ECCV'20
ONLINE
23-28 AUGUST 2020

16TH EUROPEAN CONFERENCE ON

COMPUTER VISION

WWW.ECCV2020.EU

Domain Adaptation for Visual Applications
Part 3: Beyond Classical Domain Adaptation

Tatiana Tommasi
Assistant Professor, Polytechnic University of Turin, Italy
Affiliated Researcher, Italian Institute of Technology

Outline

- An Overview on several Cross-Domain Learning Settings
  - (annotated) source data
  - (annotated) target data
  - source / target overlap

- Self-Supervision for Cross-Domain Learning

Classical Domain Adaptation

Source
(Train)

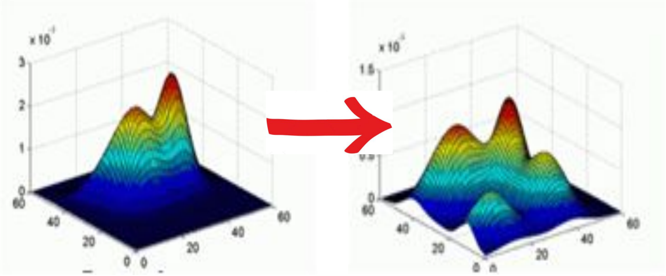

Target
(Test)

Classical Domain Adaptation

Source
(Train)

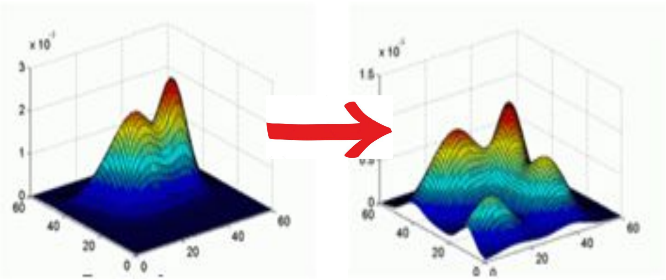

Target
(Test)

Labelled
Source Domain

Unlabelled
Target Domain

Unsupervised DA,
transductive setting

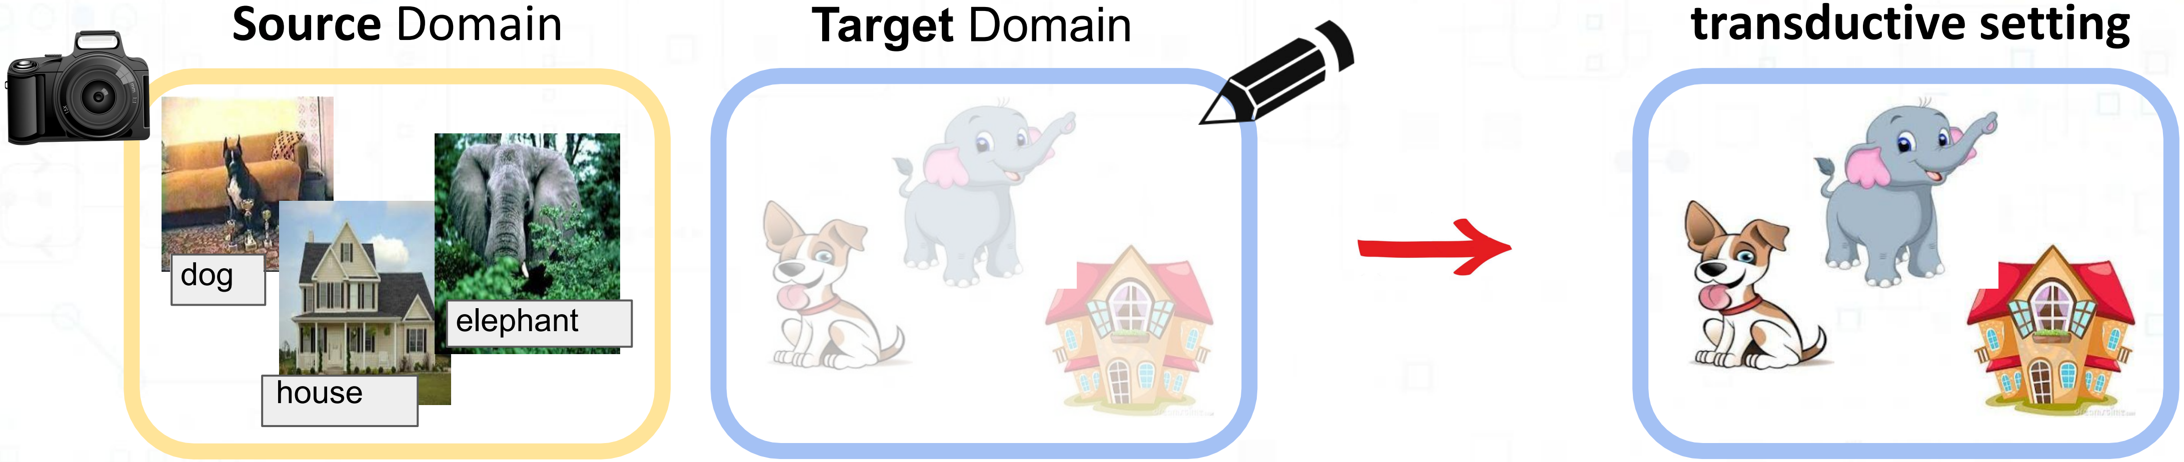

Train

Test

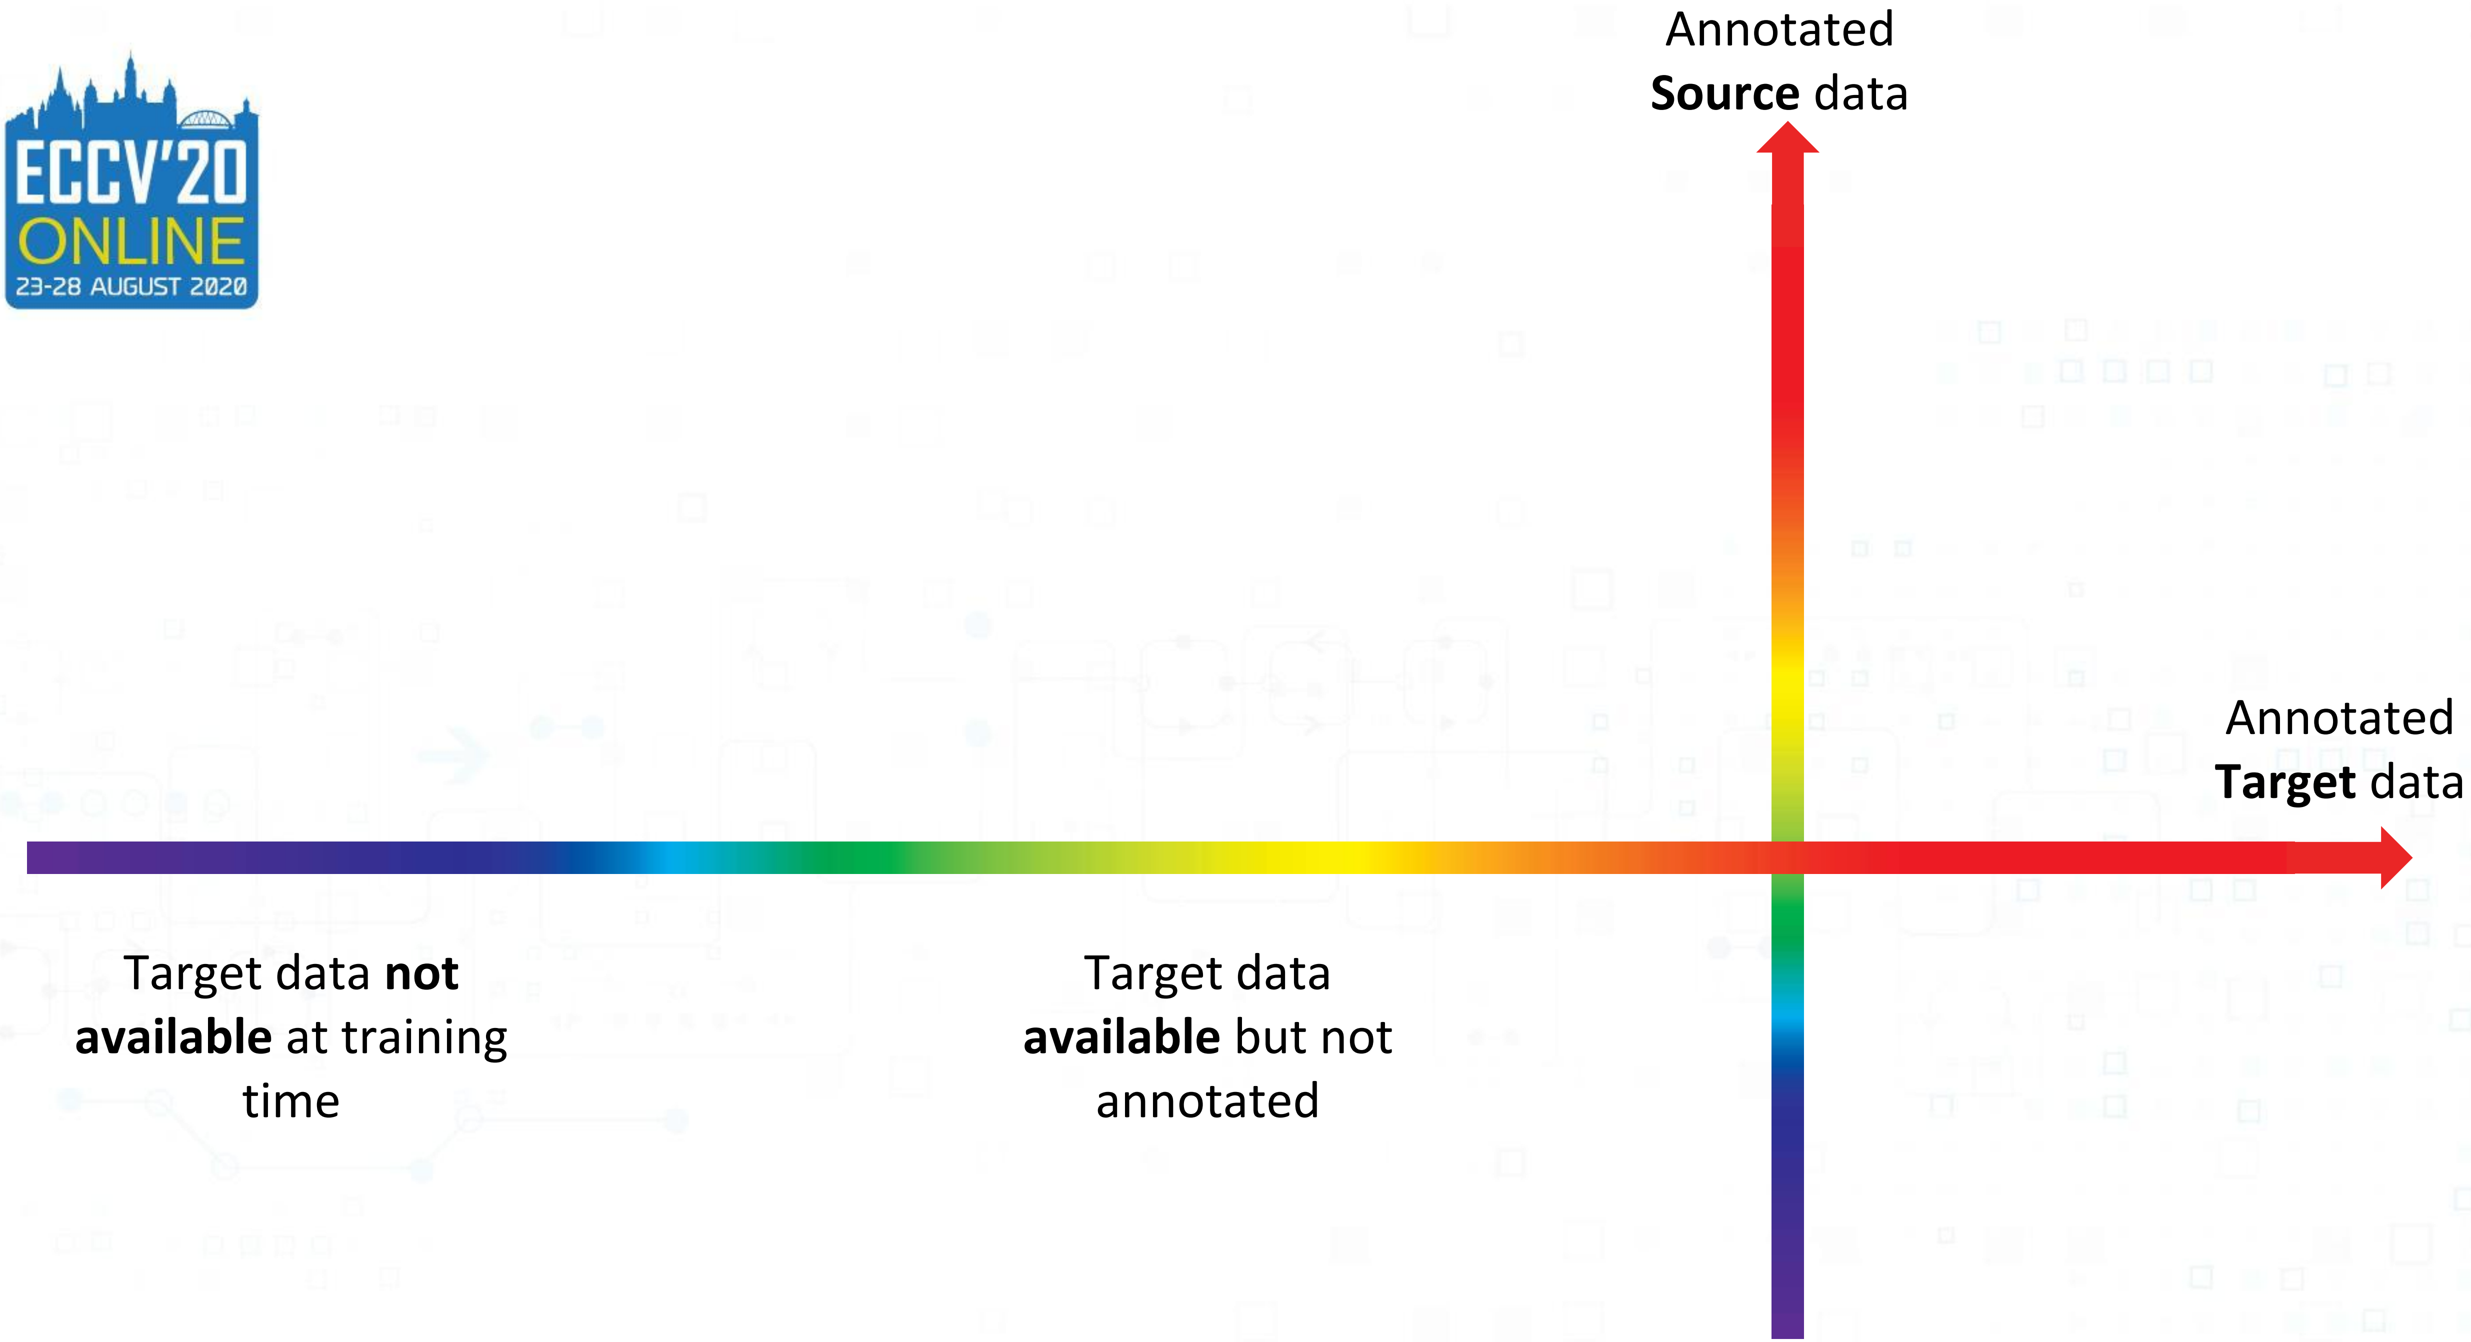
Annotated
Source data
Annotated
Target data
Target data **not available** at training time
Target data **available** but not annotated

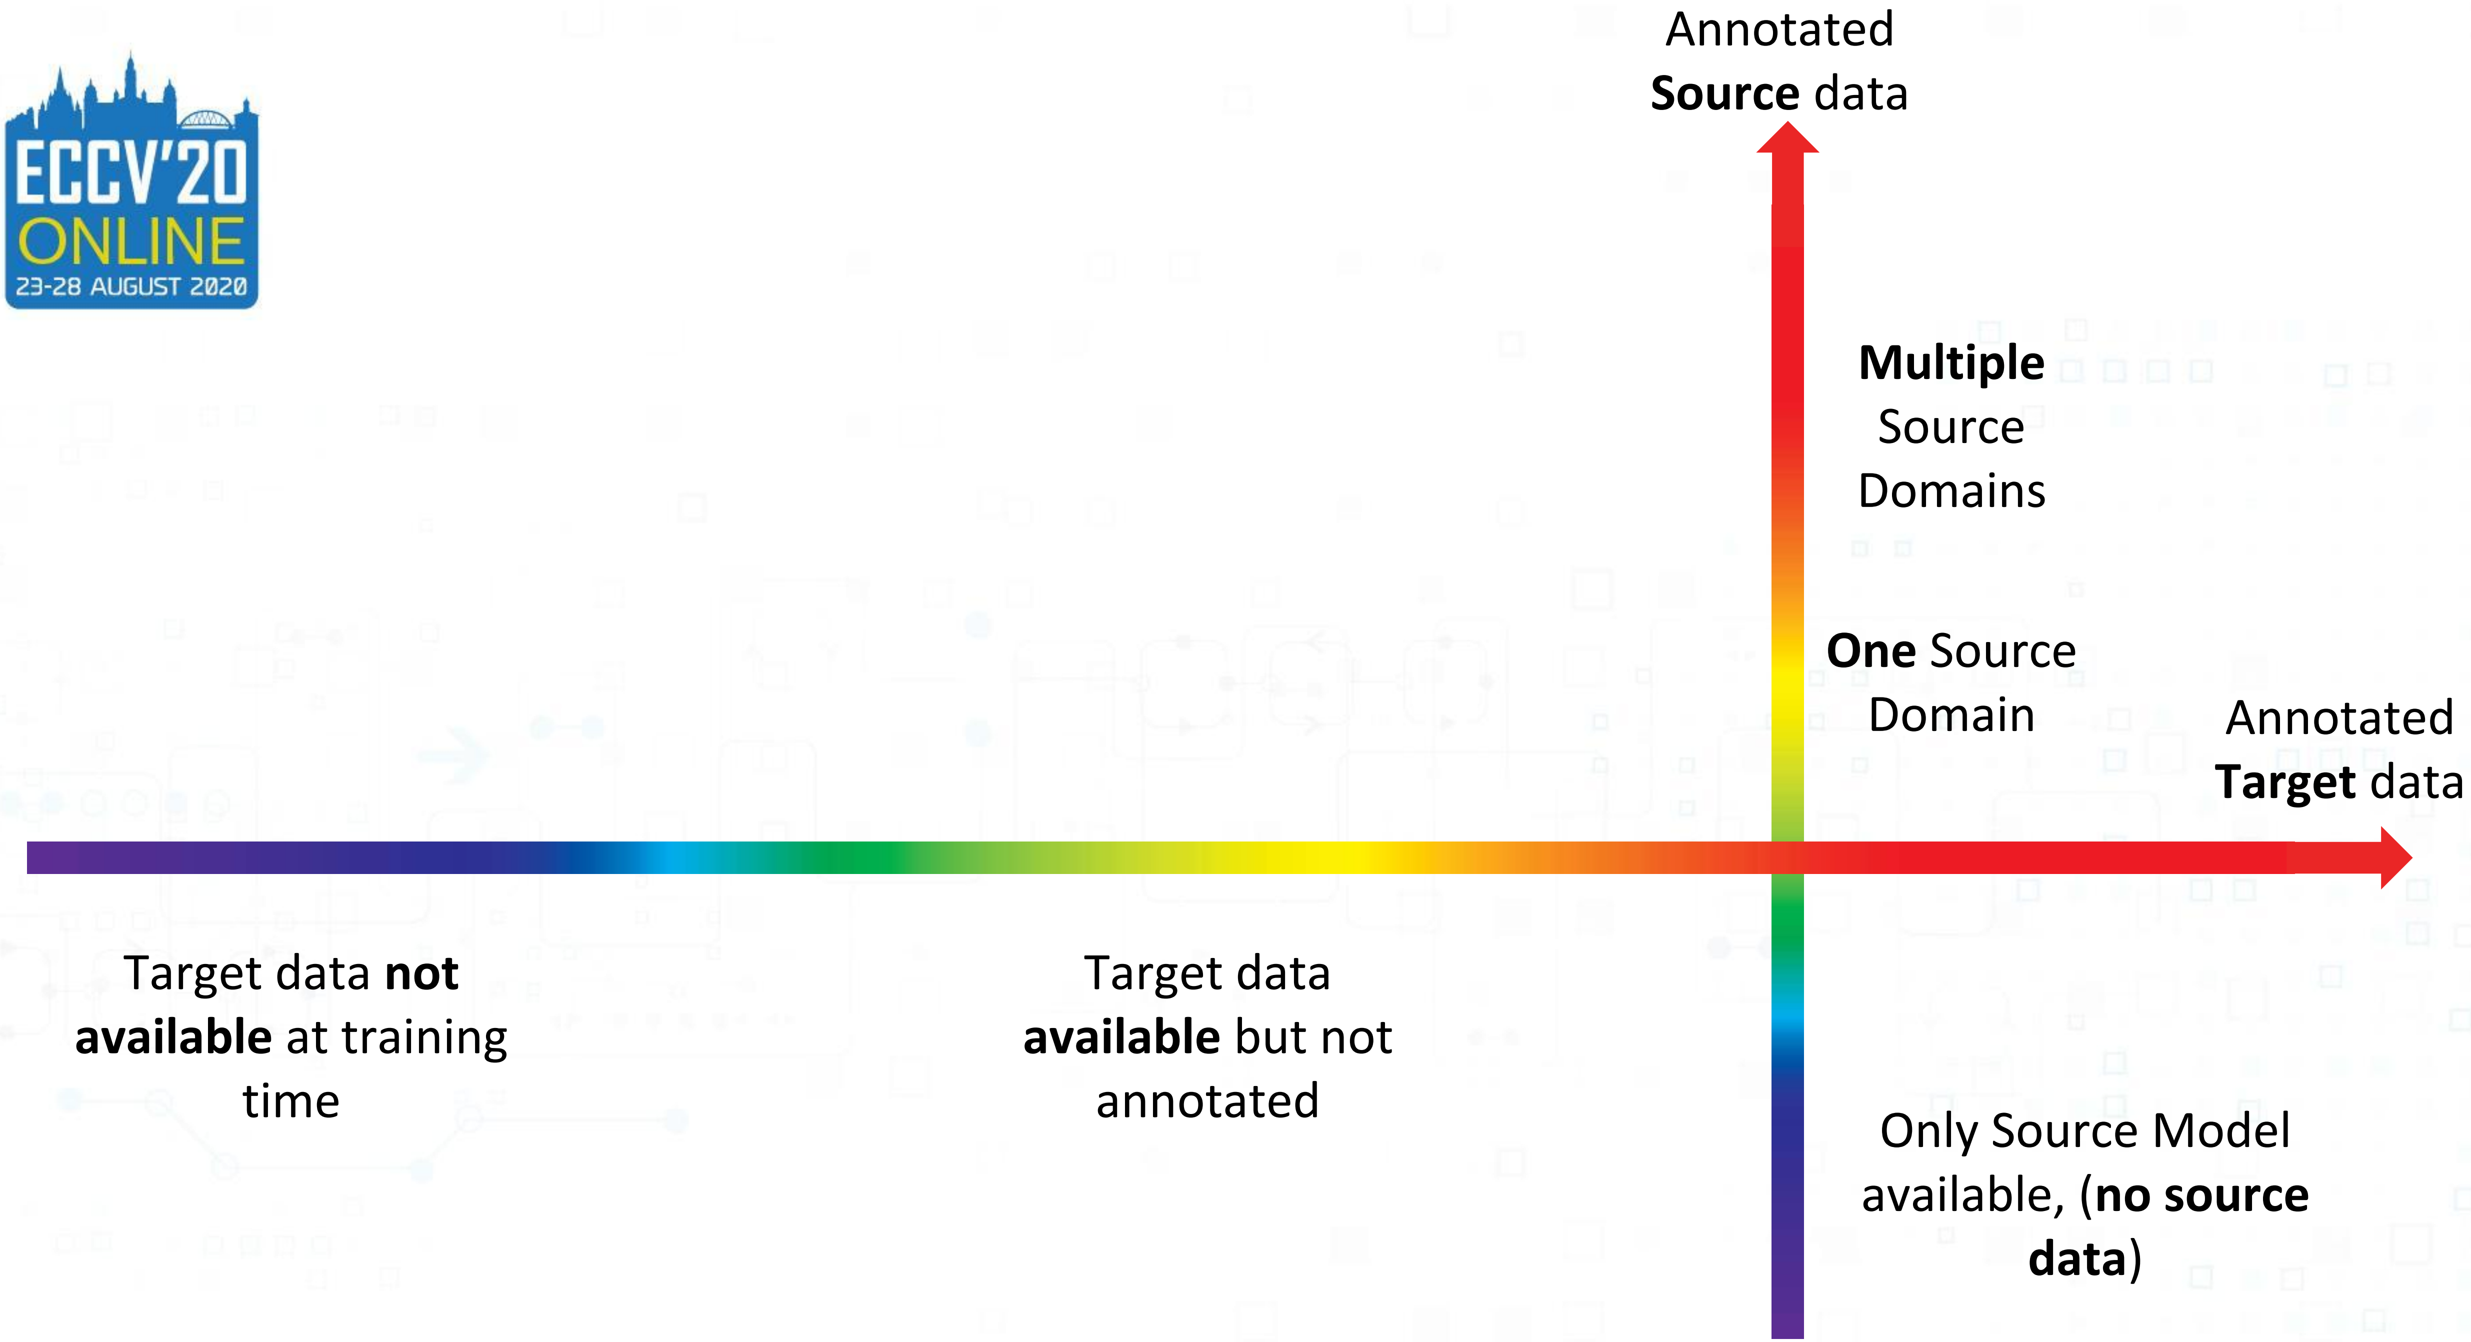
Annotated **Source** data
Multiple Source Domains
One Source Domain
Annotated **Target** data
Target data **not available** at training time
Target data **available** but not annotated
Only Source Model available, (**no source data**)

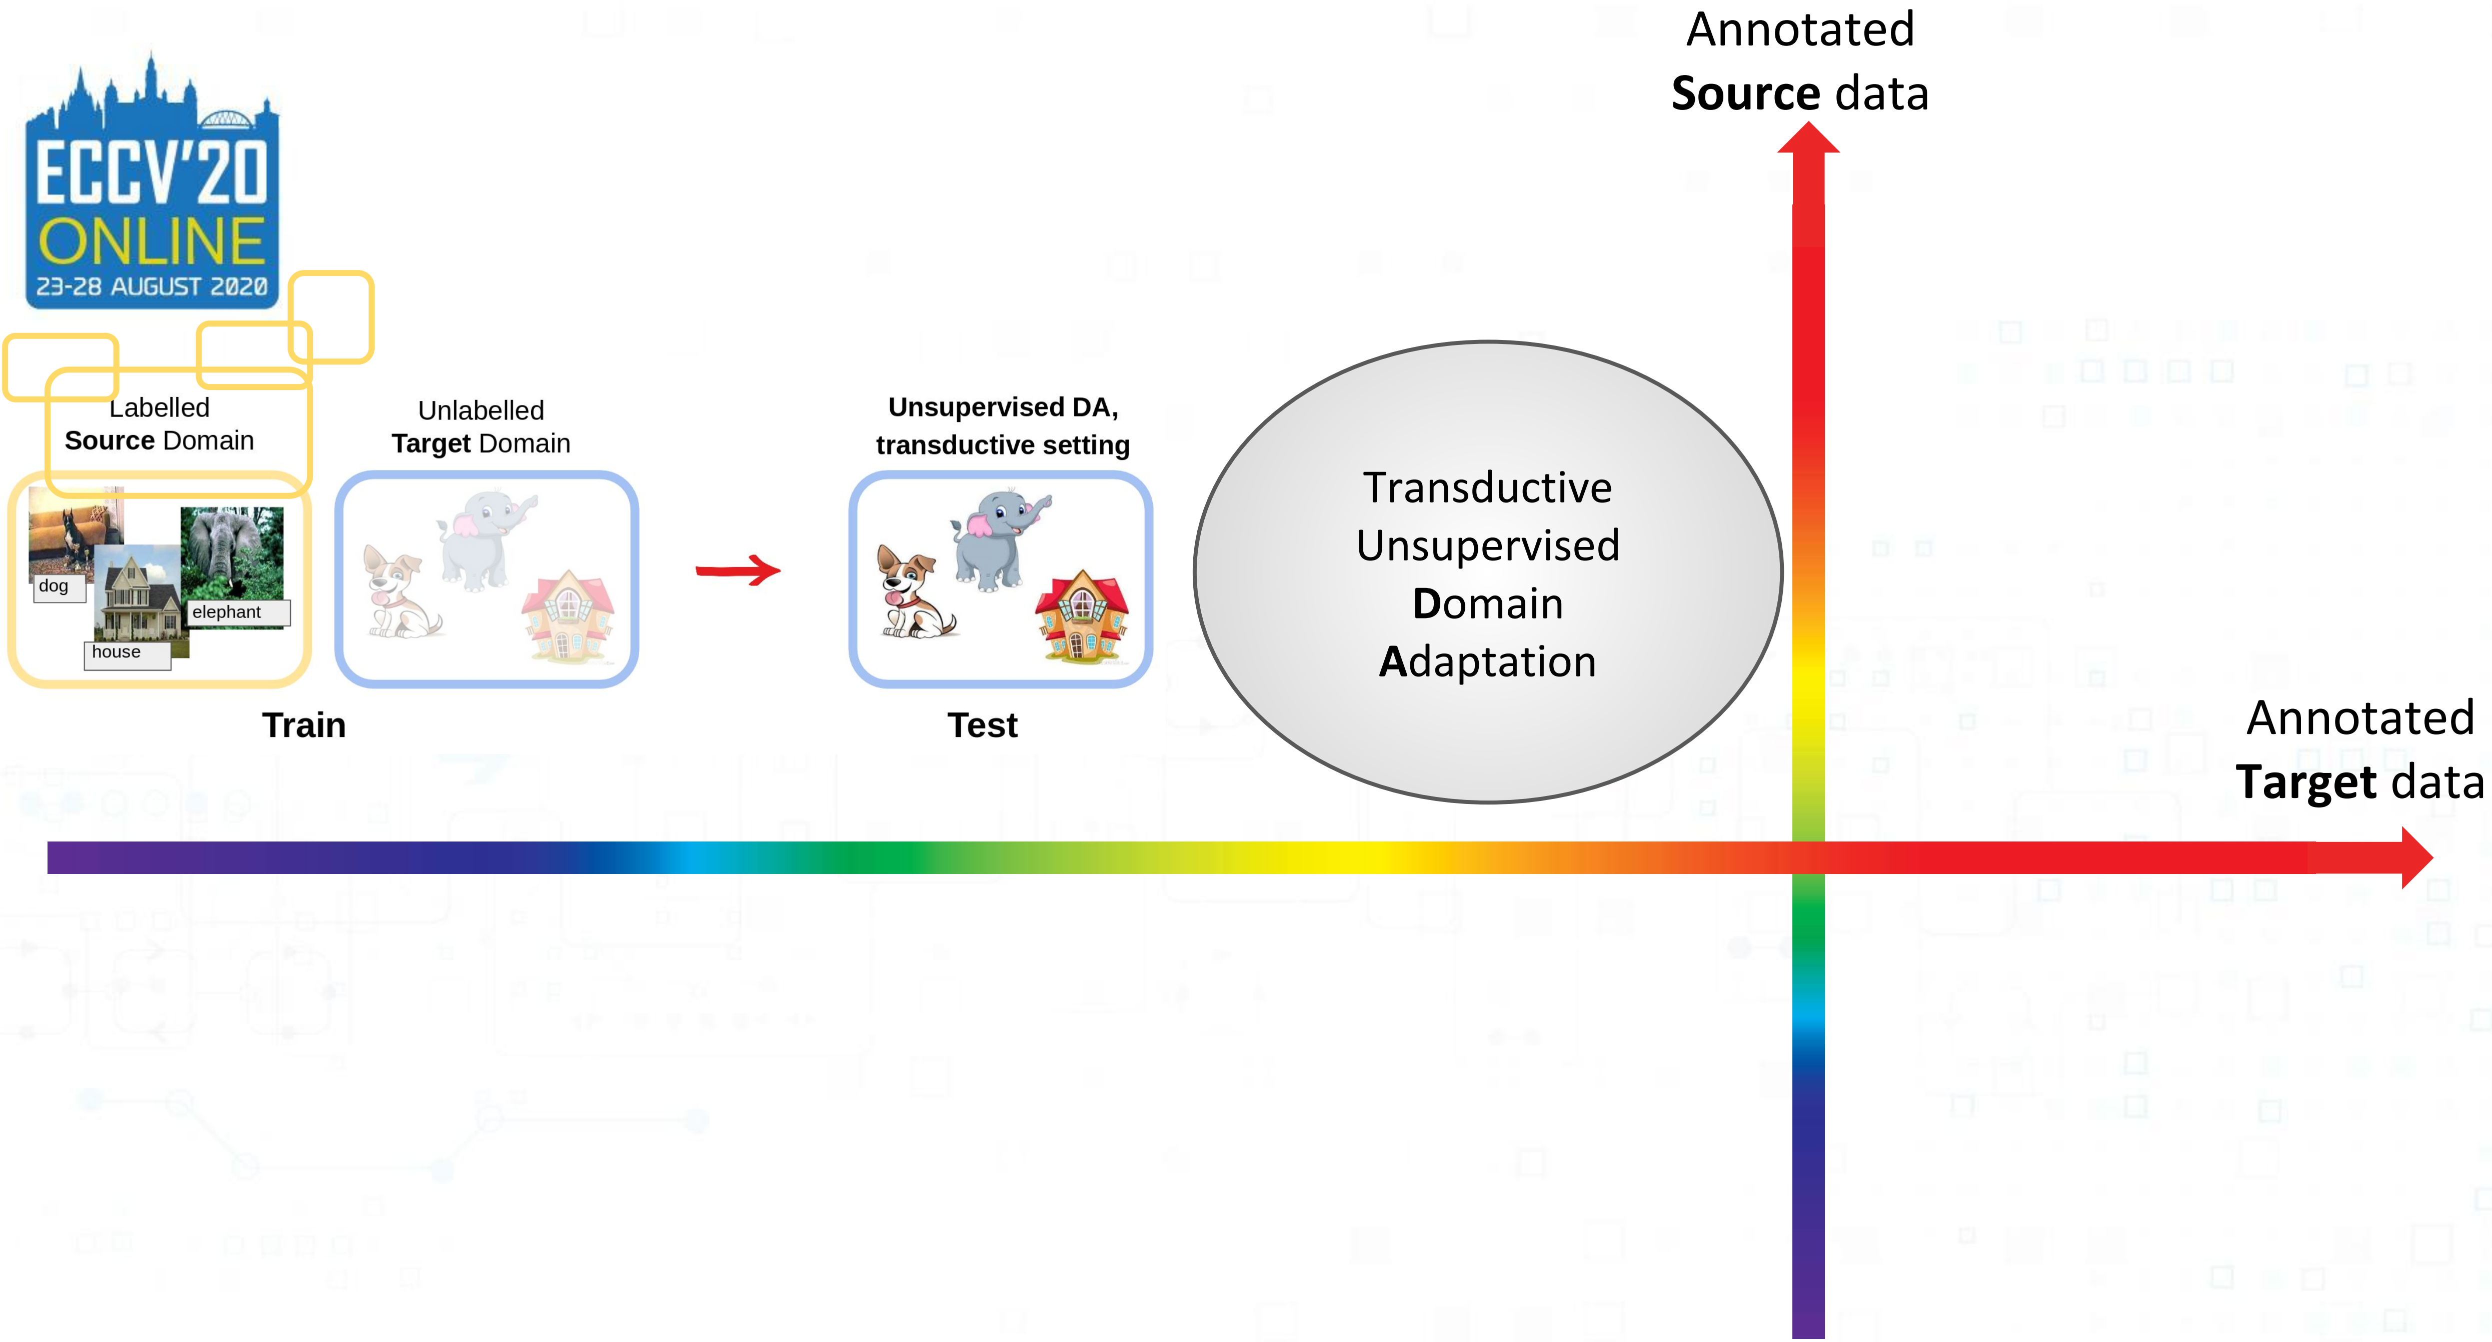
Annotated
Source data
Labelled
Source Domain
dog
elephant
house
Unlabelled
Target Domain
Train
Unsupervised DA,
transductive setting
Test
Transductive
Unsupervised
Domain
Adaptation
Annotated
Target data

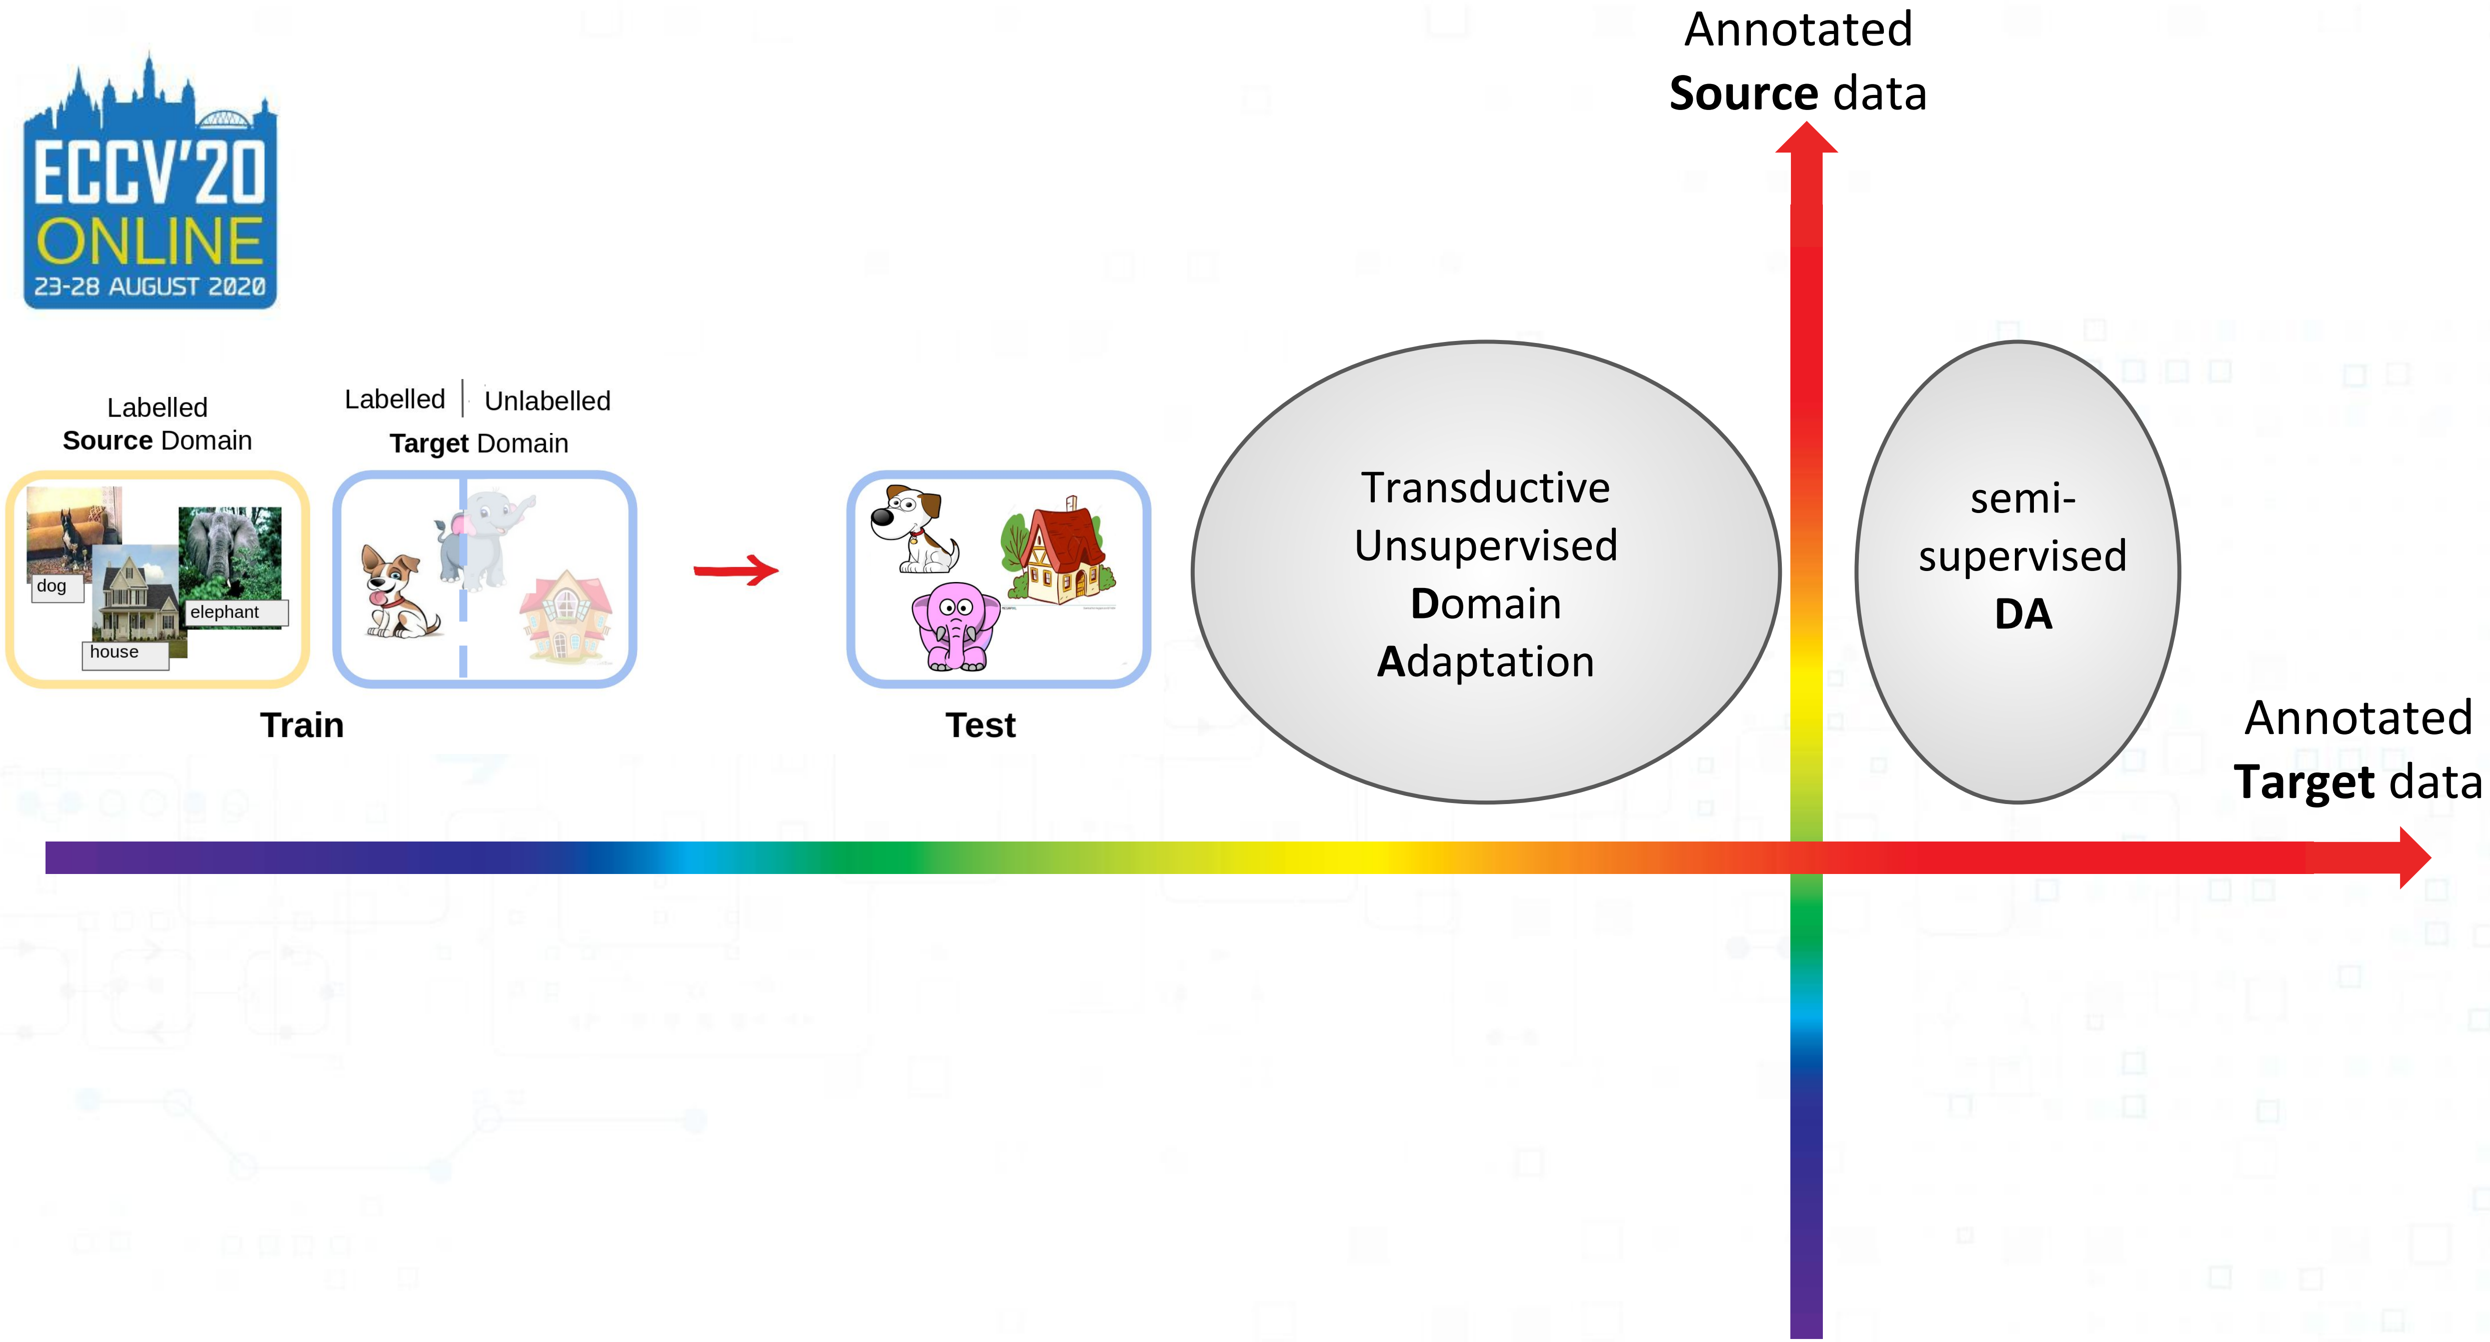
Annotated
Source data
Labelled
Source Domain
dog
elephant
house
Labelled
Unlabelled
Target Domain
Train
Test
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
Annotated
Target data

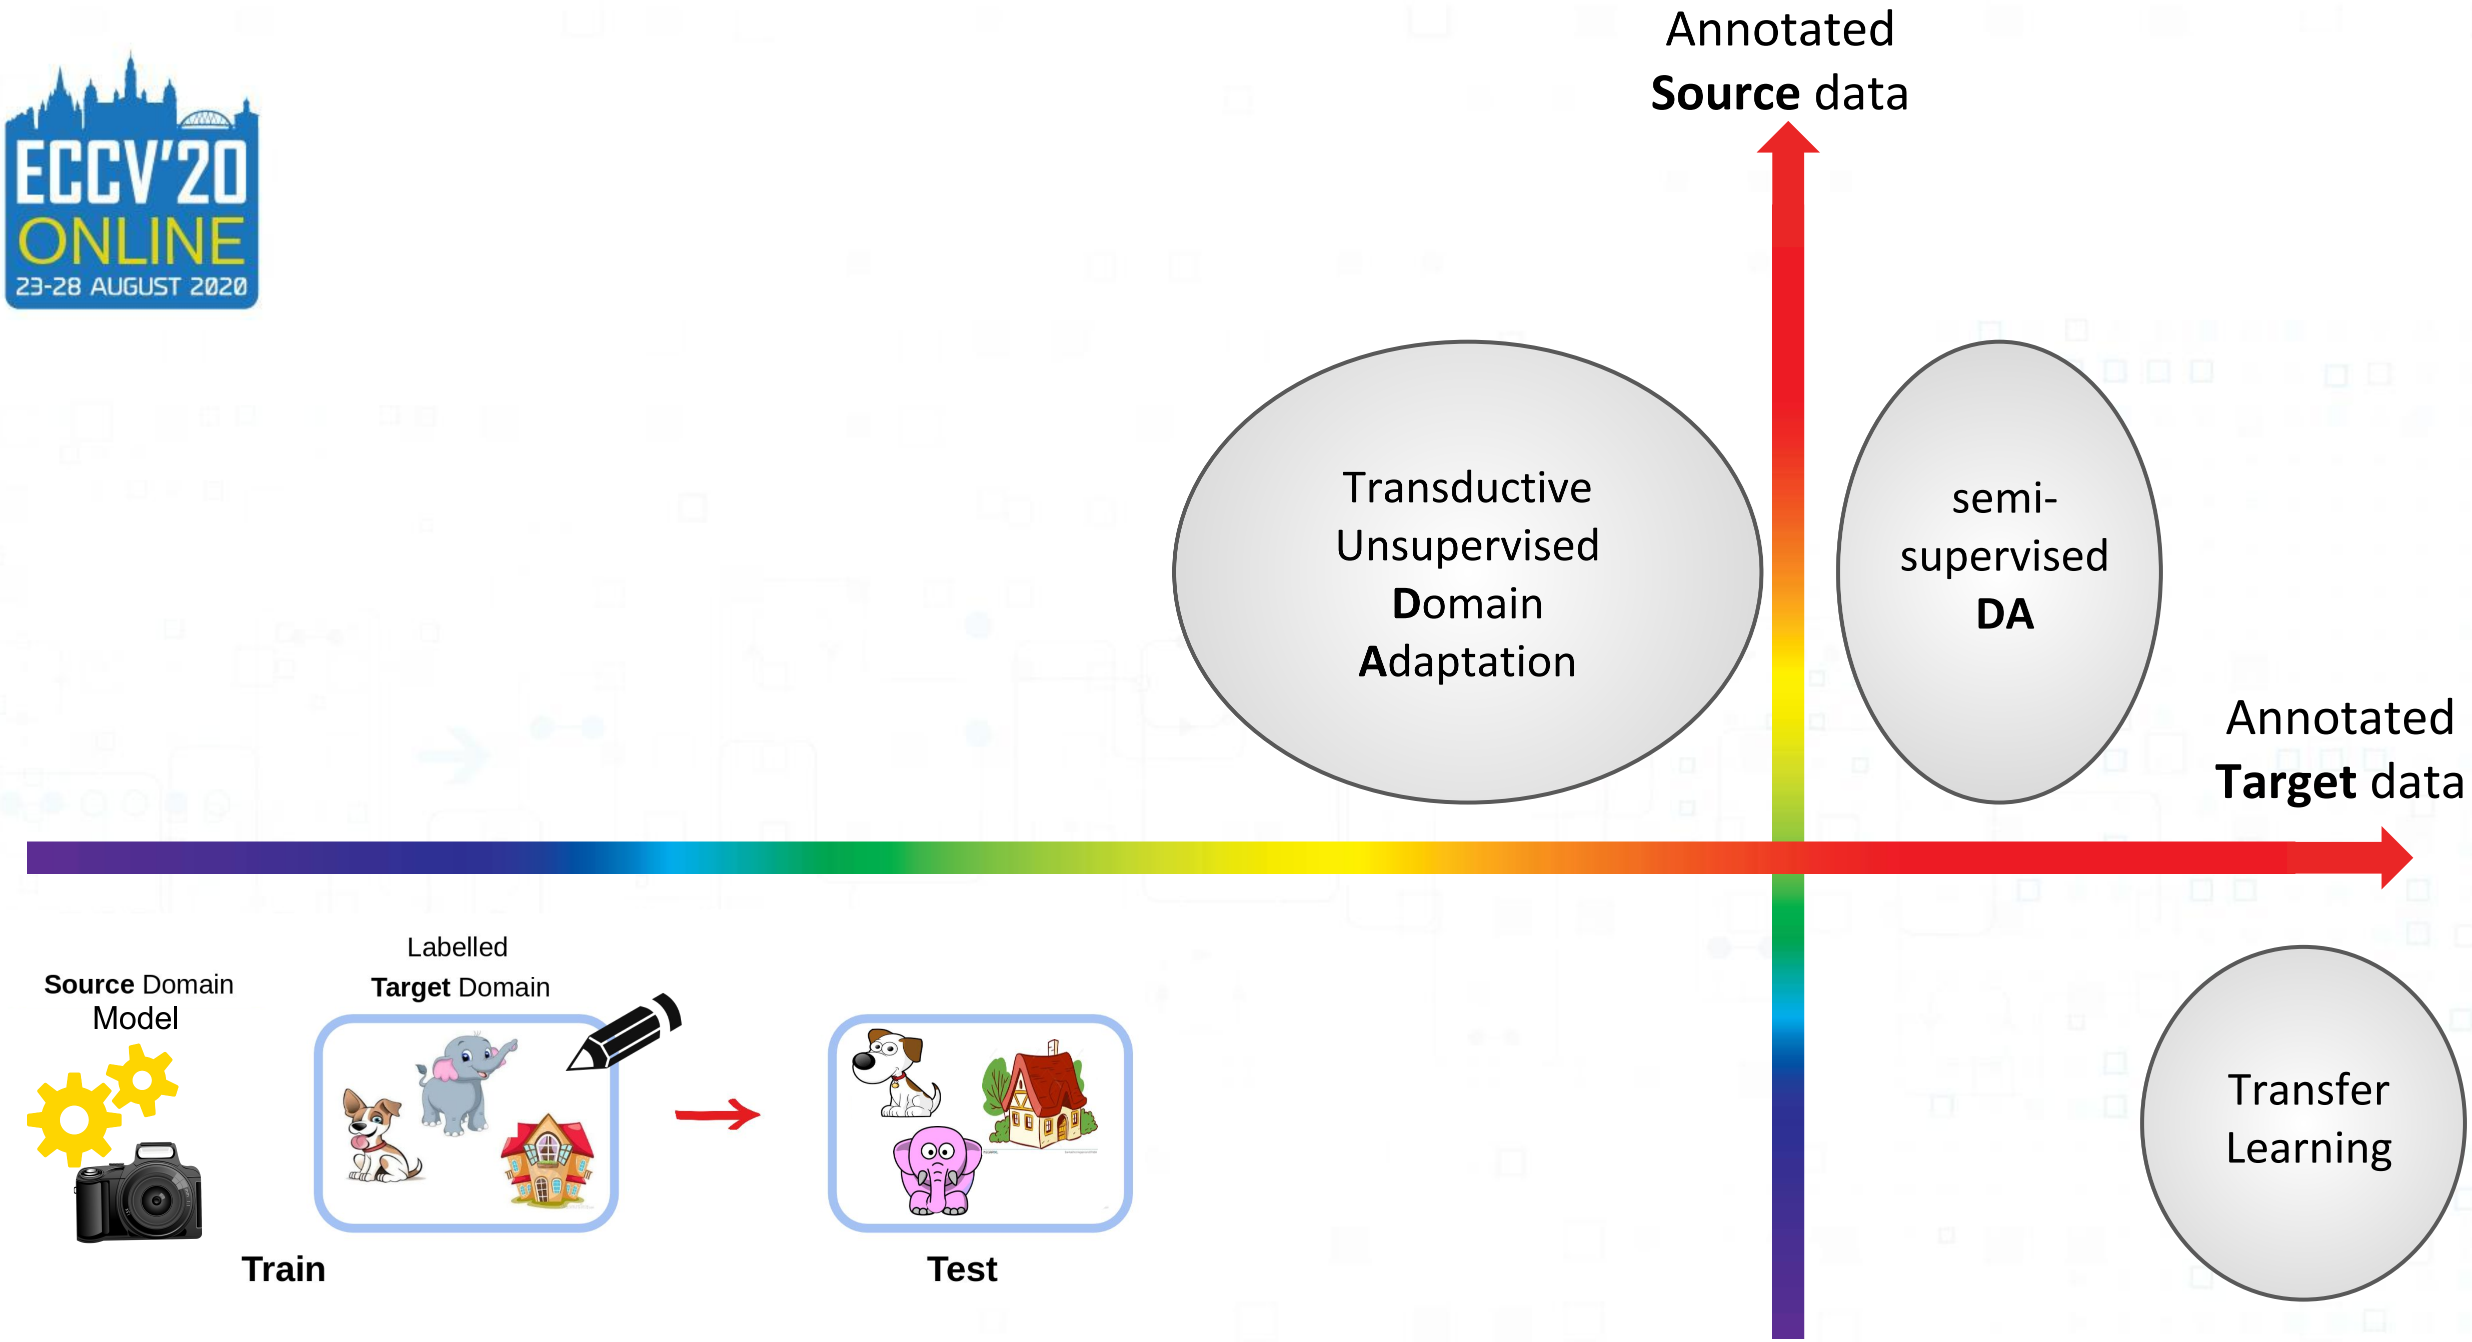
Annotated
Source data
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
Annotated
Target data
Source Domain
Model
Labelled
Target Domain
Train
Test
Transfer
Learning

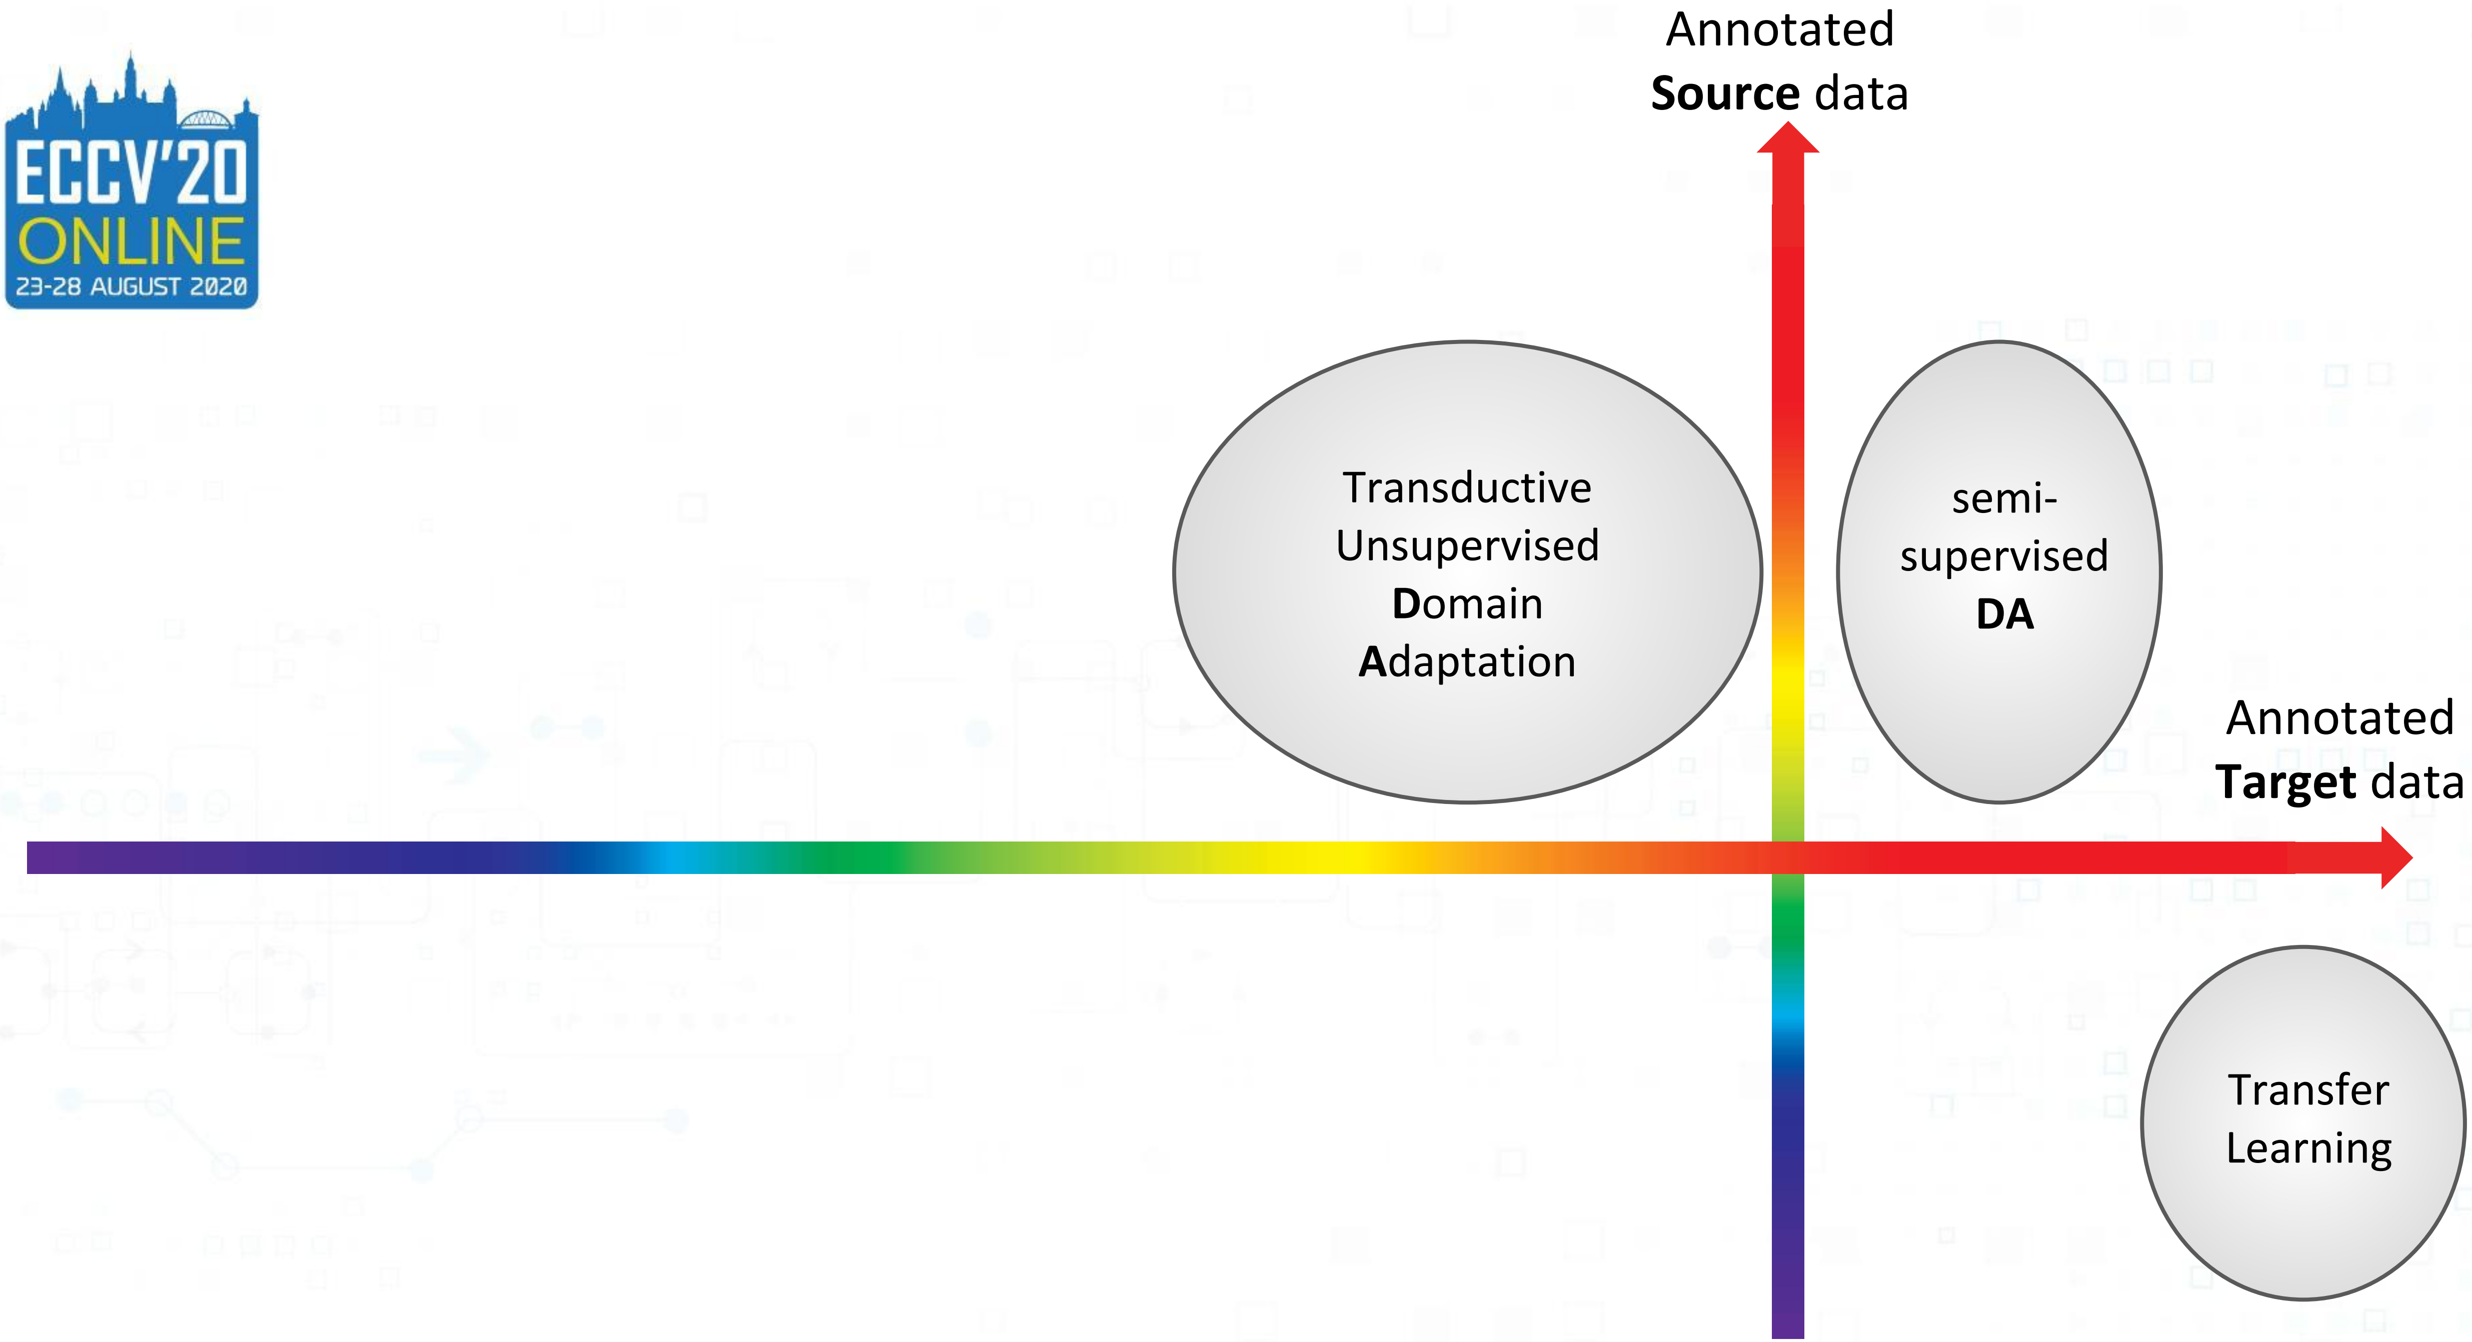
Annotated
Source data
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
Annotated
Target data
Transfer
Learning

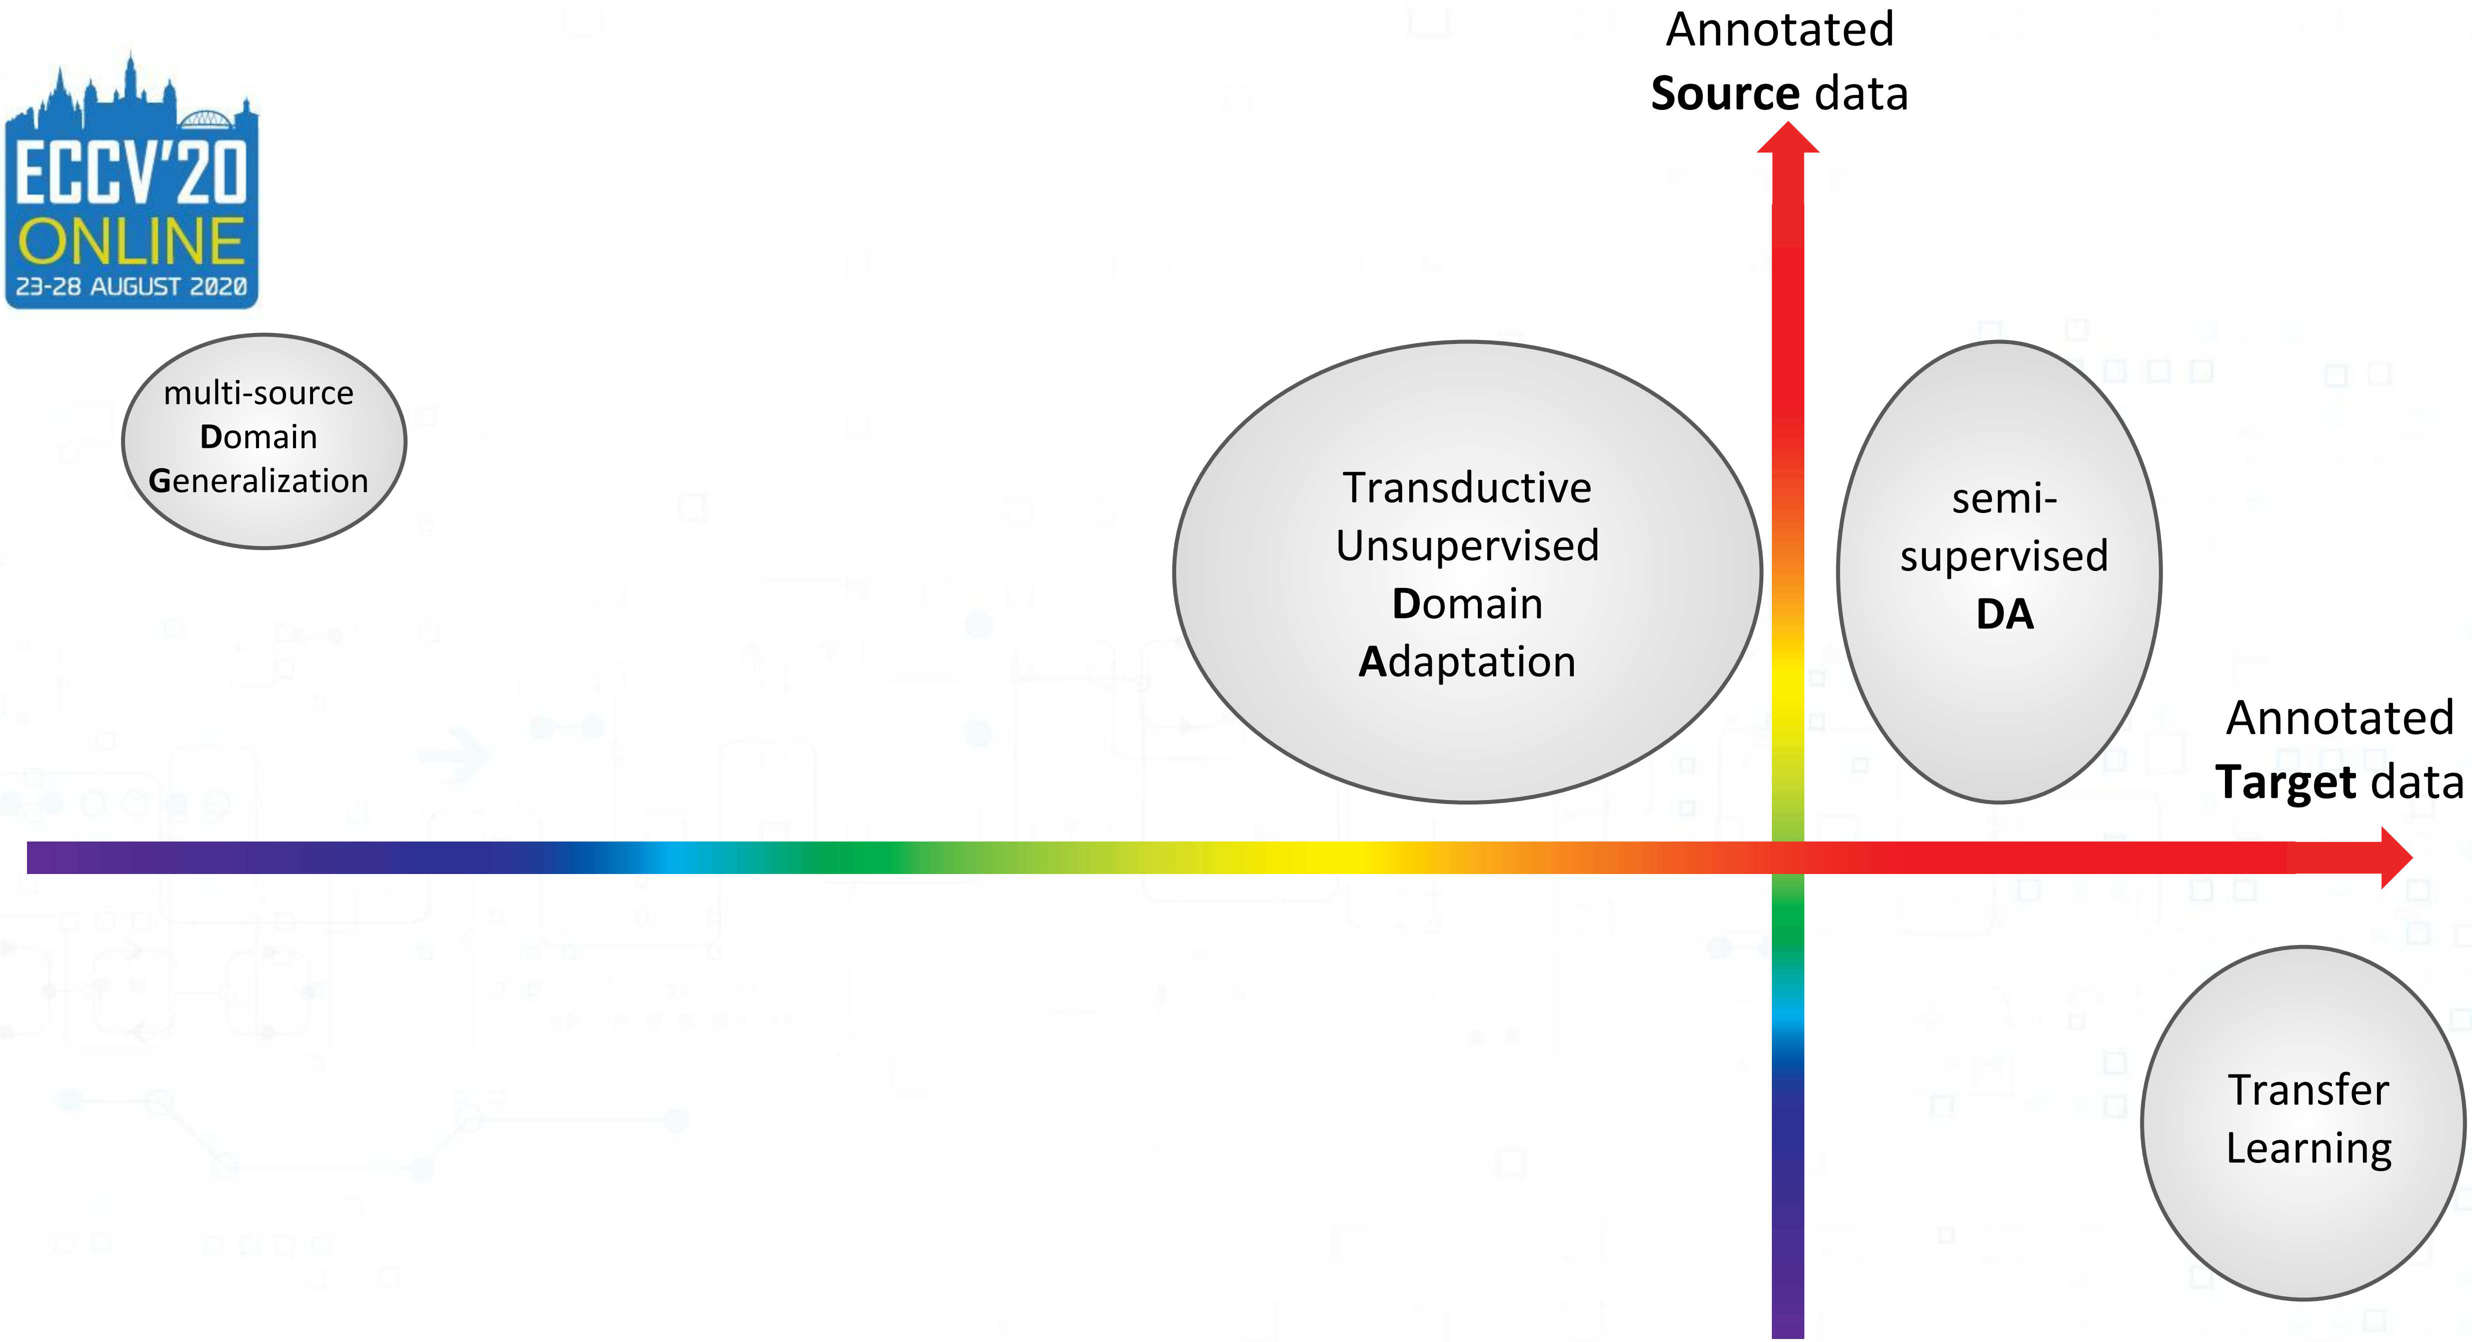
Annotated
Source data
multi-source
Domain
Generalization
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
Annotated
Target data
Transfer
Learning

ECCV'20 ONLINE 23-28 AUGUST 2020

multi-source **D**omain **G**eneralization

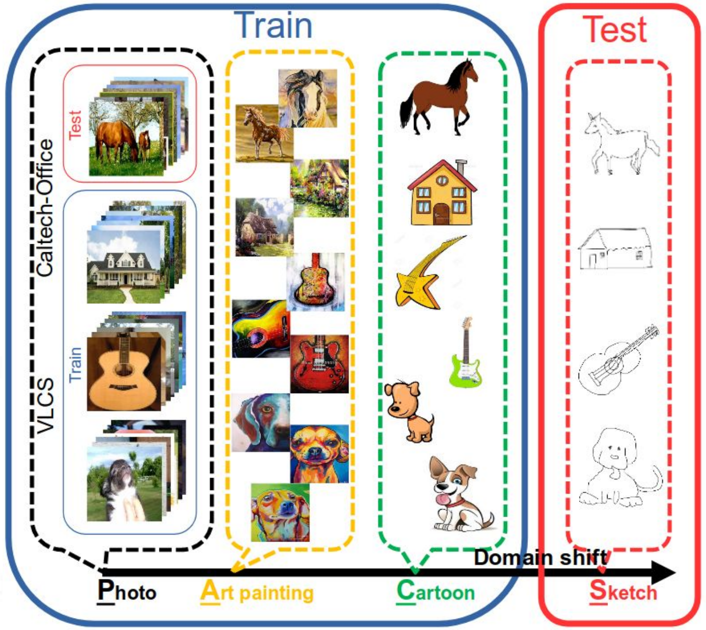

[Deeper, Broader and Artier Domain Generalization, ICCV 2017]

Annotated **Source** data

Transductive Unsupervised **D**omain **A**daptation

semi-supervised **DA**

Annotated **Target** data

Transfer Learning

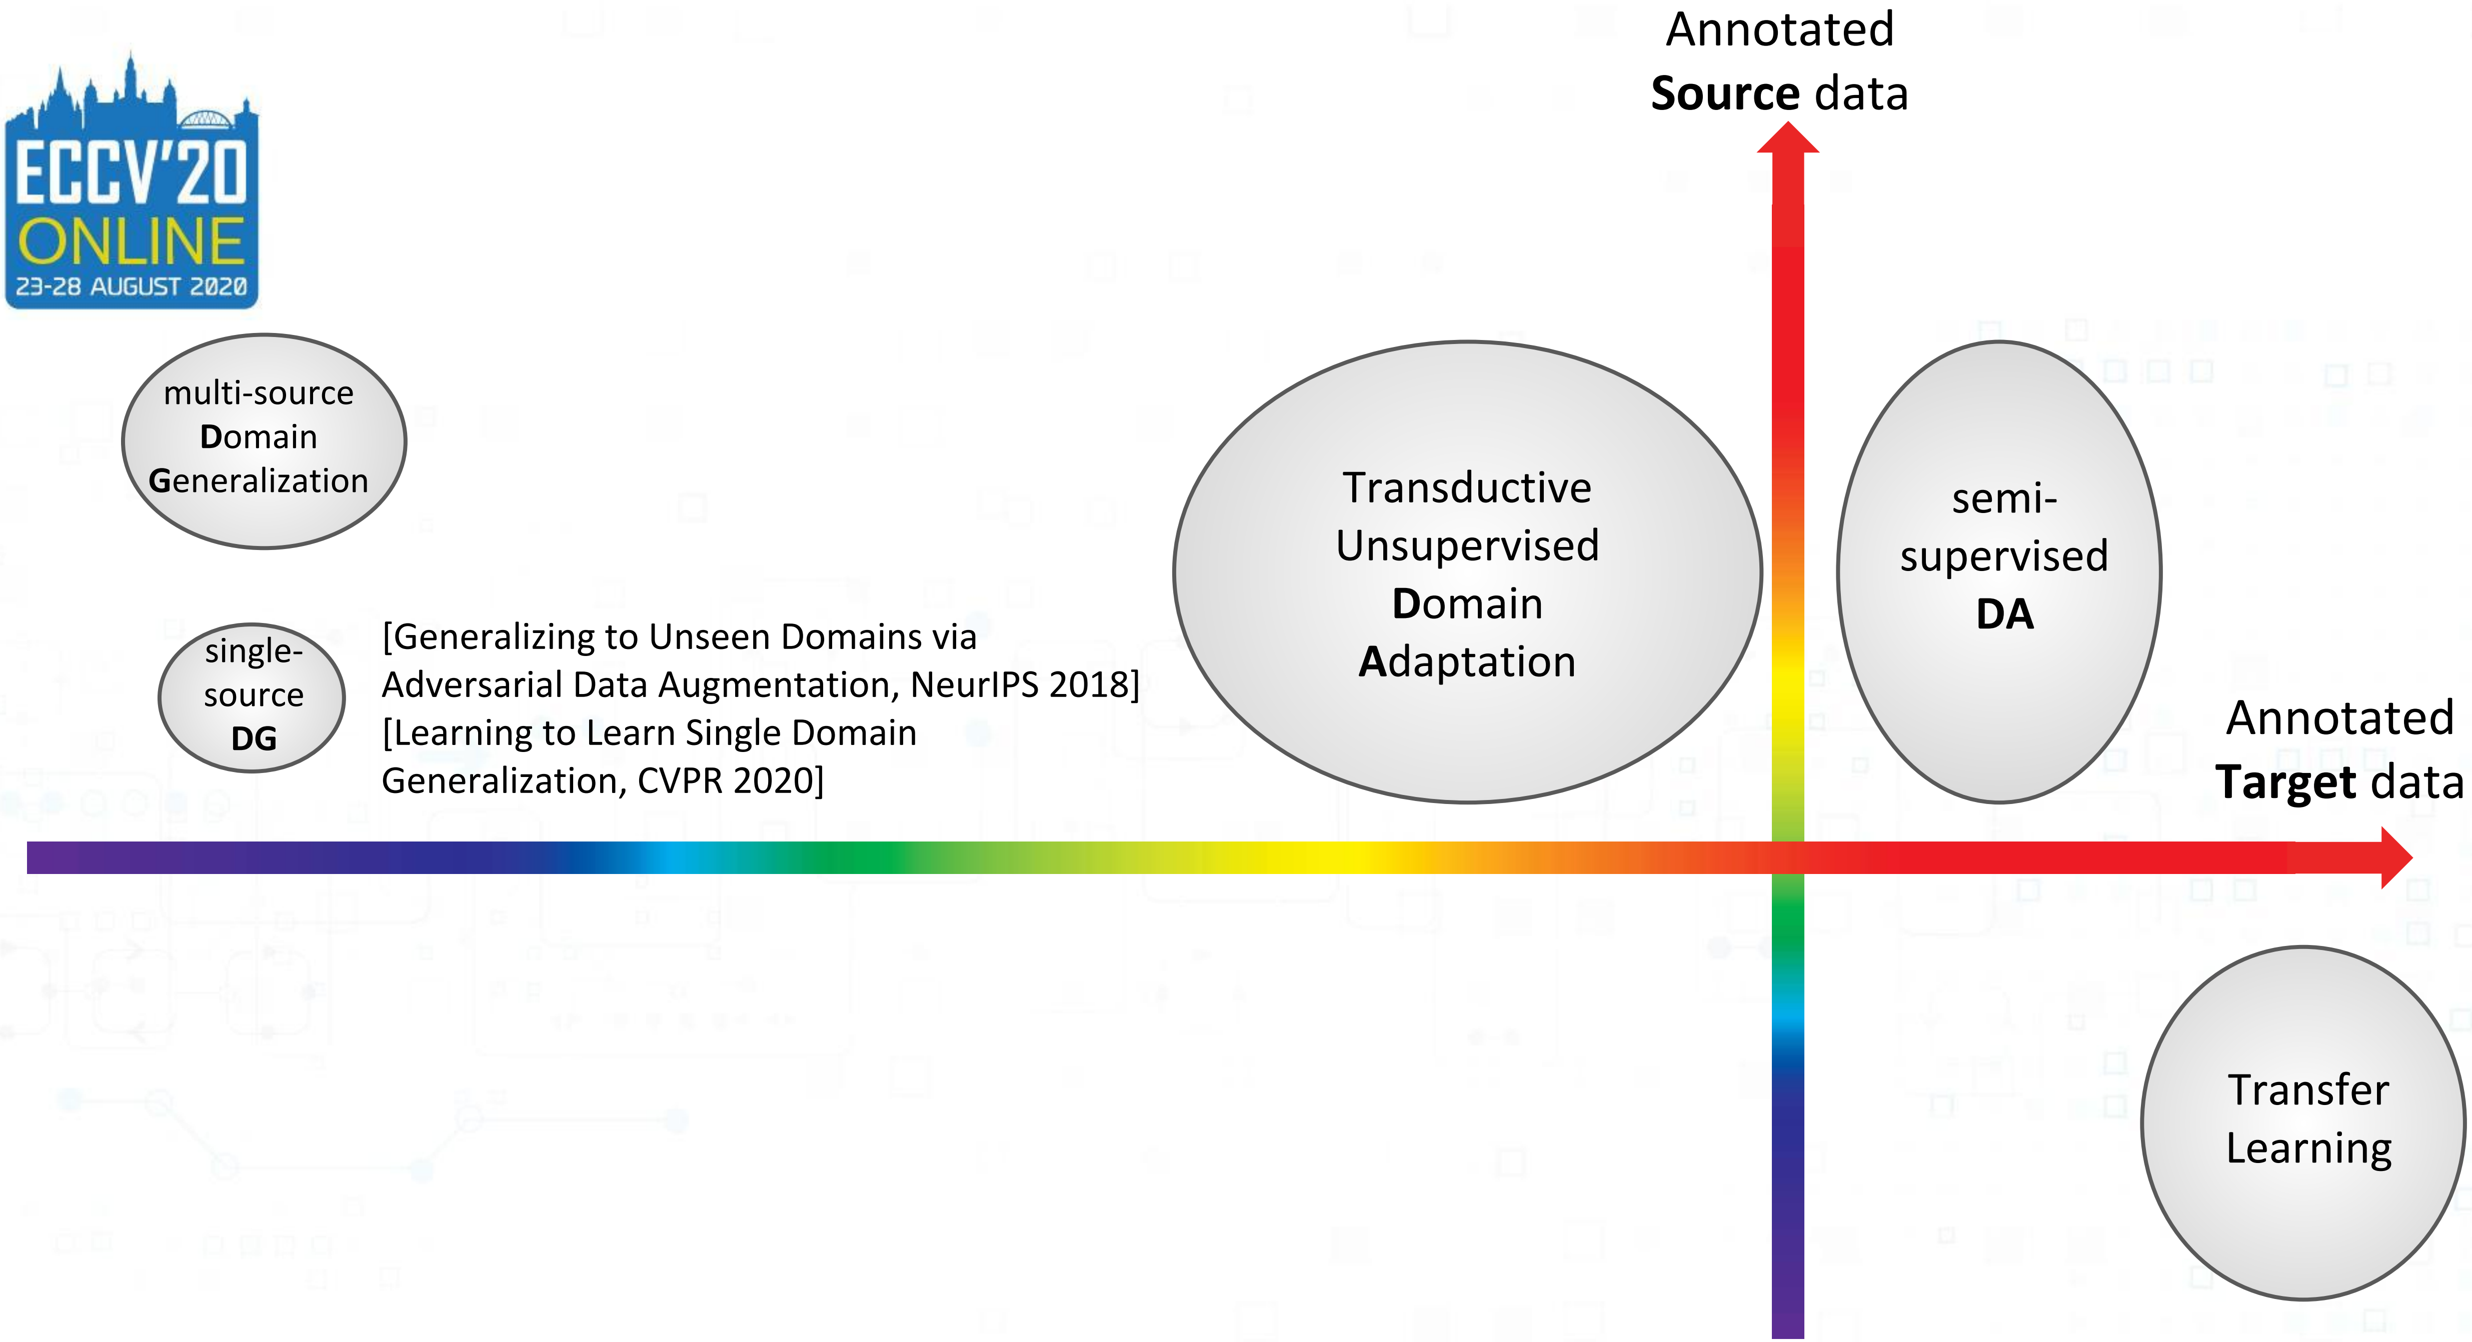
Annotated
Source data
multi-source
Domain
Generalization
single-
source
DG
[Generalizing to Unseen Domains via Adversarial Data Augmentation, NeurIPS 2018]
[Learning to Learn Single Domain Generalization, CVPR 2020]
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
Annotated
Target data
Transfer
Learning

