


Topic 2: Domain Generalization




Domain Adaptation and Generalization Visualization

(thanks to Tatiana Tomassi's ECCV 2020 Tutorial)

Classical Domain Adaptation


Source
(Train)



Target
(Test)

Classical Domain Adaptation

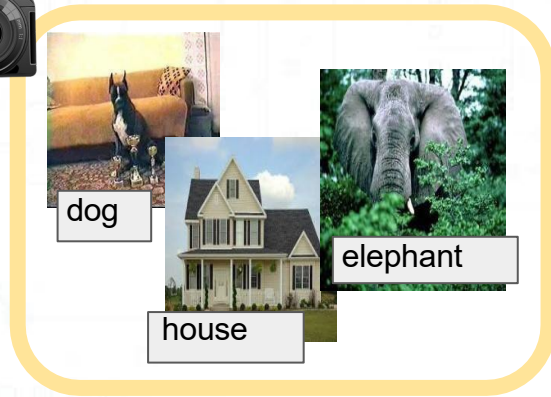
Source
(Train)



Target
(Test)




Labelled
Source Domain



Train

Unlabelled
Target Domain



Unsupervised DA,
transductive setting



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

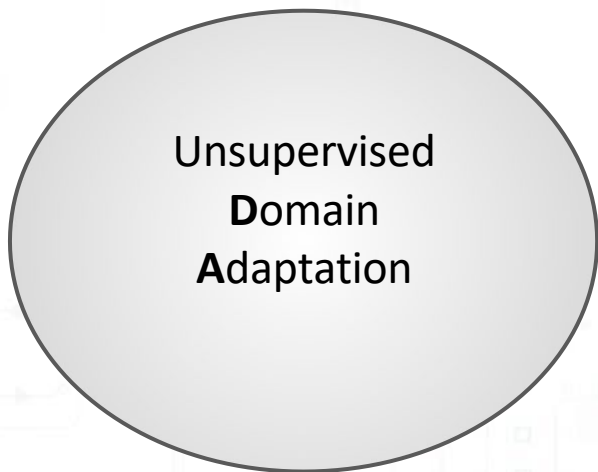
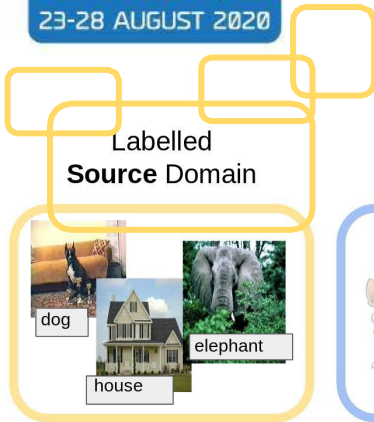
Annotated
Target data

Target data **not**
available at training
time

Target data
available but not
annotated

Only Source Model
available, (**no source**
data)






