GAN-Based Image Data Augmentation

Nathan Hu
zixia314@stanford.edu

David Liu
dliud@stanford.edu

Introduction

- Generative Adversarial Networks (GANs) are powerful generative models introduced by (Goodfellow et al.) [7] and can be trained on as little data as a single image [5].
- Lack of data makes ML hard -- data augmentation
- Prior work:
  - "Translating" images [3]
  - Generating numeric data [1]
- *Motivation: Explore using these super powerful generative models to augment more complex data sets*
- Classic Problem: Image classification of numbers in the MNIST database.

Direct Data Augmentation

- Trained the classifier on purely GAN-generated data for GANs of various sizes
- Pure synthetic data comparable to pure real data in training classifier

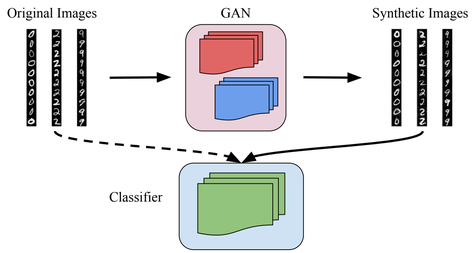

Train Size	Baseline	α = 1	α = 2	**α = 4**	α = 6
250	0.641	0.422	0.615	**0.709**	0.698
500	0.648	0.611	0.741	**0.763**	0.710
1000	0.683	0.670	0.694	**0.756**	0.738
2000	0.788	0.687	0.680	**0.793**	0.781

- Trained classifier on mixed synthetic data + real data in various ratios
- Mixed data outdoes pure real data; more noticeable for small datasets
- Unstable training losses suggest higher variance in real data

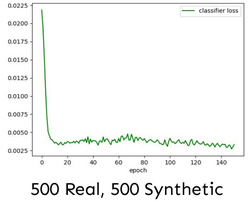

500 Real, 500 Synthetic

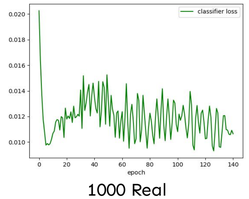

1000 Real

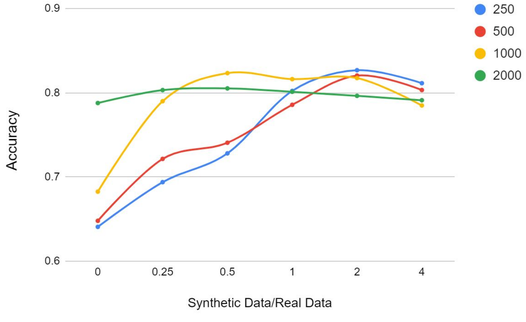

Recursive GAN Training

- Repeatedly use GANs to augment the dataset of images then used to train more GANs
- Classifier performance shows oscillating accuracies before long-term drop in performance

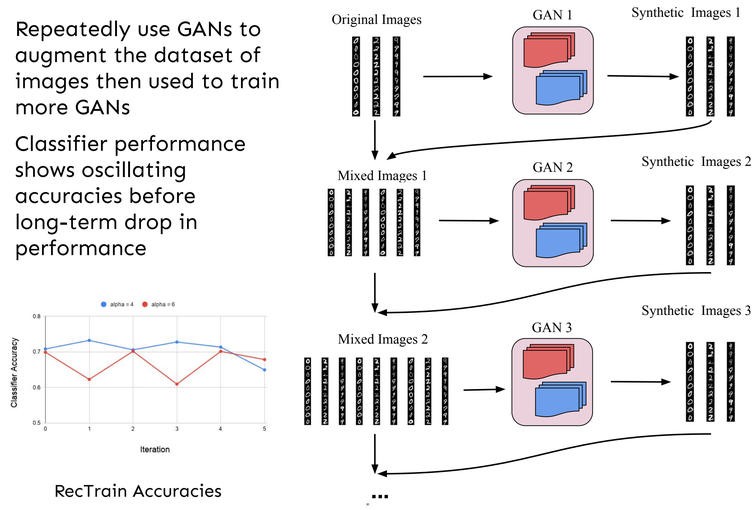

RecTrain Accuracies

Model Architectures

GAN Loss is like a Two-Player Game:

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z\sim p_z(z)}[\log(1 - D(G(z)))]$$

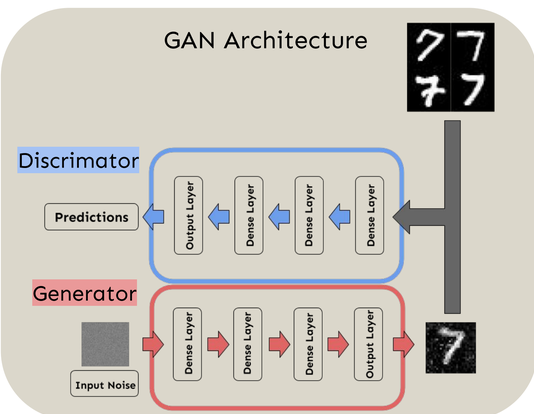

The classifier used cross entropy loss with regularization

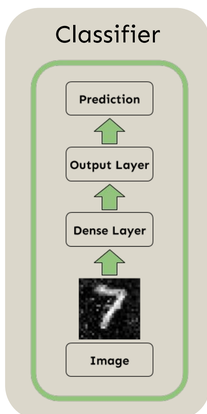

Summary and Future Work

- Summary:
  - We achieved comparable performance when training only on GAN-generated data and significant performance increases when using GAN-generated data and real data.
    - Adding GAN generated data can be more beneficial than adding more original data, and leads to more stability in training
  - Recursive training of GANs failed to yield performance increase
- Future work:
  - More fine tuning of hyperparameters when training GANs
  - Exploring other classifier architectures and generative models
  - More complex image classification tasks, ex. CIFAR 100

References:
[1] Fabio Henrique Kiyoiti dos Santos Tanaka and Claus Aranha. Data Augmentation Using GANs. In *Proceedings of Machine Learning Research*, 2019.
[2] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei Efros. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *International Conference on Computer Vision*, 2017.
[3] Veit Sandfort, Ke Yan, Perry Pickhardt, and Ronald Summers. Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. In *Nature Research*, 2019.
[4] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and Improving the Image Quality of StyleGAN. 12 2019.
[5] Tamar Rott Shaham, Tali Dekel, and Tomar Michaeli. SinGAN, Learning a Generative Model from a Single Natural Image. In *International Conference on Computer Vision*, 2019.
[6] Jiajun Shen, Peng-Jen Chen, Matt Le, Junxian He, Jiatao Gu, Myle Ott, Michael Auli, and Marc'Aurelio Ranzato. The source-target domain mismatch problem in machine translation. 09 2019.
[7] Ian Goodfellow et al. Generative Adversarial Nets. In *Conference on Neural Information Processing Systems*, 2014.
[8] Adam Byerly, Tatiana Kalganova, and Ian Dear. A Branching and Merging Convolutional Network with Homogeneous Filter Capsules. In arXiv, 2019.

Introduction

- Generative Adversarial Networks (GANs) are powerful generative models introduced by (Goodfellow et al.) [7] and can be trained on as little data as a single image [5].
- Lack of data makes ML hard -- data augmentation
- Prior work:
  - "Translating" images [3]
  - Generating numeric data [1]
- *Motivation: Explore using these super powerful generative models to augment more complex data sets*
- Classic Problem: Image classification of numbers in the MNIST database.

