Google Research

Rethinking Domain Generalization for Face Anti-spoofing: Separability and Alignment

Tag: THU-PM-379




Yiyou Sun
UW-Madison

Yaojie Liu
Google

Xiaoming Liu
Michigan State University
Google

Yixuan Li
UW-Madison

Wen-Sheng Chu
Google

Face Anti-spoofing (FAS) is Important!

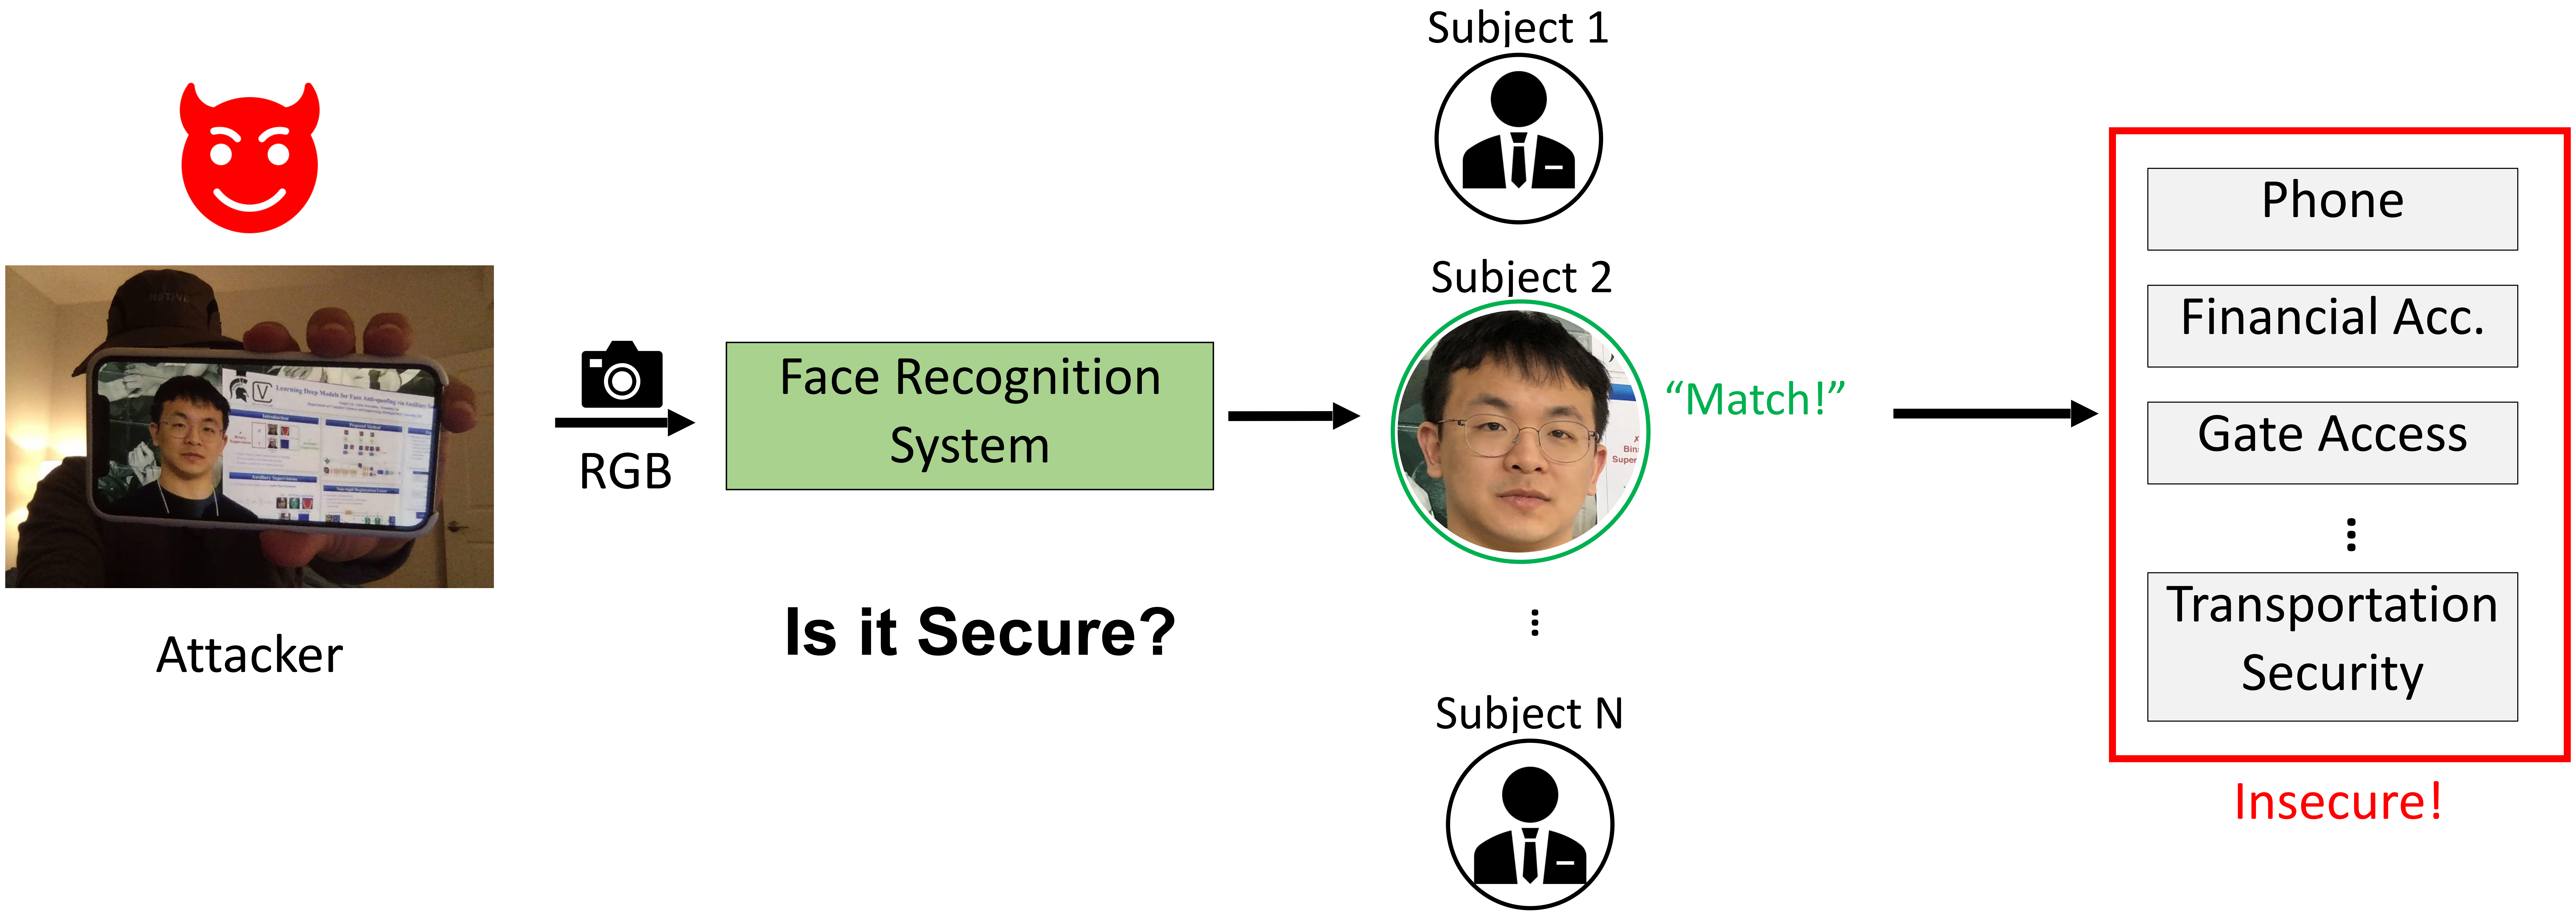

Face Anti-spoofing (FAS) is Hard in the Wild

Different Cameras

Different Environments

New Challenges: design algorithm well with **domain generalization**!

1-min Highlight: Motivation

Learns a Domain-invariant Live/Spoof Classifier

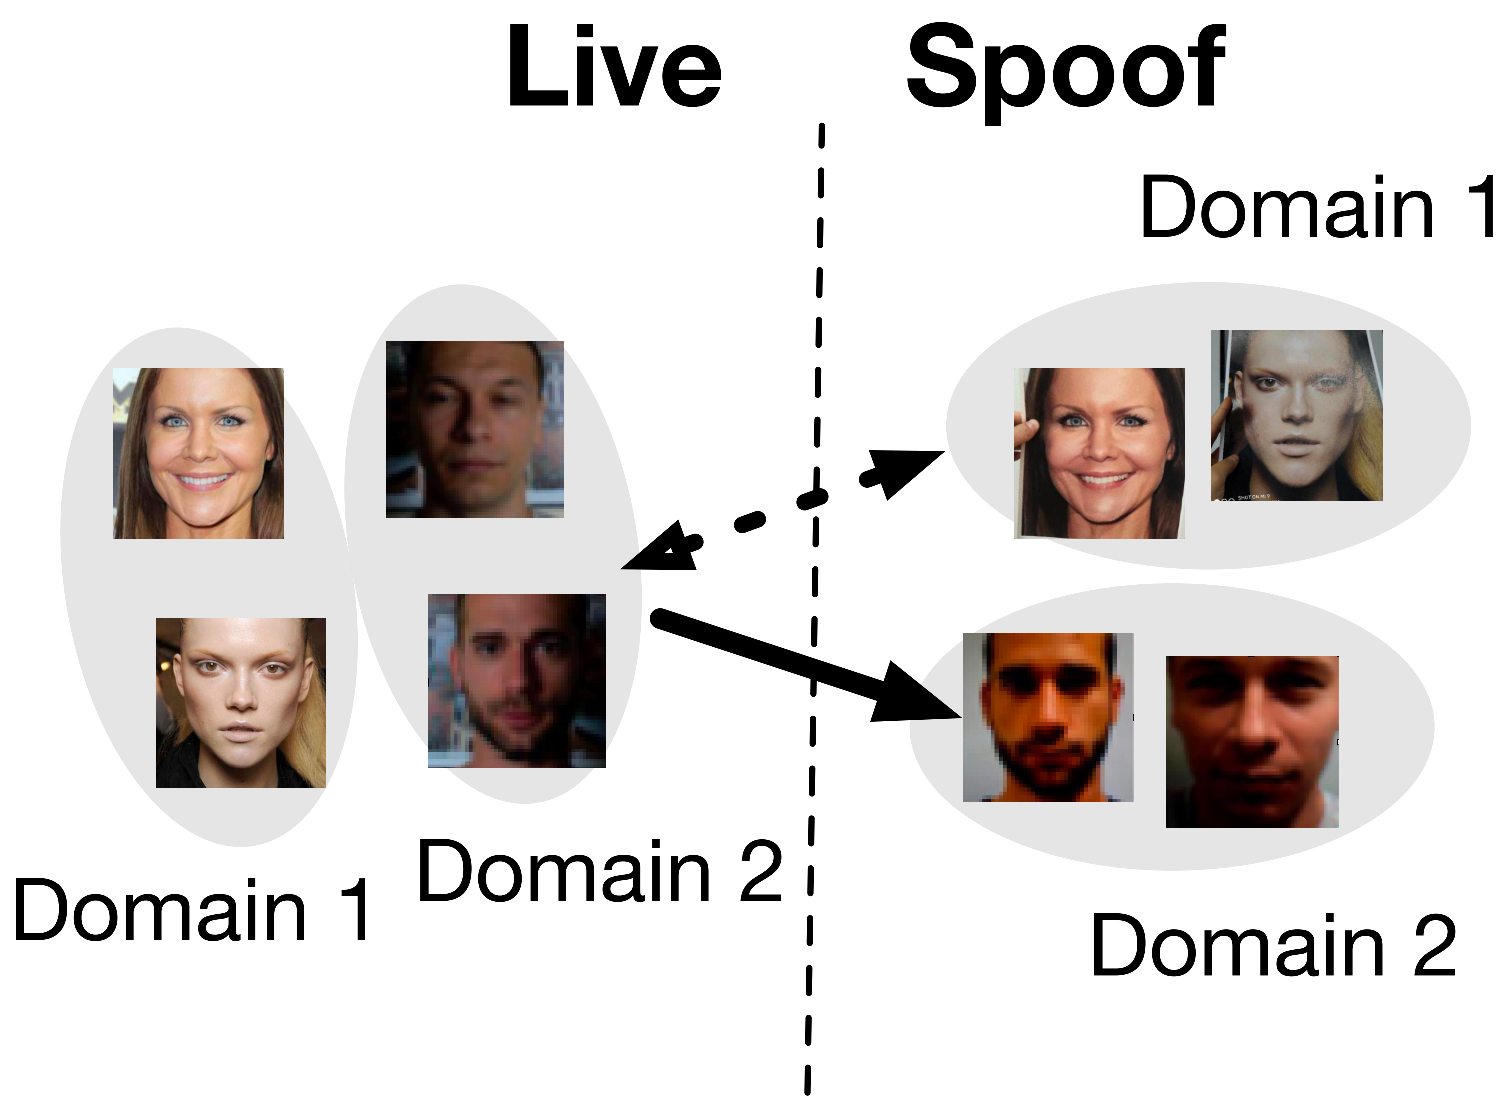

- Domain signal is **ignored**
- Live-to-spoof transition is **inconsistent**

Common Solutions:
SSDG (CVPR 2020) / SSAN (CVPR 2022)

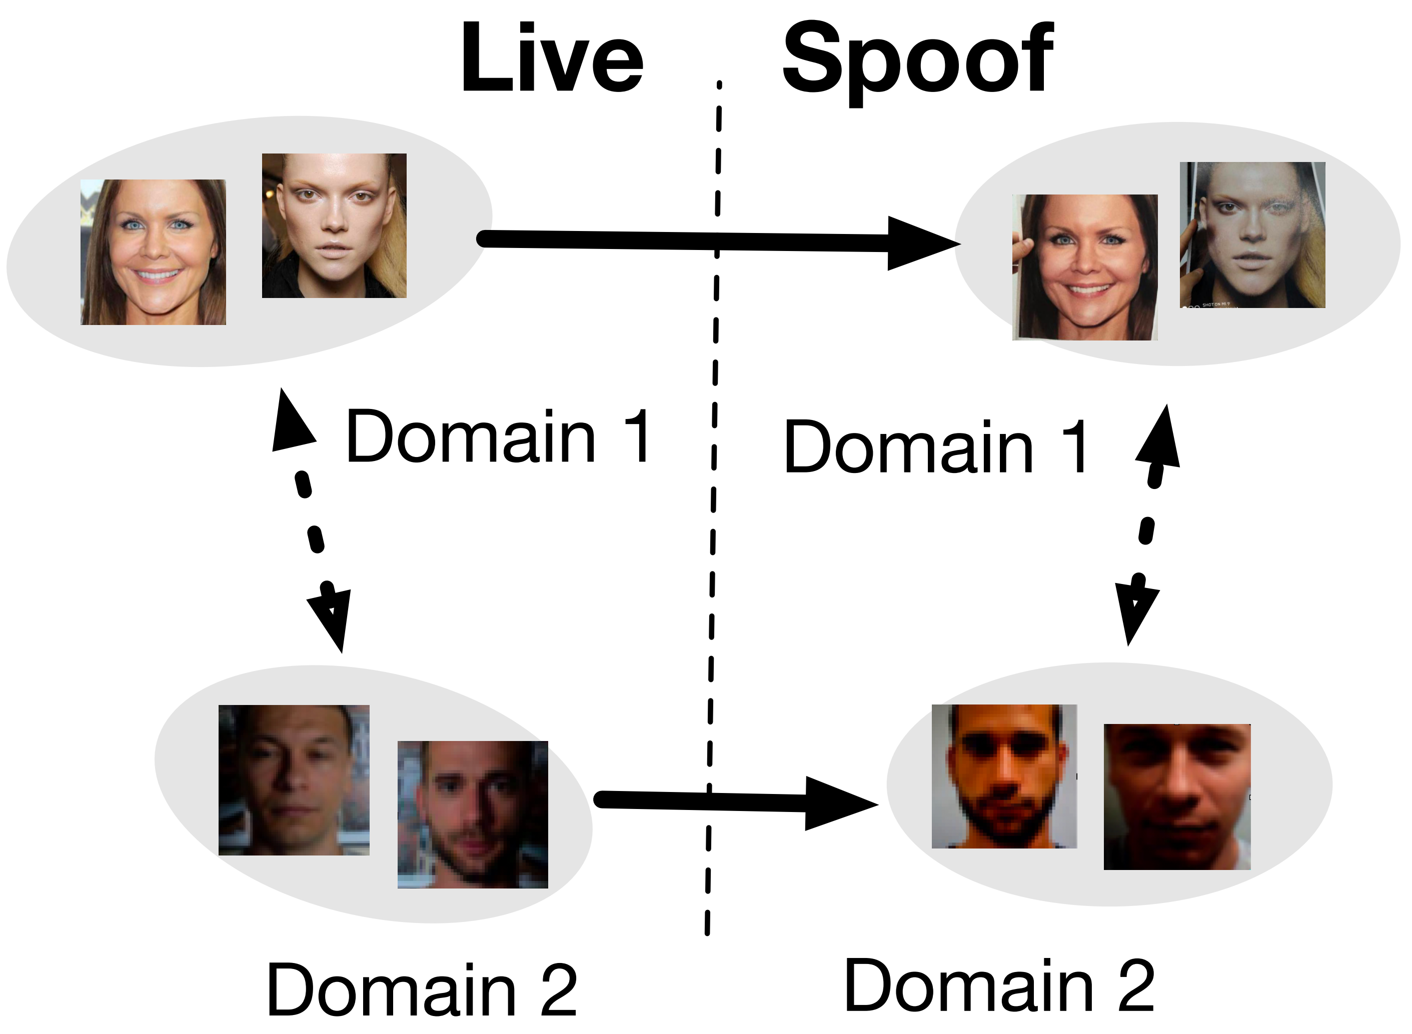

- Domain signal is **leveraged**
- Live-to-spoof transition is **aligned**

SAFAS (Our Solution)

- Jia et al., "Single-side domain generalization for face anti-spoofing," CVPR 2020.
- Wang et al., "Domain generalization via shuffled style assembly for face anti-spoofing," CVPR 2022.

1-min Highlight: Methodology

Separability (SupCon) and Alignment (IRM)

Challenge 1: Separability

Supervised Contrastive Learning (SupCon)

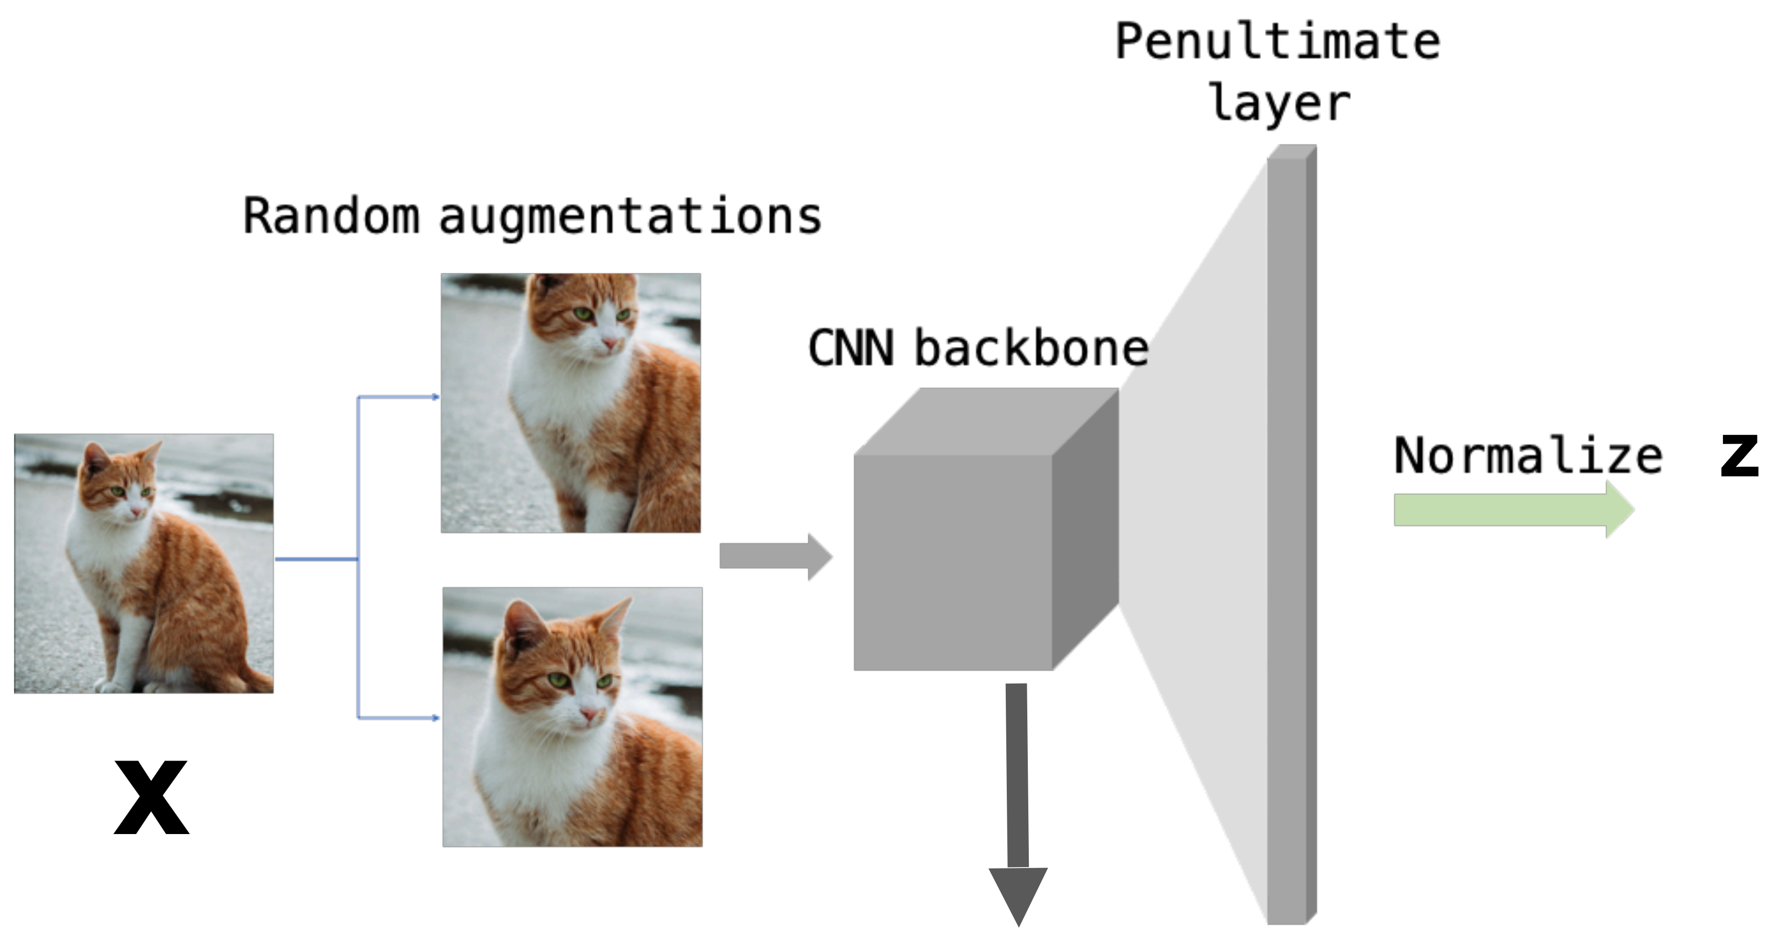

$$\phi : \mathcal{X} \mapsto \mathbb{R}^m$$

$$\mathcal{L}_{SupCon} = -\frac{1}{|\mathcal{P}(\mathbf{x})|} \sum_{\mathbf{z}^+ \in \mathcal{P}(\mathbf{x})} \log \frac{\exp\left(\mathbf{z}^\top \cdot \mathbf{z}^+ / \tau\right)}{\sum_{\mathbf{z}^- \in \mathcal{N}(\mathbf{x})} \exp\left(\mathbf{z}^\top \cdot \mathbf{z}^- / \tau\right)}$$

