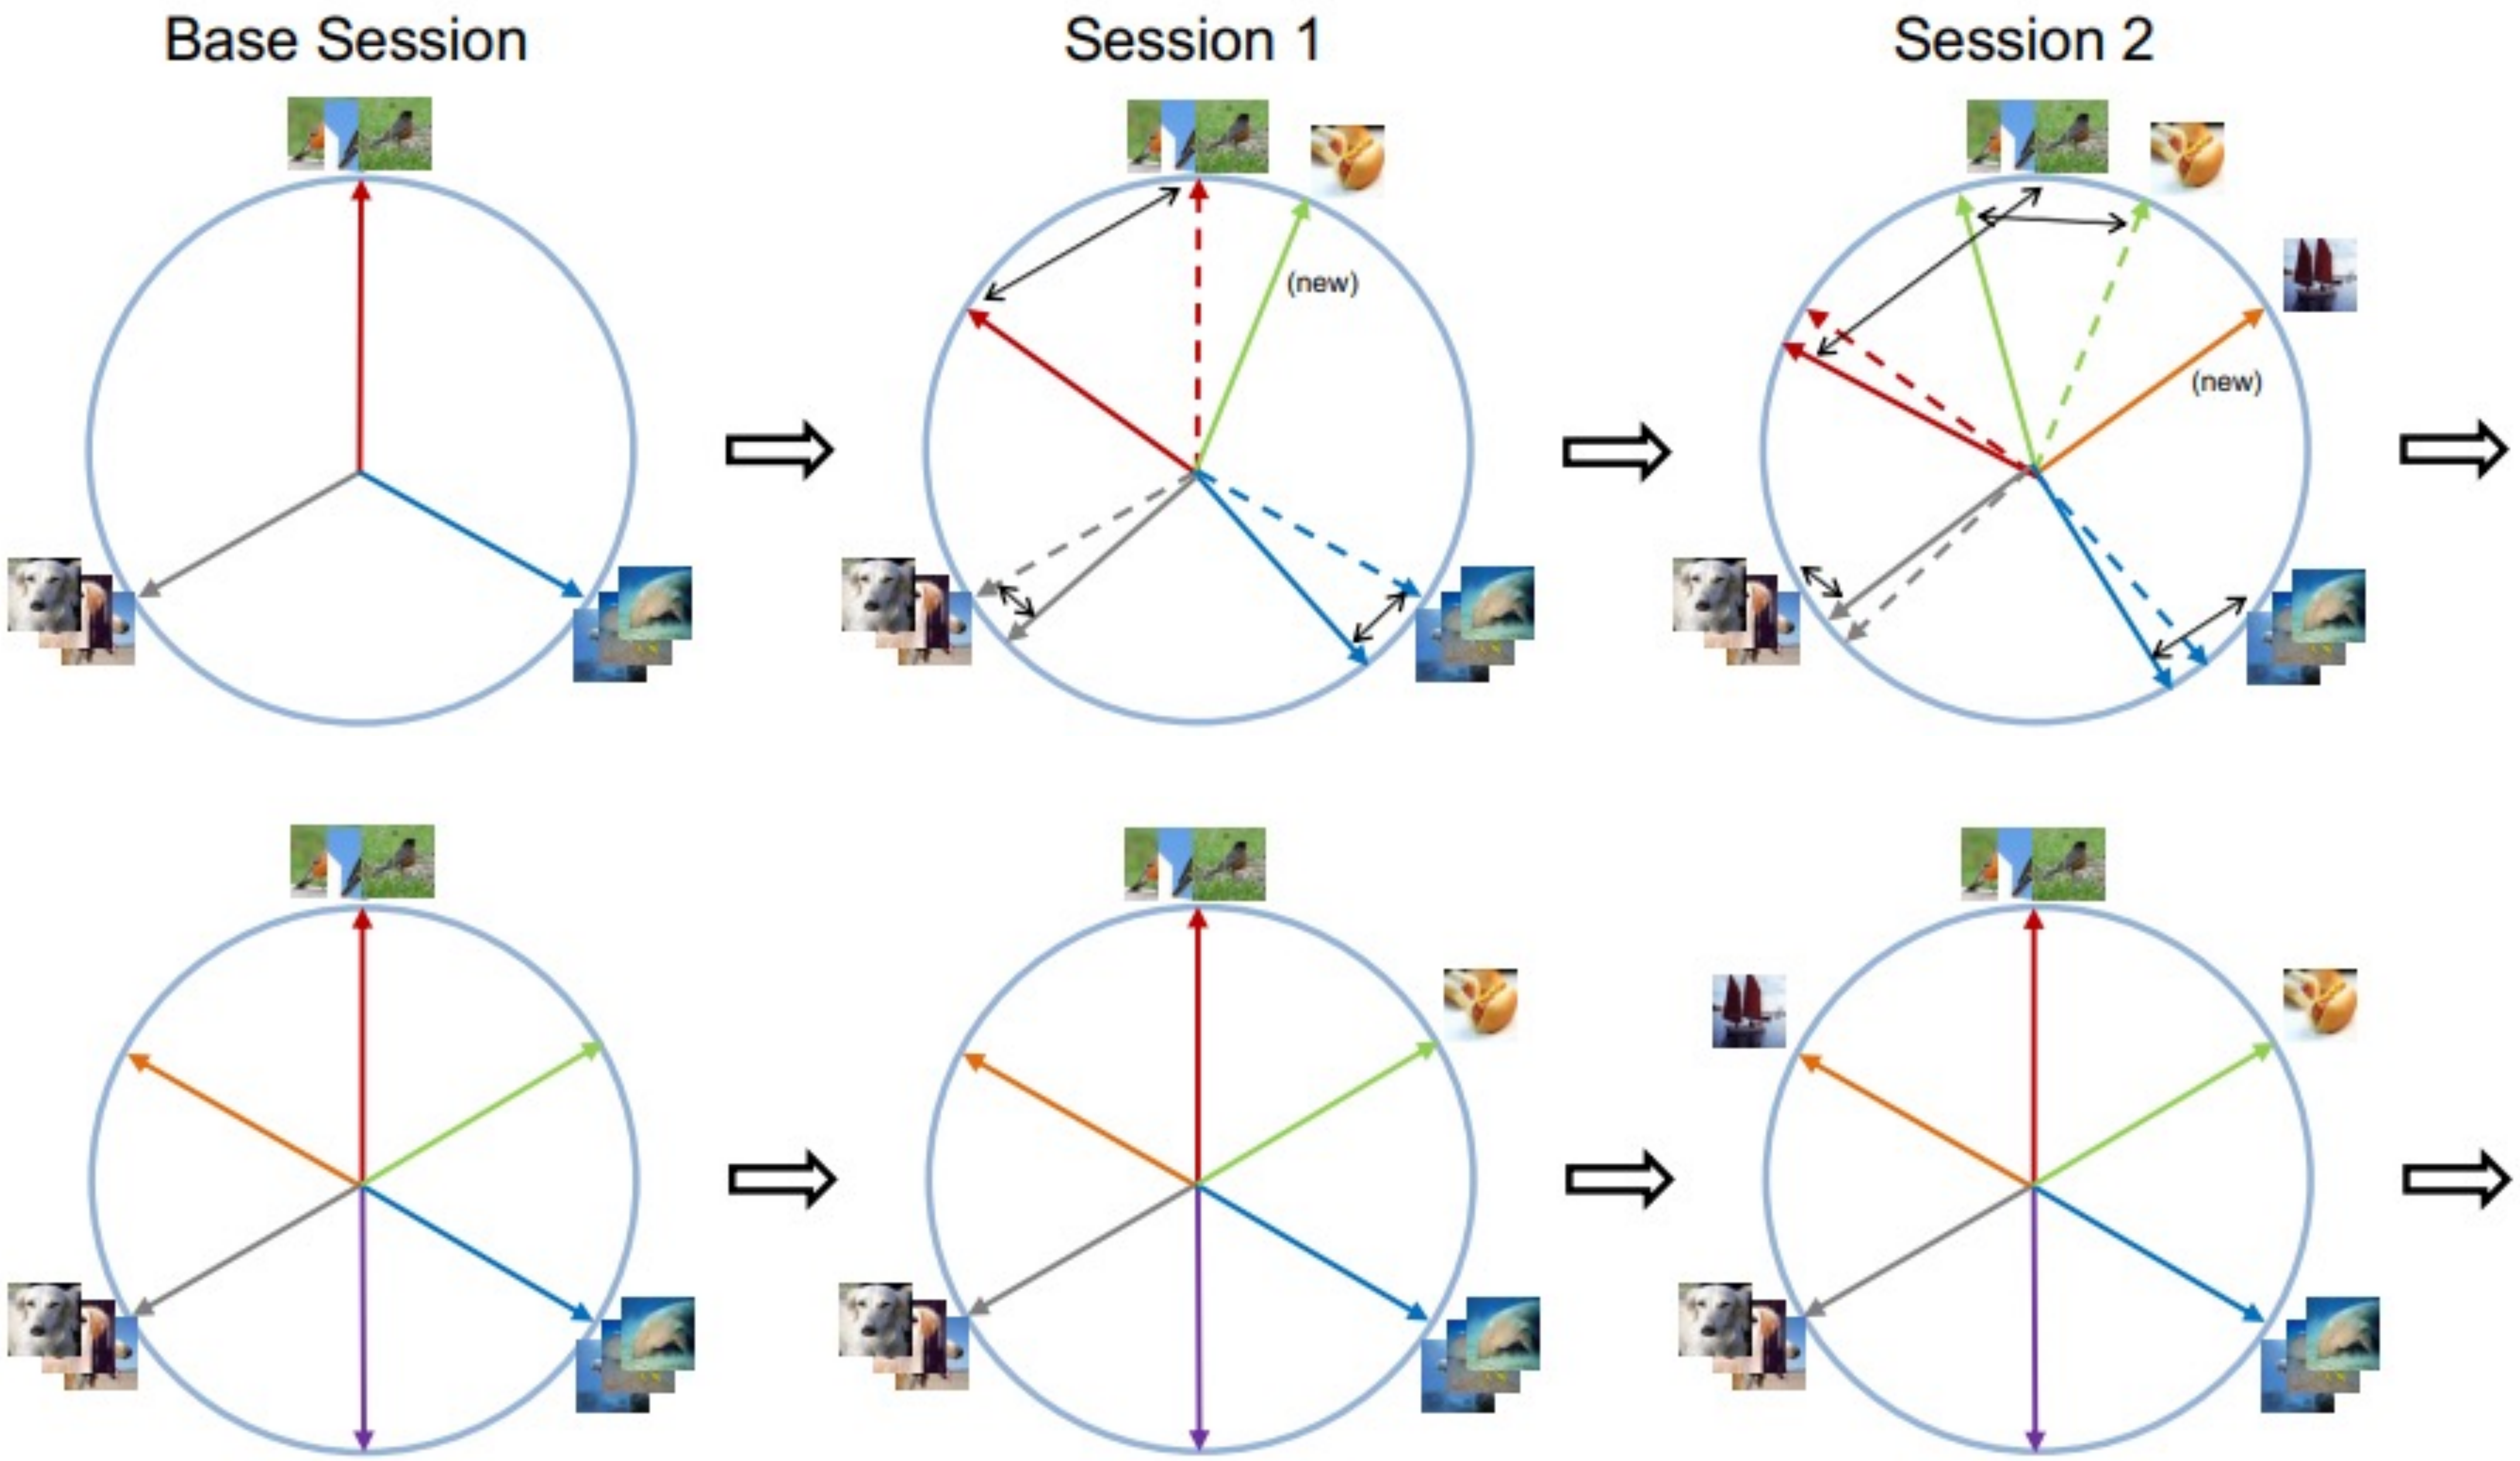

⇨ Learnable classifier prototypes

⇨ Pre-assigned classifier prototypes

credit: [Yang-ICLR2023]

Yang, et al. "Do we really need a learnable classifier at the end of deep neural network?." NeurIPS2022.
Yang, et al. "Neural Collapse Inspired Feature-Classifier Alignment for Few-Shot Class Incremental Learning." ICLR2023.

Few-Shot Learning

- Learning from base data
- Adaptation on novel data
  - FSL+ unlabeled data as SSL
  - Calibrating prototypes
  - Calibrating class distribution
- FSL in 2020s

The Key Idea of “FSL with unlabeled data as SSL”

Fast Adaptation: Sample Selection

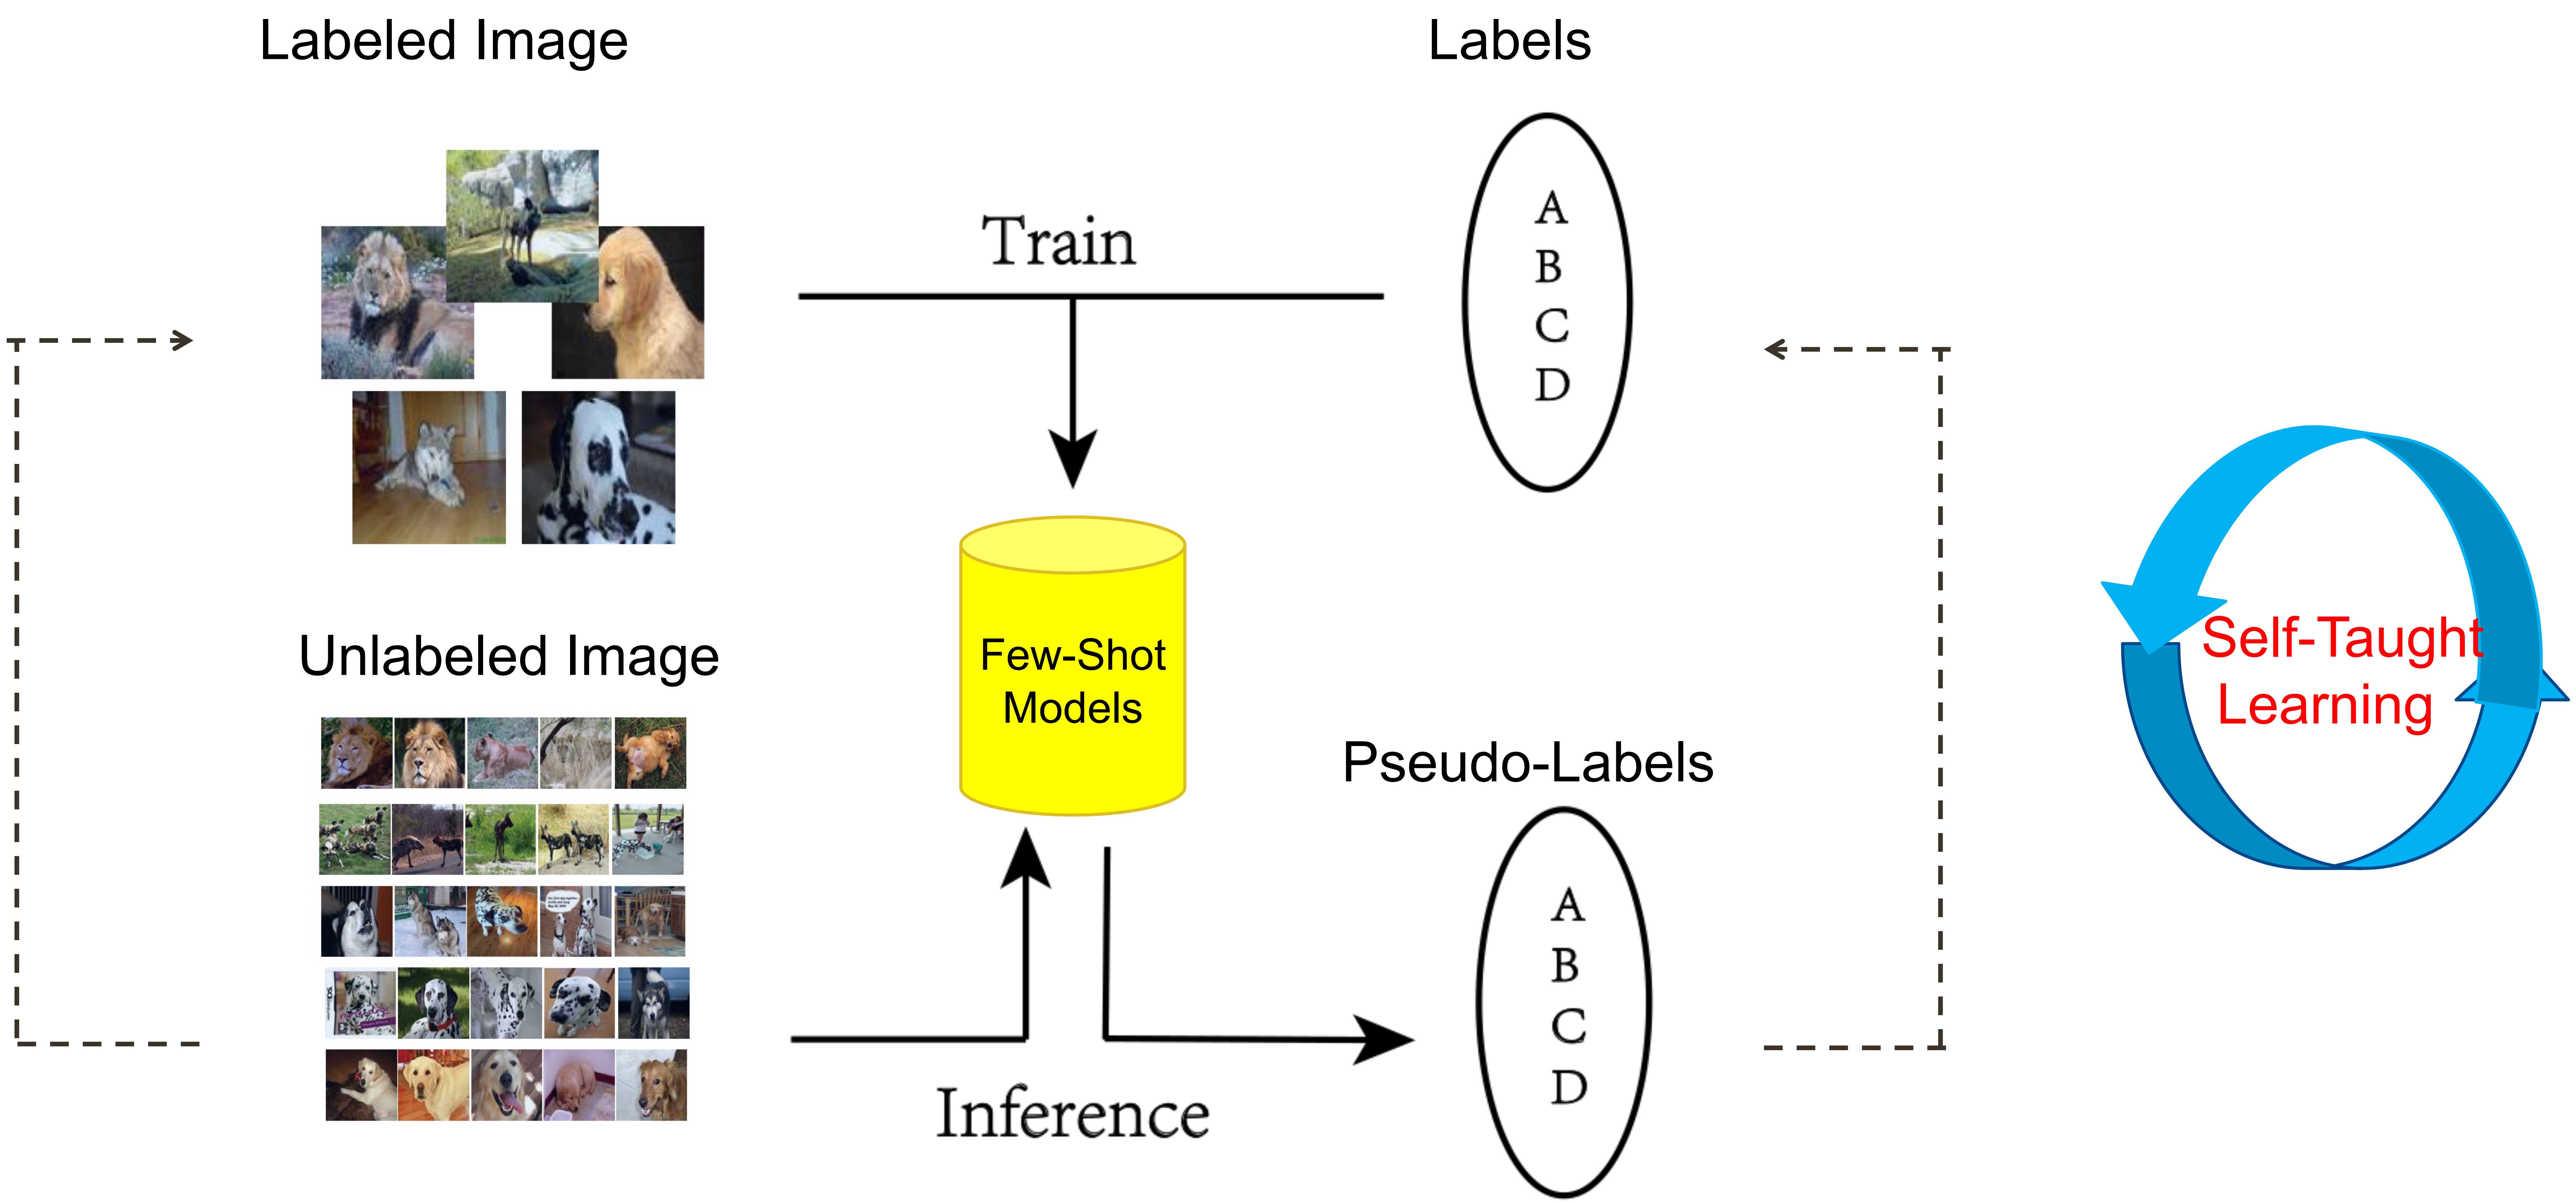

How to select credible samples to benefit learning?

