Conditional Generative Adversarial Networks

CAMERON FABBRI

Outline

- Generative Adversarial Networks (GANs)
- Wasserstein GANs
- Conditional GANs
- Experiments
- Future Work

Generative Adversarial Networks (GANs)

- Two player mini-max game
  - Generator **G**
  - Discriminator **D**
- D is trained to discriminate between a real image and a generated image.
- G is trained to fool D

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

[1]

Generative Adversarial Networks

$z = \mathcal{N}(-1, 1)$

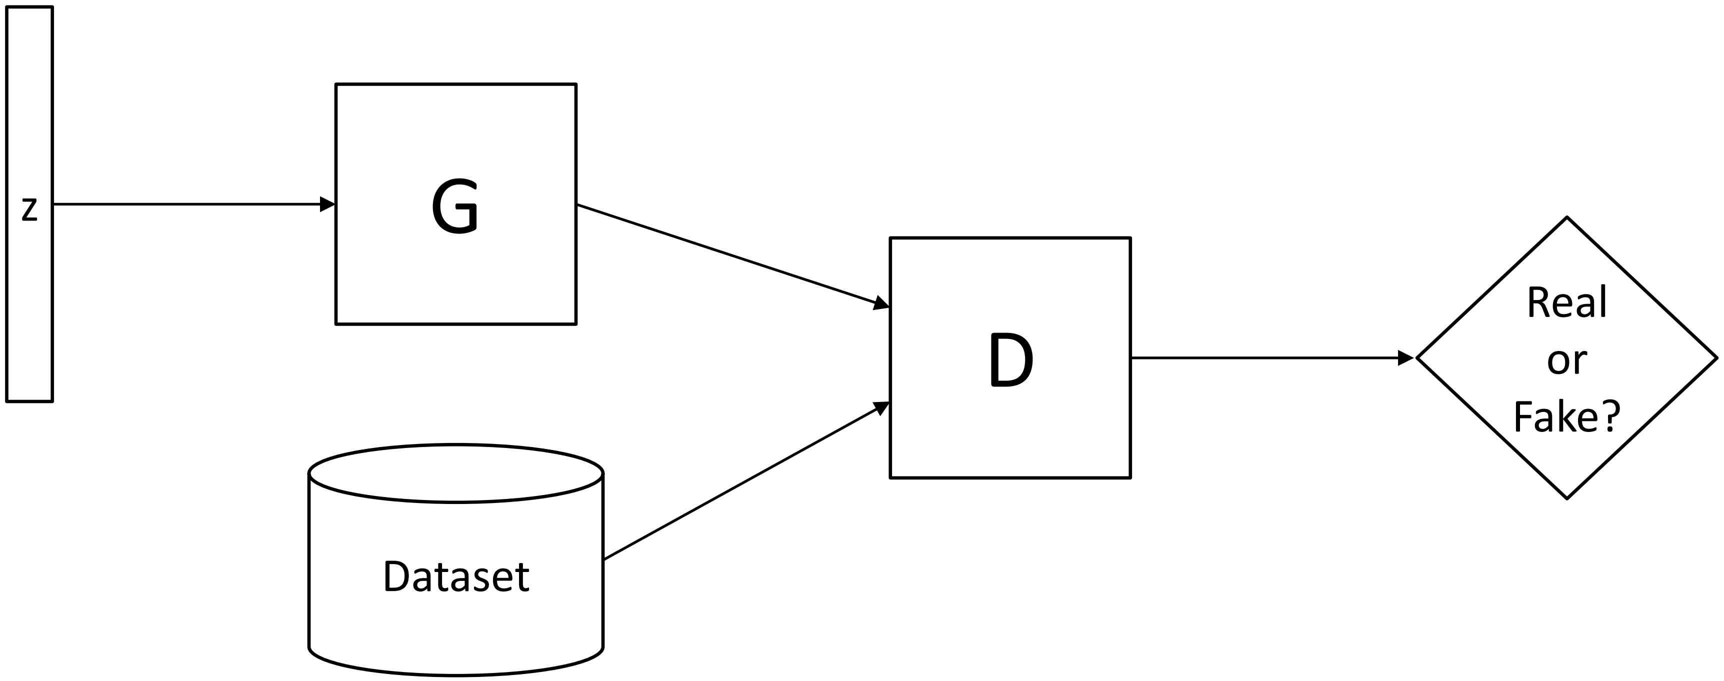

Deep Convolutional GANs (DCGANs)

