MASKGAN: BETTER TEXT GENERATION VIA FILLING IN THE ____

William Fedus, Ian Goodfellow, Andrew M. Dai

Presented by: Joey Bose

February 16, 2018

Motivation

- Maximum Likelihood RNN's are the most common generative model for sequences
- Teacher Forcing leads to unstable dynamics in the hidden states
- Professor Forcing does solve the above but does not encourage high sample quality
- GAN's have shown incredible quality samples for images but discrete nature of text makes training a generator harder.
- Reinforcement Learning framework can be leveraged to train the generator by policy gradients

Teacher Forcing

Teacher forcing is a method for quickly and efficiently training recurrent neural network models that use the **true** output rather than the model output from a prior time step as input. (Figure taken from http://www.cedar.buffalo.edu/ srihari/CSE676)

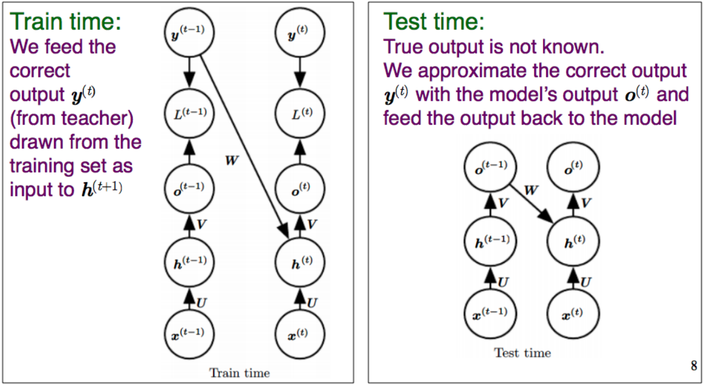

Professor Forcing

- Training RNN's such that the network is indistinguishable during training and when sequences are self-generated.
- Train a second model to match the distribution of training and self generated sequences

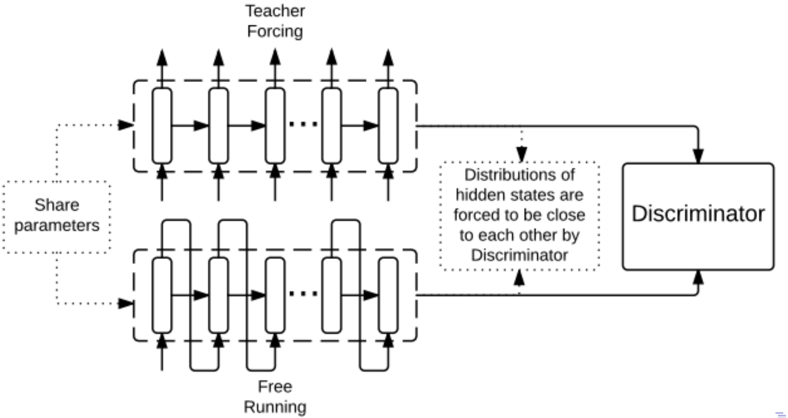

GAN's

GAN training involves jointly training both a generator and discriminator network in a min-max setting where the value function is $\min_G \max_D V(D, G)$:

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

Similarly one can add side channel information and condition the generator:

$$\mathbb{E}_{x \sim p_{data}(x)}[\log(D(x))] + \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z|y)))], \tag{2}$$

Common failure modes for GAN's include unstable training dynamics and mode collapse, which are exacerbated in a discrete setting