Safe Semi-Supervised Learning Revisited

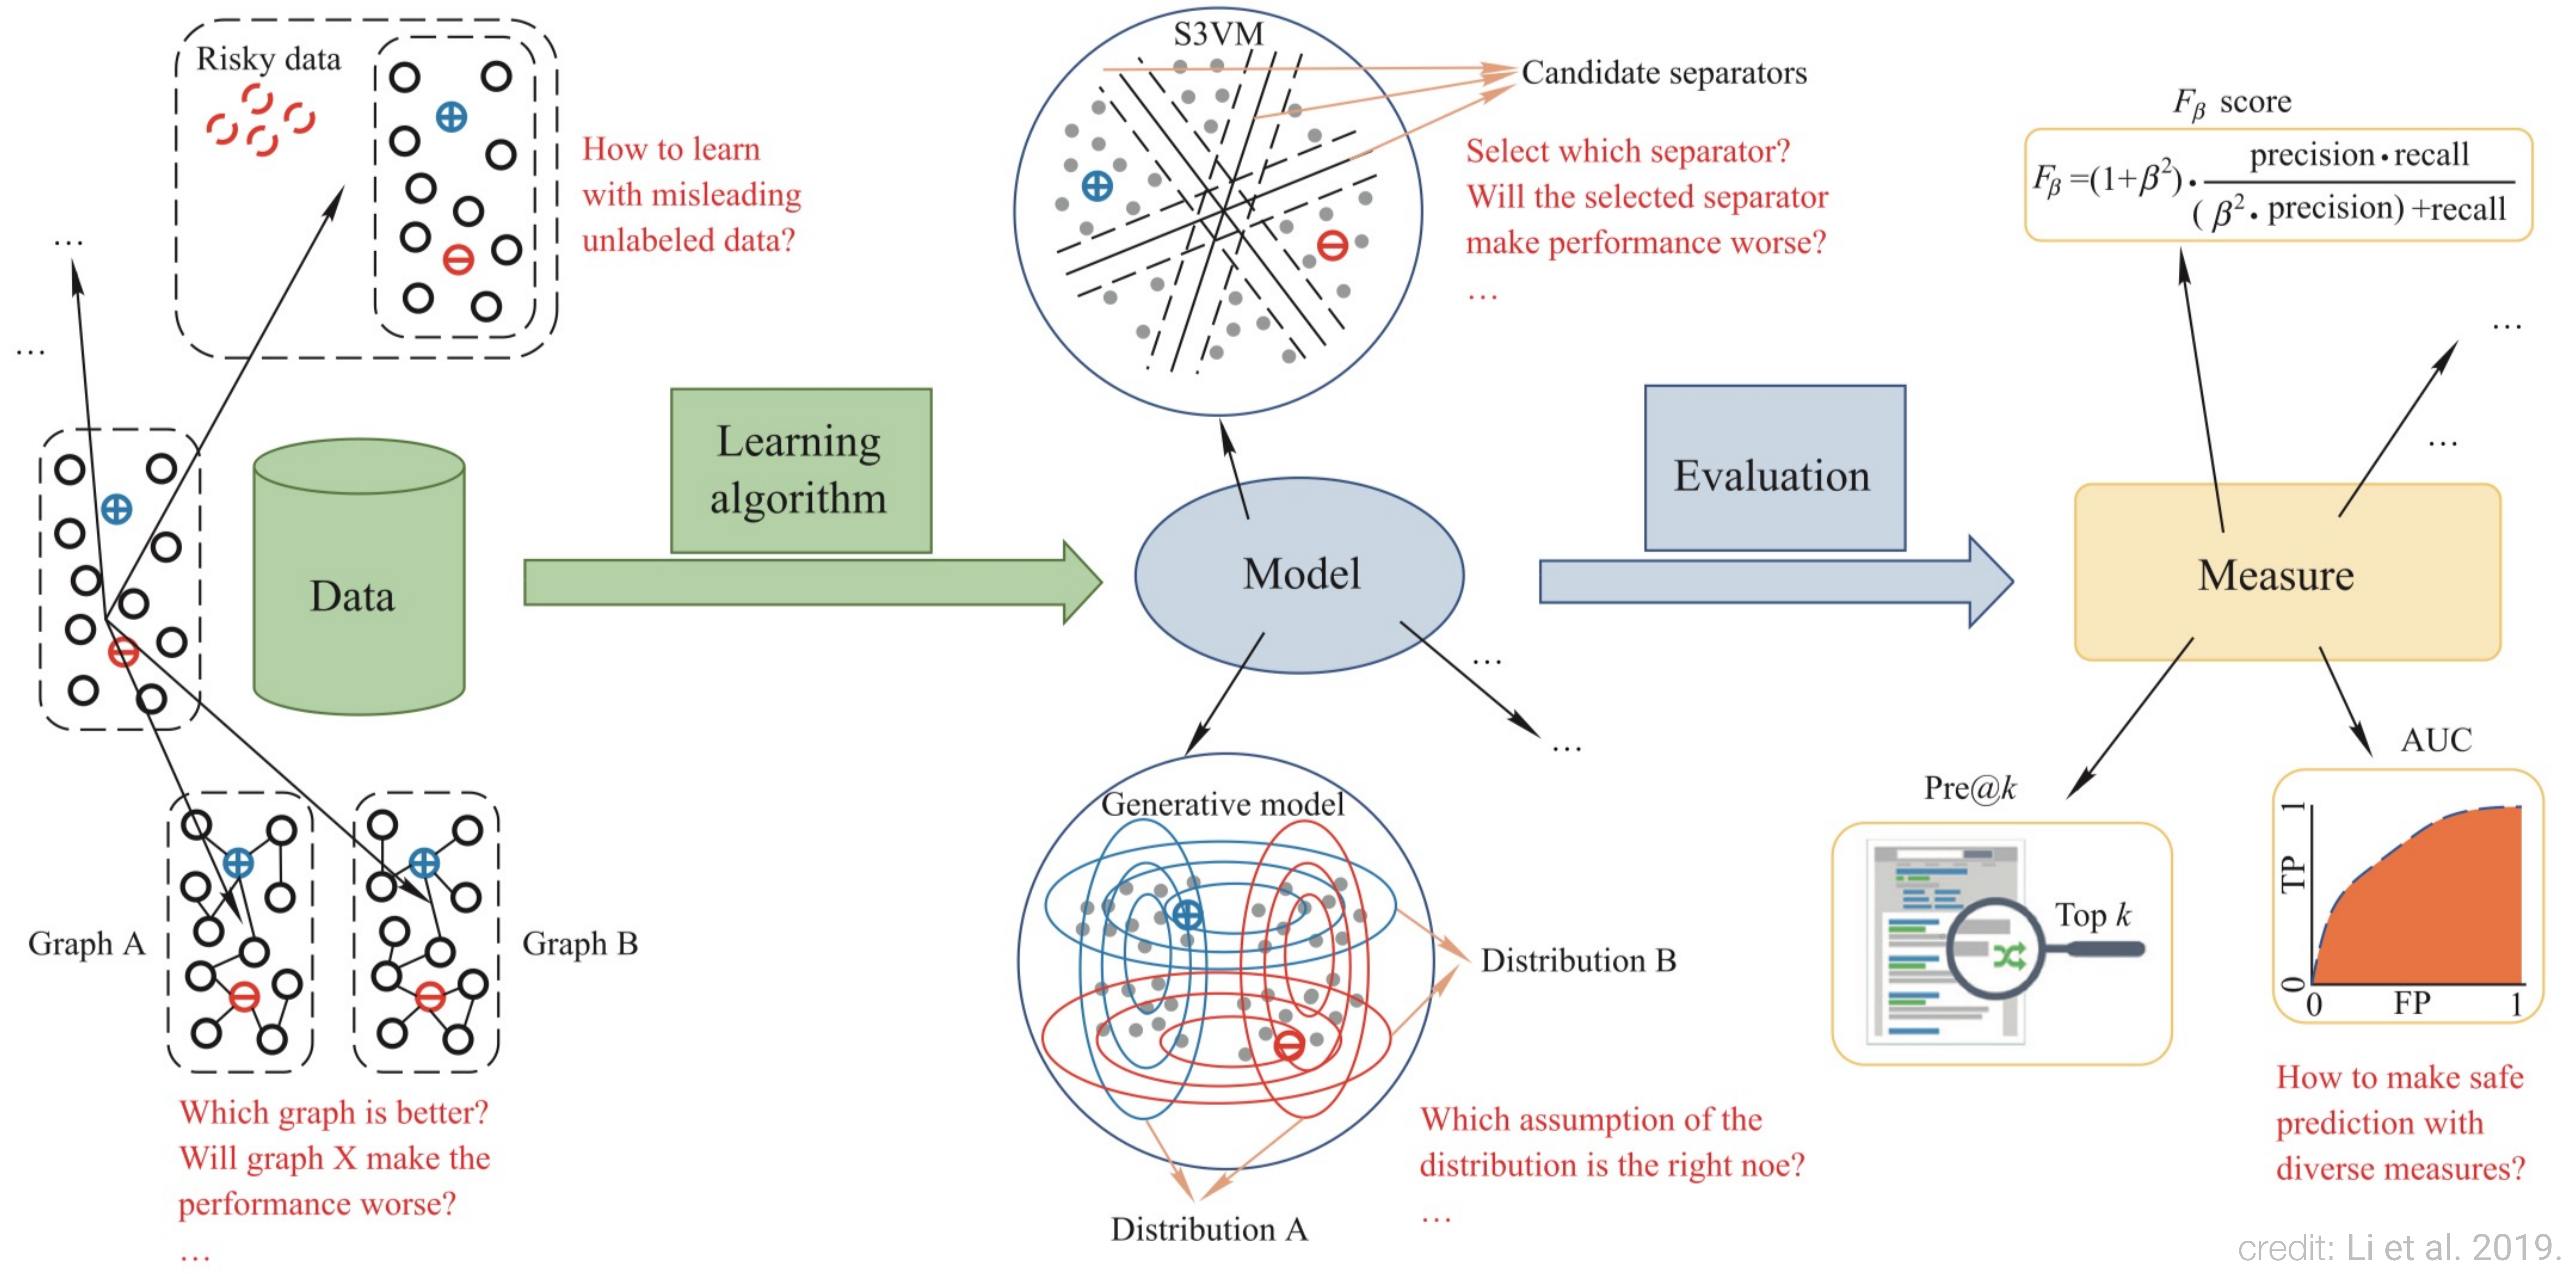

credit: Li et al. 2019.

Li, Yu-Feng, and Zhi-Hua Zhou. "Towards making unlabeled data never hurt." TPAMI 2014.
Li, Yu-Feng, Lan-Zhe Guo, and Zhi-Hua Zhou. "Towards safe weakly supervised learning." TPAMI 2019.
Li, Yu-Feng, and De-Ming Liang. "Safe semi-supervised learning: a brief introduction." FCS, 2019.

Instance Credibility Inference: Identify Noisy Data in Label Space

Linear system
with Noisy Data/Labels $Y = X\beta + \gamma$

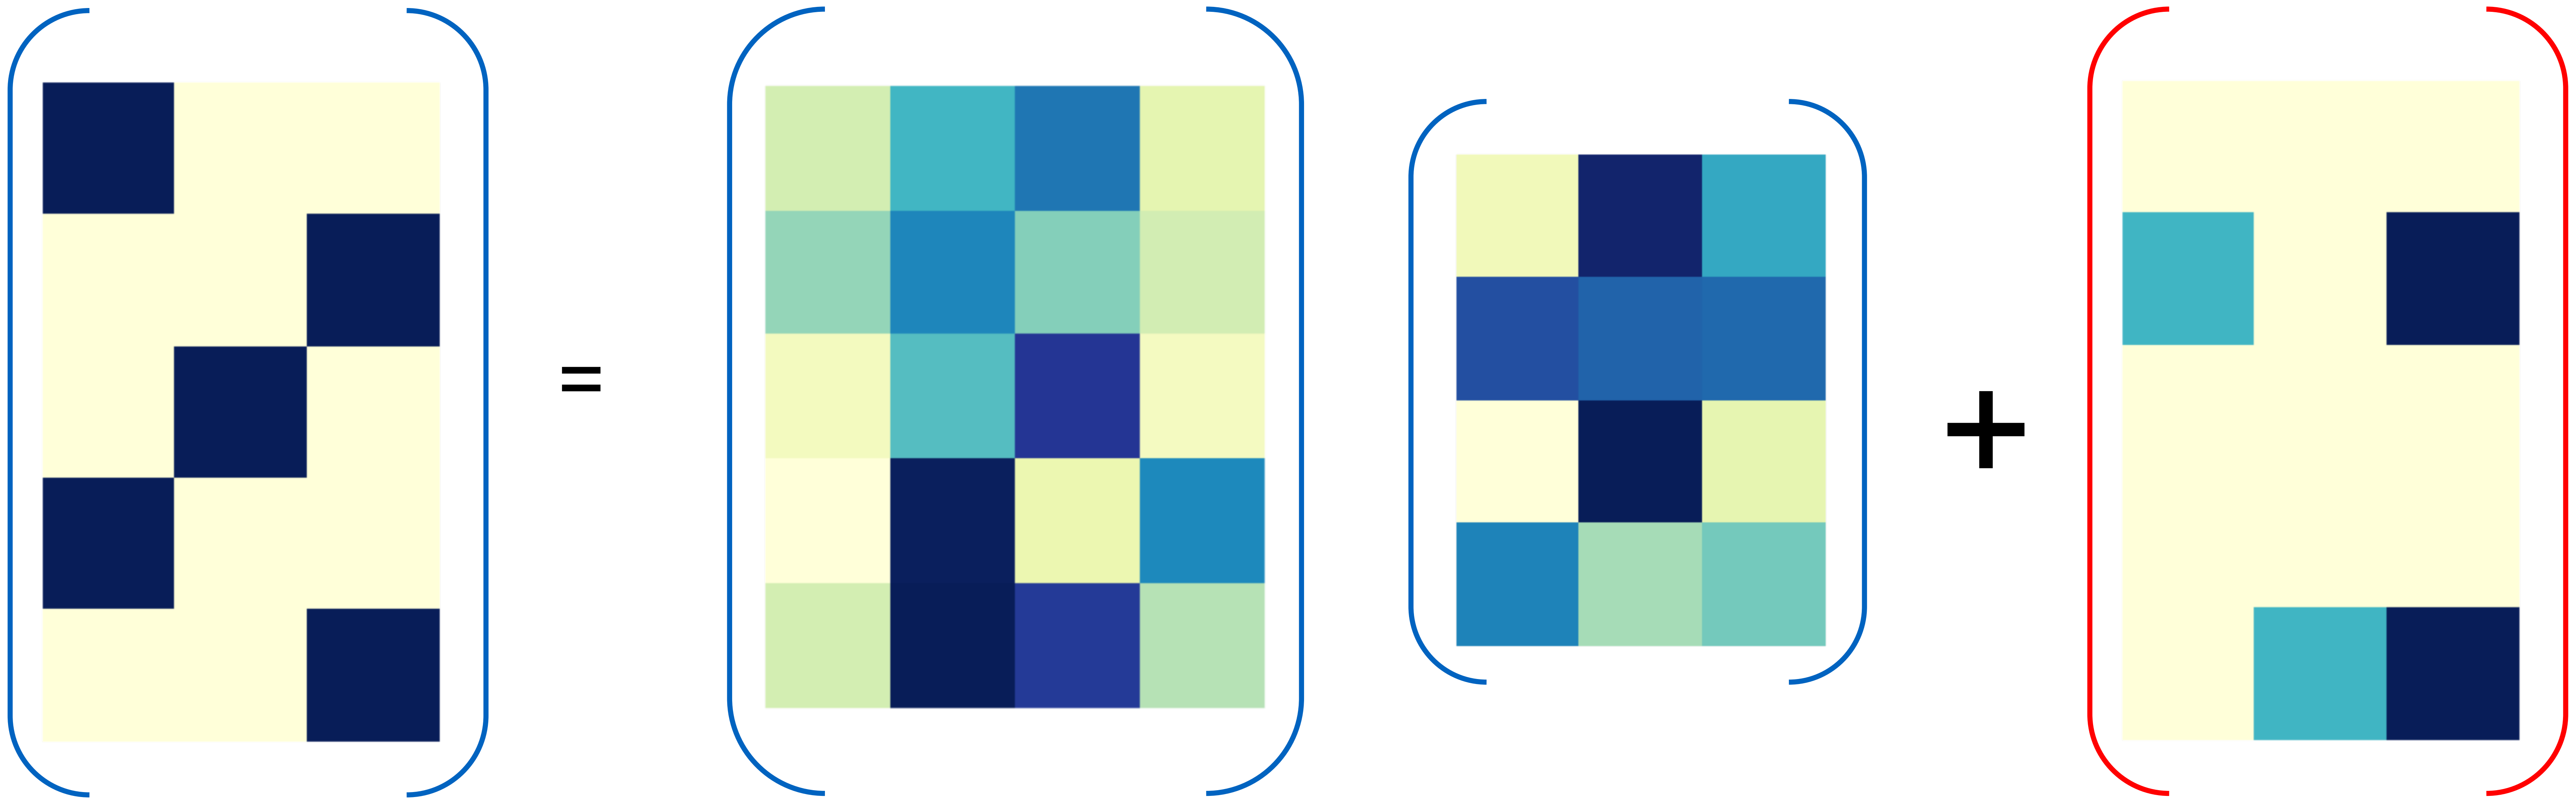

Noisy One-hot Labels $Y \in \mathbb{R}^{n\times c}$

Deep Features $X \in \mathbb{R}^{n\times d}$

Fitted Coef. $\beta \in \mathbb{R}^{d\times c}$

Noisy Data Indicator $\gamma \in \mathbb{R}^{n\times c}$

[Wright et al. TPAMI 09] [She et al. JASA 11] [Fu et al. ECCV 14, TPAMI 16.] [Fan et al. Statistical Sinica 18] [Wang et al. CVPR 20, TPAMI 21, CVPR 22]

Instance Credibility Inference: Understanding γ in Statistics

$$y = x^\top \beta + \varepsilon + \gamma$$

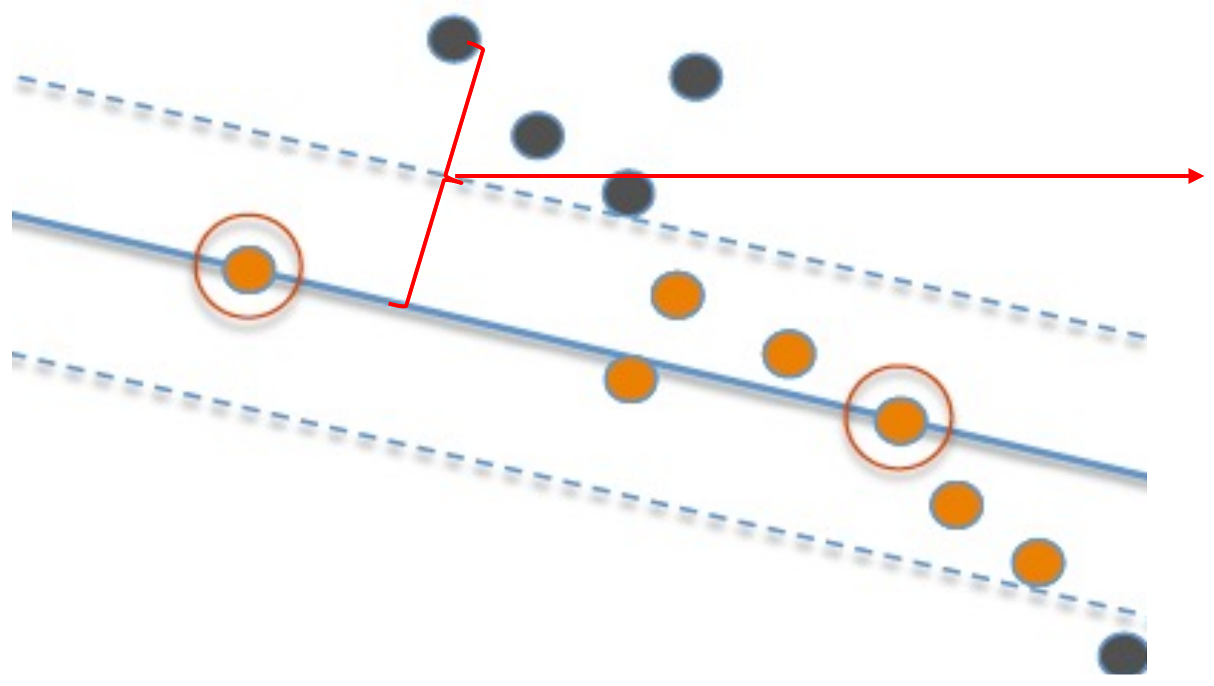

γ_i equals to the residual predict error $\gamma_i = y_i - x_i^\top \hat{\beta}$

Leave-one-out externally studentized residual:

$$t_i = \frac{y_i - \boldsymbol{x}_i^\top \hat{\beta}_{(i)}}{\hat{\sigma}_{(i)}(1+\boldsymbol{x}_i(\boldsymbol{X}_{(i)}^\top \boldsymbol{X}_{(i)})^{-1}\boldsymbol{x}_i)^{1/2}}$$

$\Leftrightarrow$ test whether $\gamma = 0$ in $\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \gamma 1_i + \boldsymbol{\varepsilon}.$

When there are multiple outliers:
masking and swamping

$$\boldsymbol{y} = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\epsilon} + \boldsymbol{\gamma}$$

She et al. Outlier detection using nonconvex penalized regression. Journal of the American Statistical Association, 2011.