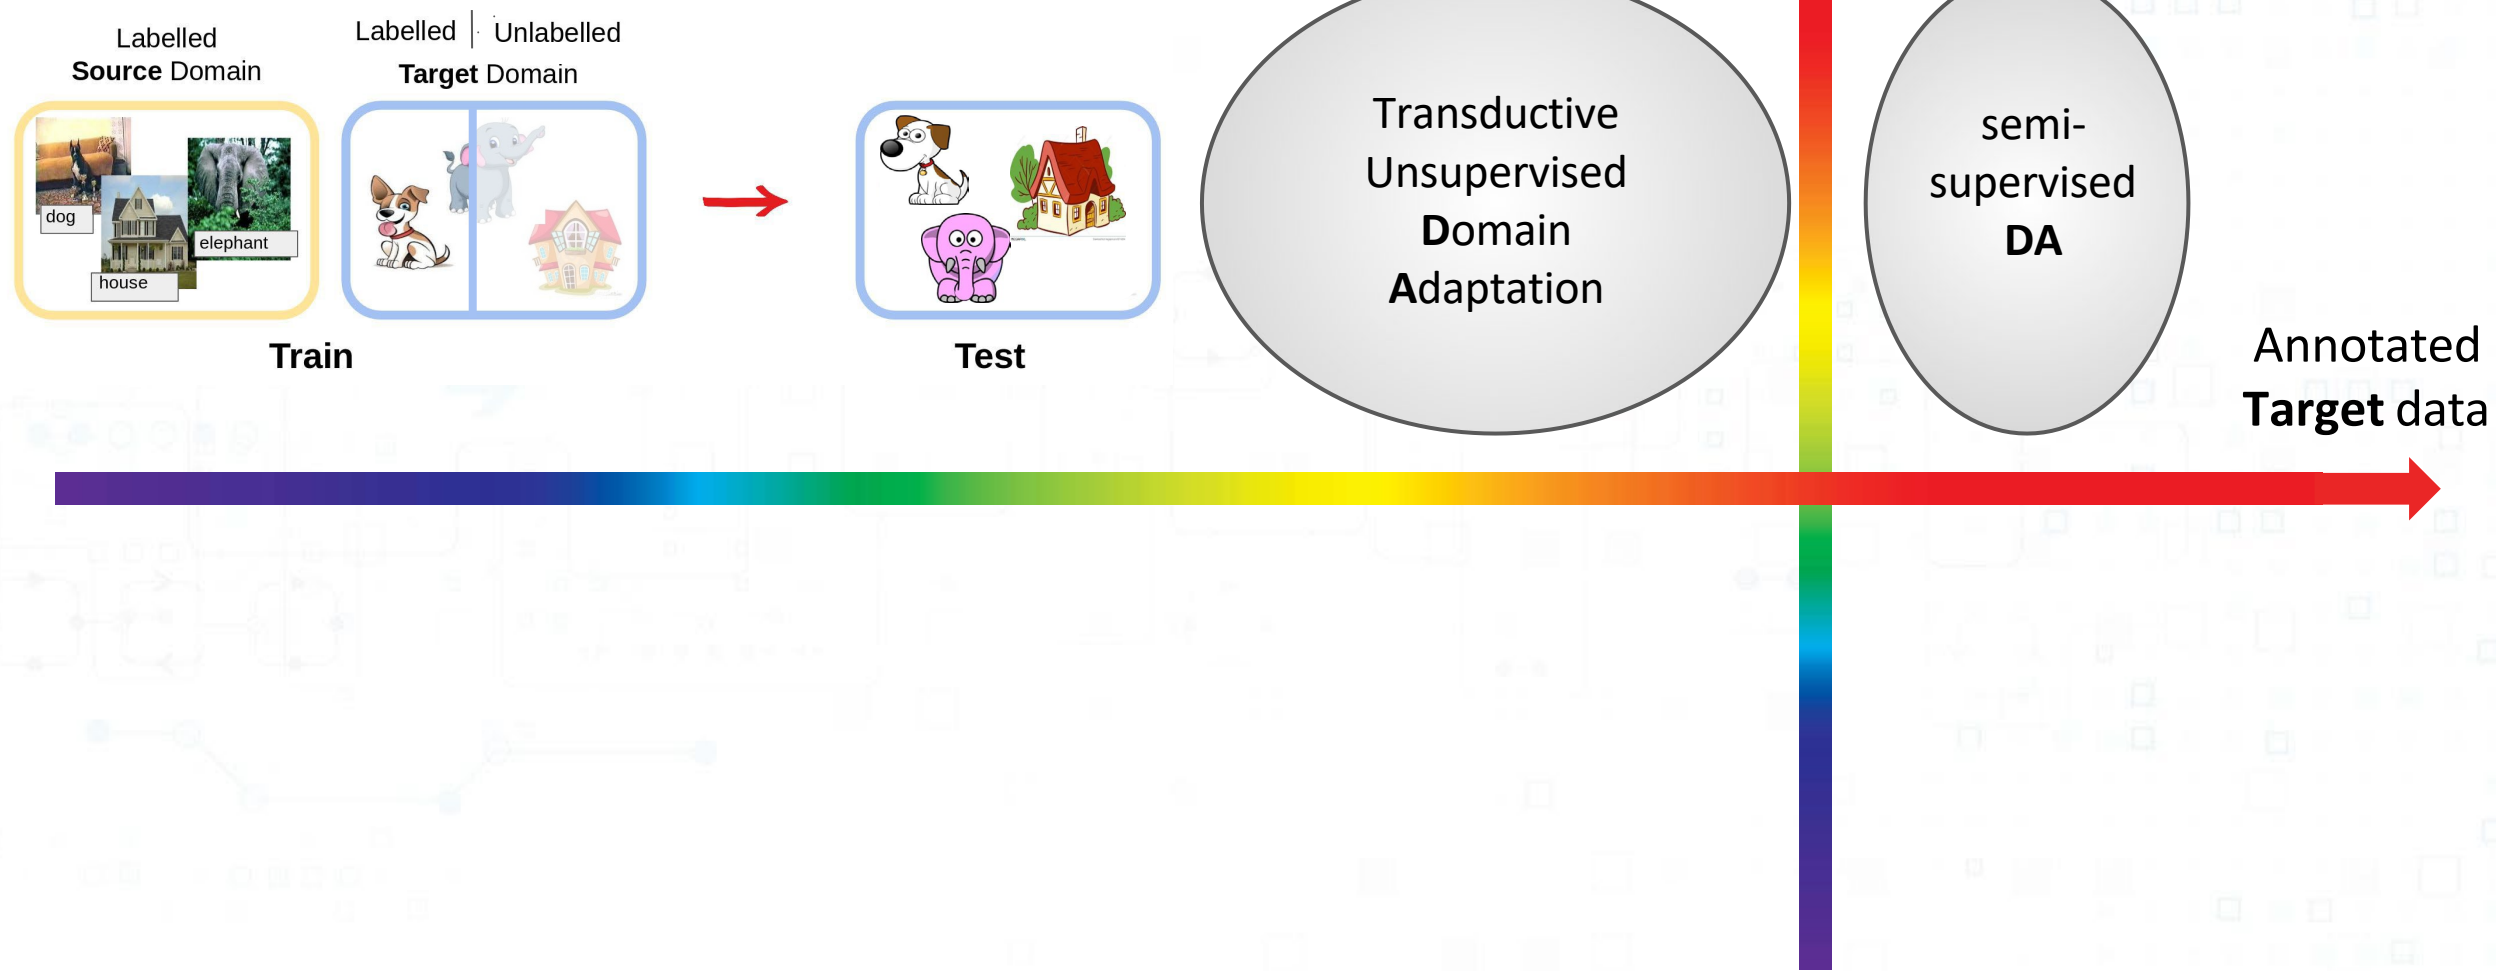
Annotated Source data

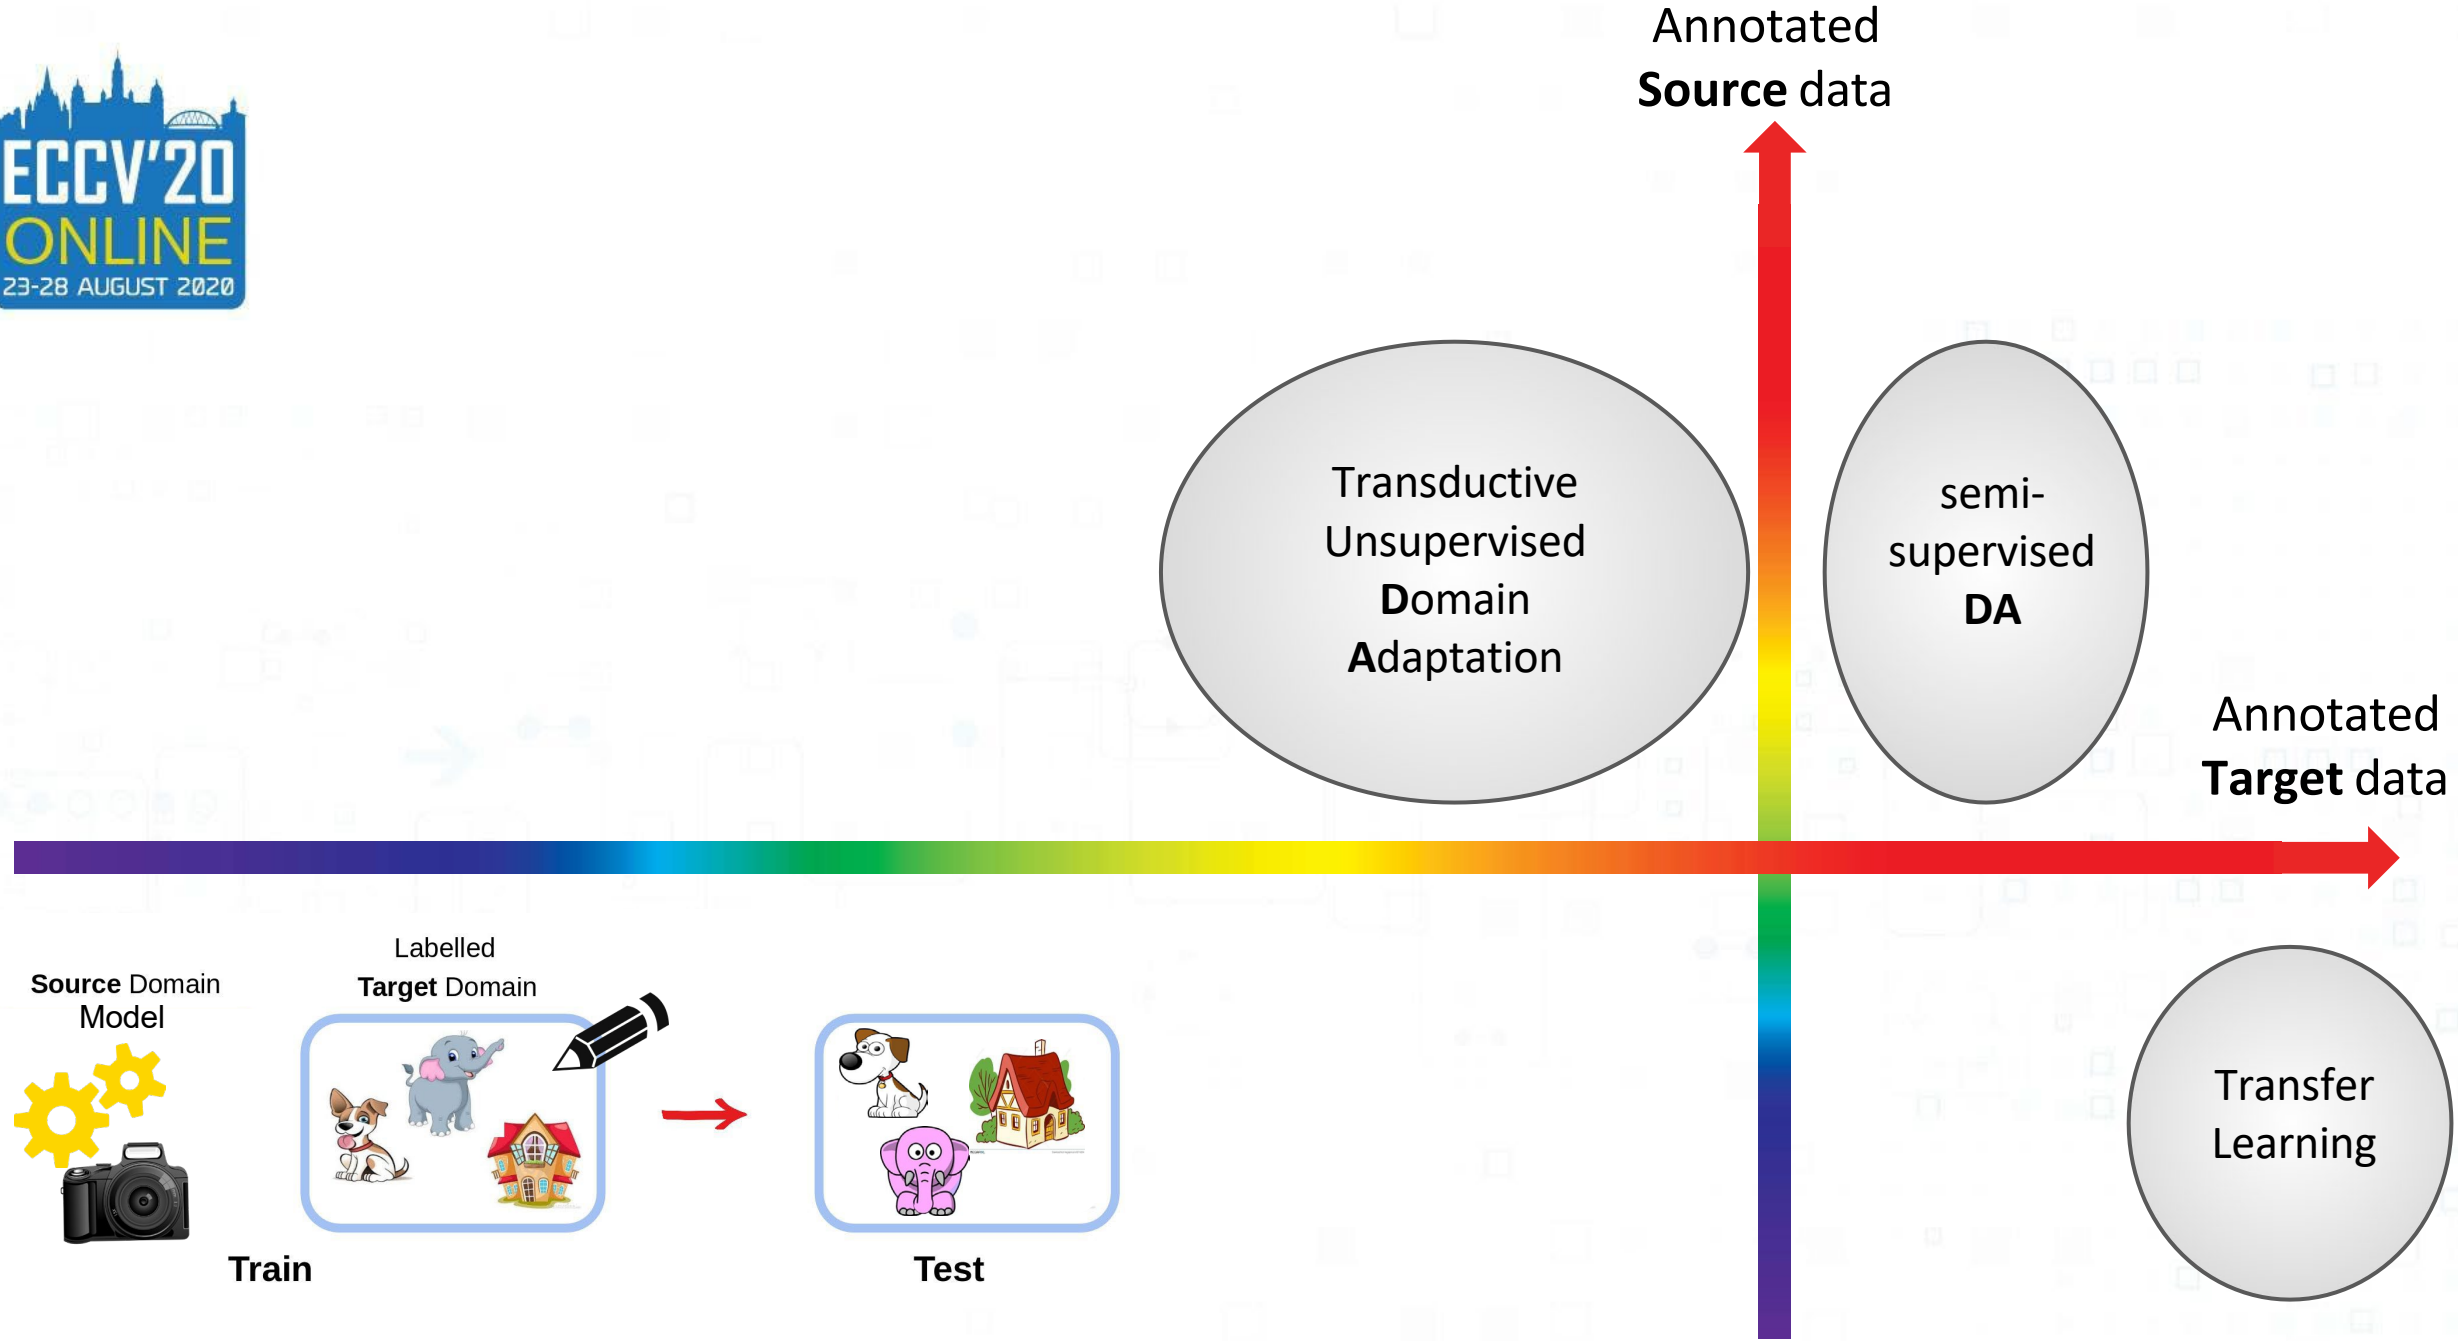
Annotated Target data

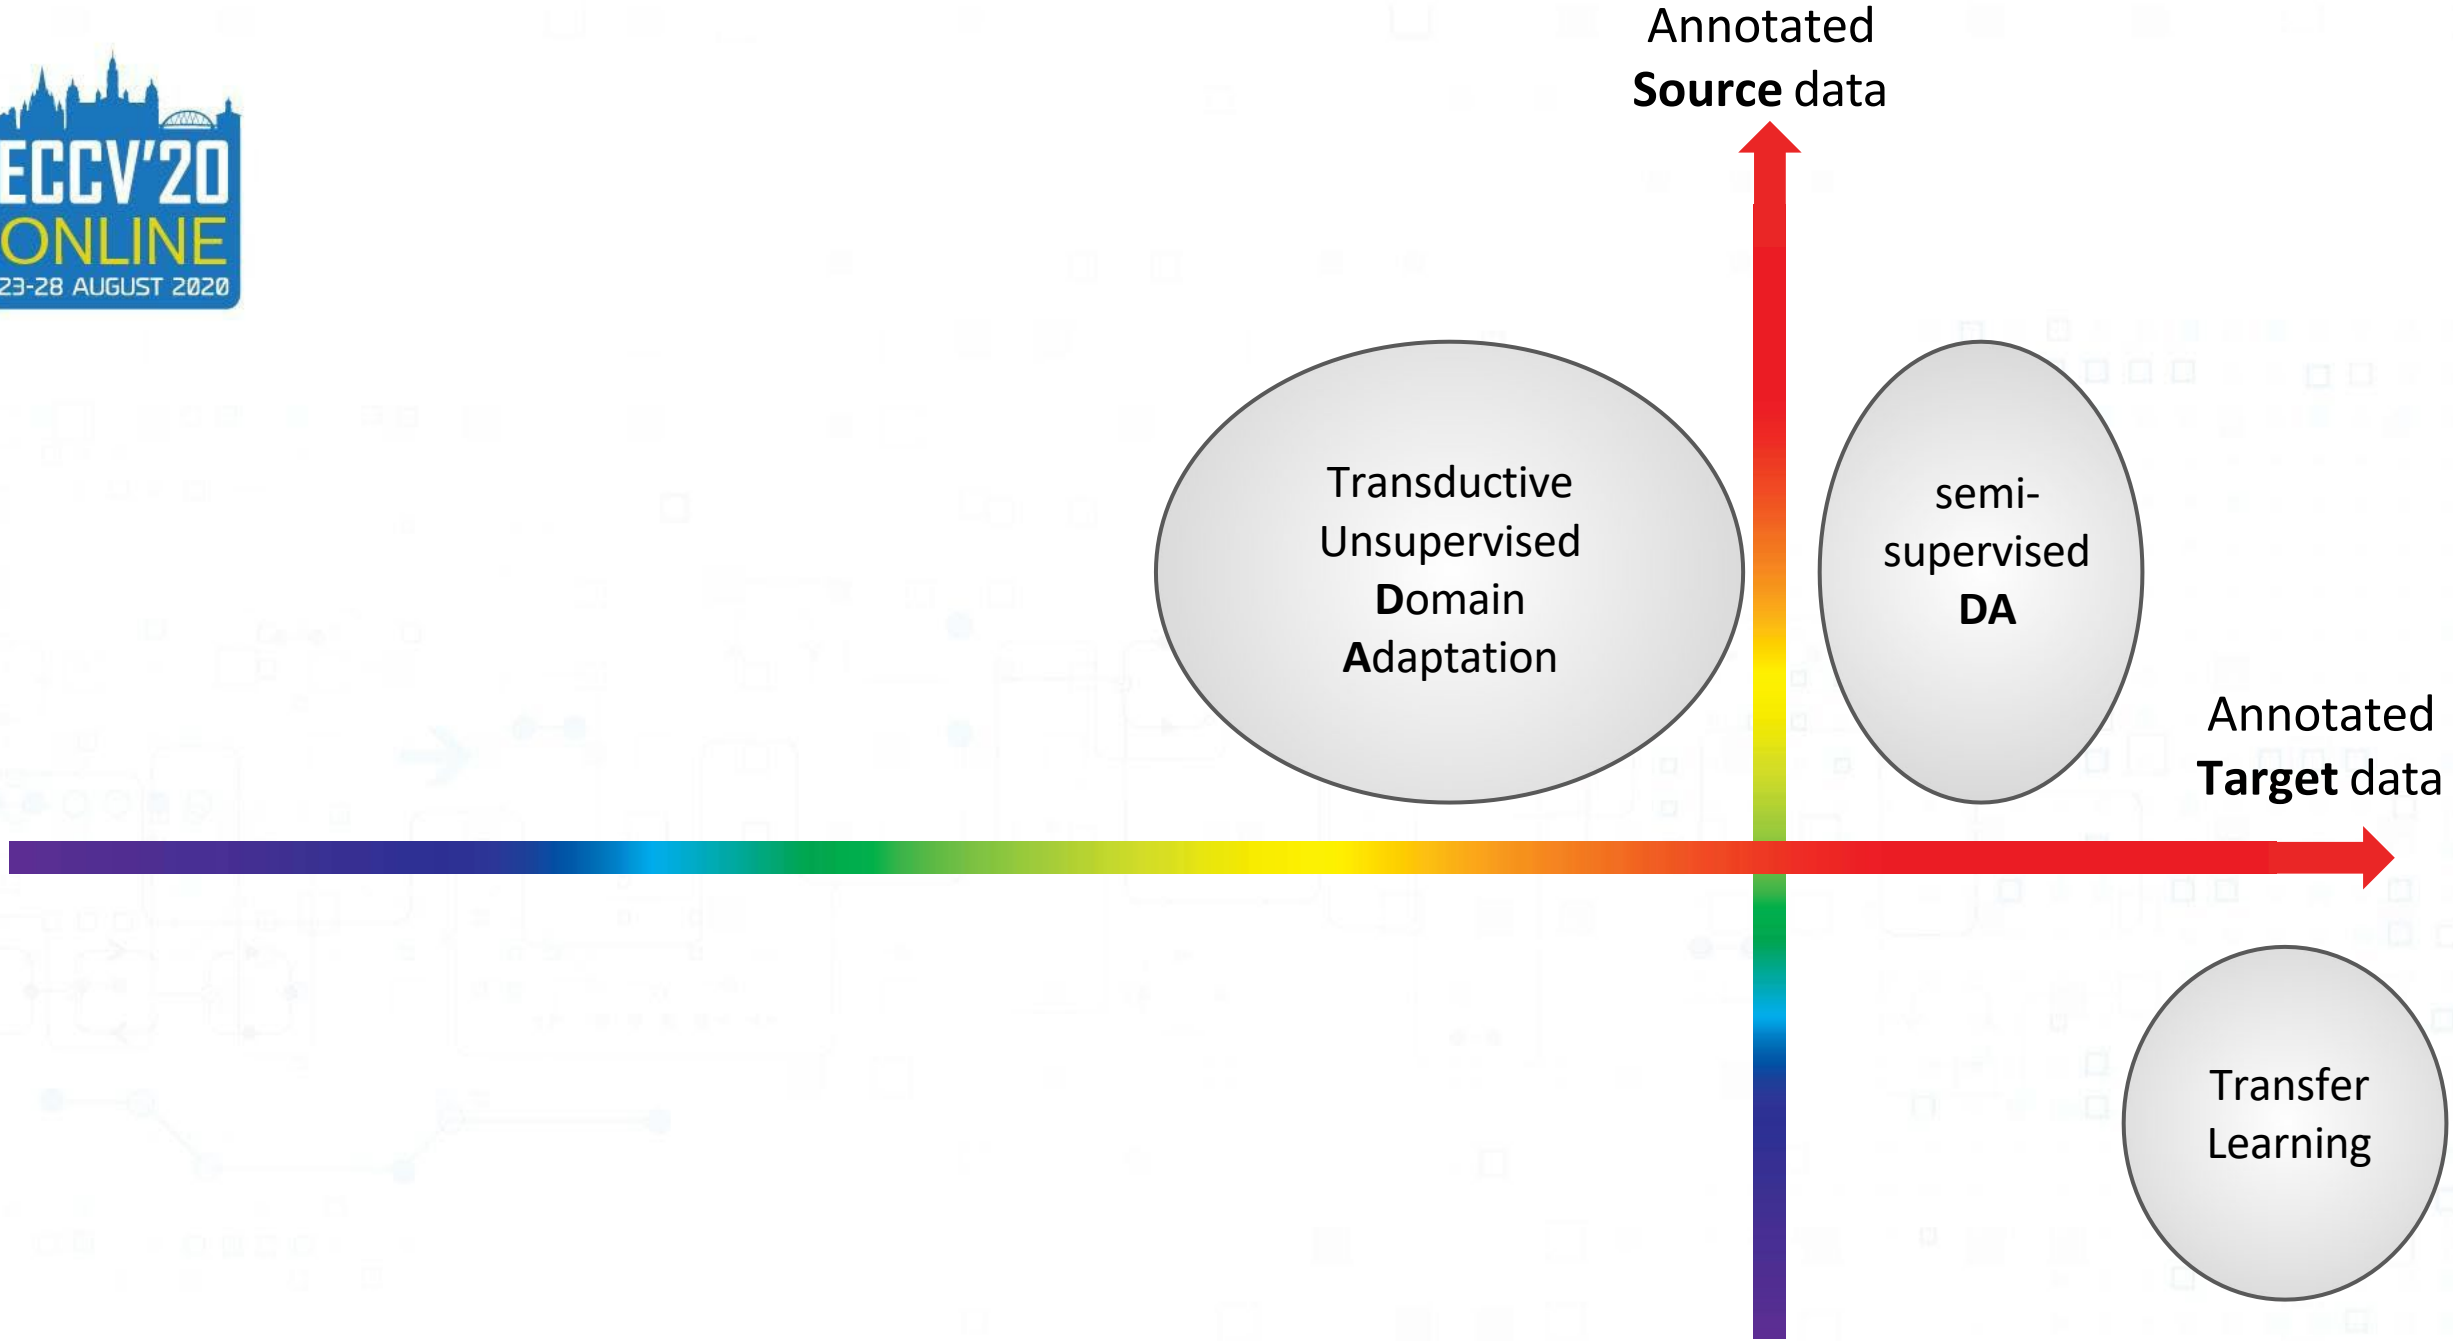
Train

Test











multi-source
**Domain
Generalization**


Transductive
Unsupervised
**Domain
Adaptation**

semi-
supervised
DA

Annotated
Source data

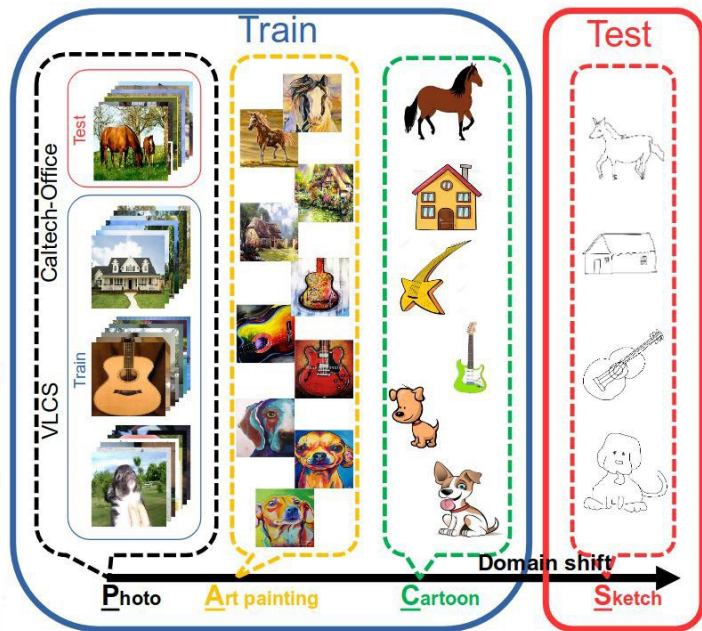
Annotated
Target data

Transfer
Learning





multi-source
Domain
Generalization



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


Transductive
Unsupervised
Domain
Adaptation

semi-
supervised
DA

Transfer
Learning

Annotated
Source data

Annotated
Target 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
Source data

Annotated
Target data


Transfer
Learning



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]