Challenge 2: Alignment

Invariant Risk Minimization (IRM)

$$\min_{\phi, \beta^*} \frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} \mathcal{R}^e(\phi, \beta^*) \rightarrow \mathcal{L}_{IRM}$$

$$\boxed{\text{s.t.} \quad \beta^* \in \arg\min_{\beta} \mathcal{R}^e(\phi, \beta) \quad \forall e \in \mathcal{E},}$$

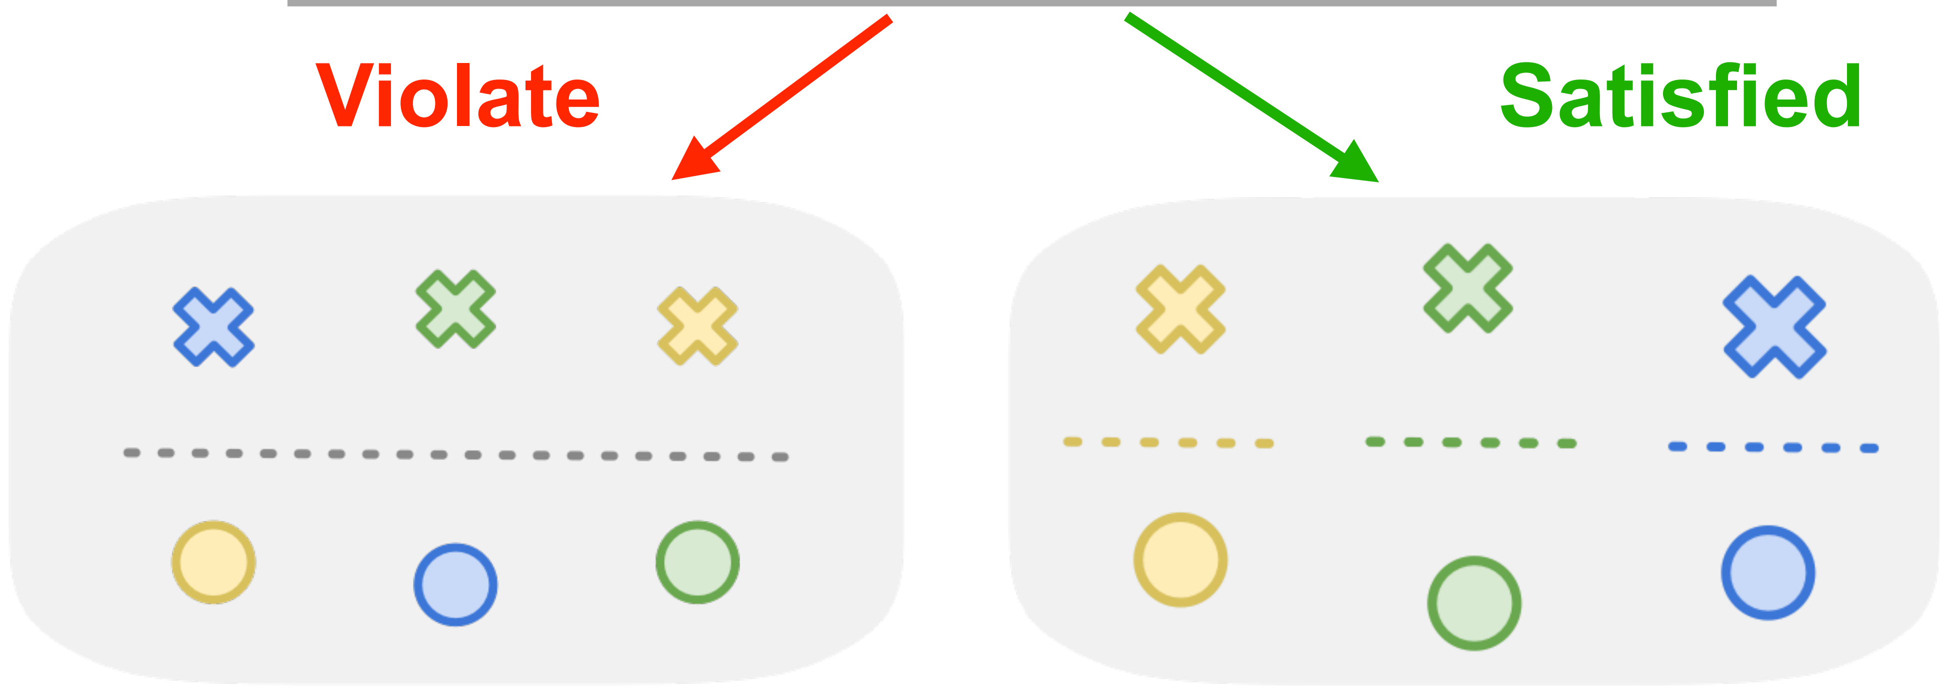

1-min Highlight: Experiment

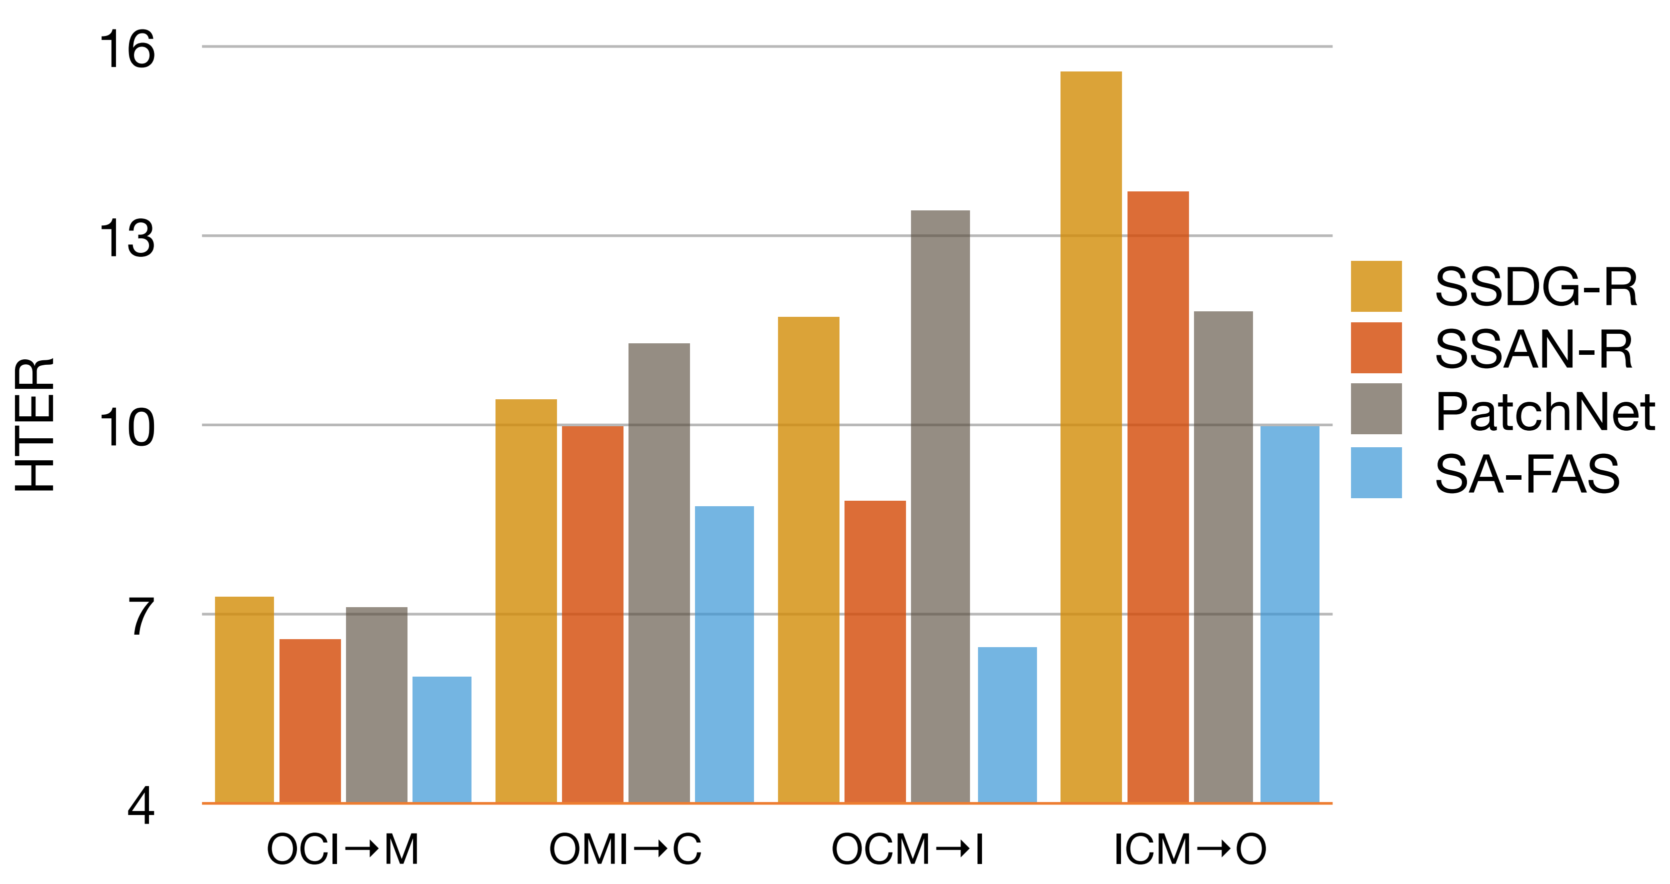

(a) Comparison with Baselines (Best Possible Performance)

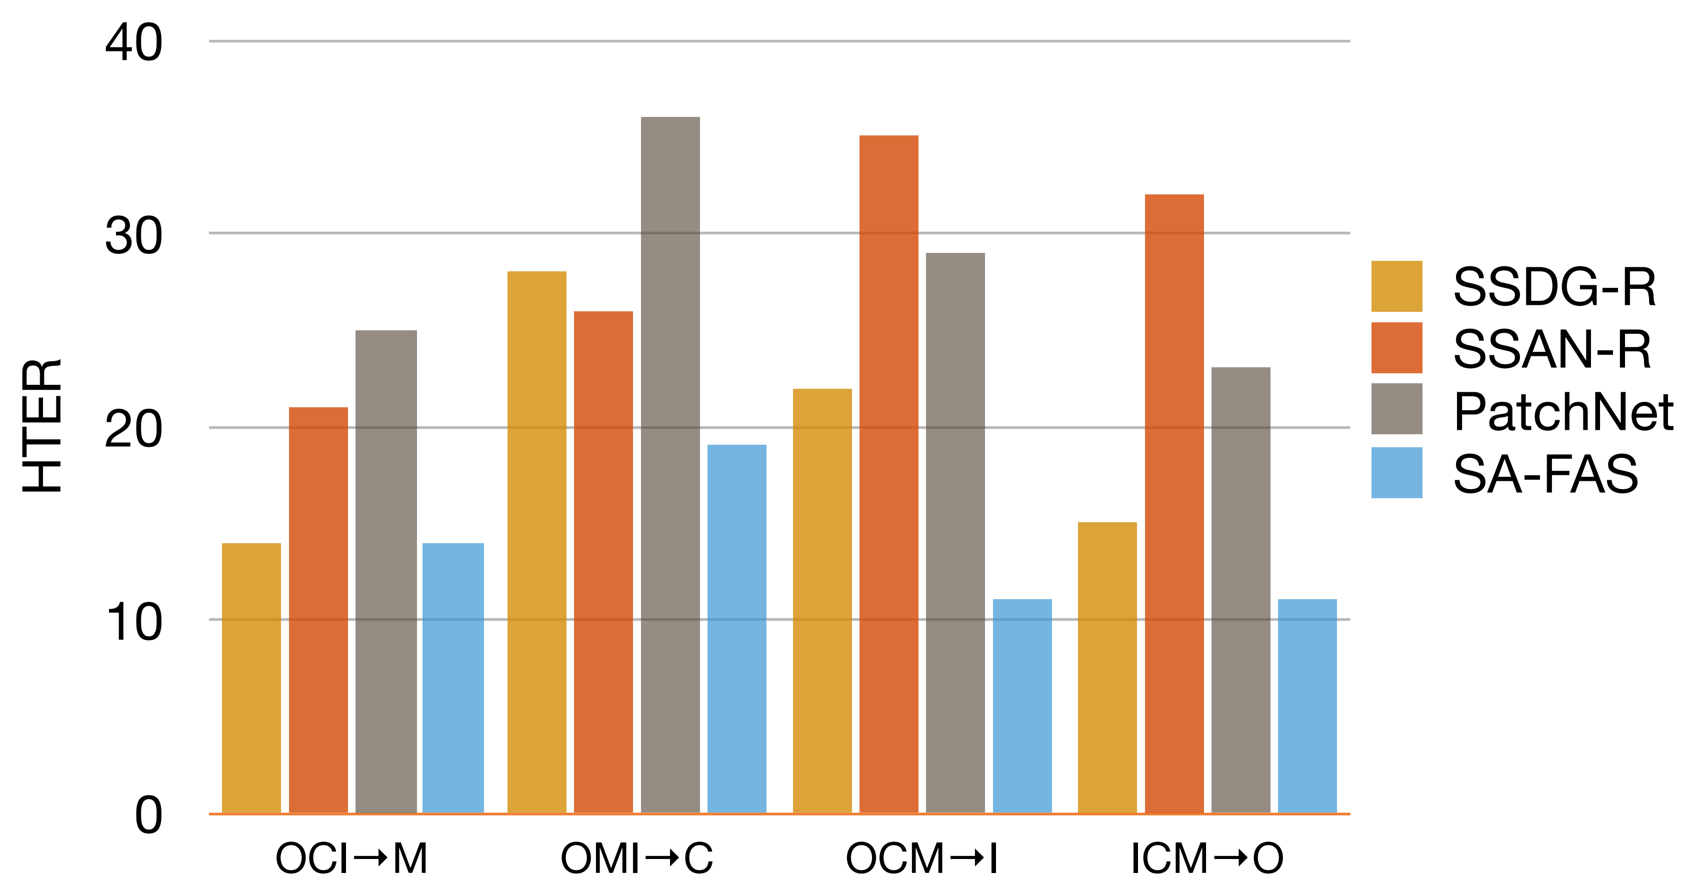

(b) Comparison with Baselines (Upon Convergence)

