Make Neural Nets Great Again

A neural net to generate Trump tweets

C. Foster, D. Gibb, C. Seinen, C. Watson

November 29, 2016

University of Victoria

Introduction

Motivations

This is an example of Natural Language Generation (NLG)

- Improve computer-human interaction
  - Data-to-text applications
  - Accessibility
- Automatic Translation

Okay, so why Donald Trump tweets?

1. Size limit
2. Twitter data is easy to get
3. Tweets provide an “interesting” challenges for NLG
4. It’s Amusing!

Natural Language Generation

Why is this a hard and interesting problem?

- Immensely context-sensitive
  - "Time flies like an arrow." vs. "Fruit flies like a banana."
- Context can span the short- and long-term.
  - "I'm from France, so I speak _____."
  - "I was born in Paris...*[three pages later]*...I always loved speaking my native language _____."

Tag in Recurrent Neural Networks!

Existing Work

1. Learning Holiness [cpury()]
    - “26:24 And the children of Israel went up from the LORD the word of the LORD, and set the LORD in the house of the LORD.”
2. Sunspring [Films()]
    - “I know it’s a consequence. Whatever you want to know about the presence of the story, I’m a little bit of a boy on the floor.”
3. Artifical Scenes from the hit 90’s television sitcom, *Friends* [Herd(2016)]
    - “Phoebe: Wow lady! You’re just gonna come over to him jumpy.
      (They start to cry.)
      Chandler: So, Phoebe likes my pants.”
4. DeepDrumpf [Hayes(2016)]

Neural Networks

Feed-Forward Neural Networks

Feed-forward neural networks:

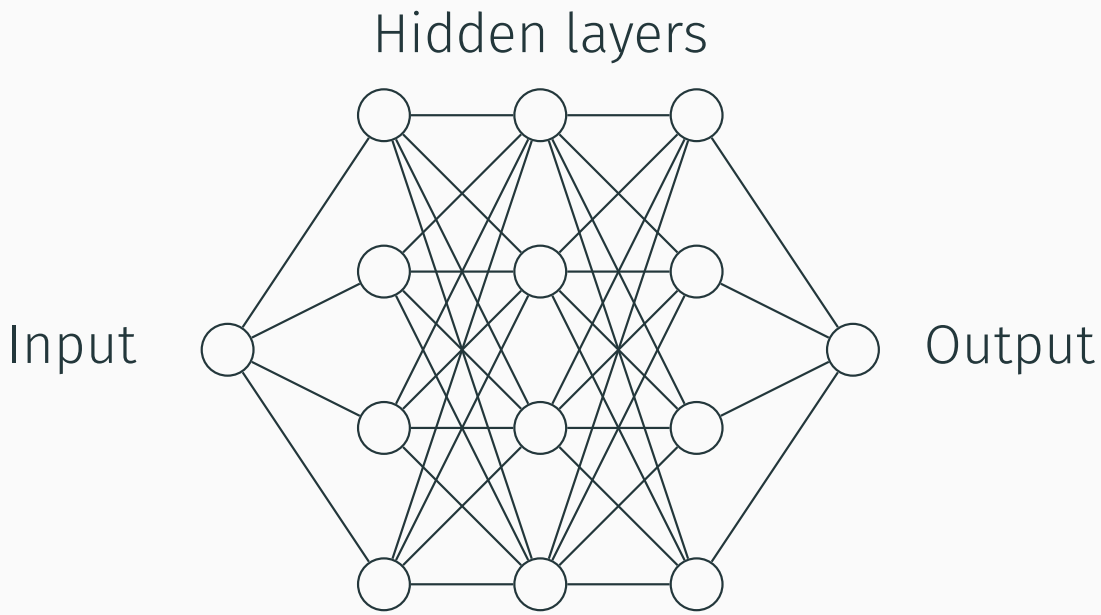

- Good for sequential structure of fixed or bounded length.
- Bad for language, where structure can be of any length.

Recurrent Neural Networks

Recurrent neural networks (RNNs) remember the previous output.