● Sample in Source Domain

○ Augmented Sample


Transductive
Unsupervised
**Domain
Adaptation**

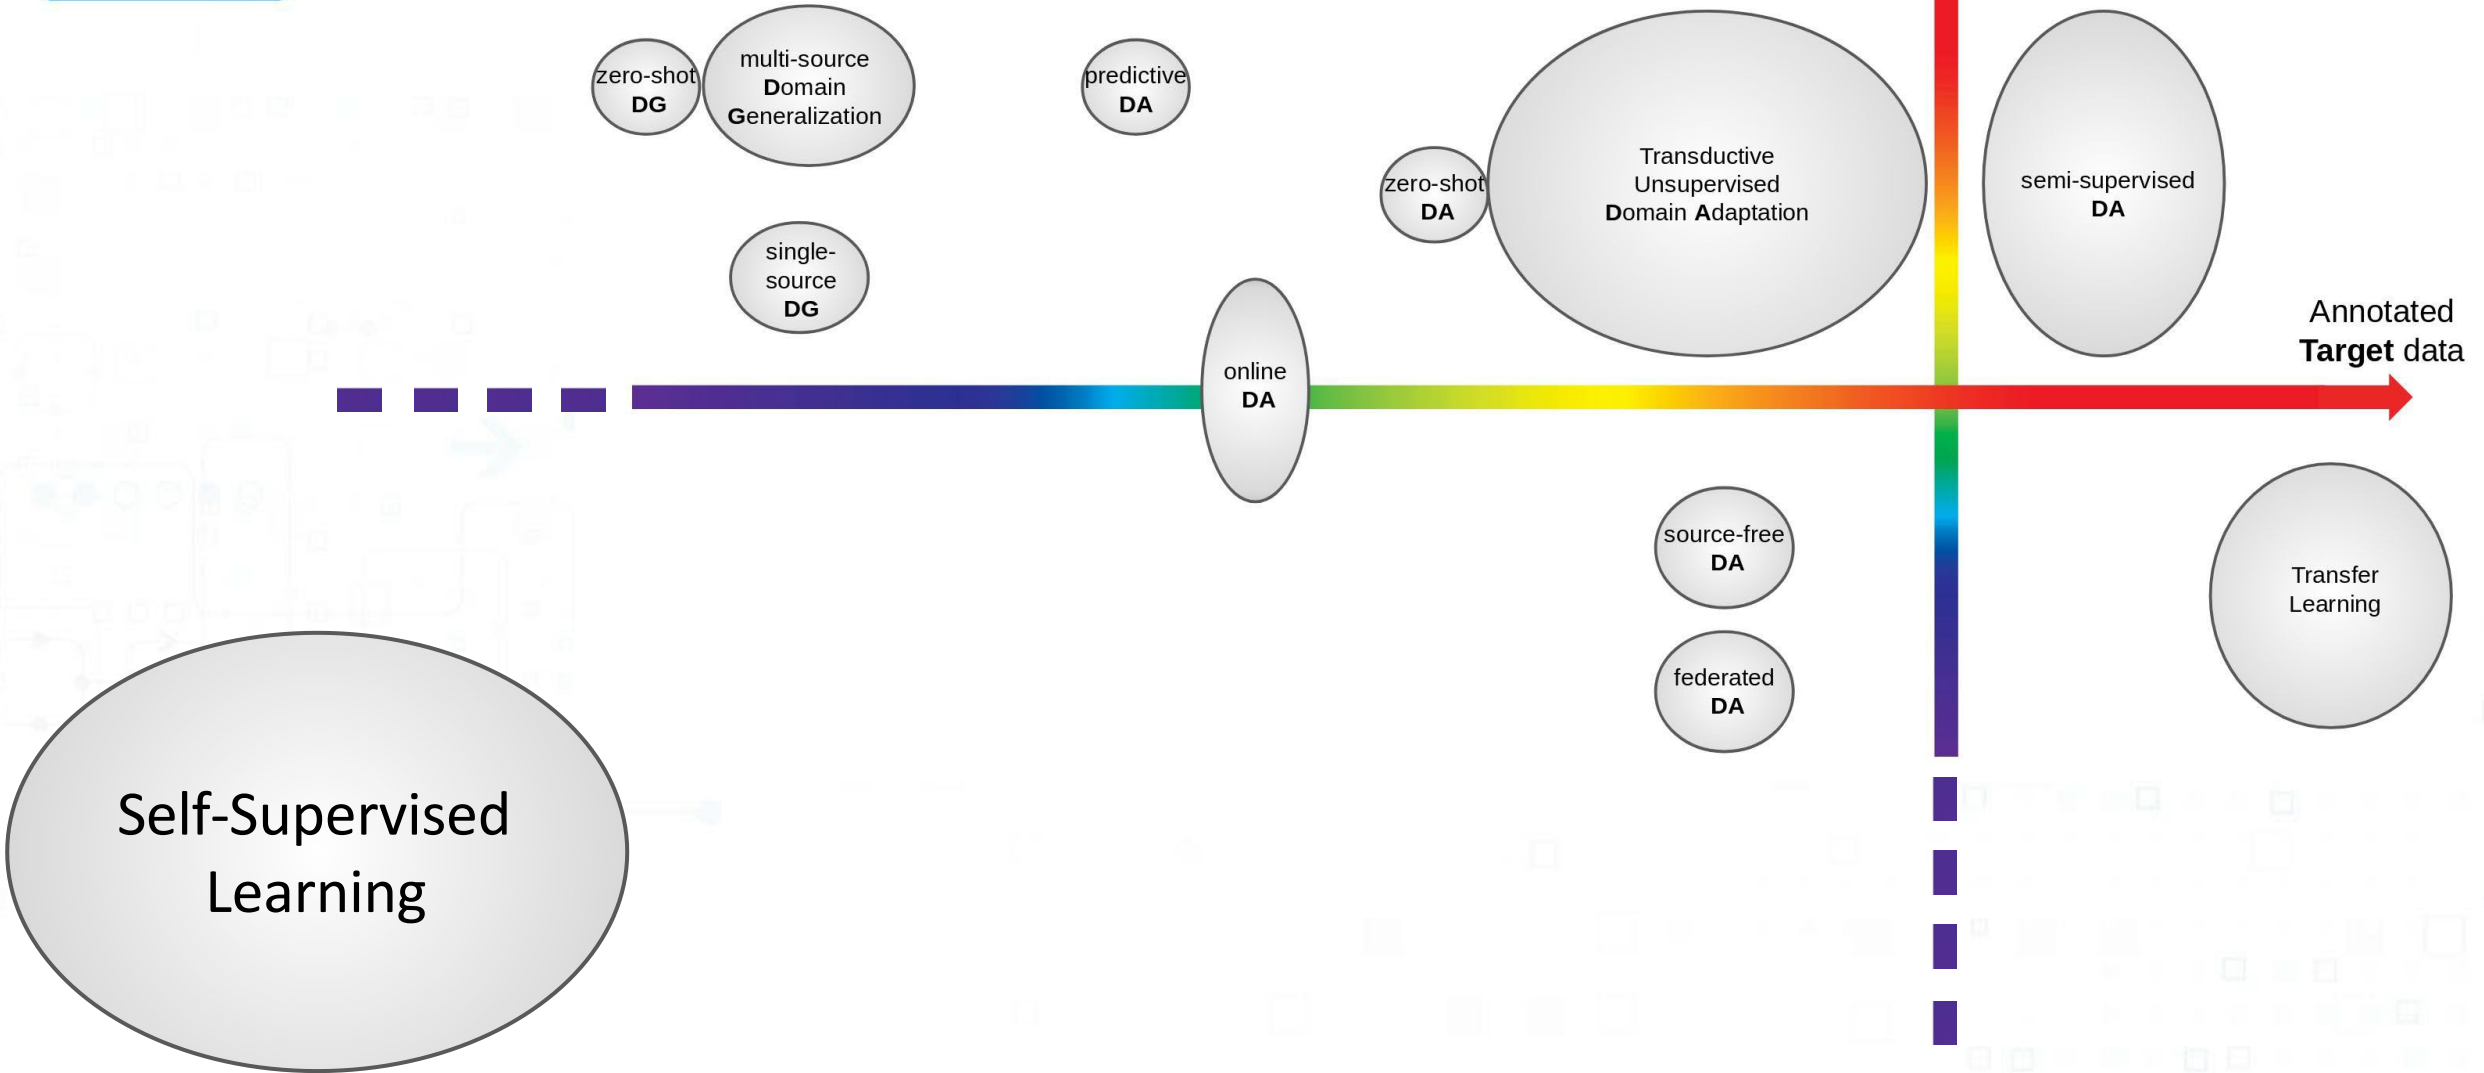
semi-
supervised
DA

Transfer
Learning

Annotated
Source data

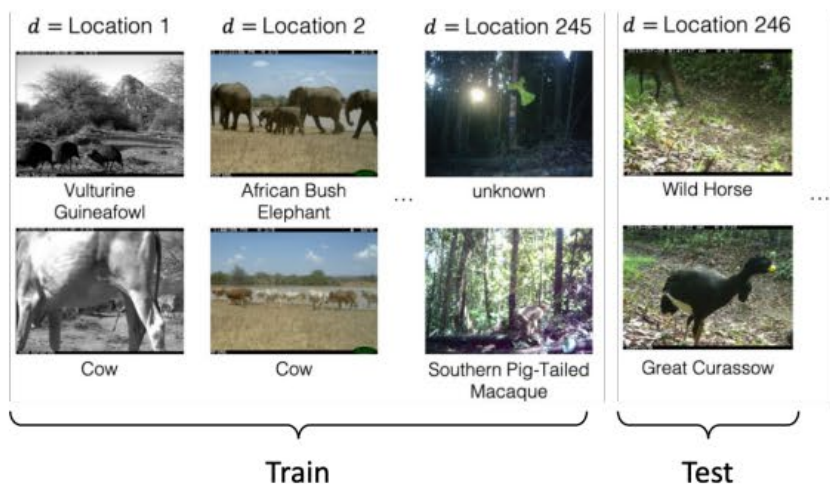
Annotated
Target data






Domain Generalization: Applications


Wildlife recognition



Tissue classification



Molecule property prediction



Code completion

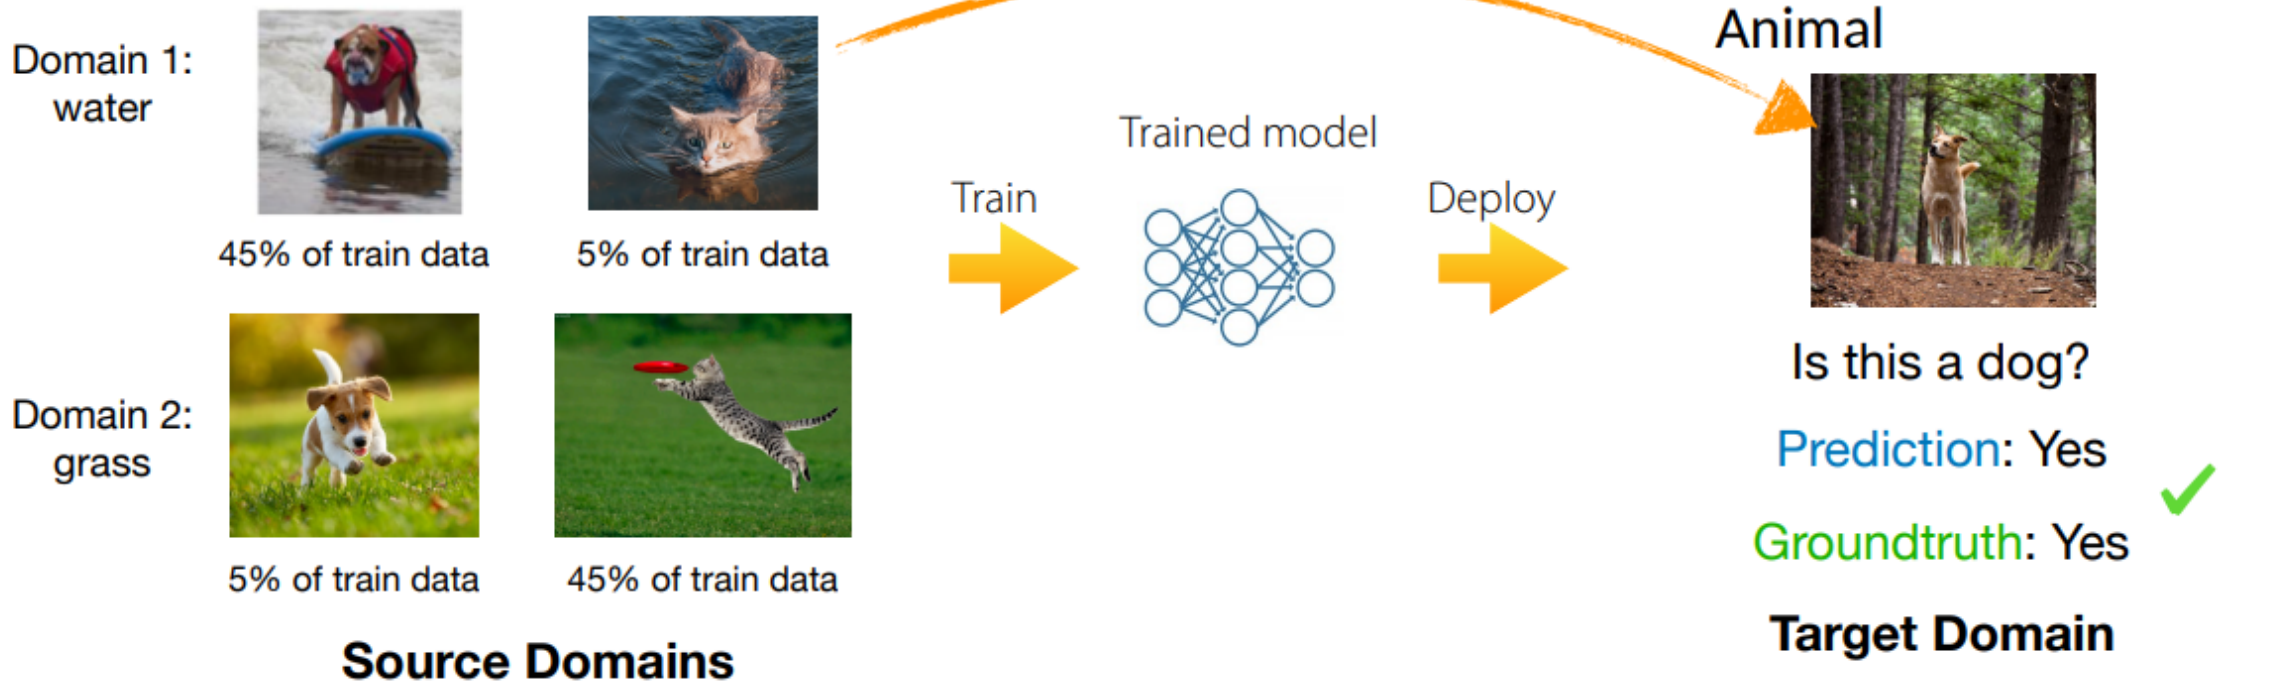
	Repository ID (d)	Source code context (x)	Next tokens (y)
Train	Repository 1	... from easyrec.gateway import EasyRec <EOL> gateway = EasyRec('tenant', 'key') <EOL> item_type = gateway. 	get_item_type
	Repository 2	... response = gateway.get_other_users() <EOL> get_params = HTTPretty. 	last_request
Test	Repository 6,001	... if e.errno == errno.EWOULDBLOCK: <EOL> continue <EOL> p = subprocess.Popen () <EOL> stdout = p. 	communicate
	:	... command = shlex.split(command) <EOL> command = map(str, command) <EOL> env = os. 	environ

How to Learn Generalizable Representations?

To overcome spurious correlation → train a neural network to learn **domain invariance**

Domain invariance: we want to learn **features that don't change across domains**

Goal: classify dog vs. cat



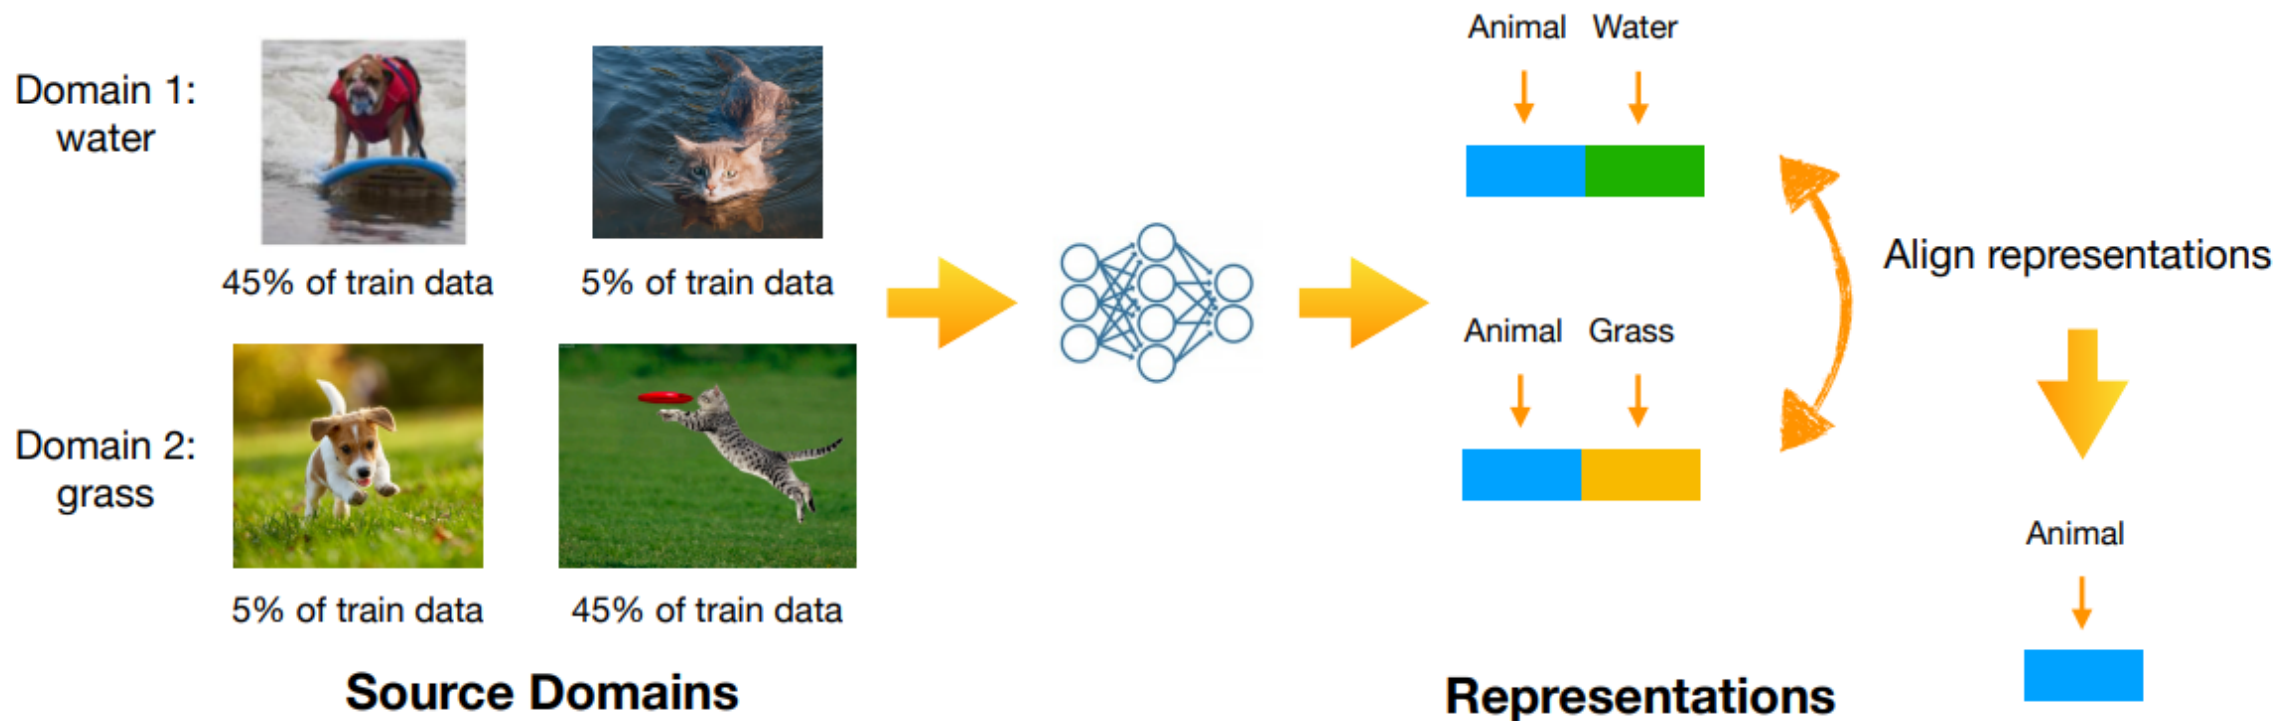
Idea #1

Regularization

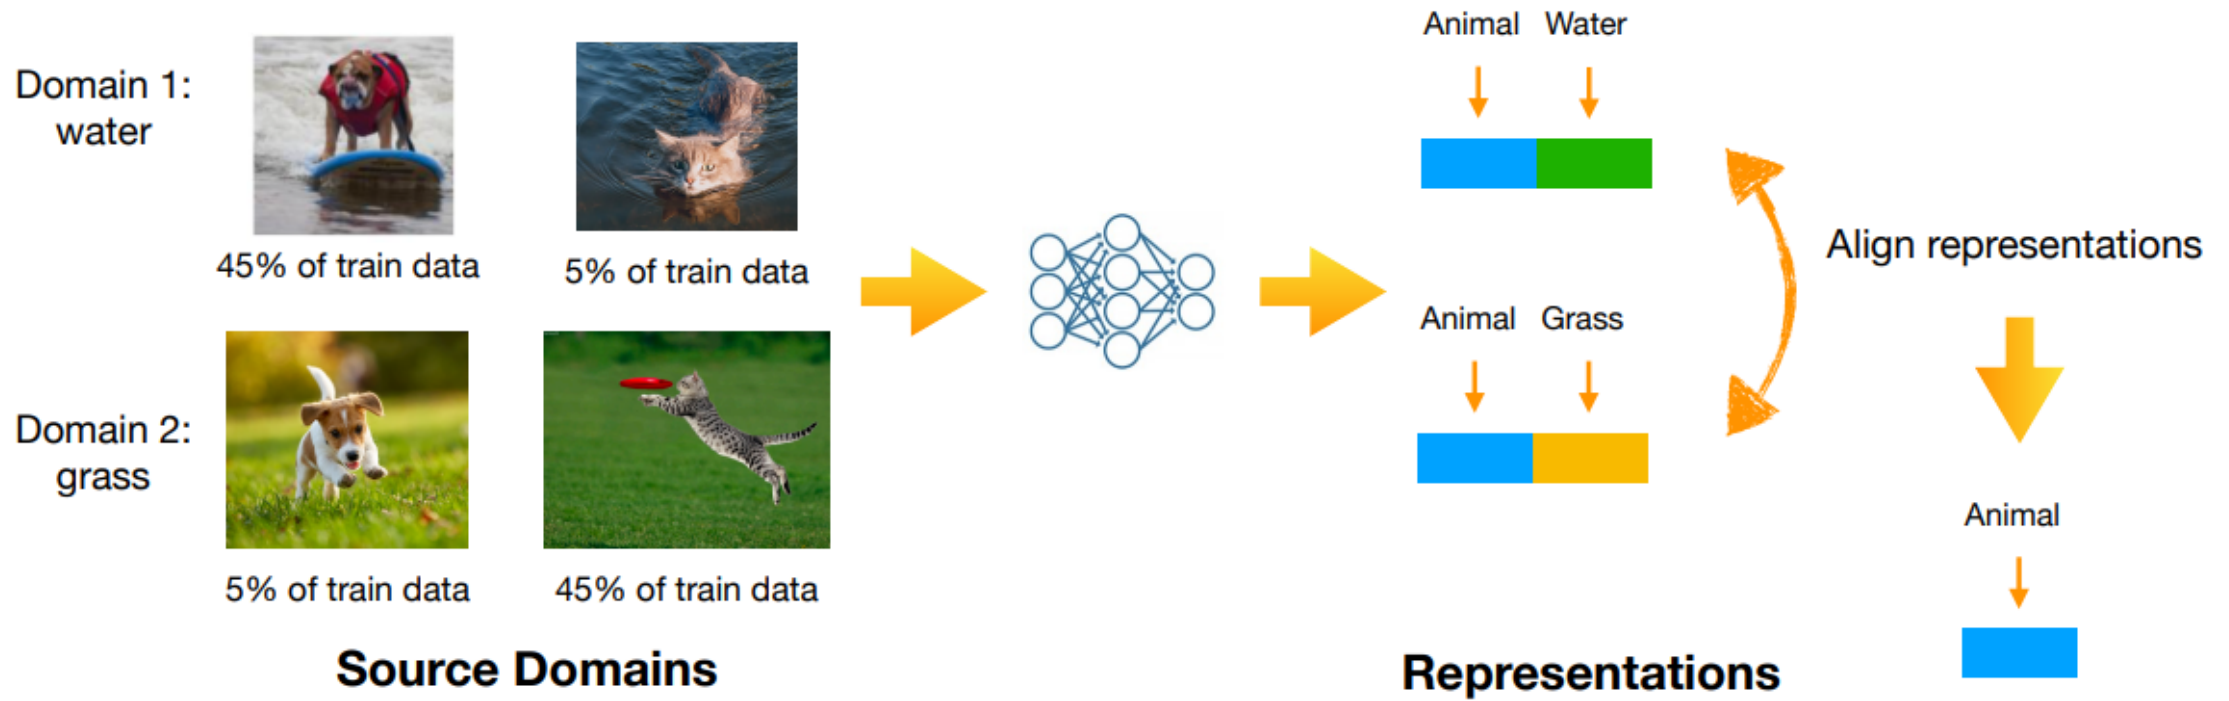
Regularization-based Method

Key idea: Use a regularizer to align representations across domains

—> get domain-invariant representation



Regularization-based Method



Label classification loss

$$\min_{\theta} \mathbb{E}_{(x,y)} [\ell(f_{\theta}(x), y)] + \lambda \mathcal{L}_{reg}$$

Average over training examples

Explicit regularizer to learn domain-invariant representation

Domain Adversarial Training (one of the student presentations)