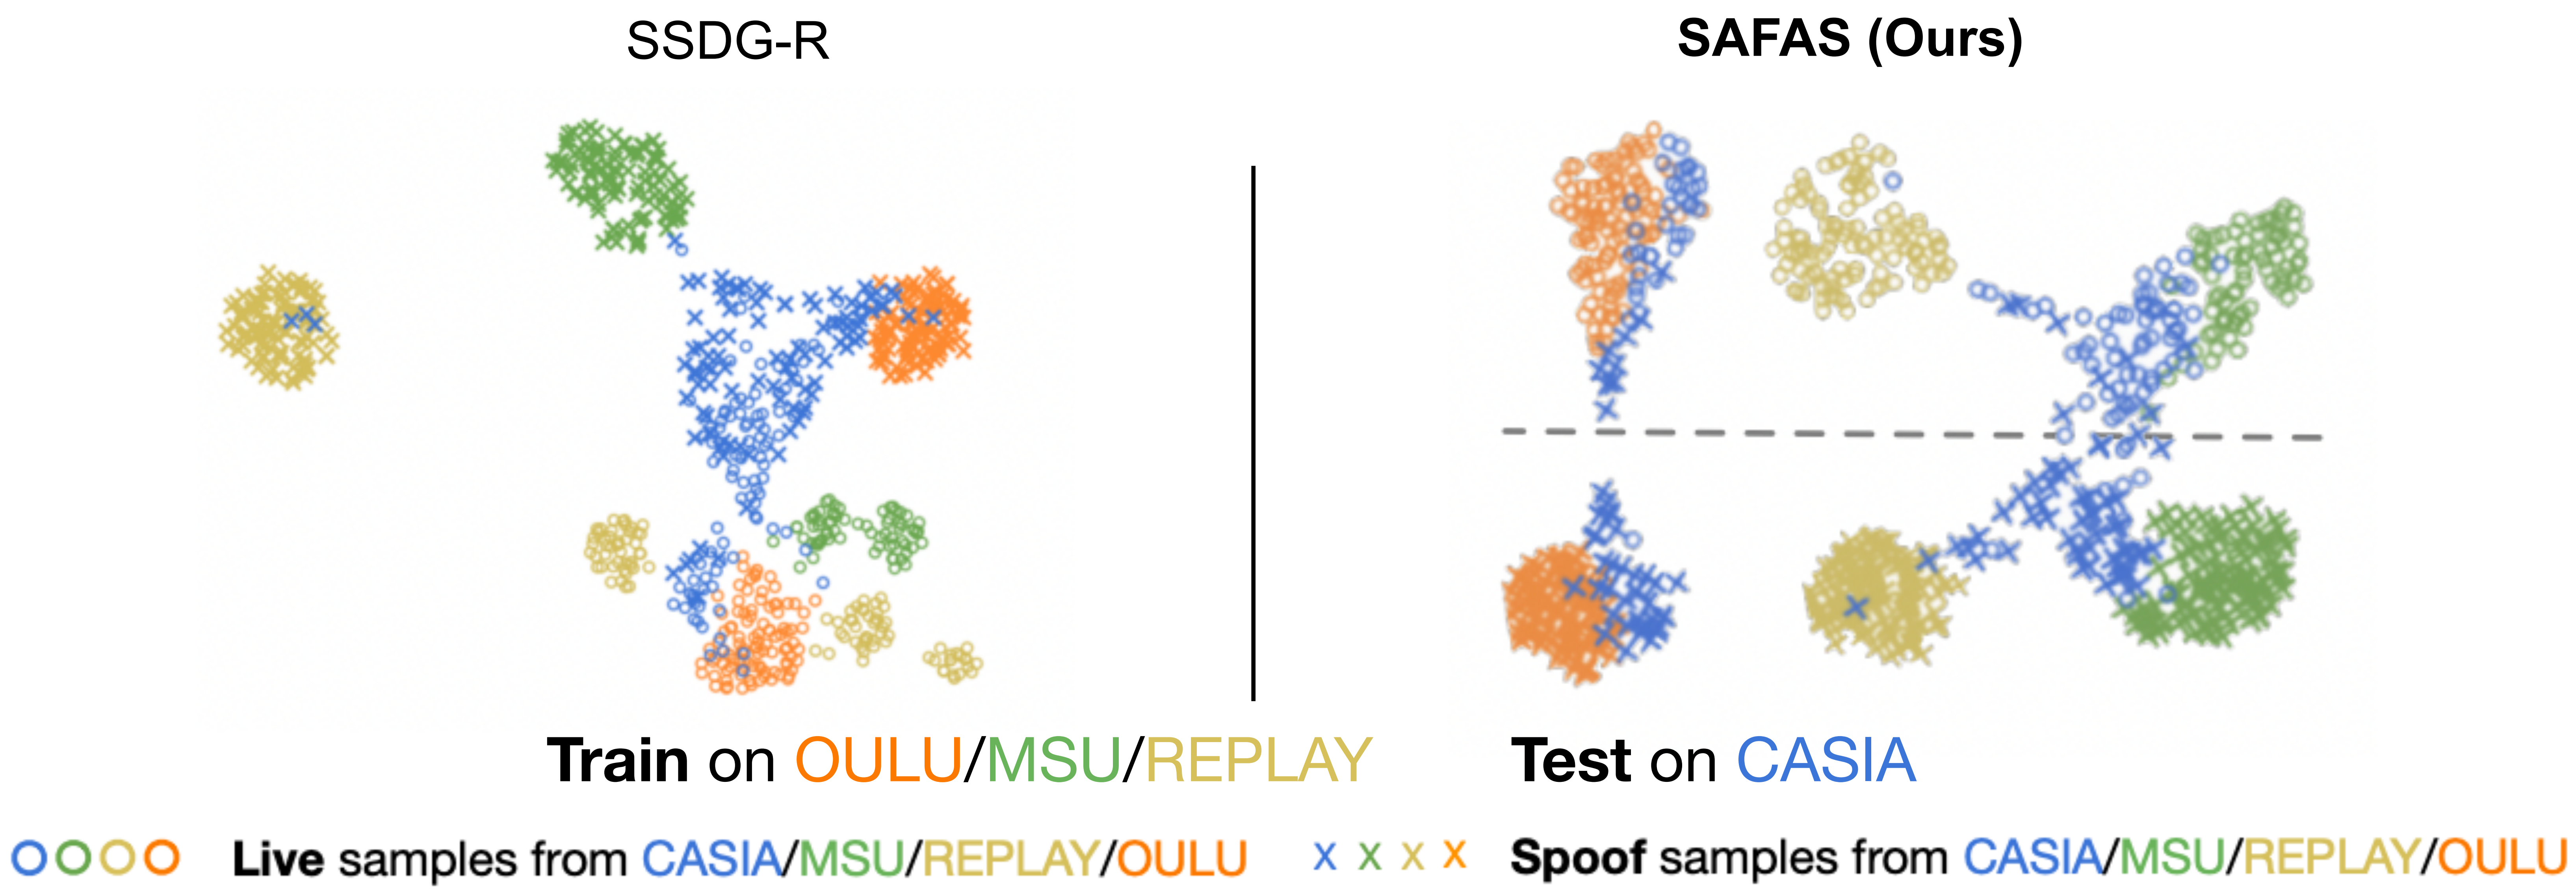

(c) Visualization of Feature Space

Problem Setting

Training Data

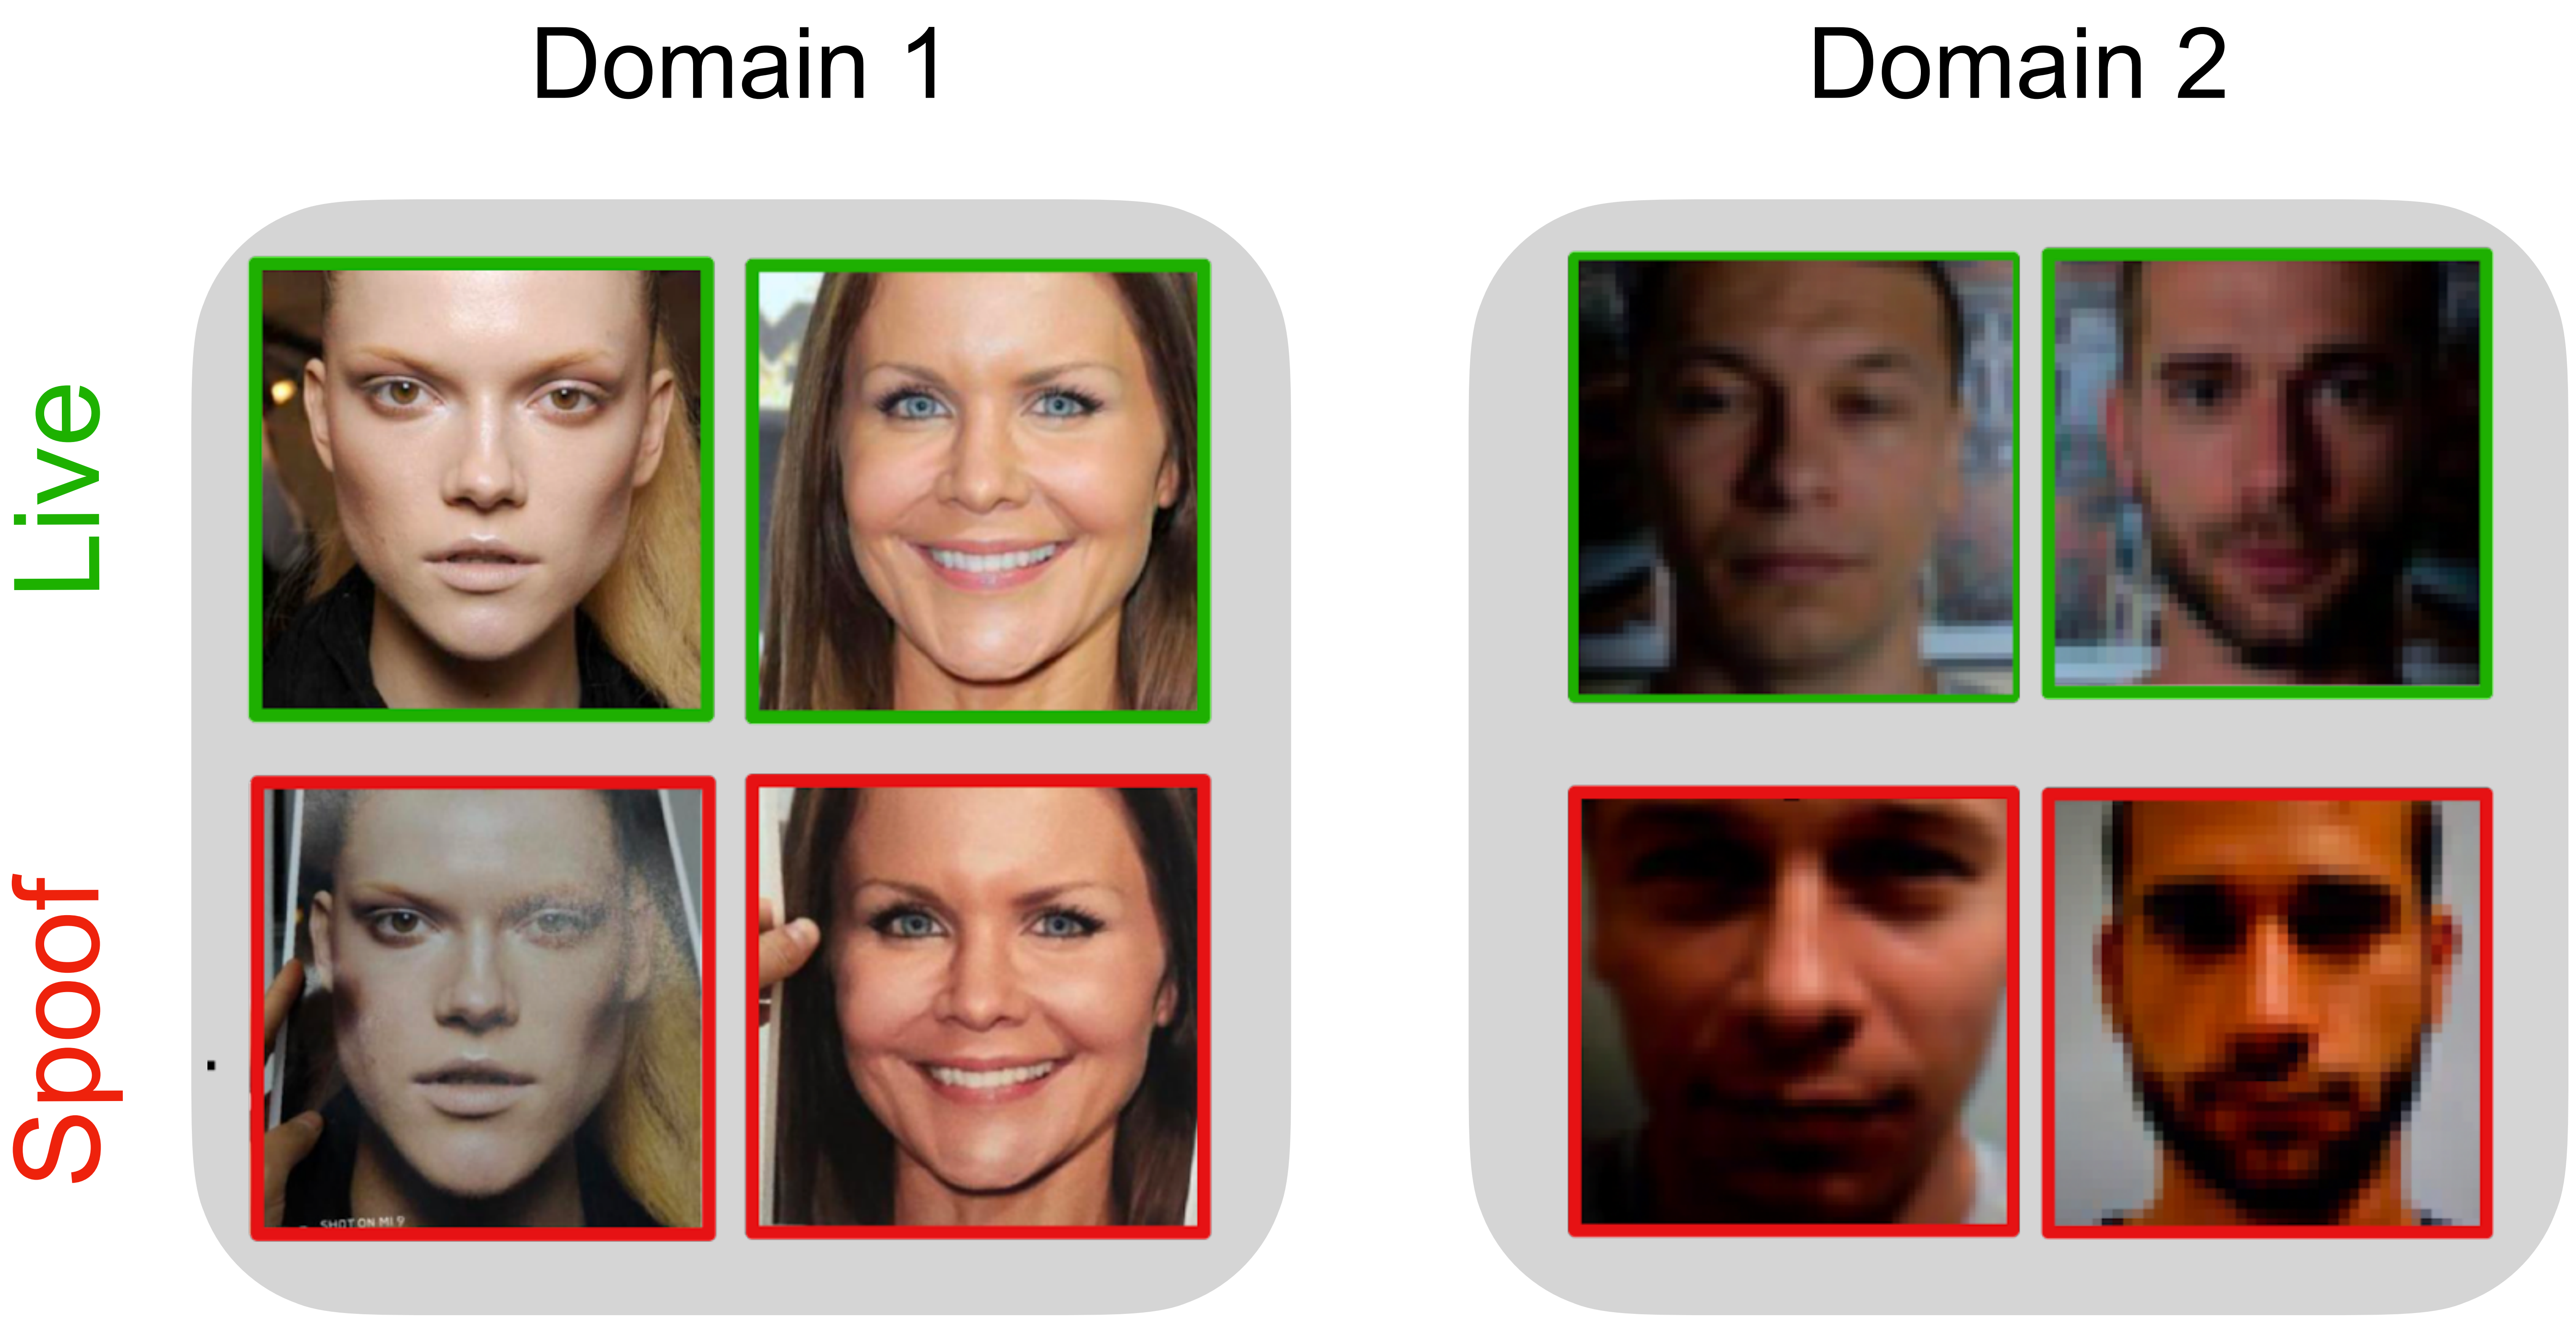

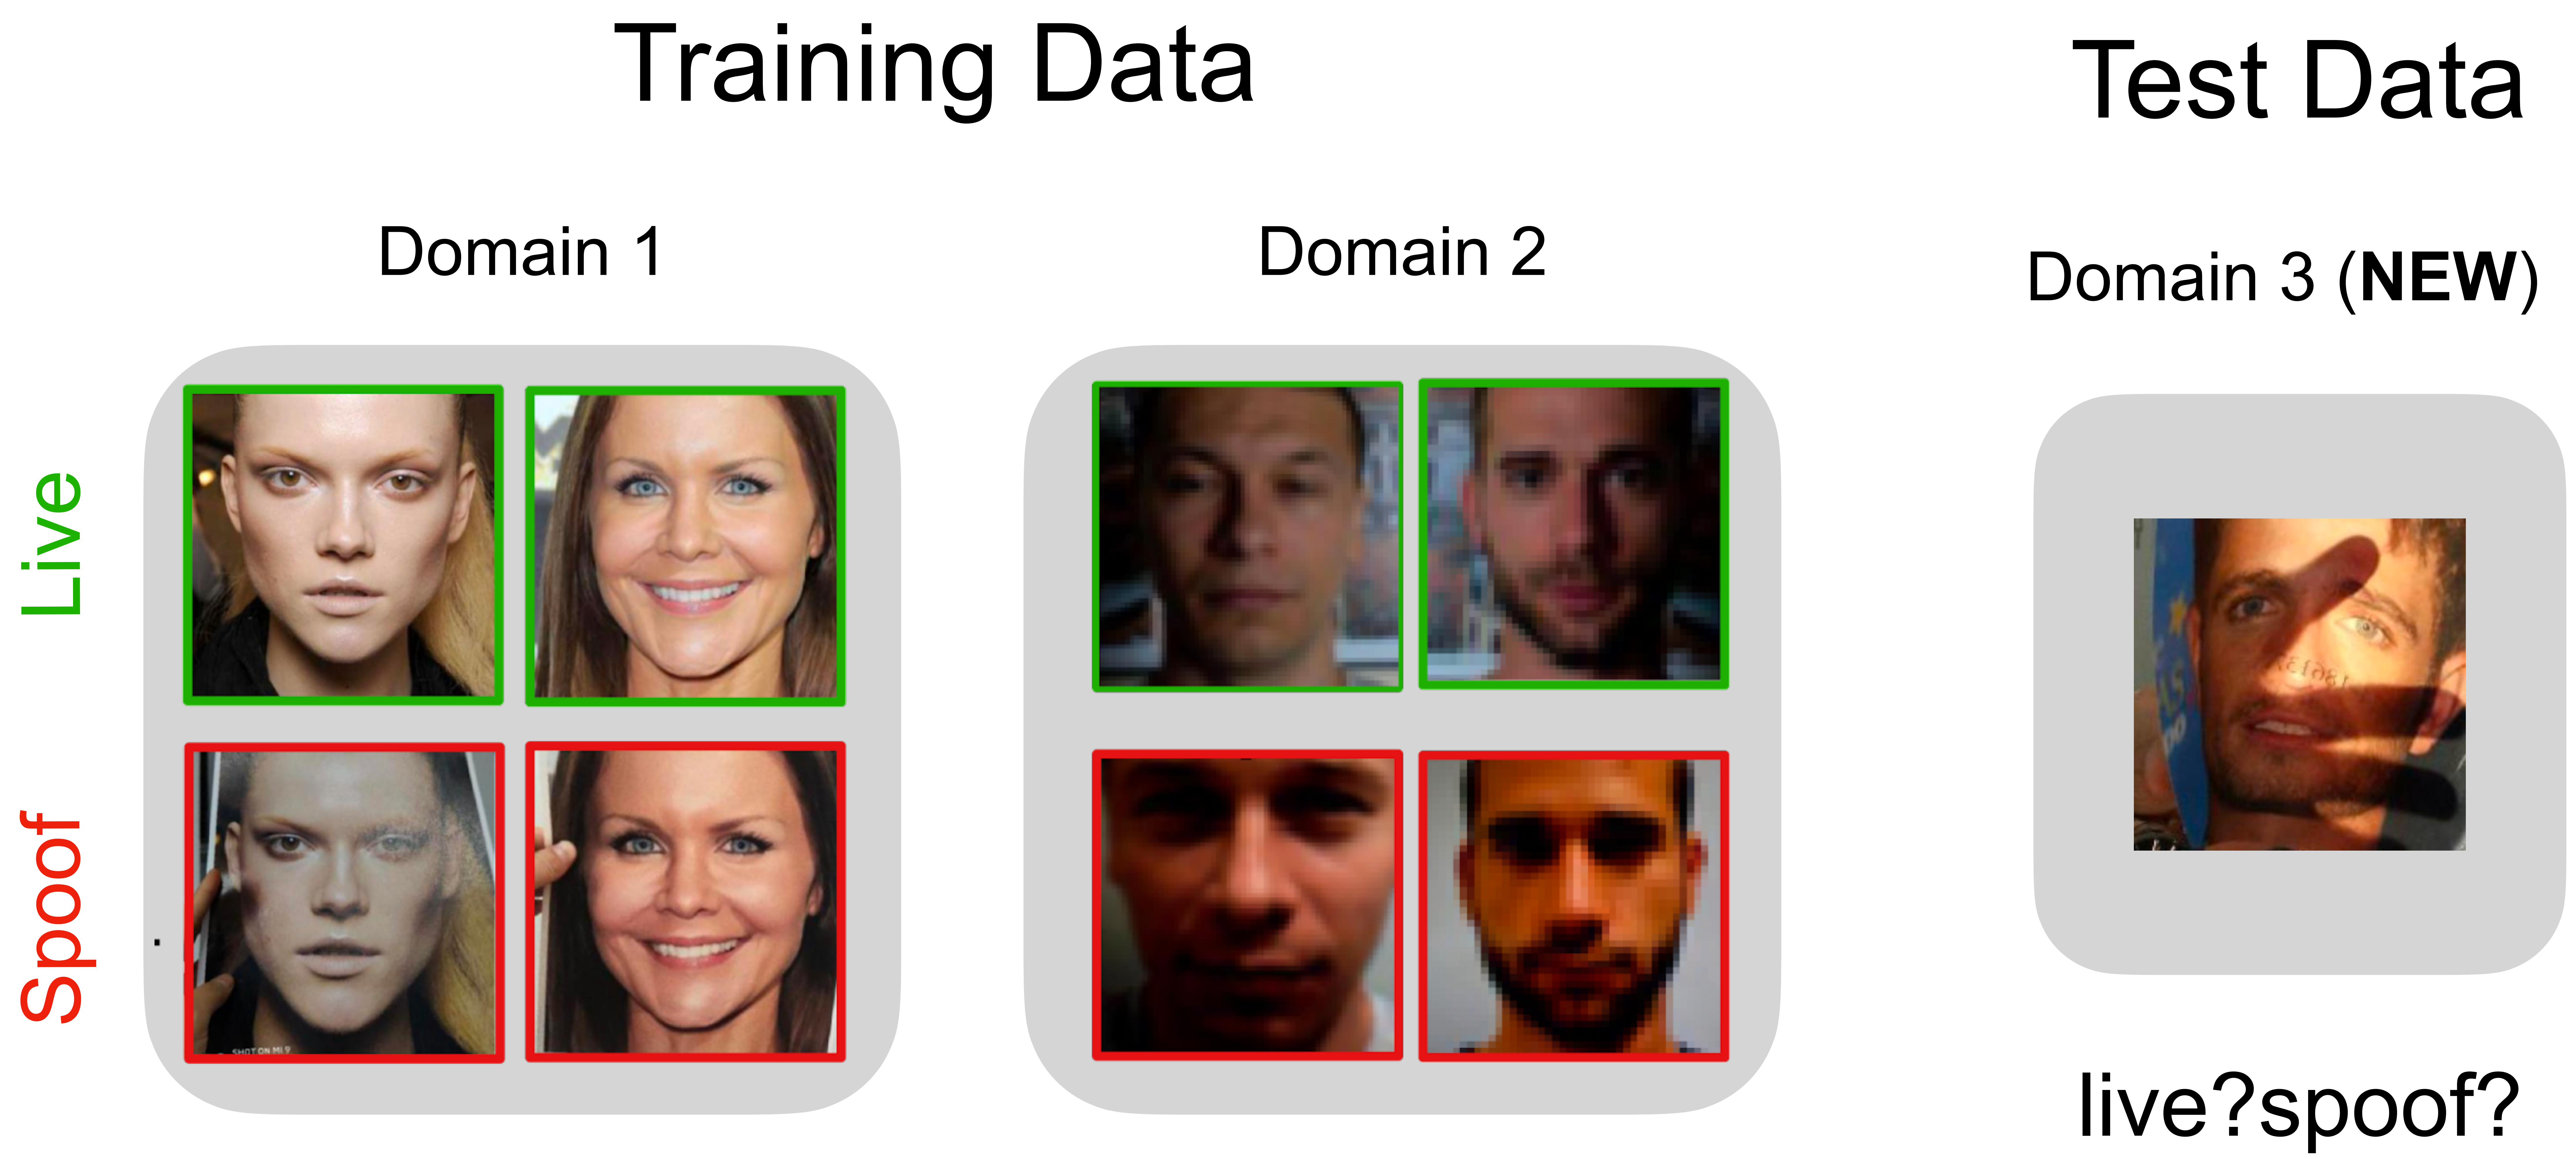
Training Data
Test Data
Domain 1
Domain 2
Domain 3 (NEW)
Live
Spoof
live?spoof?

