



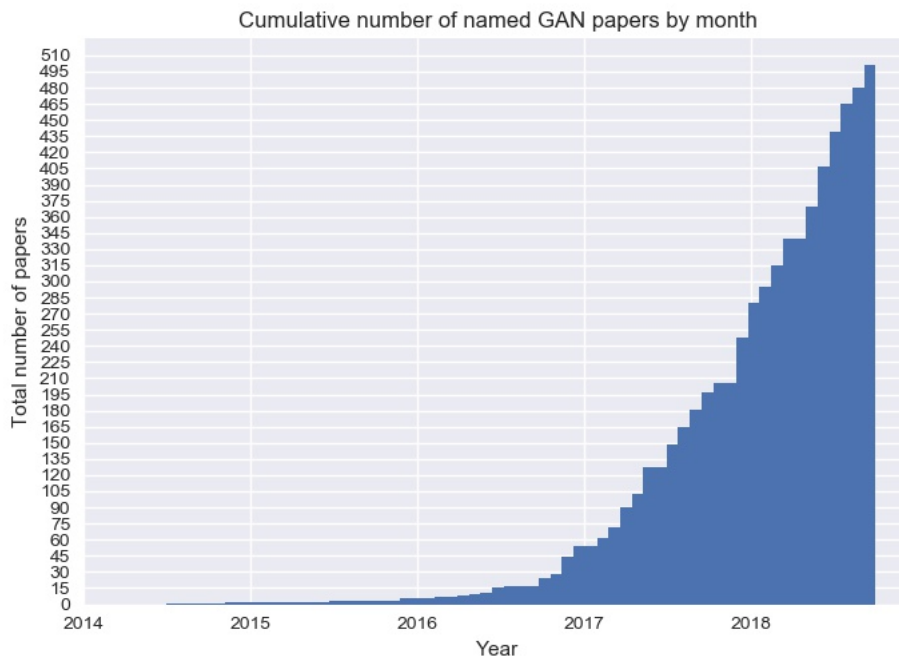
Generative Adversarial Networks

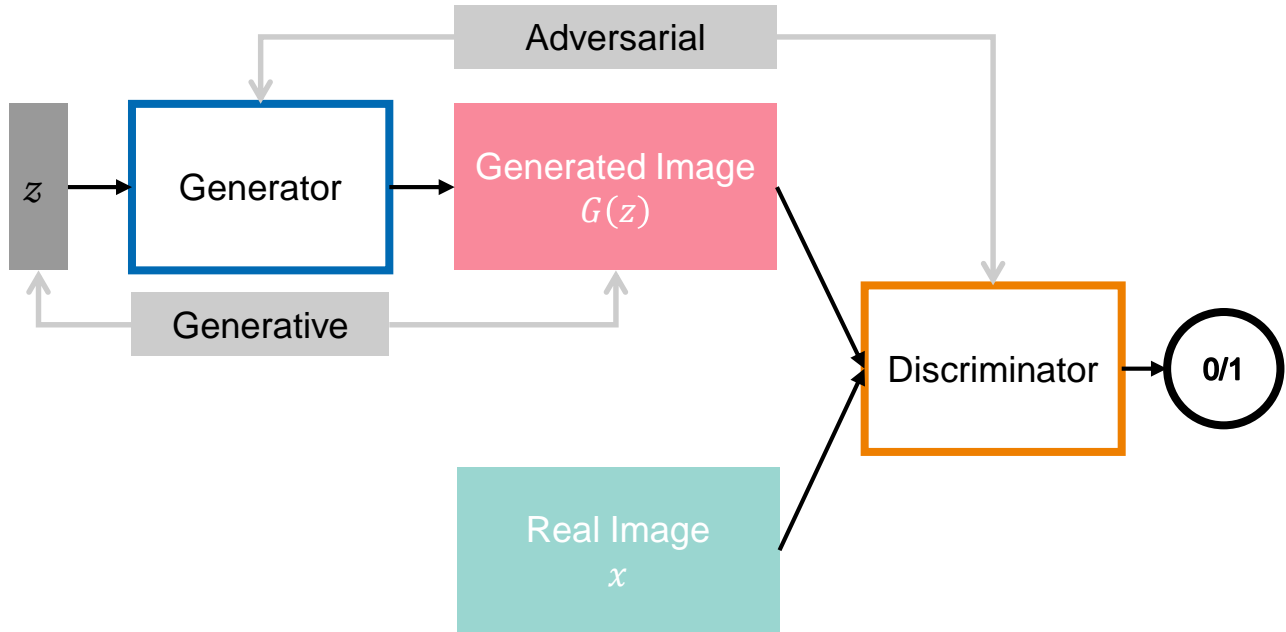
Hsien-Chin Lin

29.05.2020

- Generative Adversarial Networks (GAN)
 - Introduction
 - Algorithm
 - Example
 - Problems
 - Comparison with variational autoencoder
- Conditional GAN (cGAN)
 - Image to Image or Text to Image Task
 - Text to Text (SeqGAN)
- Adversarial Learning in Dialogue
- Conclusion

