X575: writing rengas with web services

Daniel Winterstein
Winterwell Associates
daniel@winterwell.com

Joseph Corneli
Goldsmiths College, University of London
j.corneli@gold.ac.uk

Abstract

Our software system simulates the classical collaborative Japanese poetry form, *renga*, made of linked haikus. We used NLP methods wrapped up as web services. This approach is suitable for collaborative human-AI generation, as well as purely computer-generated poetry. Evaluation included a blind survey comparing AI and human haiku. To gather ideas for future work, we examine related research in semiotics, linguistics, and computing.

1 Introduction

Computer haikus have been explored in practice at least since Lutz (1959). More recently, haikus have been used by Ventura (2016) as the testbed for a thought experiment on levels of computational creativity. As we will discuss below, the classic haiku traditionally formed the starting verse of a longer poetry jam, resulting in a poem called a *renga*. A computational exploration of renga writing allows us to return to some of the classical ideas in Japanese poetry via thoroughly modern ideas like concept blending and collaborative AI.

Ventura's creative levels range from *randomisation* to *plagiarisation*, *memorisation*, *generalisation*, *filtration*, *inception*[1] and *creation*. Further gradations and criteria could be advanced, for example, the fitness function used for filtration could be developed and refined as the system *learns*. Creativity might be assessed in a social context, as we investigate how a system *collaborates*.

While self-play was a good way for the recently developed board game-playing system AlphaGo to transcend its training data (Silver et al., 2016), we do not yet have computationally robust qualitative evaluation measures for the poetry domain, where there is no obvious "winning condition." We began by creating a program for generating haikus, trained on a small corpus. Our technical aim then was to simulate the collaborative creation of renga, i.e., linked haikus. There are several forms of renga with varying constraints (Carley, 2015), for example the 20 stanza "Nijiun" renga which alternates between two-line and three-line verses, with a focus on seasonal symbolism and rules against repetition.[2] Our initial effort was a technical success, however the rengas we produced fail to fully satisfy classical constraints. A subsequent experiment is more convincing in this regard, but still leaves room for improvement. Our discussion considers the aesthetics of the generated poems and outlines directions for future research.

2 Background

Coleridge considered poetry to be "the blossom and the fragrance of all human knowledge." AI researcher Ruli Manurung defines poetry somewhat more drily: "A *poem* is a natural language artefact which simultaneously fulfils the properties of meaningfulness, grammaticality and poeticness" (Manurung, 2004, p. 8).

The *haiku* as we know it was originally called

[1] "[I]nject[ing] knowledge into a computationally creative system without leaving the injector's fingerprints all over the resulting artifacts."

[2] http://www.renga-platform.co.uk/webpages/renga_01.htm

hokku – 発句, literally the "starting verse" of a collaboratively written poem, *hakai no renga*. Typically each of following links in a renga take the familiar 5/7/5 syllable form. Classical rengas vary in length from two to 100 links (and, rarely, even 1000). The starting verse is traditionally comprised of two images, with a *kireji* – a sharp cut – between them. The term *haiku* introduced by the 19th Century poet Masaoka Shiki supersedes the older term. Stylistically, a haiku captures a moment.

In classical renga, all of the verses after the first have additional complex constraints, such as requiring certain images to be used at certain points, but disallowing repetition, with various proximity constraints. The setting in which rengas were composed is also worth commenting on. A few poets would compose together in party atmosphere, with one honoured guest proposing the starting haiku, then the next responding, and continuing in turn, subject to the oversight of a scribe and a renga master. These poetry parties were once so popular and time consuming that they were viewed as a major decadence. Jin'Ichi et al. (1975) offers a useful overview.

Because of the way we've constructed our haiku generating system, it can take an entire haiku as its input topic – we just add the word vectors to make a topic model – and compose a response. This affords AI-to-AI collaboration, or AI-human collaboration. It can also blend two inputs – for example, the previous haiku and the current constraint from the renga ruleset (e.g., the requirement to allude to "cherry blossoms" or "the moon").

3 Implementation

Working with a small haiku corpus, we used a POS tagger to reveal the grammatical structure typical to haikus. The CMU Pronouncing Dictionary is used to count syllables of words that fill in this structure.[3] The Brown corpus was used to generate n-grams, and the generation process prefers more common constructions in haikus.[4] Wikipedia data was processed with GloVe (Pennington et al., 2014) to create a semantic vector space model of topics, based on word co-occurrences.[5] Adding a web API turned the haiku generating system into a haiku server, and facilitated subsequent work with FloWr. In short:

```
1. Haiku corpus → POS tagger →
   grammatical skeleton fragments.
2. General text corpus → n-gram model.
3. Wiki corpus → topic vectors.
4. Combine skeleton fragments to make a
   haiku template.
5. Assign syllable counts to slots.
6. Fill in the template, preferring
   n-grams and close topic matches.
7. Wrap the process with a JSON HTTP
   API
```