Prior Solutions for Cross-domain FAS

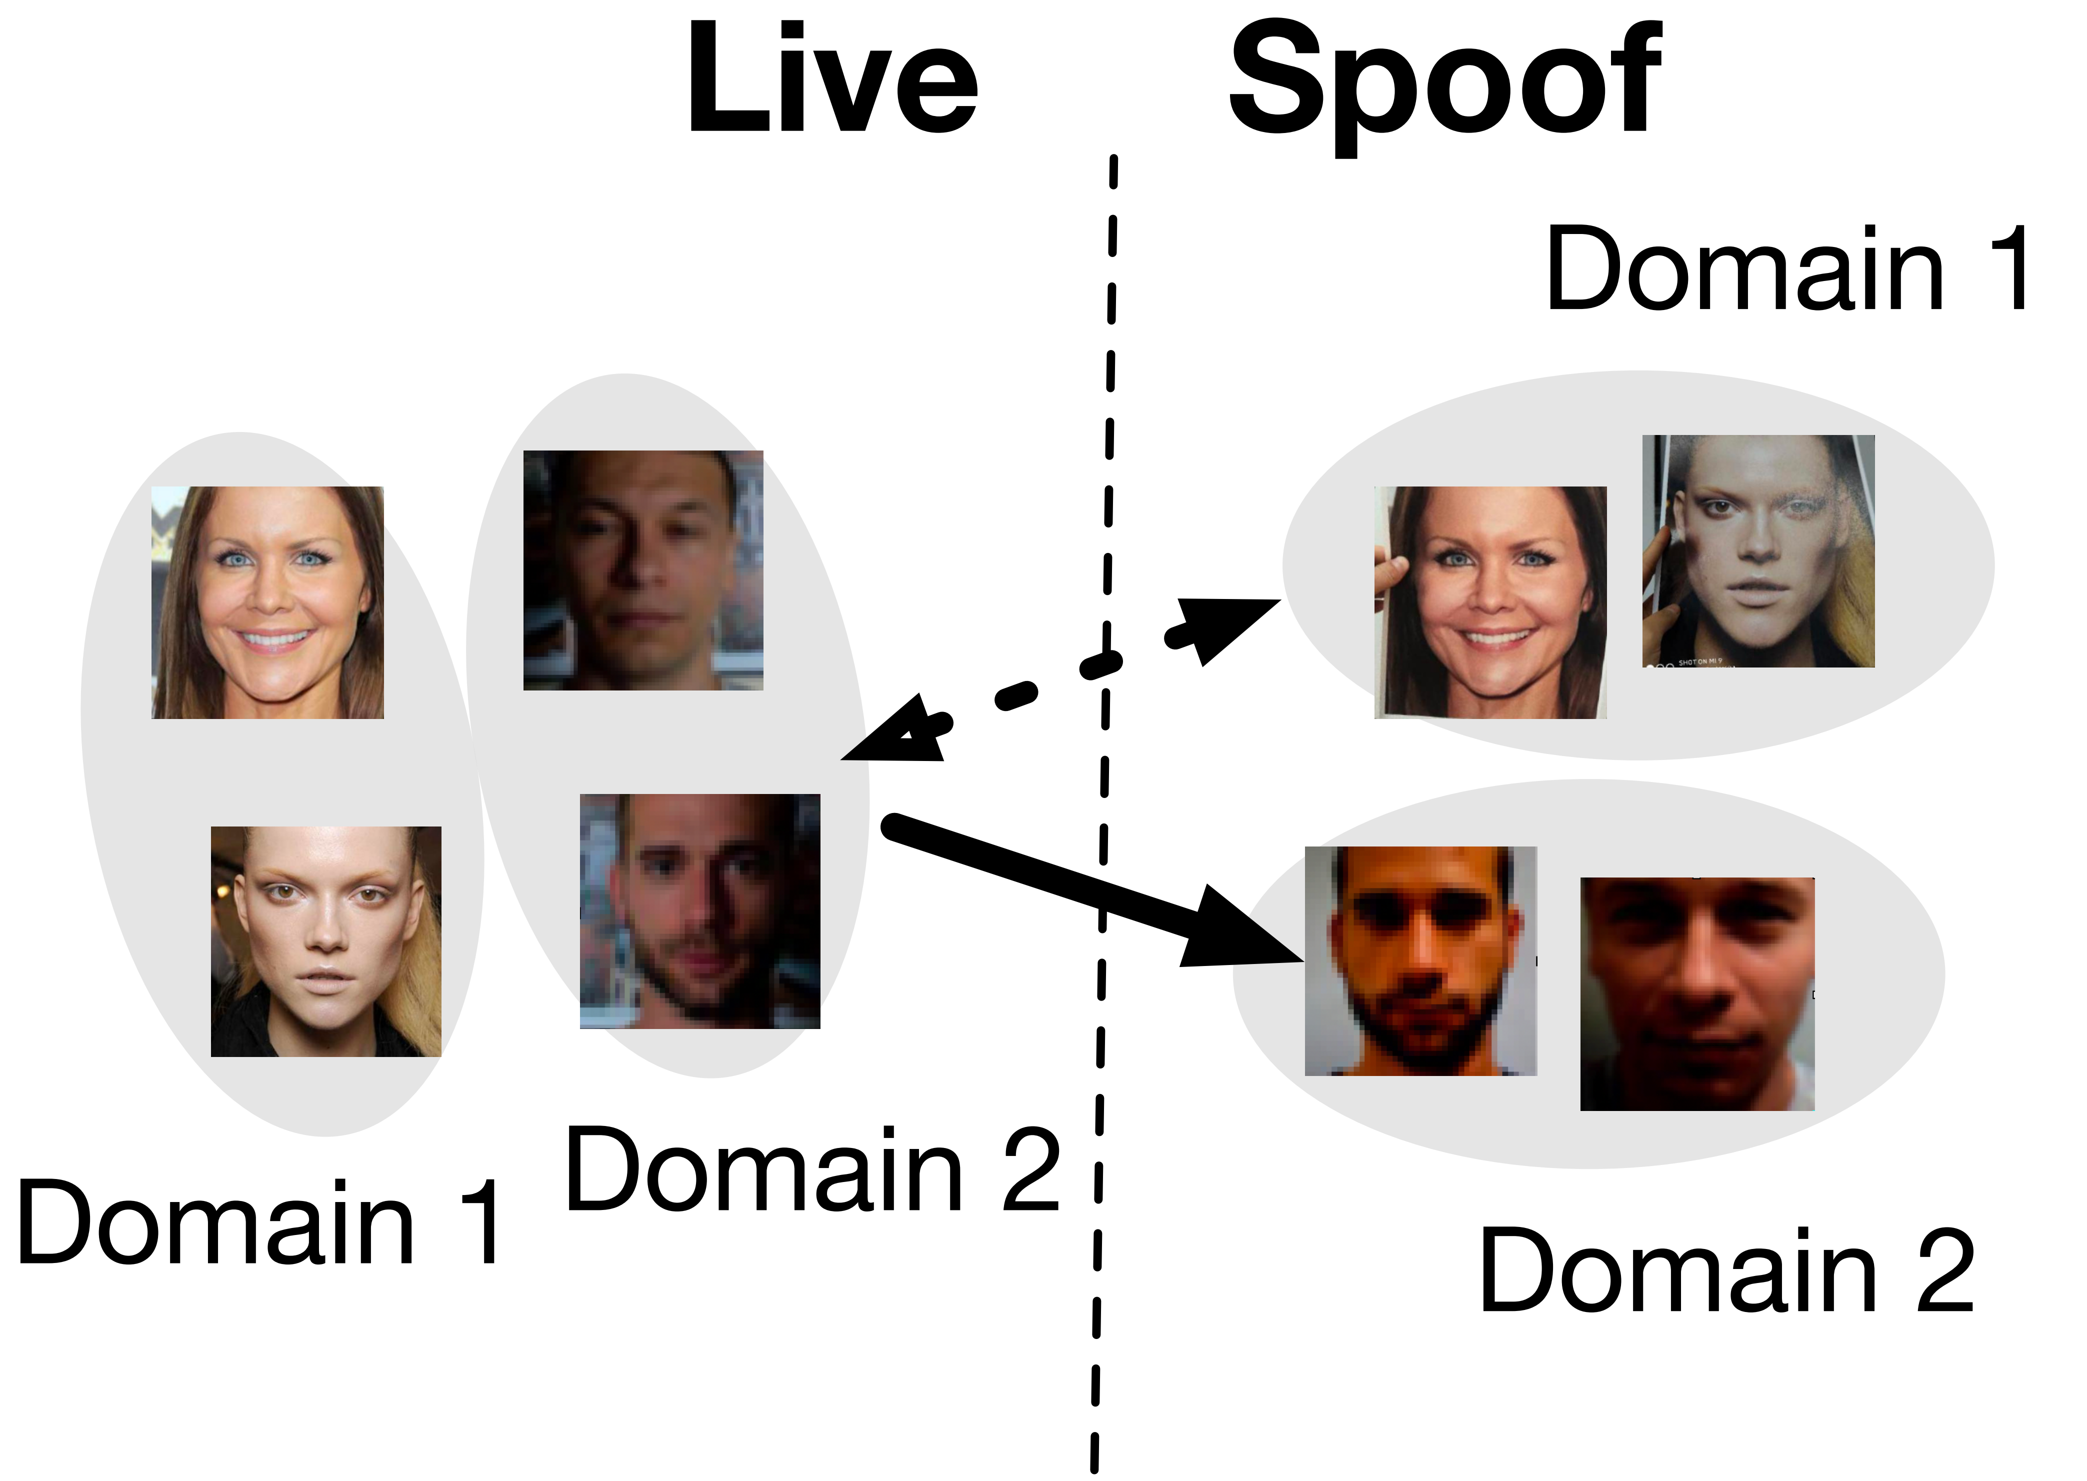

- Domain signal is **ignored**
- Live-to-spoof transition is **inconsistent**

Common Solutions: SSDG (CVPR 2020) / SSAN (CVPR 2022)

- Jia et al., "Single-side domain generalization for face anti-spoofing," CVPR 2020.
- Wang et al., "Domain generalization via shuffled style assembly for face anti-spoofing," CVPR 2022.

Our Solution

Learns a Domain-invariant Live/Spoof Classifier

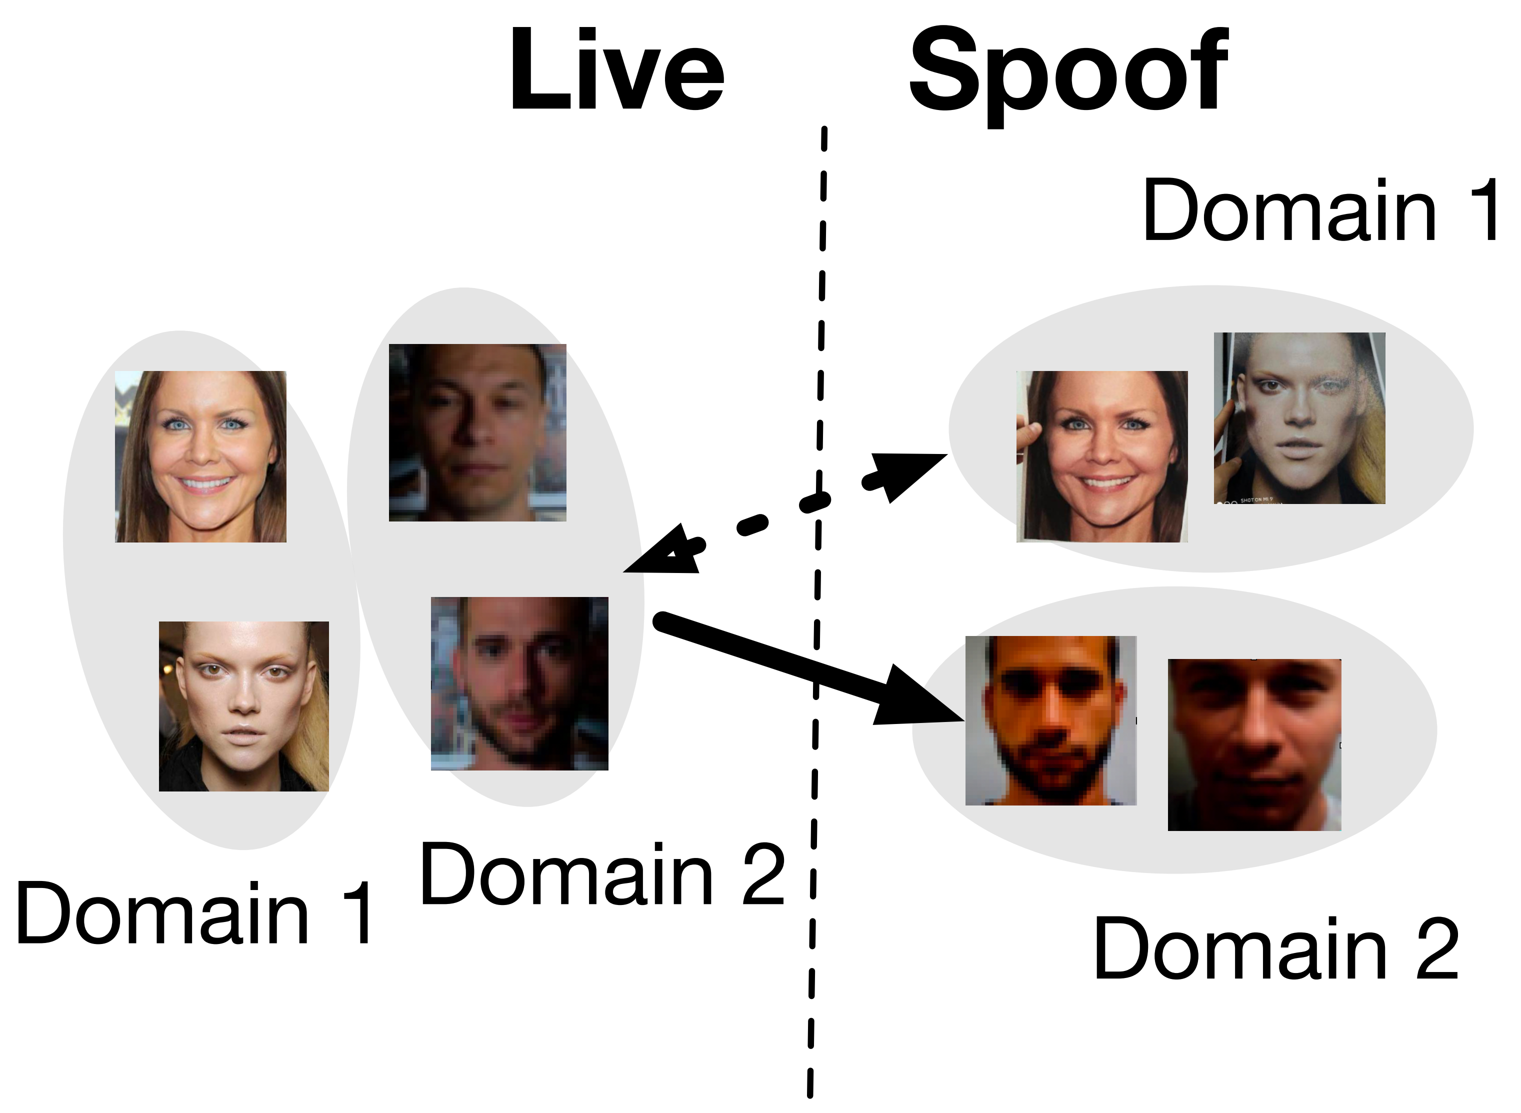

- Domain signal is **ignored**
- Live-to-spoof transition is **inconsistent**

Common Solutions:
SSDG (CVPR 2020) / SSAN (CVPR 2022)

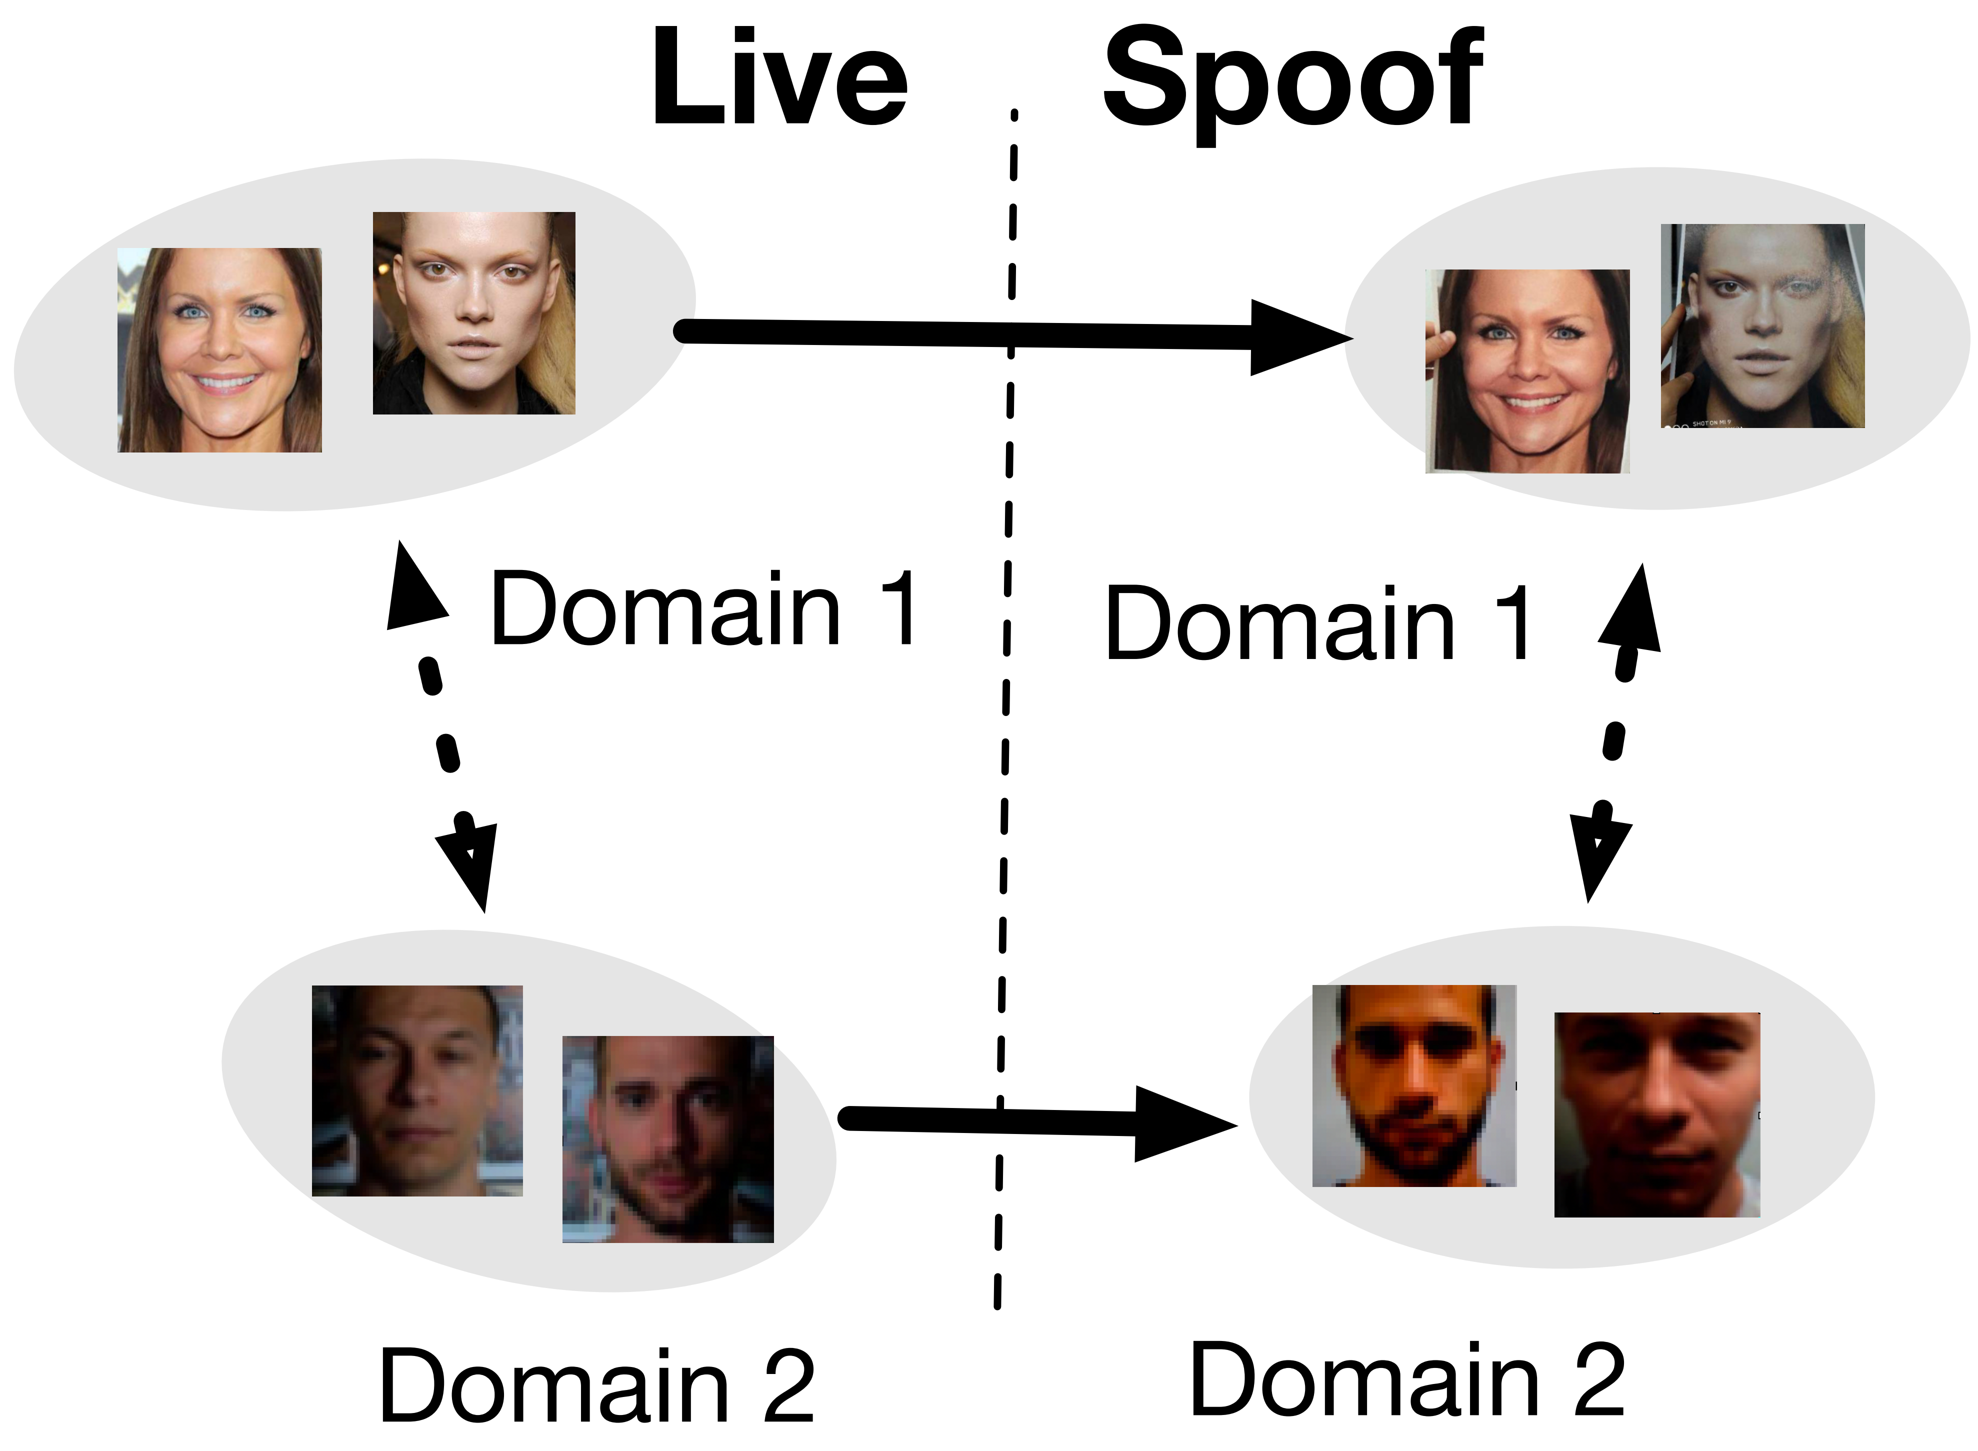

- Domain signal is **leveraged**
- Live-to-spoof transition is **aligned**

SAFAS (Our Solution)

- Jia et al., "Single-side domain generalization for face anti-spoofing," CVPR 2020.
- Wang et al., "Domain generalization via shuffled style assembly for face anti-spoofing," CVPR 2022.

How to Achieve the Ideal Feature Space?