One of the most popular research topics





- Generate realistic outputs
- Generator models the data distribution
 - given $x \sim p_{data}(x)$, find $p_{model}(x; \theta) \approx p_{data}(x)$
- Maps random noise z to semantic space
- The output should be as realistic as possible
 - No blur edges, high resolution
 - Vivid color
 - Turing test



2014



2015

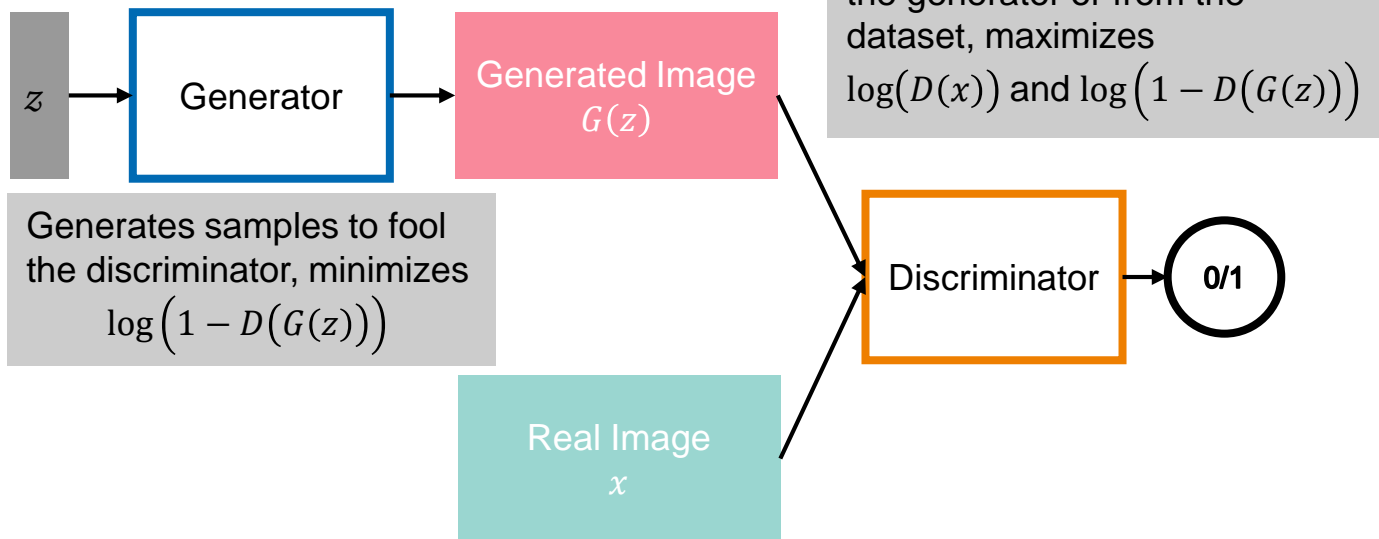


2016



2017

D and G play the two-player minimax game



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

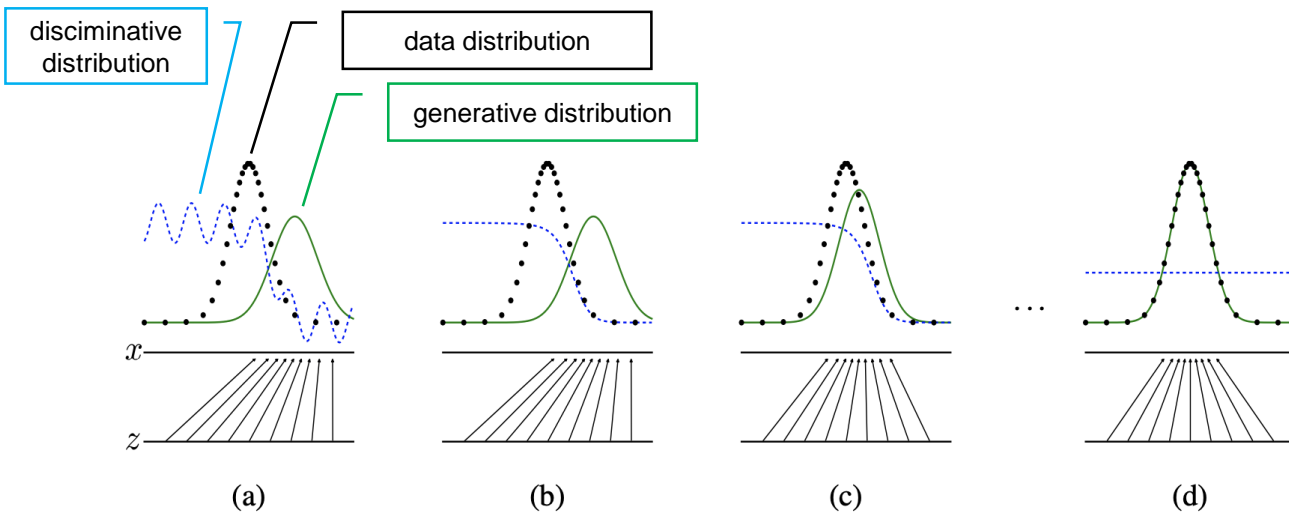
G: generator, D: discriminator

■ Update D

- for $i=1$, D-steps do
 - sample m noise samples $G(z)$ and m real samples x from dataset
 - update D, $\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m \left[\log D(x^i) + \log \left(1 - D(G(z^i)) \right) \right]$

■ Update G

- for $i=1$, G-steps do
