Ensembling with Deep Generative Views

Lucy Chai, Jun-Yan Zhu, Eli Shechtman, Phillip Isola, Richard Zhang

Massachusetts Institute of Technology · Carnegie Mellon University · Adobe

https://chail.github.io/gan-ensembling/

Overview

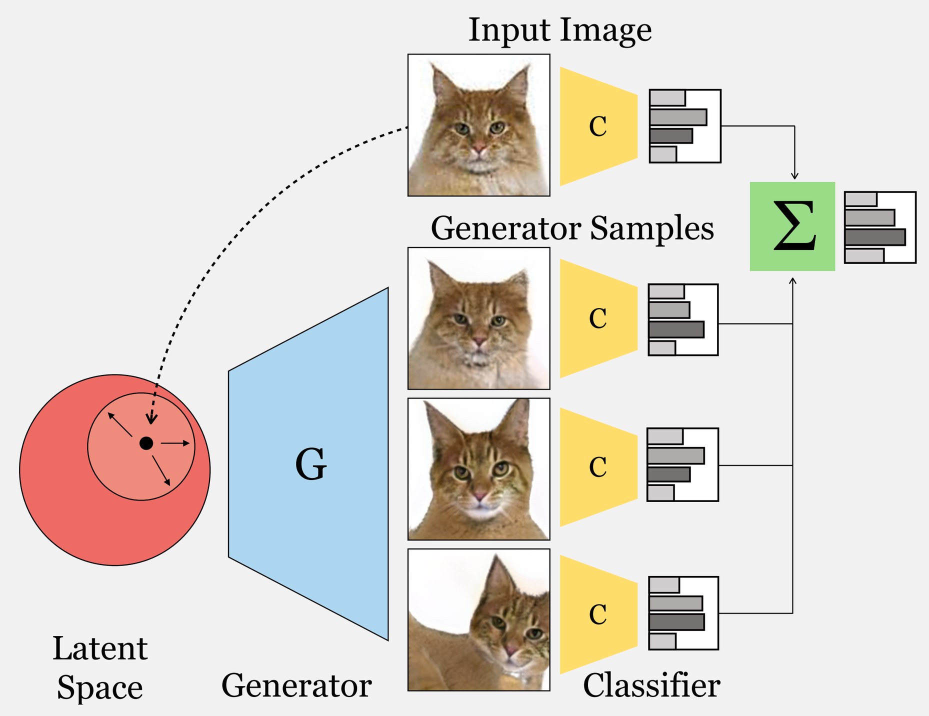

- Pretrained GAN creates "views" of an original input image
- Ensemble views for classification
- Balance original and GAN views, due to classifier sensitivity to GAN artifacts

