

Spring 2020

GPT-2 AI Poetry Generation: Writing like Donne

Kaiya Case

Kenyon College, case1@kenyon.edu

Follow this and additional works at: https://digital.kenyon.edu/dh_iphs_prog

Recommended Citation

Case, Kaiya, "GPT-2 AI Poetry Generation: Writing like Donne" (2020). *IPHS 200: Programming Humanity*. Paper 17.

https://digital.kenyon.edu/dh_iphs_prog/17

This Article is brought to you for free and open access by the Digital Humanities at Digital Kenyon: Research, Scholarship, and Creative Exchange. It has been accepted for inclusion in IPHS 200: Programming Humanity by an authorized administrator of Digital Kenyon: Research, Scholarship, and Creative Exchange. For more information, please contact noltj@kenyon.edu.

GPT-2 AI Poetry Generation: Writing like Donne

Kaiya Case

IPHS 200.0: Programming Humanity

Introduction

We often view the arts as one area of study which seems the least prone to automation. We may even consider creativity and subjective thought to be a pillars of our very humanity, the abilities which shape our culture and set us apart as intelligent beings. However, recent developments in AI have begun to challenge that notion in regards to literature. While the technology is still far from independent, the GPT-2 neural network can be trained on an existing body of text in order to emulate the linguistic style and eventually produce new, original work. The platform has been most commonly used to imitate prose—given the ease of form and larger associated corpus—but this project focuses on its application to classical poetry, and how the more concise, metrical form may translate to machine generation.

Methodology

OpenAI's GPT-2 (scaled-up successor to GPT) is a large transformer-based, natural language processing system, developed for language modeling and text-generation [3]. There have been three different versions of the model, including 124M, 335M, and 774M, with the second parameter listed as the foundation for this project. The system is able to generate text by filtering the textual input as a series of tokens, with the simple objective to predict each next word [9]. It is intended to adapt to both the style and content of the training text.

In early 2019, Gwern Branwen was able to train the 117M GPT-2 on the Project Gutenberg poetry corpus, as well as a more contemporary dataset from the Poetry Foundation [1]. Training the model on a specific (yet substantial) dataset in this manner allows for more detailed control over the generated output text. This project attempts a very similar task, by training the 335M GPT-2 model on the work of specific poets, beginning with the more accessible and uniform structure of William Shakespeare's sonnets, followed by the generally longer-form, metaphysical school of John Donne.

The GPT-2 model was trained using 6,000 epochs, at various temperatures (including 0.7, 0.8, and 0.9, with higher quantities indicating a greater level of entropy), all in an attempt to avoid directly reproducing the existing poetry, as well as the generation of sheer nonsense. It was fed a txt file containing Shakespeare's 154 sonnets, followed by a reformatted compilation of Donne's collected poetic works (including his verse letters, which continue to exhibit a deliberate meter and rhyme scheme) in a separate trial, both of which were obtained from the greater Project Gutenberg corpus [4-6].

Results and Analysis

While the lower (0.7) temperature trials resulted in a large proportion of near-direct copies, after the first few epochs these become increasingly mixed in with some more original text. While one reprints the original sonnet 122, the following generation begins the same, but branches off with relative coherency:

```
<|sonnet|>
XXXI
```

```
Thy gift, thy tables, are within my brain
Full character'd with lasting memory,
```

```
<|sonnet|>
XXXII
```

```
Thy gift, thy tables, are but an echo
Of thought, and ghost--form'd but for thy pleasure;
```

In regards to form, the length of the consistent sonnet format (composed of three quatrains and a final couplet) is relatively well reflected in the GPT-2 generations, but the strict rhyme scheme is almost entirely lost. While the syntax is confused (beyond some basic sentence structure), each defined poem is able to maintain a relatively cohesive thematic core with its selected diction (i.e. one sonnet may appear to discuss nature and art in a more philosophical sense, using words such as “foliage” and “beauty,” while another angrily addresses the reader, including “cruel” and “regret” instead). The closest the model comes to replicating rhyme is frequent consonance between the sounds of the final word of each line, or even the plain repetition of the first word in the place of a rhyme, as highlighted below:

```
<|sonnet|>
CXIII
```

```
When all nature is at her most bold,
Steel's rough gate her fierce assault o'er;
The foliage their beauty in their sport,
Bearing the lion's snout in triumph o'er;
```

