Name(s)
Alejandro Rodriguez

Project Number
38477

Project Title
BACON: Deep-Learning Powered AI for Poetry Generation with Author Linguistic Style Transfer

Abstract

Objectives/Goals
Can machines create art? Can a computer write meaningful poetry? The goal of this project is to develop a software system that automatically generates poetry (in English) with the style of any given author, what is known as linguistic style transfer.

Methods/Materials
The problem stated above is split into the following 3 sub-problems:
(1) Linguistic style modeling: this is achieved by extracting, from the author#s text corpus, the following linguistic features: (a) high-entropy n-grams: by building a TF-IDF (Term Frequency - Inverse Document Frequency) model; (b) relevant themes and topic words: by building an LDA (Latent Dirichlet Allocation) model
(2) Style transfer: use these probabilistic models to guide the automatic generation of poems. Achieved by probabilistic boosting of the previous features in the language model used by an Automatic Poem Generator (APG)
(3) Automatic poem generation: generate meaningful poetry (content) with rich aesthetic rules (form). This is built as a pipeline of a LSTM (Long Short-Term Memory) RNN (Recurrent Neural Network) -which addresses poem content generation-, and a WFST (Weighted Finite State Transducer) -which addresses poem form shaping-.

BACON, a basic software prototype that addresses the above mentioned methods, has been developed. It consists of several Python and Lua scripts, on a software platform built with several open source ML/DL modules and that extends an APG implementation developed by D. Kiela (Facebook AI Research) and J. Hopkins (Univ. Cambridge, UK).

Results
The quality and human-likeness of the poetry produced by BACON was tested through an extrinsic evaluation procedure, involving 62 human participants. The results indicate that participants were unable to tell the difference between human and computer-generated poems in any statistically significant way (Chi-squared test, p=0.85).

Conclusions/Discussion
The poetry generated by BACON approaches human poetry in aesthetic quality, while capturing some linguistic features of an author#s style. Therefore, the answer to the question posed in the project goal is positive. Future work will expand the set of features subject to linguistic style transfer.

Summary Statement
I developed a deep-learning powered AI software system that automatically generates poetry in the style of any given author.

Help Received
I initially reached out to D. Kiela (Facebook AI Research) and J. Hopkins (Univ. Cambridge, UK) for guidance about feasibility and about possible approaches. After that, the solution workflow, methods, and programming code that are described in the project were developed by me.