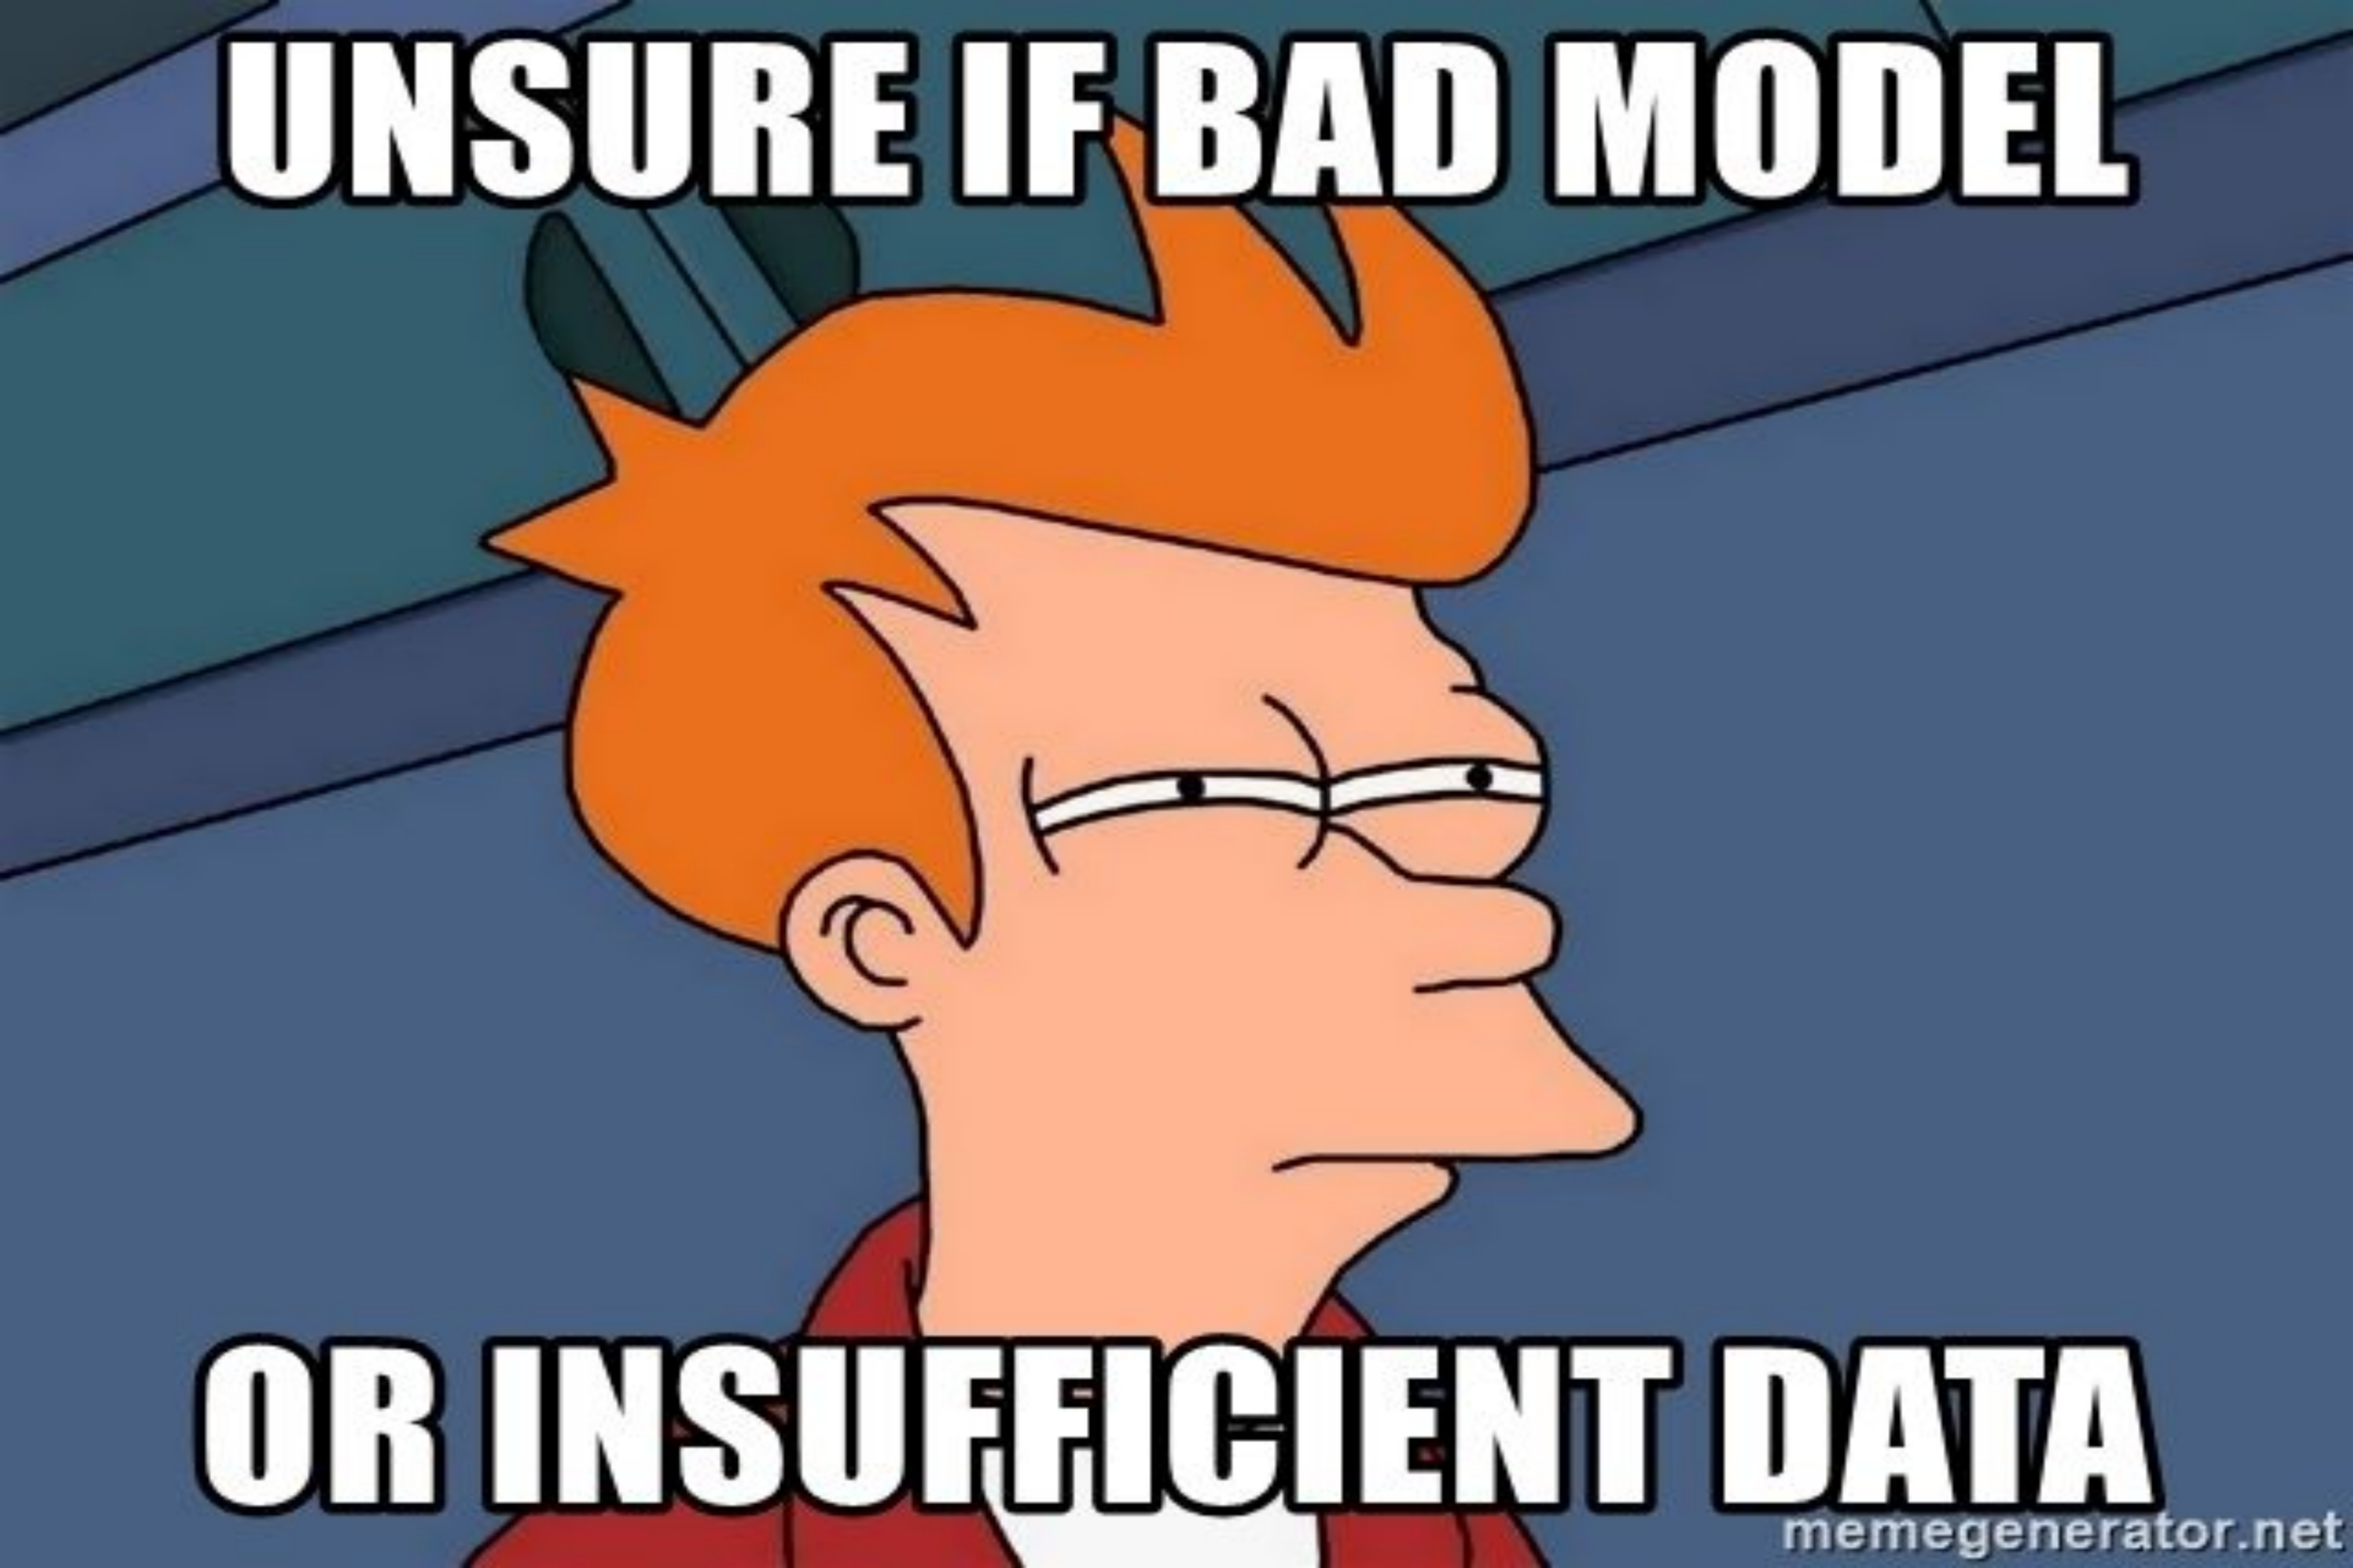
UNSURE IF BAD MODEL
OR INSUFFICIENT DATA
memegenerator.net