 - sample m noise samples $G(z)$
 - freeze D, update G, $\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m \log \left(1 - D(G(z^i)) \right)$

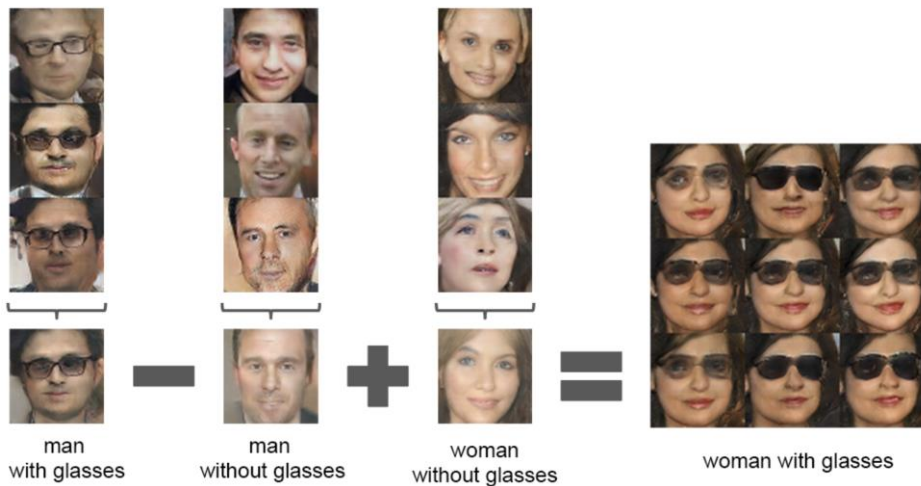


Example

- linearly interpolating (Goodfellow et al. 2014)



- Vector arithmetic for visual concepts (Radford, et al. 2015.)

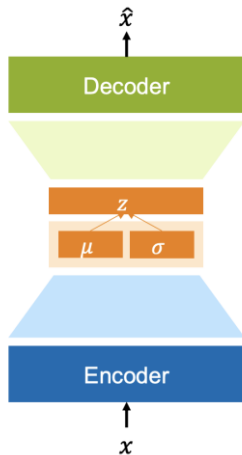


It is hard to train

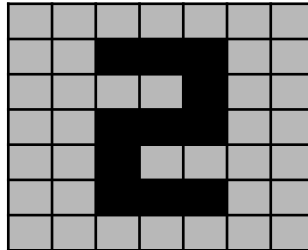
- Hard to balance the update of G and D
- Mode collapse
 - G only repeats the same image or copy the image in the real dataset
- D is too strong
 - the generator cannot get enough information to improve
- D is too weak
 - the generator will produce unrealistic images
- Put too much semantic information in one dimension

Variational Autoencoder

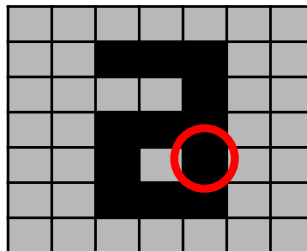
- Which information is missing in VAE's training?