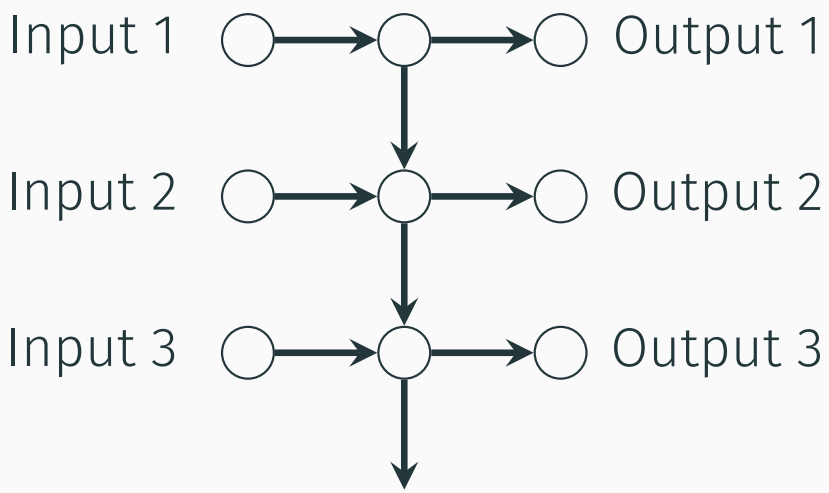

In theory, RNNs can recognize structures over any number of steps.
In practice, they do not learn long-term structure.

RNNs suffer from the *vanishing gradient problem* (VGP).

Solving the Vanishing Gradient Problem

We can avoid the VGP by never squashing the unit's memory.

Long short-term memory units (LSTMs) were introduced in 1997.

They solve the VGP, but are more complex than standard RNNs.

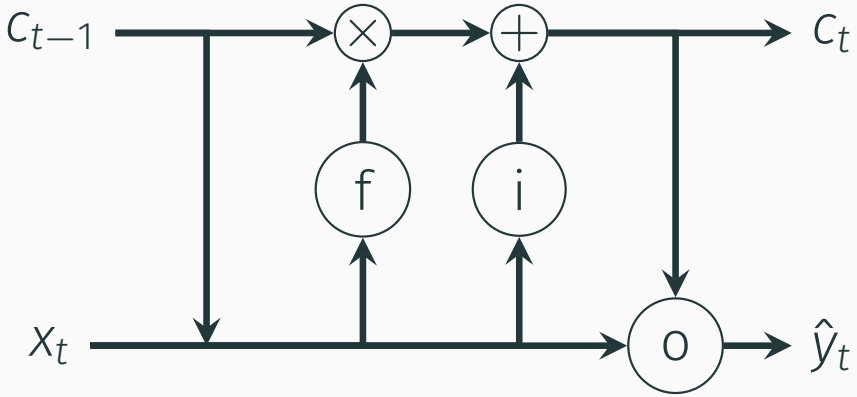

Figure 1: A typical LSTM.

Gated reccurent units (GRUs) were introduced in 2014. Simpler than LSTMs, but often perform similarly.

Generative Neural Networks

We have used neural nets as classifiers. How can they generate anything?

Example

You feed the sequence [Make, America] into the RNN.

great	Mexico	crooked	vote	...
0.8	0.03	0.1	0.05	0.01

Use these likelihoods as probabilities. Choose a word accordingly.

Feed the chosen word, crooked, into the RNN as the third input.

again	Hillary	politicians	polls	...
0.03	0.85	0.1	0.05	0.02

Continue as long as you like.

Implementation

Data Sources

We started with an initial CSV from Kaggle containing 3,000 tweets
Some simple preprocessing was required:

- Remove retweets
- Filter tweets by @realDonaldTrump
- Extract tweet textual content
- Add tweet start/end delimiters

Later in the election, we built a custom scraper using PyTweet to update our dataset with the latest tweets, giving us 3,700 tweets.

Neural Nets

Candidate networks:

1. TensorFlow LSTM network with 512 nodes and 2 layers
2. Theano GRU network with 250 nodes and 2 layers
3. KERAS LSTM network with 128 nodes and 1 layer

Our LSTM-based networks use character-level modeling
Our GRU-based network uses word-level modeling

Results

Network Performance

How do you measure performance for a generative network?

- Perfect performance is essentially just overfitting
- Instead, we measure with a ”Trump Turing Test”