The results from the Donne corpus are much harder to sift through, as they vary greatly in form, and are significantly more prone to mode collapse across the earlier epochs, even at higher temperatures where appears (albeit decreasingly) all throughout the results (around 16 instances at 0.9). Even without complete collapse, many of the later poems exhibited an unnatural amount individual word or phrase repetition:

```
He who, to one that loves heretofore, knows not where,
Loves here, and here lives, because they are both one.
```

```
He who, on the other,
Loves here, because he can, and, as we all here,
Love because there is here to love, and he cannot.
```

The vague thematic currents seen throughout Donne's body of work are once again present (and generally aggregated) within the generated diction, but the three distinct periods of his development as a writer (from his early coterie poetry to his later devotional work) tend to blend together across the generated text, and are further diluted by the influence of his verse letters. However, while blind repetition may at first clog the meaning of passages like the one above, they also allow for some thematic clarity, and short, evocative phrases like the following may then emerge:

```
Or, if being could express nothing, nothing would be
more true.
Then would love be infinite, and eternity nothing.
```

These lines feel distinctly situated within the “Un-Donne” period of Donne's career, after his controversial marriage to Anne Moore, where there is a greater gravity to the concept of love, weighed against faith and his own social undoing. Donne is generally revered for his metaphysical style and emotional intensity, both of which are logically very difficult for GPT-2 to replicate convincingly. Much of his work relies on extended metaphor, which the generator is not able to coherently maintain over the course of a poem, yet certain resonant turns of phrase do occasionally emerge from the surface level, perhaps evocative of paradox but without a centralized purpose:

```
_The Draught._
```

```
Sleep sleep old Draught, thou canst not have repuls'd
So many times, as one sleepes, though hee
Sleepe unseparable number of them all.
Such freedom doth our body enjoy,
As one free SPHERITE from many;
```

5

```
The intelligence of a single being is more,
Then most insects, which have their larres, and doe
Within the cocke, and cocke the world, our home.
Our body, which cannot sinne, and yet makes us ill,
Enjoyes yet another year, another day. 10
There is still a little youth in us,
Though not as soone as thou, who hast beene
To sea, yet doth heare, and speak his minde.
O strong and long-liv'd death, how cam'st thou in?
Vera, thy selfe ample stomach, yet didst feed upon
Things, which would in future grow to be great;
```

Conclusion

Given such a limited corpus, along with the heightened significance of individual word choice inherent to the poetic medium, little past thematic word association and relatively isolated, lucid phrases can really be expected from the current GPT-2 model. Rhyme scheme is entirely neglected. Traces of Donne's reflective conceits begin to appear in the generated results throughout the latter half of the results, but are never able to pick up coherent momentum.

References

- Alexander, Scott. “Gwern's AI-Generated Poetry.” 14 Mar. 2019, <https://slatestarcodex.com/2019/03/14/gwerns-ai-generated-poetry/>.
- Anderson, Judith H., and Jennifer C. Vaught, editors. *Shakespeare and Donne: Generic Hybrids and the Cultural Imaginary*. Fordham University Press, 2013. JSTOR, www.jstor.org/stable/j.ctt13x02h9.
- “Better Language Models and Their Implications.” *OpenAI*. 14 Feb. 2019. <https://openai.com/blog/better-language-models/>.
- “The Poems of John Donne [2 vols.] Volume I.” <http://www.gutenberg.org/files/48688/48688-h/48688-h.htm>.
- “The Poems of John Donne, Volume II (of 2).” <http://www.gutenberg.org/files/48772/48772-h/48772-h.htm>.
- “The Project Gutenberg Etext of Shakespeare's Sonnets.” *Internet Archive*. <https://archive.org/stream/shakespearesson01041gut/wssnt10.txt>.
- Parish, Allison. “A Gutenberg Poetry Corpus.” *GitHub*. 13 Aug. 2018, <https://github.com/aparrish/gutenberg-poetry-corpus>.
- Vendeville, Geoffrey. “Can a computer write a sonnet as well as Shakespeare?” *University of Toronto*. 7 Aug. 2018, <https://www.utoronto.ca/news/can-computer-write-sonnet-well-shakespeare>.
- Woolf, Max. “How To Make Custom AI-Generated Text With GPT-2.” 4 Sept. 2019, <https://minimaxir.com/2019/09/howto-gpt2/>.