Creating GAN-generated Image Variants

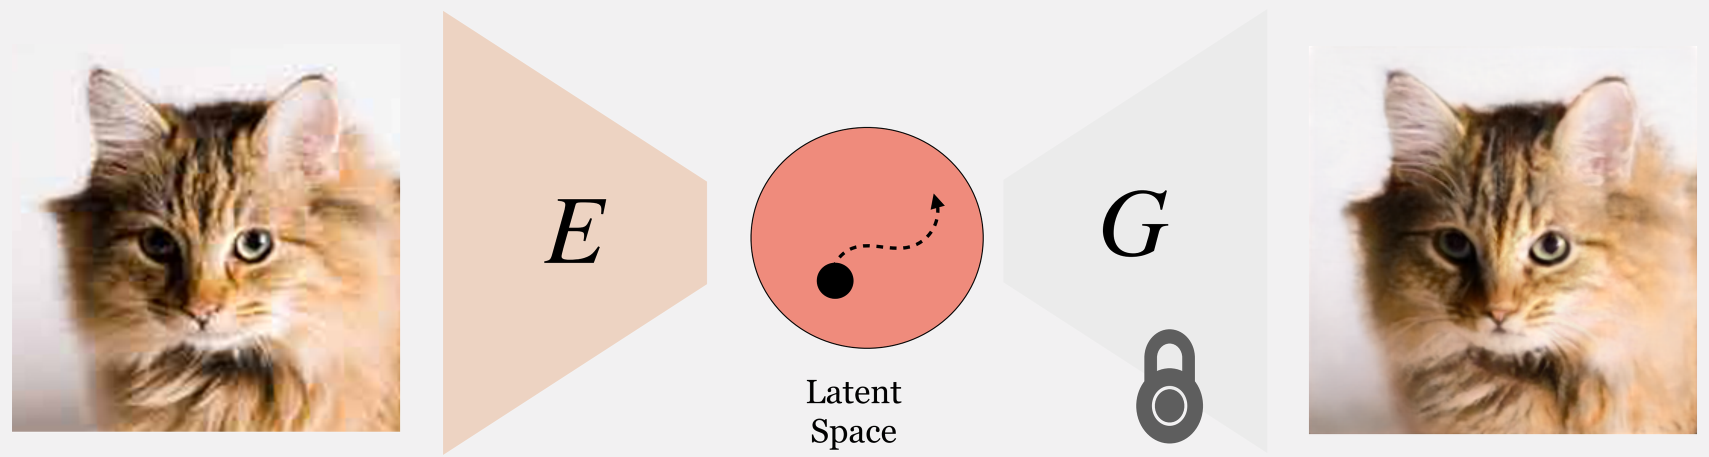

$$w^* = \arg\min_w L_{\text{img}}(x, G(w)) + \lambda L_{\text{latent}}(w, E(x))$$

- Encode and optimize to find image latent code, inversion required be fast and accurate

