

Learning Rhyming Constraints using Structured Adversaries

Harsh Jhamtani¹ Sanket Vaibhav Mehta¹ Jaime Carbonell¹ Taylor Berg-Kirkpatrick²

¹ School of Computer Science, Carnegie Mellon University

² Department of Computer Science and Engineering, University of California San Diego

{jharsh, svmehta, jgc}@cs.cmu.edu, tberg@ucsd.eng.edu

Abstract

Existing recurrent neural language models often fail to capture higher-level structure present in text: for example, rhyming patterns present in poetry. Much prior work on poetry generation uses manually defined constraints which are satisfied during decoding using either specialized decoding procedures or rejection sampling. The rhyming constraints themselves are typically not learned by the generator. We propose an alternate approach that uses a structured discriminator to learn a poetry generator that directly captures rhyming constraints in a generative adversarial setup. By causing the discriminator to compare poems based only on a learned similarity matrix of pairs of line ending words, the proposed approach is able to successfully learn rhyming patterns in two different English poetry datasets (Sonnet and Limerick) without explicitly being provided with any phonetic information.

1 Introduction

Many existing approaches to text generation rely on recurrent neural networks trained using likelihood on sequences of words or characters. However, such models often fail to capture overall structure and coherency in multi-sentence or long-form text (Bosselut et al., 2018; Holtzman et al., 2018). To rectify this, prior work has proposed losses which encourage overall coherency or other desired behavior (Li et al., 2016; Zhang and Lapata, 2017; Bosselut et al., 2018). However, most of these approaches rely on manually provided definitions of what constitutes a good or suitable structure, thereby limiting their applicability. In this paper we propose a method for English poetry generation that directly learns higher-level rhyming constraints as part of a generator without requiring strong manual intervention. Prior works on poetry generation (Oliveira, 2017; Ghazvininejad et al.,

2018) have focused mostly on ad-hoc decoding procedures to generate reasonable poetry, often relying on pruning from a set of candidate outputs to encourage desired behavior such as presence of explicitly-defined rhyming patterns.

We propose an adversarial approach to poetry generation that, by adding structure and inductive bias into the discriminator, is able to learn rhyming constraints directly from data without prior knowledge. The role of the discriminator is to try to distinguish between generated and real poems during training. We propose to add inductive bias via the choice of discriminator architecture: We require the discriminator to reason about poems through pairwise comparisons between line ending words. These learned word comparisons form a similarity matrix for the poem within the discriminator’s architecture. Finally, the discriminator evaluates the poem through a 2D convolutional classifier applied directly to this matrix. This final convolution is naturally biased to identify spatial patterns across word comparisons, which, in turn, biases learned word comparisons to pick up rhyming since rhymes are typically the most salient spatial patterns.

Recent work by Lau et al. (2018) proposes a quatrain generation method that relies on specific domain knowledge about the dataset to train a classifier for learning the notion of rhyming: that a line ending word always rhymes with exactly one more ending word in the poem. This limits the applicability of their method to other forms of poetry with different rhyming patterns. They train the classifier along with a language model in a multi-task setup. Further, at generation time, they heavily rely on rejection sampling to produce quatrains which satisfy any valid rhyming pattern. In contrast, we find that generators trained using our structured adversary produce samples that satisfy rhyming constraints with much higher frequency.