ECCV'20 ONLINE
23-28 AUGUST 2020

Annotated **Source** data

multi-source **D**omain **G**eneralization

single-source **DG**

[Generalizing to Unseen Domains via Adversarial Data Augmentation, NeurIPS 2018]
[Learning to Learn Single Domain Generalization, CVPR 2020]

Transductive Unsupervised **D**omain **A**daptation

semi-supervised **DA**

Annotated **Target** data

Transfer Learning

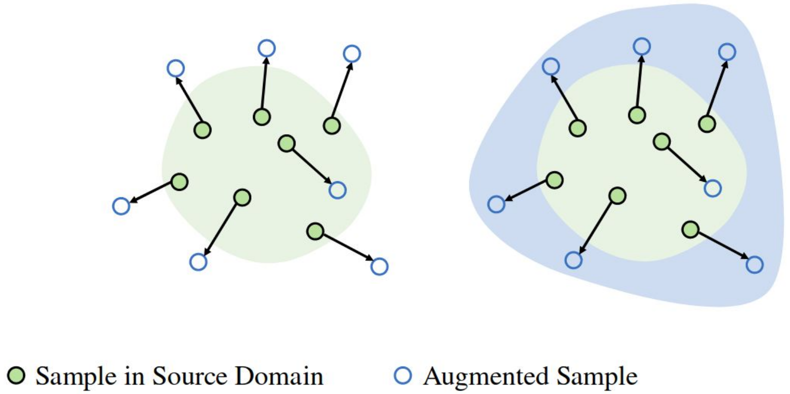

Single Source Domain Generalization

[Learning to Learn Single Domain Generalization, CVPR 2020]

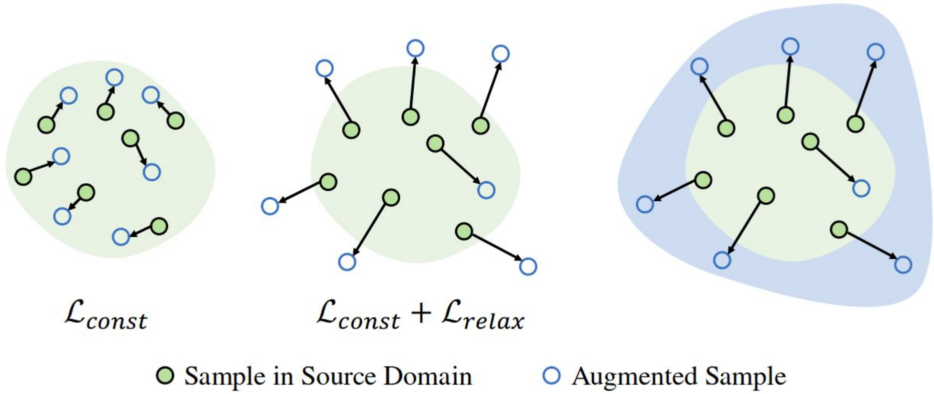

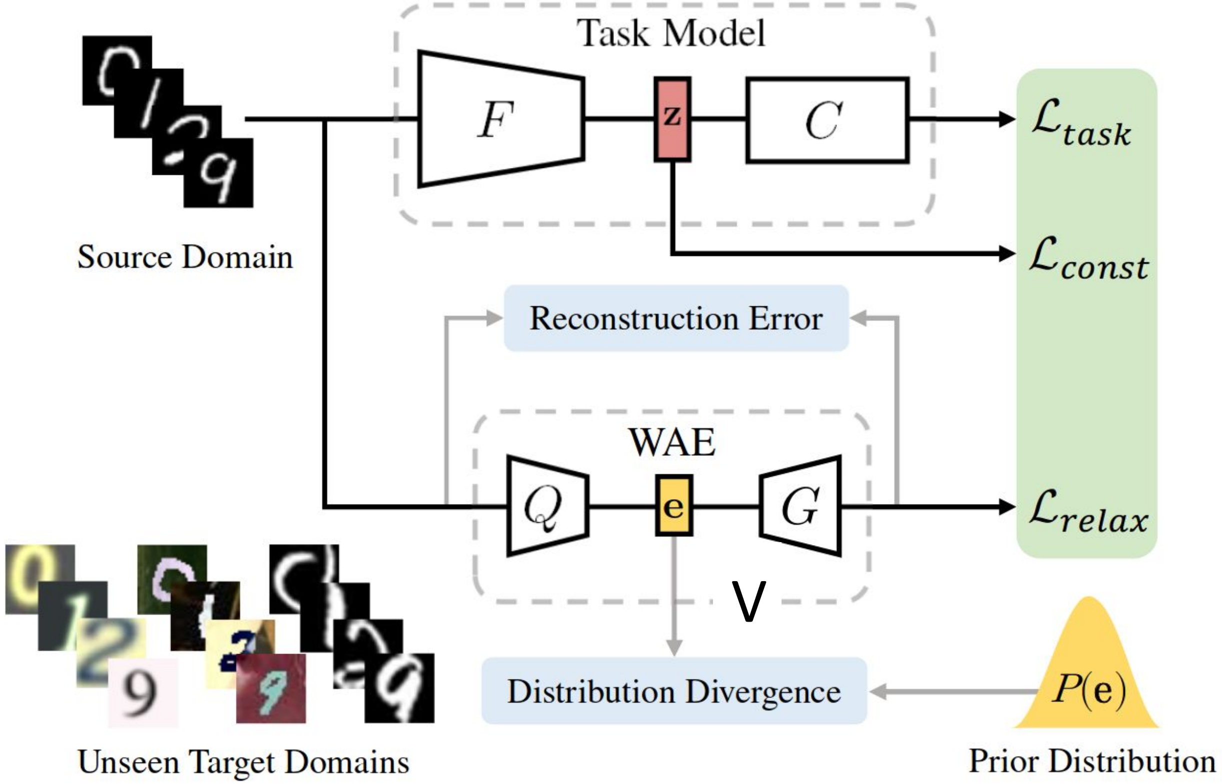

$$\mathcal{L}_{\text{const}} = \frac{1}{2}\|\mathbf{z} - \mathbf{z}^{+}\|_2^2 + \infty \cdot \mathbf{1}\{\mathbf{y} \neq \mathbf{y}^{+}\}$$

Single Source Domain Generalization

[Learning to Learn Single Domain Generalization, CVPR 2020]

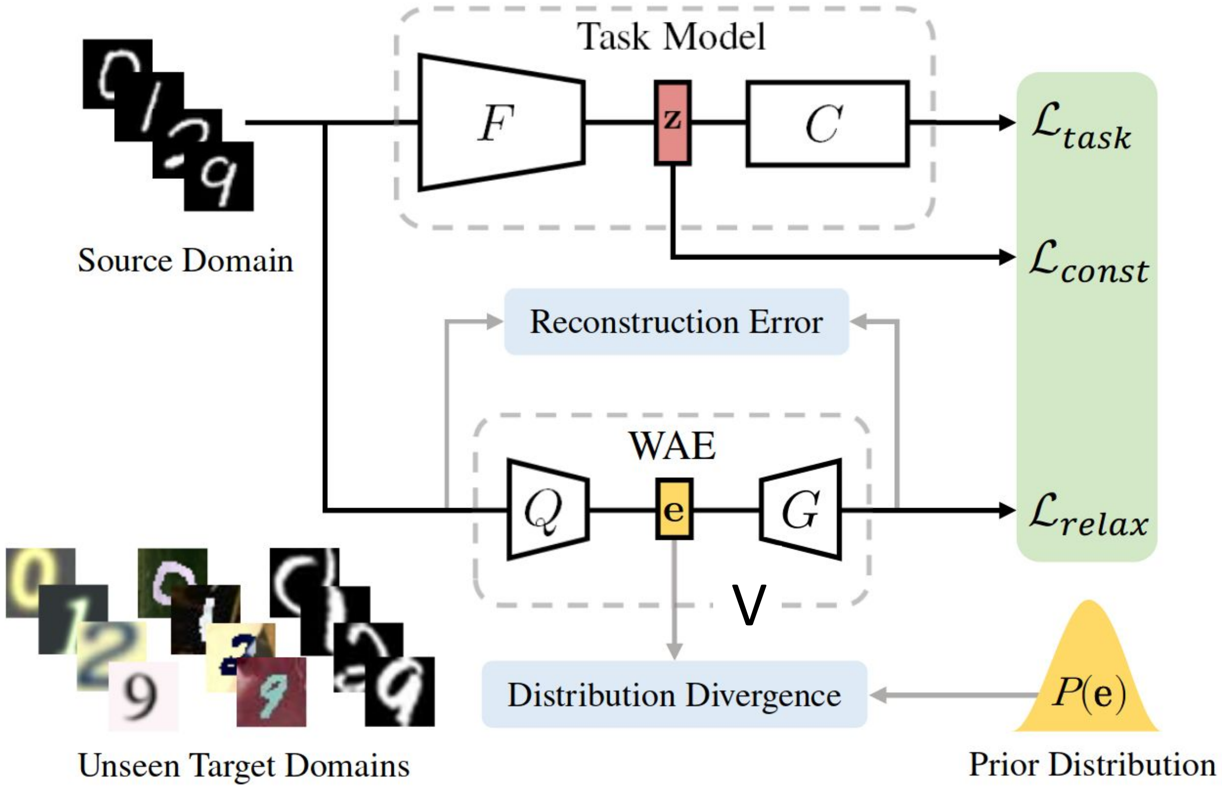

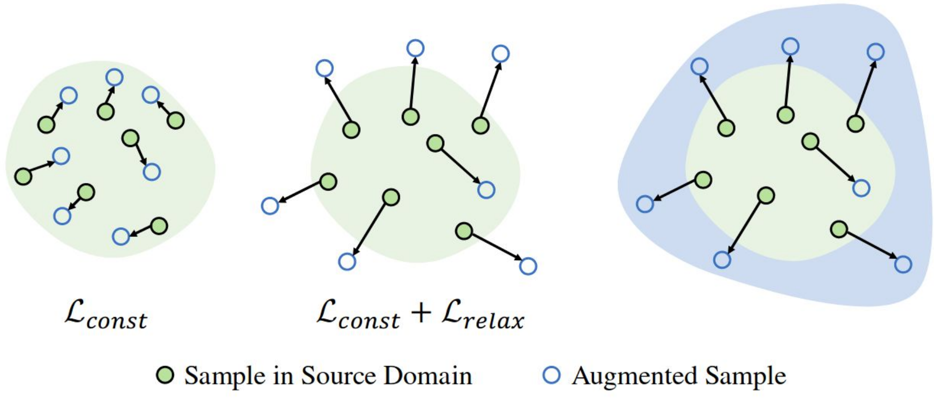

$$\mathcal{L}_{\text{const}} = \frac{1}{2}\|\mathbf{z} - \mathbf{z}^{+}\|_2^2 + \infty \cdot \mathbf{1}\left\{\mathbf{y} \neq \mathbf{y}^{+}\right\}$$

$$\mathcal{L}_{\text{relax}} = \|\mathbf{x}^{+} - V(\mathbf{x}^{+})\|^2$$

Single Source Domain Generalization

[Learning to Learn Single Domain Generalization, CVPR 2020]

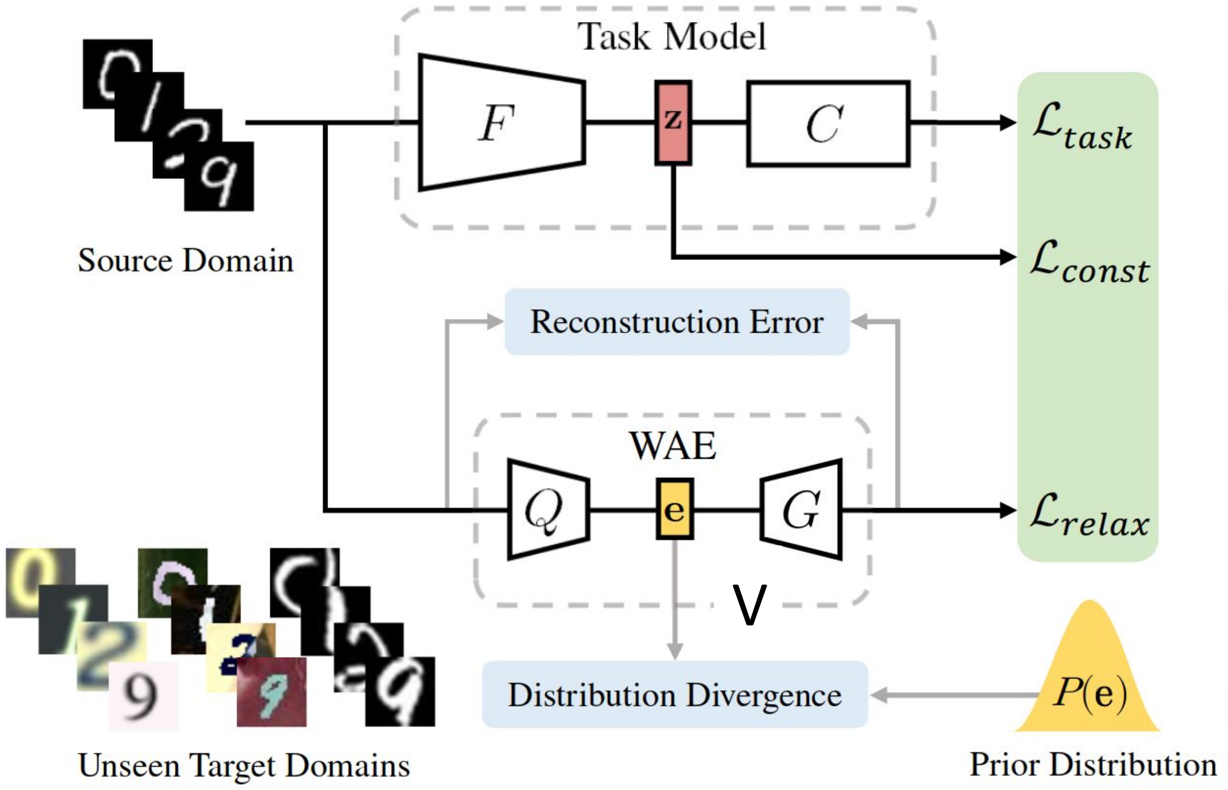

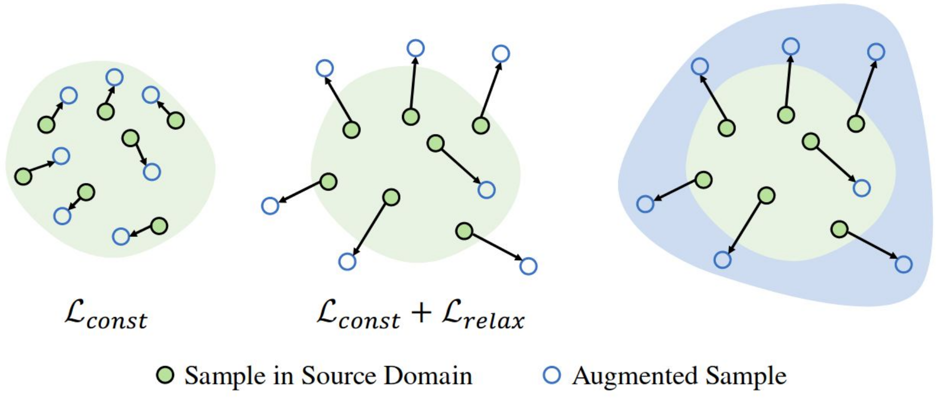

$$\mathcal{L}_{\text{const}} = \frac{1}{2}\|\mathbf{z} - \mathbf{z}^+\|_2^2 + \infty \cdot \mathbf{1}\{\mathbf{y} \neq \mathbf{y}^+\}$$

$$\mathcal{L}_{\text{relax}} = \|\mathbf{x}^+ - V(\mathbf{x}^+)\|^2$$

$$\mathcal{L}_{\text{ADA}} = \underbrace{\mathcal{L}_{\text{task}}(\theta; \mathbf{x})}_{\text{Classification}} - \alpha \underbrace{\mathcal{L}_{\text{const}}(\theta; \mathbf{z})}_{\text{Constraint}} + \beta \underbrace{\mathcal{L}_{\text{relax}}(\psi; \mathbf{x})}_{\text{Relaxation}},$$

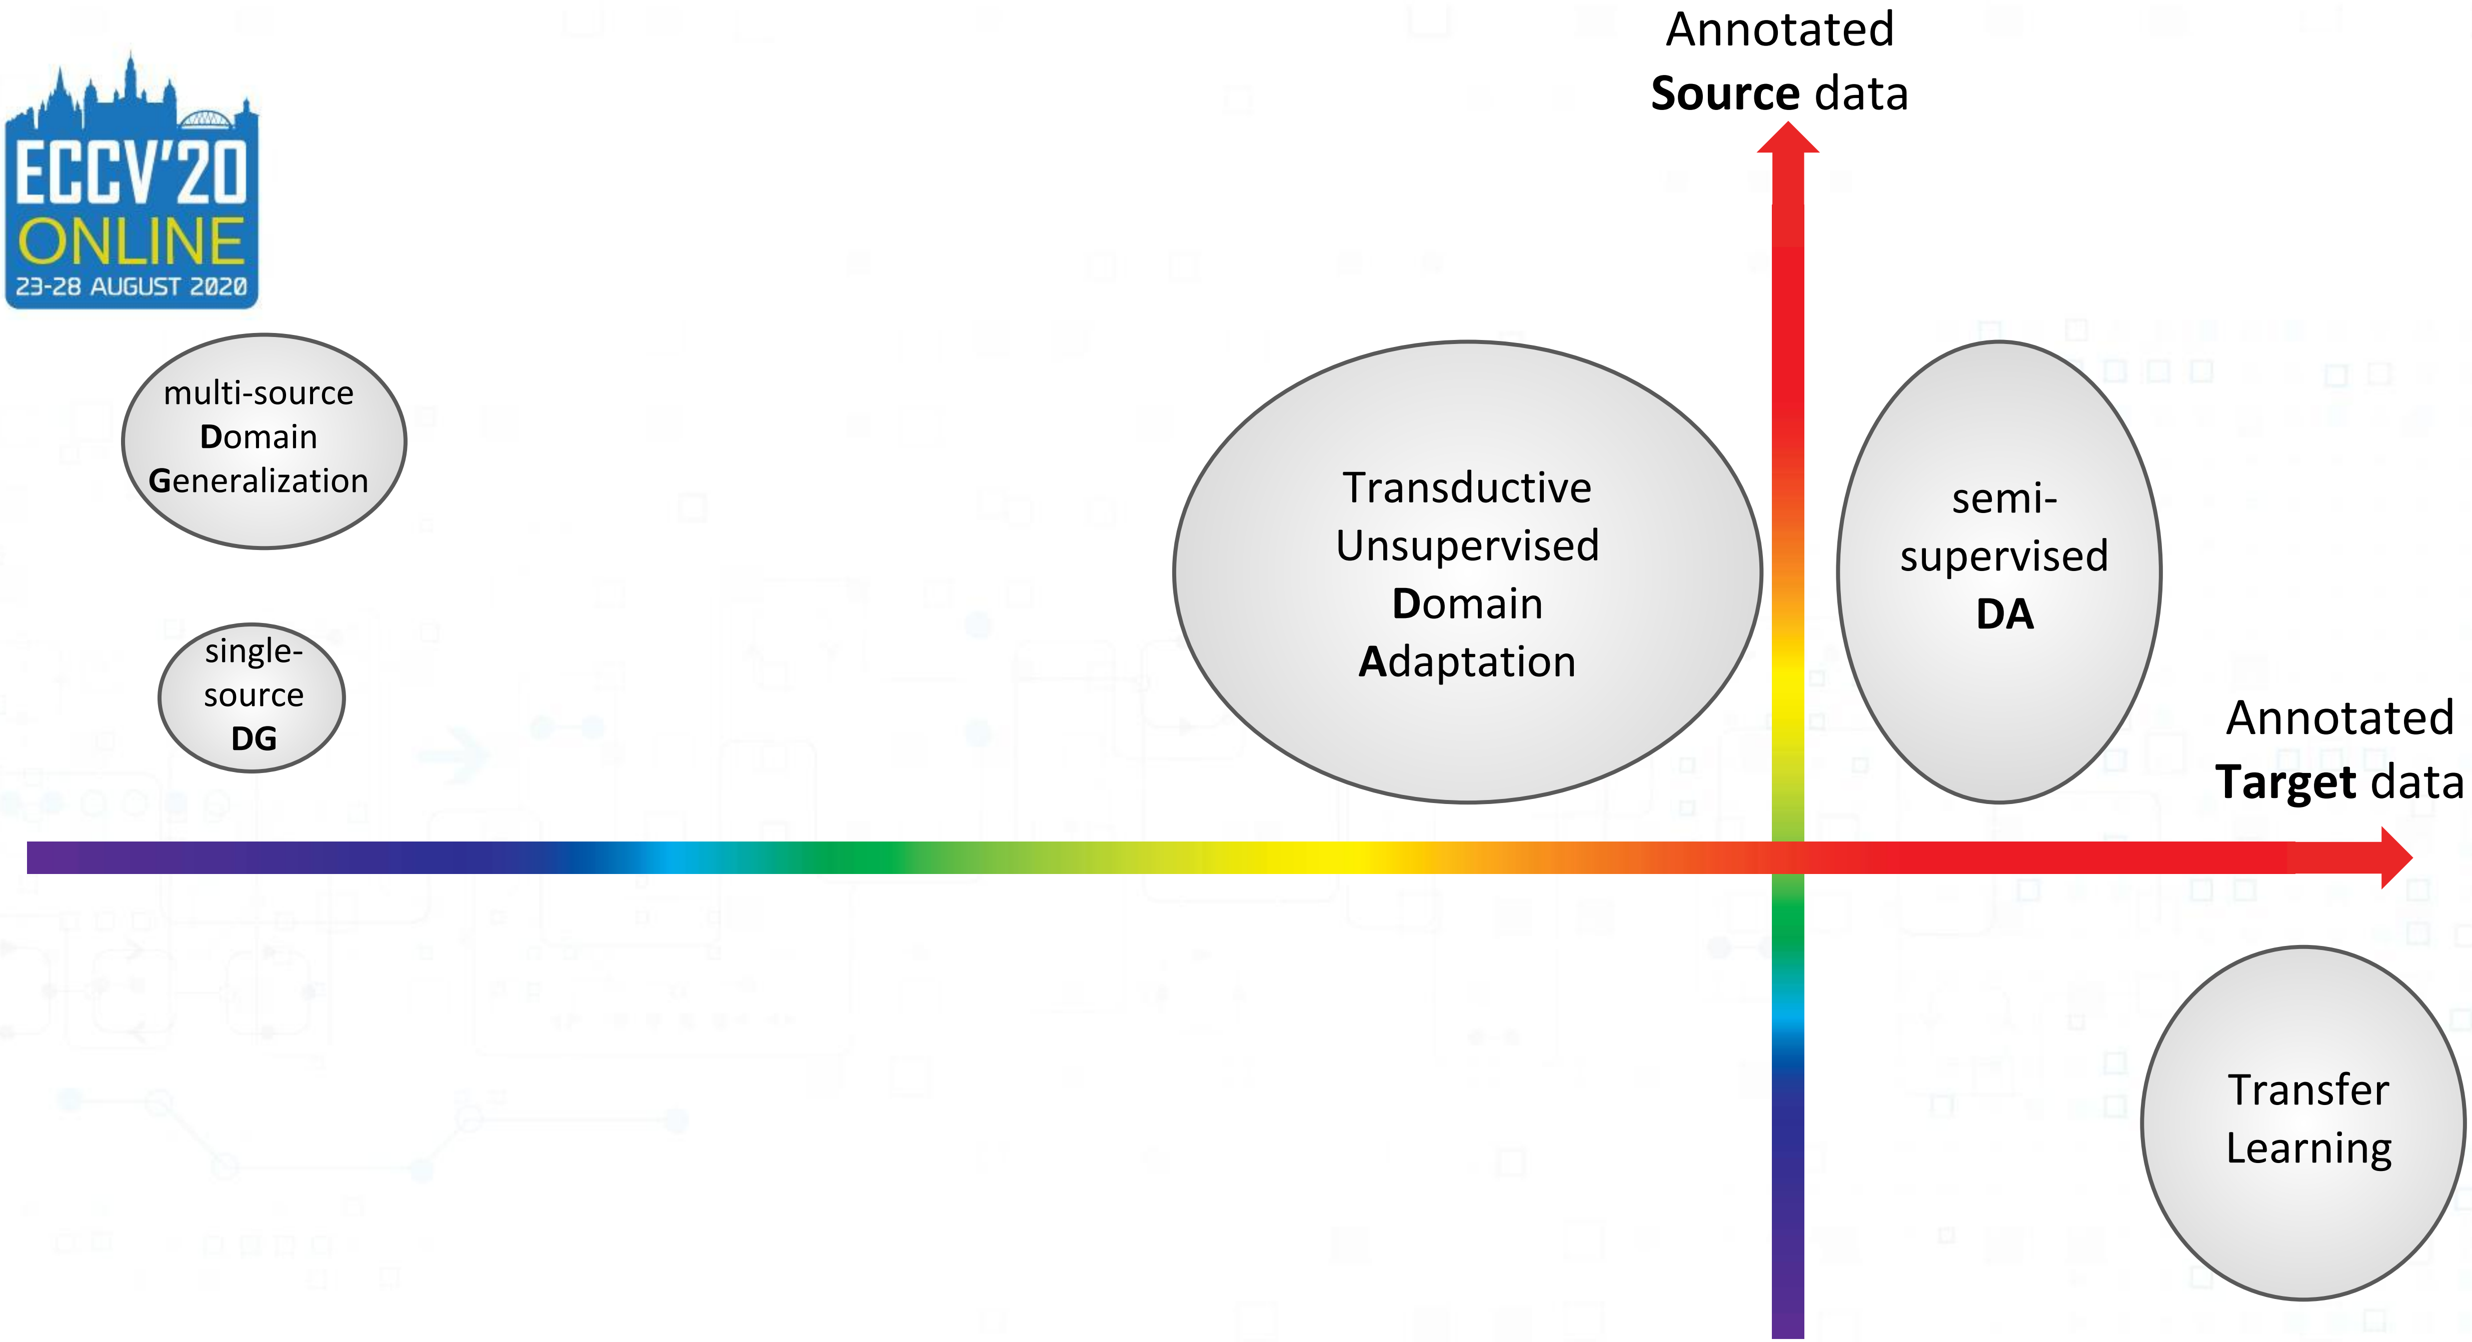
Annotated
Source data
multi-source
Domain
Generalization
single-
source
DG
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
Annotated
Target data
Transfer
Learning

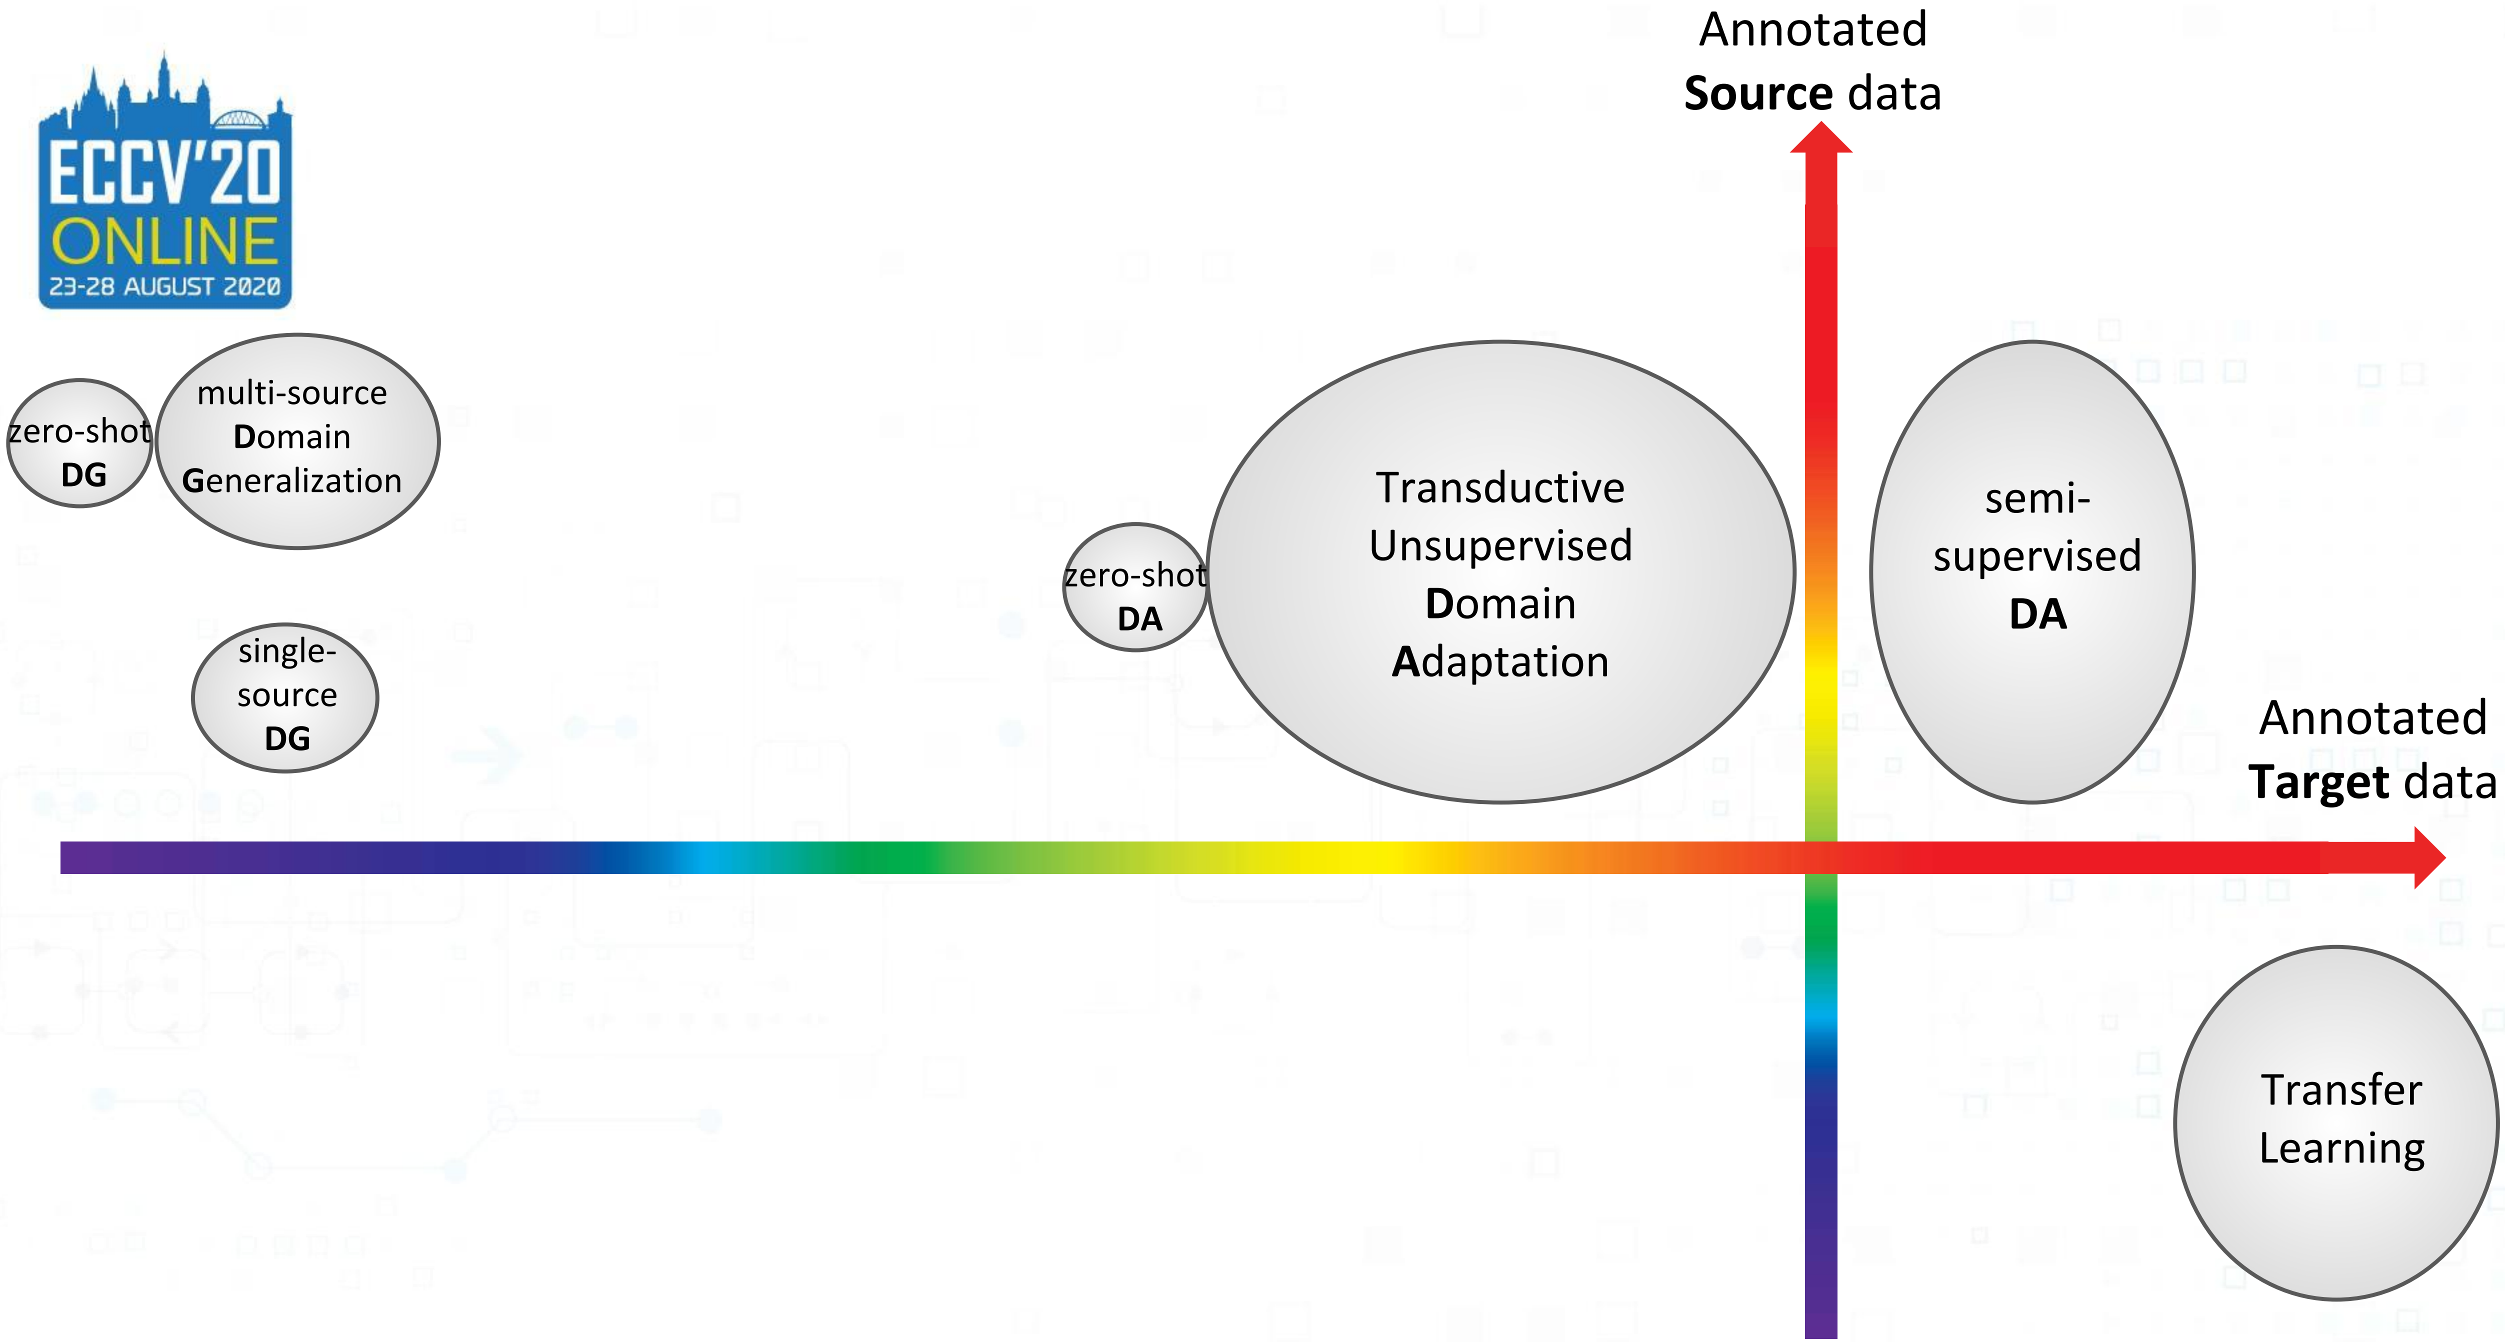

Annotated
Source data
Annotated
Target data
zero-shot
DG
multi-source
Domain
Generalization
single-
source
DG
zero-shot
DA
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
Transfer
Learning

Zero-Shot Domain Adaptation

[Zero-Shot Deep Domain Adaptation, ECCV 2018]

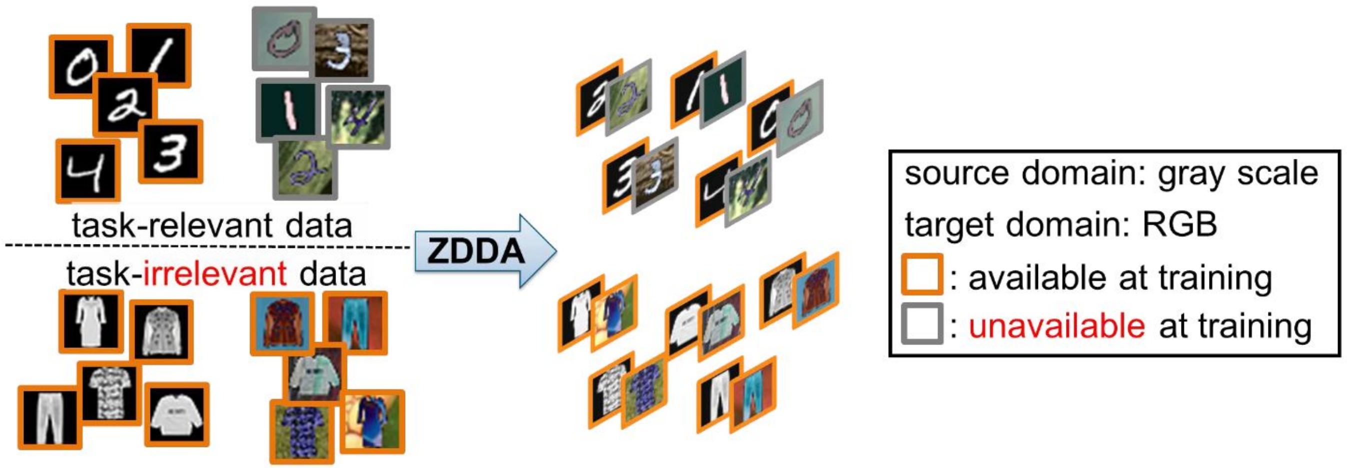

Zero-Shot Domain Adaptation

[Zero-Shot Deep Domain Adaptation, ECCV 2018]

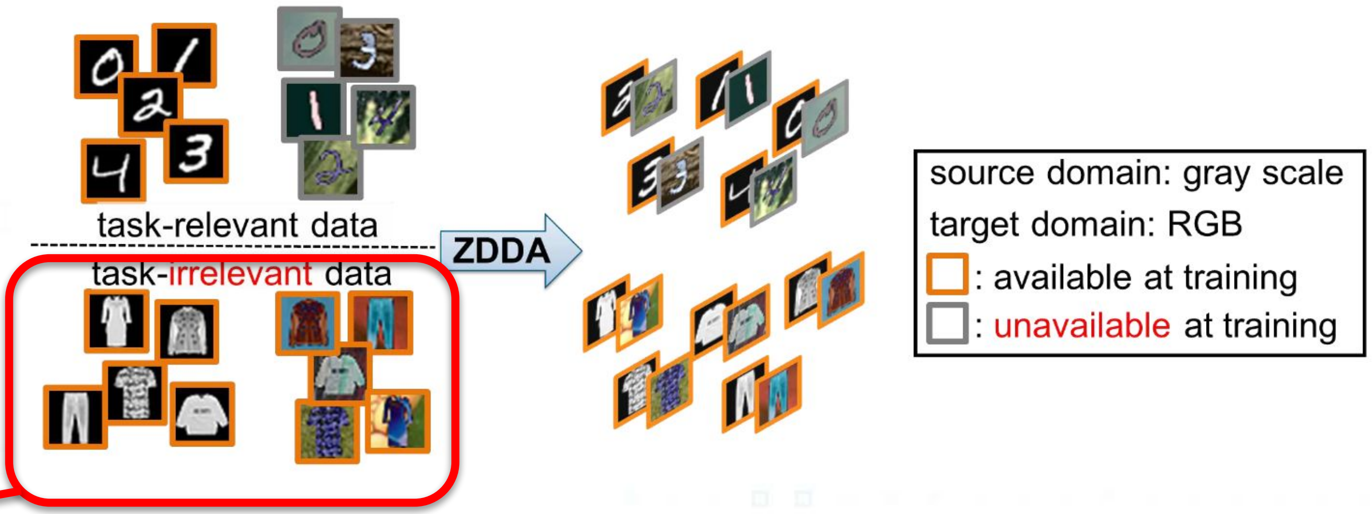

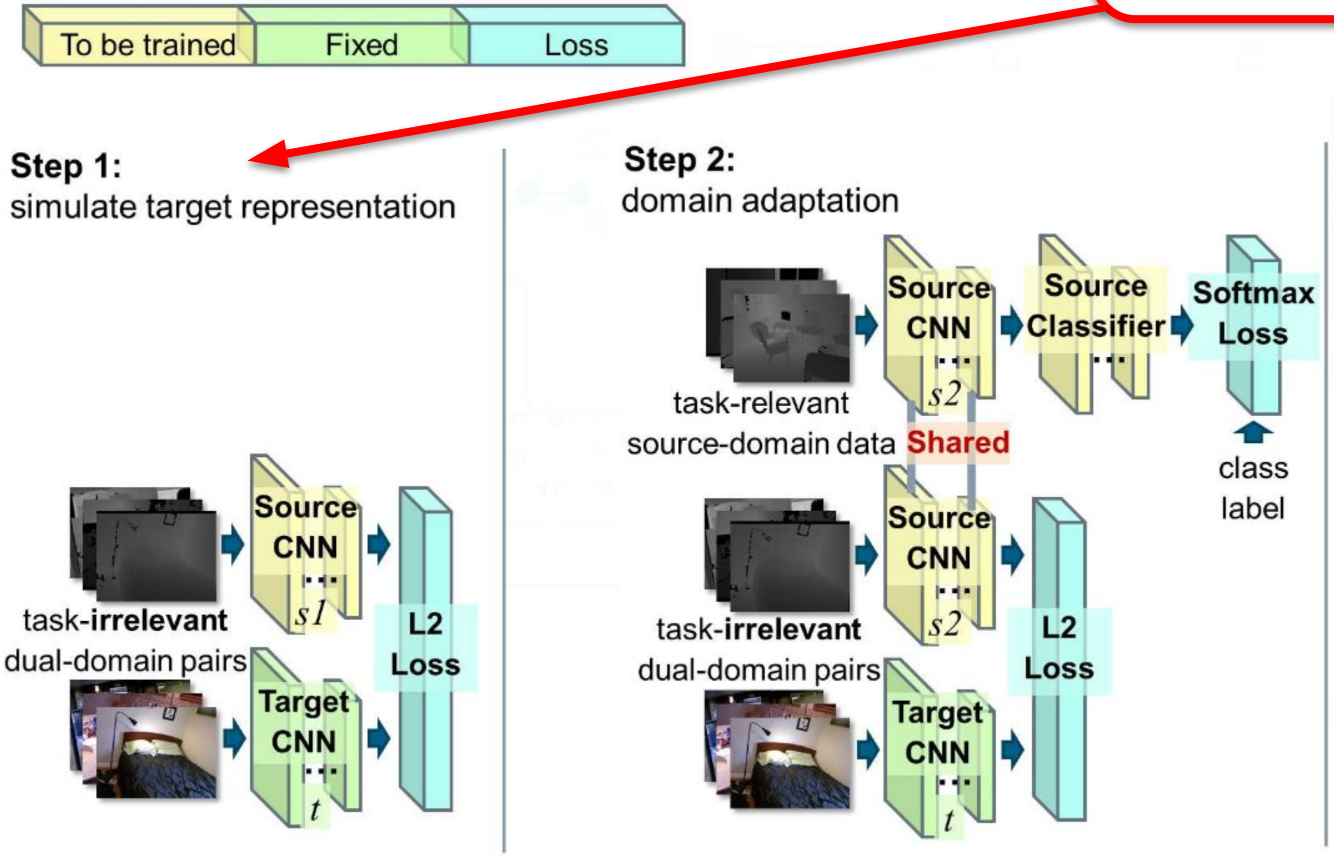

Zero-Shot Domain Adaptation

[Zero-Shot Deep Domain Adaptation, ECCV 2018]

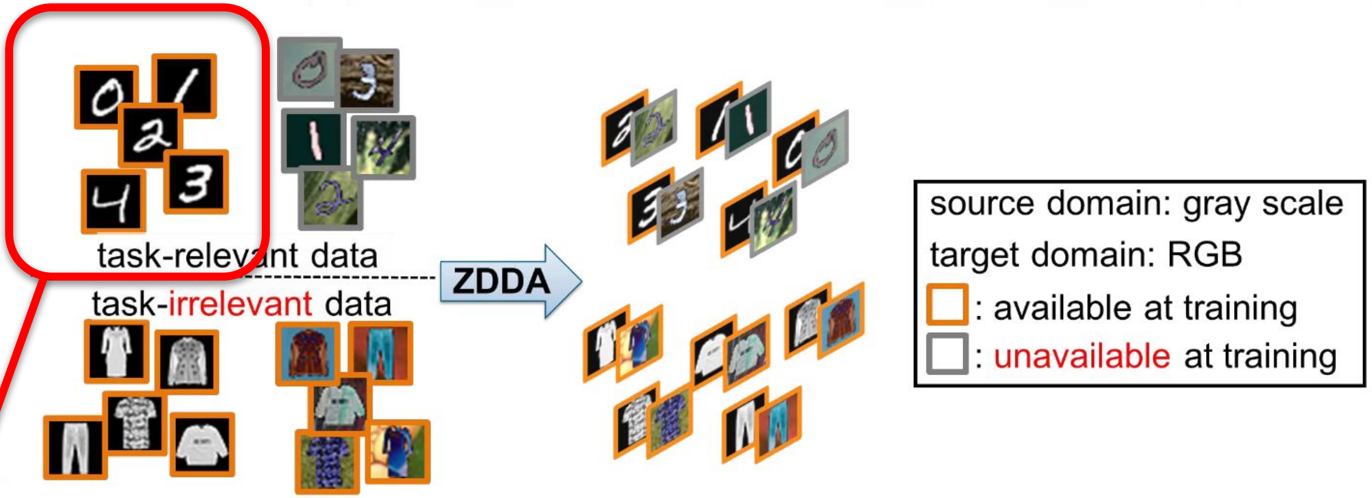

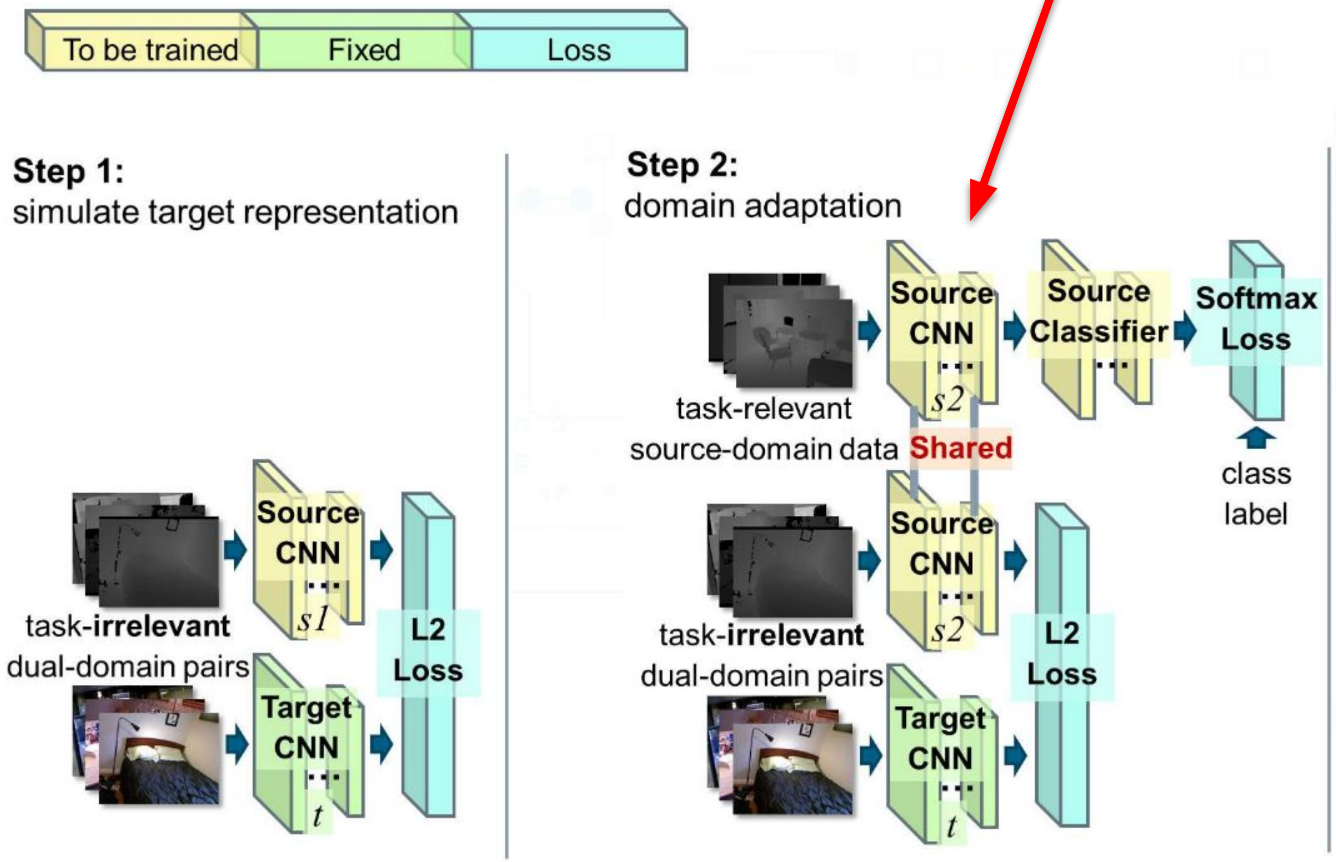

Zero-Shot Domain Adaptation

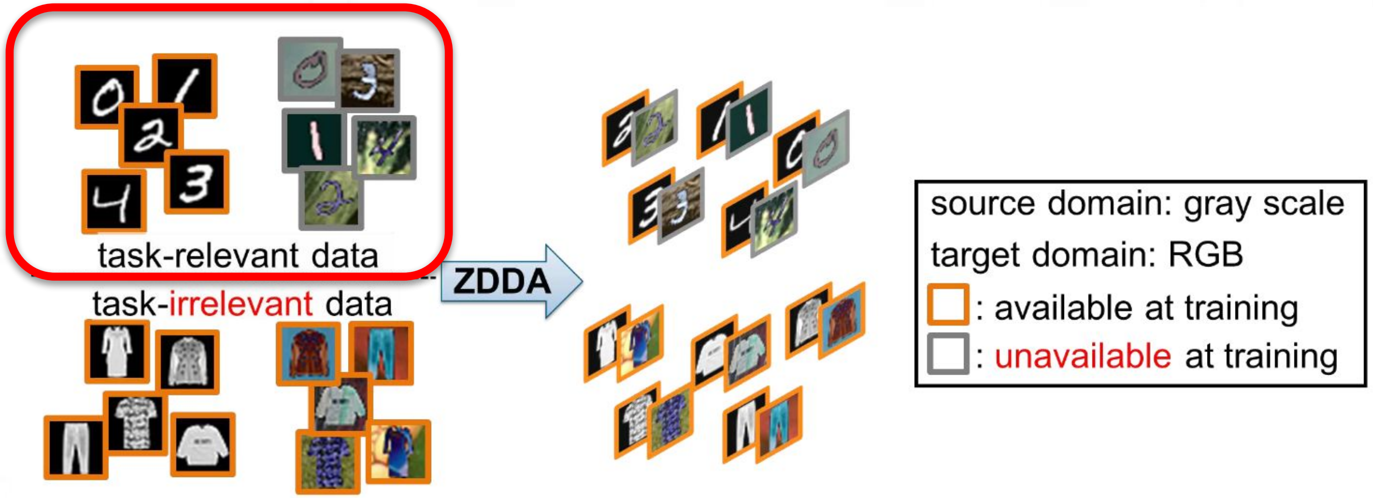

[Zero-Shot Deep Domain Adaptation, ECCV 2018]

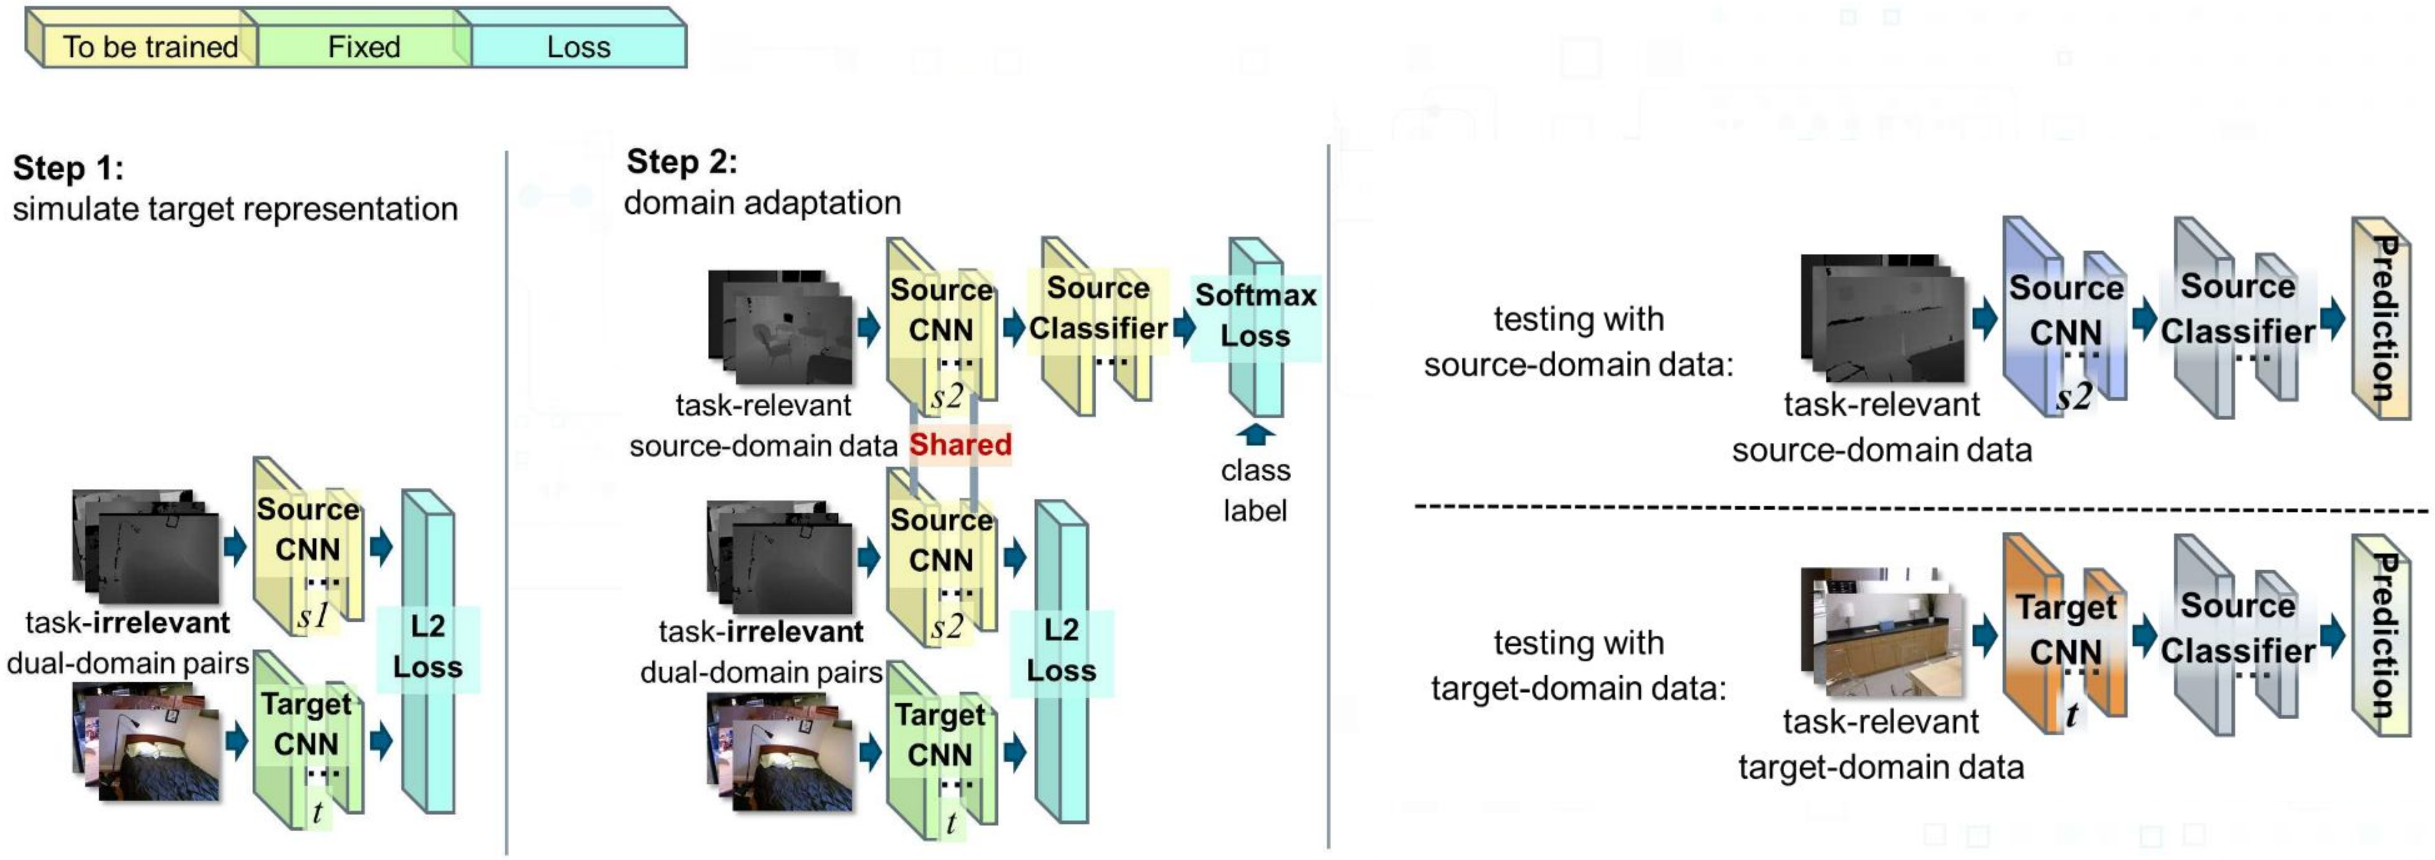

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

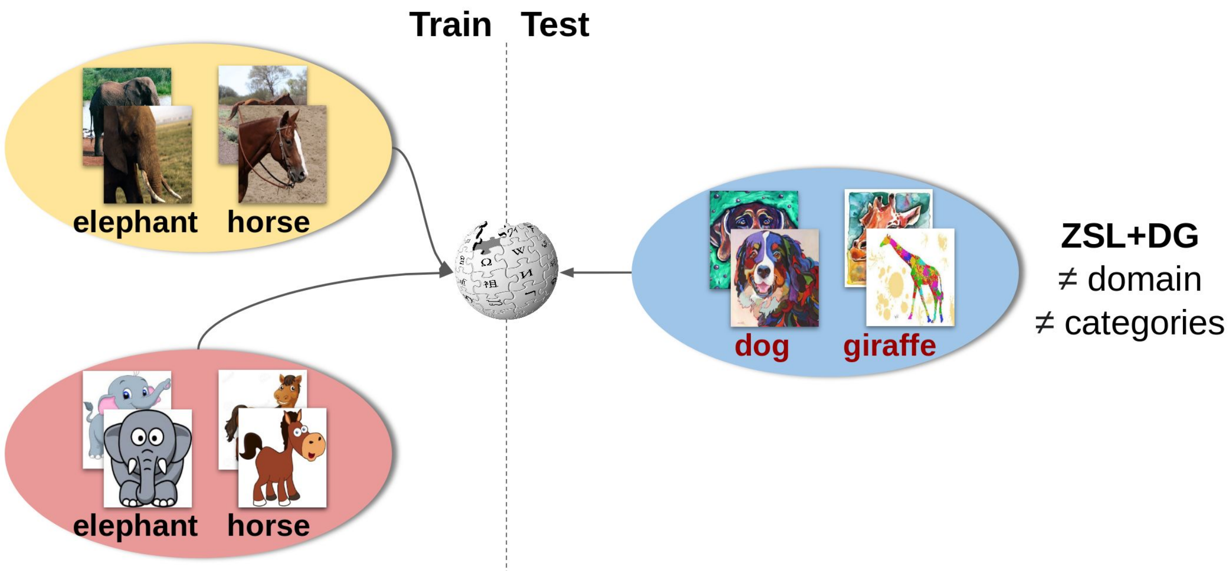

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

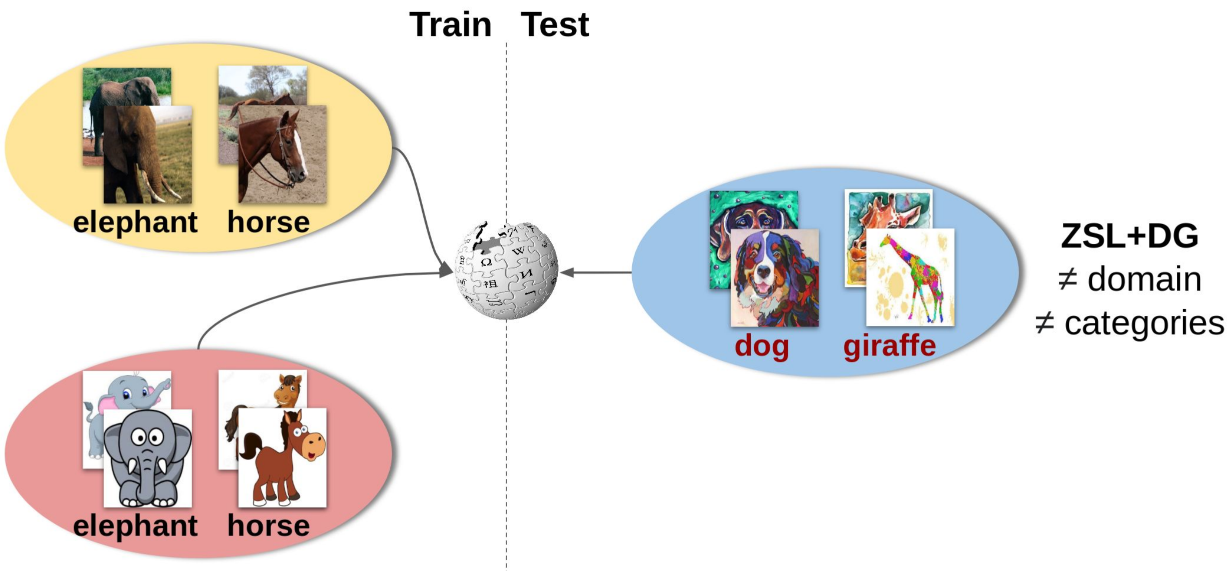

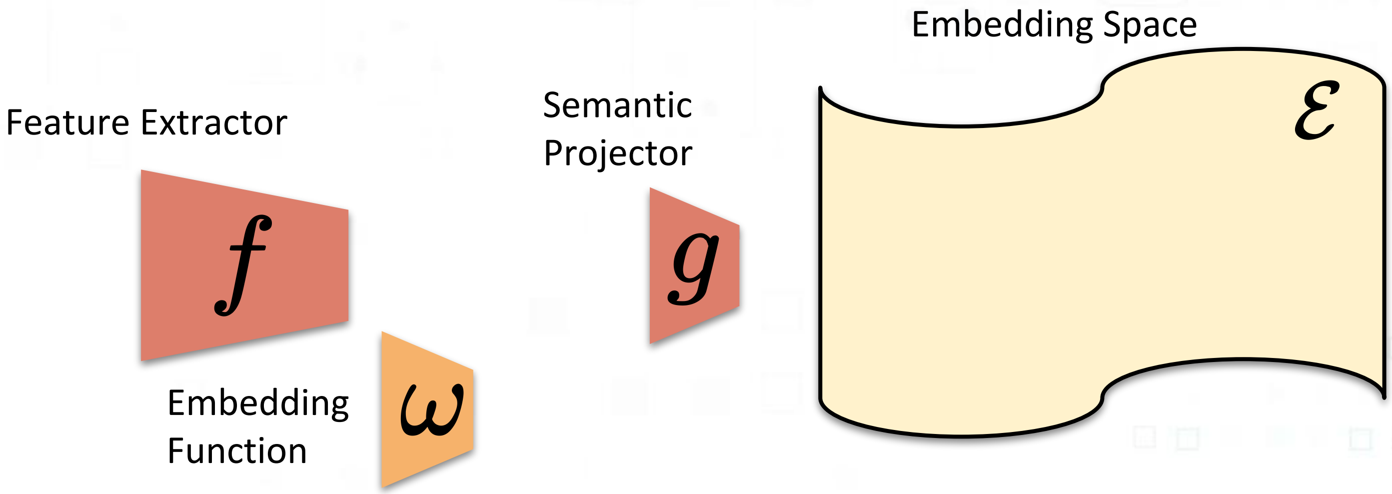

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

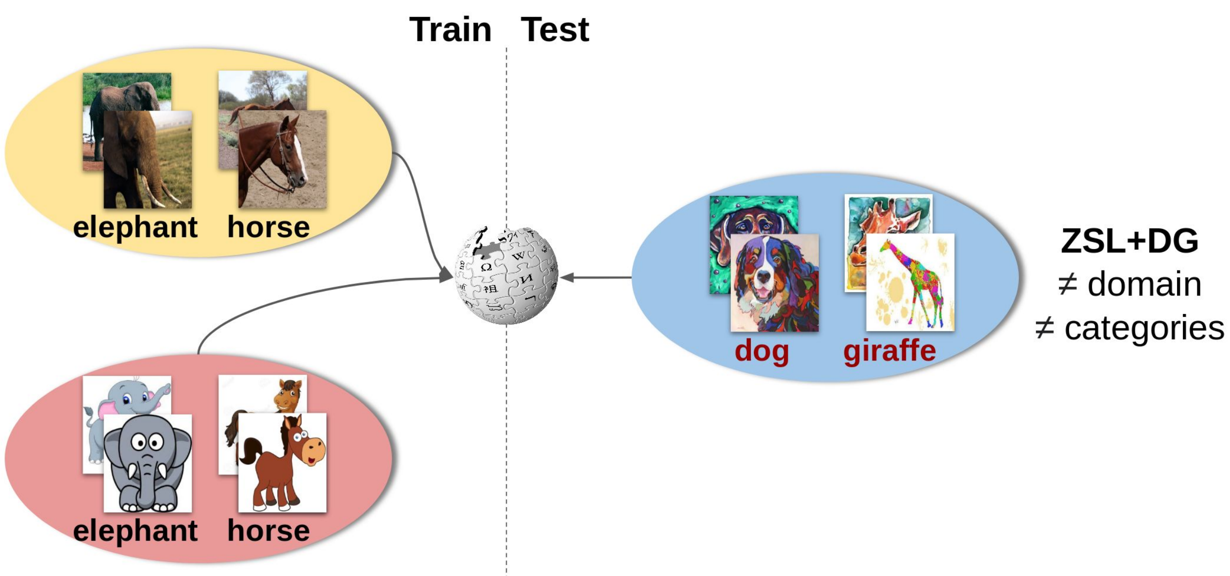

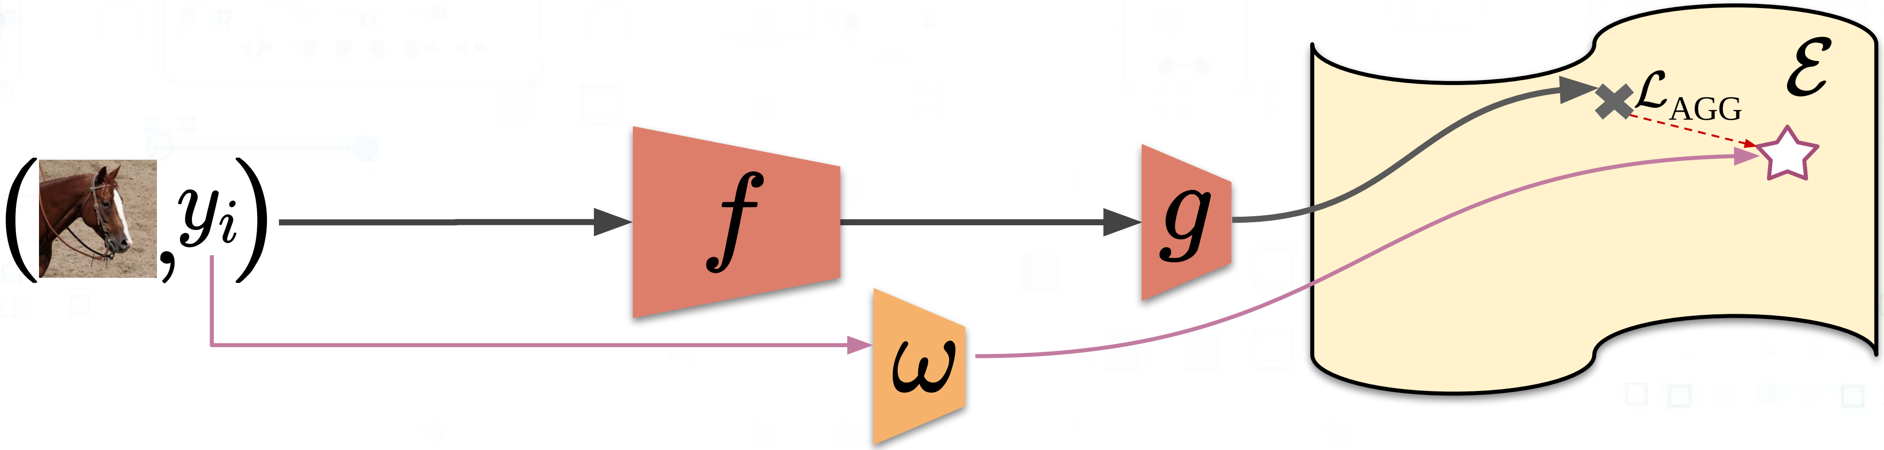

ECCV'20 ONLINE
23-28 AUGUST 2020

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

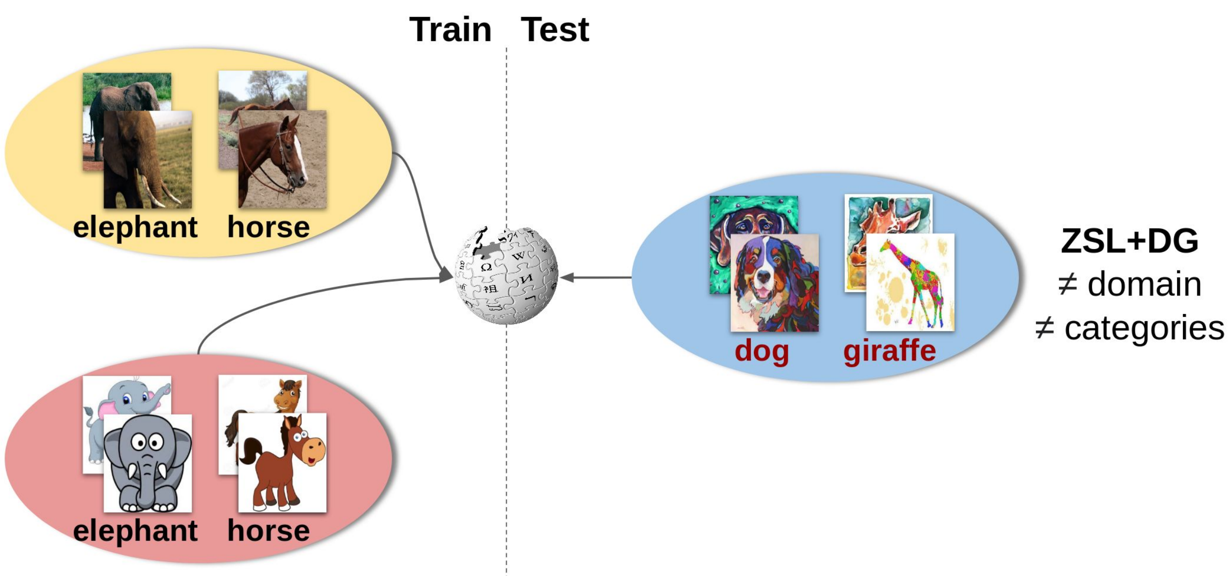

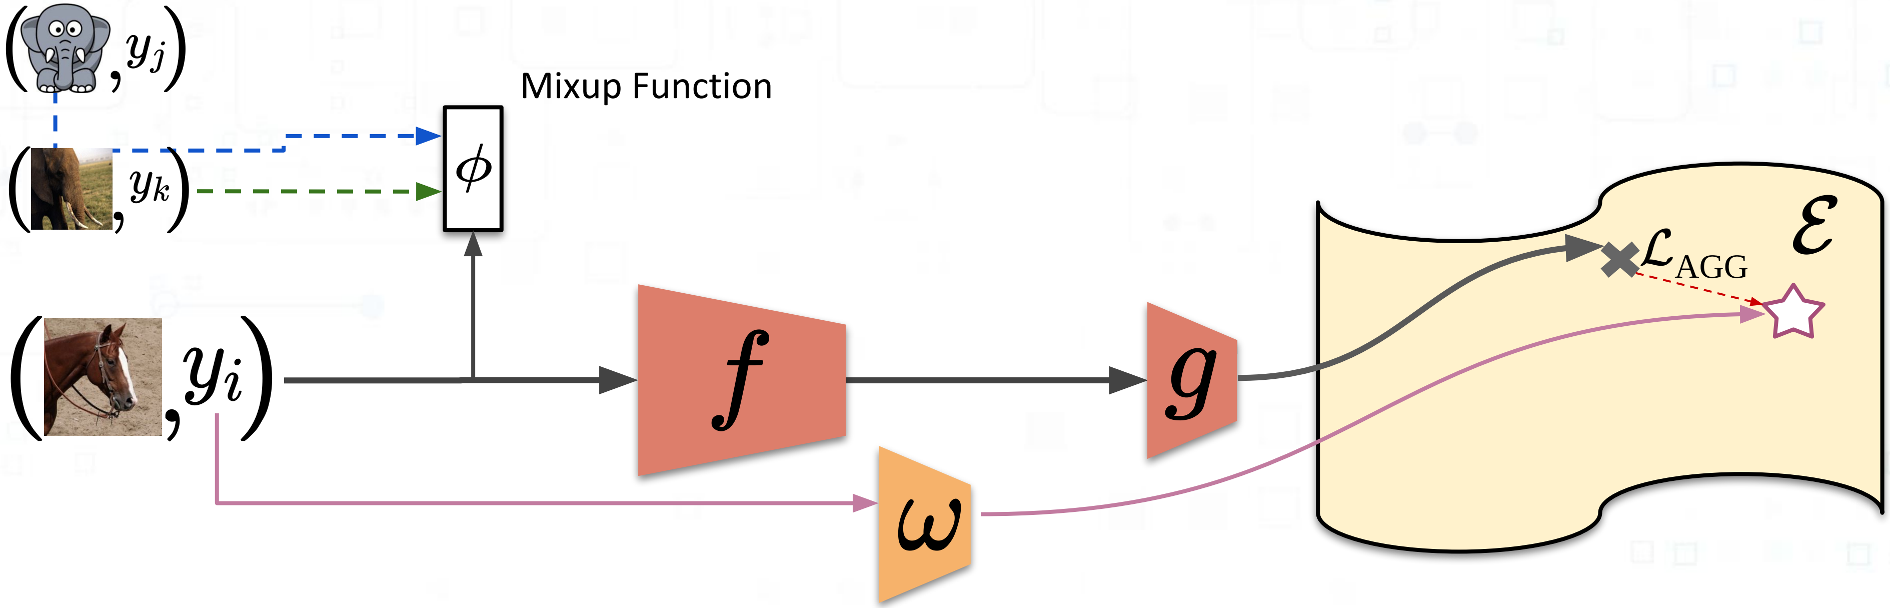

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

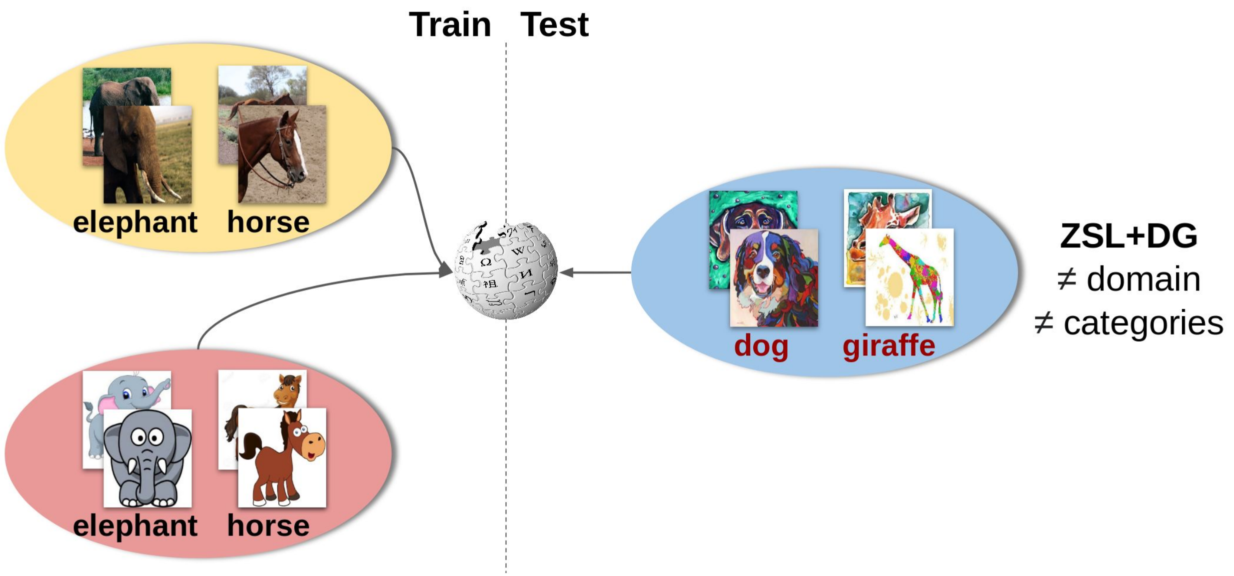

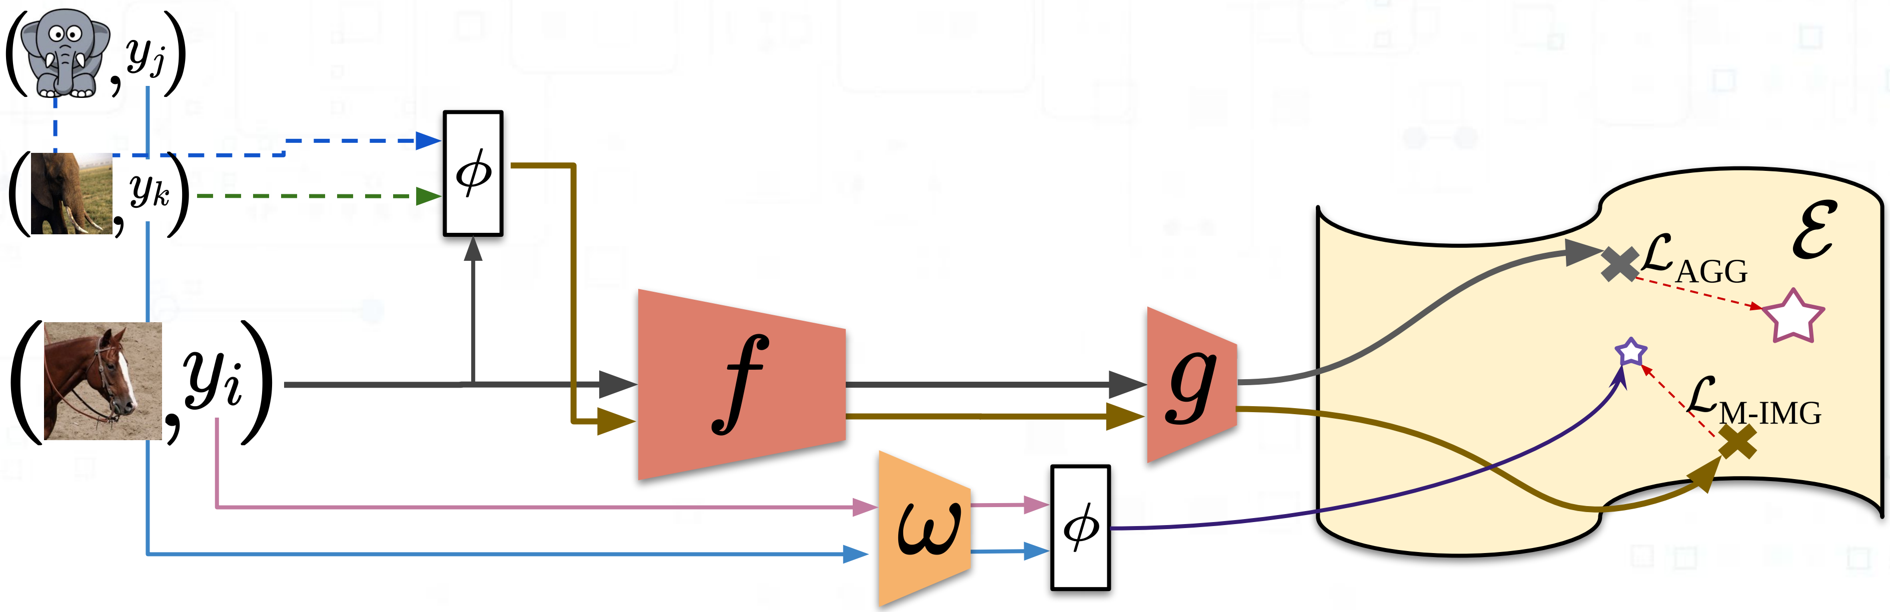

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

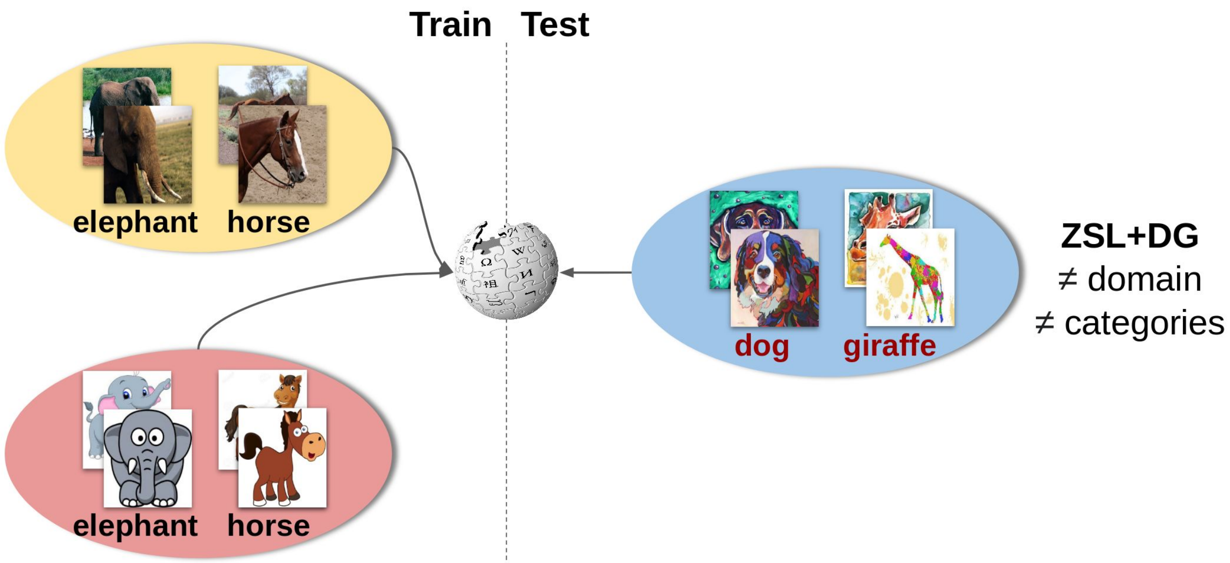

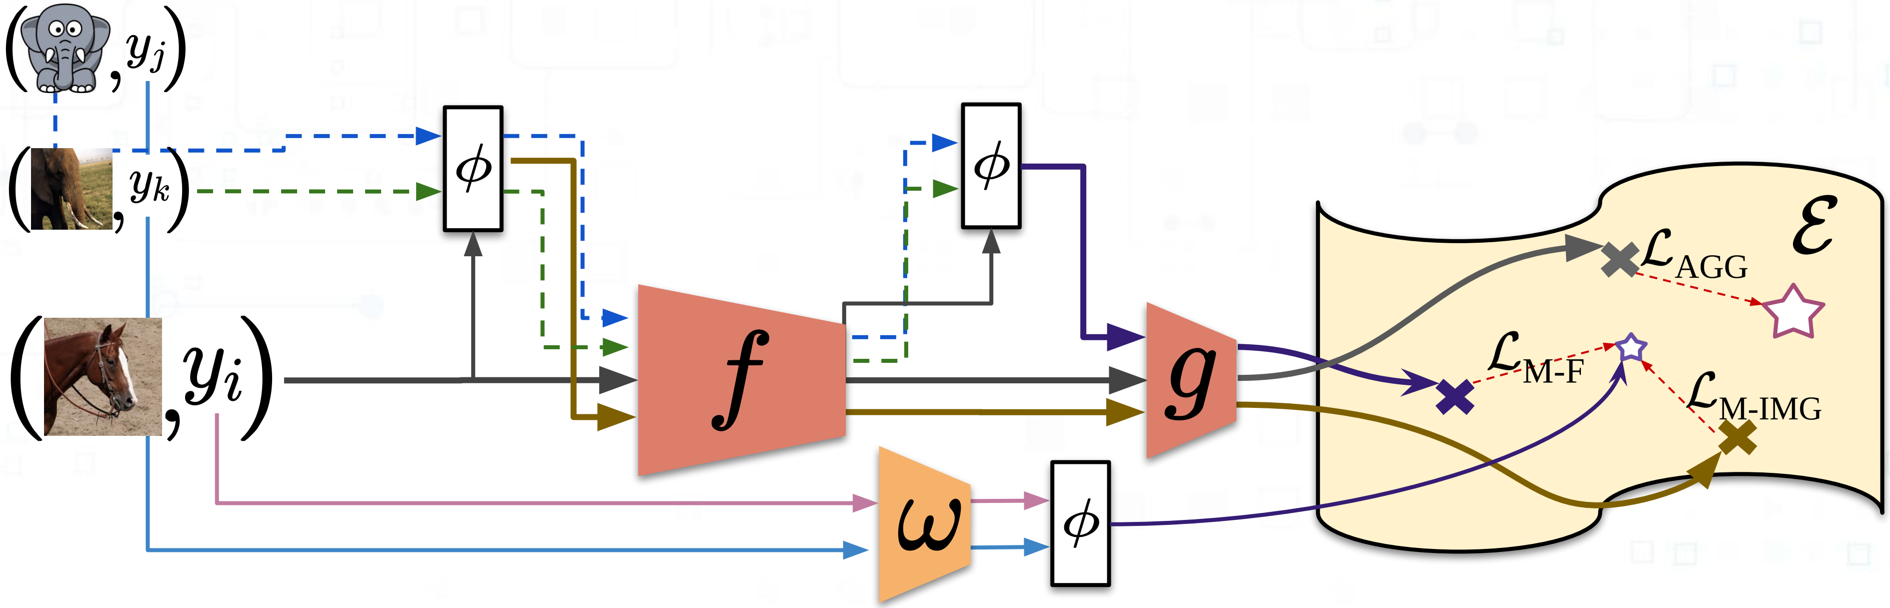

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

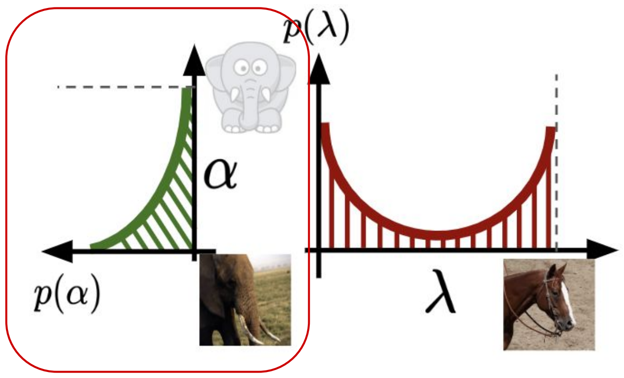

$$\phi(a_i, a_j, a_k) = \lambda a_i + (1-\lambda)\boxed{(\alpha a_j + (1-\alpha)a_k)}$$

