Caltech Implicit Competitive Regularization in GANs

Florian Schäfer, Hongkai Zheng, and Anima Anandkumar

上海交通大學 SHANGHAI JIAO TONG UNIVERSITY

How do we interpret GAN

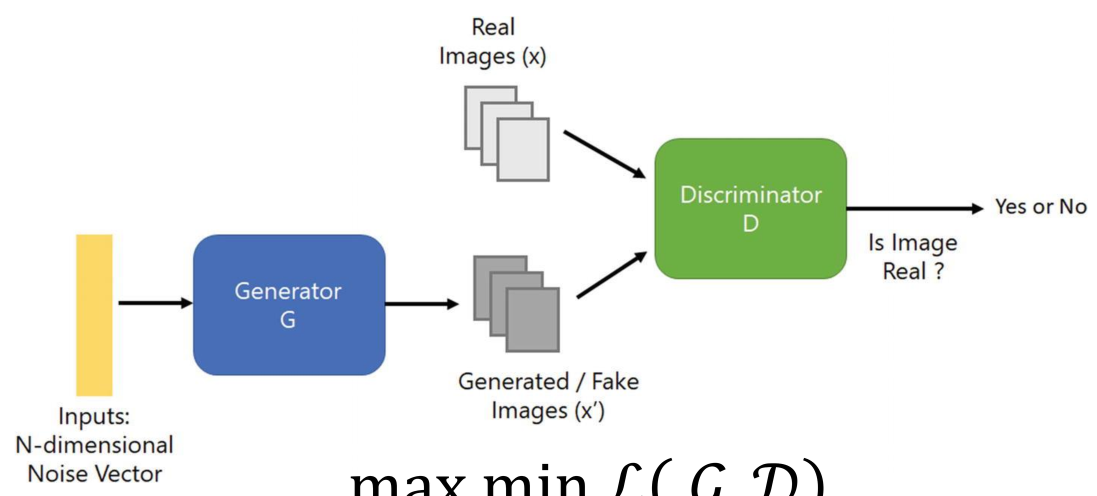

$$\max_{\mathcal{G}} \min_{\mathcal{D}} \mathcal{L}(\mathcal{G}, \mathcal{D})$$

1. **The result of simultaneous gradient descent.**
   - **Problem: not helpful for improving stability**
2. **Minimax problem.**
   - **Problem with disjoint supports of training data and generated data. WGAN needs metric on image space.**

     $\min_{\mathcal{D}} \mathcal{L}(\mathcal{D}, \mathcal{G}) = 2\log 2 - 2JSD(\mathrm{P_r}||\mathrm{P_g})$
3. **Local Nash equilibria of** $\mathcal{L}(\mathcal{G}, \mathcal{D})$.
   - **Good solutions are often not local Nash equilibria. [Berard et al. 2019]**

