A Large-Scale Study on Regularization and Normalization in GANs



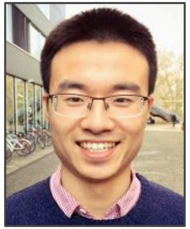

Karol Kurach*
Mario Lucic*
Xiaohua Zhai
Marcin Michalski
Sylvain Gelly

Google Research
Brain Team

Which *model* achieves a given level of *sample quality* the *fastest*?

Which *model* achieves a given level of *sample quality* the *fastest*?

Model

$$\underbrace{\text{Loss} + \text{regularization} + \text{normalization} + \text{neural architectures}}$$

- Non-saturating
- Wasserstein
- Least-squares

- Gradient penalty
- L2 regularization
- Spectral normalization
- Layer normalization
- Batch normalization

- RESNET
- DCGAN/SNDCGAN

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Which *model* achieves a given level of *sample quality* the *fastest*?

Sample quality: measured by

(a) Frechet Inception Distance
(b) Inception Score

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Which *model* achieves a given level of *sample quality* the *fastest*?

Fastest: minimizes the number of hyperparameter settings needed

$\underbrace{\hspace{10em}}$

Reported in the literature + sequential Bayesian optimization

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Experimental protocol

1. Pick a dataset (CelebA-HQ, LSUN Bedrooms, CIFAR10)
2. Pick a model (>15,000)
3. Train it on more than 260 hyperparameter settings
4. Plot the *best score* for a given budget of hyperparameters

Study: Regularization and normalization

1. *Improved Training of Wasserstein GANs*
   Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. C. (2017)

2. *On Convergence and Stability of GANs*
   Kodali, N., Abernethy, J., Hays, J., & Kira, Z (2017)

3. *Spectral Normalization for Generative Adversarial Networks*
   Miyato, T., Kataoka, T., Koyama, M., & Yoshida, Y. (2018)

4. *Layer Normalization*
   Lei Ba, J., Kiros, J. R., & Hinton, G. E. (2016)

5. *Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift*
   Ioffe, S., & Szegedy, C. (2015)

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Results: Regularization and normalization study

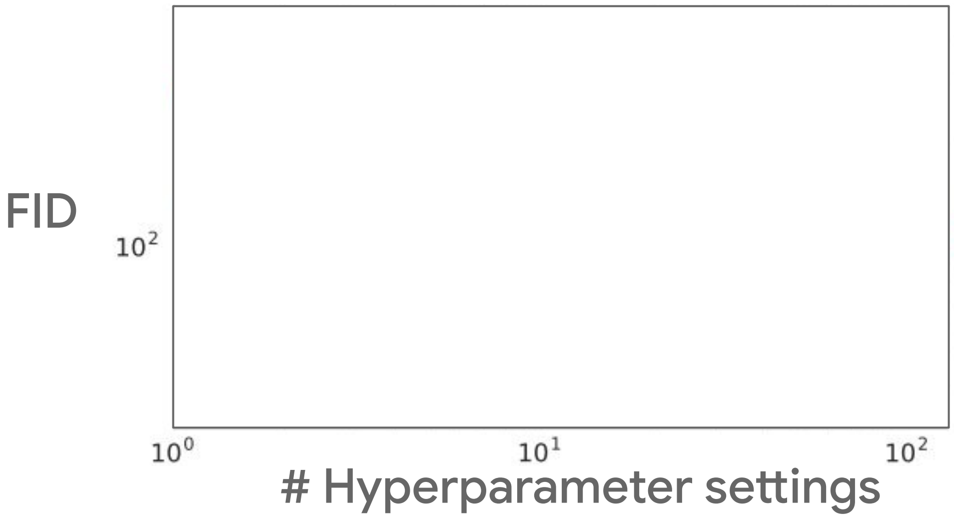

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Results: Regularization and normalization study