Instance Credibility Inference: Statistical Outlier Detection

$$y_i = x_i^\top \beta + \varepsilon + \gamma_i \longrightarrow \hat{\gamma}_i \longrightarrow O = \{i : \hat{\gamma}_i \neq 0\}$$

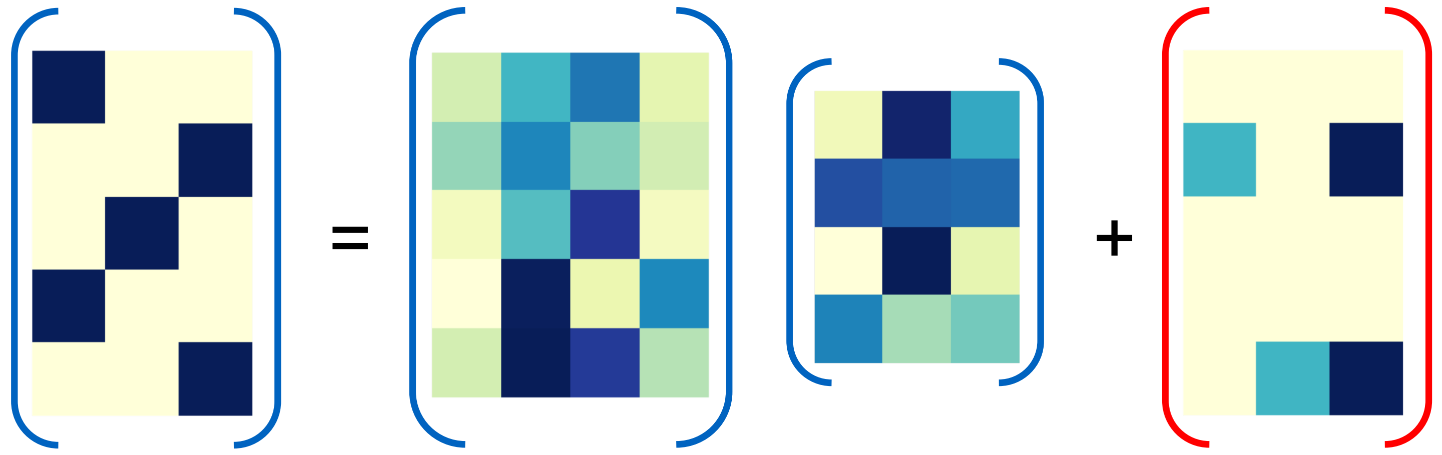

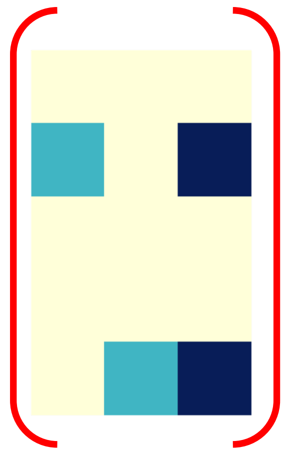

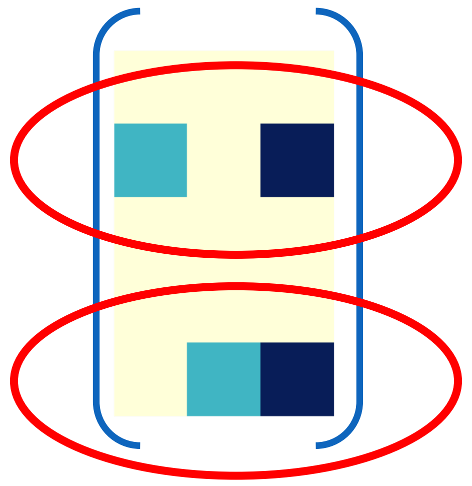

$$\underset{\boldsymbol{\beta},\boldsymbol{\gamma}}{\operatorname{argmin}} L(\boldsymbol{\beta},\boldsymbol{\gamma}) := \|\boldsymbol{Y} - \boldsymbol{X}\boldsymbol{\beta} - \boldsymbol{\gamma}\|_{\mathrm{F}}^2 + \lambda P(\boldsymbol{\gamma}) \longrightarrow \underset{\boldsymbol{\gamma}}{\operatorname{argmin}} \left\| \tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{X}}\boldsymbol{\gamma} \right\|_{\mathrm{F}}^2 + \lambda P(\boldsymbol{\gamma})$$

Wang et al. Instance Credibility Inference for Few-Shot Learning. CVPR 2020

Wang et al. How to Trust Unlabeled Data? Instance Credibility Inference for Few-Shot Learning. IEEE TPAMI 2021.

Instance Credibility Inference: Solving Gamma

$$\underset{\boldsymbol{\gamma}}{\operatorname{argmin}} \left\| \tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{X}} \boldsymbol{\gamma} \right\|_{\mathrm{F}}^{2} + \lambda P(\boldsymbol{\gamma})$$

How to select λ?

- heuristics rules $\lambda = 2.5\hat{\sigma}$?
- Cross-validation?
- Data adaptive techniques?
- AIC, BIC?

It is hard to select a proper λ.

$$\hat{\gamma} = f(\lambda).$$

$$\lambda \to \infty,\ \hat{\gamma} \to 0.$$

$$P(\boldsymbol{\gamma}) = \sum_{i=1}^{n} \|\boldsymbol{\gamma}_i\|_2,$$

$$C_i = \sup\{\lambda : \|\hat{\gamma}_i(\lambda)\| \neq 0\}$$

This can be solved by GLM-Net.

Friedman, et al.. “Regularization Paths for Generalized Linear Models via Coordinate Descent.” Journal of Statistical Software, 2010

Instance Credibility Inference: Theoretical Guarantees

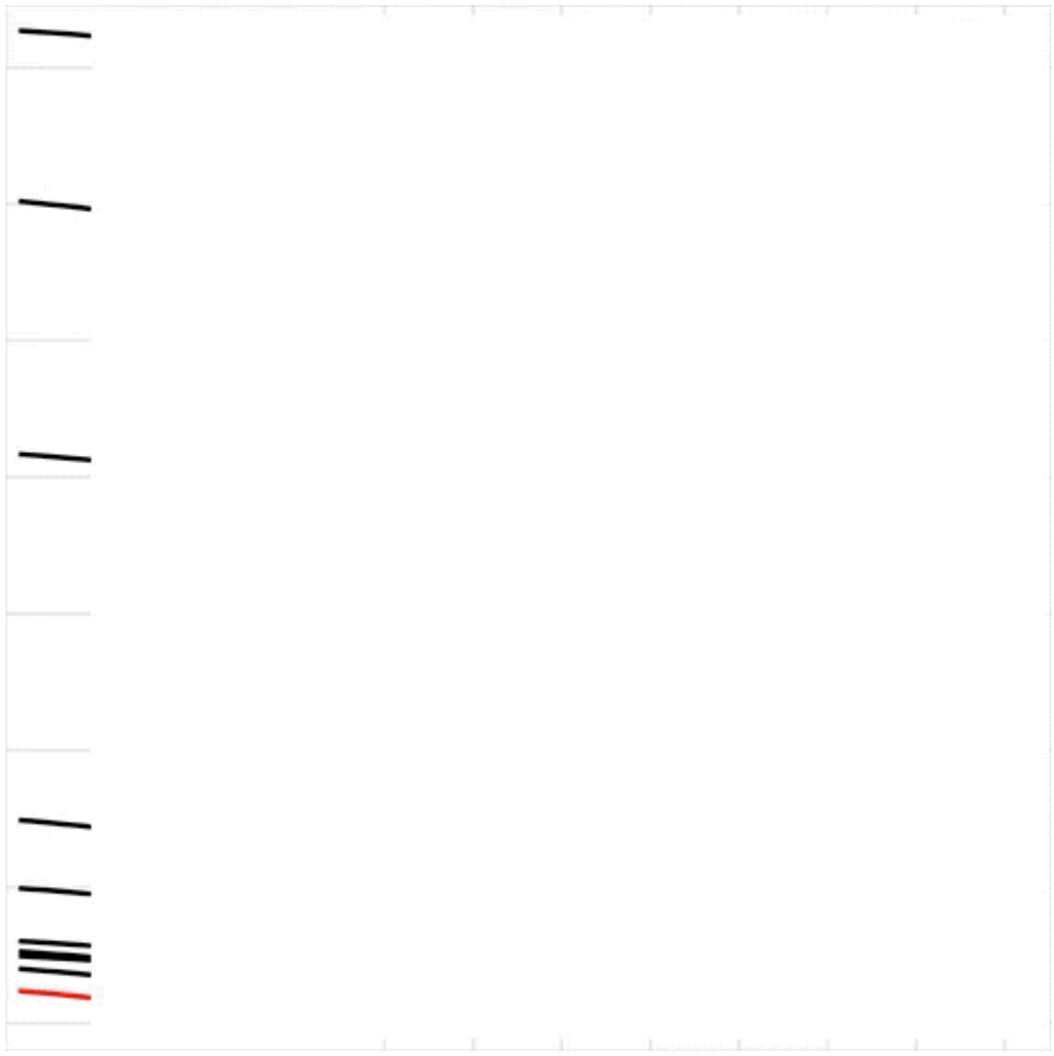

Theoretical guarantees:

1) Under the restricted eigenvalue and irrepresentability conditions, the noisy data identified by ICI is the subset of the ground-truth noisy data;
2) With further satisfied large error condition, ICI will identify all the noisy data.

See our CVPR 2022 tutorial for details.
https://sparse-learning.github.io

Wang et al. Instance Credibility Inference for Few-Shot Learning. CVPR 2020
Wang et al. How to Trust Unlabeled Data? Instance Credibility Inference for Few-Shot Learning. IEEE TPAMI 2021.

The Key Idea of “Calibrating prototypes and class distribution”

- Bias towards base tasks

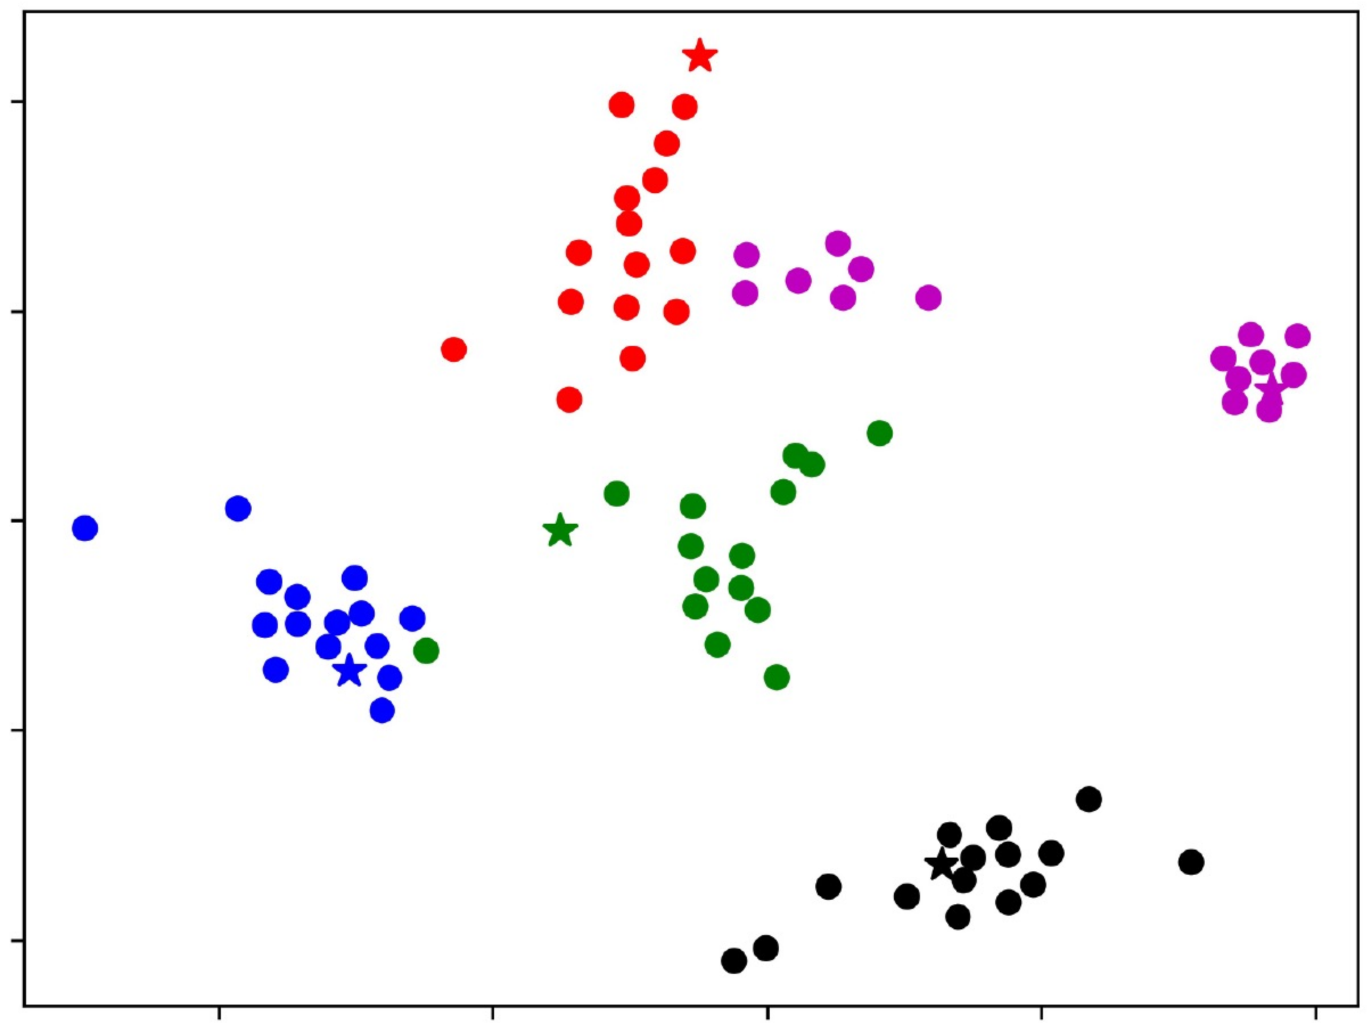

- Overfitting on few-shot samples

How to calibrate novel task to make robust estimation?

Adaptation on novel data: Calibration

The representation learned on base dataset are biased.

We can calibrate novel data to reduce the bias.

Li, et al. “Ranking distance calibration for cross-domain few-shot learning.” CVPR, 2022.