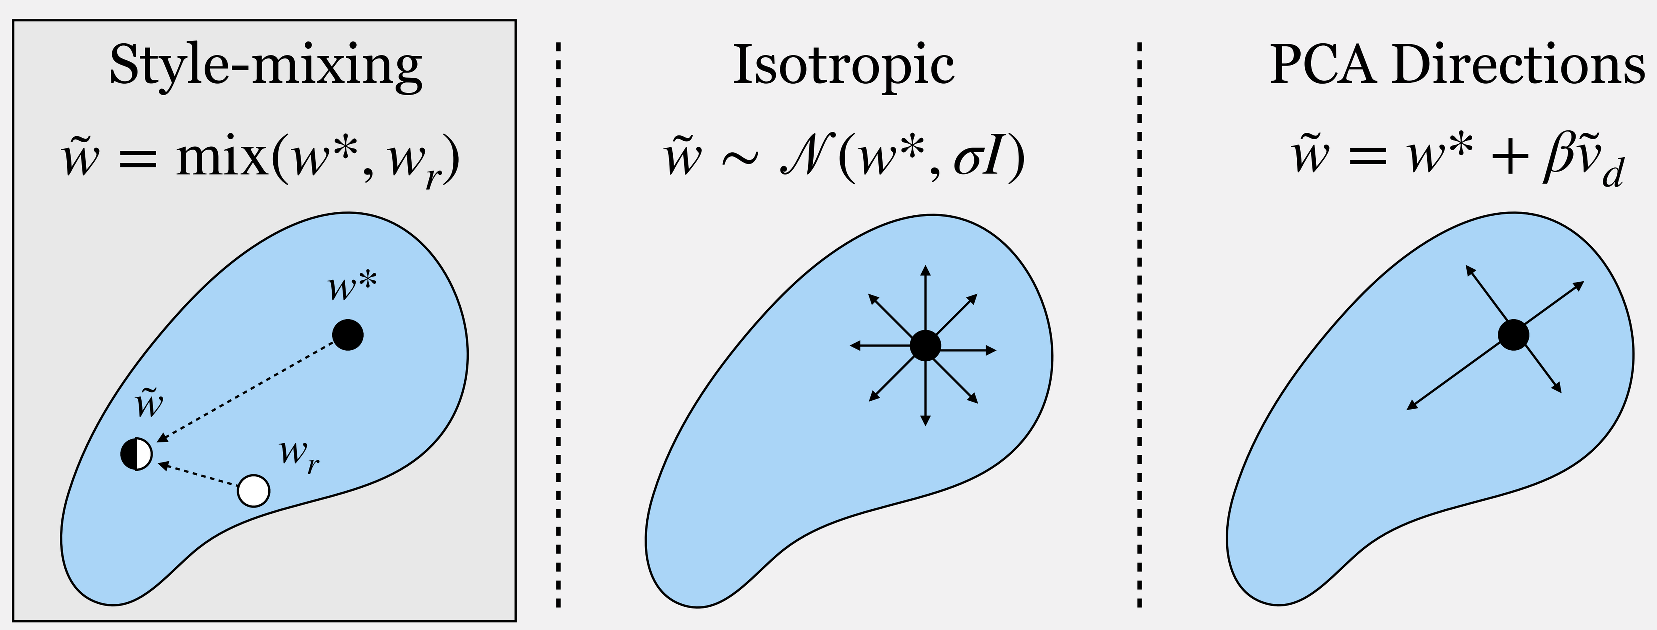

- Perturb optimized latent code in StyleGAN2 W+ space
- Style-mixing: swaps a random latent code at specified generator layers
- Style-mixing is more robust compared to isotropic or PCA perturbation