- Developed an architecture combining GANs and deep learning

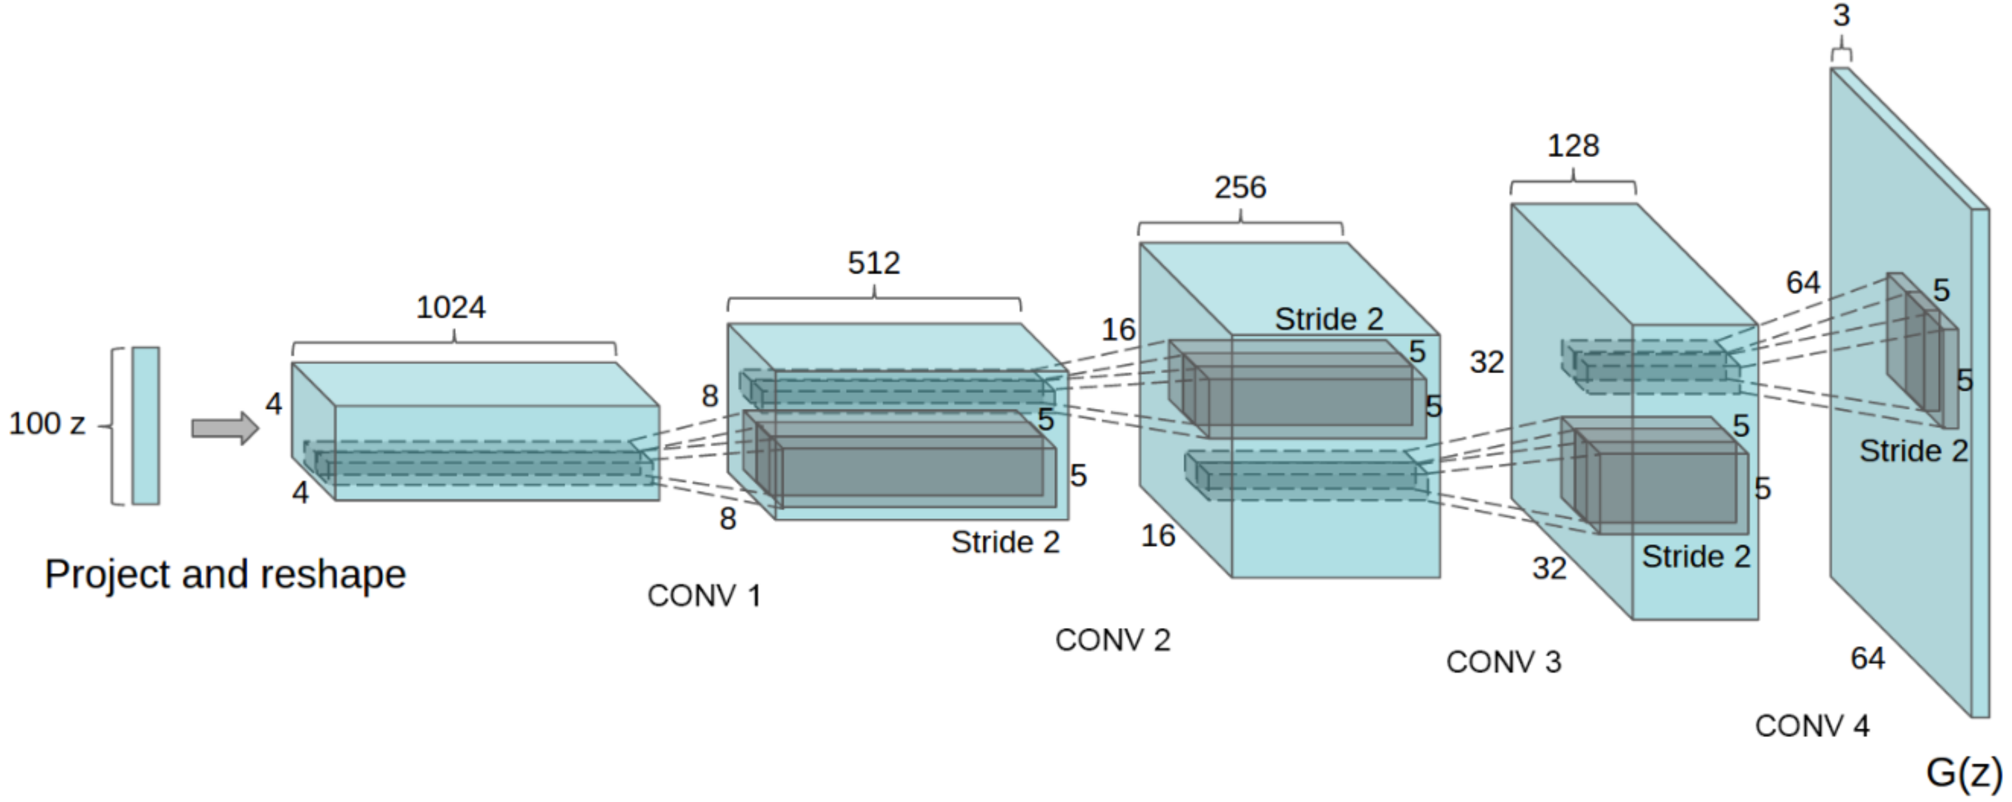

[2]

Wasserstein GANs (WGAN)

- The original GAN used D as a classifier, minimizing the Jensen-Shannon Divergence:

$$JS(\mathbb{P}_r, \mathbb{P}_g) = KL(\mathbb{P}_r || \mathbb{P}_m) + KL(\mathbb{P}_g || \mathbb{P}_m)$$

- The Wasserstein GAN applies the Kantorovich-Rubinstein duality to approximate the Wasserstein-1 distance:

$$\min_G \max_{D \in \mathcal{D}} \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathbb{P}_r} \big[D(\boldsymbol{x}) \big] - \mathop{\mathbb{E}}_{\tilde{\boldsymbol{x}} \sim \mathbb{P}_g} \big[D(\tilde{\boldsymbol{x}})) \big]$$

[3]

Improved Wasserstein GANs (WGAN-GP)

- A differentiable function is 1-Lipschitz if and only if it has gradients with norm less than or equal to 1 everywhere
  - Idea: Enforce the Lipschitz-1 constraint by penalizing the gradient norm of the discriminator with respect to the input – not tractable.
  - Instead penalize the square distance of samples from 1.

$$L = \underbrace{\underset{\tilde{\boldsymbol{x}} \sim \mathbb{P}_g}{\mathbb{E}} \left[D(\tilde{\boldsymbol{x}}) \right] - \underset{\boldsymbol{x} \sim \mathbb{P}_r}{\mathbb{E}} \left[D(\boldsymbol{x}) \right]}_{\text{Original critic loss}} + \underbrace{\lambda \underset{\hat{\boldsymbol{x}} \sim \mathbb{P}_{\hat{\boldsymbol{x}}}}{\mathbb{E}} \left[\left(\|\nabla_{\hat{\boldsymbol{x}}} D(\hat{\boldsymbol{x}})\|_2 - 1 \right)^2 \right]}_{\text{Our gradient penalty}} .$$

[4]

Conditional GANs (cGANs)

- Introduce some additional information into the generator and discriminator.