Ranking Distance Calibration

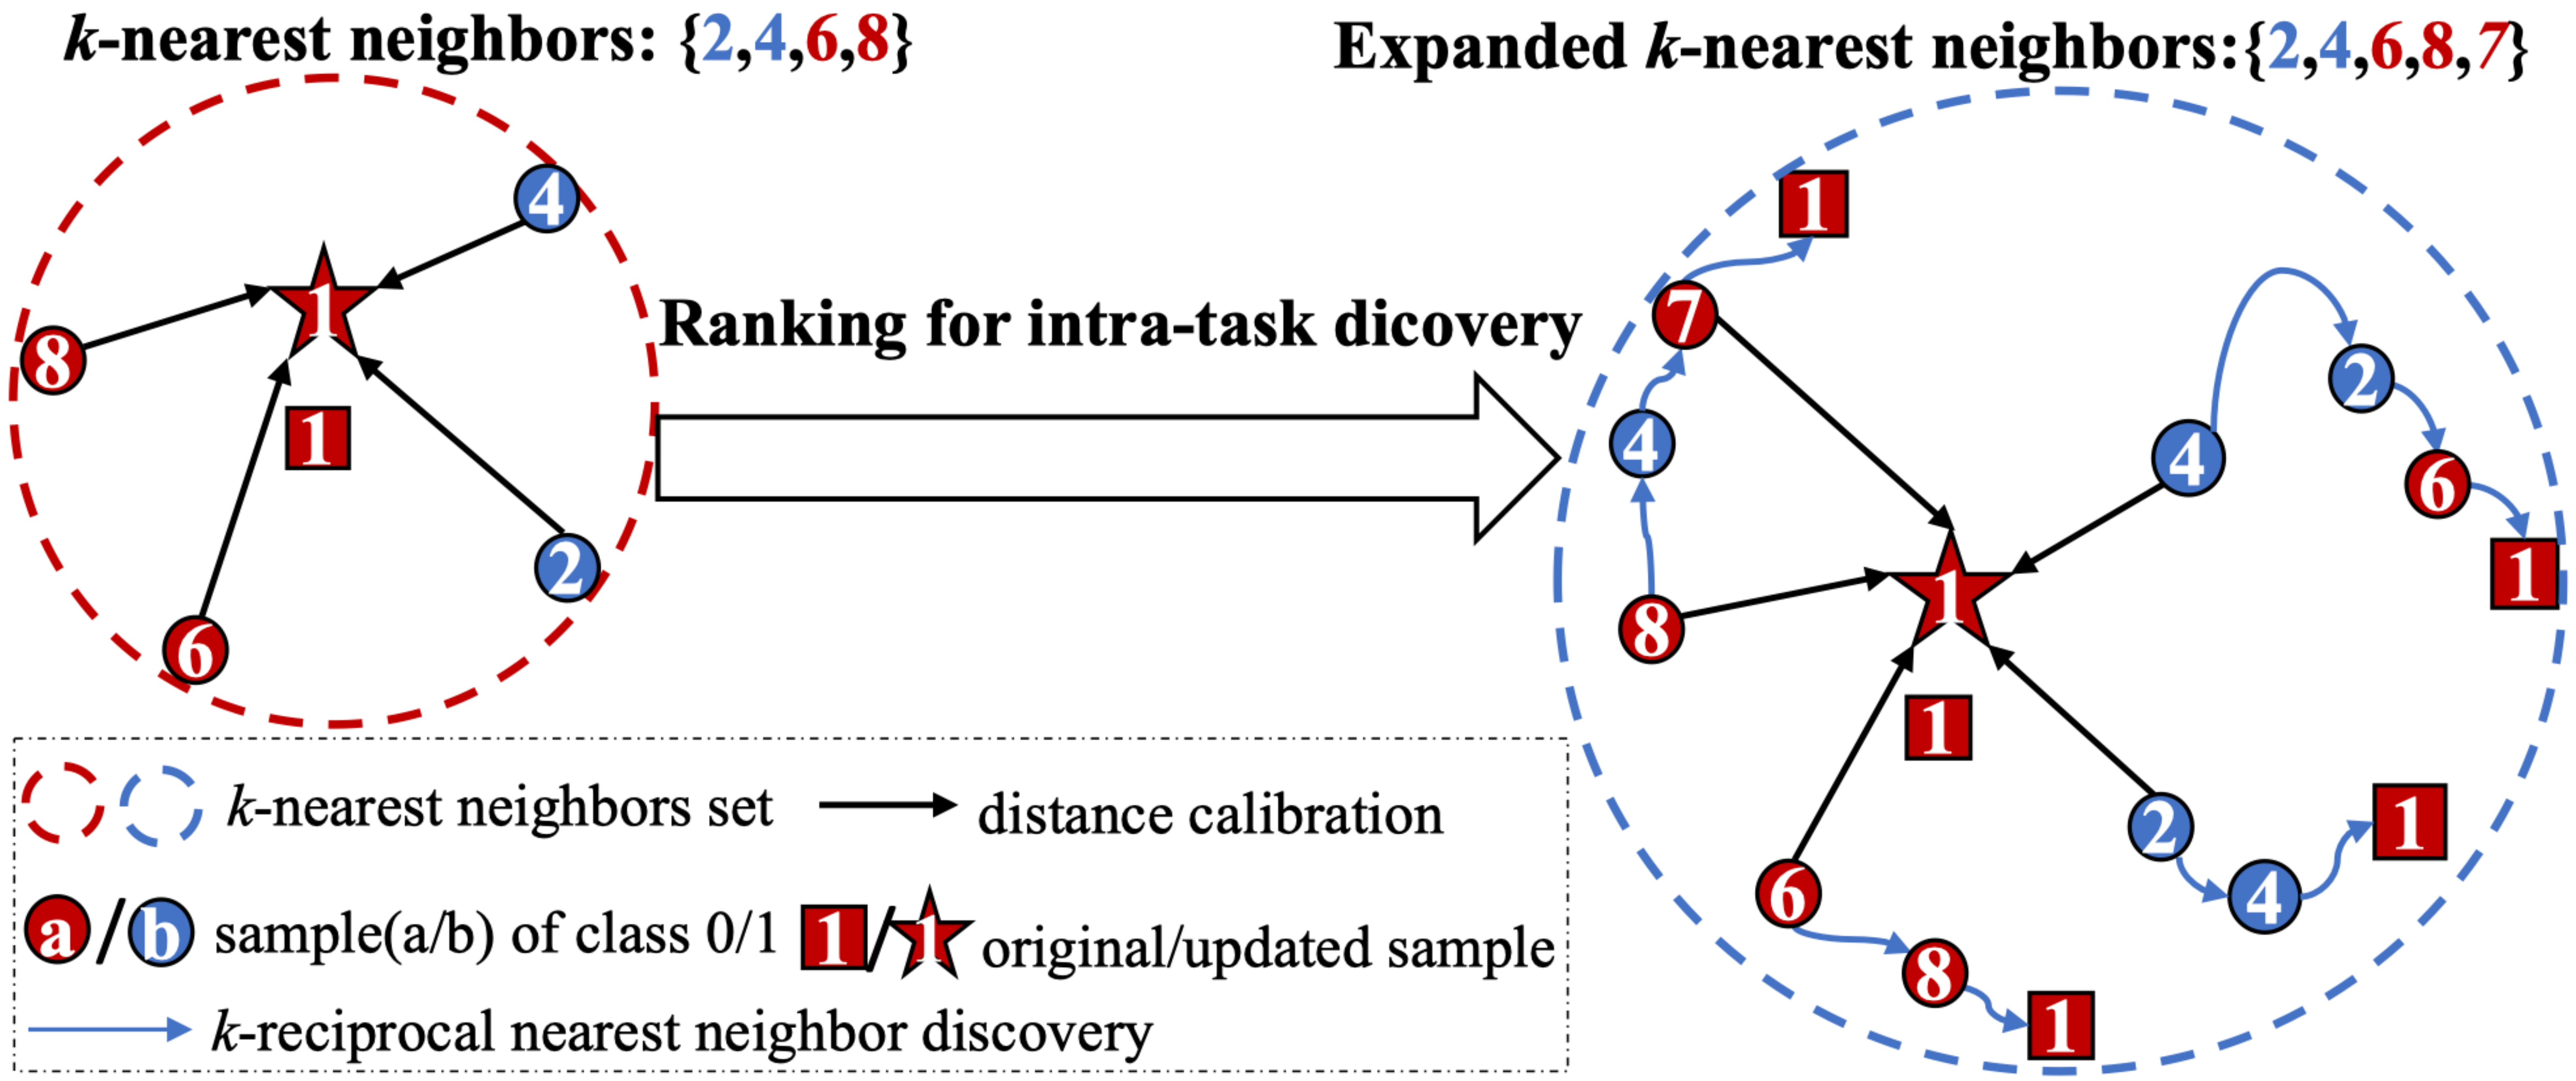

Discover likely positive samples and calibrate their pairwise distances

Novel Logit Distribution Calibration

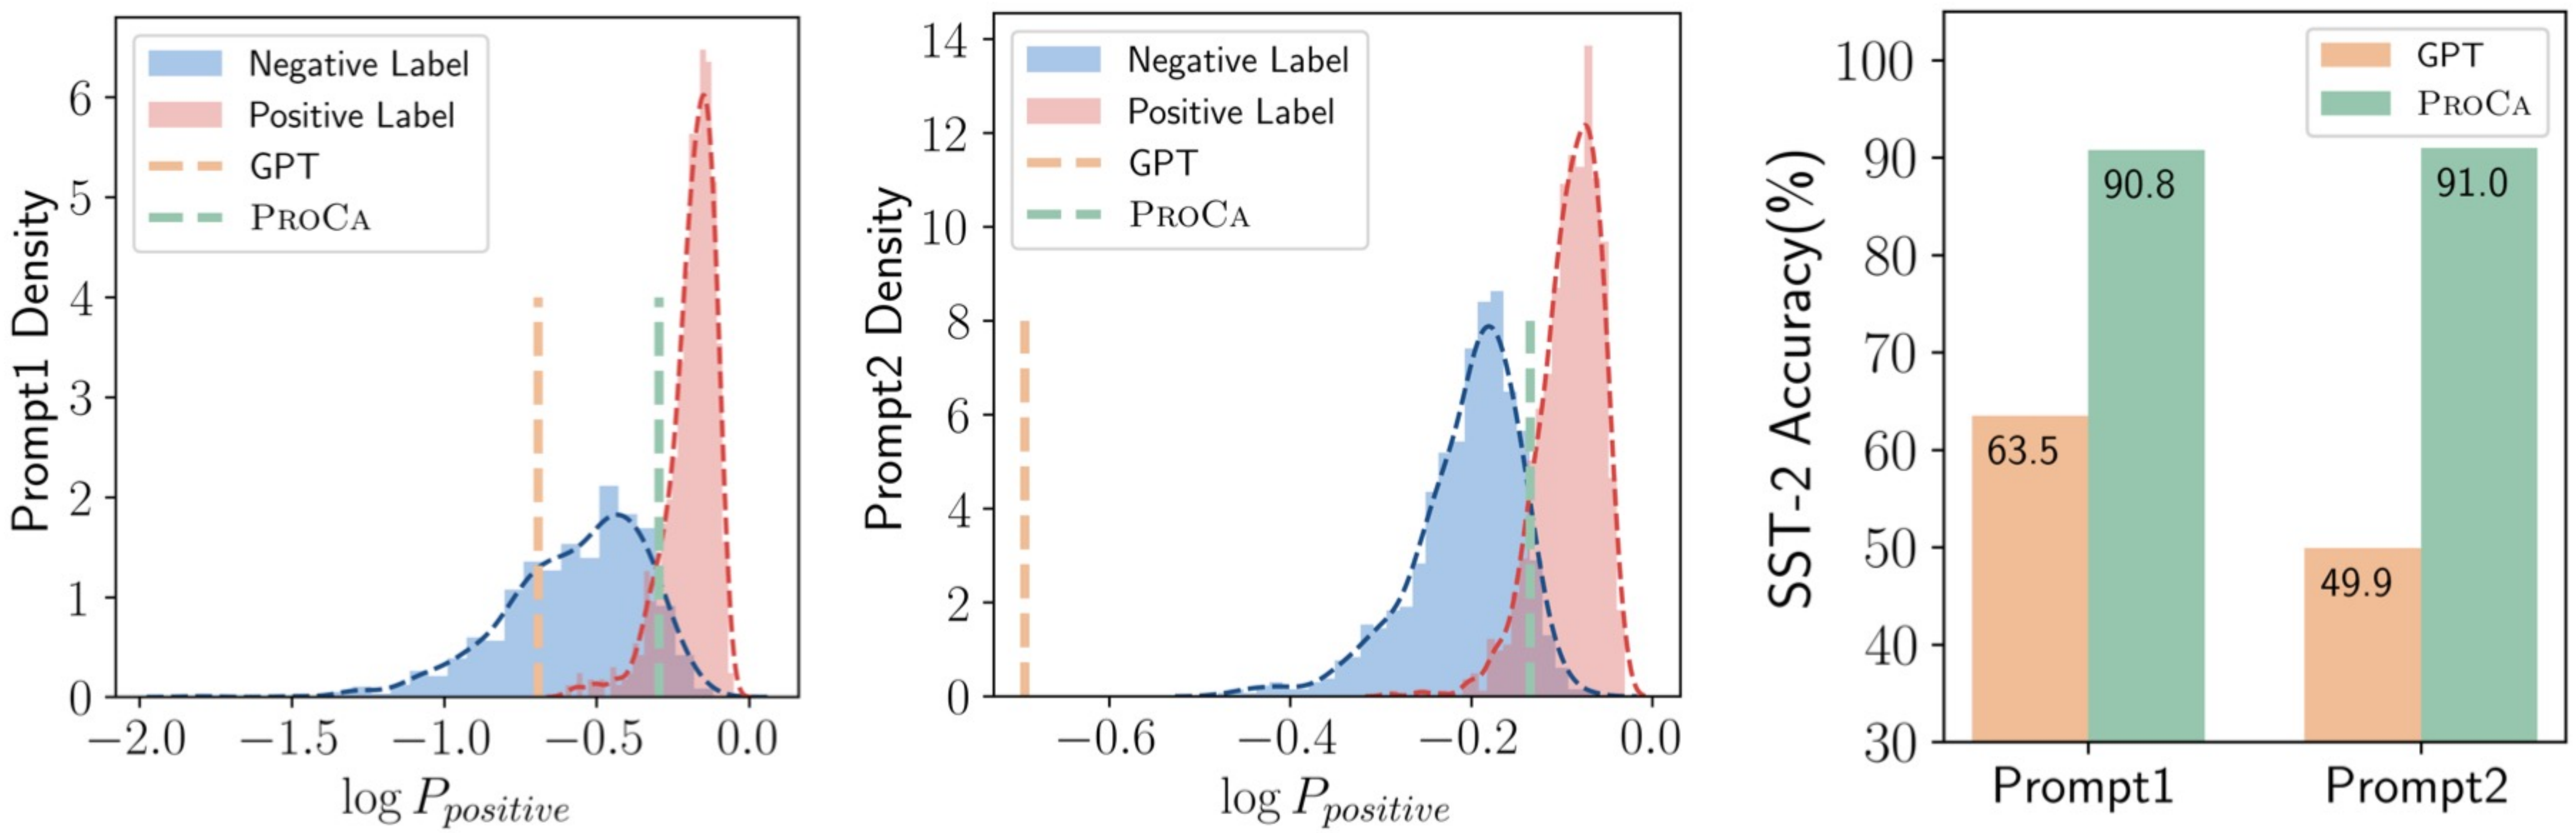

credit: [Han, et al. 2023]

Learning a class-dependent threshold based on Gaussian mixture model of category logits distribution.

Han, et al. "Prototypical calibration for few-shot learning of language models." ICLR, 2023.

Novel Distribution Calibration

Transfer statistics from base classes to calibrate novel distribution.

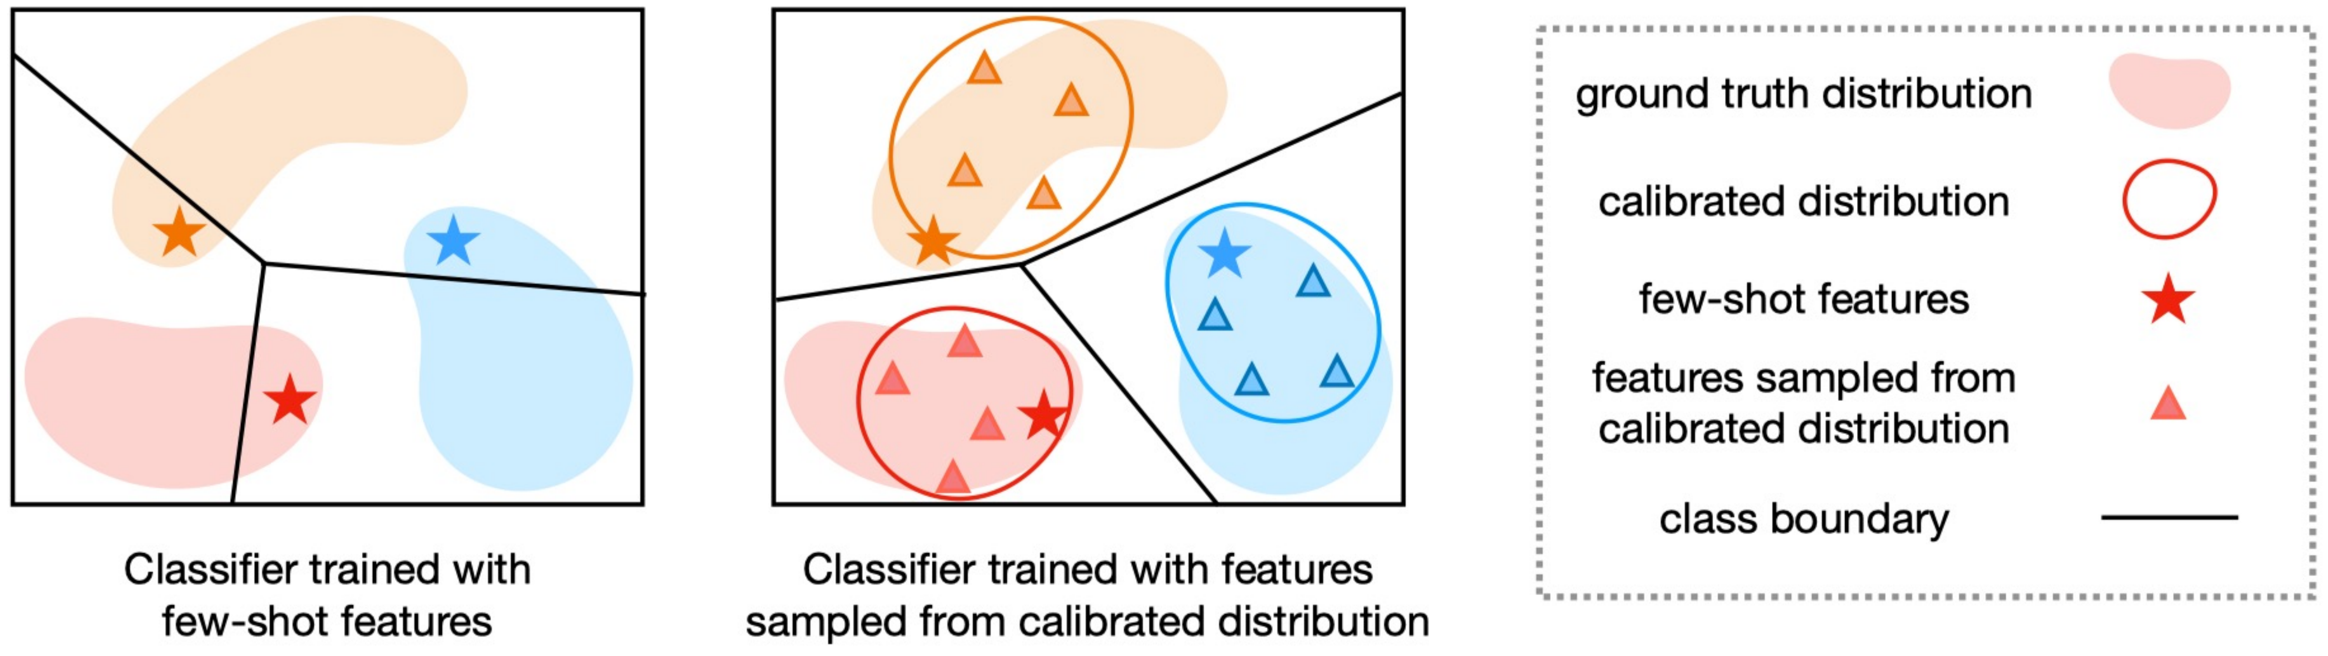

credit: [Yang, et al. 2021]

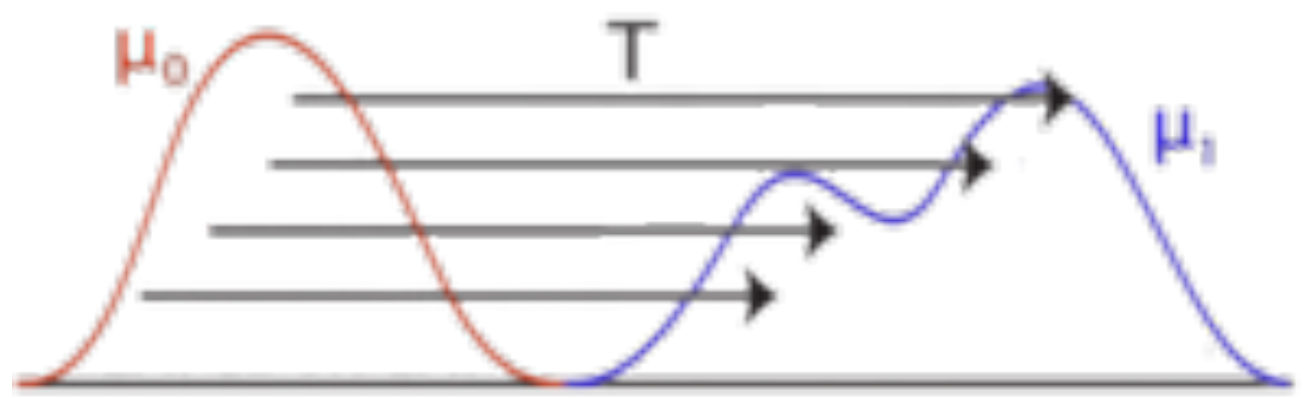

credit: wikipedia

The relationship between base classes can be modeled via Euclidean distance.

And can be modeled via more holistically, for example using optimal transport.

Yang, et al. "Free lunch for few-shot learning: Distribution calibration." ICLR, 2021.
Guo, et al. “Adaptive Distribution Calibration for Few-Shot Learning with Hierarchical Optimal Transport.” NeurIPS, 2022.

Bayesian Approach on Uncertainty Calibration

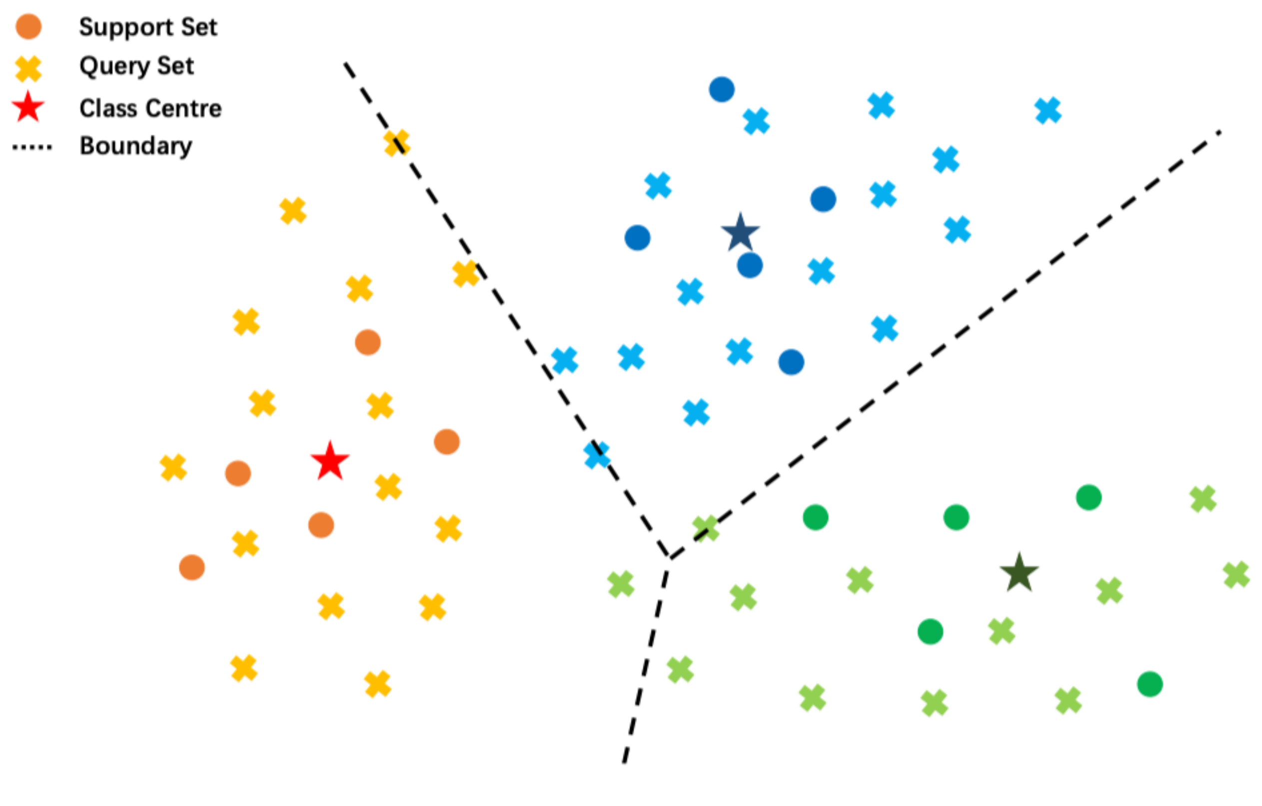

(a)

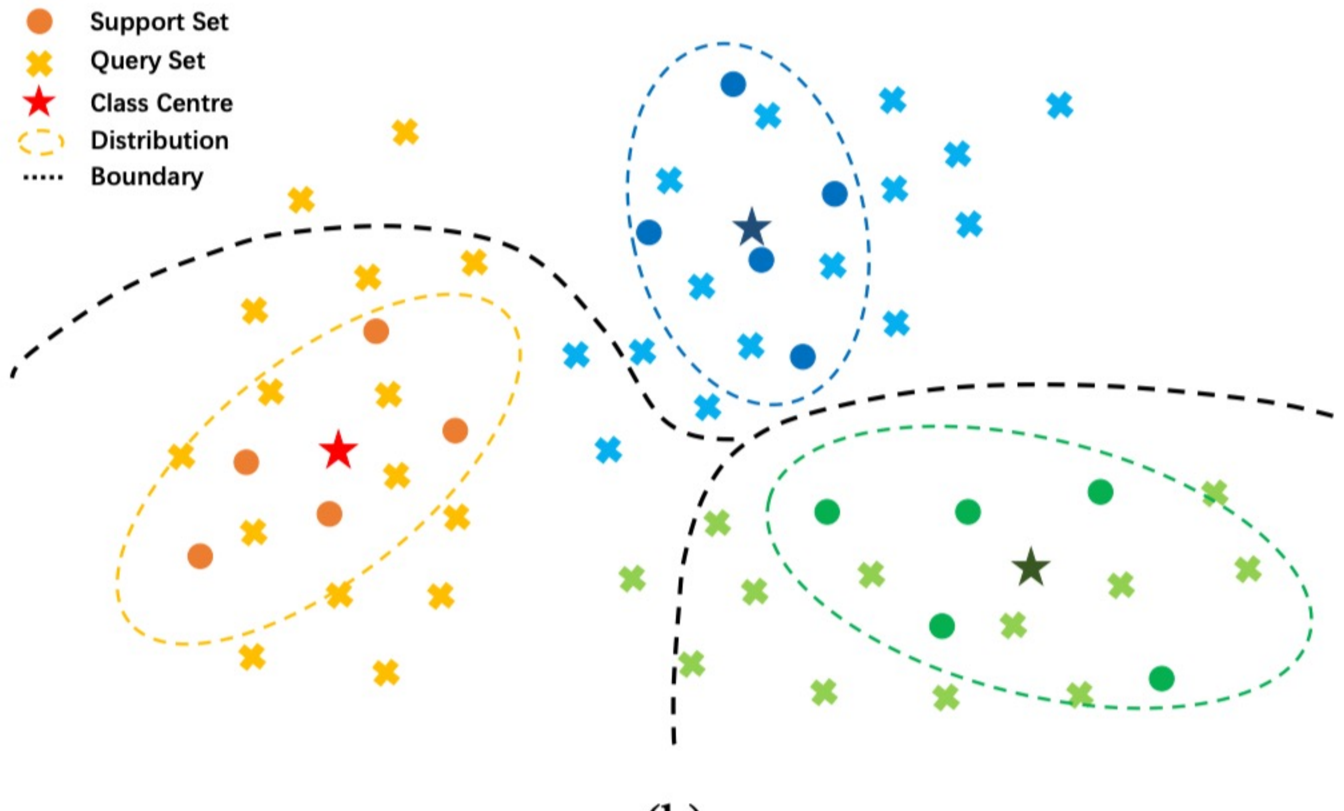

(b)