Visualizing GAN Outputs

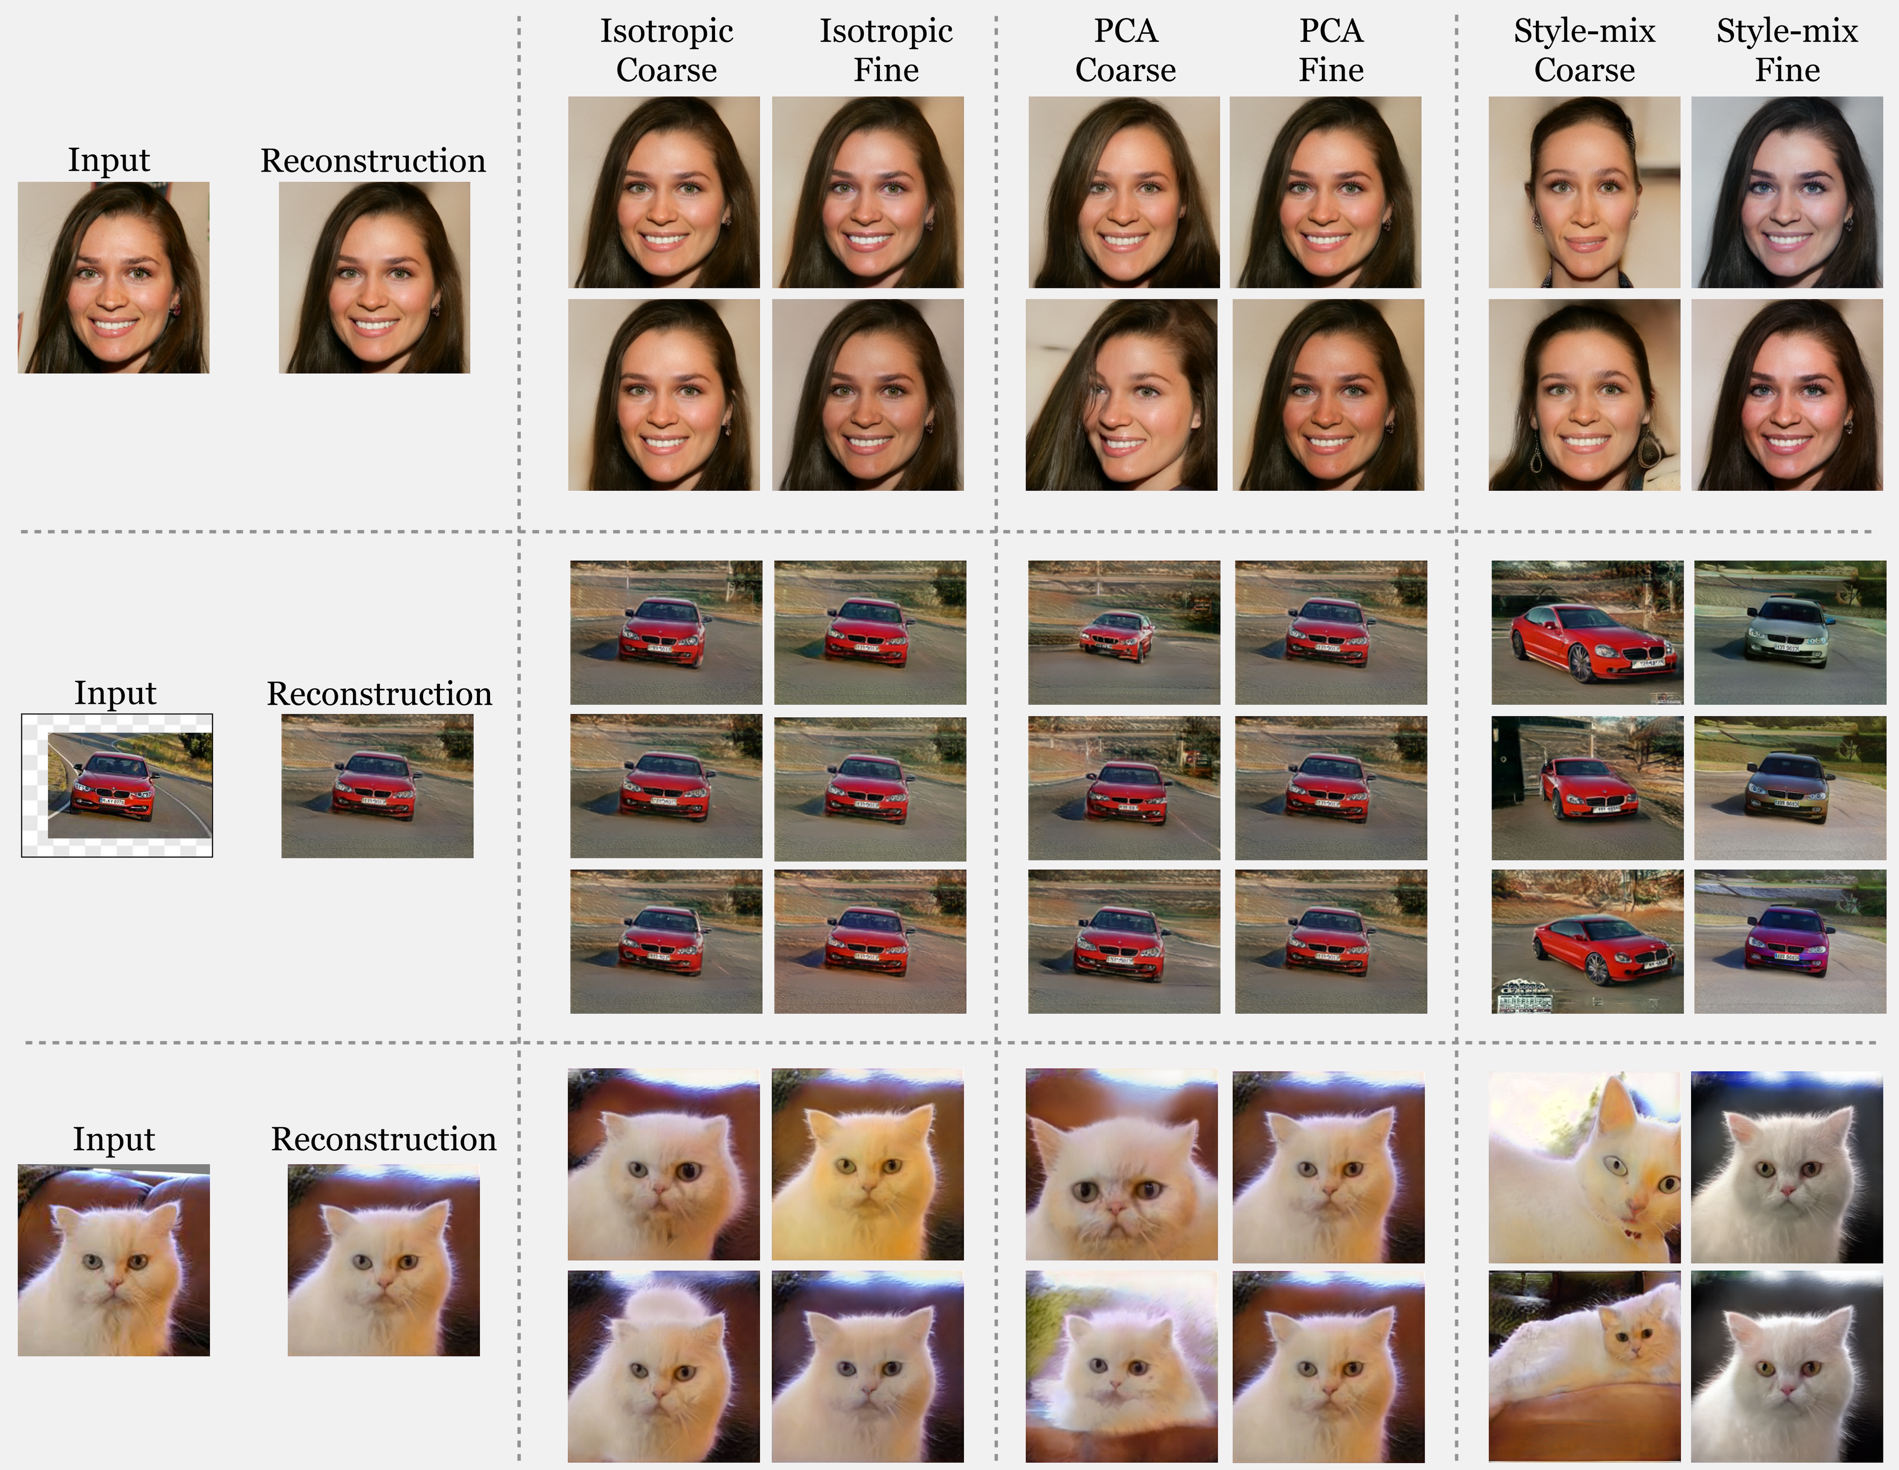

- Perturb coarse or fine layers of StyleGAN2 (middle layers change identity)
- Coarse layers change shape, fine layers change coloring
- Style-mixing causes larger visual changes compared to isotropic or PCA perturbations

