Robust Projection onto Image Manifolds

Rushil Anirudh,
Jayaraman J Thiagarajan,
Bhavya Kailkhura,
Timo Bremer

CASC
Center for Applied Scientific Computing

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

LLNL-PRES-774227-DRAFT

Lawrence Livermore National Laboratory

Motivation

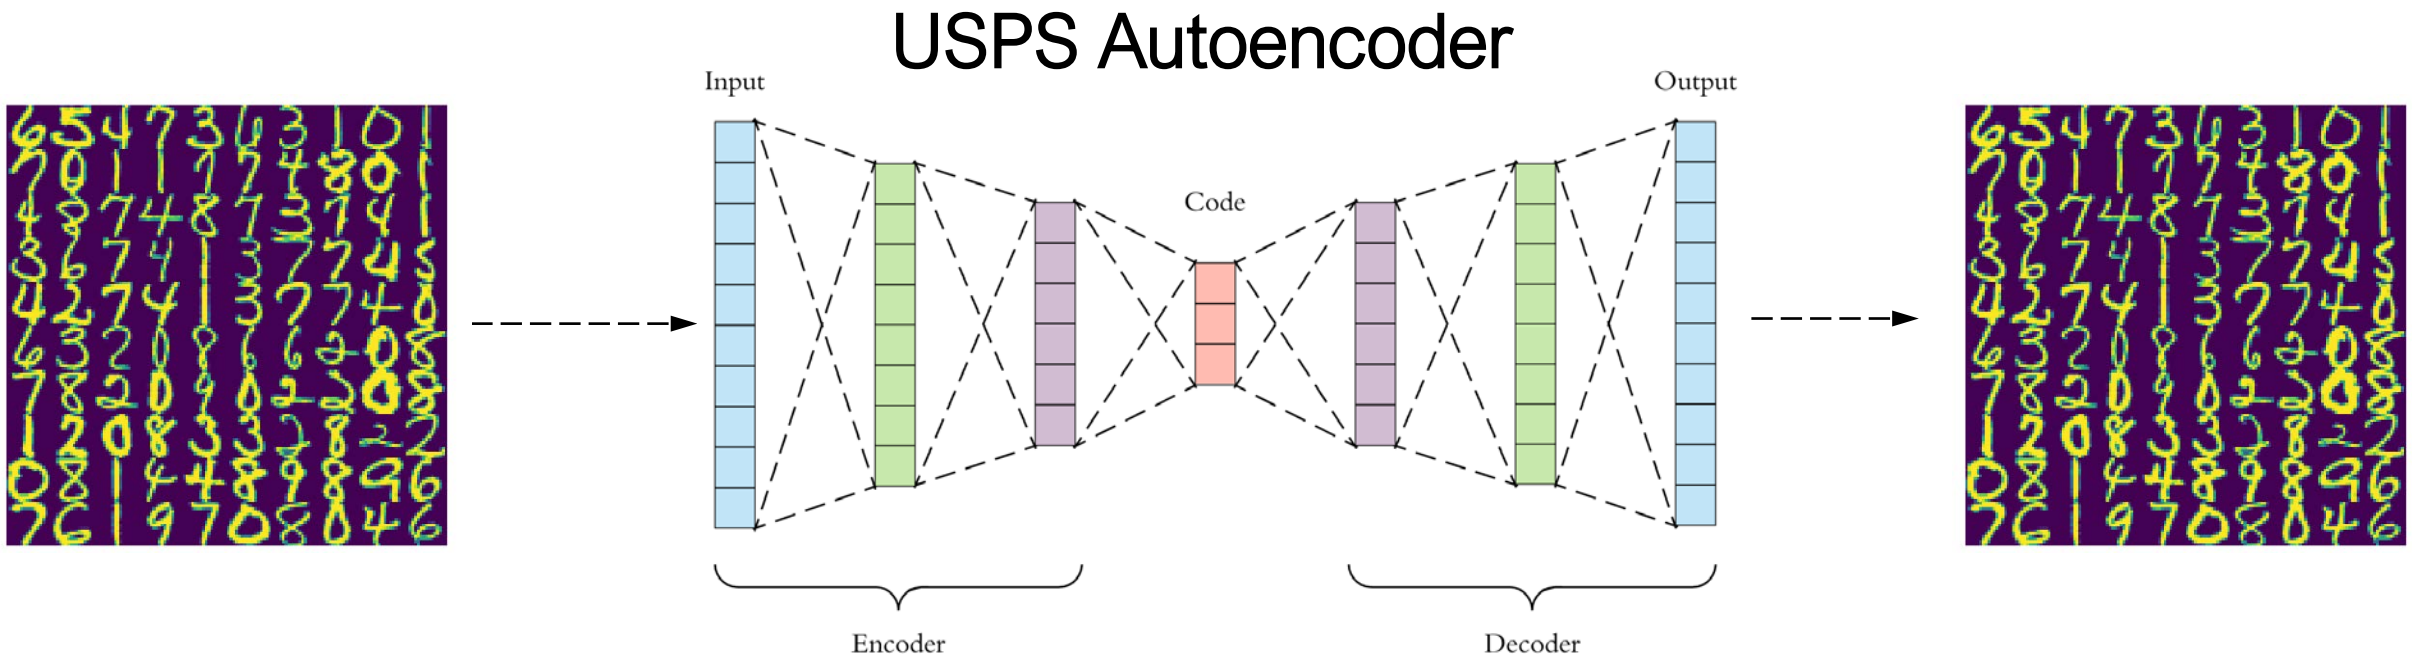

Reconstructing when dataset has changed?

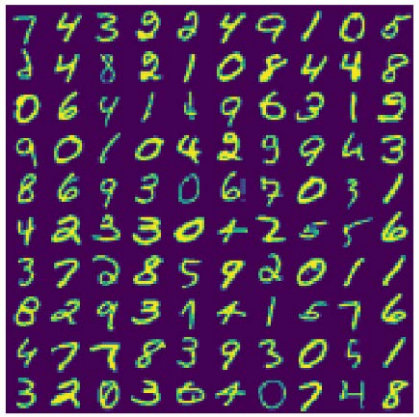

Images from MNIST dataset

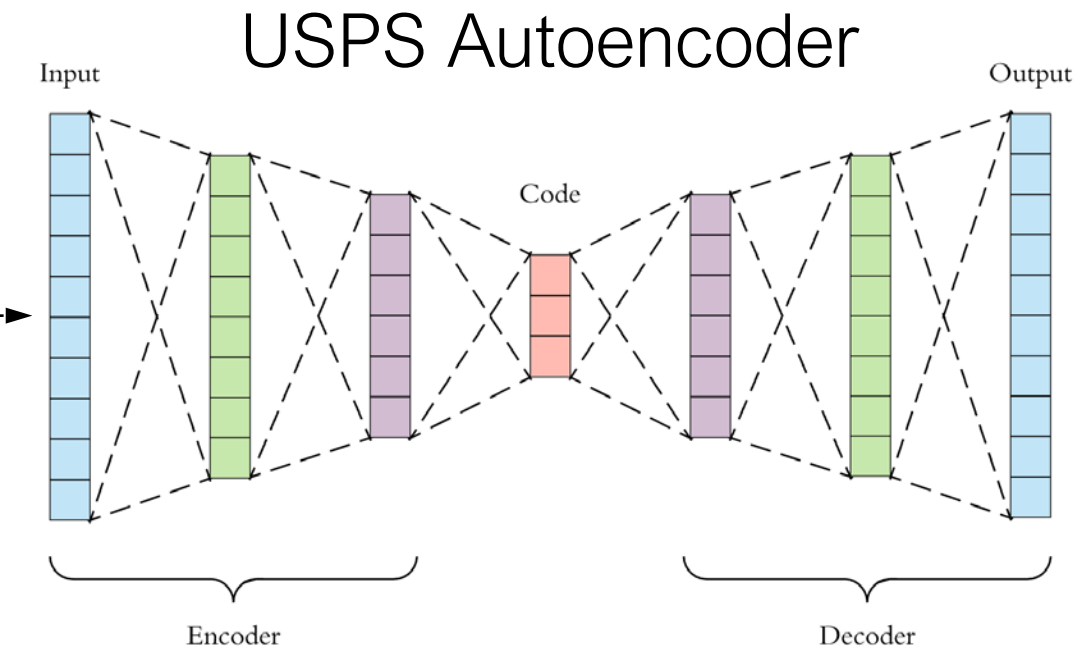

?