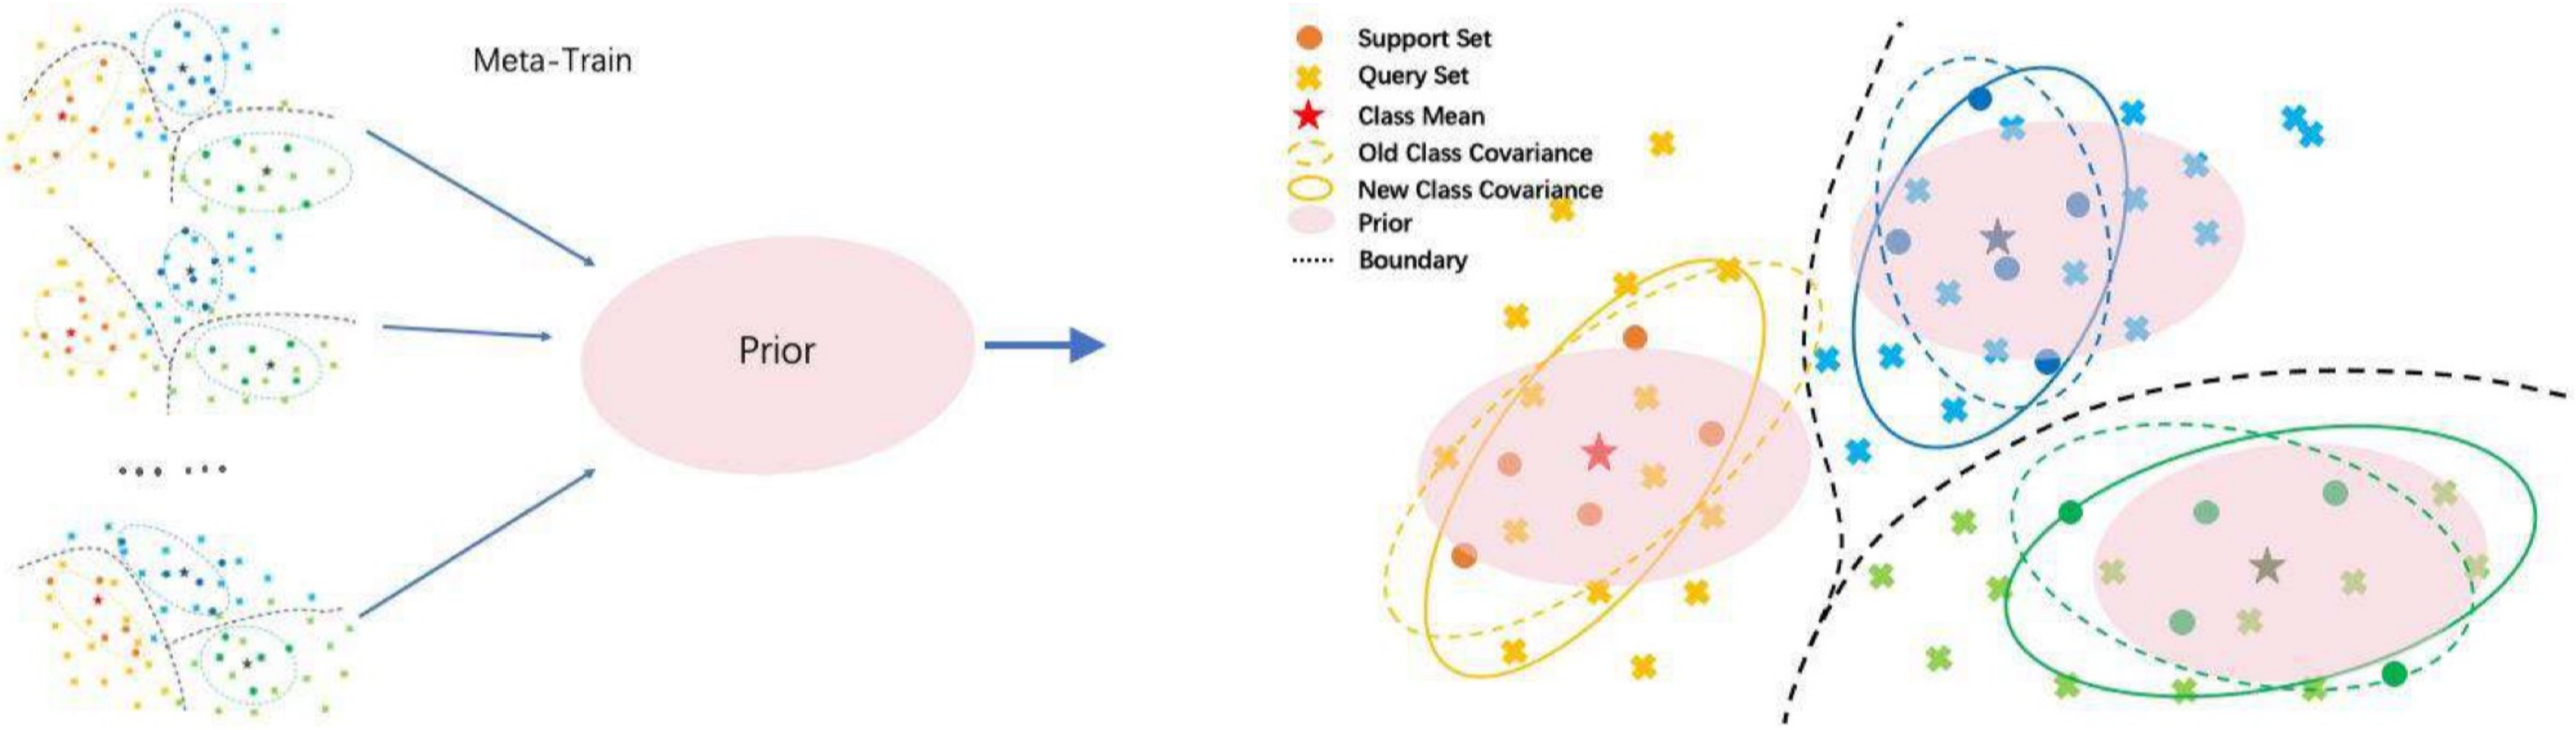

(c)

Zhang, et al. "Shallow bayesian meta learning for real-world few-shot recognition." ICCV. 2021.

Few-Shot Learning

- Learning from base data
- Adaptation on novel data
- FSL in 2020s

  Updating backbones

  Cross-domain FSL

  Foundation models as FSL learners

The Key Idea of "Updating backbones"

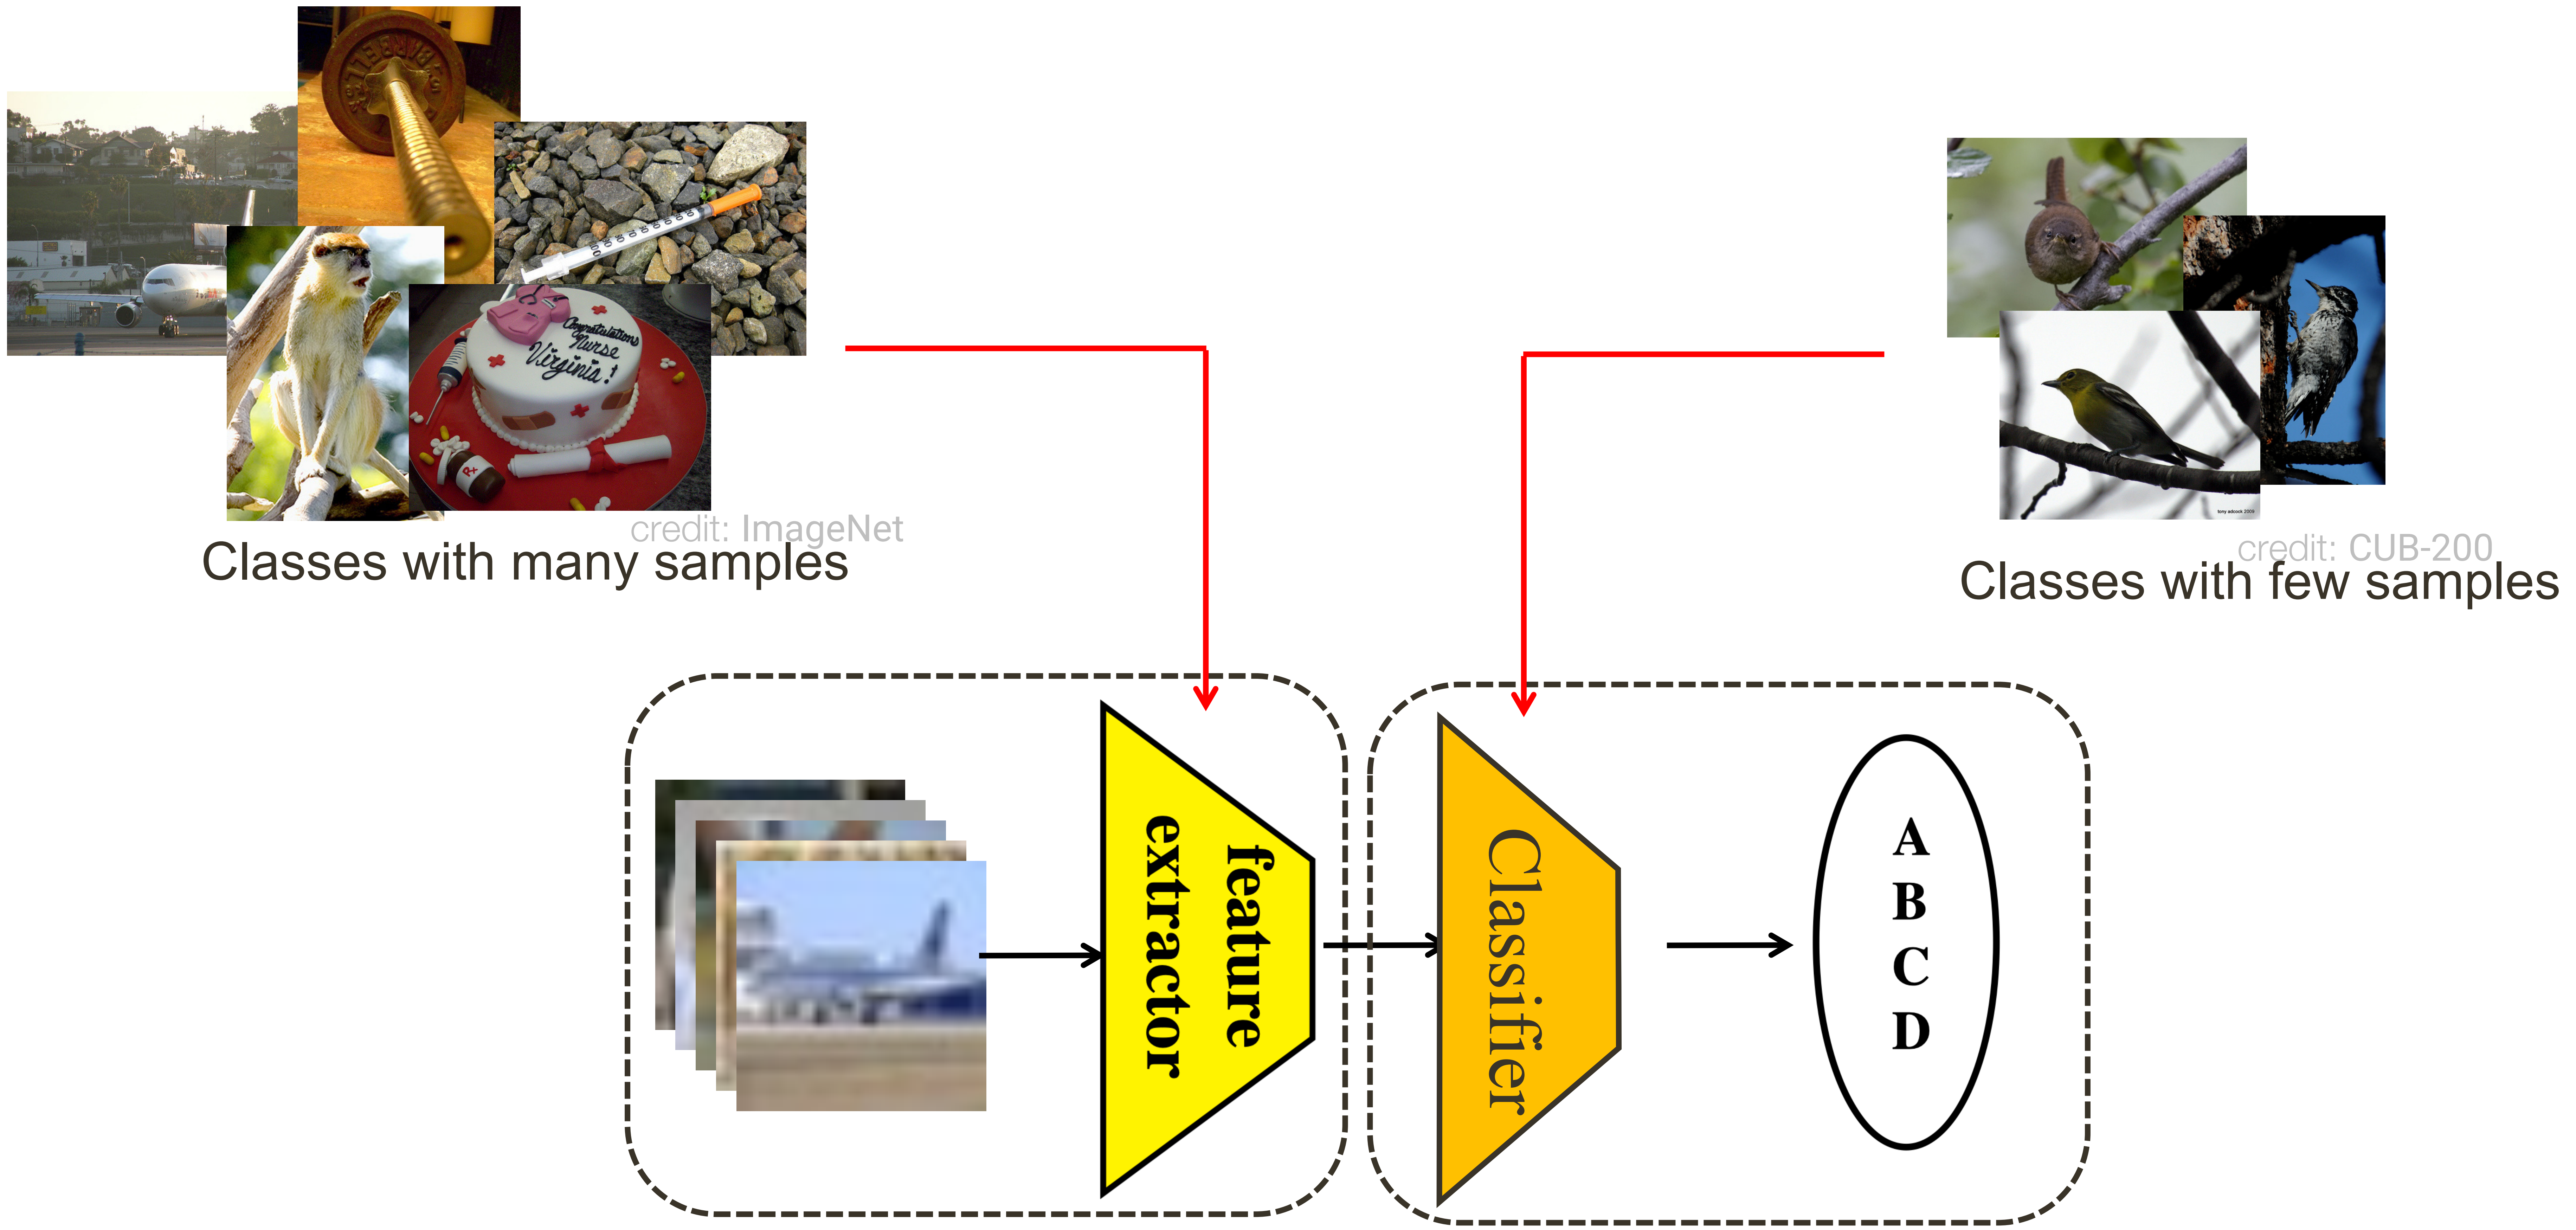

Can we enhance few-shot learning with recent development of deep learning?

Pushing the Limits of Simple Pipelines for Few-Shot Learning

What pipeline do we need in 2020s?

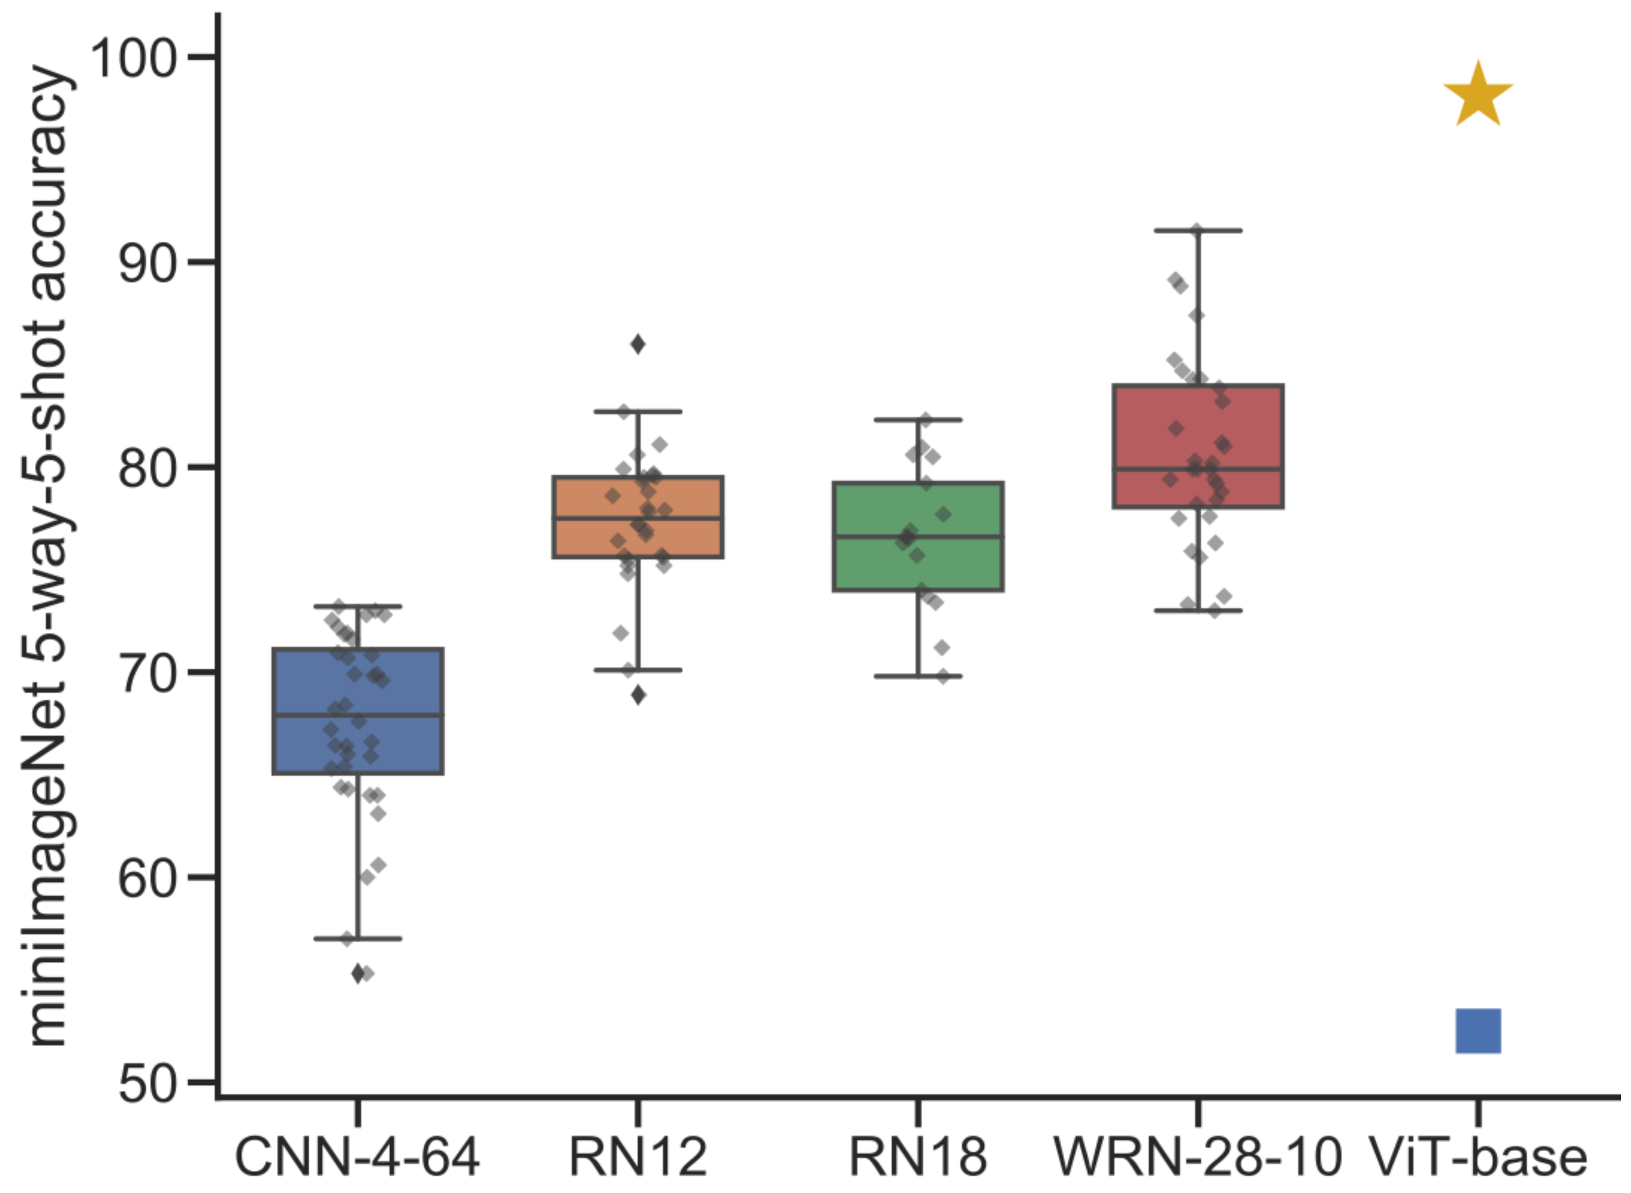

Larger backbone is better.

