Target-Aware Generative Augmentations for Single-Shot Adaptation

Kowshik Thopalli
Lawrence Livermore National Laboratory


Rakshith Subramanyam
ASU


Pavan Turaga
ASU


Jay Thiagarajan
LLNL

Predictive Models Tend To Break When There Is Shift Between Source And Target Distributions

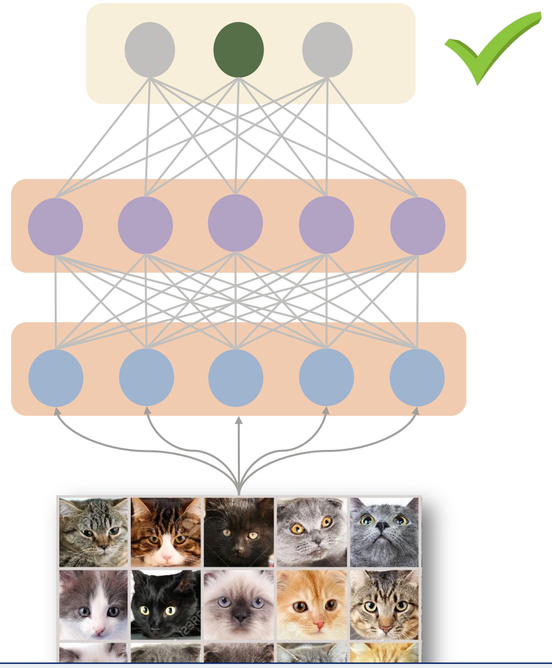

In conventional supervised learning we assume
$p_s(x) = p_t(x)$

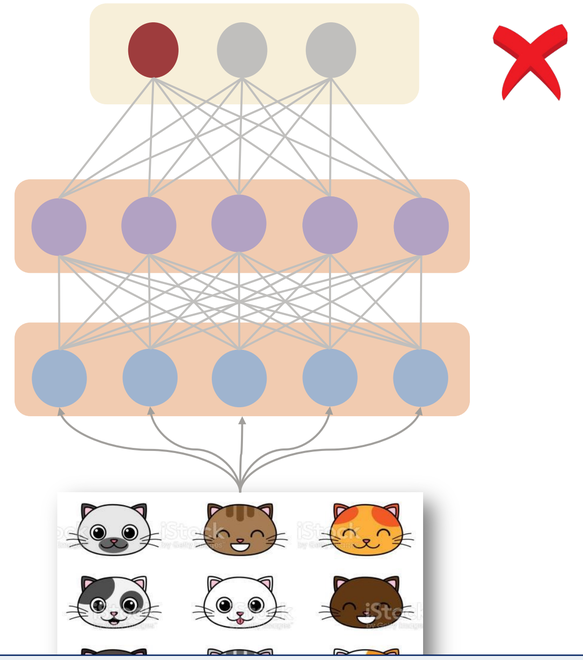

In practice $p_s(x) \neq p_t(x)$

A Popular Solution: Just fine-tune the model with additional data from the target domain.
But wait.. What if We do not have Enough Data?

Source distribution $p_s(x)$

Target distribution $p_t(x)$

Synthetic Data Augmentations Are Known to be Insufficient for Real-World Distribution Shifts

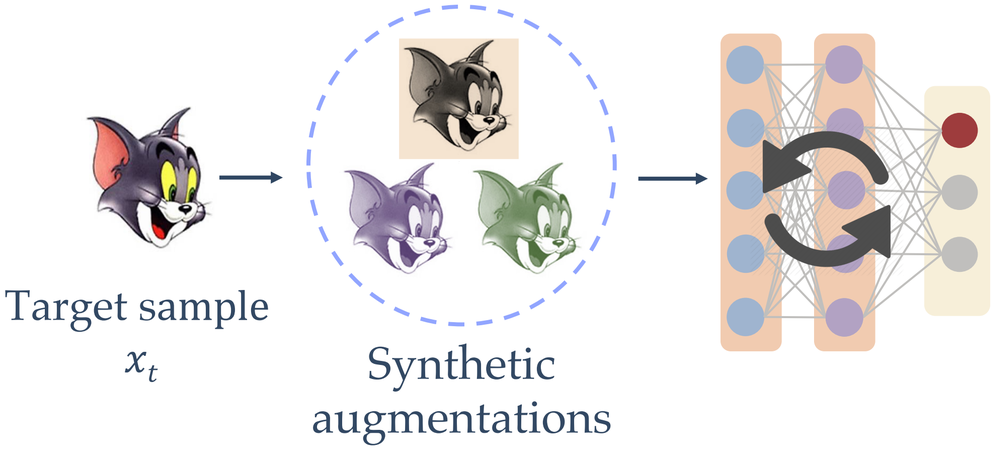

MEMO (2022)
Test Time Robustness via Adaptation and Augmentation

CATTAN (2022):
Geometric Alignment Improves Test-time Adaptation

How Should We Use Pre-Trained Generative Models to Fine-Tune Classifiers with Extremely Limited Data?