Conditional GAN

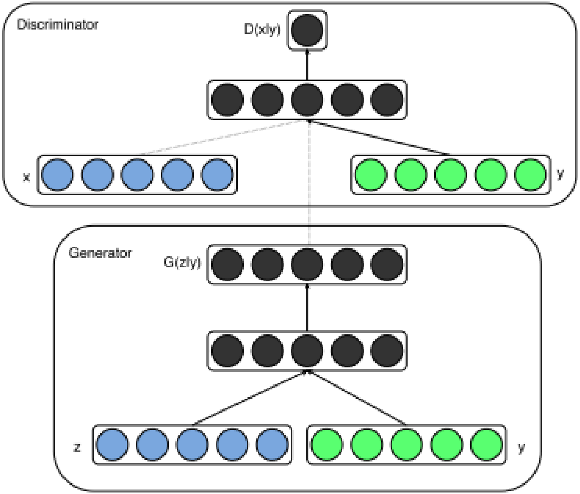

SeqGAN

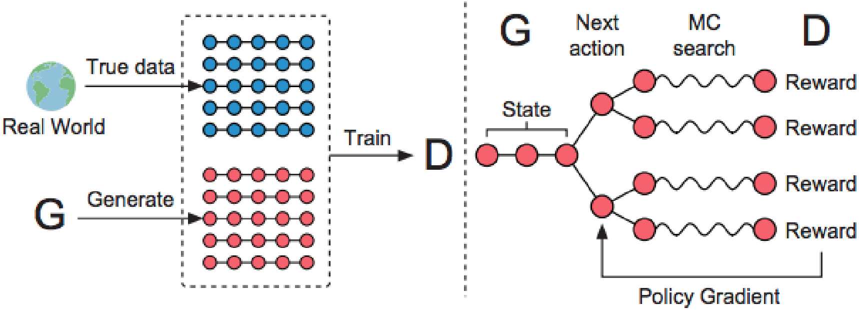

True data
Real World
G
Generate
Train
D
G
Next action
MC search
D
State
Reward
Reward
Reward
Reward
Policy Gradient

MaskGAN problem setup

Start with a ground truth discrete sequence $x = (x_1, ..x_T)$ and a binary mask of the same length, $m = (m_1, ..., m_T)$. Applying the mask on the input sequence creates, $m(x)$, a sequence with blanks:

For example:

$$x = a, b, c, d, e$$
$$m = 0, 1, 1, 0, 0$$
$$m(x) = a, \_, \_, d, e$$

The goal of the generator is then to autoregressively fill in the missing tokens conditioned on the previous tokens and the **mask**

MaskGAN Generator

The Generator for MaskGAN is Seq2Seq network and models the following distribution:

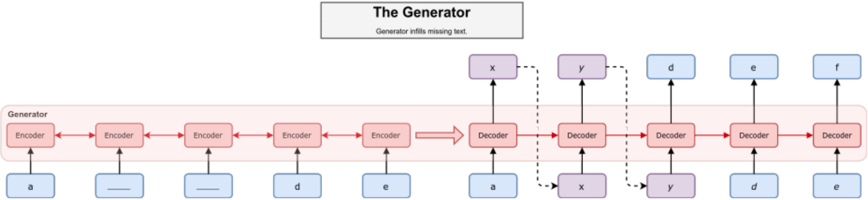

$$P(\hat{x_1}, ..., \hat{x_T} | m(x)) = \prod_{t=1}^{T} P(\hat{x_t} | \hat{x_1}, ..., \hat{x_{t-1}}, m(x)), \tag{3}$$

$$G(x_t) = P(\hat{x_t} | \hat{x_1}, ..., \hat{x_{t-1}}, m(x)), \tag{4}$$

MaskGAN Discriminator

The discriminator is also a Seq2Seq network and receives the filled in sequence and **m(x)**.

Failure Mode, Example

G: **The director director guided the series**
Possible Label 1: **The *associate* director guided the series** Possible Label 2: **The director *expertly* guided the series**

This gives an ambiguous signal to Generator when $m(x)$ is not given as the generator doesn't know which of the directors tokens is wrong. The discriminator computes $P(\tilde{x}_t = x_t^{real} | \tilde{x}_{0:T}, m(x))$

MaskGAN Generator Training using RL

- Reward is *log* of discriminator output $r_t = logP(\tilde{x_t} = x_t^{real} | \tilde{x_{0:T}}, m(x))$
- Additional Critic Head of the discriminator and estimates the value function $R_t = \sum_{s=t}^{T} \gamma^s r_s$
- The generator maximizes expected reward $\mathbb{E}_{G(\theta)}[R]$ using the REINFORCE gradient estimator with the Critic as a baseline

$$\nabla_\theta \mathbb{E}[R] = (R_t - b_t) \nabla_\theta logG_\theta(\tilde{x_t}) \quad (5)$$

The gradient to the generator for token $\tilde{x_t}$ depends on all discounted future rewards assigned by the discriminator. The discriminator update is the same as vanilla GAN.