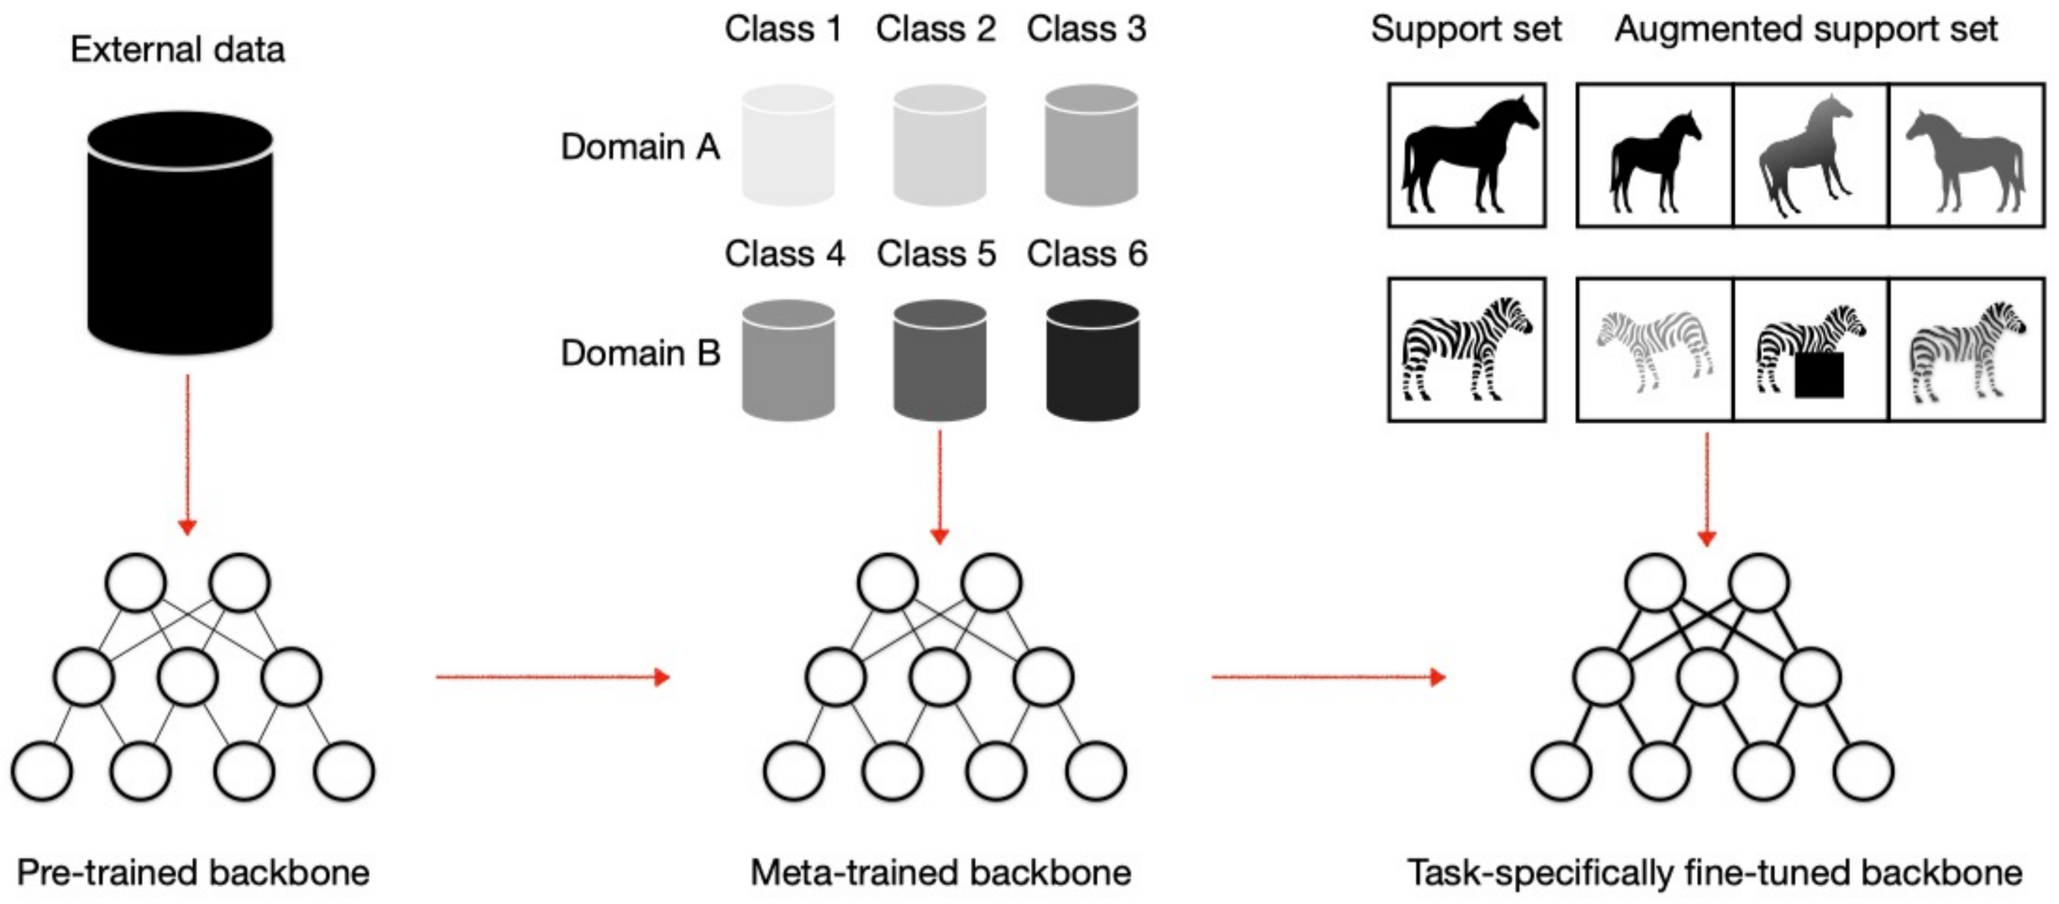

Both pre-train and fine-tuning are beneficial.

Hu,et al. "Pushing the limits of simple pipelines for few-shot learning: External data and fine-tuning make a difference." CVPR, 2022.

Efficient Transformer Tuning

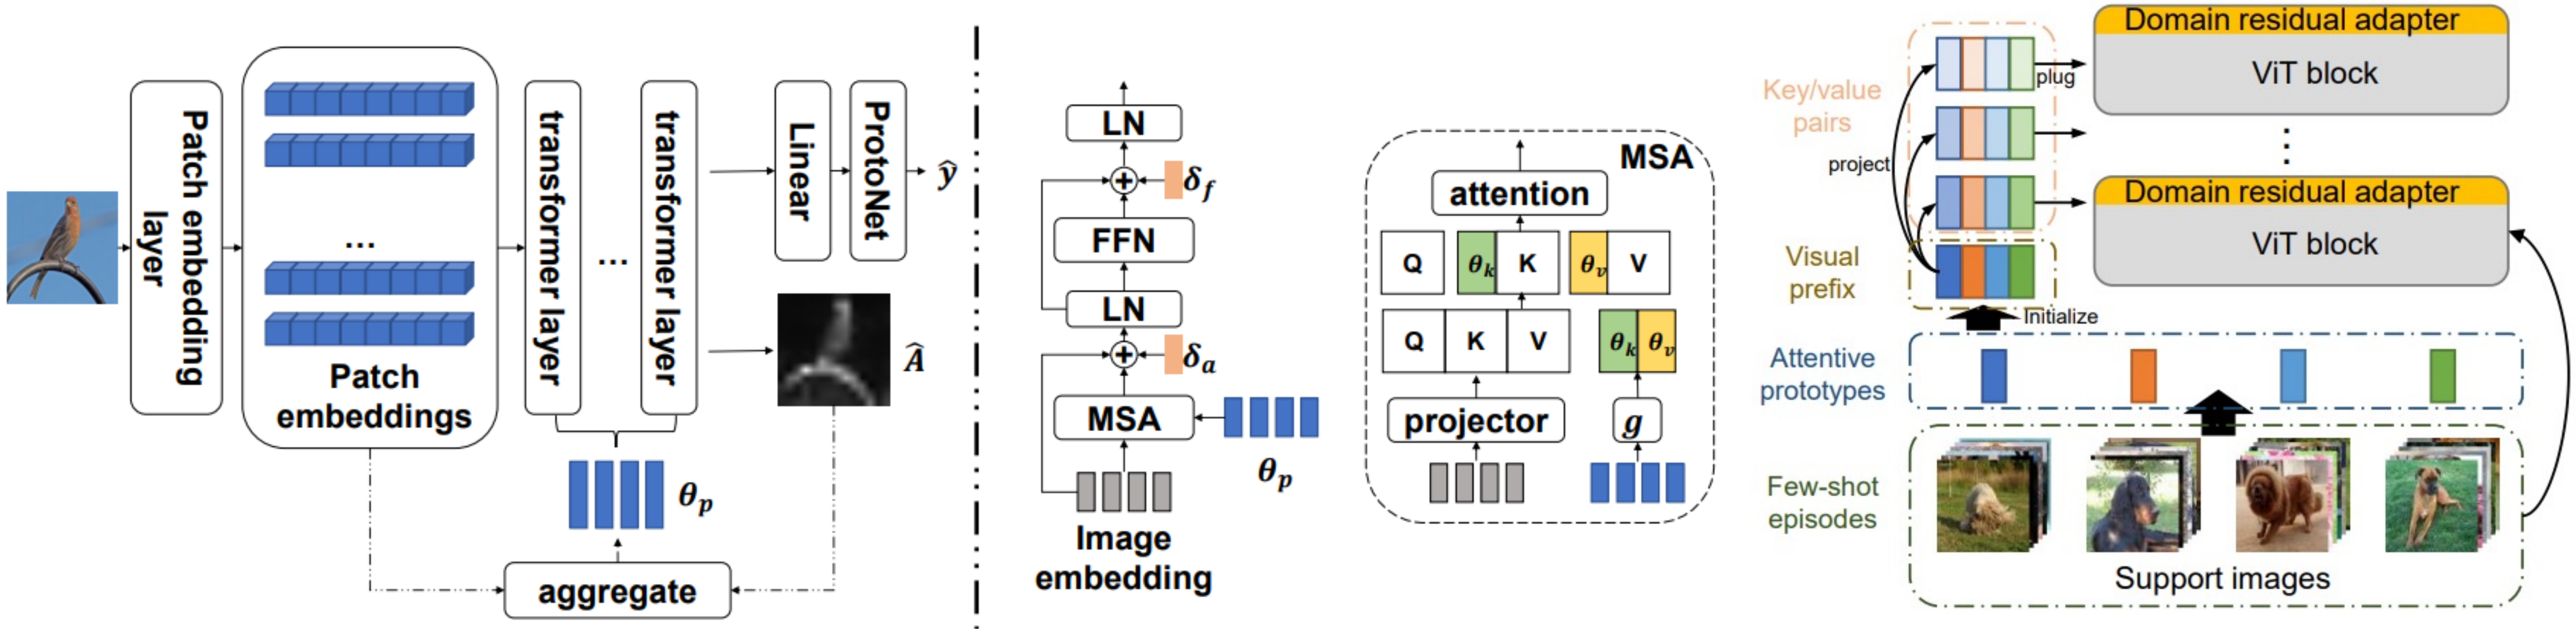

Fewer learnable parameters, while being flexible and effective enough.

Xu et al. Exploring Efficient Few-shot Adaptation for Vision Transformers. TMLR, 2022.

Cross-Domain Few-Shot Learning (CD-FSL)

Improve FSL models across different domains.

Target Data

Source Data

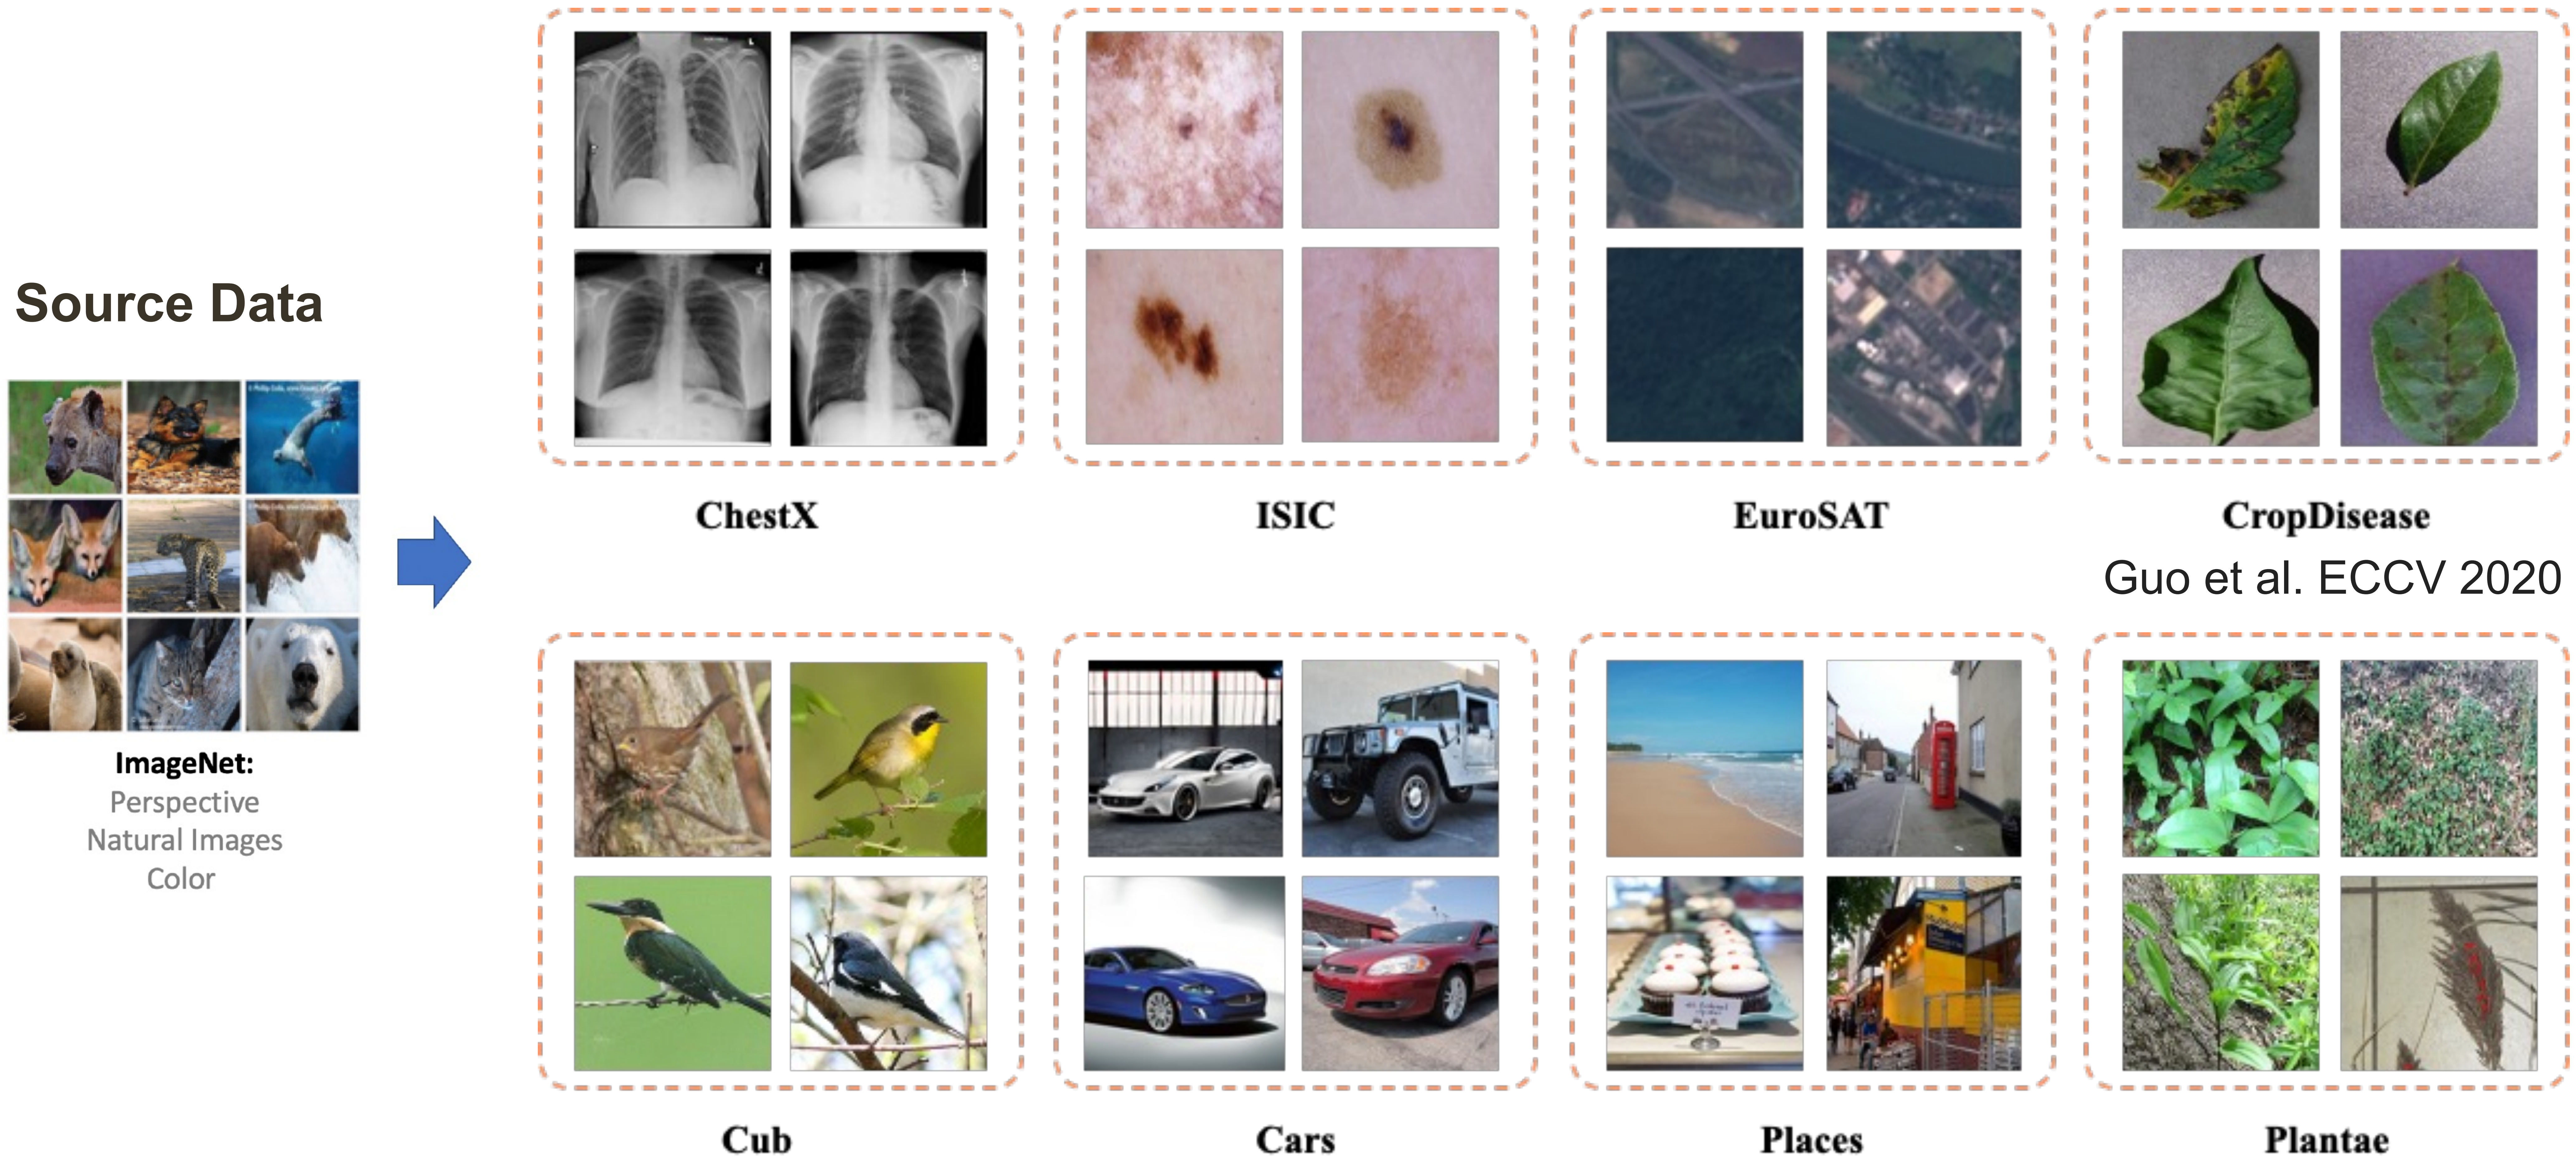

BSCD-FSL Benchmark

Guo et al. ECCV 2020

FWT-CDFSL Benchmark

Tseng et al. ICLR 2020

Guo, et al. A Broader Study of Cross-Domain Few-Shot Learning. ECCV2020
Tseng et al. Cross-Domain Few-Shot Classification via Learned Feature-Wise Transformation ICLR 2020

Cross-Domain Few-Shot Learning (CD-FSL)

Core Questions:

1. How to improve CD-FSL with only a source dataset as training data?
2. Could we utilize some target data for further promoting the CD-FSL?