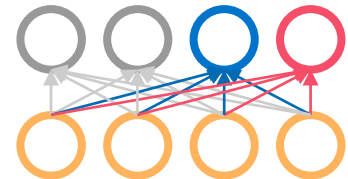
real example



1-pixel error

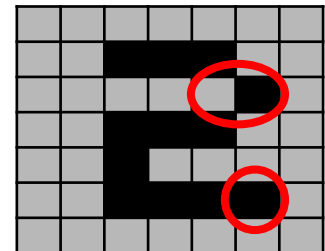


output layer



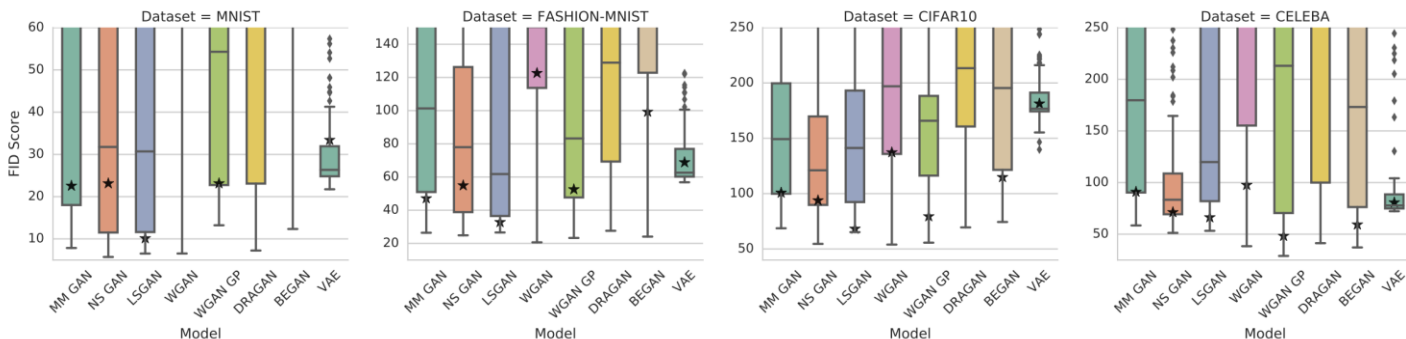
hidden layer

3-pixel error



Comparison with VAE

- GANs are more sensitive
- GANs perform better than VAE if we fully optimize the model
- VAE is more stable
- Fréchet Inception Distance (FID)
 - Measure the quality of generated samples, the lower is better

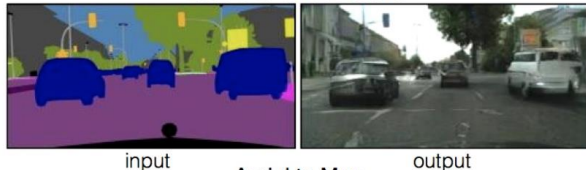


- Maps a random noise space to the semantic space
- Generates vivid outputs
- Sensitive to parameter choosing
- Hard to train, not stable

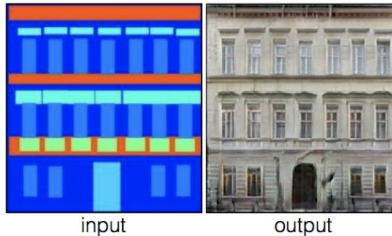
We want to generate samples based on some given information

- Image to image

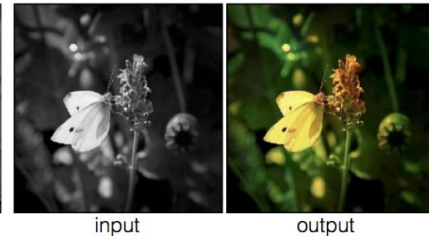
Labels to Street Scene



Labels to Facade



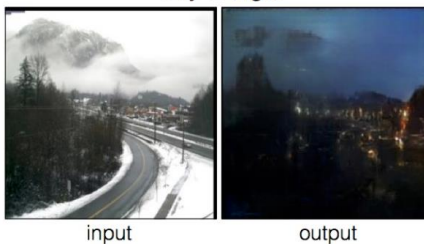
BW to Color



Aerial to Map



Day to Night



Edges to Photo



Text to Image

this small bird has a pink breast and crown, and black primaries and secondaries.



this magnificent fellow is almost all black with a red crest, and white cheek patch.



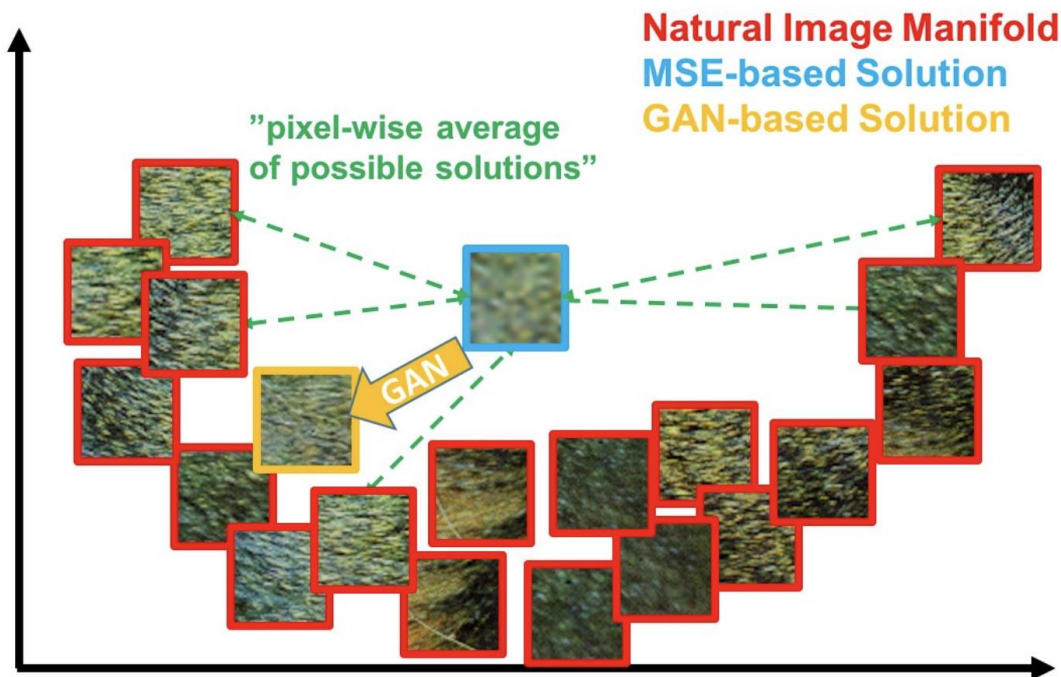
the flower has petals that are bright pinkish purple with white stigma