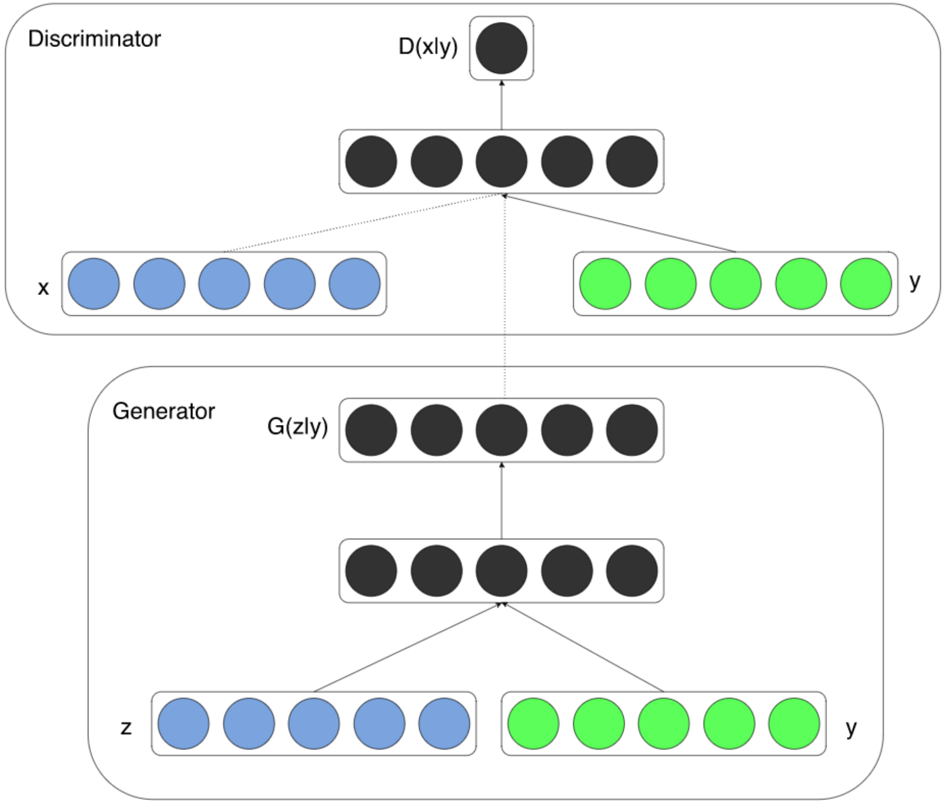

[2]

Experiments

- Implement the Improved Wasserstein method in the conditional GAN setting.

