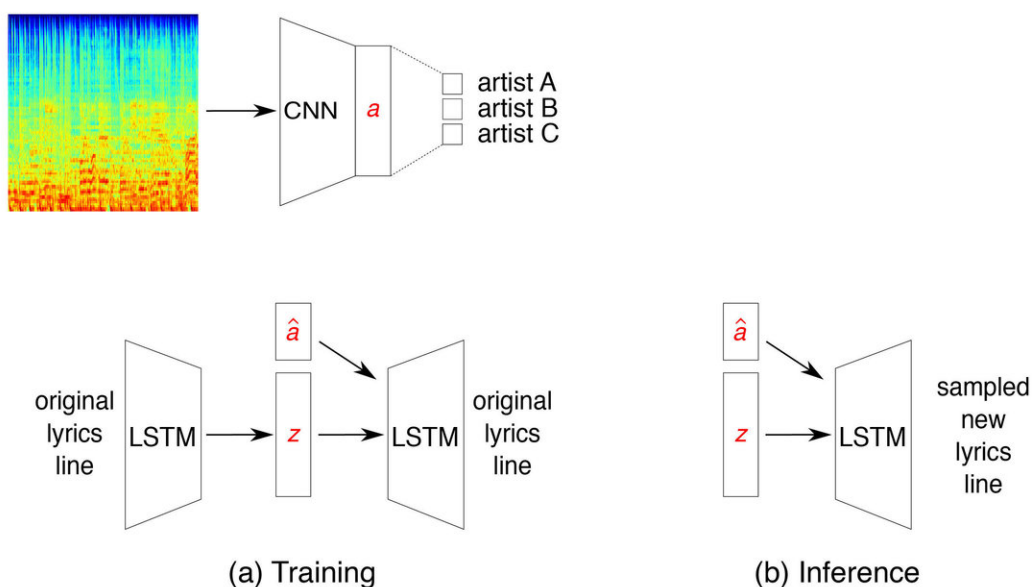


A system to generate new song lyrics that match the style of specific artists

January 8 2019, by Ingrid Fadelli



An overview of the approach by Vechtomova et al. First, a CNN is implemented to classify artists based on spectrogram images, thereby learning artist embeddings. Then, a VAE is trained to reconstruct lines from song lyrics, conditioned on the pre-trained artist embeddings. At inference time, in order to generate lyrics in the style of a desired artist, the researchers sample z from the latent space and decode it conditioned on the embedding of that artist. Credit: Vechtomova et al.

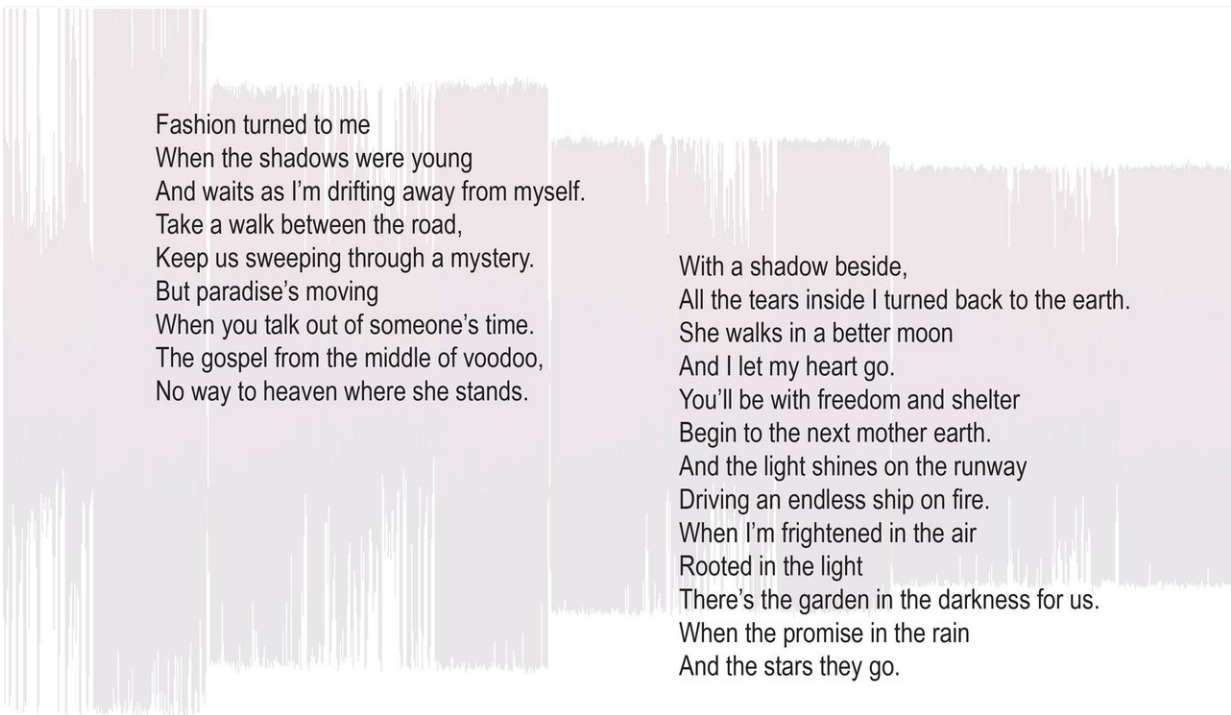
Researchers at the University of Waterloo, Canada, have recently developed a system for generating song lyrics that match the style of particular music artists. Their approach, outlined in [a paper pre-published on arXiv](#), uses a variational autoencoder (VAE) with artist embeddings and a CNN classifier trained to predict artists from MEL spectrograms of their song clips.