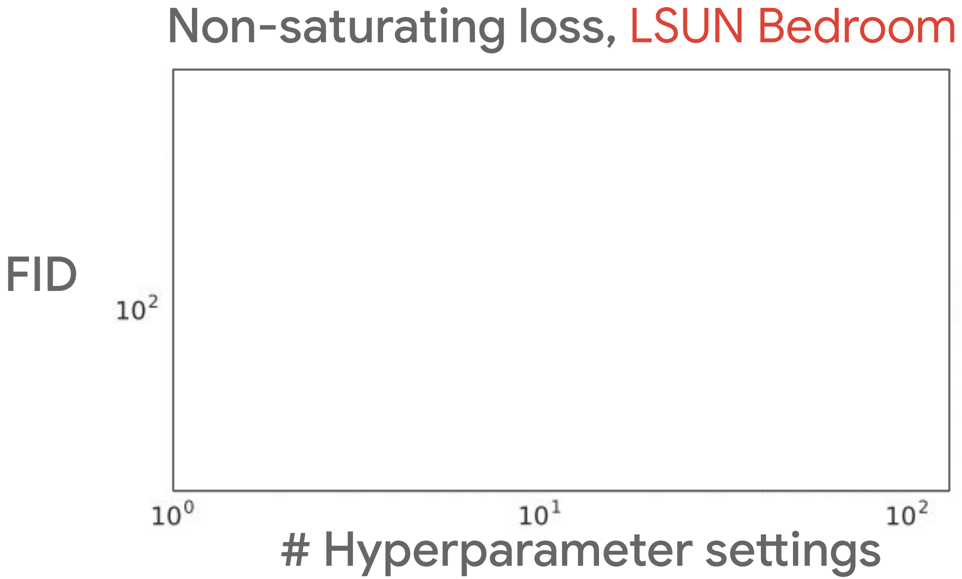

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Results: Regularization and normalization study

Non-saturating loss, LSUN Bedroom

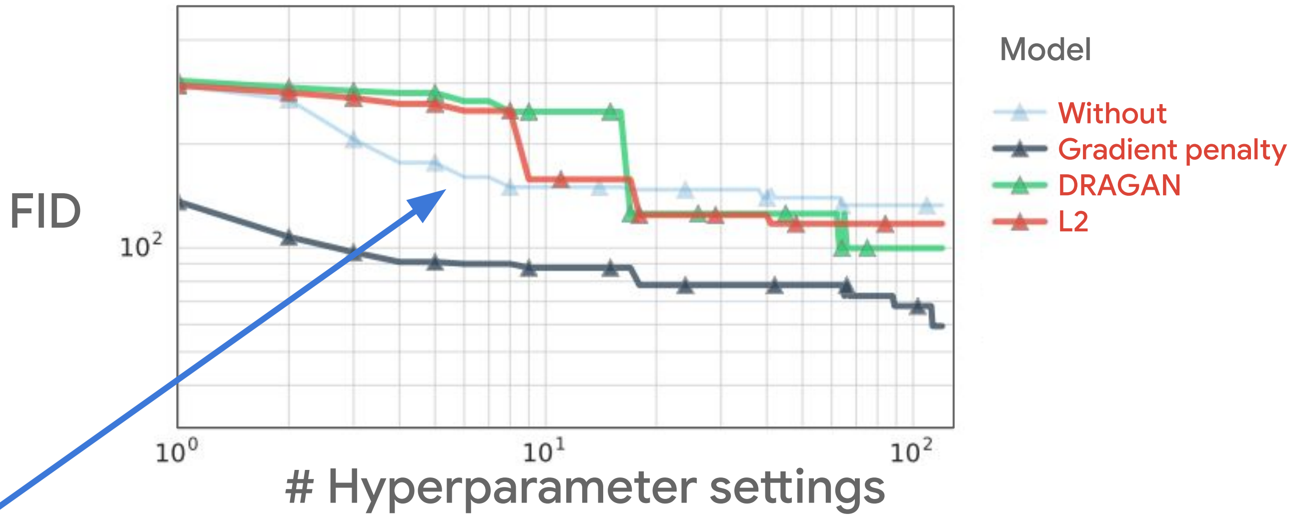

No regularization and normalization

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Results: Regularization and normalization study