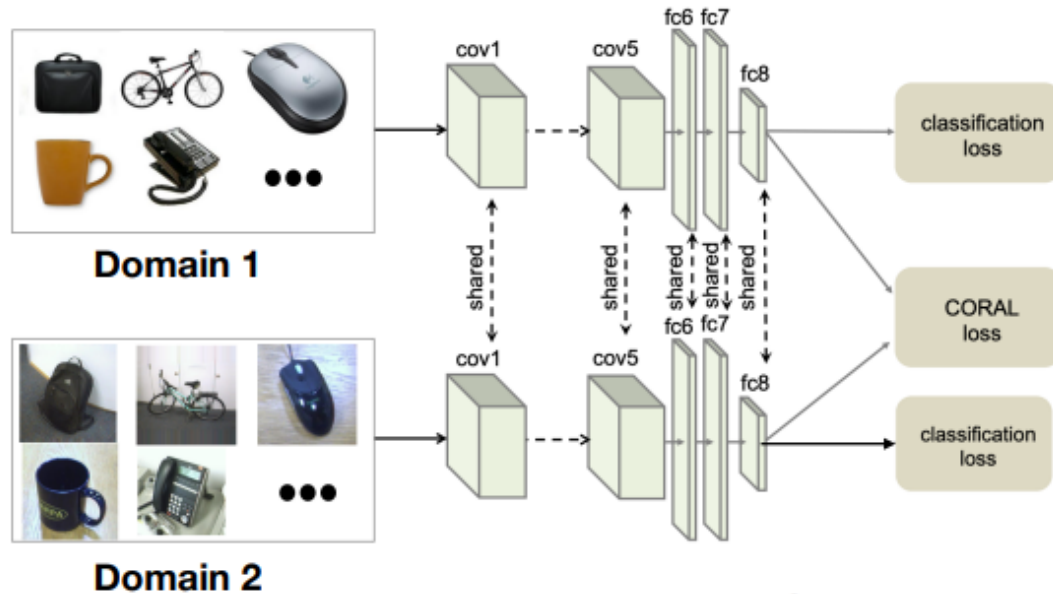
Tzeng et al. Deep Domain Confusion. arXiv '14

Ganin et al. Domain-Adversarial Training of Neural Networks. JMLR '16

Alternative Approach — CORAL

Key idea: directly aligning representations between different domains with some similarity metrics

CORAL: Correlation Alignment for Domain Adaptation (usually also used in DG)



Notations $\mathbf{X}_1 \in \mathbb{R}^{n_1 \times k}$ $\mathbf{X}_2 \in \mathbb{R}^{n_2 \times k}$

k : num of features $\mu_1 = \frac{1}{n_1} \mathbf{1}^T \mathbf{X}_1 \in \mathbb{R}^{1 \times k}$ $\mu_2 = \frac{1}{n_2} \mathbf{1}^T \mathbf{X}_2 \in \mathbb{R}^{1 \times k}$

Calculate covariance matrices $C_1 = \frac{1}{n_1 - 1} \sum_{i=1}^{n_1} (\mathbf{X}_1 - \mu_1)^T (\mathbf{X}_1 - \mu_1)$

$C_2 = \frac{1}{n_2 - 1} \sum_{i=1}^{n_2} (\mathbf{X}_2 - \mu_2)^T (\mathbf{X}_2 - \mu_2)$

CORAL loss $\mathcal{L}_{coral} = \frac{1}{4k^2} \|C_1 - C_2\|_F^2$


Classification loss

$$\mathcal{L} = \sum_{j=1}^{n_1+n_2} \mathcal{L}_c(f_\theta(x_j), y_j) + \lambda \mathcal{L}_{coral}$$

Explicit regularizer to learn domain-invariant representation

Results

	ERM	CORAL	DANN
OfficeHome	66.5%	68.7%	65.9%
DomainNet	40.9%	41.5%	38.3%
iWildCam	30.8%	32.7%	n/a



Idea #2

Data Augmentation

Recap: Spurious Correlation

Recap: spurious correlation between domains and labels

Goal: classify dog vs. cat

Domain 1:
water




45% of train data



5% of train data

Domain 2:
grass



5% of train data



45% of train data

Source Domains

Train



Trained model




Deploy



Spurious information

Grass



Is this a dog?

Prediction: No

Groundtruth: Yes



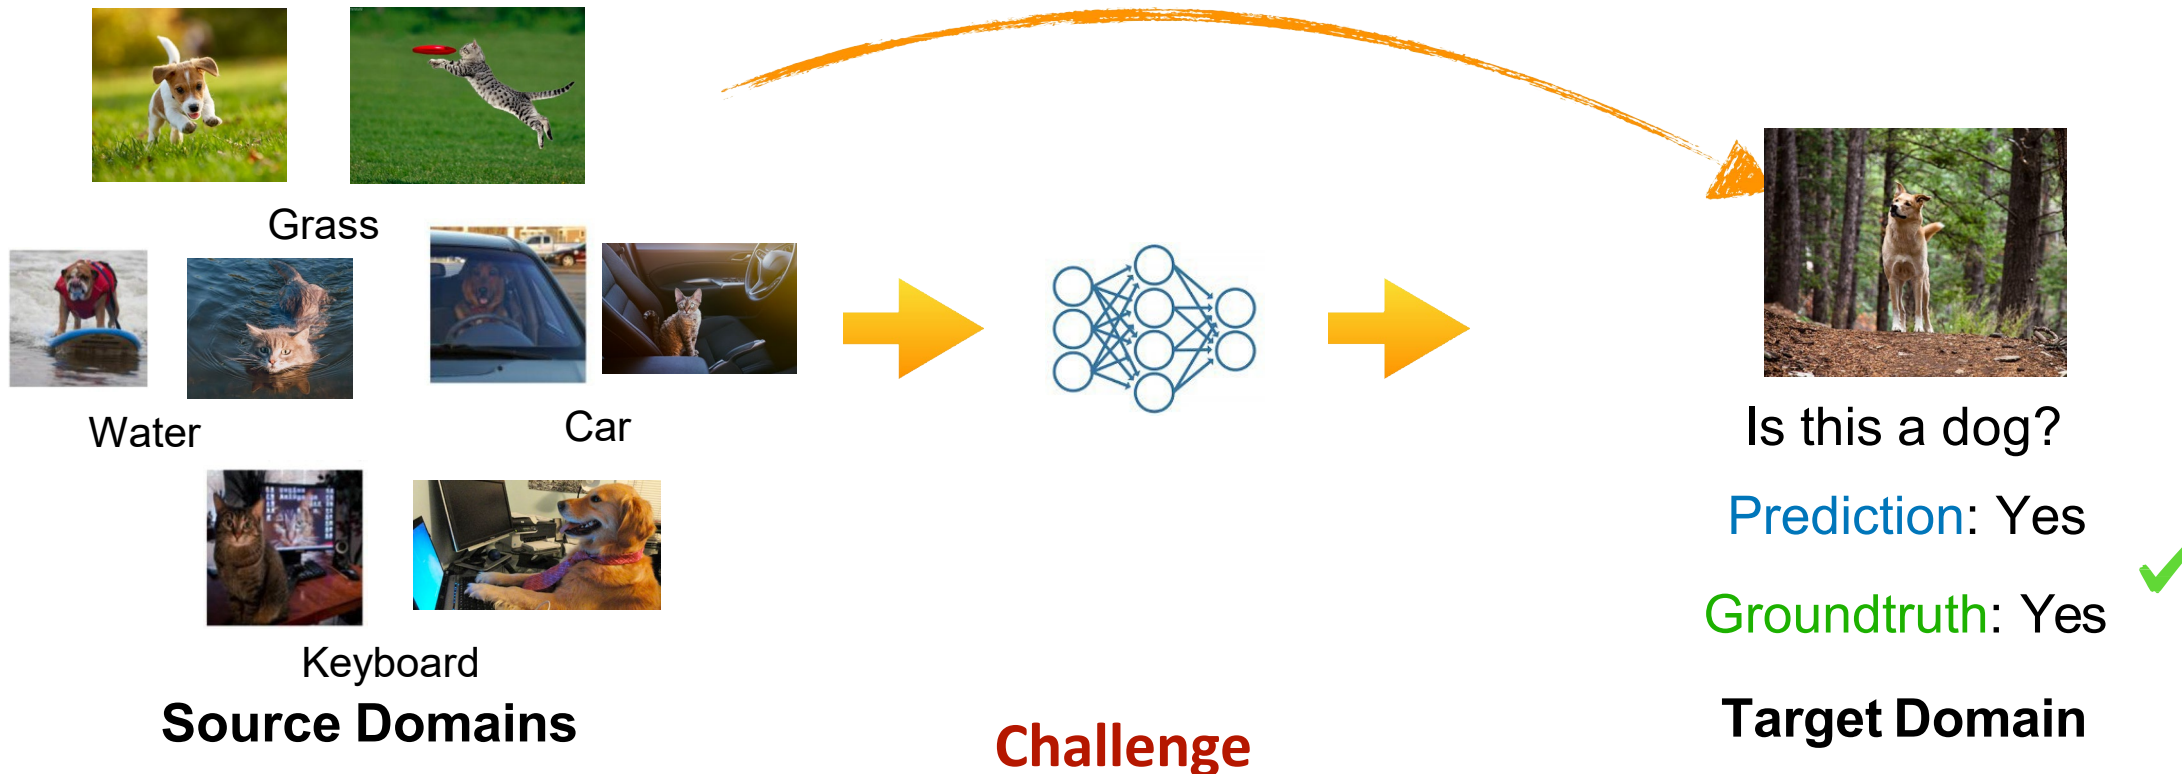
Target Domain

Data Augmentation

If we can collect more data

Question: Will the network still associate dogs with water background in source domains?

NO! There are many more backgrounds. We can't recognize dogs only with grass background.



We can not collect more data → → Let's generate data!

Data Augmentation

Generating data with **simple operators**

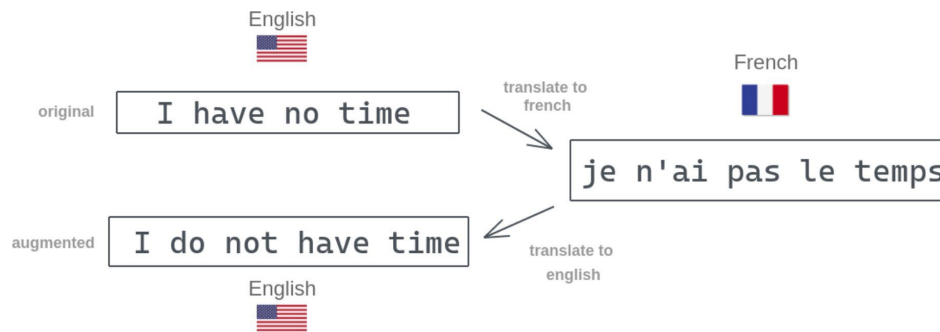



Figure: Back Translation

Requires knowledge of the problem domain

Any general approaches?

<https://amitnesh.com/2020/02/back-translation-in-google-sheets/>

Data Augmentation – Mixup

Interpolating training examples

A learning model

$$\mathcal{D}_{tr} = \{x_i, y_i\}_{i=1}^N \rightarrow \text{Classifier},$$

Mixup

$$\tilde{\mathcal{D}}_{tr} = \{\tilde{x}_i, \tilde{y}_i\}_{i=1}^N \rightarrow \text{Classifier},$$

where

$$\tilde{x}_i = \lambda x_i + (1 - \lambda)x_j, \tilde{y}_i = \lambda y_i + (1 - \lambda)y_j$$

$$\lambda \sim \text{Beta}(\alpha, \beta)$$


Generating some virtual examples between two classes



[1.0, 0.0]
cat



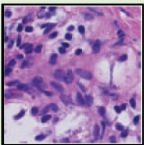

[0.0, 1.0]
dog



[0.7, 0.3]
cat dog


Data Augmentation – Mixup

Mixup can improve the performance on domain generalization

	Empirical Risk Minimization	mixup
 Camelyon17	70.3%	71.2%
 FMoW	32.8%	34.2%

2
9

But it is not always good!



RxRx1

29.9%

26.5%

Original mixup only focuses on data augmentation instead of learning domain invariance.

How to Improve it?

Regularization-based v.s. Augmentation-based Methods

Regularization-based Method

- + General to all kinds of data and networks
- + Some theoretical guarantee
- Rely on the design of regularizers

Augmentation-based Method

- + Easy to understand and simple to implement
- + No need to worry about how to design regularizers
- Largely limited to classification

Discovering Adversarial Data Augmentation

(Series of Work by Tejas)

Problem Setting: Single Source Domain Generalization




*How can classifiers trained on one domain generalize to other **unseen domains**?*

- Given:



Labeled training data from "Source" Domain

- **Not Given:**




Target exemplars; knowledge of unseen domains

SSDG: Single Source Domain Generalization

*How can classifiers trained on one domain generalize to other **unseen domains**?*

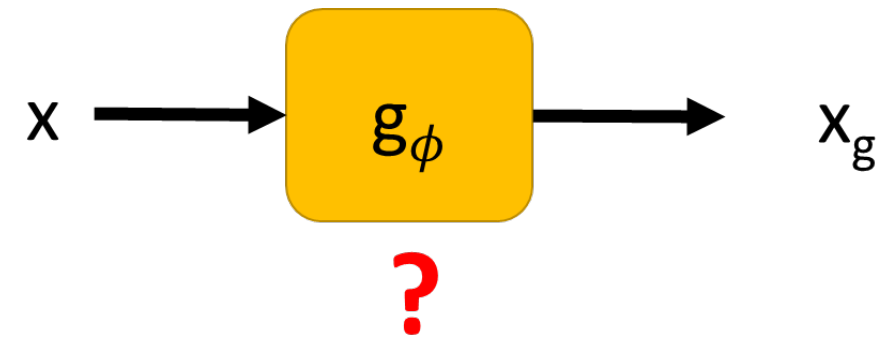
- For SSDG, Data Augmentation is crucial !!!
 - *To increase diversity of training data*
 - *To simulate new domains*
 - *To cover distributions that may be encountered at test-time*



Data Augmentation is Crucial ... But which augmentation?

How can classifiers trained on one domain generalize to other **unseen domains**?

- For SSDG, Data Augmentation is crucial !!!
 - *To increase diversity of training data*
 - *To simulate new domains*
 - *To cover distributions that may be encountered at test-time*



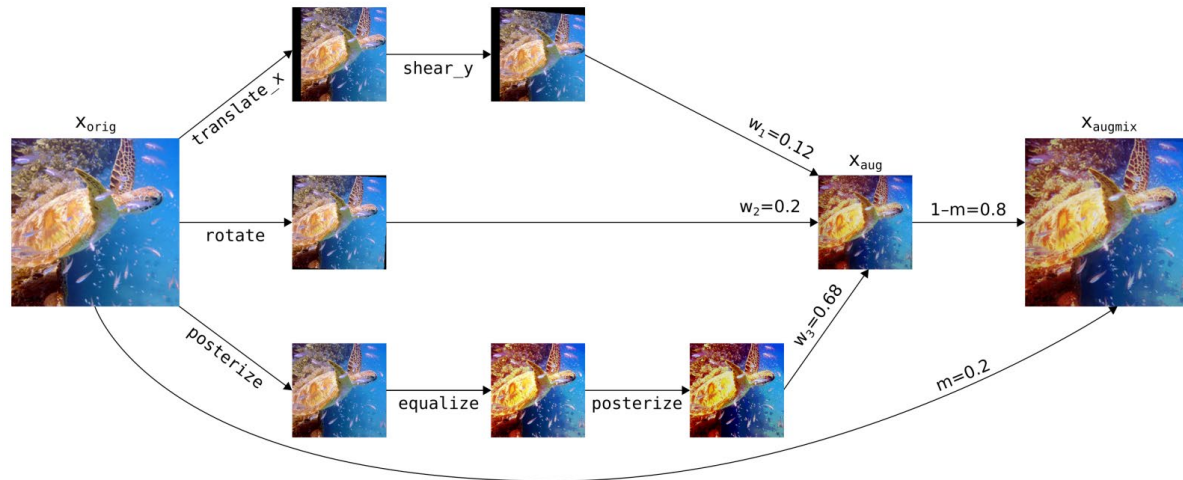
How do we know which data augmentations will be useful?
(We don't have access to the test domains)

Data Augmentation is Crucial ... But which augmentation?