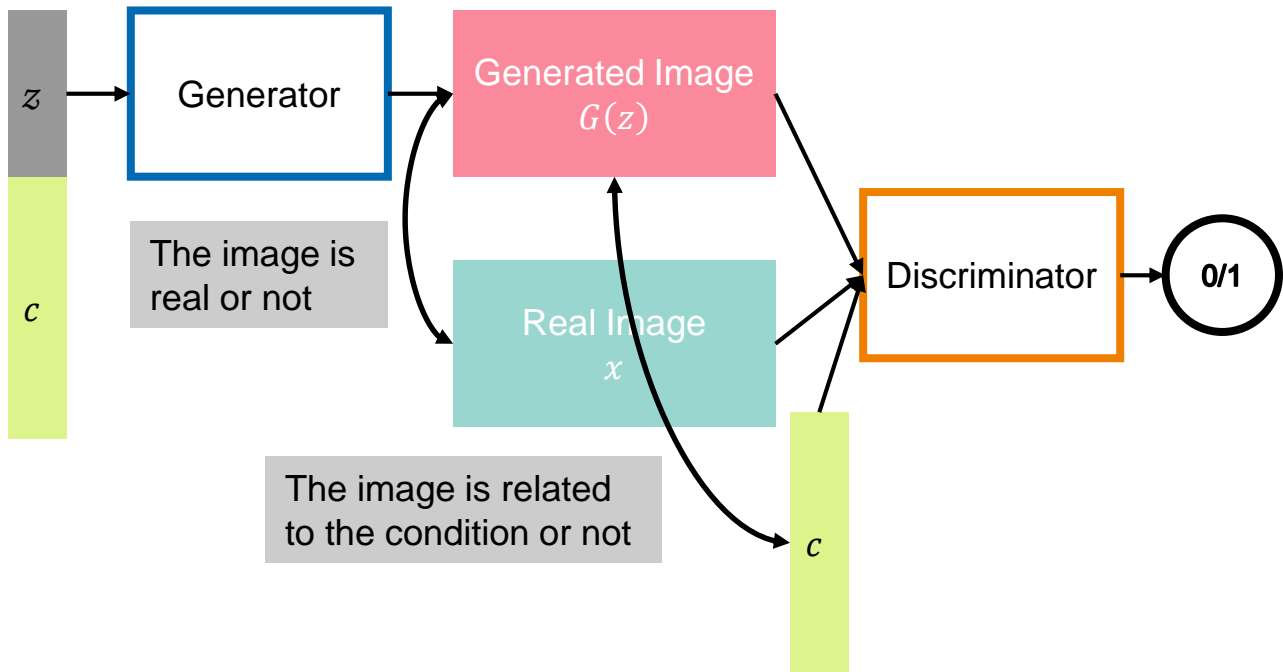


this white and yellow flower have thin white petals and a round yellow stamen



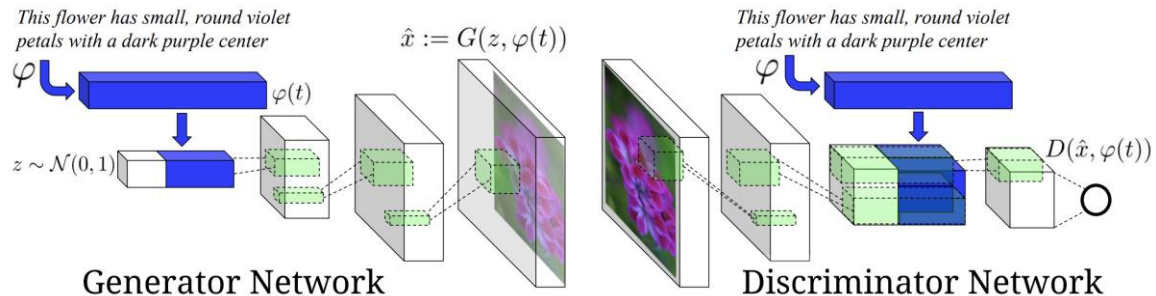
Why can't we use supervise learning?