Adjusting the Ensemble Weight

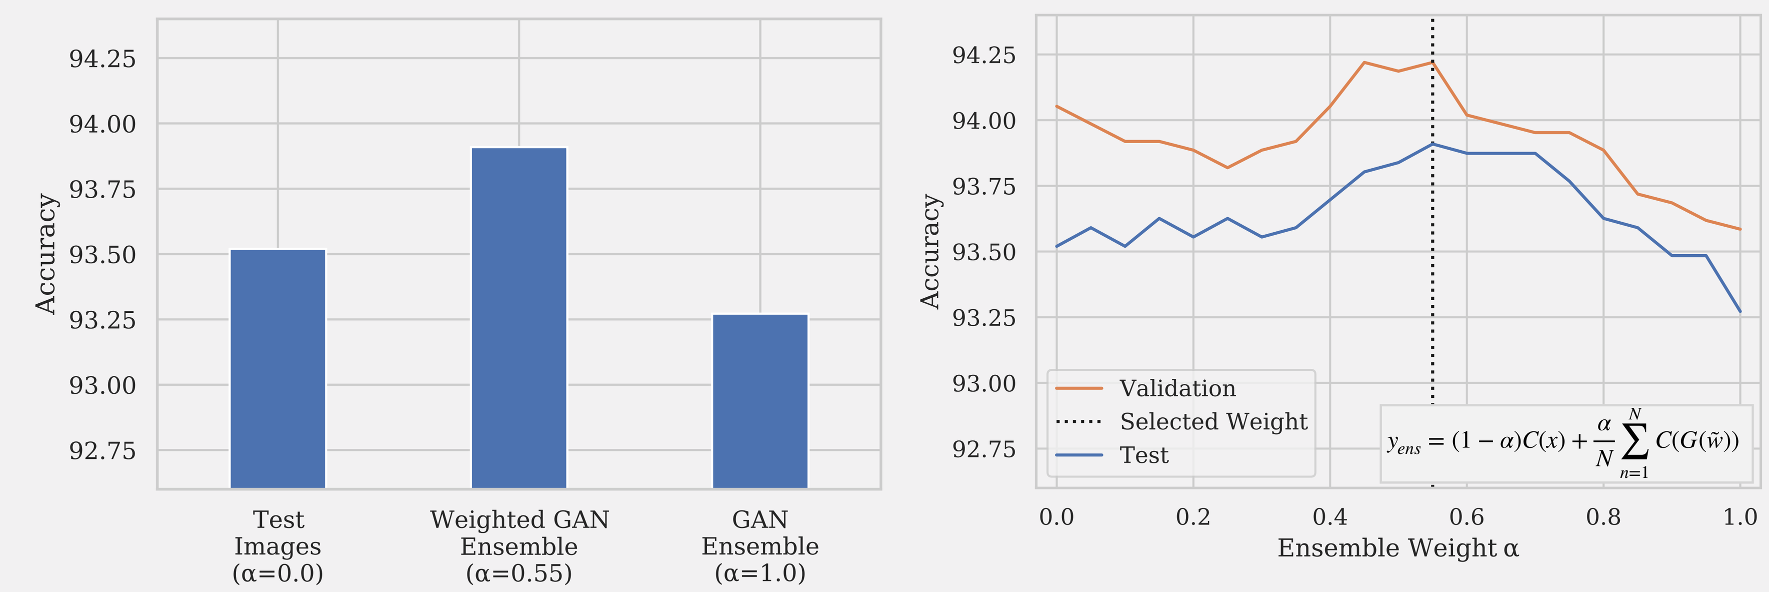

- Only using GAN views can decrease accuracy
- Softly weight between the dataset original and GAN views
- Weight selected based on validation data