Domain Picking

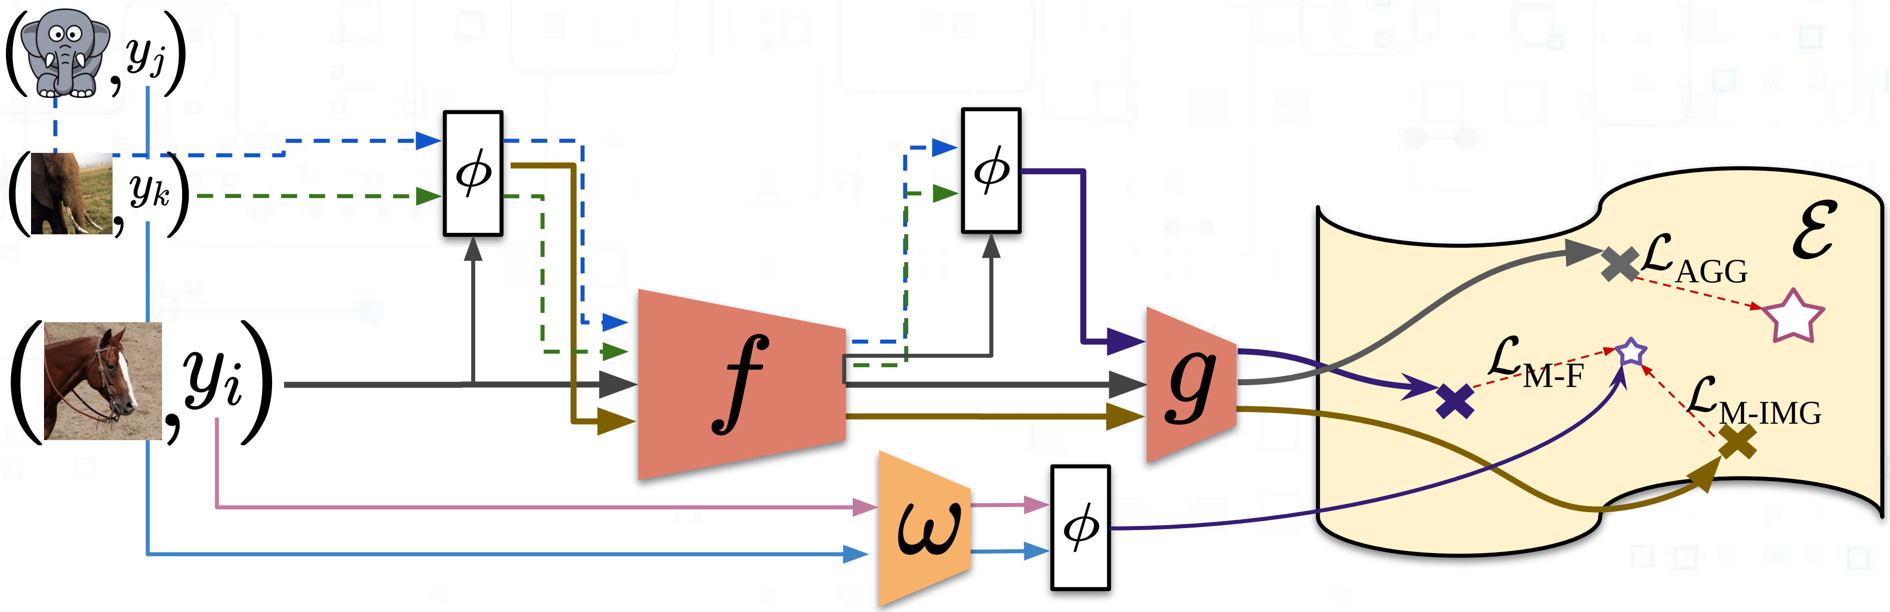

ECCV'20 ONLINE 23-28 AUGUST 2020

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

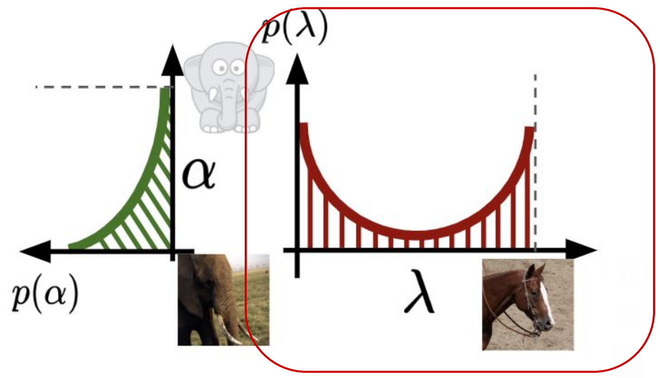

$$\phi(a_i, a_j, a_k) = \boxed{\lambda a_i + (1-\lambda)}(\alpha a_j + (1-\alpha) a_k)$$

Actual Mixing

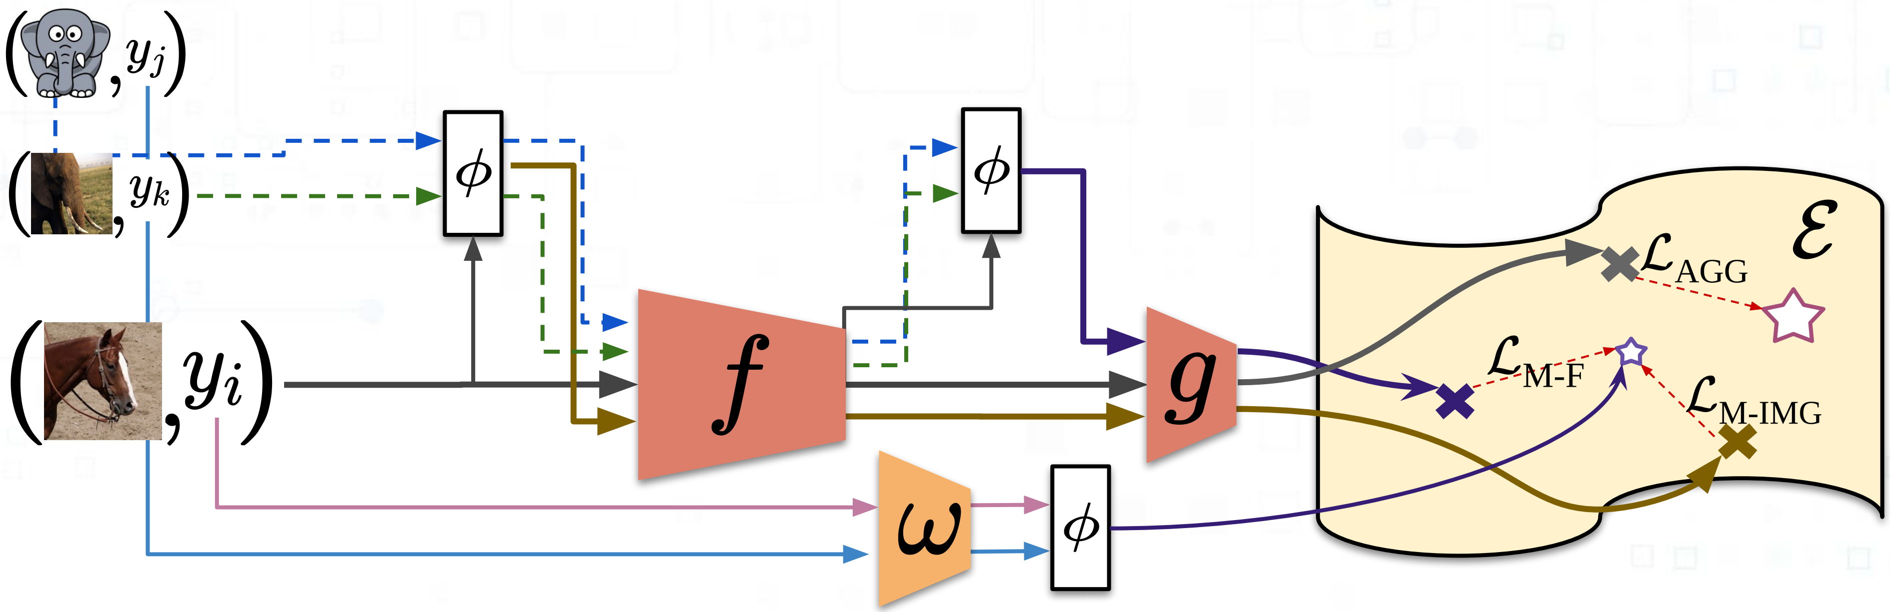

Zero-Shot Domain Generalization

[Towards Recognizing Unseen Categories in Unseen Domains, ECCV 2020]

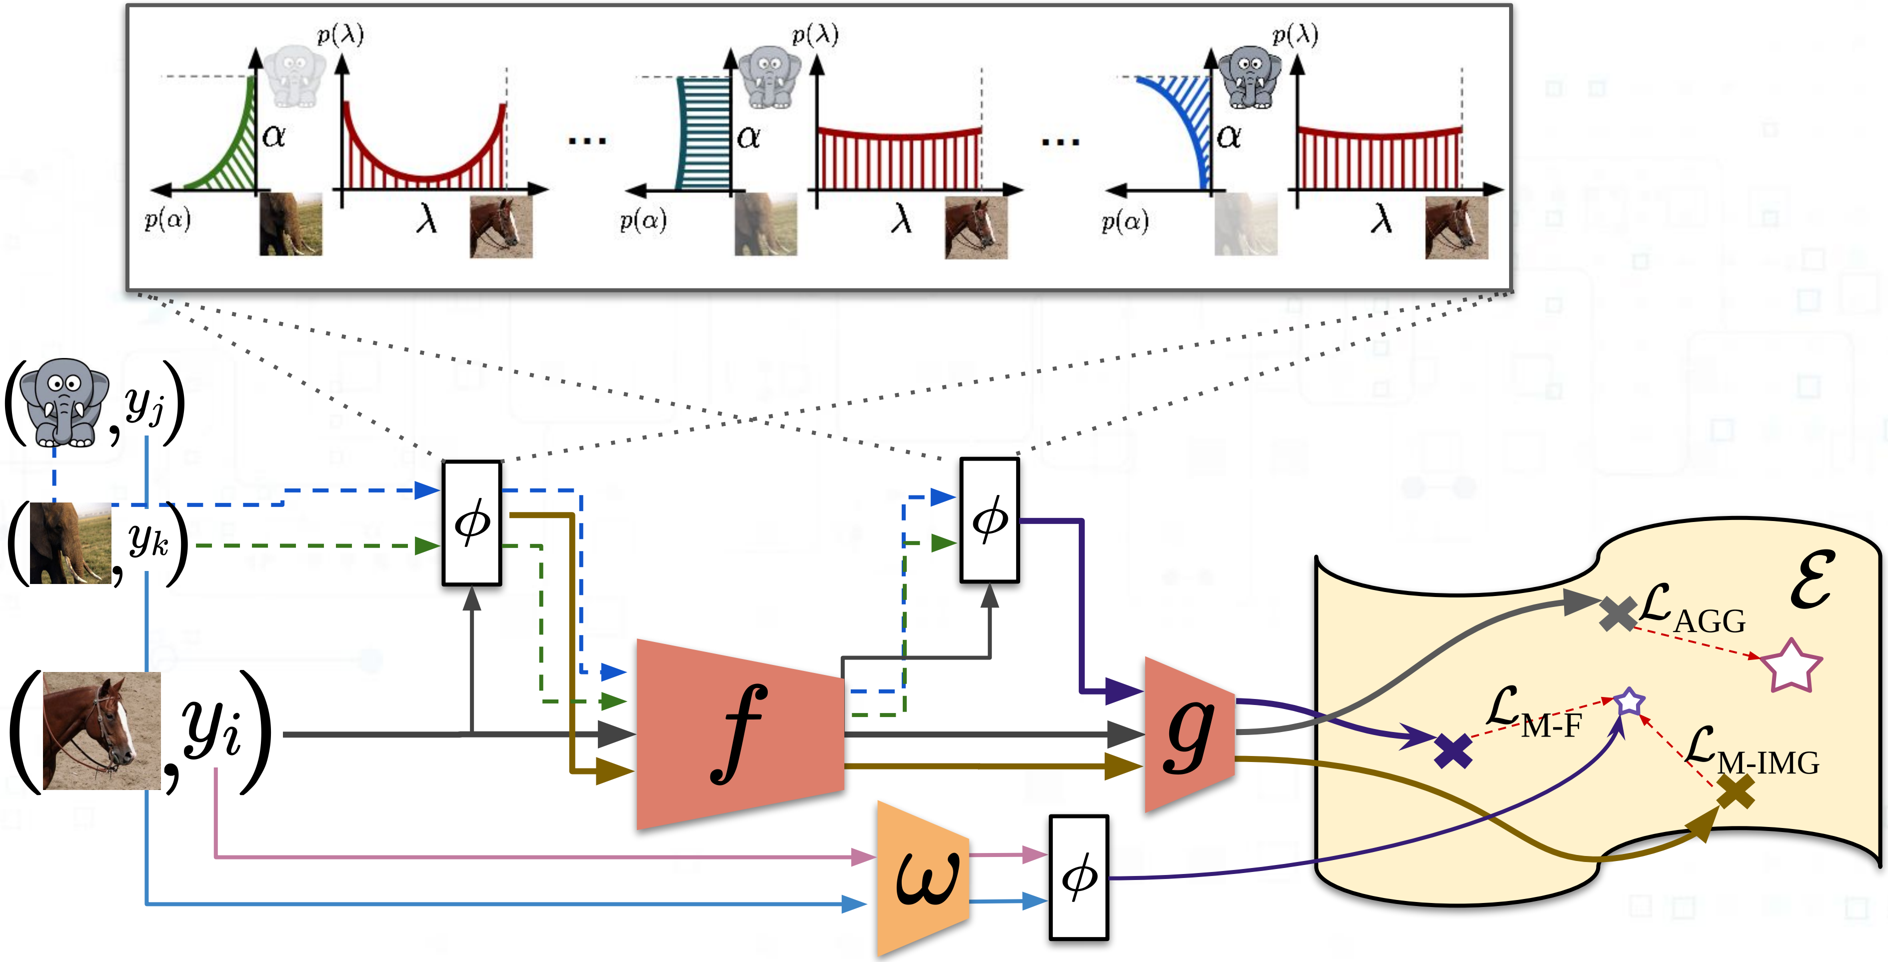

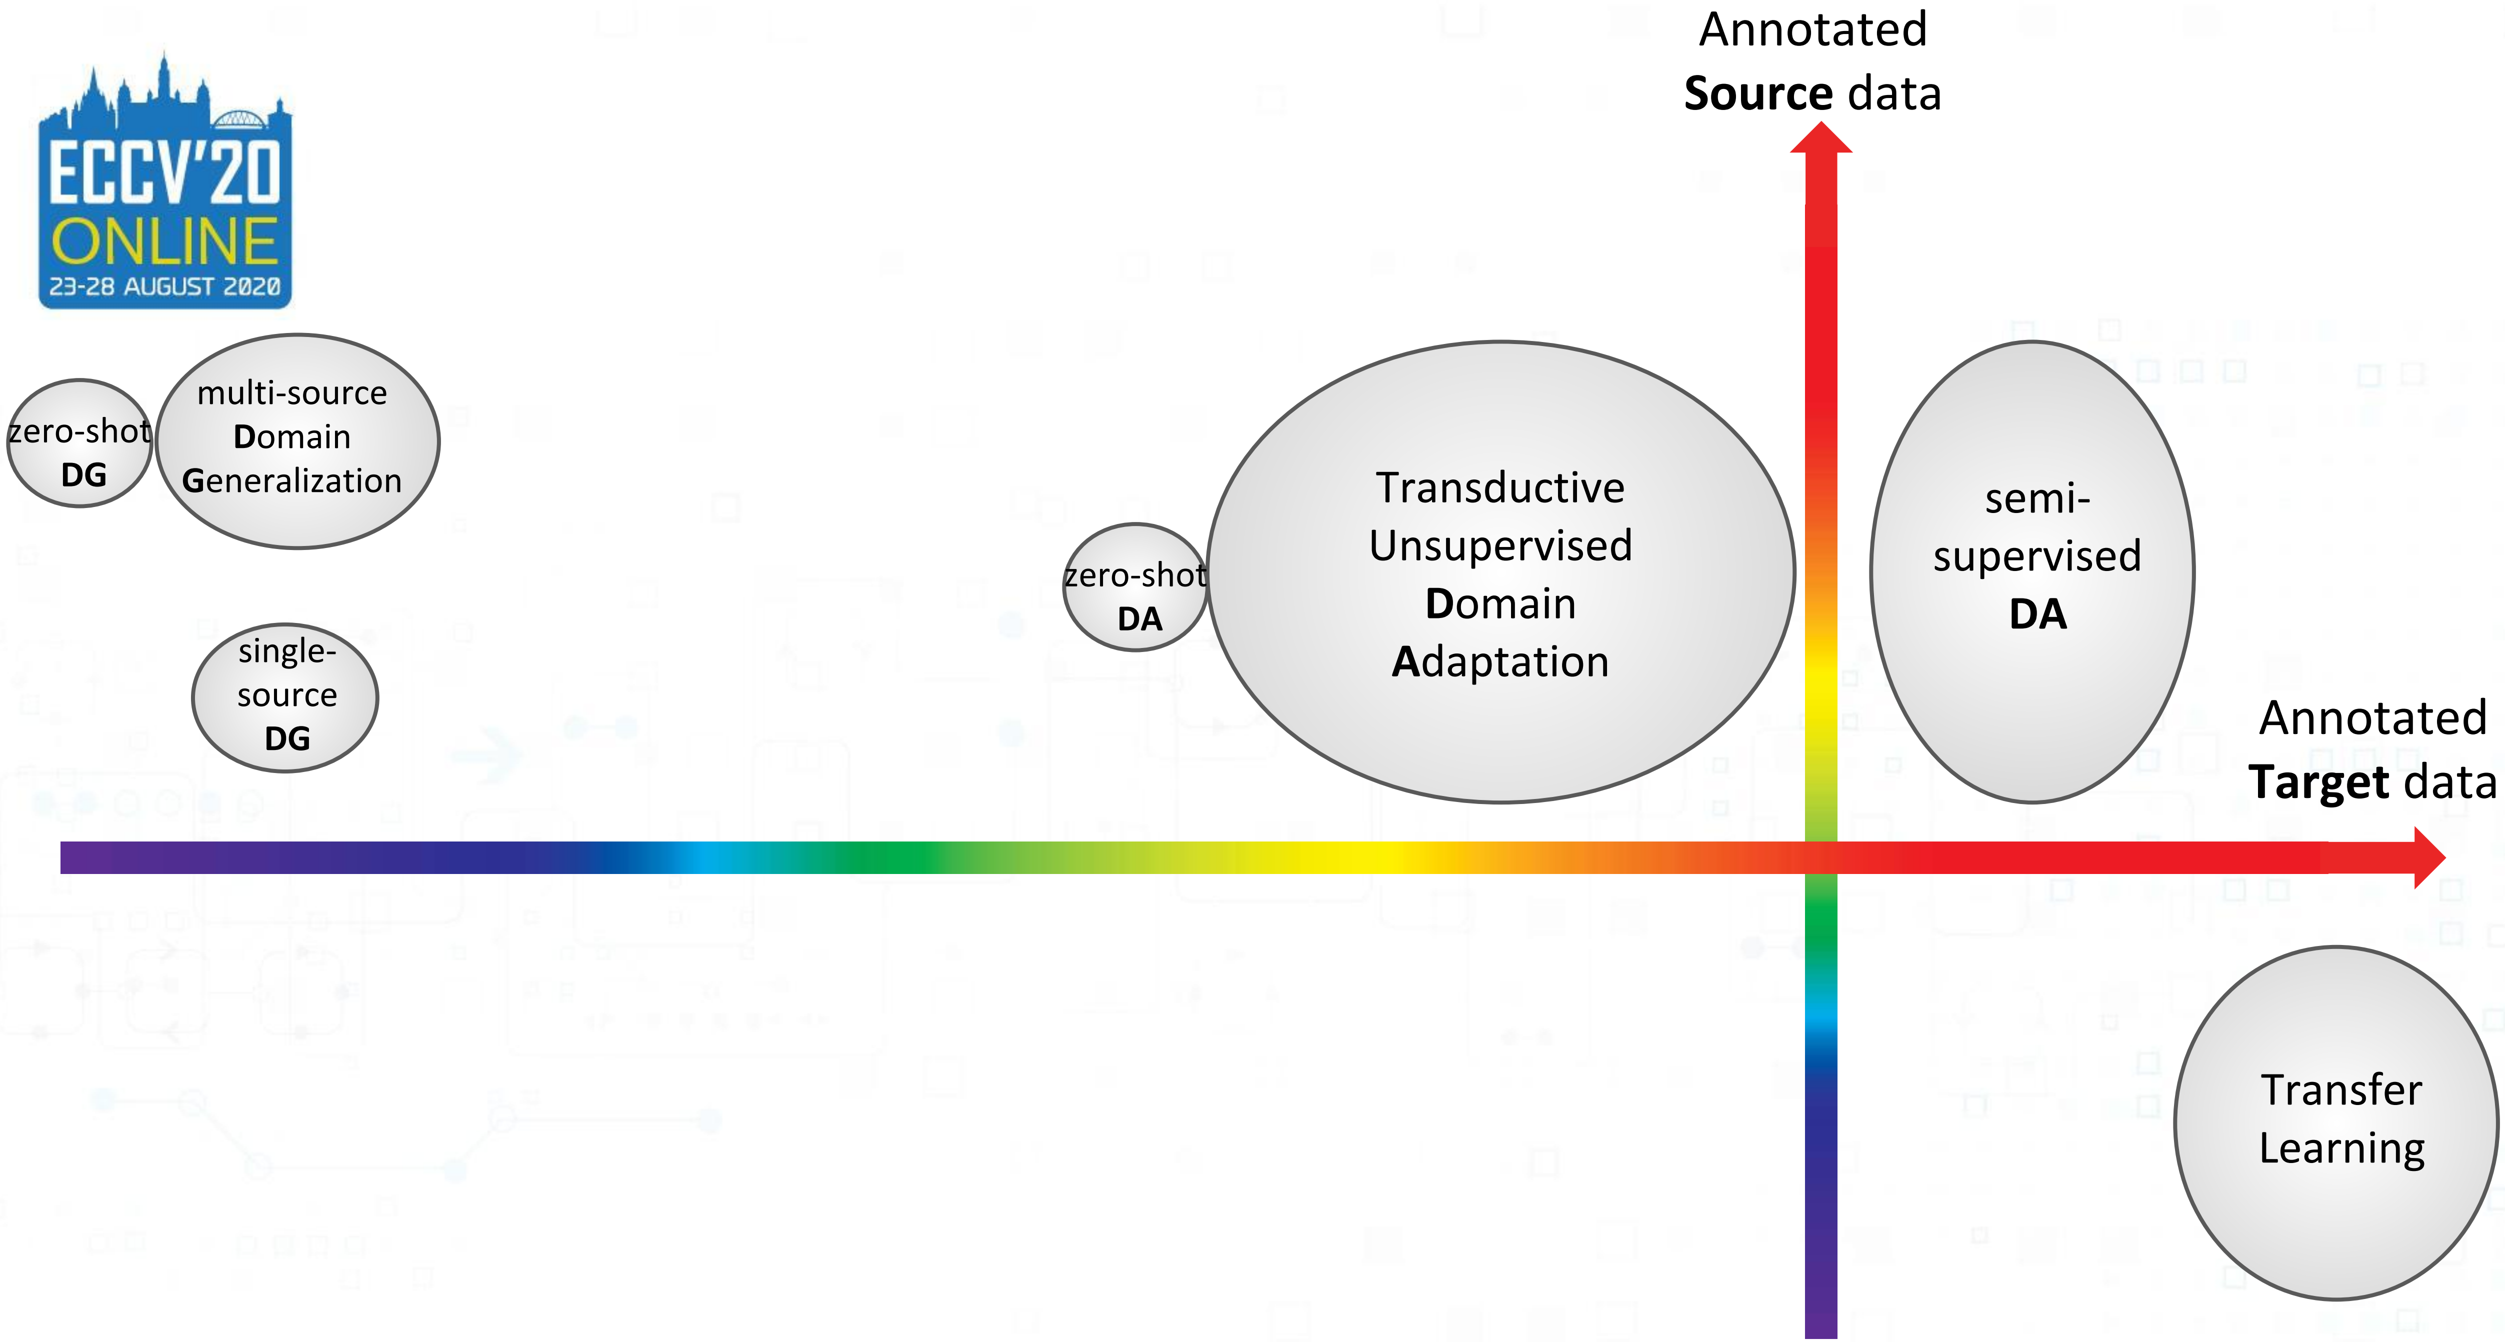

Annotated
Source data
multi-source
Domain
Generalization
zero-shot
DG
single-source
DG
zero-shot
DA
Transductive
Unsupervised
Domain
Adaptation
semi-supervised
DA
Annotated
Target data
Transfer
Learning

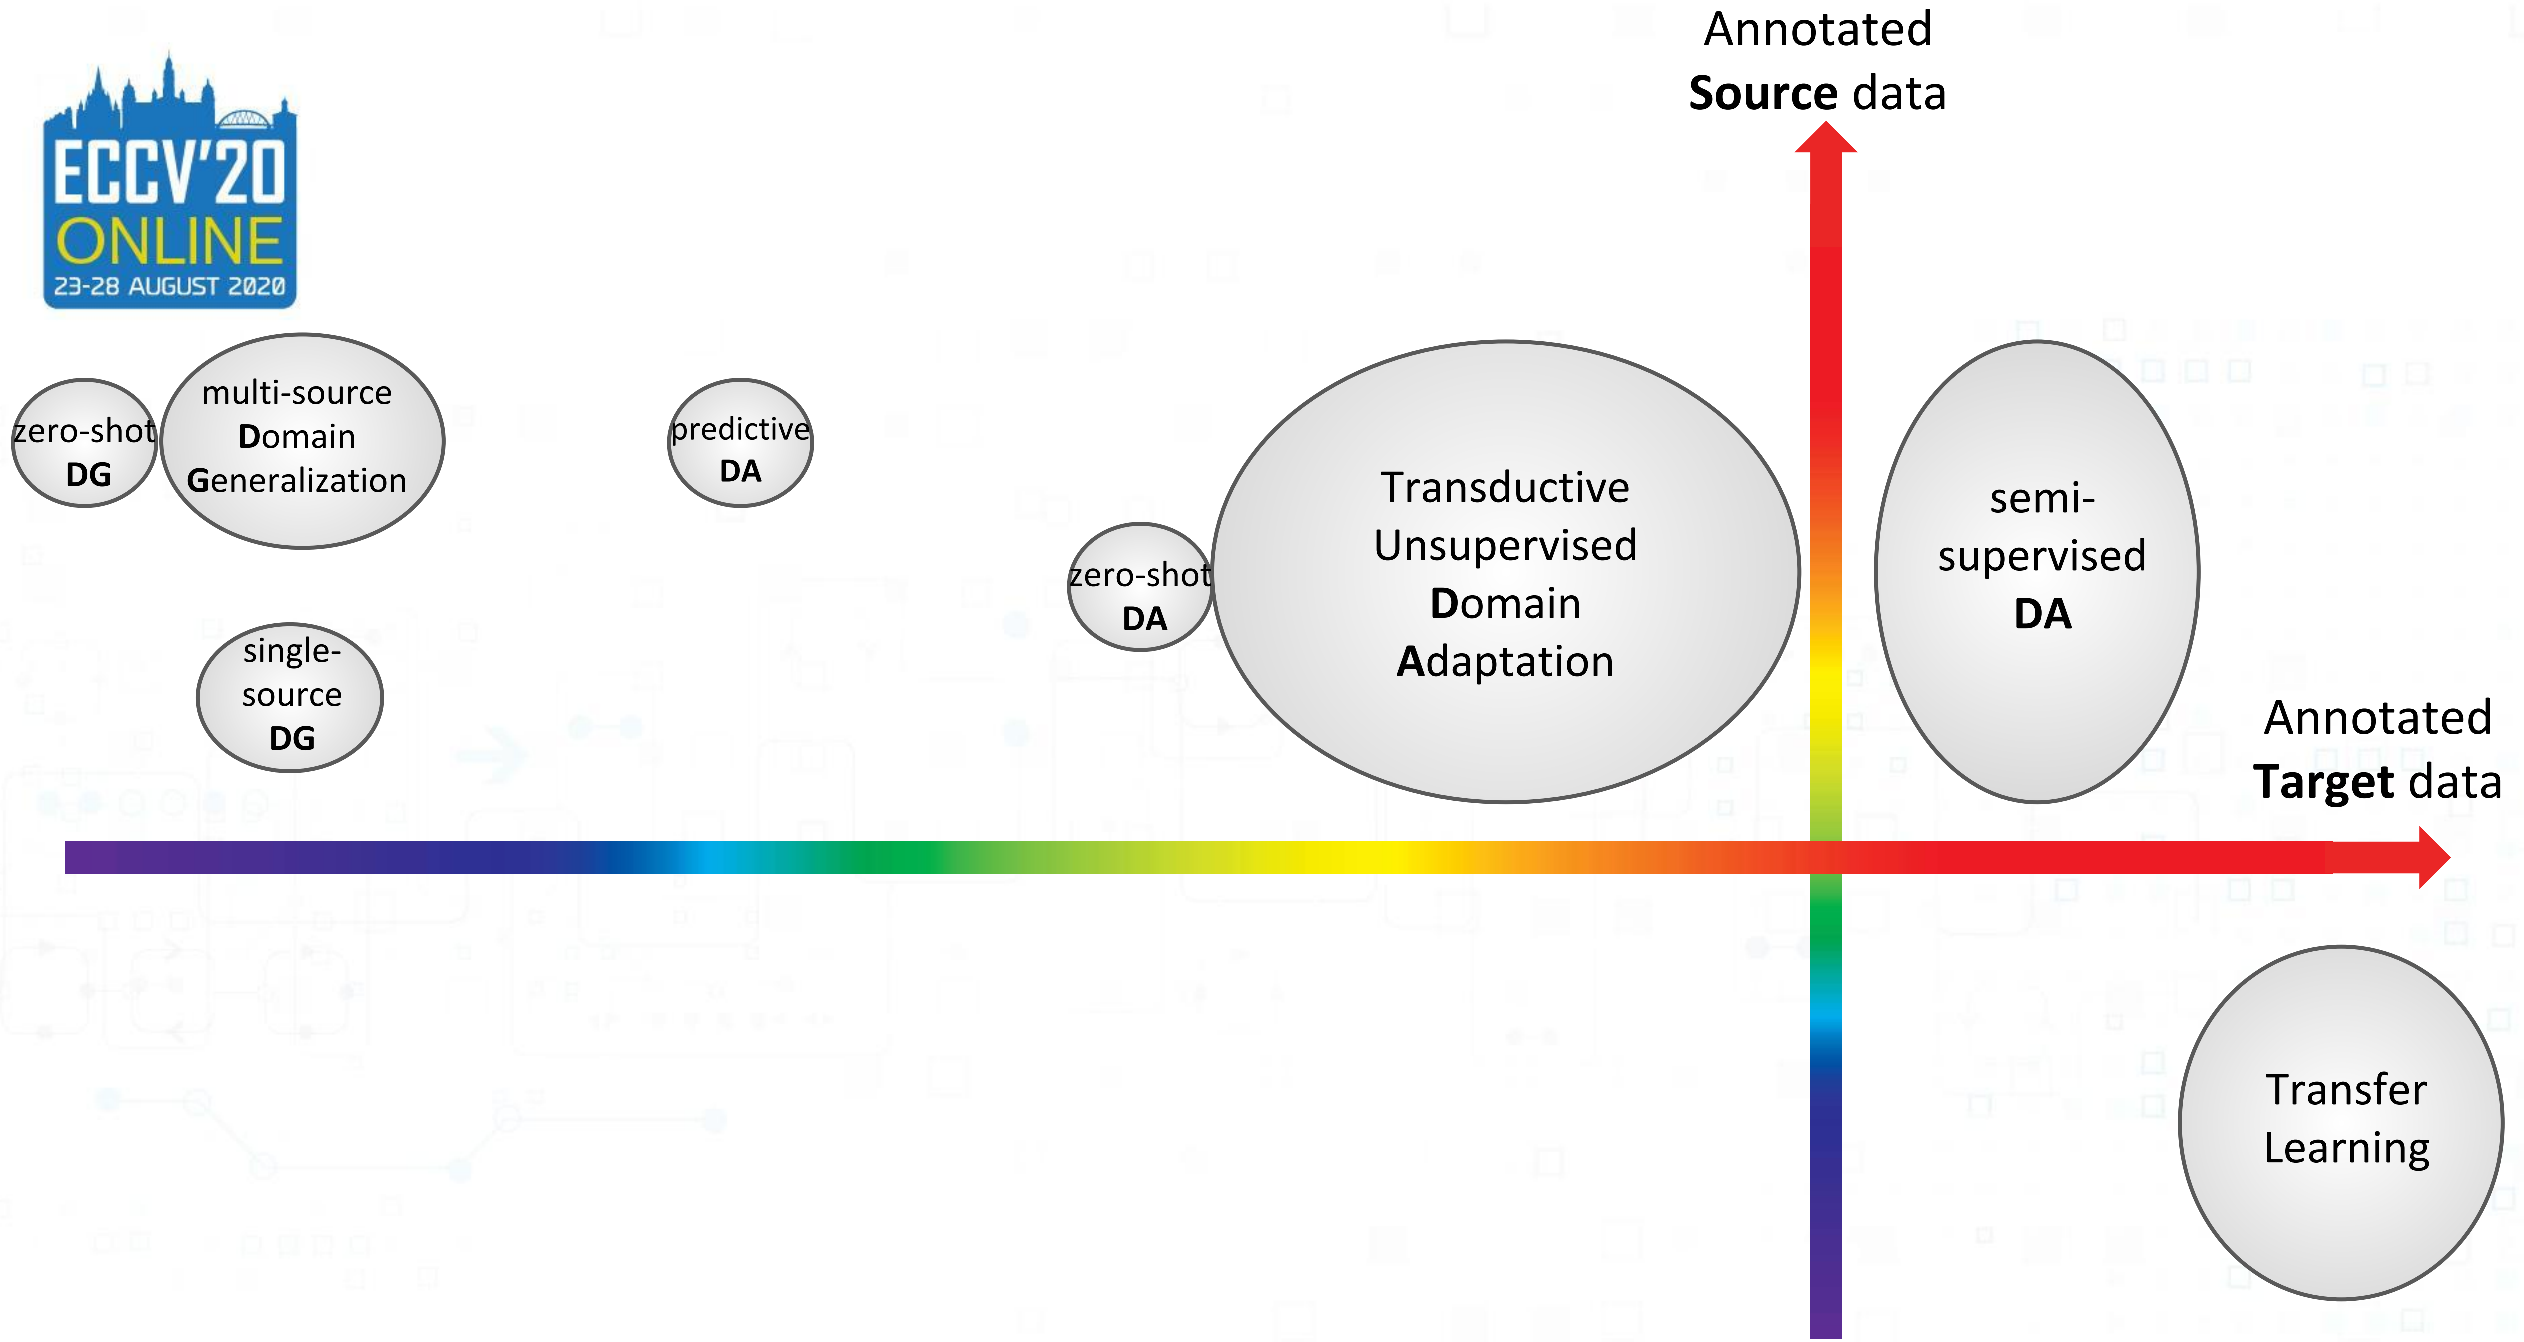
Annotated
Source data
zero-shot
DG
multi-source
Domain
Generalization
predictive
DA
zero-shot
DA
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
single-
source
DG
Annotated
Target data
Transfer
Learning

Predictive DA

[Multivariate Regression on the Grassmannian for Predicting Novel Domains, CVPR 2016]
[AdaGraph: Unifying Predictive and Continuous Domain Adaptation through Graphs, CVPR 2019]

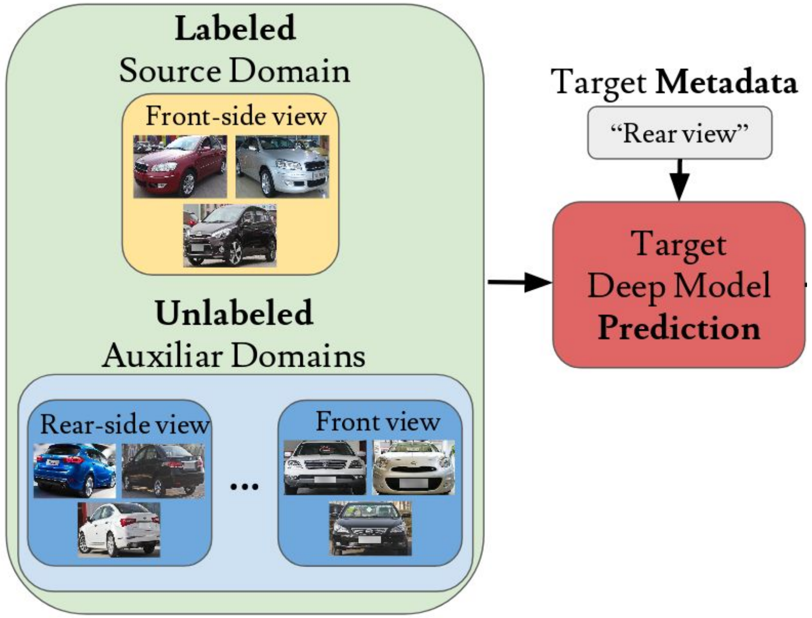

Predictive DA

[Multivariate Regression on the Grassmannian for Predicting Novel Domains, CVPR 2016]
[AdaGraph: Unifying Predictive and Continuous Domain Adaptation through Graphs, CVPR 2019]

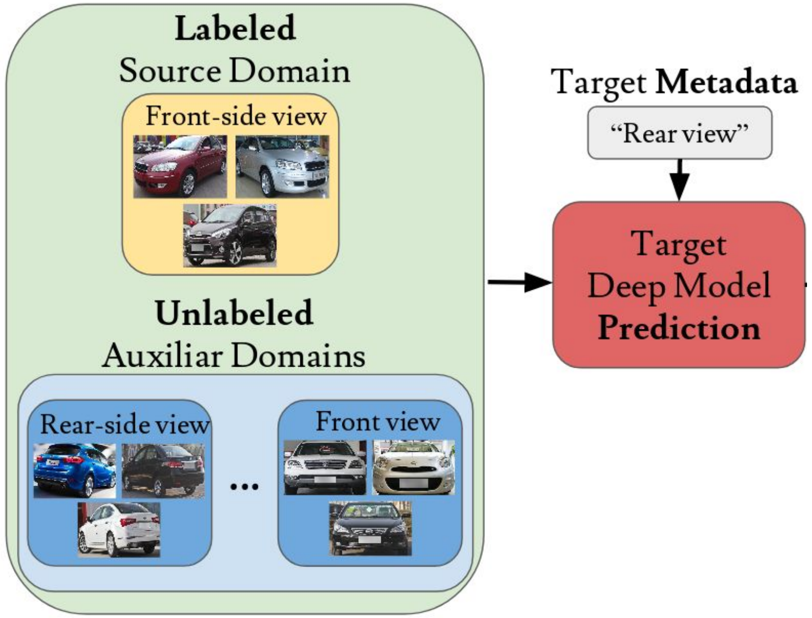

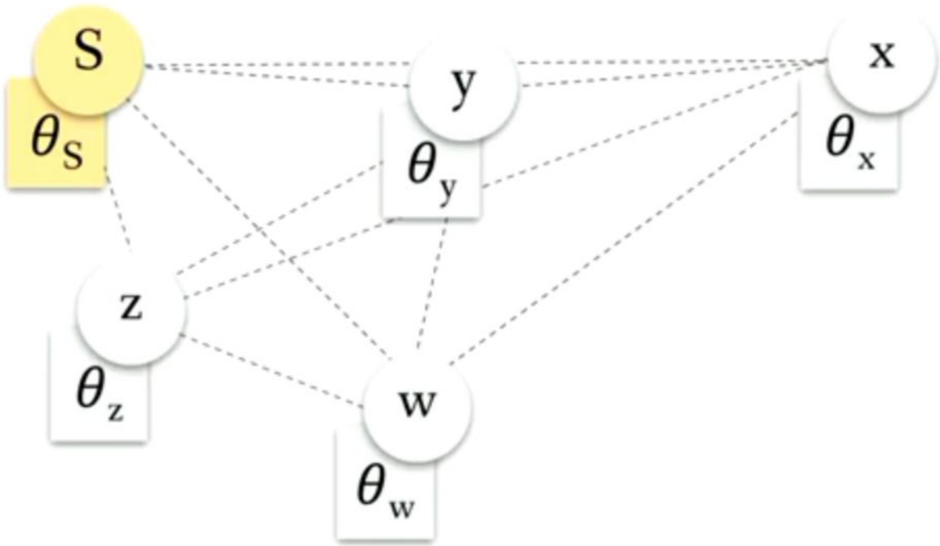

Predictive DA

[Multivariate Regression on the Grassmannian for Predicting Novel Domains, CVPR 2016]
[AdaGraph: Unifying Predictive and Continuous Domain Adaptation through Graphs, CVPR 2019]

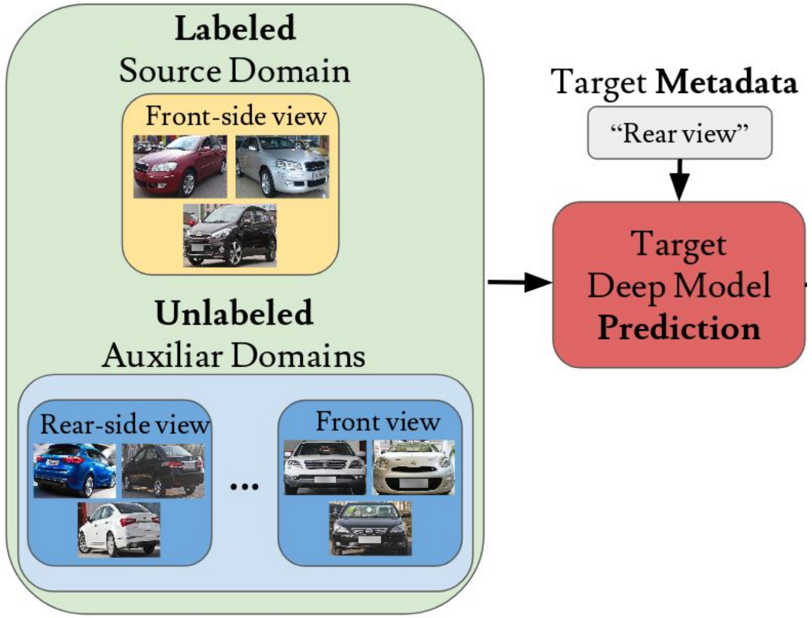

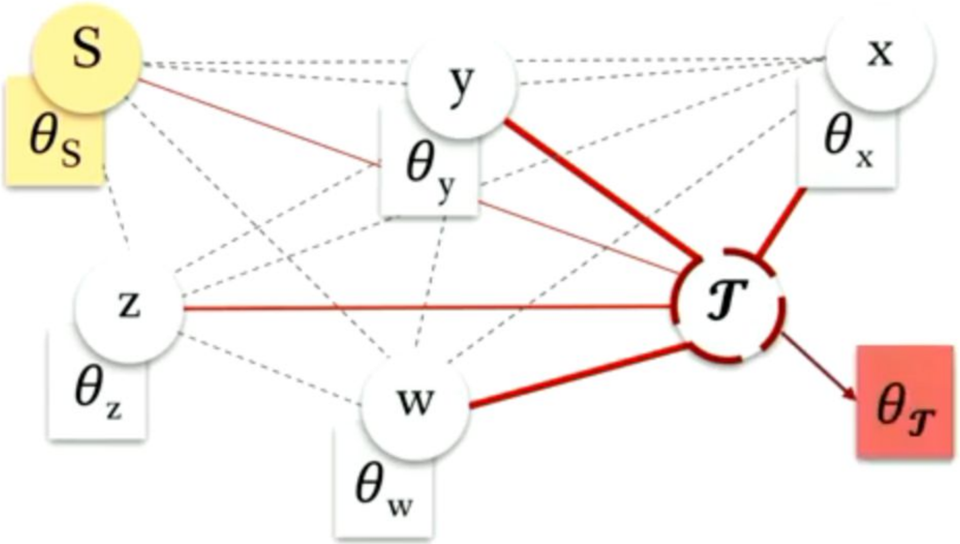

Predictive DA

[Multivariate Regression on the Grassmannian for Predicting Novel Domains, CVPR 2016]
[AdaGraph: Unifying Predictive and Continuous Domain Adaptation through Graphs, CVPR 2019]

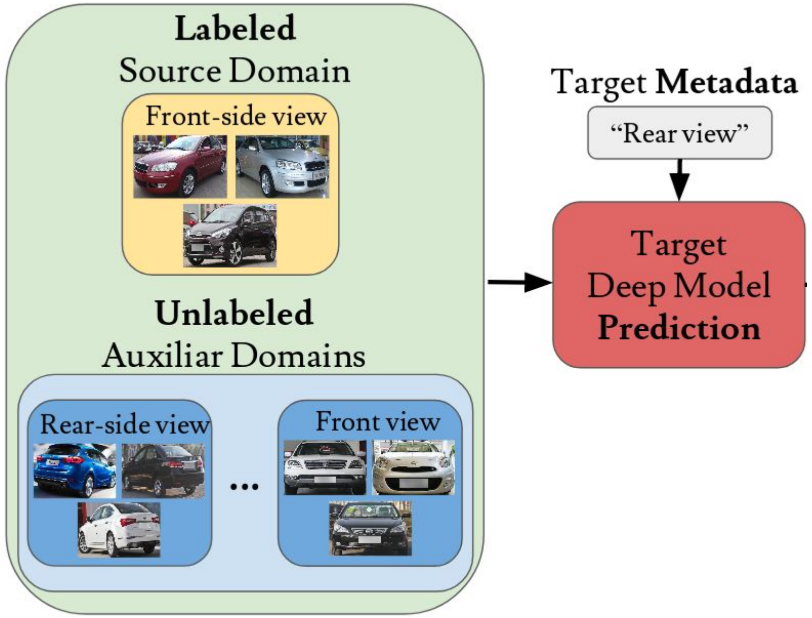

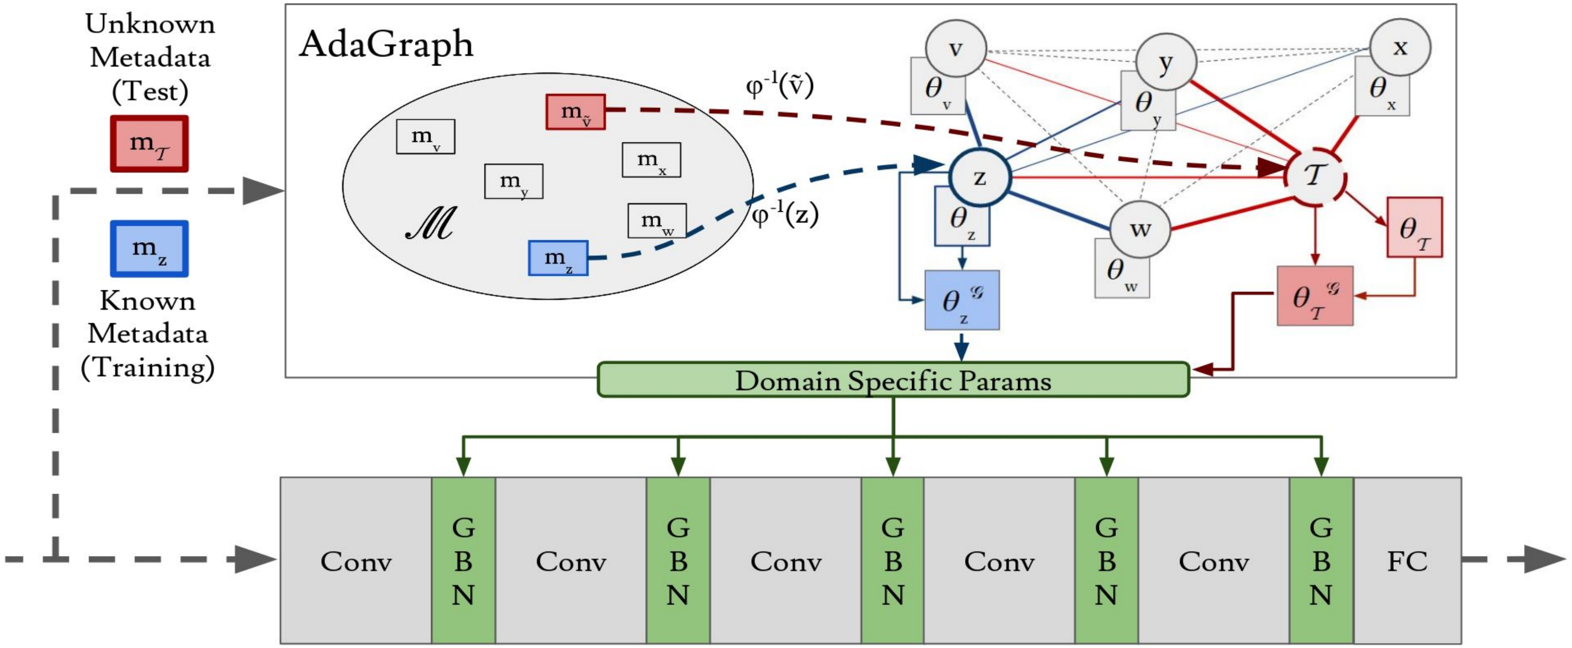

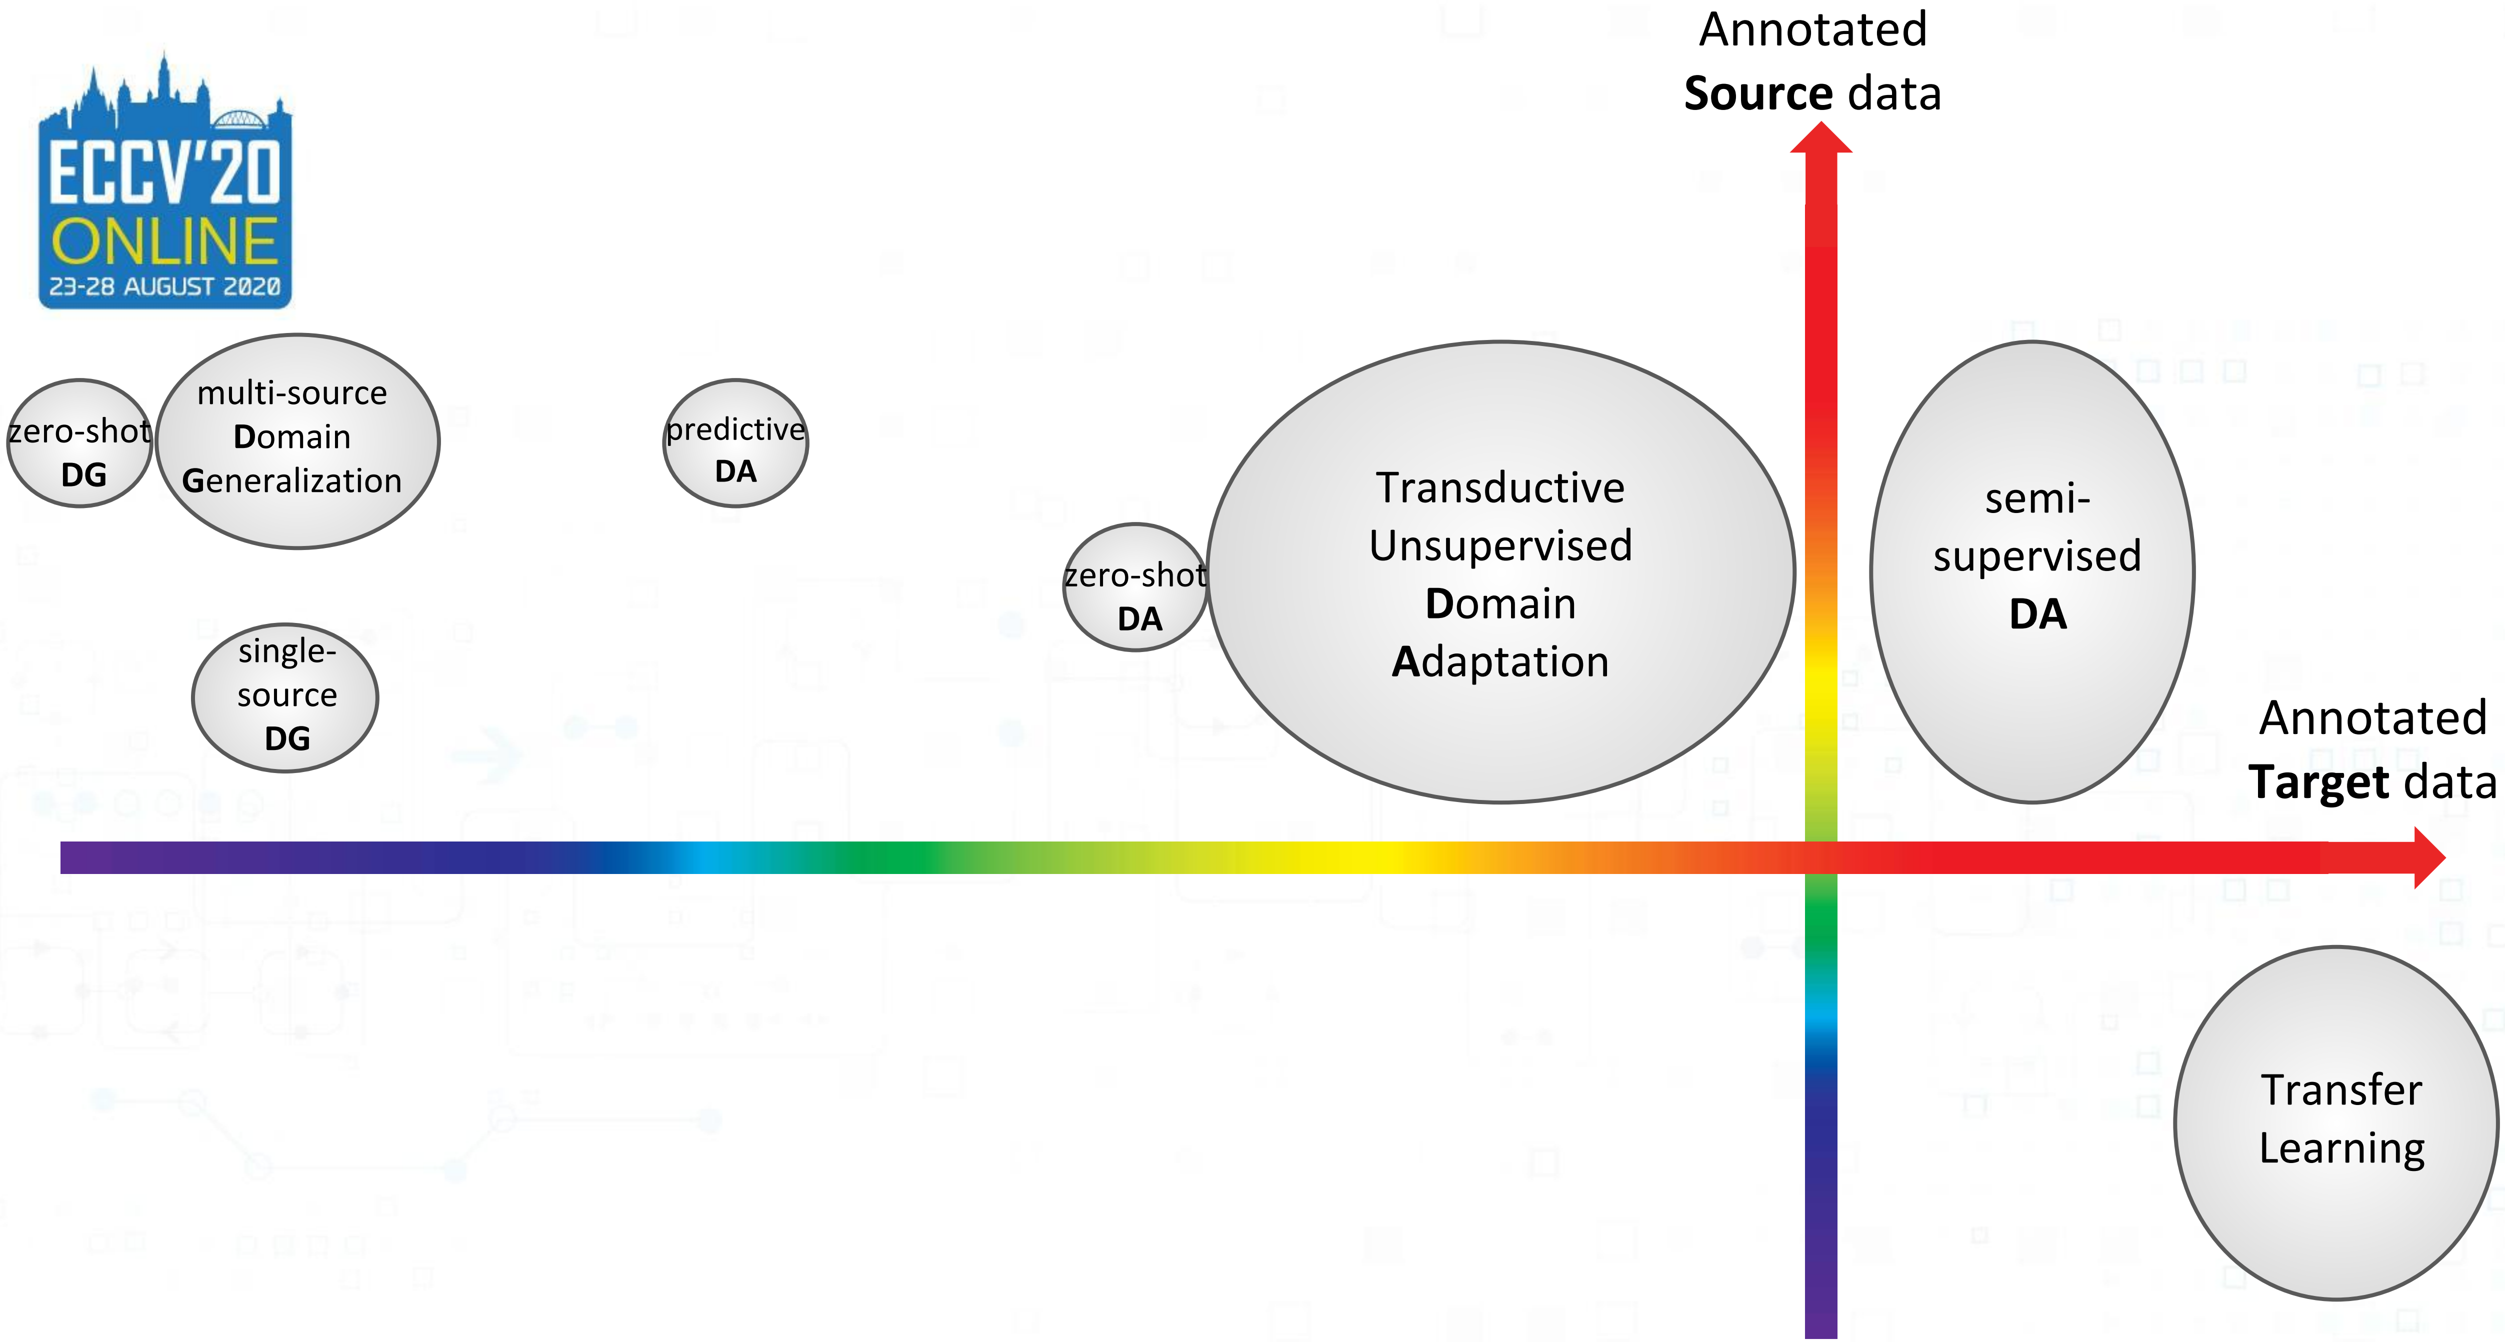

Annotated **Source** data
Annotated **Target** data
zero-shot **DG**
multi-source **D**omain **G**eneralization
single-source **DG**
predictive **DA**
zero-shot **DA**
Transductive Unsupervised **D**omain **A**daptation
semi-supervised **DA**
Transfer Learning

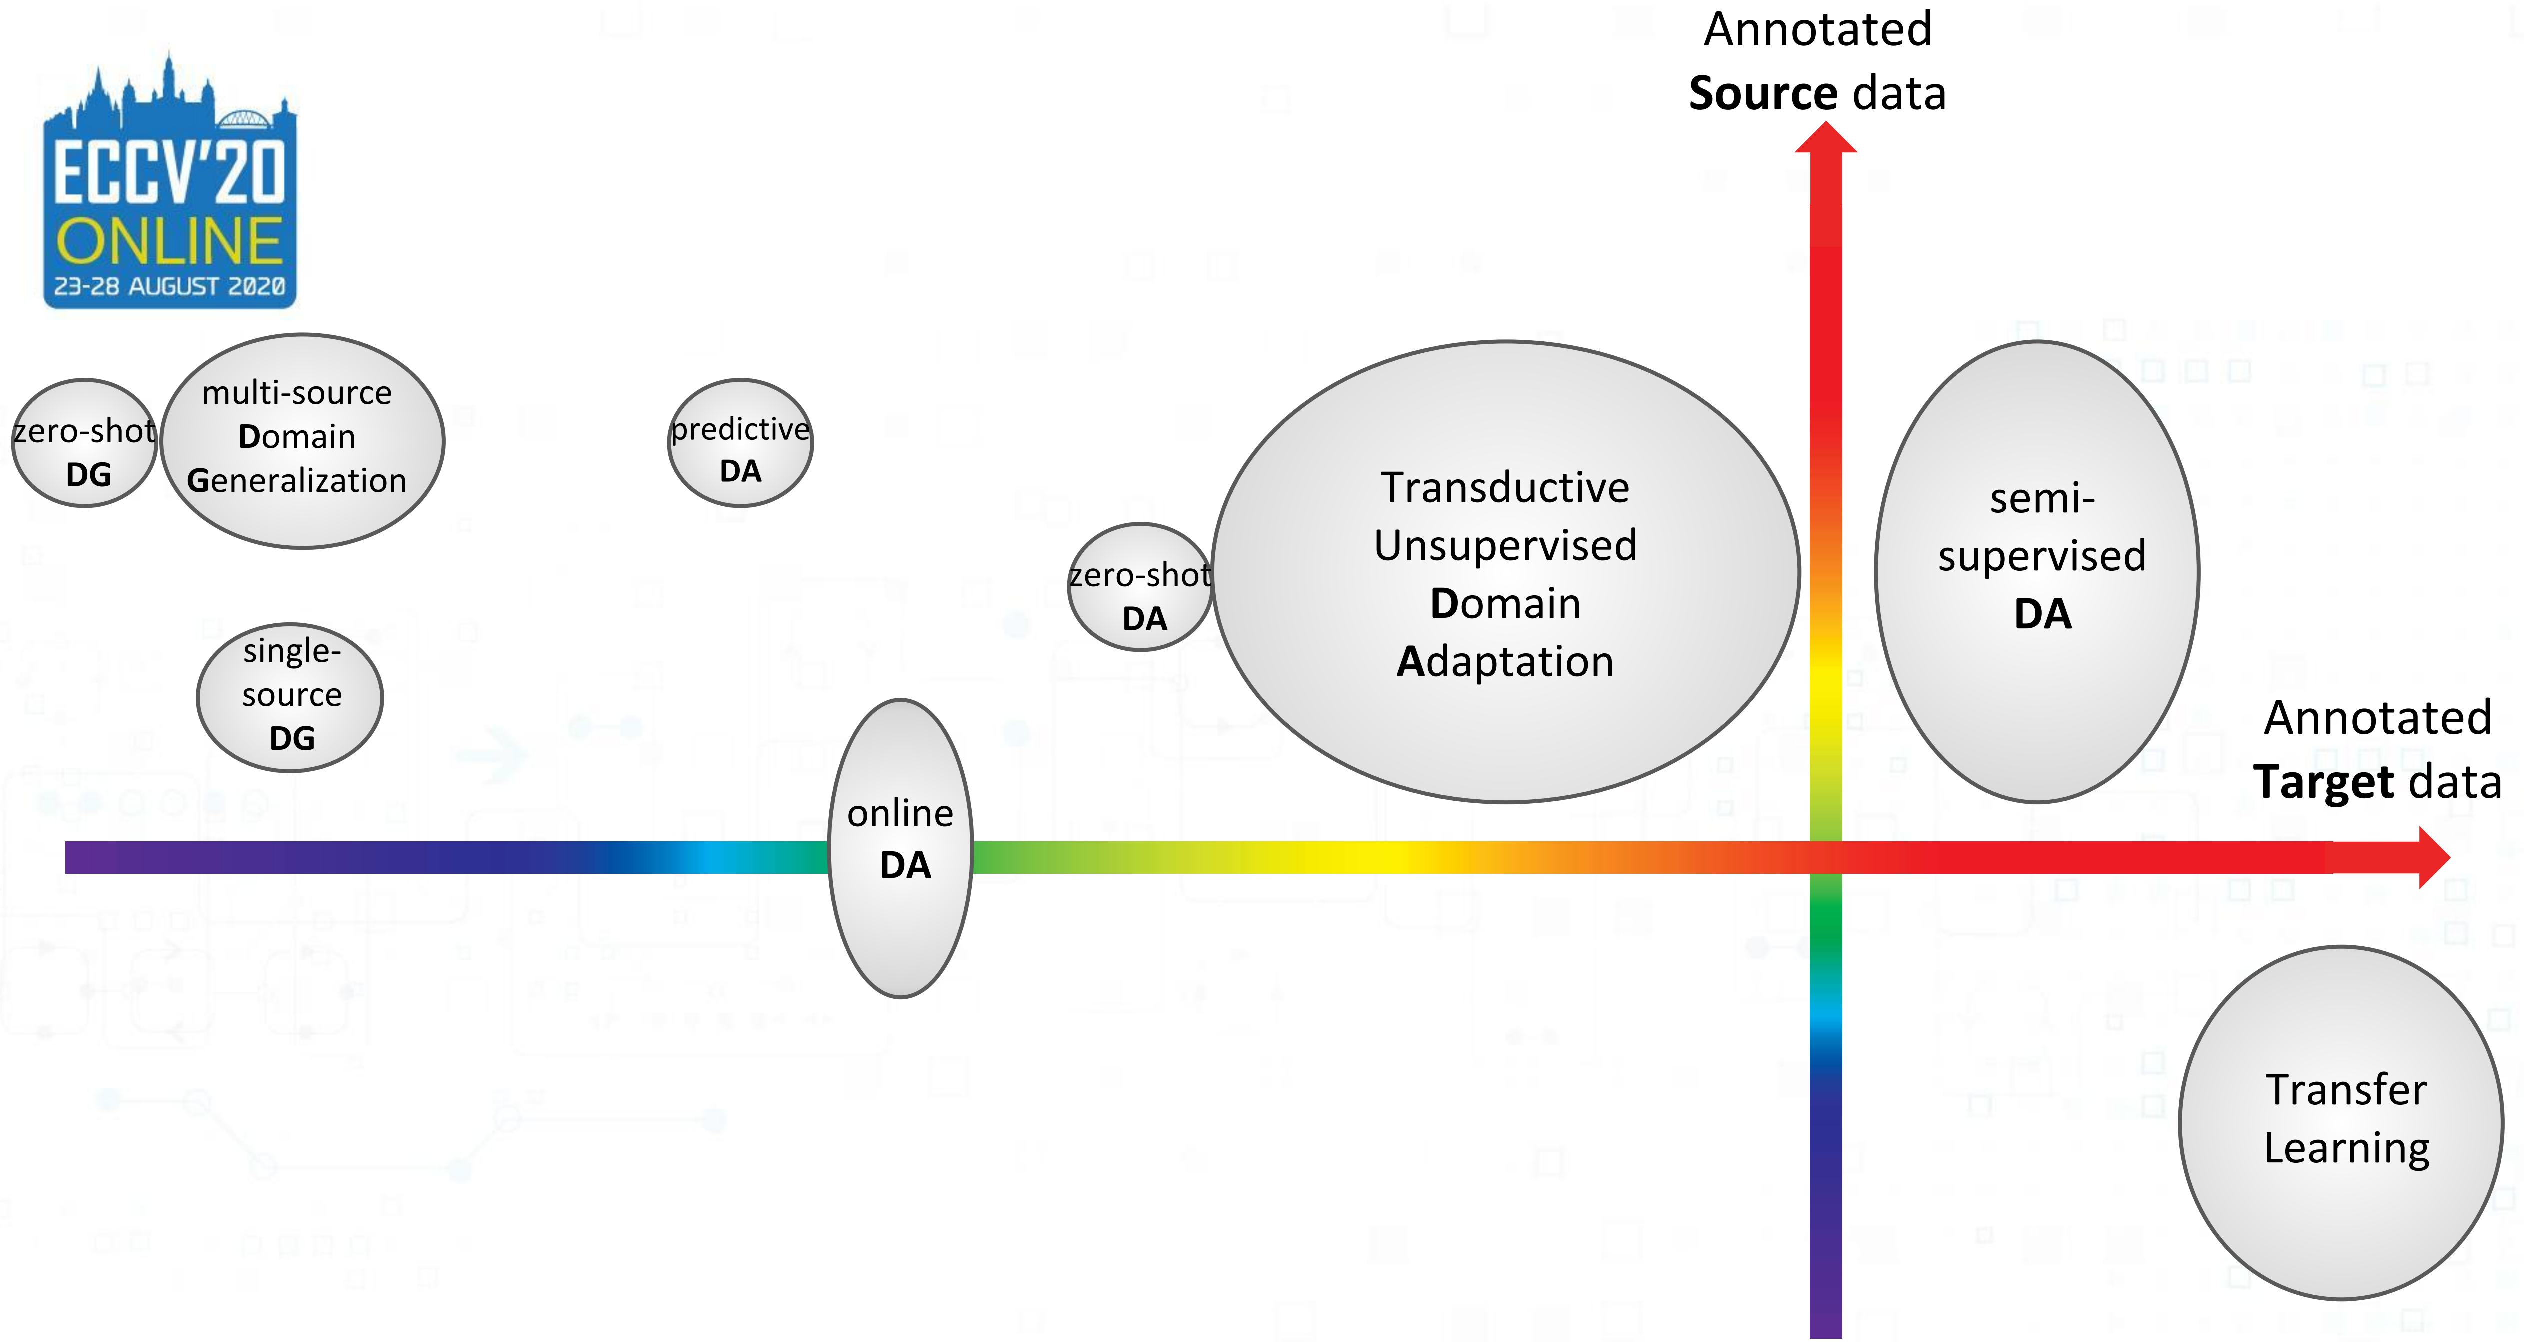
Annotated
Source data
Annotated
Target data
zero-shot
DG
multi-source
Domain
Generalization
single-
source
DG
predictive
DA
zero-shot
DA
Transductive
Unsupervised
Domain
Adaptation
online
DA
semi-
supervised
DA
Transfer
Learning

Online DA

[Kitting in the Wild through Online Domain Adaptation, ICRA 2018]

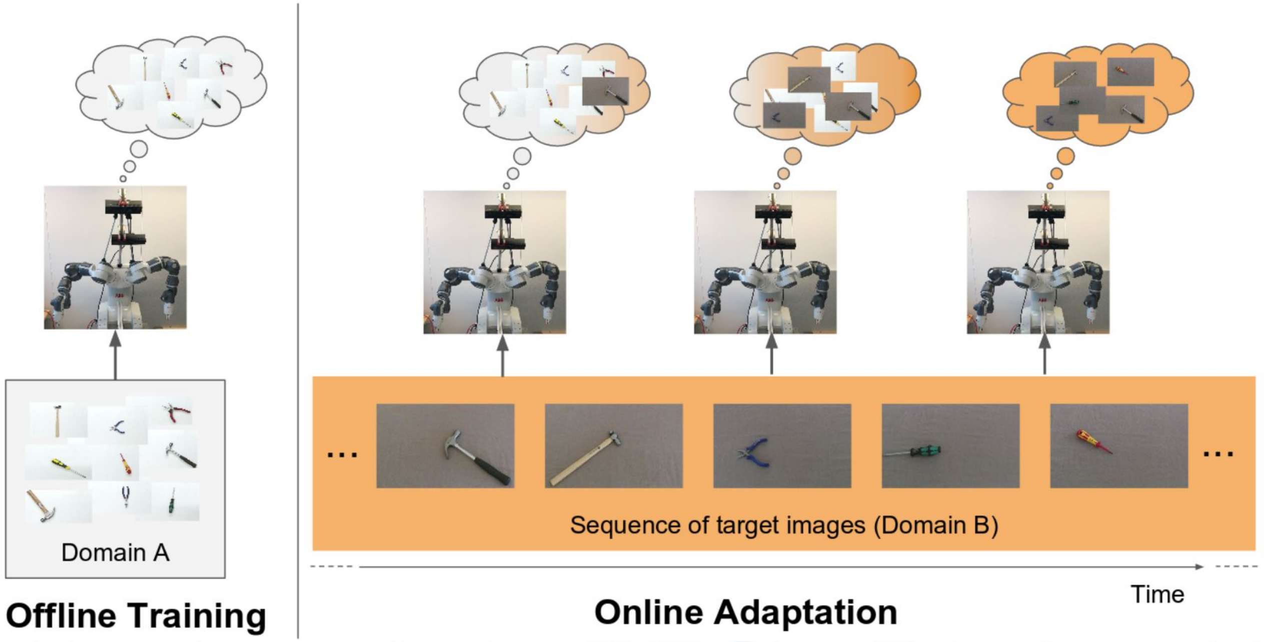

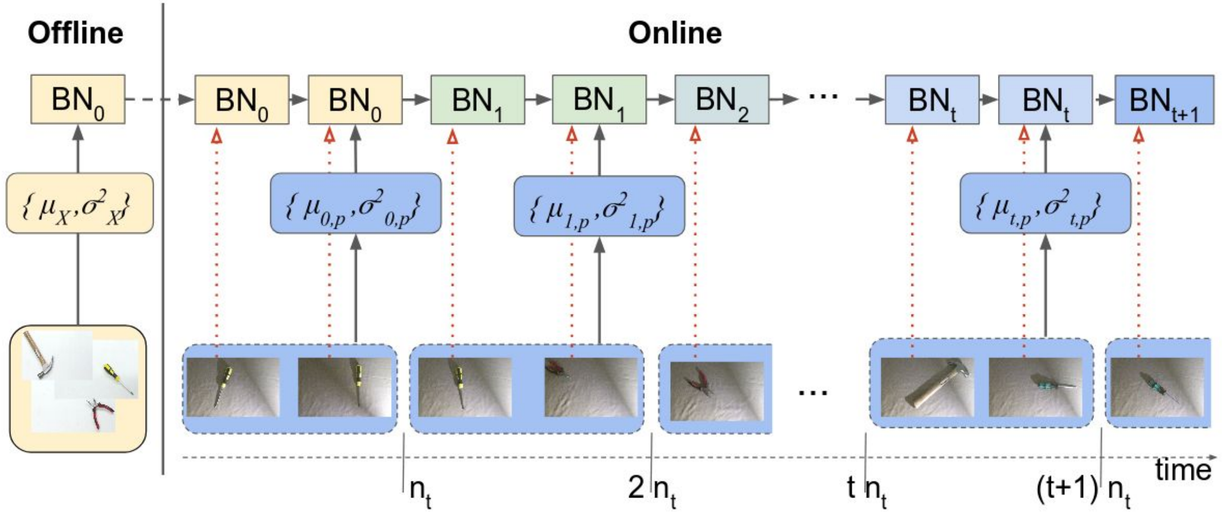

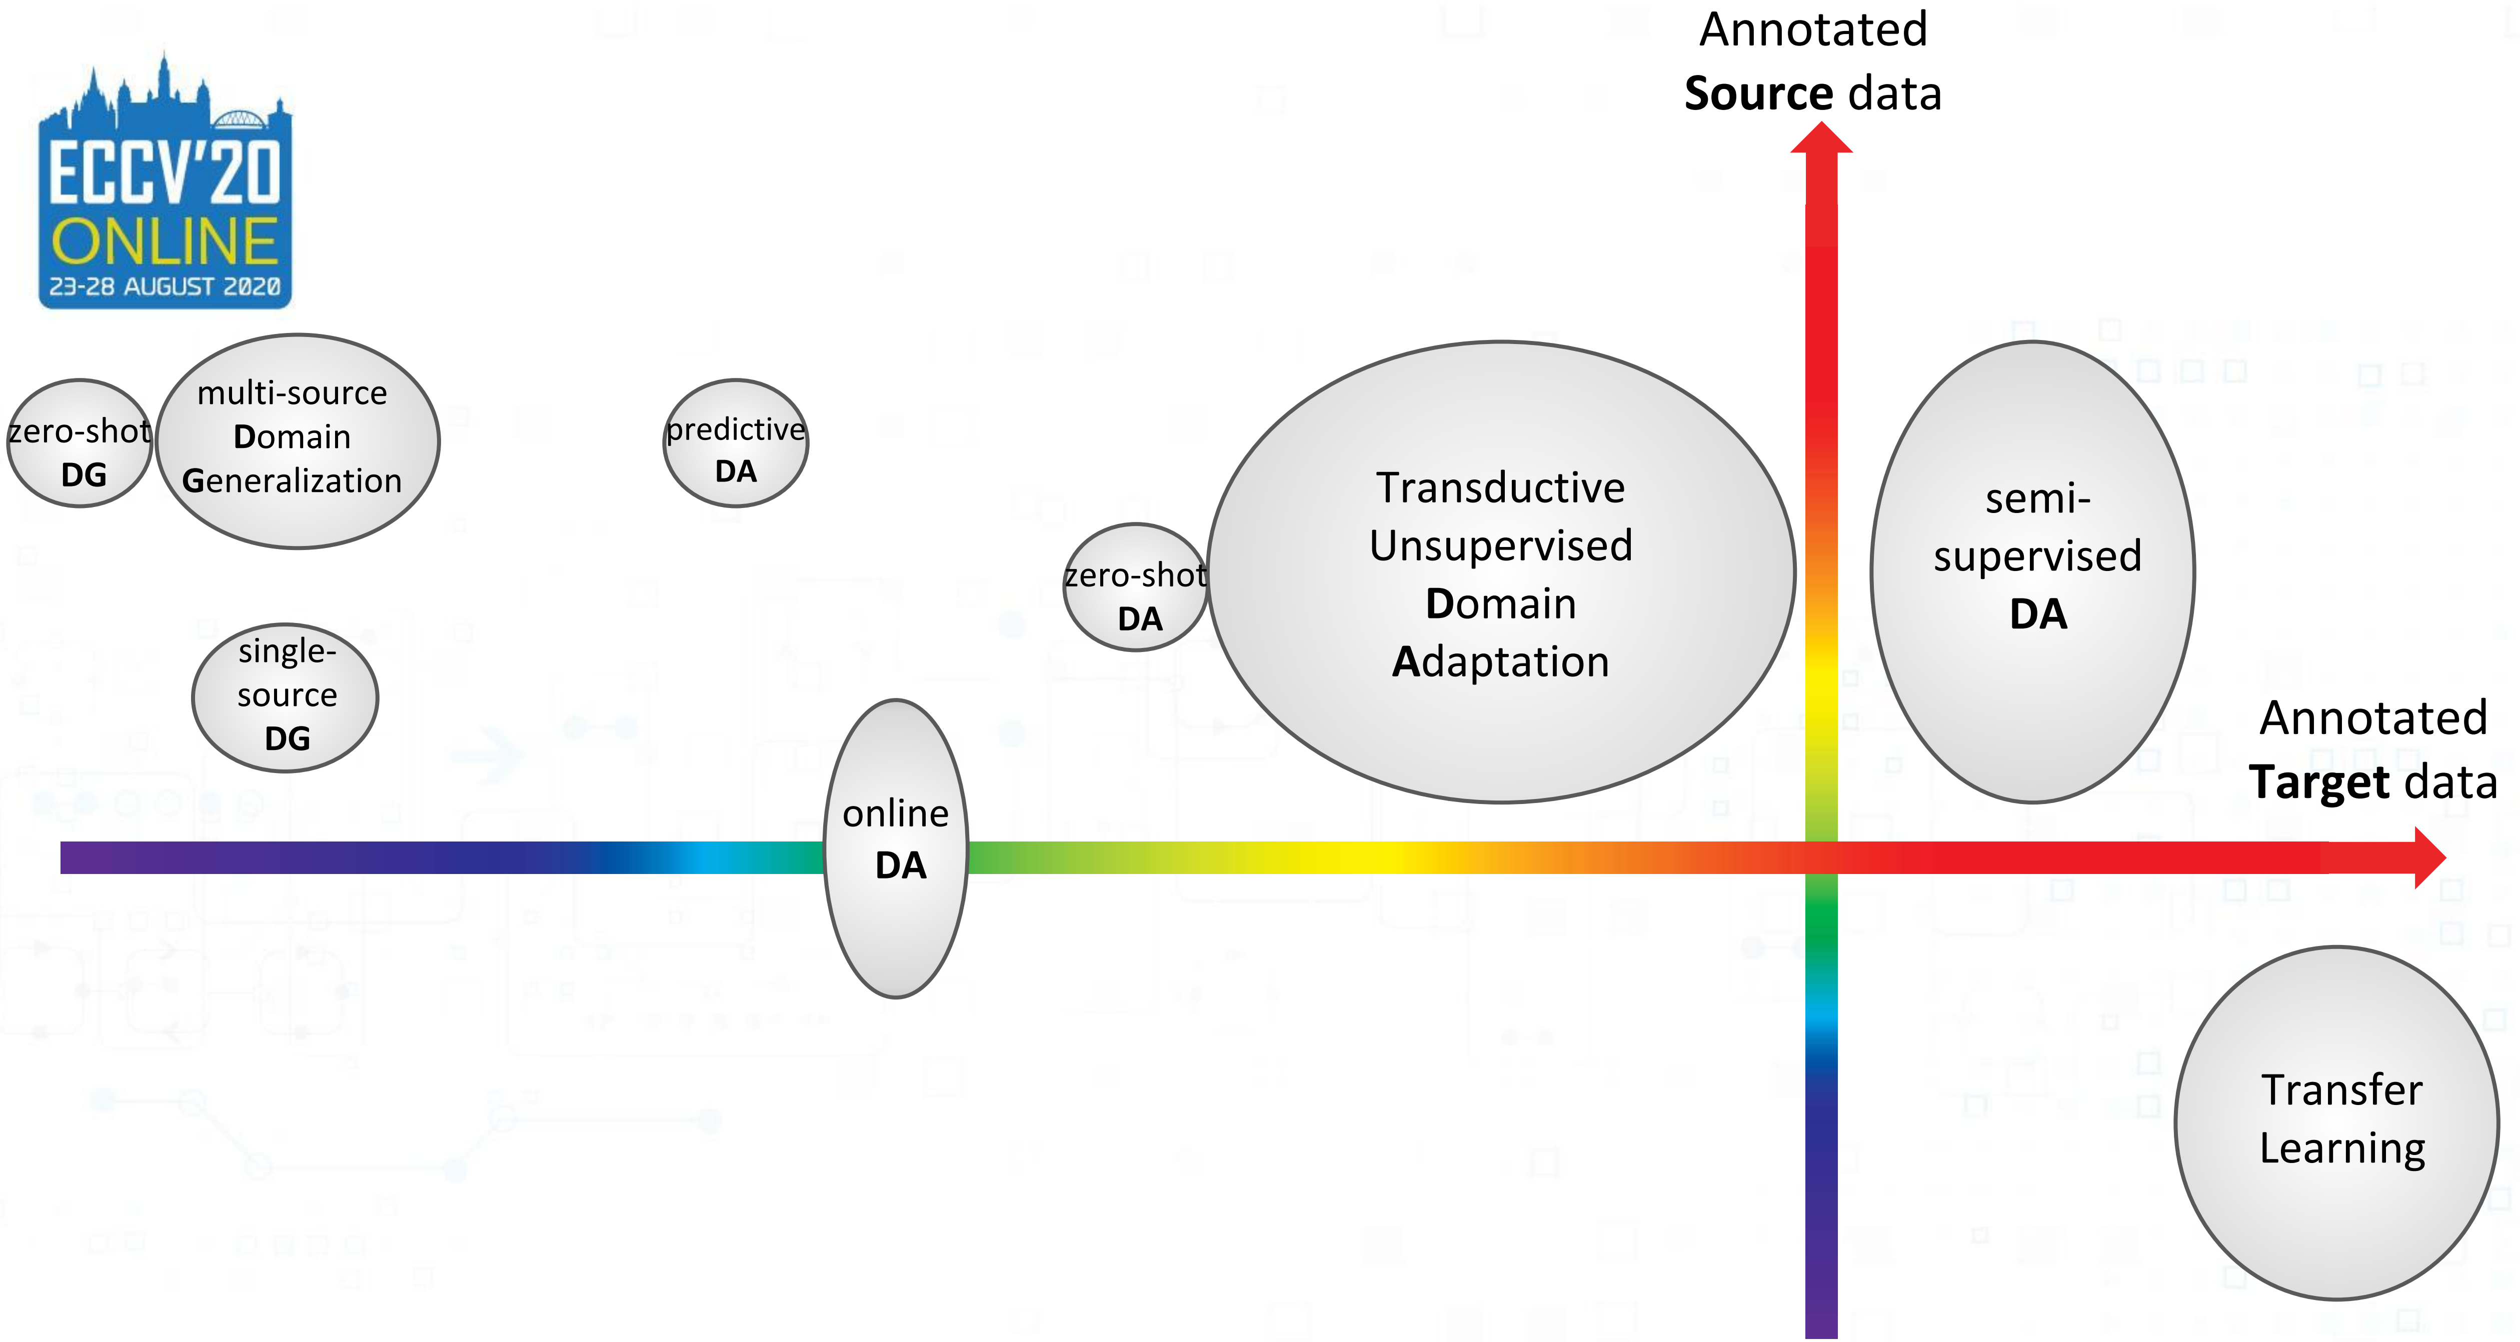
Annotated
Source data
Annotated
Target data
zero-shot
DG
multi-source
Domain
Generalization
single-
source
DG
predictive
DA
zero-shot
DA
online
DA
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
Transfer
Learning