4 Experiments

I. Initial evaluation of haikus Following Manurung's definition of poetry, above, we would like to assess: (1) whether a given haiku makes sense and how well it fits the topic, (2) whether it fits the form, i.e., is it a valid haiku?, and (3), the beauty of the writing, the emotion it evokes. Details of a survey-based blind comparison of human and computer-written haikus were written up by Aji (2015). The system was then extended with multiple inputs, in some cases producing interesting blends: e.g., the following in response to "frog pond" and "moon":

that gull in the dress –
vivacious in statue
from so many ebbs

II. Generation of rengas Here are two rengas generated by wrapping the haiku API inside the FloWr flowchart system (Charnley et al., 2016):

fertile forefingers *took orchard for my lather* *brackish was cherished*	*that vase in the quilt –* *the effeminate of names* *with a colored juice*
toddler of strong bet *foaling feels to a good tooth* *thriving like a paw*	*cases of sibyl* *and a stylish curators* *from downed in the aim*
a drawer straight inside *under the slicked interim* *to shrink the safe cute*	*figures of digress* *and a sumac excises* *from key in the ribbed*
readjusted blots *in the creativity –* *one child at a love*	*cluster for icebergs –* *and a waging everglades* *from huge in the drug*

In each case, the prompt for the first link is "flower blossom" and each link is passed on to

[3] http://www.speech.cs.cmu.edu/cgi-bin/cmudict

[4] https://en.wikipedia.org/wiki/Wikipedia:Database_download

[5] http://nlp.stanford.edu/projects/glove/

the next link along with a secondary prompt. The secondary links are "moon," "autumn," and "love," respectively. For the first renga, we designed a flowchart that selects the "most positive" haiku from the ten that the haiku API returns, using the AFINN word list.[6] In the second renga, we designed a flowchart to select the haiku with the lowest word variety (computed in terms of Levenshtein distance).

III. Tuning the parameters for the Nijiun form

national forces
on the government afghan
because no-one armed

military government
the war on the armed afghan

a soviet troops
upcoming in the winter
military armed

troops fighting military
in the administration

blooms operations
the military planted
of the flower war

the army of government
we die the war civilians

We made improvements to the use of the Brown corpus to utilise n-grams for word-flow and sense, as well as tuning the weightings given to sense and topic. We implemented the injection of topics via by blending, as per classical constraints (e.g., required seasonal themes like "winter," or "flowers" in the penultimate link). At left, we quote the closing links of the first Nijiun renga generated by our software.

5 Discussion and Related Work

Towards automated evaluation Some aspects of the evaluation dimensions are built into the way the poems are constructed.

Form: the haiku-generating subsystem guarantees that the requirements of a grammatical skeleton are met, and the 5/7/5 syllable pattern is guaranteed (up to the accuracy of the CMU Pronouncing Dictionary). Surface form scales up well for rengas.

Sense: the haiku generating subsystem uses an n-gram model of text likelihood, which will yield a higher score for constructions that match frequently observed phrases. In our first round of experiments with rengas, sense tended to degrade quickly. Our subsequent adaptations to the renga generation algorithm prioritise greater continuitity between links.

Topic: we used a vector model of the topic word(s), and can measure the distance to the vector given by the sum of the words in the poem.

[6] `http://neuro.imm.dtu.dk/wiki/AFINN`

Emotion: In our experiment with FloWr, we used a quite simple method, filtering a list for the "most positive" haikus. Mohammad (2016) surveys more recent work in NLP on modelling emotion, which could be exploited in future work.

Beauty: Waugh (1980) points out that language is based on a "hierarchy of signs . . . of ascending complexity, but also one of ascending freedom or creativity," and also remarks that a "poem provides its own 'universe of discourse.'" To some extent these criteria pull in opposite directions: towards complexity, and towards coherence, respectively. Our first rengas could not be reasonably described as a 'universe of discourse' but rather, a 'universe of random nonsense'. This is improved in the subsequent experiment. Traditional rengas forbid repetition, and discourage overt reflection on themes like death, war, illness, impermanence, religion and sex (Carley, 2015, p. 89). Thus, despite being coherent, the repetitive "military" theme in the final example above is not appropriate to classical constraints. A reader may identify some fortuitous resonances, e.g., "the flower war" is interesting within the "afghan" context established in earlier links – but the system does not yet recognise these features.