- Existing Data Augmentation techniques
 - Introduce a strong preference towards certain types of diversity


AugMix (Hendrycks et al. ICLR 2019)

Combination of geometric transforms and image filters




RandConv (Xu et al. ICLR 2021)

Convolve image with random filter



Data Augmentation is Crucial ... But which augmentation?

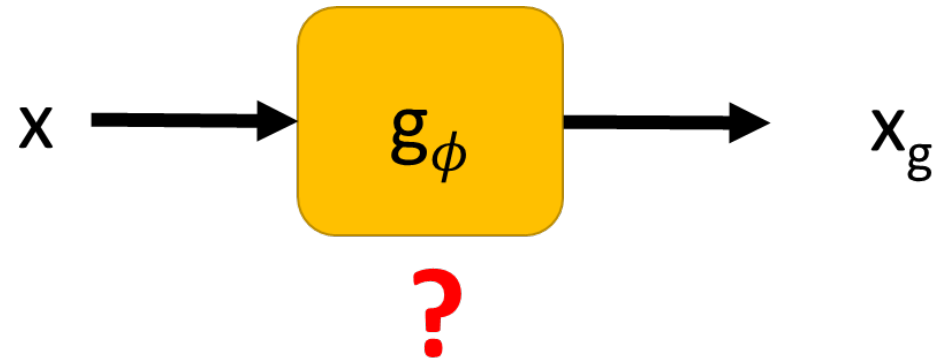
- Existing Data Augmentation techniques
 - Introduce a strong preference towards certain types of diversity
 - Mixed results on different domain shift datasets



Rethink Data-Augmentation

Go Beyond STATIC, PRE-DEFINED Augmentations

Our Solution: *Discover Data Transformations During Training*



How?

Classifier's failures are informative – leverage them for guiding data augmentation
Tune parameters of a image-to-image network $g()$ to transform images

Key Finding

Data transformations discovered during training are more effective

Instead of using pre-defined static augmentations

"ALT"

Generalizing to Domain Shift *via* Adversarially Learned Transformations

Tejas Gokhale, Rushil Anirudh, Jay Thiagarajan, Bhavya Kailkhura, Chitta Baral, Yezhou Yang

ASU Arizona State
University

Lawrence Livermore
National Laboratory

WACV 2023



ALT: Adversarially Learned Transformations




Image Transformation Network

$g()$

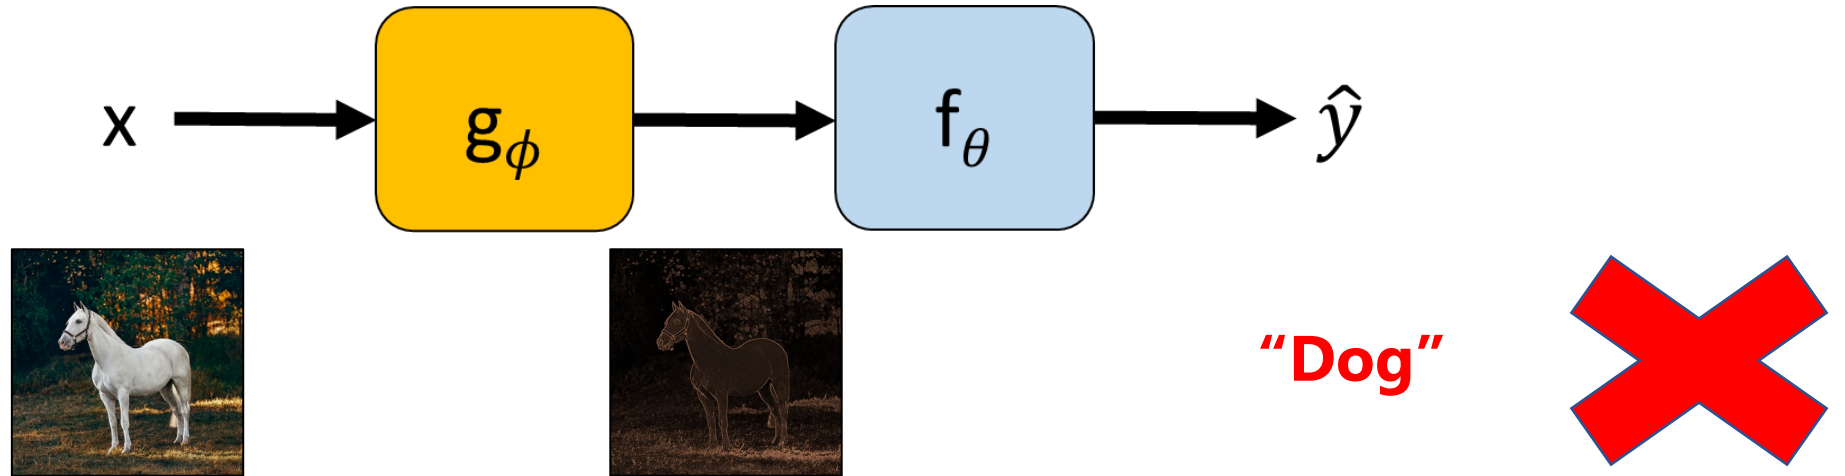
With Parameters ϕ

Classifier

$f()$

With Parameters θ

ALT: Adversarially Learned Transformations




For each batch learn perturbations of ϕ to maximize classifier loss

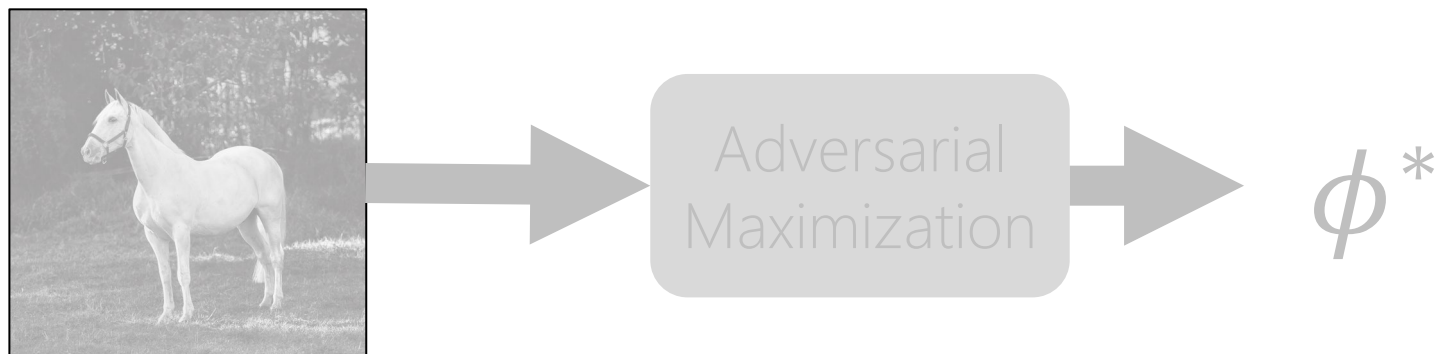
$$\max_{\phi} \mathcal{L}_{BCE}(f(g(x; \phi); \theta), y)$$

$$\phi \leftarrow \phi + \nabla(L_{cls}(f(g(x; \phi)), y) - L_{TV}(x_g))$$


Pre-Training
Phase




Learn Transformations



Generate
Augmentations



Pre-Training
Phase



θ




Learn Transformations




ϕ^*

$$\max_{\phi} \mathcal{L}_{BCE}(f(g(x; \phi); \theta), y)$$

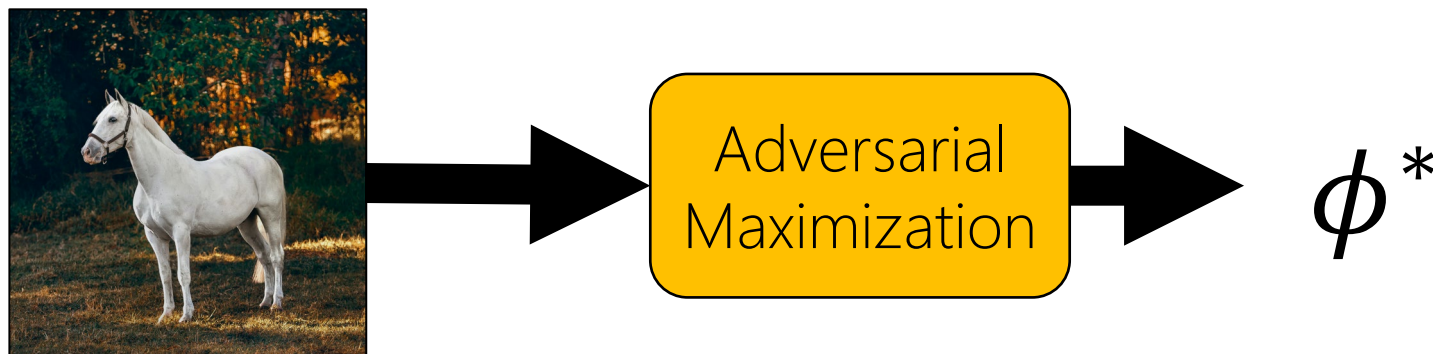
Generate
Augmentations




Pre-Training
Phase



Learn Transformations



Generate
Augmentations




These transformed images are used for training

Enforcing Consistency on Classifier's Predictions

$$p = f(\text{img}_{\text{white horse}})$$

$$p_g = f(\text{img}_{\text{dark horse}})$$




$$p_{mix} = \frac{p + p_g}{2}$$

$$L_{consistency} = D_{KL}(p_{mix}|p) + D_{KL}(p_{mix}|p_g)$$

Improving Diversity with ALT


ALT

via adversarial perturbations of ϕ



STATIC AUGMENTATION

via static data augmentations
e.g. AugMix, RandConv



Use ALT in conjunction with static data augmentations from previous work


This further boosts performance compared to using $g()$ only

Enforcing Consistency on Classifier's Predictions

$$p = f(\text{img}_c)$$

$$p_g = f(\text{img}_g)$$

$$p_r = f(\text{img}_r)$$



$$p_{mix} = \frac{p_c + p_g + p_r}{3}$$

$$L_{consistency} = D_{KL}(p_{mix}|p) + D_{KL}(p_{mix}|p_g) + D_{KL}(p_{mix}|p_r)$$

Results: (Style Shift)

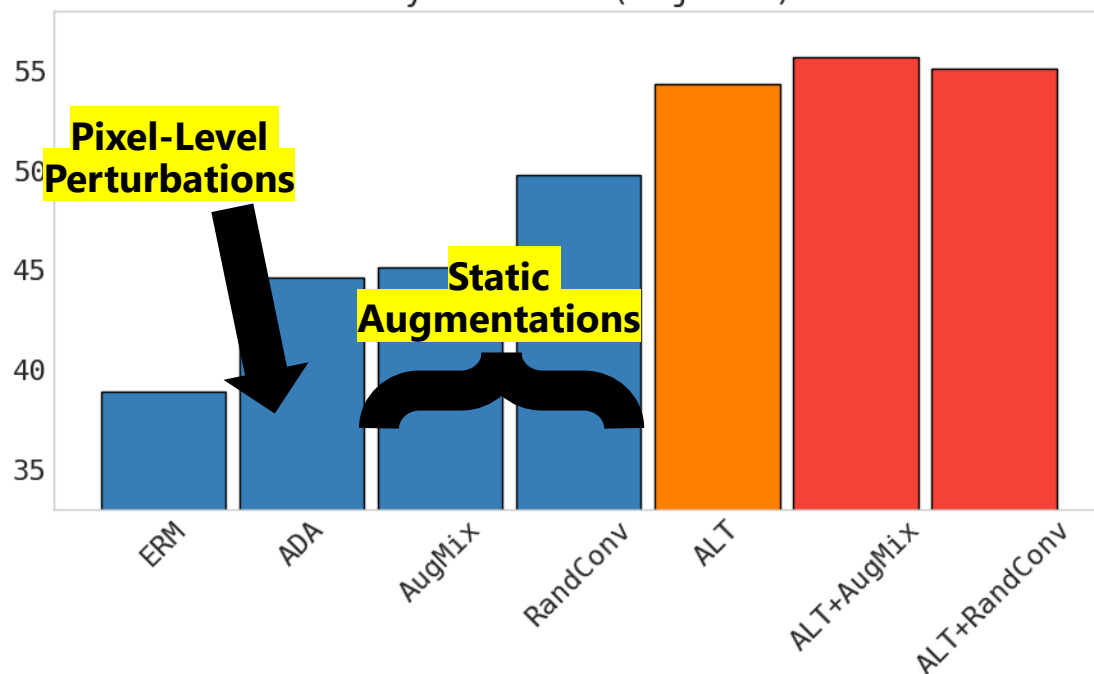
Object Classification



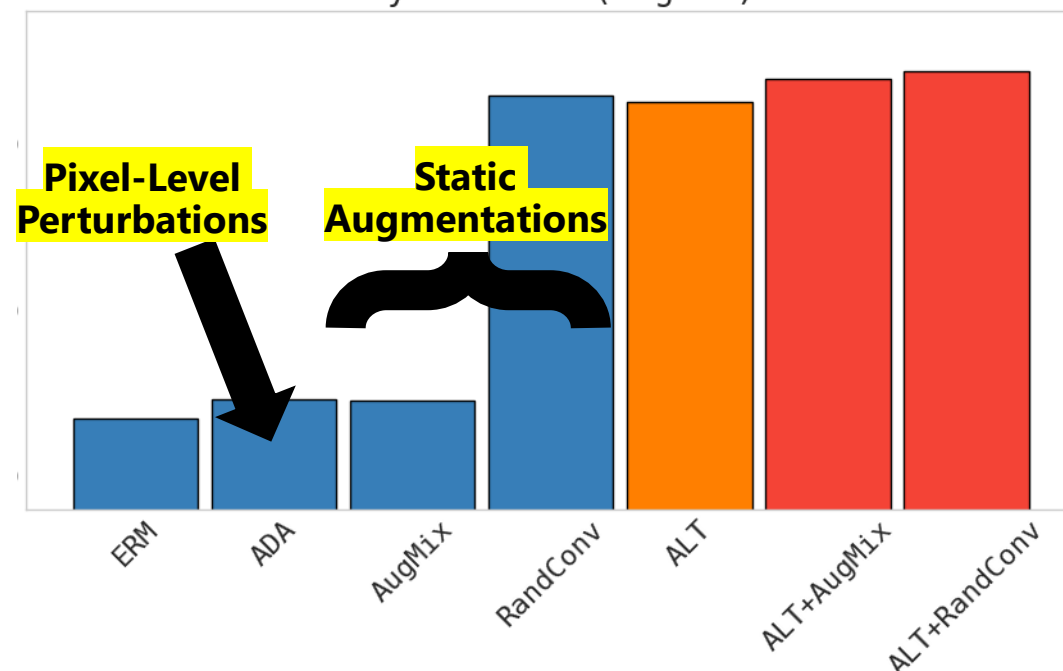
Digit Classification




Style Shift (Objects)