ECCV'20 ONLINE 23-28 AUGUST 2020

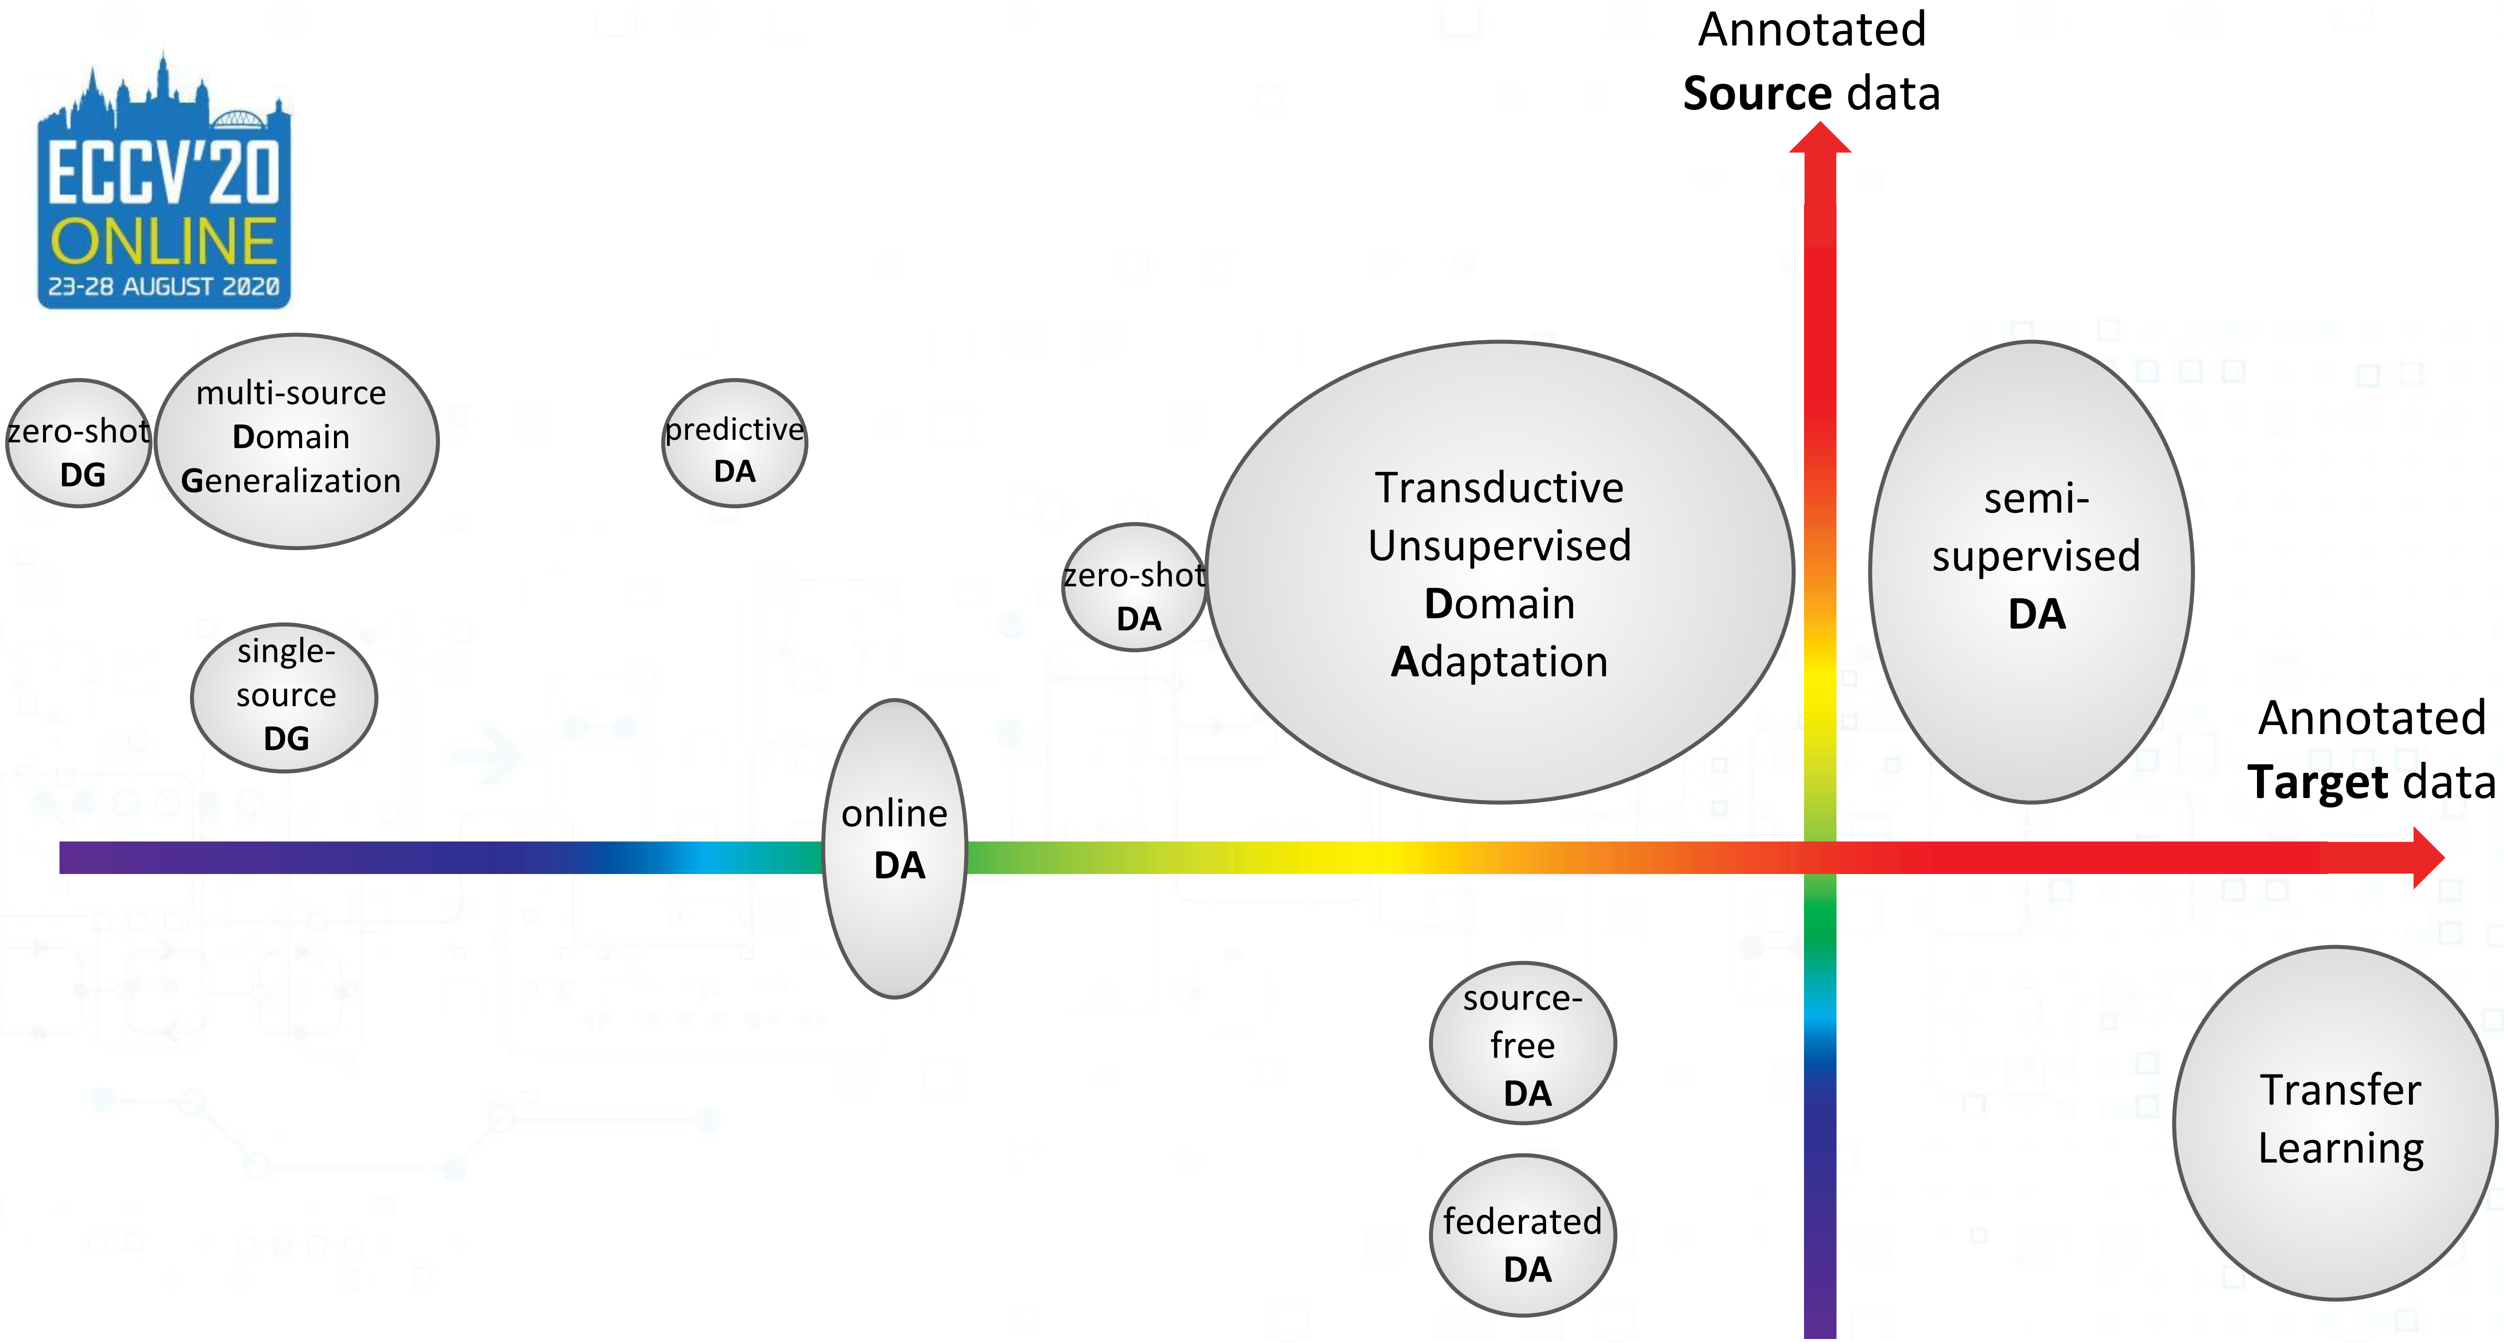

Source-Free DA

[Universal Source-Free Domain Adaptation, CVPR 2020]
[Do We Really Need to Access the Source Data? Source Hypothesis Transfer for Unsupervised Domain Adaptation, ICML 2020]

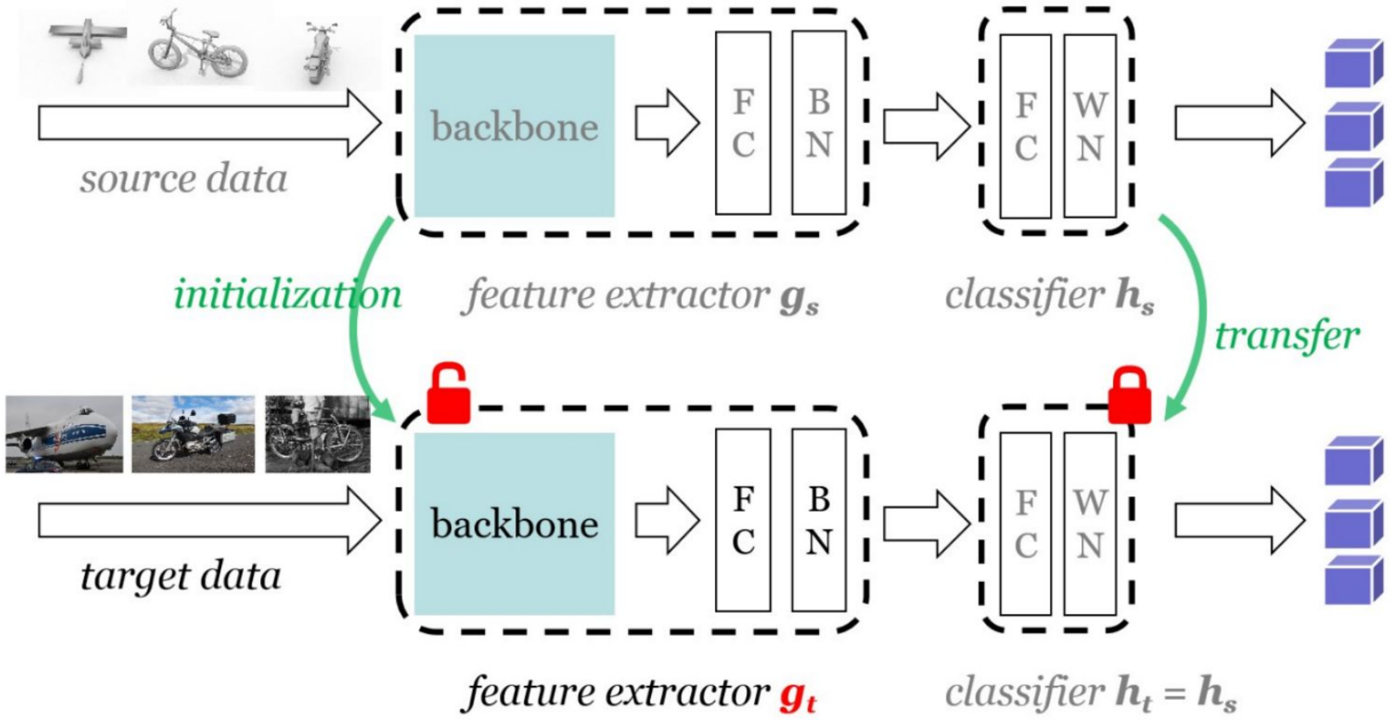

Cross Entropy Loss

Source-Free DA

[Universal Source-Free Domain Adaptation, CVPR 2020]
[Do We Really Need to Access the Source Data? Source Hypothesis Transfer for Unsupervised Domain Adaptation, ICML 2020]

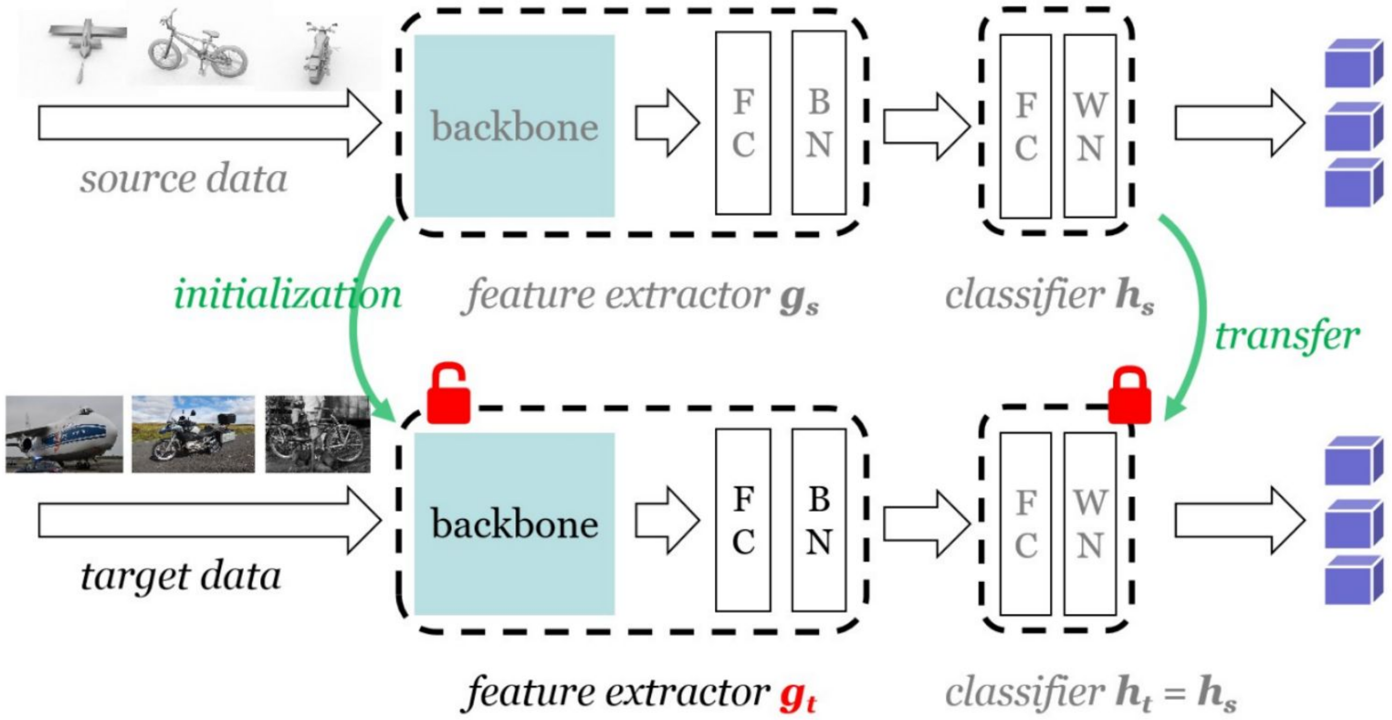

Cross Entropy Loss

(1) Information Maximization Loss
It makes the target outputs
- individually certain (as entropy minimization)
- globally diverse

(2) Target Clustering & Nearest Centroid Classifier

Federated DA

[Federated Adversarial Domain Adaptation, ICLR 2020]

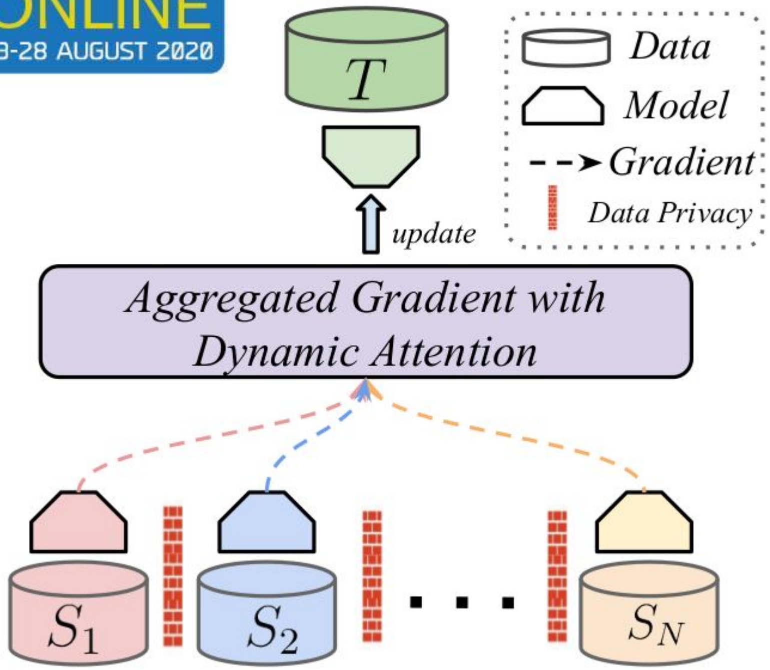

- the models on different nodes have different convergence rates

Federated DA

[Federated Adversarial Domain Adaptation, ICLR 2020]

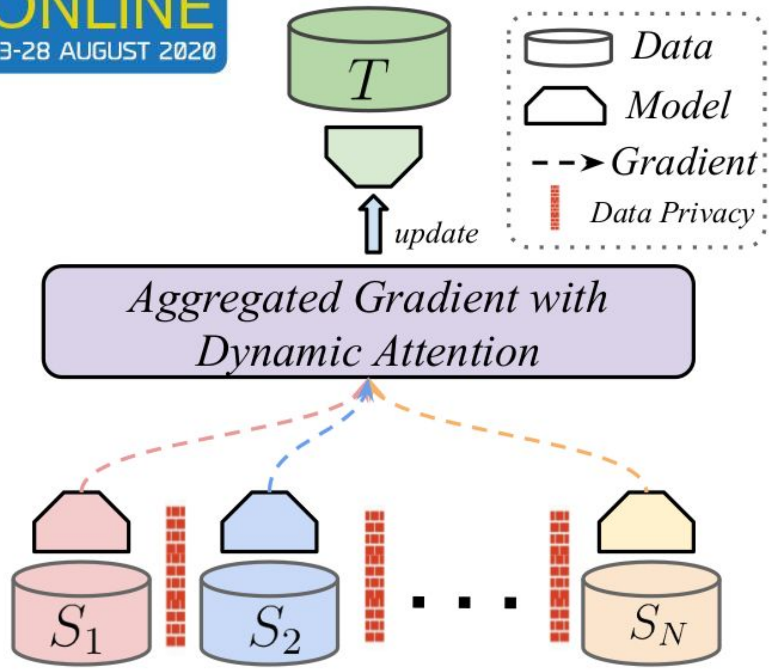

- the models on different nodes have different convergence rates
- increase the weight of those nodes whose gradients are **beneficial** to the target domain
- measure how well the target features can be **clustered** (cluster gap statistics gain)

Federated DA

[Federated Adversarial Domain Adaptation, ICLR 2020]

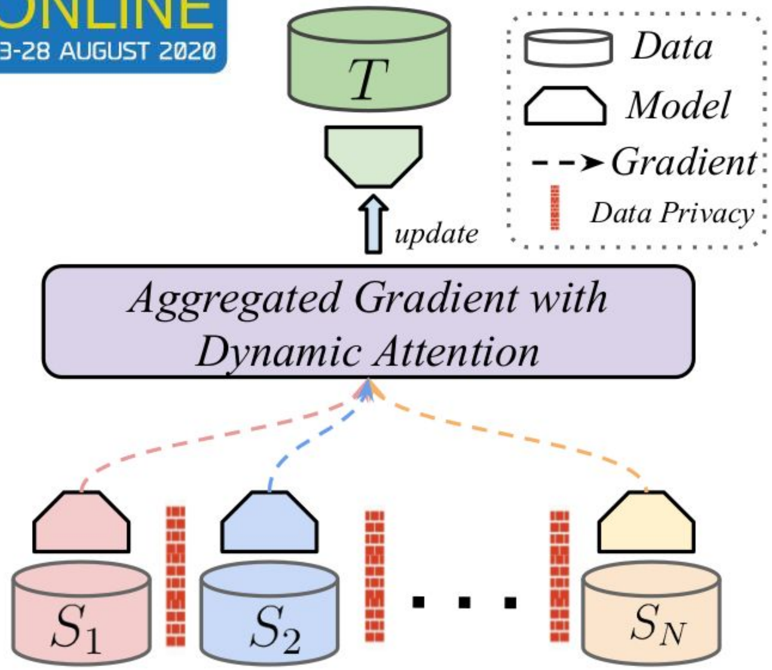

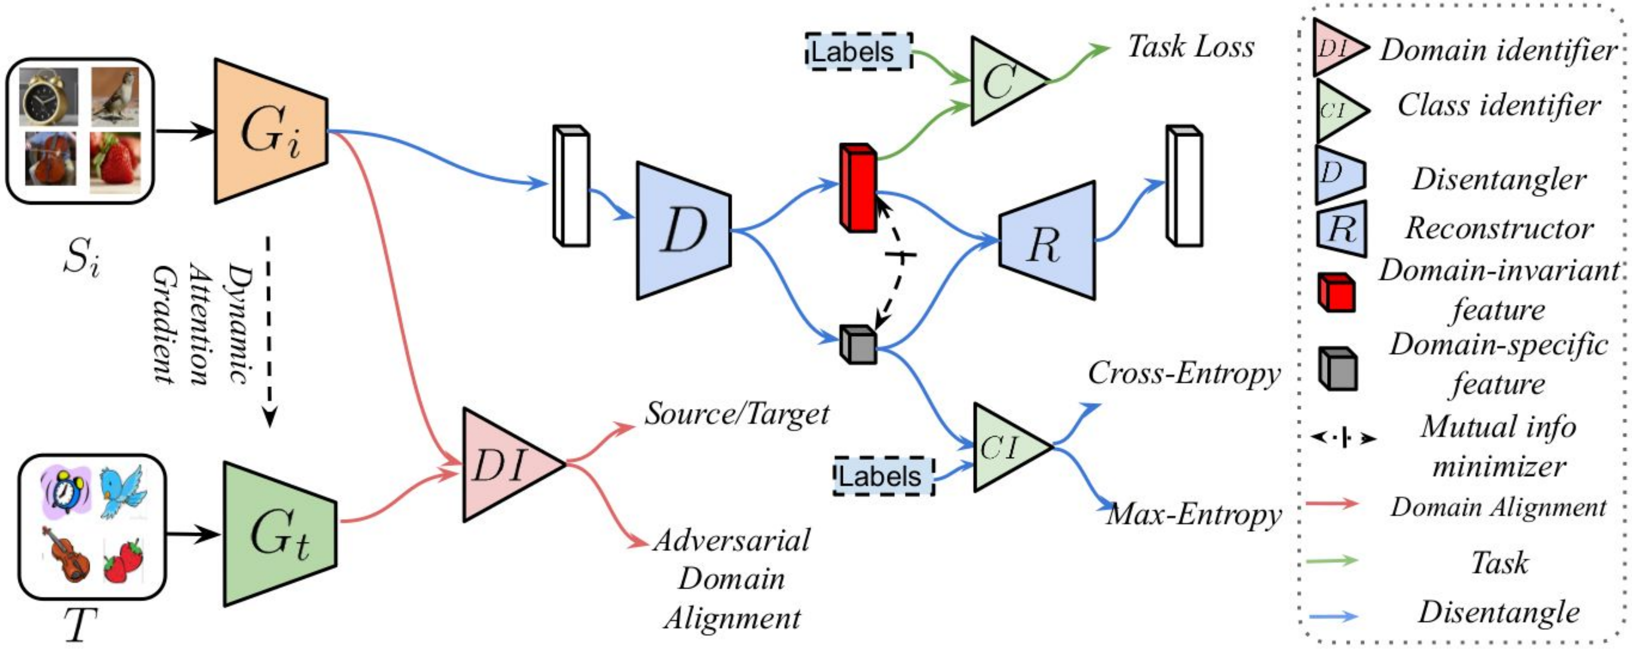

- the models on different nodes have different convergence rates
- increase the weight of those nodes whose gradients are **beneficial** to the target domain
- measure how well the target features can be **clustered** (cluster gap statistics gain)

Federated DA

[Federated Adversarial Domain Adaptation, ICLR 2020]

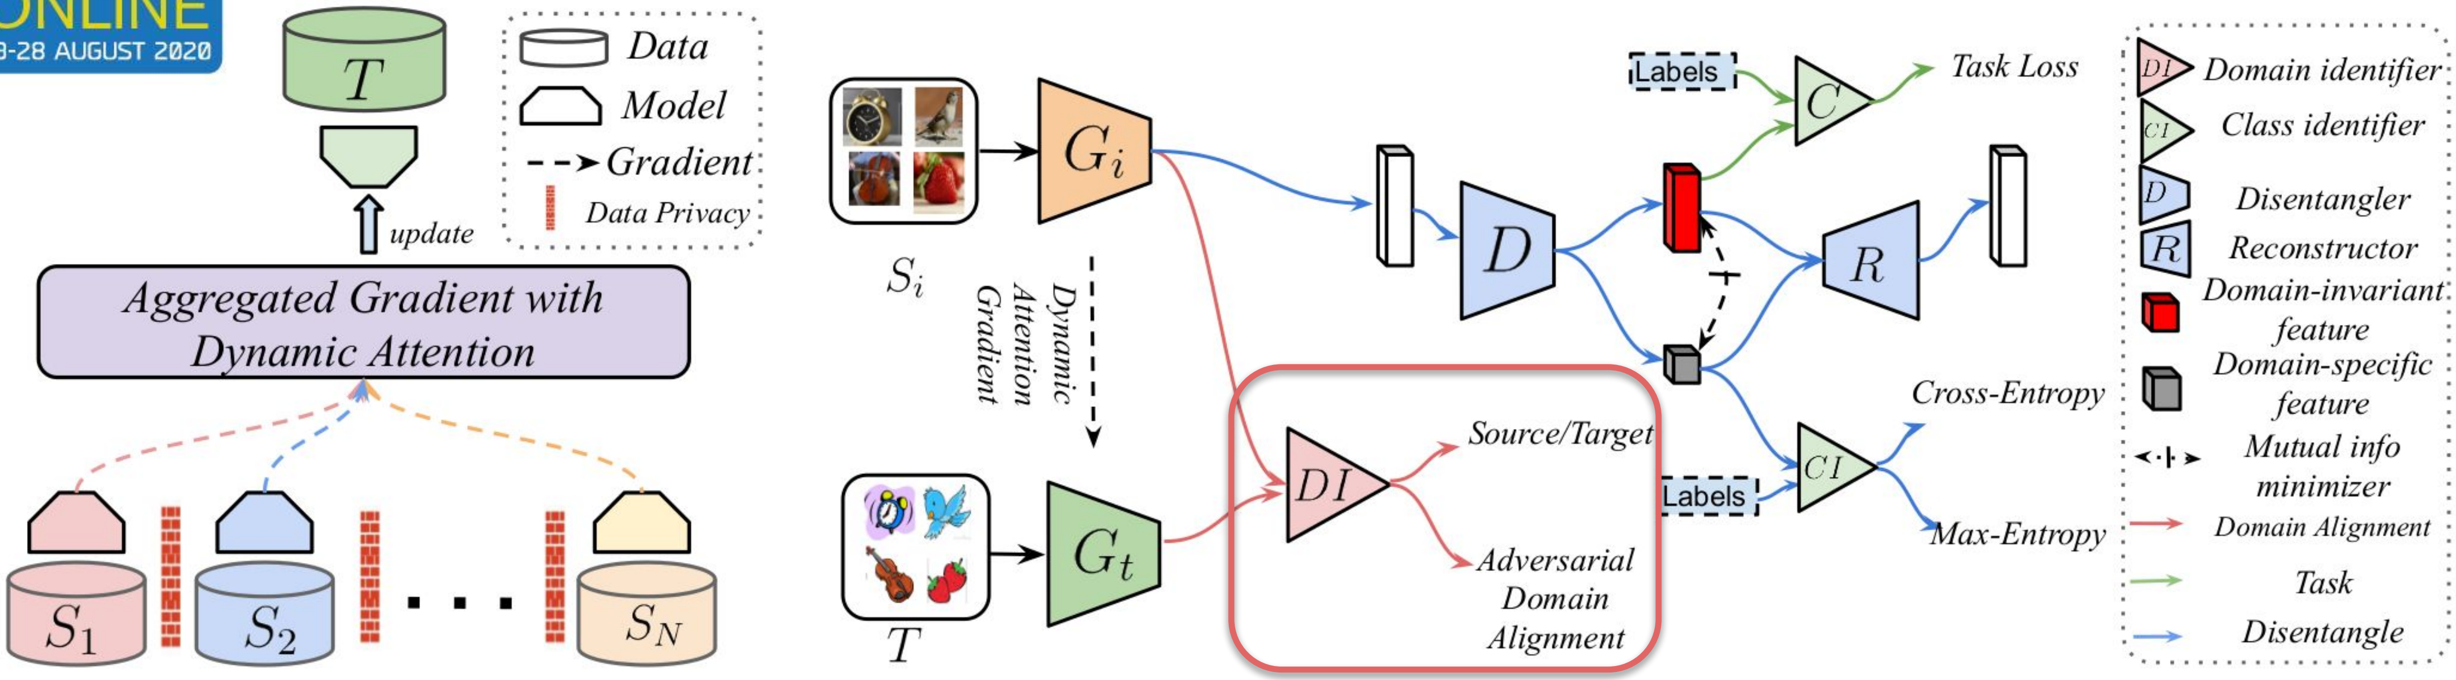

- the models on different nodes have different convergence rates
- increase the weight of those nodes whose gradients are **beneficial** to the target domain
- measure how well the target features can be **clustered** (cluster gap statistics gain)

Federated DA

[Federated Adversarial Domain Adaptation, ICLR 2020]

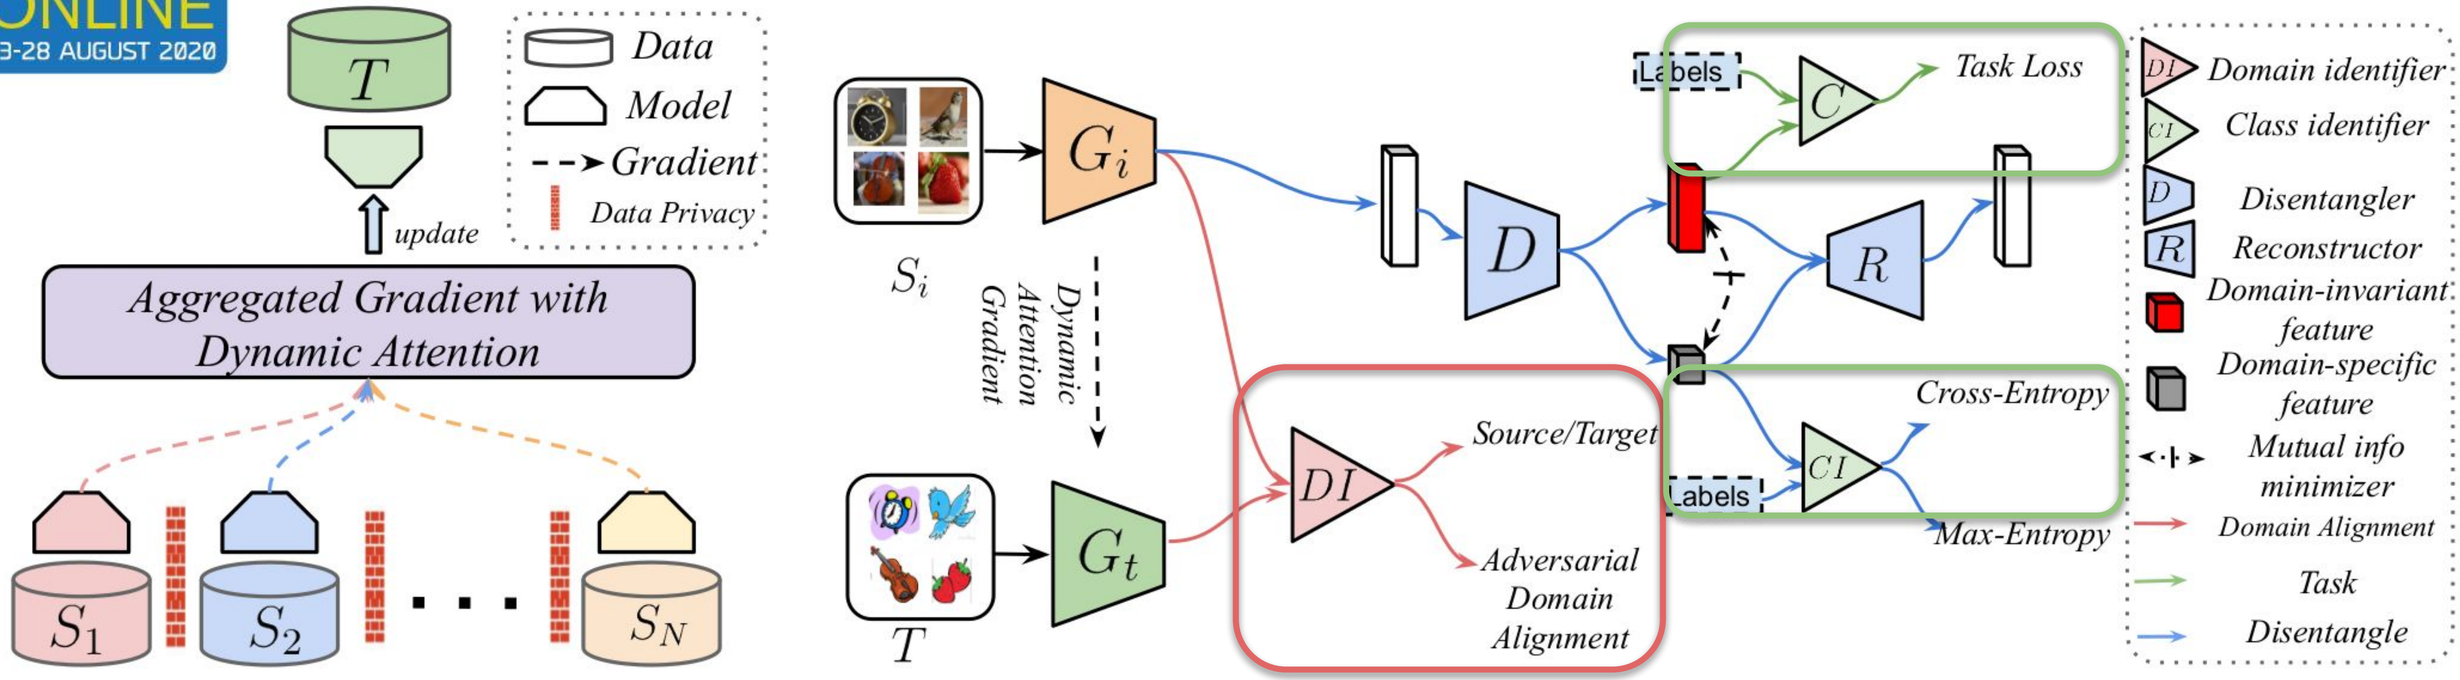

- the models on different nodes have different convergence rates
- increase the weight of those nodes whose gradients are **beneficial** to the target domain
- measure how well the target features can be **clustered** (cluster gap statistics gain)

Federated DA

[Federated Adversarial Domain Adaptation, ICLR 2020]

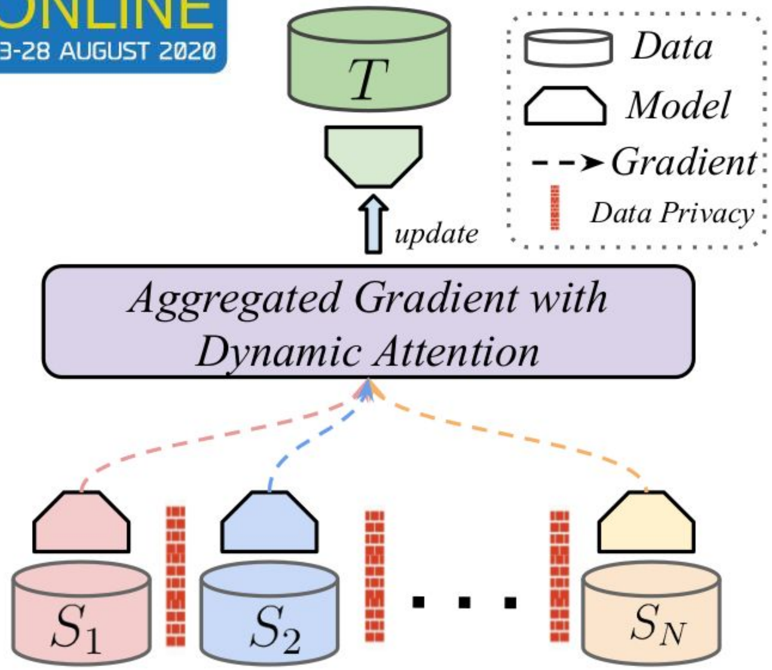

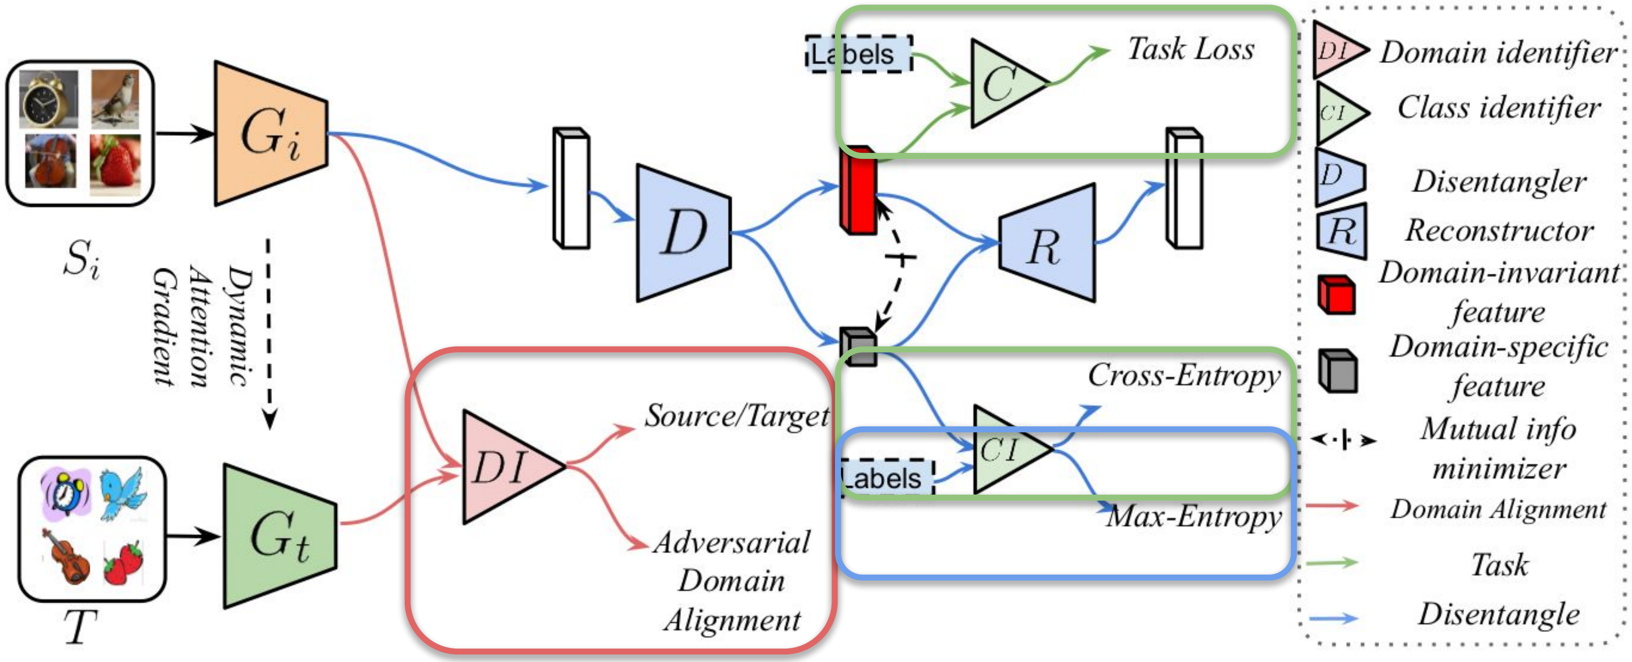

- the models on different nodes have different convergence rates
- increase the weight of those nodes whose gradients are **beneficial** to the target domain
- measure how well the target features can be **clustered** (cluster gap statistics gain)

Federated DA

[Federated Adversarial Domain Adaptation, ICLR 2020]

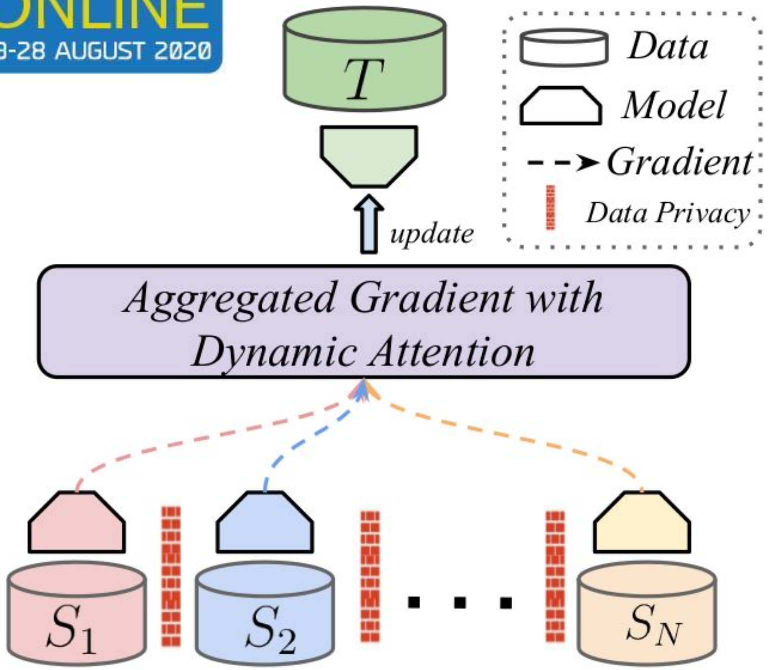

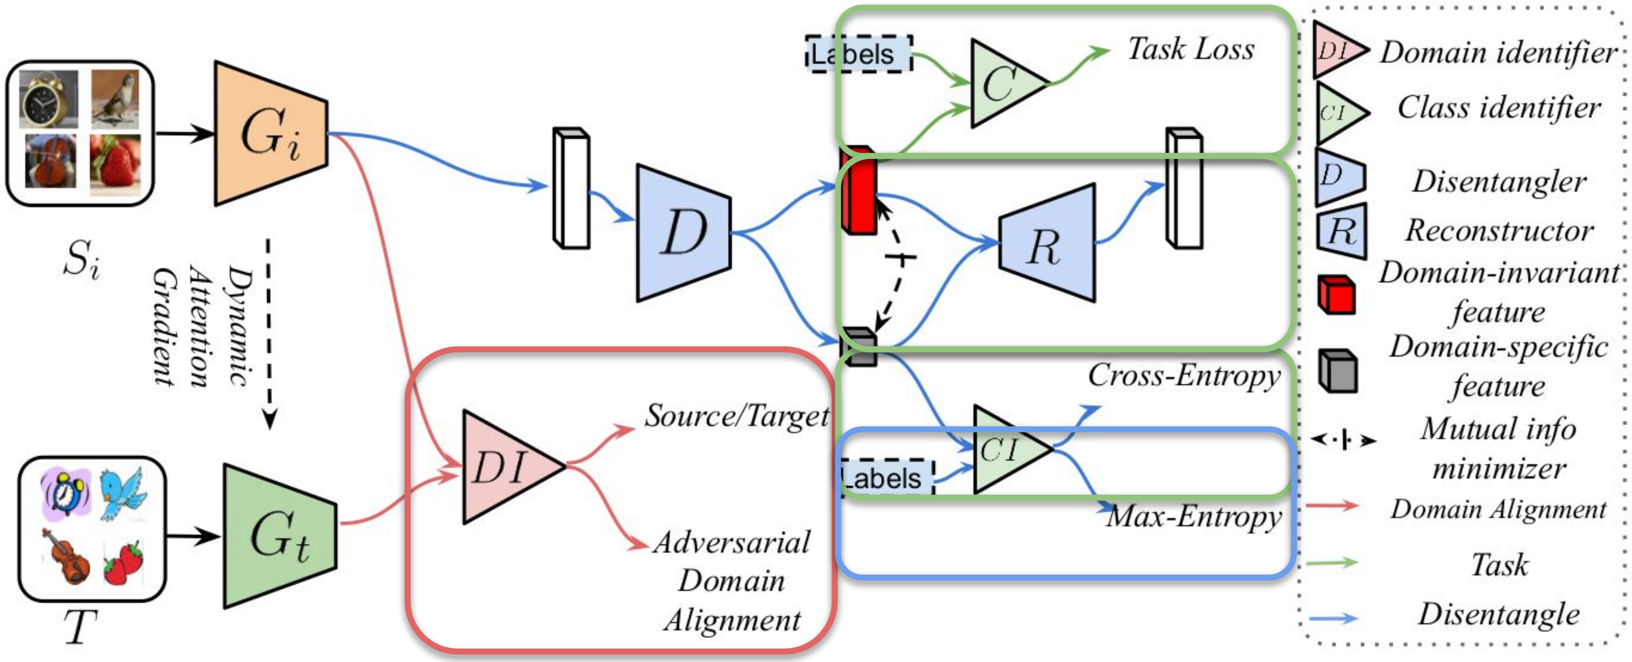

- the models on different nodes have different convergence rates
- increase the weight of those nodes whose gradients are **beneficial** to the target domain
- measure how well the target features can be **clustered** (cluster gap statistics gain)

Federated DA

[Federated Adversarial Domain Adaptation, ICLR 2020]

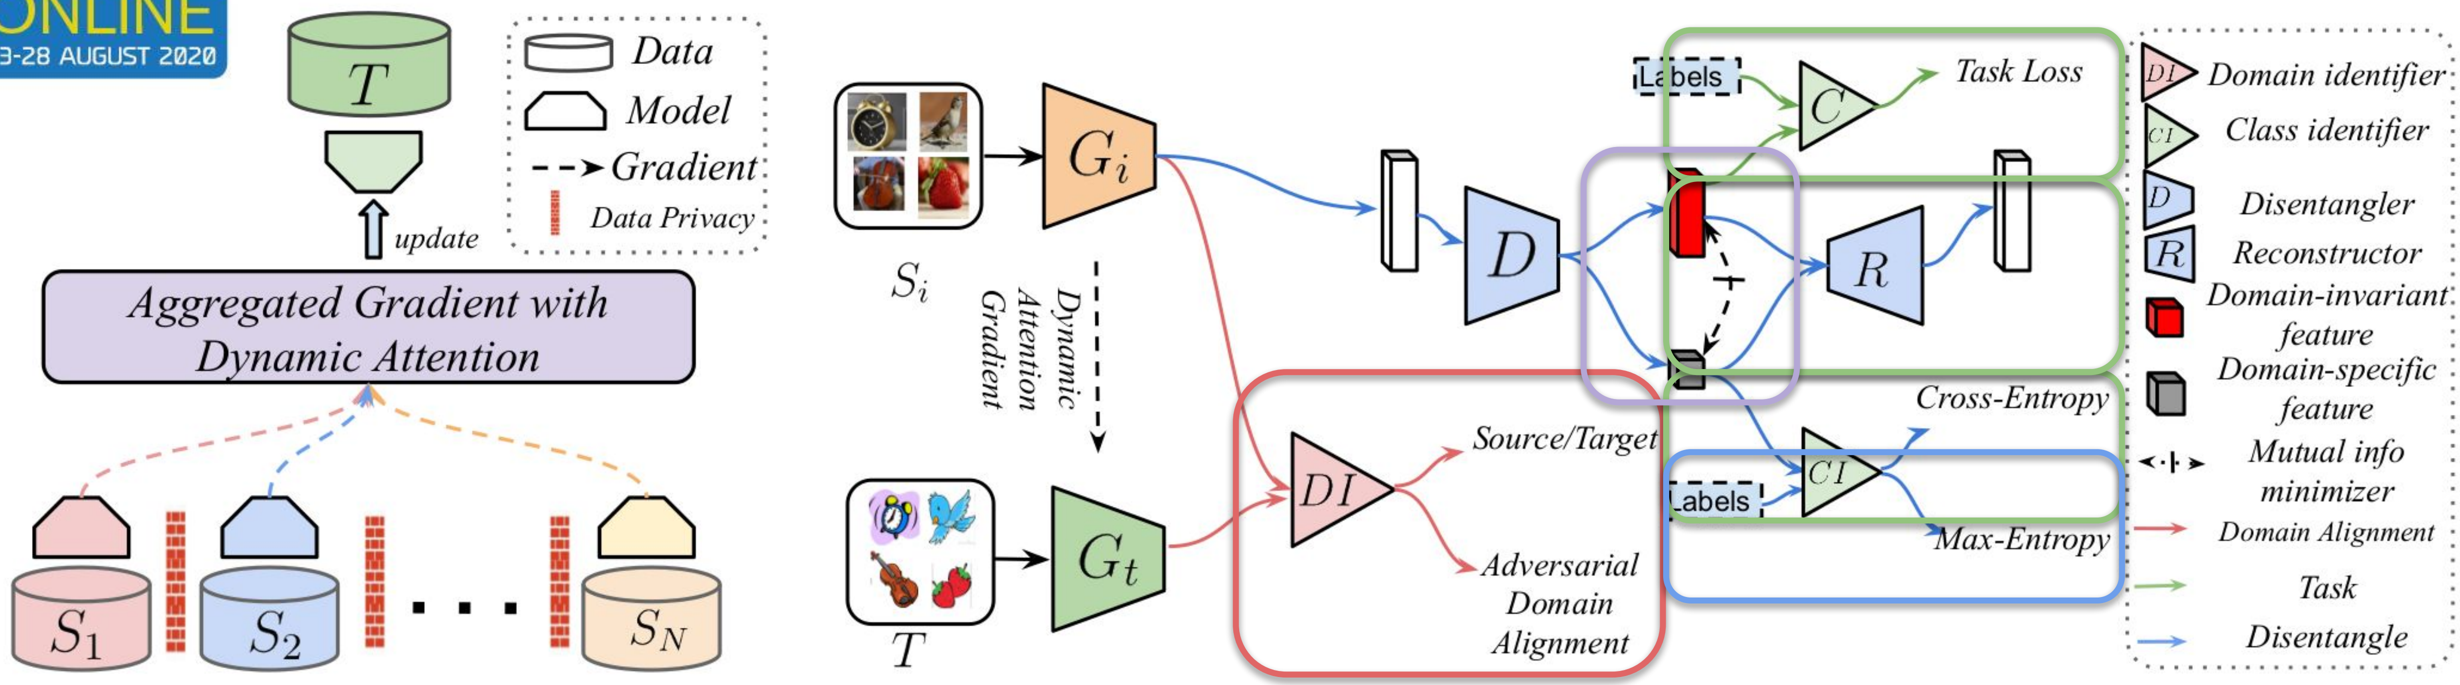

- the models on different nodes have different convergence rates
- increase the weight of those nodes whose gradients are **beneficial** to the target domain
- measure how well the target features can be **clustered** (cluster gap statistics gain)

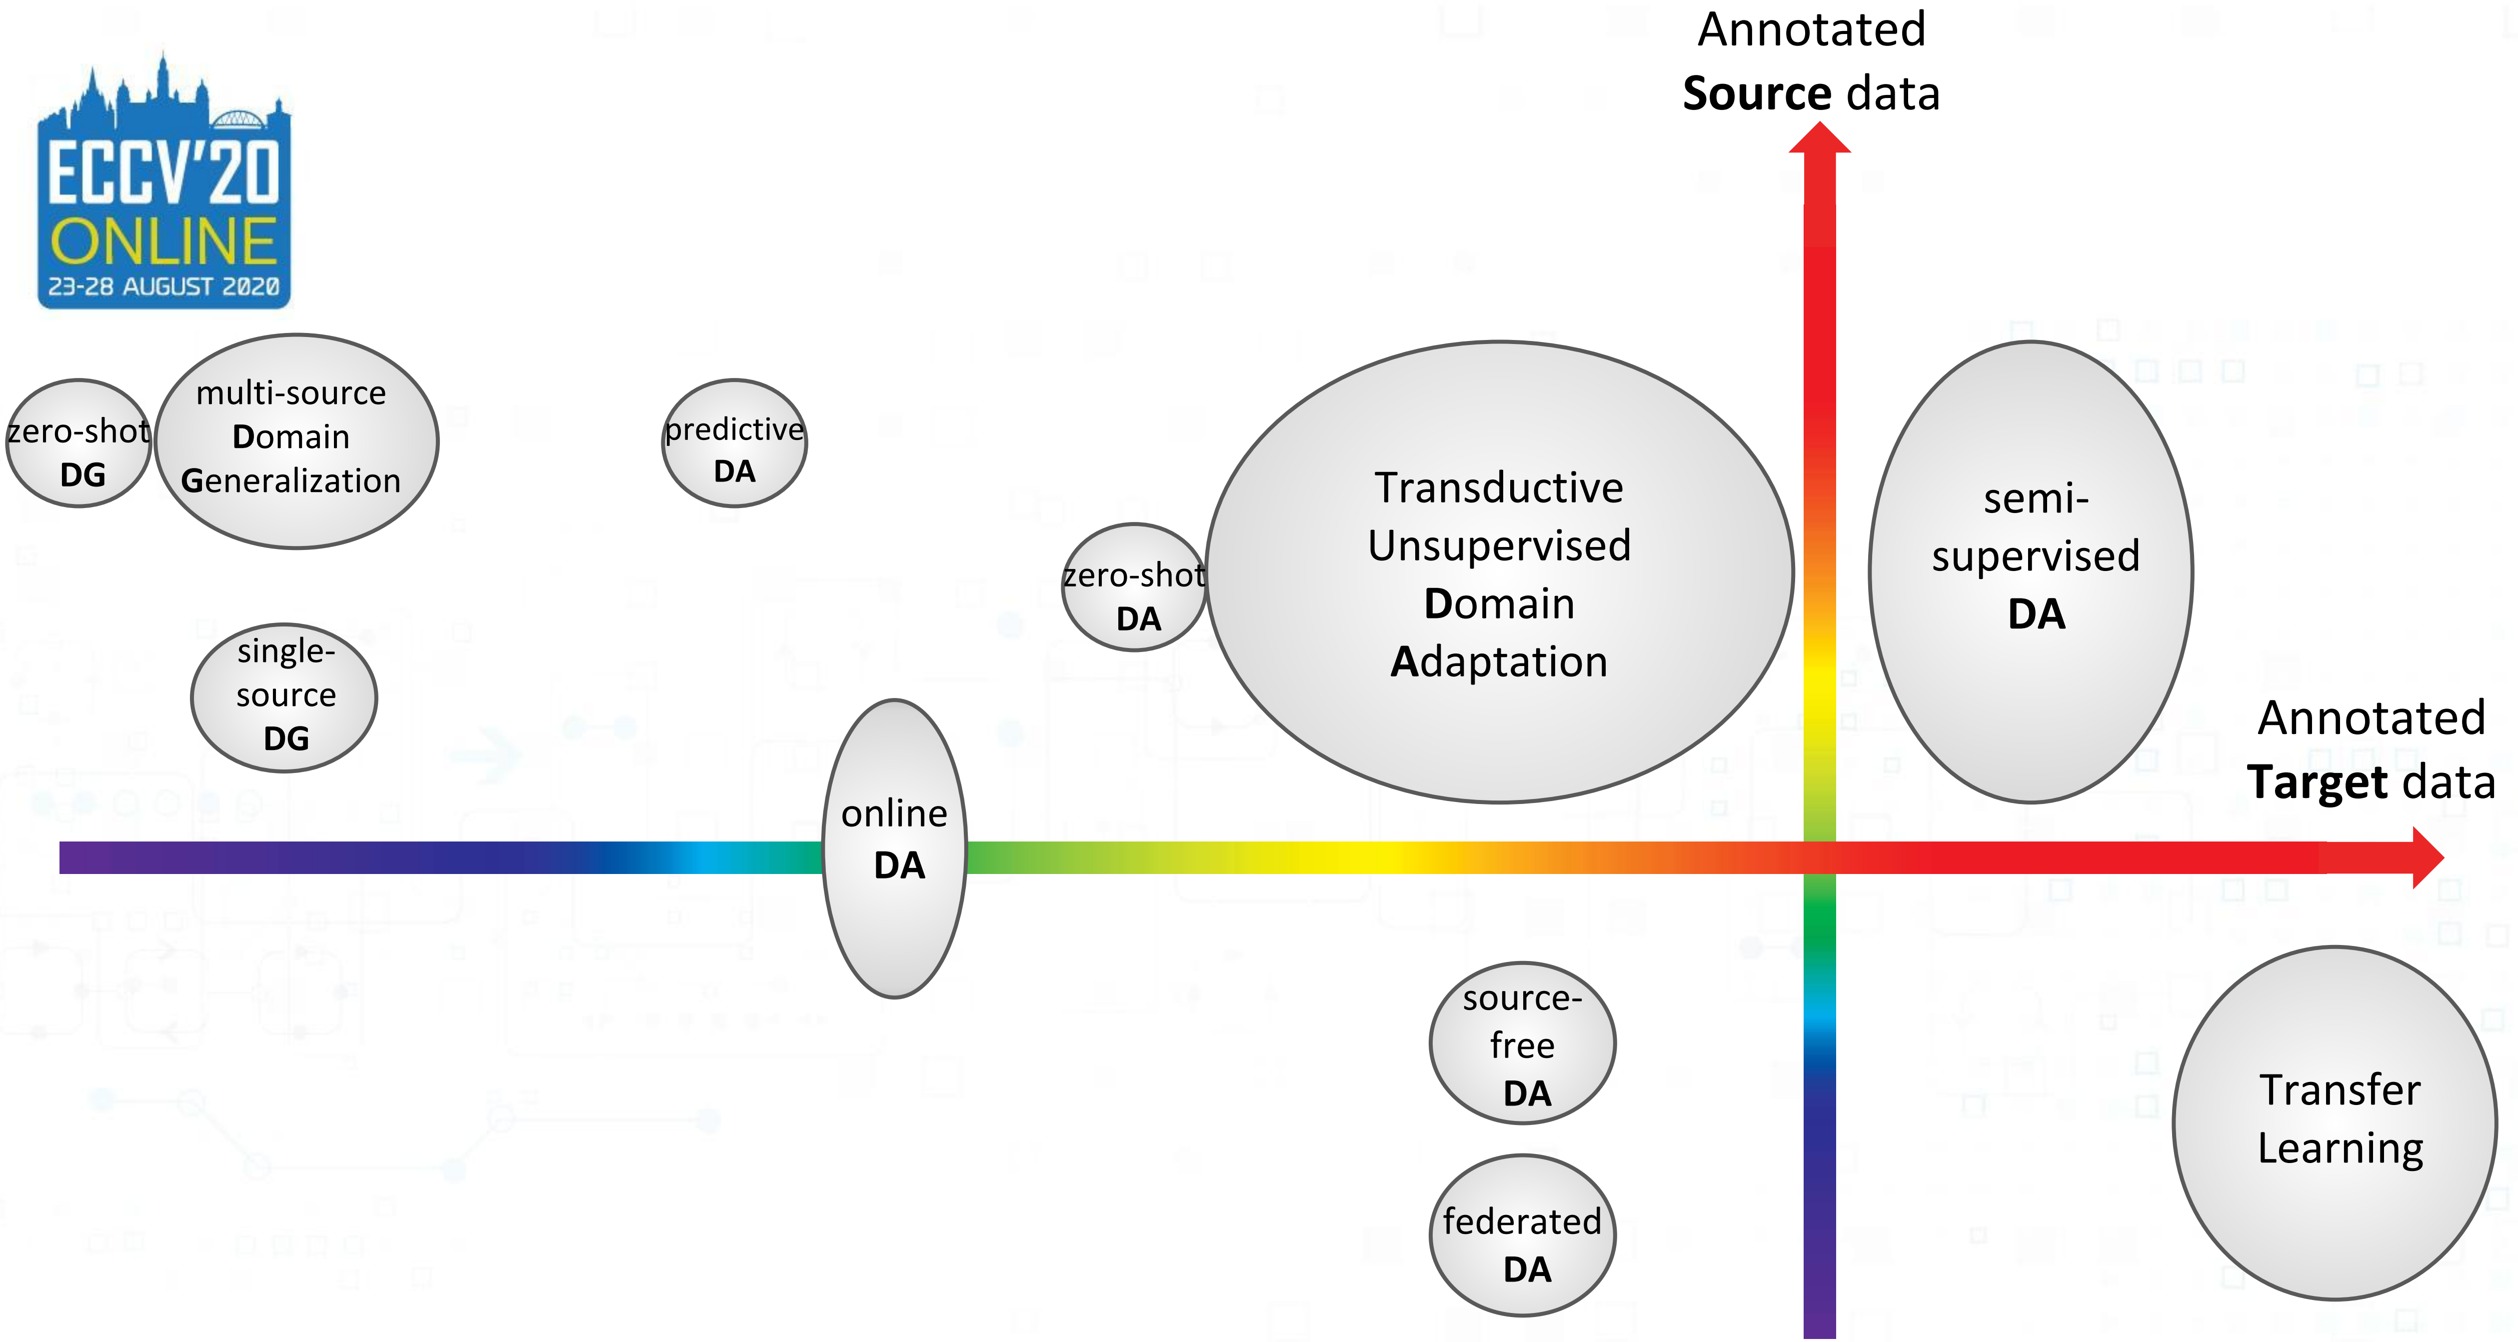

Annotated
Source data
Annotated
Target data
zero-shot
DG
multi-source
Domain
Generalization
predictive
DA
zero-shot
DA
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
single-
source
DG
online
DA
source-
free
DA
federated
DA
Transfer
Learning

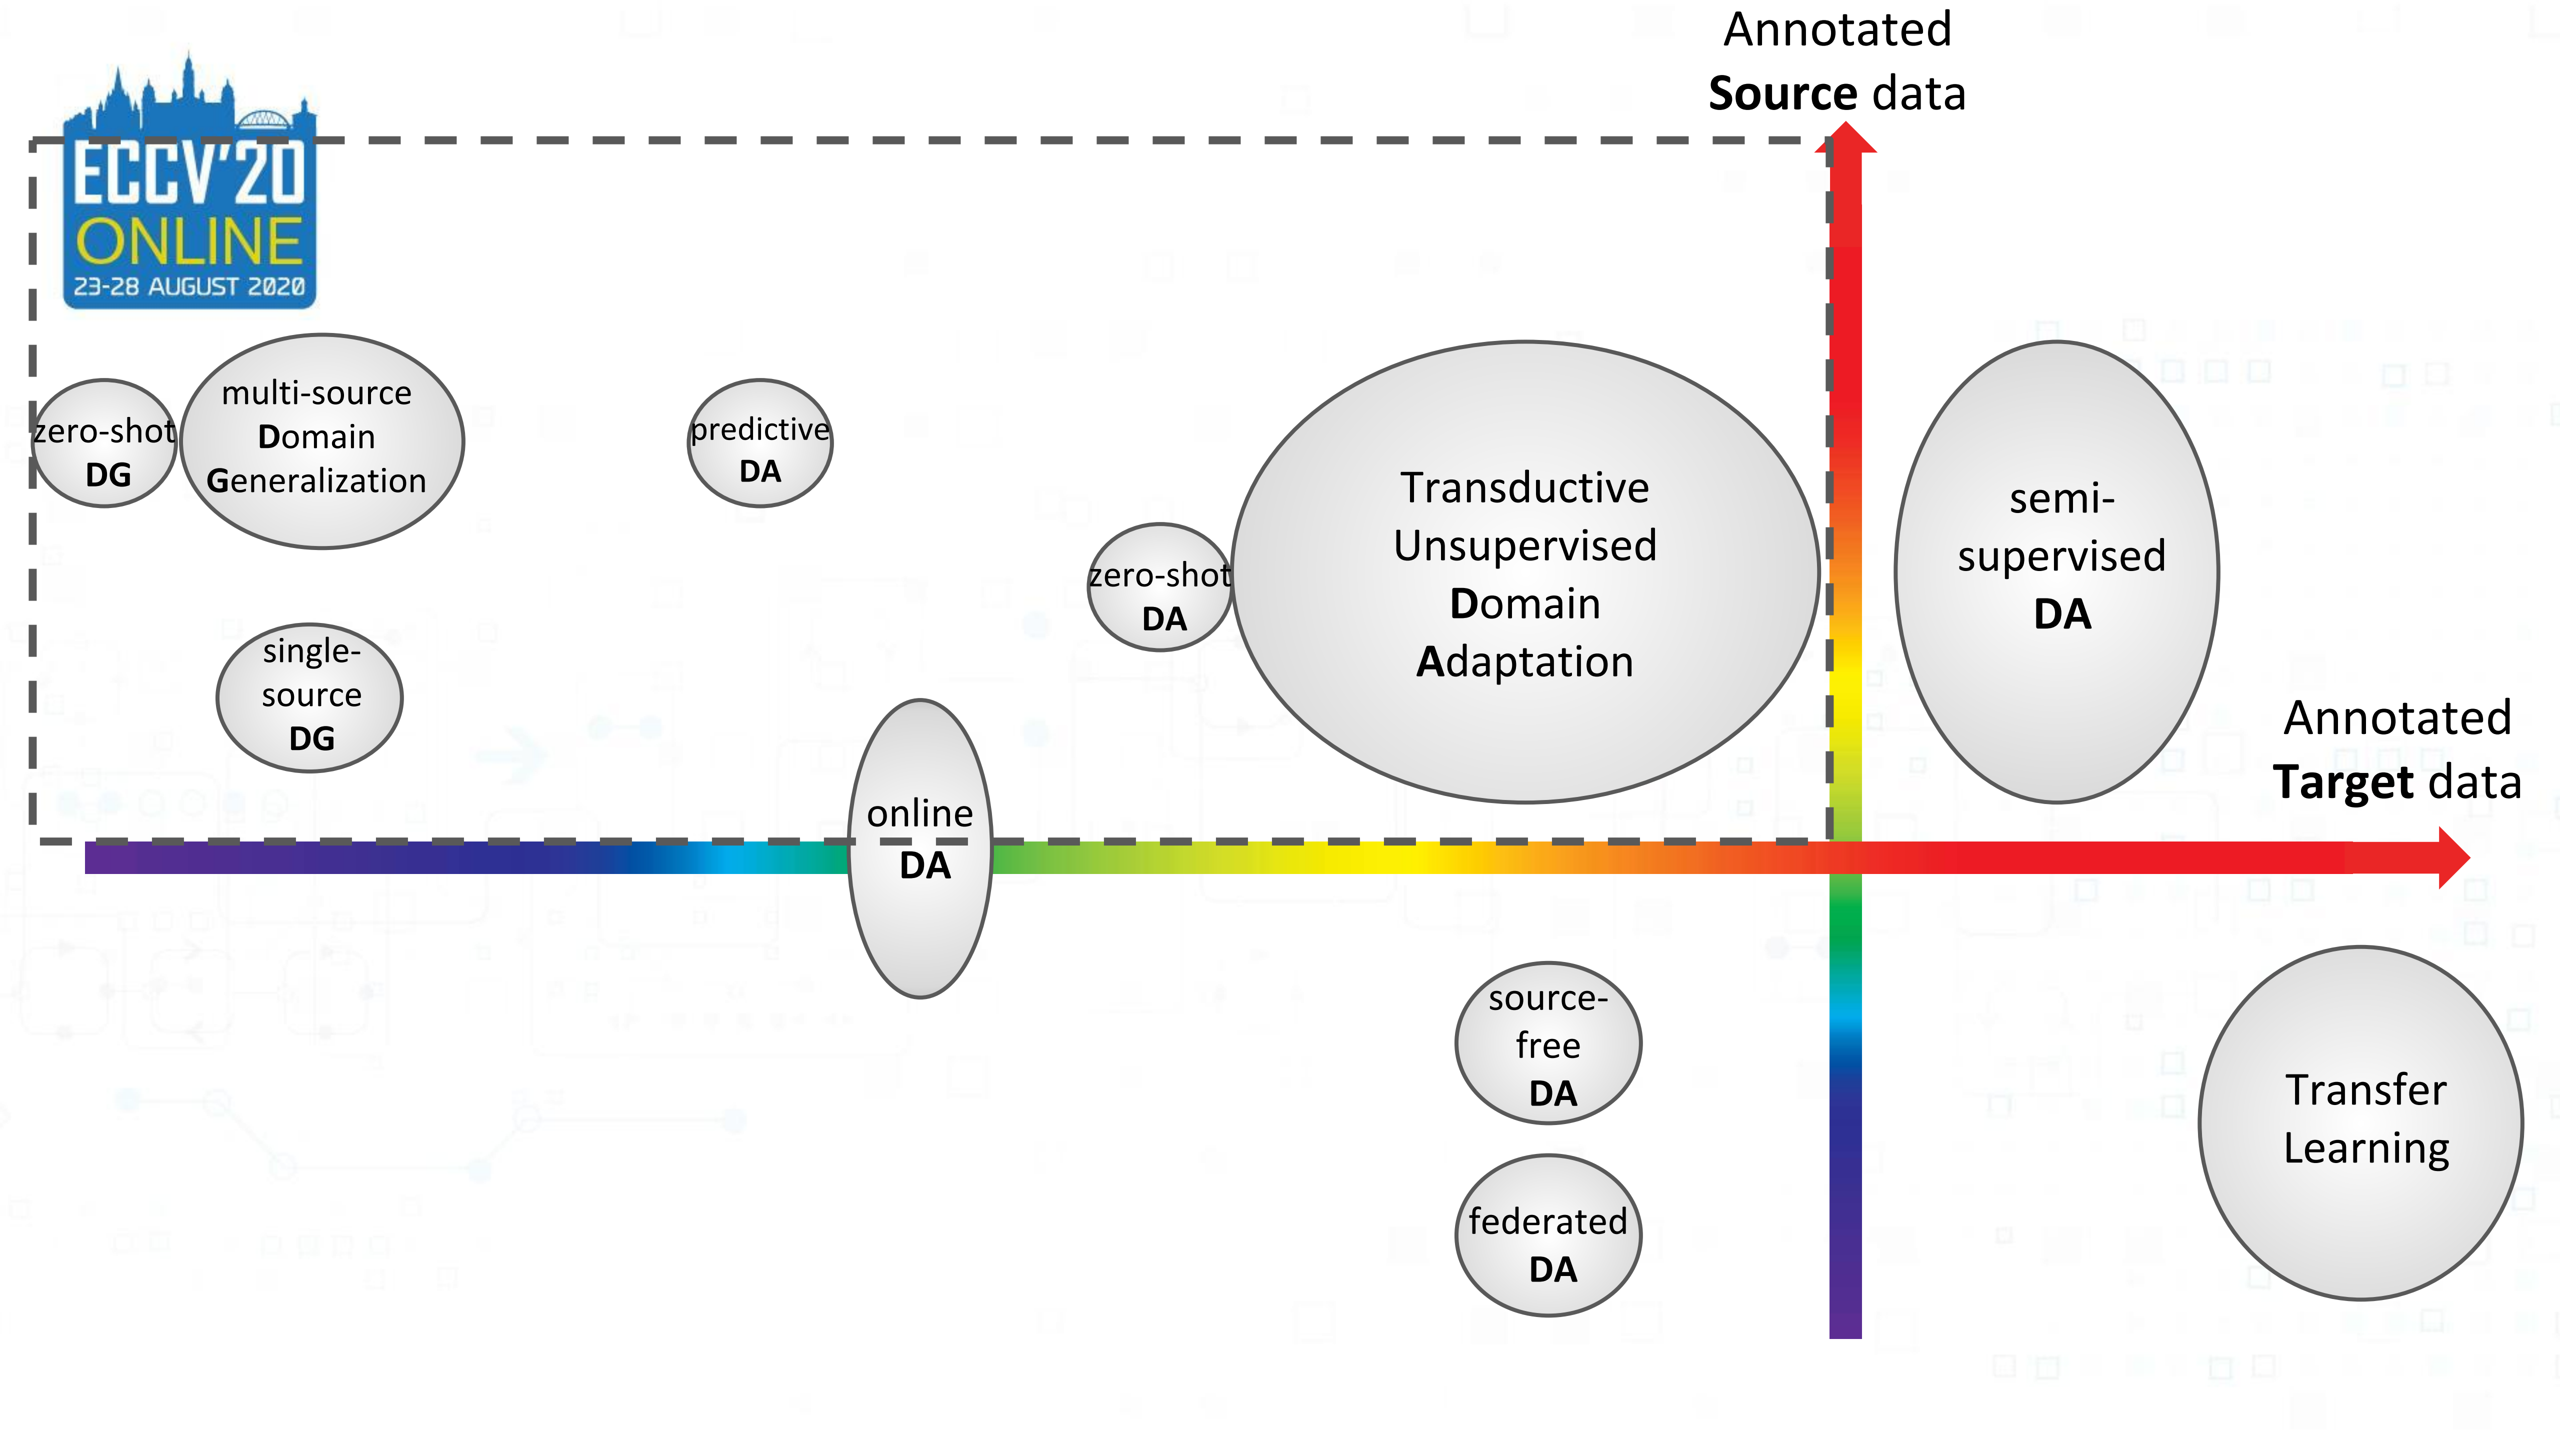

ECCV'20
ONLINE
23-28 AUGUST 2020
Annotated
Source data
Annotated
Target data
zero-shot
DG
multi-source
Domain
Generalization
predictive
DA
single-
source
DG
zero-shot
DA
Transductive
Unsupervised
Domain
Adaptation
semi-
supervised
DA
online
DA
source-
free
DA
federated
DA
Transfer
Learning

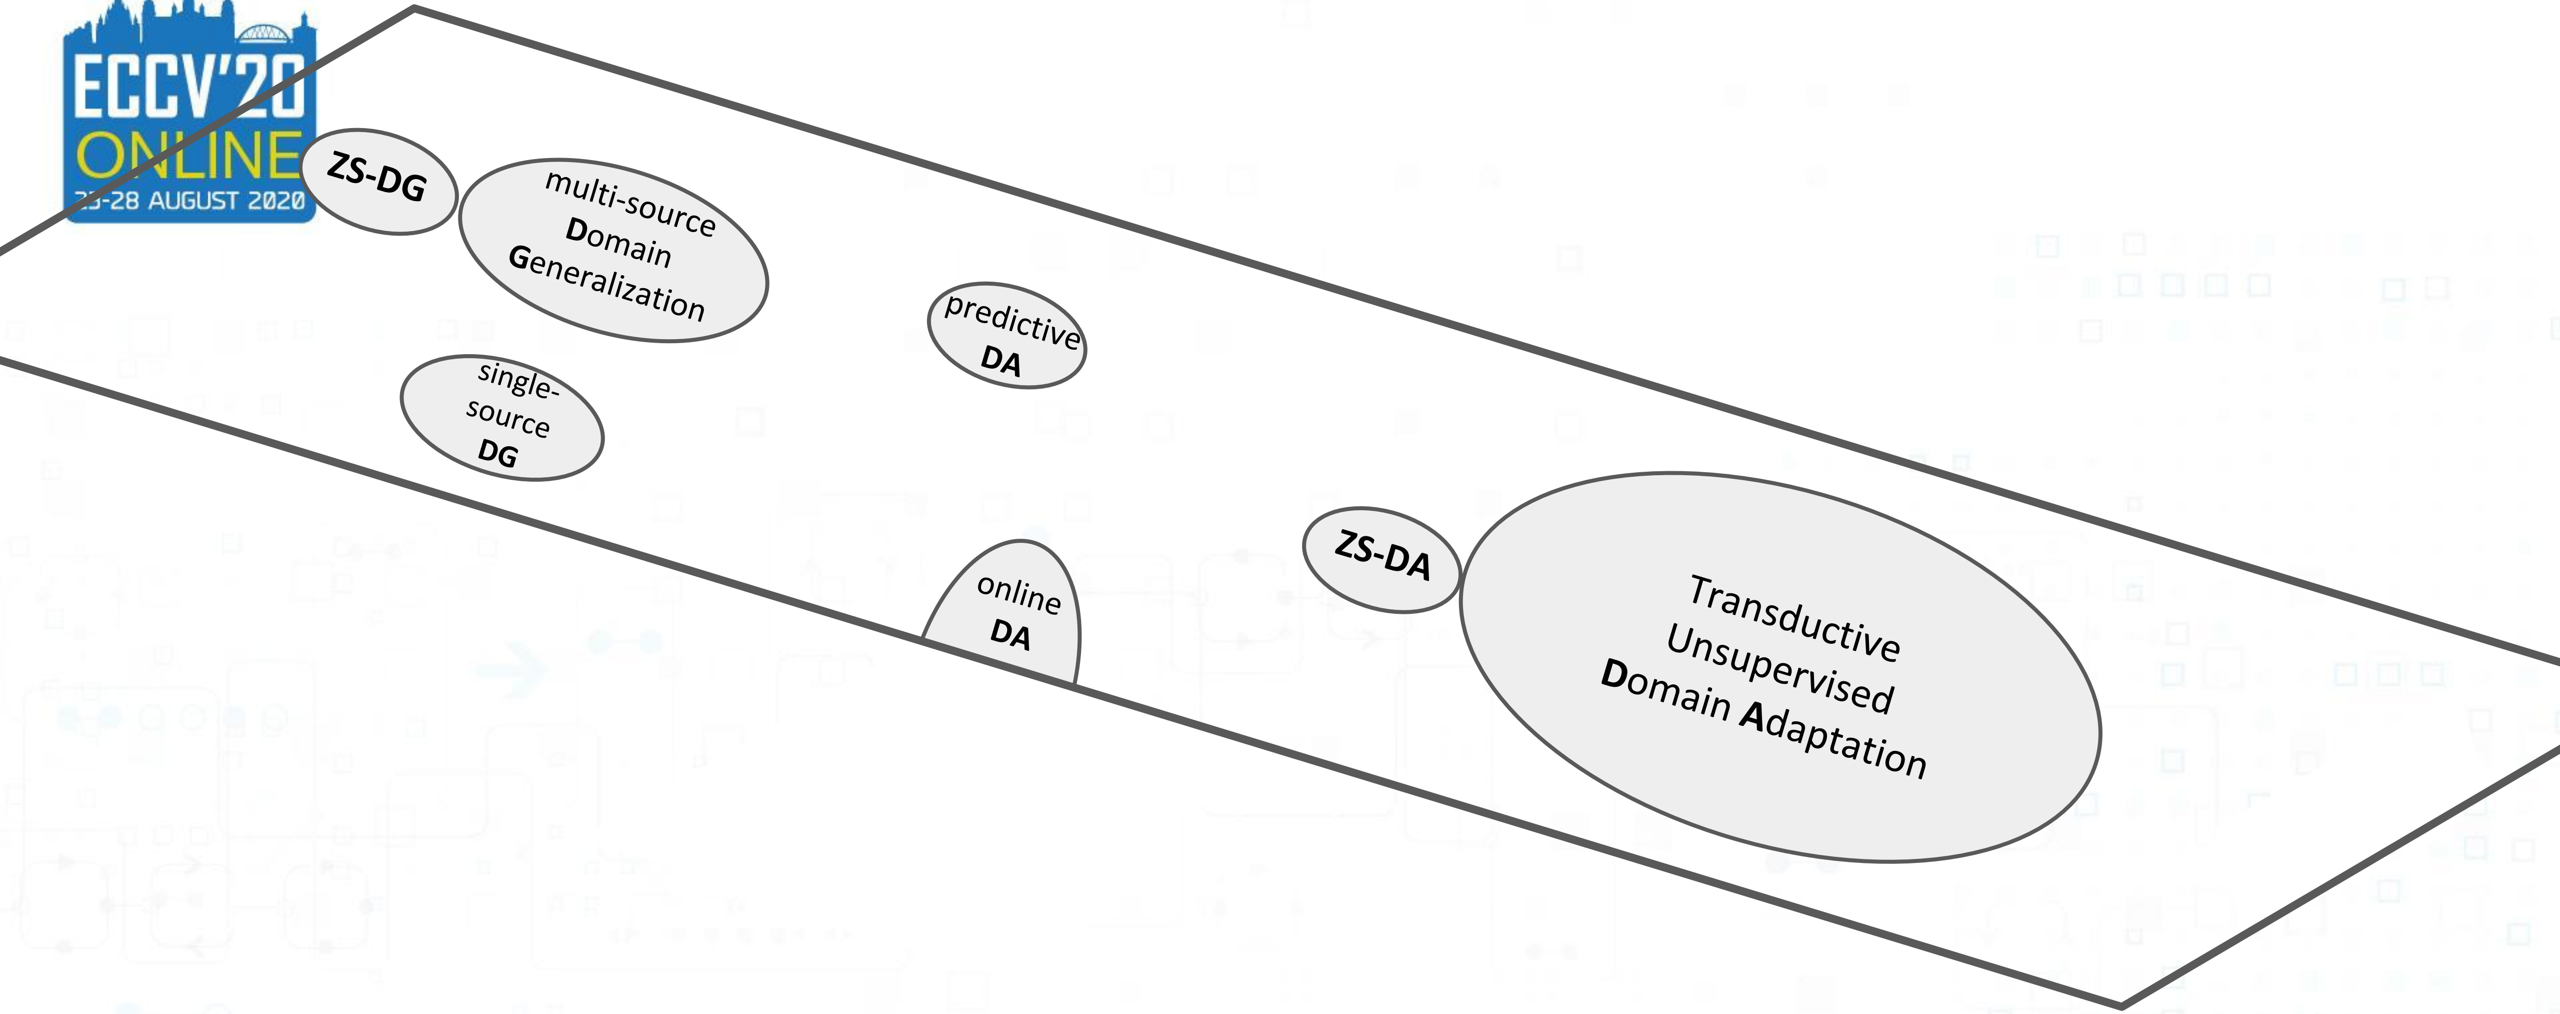
ZS-DG
multi-source Domain Generalization
single-source DG
predictive DA
online DA
ZS-DA
Transductive Unsupervised Domain Adaptation