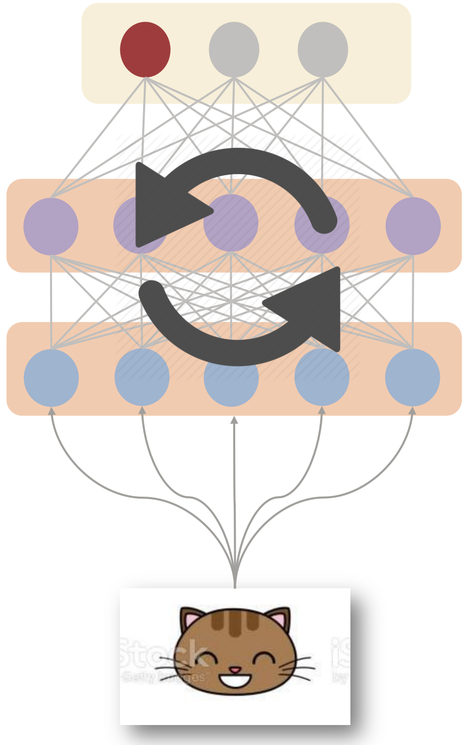

Single-Shot Target sample

x_t

OUR SOLUTION:

Leverage pre-trained generative models to characterize distribution shifts and devise novel sampling strategies to create synthetic target data

Latent space of Generative Model

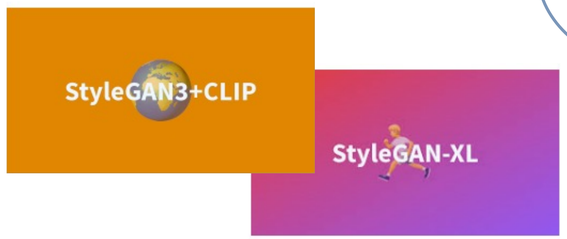

SiSTA: Single-Shot Target-Aware Generative Augmentation For Adapting Source Classifiers