Properties of GAN Ensembling

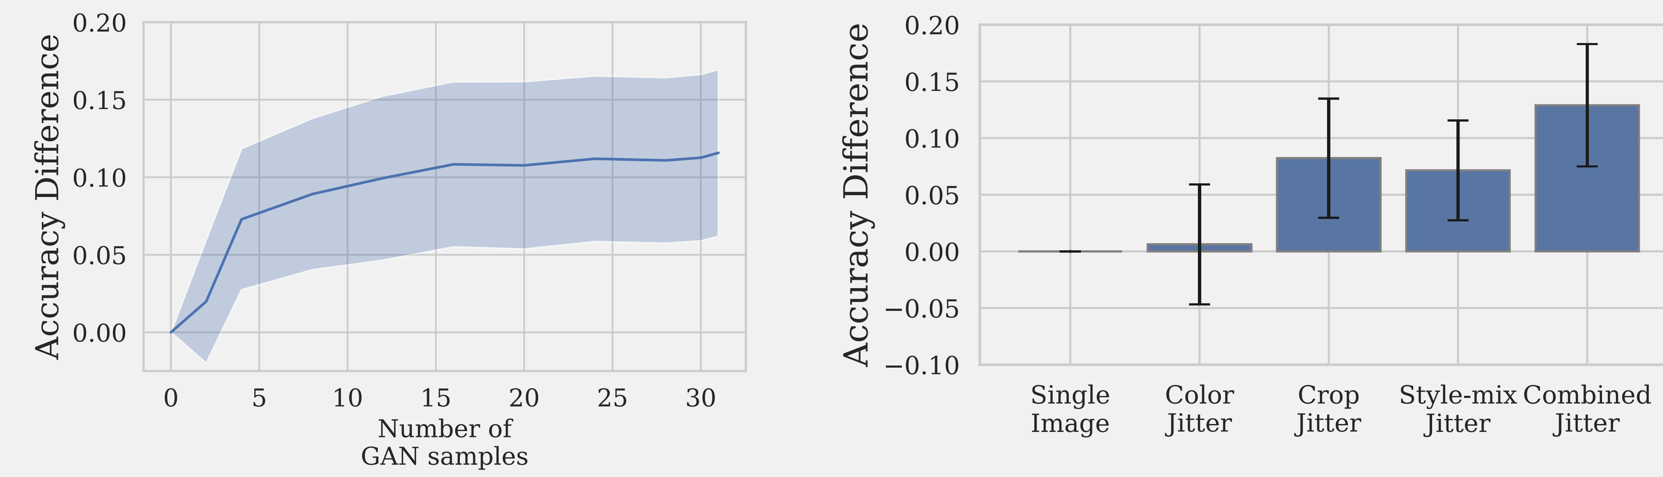

- Showing average over 40 CelebAHQ binary attributes
- Increasing GAN ensemble size improves classification but saturates
- GAN style-mixing and traditional image cropping perform similarly
- Combining traditional and GAN views offers additional improvement