Why a rational GAN would be stable

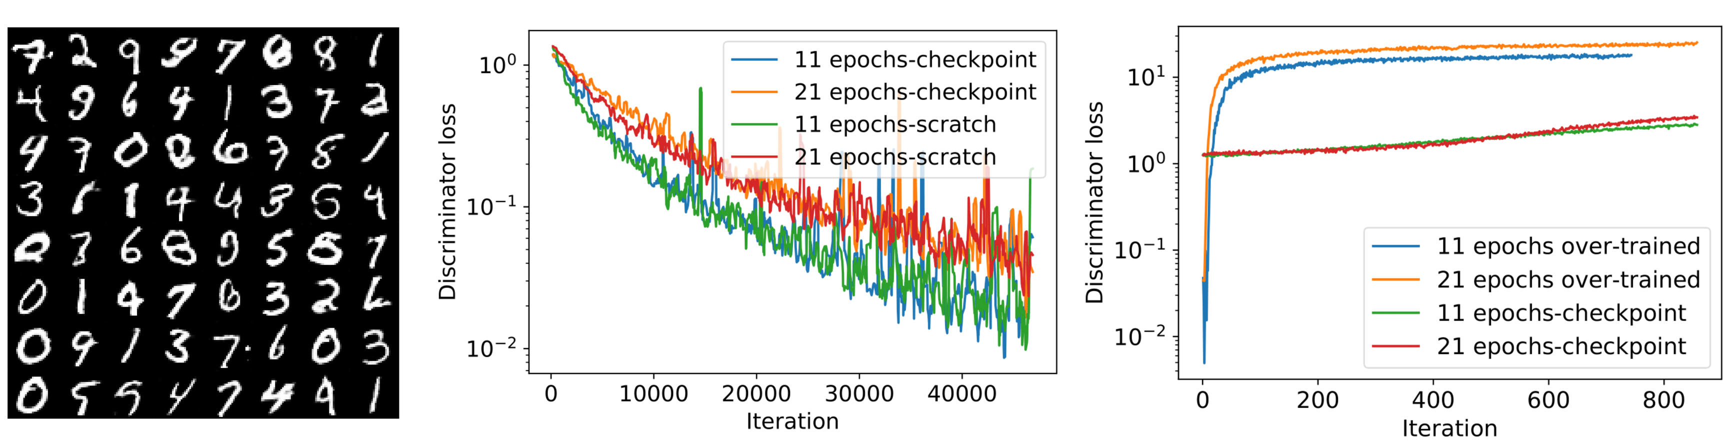

Train GAN for 21 epochs to get a 'checkpoint'. **Generates good images.** → Fix generator, only train the discriminator **Loss keeps dropping to 0.** → Fix discriminator, only train the generator. **Loss increases dramatically!**

- For fixed generator, discriminator can always improve → instability
- But by doing so, it becomes vulnerable to counter-attack of generator
- → Rational discriminator will avoid overtraining.

We can stabilize GANs by modeling strategic behavior of agents!

What is the game behind GANs?

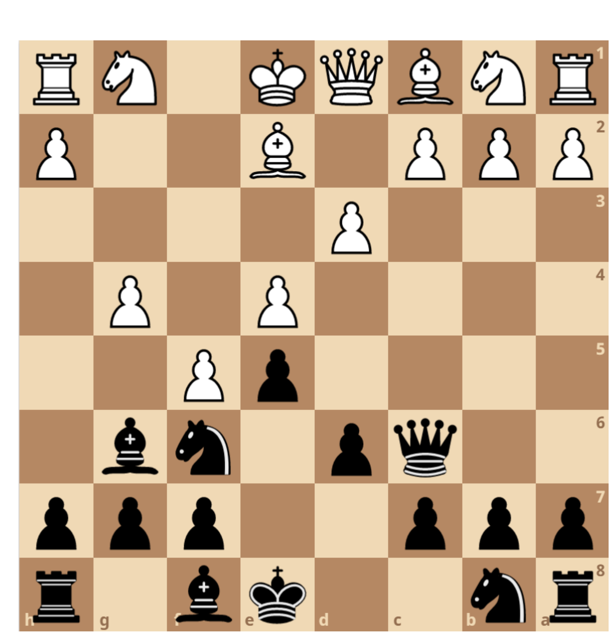

"The goal is to capture the other player's king" is not enough to model chess.

To model GANs, we need to specify:
- What **information** do the players have access to?
- Which **moves** are allowed?
- How do they **predict** their opponents?

One loss, three games

For $\alpha \gg 1$, consider the loss

$$\min_x \max_y -\exp(x^2) - \alpha xy + \exp(y^2)$$

No Nash equilibrium, but different rules lead to widely different games:

The global game:	The myopic game:	The predictive game:
1. Vision range: **global**.	1. Vision range: **1**.	1. Vision range: **1**.
2. Step size: **1**.	2. Step size: **1**.	2. Step size: **1**.
3. Predict opponent: **NO**.	3. Predict opponent: **NO**.	3. Predict opponent: **YES**.