CD-FSL

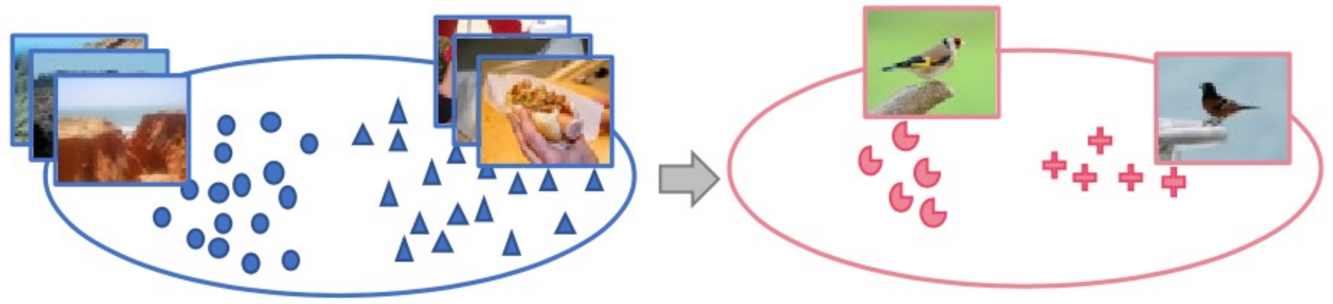

Overview:

Task	Cross-Domain Few-Shot Learning	
Setting	single-source setting	guided with few-labeled target data
	Promote	
Challenge	model generalization ability is limited	source and target data are imbalanced **(600:5)**
Insight	synthesize easy styles; synthesize hard styles	learn from mixed data; learn from individual knowledge
	upgrade	upgrade

Cross-Domain Few-Shot Learning

FSL across different Domains

StyleAdv: Adversarial style training

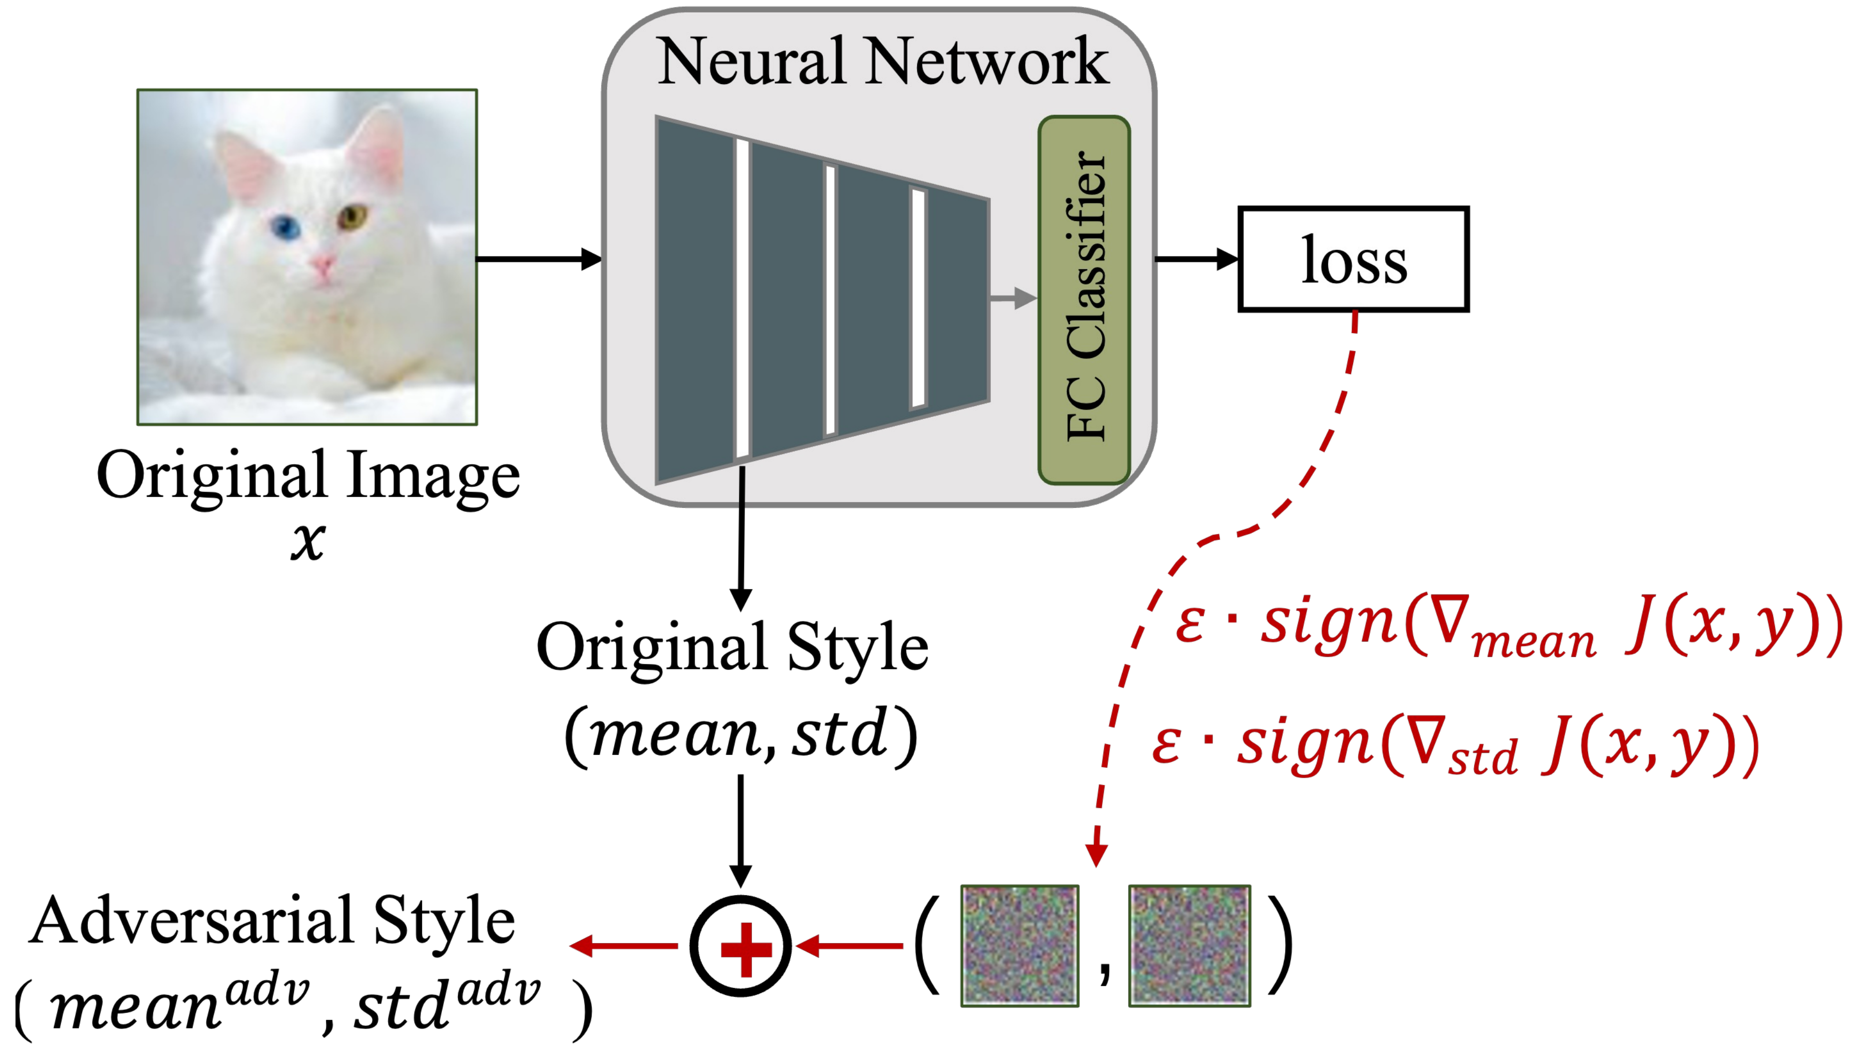

Novel Style Attack

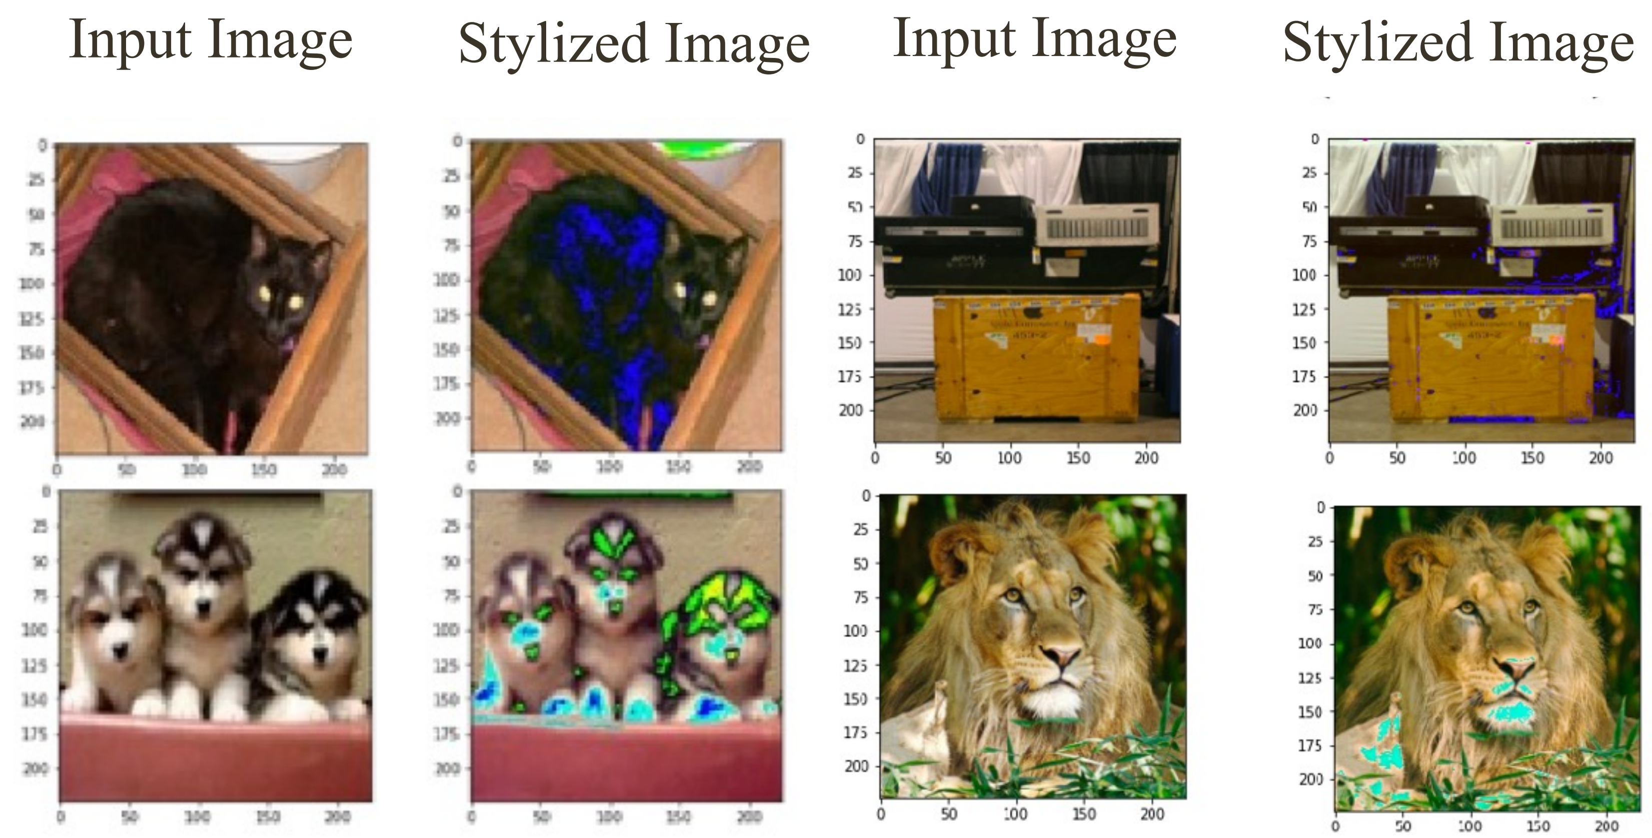

- "virtual" ,"hard" styles
- improve the generalization ability

Visualization Result

Fu et al. "StyleAdv: Meta Style Adversarial Training for Cross-Domain Few-Shot Learning." CVPR. 2023.

Cross-Domain Few-Shot Learning with Few-Labeled Target Data

Boost CD-FSL models with few labeled target data

Source Training Set (many examples) | Target Training Set (few examples) | Target Testing Set (few examples)

disjoint

- more realistic
- good performance
- don't violate the CD-FSL setting

Challenges

>> how to make use of the imbalanced source training set and target training set? (600:5)

>> how to narrow the domain gap between the source domain and target domain?

Fu et al. "Generalized Meta-FDMixup: Cross-Domain Few-Shot Learning Guided by Labeled Target Data." *TIP. 2022.*
Fu et al. "ME-D2N: Multi-Expert Domain Decompositional Network for Cross-Domain Few-Shot Learning." *ACM MM*. 2022.
Zhuo et al. "TGDM: Target Guided Dynamic Mixup for Cross-Domain Few-Shot Learning." *ACM MM*. 2022.

Foundation Models as Few-Shot Learners

Quite a lot of works on foundation models, to name a few:

GPT-4

Q: How many prime numbers are there between 150 and 250?

A: There are 13 prime numbers between 150 and 250.

credit: [Bubeck, et al. 2023]

Text-to-Text

Text-to-Image

credit: [Rombach et al. 2022]

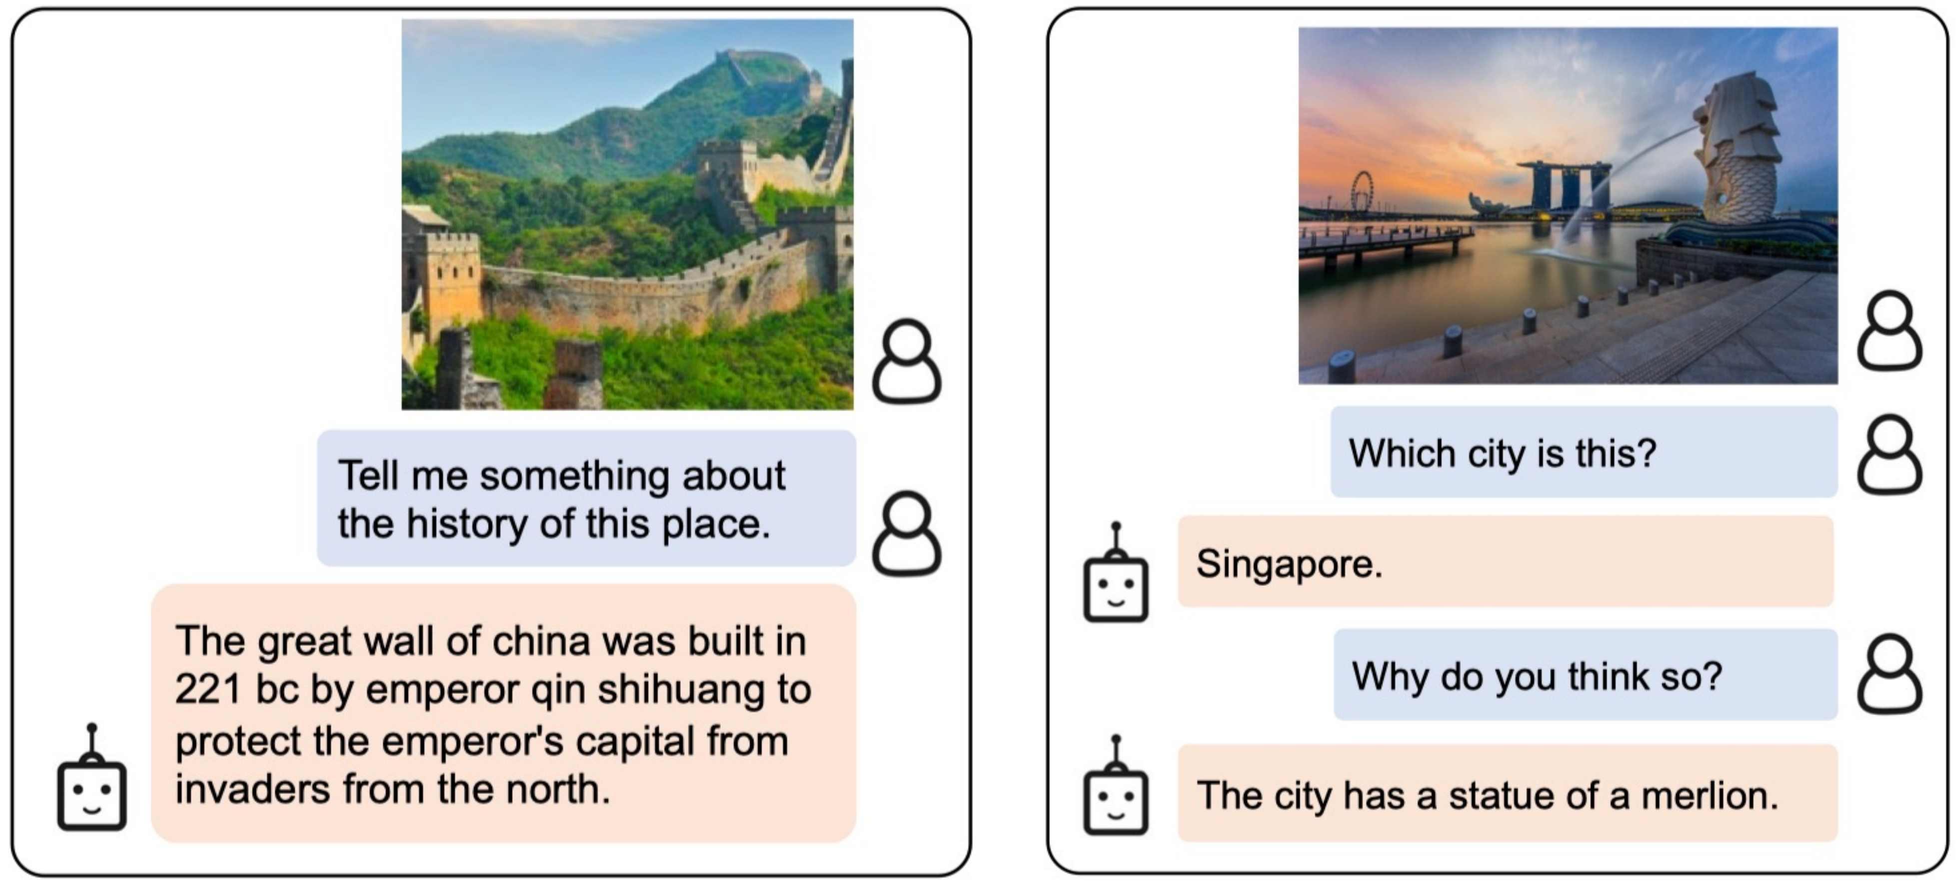

Image-to-Text

credit: [Li et al. 2023]

Image-to-Image

Bubeck, Sébastien, et al. "Sparks of artificial general intelligence: Early experiments with gpt-4." 2023.
Rombach, Robin, et al. "High-resolution image synthesis with latent diffusion models." CVPR. 2022.
Li, Junnan, et al. "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models." 2023.
Cao, Chenjie, et al. "A Unified Prompt-Guided In-Context Inpainting Framework for Reference-based Image Manipulations." 2023.