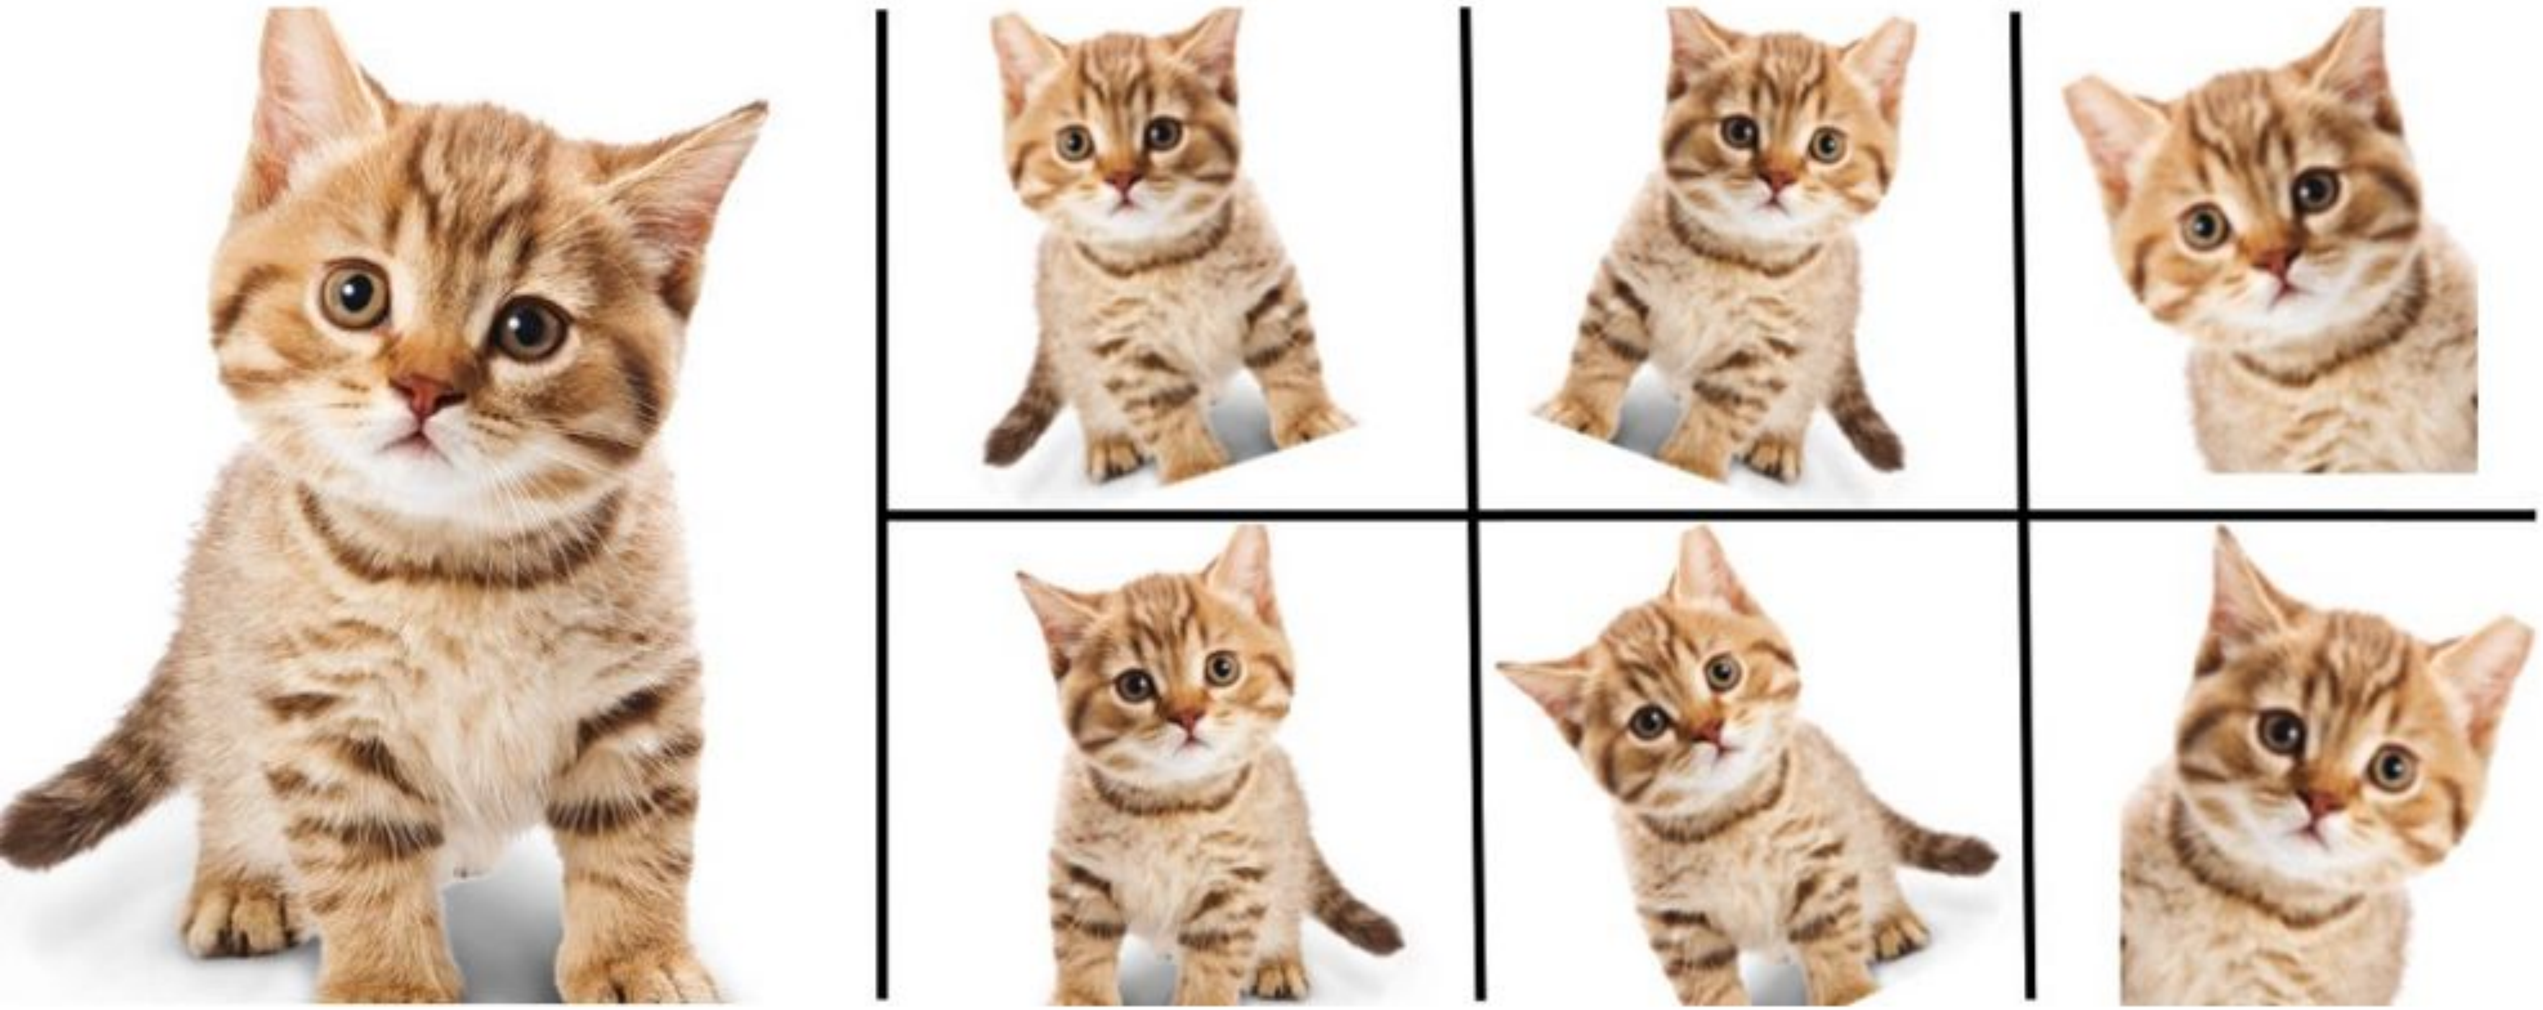

Enlarge your Dataset

Introduction

- Generative Adversarial Networks (GANs) are powerful generative models introduced by (Goodfellow et al.) [7] and can be trained on as little data as a single image [5].
- Lack of data makes ML hard -- data augmentation
- Prior work:
  - "Translating" images [3]
  - Generating numeric data [1]
- *Motivation: Explore using these super powerful generative models to augment more complex data sets*
- Classic Problem: Image classification of numbers in the MNIST database.

Model Architectures

GAN Loss is like a Two-Player Game:

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z\sim p_z(z)}[\log(1 - D(G(z)))]$$