Pretraining and Experimental Setup

There are two stages of pretraining:

- Pretrain a language model with Maximum Likelihood and then use its weights to initialize Seq2Seq
- Pretrain Seq2Seq model on infilling task using Maximum Likelihood
- Select model with lowest validation perplexity on infilling task over 500 runs
- Authors report that adding a critic decreases variance by an order of magnitude
- Authors claim pretraining is not necessary for best results but significantly reduces computation time

Evaluation Metrics

Bleu Score for MaskGAN

Compare the number of unique n-grams produced by the Generator with the n-grams of the validation corpus and count the number of matches. Then compute the geometric average.

- Reducing validation perplexity does not equate to high sample quality.
- Generator can explore off-manifold and find alternative valid samples.
- This is justified as it is known that perplexity is not the best metric for language models

Experiments

Two modes of operation

1. Conditional Mode
2. Free-Running/Unconditional

Conditional Mode is the Regular MaskGAN with deterministic / stochastic mask applied to the input sequence. Free-Running mode is when the entire input sequence is masked reducing this task to traditional language modeling

Conditional samples: PTB

Ground Truth	**the next day 's show <eos> interactive telephone technology has taken a new leap in <unk> and television programmers are**
MaskGAN	the next day 's show <eos> interactive telephone technology has taken a new leap in its retail business <eos> a
MaskMLE	the next day 's show <eos> interactive telephone technology has taken a new leap in the complicate case of the

Table 1: Conditional samples from PTB for both MaskGAN and MaskMLE models.

Unconditional samples: PTB

MaskGAN	oct. N as the end of the year the resignations were approved <eos> the march N N <unk> was down

Table 2: Language model (unconditional) sample from PTB for MaskGAN.

Conditional samples: IMDB

Ground Truth	**Pitch Black was a complete shock to me when I first saw it back in 2000 In the previous years I**
MaskGAN	Pitch Black was a complete shock to me when I first saw it back in 1979 I was really looking forward
MaskMLE	Black was a complete shock to me when I first saw it back in 1969 I live in New Zealand

Table 3: Conditional samples from IMDB for both MaskGAN and MaskMLE models.

Unconditional samples: IMDB

MaskGAN	**Positive**: Follow the Good Earth movie linked Vacation is a comedy that credited against the modern day era yarns which has helpful something to the modern day s best It is an interesting drama based on a story of the famed

Table 4: Language model (unconditional) sample from IMDB for MaskGAN.

Perplexity of Generated Samples and Mode Collapse

Test the perplexity of MaskGAN and MaskMLE under the pretrained language model used to initialize it.

Model	Perplexity of IMDB samples under a pretrained LM
MaskMLE	273.1 ± 3.5
MaskGAN	108.3 ± 3.5

Evidence of Mode Dropping

Model	% Unique bigrams	% Unique trigrams	% Unique quadgrams
LM	40.6	75.2	91.9
MaskMLE	43.6	77.4	92.6
MaskGAN	38.2	70.7	88.2

Human Evaluation on Mechanical Turk

Preferred Model	Grammaticality %	Topicality %	Overall %
LM	15.3	19.7	15.7
MaskGAN	59.7	58.3	58.0
LM	20.0	28.3	21.7
MaskMLE	42.7	43.7	40.3
MaskGAN	49.7	43.7	44.3
MaskMLE	18.7	20.3	18.3
Real samples	78.3	72.0	73.3
LM	6.7	7.0	6.3
Real samples	65.7	59.3	62.3
MaskGAN	18.0	20.0	16.7

Table 7: A Mechanical Turk blind heads-up evaluation between pairs of models trained on IMDB reviews. 100 reviews (each 40 words long) from each model are unconditionally sampled and randomized. Raters are asked which sample is preferred between each pair. 300 ratings were obtained for each model pair comparison.

Failure Cases

1. Mode Dropping is less extreme than SeqGAN but still noticeable.

   It is a very funny film that is very funny It s a very funny movie and it s charming It

2. Matching Syntax at Boundaries

   Cartoon is one of those films me when I first saw it back in 200

3. Loss of Global Context

   This movie is terrible The plot is ludicrous The title is not more interesting and original This is a great movie
   Lord of the Rings was a great movie John Travolta is brilliant

The End