- Control digit generation using MNIST

- Alter visual attributes in faces using the CelebA dataset

- Explore the latent space

MNIST Experiments

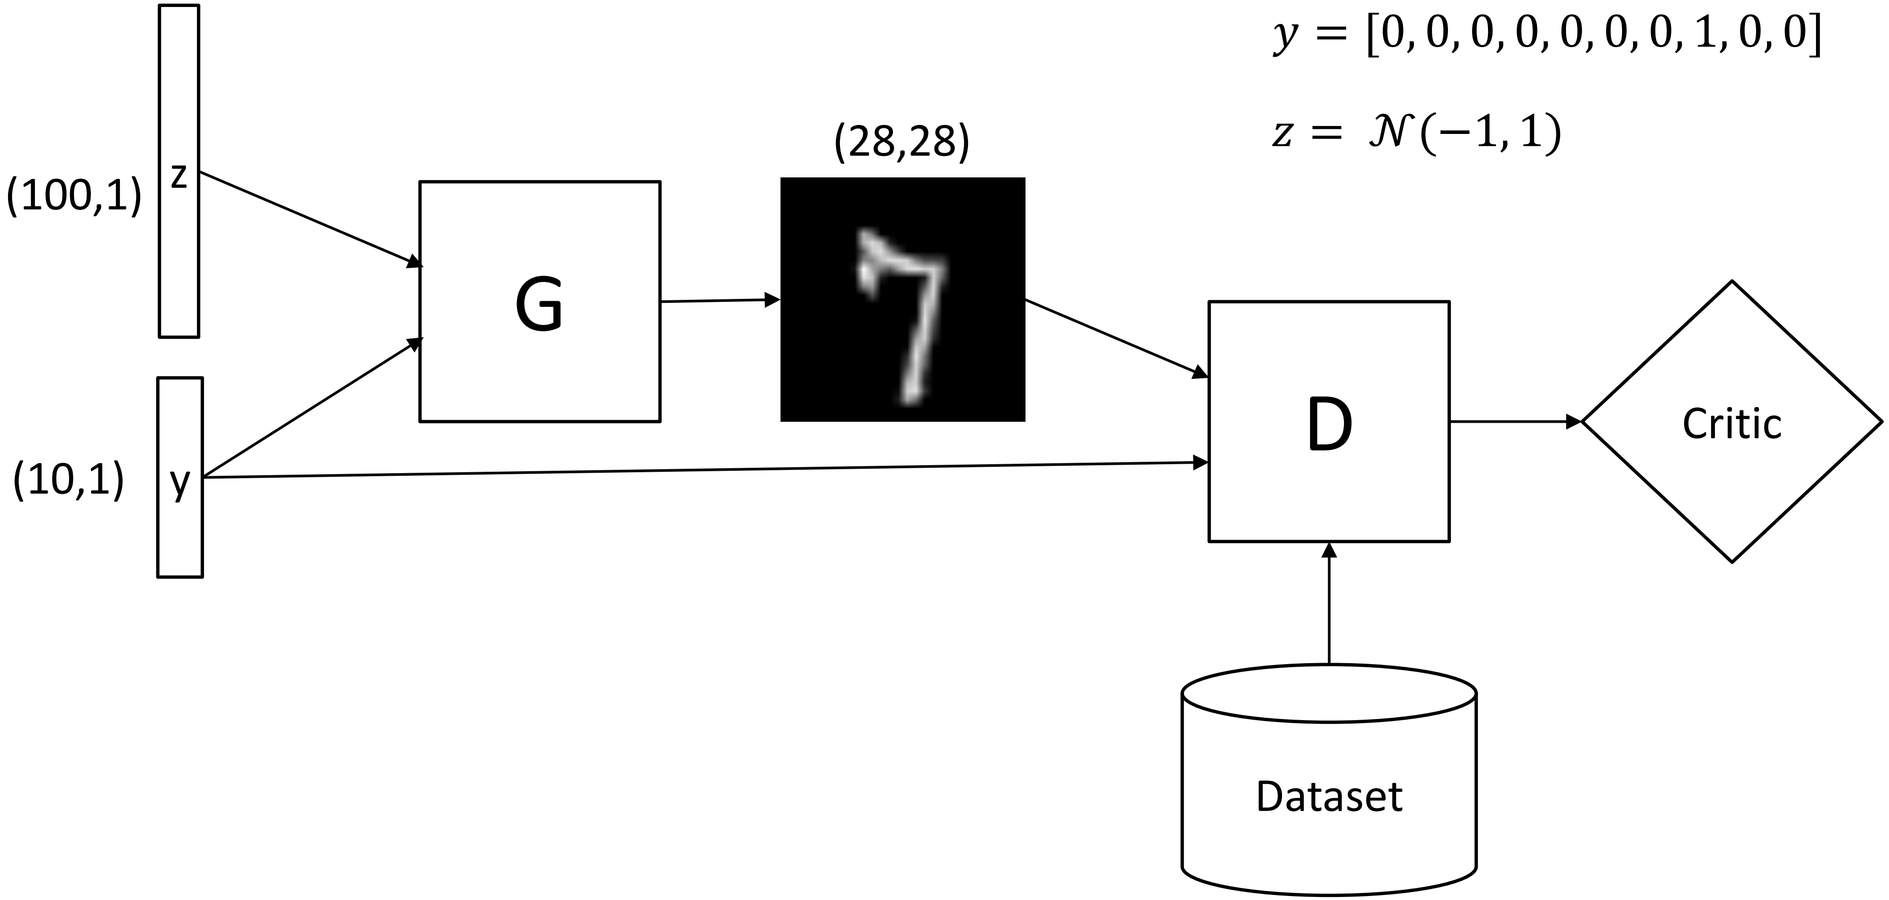

$y = [0, 0, 0, 0, 0, 0, 0, 1, 0, 0]$

$z = \mathcal{N}(-1, 1)$

Label Alteration

- Each row has the same *z* vector
- Each column has the same *y* vector

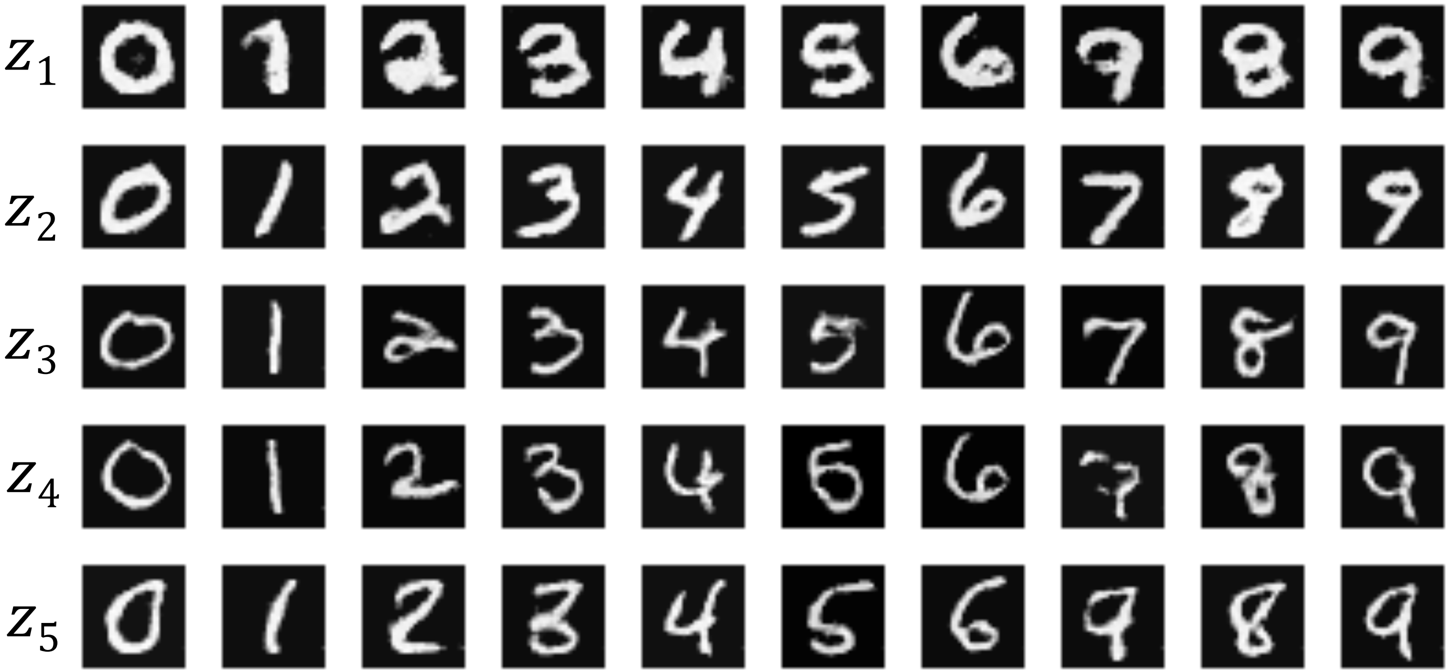

Interpolation

- Interpolate between both *z* and *y* vectors

