Generative Text Style Transfer for Improved Language Sophistication

https://youtu.be/0mg6eoe5q6g

Robert Schmidt, Spencer Braun

rschm@stanford.edu, spencerb@stanford.edu

Motivation

- Becoming common to use automated grading in practical settings
- Can a language model learn what makes a text sophisticated?
- Can it reproduce sophistication?
- We set out to see how state-of-the-art style transfer models would perform on this nuanced task

Data

Naive Data

- Student essay data and simple essays for grade schoolers.
- Essays had already replaced identifying names, locations, and numbers by tags generated by the Stanford Named Entity Recognition (NER) Tagger.

Untagged:
"Many people believe computers are bad but how can you make friends if you can never talk to them?"

Tagged:
"<person>, the owner of <organization> said that the internet saved her resturant."

Sophisticated Data

- Literature, essays, and non-fiction texts from Project Gutenberg and the Oxford Text Archive
- Formed 3 different corpora of texts with different author compositions. Varied preprocessing as well:

Unprocessed:
"Kant's solution of the problem, though not valid in my opinion, is interesting."

Processed, No Punctuation:
"kants solution of the problem though not valid in my opinion is interesting"

Processed, Tagged:
"<person> solution of the problem, though not valid in my opinion, is interesting."

Models

Approach 1:
Adversarial Autoencoder

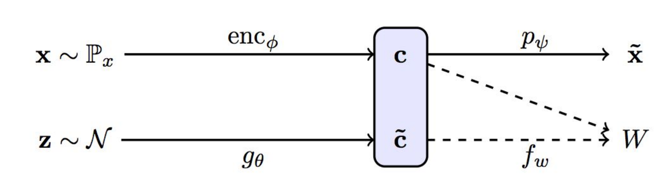

Approach 2:
Style Transformer

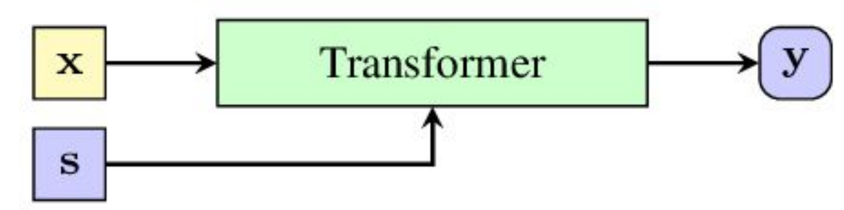

Results

Metrics weighing content similarity with style dissimilarity

Training Corpus	Model Settings	BLEU	PINC	PPL	F-K Ease	F-K Grade
Corpus 1		0.432	0.528	1044.24	81.97	5.60
Corpus 1	LR = 0.001	0.228	0.730	1638.63	67.08	9.10
Corpus 2		0.361	0.626	629.92	80.62	6.00
Corpus 2	NP	0.224	0.751	616.94	62.01	10.70
Corpus 3		0.381	0.587	282.45	75.20	6.80
Corpus 3	NP, GloVe	0.097	0.838	73.09	103.63	2.50
Corpus 3	GloVe	0.074	0.843	53.36	88.74	3.70
Corpus 3	Deep	0.514	0.434	390.92	72.16	6.40

Analysis

Significant style transfer, little content preservation:

Reference: "the sign should say what content it has so the people can stay away from the room"
Transferred: "the sky must begin what content it has so the drying can inadequate eagerly from the room"

Pre-trained embeddings lost content preservation and sentence structure:

Reference: "computers also allow people to store important files and pictures in places that wont get lost"
Transferred: "thus and in which the other people who must to the other people in which and that"

Conclusion

Key Findings

- The transformer trained on Corpus 3, containing tagged data, had best outcomes, with deeper networks adding to the differentiation in vocabulary
- More work needed to preserve content between reference and transferred texts

Future Study

- Improved grammar attention mechanisms
- Incorporation of more recent language models: BERT, GPT-2
- Parallel text supplementary networks for idiomatic phrases
- Retraining with increased computational resources

References:
[1] J. Zhao, Y. Kim, K. Zhang, A. M. Rush, and Y. LeCun, "Adversarially regularized autoencoders," 2017.
[2] N. Dai, J. Liang, X. Qiu, and X. Huang, "Style transformer: Unpaired text style transfer without disentangled latent representation,"Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019.