Two Key Challenges

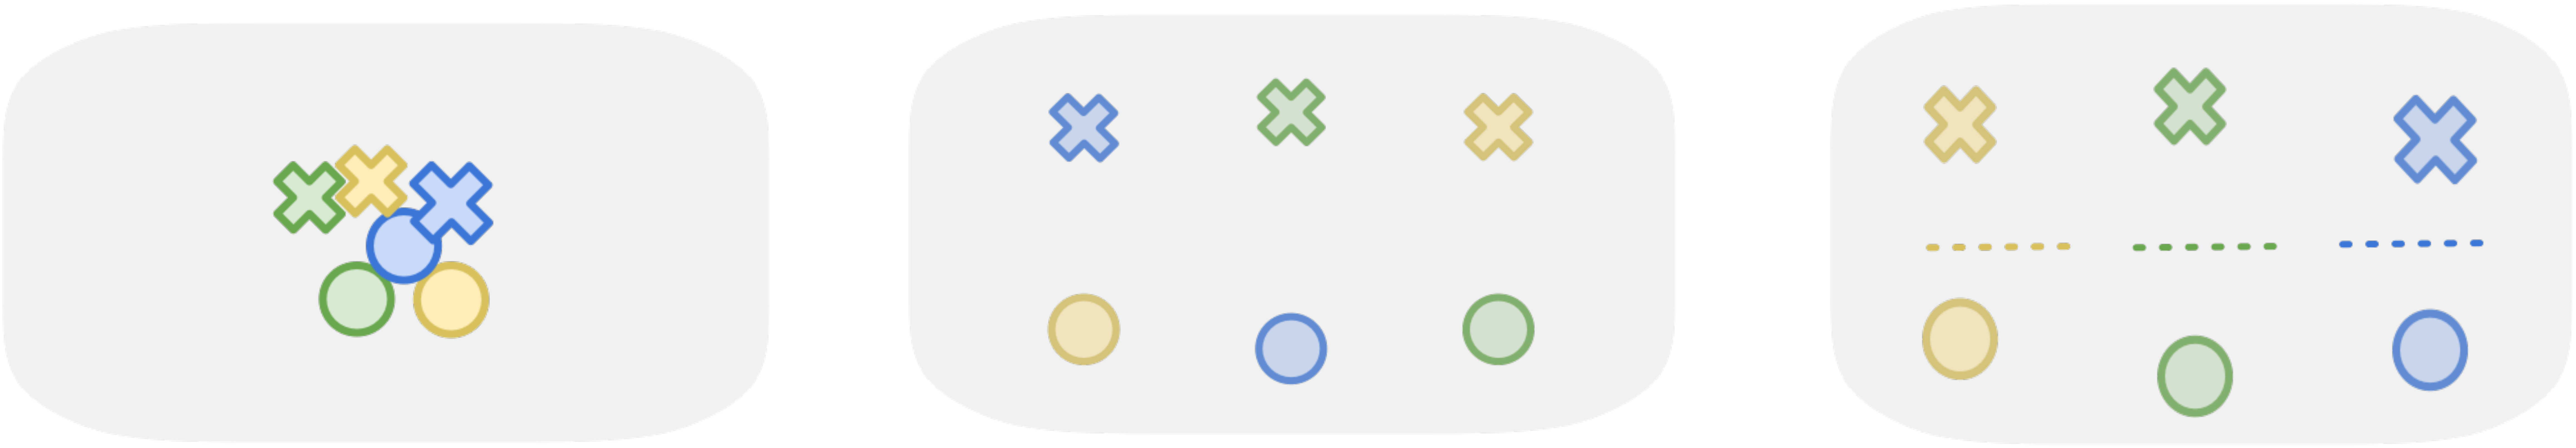

Original Feature Space

Challenge 1:
Separate Feature Space

Challenge 2:
Align Live-to-spoof Transition

Live training Domain $e^{(1)}$ /$e^{(2)}$ /$e^{(3)}$ **Spoof** trainging domain $e^{(1)}$ /$e^{(2)}$ /$e^{(3)}$

Methodology

Set Up

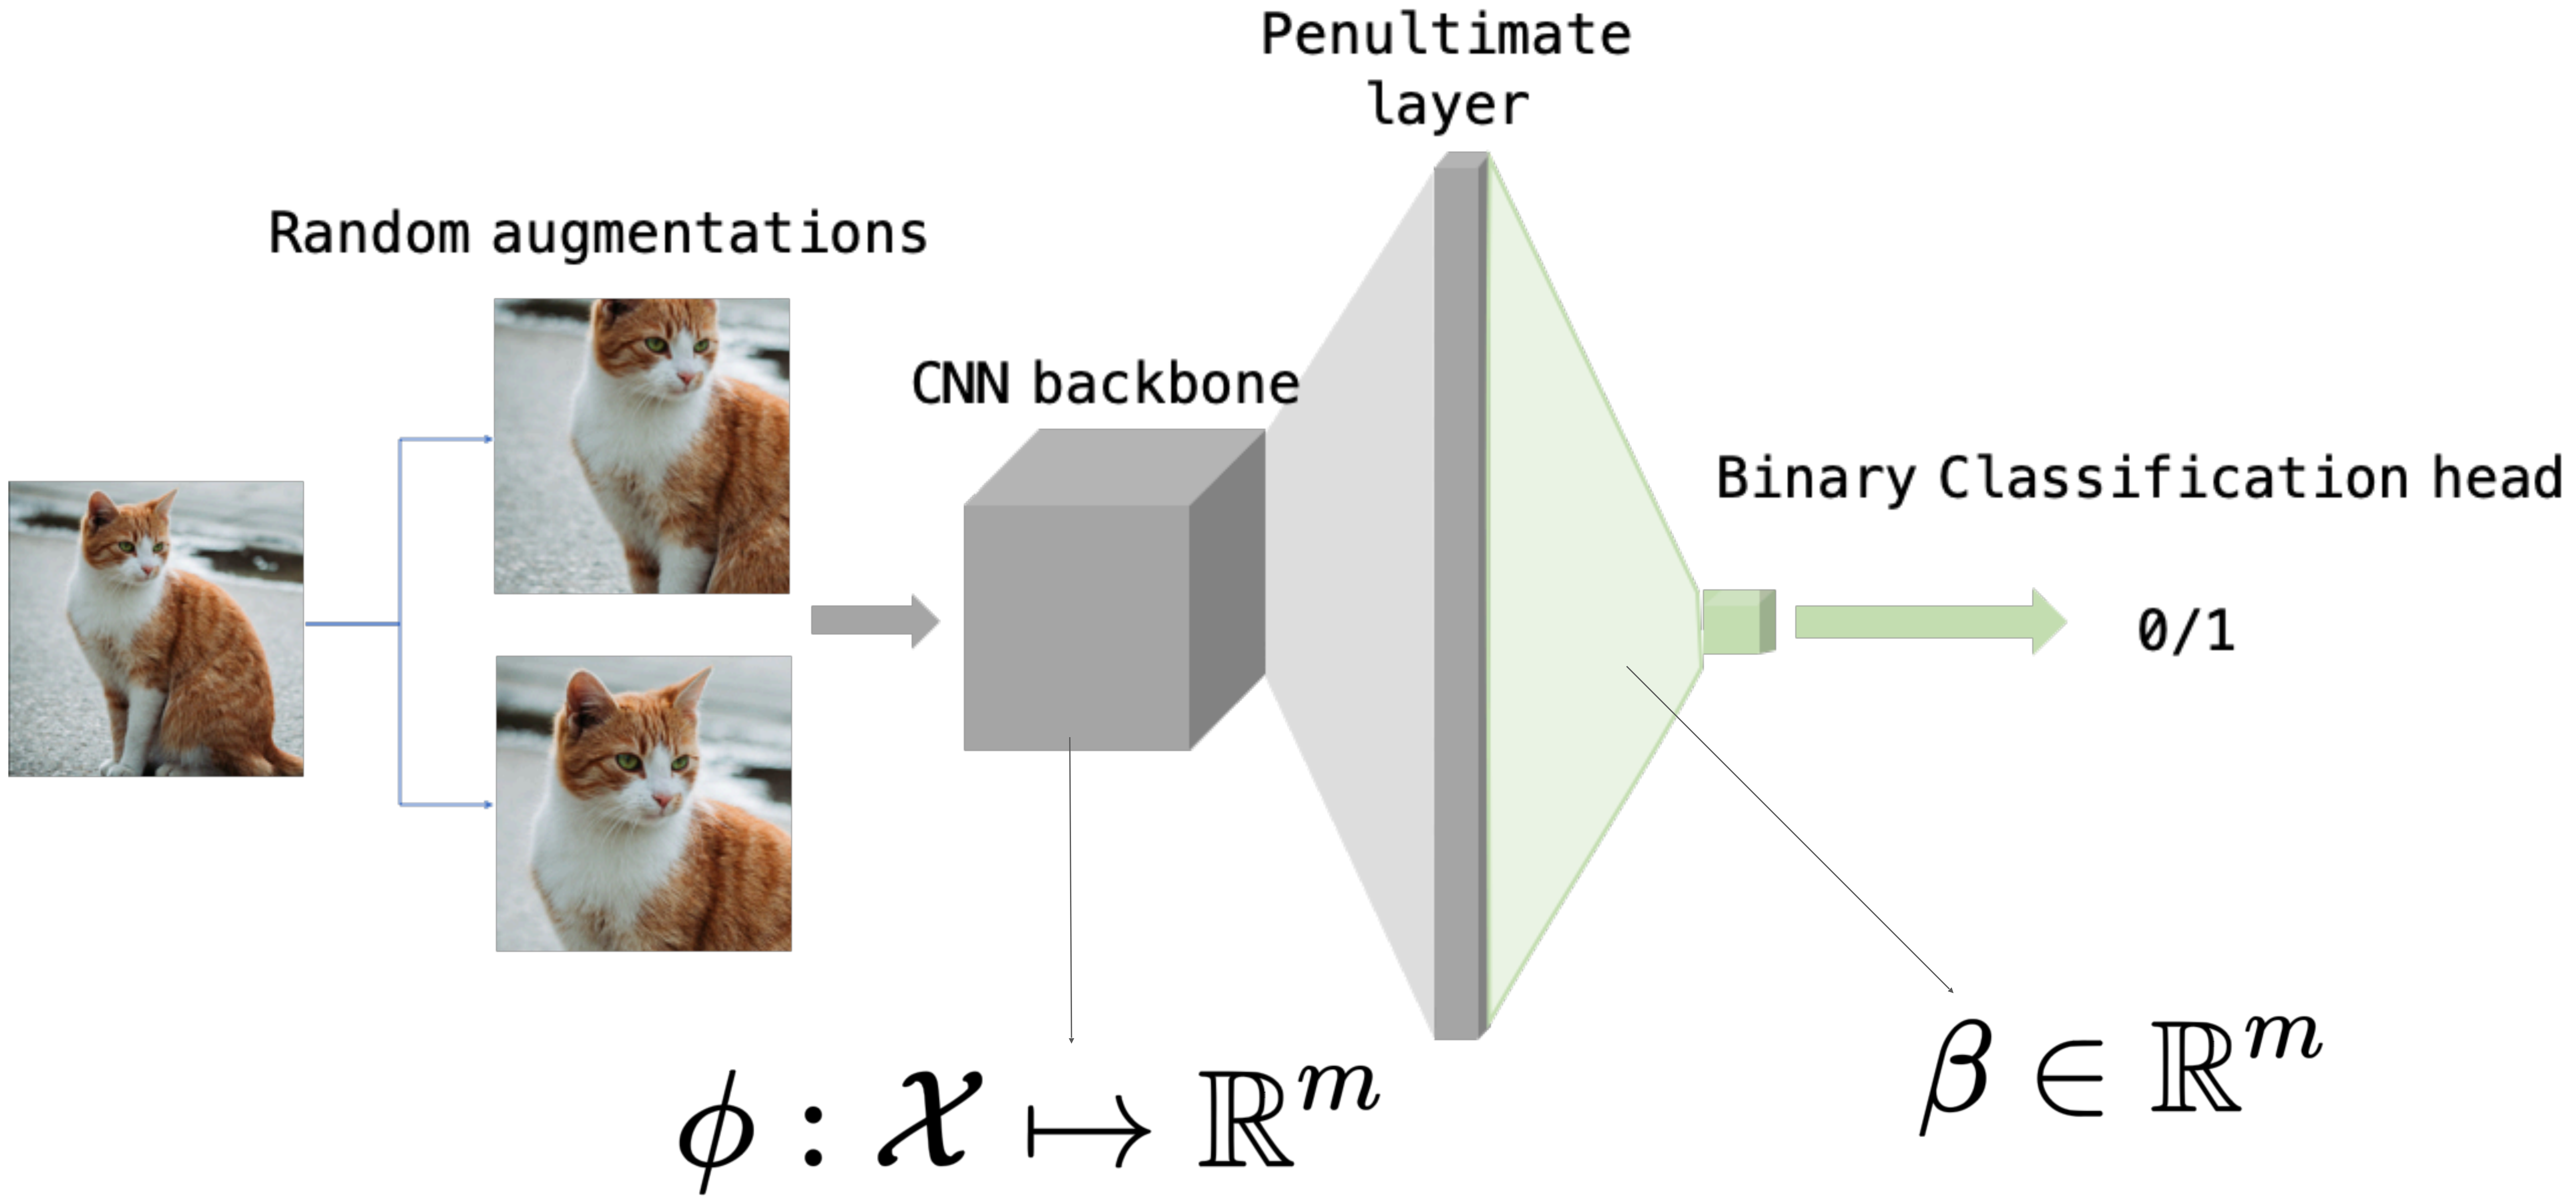
Penultimate layer
Random augmentations
CNN backbone
Binary Classification head
0/1
$\phi : \mathcal{X} \mapsto \mathbb{R}^m$
$\beta \in \mathbb{R}^m$