Style Shift (Digits)



(Subpopulation Shift) Animal Classification



- Trained on one set of sub-species
- Tested on a different set of sub-species

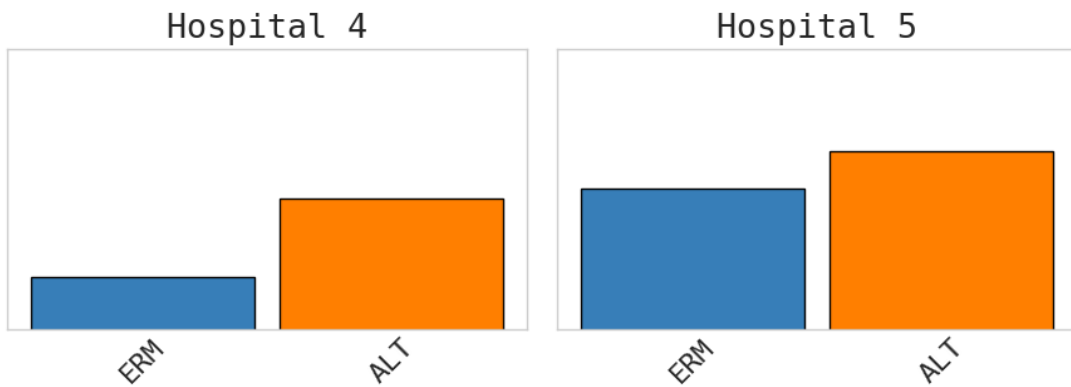
(apes: gibbon/orangutan)
(apes: gorilla/chimpanzee)

ALT improves robustness to Subpopulation Shift

Results: Application to Societal Challenges


(Hospital Shift) Tumor Classification

Train			Val (OOD)	Test (OOD)
	d = Hospital 1	d = Hospital 2	d = Hospital 3	d = Hospital 4
y = Normal				
y = Tumor				



(Terrain Shift) Land-Use Classification

Train 2002--2013			Test		
Satellite Image (x)					
Year / Region (d)	2002 / Americas	2009 / Africa	2012 / Europe	2016 / Americas	2017 / Africa
Building / Land Type (y)	shopping mall	multi-unit residential	road bridge	recreational facility	educational institution



What if knowledge about unseen domains is available?

How can we leverage that knowledge

to discover image transformations?

"AGAT"

Generalizing to Domain Shift *via* Attribute Guided Adversarial Training

Tejas Gokhale, Rushil Anirudh, Bhavya Kailkhura, Jay Thiagarajan, Chitta Baral, Yezhou Yang

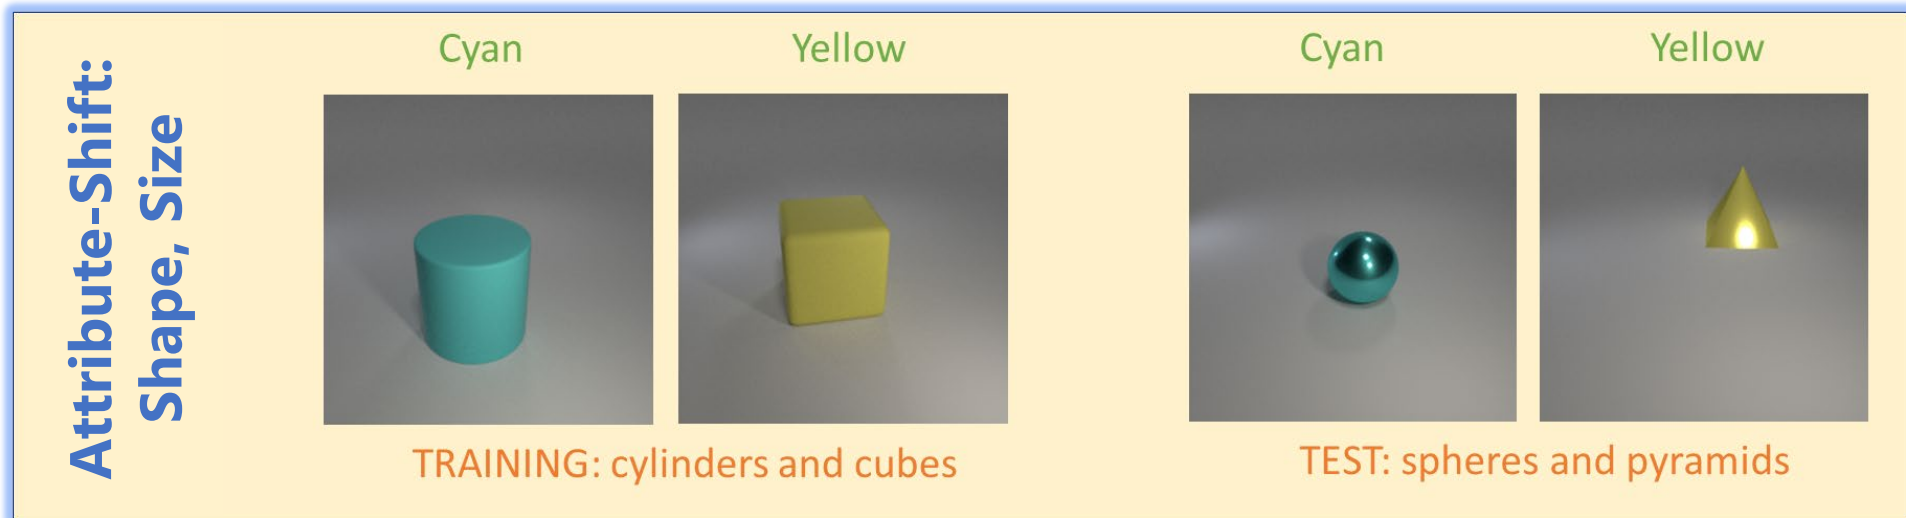


AAAI 2021



Using Attribute Knowledge for Domain Generalization

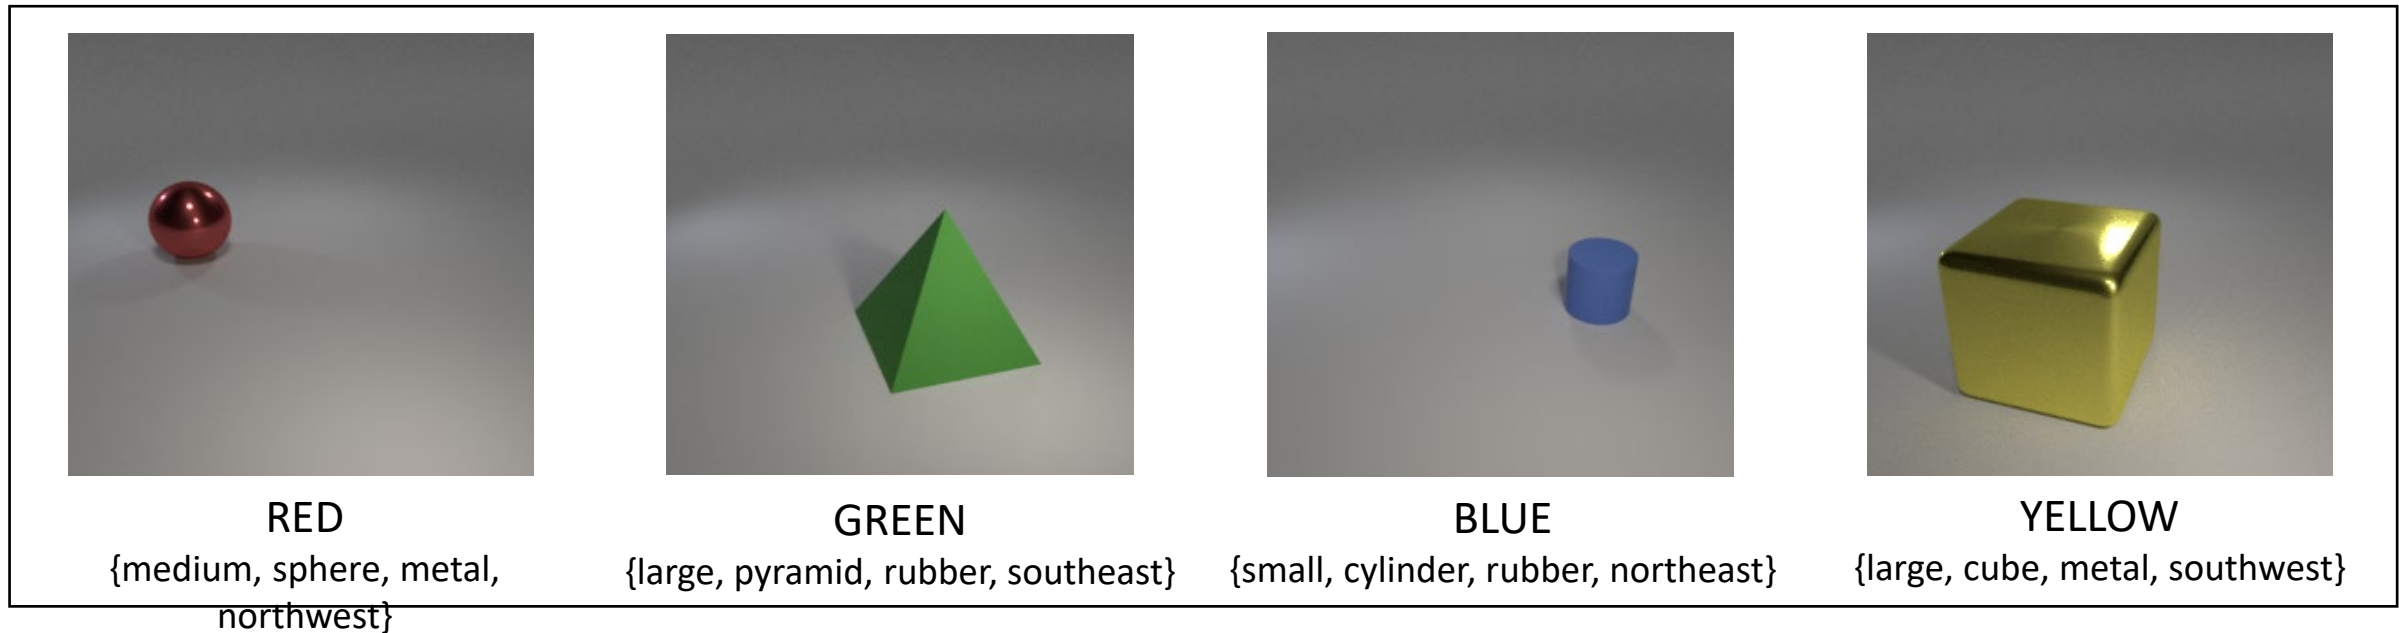
In real-world scenarios, test examples can vary along attributes
Size, Shape, Materials, Geometric Parameters, Lighting, ...



In ALT, we assumed no access to such attributes

How can we leverage attributes to learn useful image transformations?

CLEVR-Singles: A Dataset for Studying Attribute-Level Domain Shift



- Photorealistic rendering of single objects. Controlled setting for studying attribute-level domain shift
- (Classification) task attribute: **Color**; Task-invariant Attributes: **Size, Shape, Material, Position**


Size	Small, Medium, Large	3
Shape	Sphere, Cylinder, Cube, Pyramid	4
Material	Rubber, Metal	2
Position	NW, SW, NE, SE	4
Color	Red, Blue, Green, Yellow, Cyan, Purple, Grey, Brown	8

CLEVR-Singles: A Dataset for Studying Attribute-Level Domain Shift


Create train—test dataset splits s.t. attribute combinations at test time are not seen during training
e.g. unseen **Material + Position** combinations

TRAINING


Metal objects far from the camera
Rubber objects close to the camera




Metal, far



Metal, far



Rubber, close




Rubber, close

TESTING

Metal objects close to the camera
Rubber objects far from the camera




Metal, close



Metal, close







Rubber, far



Rubber, far

Problem Setting

TRAINING DATASET



Attribute Set

e.g. ["Size", "Shape", "Material", "Position"]

- **Unknown:**

- which attributes will change at test time
- by what magnitude
- in what combination

- **No Access To:**

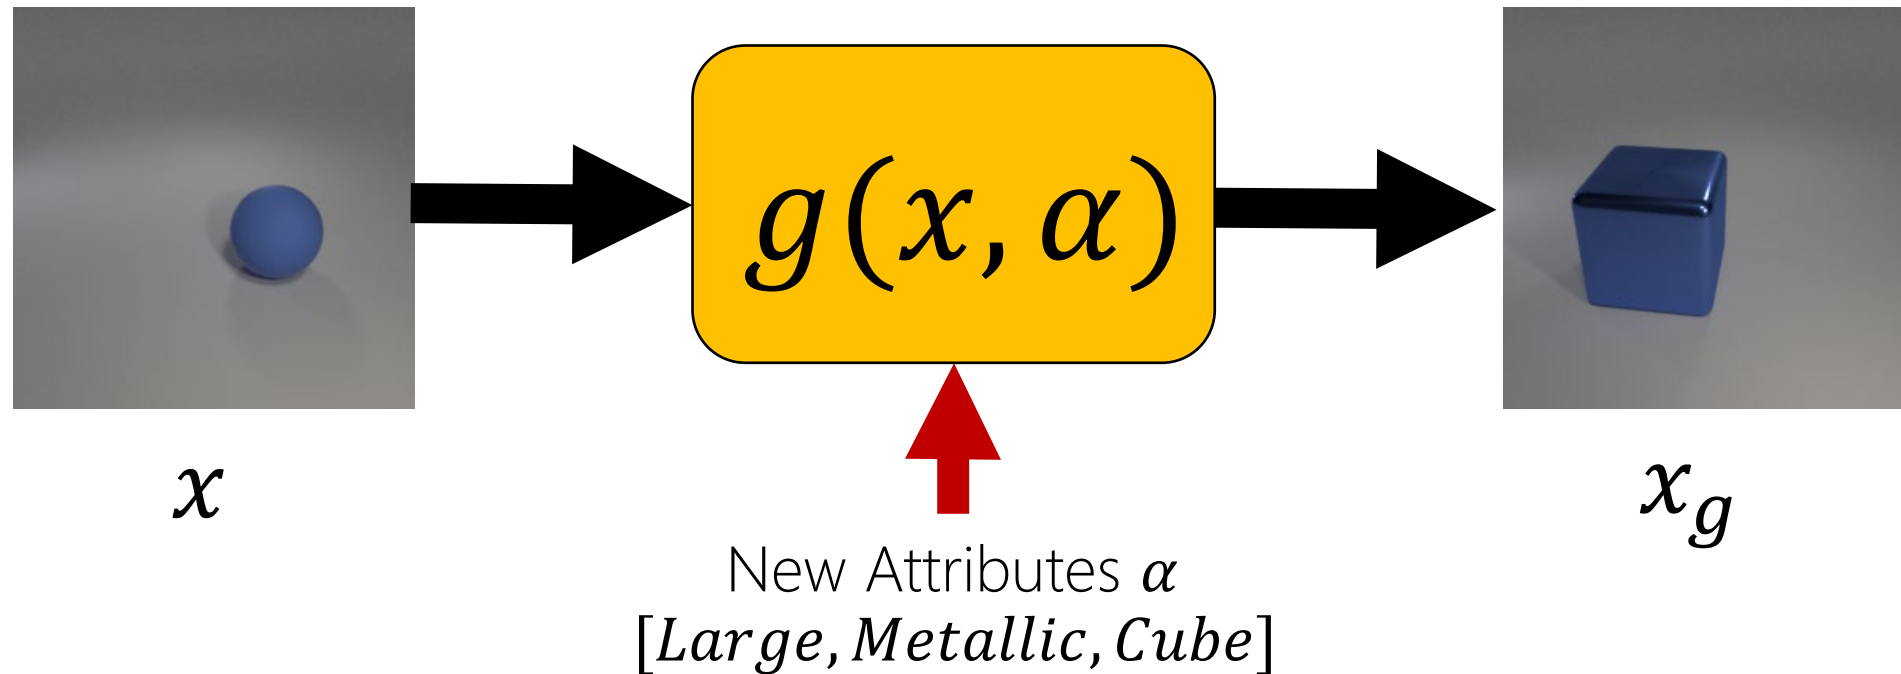
- Validation set
- exemplars representing attribute shift

