Jiwei Li, NLP Researcher

By Pragya Arora & Piyush Ghai

Introduction

- Graduated from Stanford University in 2017
- Advised by Prof. Dan Jurafsky
- Closely worked with Prof. Eduard Hovy from CMU and Prof. Alan Ritter from OSU
- Affiliated with The Natural Language Processing Group at Stanford University

Research Interests

- Jiwei's research interests focus on computational semantics, language generation and deep learning. His recent work explores the feasibility of developing a framework and methodology for computing the informational and processing complexity of NLP applications and tasks.
- His PhD thesis was on "*Teaching Machines to Converse*"
- Has over 1200[1] citations on Google Scholar.
- Has over 38[1] scholarly publications.

1 : Google Scholar Site

Teaching Machines to Converse

- Jiwei's primary research focus and his thesis work was on conversational models for machines.
- Some of his publications in this domain are :
  - Deep Reinforcement learning for dialogue generation [2016], J Li, W Monroe, A Ritter, M Galley, J Dao, D Jurafsky
  - A persona based neural conversation model [2016], J Li, M Galley, C Brockett, GP Spithourakis, J Gao, B Dolan
  - Adverserial Learnig for Neural Dialogue Generation [2017], J Li, W Monroe, T Shi, A Ritter, D Jurafsky

Adverserial Learning for Neural Dialogue Generation

Co-Authors

- Will Monroe, PhD Student @Stanford
- Tianlin Shi, PhD Student @Stanford
- Sebastien Jean, PhD Student @NYU Courant
- Alan Ritter, Assistant Professor, Dept of CSE, Ohio State University
- Dan Jurafsky, Professor, Dept of CSE, Stanford University

Goal

“To train and produce sequences that are indistinguishable from human-generated dialogue utterances”.

This paper trended on social media as well...

Adversarial Models

It’s a Min-Max game between a Generator & Discriminator

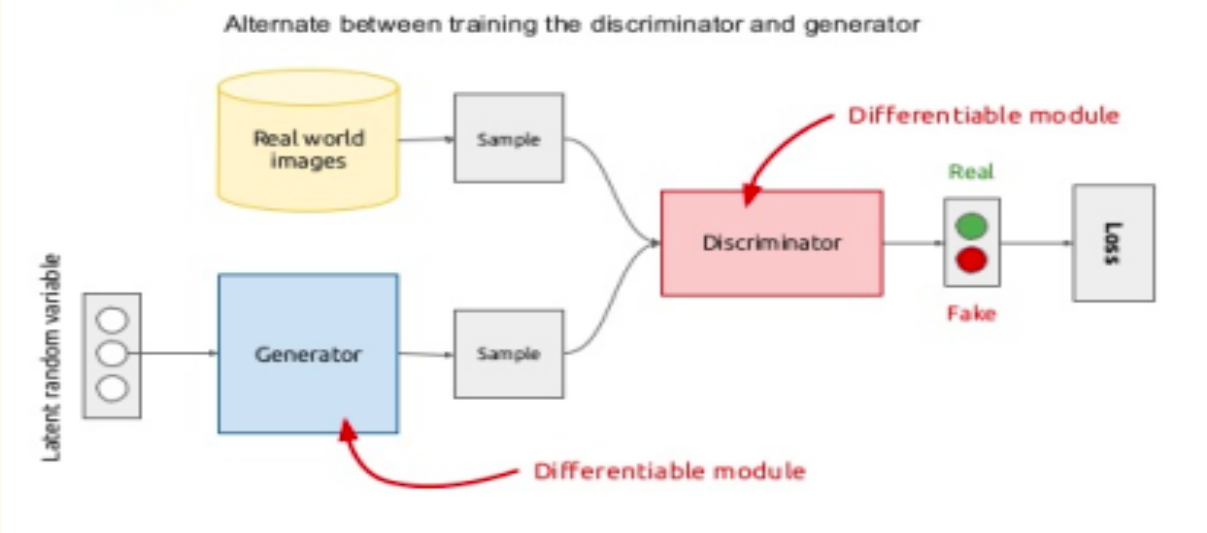

Model Used

- Earlier REINFORCE Algorithm was used, which had it’s own drawbacks.
    - The expectation of reward is approximated by only one sample and reward associate with it is used for all the samples.

Input : *What’s your name*
human : *I am John*
machine : *I don’t know*

- Vanilla REINFORCE will assign the same negative weight for all the tokens - [I, don’t, know], even though [I] matched with the human utterance.

REGS - Reward Generation for Every Step

- They reward the sequence generated at intermediate steps as well.
- They essentially train their discriminator for rewarding partially decoded sequences.
- They also use Teacher Forcing as well, where the human responses are also fed to the generator, with a positive reward. This helps it to overcome the problems where it can get stuck in Minimas and it would not know which update steps to take.

For number of training iterations **do**
. **For** i=1,D-steps **do**
. Sample (X,Y) from real data
. Sample $\hat{Y} \sim G(\cdot|X)$
. Update D using (X, Y) as positive examples and $(X, \hat{Y})$ as negative examples.
. **End**
.
. **For** i=1,G-steps **do**
. Sample (X,Y) from real data
. Sample $\hat{Y} \sim G(\cdot|X)$
. Compute Reward r for $(X, \hat{Y})$ using D.
. Update G on $(X, \hat{Y})$ using reward r
. Teacher-Forcing: Update G on (X, Y)
. **End**
End

Results

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 .
Input	so i had the doctors test sammy 's response to conditioning .
Vanilla-MLE	sammy wrote the test sammy wrote the test .
Vanilla-Sample	objects pick objects up objects objects objects objects objects
REINFORCE	i 'm not sure that 's relevant .
REGS Monte Carlo	so he took the pills .
Input	what are you going to do ?
Vanilla-MLE	i 'm going to the movies .
Vanilla-Sample	get him outta here first !
REINFORCE	i 'm going to get you .
REGS Monte Carlo	i 'm going to get you a new car .
Input	they fear your power your intellect .
Vanilla-MLE	you 're the only one who knows what 's going on .
Vanilla-Sample	when they are conquered and you surrender they will control all of us .
REINFORCE	i 'm afraid i 'm not ready yet .
REGS Monte Carlo	i 'm not afraid of your power .