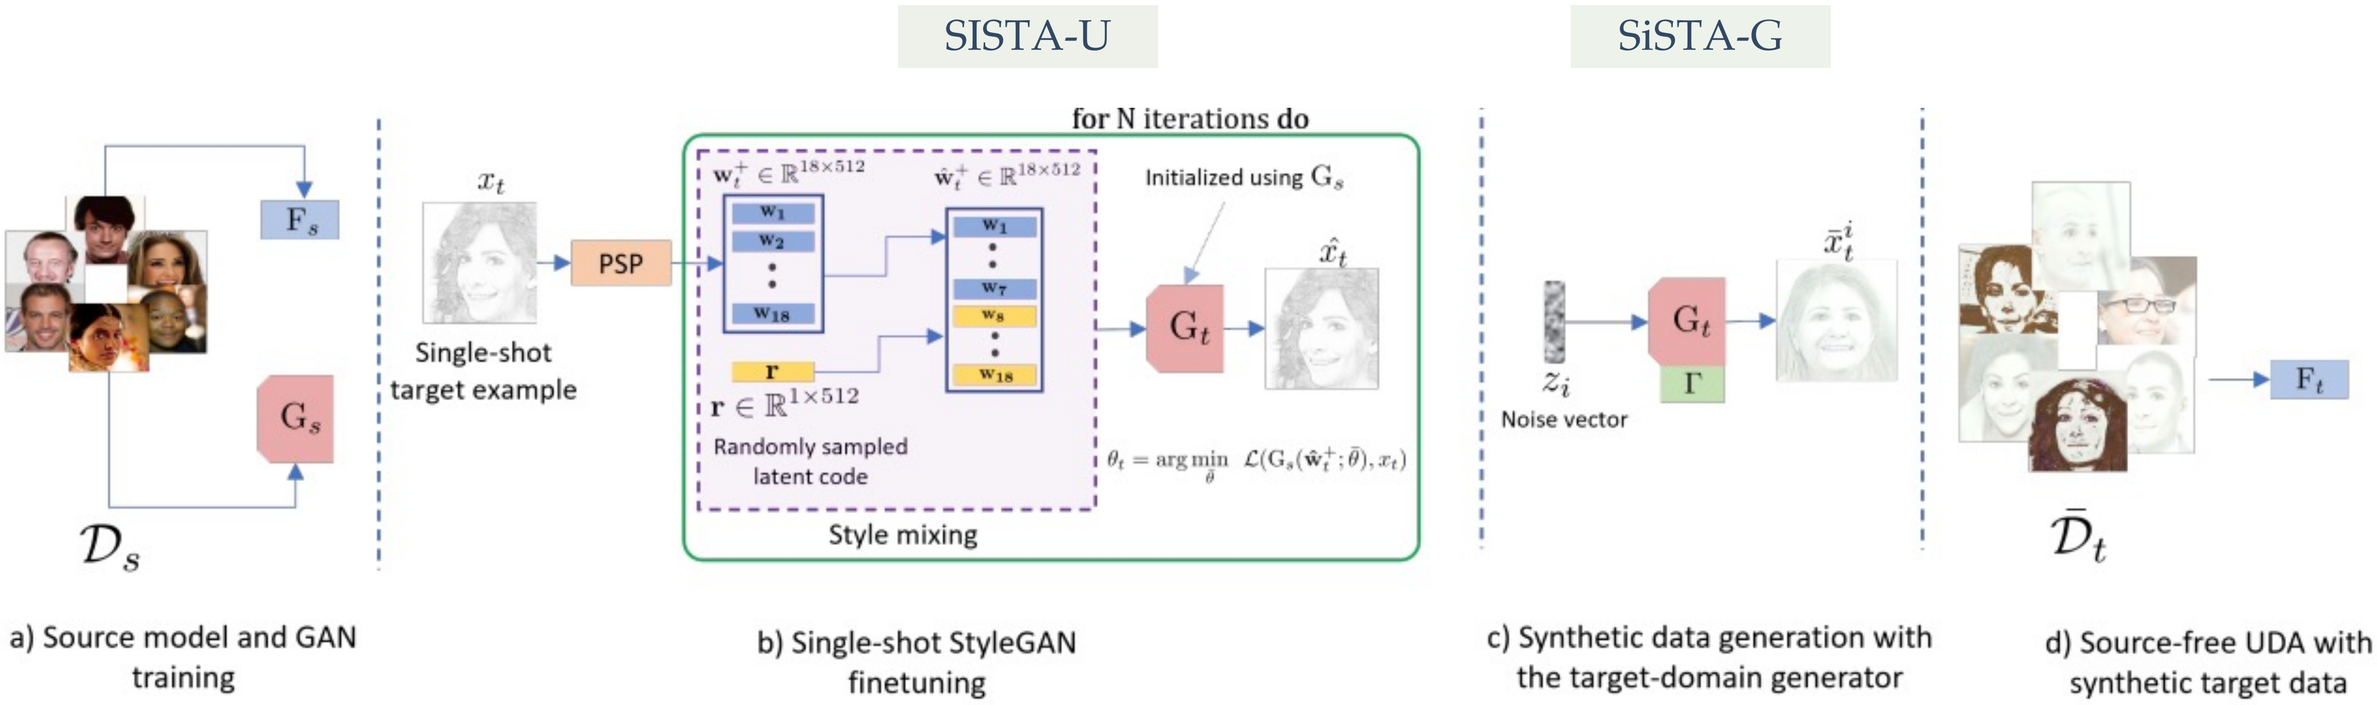

a) Source model and GAN training

b) Single-shot StyleGAN finetuning

c) Synthetic data generation with the target-domain generator

d) Source-free UDA with synthetic target data

Target GAN parameters

Source generator

Latent style code

$$\Theta_t = \arg\min_{\bar{\Theta}} \sum_{\ell} \| H_s^{\ell}(G_s(\mathbf{w}_t^+;\bar{\Theta})) - H_s^{\ell}(x_t) \|$$

Activations of layer ℓ of the source discriminator H_s

Target image

Two Sampling Strategies To Create Intermediate Representations

SiSTA Prune-Zero:
Zero-out activations in a style layer which are below a threshold

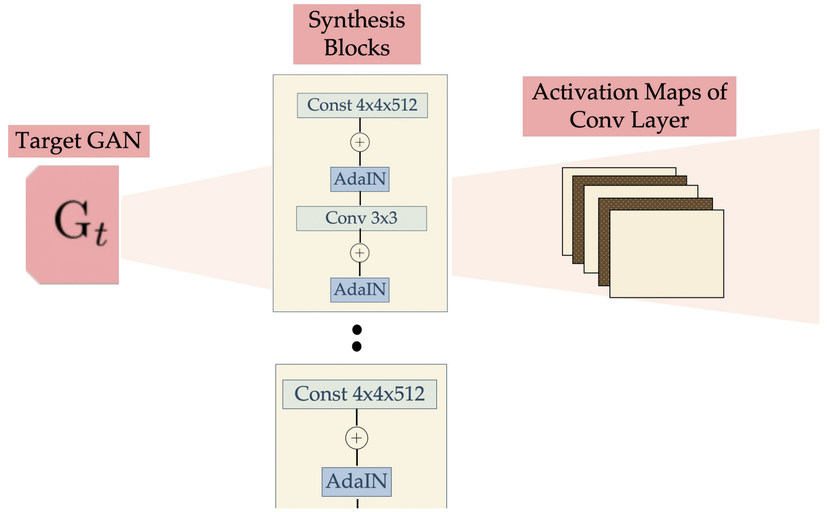

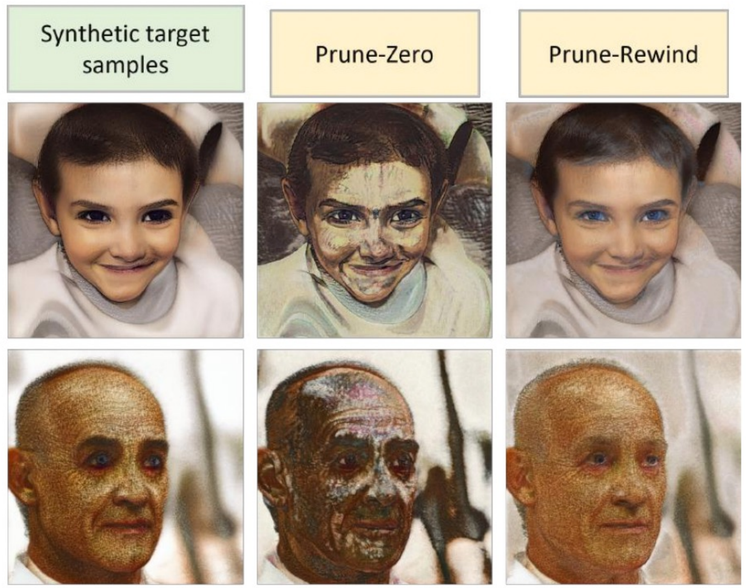

SiSTA Prune-Rewind:
Replace activations below a threshold in style layers with corresponding activations from source GAN

