

Raymond Yeo
EAST 0303
Prof. Kai Jun Chen
December 13 2022

Artificial Intelligence and Chinese Poetry Generation

Introduction

Recent advancements in Artificial Intelligence (AI) have pervaded, and at times seamlessly integrated with, virtually all areas of human culture, ranging from medicine to manufacturing to entertainment. Its widespread acceptance in a variety of industries has one of the most straightforward explanations: AI improves human lives. At the same time, we must acknowledge that any sort of AI decision-maker is designed to perform well on the inputs it was trained on, which can lead to biases. Due to such widespread usage, it has become futile to handpick specific examples of how AI is used and is beneficial. Akin to explaining all the ways water can be used, it is unproductive to pick and choose arbitrary use cases of AI. Thus, an adequate understanding of our current technological state is that AI has been publicly accepted as improving human lives, holds inherent limitations on how it makes decisions because it is not sentient, and is integrated into nearly all processes within all industries.

However, production of art seems to be one of the few areas of expertise in which AI has not been able to assert itself into. There exists a general disdain toward machine-generated artwork, driven by accusations of devaluation of the human experience and craftsmanship. Interestingly, such perspectives can be contradictory, as both the production and evaluation of art is a subjective, often private, manner. AI programs produce novel art and have, quite frankly, progressed to points in which their work can be indistinguishable from human work. For

example, an AI generated piece recently won a prize at Colorado State Fair’s annual art competition.¹

I focus my discussion on a particular subset of AI art generation: natural language processing (NLP) for generating Chinese poetry. NLP refers to the creation of programs in which language can be “understood,” where understanding is characterized by a learned pattern exhibited by the text the model was trained on. There are many algorithms and mathematical techniques, increasing in both conceptual complexity and execution time, that can be used to design NLP programs. Of such methodologies, neural networks, one of the most advanced techniques within the field of AI, are perhaps the most important.


The goal of this paper is to be both a gentle introduction to the applications of AI in literature generation and a meditation on how to come to terms with it. Through talking about the architecture of Chinese poetry generators, I hope to illustrate the assumptions that are being made by these models, and how they fit into our understanding of what poetry should be.

Neural Networks

Neural networks mimic the architecture of the human brain; a series of decision-makers called nodes (i.e. neurons) take in data, run their own mathematical operation, and choose to “fire” or not. This process, shown in the diagram below, can go on for several layers before producing a final output, representing how complex the neural network is. The power of neural networks lies in their ability to learn from new data and constantly adapt each node’s decision rules. The layers of nodes, paired with a multi-step learning and decision-making process, is similar to how the human brain learns and processes new information. In contrast, traditional AI

¹ Roose, Kevin. “An A.I.-Generated Picture Won an Art Prize.”
<https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>.

algorithms are static and resemble a cause-and-effect chart rather than a dynamic, learning system.



2

The reason for this brief interlude into computer science and mathematics is to create a foundation to discuss both the effectiveness and ethics of machine poetry generation. From a purely technical perspective, one can easily argue that the marriage between NLP and neural networks is a near perfect recipe for effectively modeling abstract concepts that are based on experience. To a certain extent, there must be some degree of appreciation that a computer program is somehow “learning” how to write poems by following a set of rules it has miraculously gleaned from a heap of training data. At any point, we can peel back a neural network’s layers and observe the multitude of impressions it has accumulated and the exact weighting of these. In other words, while it may not be pretty, a neural network can remarkably tell us exactly how to write a quality poem.

² IBM. “Neural Networks.” <https://www.ibm.com/cloud/learn/neural-networks>.

Who is the Poet?

Poetry is truly a dance between complete randomness and total control; certain lines of a poem are unintelligible, while others are so vivid they last with us for a lifetime. Unlike academic writing, we crave such ambiguity and play from poems. It's not just about understanding all the nuances and character choices, but rather the entire experience of starting a poem with a certain emotional and mental state, and freeing yourself to be temporarily molded into something different.

The role of the poet is worth noting in this process. In a sense, they are the only individuals who can understand and relate with all aspects of the poem. There is an inherent beauty in the process of readers trying to decipher as much of another human's emotions and thoughts. That, in and of itself, can be enough of a romantic reason for disregarding machine generated poems. However, it is also flawed to think poems are raw, unadulterated thoughts. Rather, all poems are thoughts that have been forced into rigid word definitions, a finite vocabulary, predetermined poetry patterns, all while attempting some semblance of aesthetic writing.