Some paths forward Wiggins and Forth (2015) use hierarchical models in a system that builds a formative evaluation as it composes or reads sentences, judging how well they match learned patterns. While this seems to have more to do with constraints around typicality, per Waugh, there is room for creativity within hierarchies. Hoey (2005) makes a convincing argument that satisfying lexical constraints while violating some familiar patterns may come across as interesting and creative.

Word similarities can be found using GloVe: this would presumably produce links with more coherent meanings, compared to the edit distance-based measure we used. Ali Javaheri Javid et al. (2016) use *information gain* to model the aesthetics of cellular automata. Can these ideas be combined to model evolving topic salience, complexity, and coherence?

If the system provided a *razo* (the troubadours' jargon for "rationale"; see Agamben (1999, p. 79)), we could debug that, and perhaps involve additional AI systems in the process (Corneli et al., 2015).

6 Conclusion

In terms of Ventura's hierarchy of creative levels, the haiku system appears to be in the "generalisation" stage. Our renga-writing experiments with FloWr brought in a "filtration" aspect. The research themes discussed above point to directions for future work in pursuit of the "inception" and "creativity" stages.

Some previous work with haiku, e.g. Netzer et al. (2009) and Rzepka and Araki (2015), have addressed the problem of *meaning*. The renga form brings these issues to the fore. We hope this early work has motivated further interest in this challenging and enjoyable poetic form that – like other less constrained forms of dialogue – combines themes of natural language generation and understanding. One natural next step would be a series of experiments in collaborative human-AI generation of rengas. Our haiku software is available for future experiments.[7]

Acknowledgement

This research was supported by the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open Grant number 611553 (COINVENT).

References

Giorgio Agamben. 1999. *The end of the poem: Studies in poetics*. Stanford University Press.

Alham Fikri Aji. 2015. Automated haiku generation based on word vector models. Master's thesis. University of Edinburgh.

Mohammad Ali Javaheri Javid, Tim Blackwell, Robert Zimmer, and Mohammad Majid al Rifaie. 2016. Analysis of information gain and Kolmogorov complexity for structural evaluation of cellular automata configurations. *Connection Science*, 28(2):155–170.

John Carley. 2015. *Renku Reckoner*. Darlington Richards Press.

John Charnley, Simon Colton, Maria Teresa Llano, and Joseph Corneli. 2016. The FloWr Online Platform: Automated Programming and Computational Creativity as a Service. In Amílcar Cardoso, François Pachet, Vincent Corruble, and Fiammetta Ghedini, editors, *Proceedings of the Seventh International Conference on Computational Creativity, ICCC 2016*.

[7] https://github.com/winterstein/HaikuGen

Joseph Corneli, Anna Jordanous, Rosie Shepperd, Maria Teresa Llano, Joanna Misztal, Simon Colton, and Christian Guckelsberger. 2015. Computational poetry workshop: Making sense of work in progress. In S. Colton, H. Toivonen, M. Cook, and D. Ventura, editors, *Proceedings of the Sixth International Conference on Computational Creativity, ICCC 2015*.

Michael Hoey. 2005. *Lexical priming: A new theory of words and language*. Psychology Press.

Konishi Jin'Ichi, Karen Brazell, and Lewis Cook. 1975. The Art of Renga. *Journal of Japanese Studies*, pages 29–61.

Theo Lutz. 1959. Stochastische texte. *Augenblick*, 4(1):3–9.

Hisar Maruli Manurung. 2004. *An evolutionary algorithm approach to poetry generation*. Ph.D. thesis. University of Edinburgh.

Saif M. Mohammad. 2016. Sentiment Analysis: Detecting Valence, Emotions, and Other Affectual States from Text. In Herb Meiselman, editor, *Emotion Measurement*. Elsevier.

Yael Netzer, David Gabay, Yoav Goldberg, and Michael Elhadad. 2009. Gaiku: Generating haiku with Word Associations Norms. In Anna Feldman and Birte Loenneker-Rodman, editors, *Proceedings of the Workshop on Computational Approaches to Linguistic Creativity*, pages 32–39. ACL.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

Rafal Rzepka and Kenji Araki. 2015. Haiku Generator That Reads Blogs and Illustrates Them with Sounds and Images. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 2496–2502. AAAI Press.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489.

Dan Ventura. 2016. Mere Generation: Essential Barometer or Dated Concept? In Amílcar Cardoso, François Pachet, Vincent Corruble, and Fiammetta Ghedini, editors, *Proceedings of the Seventh International Conference on Computational Creativity, ICCC 2016*.

Linda R Waugh. 1980. The poetic function in the theory of Roman Jakobson. *Poetics Today*, 2(1a):57–82.

Geraint A Wiggins and Jamie Forth. 2015. IDyOT: a computational theory of creativity as everyday reasoning from learned information. In *Computational Creativity Research: Towards Creative Machines*, pages 127–148. Springer.