Challenge 1: Feature's Separability

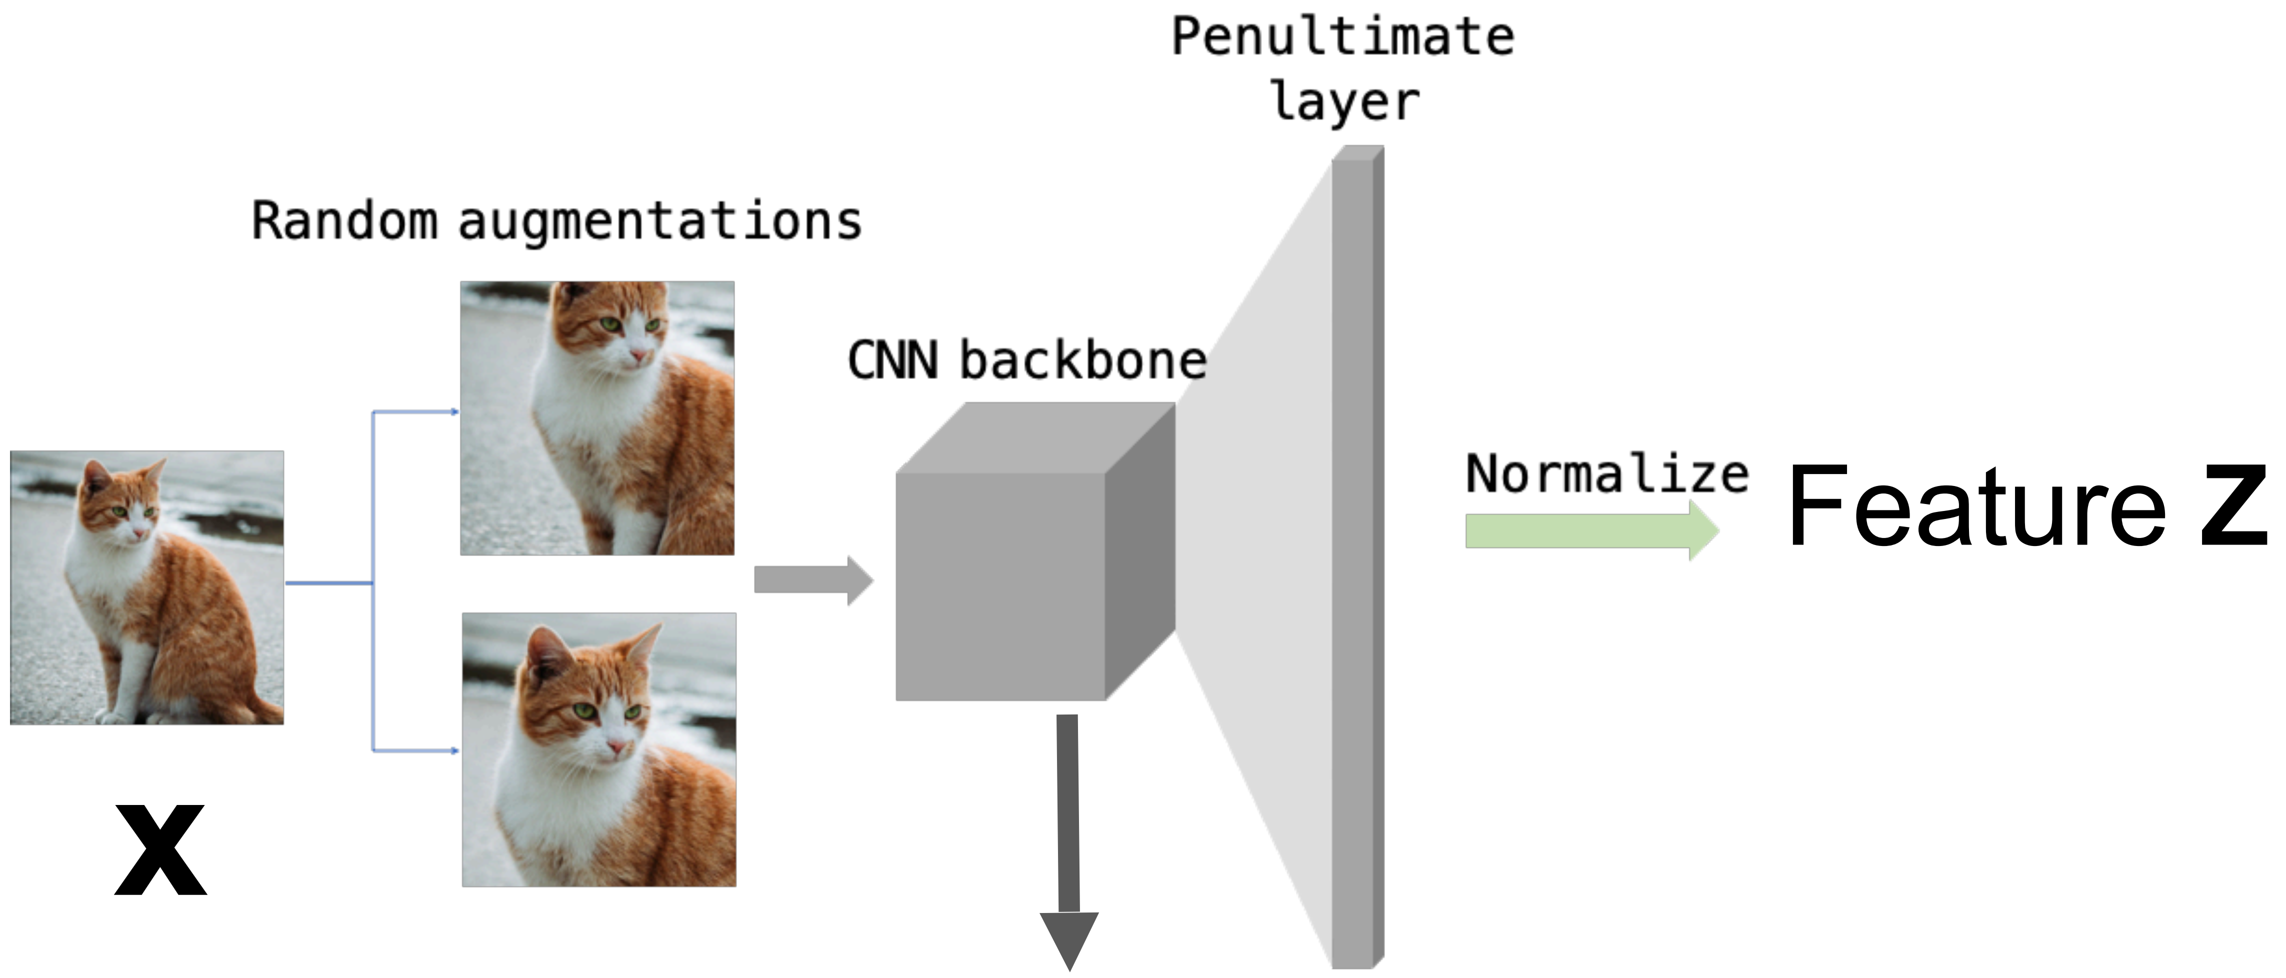

$$\phi : \mathcal{X} \mapsto \mathbb{R}^m$$

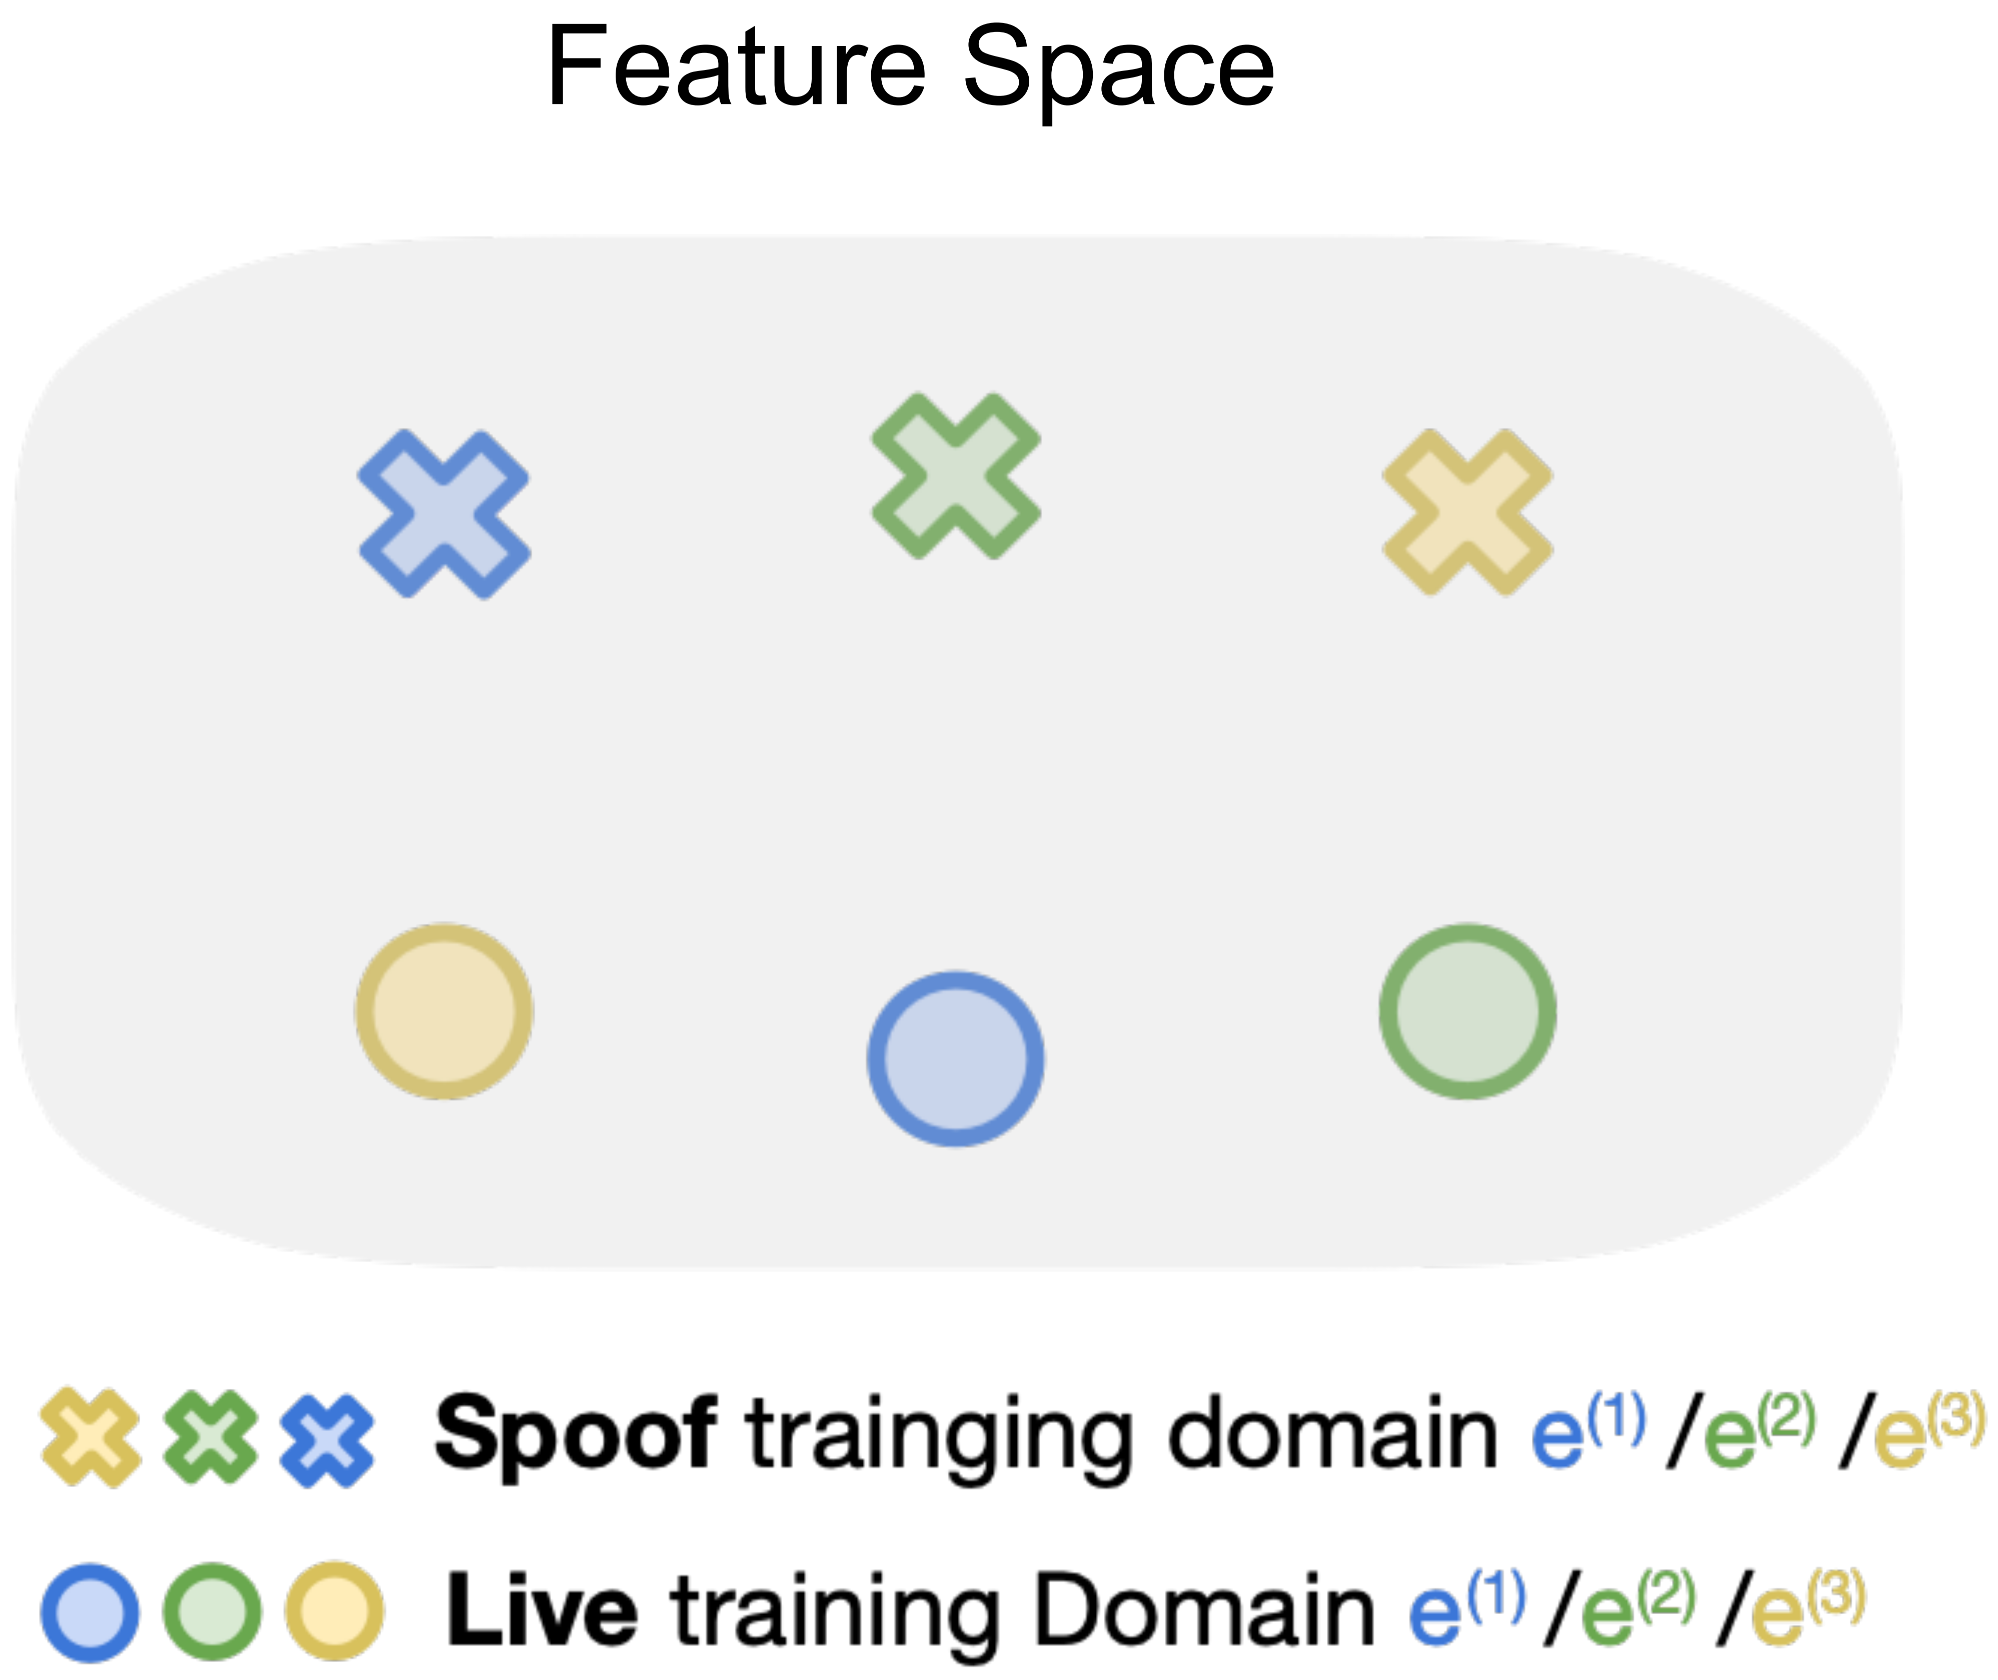

Supervised Contrastive Learning (SupCon)

$\mathcal{P}(\mathbf{x})$ contains features of samples with same live/spoof label and domain as x.

$\mathcal{N}(\mathbf{x})$ contains features of samples except x.

$$\mathcal{L}_\phi(\mathbf{x}; \tau, \mathcal{P}(\mathbf{x}), \mathcal{N}(\mathbf{x})) = -\frac{1}{|\mathcal{P}(\mathbf{x})|} \sum_{\mathbf{z}^+ \in \mathcal{P}(\mathbf{x})} \log \frac{\exp\left(\mathbf{z}^\top \cdot \mathbf{z}^+ / \tau\right)}{\sum_{\mathbf{z}^- \in \mathcal{N}(\mathbf{x})} \exp\left(\mathbf{z}^\top \cdot \mathbf{z}^- / \tau\right)}$$

Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C. and Krishnan, D. Supervised contrastive learning. *NeurIPS 2020*

Supervised Contrastive Learning (SupCon) Enhances the Separability

No SupCon

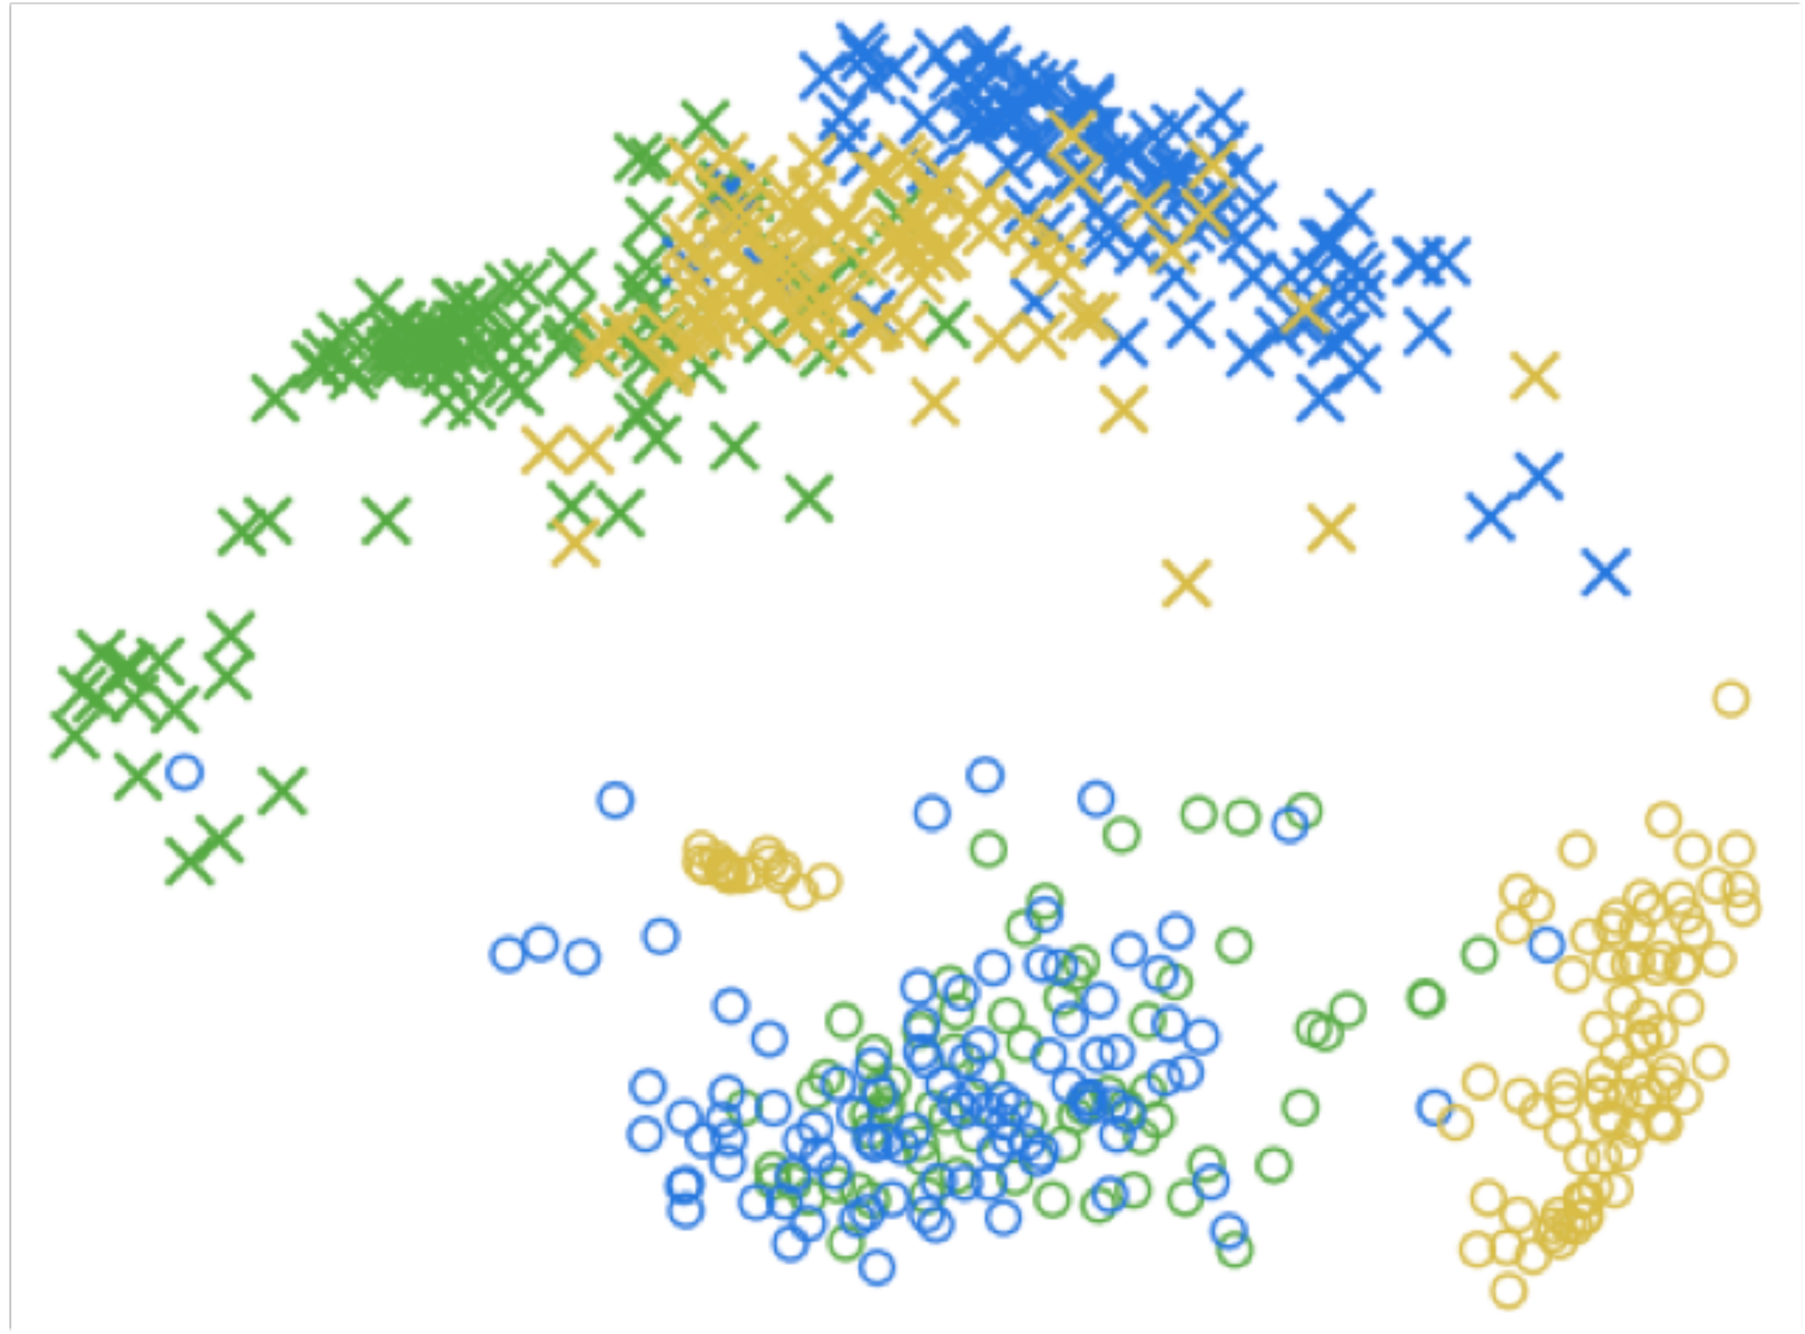

With SupCon
(More Separable between domains)