Reconstructing when dataset has changed?

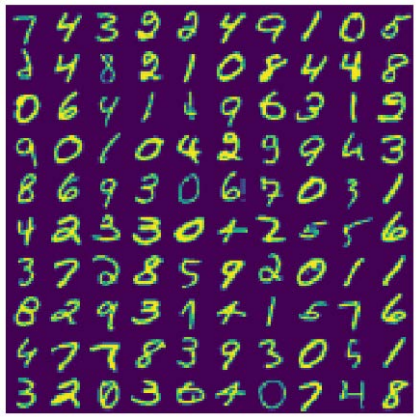

Images from MNIST dataset

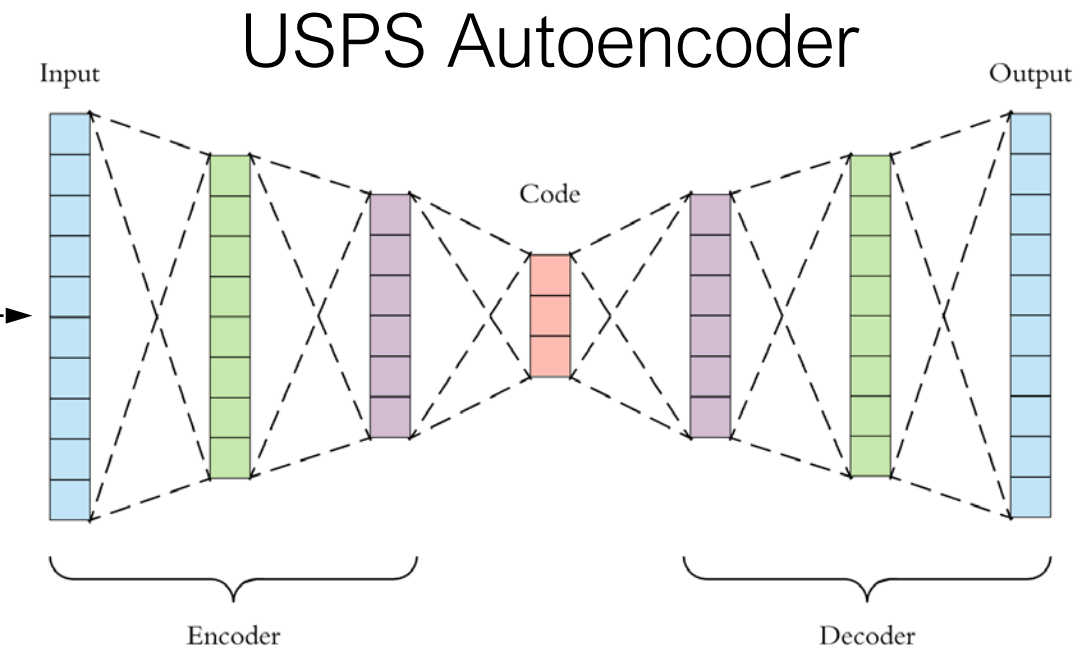

Autoencoder cannot handle the distribution shift

Projection 101: How do we project onto a manifold?

CASC
Center for Applied Scientific Computing

A: Encoder + decoder model (direct inference)