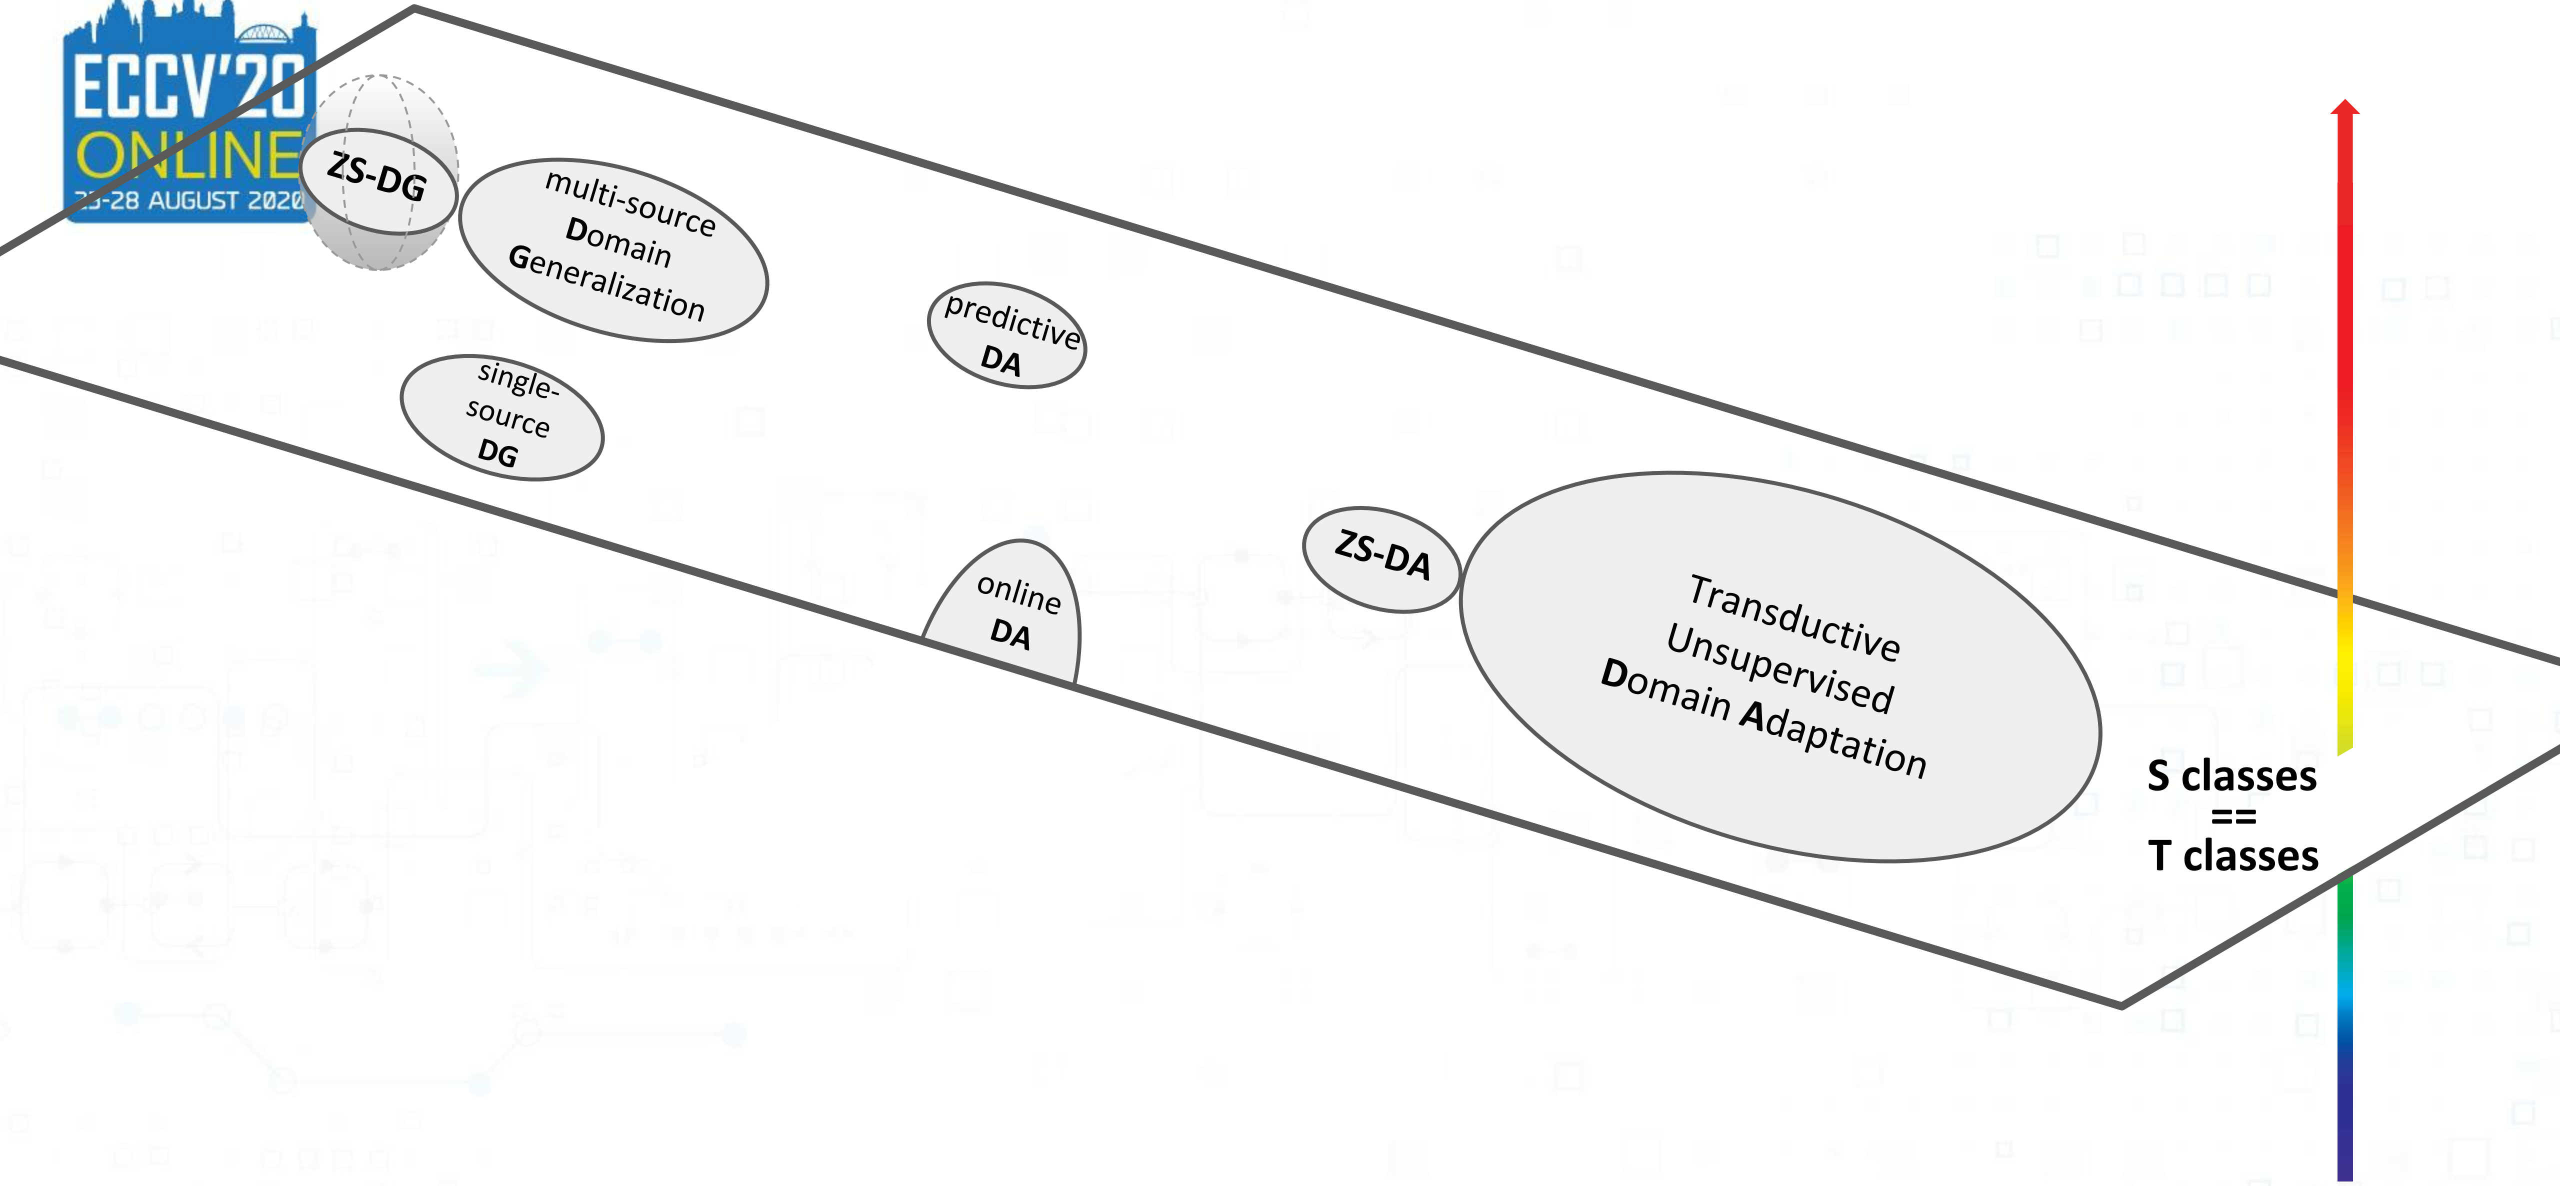

ZS-DG
multi-source Domain Generalization
single-source DG
predictive DA
online DA
ZS-DA
Transductive Unsupervised Domain Adaptation
S classes == T classes

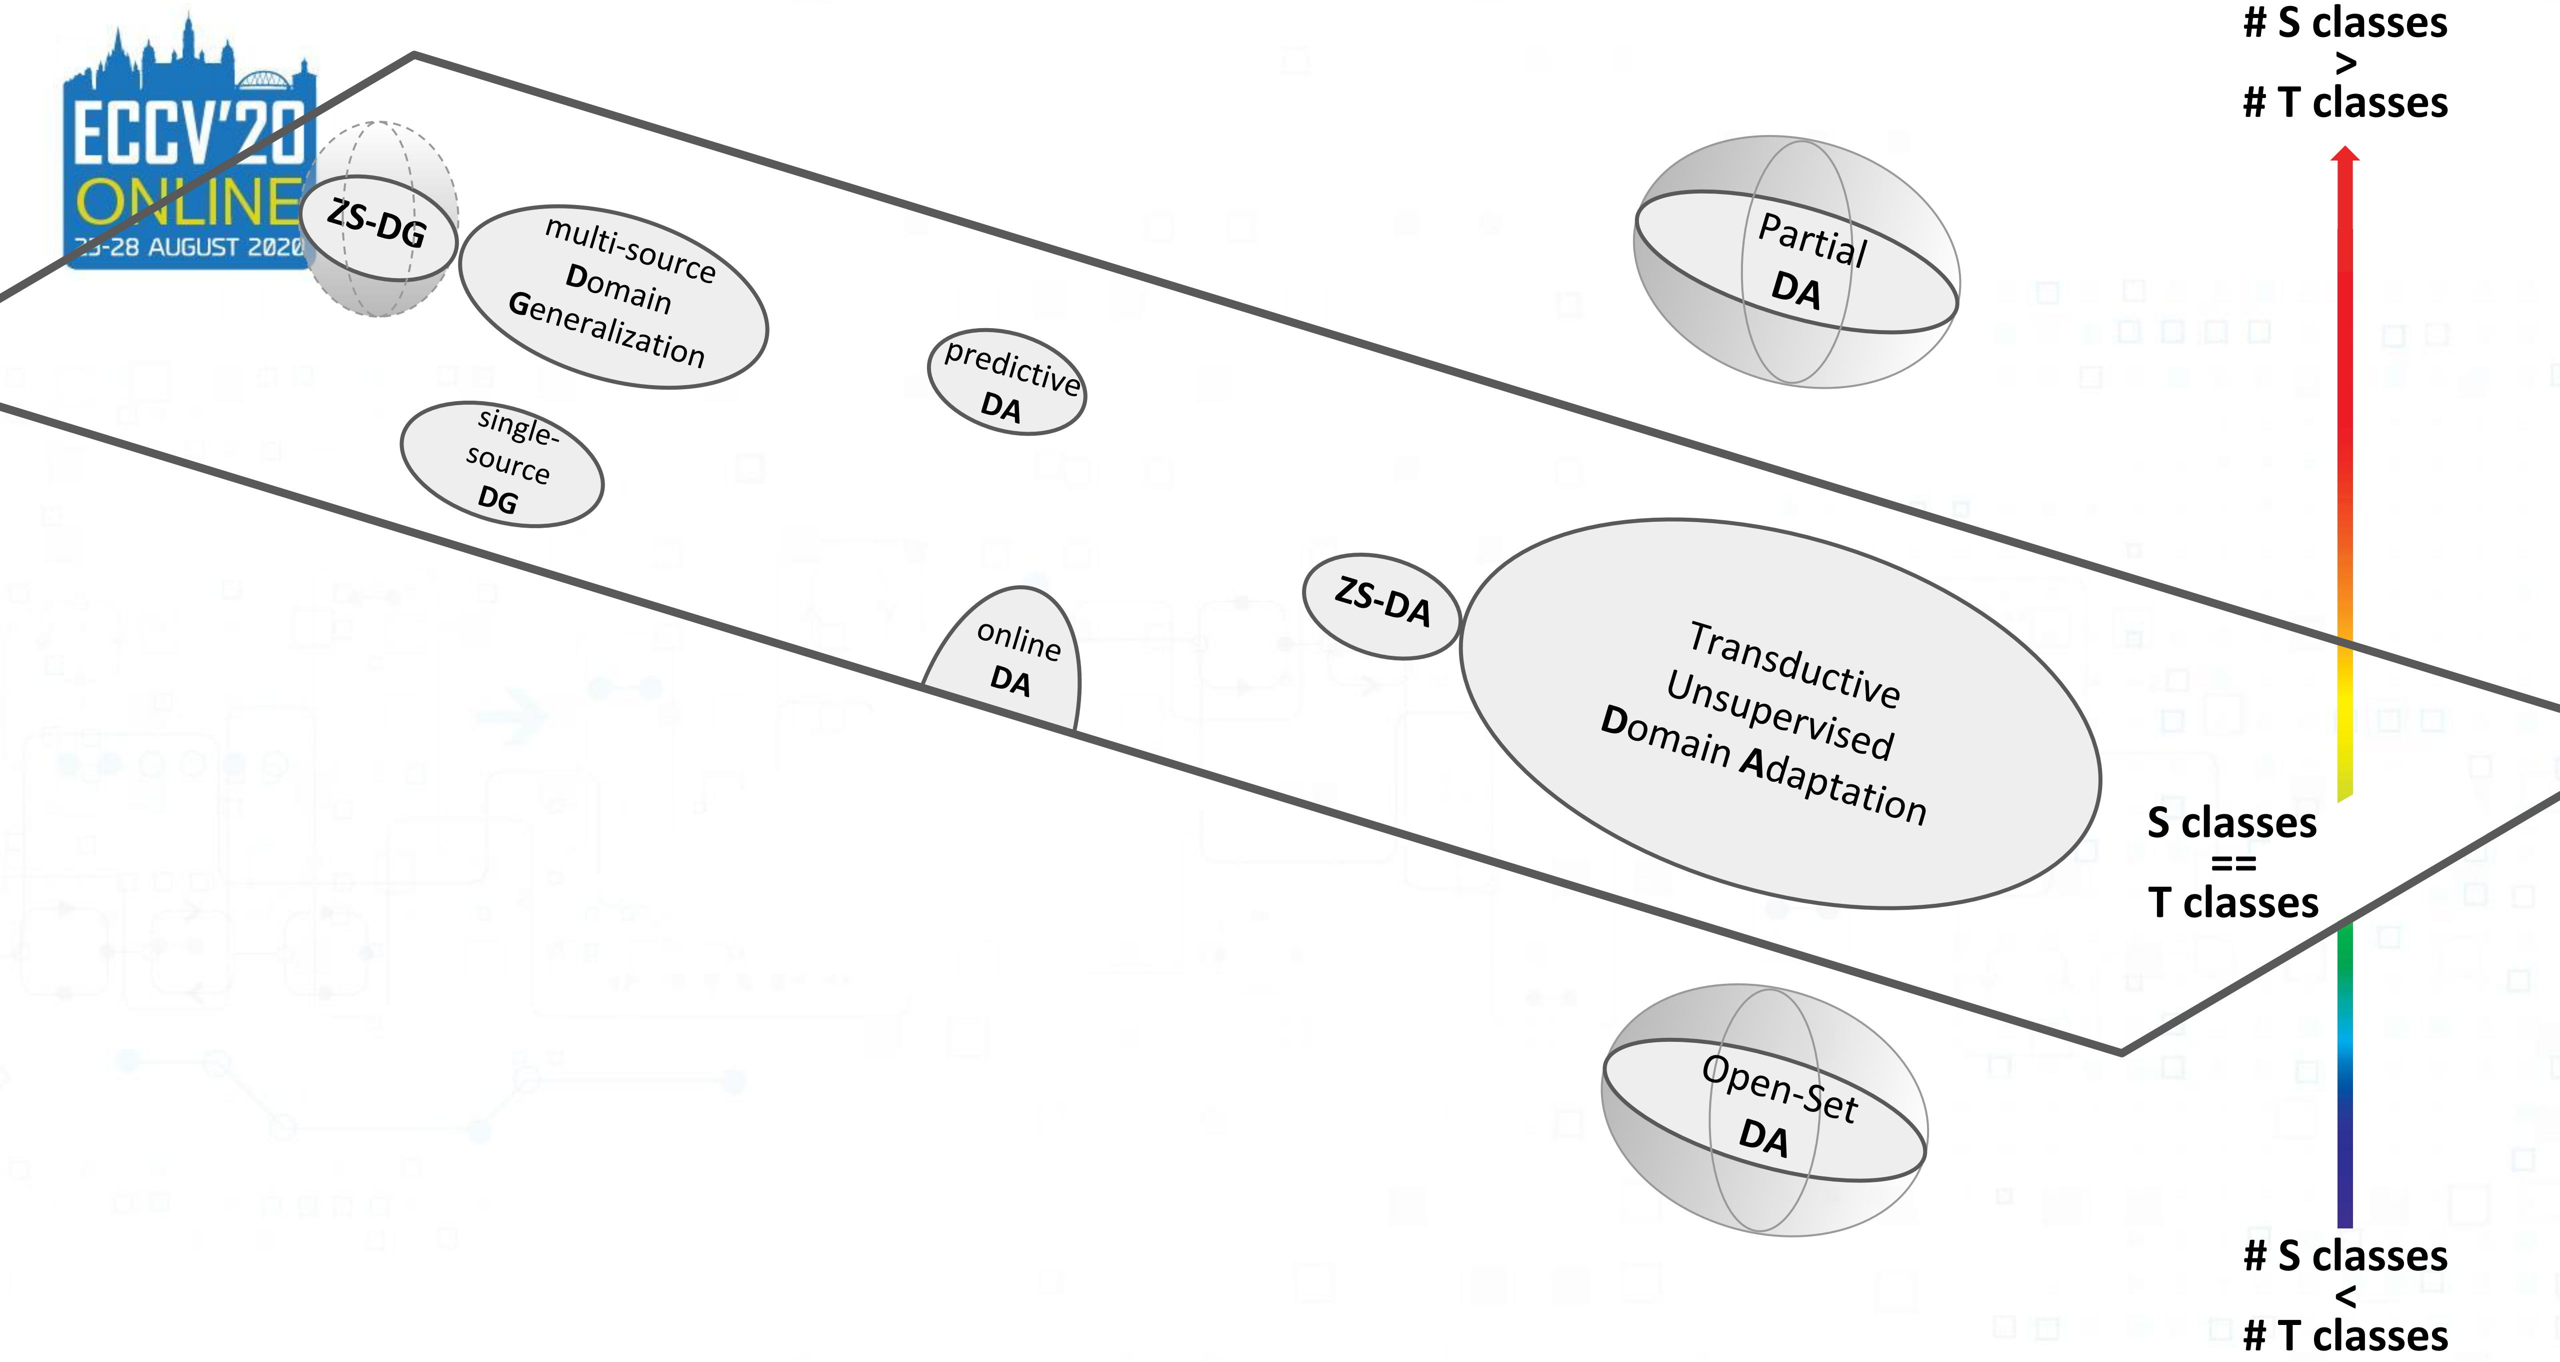
ECCV'20 ONLINE
23-28 AUGUST 2020
S classes
>
T classes
ZS-DG
multi-source
Domain
Generalization
single-
source
DG
predictive
DA
online
DA
ZS-DA
Partial
DA
Transductive
Unsupervised
Domain Adaptation
S classes
==
T classes
Open-Set
DA
S classes
<
T classes

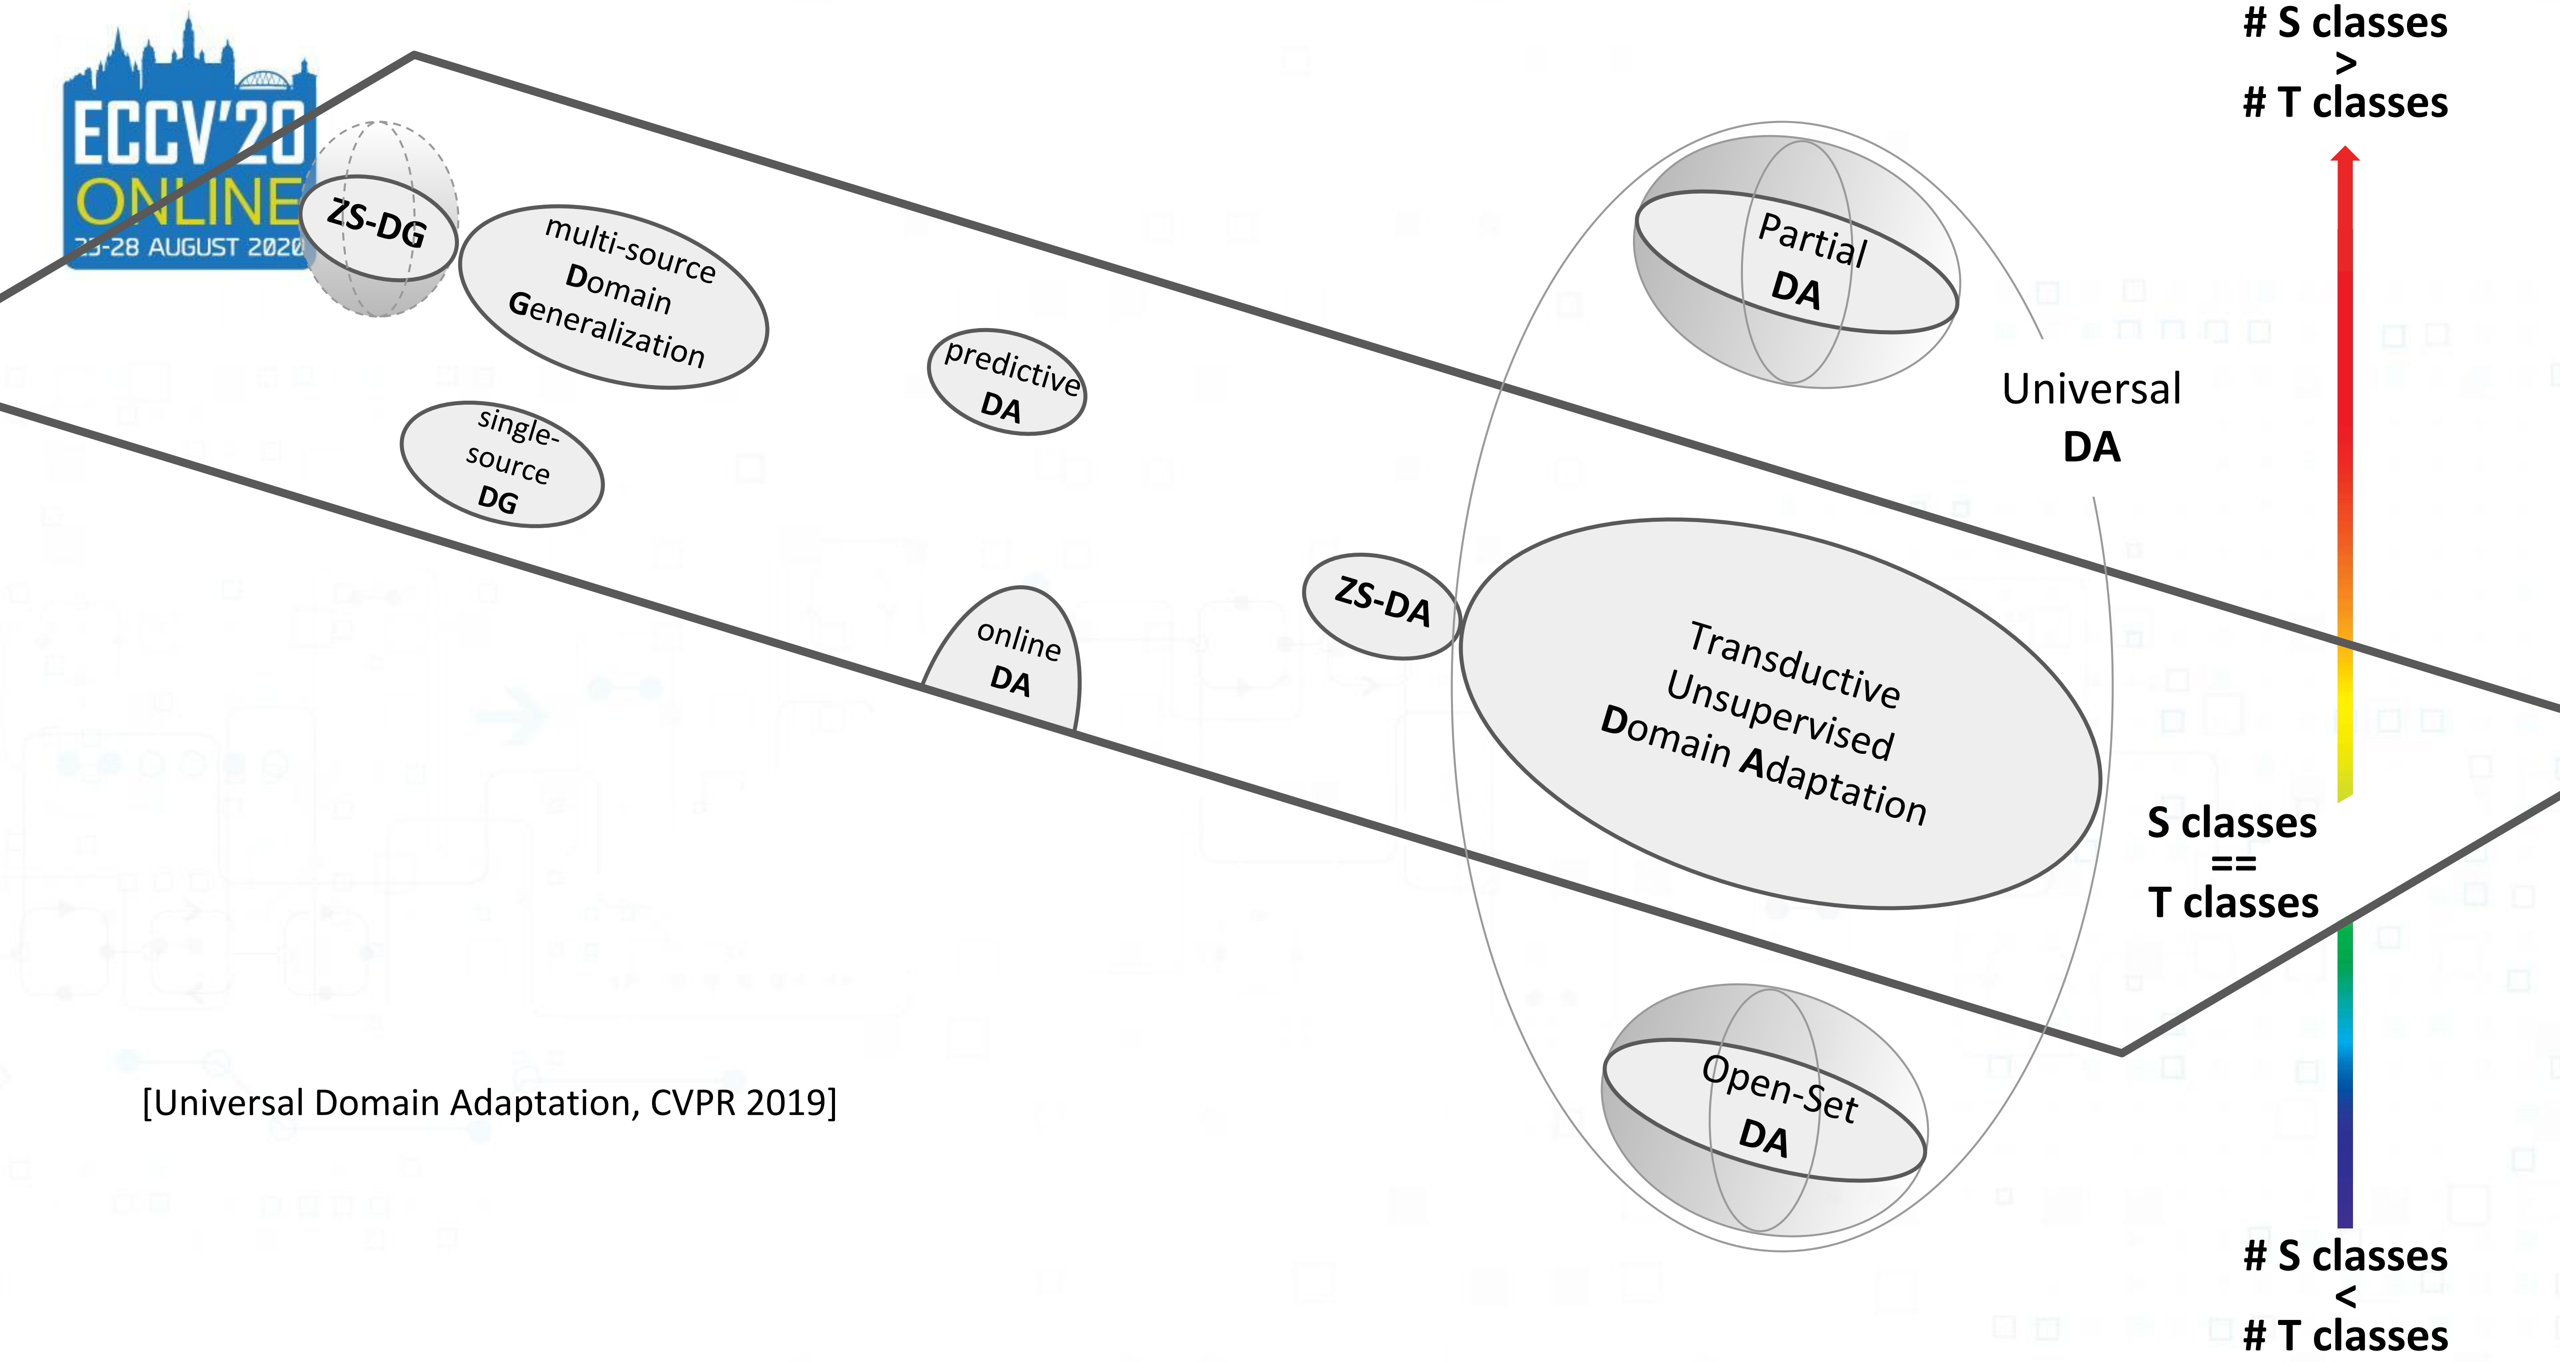

[Universal Domain Adaptation, CVPR 2019]

Partial DA

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]

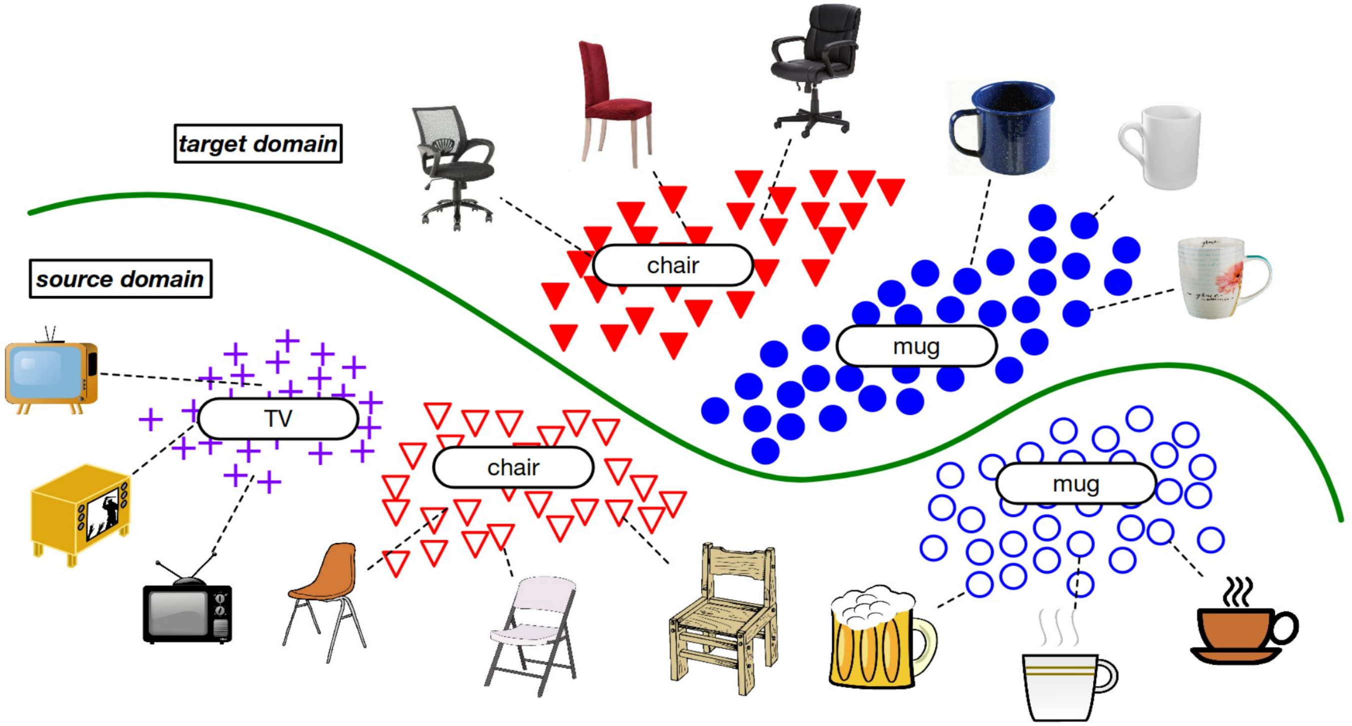

Partial DA

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]

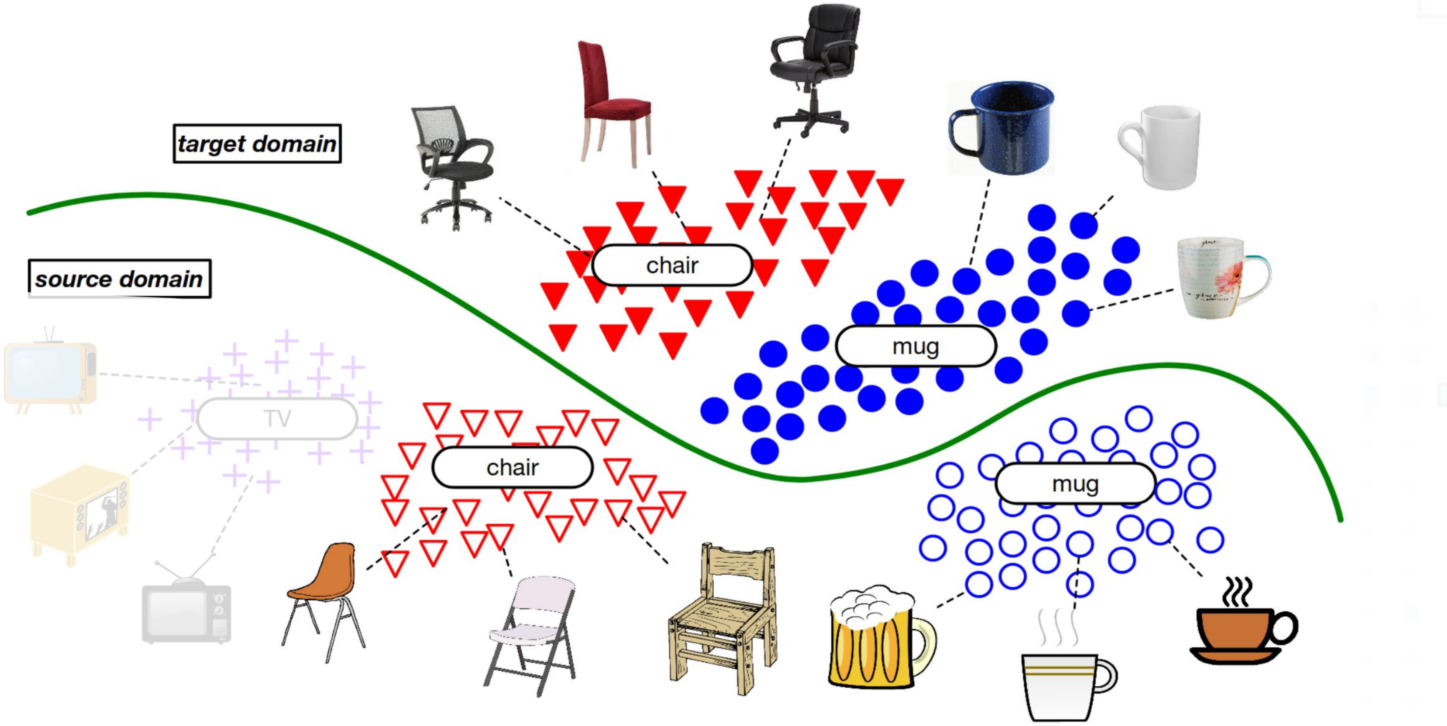

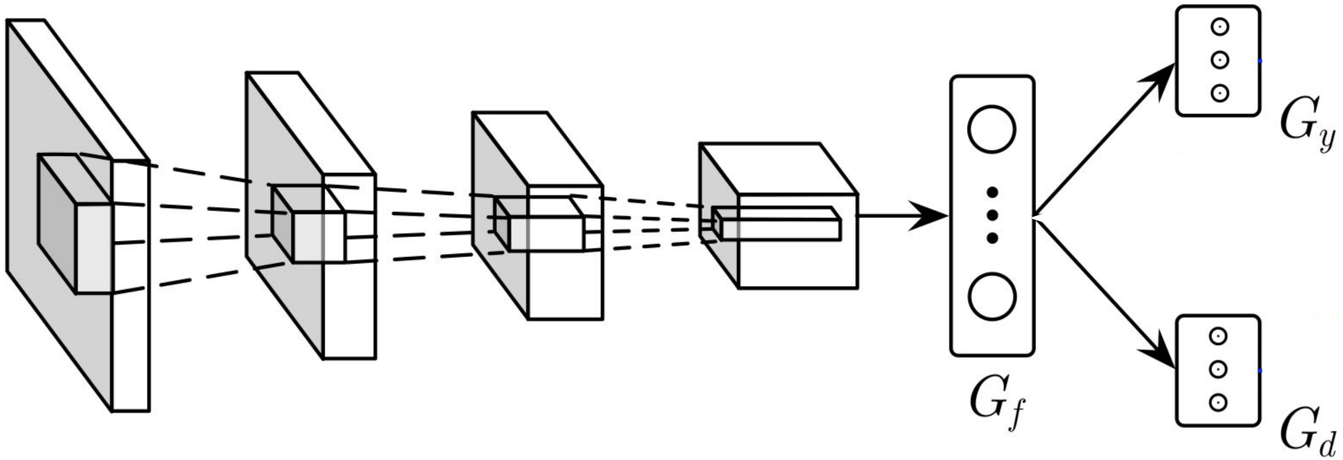

Partial DA

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]

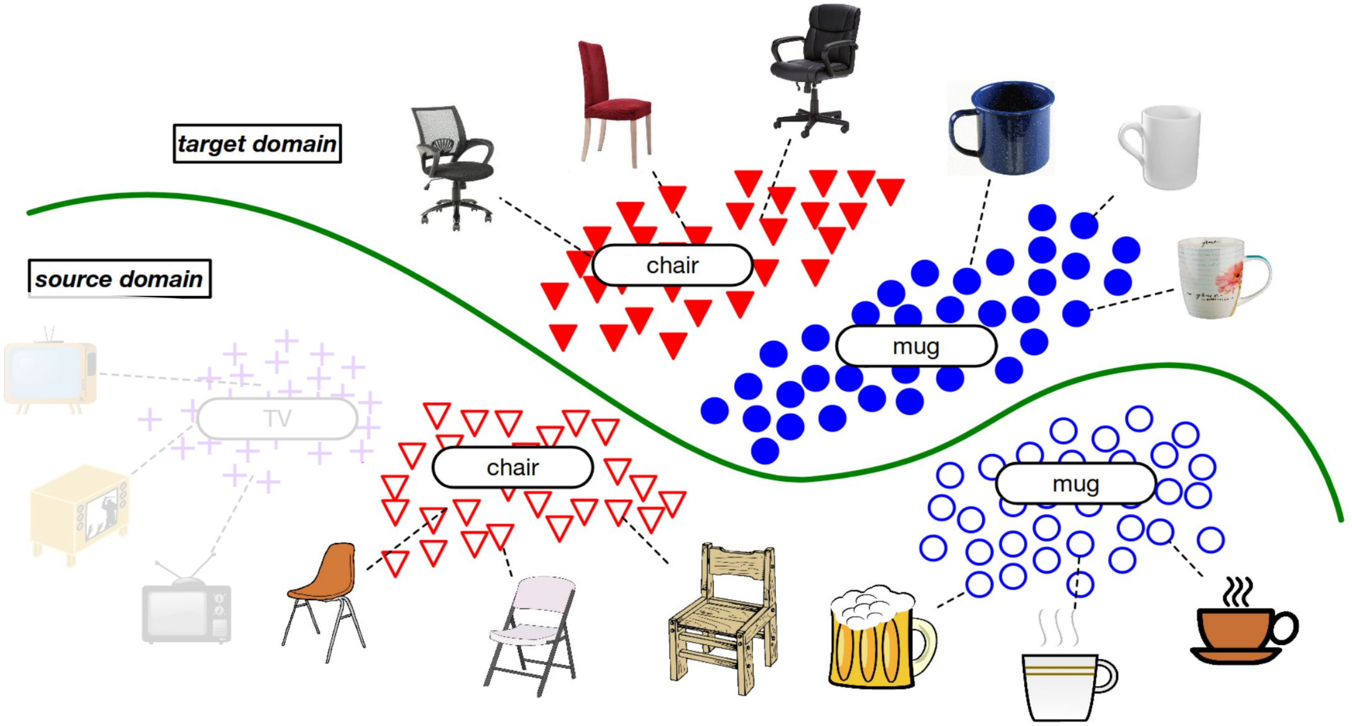

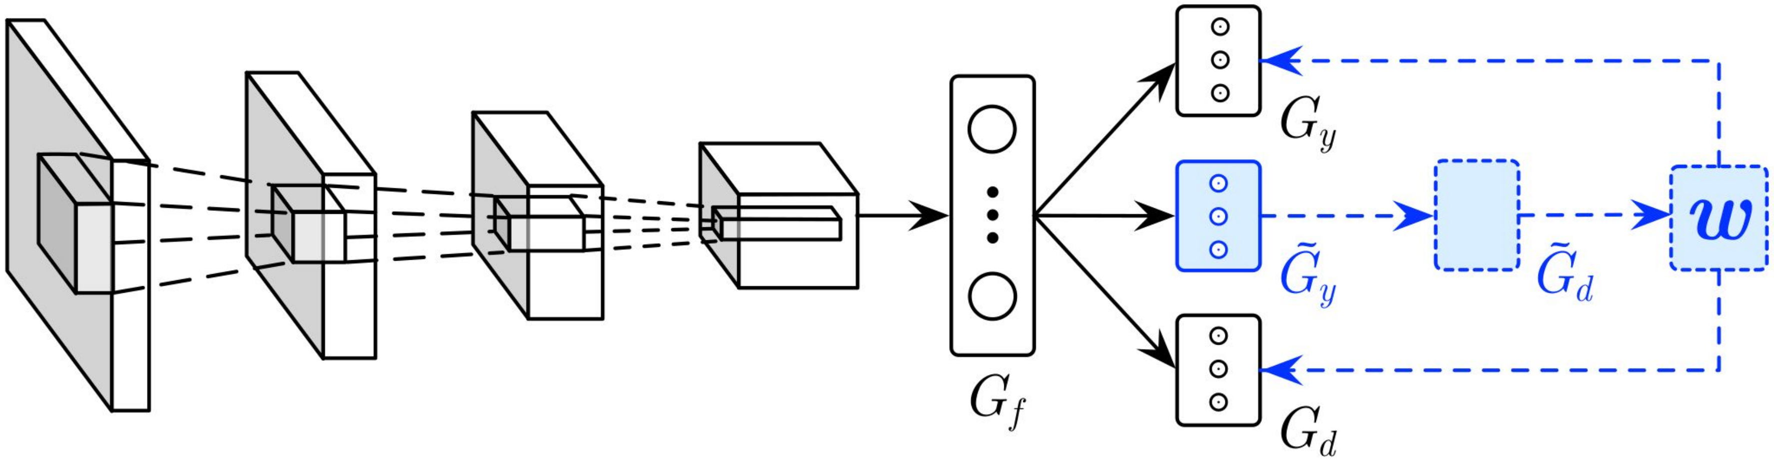

Partial DA

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]
[A Balanced and Uncertainty-aware Approach for Partial Domain Adaptation, ECCV 2020]

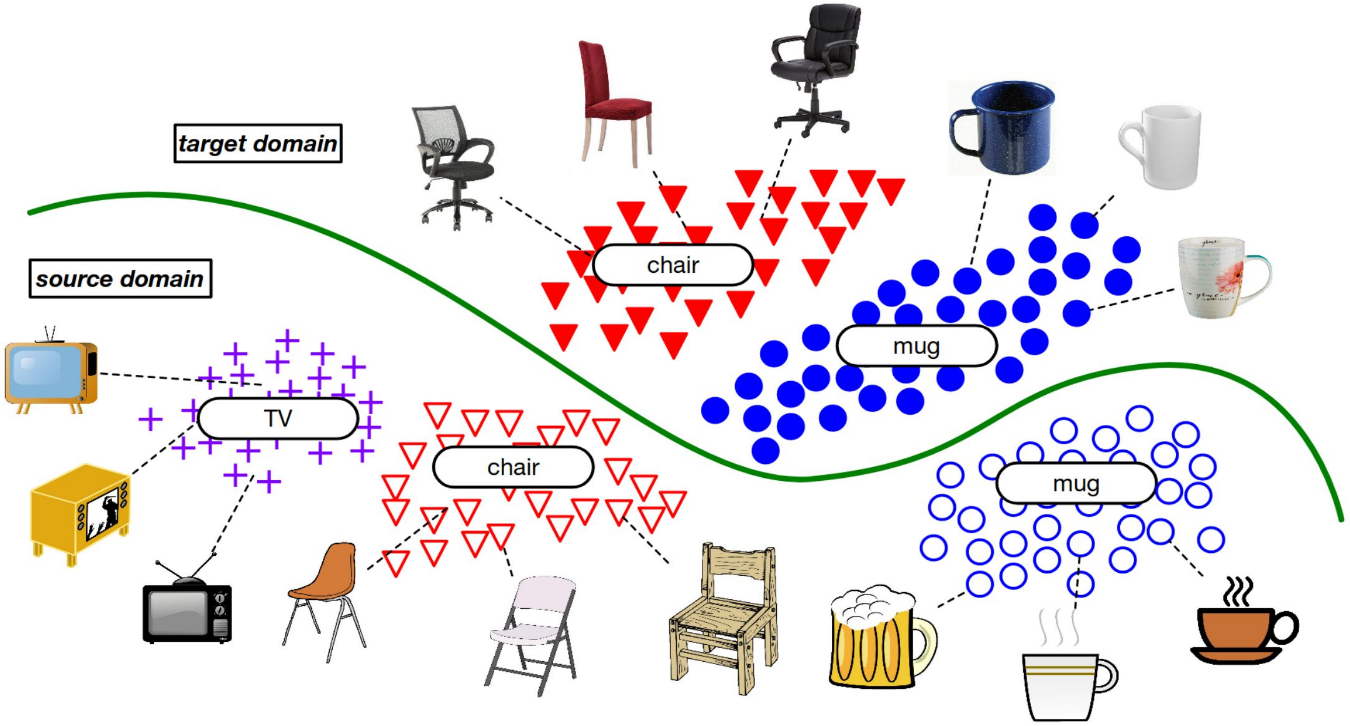

Partial DA

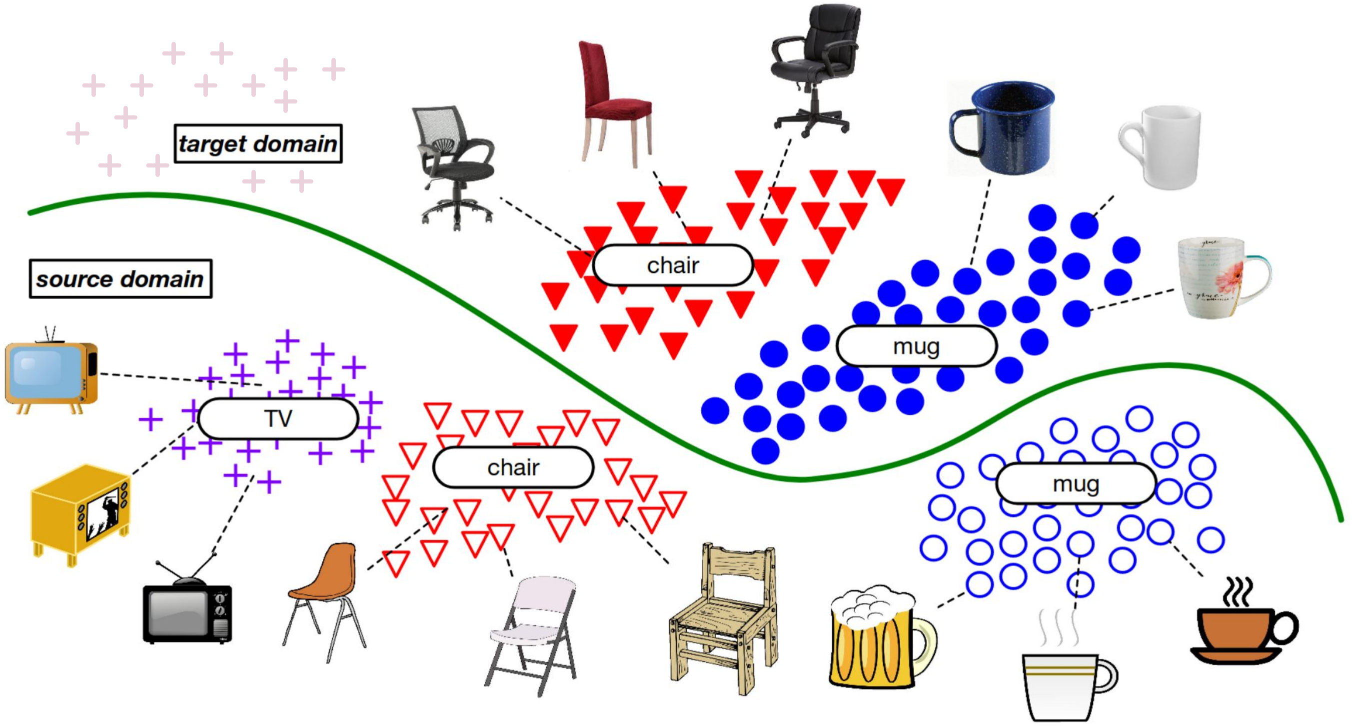

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]
[A Balanced and Uncertainty-aware Approach for Partial Domain Adaptation, ECCV 2020]

Partial DA

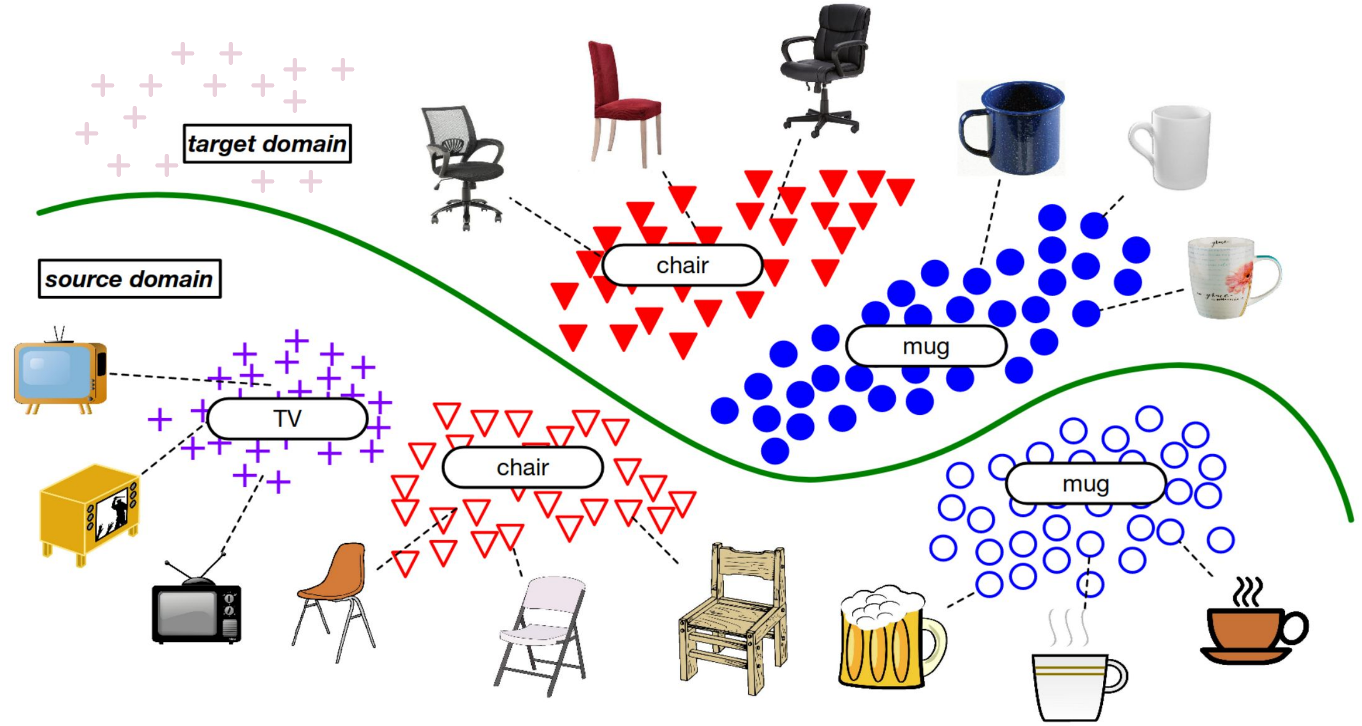

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]
[A Balanced and Uncertainty-aware Approach for Partial Domain Adaptation, ECCV 2020]

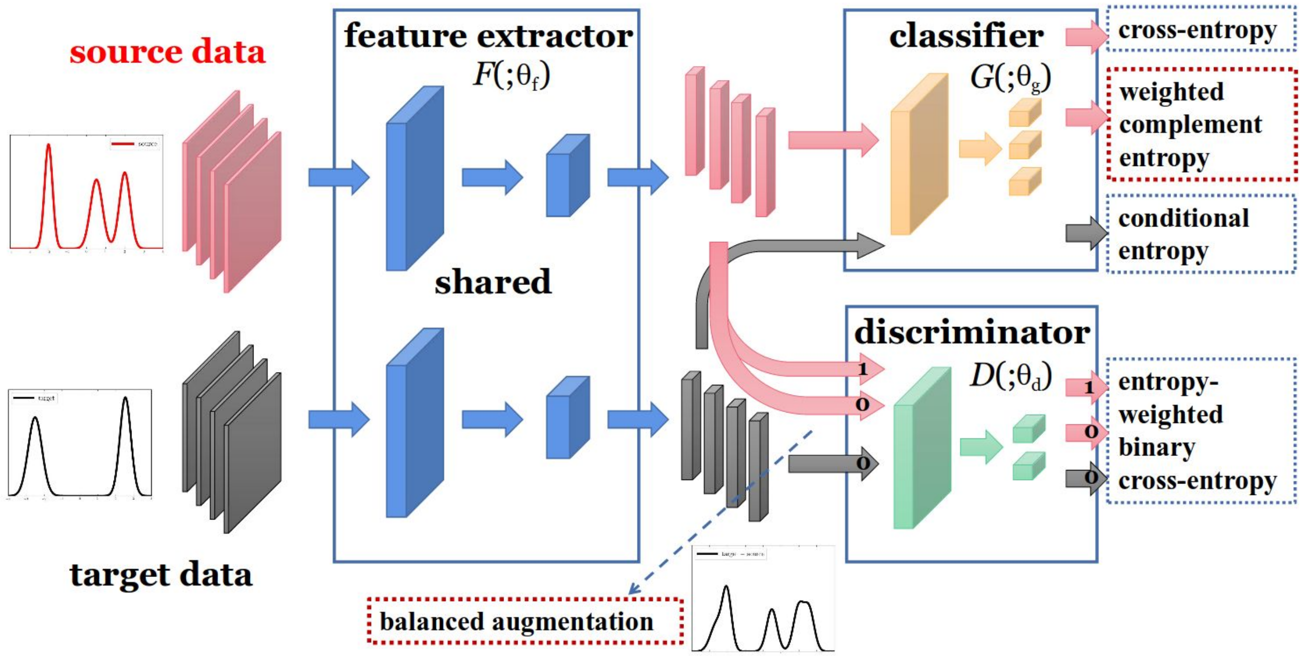

Partial DA

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]
[A Balanced and Uncertainty-aware Approach for Partial Domain Adaptation, ECCV 2020]

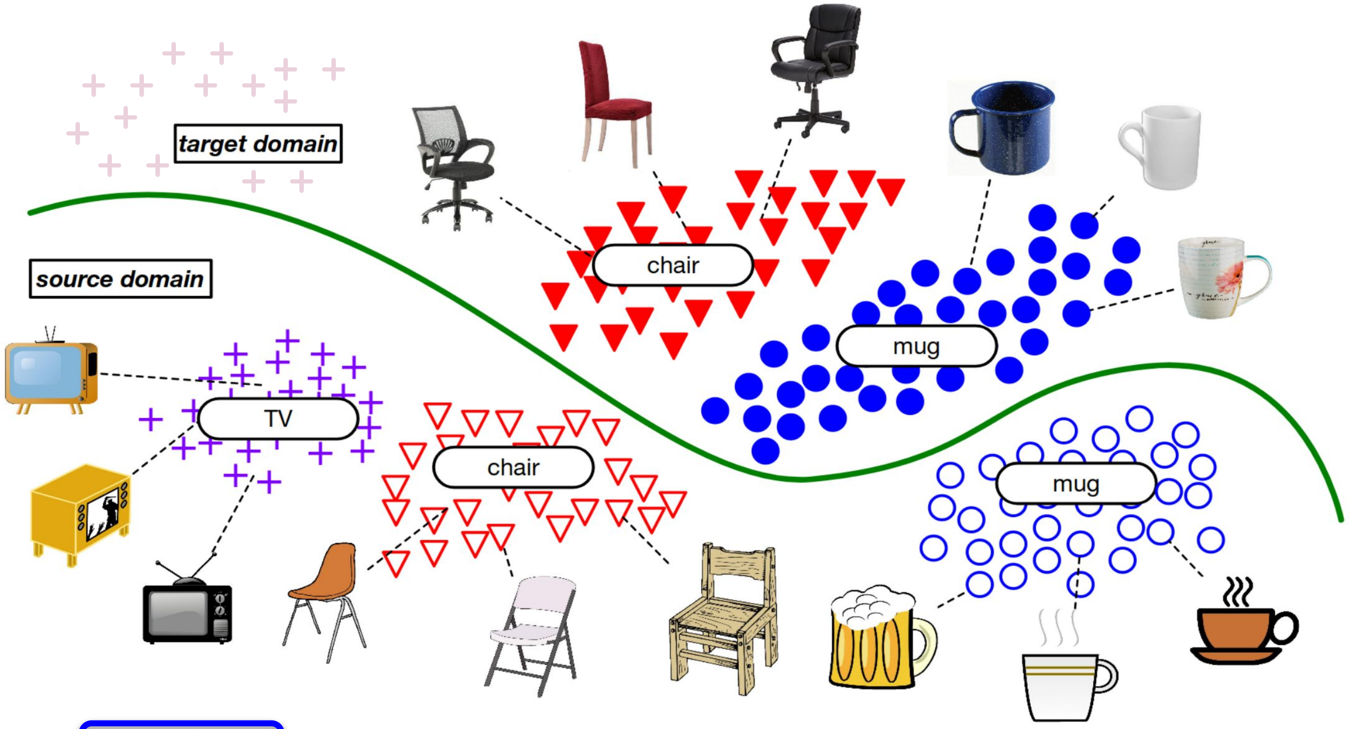

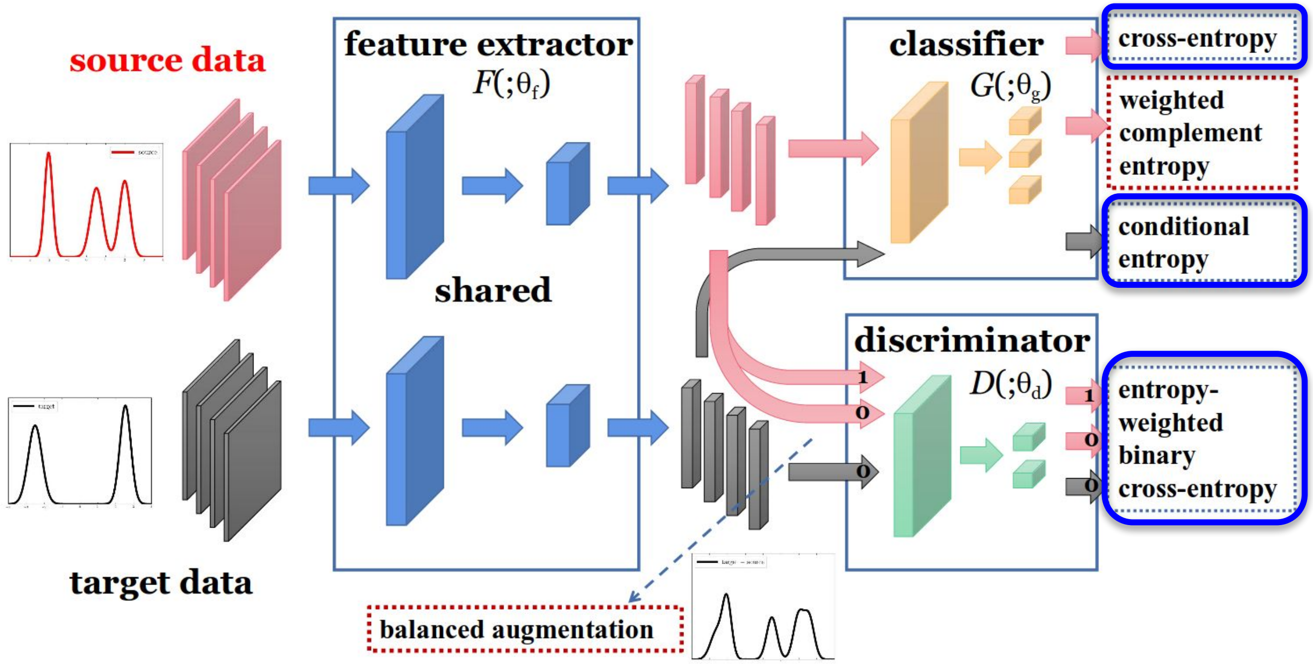

- Entropy-weighted adversarial domain discriminator (e-DANN)

Partial DA

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]
[A Balanced and Uncertainty-aware Approach for Partial Domain Adaptation, ECCV 2020]

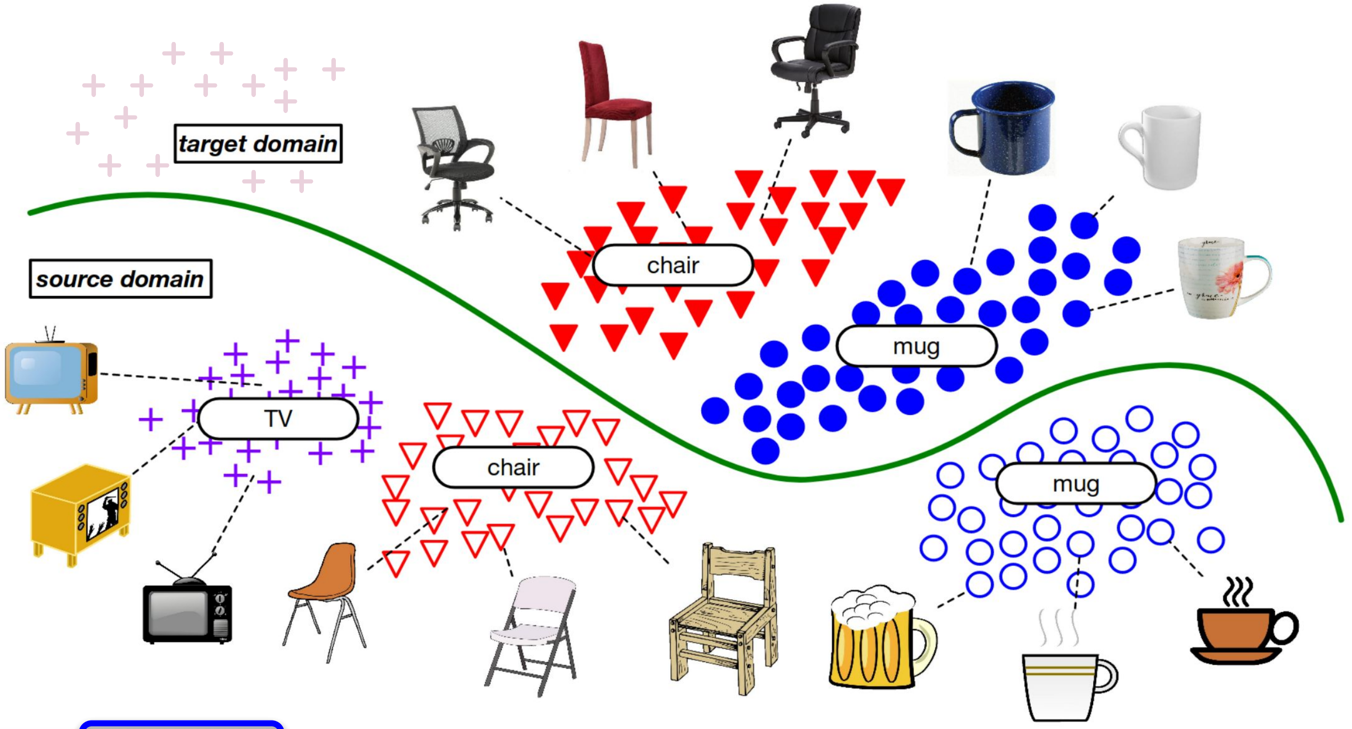

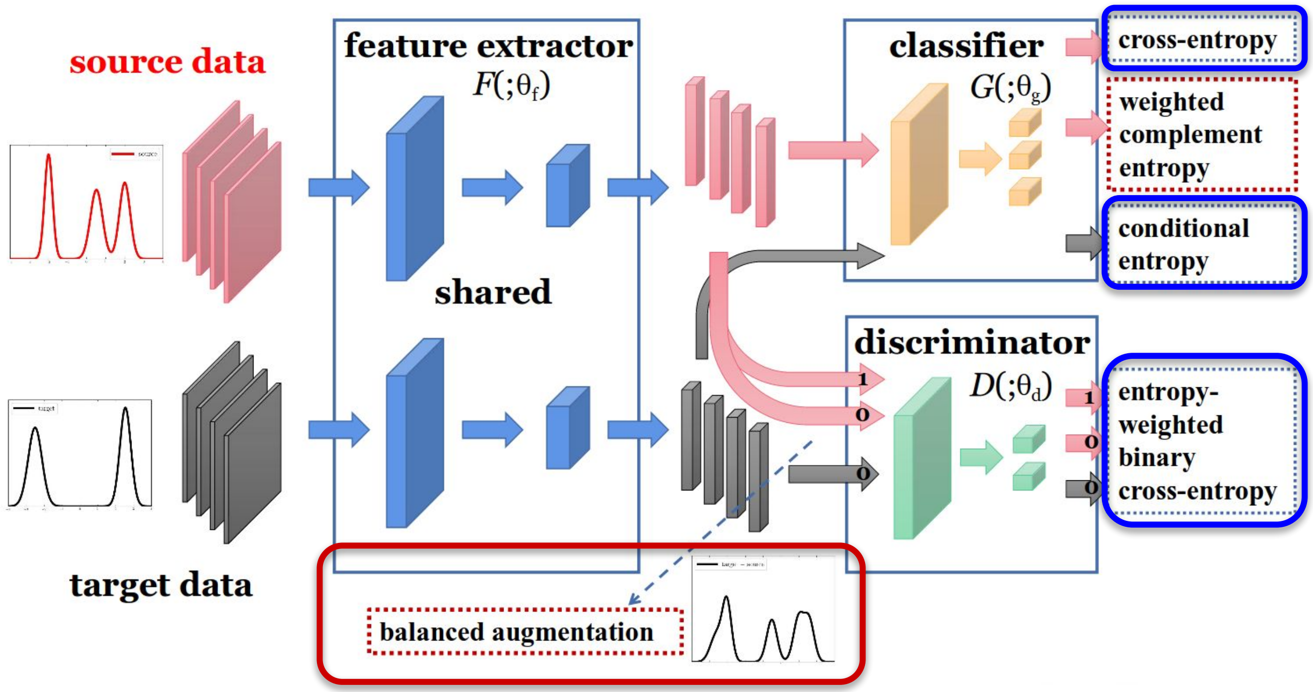

- Entropy-weighted adversarial domain discriminator (e-DANN)
- Borrow a fraction of the source sample per class and consider them as target

Partial DA

[Partial Adversarial Domain Adaptation, ECCV 2018]
[Learning to transfer examples for partial domain adaptation, CVPR 2019]

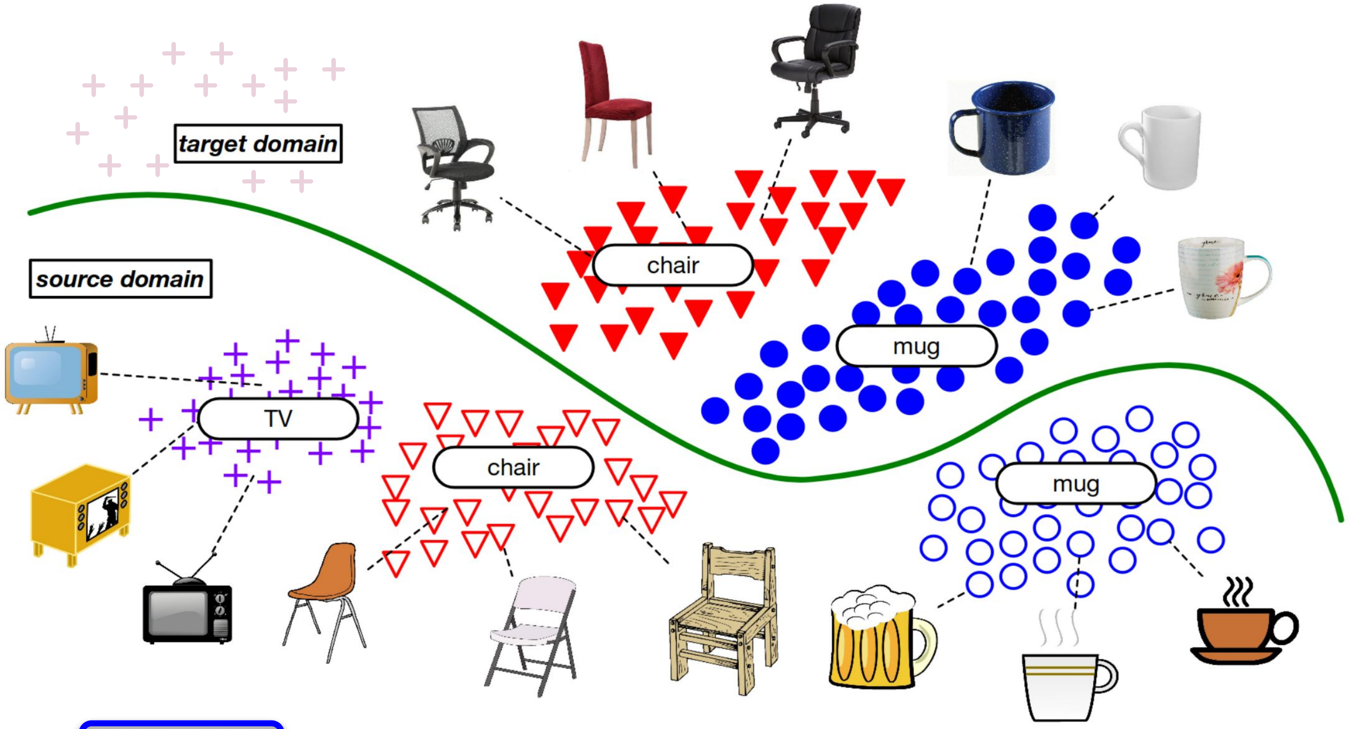

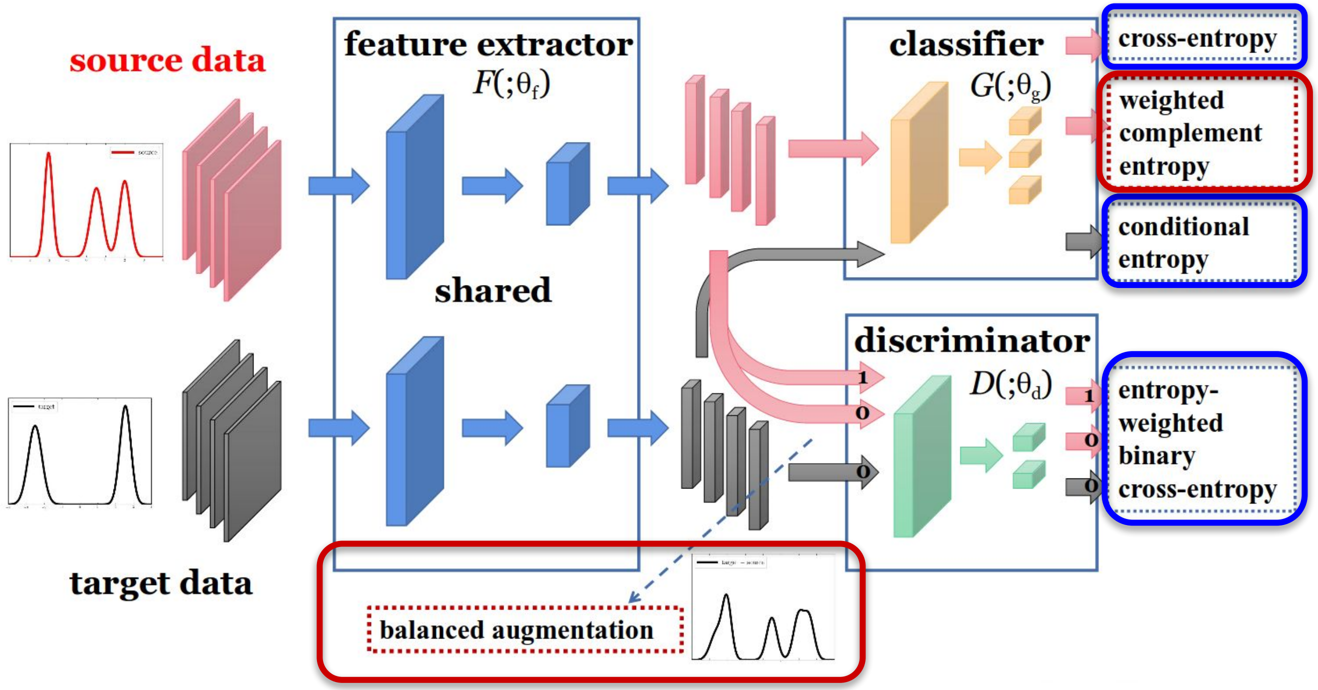

- Entropy-weighted adversarial domain discriminator (e-DANN)
- Borrow a fraction of the source sample per class and consider them as target
- Encourage uniform and low prediction scores for incorrect classes of the source

Open-Set DA

[Open Set Domain Adaptation by Backpropagation, ECCV 2018]
[Separate to Adapt: Open Set Domain Adaptation via Progressive Separation, CVPR 2019]

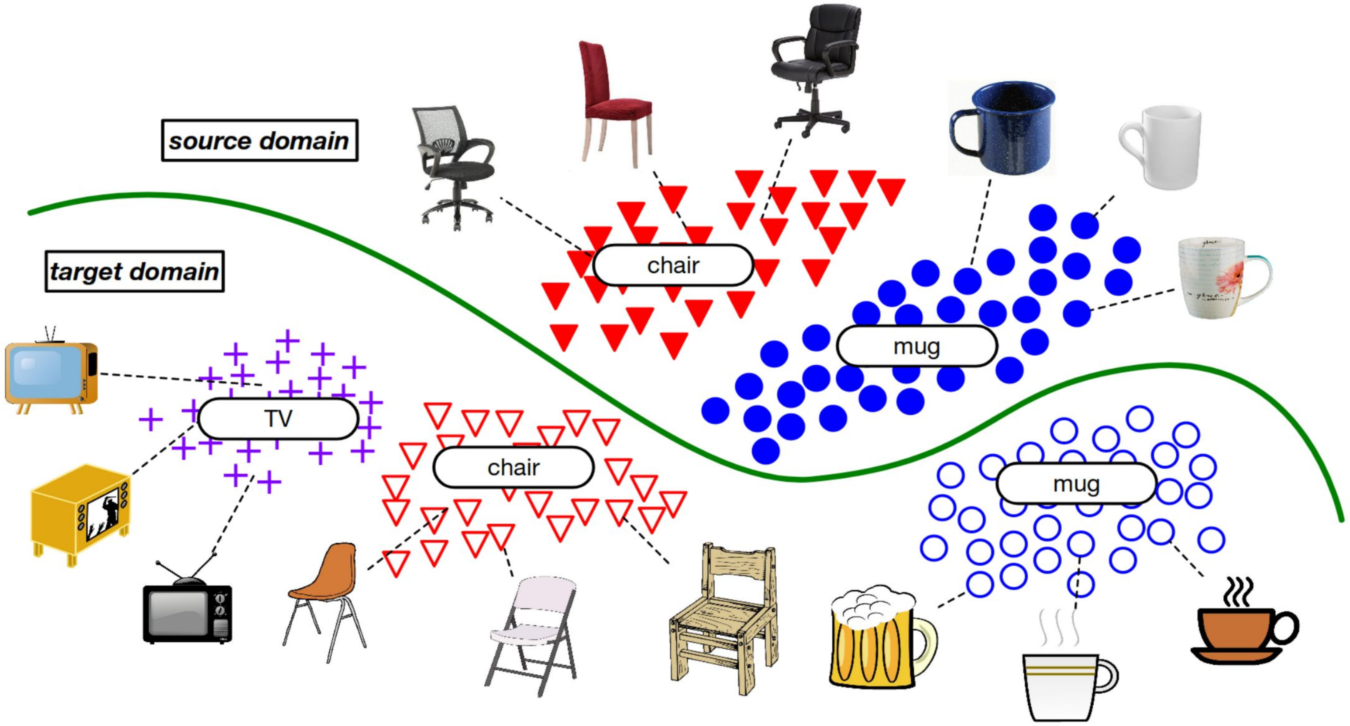

Open-Set DA

[Open Set Domain Adaptation by Backpropagation, ECCV 2018]
[Separate to Adapt: Open Set Domain Adaptation via Progressive Separation, CVPR 2019]

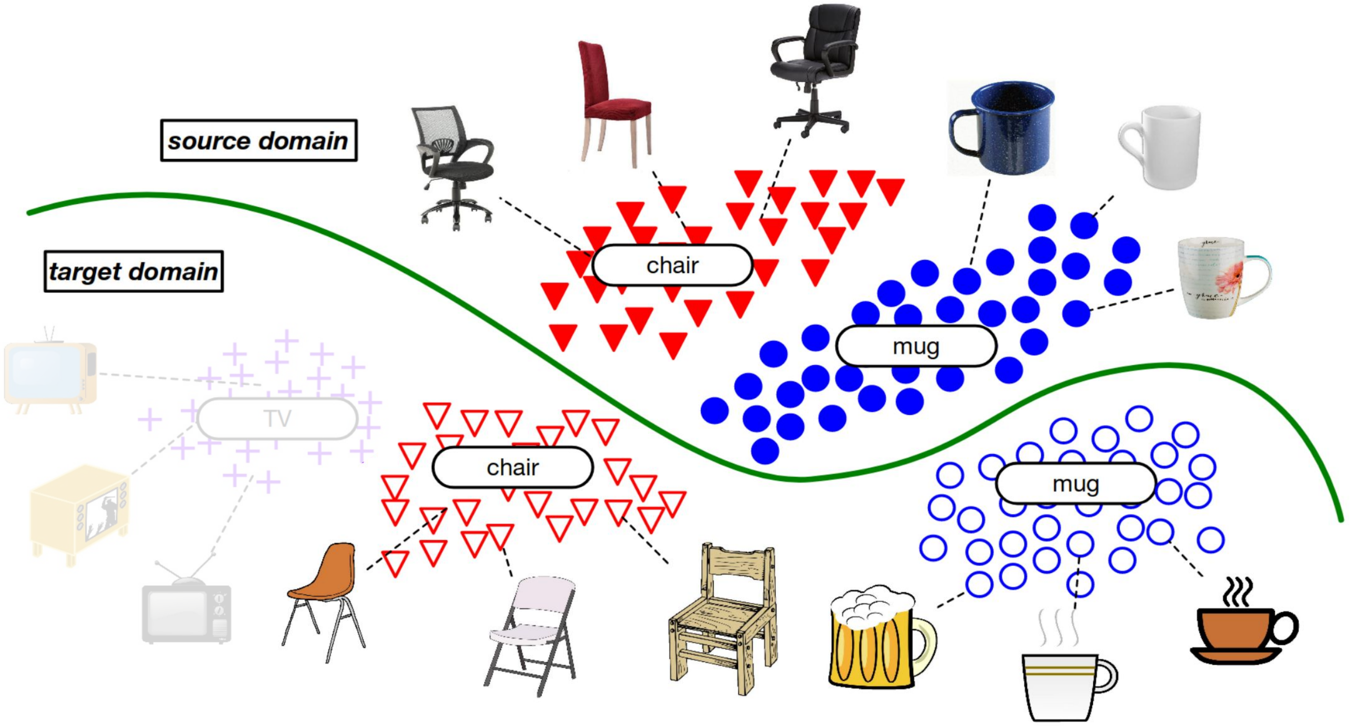

Open-Set DA

[Open Set Domain Adaptation by Backpropagation, ECCV 2018]
[Separate to Adapt: Open Set Domain Adaptation via Progressive Separation, CVPR 2019]

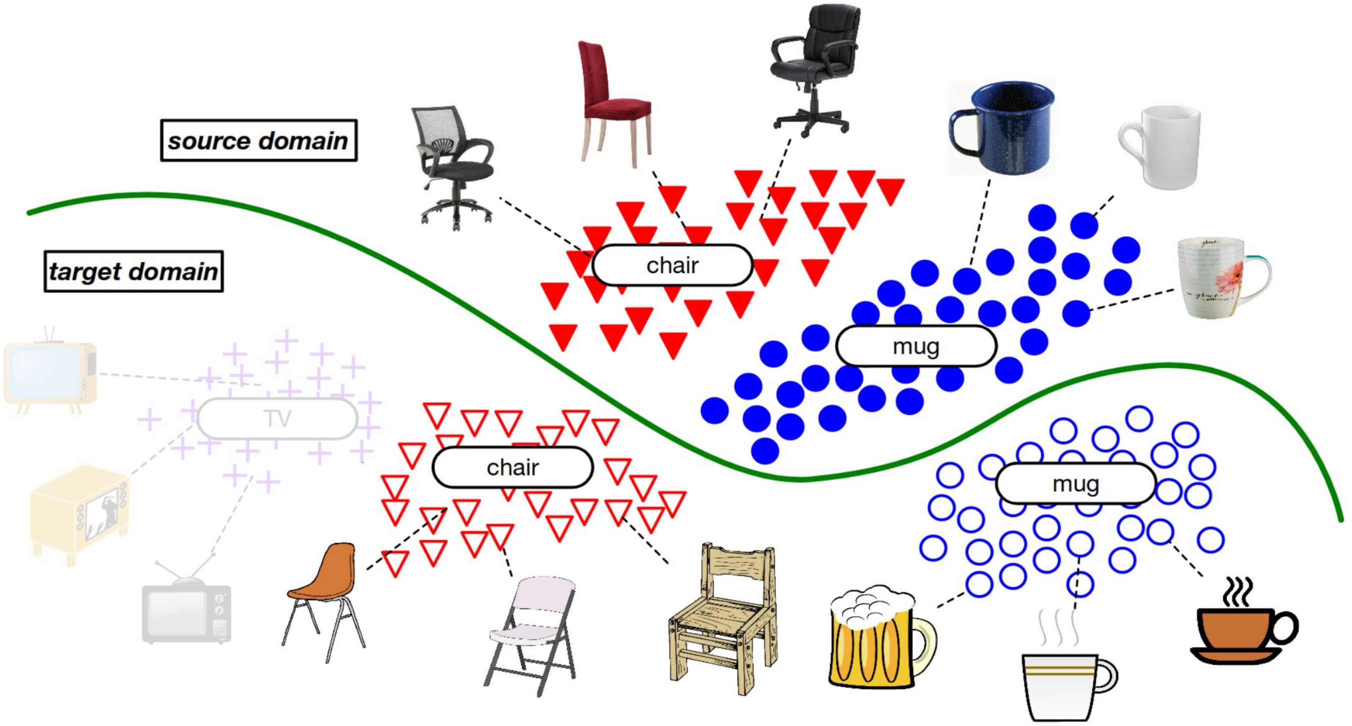

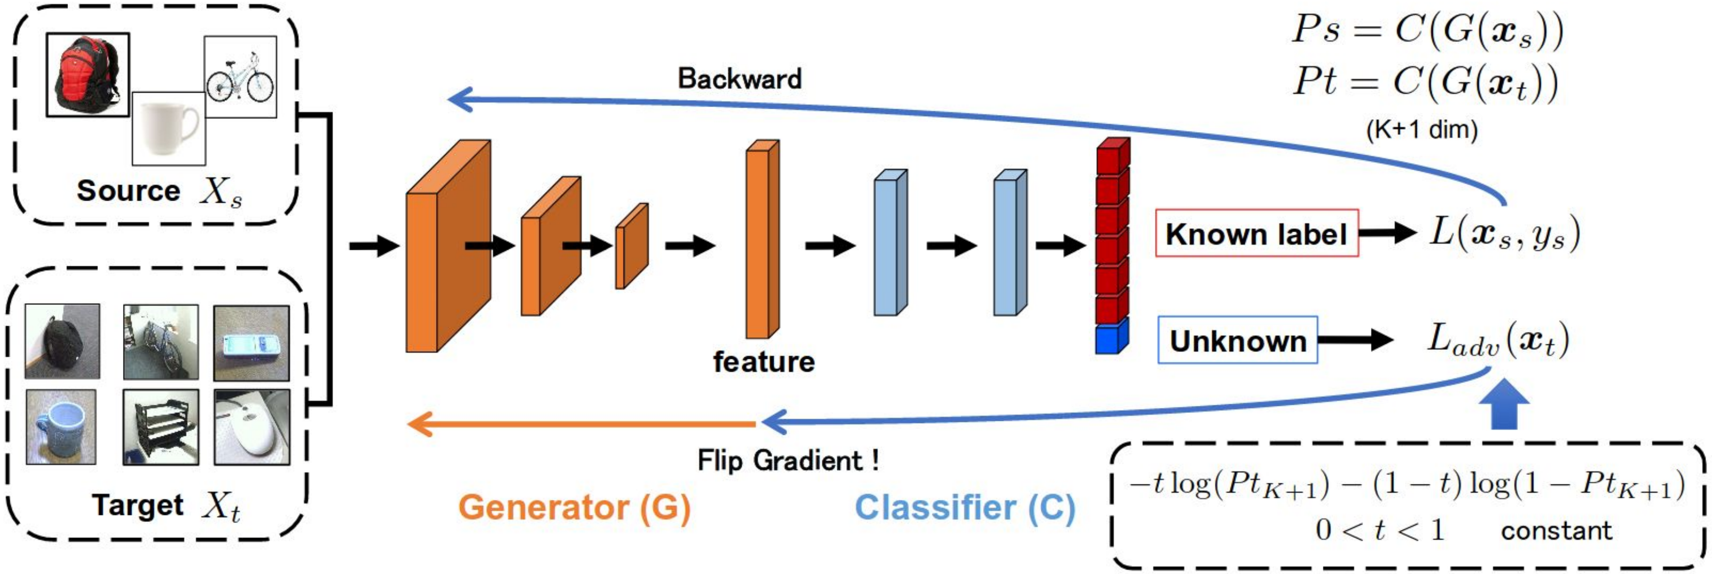

Open-Set DA

[Open Set Domain Adaptation by Backpropagation, ECCV 2018]
[Separate to Adapt: Open Set Domain Adaptation via Progressive Separation, CVPR 2019]