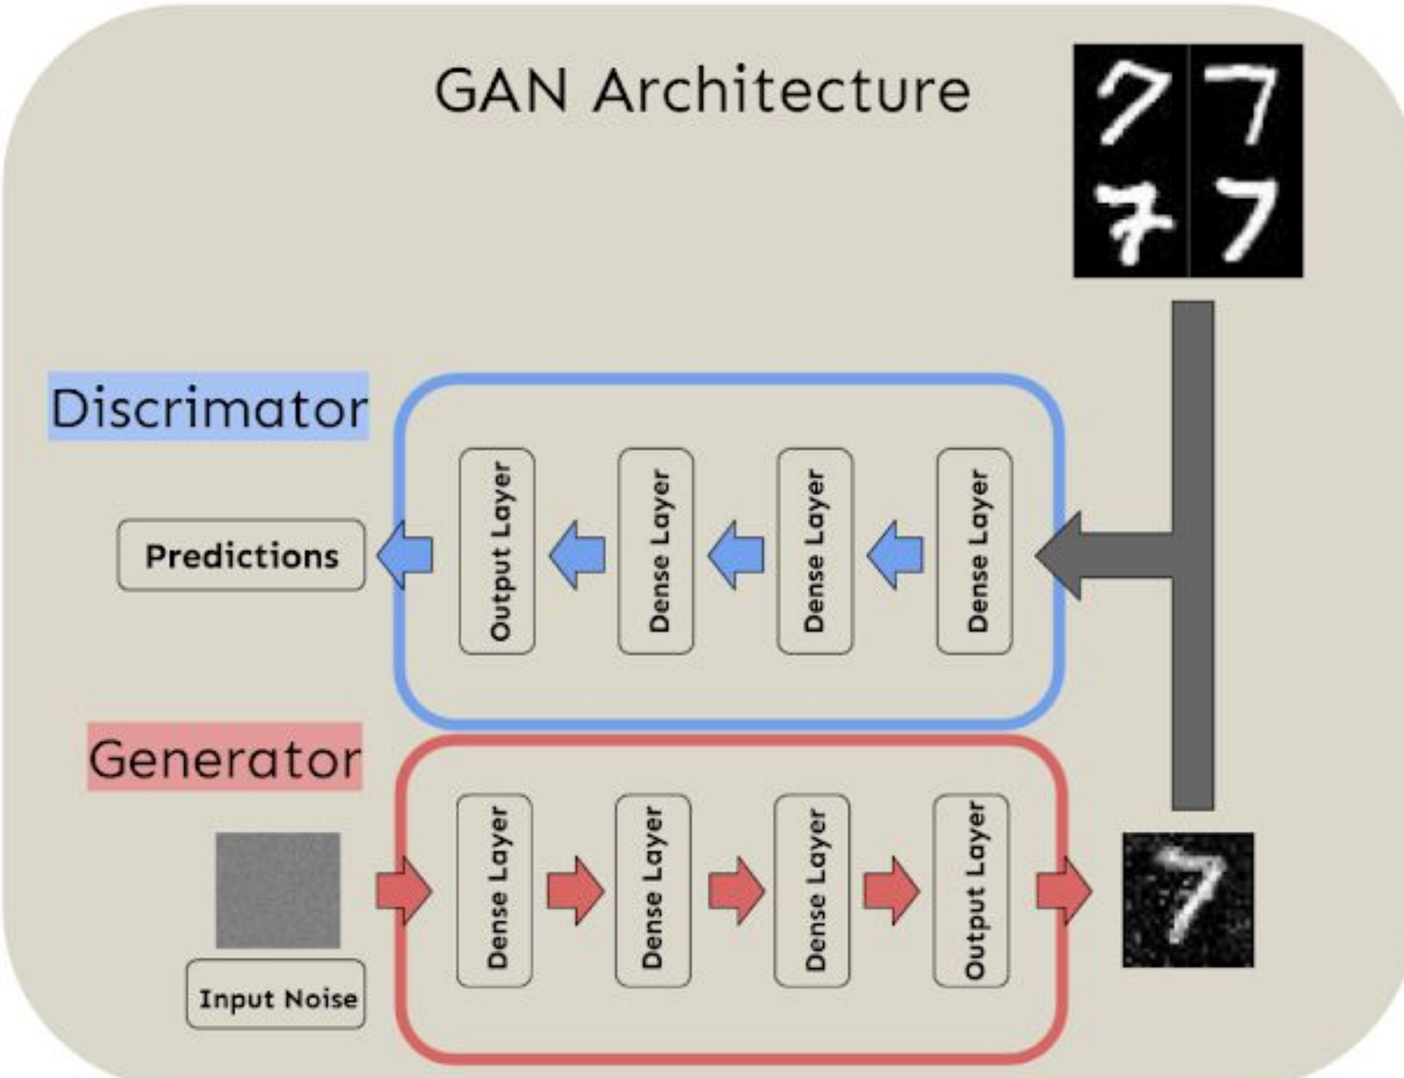

The classifier used cross entropy loss with regularization

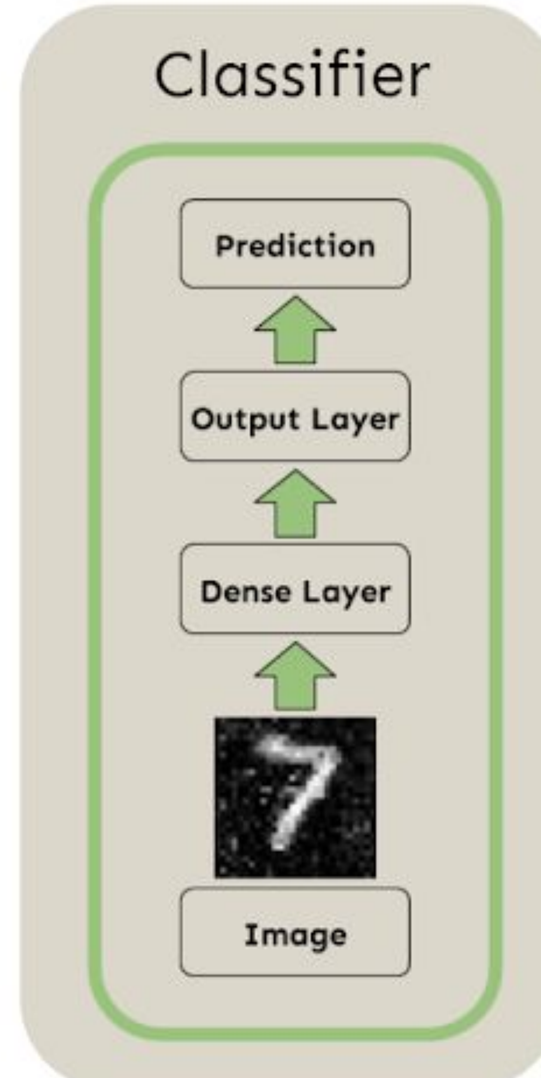

Direct Data Augmentation

- Trained the classifier on purely GAN-generated data for GANs of various sizes
- Pure synthetic data comparable to pure real data in training classifier