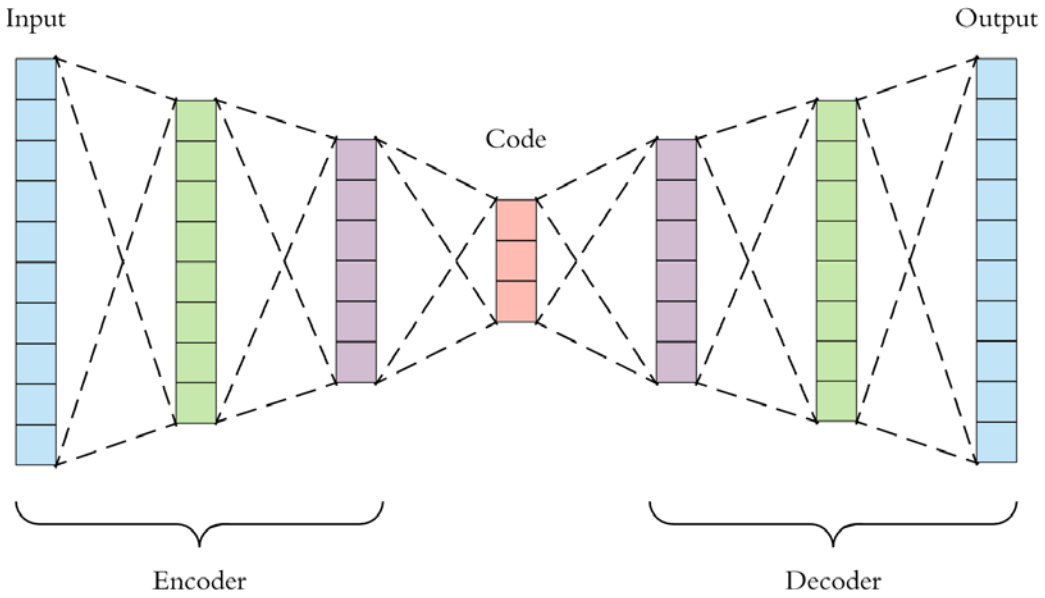

B: Extrinsic Mapping onto the manifold (in-direct inference)

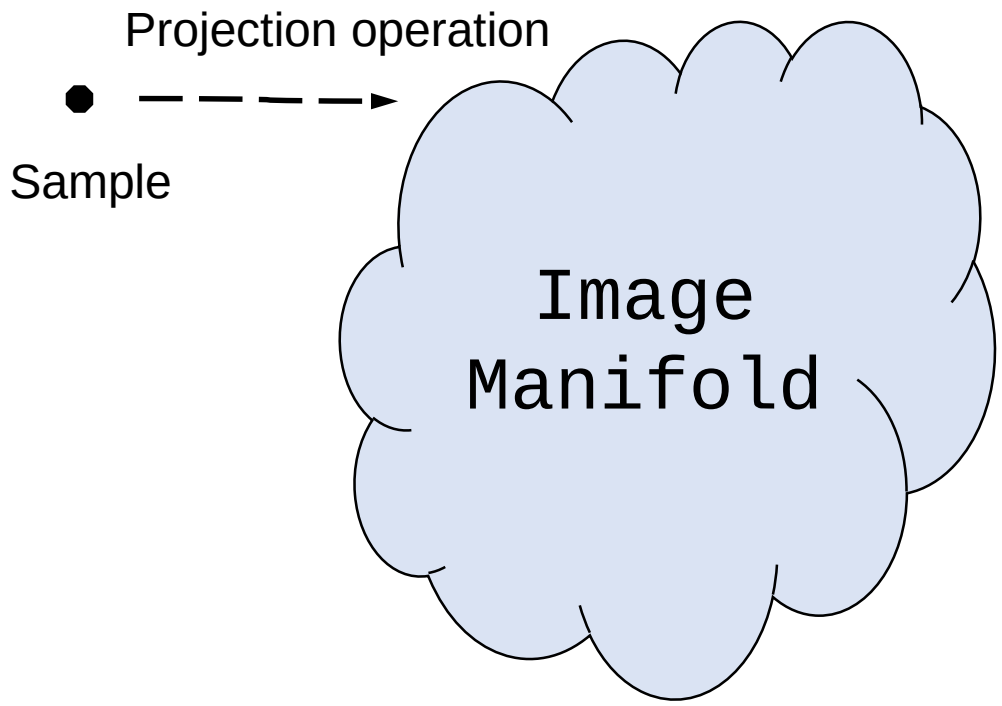

In-direct inference: PGD optimization

$$z^* = \arg\min_{z \in \mathrm{R}^d} ||Y - \mathcal{G}(z^*)||$$

Final projection is given by $\mathcal{G}(z^*)$

Minimize decoder output

Projected Gradient Descent: Walking in the latent space to minimize projection error

Why is it not robust?