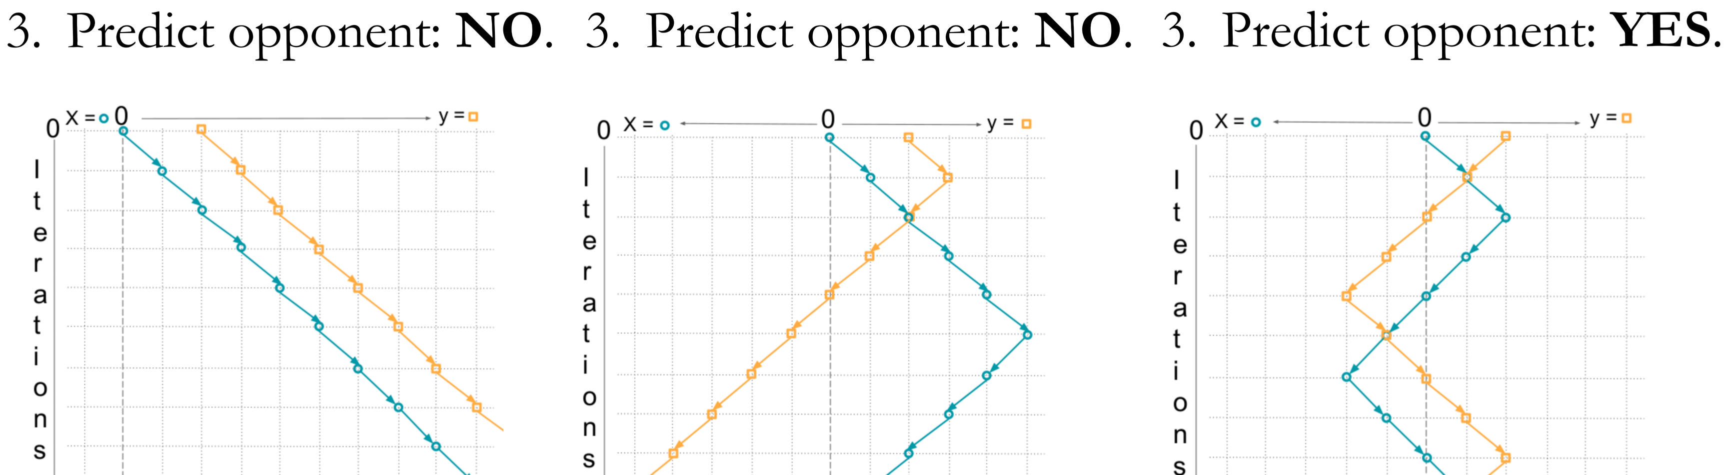

Framework of modeling the player's actions

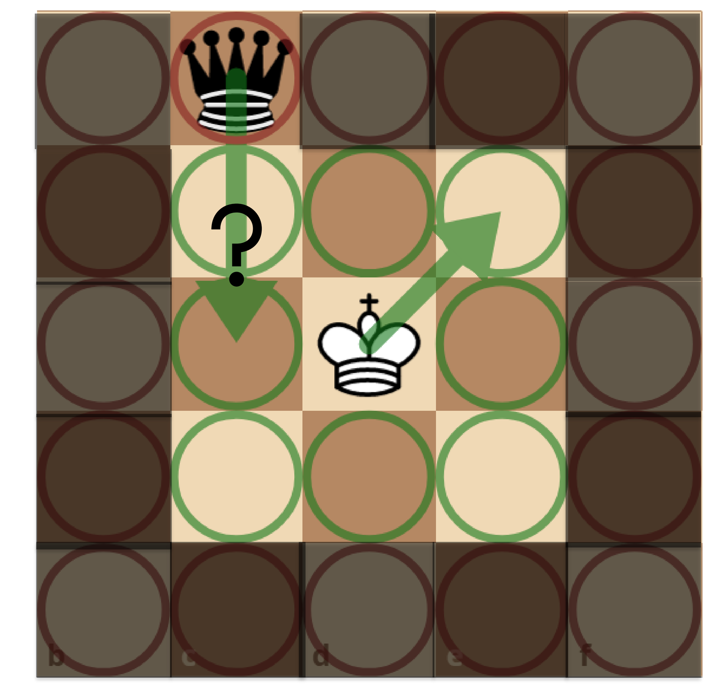

We model the players as choosing their actions based on three components:

1. The **belief** that the agents have about the loss.
2. The **uncertainty** that is incorporated into agents' decisions.
3. The **anticipation** of the action of the adversary when making decision.