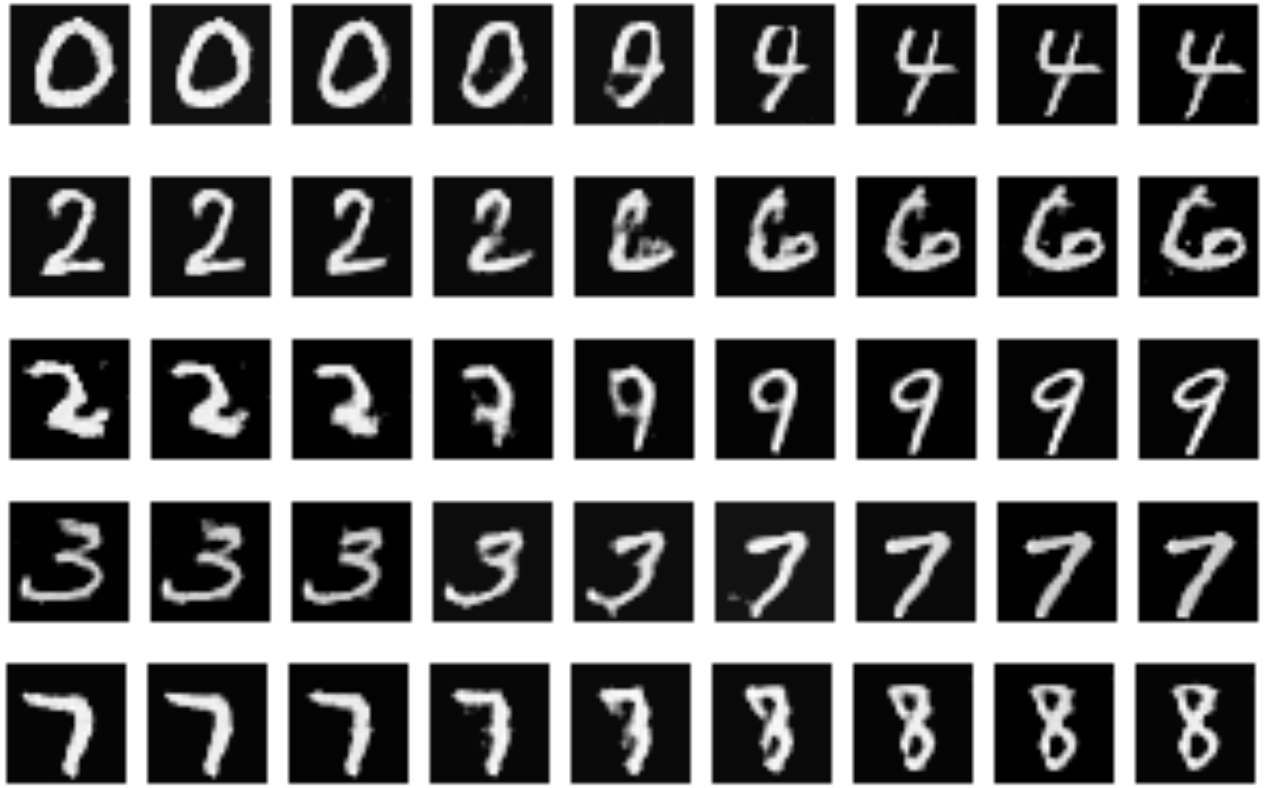

CelebA Experiments

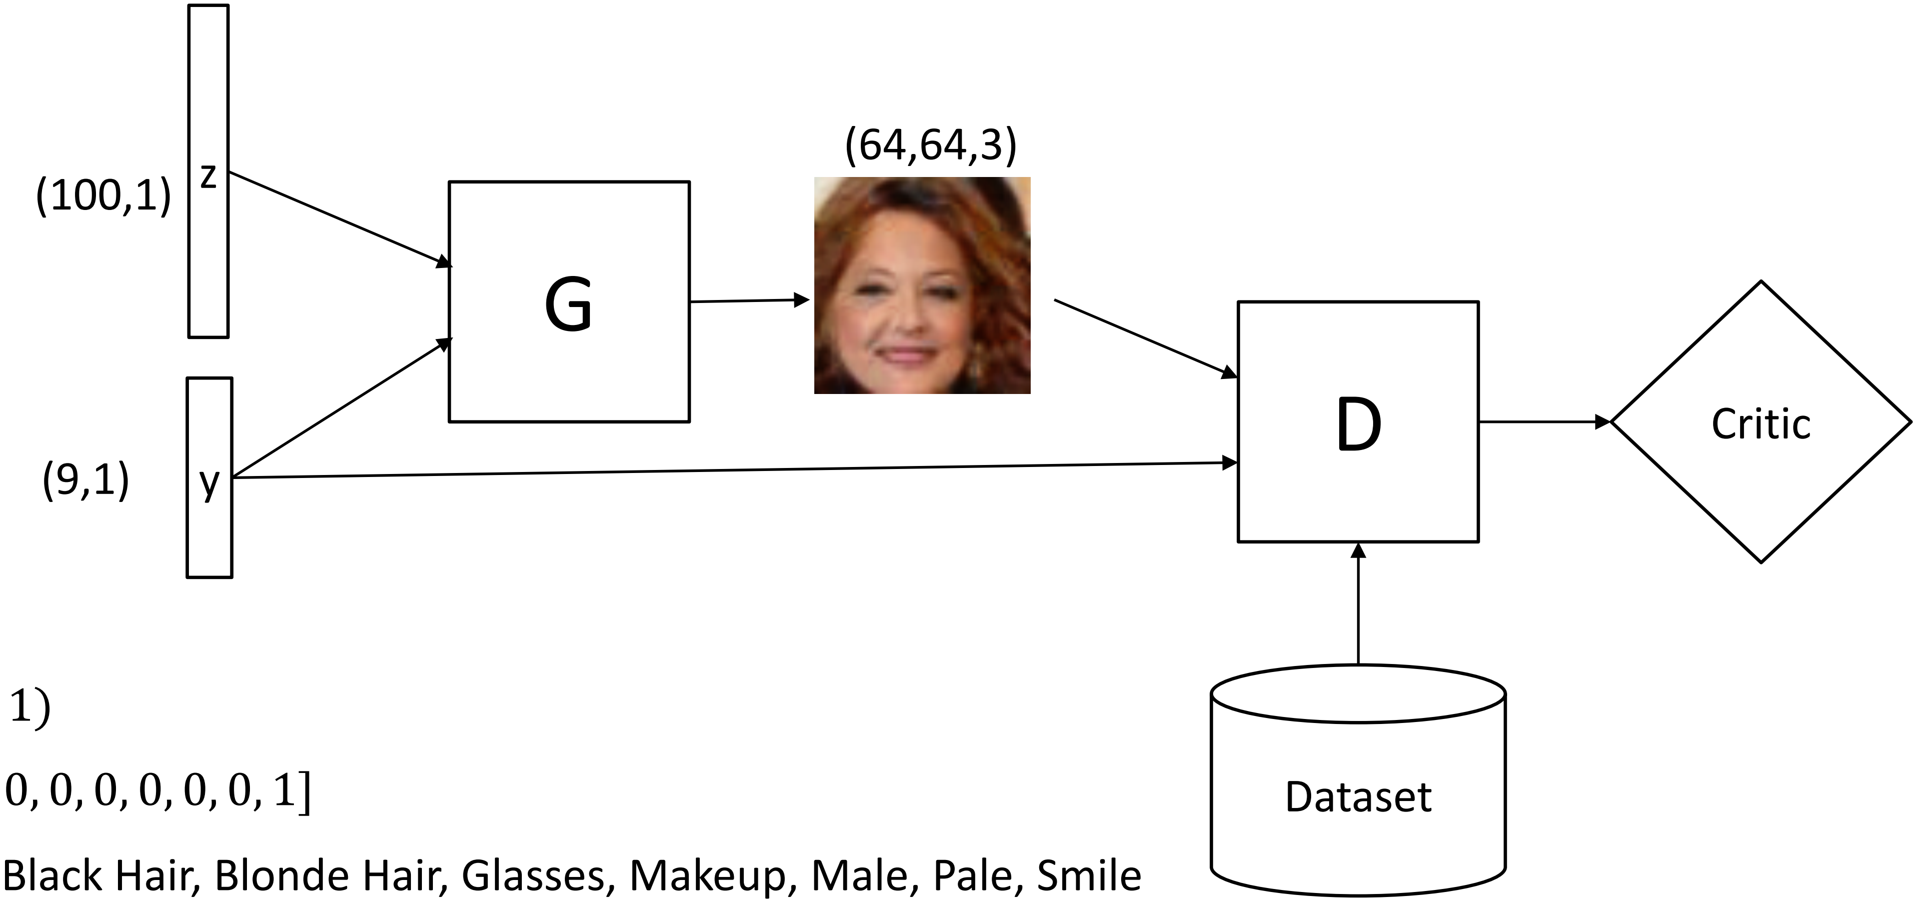

$z = \mathcal{N}(-1, 1)$

$y = [0, 0, 0, 0, 0, 0, 0, 0, 0, 1]$

Bald, Bangs, Black Hair, Blonde Hair, Glasses, Makeup, Male, Pale, Smile

Attribute Alteration

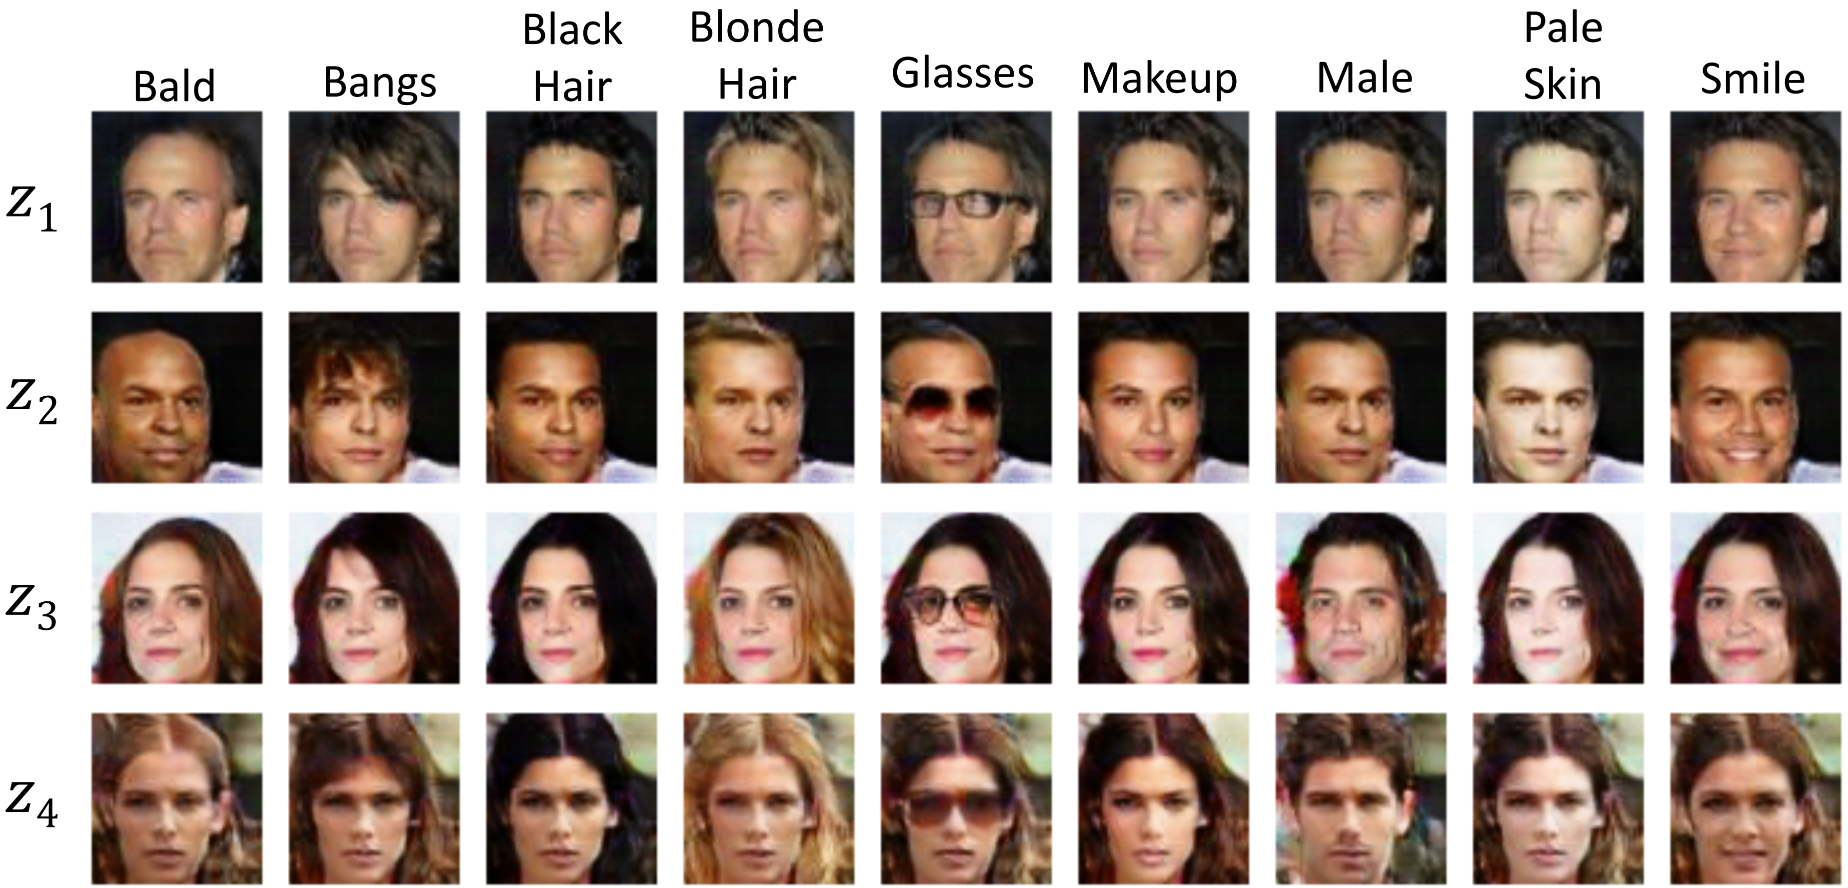

Interpolation

Bald → Not Bald

Black Hair → Blonde Hair

Female → Male

Glasses, Smiling, Bangs →
No Glasses, Not Smiling,