Take a simple example where the image to be projected is corrupted in an unknown fashion:

$$Y \rightarrow \mathcal{F}^{?}(Y)$$

$$z^{*} = \arg\min_{z \in \mathrm{R}^{d}} ||\mathcal{F}^{?}(Y) - \mathcal{G}(z)||$$

Unless the loss function is robust to the corruption “F”, **this optimization will fai**l.

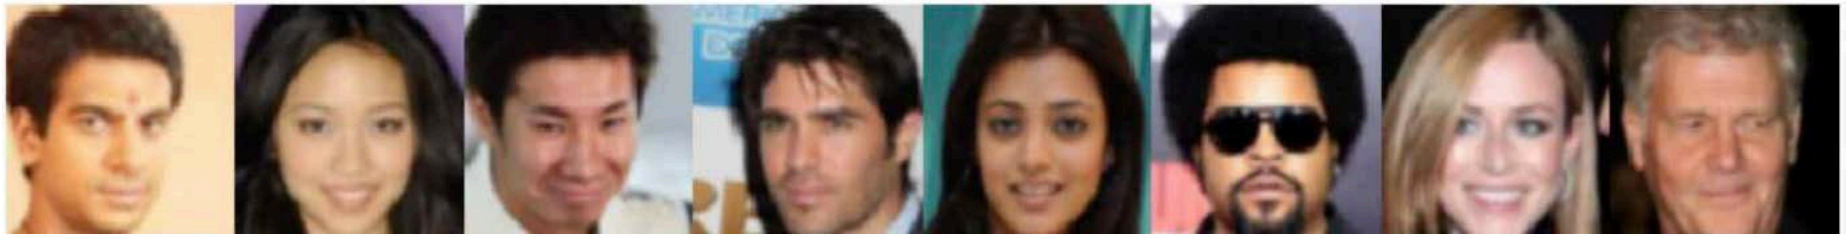

Examples from the face manifold

$\mathcal{F}^{?}(Y)$

Corrupted observations with no knowledge of corruption

$\mathcal{G}(z^{*})$

How do we make it more Robust?

The unknown corruption is causing the issue, so we can try to estimate it.

$$z^*, \hat{f}^* = \arg\min_{z \in \mathrm{R}^d} ||\mathcal{F}^?(Y) - \hat{f}(\mathcal{G}(z))||$$