Alternating gradient descent is related to the **myopic game** with updates:

$$\min_x f(x_k, y_k) + (x - x_k)^T \nabla_x f(x_k, y_k) + \frac{1}{2\eta}||x - x_k||^2$$

$$\max_y f(x_k, y_k) + (y - y_k)^T \nabla_y f(x_k, y_k) - \frac{1}{2\eta}||y - y_k||^2$$

Competitive gradient descent (CGD) is related to the **predictive game** with updates:

$$\min_x f + (x - x_k)^T \nabla_x f + (x - x_k)^T \frac{\partial^2 f}{\partial x \partial y}(\boxed{y} - y_k) + (\boxed{y} - y_k)^T \nabla_y f + \frac{1}{2\eta}||x - x_k||^2$$

$$\max_y f + (\boxed{x} - x_k)^T \nabla_x f + (\boxed{x} - x_k)^T \frac{\partial^2 f}{\partial x \partial y}(y - y_k) + (y - y_k)^T \nabla_y f - \frac{1}{2\eta}||y - y_k||^2$$

The updates of CGD are obtained as the Nash equilibrium of the game:

$$x_{k+1} = x_k - \eta\left(\mathrm{Id} - \eta^2 D_{xy}^2 f\, D_{yx}^2 g\right)^{-1}\left(\nabla_x f - \eta D_{xy}^2 f\, \nabla_y g\right)$$

$$y_{k+1} = y_k - \eta\left(\mathrm{Id} - \eta^2 D_{yx}^2 g\, D_{xy} f\right)^{-1}\left(\nabla_y g - \eta D_{yx}^2 g\, \nabla_x f\right)$$

Implicit Competitive Regularization

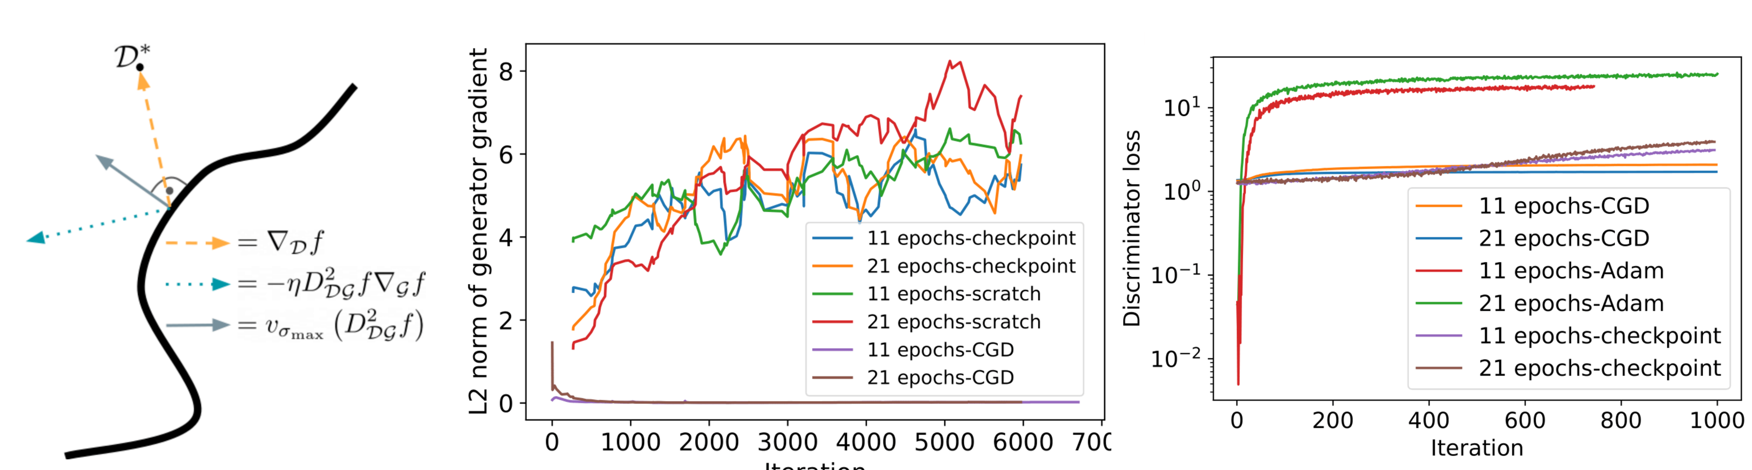

CGD prefers updates that are robust to the other player's actions.

When training the discriminator using CGD, it reduces the generator's gradient.

The discriminator trained with CGD is more robust than checkpoint.

Using CGD, the discriminator does not overtrain.