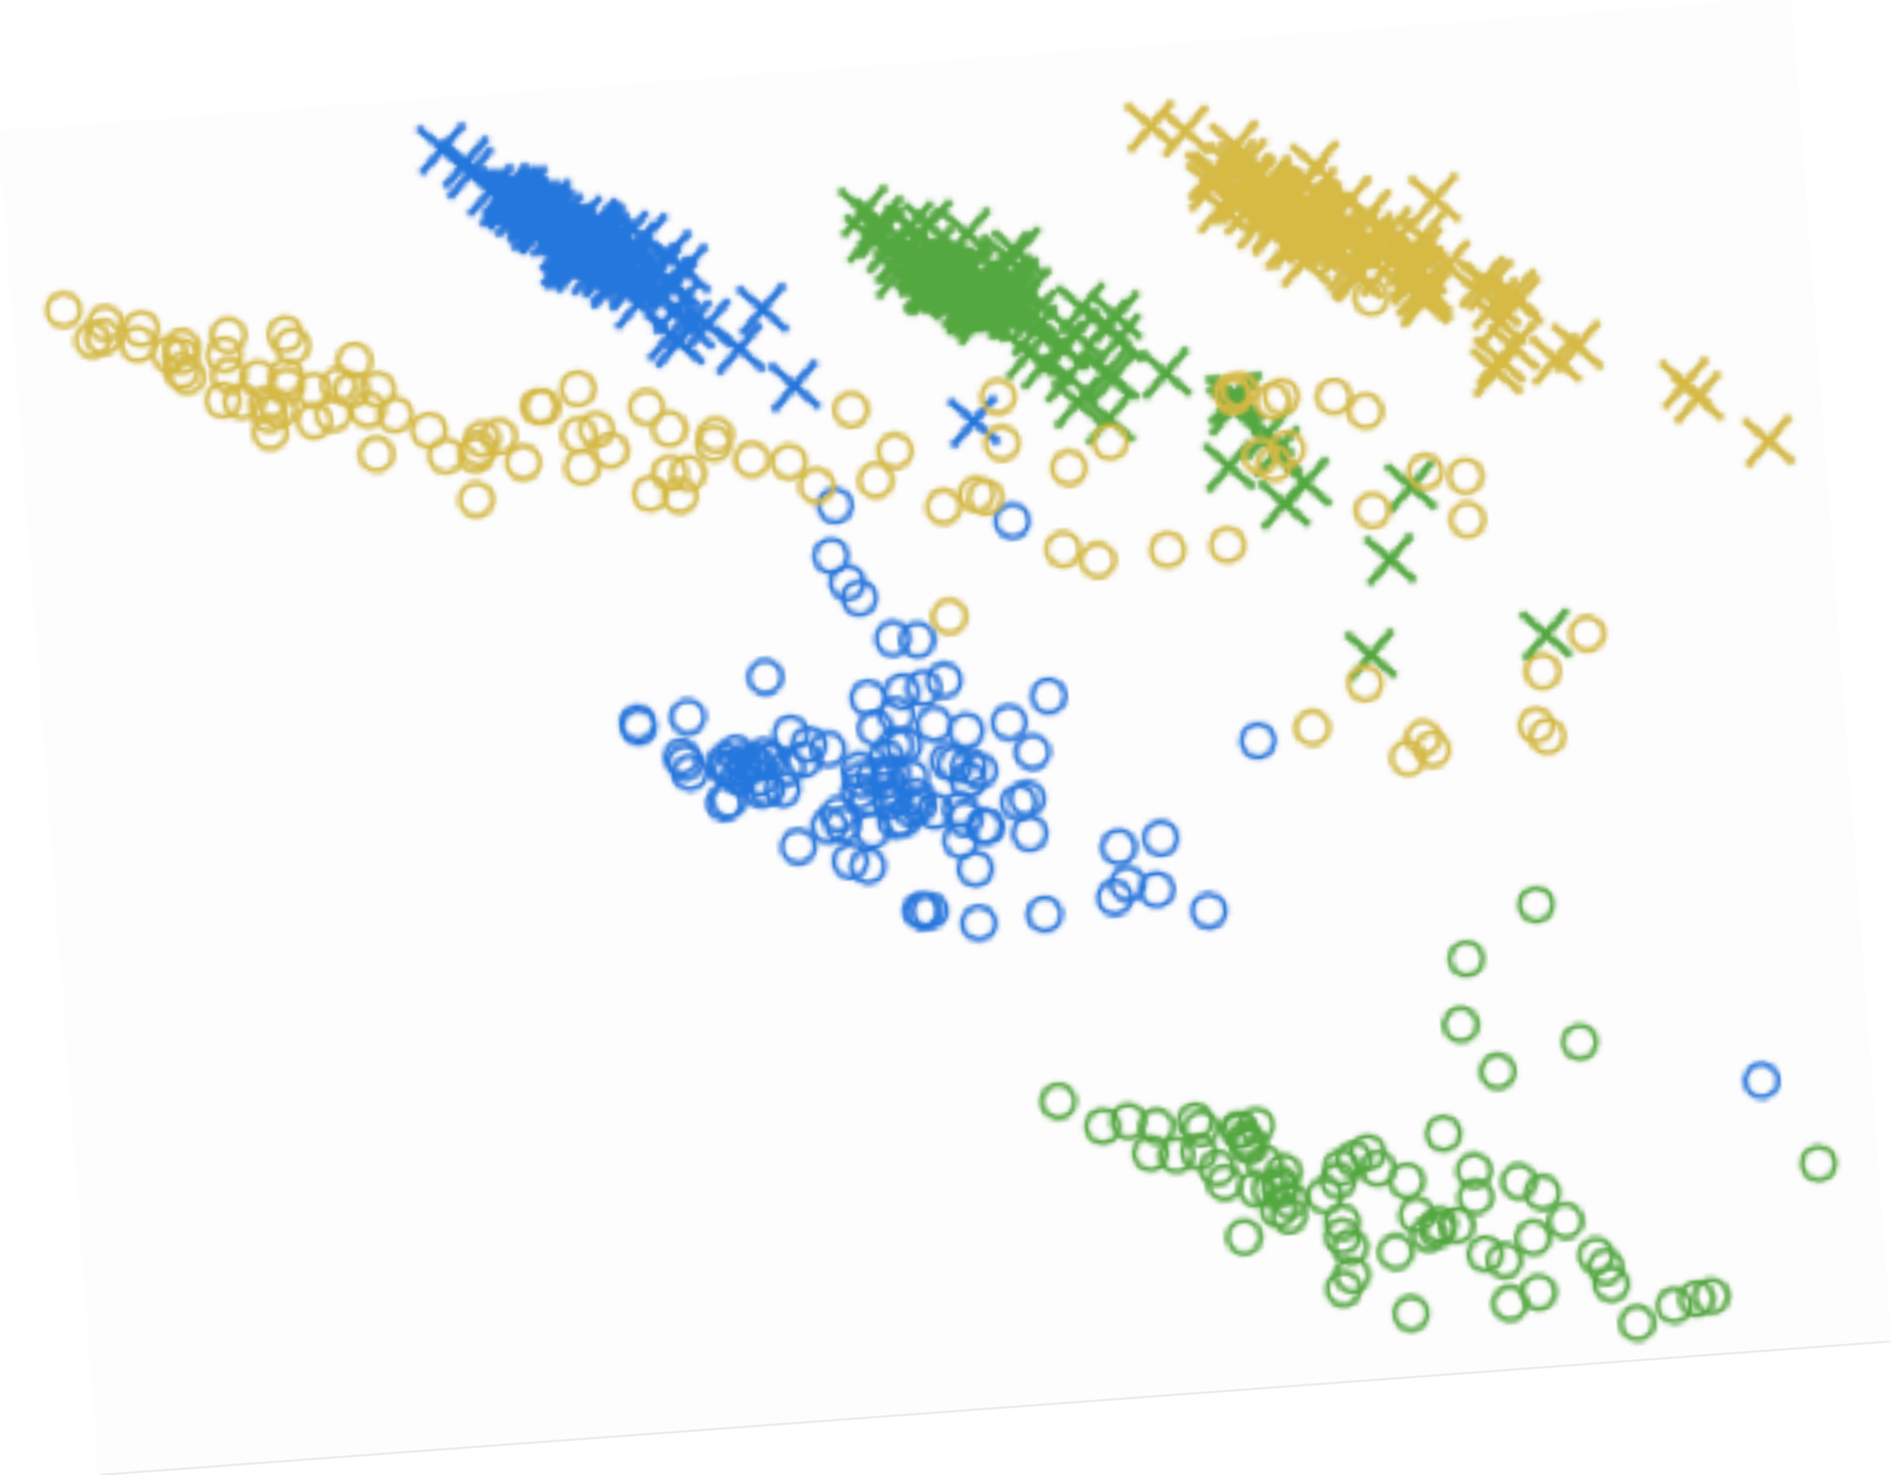

○○○ **Live** training Domain CASIA/MSU/REPLAY

X X X **Spoof** trainging domain CASIA/MSU/REPLAY

Challenge 2: Feature's Alignment

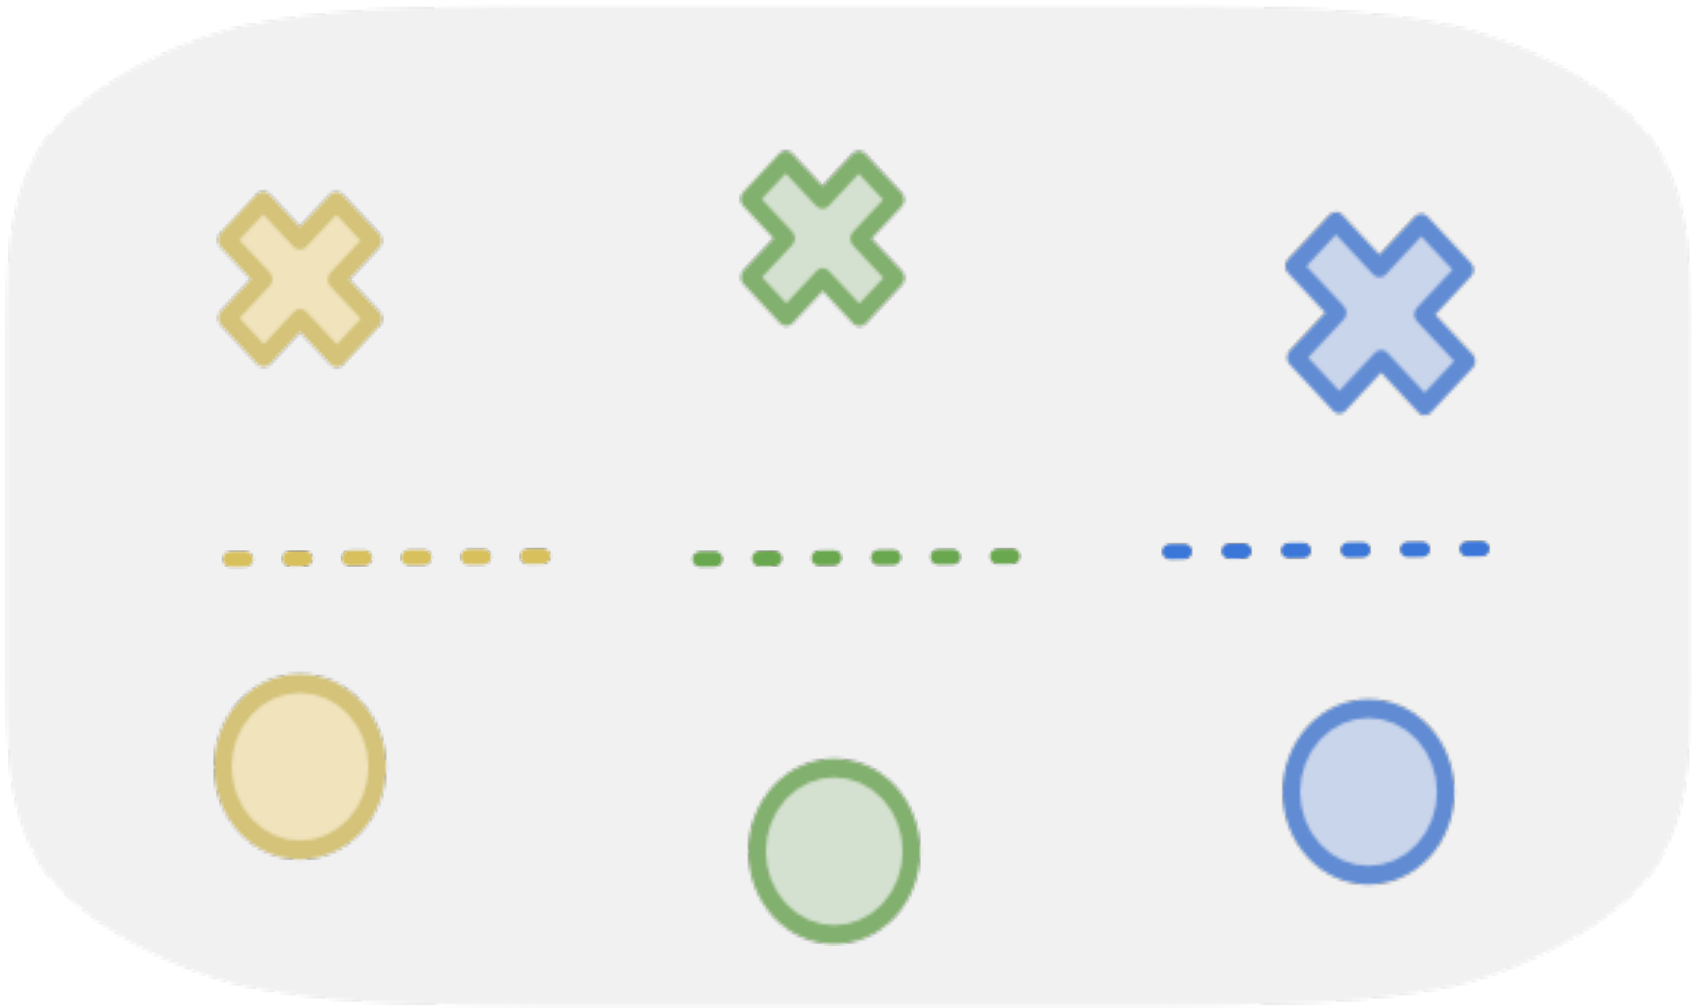

Challenge 2:
Align Live/spoof Direction

Live training Domain $e^{(1)}$ /$e^{(2)}$ /$e^{(3)}$ **Spoof** trainging domain $e^{(1)}$ /$e^{(2)}$ /$e^{(3)}$

Alignment Target: Invariant Risk Minimization (IRM)

Notations:

- Environments: $\mathcal{E} = \left\{ e^{(1)}, e^{(2)}, .., e^{E} \right\}$
- Training dataset: $\mathcal{D} = \{(\mathbf{x}_i, y_i, e_i)\}_{i=1}^{N}$
- Risk in Env. e: $\mathcal{R}^e(\phi, \beta) \doteq \mathbb{E}_{(\mathbf{x}_i, y_i, e_i = e) \sim \mathcal{D}} \ell(f(\mathbf{x}), y)$

Objective:

$$\min_{\phi, \beta^*} \frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} \mathcal{R}^e(\phi, \beta^*) \rightarrow \mathcal{L}_{IRM}$$

$$\text{s.t.} \quad \beta^* \in \arg\min_{\beta} \mathcal{R}^e(\phi, \beta) \quad \forall e \in \mathcal{E},$$

Alignment Target: Invariant Risk Minimization (IRM)

Notations:

- Environments: $\mathcal{E} = \left\{e^{(1)}, e^{(2)}, .., e^{E}\right\}$
- Training dataset: $\mathcal{D} = \{(\mathbf{x}_i, y_i, e_i)\}_{i=1}^{N}$
- Risk in Env. e: $\mathcal{R}^e(\phi, \beta) \doteq \mathbb{E}_{(\mathbf{x}_i, y_i, e_i = e) \sim \mathcal{D}} \ell(f(\mathbf{x}), y)$

Objective:

$$\boxed{\min_{\phi, \beta^*} \frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} \mathcal{R}^e(\phi, \beta^*) \rightarrow \text{ERM Loss}}$$

What if the constraint is unsatisfied?

$$\min_{\phi,\beta^*} \frac{1}{|\mathcal{E}|} \sum_{e\in\mathcal{E}} \mathcal{R}^e(\phi,\beta^*) \rightarrow \mathcal{L}_{IRM}$$

$$\text{s.t.} \quad \cancel{\beta^* \in \arg\min_{\beta} \mathcal{R}^e(\phi,\beta) \quad \forall e \in \mathcal{E},}$$

What if the constraint is unsatisfied?

$$\min_{\phi,\beta^*} \frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} \mathcal{R}^e(\phi, \beta^*) \rightarrow \mathcal{L}_{IRM}$$

$$\text{s.t.} \quad \cancel{\beta^* \in \arg\min_{\beta} \mathcal{R}^e(\phi, \beta) \quad \forall e \in \mathcal{E},}$$

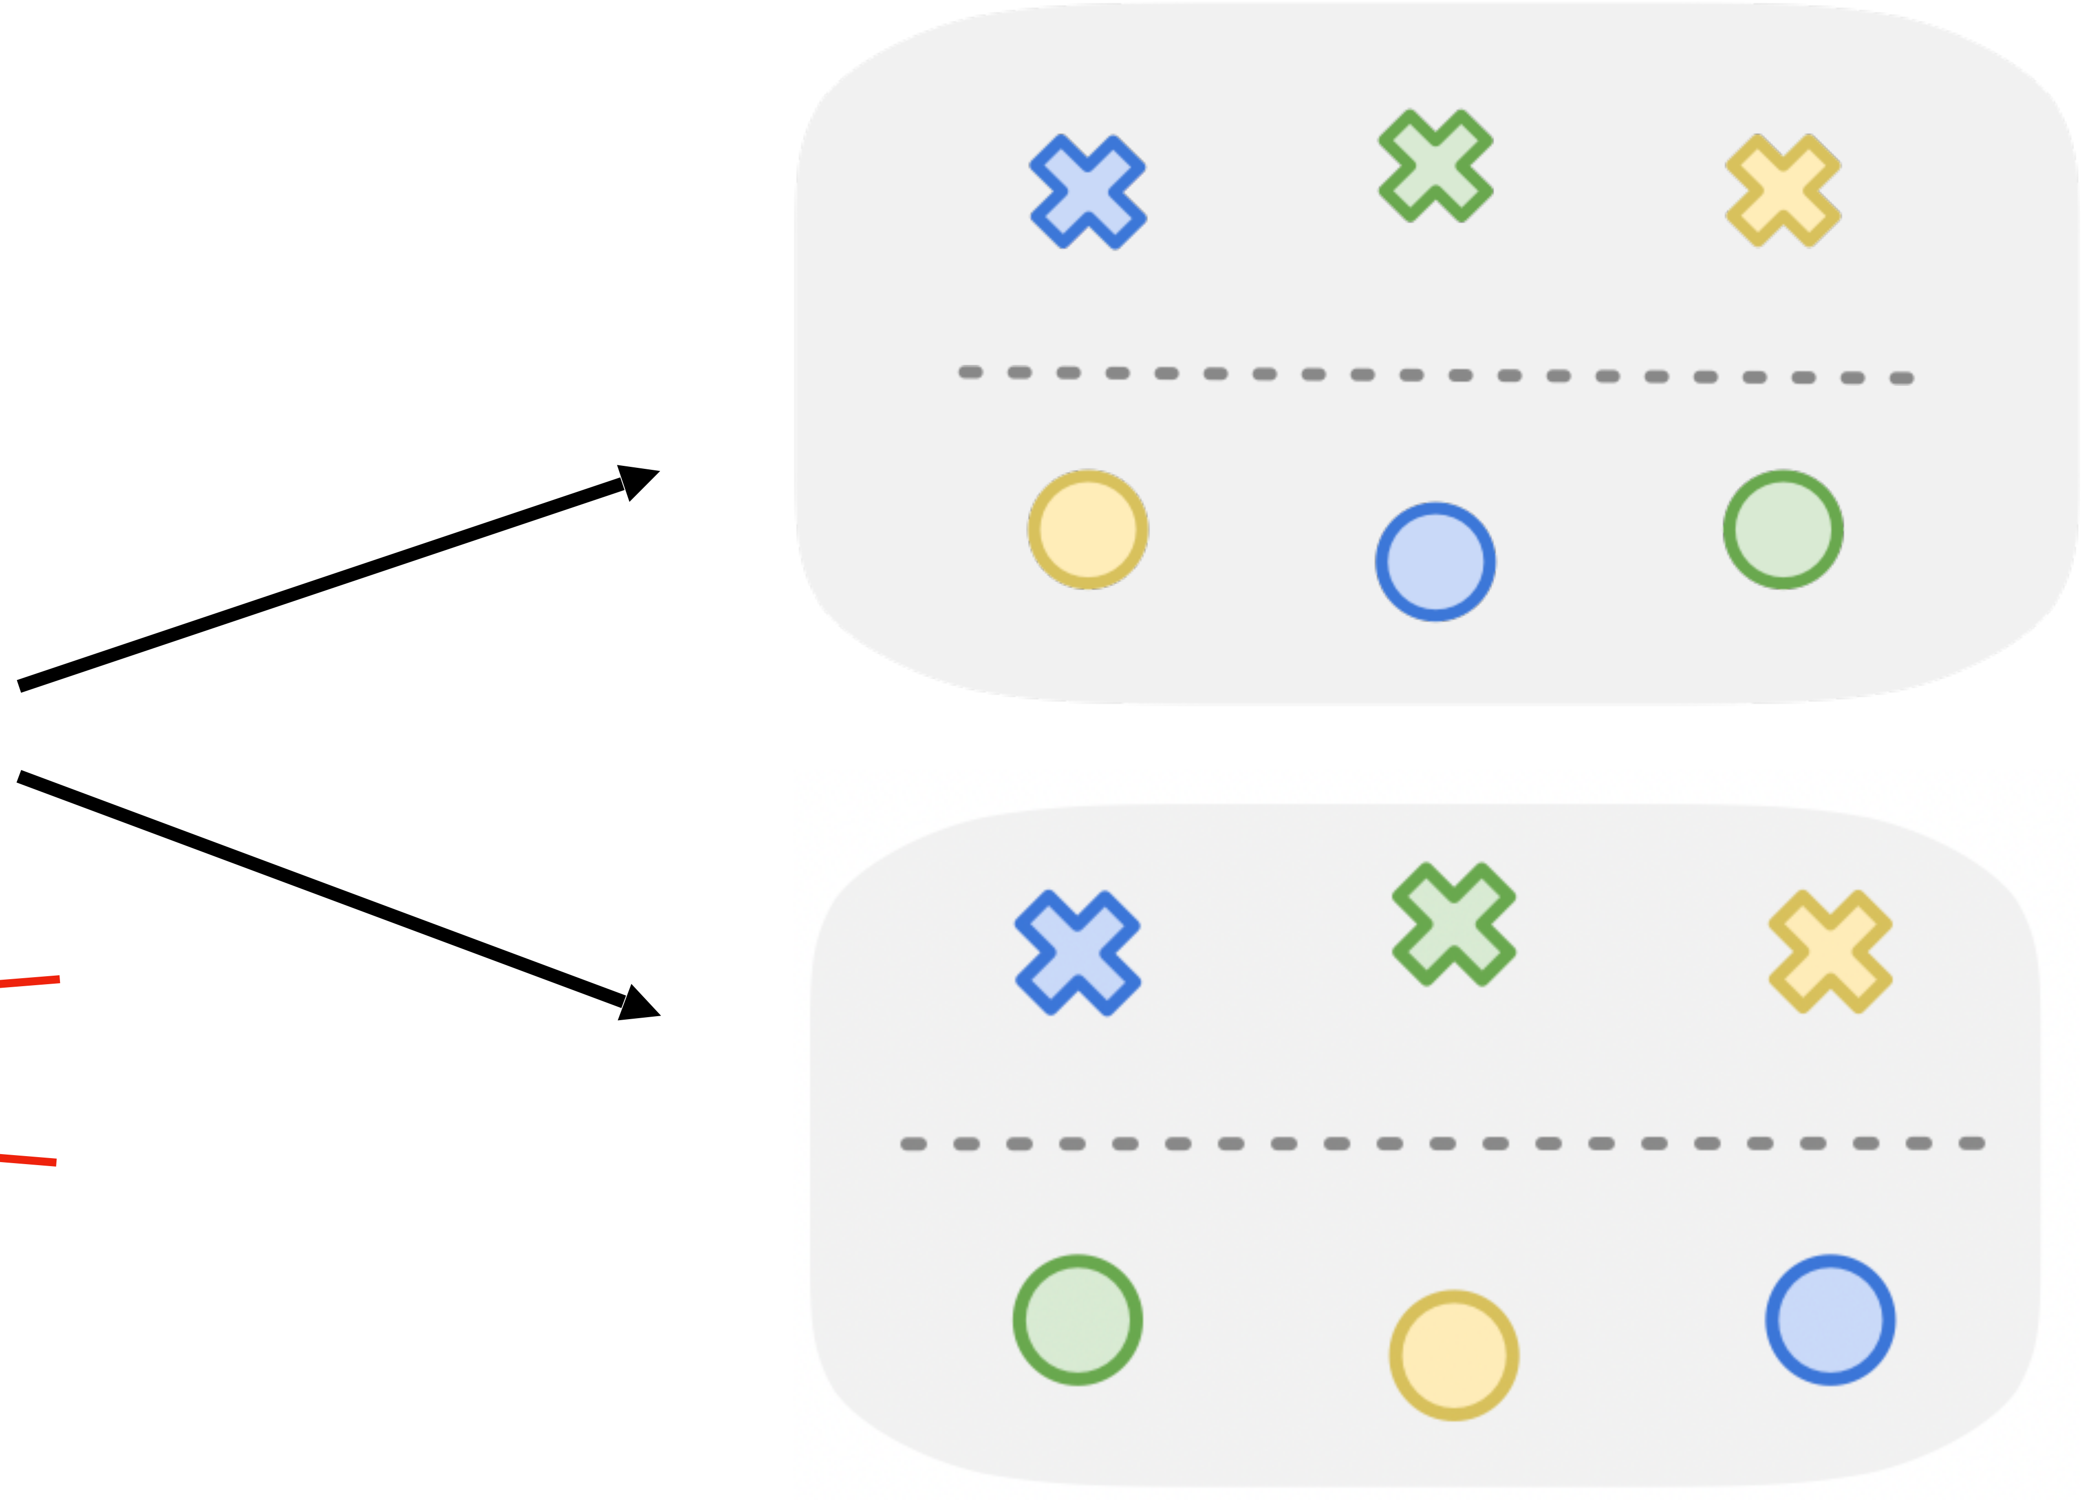

Result: The feature in different domains can be arbitrarily shuffled. Therefore, **no alignment!**

What if the constraint is satisfied?

$$\min_{\phi,\beta^*} \frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} \mathcal{R}^e(\phi, \beta^*) \rightarrow \mathcal{L}_{IRM}$$

$$\text{s.t.} \quad \beta^* \in \arg\min_{\beta} \mathcal{R}^e(\phi, \beta) \quad \forall e \in \mathcal{E},$$

$$\beta^* = \arg\min_{\beta} \mathcal{R}^{e^{(1)}}(\phi, \beta)$$

$$\beta^* = \arg\min_{\beta} \mathcal{R}^{e^{(2)}}(\phi, \beta)$$

$$\beta^* = \arg\min_{\beta} \mathcal{R}^{e^{(3)}}(\phi, \beta)$$

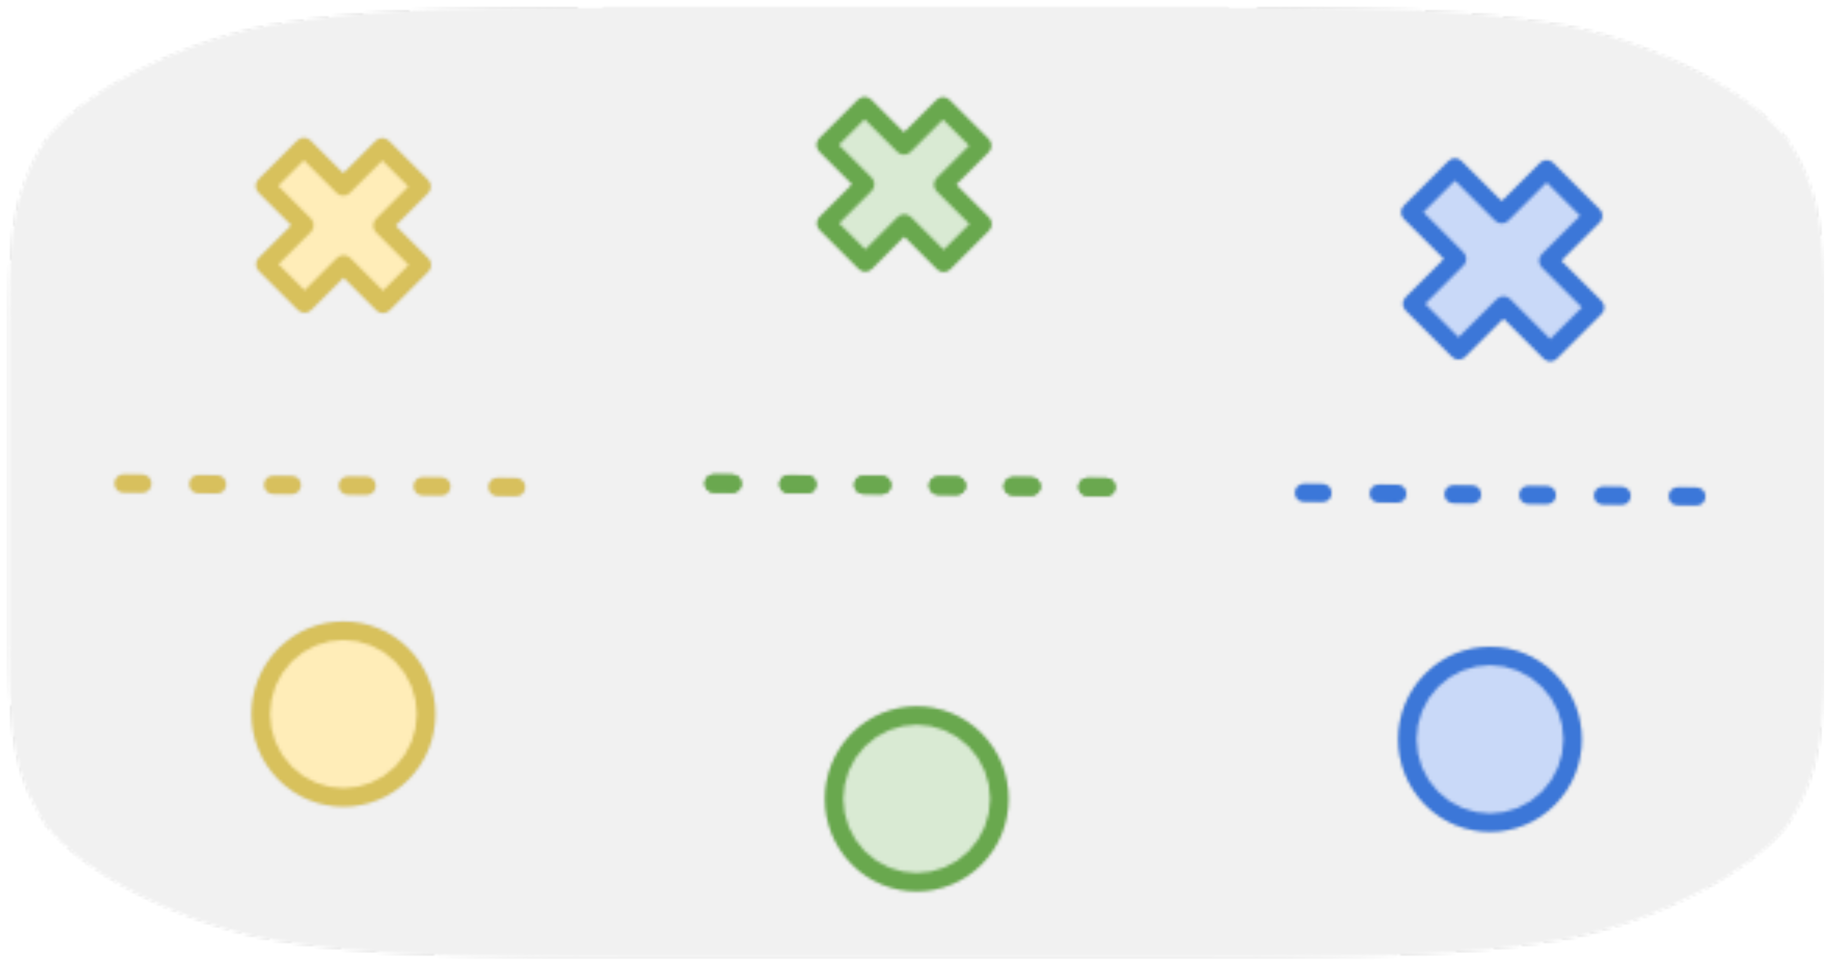

Result: Live-to-spoof transition is **aligned!**

Optimize Invariant Risk Minimization (IRM) by Projected Gradient Descent (PGD)

Theorem 1. ***(PG-IRM objective)*** *For all $\alpha \in (0,1)$, the IRM objective is equivalent to the following objective:*

$$\min_{\phi,\beta_{e(1)},\ldots,\beta_{e(E)}} \frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}} \mathcal{R}^e(\phi,\beta_e) \rightarrow \mathcal{L}_{align} \tag{5}$$
$$\textit{s.t. } \forall e \in \mathcal{E}, \exists \beta_e \in \Omega_e(\phi), \beta_e \in \Upsilon_\alpha(\beta_e),$$

*where the parametric constrained set for each environment is simplified as $\Omega_e(\phi) = \arg\min_{\beta} \mathcal{R}^e(\phi,\beta)$, and we define the α-**adjacency set**:*

$$\Upsilon_\alpha(\beta_e) = \{v | \max_{e'\in\mathcal{E}\setminus e} \min_{\beta_{e'}\in\Omega_{e'}(\phi)} \|v - \beta_{e'}\|_2 \tag{6}$$
$$\leq \alpha \max_{e'\in\mathcal{E}\setminus e} \min_{\beta_{e'}\in\Omega_{e'}(\phi)} \|\beta_e - \beta_{e'}\|_2\} \tag{7}$$

Algorithm 2 PG-IRM

Initialize $\phi, \beta_{e^{(1)}}, \ldots, \beta_{e^{(E)}}$, learning rate γ, alignment parameter α, alignment starting epoch T_a.
for t in 0, 1, ..., **do**
  Run forward pass and calculate the gradient.
  for $e \in \mathcal{E}$ **do**
    $\tilde{\beta}_e^{t+1} = \beta_e^t - \gamma \nabla_{\beta_e^t} \mathcal{L}_{PG\text{-}IRM}$
    $\alpha' := 1 - \mathbf{1}_{t>T_a}(1-\alpha)$
    select $\beta_{\bar{e}}^t$ with $\bar{e} = \underset{e'\in\mathcal{E}\setminus e}{\text{argmax}} \|\tilde{\beta}_e^{t+1} - \beta_{e'}^t\|_2$
    $\beta_e^{t+1} = \alpha' \tilde{\beta}_e^{t+1} + (1-\alpha')\beta_{\bar{e}}^t$
  end for
  Update $\phi^{t+1} = \phi^t - \gamma \nabla_{\phi^t} \mathcal{L}_{PG\text{-}IRM}$.
end for

More details are in the paper!