"The motivation for this project came from my personal interest," Olga Vechtomova, one of the researchers who carried out the study, told TechXplore. "Music is a passion of mine, and I was curious about whether a machine can generate lines that sound like the lyrics of my favourite music artists. While working on text generative models, my research group found that [neural networks](#) can generate some impressive lines of text. The natural next step for us was to explore whether a machine could learn the 'essence' of a specific music artist's lyrical style, including choice of words, themes and sentence structure, to generate novel lyrics lines that sound like the artist in question."

The system developed by Vechtomova and her colleagues is based on a neural network model called variational autoencoder (VAE), which can learn by reconstructing original lines of text. In their study, the researchers trained their model to generate any number of new, diverse and coherent lyric lines.

"To generate lines in the style of a given artist, we conditioned the generation on an artist embedding (i.e. a multi-dimensional vector of real numbers), learned by a separate neural network, which is trained to classify spectrograms of music audio clips by artists," Vechtomova said. "We then use the artist embeddings to condition the generation of lyrics lines in the style of each artist. The motivation behind this is that we want the differences between artist embeddings to reflect the differences in their lyrical as well as musical styles."

In a series of preliminary evaluations, the system developed by Vechtomova and her colleagues performed remarkably well. Their findings suggest that artist embeddings are useful for generating lyrics that match an artist's style. Many lines generated by the model were unmistakably aligned with the artist it was conditioned on, reflecting the themes generally addressed in his/her music.



Two poems generated by the system and included in the collection submitted to the NeurIPS 2018 Workshop on ML for Creativity and Design. Vechtomova created each poem by selecting lines generated by the VAE and arranging them in an artistically meaningful way. No editing was done to the individual lines, except for adding capitalization and punctuation marks. Credit: Vechtomova.

"While the generated lines often contain the words of an artist, these are

used in an interesting new way, expressing novel thoughts not found in the original lyrics," Vechtomova explained. "Some of the generated lines convey new and powerful poetic imagery, expressed using stylistic devices such as metaphors and oxymorons, while remaining true to the style of the artist."

In the future, the system created by Vechtomova and her colleagues could be used to inspire artists who are composing lyrics for new songs. Rather than replacing lyric composers, the researchers hope that it will provide new ideas, which artists can then mould, build upon and develop independently.

"The system is not meant to replace a music artist, but to be used as a source of inspiration during the songwriting process," Vechtomova said. "In the music world, this could be analogous to a synthesizer that can generate an infinite number of sounds, from which an artist then creates a song. Similarly, this tool can generate an infinite number of novel lines that artists can use in any way they like to compose their own lyrics."

As part of a different project, Vechtomova used the same system to generate intriguing poetry in the style of different music artists. The resulting collection of poems was accepted as an artwork at the [NeurIPS 2018 Workshop on ML for Creativity and Design](#).

"In the future, we plan to work on models that can learn new themes and vocabulary from additional sources, and use them to generate lyrics in the [style](#) of a given [artist](#)," Vechtomova said. "I would also like to explore how such a system could potentially be used by [music](#) artists as a source of inspiration."

More information: Generating lyrics with variational autoencoder and multi-modal artist embeddings. arXiv: 1812.08318 [cs.CL]. arxiv.org/abs/1812.08318

www.aiartonline.com/community/olga-vechtomova/

© 2019 Science X Network

Citation: A system to generate new song lyrics that match the style of specific artists (2019, January 8) retrieved 10 April 2024 from <https://techxplore.com/news/2019-01-song-lyrics-style-specific-artists.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.