Reed et al, 2016

Model structure

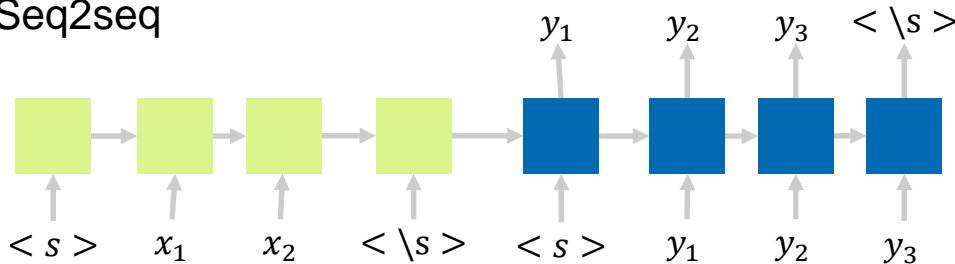


Example

- Interpolating
- Fixed random noise
- Fixed sentence embedding



■ Seq2seq



■ Tend to generate an “average” response

- “I don’t know”.

- “I’m sorry.”

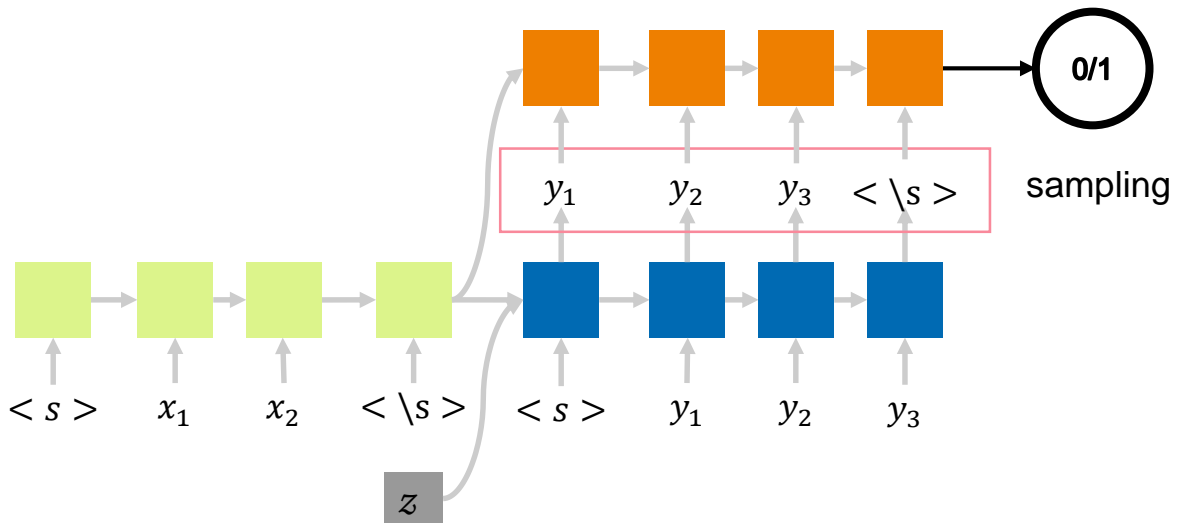
$$P(Y) = \prod_{i=1}^n P(y_i | y_{1:i-1}, X)$$

■ How do we fit this to the cGAN structure?

Directly fit into the cGAN framework?