Foundation Models as Few-Shot Learners: Flamingo

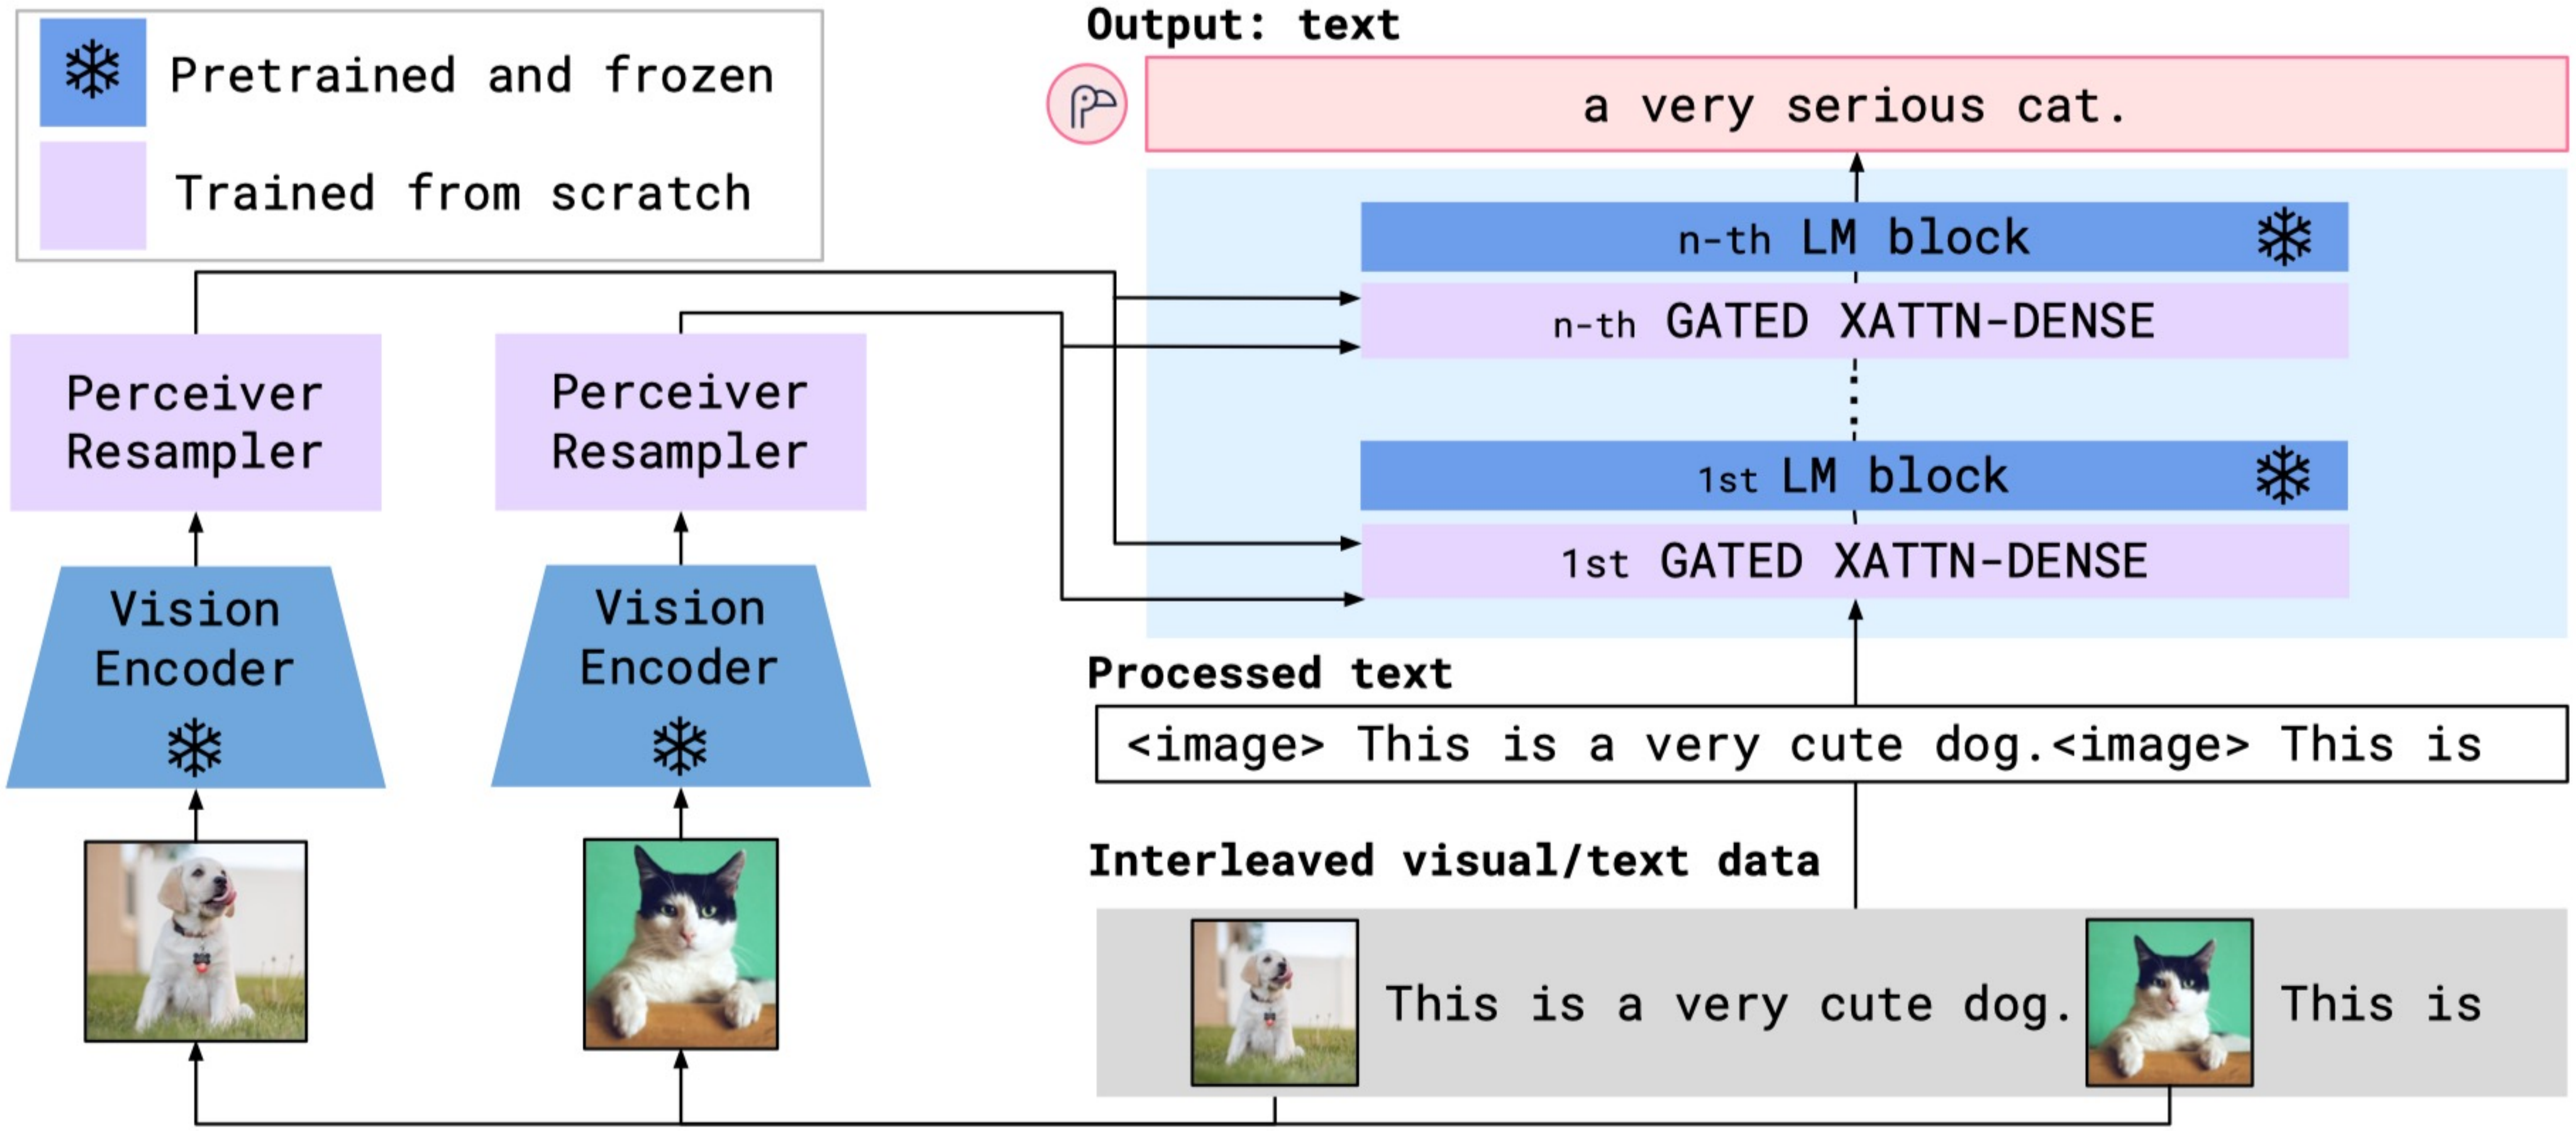

credit: [Alayrac et al. 2022]

Directly solving few-shot learning problems via the foundation model.

Alayrac, et al. "Flamingo: a visual language model for few-shot learning." NeurIPS, 2022.

In-context Learning

Prompt:

```
This is awesome! // Negative
This is bad! // Positive
Wow that movie was rad! // Positive
What a horrible show! //
```

Output:

```
Negative
```

credit: promptingguide.ai

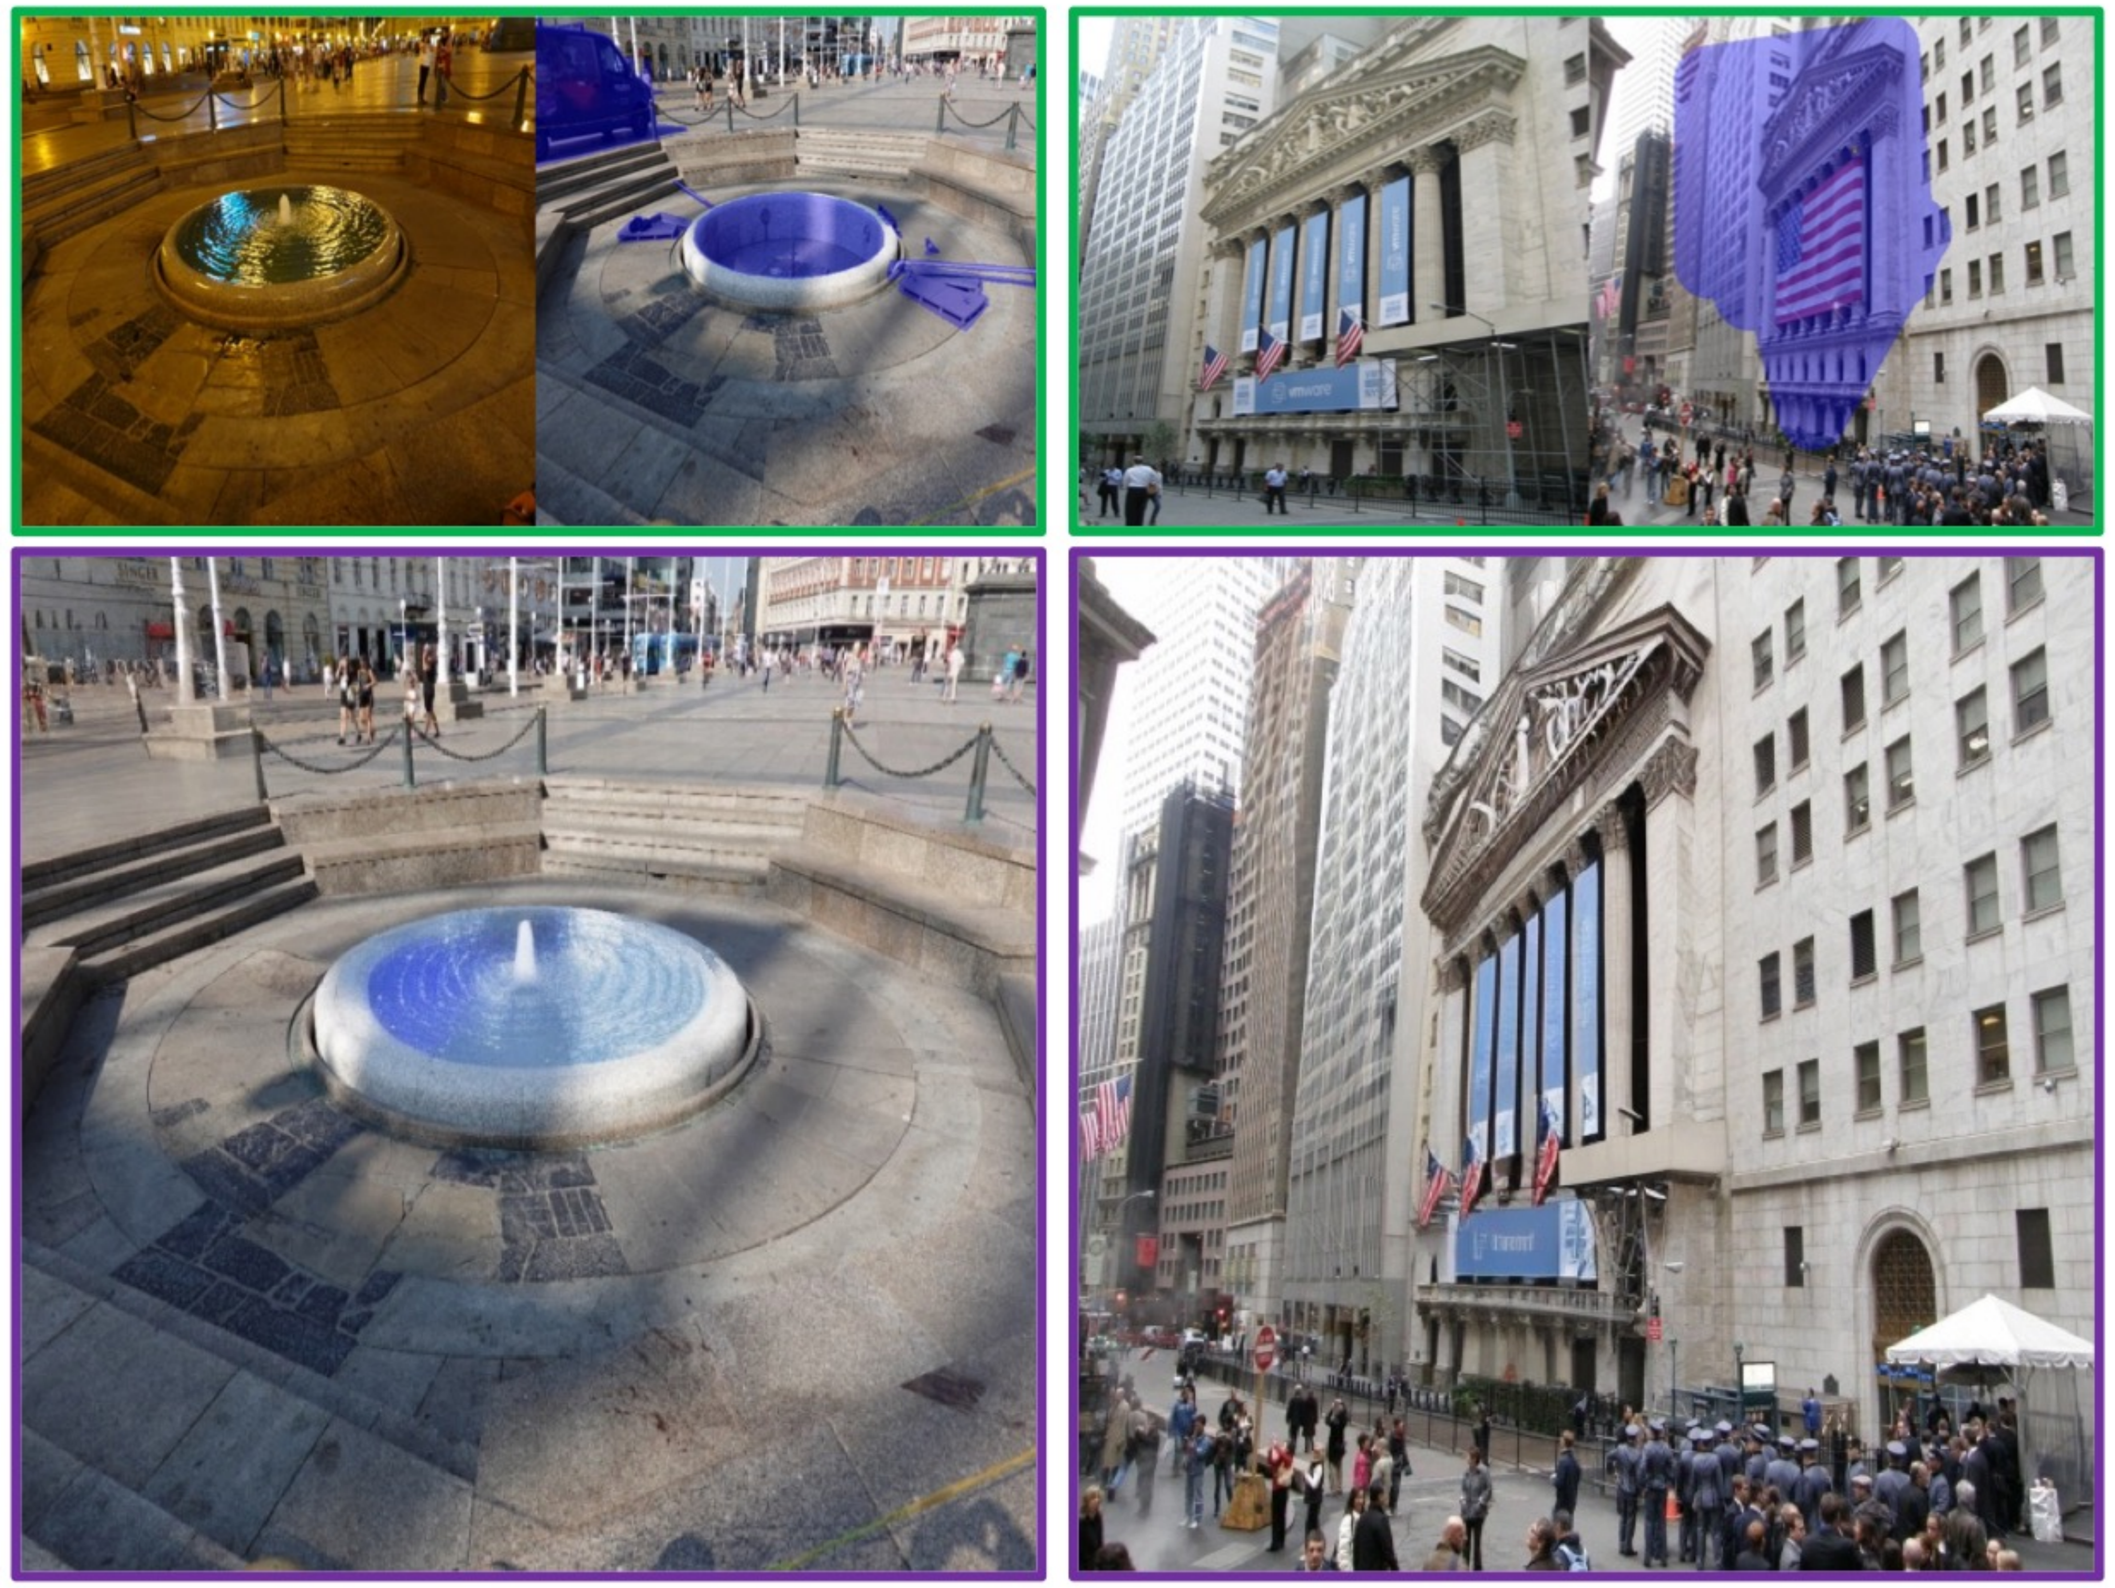

Take Home Message

To learn few-shot learning by statistical methods, we could,

- For learning from base data:
  - Learn causal-features that are truly transferable to novel tasks;
  - Model the similarity measurement in a more statistical way;
  - Utilize neural collapse to benefit few-shot learning.

- For adaptation on novel data:
  - Ensure the safety of learning with unlabeled data via statistical outlier detection;
  - Calibrate prototypes to reduce the bias;
  - Calibrate class distribution to prior knowledge;

- And for few-shot learning in 2020s:
  - Benefit few-shot learning with deeper architecture and more powerful pipeline;
  - Tackle more challenging cross-domain few-shot learning;
  - Adopt foundation models as few-shot learners.

Dr. Yanwei Fu

Yikai Wang

Chengming Xu

Yuqian Fu

Team to work on this talk in our group.

Yikai Wang

I’m looking for short-term visiting, or Potential Post-doc position (2024).

If there’s any chance, please email me:

yikaiwang19@fudan.edu.cn

THANKS

Paper List: Few-Shot Learning Revisited

- [Fu-SPM2018] Fu, Yanwei, et al. “Recent advances in zero-shot recognition: Toward data-efficient understanding of visual content.” IEEE SPM, 2018.
- [Li-IJCAI2021] Li, Pan, et al. “Regularising Knowledge Transfer by Meta Functional Learning .” IJCAI, 2021.
- [Lin-ICRA2022] Lin et al. I Know What You Draw: Learning Grasp Detection Conditioned on a Few Freehand Sketches. ICRA 2022
- [Pixel2mesh] Wang et al. Pixel2Mesh: 3D Mesh Model Generation via Image Guided Deformation. IEEE TPAMI 2020
- [In-context] Cao et al. A Unified Prompt-Guided In-Context Inpainting Framework for Reference-based Image Manipulations. Arxiv 2023
- [Zheng-CVPR2021] Zheng, et al. "Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers." CVPR. 2021.
- [Meta-Dataset] Triantafillou, et al. "Meta-Dataset: A Dataset of Datasets for Learning to Learn from Few Examples." ICLR2020.

Paper List: Learning from base data

- [Lake-CS2015] Lake, et al. "Human-level concept learning through probabilistic program induction." COGNITIVE SCIENCE. 2015
- [Lake-NeurIPS2013] Lake, et al. “One-shot learning by inverting a compositional causal process”, NeurIPS 2013
- [Yue-NeurIPS2020] Yue, et al. "Interventional few-shot learning." NeurIPS2020.
- [Xu-TPAMI2022] Xu, et al. "PatchMix Augmentation to Identify Causal Features in Few-shot Learning." TPAMI2022.
- [Vinyals-NeurIPS2016] Vinyals et al. “Matching Networks for One Shot Learning.” NeurIPS 2016.
- [Sung-CVPR2018] Sung et al. ” Learning to Compare: Relation Network for Few-Shot Learning”, CVPR 2018.
- [Zhang-CVPR2020] Zhang et al. "Deepemd: Few-shot image classification with differentiable earth mover's distance and structured classifiers." CVPR2020.
- [Zhang-TPAMI2021] Zhang, et al. "Deepemd: Differentiable earth mover's distance for few-shot learning." TPAMI2022.
- [Xie-CVPR2022] Xie et al. "Joint distribution matters: Deep brownian distance covariance for few-shot classification." CVPR2022.
- [Nguyen-ECCV2020] Nguyen et al. "Sen: A novel feature normalization dissimilarity measure for prototypical few-shot learning networks." ECCV 2020.
- [Papyan-PNAS2020] Papyan, Han, and Donoho. "Prevalence of neural collapse during the terminal phase of deep learning training." PNAS2020.
- [Liu-AISTATS2023] Liu, et al. "Inducing Neural Collapse in Deep Long-tailed Learning. " AISTATS2023.
- [Yang-NeurIPS2022] Yang, et al. "Do we really need a learnable classifier at the end of deep neural network?." NeurIPS2022.
- [Yang-ICLR2023] Yang, et al. "Neural Collapse Inspired Feature-Classifier Alignment for Few-Shot Class Incremental Learning." ICLR2023.
- [Galanti-ICLR2021] Galanti, et al. "On the role of neural collapse in transfer learning." ICLR2021.

Paper List: Adaptation on novel data

- [Li-TPAMI2014] Li, et al. "Towards making unlabeled data never hurt.” TPAMI 2014.
- [Li-TPAMI2019] Li, et al. "Towards safe weakly supervised learning.” TPAMI 2019.
- [Li-FCS2019] Li, et al. "Safe semi-supervised learning: a brief introduction.” FCS, 2019.
- [Wright-TPAMI2019] Wright et al. “Robust face recognition via sparse representation.” TPAMI 2009.
- [She-JASA2011] She et al. Outlier detection using nonconvex penalized regression. Journal of the American Statistical Association, 2011.
- [Fu-ECCV2014] Fu et al. “Interestingness Prediction by Robust Learning to Rank.” ECCV 2014,
- [Fu-TPAMI2016] Fu et al. “Robust Estimation of Subjective Visual Properties from Crowdsourced Pairwise Labels.” TPAMI 2016.
- [Fan-SS2018] Fan et al. “Partial Consistency with Sparse Incidental Parameters”. Statistical Sinica 2018.
- [Wang-CVPR2020] Wang et al. Instance Credibility Inference for Few-Shot Learning. CVPR 2020
- [Wang-TPAMI2021] Wang et al. How to Trust Unlabeled Data? Instance Credibility Inference for Few-Shot Learning. IEEE TPAMI 2021.
- [Friedman-JSS2010] Friedman, et al.. “Regularization Paths for Generalized Linear Models via Coordinate Descent.” Journal of Statistical Software, 2010
- [Li-CVPR2022] Li, et al. “Ranking distance calibration for cross-domain few-shot learning.” CVPR, 2022.
- [Han-ICLR2023] Han, et al. "Prototypical calibration for few-shot learning of language models." ICLR, 2023.
- [Yang-ICLR2021] Yang, et al. "Free lunch for few-shot learning: Distribution calibration." ICLR, 2021.
- [Guo-NeurIPS2022] Guo, et al. “Adaptive Distribution Calibration for Few-Shot Learning with Hierarchical Optimal Transport.” NeurIPS, 2022.
- [Zhang-ICCV2021] Zhang, et al. "Shallow bayesian meta learning for real-world few-shot recognition." ICCV. 2021.

Paper List: FSL in 2020s

- [Hu-CVPR2022] Hu,et al. "Pushing the limits of simple pipelines for few-shot learning: External data and fine-tuning make a difference." CVPR, 2022.
- [Xu-TMLR2022] Xu et al. Exploring Efficient Few-shot Adaptation for Vision Transformers. TMLR, 2022.
- [Fu-CVPR2023] Fu et al. "StyleAdv: Meta Style Adversarial Training for Cross-Domain Few-Shot Learning." CVPR. 2023.
- [Fu-TIP2022] Fu et al. “Generalized Meta-FDMixup: Cross-Domain Few-Shot Learning Guided by Labeled Target Data.” TIP. 2022.
- [Fu-MM2022] Fu et al. "ME-D2N: Multi-Expert Domain Decompositional Network for Cross-Domain Few-Shot Learning." ACM MM. 2022.
- [Zhuo-MM2022] Zhuo et al. "TGDM: Target Guided Dynamic Mixup for Cross-Domain Few-Shot Learning." ACM MM. 2022.
- [Bubeck-2023] Bubeck, et al. “Sparks of artificial general intelligence: Early experiments with gpt-4.” 2023.
- [Rombach-CVPR2022] Rombach, et al. "High-resolution image synthesis with latent diffusion models." CVPR. 2022.
- [Li-2023] Li, Junnan, et al. "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models." 2023.
- [In-Context] Cao, et al. "A Unified Prompt-Guided In-Context Inpainting Framework for Reference-based Image Manipulations." 2023.
- [Alayrac-NeurIPS2022] Alayrac, et al. "Flamingo: a visual language model for few-shot learning." NeurIPS, 2022.