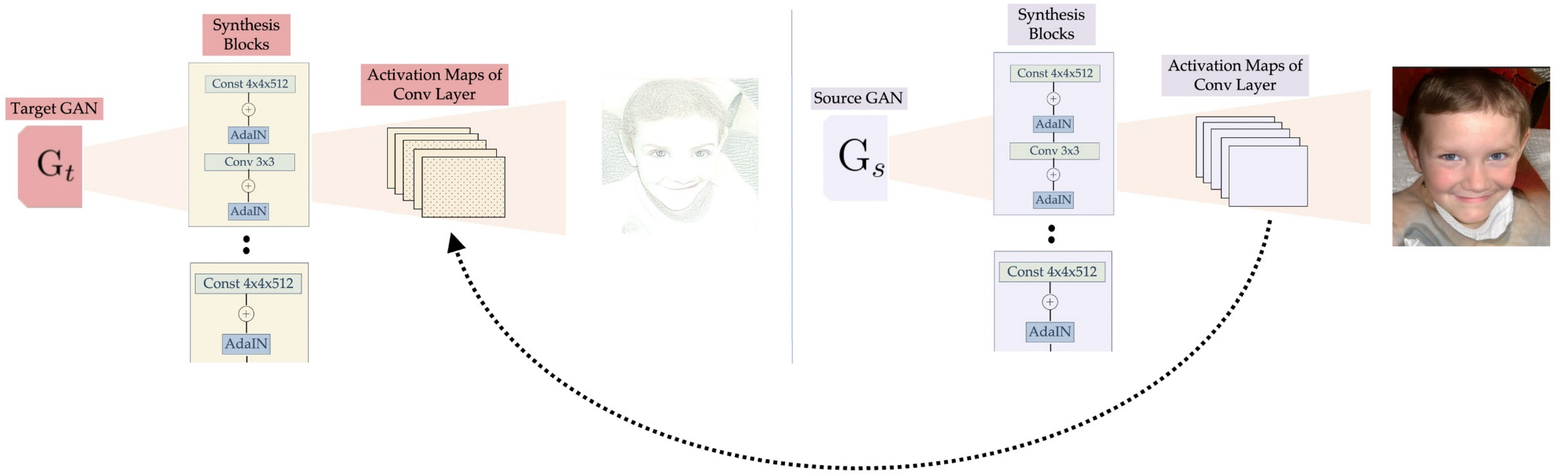

Attribute Detection from Face Images – A Fine-Grained Detection Task that is Challenging to Generalize Even Under Simple Shifts