Ensembling on Multi-class Tasks

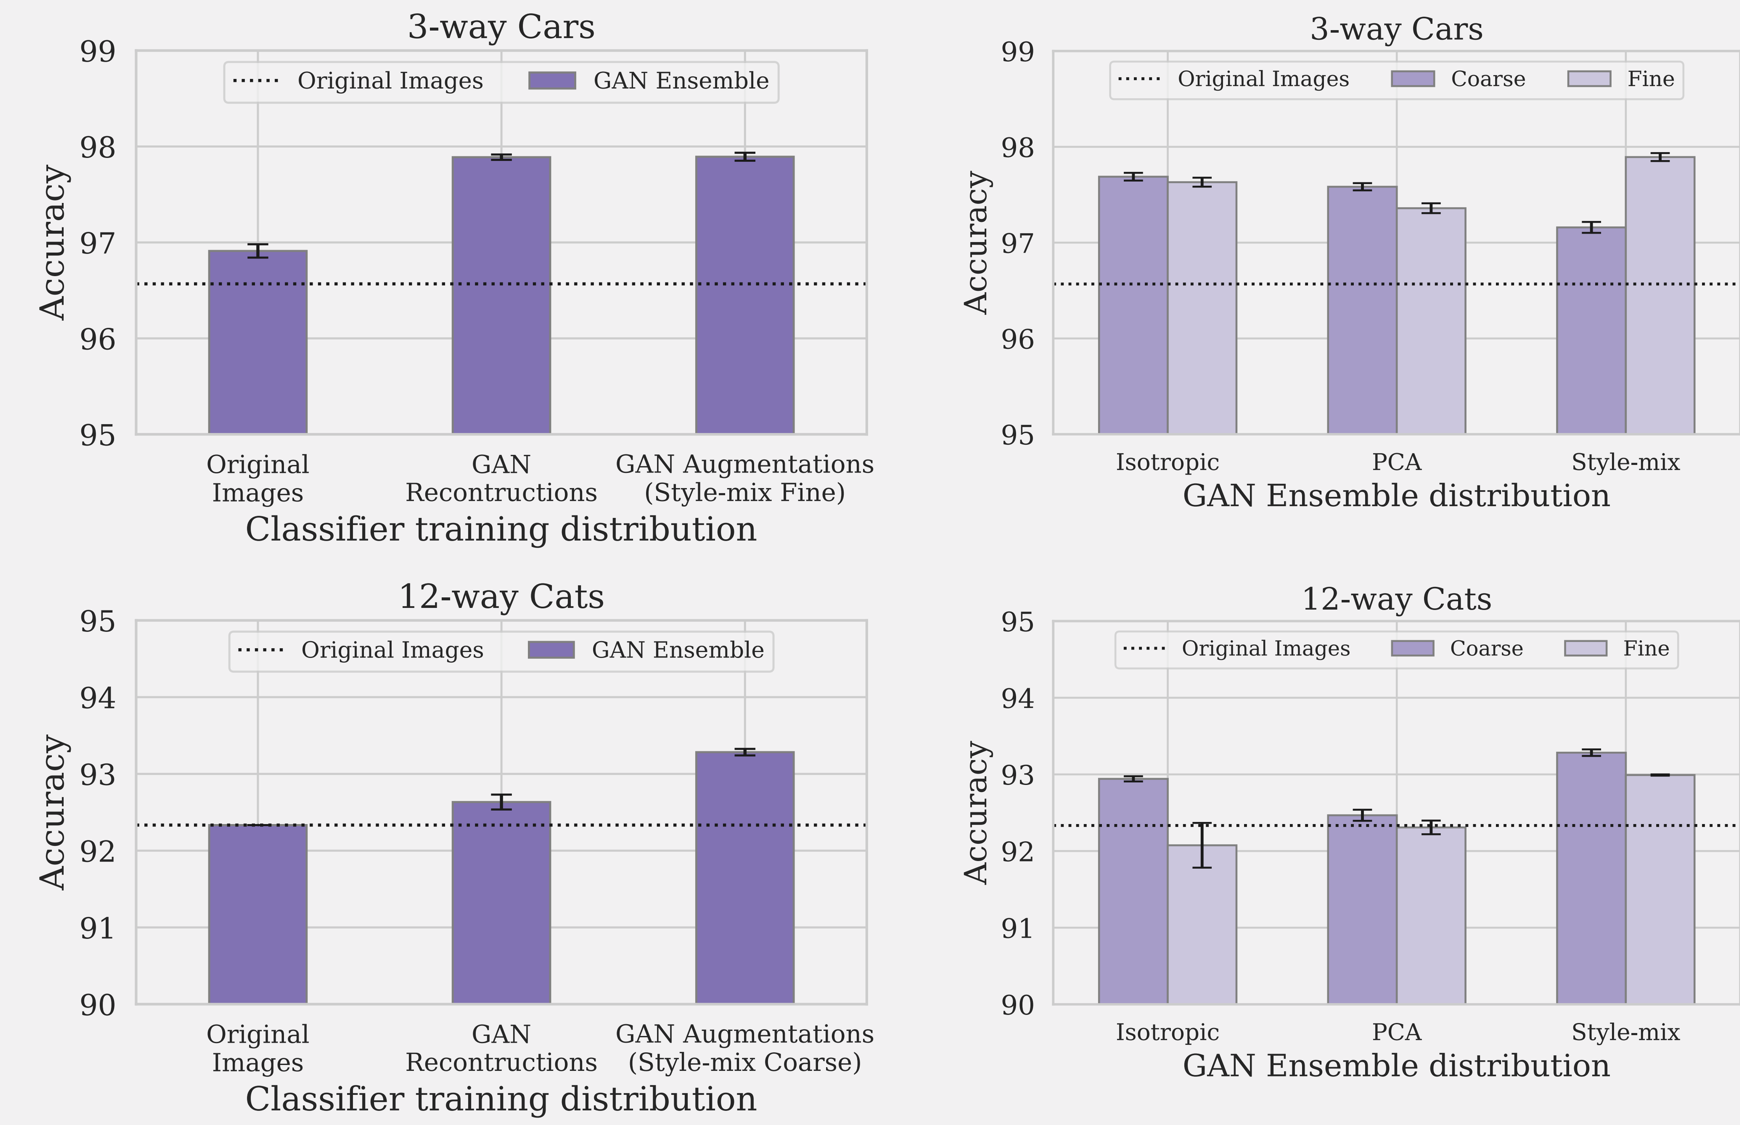

- Finetuning the classifier on GAN outputs improves ensembling benefit
- Style-mixing perturbation at test time is best

Limitations of Ensembling with GANs

- GAN reconstruction capability: must preserve discriminative attributes
- GAN inversion efficiency: optimization time and reconstruction tradeoff
- Classifier sensitive to GAN artifacts
- Currently limited to simple tasks and small structured datasets
- Similar techniques may yield greater benefits in the future as GANs improve