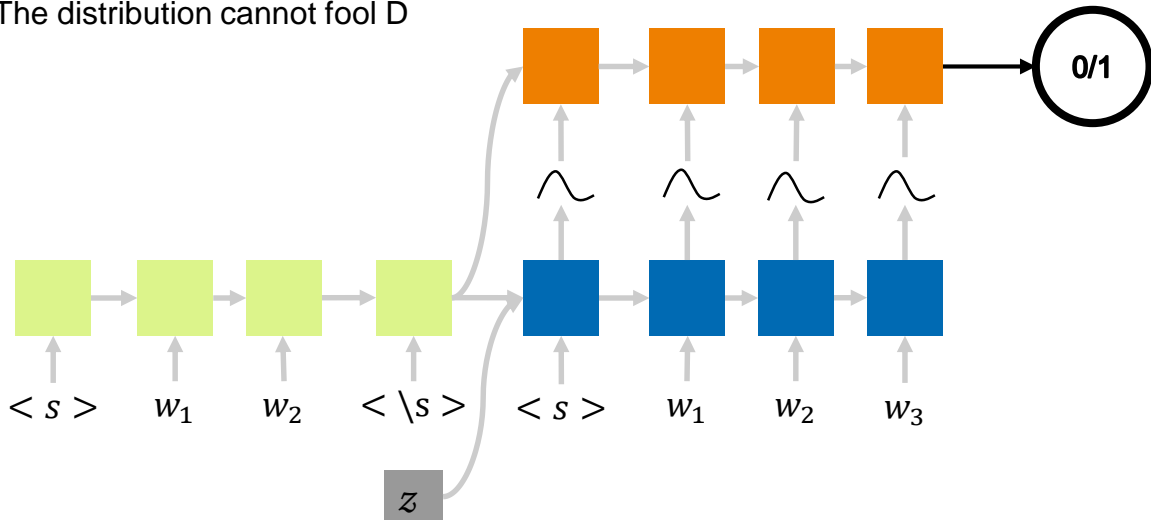
- Cannot update G because the sampling process is not differentiable

- `gray` -> `gray + 0.1`
- `cat` -> `cat + 0.1` ?



Pass distribution instead of sampling token

- The real data is discrete (one hot representation)
- The distribution cannot fool D

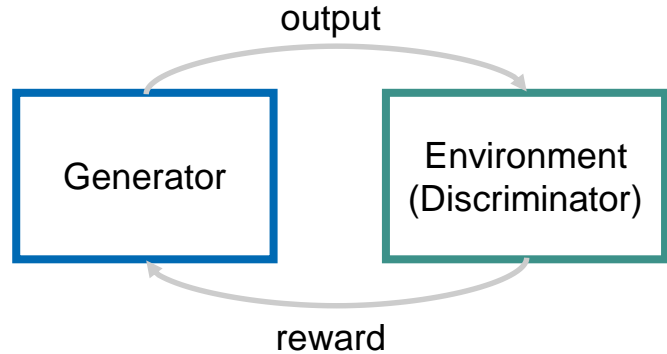


- Use reinforcement learning!

Using reinforcement learning to update

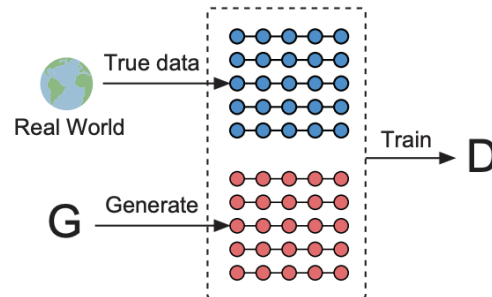
■ Update G

- D is the “environment”
- Use policy gradient to update G



■ Update D

- use supervise learning

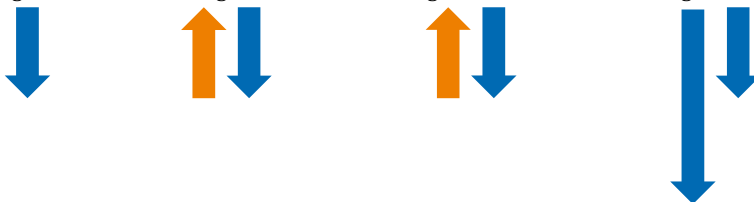


SeqGAN with policy gradient

- Without intermediate reward
- the generator (policy): $G_\theta(y_t|Y_{1:t-1})$
- maximize expected end reward $J_\theta = \mathbb{E}[R_T|s_0, \theta] = \sum_{y_1 \in \mathcal{Y}} G_\theta(y_1|s_0) \cdot Q_{D_\varphi}^{G_\theta}(s_0, y_1)$
 - R_T comes from discriminator D_φ
 - $Q_{D_\varphi}^{G_\theta}(s, a)$ is the action-value function, the expected accumulative reward
- $Q_{D_\varphi}^{G_\theta}(s = Y_{1:T-1}, a = y_T) = D_\varphi(Y_{1:T})$
 - reward only at the end of the sentence

Use RL to train the sequential generator

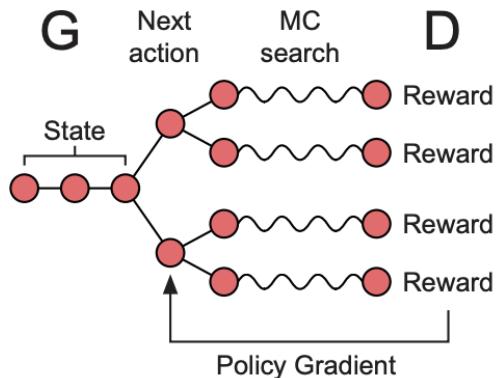
- C: "What's your name?" G: "I am fine."
- $D(C, G)$ is negative, update θ_g to decrease $\log P_{\theta_g}(G|C)$
- $\log P_{\theta_g}(G|C) = \log P_{\theta_g}(x_1|C) + \log P_{\theta_g}(x_2|C, x_1) + \log P_{\theta_g}(x_3|C, x_{1:2})$