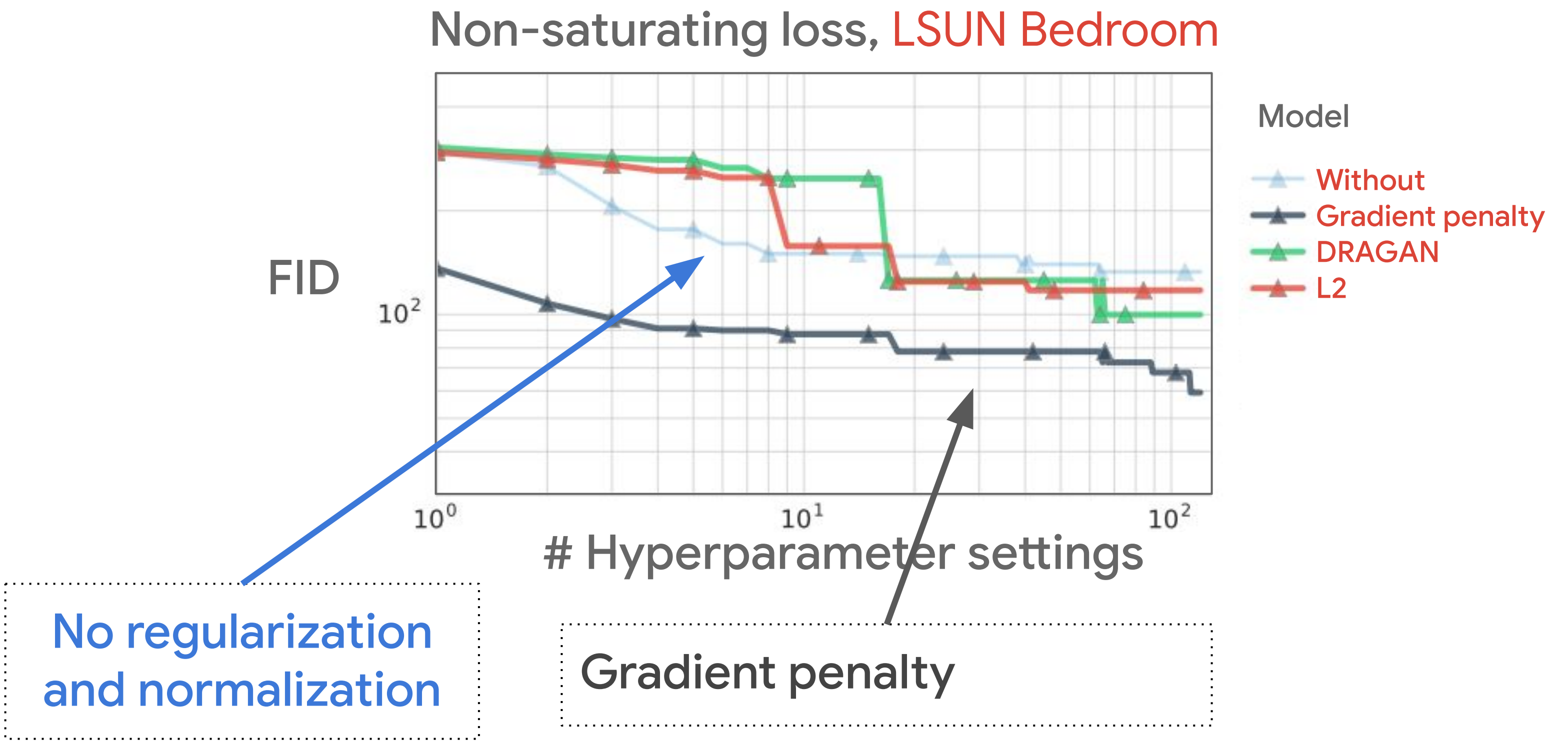

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Results: Regularization and normalization study