Experiment

Cross-Domain FAS Experiment Settings

Model

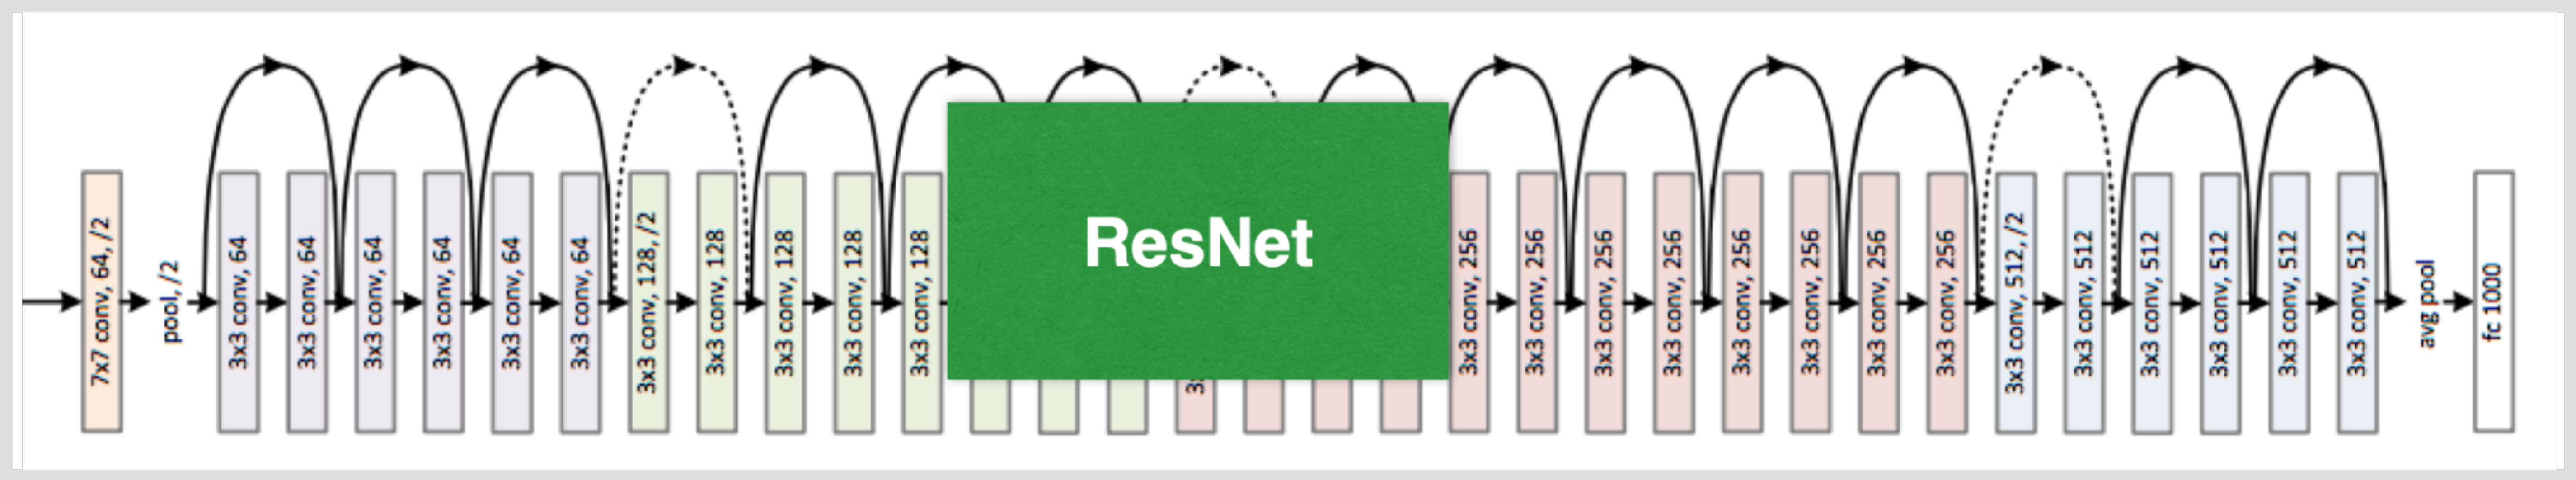

CASIA

OULU

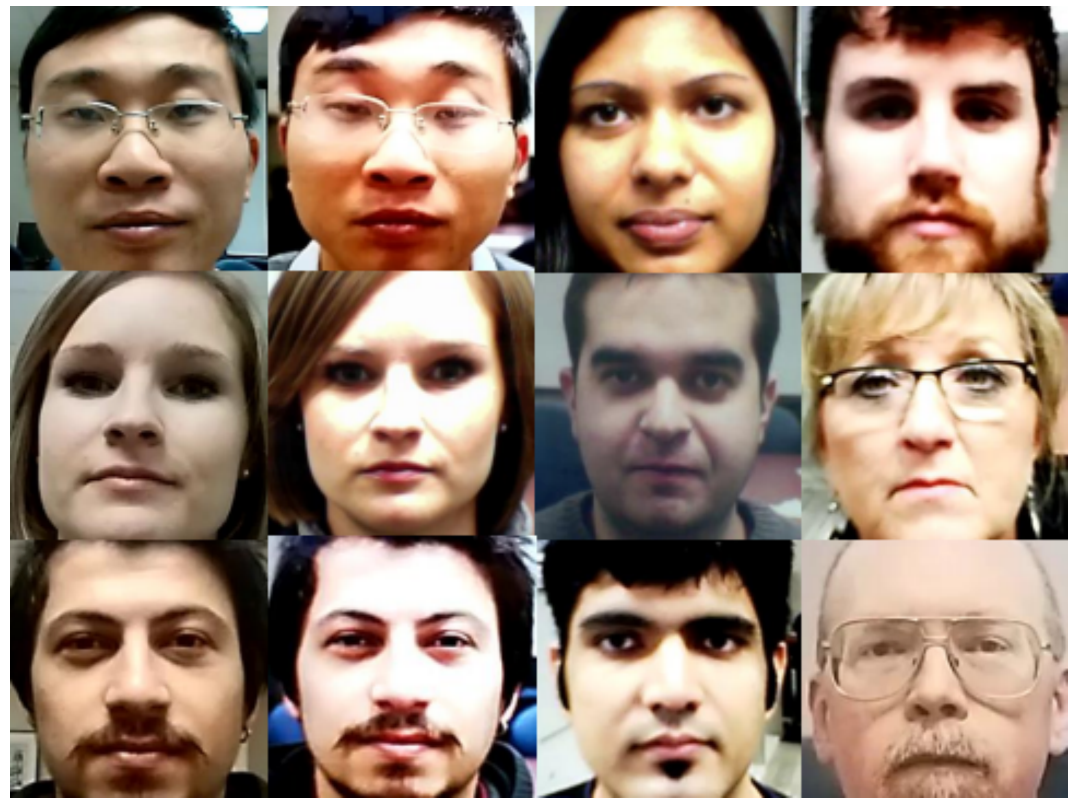

MSU

Replay

Leave-one-out protocol**:** Train on three, test on one**.**

Our Framework Establishes the Competitive Performance

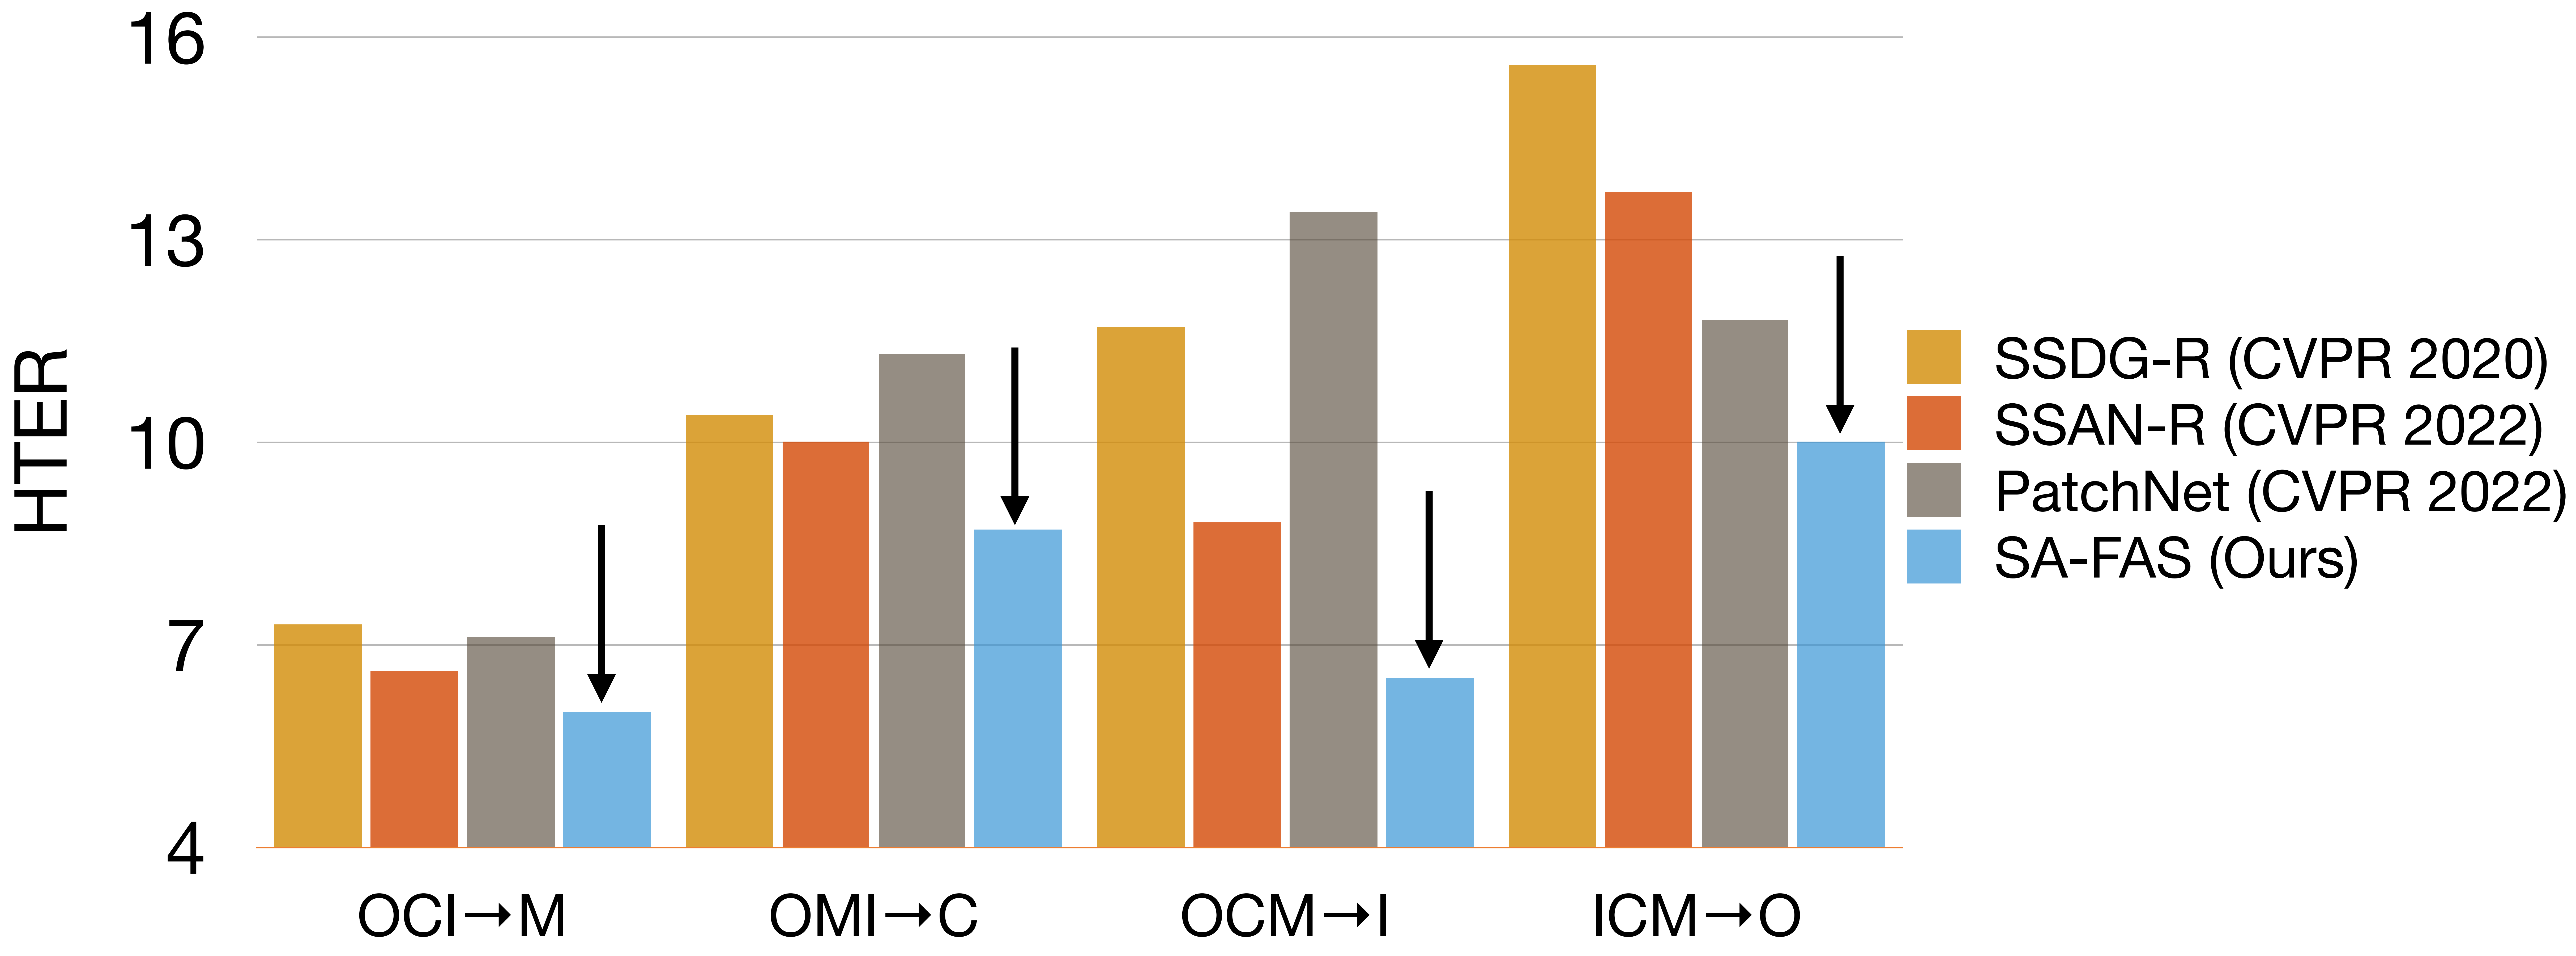

- Jia et al., "Single-side domain generalization for face anti-spoofing," CVPR 2020.
- Wang et al., "Domain generalization via shuffled style assembly for face anti-spoofing," CVPR 2022.
- Wang et al., "PatchNet: A simple face anti-spoofing framework via fine-grained patch recognition," CVPR 2022.

Comparison with Baseline Methods (Upon Convergence)

- Jia et al., "Single-side domain generalization for face anti-spoofing," CVPR 2020.
- Wang et al., "Domain generalization via shuffled style assembly for face anti-spoofing," CVPR 2022.
- Wang et al., "PatchNet: A simple face anti-spoofing framework via fine-grained patch recognition," CVPR 2022.

Visualization on Feature Space

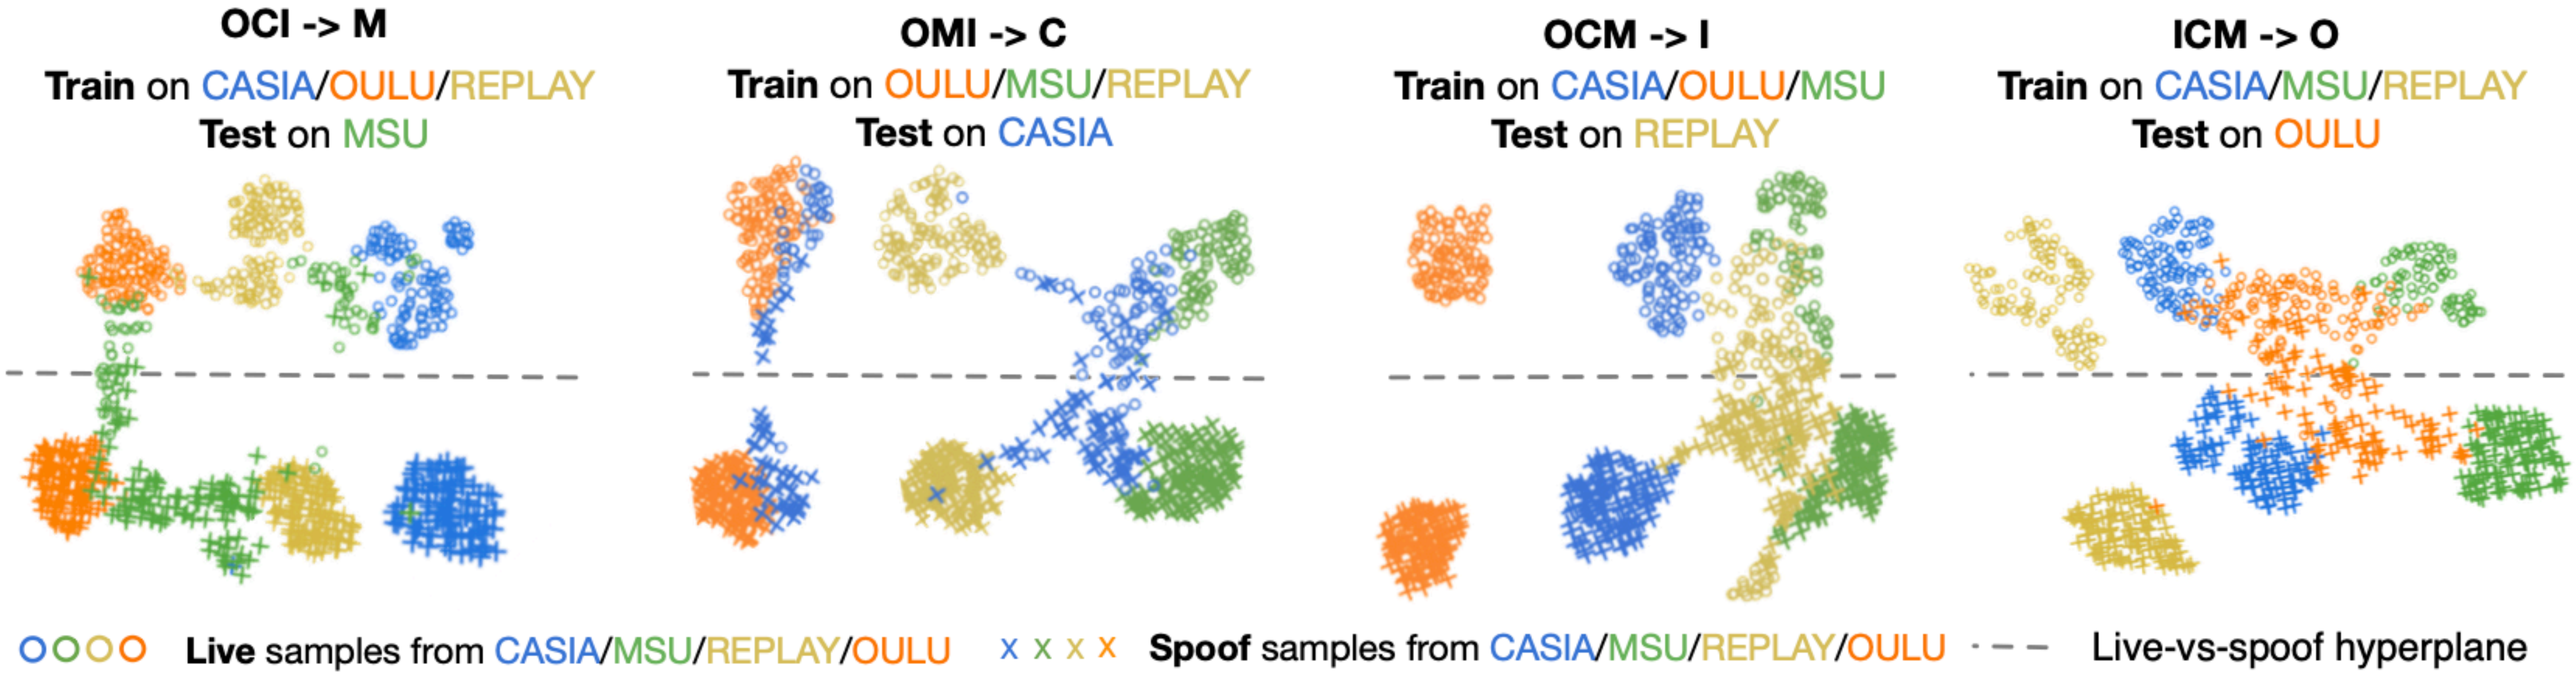

Summary

1. We offer a new perspective for cross-domain FAS by designing the feature space based on **separability** and **alignment**.

2. We first exploit the domain-variant representation learning by combining contrastive learning and optimizing invariant risk minimization (IRM) through the projected gradient algorithm for cross-domain FAS.

Thank you!

Code available at https://github.com/sunyiyou/SAFAS.

Google Research