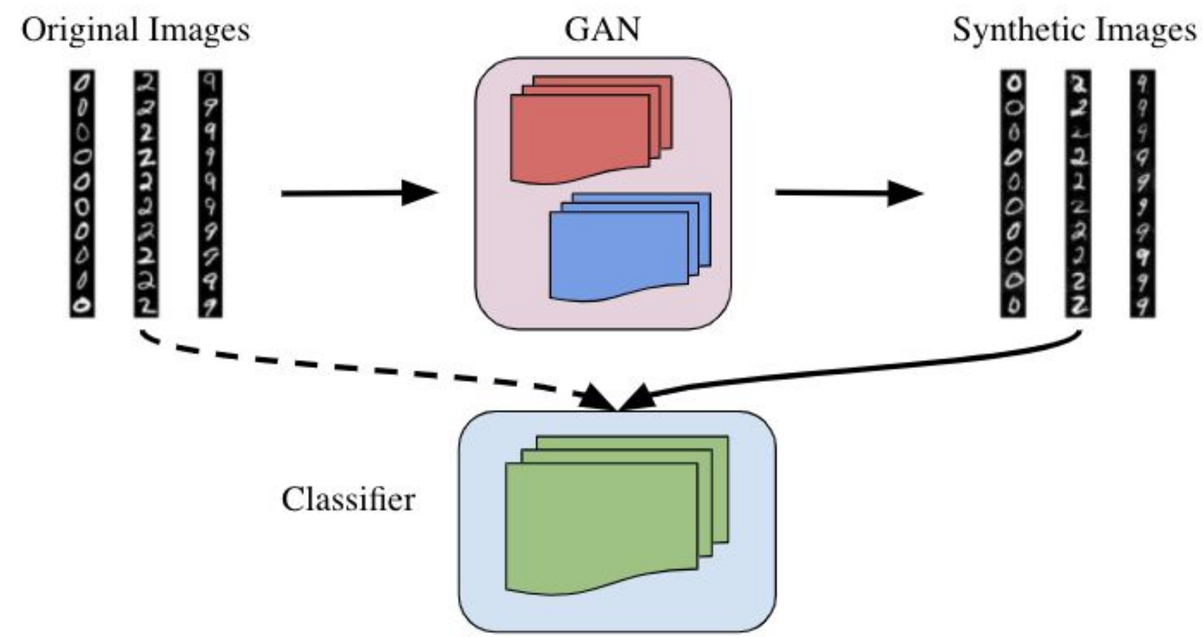

Train Size	Baseline	$\alpha = 1$	$\alpha = 2$	**$\alpha = 4$**	$\alpha = 6$
250	0.641	0.422	0.615	**0.709**	0.698
500	0.648	0.611	0.741	**0.763**	0.710
1000	0.683	0.670	0.694	**0.756**	0.738
2000	0.788	0.687	0.680	**0.793**	0.781

- Trained classifier on mixed synthetic data + real data in various ratios
- Mixed data outdoes pure real data; more noticeable for small datasets
- Unstable training losses suggest higher variance in real data

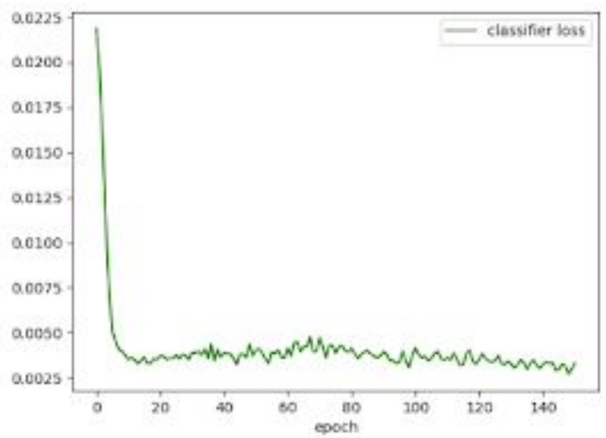

500 Real, 500 Synthetic

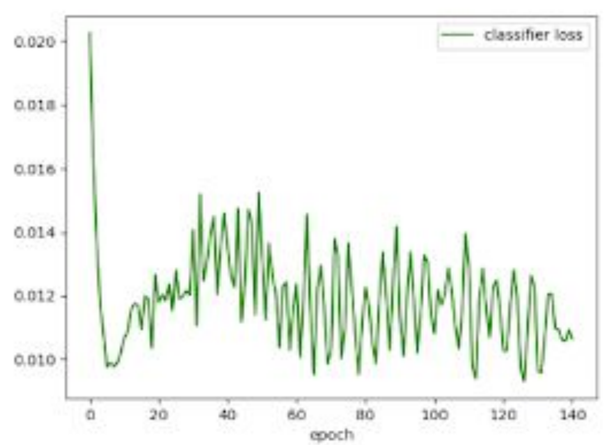

1000 Real

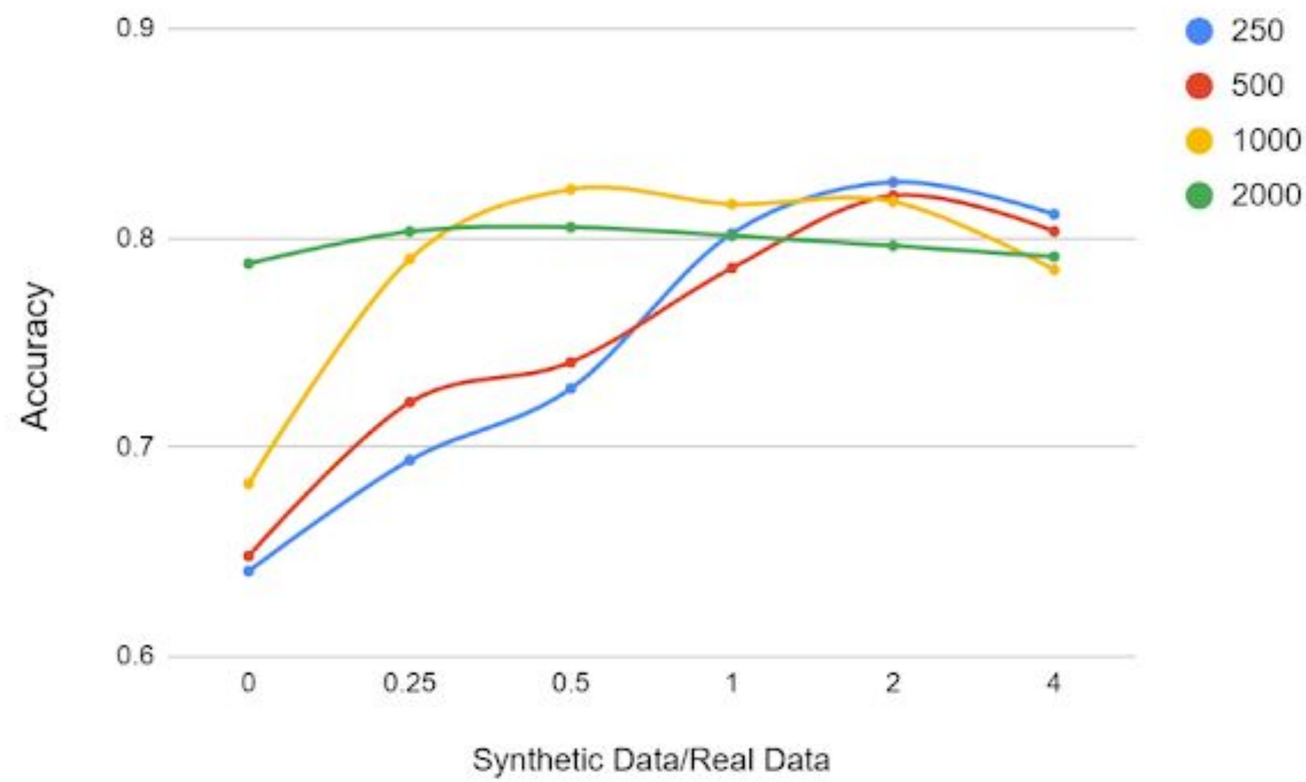

Recursive GAN Training

- Repeatedly use GANs to augment the dataset of images then used to train more GANs
- Classifier performance shows oscillating accuracies before long-term drop in performance