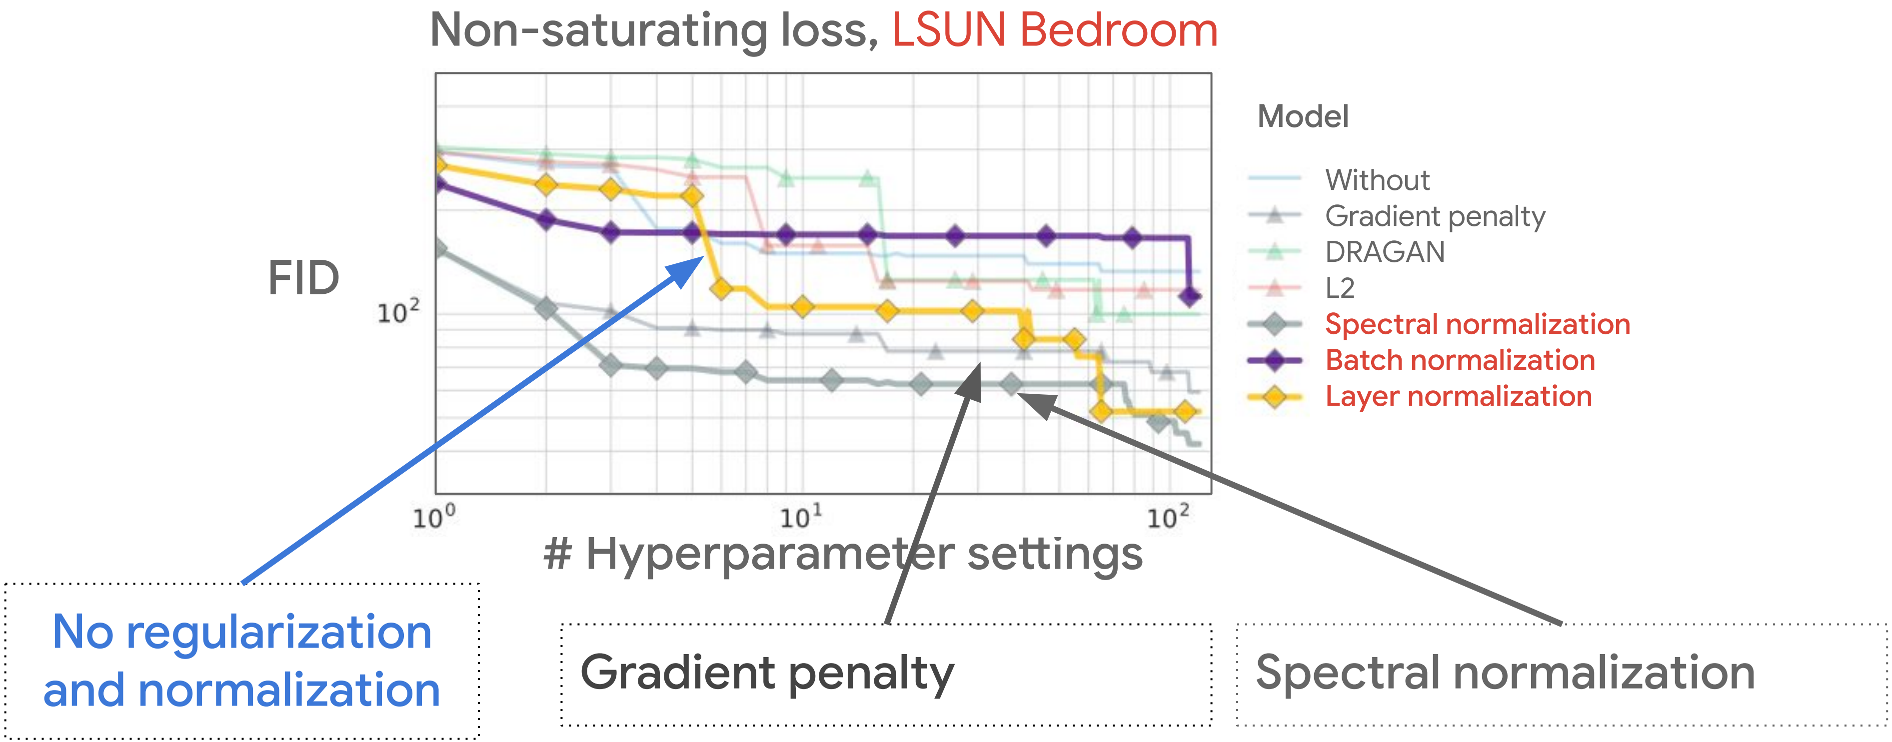

A Large-Scale Study on Regularization and Normalization in GANs, ICML 2019

Results: Regularization and normalization study

1. Optimization remains a key challenge
2. Spectral normalization is the most effective technique
    - No dataset-specific hyperparameter tuning is required
    - Per-iteration overhead is minimal
3. Gradient-based regularization from *Gulrajani et al.*
    - Useful, but necessitates hyperparameter tuning
    - Has a significantly higher per-iteration cost than spectral normalization

Many Paths to Equilibrium: GANs do not Need to Decrease a Divergence at Every Step

Fedus, W., Rosca, M., Lakshminarayanan, B., Dai, A.M., Mohamed, S., & Goodfellow, I. (2017)

Varying losses and neural architectures

1. *Insights transfer across loss functions*
   - *Wasserstein GAN*
     Arjovsky, M., Chintala, S., & Bottou, L. (2017)
   - *Least squares Generative Adversarial Networks.*
     Mao, X., Li, Q., Xie, H., Lau, R. Y., Wang, Z., & Paul Smolley, S. (2017)

2. *Insights transfer across neural architectures*
   - *(SN)-DCGAN*
     Radford, A., Metz, L., & Chintala, S. (2015)
   - *Deep Residual Learning for Image Recognition*
     He, K., Zhang, X., Ren, S., & Sun, J. (2016)

Summary

There is a lot of room for progress -- optimization remains a key challenge

Spectral normalization is currently the most effective technique*

Gradient-based regularization is useful, but necessitates more tuning

Future work

1. Customized architectures (e.g. BigGAN, StyleGAN)
2. Self-attention and self-modulation mechanisms
3. Improved quantitative evaluation measures
   - *Assessing Generative Models via Precision and Recall*
     Sajjadi, M. S., Bachem, O., Lucic, M., Bousquet, O., & Gelly, S. (2018)
   - *Improved Precision and Recall Metric for Assessing Generative Models*
     Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., Aila, T. (2019)

Resources

Code, pretrained models and Colab available at:

github.com/google/compare_gan

Check out our **poster #9** tonight (Jun 12th) 6:30-9:00 pm!