Web Application

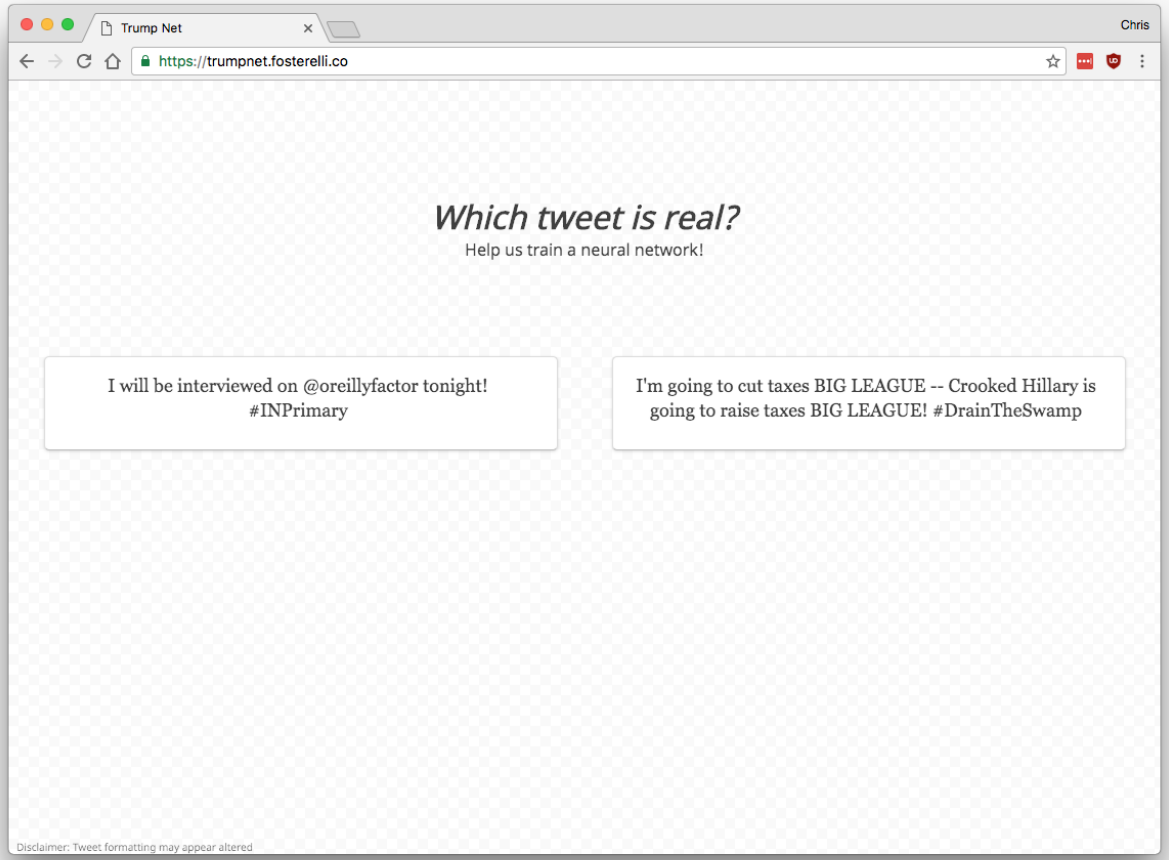

https://trumpnet.fosterelli.co

Implemented in Node.js with React and RethinkDB

Results

- Received over 3,500 responses
- GRU has a performance of 83.5%
- LSTM has a performance of 82.7%
- We could fool users approximately 17% of the time
- Users are approximately 33% better than randomly guessing

Example Tweets

Donald J. Trump @realDonaldTrump · 19h

I will MAKE AMERICA SAFE AND GREAT AGAIN! #ImWithYou #AmericaFirst https://t.co/cgoqZLDXDV

26K 81K

Donald J. Trump @realDonaldTrump · 19h

Great to be back in Arizona! #MakeAmericaGreatAgain #Trump2016 #VoteTrumpNC

26K 81K

Donald J. Trump @realDonaldTrump · 20h

Join me in Cleveland, Ohio tomorrow at 3:00pm. Can't wait to help #ObamaCare. #VoteTrumpNH #MakeAmericaGreatAgain #FITN https://t.co/GtPk1QPzIm

26K 81K

Example Tweets

Donald J. Trump @realDonaldTrump · 19h

Thank you Ohio! Go to the bosses today - we can replace Obamacare and "choice" - a total lie!

26K 81K

Donald J. Trump @realDonaldTrump · 19h

The failing @nytimes is truly one of the worst post. They Alicia clue lies and never even call to fact playing. Really bad people!

26K 81K

Donald J. Trump @realDonaldTrump · 19h

.@SenTedCruz Ted - I will make our economy she would lose!

26K 81K

Donald J. Trump @realDonaldTrump · 19h

The system is totally rigged and corrupt media and her government protection process. People get it!

26K 81K

MAKE NEURAL NETS GREAT AGAIN!

#UVic2016

References

cpury.
Learning holiness.
URL http://cpury.github.io/learning-holiness/.

Therefore Films.
Sunspring.
URL http://www.thereforefilms.com/sunspring.html.

Andy Herd, jan 2016.
URL https://twitter.com/_pandy/status/689209034143084547.

Bradley Hayes.
Postdoc develops twitterbot that uses ai to sound like donald trump, Mar 2016.
URL http://www.csail.mit.edu/deepdrumpf.