Goal

Train a classifier that can generalize to attribute-level domain shift

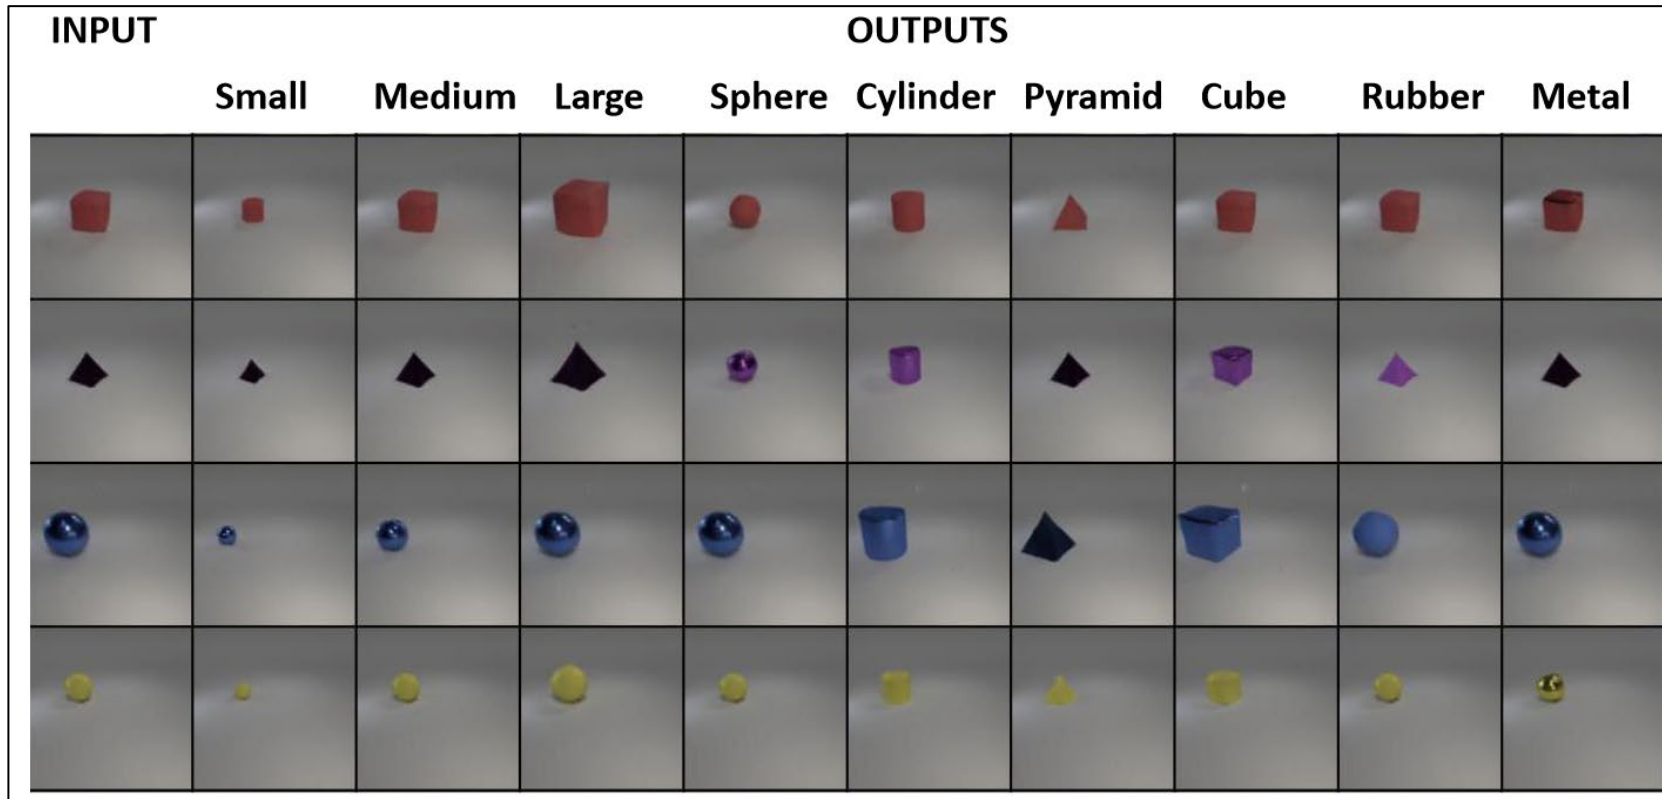
Attribute-Guided Adversarial Training

- Parameterize input space by attributes α
 - Train a Generative Model conditioned on the attributes
- Maximize exploration of input space by learning attribute-level transformations

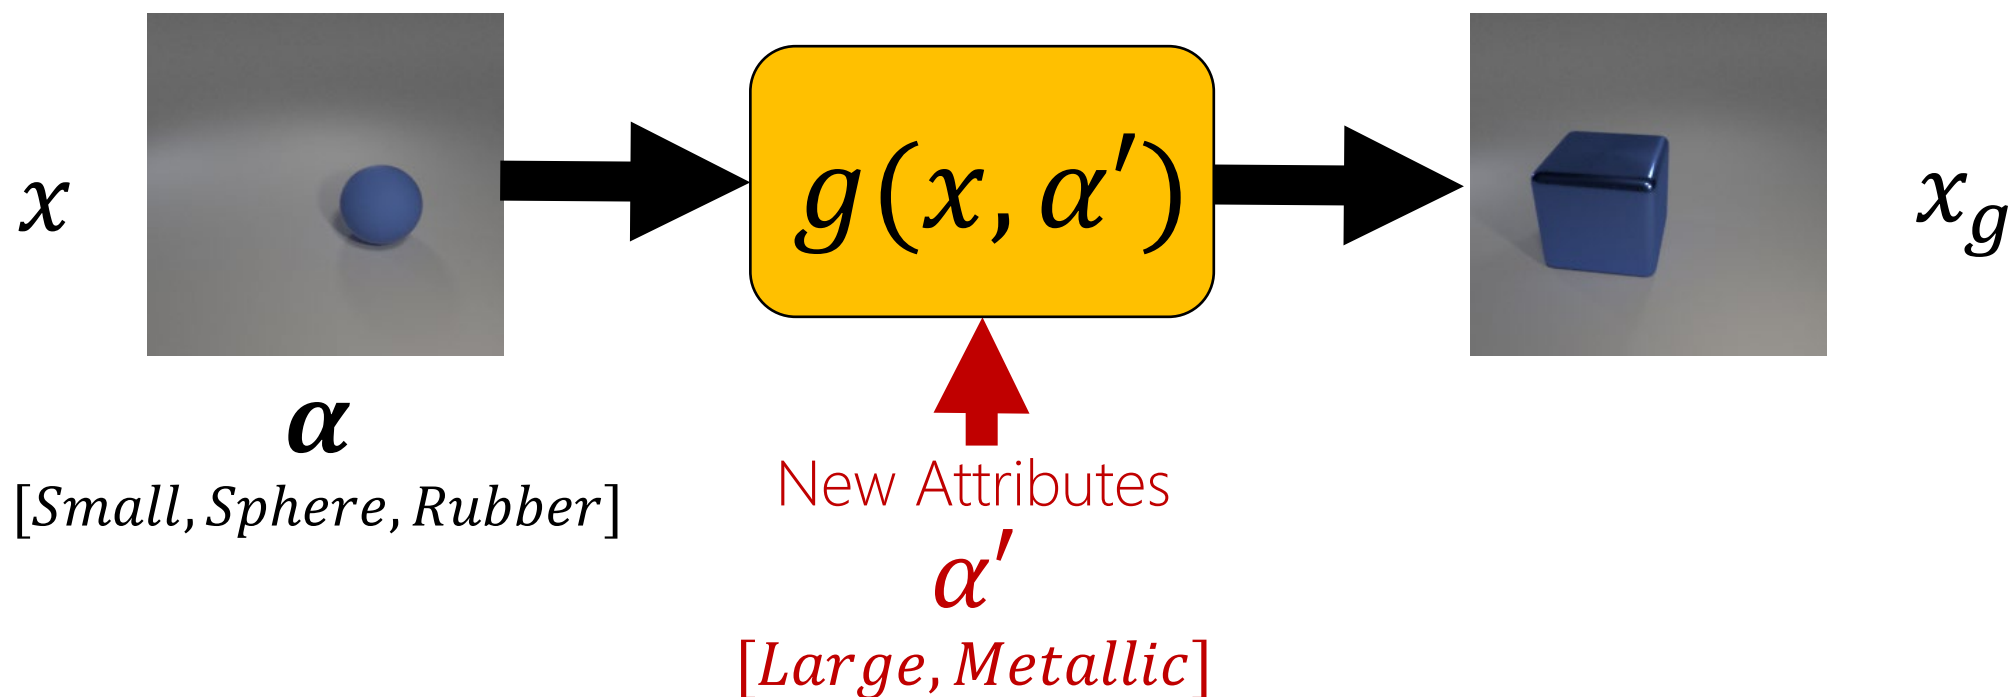


Attribute-Guided Adversarial Training

- Desirable Properties of Generative Function $g()$
 - Generate **plausible** and **diverse** perturbations of attributes
 - Reflect a **larger coverage of attribute space than training data**
 - Generate **novel attribute combinations**




Attribute-Guided Adversarial Training



Discover attribute combinations that are adversarial to the classifier.

$$\max_{\alpha'} \ell(f(g(x, \alpha')), y)$$

Attribute-Guided Adversarial Training



New Attributes

α'


[*Large, Metallic*]

Discover attribute combinations that are adversarial to the classifier.


AND explore new regions in the attribute space

$$\max_{\alpha'} \ell(f(g(x, \alpha')), y) + \gamma \|\alpha - \alpha'\|_2$$

Pre-Training Phase




Classifier
Loss




θ

Discover Adversarial Attributes




Adversarial
Maximization




α^*


Generate Augmentations




$g(x, \alpha^*)$




Train with Augmented Data



+



Classifier
Loss



θ^*


AGAT is effective for Discrete Attribute-Shift

- CLEVR-Singles
- Attributes: Size, Shape, Material, Position
- Train—Test split s.t.
 - Limited attribute combinations are observed in during training
 - Performance is evaluated on all combinations


TASK: color classification

TRAINING


Metal objects far from the camera
Rubber objects close to the camera



Metal, far



Metal, far



Rubber, close




Rubber, close

TESTING

Metal objects close to the camera
Rubber objects far from the camera




Metal, close




Metal, close



Rubber, far



Rubber, far




AGAT is effective for Geometric Shift

Training




Testing

(Rotation/Translation/Scaling)



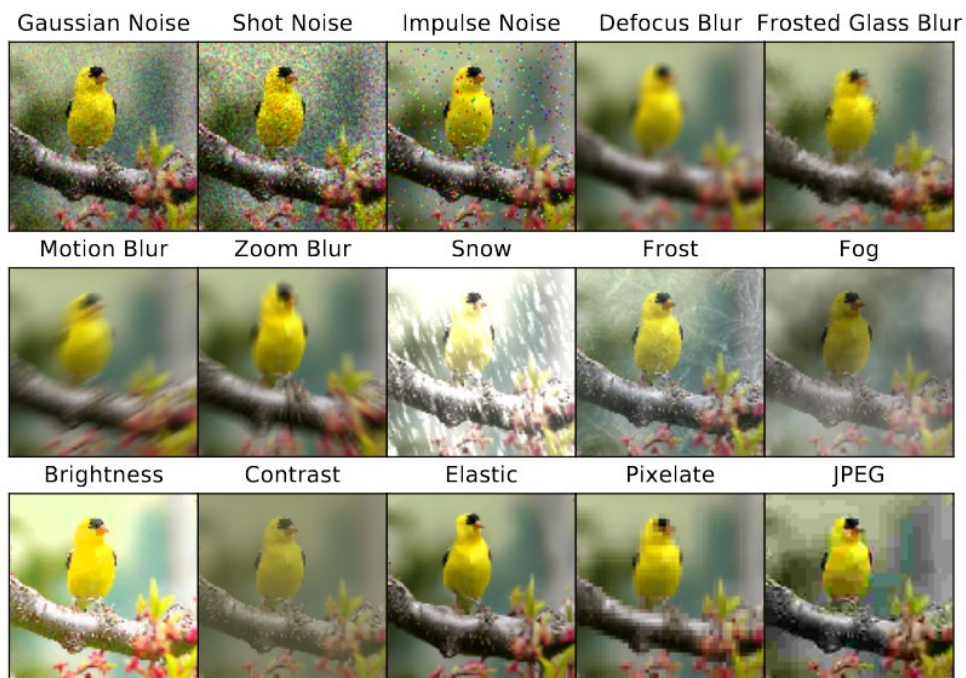
- Training data: MNIST digits
- Test data: **unknown rotation, translation, scaling**
- AGAT outperforms prior work by a large margin



Results (3): AGAT is Effective for Natural Corruptions


Commonly occurring corruptions

(CIFAR-10-C, *Hendrycks et al. ICLR 2019*)



- No access to specific attributes
- We use pseudo-attributes (Gaussian coefficients α_1, α_2)

$$x_g = \frac{1}{\alpha_1 \sqrt{2\pi}} e^{-\frac{x^2}{2\alpha_1^2}} + n; \quad \text{where } n \sim \mathcal{N}(0, \alpha_2)$$



Adversarially Discovering Image Transformations

Without domain knowledge: **ALT**

With attribute-level knowledge: **AGAT**

Strong results for many types of distribution shift