Improved Inception Score on CIFAR10

- **CGD + WGAN loss achieves highest inception score** (IS) (Figure 2).
- **CGD outperforms Adam under all the different settings** (Figure 3 plots the IS difference between CGD and Adam over all the settings).
- **CGD + WGAN loss without gradient penalty outperforms the WGAN-GP.**
- Tensorflow IS are reported in Figure4, 5, others are Pytorch IS.

Generated image

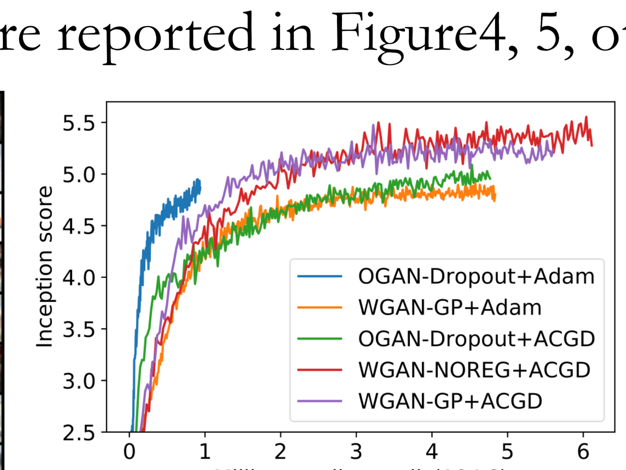

Figure 2

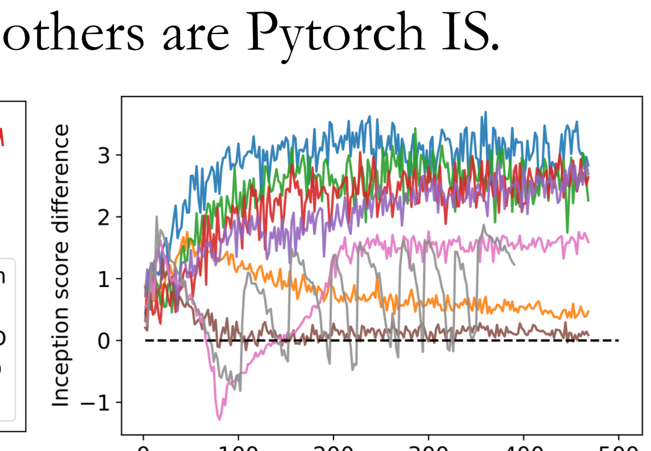

Figure 3

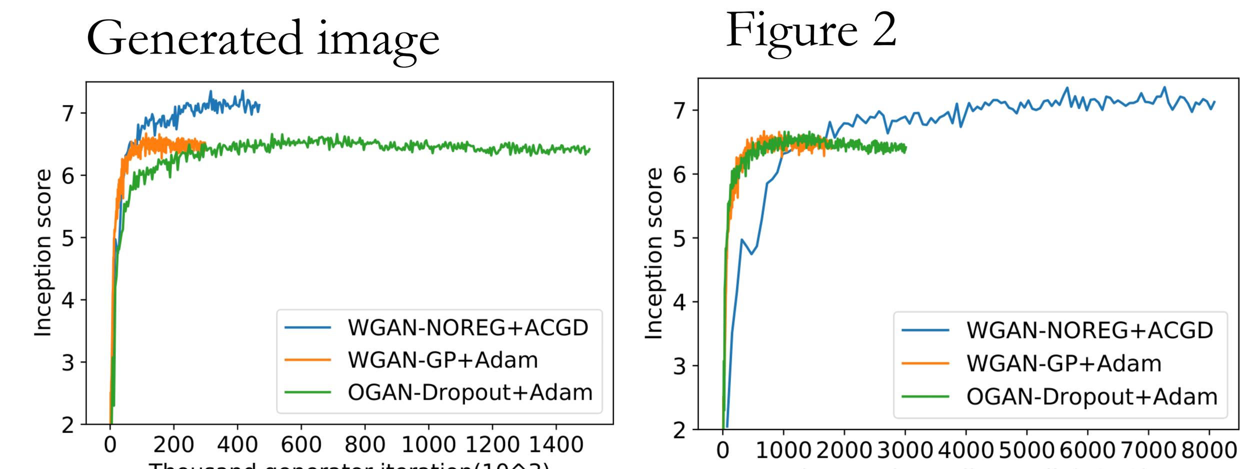

Tensorflow inception score

Paper and code