Water-Color

Color-Sketch

Pencil-Sketch

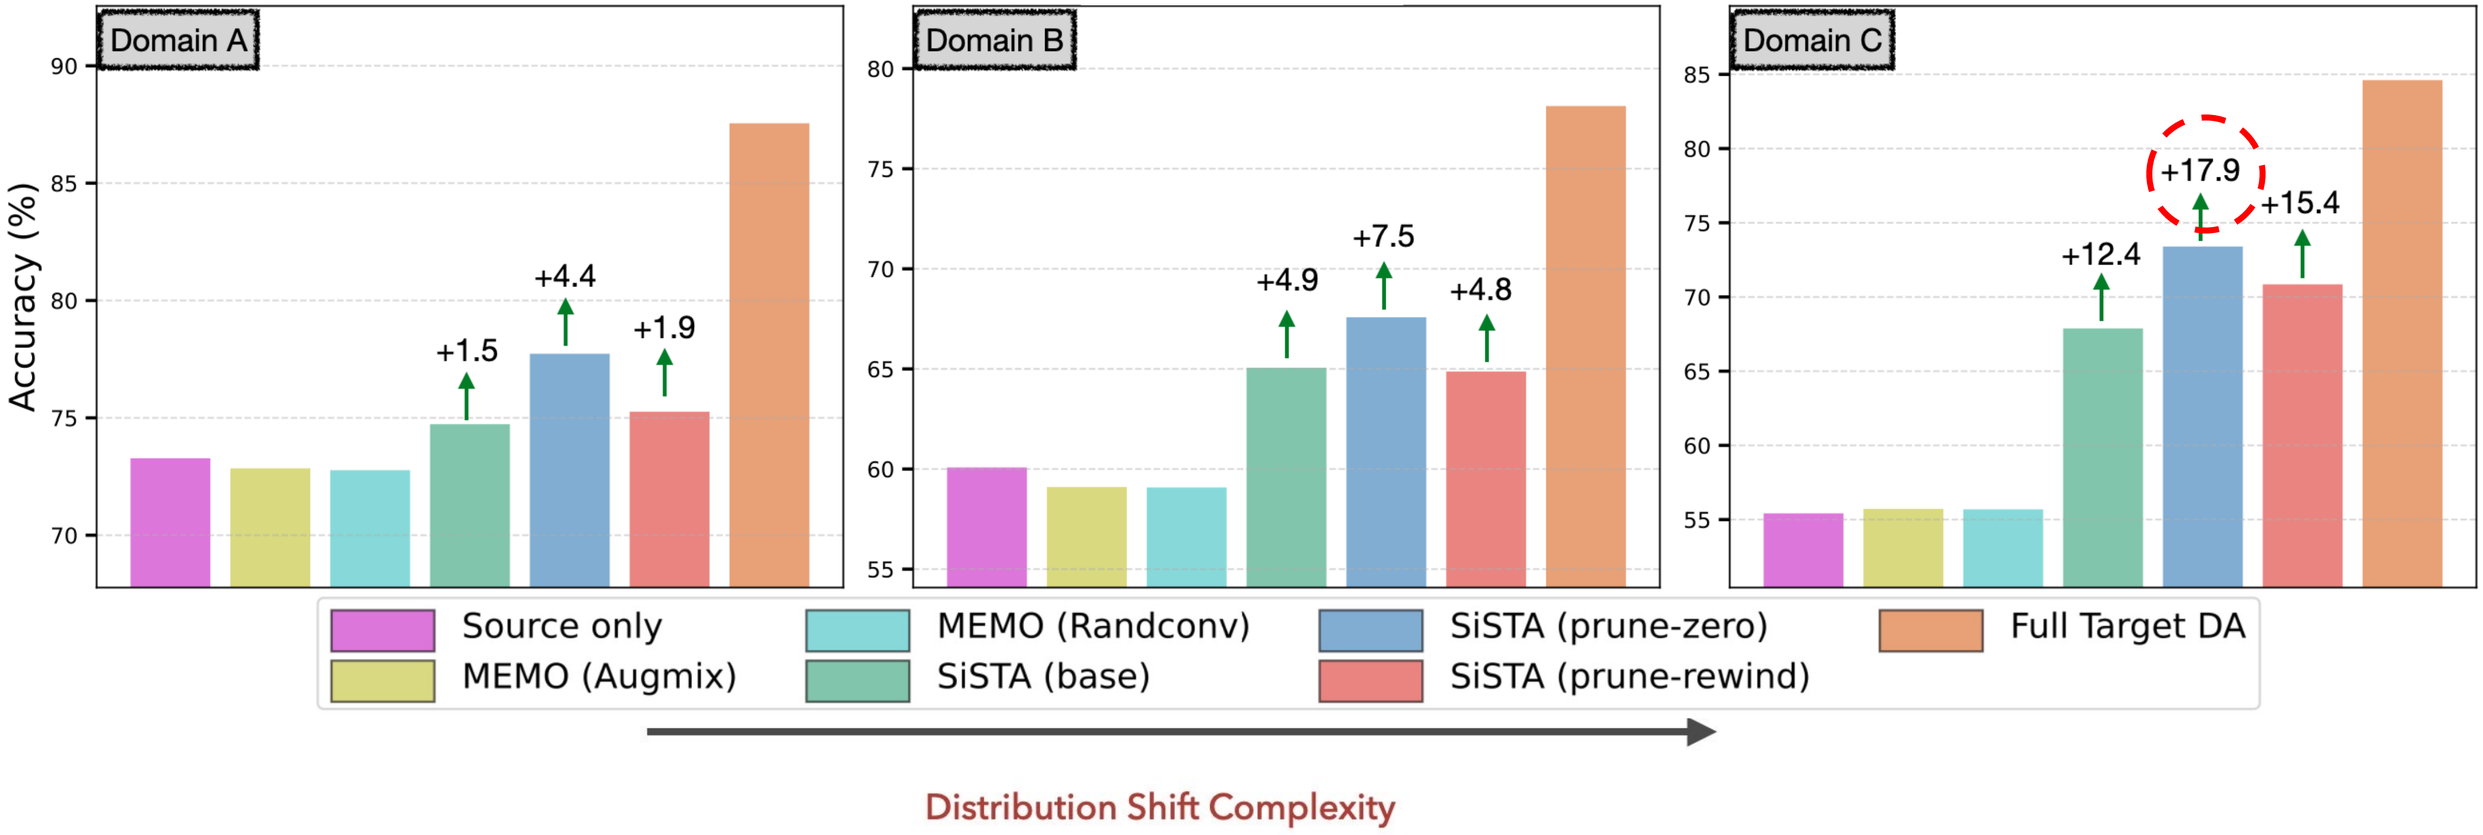

Employing Toolbox Augmentations on the Synthetic Target Images Matches the Full Target DA Performance