Final projection is given by

$$\mathcal{G}(z^*)$$

A shallow neural network is "trained" to estimate the corruption

Robustness Experiments

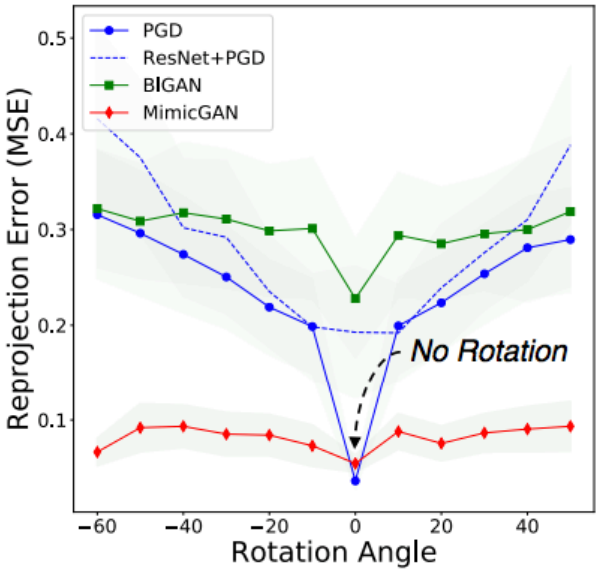

(a) Robustness to rotations

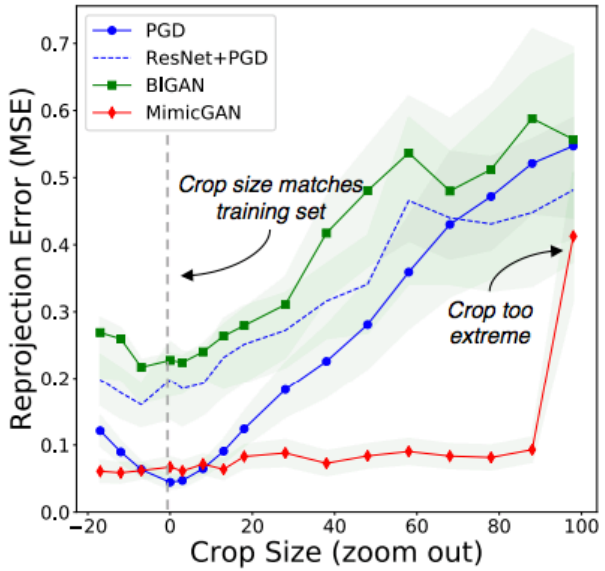

(b) Robustness to scale

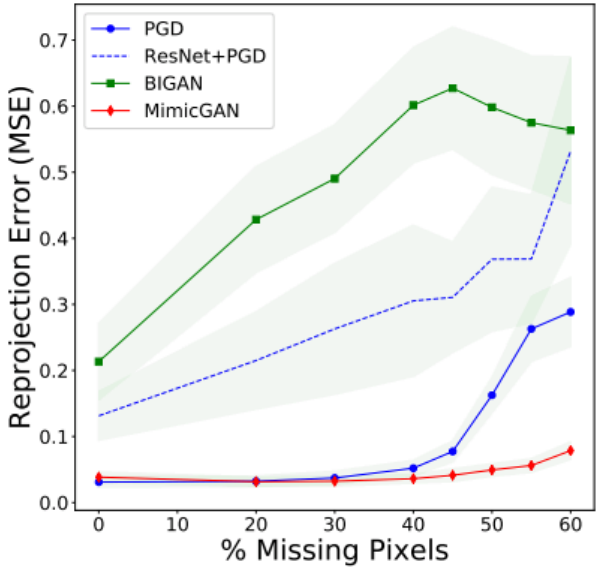

(c) Projections under missing pixels

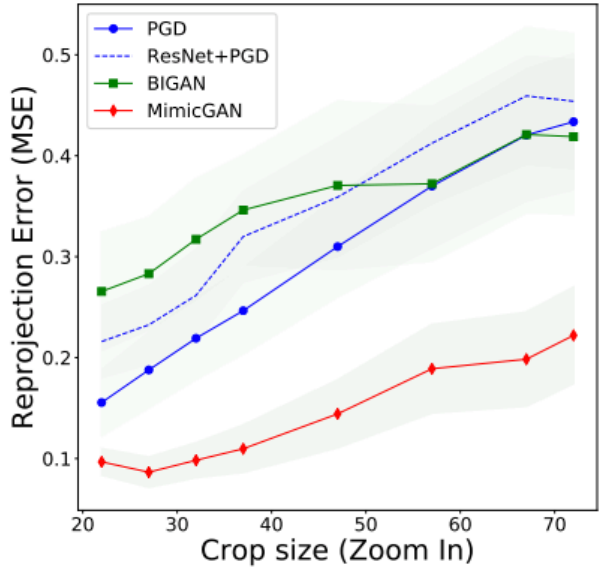