Assuming that we don't consider AI programs to be true poets, we are met with the following question: If there is beauty in having just a singular author who understands the work through and through, what if there is no true author? AI generated poems are the world's first set of comprehensible poems in which no one truly understands all the nuances about. If one believes that a lack of a true author violates poetry, then we can reason they also place a stronger emphasis on the poet-reader bond established through literary work. Art is a cry to be heard, where artists package their emotions into abstractions and symbols, in the hopes that their work will resonate with others. In that case, there is a requirement for there to be an author; an identity

behind the work. Others may believe that a lack of a true author amplifies the poetic work, revealing our natural tendencies to project emotions onto art. It can be difficult to truly relate to a poet: they are often from a different century, continent, religion, culture, sexuality, gender, and socioeconomic status from you. In that case, art can be treated as a vessel for individuals to come to terms with their own, unique emotions. This projection onto art and poetry requires no need to take a step back and view the work from the artists' true perspective. Art and poetry become utilitarian.

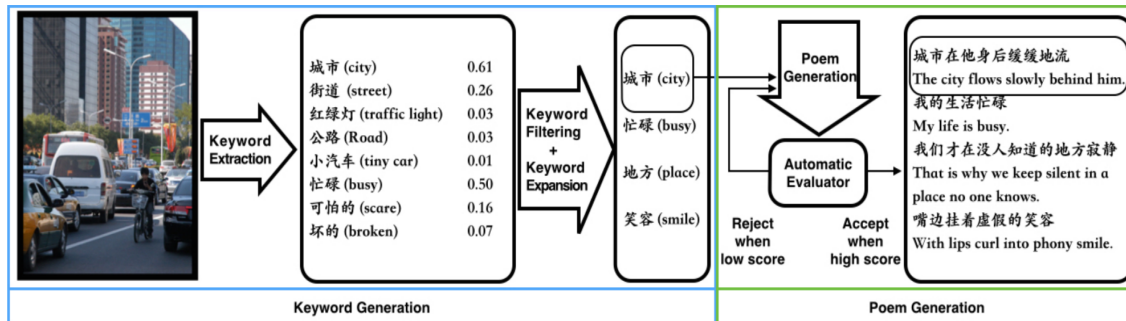
Given these frameworks and questions to consider, I now address two specific advancements in AI generated Chinese poetry: Microsoft XiaoIce and the Novel Polishing Framework.

Microsoft XiaoIce

In 2017, a poetry collection titled, “Sunshine Misses Windows” was published. The work was widely discussed, for this marked the first-ever poetry collection written by an AI named “XiaoIce”. Like all neural networks, XiaoIce had to be trained in order to learn how to write poems.³ To do so, researchers opted to teach XiaoIce using modern poems written by over a select five-hundred Chinese poets within the past century. In order to gain inspiration for writing, XiaoIce takes in image input and extracts key features from it through a computer vision neural network. For example, a wintery photo of the Alps may be extracted into words such as “mountain”, “snow”, and “towering”. XiaoIce then takes these words and fits them into its

³ People’s Daily Online. “First AI-authored collection of poems published in China.” <http://en.people.cn/n3/2017/0531/c90000-9222463.html>.

understanding of what poetry is, to create an AI generated poem.⁴ The keyword to poem generation flow is shown in the image below.



5

The XiaoIce researchers make several convenient assumptions to support their technical decisions. First is that if an AI model is provided with an image, it must generate a meaningful poem out of it. The underlying assumption here is that any picture can be crafted into a literary work, even if it is rejected multiple times by the automatic evaluator that XiaoIce uses. In practice this may be true, but it seems to violate the very nature of poetry. The thought of forcing a poet to sit down and produce quality work based on an image that they cannot find the words for is cruel and unartistic. Perhaps there should exist a stop mechanism which terminates the generation process for uninspiring images. There should be no requirement for a program to be able to produce something from anything.

From my understanding of XiaoIce’s research paper and prior knowledge on image extraction, there is a significant difference in a machine’s ability to identify nouns and derive adjectives. This is consistent with our conceptual understanding of the world: it’s easy to determine whether there is a cat in a picture or not, but difficult to describe the mood of the room that it is lounging in. Thus, an inherent limitation in the model is that it is biased toward

⁴ Cheng, Wen-Feng et al., 2018. “Image Inspired Poetry Generation in XiaoIce.” <https://arxiv.org/pdf/1808.03090.pdf>.