Water-Color

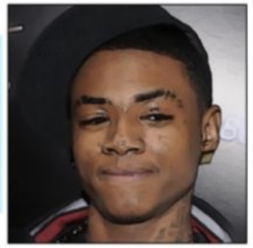

Color-Sketch

Pencil-Sketch

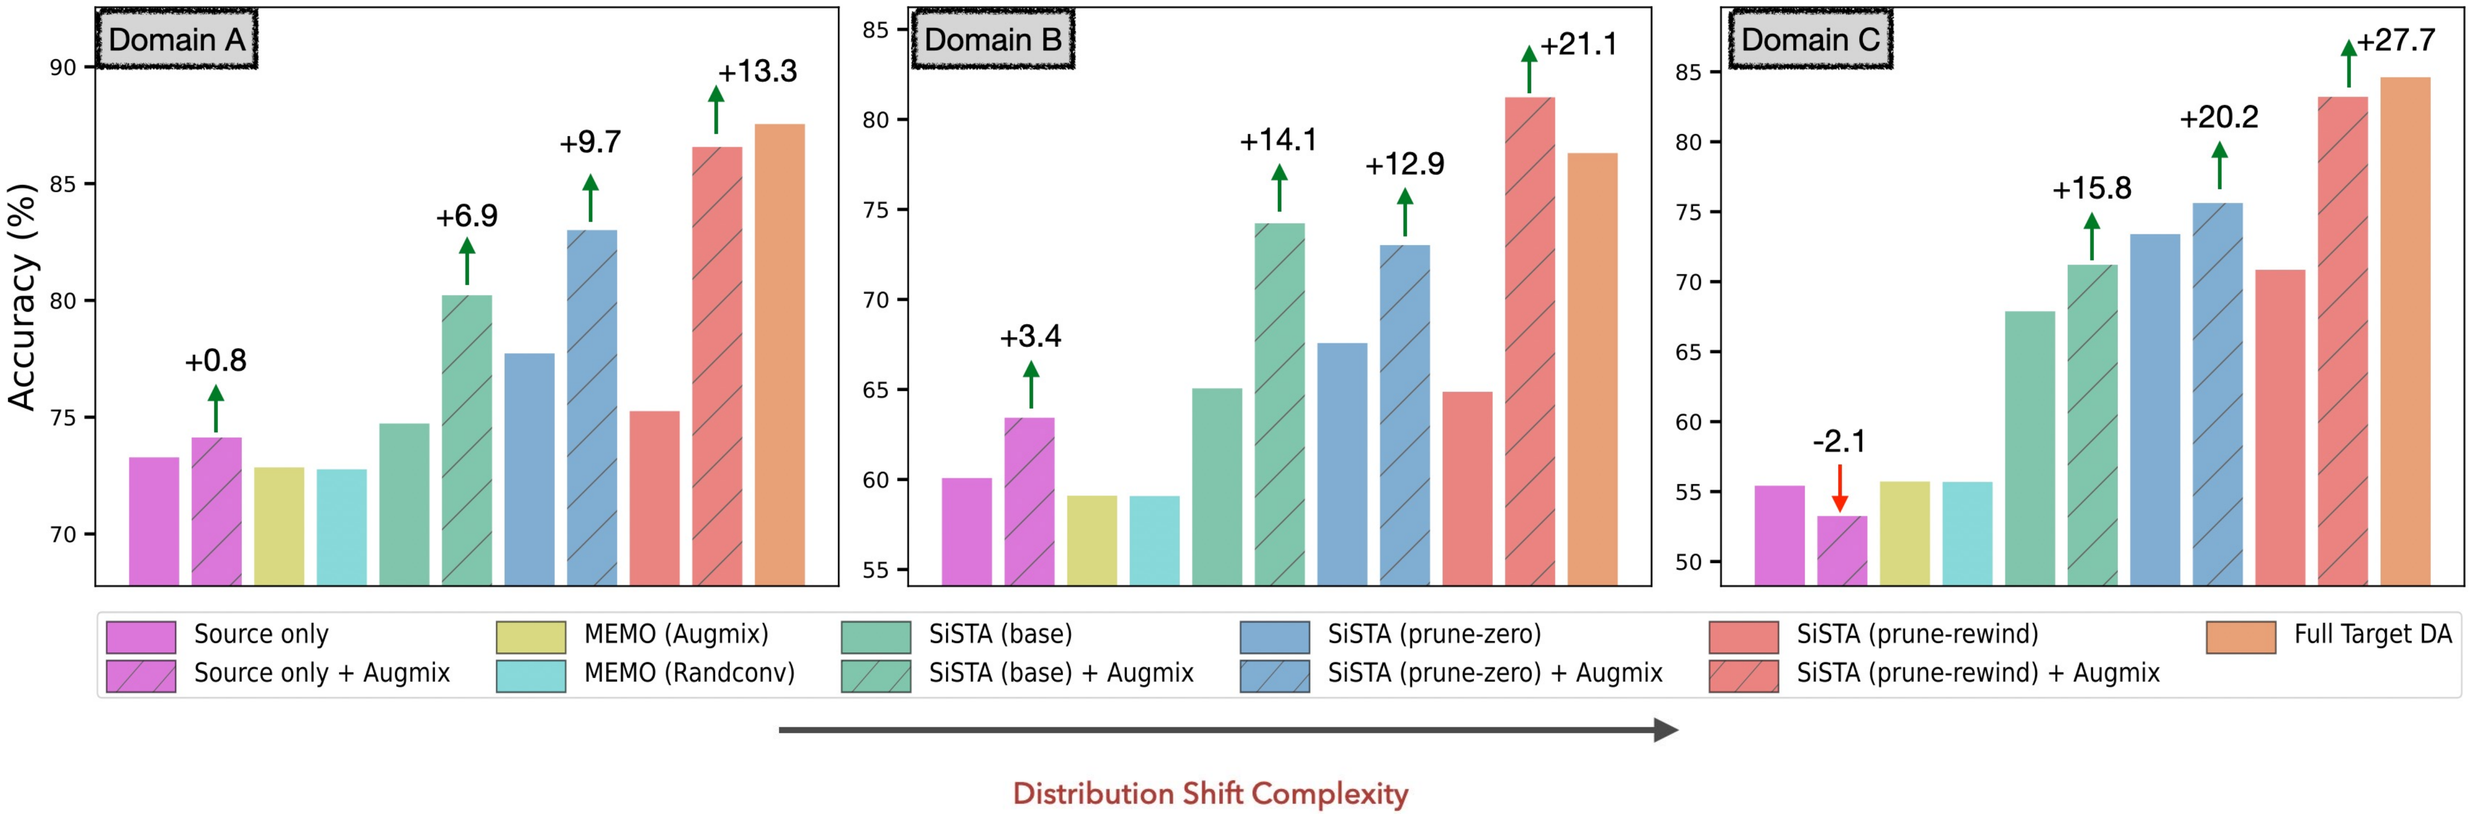

SiSTA Is Effective In Multi-Class Classification

CIFAR 10: Target Domain 1-shot Data

SiSTA-mCG

Class-Conditional Generator

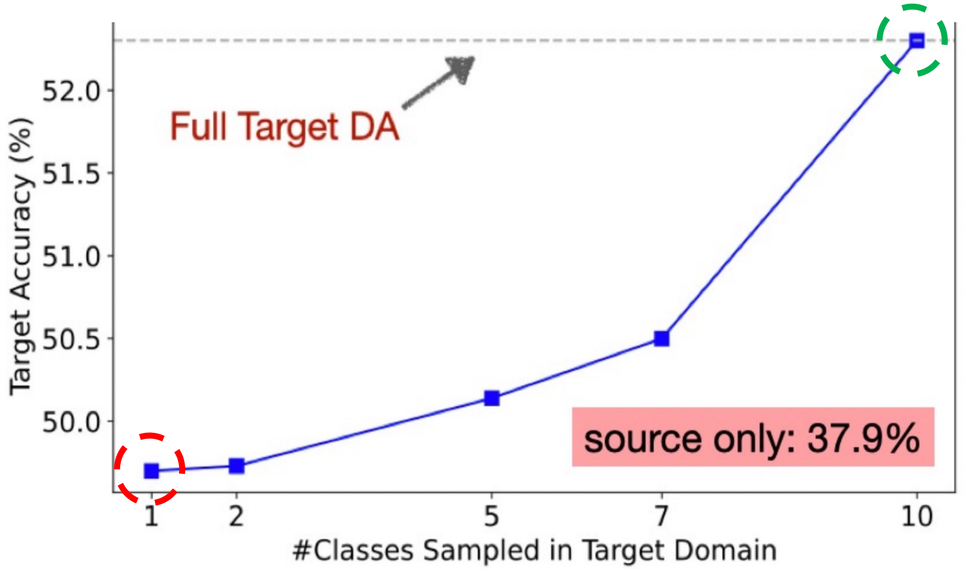

And The Benefits Continue to Persist Even On Large Scale Benchmarks

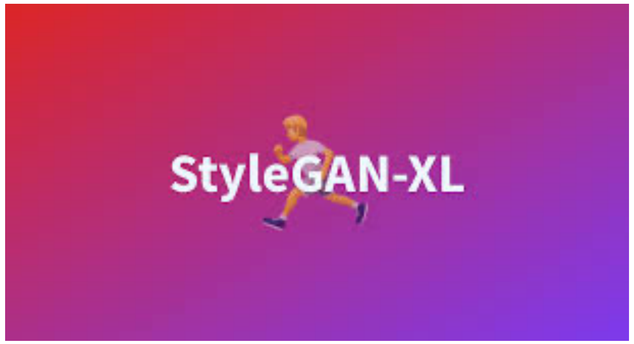

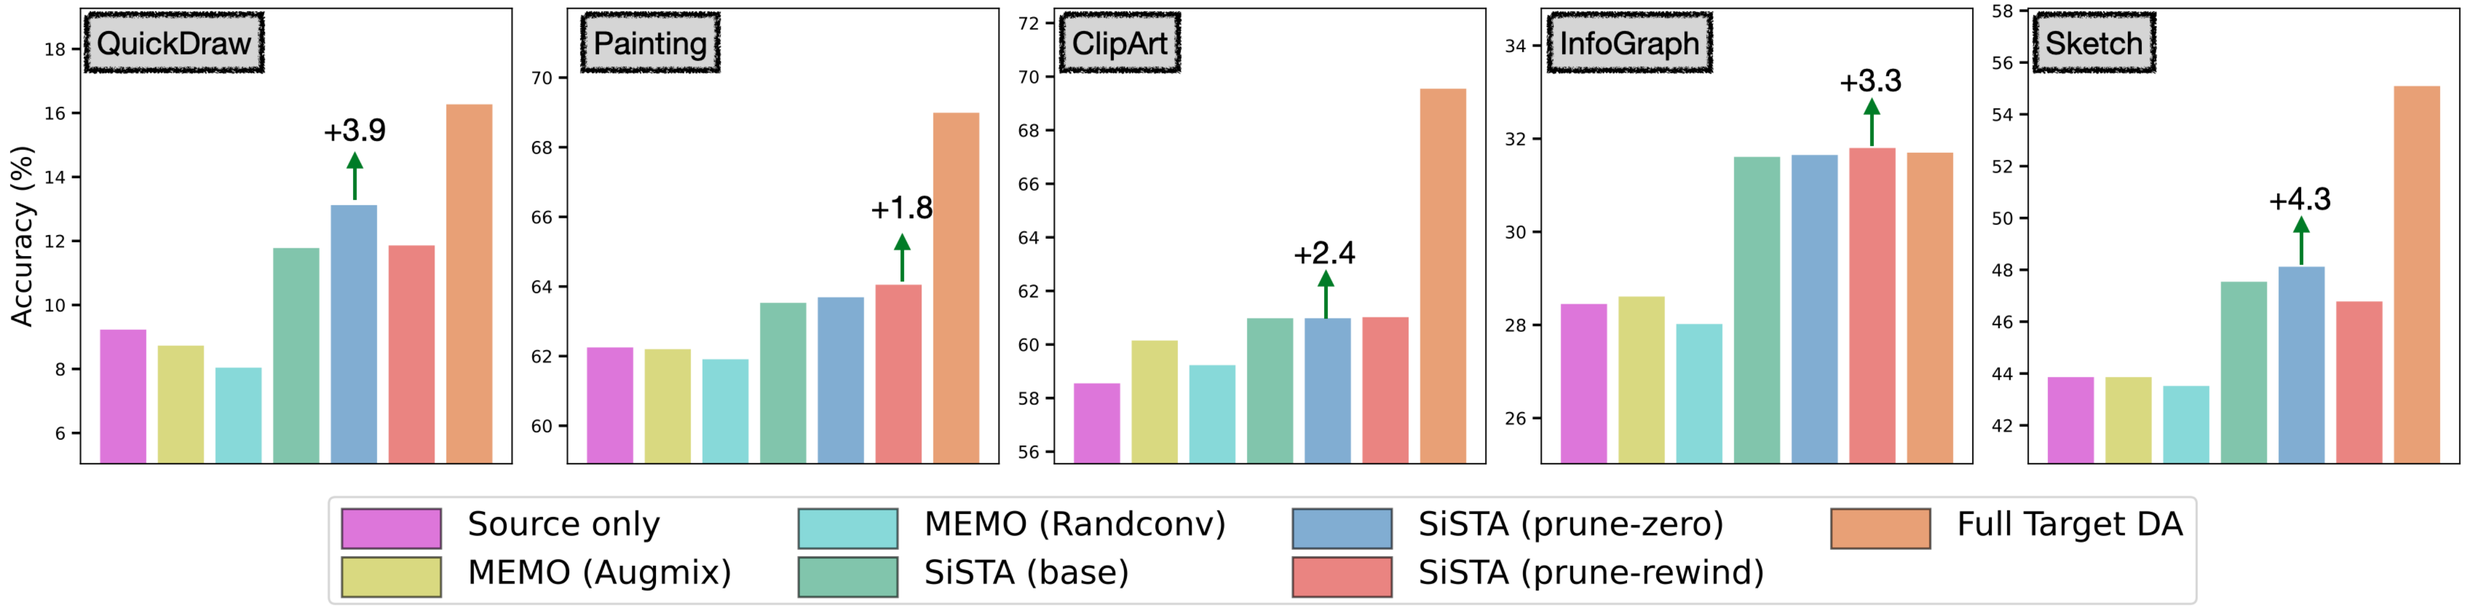

SiSTA:Target-Aware Generative Augmentations for Single-Shot Adaptation

Kowshik Thopalli * [2] Rakshith Subramanyam * [1] Pavan Turaga [1] Jayaraman J. Thiagarajan [2]

[1]Arizona State University [2]Lawrence Livermore National Laboratory

Paper Slides Poster Github

Summary

We address the challenge of adapting models from a source domain to a target domain, given the limited generalization ability of deep neural networks. Existing techniques rely on synthetic data augmentations when target data is scarce, but they struggle with significant distribution shifts. To overcome this, we propose SiSTA (Single-Shot Target Augmentations), which fine-tunes a generative model using a single target sample and employs innovative sampling strategies to generate synthetic target data. SiSTA outperforms existing methods in binary and multi-class problems, handles various distribution shifts effectively, and achieves performance comparable to models trained on full target datasets.

Paper

Website

Code

Come talk to us at the
Poster Session #5
July 27th 10:30 AM – 12:00 PM HST
Exhibit Hall 1

kowshik_thopalli@llnl.gov

rakshith.subramanyam@asu.edu