- However, "I am John." is a positive example.
- How to get the intermediate reward? Use Monte Carlo search

N-time Monte Carlo search

- $Q_{D_\phi}^{G_\theta}(s = Y_{1:T-1}, a = y_T) = D_\phi(Y_{1:T})$
- $\{Y_{1:T}^1, \dots, Y_{1:T}^N\} = MC^{G_\beta}(Y_{1:t}; N)$
- $Q_{D_\phi}^{G_\theta}(s = Y_{1:T-1}, a = y_T) = \begin{cases} \frac{1}{N} \sum_{n=1}^N D_\phi(Y_{1:T}^n), & Y_{1:T}^n \in MC^{G_\beta}(Y_{1:t}; N) \text{ for } t < T \\ D_\phi(Y_{1:t}), & \text{for } t = T \end{cases}$



- Input “condition” to G and D
- For sequential generator
 - use reinforcement learning to update G
 - utilise Monte Carlo search
 - more computational power
 - more unstable

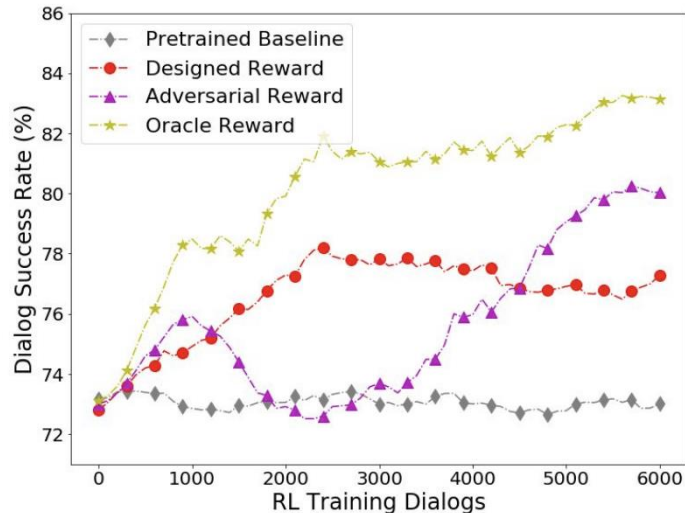
Train the dialogue agent (generator) (Li, et al, 2017)

- Open-domain dialogue
- with the dialogue history and the user utterance
- to generate the system response similar to human response

Input	tell me ... how long have you had this falling sickness ?
Vanilla-MLE	i 'm not a doctor .
Vanilla-Sample	well everything you did was totally untrue .
REINFORCE	i don 't know how long it 's been .
REGS Monte Carlo	A few months, I guess .

Estimate the reward function (discriminator) (Liu, et al, 2018)

- Task-Oriented
- the reward function can be learned as a discriminator
- Oracel
 - +1 success, +0 fail
- Human design:
 - +1 for each correct informable slot
 - if all informable slots are correct, +1 for each success requestable slot



■ Pros

- Powerful to generate photo-like images
- Model the data distribution
- Learn the representation of semantic space by mapping the noise
- cGANs have various generation condition

■ Cons

- Training and tuning GANs is not trivial
- Not stable
- Require a huge amount of computational power

■ Potentials

- GANs on the natural language (sequential generating) still need to improve

Thank you

Generative Adversarial Networks

- Goodfellow, Ian, et al. "Generative adversarial nets." *Advances in neural information processing systems*. 2014.
- Radford, Alec, Luke Metz, and Soumith Chintala. "Unsupervised representation learning with deep convolutional generative adversarial networks." *arXiv preprint arXiv:1511.06434* (2015).
- Lucic, Mario, et al. "Are gans created equal? a large-scale study." *Advances in neural information processing systems*. 2018.

Conditional GAN

- Isola, Phillip, et al. "Image-to-image translation with conditional adversarial networks." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017.
- Reed, Scott, et al. "Generative adversarial text to image synthesis." *arXiv preprint arXiv:1605.05396* (2016).
- Ledig, Christian, et al. "Photo-realistic single image super-resolution using a generative adversarial network." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017.
- Yu, Lantao, et al. "Seqgan: Sequence generative adversarial nets with policy gradient." *Thirty-First AAAI Conference on Artificial Intelligence*. 2017.

Adversarial Learning in Dialogue

- Li, Jiwei, et al. "Adversarial learning for neural dialogue generation." *arXiv preprint arXiv:1701.06547* (2017).
- Liu, Bing, and Ian Lane. "Adversarial Learning of Task-Oriented Neural Dialog Models." *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*. 2018.