No Bangs, Pale

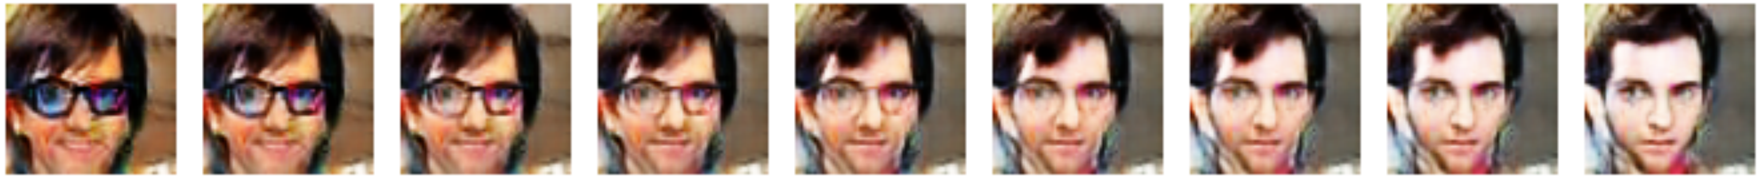

CelebA Experiments using WGAN-GP

Male, Black Hair, Smiling

Female, Smiling

Female, Makeup

Male, Blonde Hair

CelebA Experiments using GAN

Male, Black Hair, Smiling

Female, Smiling

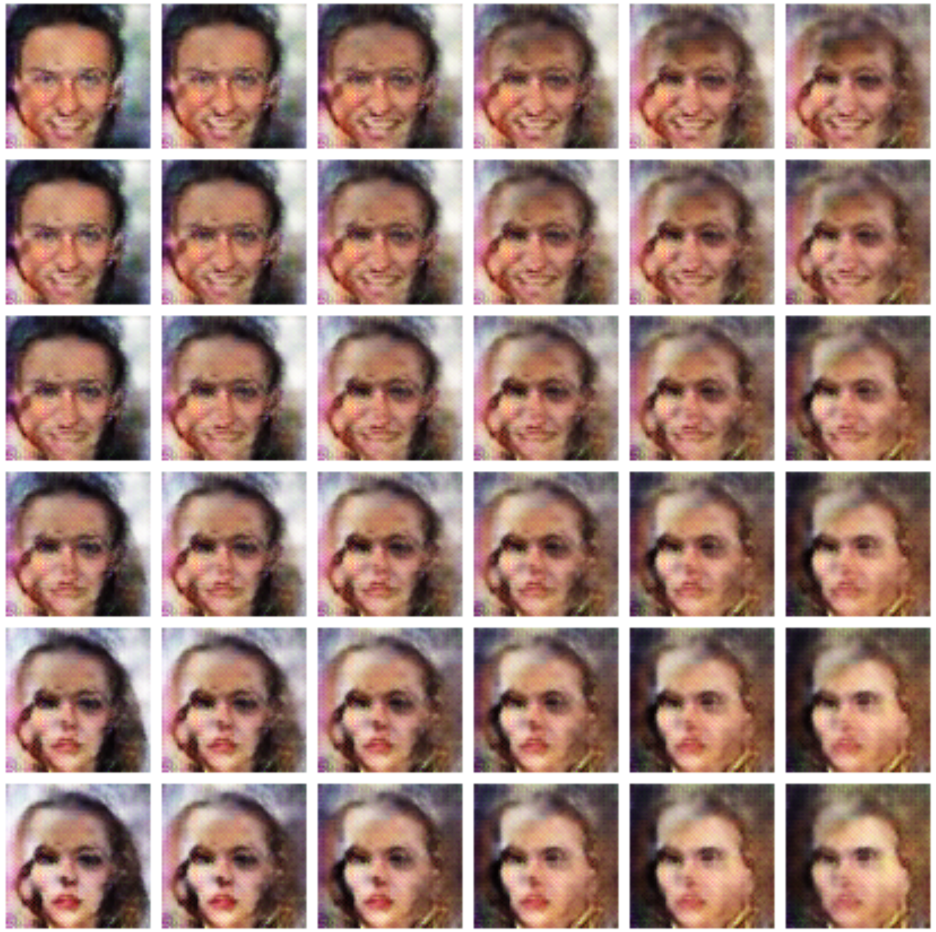

Female, Makeup

Male, Blonde Hair

Conclusion

- Generative Adversarial Networks
  - Wasserstein Method
- Conditional GANs
- Experiments on two datasets
- Future Work
  - Try on other datasets
  - Compare with other GAN methods
    - Least Squares GANs
    - Energy Based GANs

References

- [1] Goodfellow, Ian, et al. "Generative adversarial nets." *Advances in neural information processing systems*. 2014.
- [2] Mirza, Mehdi, and Simon Osindero. "Conditional generative adversarial nets." *arXiv preprint arXiv:1411.1784* (2014).
- [3] Arjovsky, Martin, Soumith Chintala, and Léon Bottou. "Wasserstein gan." *arXiv preprint arXiv:1701.07875* (2017).
- [4] Gulrajani, Ishaan, et al. "Improved training of wasserstein gans." *arXiv preprint arXiv:1704.00028* (2017).
- [5] Radford, Alec, Luke Metz, and Soumith Chintala. "Unsupervised representation learning with deep convolutional generative adversarial networks." *arXiv preprint arXiv:1511.06434* (2015).

Thank you

Questions?

Code: https://github.com/cameronfabbri/cWGANs