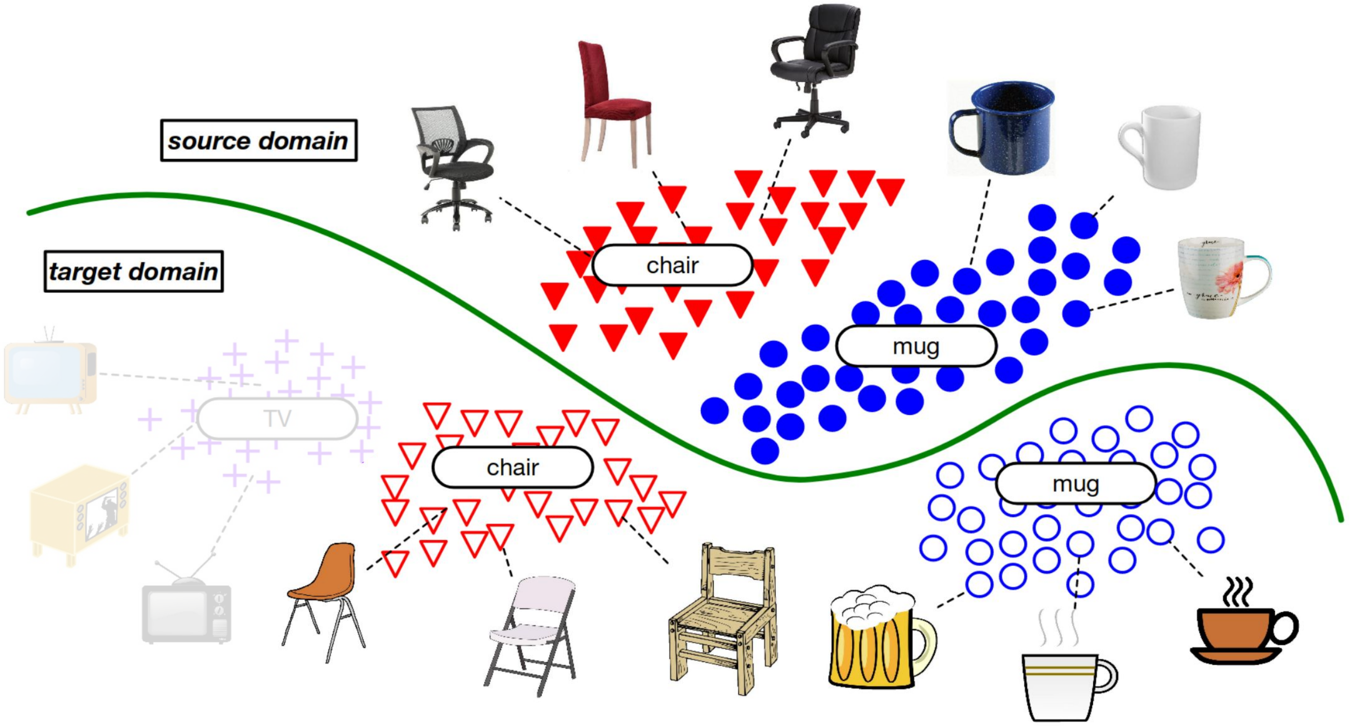

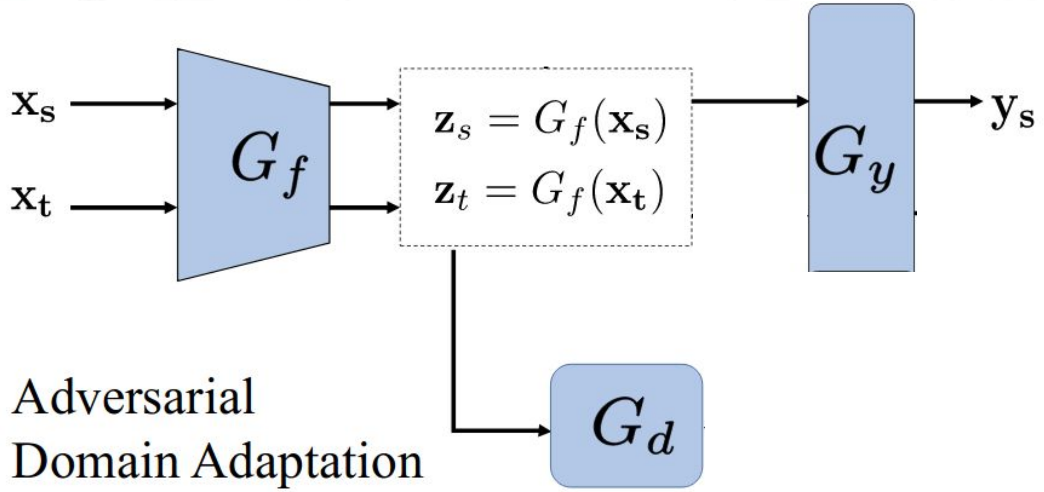

Open-Set DA

[Open Set Domain Adaptation by Backpropagation, ECCV 2018]
[Separate to Adapt: Open Set Domain Adaptation via Progressive Separation, CVPR 2019]

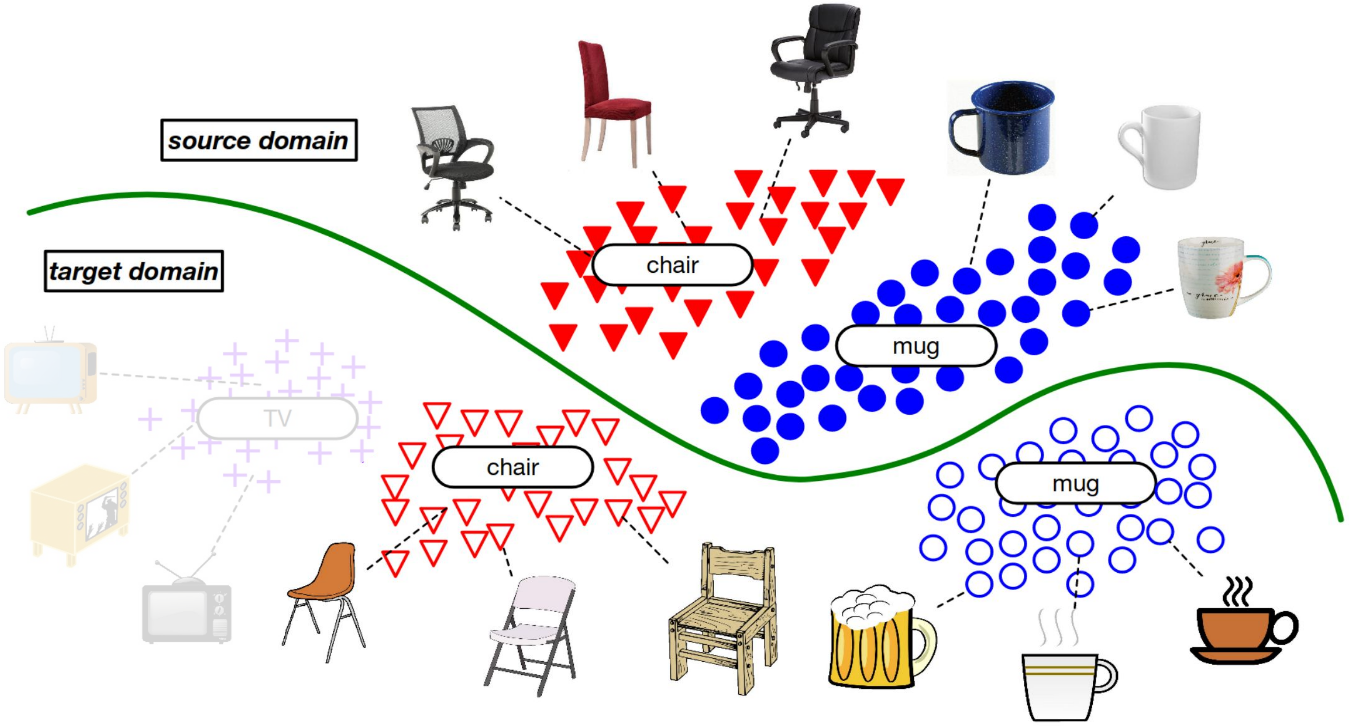

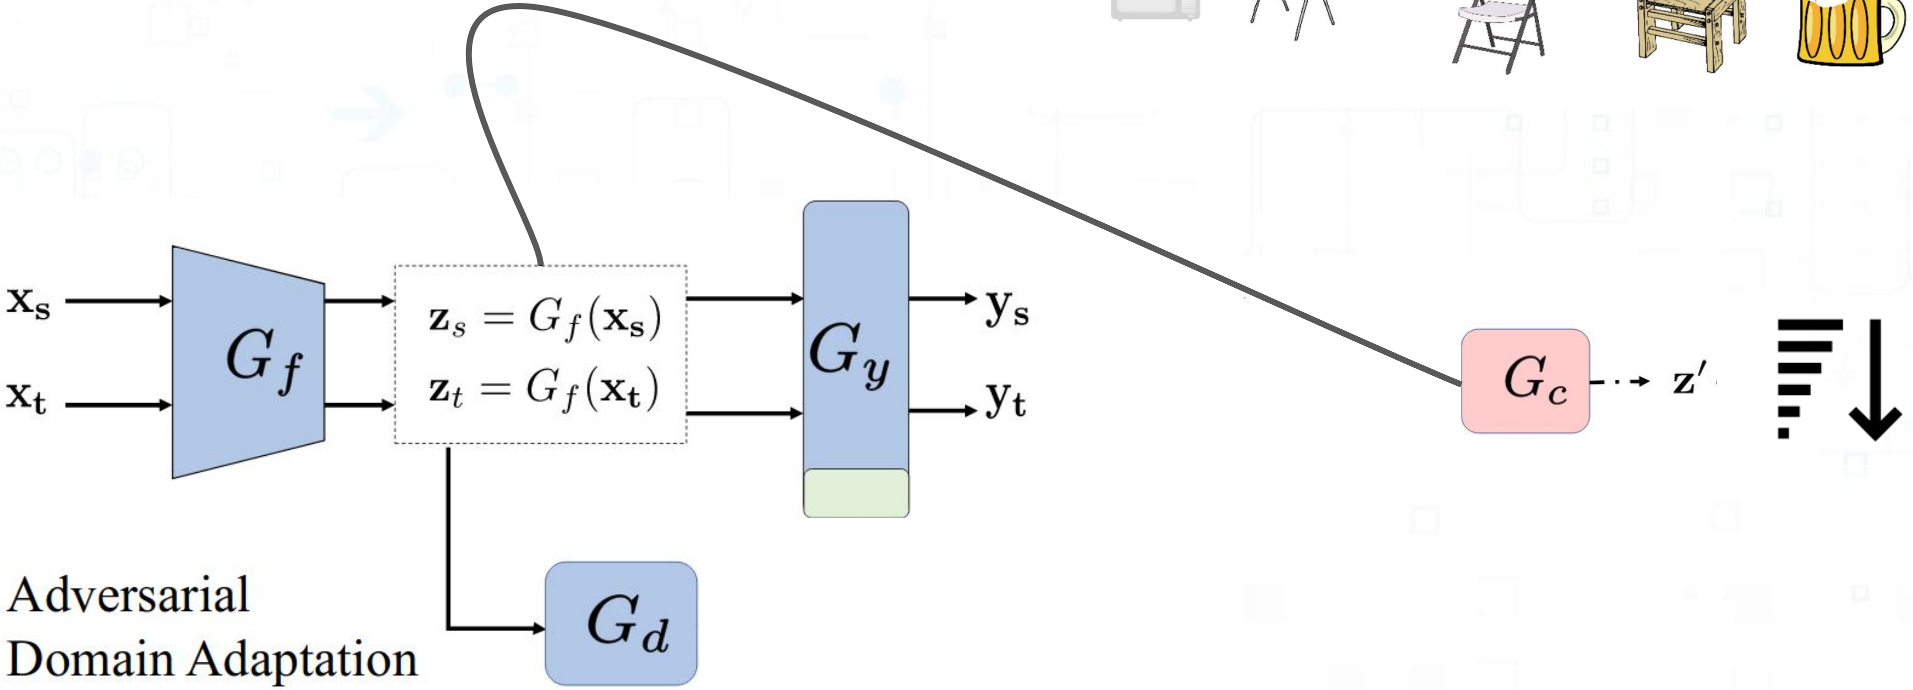

Open-Set DA

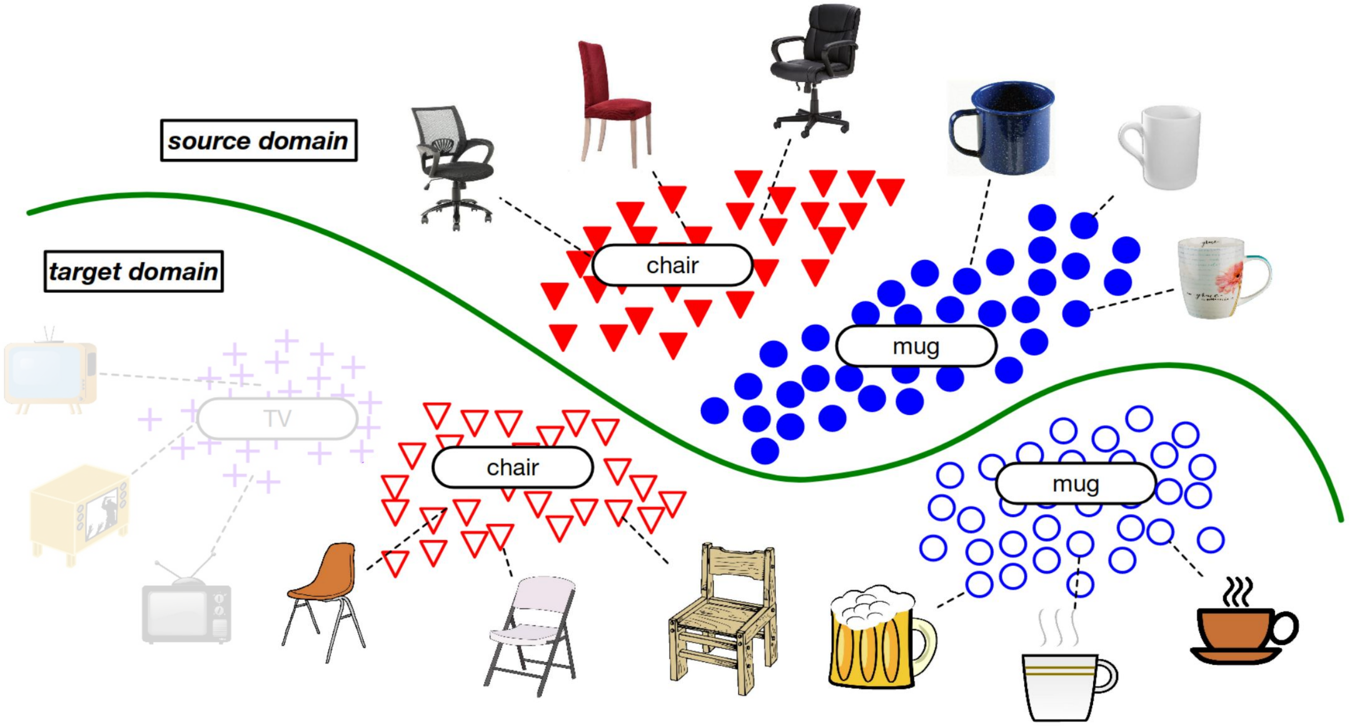

[Open Set Domain Adaptation by Backpropagation, ECCV 2018]
[Separate to Adapt: Open Set Domain Adaptation via Progressive Separation, CVPR 2019]

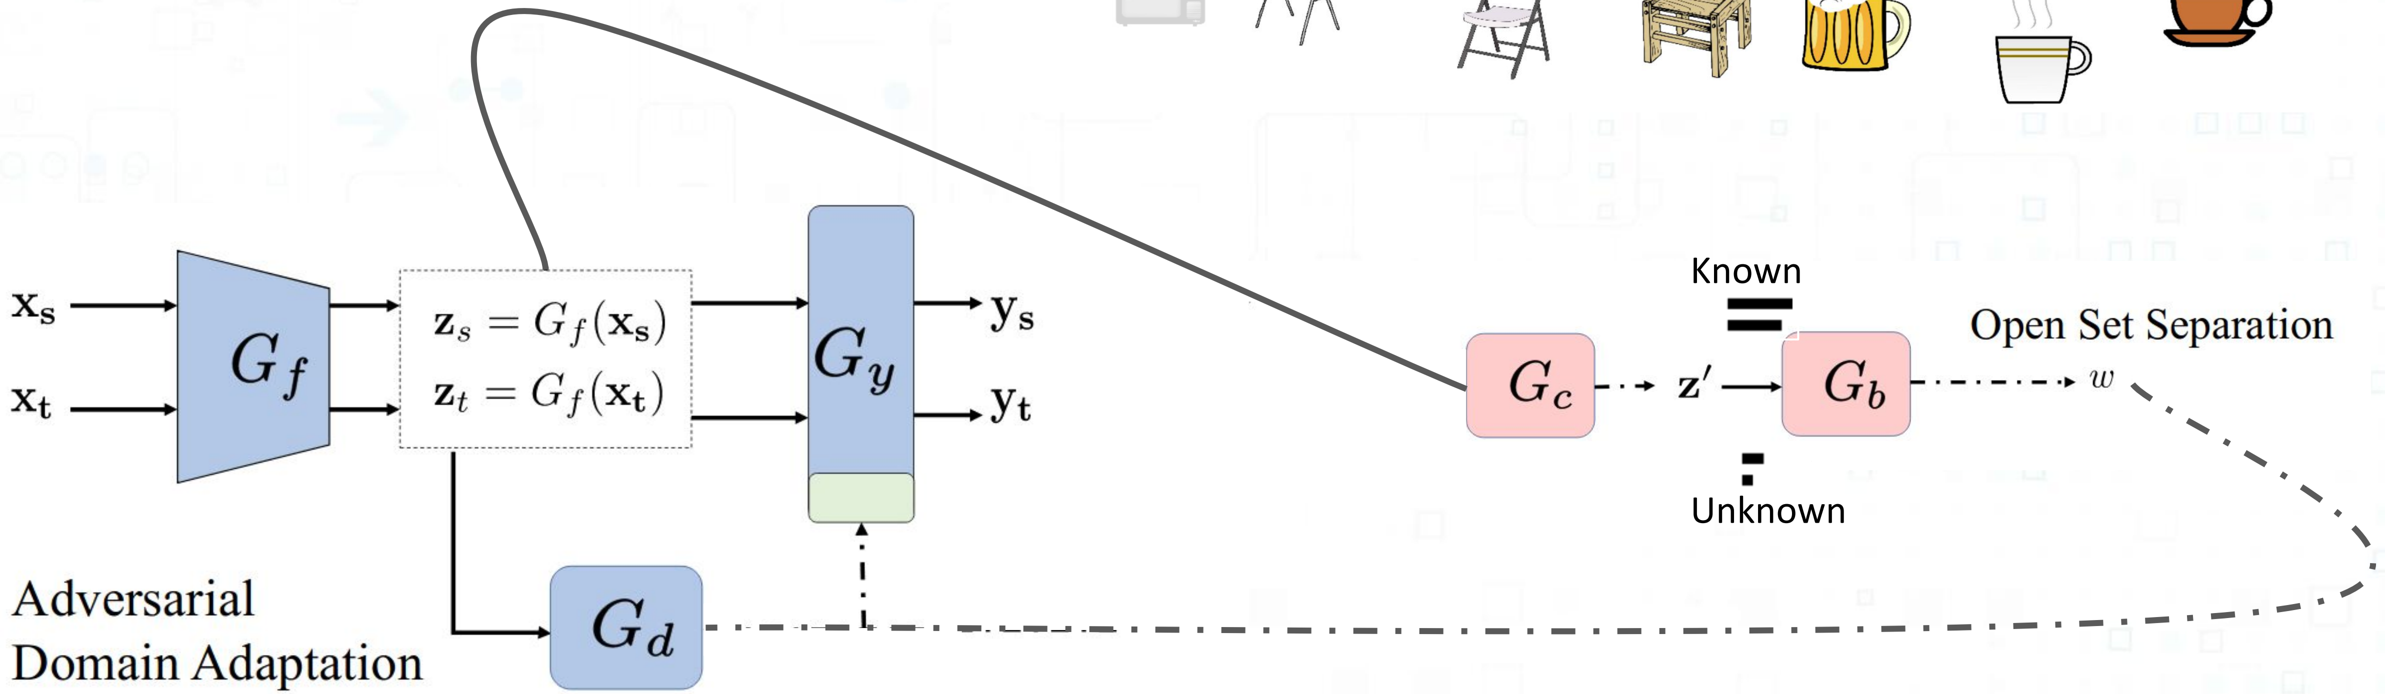

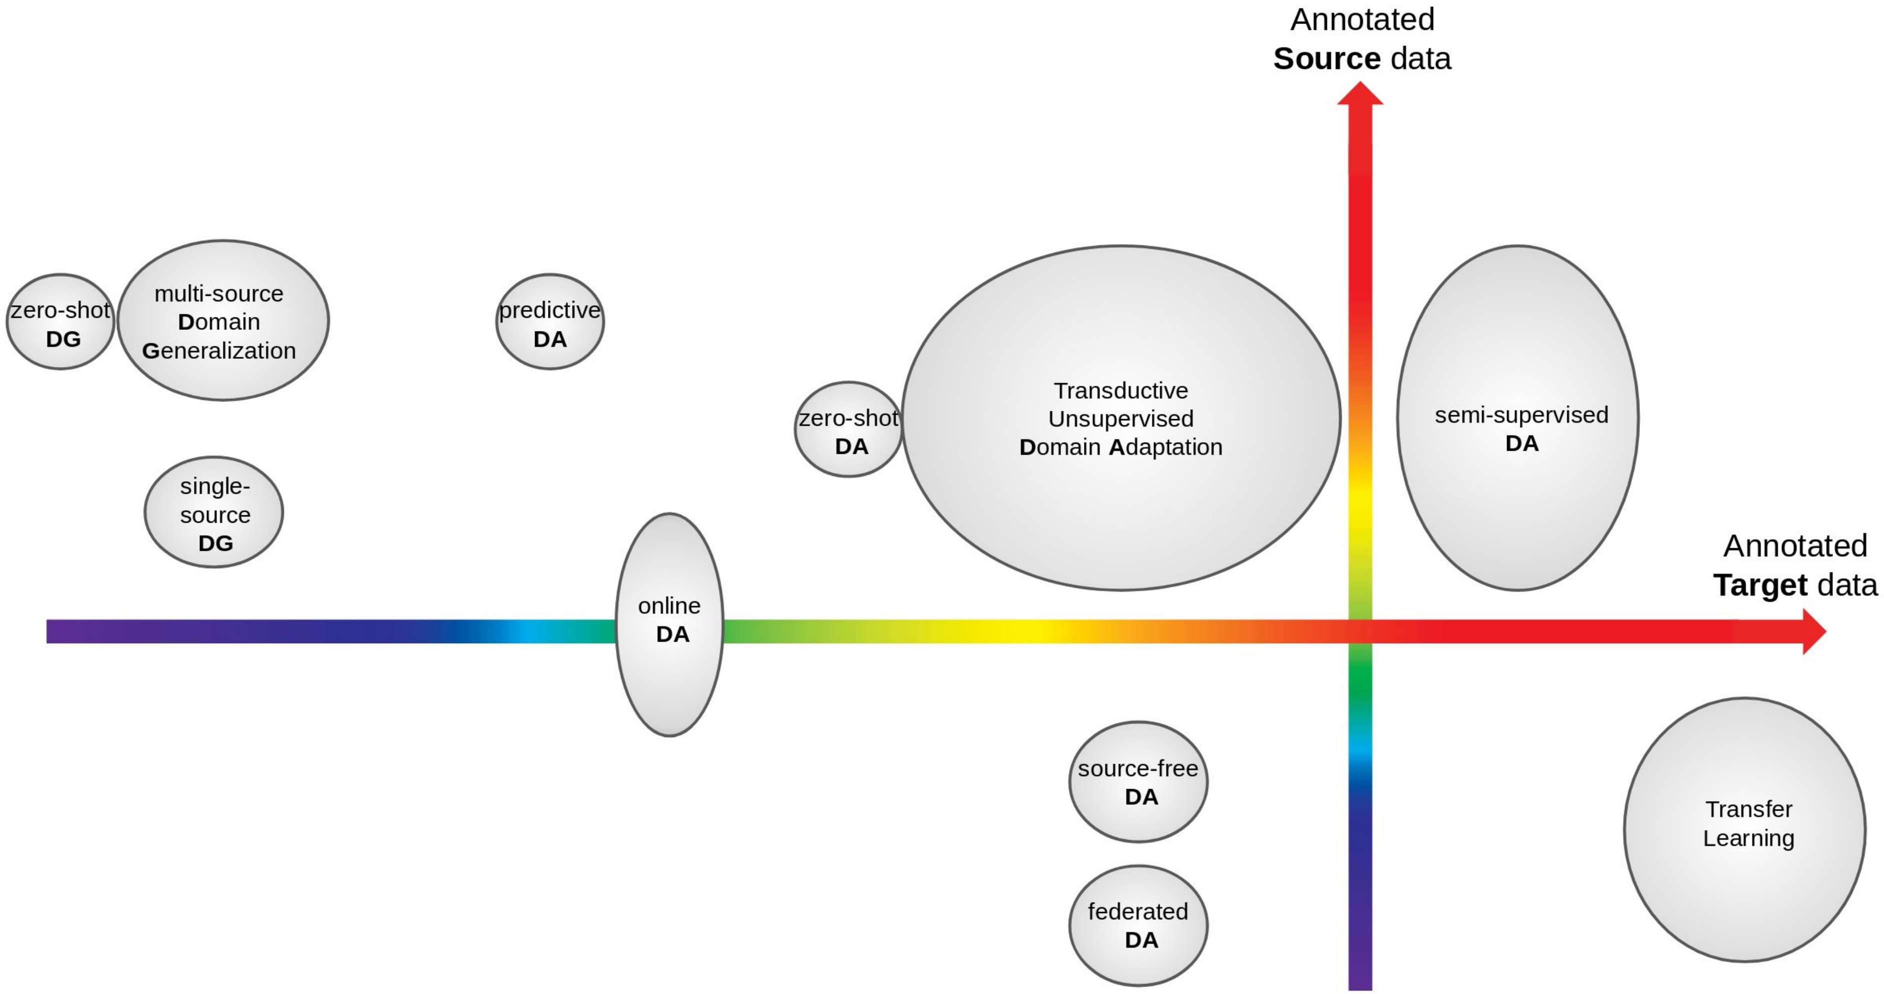
Annotated
Source data
Annotated
Target data
zero-shot
DG
multi-source
Domain
Generalization
predictive
DA
single-
source
DG
zero-shot
DA
Transductive
Unsupervised
Domain **A**daptation
semi-supervised
DA
online
DA
source-free
DA
federated
DA
Transfer
Learning

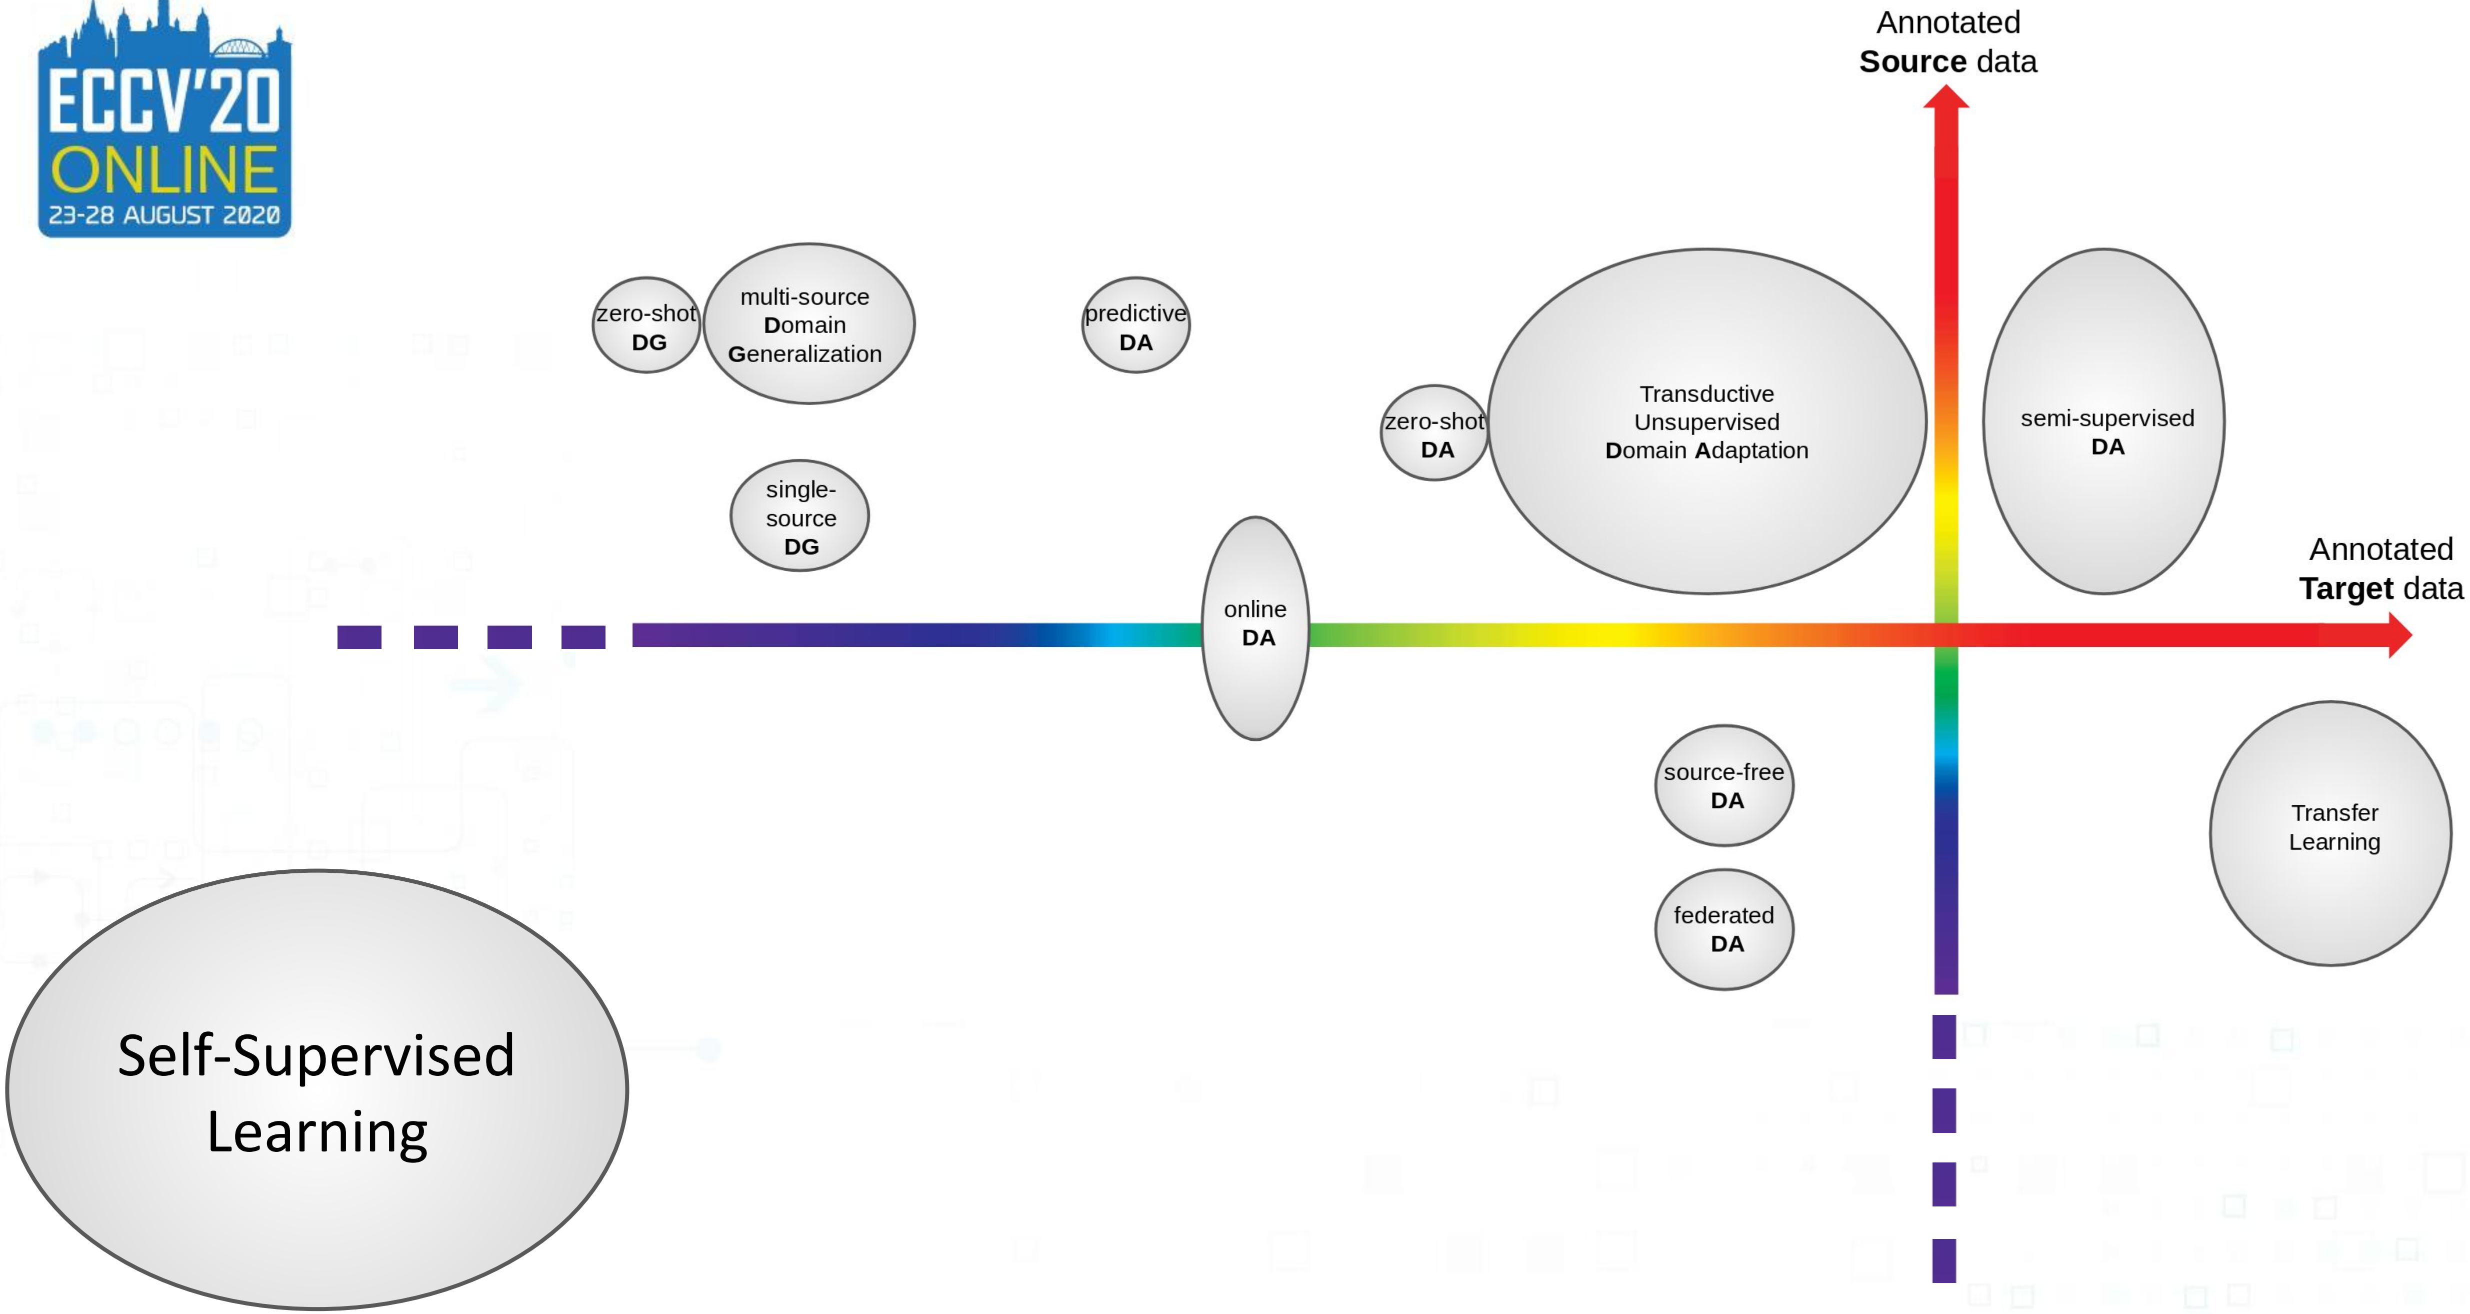

Annotated
Source data
Annotated
Target data
zero-shot
DG
multi-source
Domain
Generalization
predictive
DA
single-
source
DG
zero-shot
DA
Transductive
Unsupervised
Domain Adaptation
semi-supervised
DA
online
DA
source-free
DA
federated
DA
Transfer
Learning
Self-Supervised
Learning

Solve Jigsaw Puzzles

- Decompose an image in patches
- Shuffle them = remove their spatial co-location
- Ask a network to recompose the original image

Self-Supervised Learning

[Unsupervised learning of visual representations by solving jigsaw puzzles, CVPR 2016]

Recognize Image Orientation

- Rotate the image = remove the original orientation
- Ask a network to predict the rotation angle

Self-Supervised Learning

[Unsupervised Representation Learning by Predicting Image Rotations, ICLR 2018]

Self-Supervised Learning

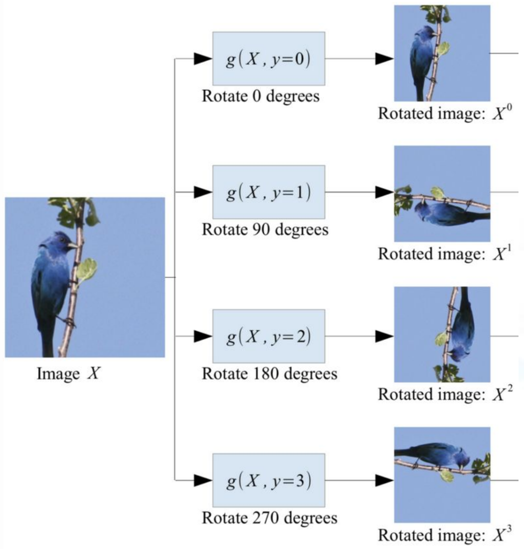

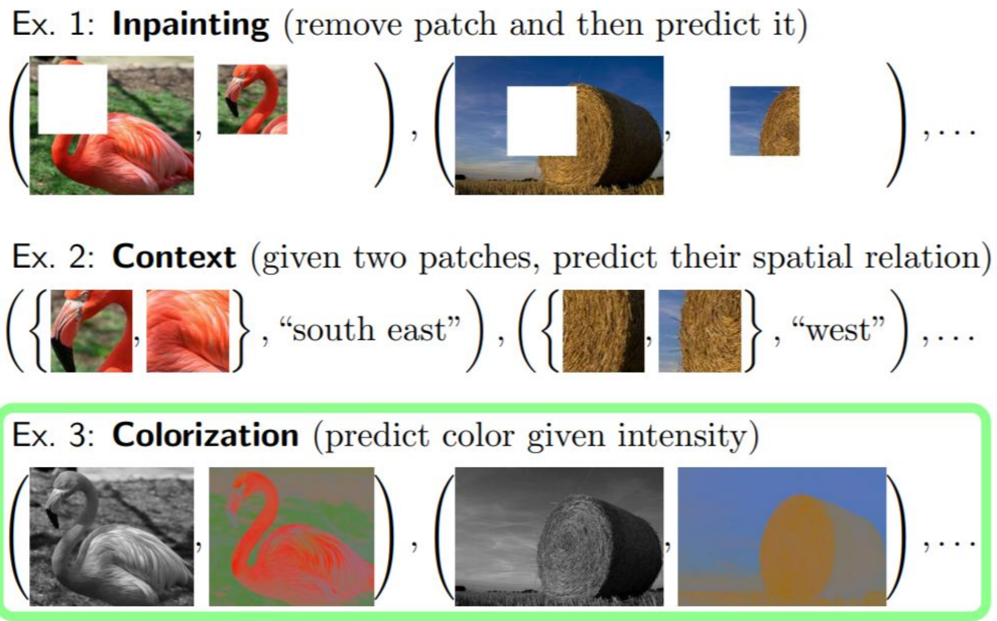

Self-Supervised Learning

[Steering Self-Supervised Feature Learning Beyond Local Pixel Statistics, CVPR 2020]
[Unsupervised learning of visual representations by solving jigsaw puzzles, CVPR 2016]
[Unsupervised representation learning by predicting image rotations, ICLR 2018]
[Colorization as a proxy task for visual understanding, CVPR 2017]
[Self-supervised feature learning by learning to spot artifacts, CVPR 2018]
[Colorization as a Proxy Task for Visual Understanding, CVPR 2017]
[Self-Supervised Feature Learning by Learning to Spot Artifacts, CVPR 2018]

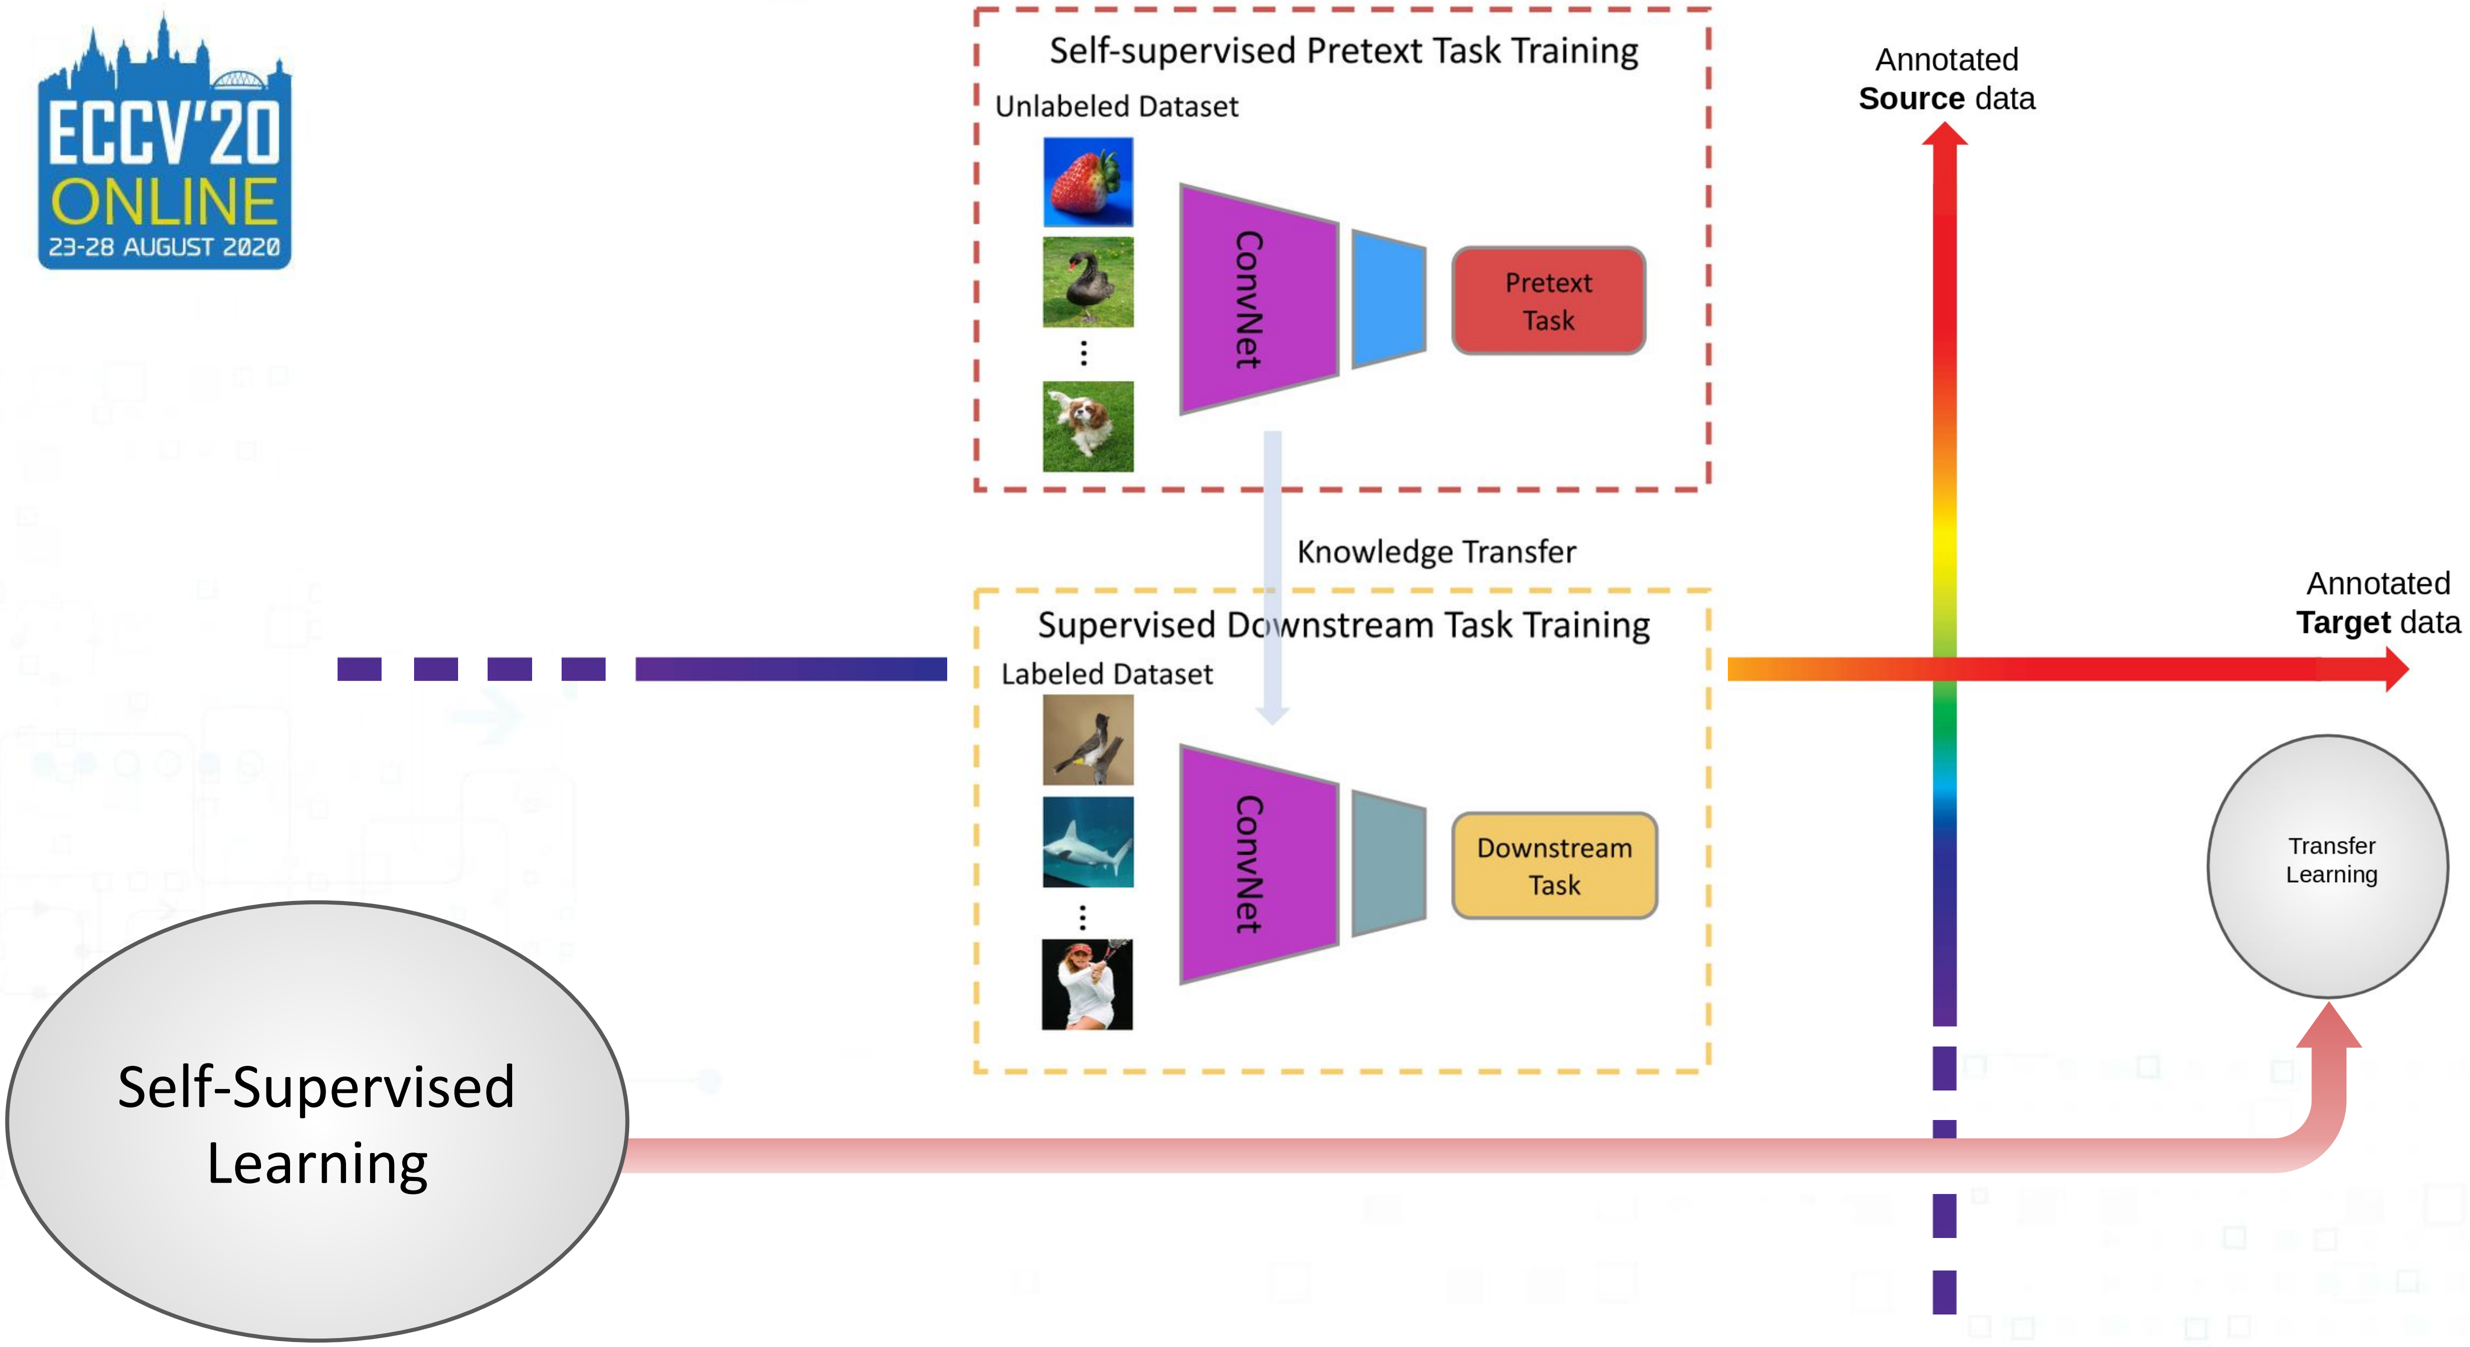

ECCV'20 ONLINE
23-28 AUGUST 2020
Self-supervised Pretext Task Training
Unlabeled Dataset
ConvNet
Pretext Task
Knowledge Transfer
Supervised Downstream Task Training
Labeled Dataset
ConvNet
Downstream Task
Annotated Source data
Annotated Target data
Transfer Learning
Self-Supervised Learning

[Domain Generalization by Solving Jigsaw Puzzles, CVPR 2019]
[Self-Supervised Learning Across Domains, ArXiv 2020]

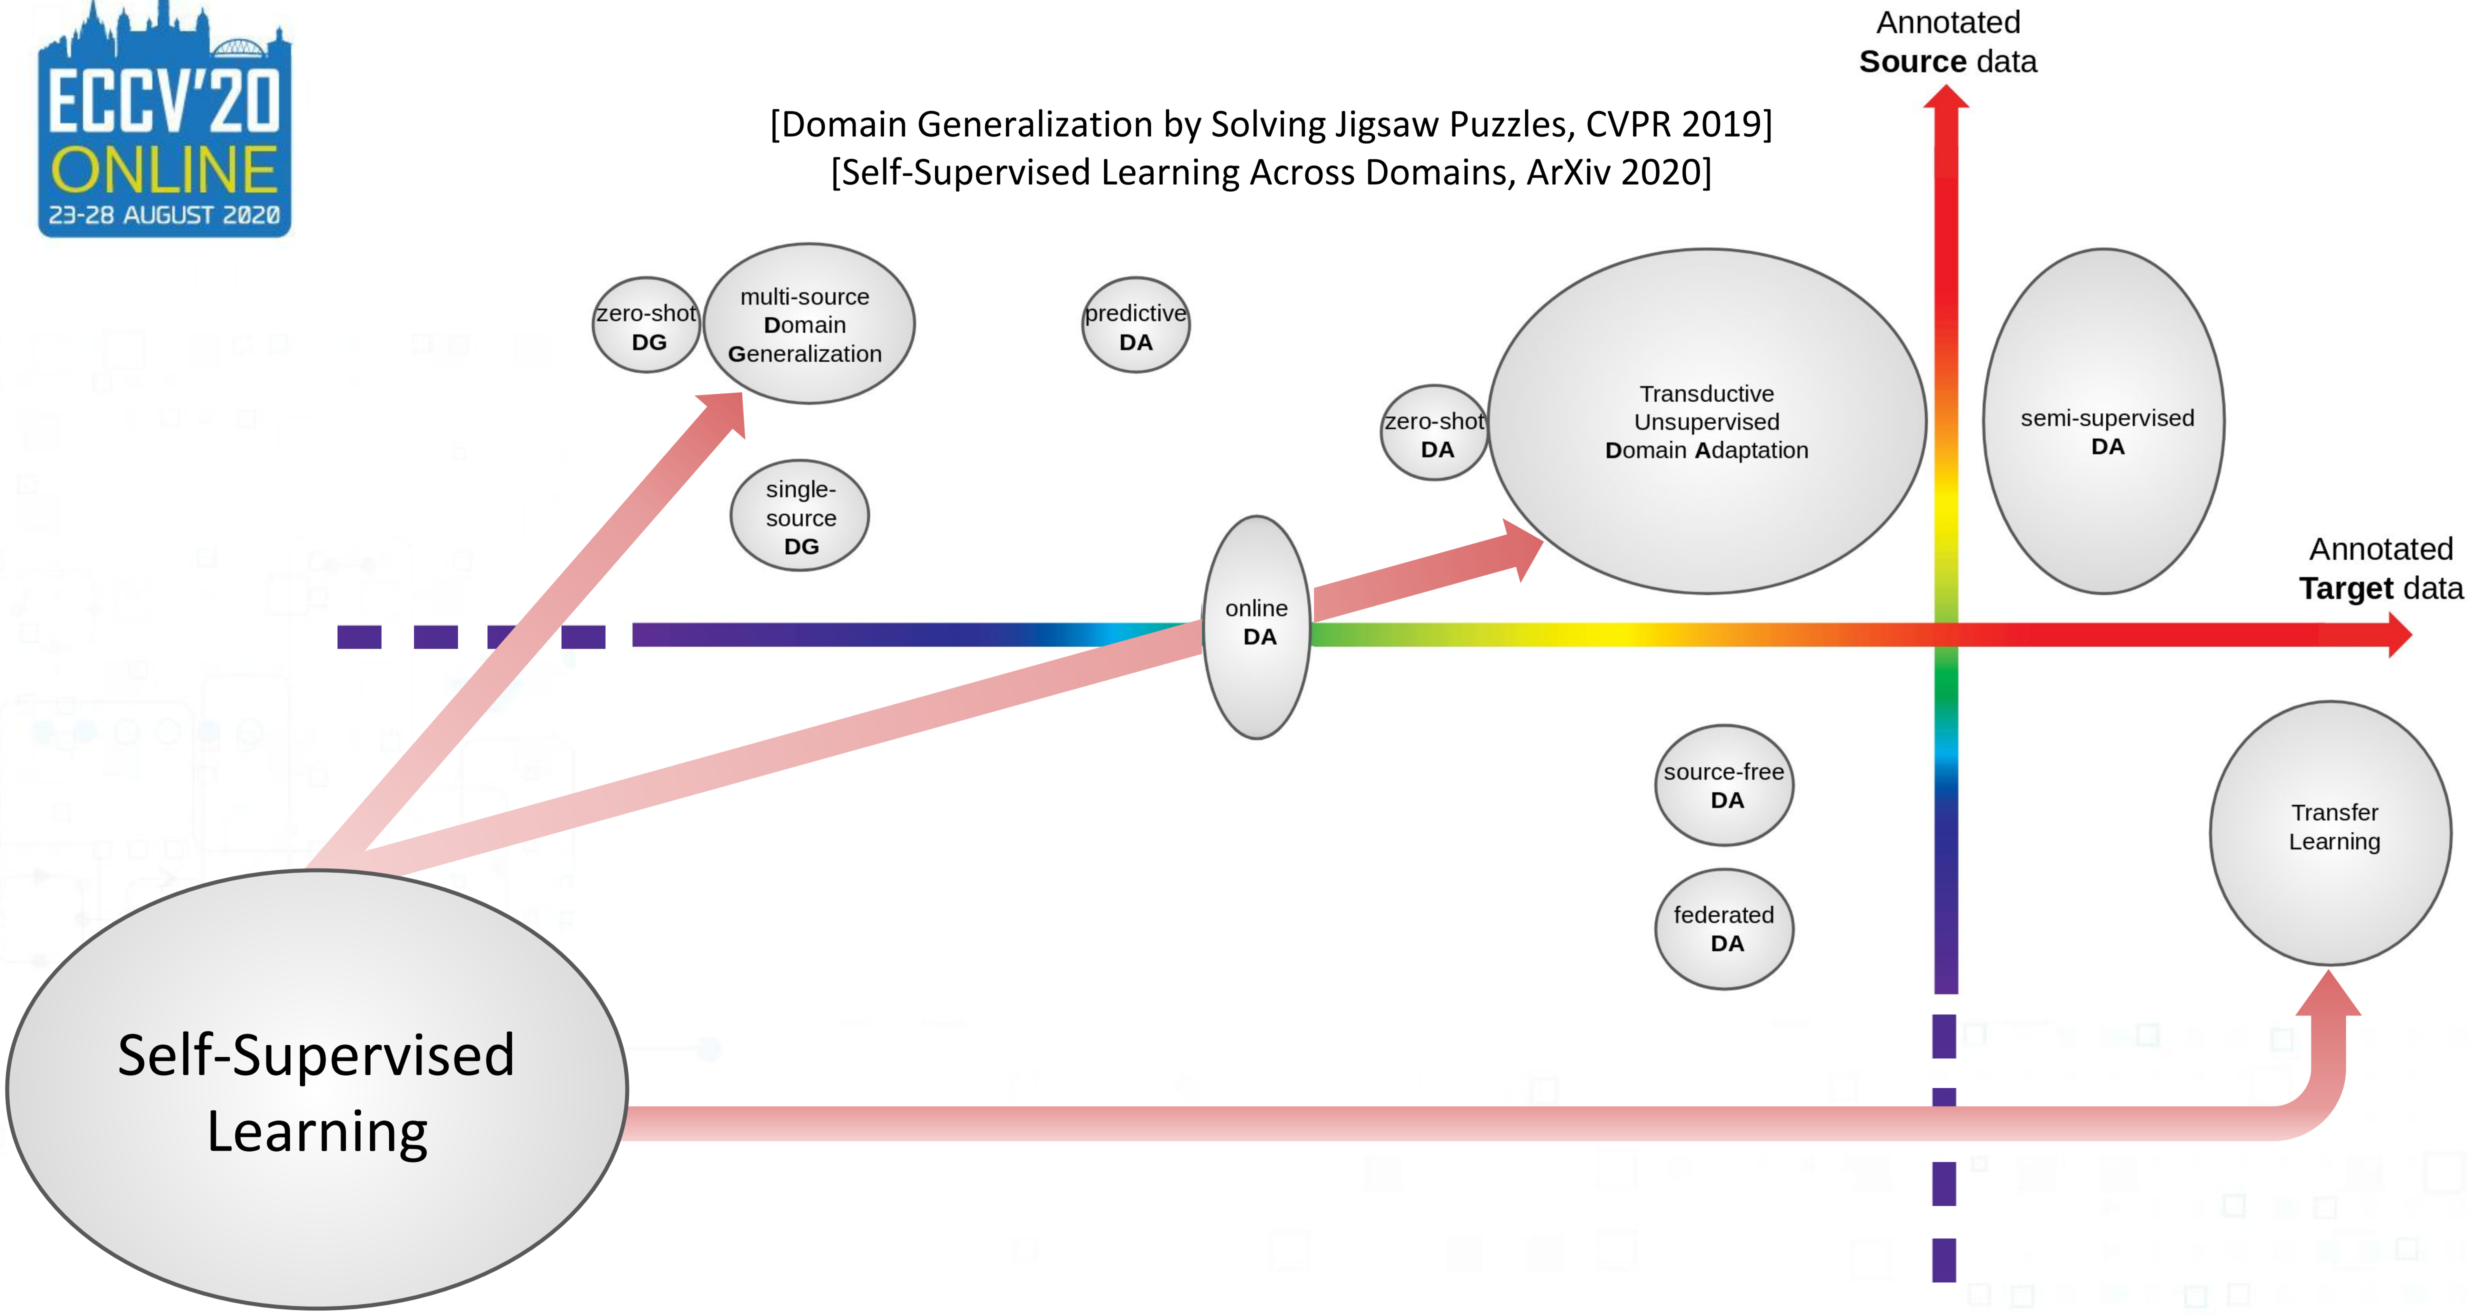

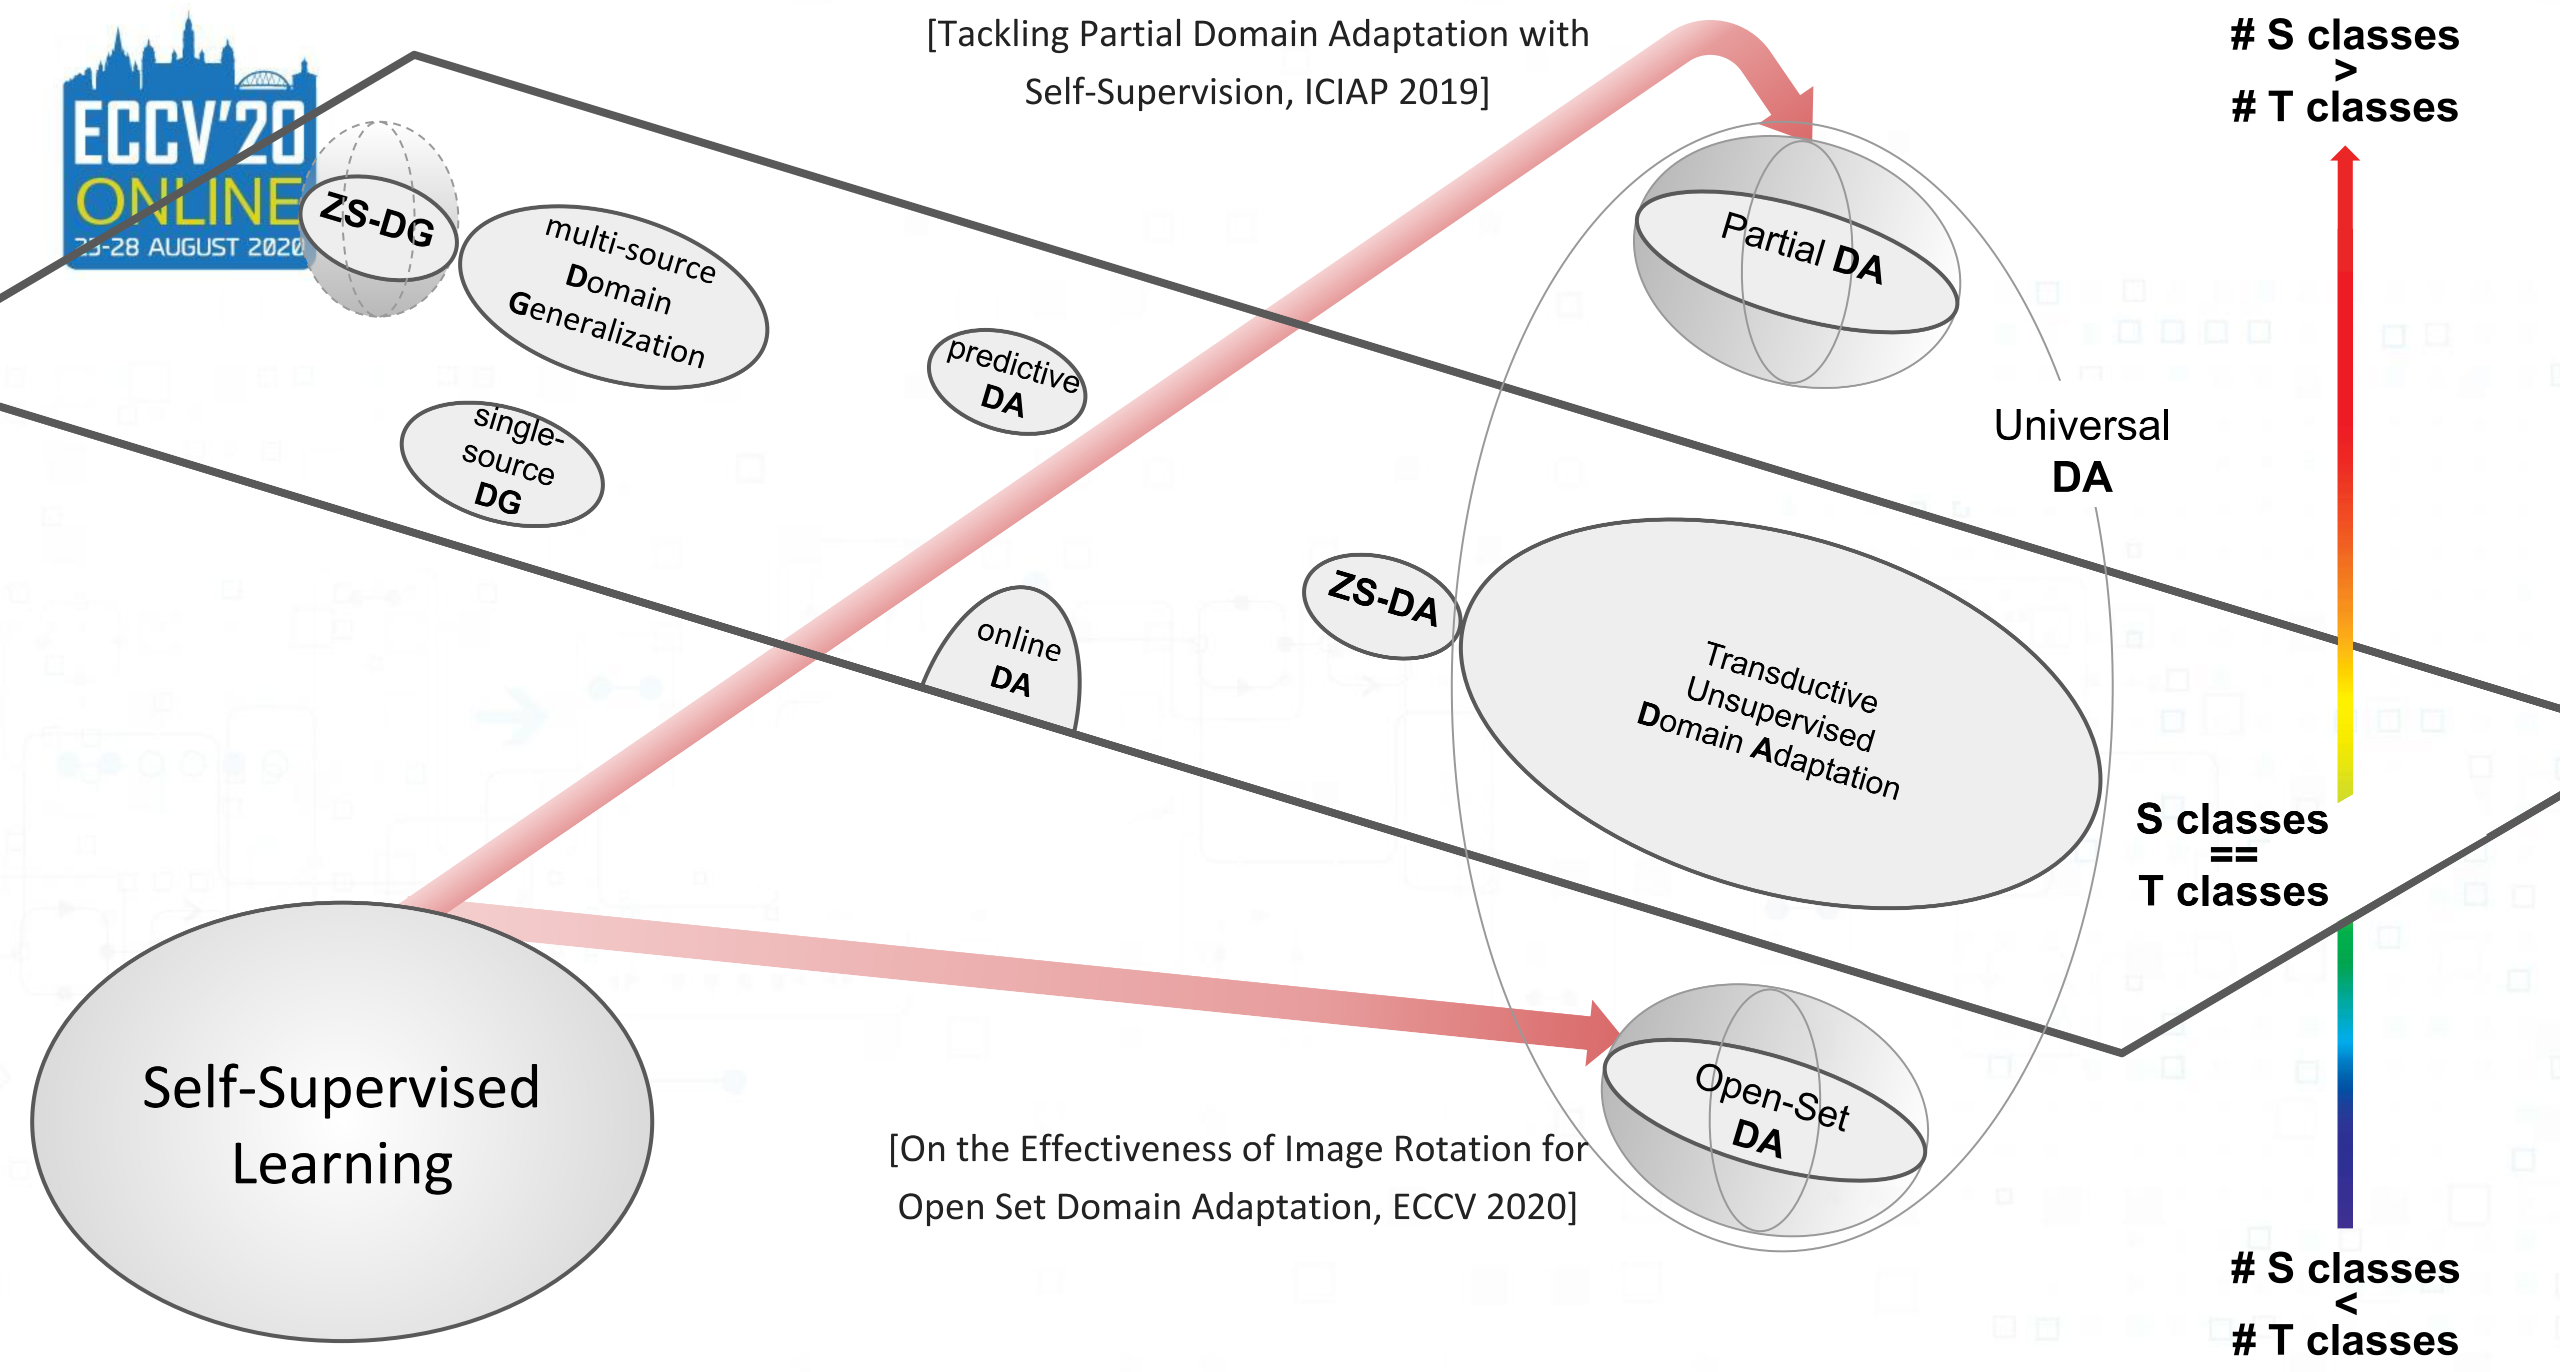
ECCV'20 ONLINE
23-28 AUGUST 2020
[Tackling Partial Domain Adaptation with
Self-Supervision, ICIAP 2019]
S classes
>
T classes
ZS-DG
multi-source Domain Generalization
predictive DA
single-source DG
Partial DA
Universal DA
ZS-DA
online DA
Transductive Unsupervised Domain Adaptation
S classes
==
T classes
Self-Supervised Learning
Open-Set DA
[On the Effectiveness of Image Rotation for
Open Set Domain Adaptation, ECCV 2020]
S classes
<
T classes

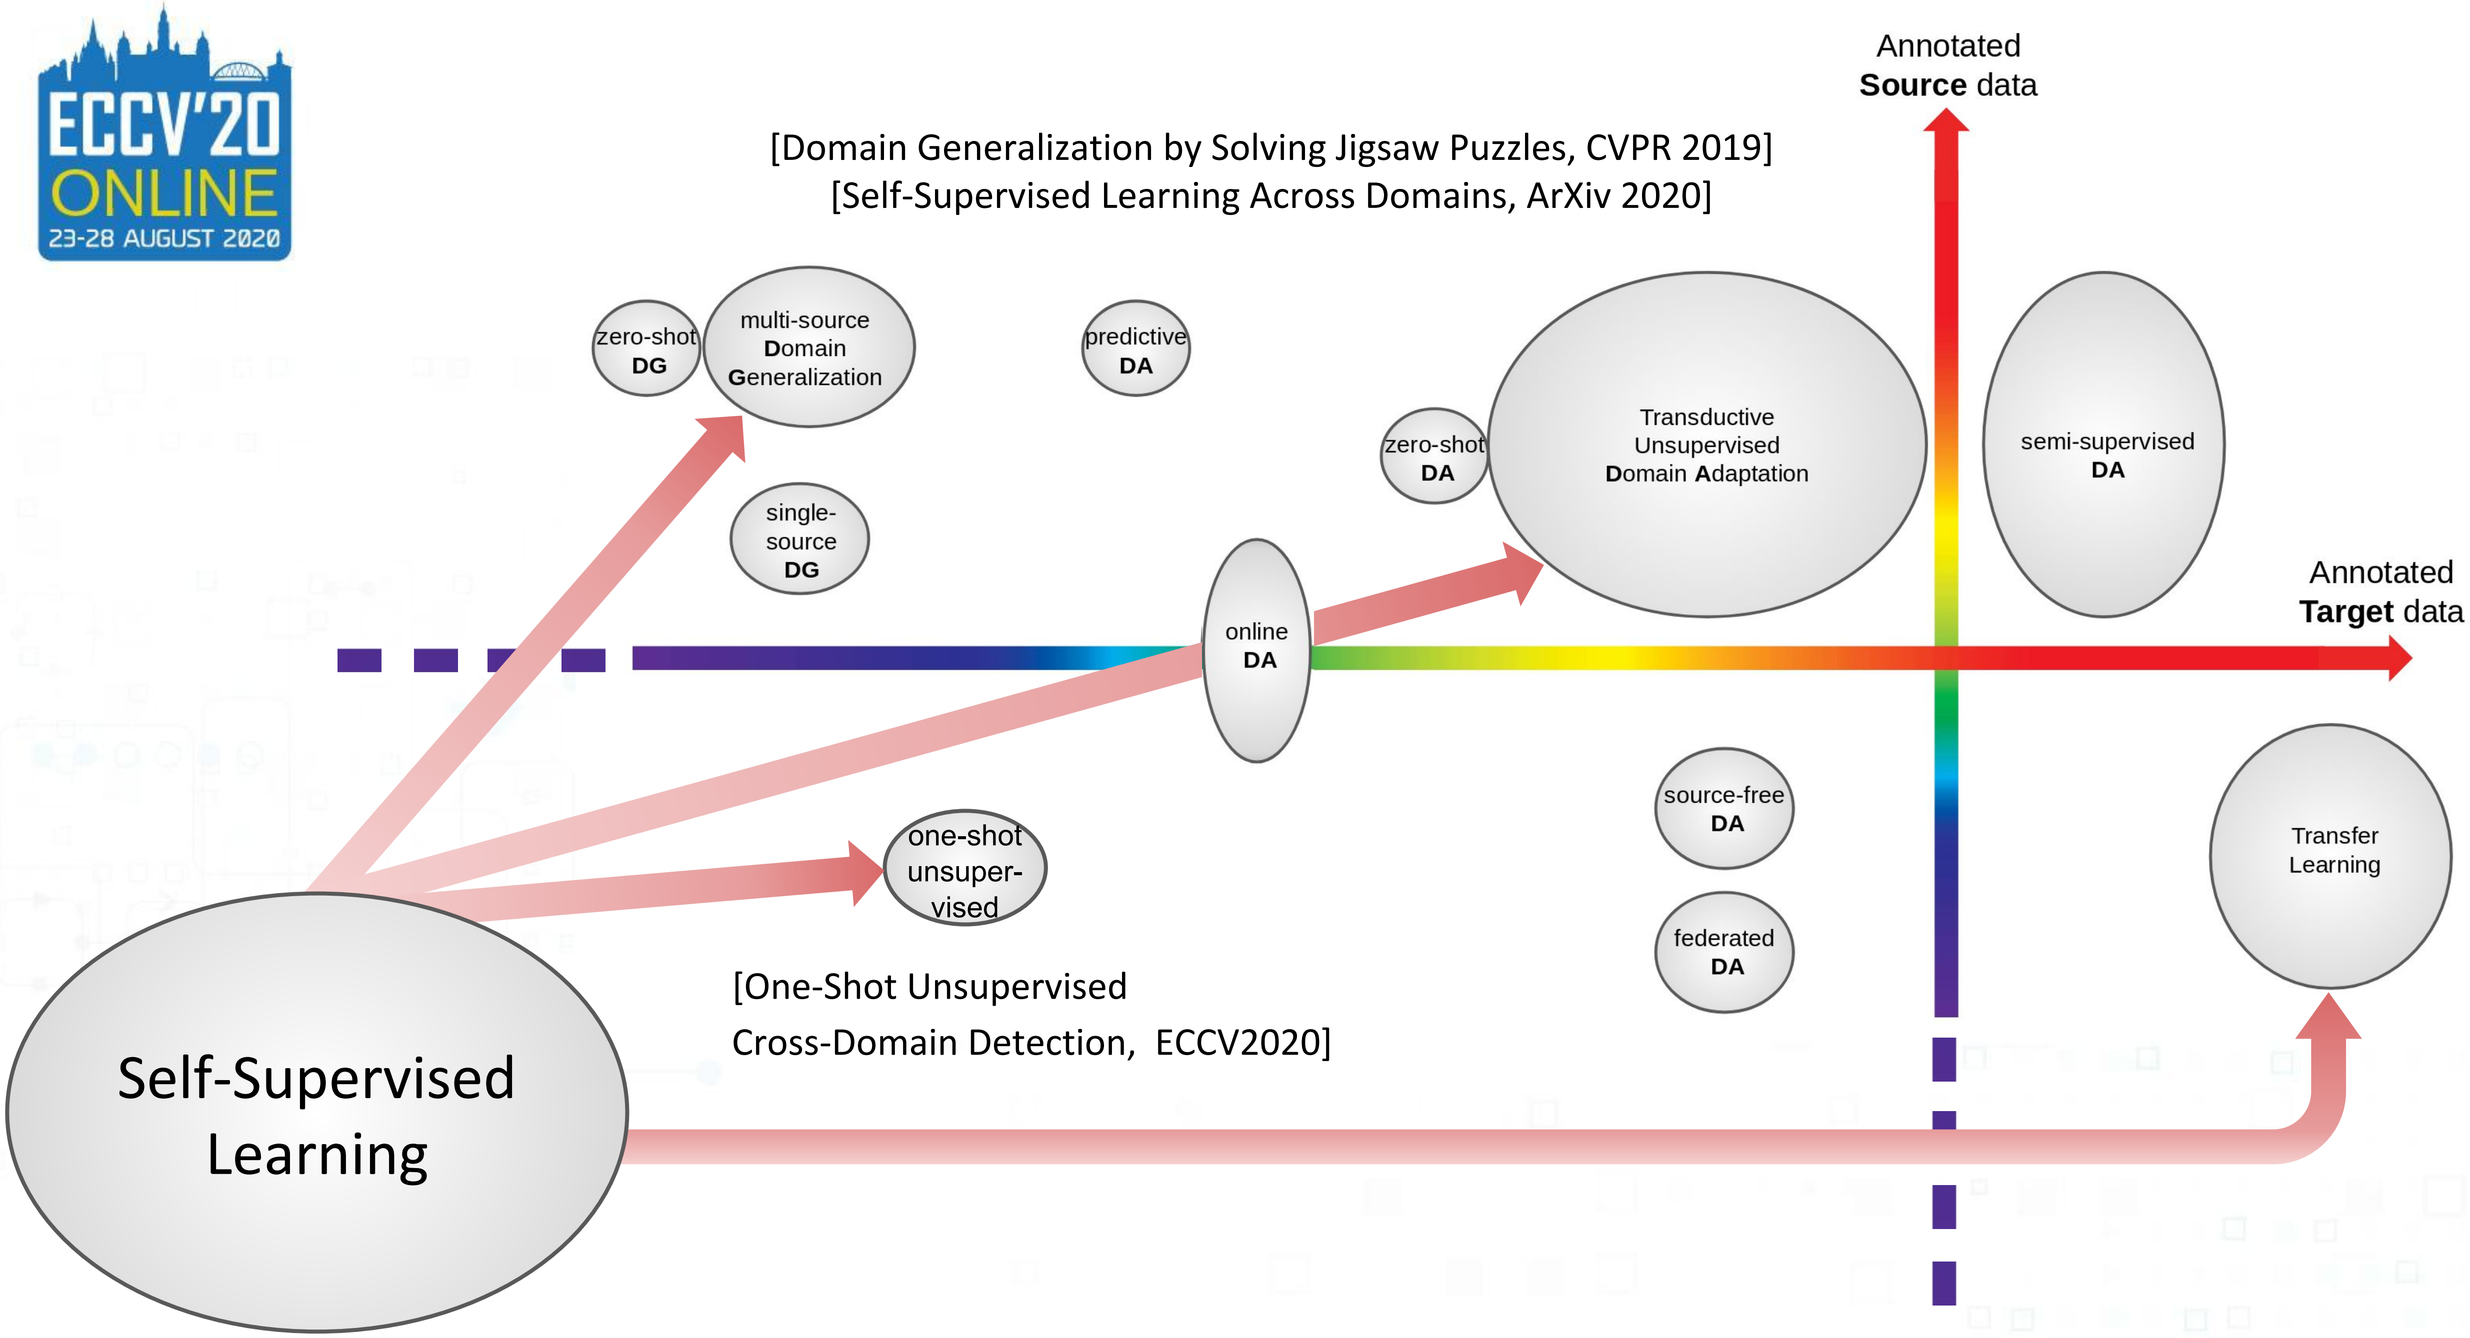

ECCV'20 ONLINE
23-28 AUGUST 2020
[Domain Generalization by Solving Jigsaw Puzzles, CVPR 2019]
[Self-Supervised Learning Across Domains, ArXiv 2020]
Annotated Source data
Annotated Target data
zero-shot DG
multi-source Domain Generalization
single-source DG
predictive DA
zero-shot DA
Transductive Unsupervised Domain Adaptation
semi-supervised DA
online DA
one-shot unsuper-vised
source-free DA
federated DA
Transfer Learning
Self-Supervised Learning
[One-Shot Unsupervised Cross-Domain Detection, ECCV2020]