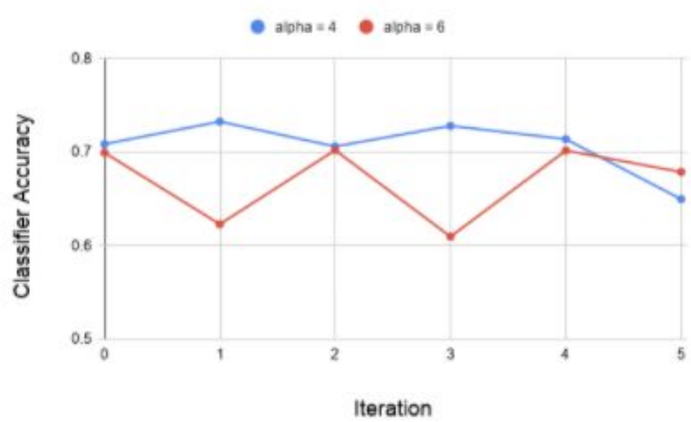

RecTrain Accuracies

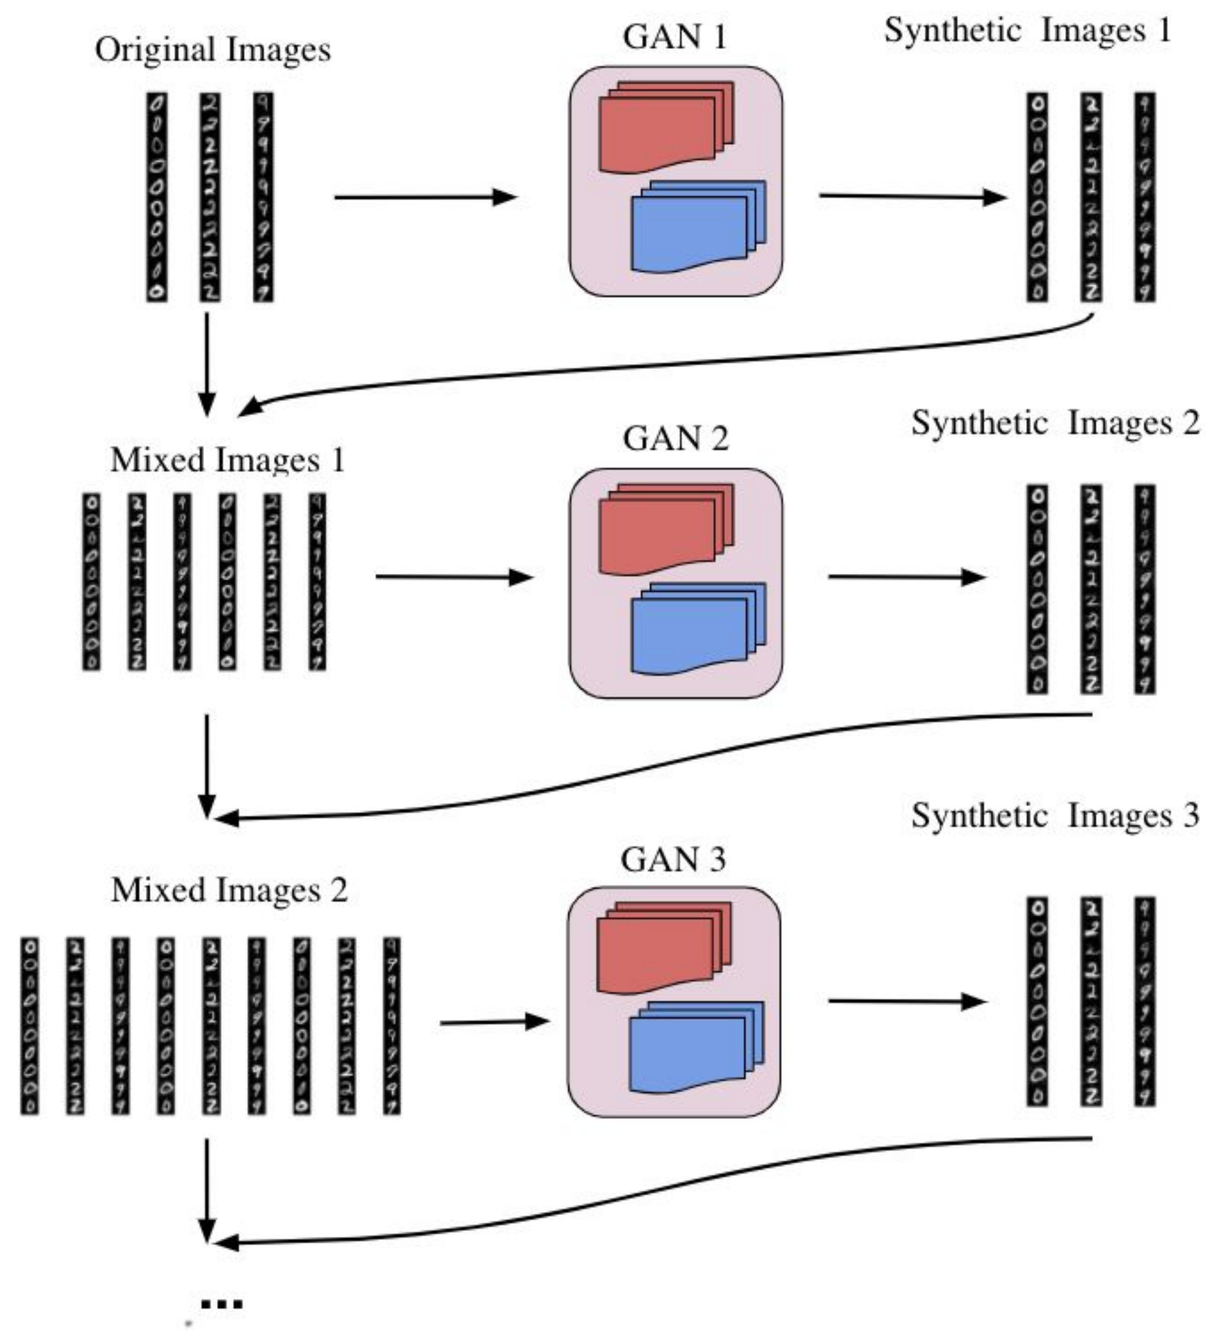

Summary and Future Work

- Summary:
  - We achieved comparable performance when training only on GAN-generated data and significant performance increases when using GAN-generated data and real data.
    - Adding GAN generated data can be more beneficial than adding more original data, and leads to more stability in training
  - Recursive training of GANs failed to yield performance increase

- Future work:
  - More fine tuning of hyperparameters when training GANs
  - Exploring other classifier architectures and generative models
  - More complex image classification tasks, ex. CIFAR 100

References:
[1] Fabio Henrique Kiyoiti dos Santos Tanaka and Claus Aranha. Data Augmentation Using GANs. In *Proceedings of Machine Learning Research*, 2019.
[2] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei Efros. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *International Conference on Computer Vision*, 2017.
[3] Veit Sandfort, Ke Yan, Perry Pickhardt, and Ronald Summers. Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. In *Nature Research*, 2019.
[4] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and Improving the Image Quality of StyleGAN. 12 2019.
[5] Tamar Rott Shaham, Tali Dekel, and Tomar Michaeli. SinGAN, Learning a Generative Model from a Single Natural Image. In *International Conference on Computer Vision*, 2019.
[6] Jiajun Shen, Peng-Jen Chen, Matt Le, Junxian He, Jiatao Gu, Myle Ott, Michael Auli, and Marc'Aurelio Ranzato. The source-target domain mismatch problem in machine translation. 09 2019.
[7] Ian Goodfellow et al. Generative Adversarial Nets. In *Conference on Neural Information Processing Systems*, 2014.
[8] Adam Byerly, Tatiana Kalganova, and Ian Dear. A Branching and Merging Convolutional Network with Homogeneous Filter Capsules. In arXiv, 2019.