⁵ Ibid.

acknowledging these features it is most confident in observing. The researchers share that, “we choose to use the keywords that have not only high confidence in image recognition but also enough occurrences in training corpus.”⁶ From a technical perspective, this is a logical and safe decision for the program to make. At the same time, one cannot help but shake their head at such a lost opportunity to take a bolder learning approach. Unlike humans, computers can analyze each individual pixel in an image and find patterns we cannot fathom obtaining. To the idealist, poems bringing forth such insights would be otherworldly.

Due to the technical nature of the XiaoIce research paper, there are two important questions which cannot be answered by my understanding. First is whether or not XiaoIce can address objects it recognizes in images without directly referencing them. To illustrate this point, we can reference the keyword to poem generation flow once again. It is reasonable to assume a talented human poet would be able to write some metaphor of a mouse in a maze to allude to the hustle-and-bustle of city life. Can XiaoIce do the same? Additionally, I wonder how deterministic the XiaoIce model is. Will the same image always produce the same set of poems? If they differ, then do they address completely different parts of the image, or are low-confidence adjectives replaced? My intuition is that the model would opt for the latter approach, further exacerbating the lost potential for creating truly novel poetry.

Polishing Framework

Art is an iterative process. Poets have several reasons for changing up their initial drafts. They may feel as though they were not capturing the right emotions, overemphasizing a trivial detail, or using an imperfect word. There can be infinite reasons for altering one’s work, with a simple “because I felt like changing it” being a valid justification as well.

⁶ Ibid.

The paper, “An Iterative Polishing Framework based on Quality Aware Masked Language Model for Chinese Poetry Generation,” offers a novel methodology for “polishing” poems to improve them. Researchers here opted to train their model on traditional quatrains written from pre-modern China; in particular, the Tang, Song, Yuan, Ming, and Qing Dynasties. Once initial drafts of poems were generated, they were evaluated by a separate Quality-Aware Masked Language Model (QAMLM), which identified improper characters in the context of the entire poem.⁷

Constraining the model to four line poems with five to seven characters truly limits the number of potential outputs. Even with such structural limitations, the model was trained on 130,525 poems with a total 905,790 number of poem lines.⁸ Therefore, it is very likely that computer generated poems would have significant overlap with pre-existing works. While critics may be quick to accuse the algorithm for plagiarism, a more optimistic perspective could be that the AI was able to produce a poem that blended together the semantically shared emotions of multiple individuals. Yet, I believe this to be more of an unintentional consequence; the program only cares for producing what it considers to be a high quality poem. It would be a different story if the algorithm purposefully sought to extract and connect poets based on their shared sentiments.

An assumption that I would like to challenge is the motivation behind implementing a polishing framework. The researchers write that, “Our model only polishes a small part of low-quality characters instead of rewriting the whole text. The polishing process implemented by our model is independent of the draft generation process, which can polish the draft generated by

⁷ Deng, Liming et al., 2019. “An Iterative Polishing Framework based on Quality Aware Masked Language Model for Chinese Poetry Generation.” <https://arxiv.org/pdf/1911.13182.pdf>.

⁸ Ibid.

another separate model.”⁹ For human writers, the poetry writing and editing process is a continuous one that requires complete understanding of one’s work. The polishing framework in the paper asserts that poems can be objectively improved, even with a disconnect between the poem writer and editor. In other words, there is a singular “perfect” poem that exists and that the AI polisher seeks to obtain. I do not argue that it is not possible to find an objectively better written poem. Rather, the flaw in this methodology lies in the fact that perfection is not what we seek in poems. We recognize there is a frustration in getting words down onto paper; as humans we know how difficult that can be. There is a mutual understanding when reading poems that while the poet did all they could to convey their feelings, there is a chance they were unable to capture all of their emotions. These are not flaws in a poem, but rather opportunities for the reader to fill in the gaps with their own impressions.

Conclusion

In both XiaoIce and the Polishing Framework paper, the poetry generation models were built and trained on Chinese poems. On the other hand, humans are “trained” on life experiences. These life experiences can lead people to produce poetry as a mechanism for reflecting and coping with their emotions. Thus, the inability for these neural networks to be trained on human experiences, the inability for them to live a human life, is arguably the most fatal flaw. To deal with this “lack of life”, machines must filter through infinite perspectives provided to them in training; each perspective a bundle of a lifetime's worth of experiences. From this, the machine chooses what is most convenient, what is most “safe”, what would yield the most natural sounding verse. Humans do not have that sort of luxury. We possess a singular perspective, and in poetry we express the experiences and emotions that we feel we must share with the world.

⁹ Ibid.