Self-Supervision + Domain Generalization

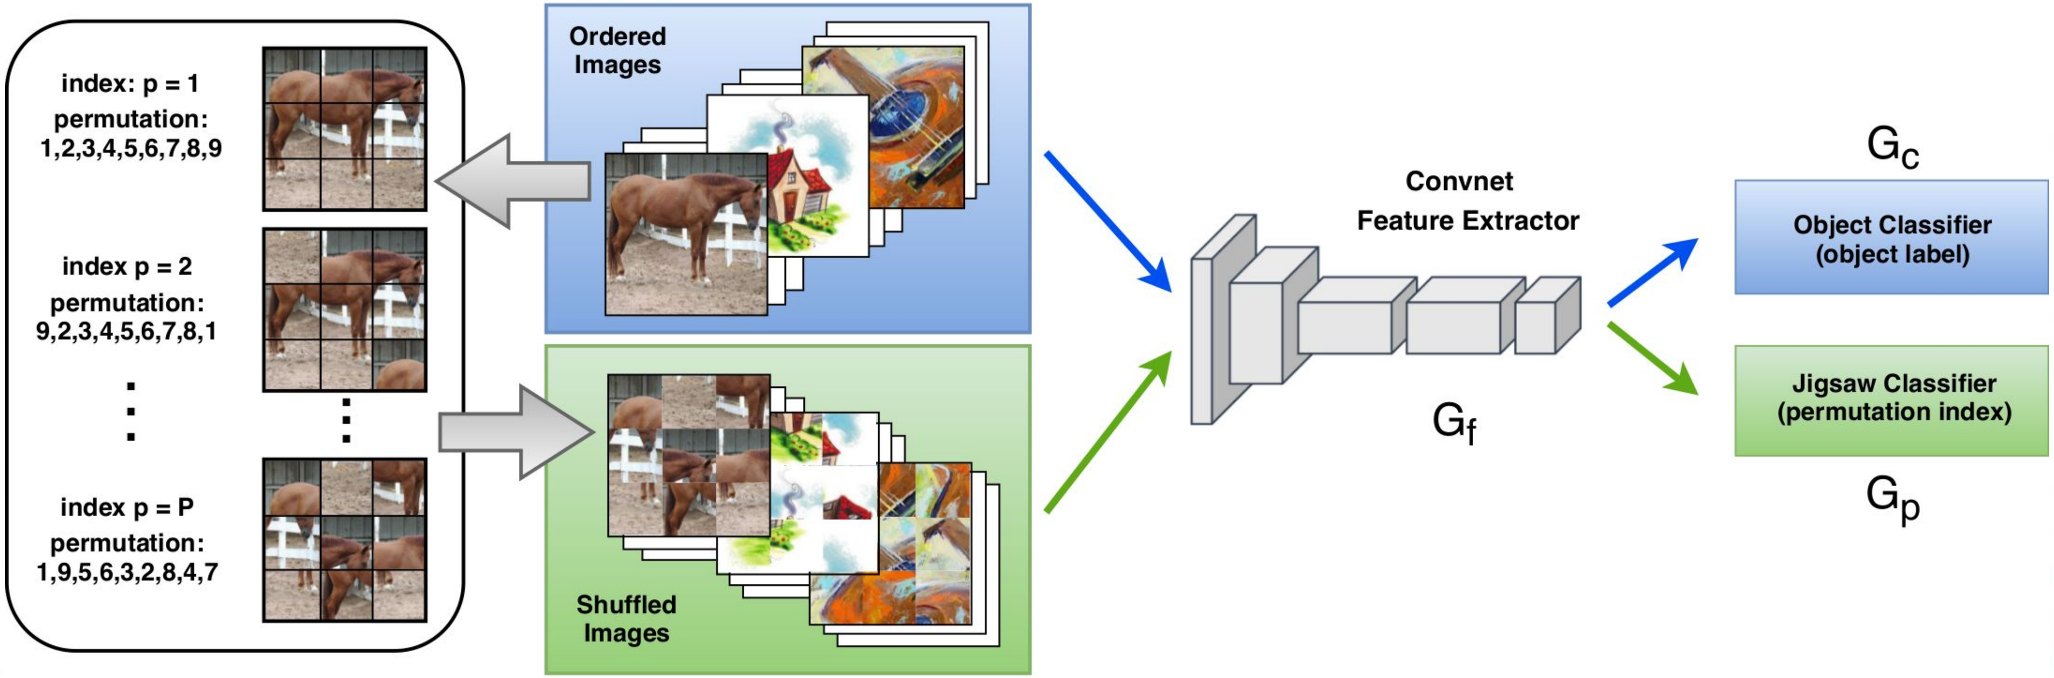

$$\arg\min_{\theta_f,\theta_c,\theta_p} \frac{1}{n^s}\sum_{i=1}^{n^s}\mathcal{L}_c(G_c(G_f(\mathbf{x}_i^s)), y_i^s)+$$
$$\alpha_s \frac{1}{K^s}\sum_{k=1}^{K^s}\mathcal{L}_p(G_p(G_f(\mathbf{z}_k^s)), p_k^s)$$

[Domain Generalization by Solving Jigsaw Puzzles, CVPR 2019]
[Self-Supervised Learning Across Domains, ArXiv 2020]

Self-Supervision + Domain Generalization

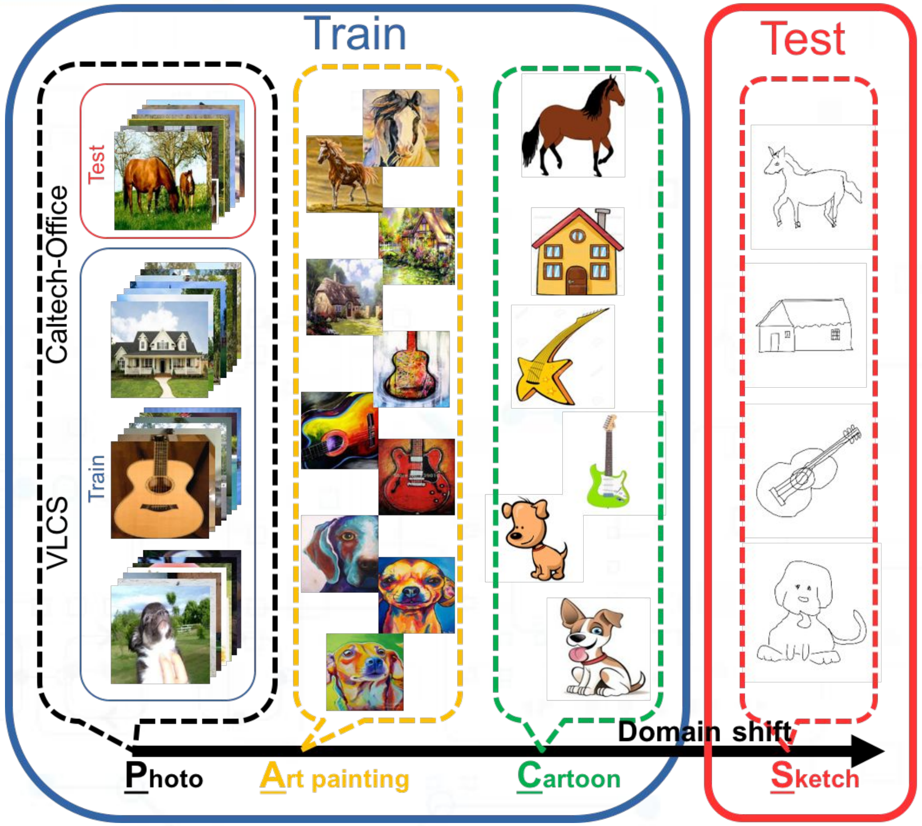

[Domain Generalization by Solving Jigsaw Puzzles, CVPR 2019]
[Self-Supervised Learning Across Domains, ArXiv 2020]
[Deeper, Broader and Artier Domain Generalization, ICCV 2017]

PACS	art_paint.	cartoon	sketches	photo	Avg.
Alexnet					
DeepAll	63.30	63.13	54.07	87.70	67.05
TF	62.86	66.97	57.51	89.50	69.21
DeepAll	64.44	72.07	58.07	87.50	70.52
D-SAM	63.87	70.70	64.66	85.55	71.20
DeepAll	63.40	66.10	56.60	88.50	68.70
Epi-FCR	64.70	72.30	65.00	86.10	72.00
DeepAll	64.91	64.28	53.08	86.67	67.24
MLDG	66.23	66.88	58.96	88.00	70.01
DeepAll	67.21	66.12	55.32	88.47	69.28
MetaReg	69.82	70.35	59.26	91.07	72.62
DeepAll	63.30	63.10	54.00	87.70	67.03
PAR	68.70	70.50	64.60	90.40	73.54
DeepAll	68.09	70.23	61.80	88.86	72.25
MMLD	66.99	70.64	67.78	89.35	73.69
DeepAll	66.50	69.65	61.42	89.68	71.81±0.26
Jigsaw	67.76	70.79	64.01	89.64	73.05±0.20
Rotation	69.43	69.40	65.20	89.17	73.30±0.47
Jigsaw+Rotation	69.70	71.00	66.00	89.60	74.08±0.32

Self-Supervision + Domain Generalization

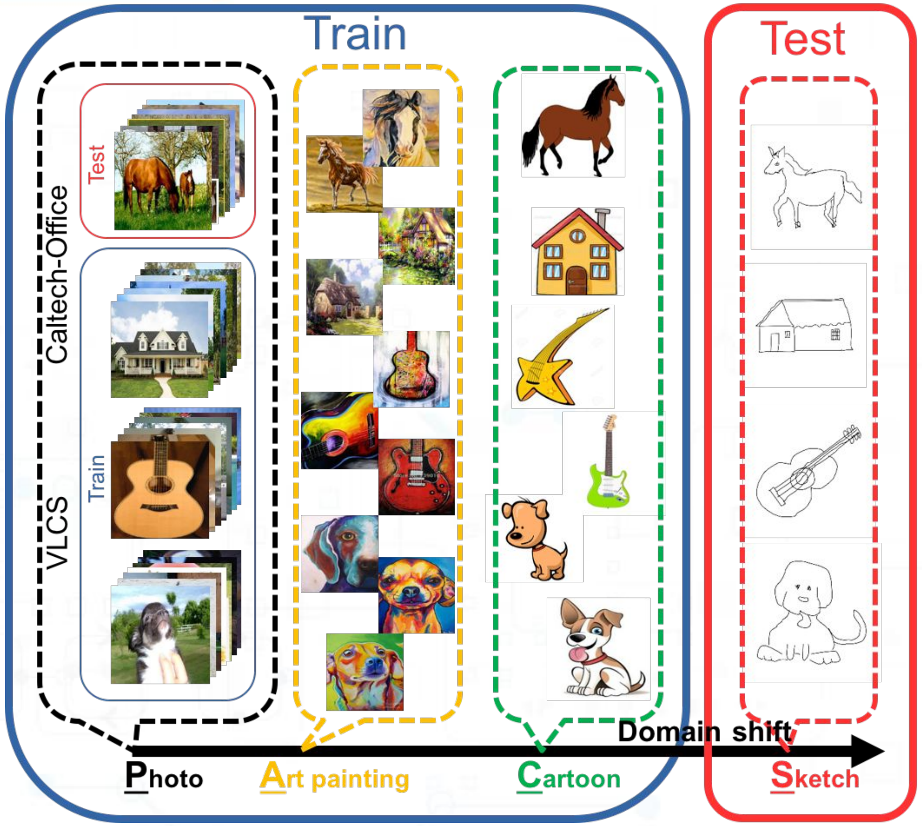

[Domain Generalization by Solving Jigsaw Puzzles, CVPR 2019]
[Self-Supervised Learning Across Domains, ArXiv 2020]
[Deeper, Broader and Artier Domain Generalization, ICCV 2017]

PACS	art_paint.	cartoon	sketches	photo	Avg.
Alexnet					
DeepAll	63.30	63.13	54.07	87.70	67.05
TF	62.86	66.97	57.51	89.50	69.21
DeepAll	64.44	72.07	58.07	87.50	70.52
D-SAM	63.87	70.70	64.66	85.55	71.20
DeepAll	63.40	66.10	56.60	88.50	68.70
Epi-FCR	64.70	72.30	65.00	86.10	72.00
DeepAll	64.91	64.28	53.08	86.67	67.24
MLDG	66.23	66.88	58.96	88.00	70.01
DeepAll	67.21	66.12	55.32	88.47	69.28
MetaReg	69.82	70.35	59.26	91.07	72.62
DeepAll	63.30	63.10	54.00	87.70	67.03
PAR	68.70	70.50	64.60	90.40	73.54
DeepAll	68.09	70.23	61.80	88.86	72.25
MMLD	66.99	70.64	67.78	89.35	73.69
DeepAll	66.50	69.65	61.42	89.68	71.81±0.26
Jigsaw	67.76	70.79	64.01	89.64	73.05±0.20
Rotation	69.43	69.40	65.20	89.17	73.30±0.47
Jigsaw+Rotation	69.70	71.00	66.00	89.60	74.08±0.32

Self-Supervision + Domain Generalization

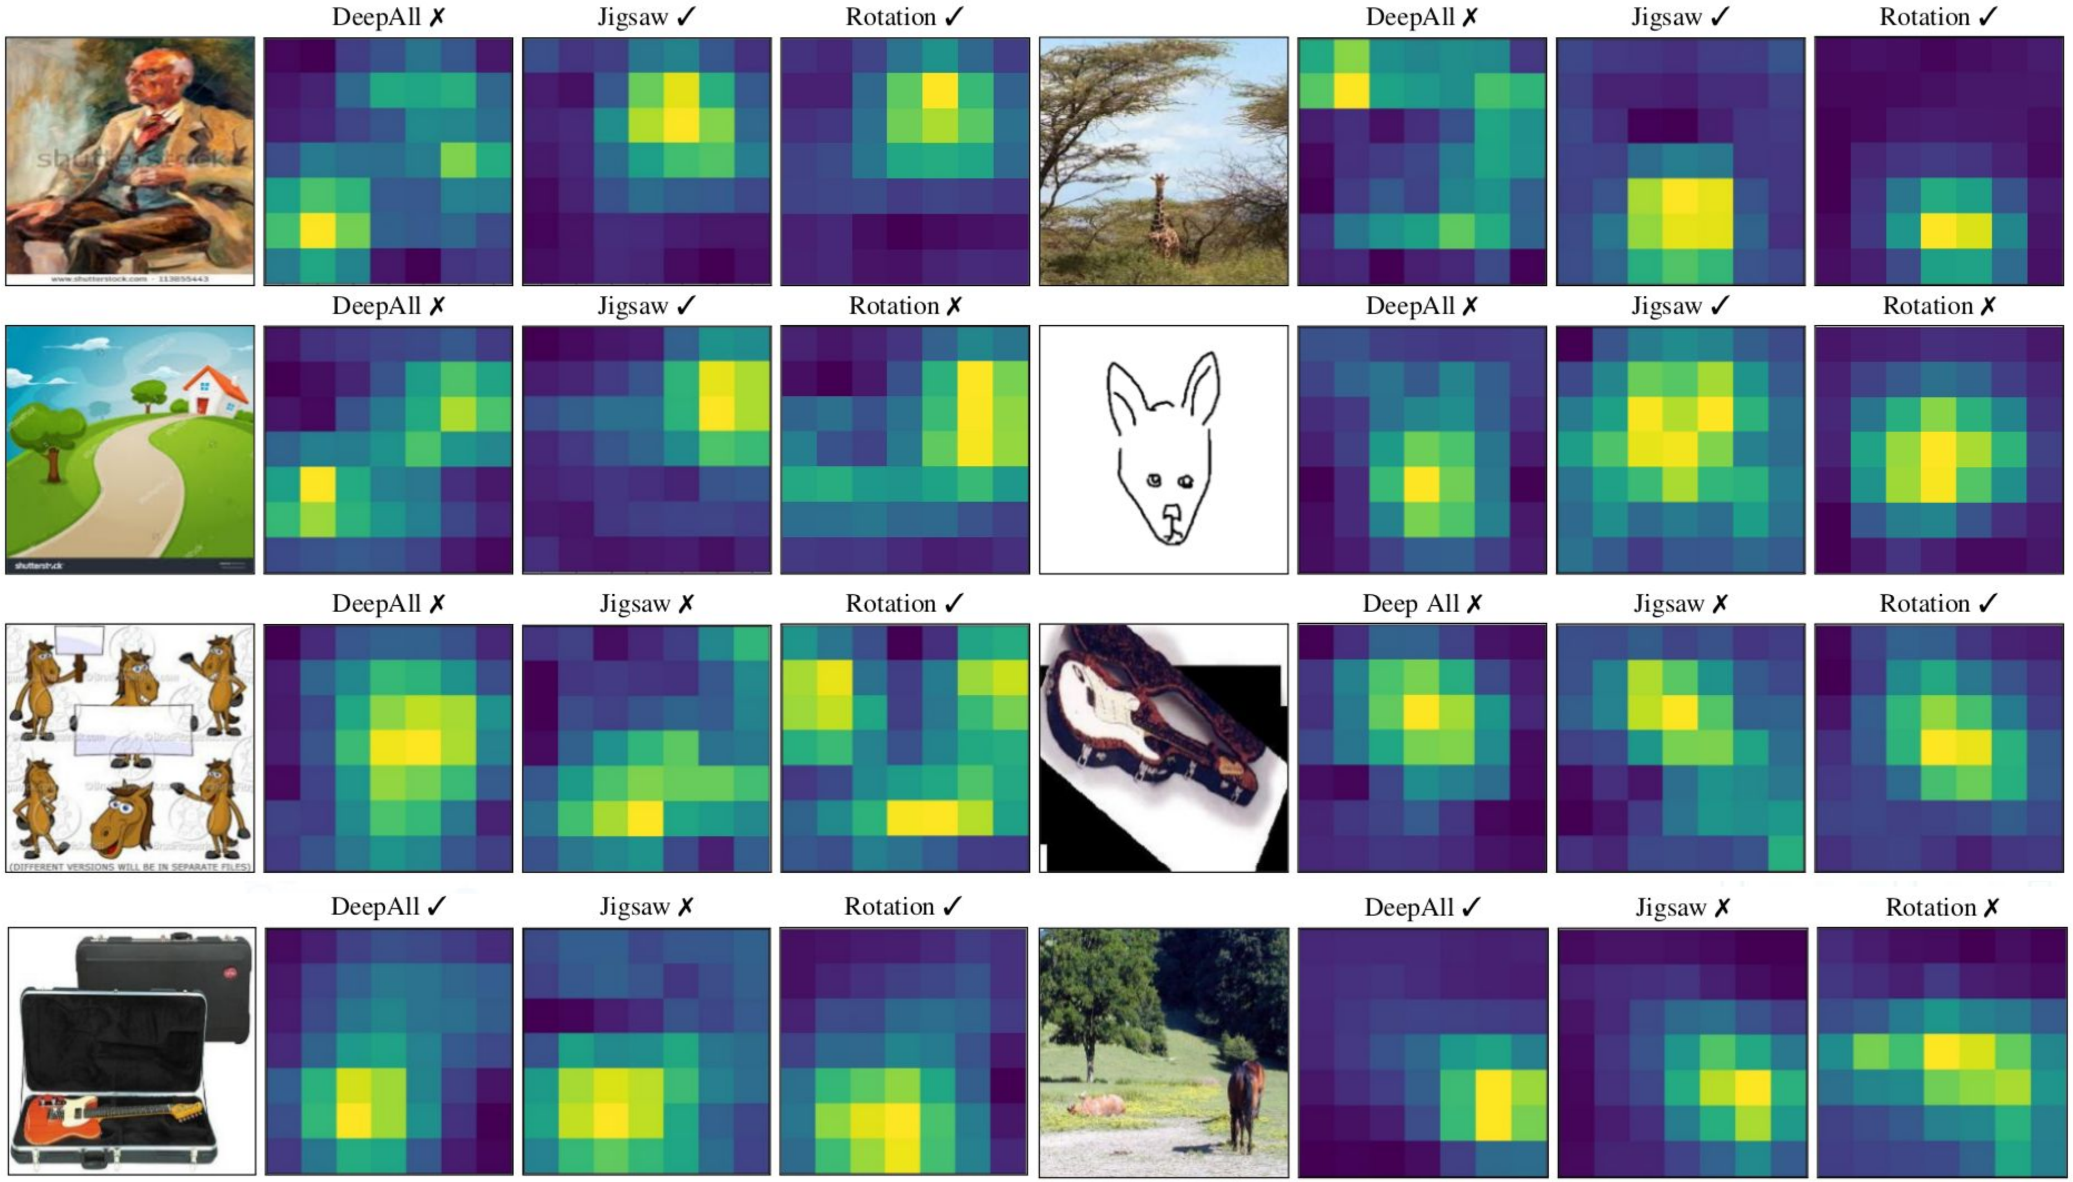

Self-Supervision + Domain Generalization

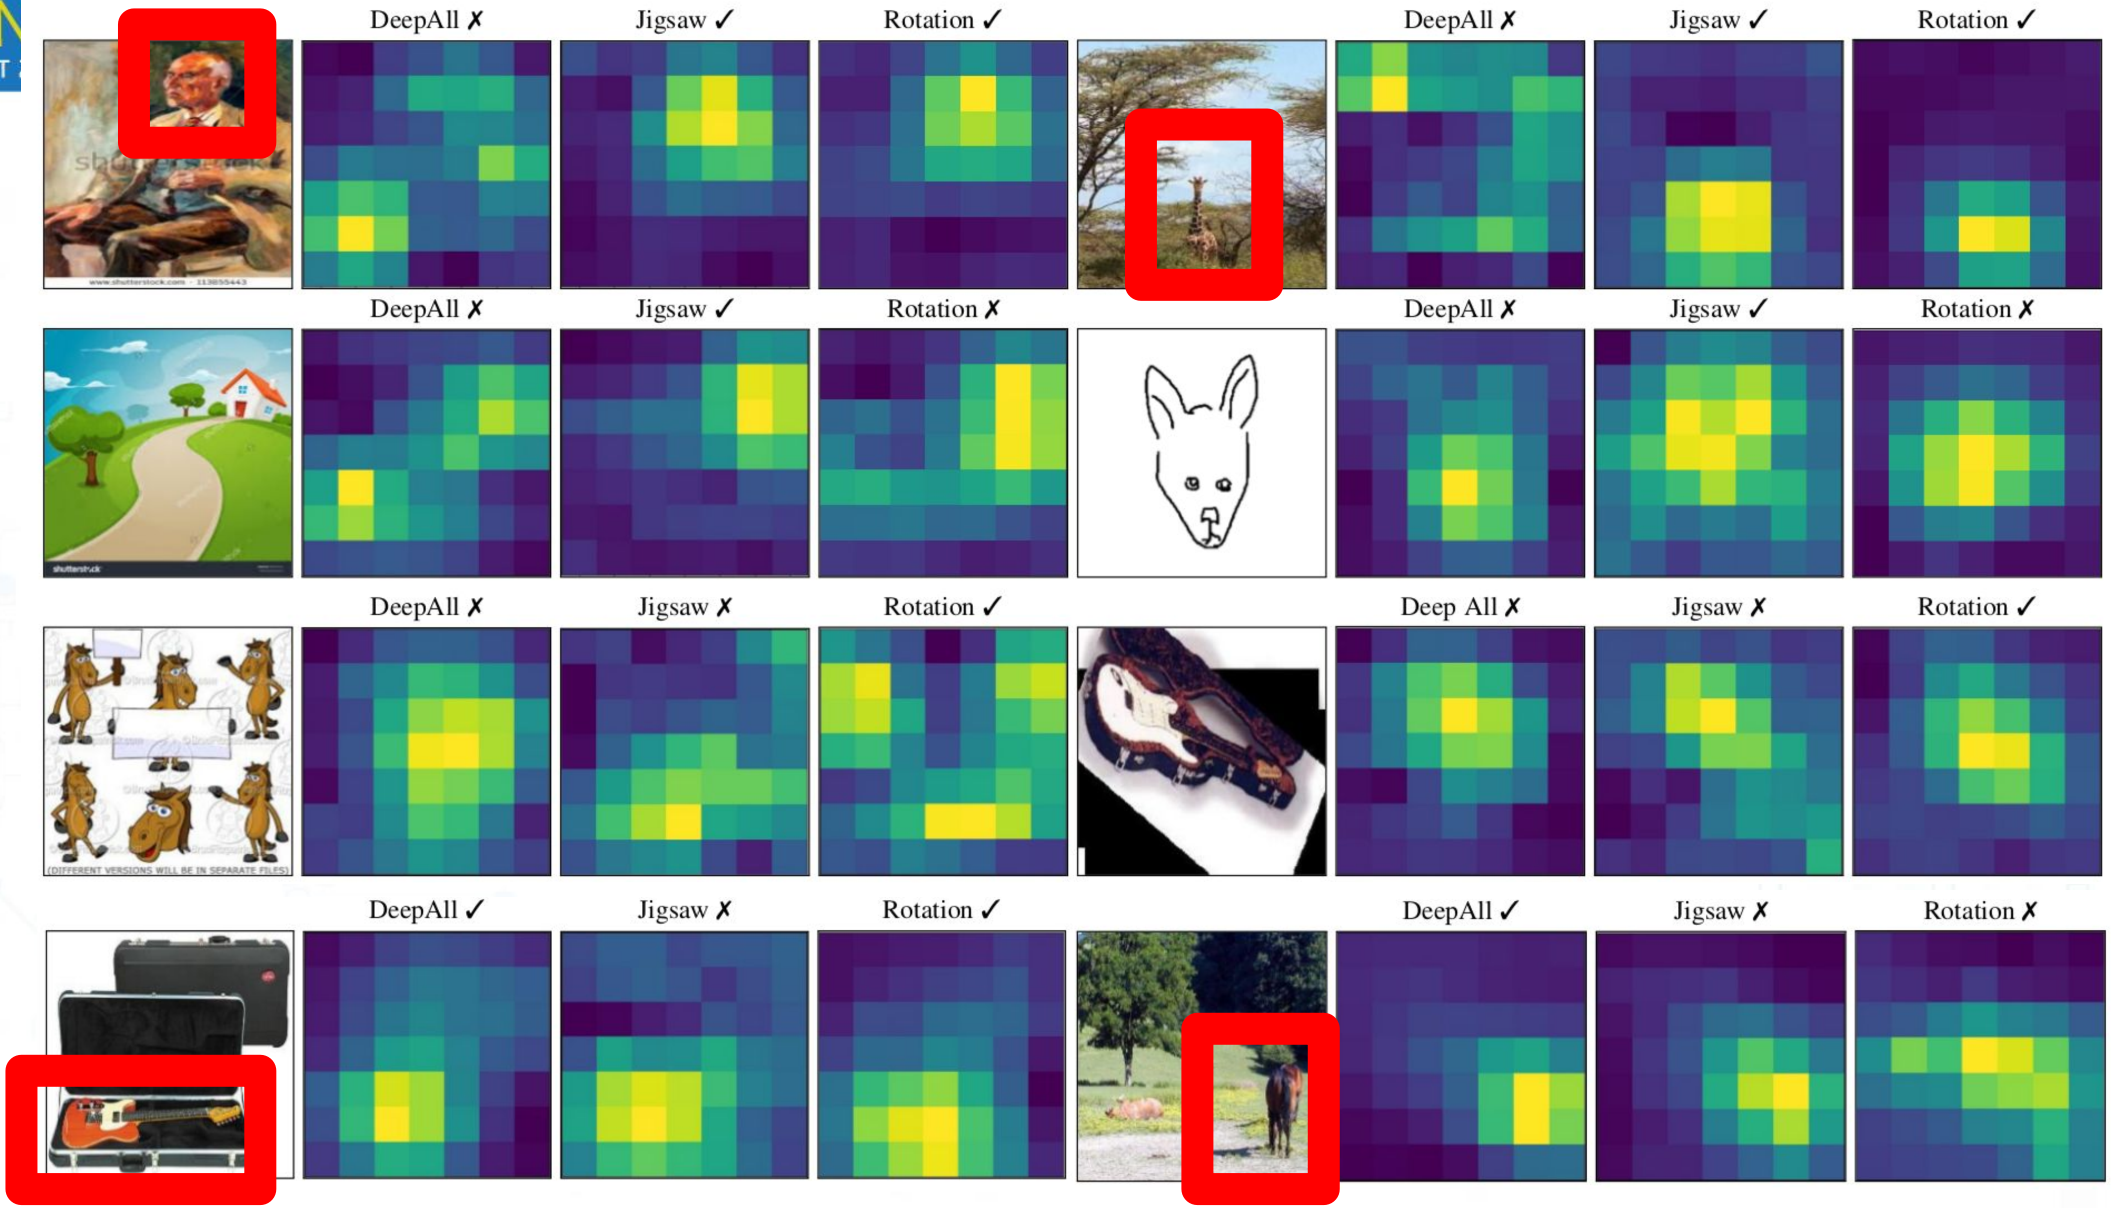

Self-Supervision + Domain Generalization

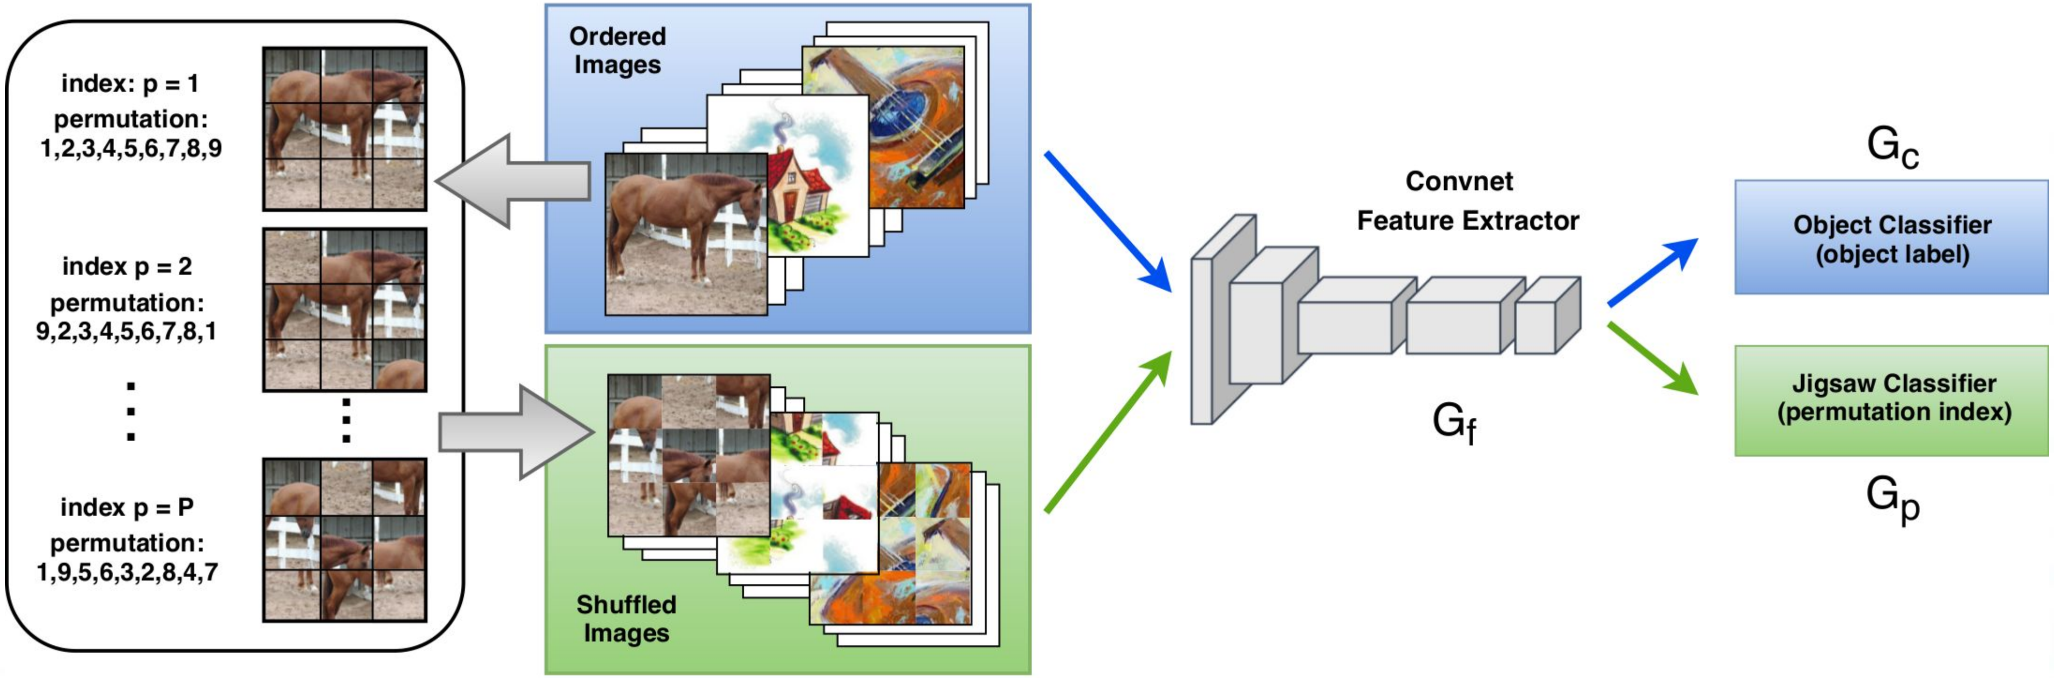

$$\arg \min_{\theta_f, \theta_c, \theta_p} \quad \frac{1}{n^s} \sum_{i=1}^{n^s} \mathcal{L}_c(G_c(G_f(\mathbf{x}_i^s)), y_i^s) + \alpha_s \frac{1}{K^s} \sum_{k=1}^{K^s} \mathcal{L}_p(G_p(G_f(\mathbf{z}_k^s)), p_k^s)$$

[Domain Generalization by Solving Jigsaw Puzzles, CVPR 2019]
[Self-Supervised Learning Across Domains, ArXiv 2020]

Self-Supervision + Domain Adaptation

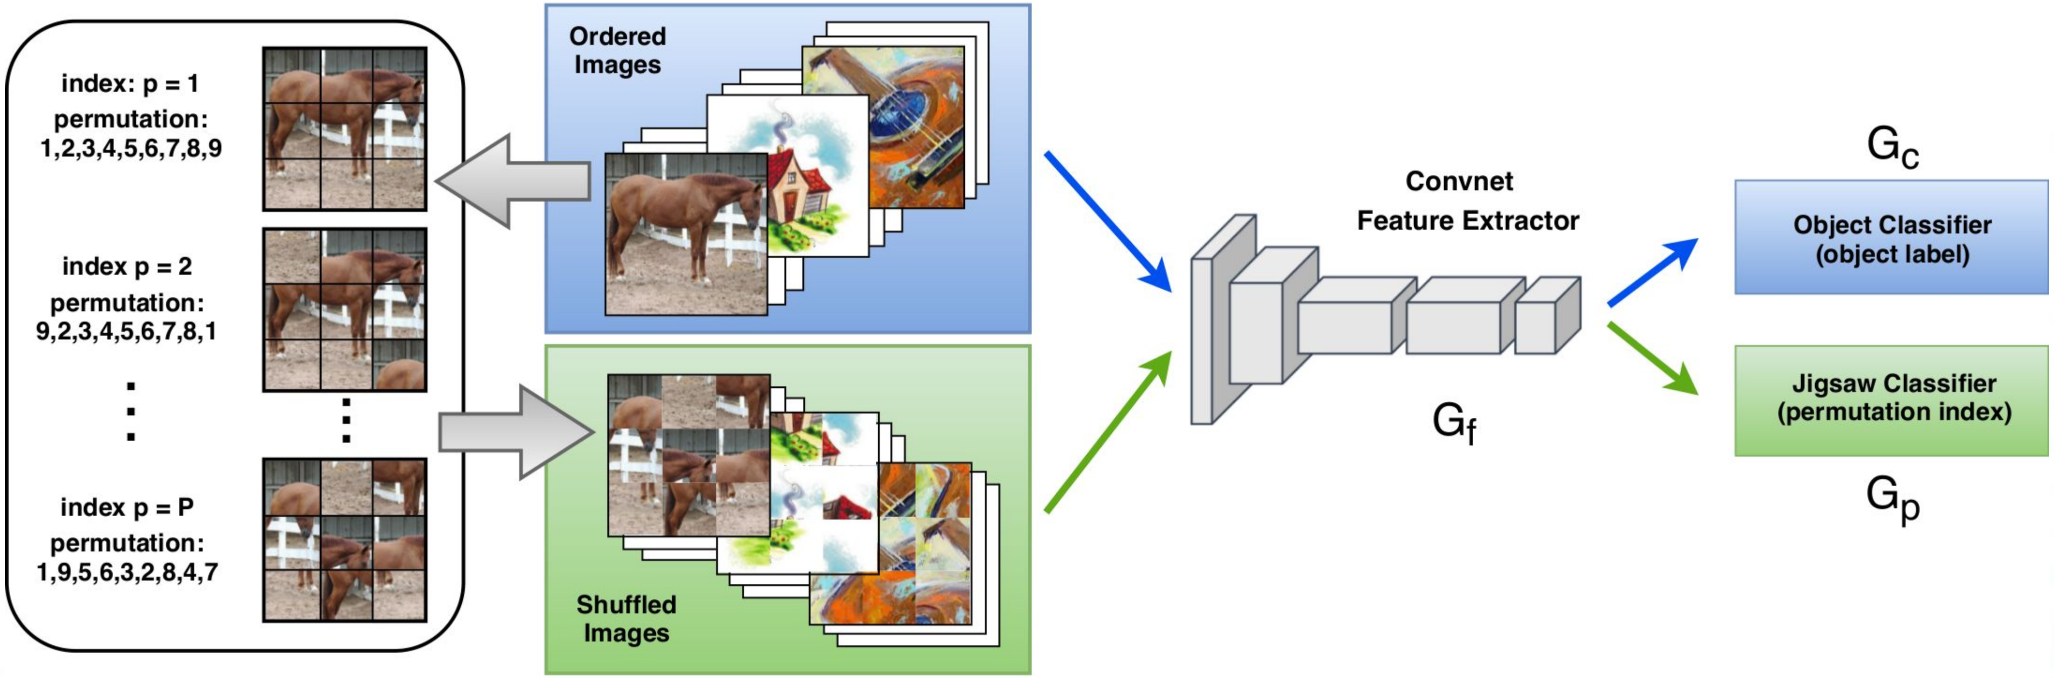

$$\arg\min_{\theta_f,\theta_c,\theta_p} \frac{1}{n^s}\sum_{i=1}^{n^s}\mathcal{L}_c(G_c(G_f(\mathbf{x}_i^s)),y_i^s)+$$
$$\alpha_s\frac{1}{K^s}\sum_{k=1}^{K^s}\mathcal{L}_p(G_p(G_f(\mathbf{z}_k^s)),p_k^s)+$$
$$\alpha_t\frac{1}{K^t}\sum_{k=1}^{K^t}\mathcal{L}_p(G_p(G_f(\mathbf{z}_k^t)),p_k^t).$$

[Domain Generalization by Solving Jigsaw Puzzles, CVPR 2019]
[Self-Supervised Learning Across Domains, ArXiv 2020]

Self-Supervision + Domain Adaptation

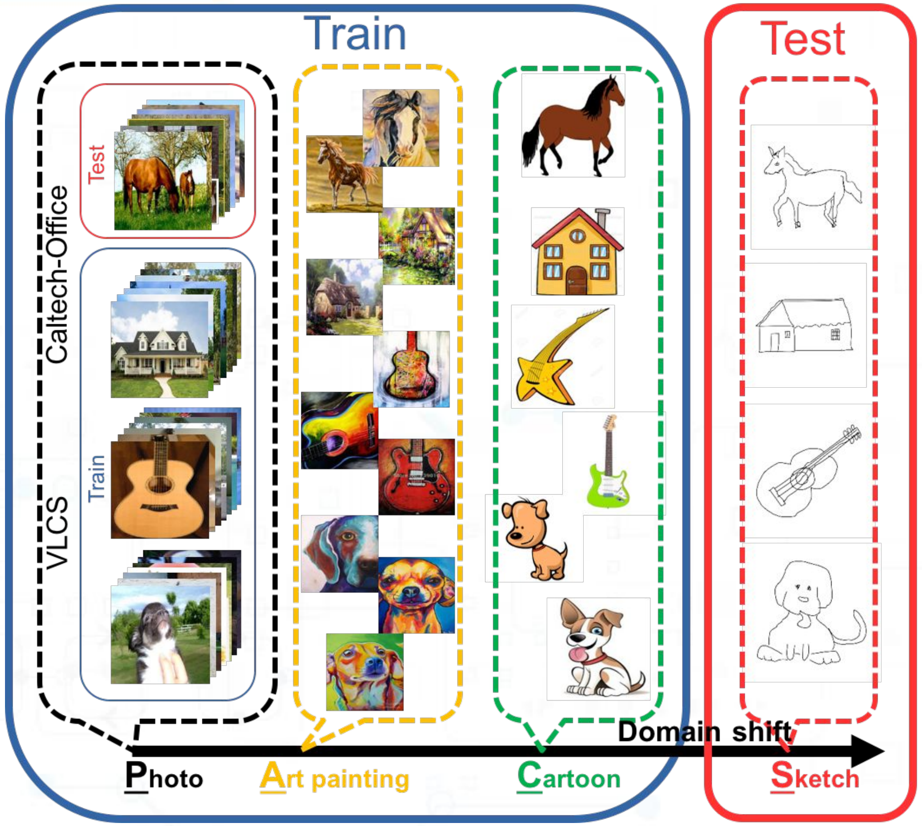

PACS-DA	art_paint.	cartoon	sketches	photo	Avg.
Resnet-18					
DeepAll	74.70	72.40	60.10	92.90	75.03
Dial	87.30	85.50	66.80	97.00	84.15
DDiscovery	87.70	86.90	69.60	97.00	85.30
DeepAll	76.17	73.58	55.65	96.07	75.37±0.42
HAFN	84.95	79.64	64.24	97.70	81.63±0.50
SAFN	86.78	82.72	60.26	98.26	82.01±0.32
SAFN+ENT	89.22	87.39	60.02	98.14	83.69±0.17
DeepAll	77.83	74.26	65.81	95.71	78.40±0.28
Jigsaw$_{\alpha_s=\alpha_t=0.7}$	84.49	82.07	79.86	97.98	86.10±0.26
Rotation$_{\alpha_s=\alpha_t=0.8}$	89.97	82.60	82.00	98.07	88.16±0.51
Jigsaw+Rotation	90.87	82.77	83.80	98.37	88.95±0.38

[Domain Generalization by Solving Jigsaw Puzzles, CVPR 2019]
[Self-Supervised Learning Across Domains, ArXiv 2020]
[Deeper, Broader and Artier Domain Generalization, ICCV 2017]

Self-Supervision + 3D Domain Adaptation

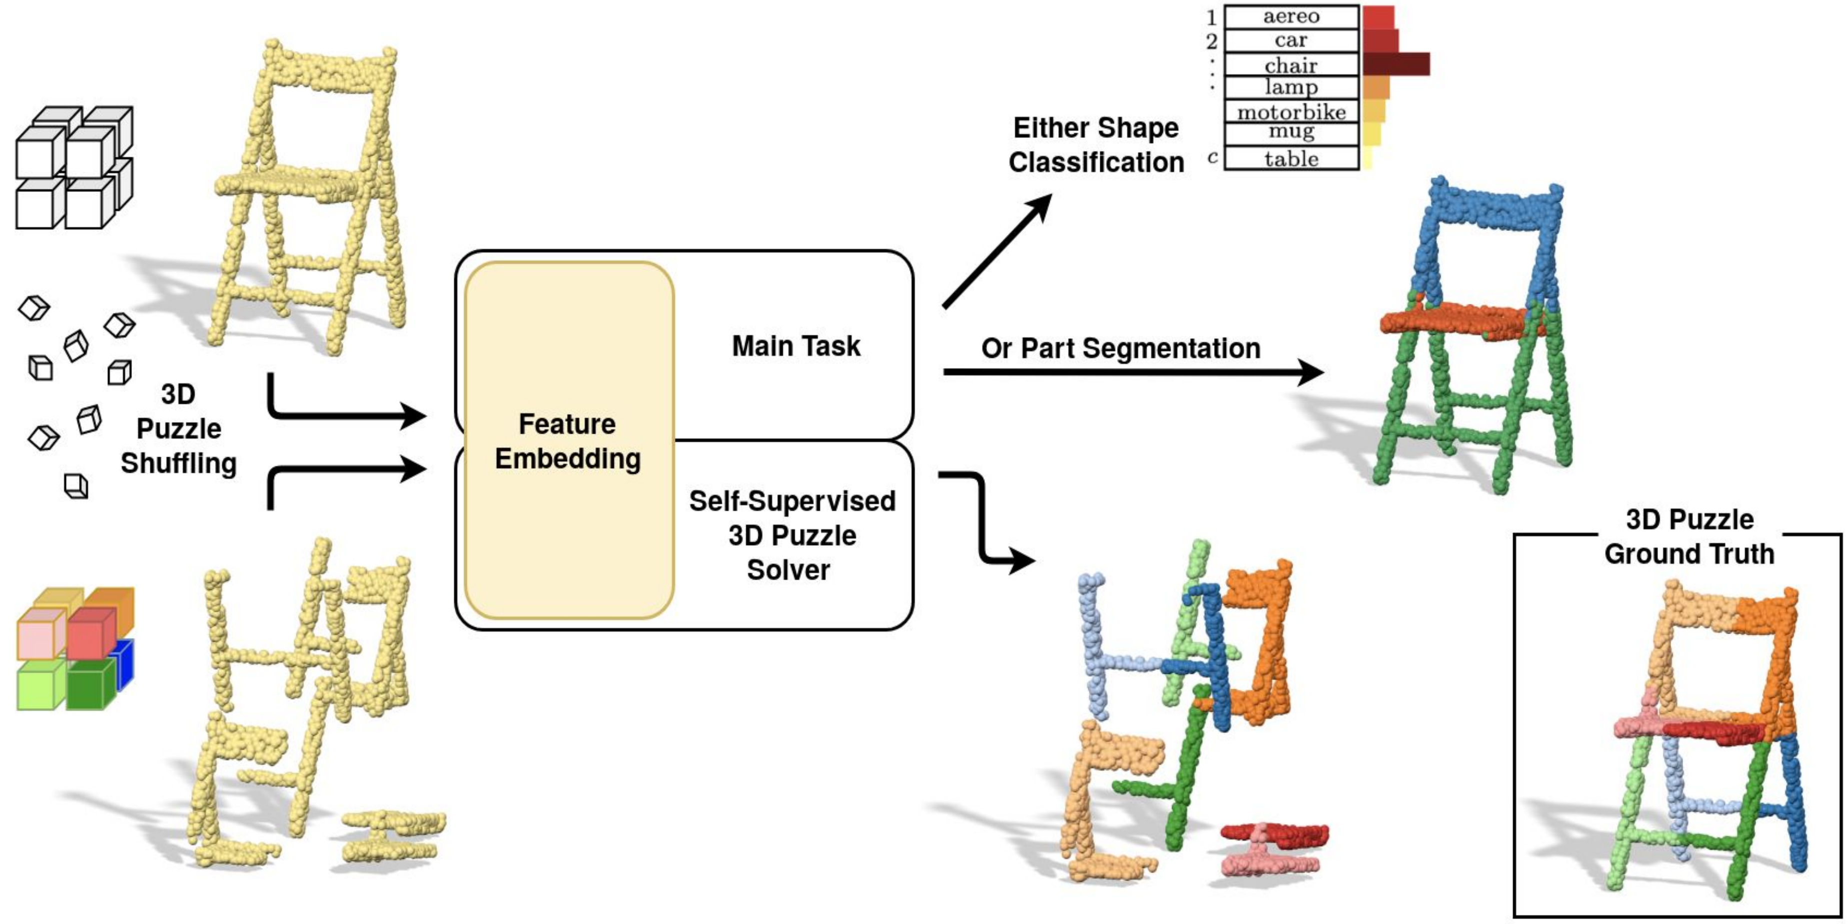

[Self-Supervision for 3D Real-World Challenges, TASK-CV Workshop ECCV 2020]

Self-Supervision + Open-Set DA

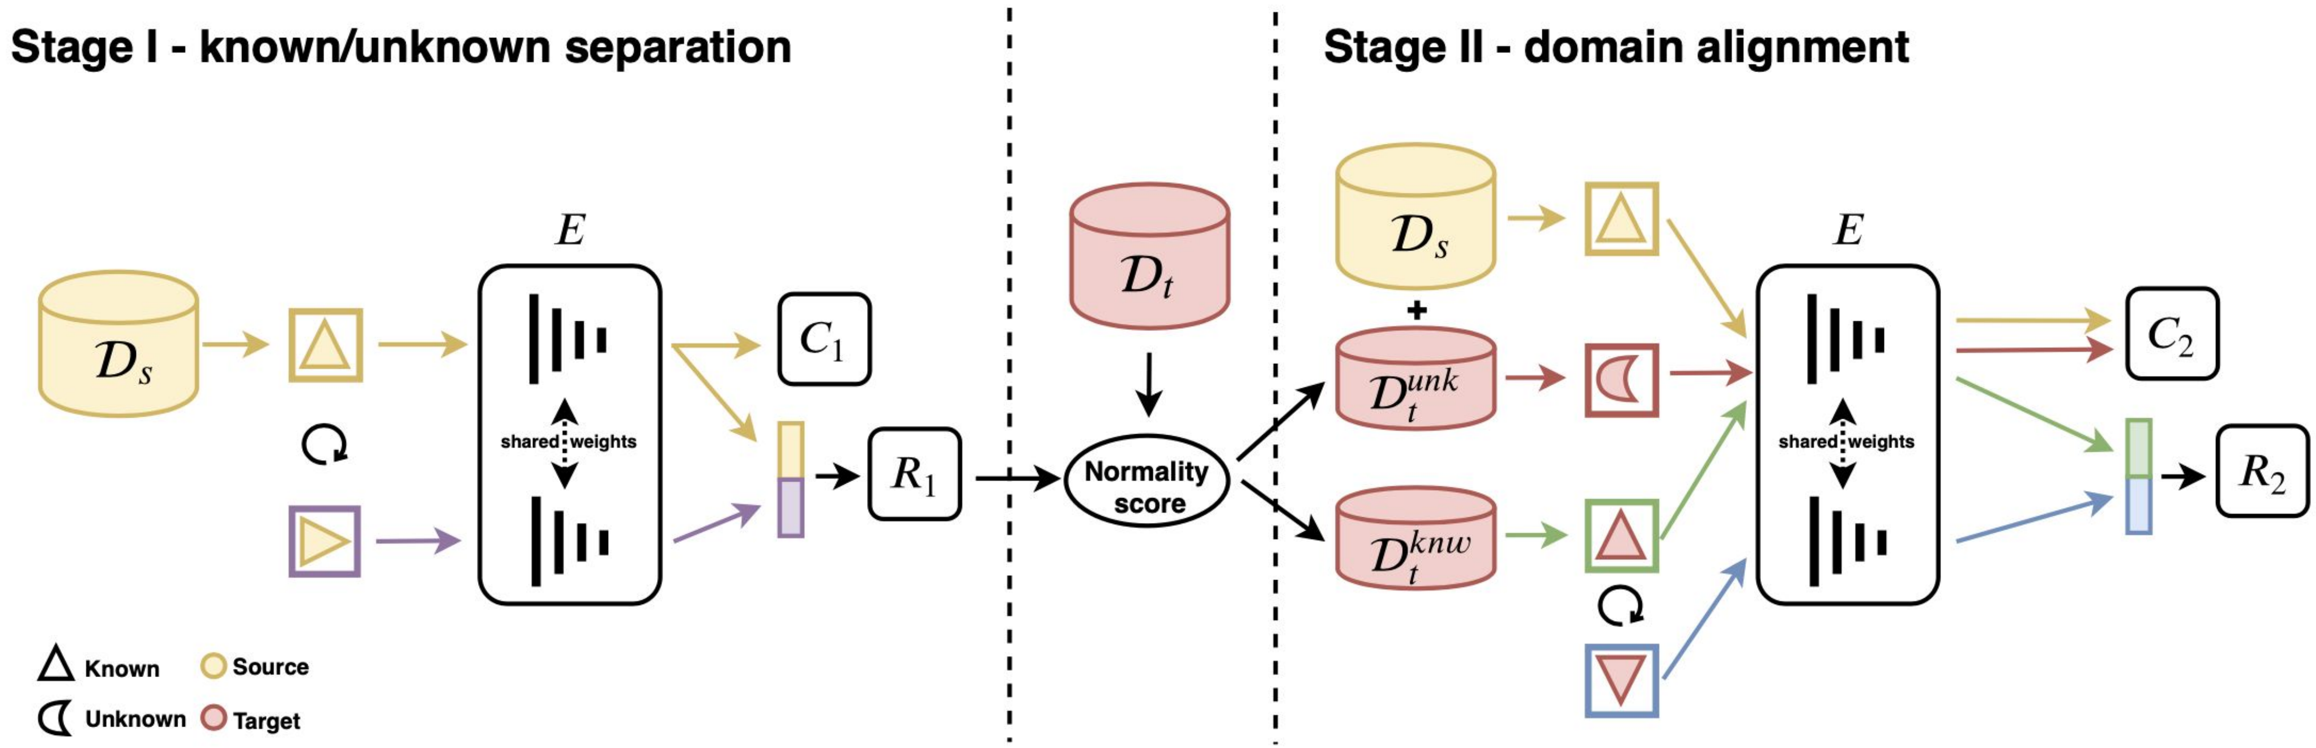

[On the Effectiveness of Image Rotation for Open Set Domain Adaptation, ECCV 2020]

Stage I - Known/Unknown Separation

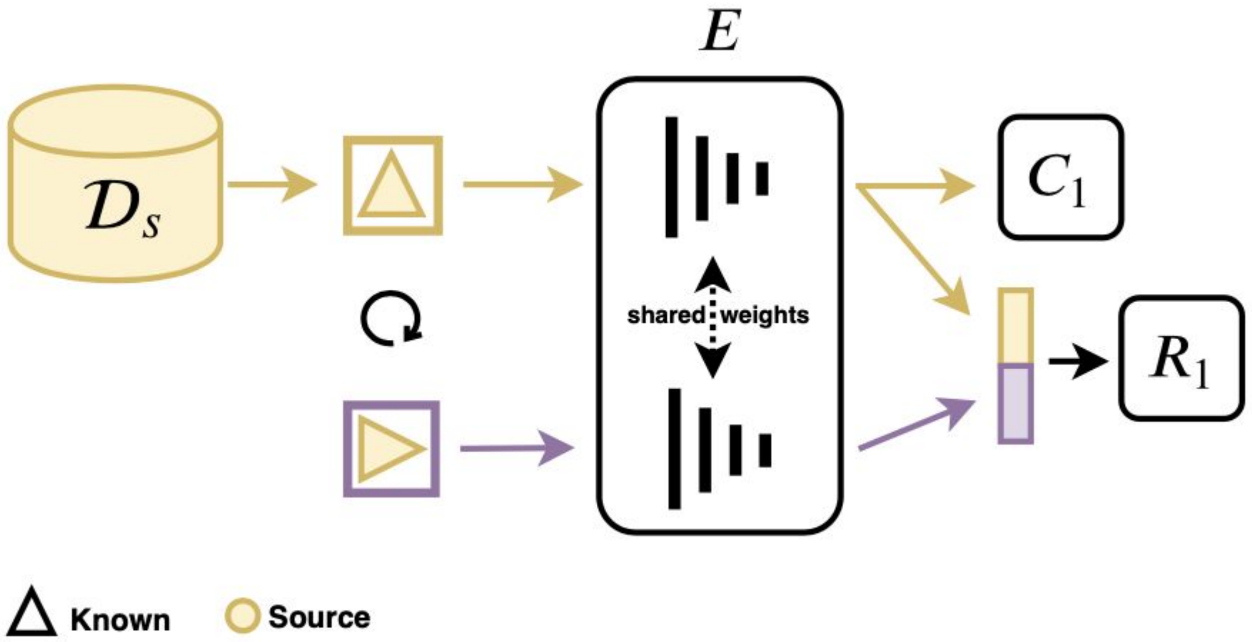

Stage I - Known/Unknown Separation

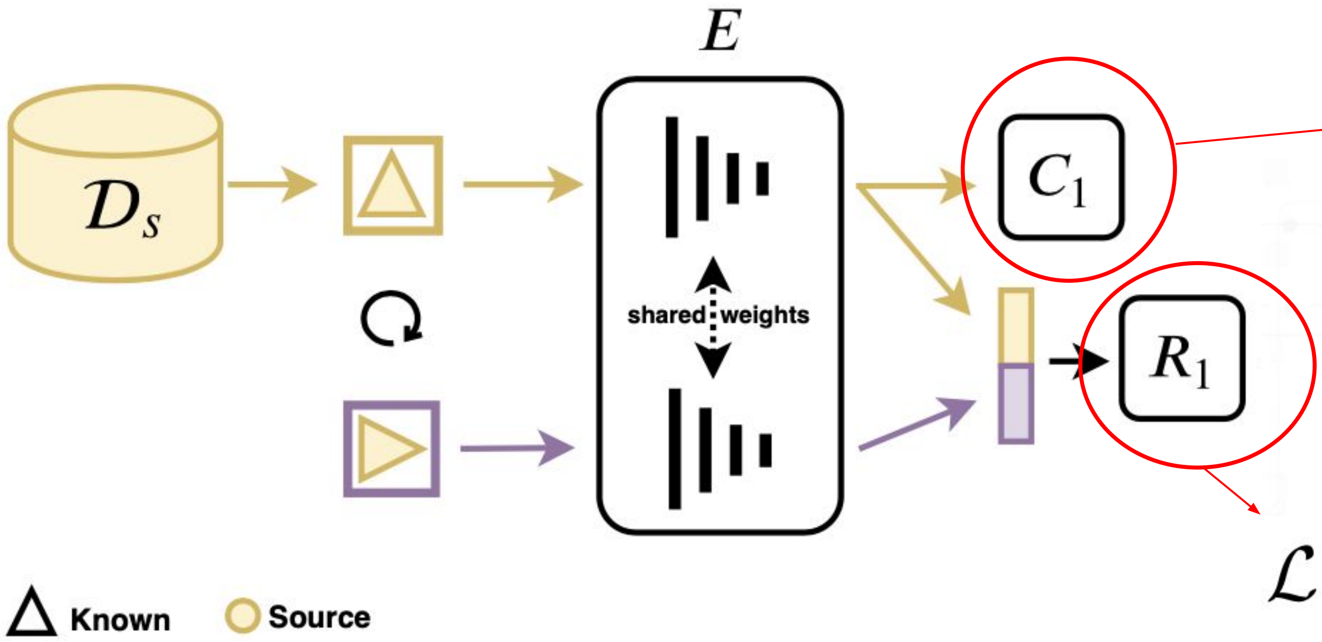

$$\mathcal{L}_{C_1} = -\sum_{j \in \mathcal{D}_s} \boldsymbol{y}_j^s \cdot \log(\hat{\boldsymbol{y}}_j^s)$$

$$\mathcal{L}_{R_1} = \sum_{j \in \tilde{\mathcal{D}}_s} -\lambda_{1,1} \boldsymbol{z}_j^s \cdot \log(\hat{\boldsymbol{z}}_j^s)$$

$$+ \lambda_{1,2} ||\boldsymbol{v}_j^s - \gamma(\boldsymbol{z}_j^s)||_2^2$$

Stage I - Known/Unknown Separation

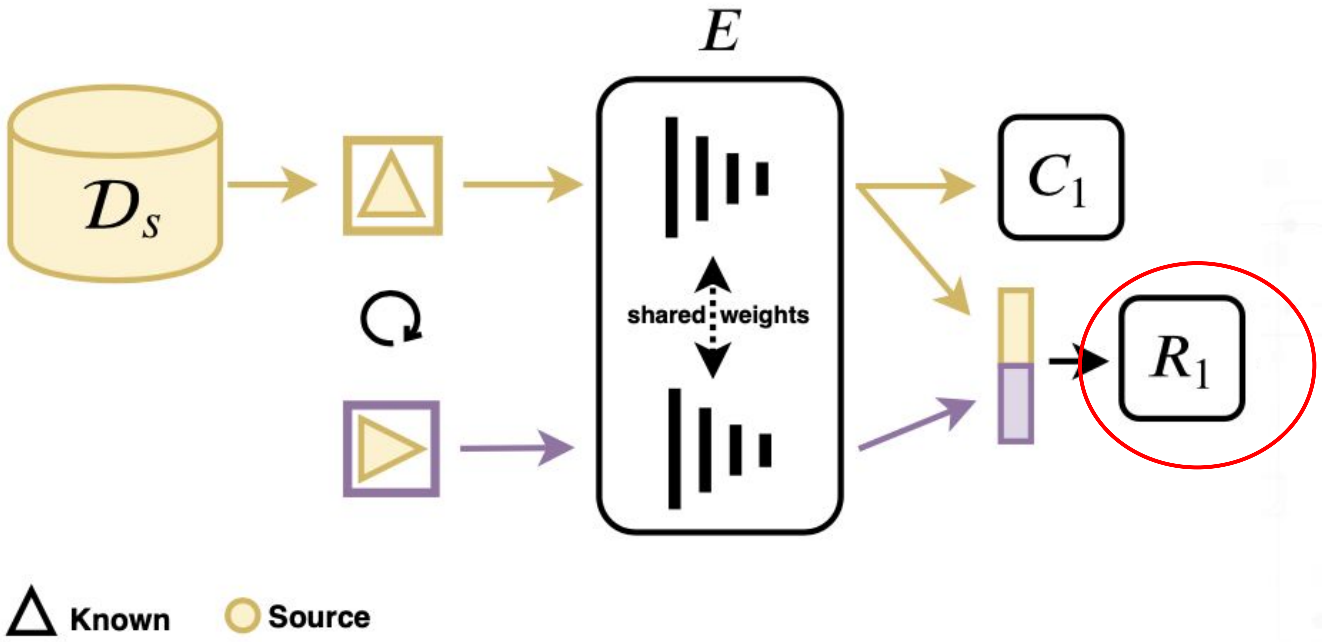

Multi-Rotation Classifier

$(4 \times |\mathcal{C}_s|)$ class labels

Stage I - Known/Unknown Separation

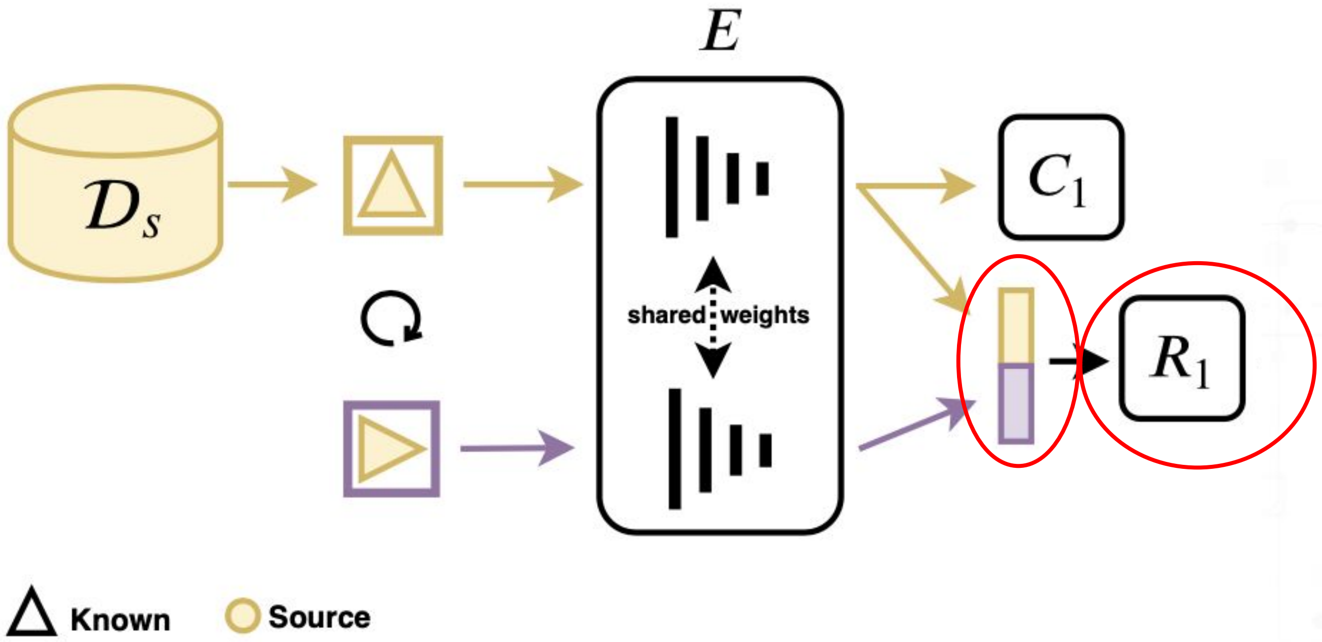

Relative Multi-Rotation Classifier

$(4 \times |\mathcal{C}_s|)$ class labels

Normality Score

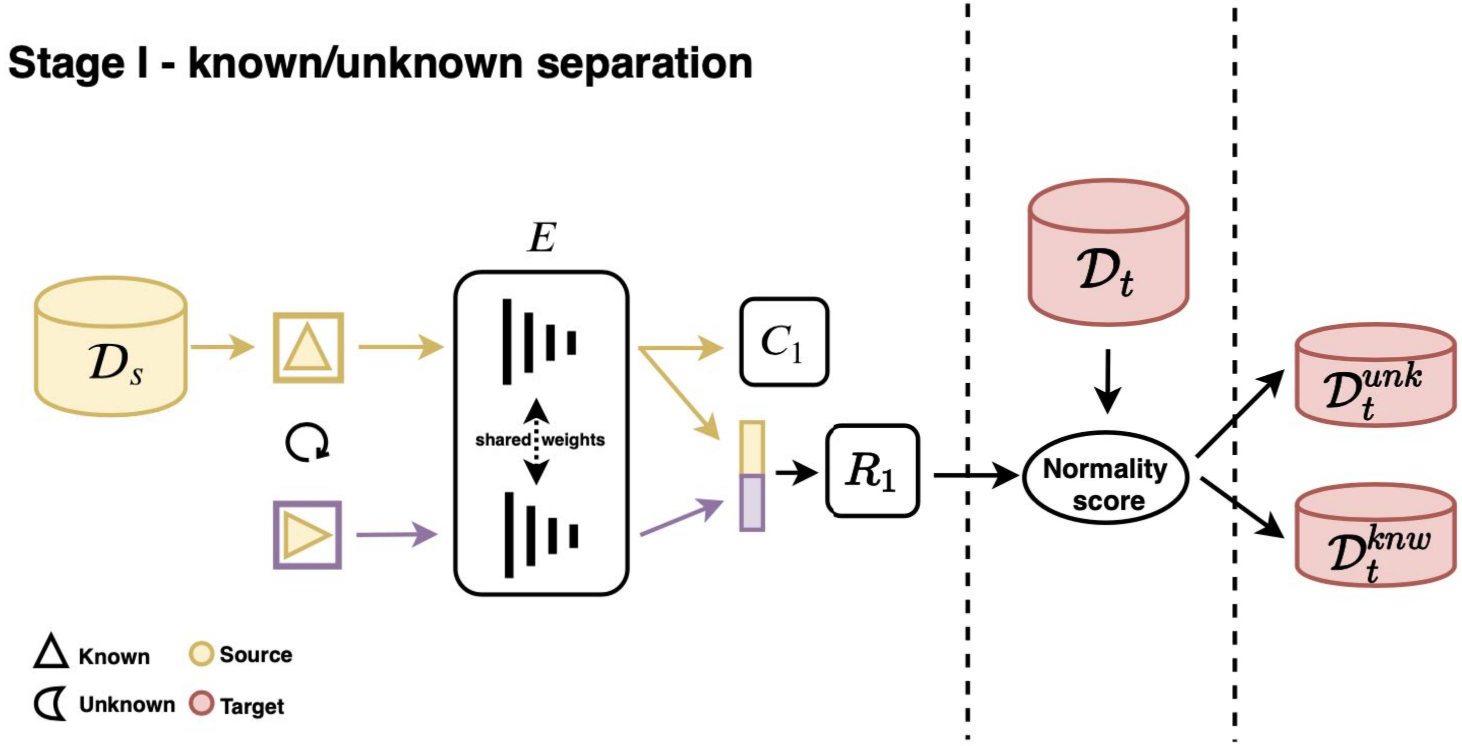

Normality Score

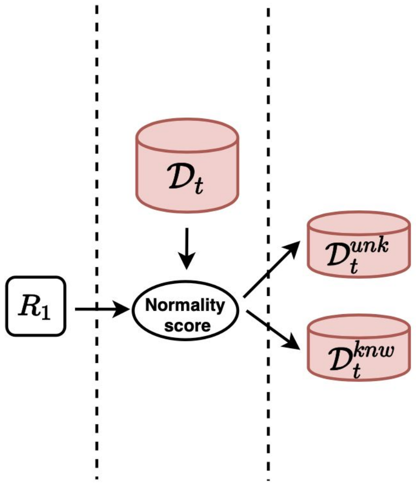

Known Classes = { A, B, C}

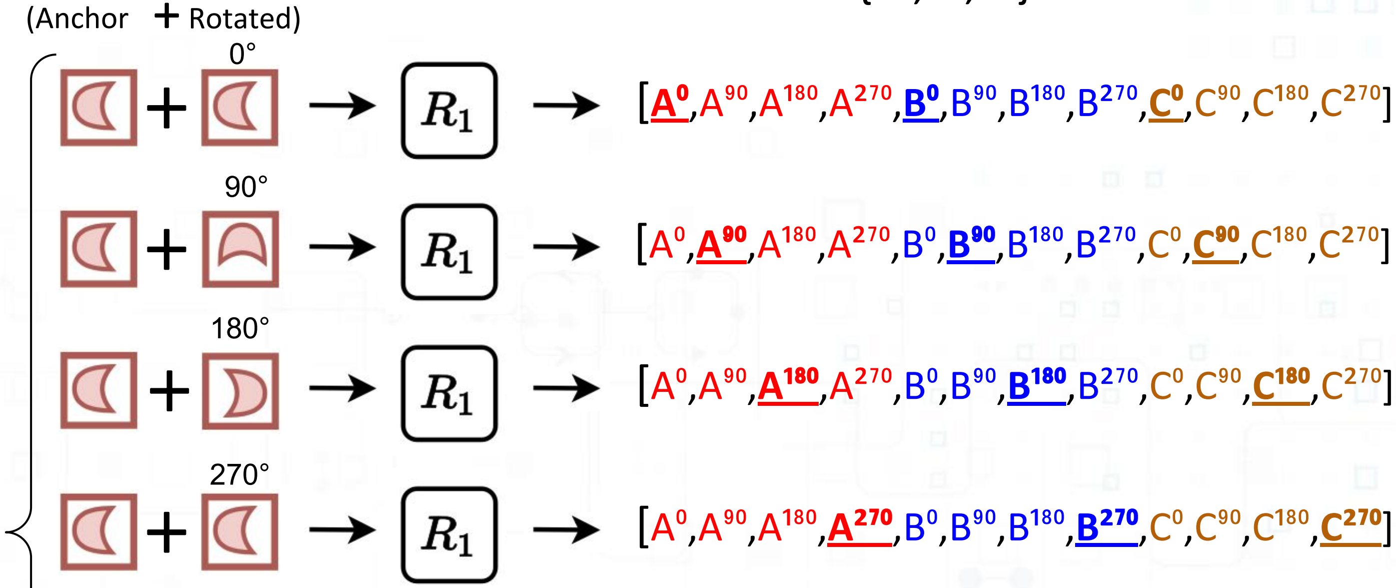

For one sample T

(The evaluation is done for each Target sample)

Score T = $[\underline{\mathbf{A}^0}+\underline{\mathbf{A}^{90}}+\underline{\mathbf{A}^{180}}+\underline{\mathbf{A}^{270}}, \underline{\mathbf{B}^0}+\underline{\mathbf{B}^{90}}+\underline{\mathbf{B}^{180}}+\underline{\mathbf{B}^{270}}, \underline{\mathbf{C}^0}+\underline{\mathbf{C}^{90}}+\underline{\mathbf{C}^{180}}+\underline{\mathbf{C}^{270}}]$

Normality Score

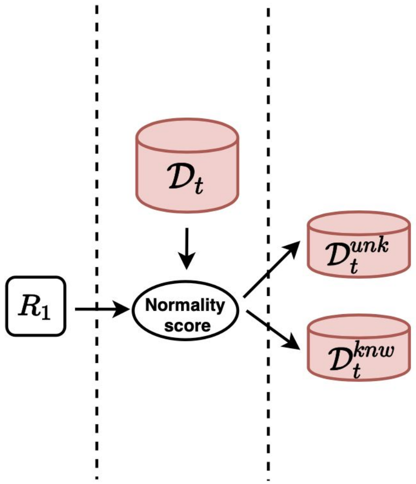

Known Classes = { A, B, C}

For one sample T

(The evaluation is done for each Target sample)

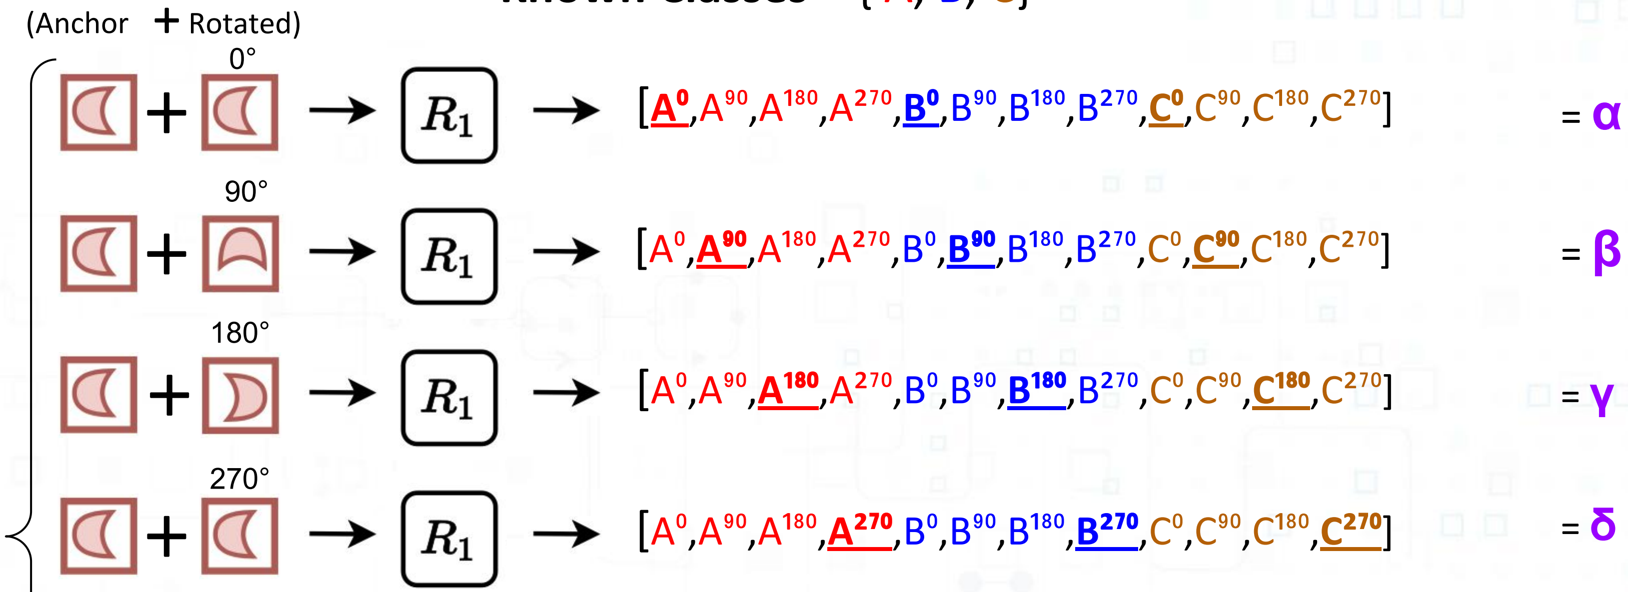

Score T = $[A^{0}+A^{90}+A^{180}+A^{270},B^{0}+B^{90}+B^{180}+B^{270},C^{0}+C^{90}+C^{180}+C^{270}]$

Entropy Score T = mean(**Entropy(α)+Entropy(β)+Entropy(γ)+Entropy(δ)**)

Normality Score

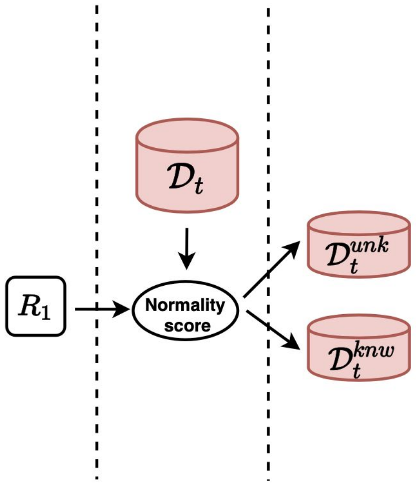

For one sample T

(The evaluation is done for each Target sample)

Known Classes = { A, B, C}

(Anchor + Rotated)

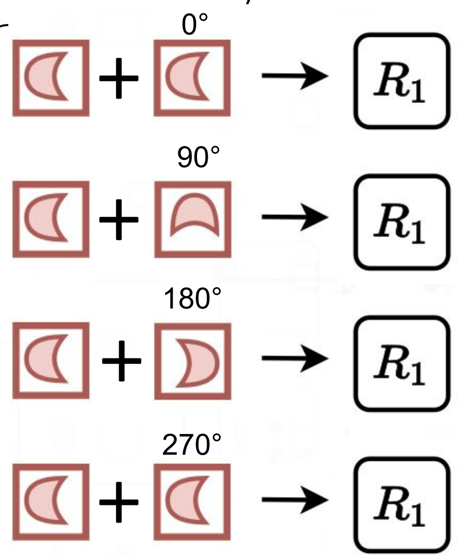

$[\underline{\mathbf{A^0}}, A^{90}, A^{180}, A^{270}, \underline{\mathbf{B^0}}, B^{90}, B^{180}, B^{270}, \underline{\mathbf{C^0}}, C^{90}, C^{180}, C^{270}] = \alpha$

$[A^0, \underline{\mathbf{A^{90}}}, A^{180}, A^{270}, B^0, \underline{\mathbf{B^{90}}}, B^{180}, B^{270}, C^0, \underline{\mathbf{C^{90}}}, C^{180}, C^{270}] = \beta$

$[A^0, A^{90}, \underline{\mathbf{A^{180}}}, A^{270}, B^0, B^{90}, \underline{\mathbf{B^{180}}}, B^{270}, C^0, C^{90}, \underline{\mathbf{C^{180}}}, C^{270}] = \gamma$

$[A^0, A^{90}, A^{180}, \underline{\mathbf{A^{270}}}, B^0, B^{90}, B^{180}, \underline{\mathbf{B^{270}}}, C^0, C^{90}, C^{180}, \underline{\mathbf{C^{270}}}] = \delta$

Score T = $[\underline{\mathbf{A^0}}+\underline{\mathbf{A^{90}}}+\underline{\mathbf{A^{180}}}+\underline{\mathbf{A^{270}}}, \underline{\mathbf{B^0}}+\underline{\mathbf{B^{90}}}+\underline{\mathbf{B^{180}}}+\underline{\mathbf{B^{270}}}, \underline{\mathbf{C^0}}+\underline{\mathbf{C^{90}}}+\underline{\mathbf{C^{180}}}+\underline{\mathbf{C^{270}}}]$

Entropy Score T = mean(**Entropy(α)+Entropy(β)+Entropy(γ)+Entropy(δ)**)

Normality Score T = max$\{$ max$_{\{A,B,C\}}$ (Score T), (1-Entropy Score T)$\}$

Normality Score

The **Normality Score** gives the probability that each Target sample is from a **Known Class**.

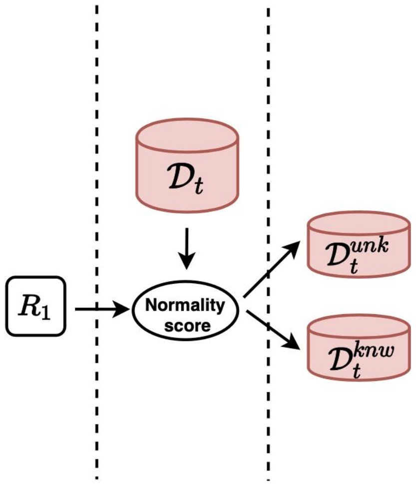

$$\begin{cases} \boldsymbol{x}^t \in \mathcal{D}_t^{knw} & \text{if} \quad \mathcal{N}(\boldsymbol{x}^t) > \bar{\mathcal{N}} \\ \boldsymbol{x}^t \in \mathcal{D}_t^{unk} & \text{if} \quad \mathcal{N}(\boldsymbol{x}^t) < \bar{\mathcal{N}} \end{cases}$$

$$\bar{\mathcal{N}} = \frac{1}{N_t} \sum_{j=1}^{N_t} \mathcal{N}_j$$

The Threshold IS NOT an hyperparameter.

Stage II - Domain Alignment

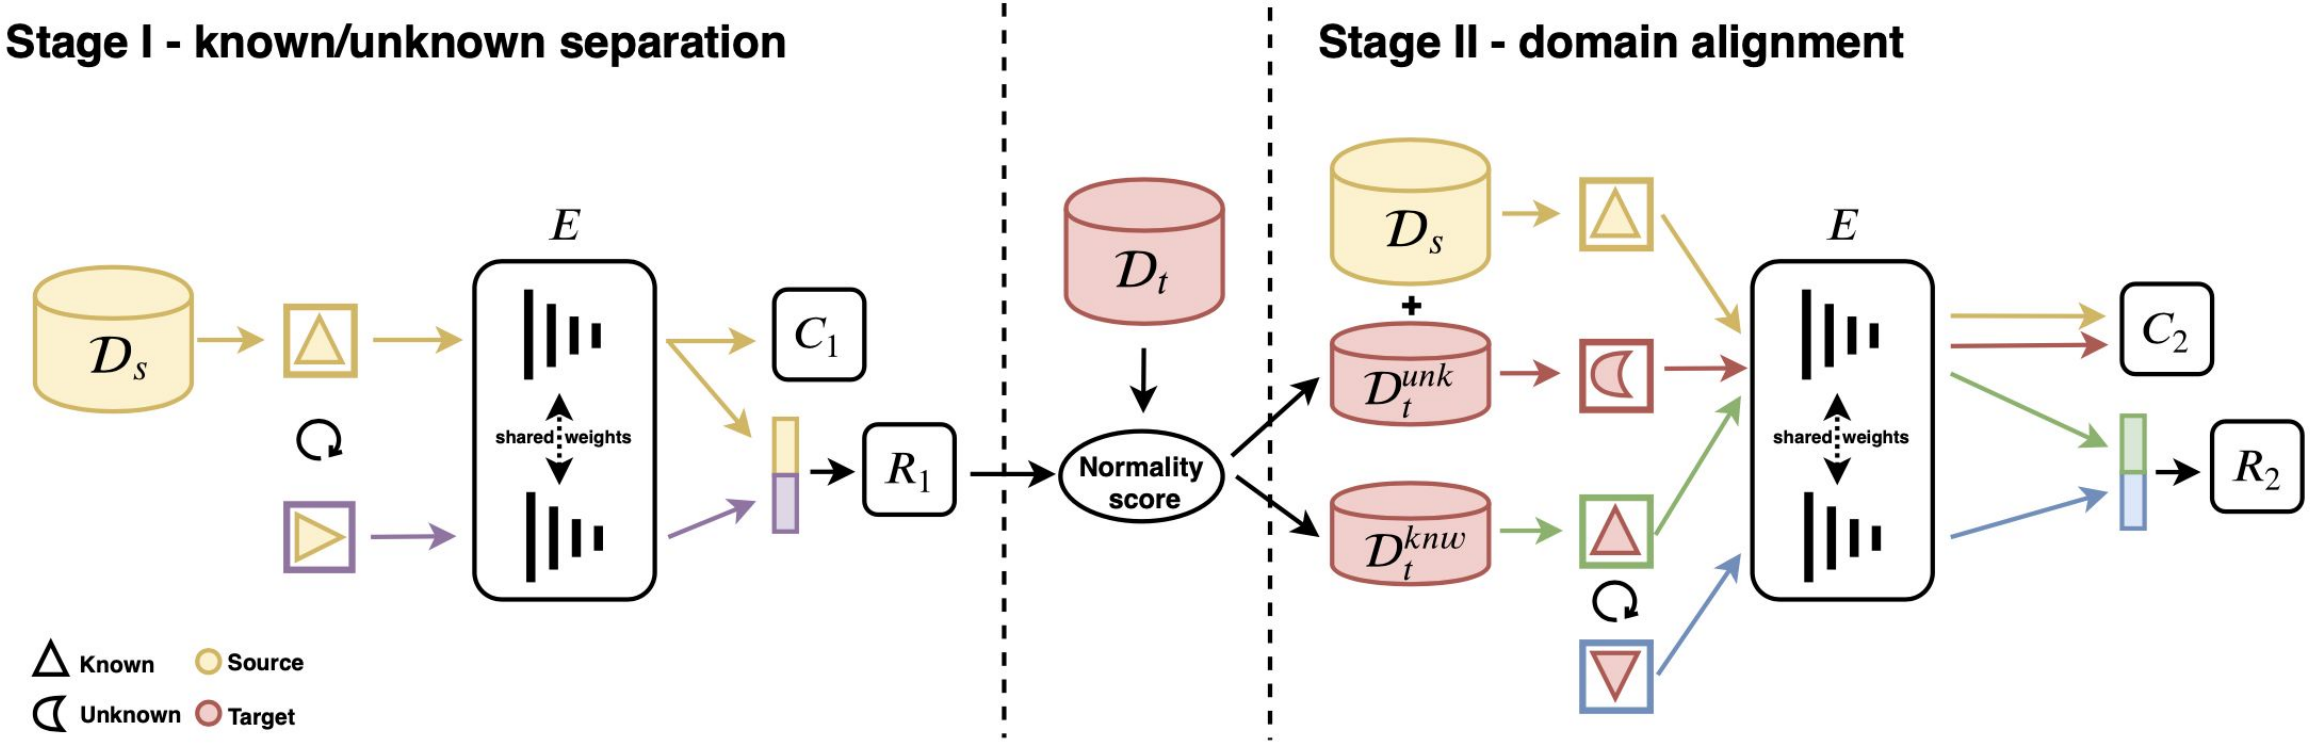

[On the Effectiveness of Image Rotation for Open Set Domain Adaptation, ECCV 2020]

Stage II - Domain Alignment

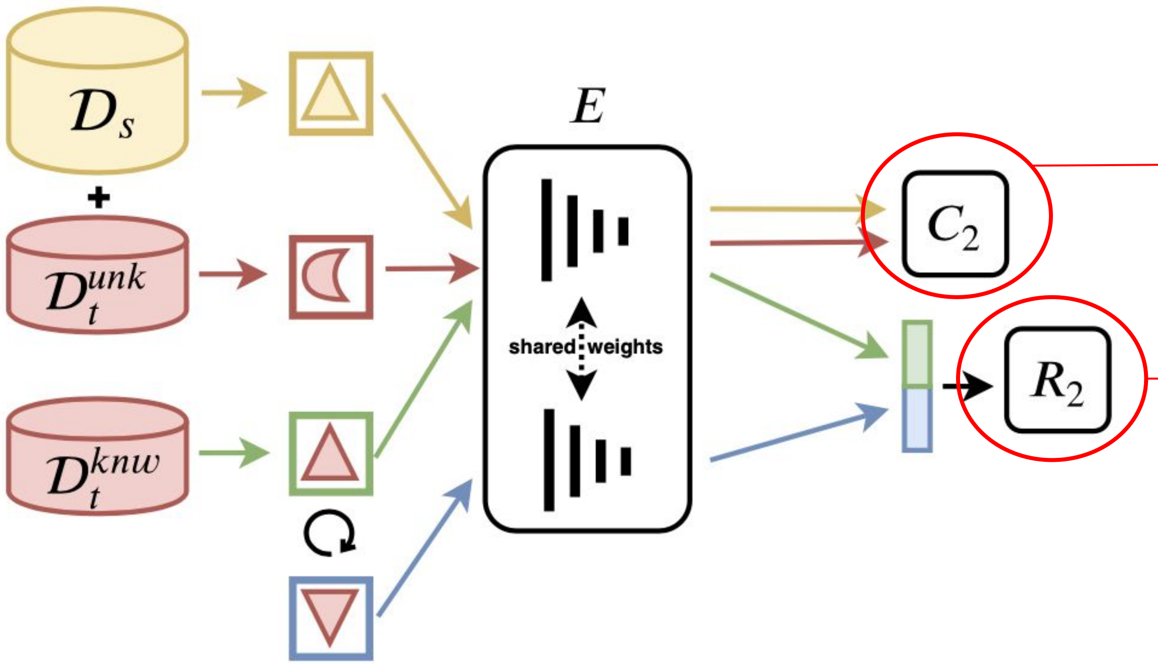

Cs+1 classes

Relative Rotation Classifier for the **Domain Alignment** (4 classes)

$$\mathcal{L}_{R_2} = -\lambda_{2,2} \sum_{j \in \mathcal{D}_t^{knw}} \boldsymbol{q}_j \cdot \log(\hat{\boldsymbol{q}}_j)$$

New Open-Set DA Metrics

Number of
Known Classes

$$OS = \frac{|\mathcal{C}_s|}{|\mathcal{C}_s|+1} \times OS^* + \frac{1}{|\mathcal{C}_s|+1} \times UNK$$

Measure of the
overall performance

Class accuracy over
the **Known Classes**

Class accuracy over
the **Unknown Class**

$$HOS = 2\frac{OS^* \times UNK}{OS^* + UNK}$$

<u>Harmonic mean</u> of **OS*** and **UNK**

It provides a high score only if the algorithm performs well both on known and on unknown samples, independently of |Cs|.