GAN-Based Image Data Augmentation

Nathan Hu
zixia314@stanford.edu

David Liu
dliud@stanford.edu

Introduction

- Generative Adversarial Networks (GANs) are powerful generative models introduced by (Goodfellow et al.) [7] and can be trained on as little data as a single image [5].
- Lack of data makes ML hard -- data augmentation
- Prior work:
  - "Translating" images [3]
  - Generating numeric data [1]
- *Motivation: Explore using these super powerful generative models to augment more complex data sets*
- Classic Problem: Image classification of numbers in the MNIST database.

Direct Data Augmentation

- Trained the classifier on purely GAN-generated data for GANs of various sizes
- Pure synthetic data comparable to pure real data in training classifier

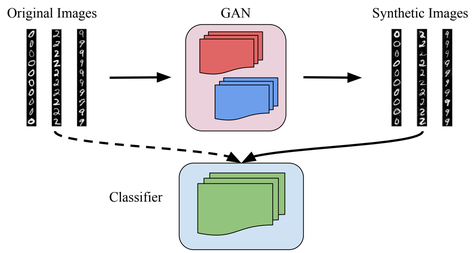

Train Size	Baseline	α = 1	α = 2	**α = 4**	α = 6
250	0.641	0.422	0.615	**0.709**	0.698
500	0.648	0.611	0.741	**0.763**	0.710
1000	0.683	0.670	0.694	**0.756**	0.738
2000	0.788	0.687	0.680	**0.793**	0.781

- Trained classifier on mixed synthetic data + real data in various ratios
- Mixed data outdoes pure real data; more noticeable for small datasets
- Unstable training losses suggest higher variance in real data

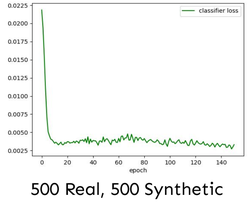

500 Real, 500 Synthetic

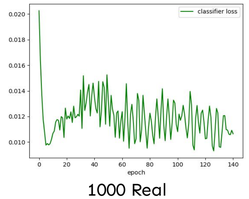

1000 Real

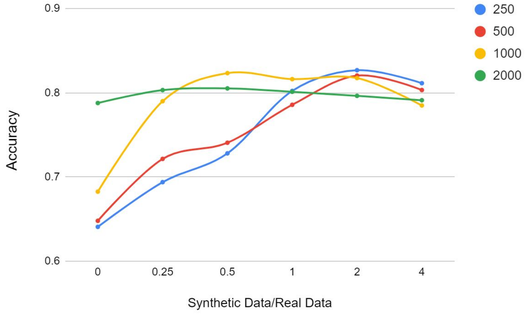

Recursive GAN Training

- Repeatedly use GANs to augment the dataset of images then used to train more GANs
- Classifier performance shows oscillating accuracies before long-term drop in performance

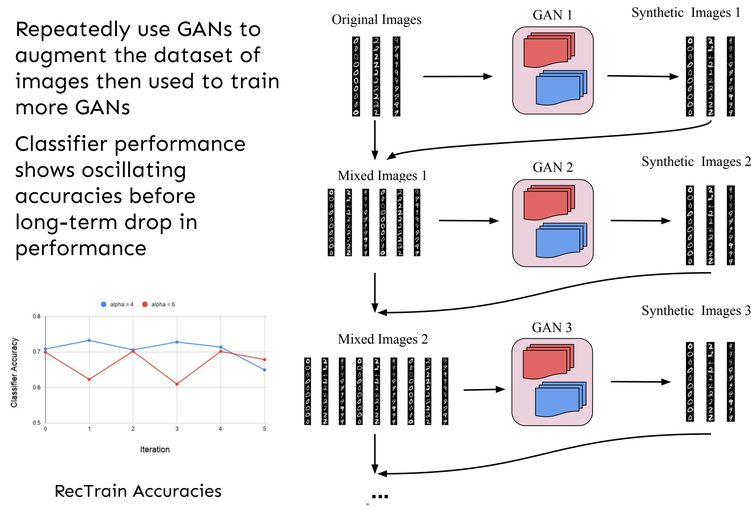

RecTrain Accuracies

Model Architectures

GAN Loss is like a Two-Player Game:

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z\sim p_z(z)}[\log(1 - D(G(z)))]$$

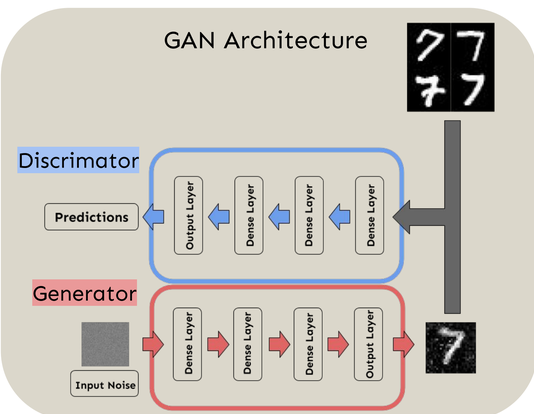

The classifier used cross entropy loss with regularization

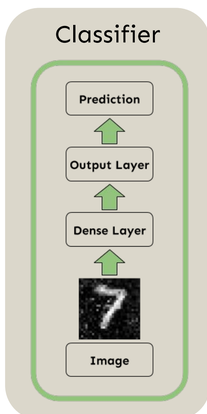

Summary and Future Work

- Summary:
  - We achieved comparable performance when training only on GAN-generated data and significant performance increases when using GAN-generated data and real data.
    - Adding GAN generated data can be more beneficial than adding more original data, and leads to more stability in training
  - Recursive training of GANs failed to yield performance increase
- Future work:
  - More fine tuning of hyperparameters when training GANs
  - Exploring other classifier architectures and generative models
  - More complex image classification tasks, ex. CIFAR 100

References:
[1] Fabio Henrique Kiyoiti dos Santos Tanaka and Claus Aranha. Data Augmentation Using GANs. In *Proceedings of Machine Learning Research*, 2019.
[2] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei Efros. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *International Conference on Computer Vision*, 2017.
[3] Veit Sandfort, Ke Yan, Perry Pickhardt, and Ronald Summers. Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. In *Nature Research*, 2019.
[4] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and Improving the Image Quality of StyleGAN. 12 2019.
[5] Tamar Rott Shaham, Tali Dekel, and Tomar Michaeli. SinGAN, Learning a Generative Model from a Single Natural Image. In *International Conference on Computer Vision*, 2019.
[6] Jiajun Shen, Peng-Jen Chen, Matt Le, Junxian He, Jiatao Gu, Myle Ott, Michael Auli, and Marc'Aurelio Ranzato. The source-target domain mismatch problem in machine translation. 09 2019.
[7] Ian Goodfellow et al. Generative Adversarial Nets. In *Conference on Neural Information Processing Systems*, 2014.
[8] Adam Byerly, Tatiana Kalganova, and Ian Dear. A Branching and Merging Convolutional Network with Homogeneous Filter Capsules. In arXiv, 2019.