(d) Projections under missing context

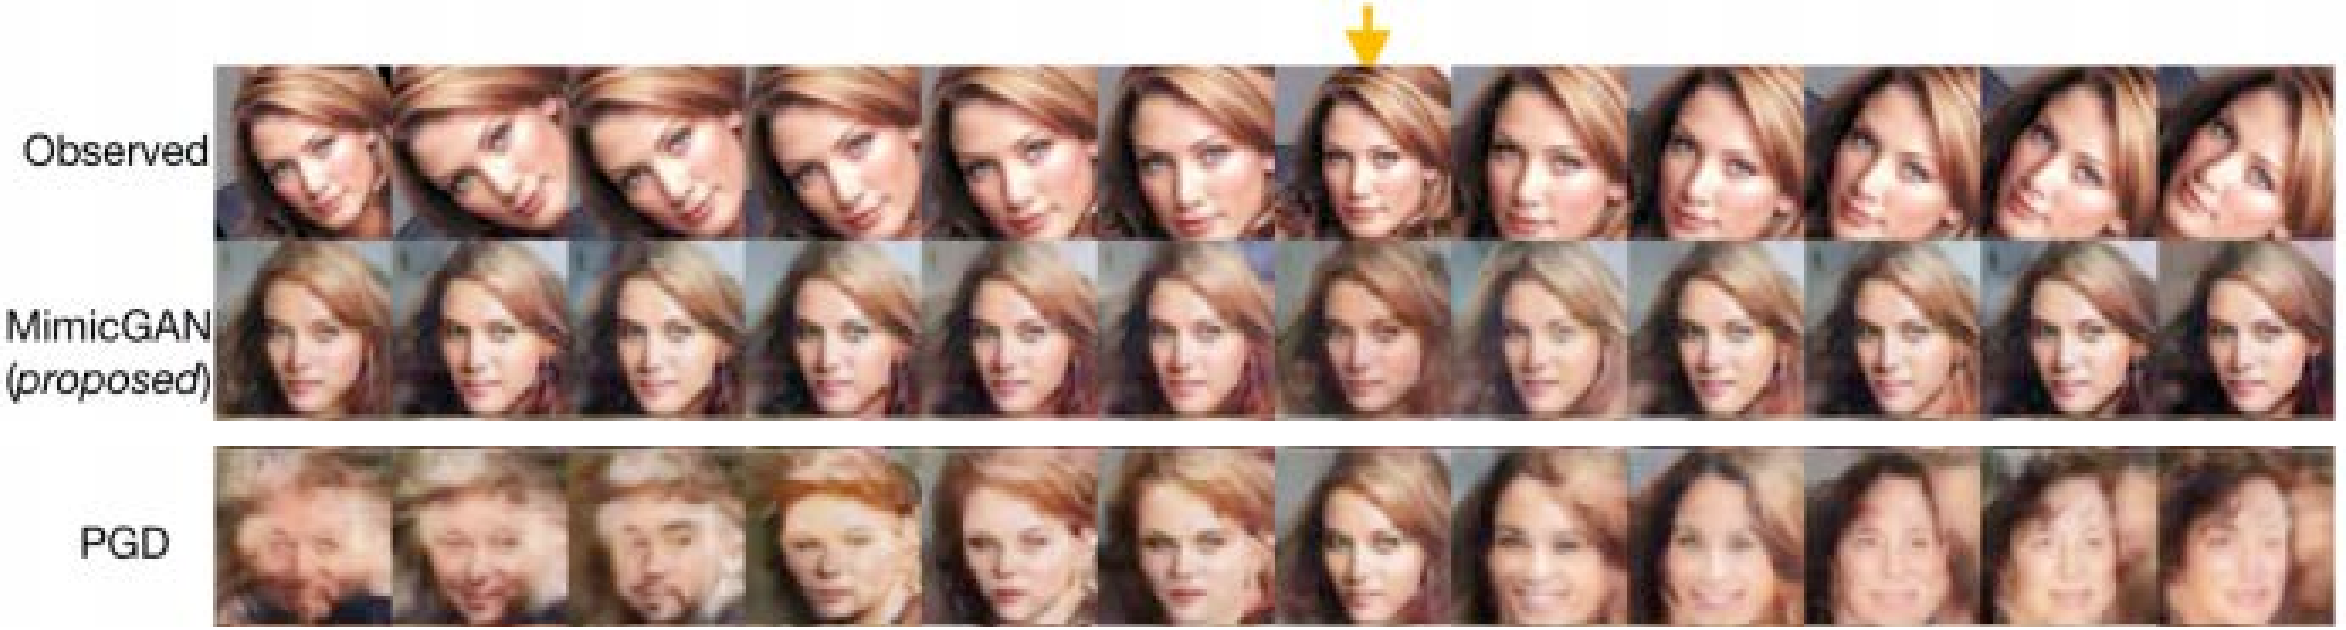
Observed
MimicGAN
(proposed)
PGD

Projected onto MNIST manifold with MimicGAN ← *USPS Target* → *Projected onto MNIST manifold with PGD*

Projected onto USPS manifold with MimicGAN ← *MNIST Target* → *Projected onto USPS manifold with PGD*

(b) MNIST→USPS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

LLNL-PRES-774227-DRAFT