Results on Office-Home Dataset

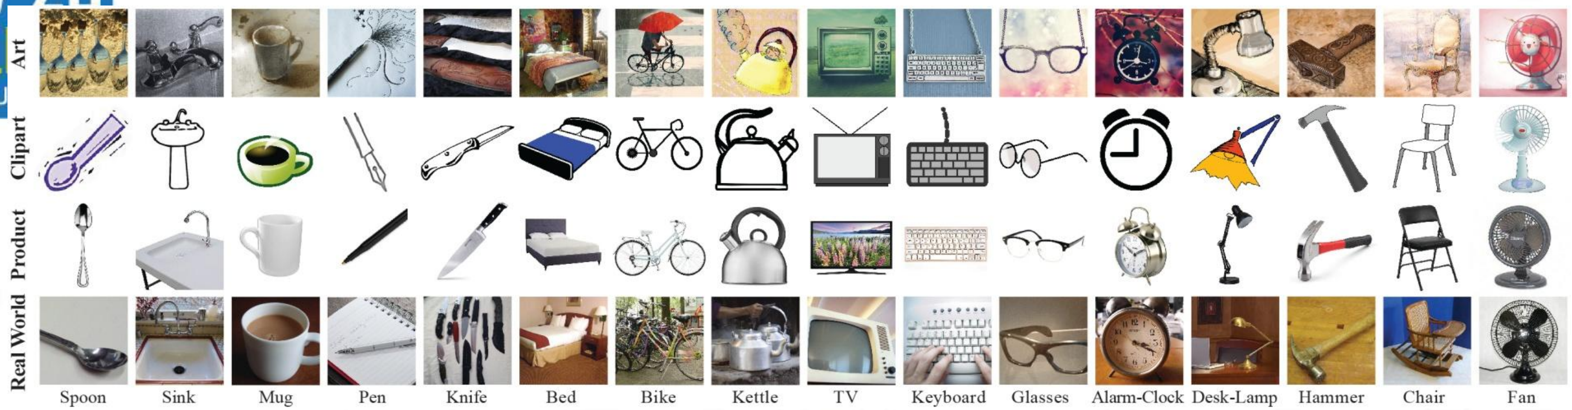

25 Known Classes
40 Unknown Classes

Office-Home

	Pr → Rw			Pr → Cl			Pr → Ar			Ar → Pr			Ar → Rw			Ar → Cl		
	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**
STA$_{sum}$ [24]	78.1	63.3	69.7	44.7	71.5	55.0	55.4	73.7	63.1	68.7	59.7	63.7	81.1	50.5	62.1	50.8	63.4	56.3
STA$_{max}$	76.2	64.3	69.5	44.2	67.1	53.2	54.2	72.4	61.9	68.0	48.4	54.0	78.6	60.4	68.3	46.0	72.3	55.8
OSBP [33]	76.2	71.7	73.9	44.5	66.3	53.2	59.1	68.1	**63.2**	71.8	59.8	65.2	79.3	67.5	72.9	50.2	61.1	55.1
UAN [45]	84.0	0.1	0.2	59.1	0.0	0.0	73.7	0.0	0.0	81.1	0.0	0.0	88.2	0.1	0.2	62.4	0.0	0.0
ROS	70.8	78.4	**74.4**	46.5	71.2	**56.3**	57.3	64.3	60.6	68.4	70.3	**69.3**	75.8	77.2	**76.5**	50.6	74.1	**60.1**

	Rw → Ar			Rw → Pr			Rw → Cl			Cl → Rw			Cl → Ar			Cl → Pr			Avg.		
	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**	OS*	UNK	**HOS**
STA$_{sum}$	67.9	62.3	65.0	77.9	58.0	66.4	51.4	57.9	54.2	69.8	63.2	66.3	53.0	63.9	57.9	61.4	63.5	62.5	63.4	62.6	61.9±2.1
STA$_{max}$	67.5	66.7	67.1	77.1	55.4	64.5	49.9	61.1	54.5	67.0	66.7	66.8	51.4	65.0	57.4	61.8	59.1	60.4	61.8	63.3	61.1±0.3
OSBP	66.1	67.3	66.7	76.3	68.6	72.3	48.0	63.0	54.5	72	69.2	**70.6**	59.4	70.3	**64.3**	67.0	62.7	64.7	64.1	66.3	64.7±0.2
UAN	77.5	0.1	0.2	85.0	0.1	0.1	66.2	0.0	0.0	80.6	0.1	0.2	70.5	0.0	0.0	74.0	0.1	0.2	75.2	0.0	0.1±0.0
ROS	67	70.8	**68.8**	72	80	**75.7**	51.5	73	**60.4**	65.3	72.2	68.6	53.6	65.5	58.9	59.8	71.6	**65.2**	61.6	72.4	**66.2± 0.3**

[Deep Hashing Network for Unsupervised Domain Adaptation, CVPR 2017]

...with less known classes

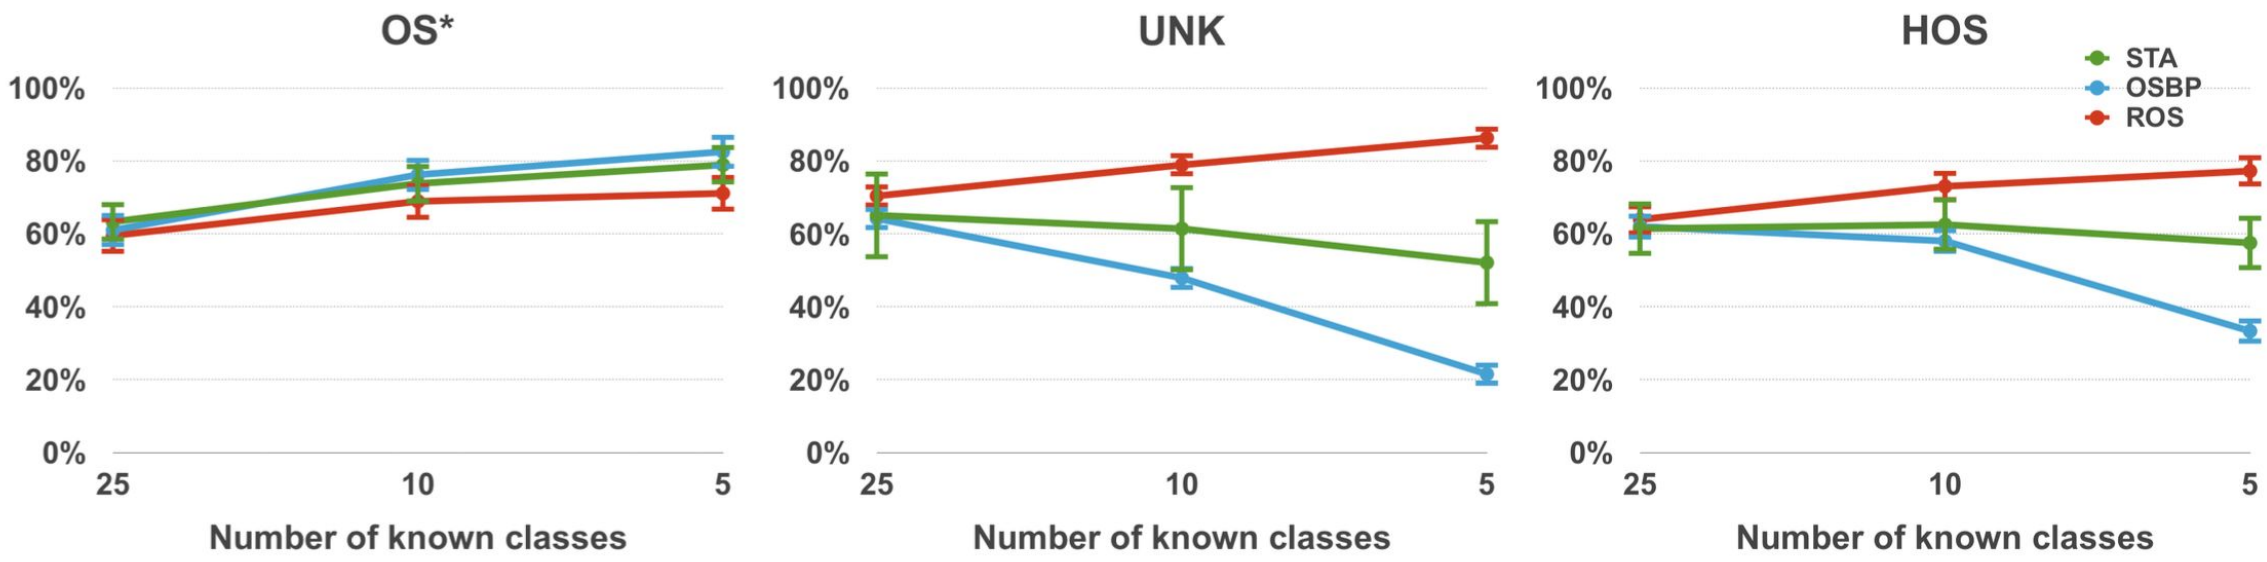

ECCV'20 ONLINE 23-28 AUGUST 2020

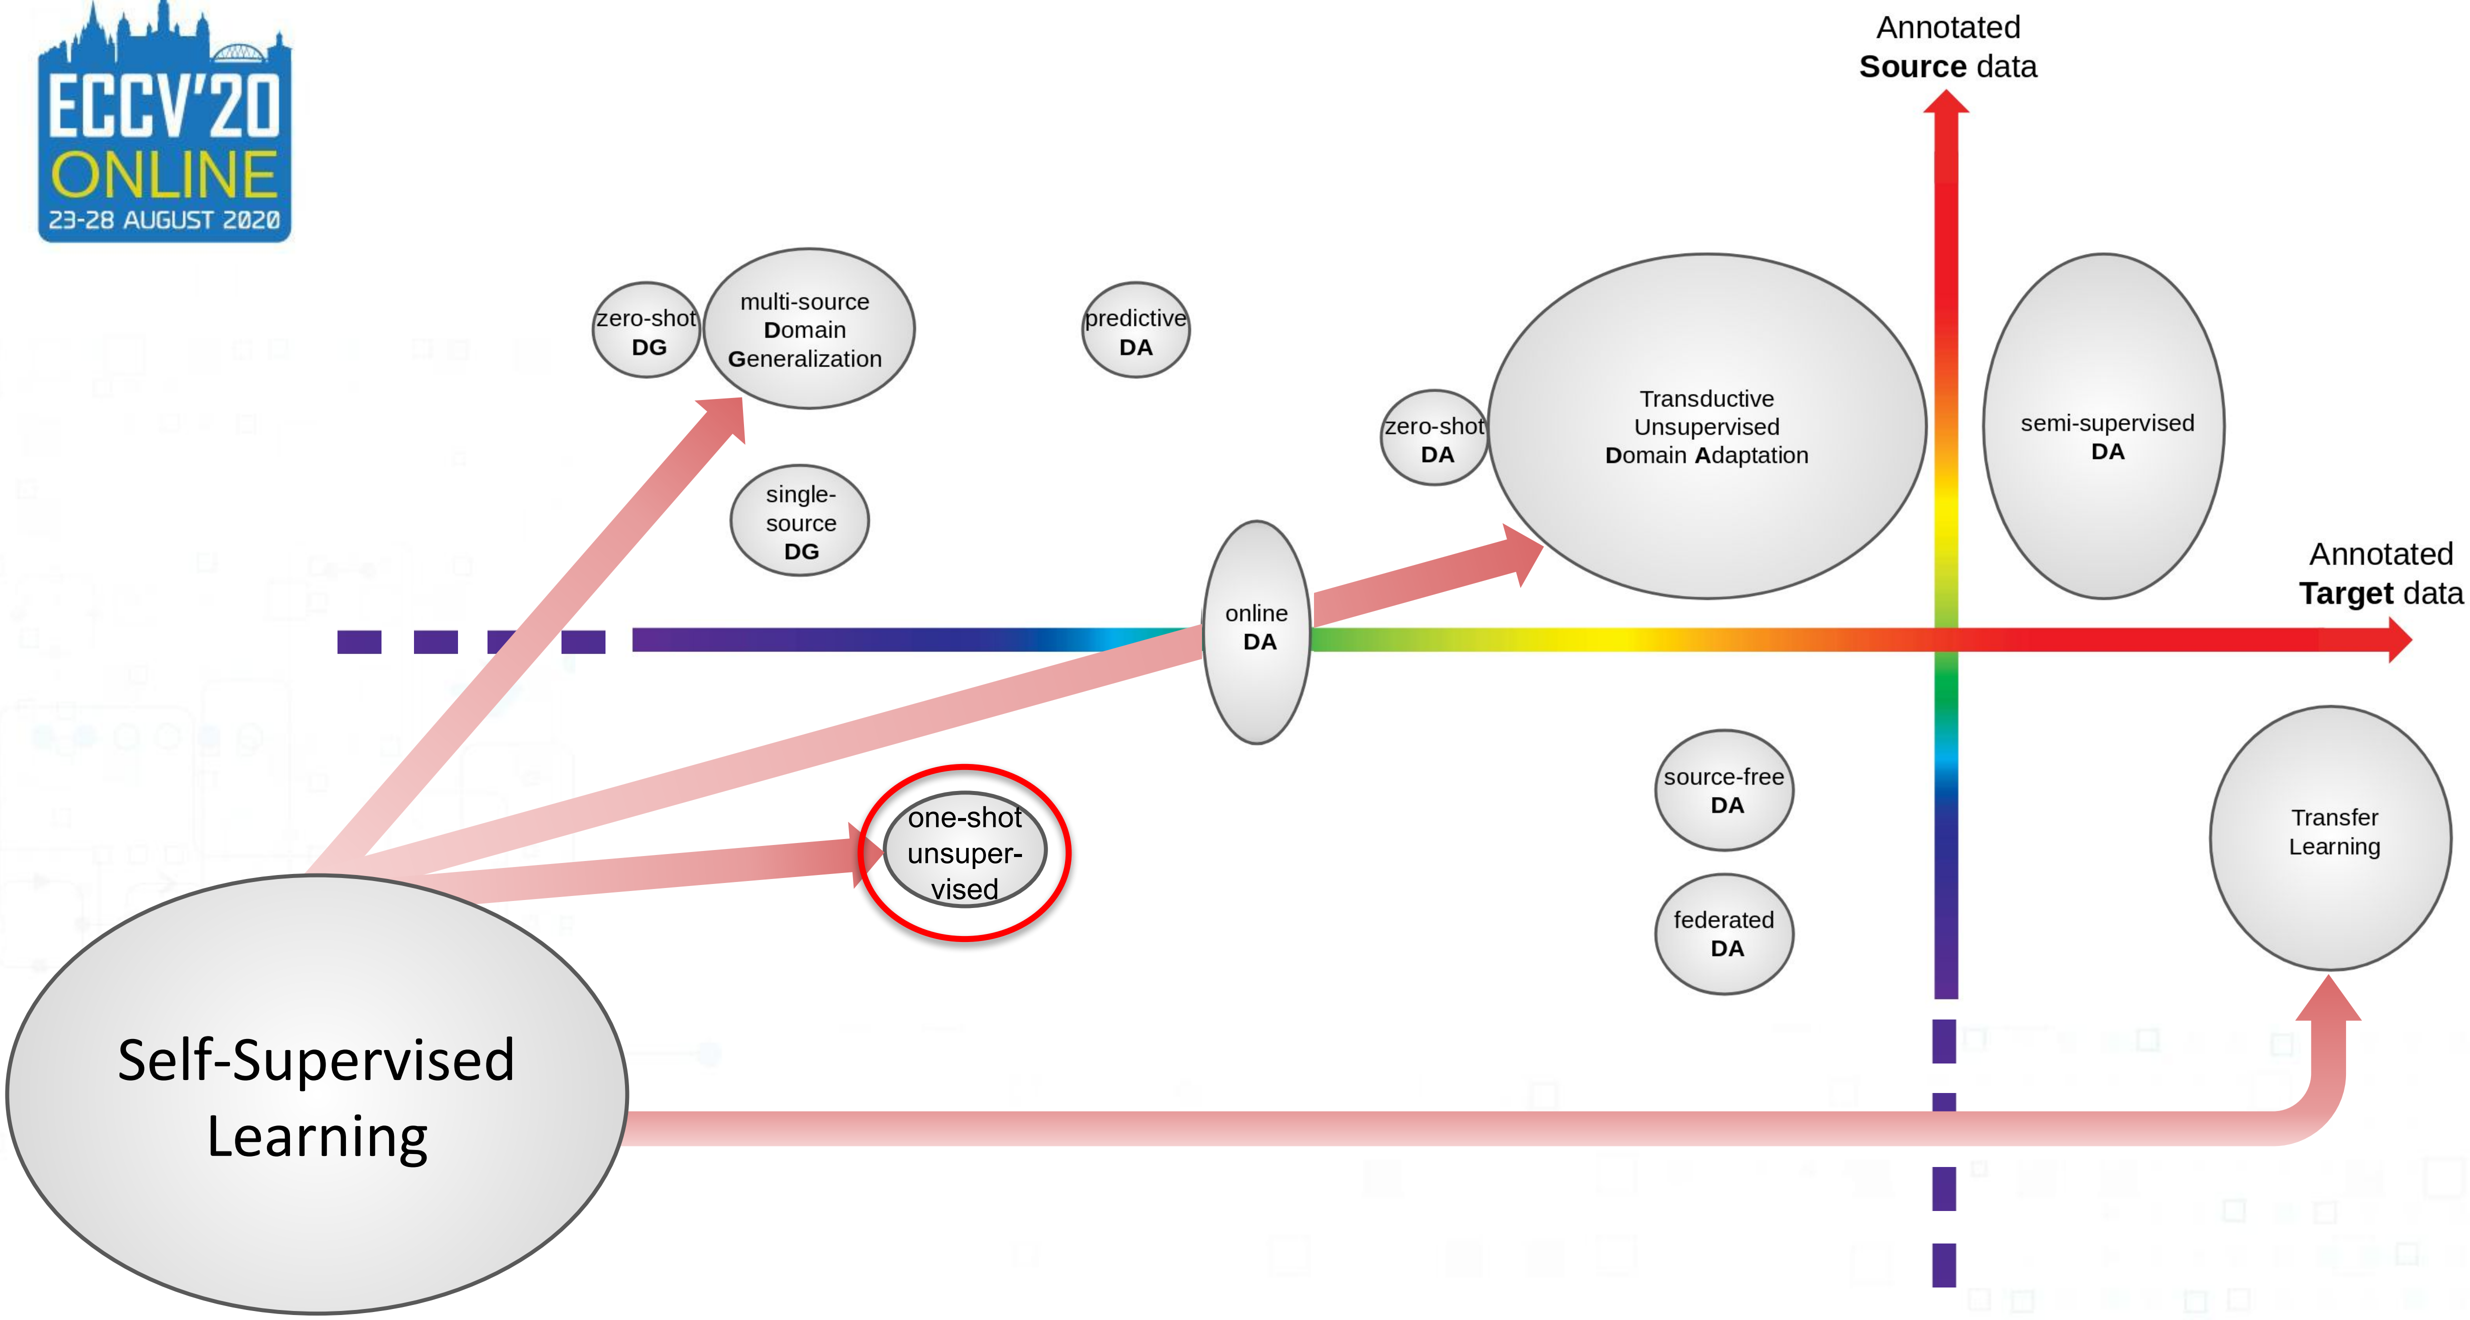

Self-Supervision + One-Shot Unsupervised Cross-Domain Detection

TRAINING

TEST

Source - labeled

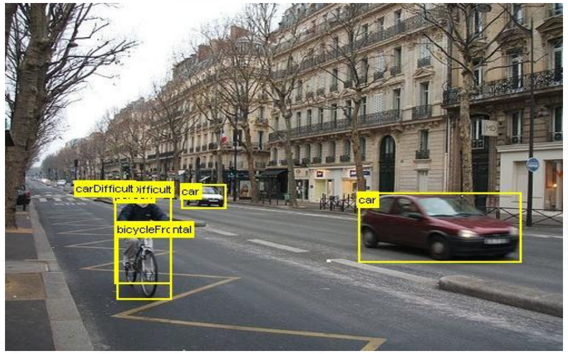

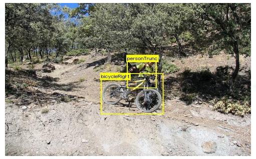

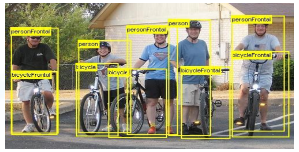

Target - unlabeled

?

Cross-Domain Object Detection

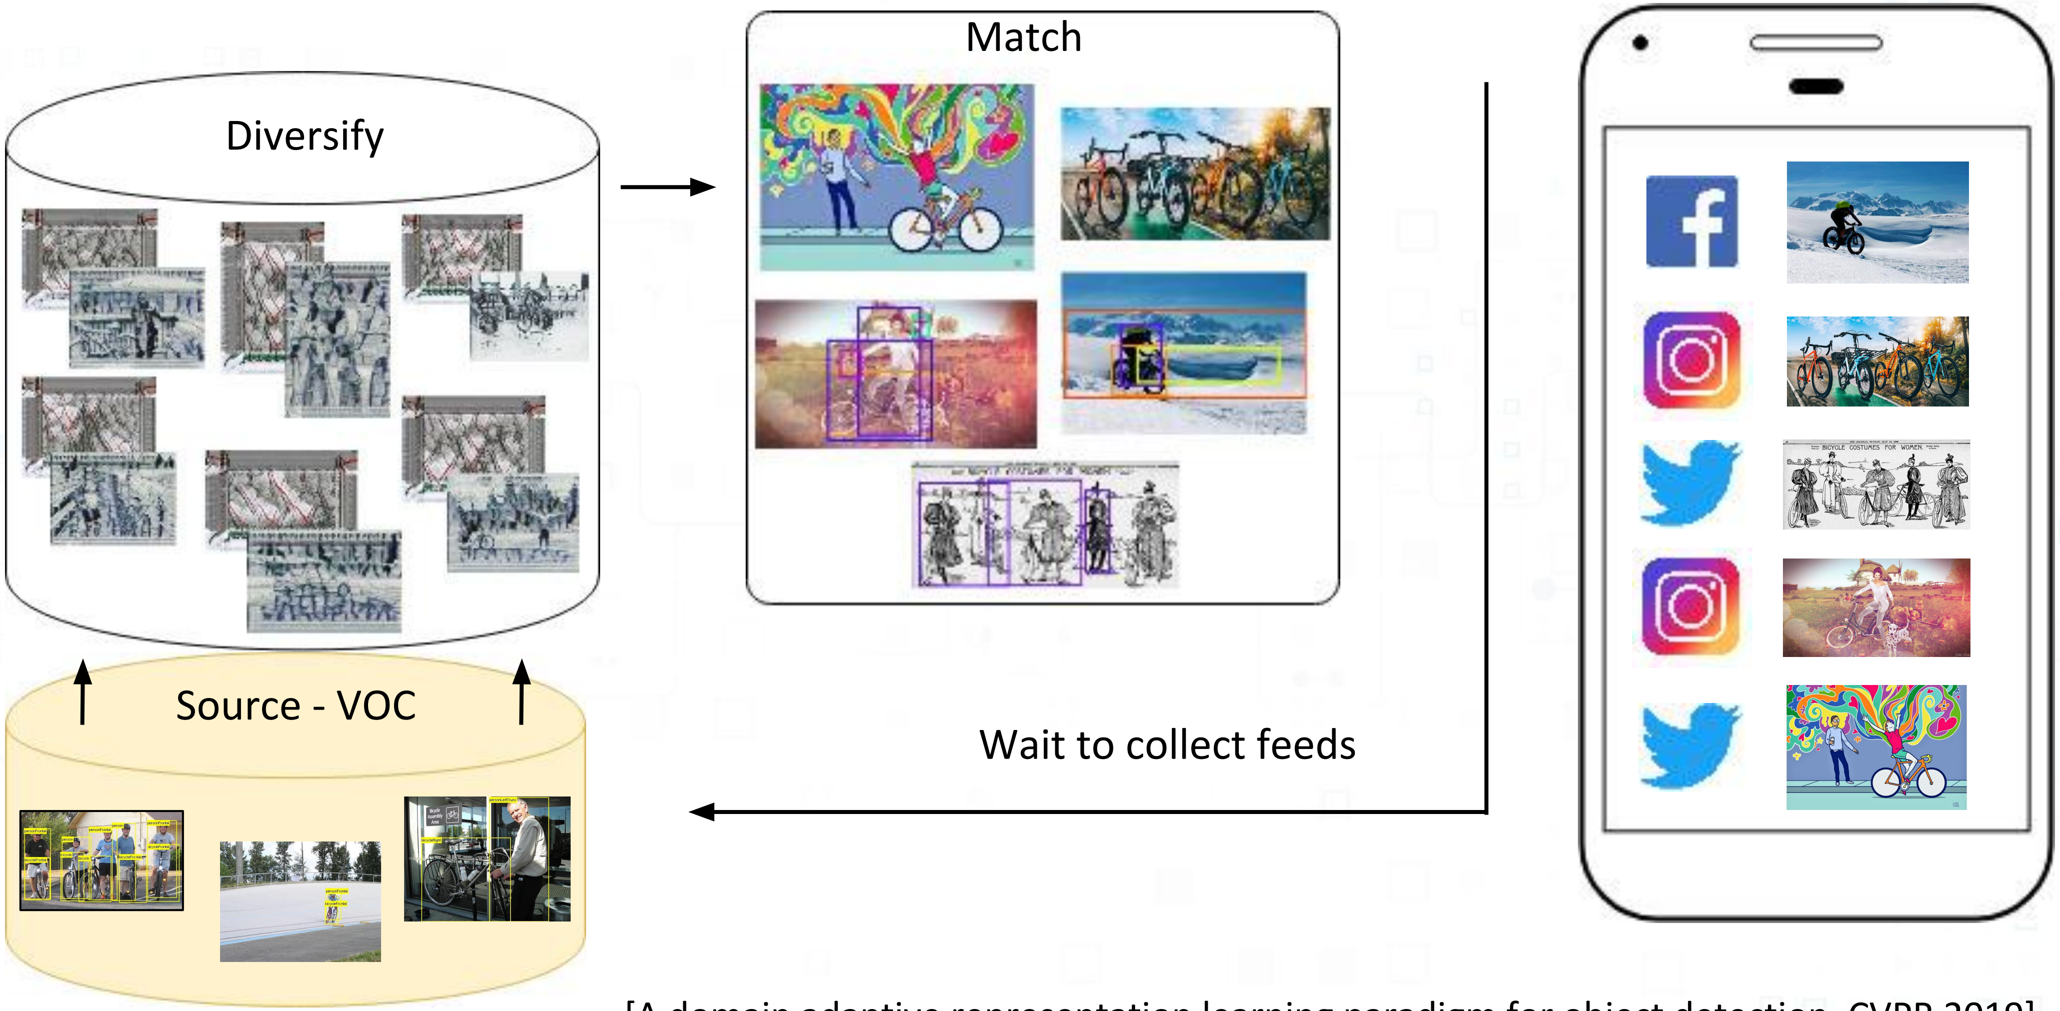

[A domain adaptive representation learning paradigm for object detection, CVPR 2019]

One-Shot Unsupervised Cross-Domain Detection

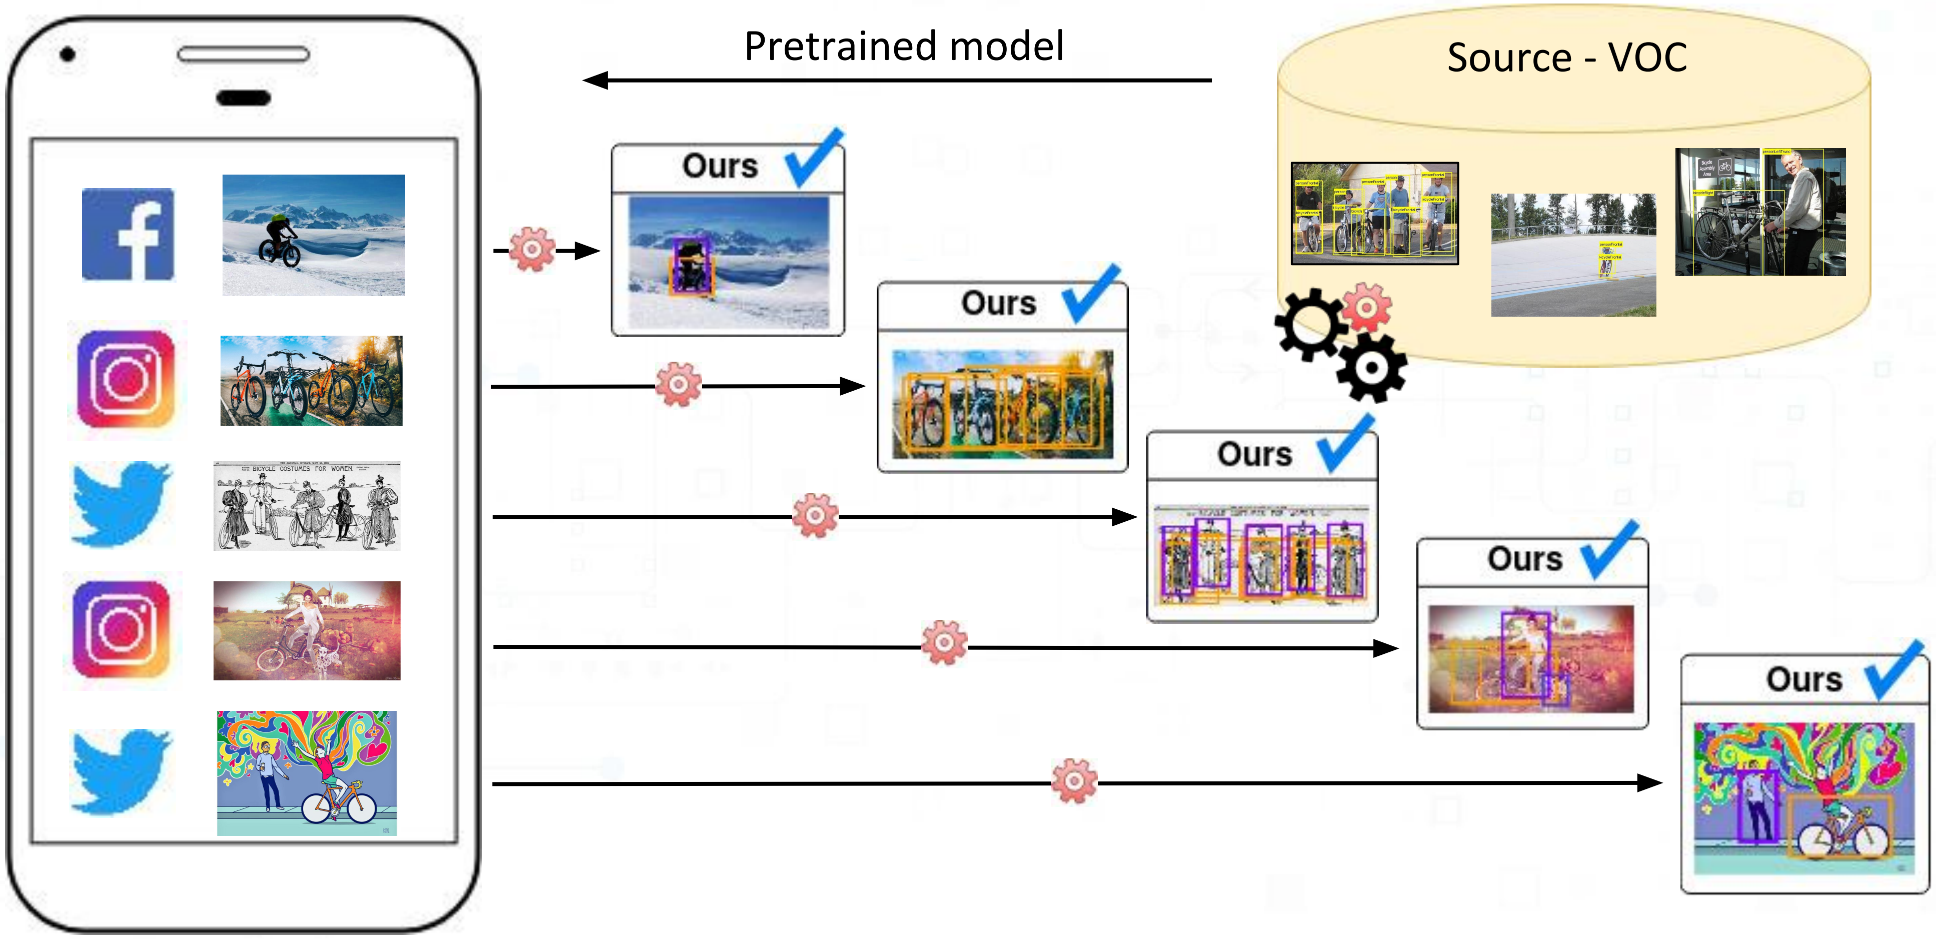

Per-sample unsupervised adaptation

One-Shot Unsupervised Cross-Domain Detection

[One-Shot Unsupervised Cross-Domain Detection, ECCV2020]

OSHOT

1. start from FasterCNN

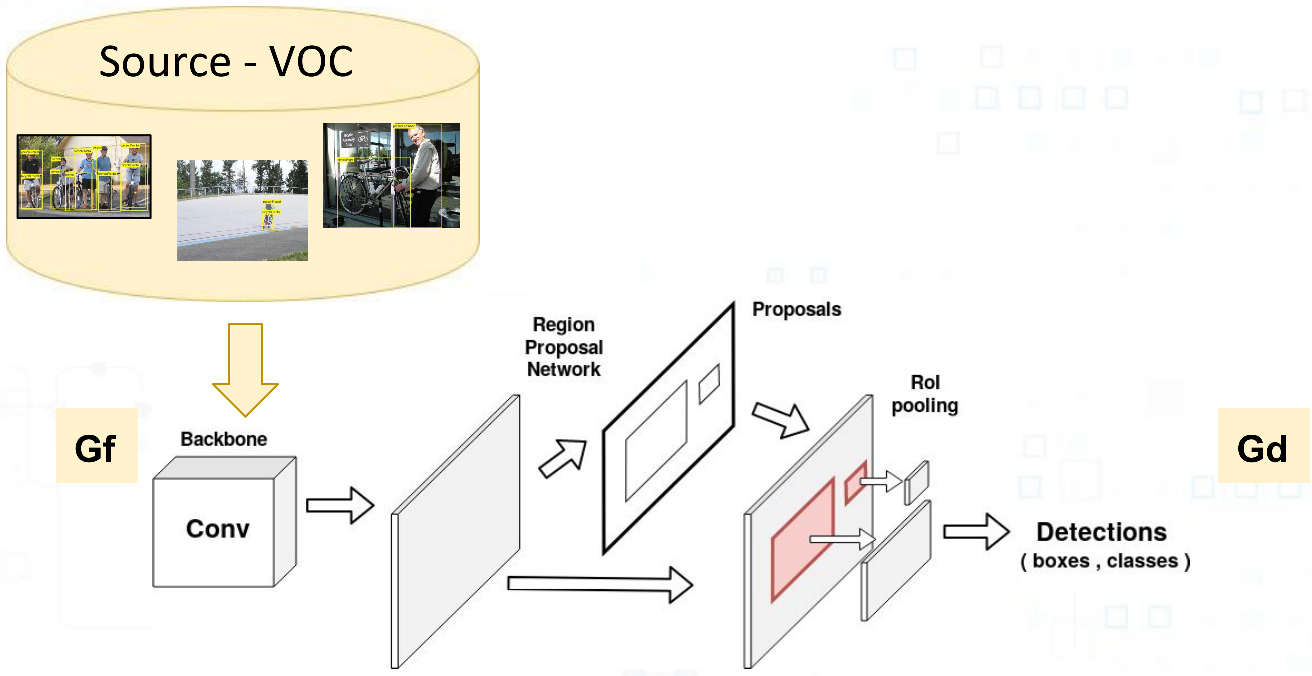

[Towards real-time object detection with region proposal networks, NIPS 2015]

One-Shot Unsupervised Cross-Domain Detection

[One-Shot Unsupervised Cross-Domain Detection, ECCV2020]

OSHOT

1. start from FasterCNN
2. add a self-supervised task: rotation recognition

$$\underset{\theta_f,\theta_d,\theta_r}{\operatorname{argmin}} \sum_{i=1}^{N} \mathcal{L}_d(G_d(G_f(x_i^s)), y_i^s) + \lambda \sum_{j=1}^{M} \mathcal{L}_r(G_r(G_f(R(x^s)_j)), q_j^s)$$

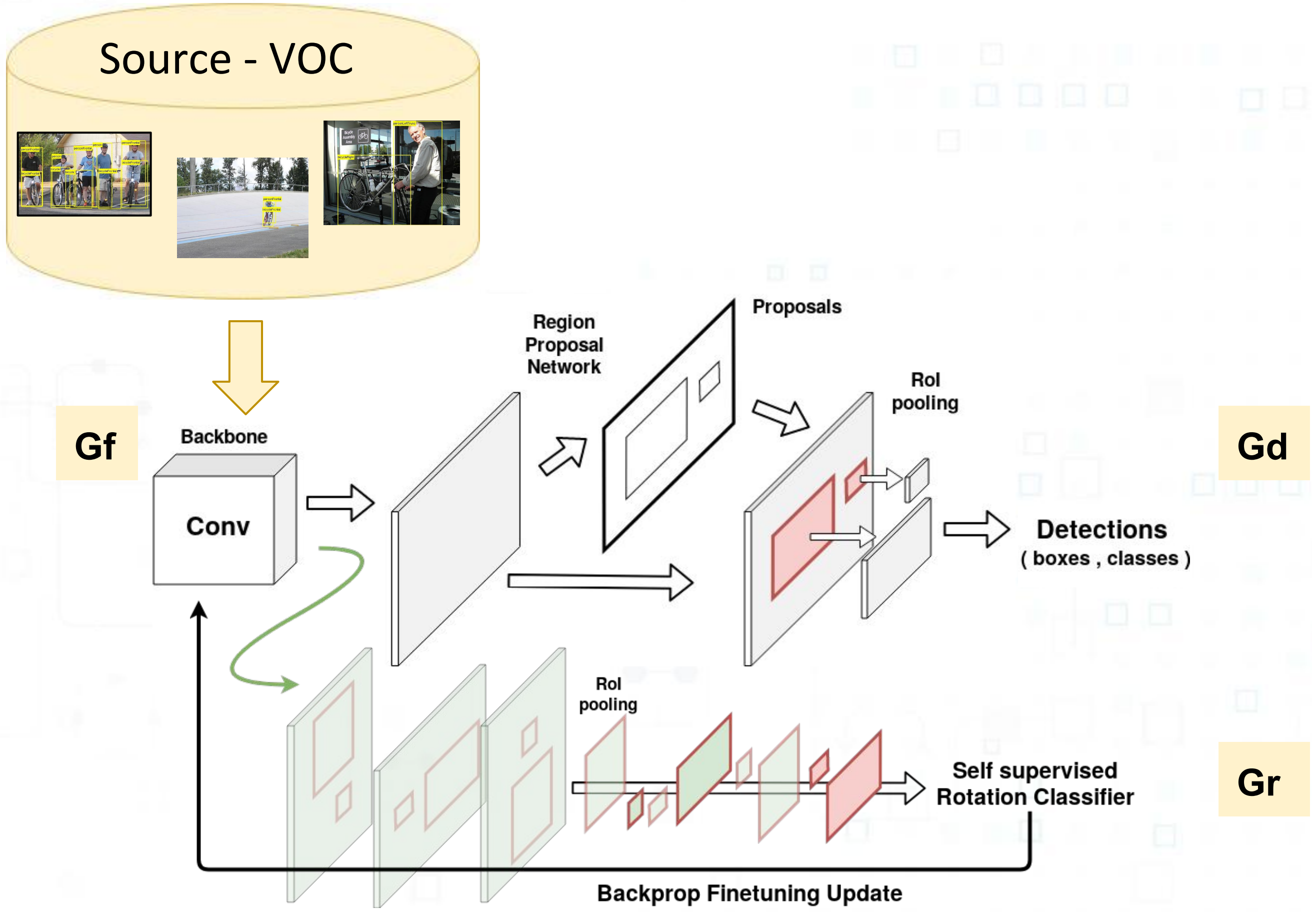

[Towards real-time object detection with region proposal networks, NIPS 2015]

One-Shot Unsupervised Cross-Domain Detection

[One-Shot Unsupervised Cross-Domain Detection, ECCV2020]

OSHOT

1. start from FasterCNN
2. add a self-supervised task: rotation recognition
3. no need to further access the source

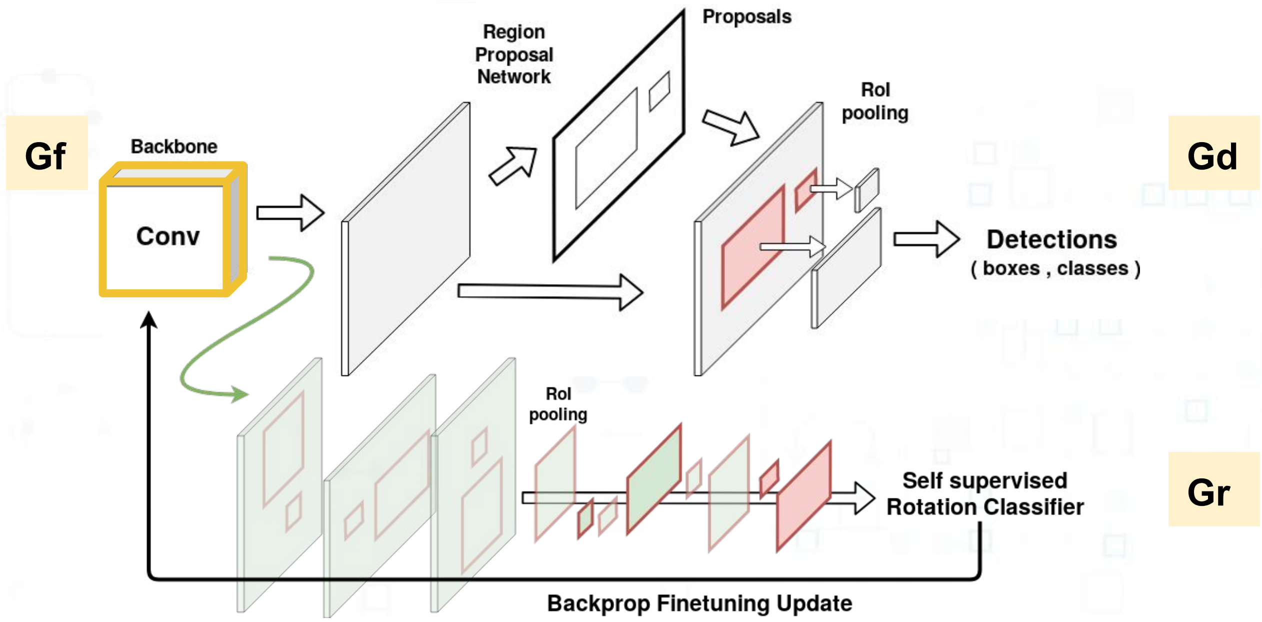

[Towards real-time object detection with region proposal networks, NIPS 2015]

One-Shot Unsupervised Cross-Domain Detection

[One-Shot Unsupervised Cross-Domain Detection, ECCV2020]

OSHOT

1. start from FasterCNN
2. add a self-supervised task: rotation recognition
3. no need to further access the source
4. sample-guided self-supervised fine-tuning

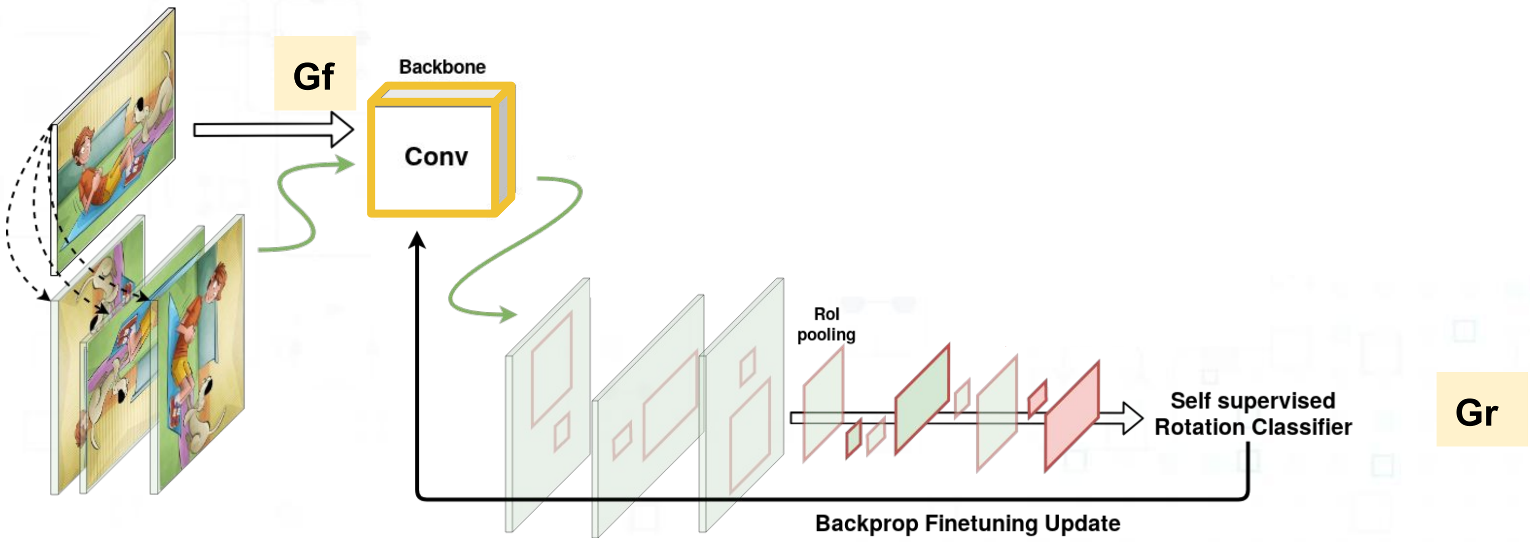

[Towards real-time object detection with region proposal networks, NIPS 2015]

One-Shot Unsupervised Cross-Domain Detection

[One-Shot Unsupervised Cross-Domain Detection, ECCV2020]

OSHOT

$$y^t = G_d(G_f(x^t))$$

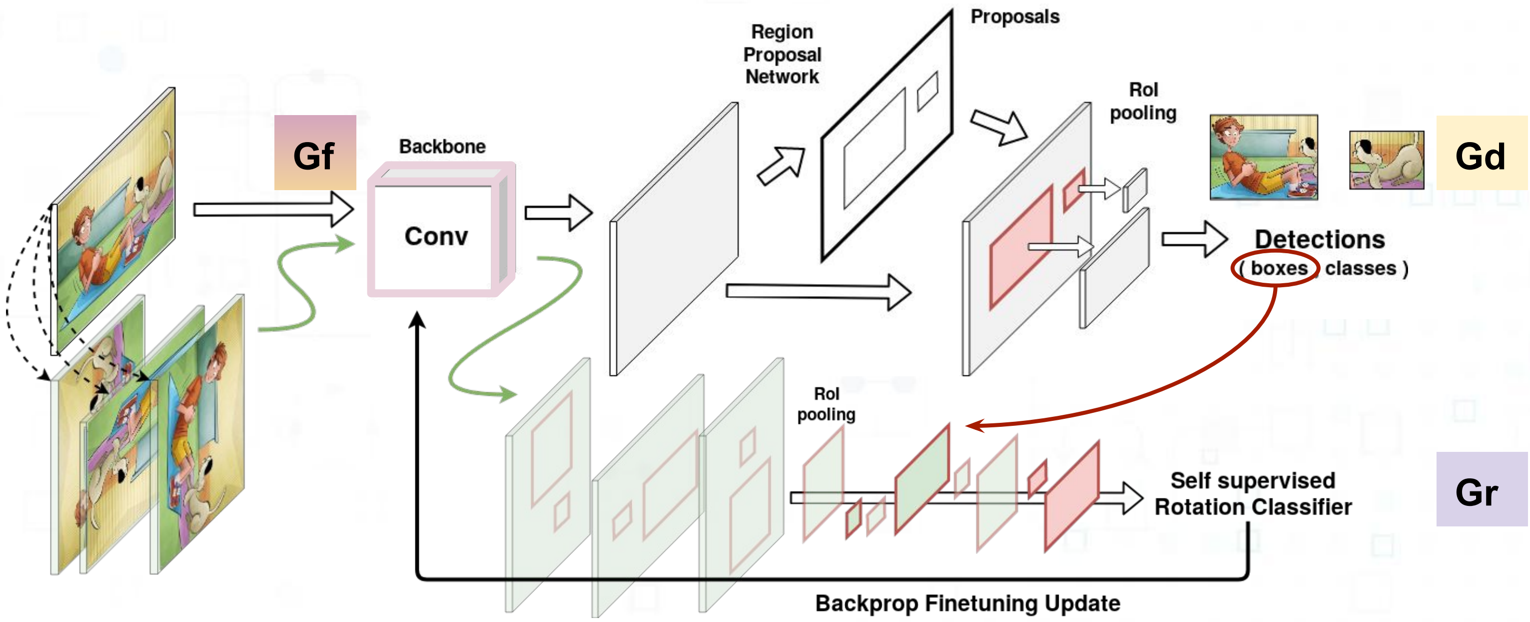

[Towards real-time object detection with region proposal networks, NIPS 2015]

Adapting to a stream of Social Images

[One-Shot Unsupervised Cross-Domain Detection, ECCV2020]

One-Shot Target			
Method	person	bicycle	mAP
FRCNN	67.7	56.6	62.1
OSHOT $(\gamma = 0)$	72.1	52.8	62.4
OSHOT $(\gamma = 30)$	69.4	59.4	**64.4**
Full *Target*			
DivMatch [28]	63.7	51.7	57.7
SW [42]	63.2	44.3	53.7

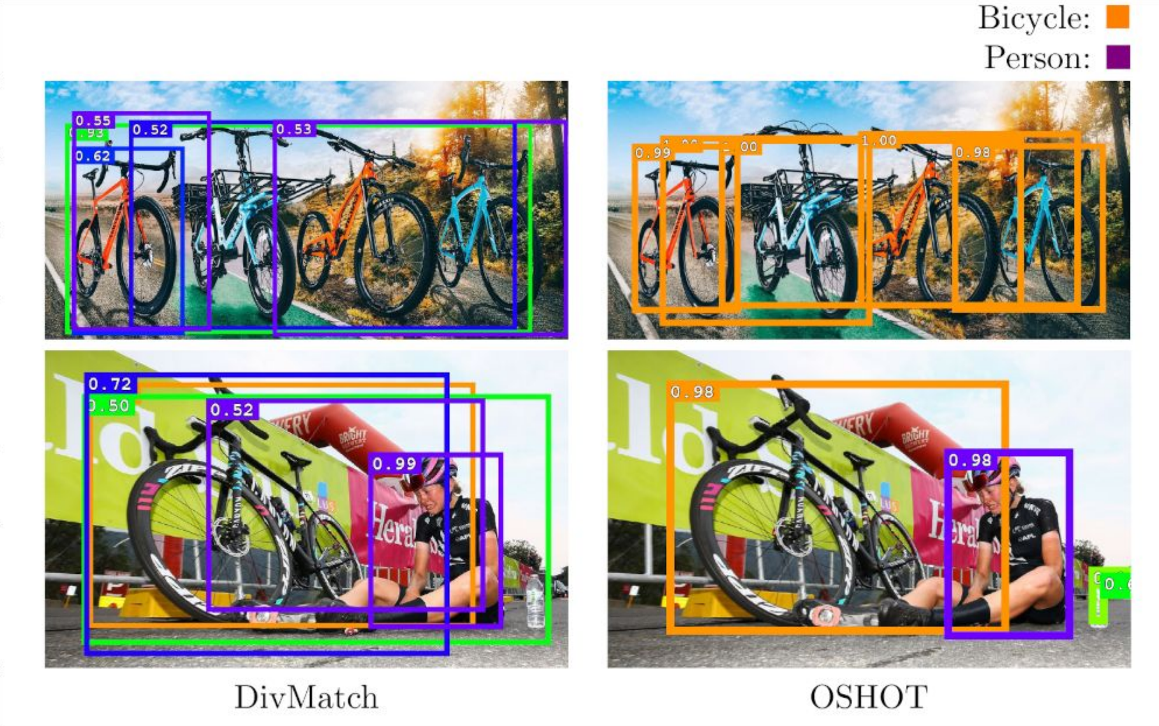

[28] [Diversify and match: A domain adaptive representation learning paradigm for object detection, CVPR 2019]
[42] [Saito et al: Strong-weak distribution alignment for adaptive object detection, CVPR 2019]

From Real to Artistic Images

[One-Shot Unsupervised Cross-Domain Detection, ECCV2020]

(a) VOC → Clipart

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
***One-Shot** Target*																					
FRCNN	18.5	43.3	20.4	13.3	21.0	47.8	29.0	16.9	28.8	12.5	19.5	17.1	23.8	40.6	34.9	34.7	9.1	18.3	40.2	38.0	26.4
OSHOT ($\gamma = 0$)	23.1	55.3	22.7	21.4	26.8	53.3	28.9	4.6	31.4	9.2	27.8	9.6	30.9	47.0	38.2	35.2	11.1	20.4	36.0	33.6	28.3
OSHOT ($\gamma = 10$)	25.4	61.6	23.8	21.1	31.3	55.1	31.6	5.3	34.0	10.1	28.8	7.3	33.1	59.9	44.2	38.8	15.9	19.1	39.5	33.9	31.0
OSHOT ($\gamma = 30$)	25.4	56.0	24.7	25.3	36.7	58.0	34.4	5.9	34.9	10.3	29.2	11.8	46.9	70.9	52.9	41.5	21.1	21.0	38.5	31.8	**33.9**
***Ten-Shot** Target*																					
DivMatch [28]	19.5	57.2	17.0	23.8	14.4	25.4	29.4	2.7	35.0	8.4	22.9	14.2	30.0	55.6	50.8	30.2	1.9	12.3	37.8	37.2	26.3
SW [42]	21.5	39.9	21.7	20.5	32.7	34.1	25.1	8.5	33.2	10.9	15.2	3.4	32.2	56.9	46.5	35.4	14.7	15.2	29.2	32.0	26.4

(b) VOC → Comic

Method	bike	bird	car	cat	dog	person	mAP
***One-Shot** Target*							
FRCNN	25.2	10.0	21.1	14.1	11.0	27.1	18.1
OSHOT ($\gamma = 0$)	26.9	11.6	22.7	9.1	14.2	28.3	18.8
OSHOT ($\gamma = 10$)	35.5	11.7	25.1	9.1	15.8	34.5	22.0
OSHOT ($\gamma = 30$)	35.2	14.4	30.0	14.8	20.0	46.7	**26.9**
***Ten-Shot** Target*							
DivMatch [28]	27.1	12.3	26.2	11.5	13.8	34.0	20.8
SW [42]	21.2	14.8	18.7	12.4	14.9	43.9	21.0

(c) VOC → Watercolor

Method	bike	bird	car	cat	dog	person	mAP
***One-Shot** Target*							
FRCNN	62.5	39.7	43.4	31.9	26.7	52.4	42.8
OSHOT ($\gamma = 0$)	70.2	46.7	45.5	31.2	27.2	55.7	46.1
OSHOT ($\gamma = 10$)	70.2	46.7	48.1	30.9	32.3	59.9	48.0
OSHOT ($\gamma = 30$)	77.1	44.7	52.4	37.3	37.0	63.3	**52.0**
***Ten-Shot** Target*							
DivMatch [28]	64.6	44.1	44.6	34.1	24.9	60.0	45.4
SW [42]	66.3	41.1	41.1	30.5	20.5	52.3	42.0

[28] [Diversify and match: A domain adaptive representation learning paradigm for object detection, CVPR 2019]
[42] [Saito et al: Strong-weak distribution alignment for adaptive object detection, CVPR 2019]

To Summarize...

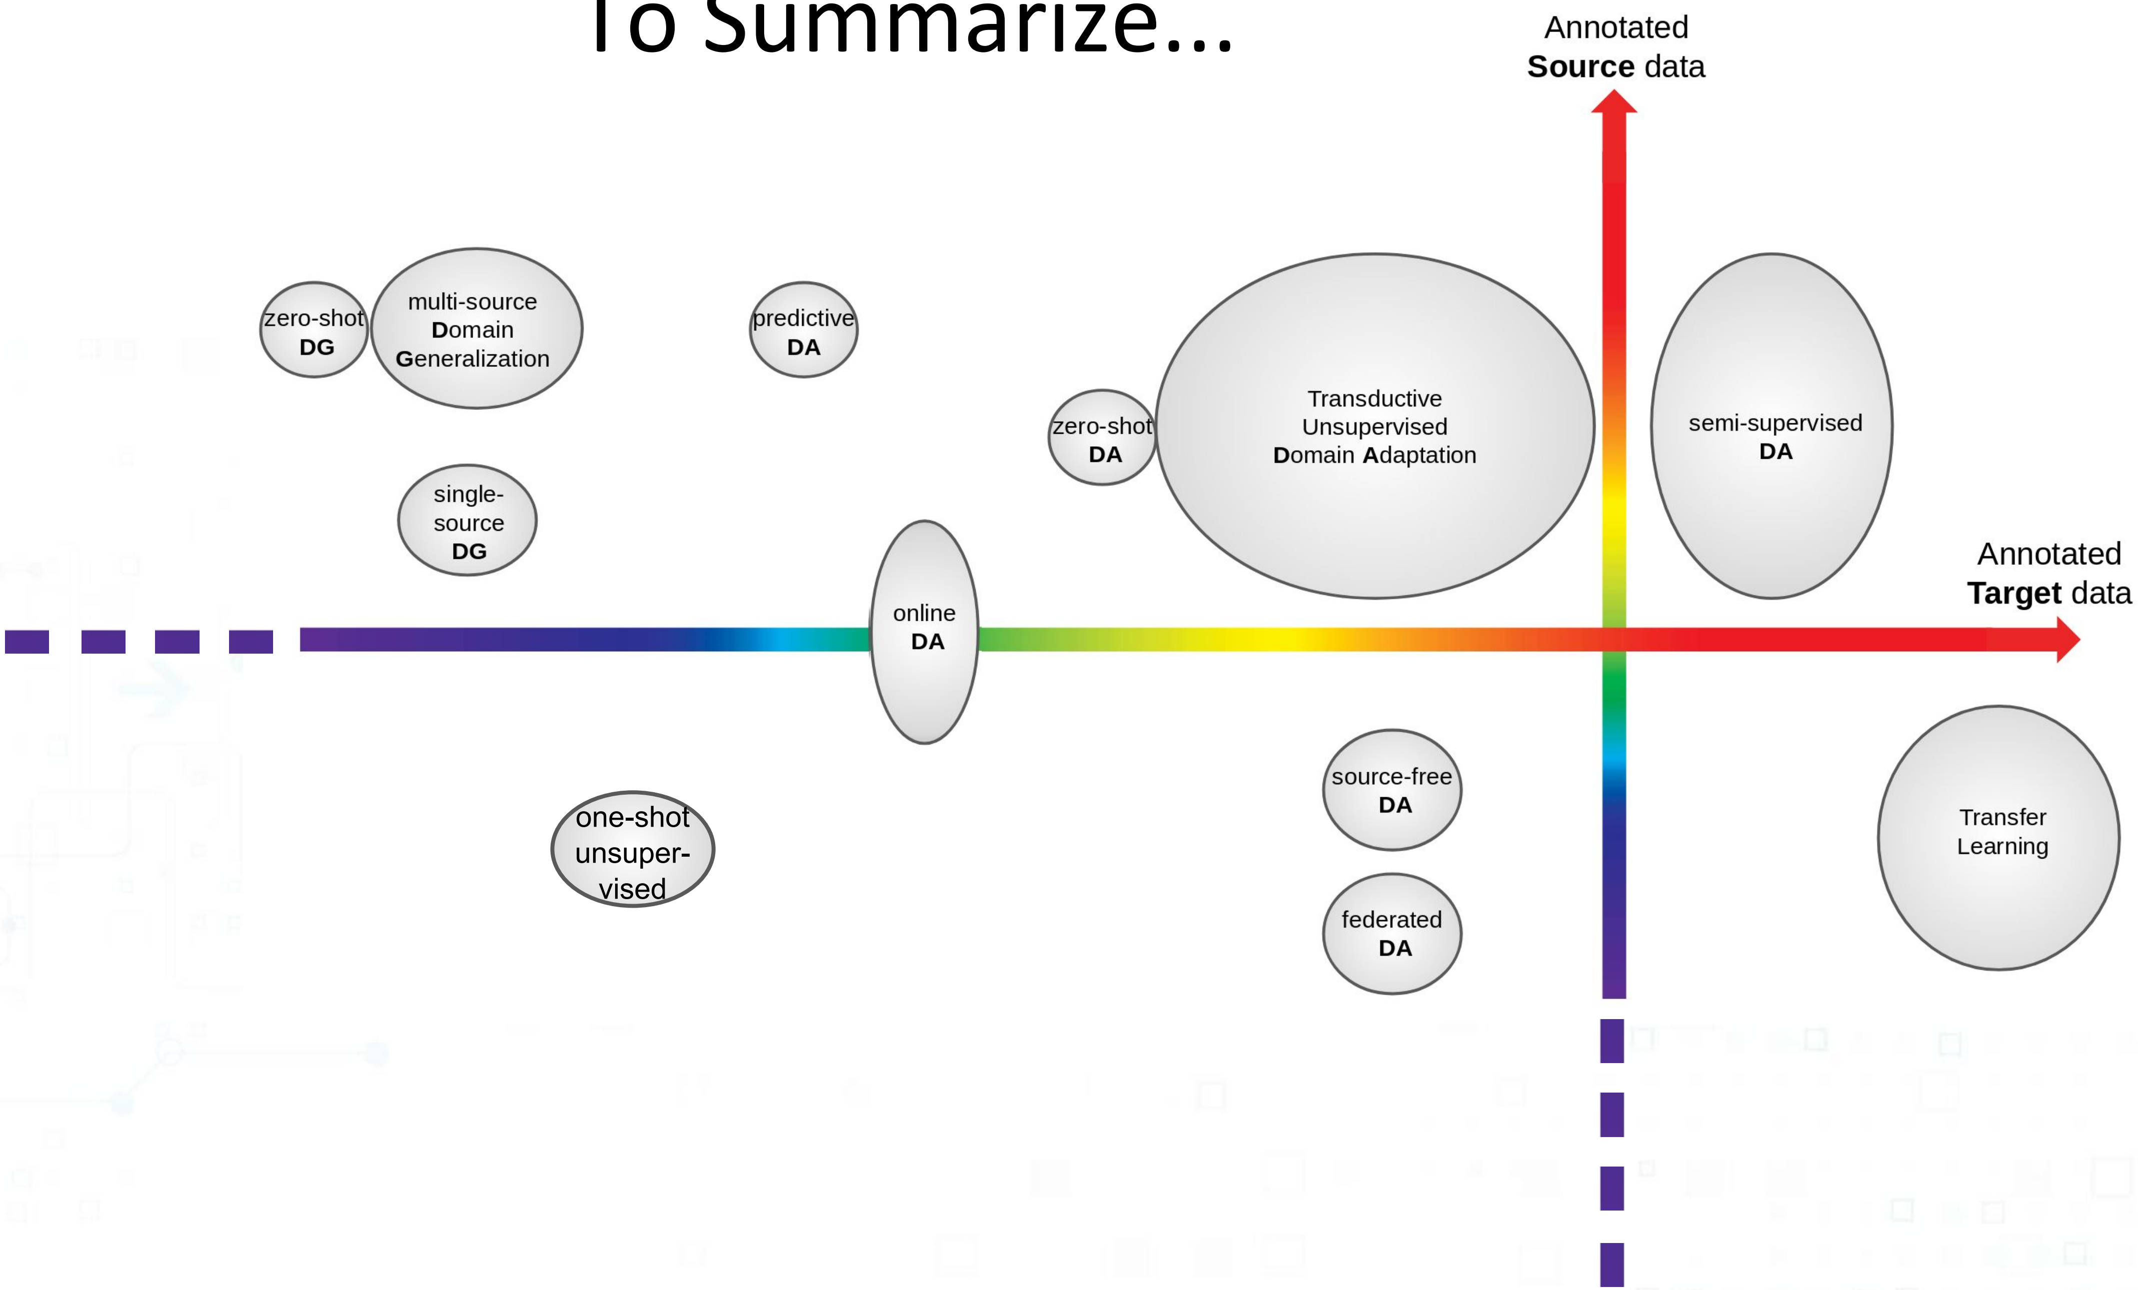

To Summarize...

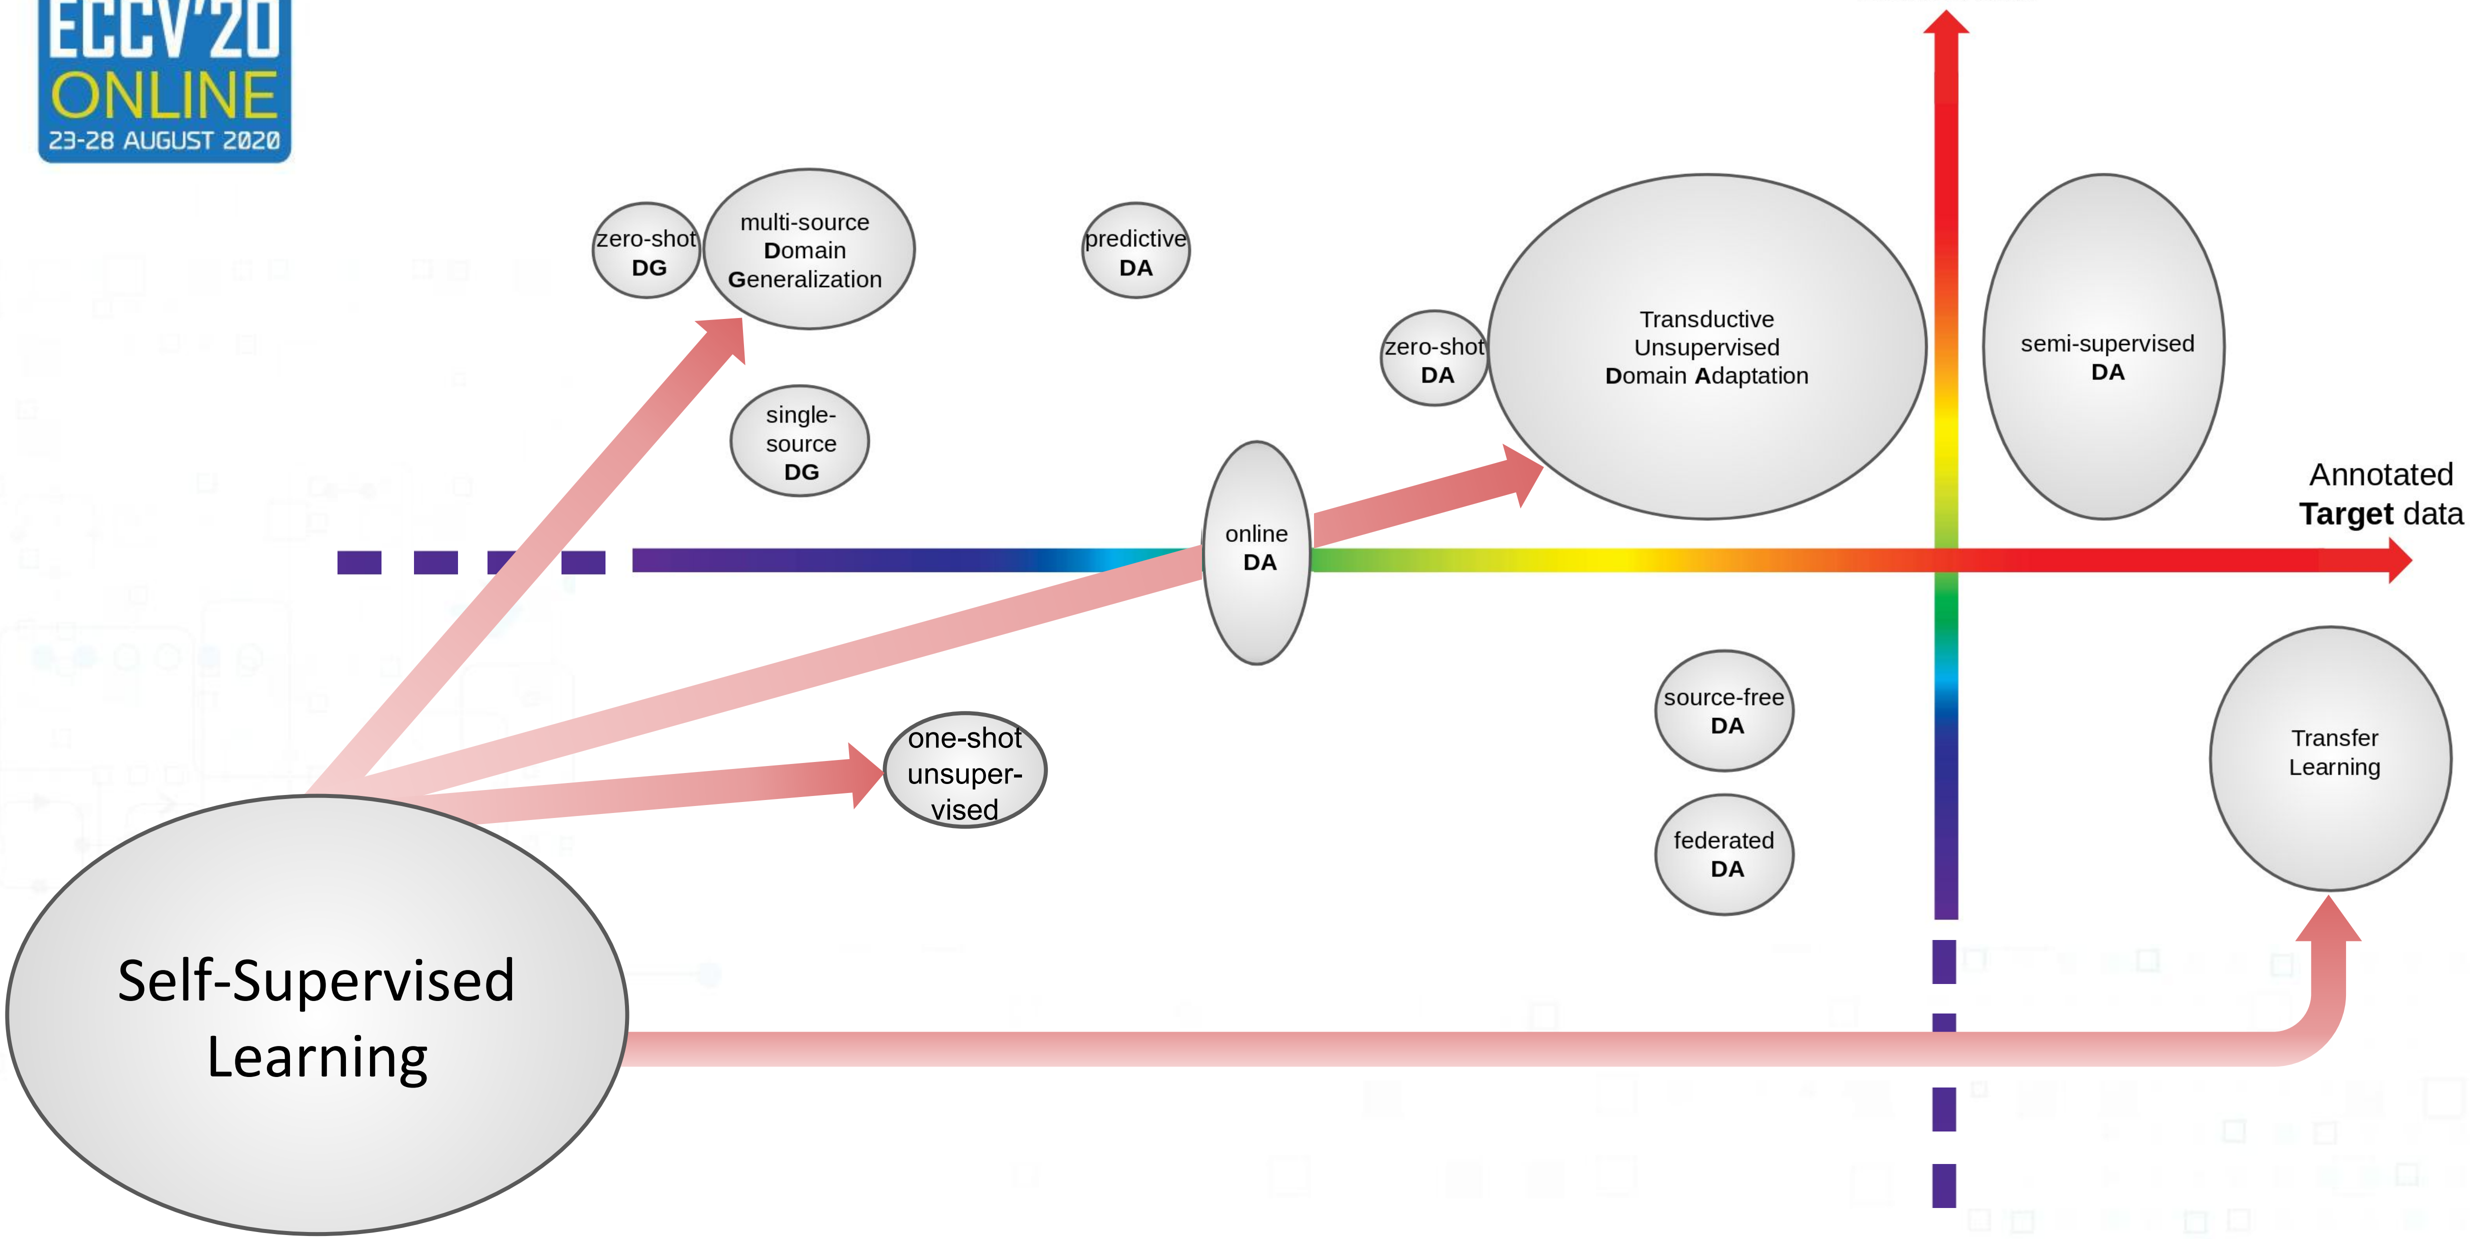

To Summarize...

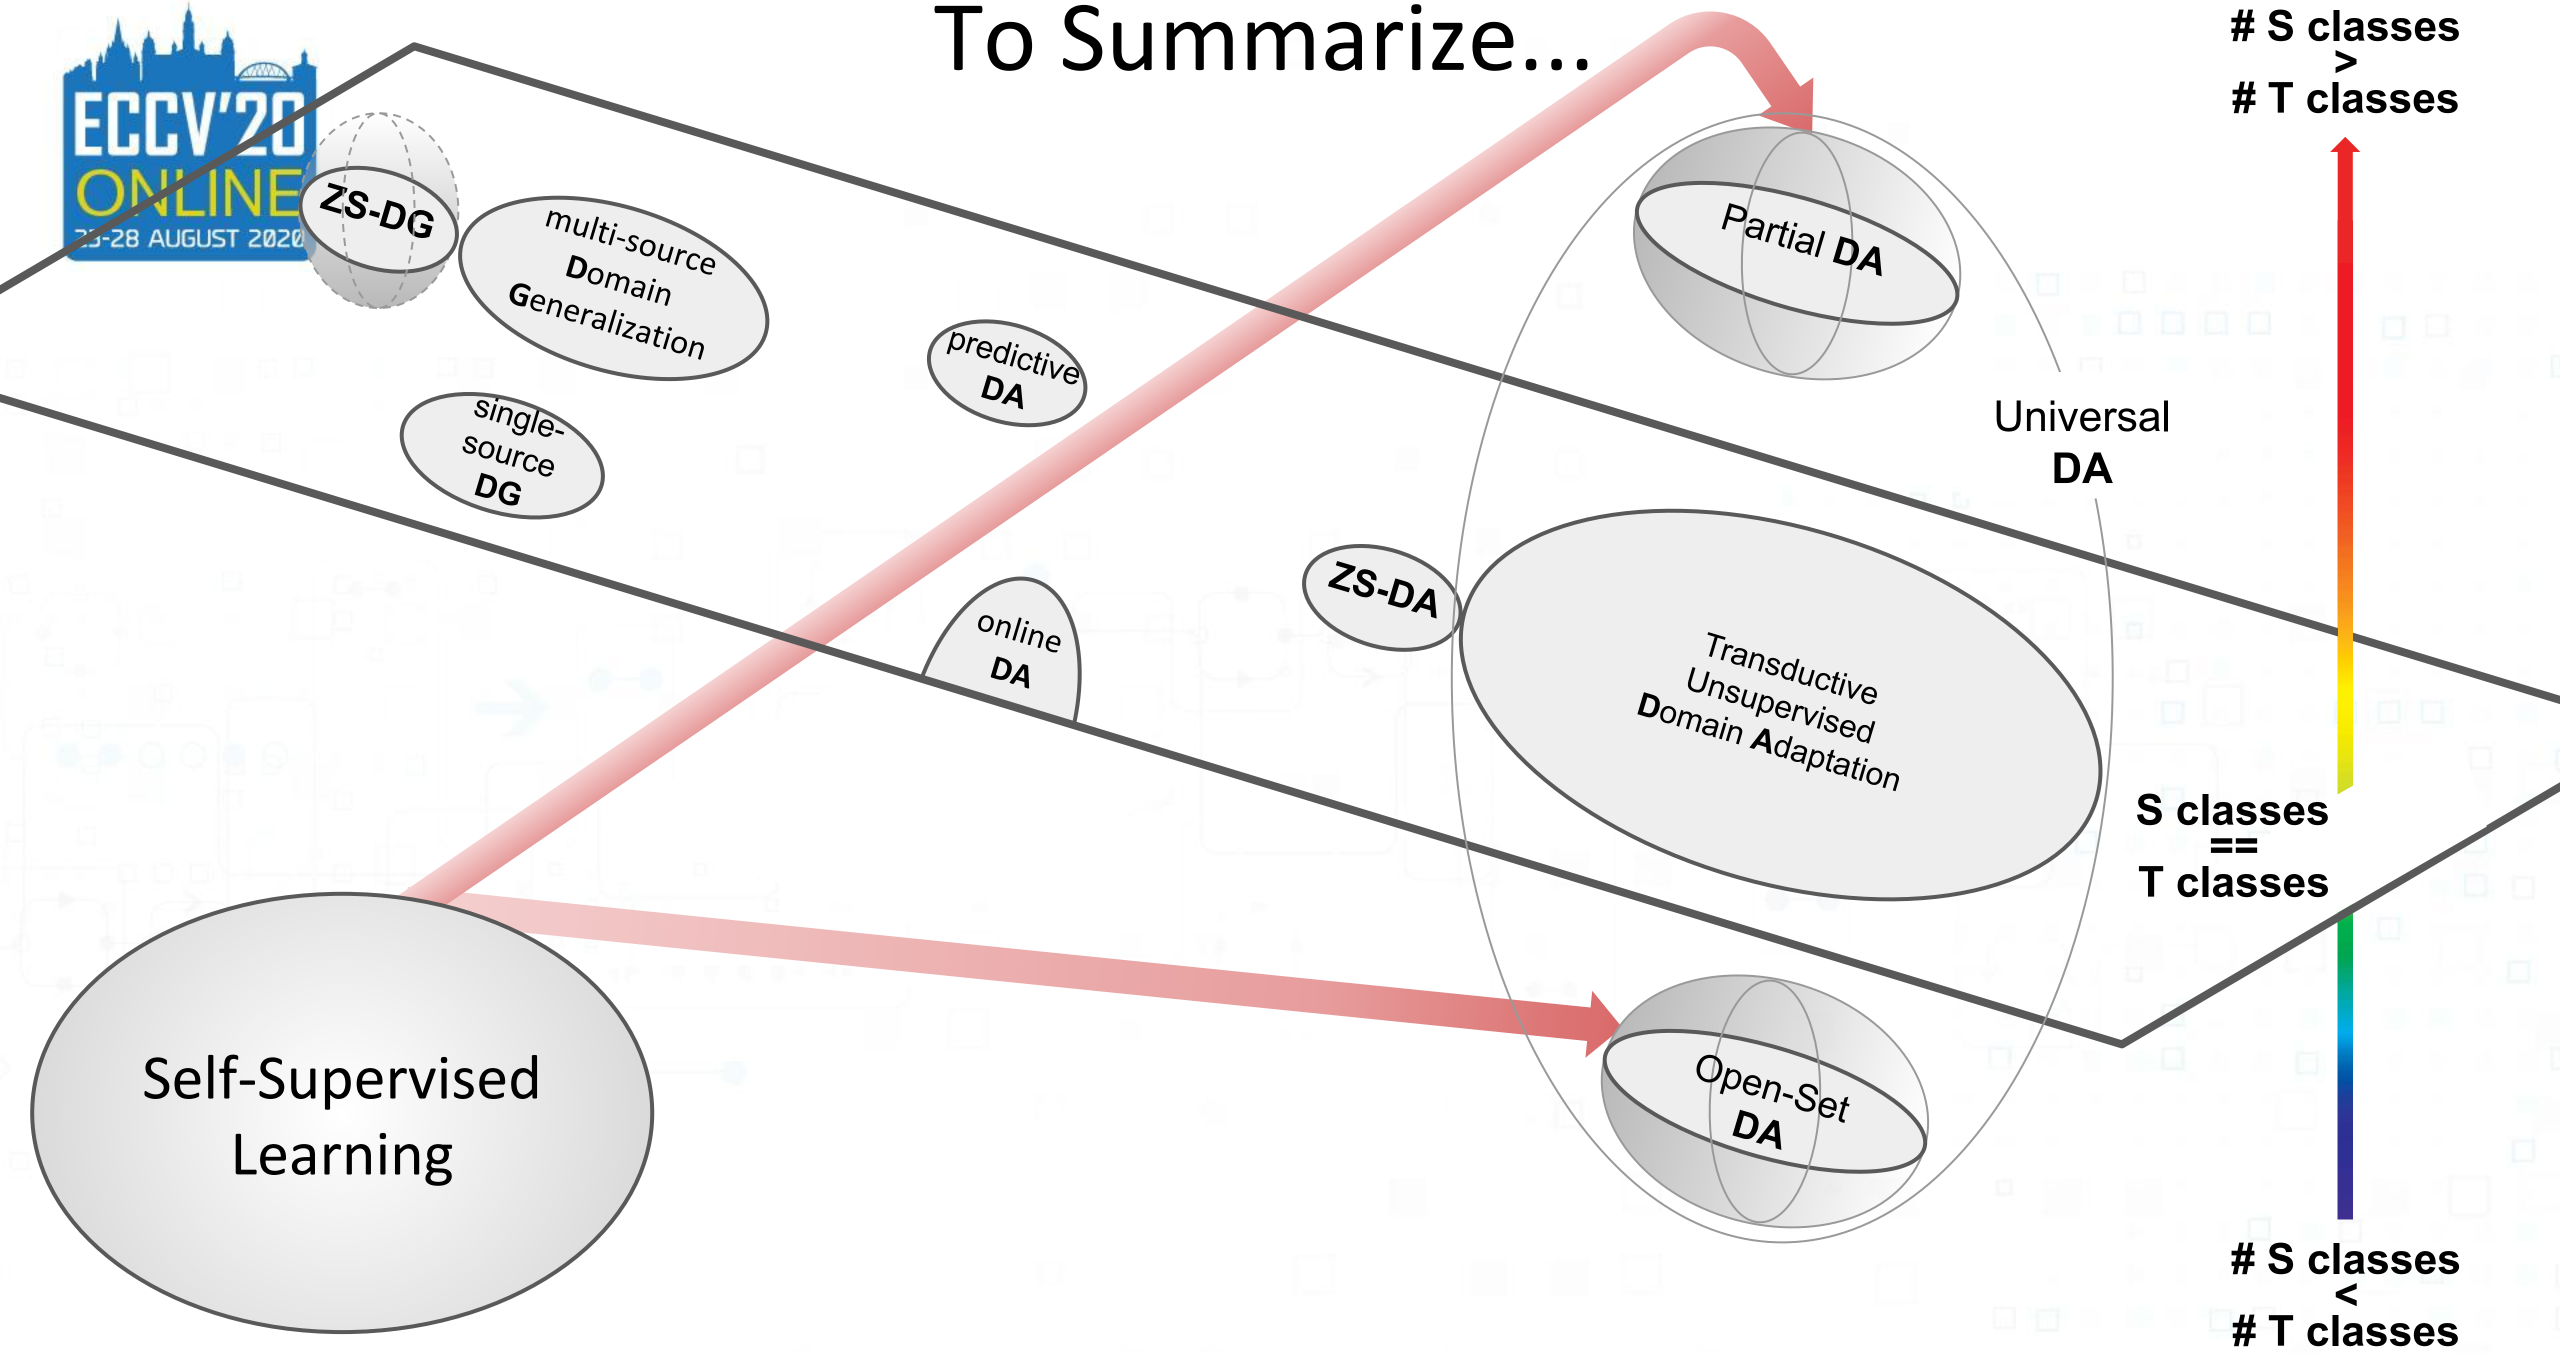

Thanks for your attention

Domain Adaptation for Visual Applications
Part 3: Beyond Classical Domain Adaptation

Tatiana Tommasi
tatiana.tommasi@polito.it