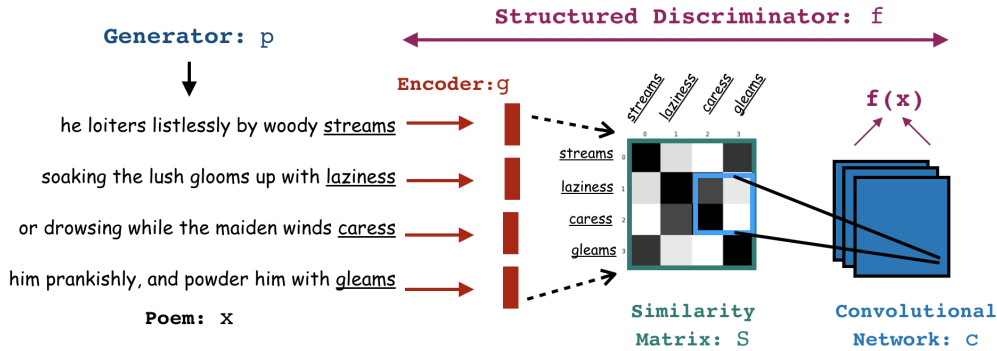


Figure 1: Model Overview: We propose a structured discriminator to learn a poetry generator in a generative adversarial setup. Similarities between pairs of end-of-line words are obtained by computing cosine similarity between their corresponding representations, produced by a learned character-level LSTM encoder. The discriminator operates on the resulting matrix S representing pair-wise similarities of end words. The proposed discriminator learns to identify rhyming word pairs as well as rhyming constraints present in the dataset without being provided phonetic information in advance.

Our main contributions are as follows: We introduce a novel structured discriminator to learn a poetry generation model in a generative adversarial setup. We show that the discriminator induces an accurate rhyming metric and the generator learns natural rhyming patterns without being provided with phonetic information. We successfully demonstrate the applicability of our proposed approach on two datasets with different structural rhyming constraints. Our poem generation model learned with the structured discriminator is more sampling efficient compared to prior work – many fewer generation attempts are required in order to obtain a valid sample which obeys the rhyming constraints of the corresponding poetry dataset.

2 Method

Many forms of poetry make use of rhyming patterns on line-ending words (Oliveira, 2017). Therefore, to characterize a rhyming poem, a model needs (1) to know what it means to rhyme (2) to identify the specific permissible rhyming patterns for a particular poem type. For example, a limerick is a 5 line poem with a rhyming constraint of the type AABBA, i.e. the ends of the first, second, and fifth lines rhyme. We consider an adversarial learning setup with a hierarchical language model and a structured discriminator, such that the discriminator is trained to distinguish between generated examples and training examples, and the generator is trained to *fool* the discriminator. Our novel structured discriminator operates on a matrix which encodes a *learned* pair-wise similarity function of the line ending words. We refer to our model as **RHYME-GAN**.

2.1 Neural Generation Model

Our generator is a hierarchical neural language model (Figure 1) that first generates a sequence of line-ending words, and thereafter generates the poem’s lines conditioned on the ending words. We use recurrent neural networks for ending word generation as well line generation conditioned on ending words. Following prior work (Lau et al., 2018), we generate words in each line in reverse order (i.e. right to left), and begin generation with the last line first. Let \hat{x} represent a sample from the current generator distribution, denoted by p_θ , where θ represents the generator parameters. We initialize the word embeddings in the generator with pre-trained word embeddings (Lau et al., 2018) trained on a separate non-sonnet corpus.

2.2 Structured Discriminator

We introduce a structured discriminator, denoted by function $f_\phi(x)$, which outputs the probability that x is a sample from the dataset as opposed to generated. Our architecture defines an intermediate matrix $S \in R^{T \times T}$, where T denotes the number of lines in the poem, which encodes pair-wise similarities between line ending words in order to capture rhyming structure. The discriminator’s output is determined by a two layer 2D convolutional neural network applied to S . Convolutional neural networks have been shown to capture local as well as global patterns in 2D data – for example, images. Thus, our discriminator is composed of two main components: computation of a ma-

trix S , and a convolutional neural network to classify the computed matrix S . The pair-wise computation provides a useful inductive bias to identify the notion of rhyming, whereas the convolutional network is a suitable choice to capture overall rhyming patterns.

More specifically, let the words at the ends of lines in x be denoted by e . The number of ending words will be same as the number of lines in x , which we denote as T . We encode each ending word using a character-level LSTM (Hochreiter and Schmidhuber, 1997) denoted by g_{ϕ_g} , and use the last hidden state of the LSTM as a vector representation of the word. We let S_{ij} be the cosine similarity between the representations of ending words e_i, e_j , given by following equation:

$$S_{ij} = \frac{g(e_i)g(e_j)}{|g(e_i)||g(e_j)|} \quad (1)$$

The matrix S is passed through a convolutional neural network composed with a linear layer, together denoted by c_{ϕ_c} . The final output is passed through a sigmoid non-linearity, so that $f_{\phi}(x) \in [0, 1]$. The value of $f_{\phi}(x)$ represents the discriminator’s assessment of the probability that datum x belongs to the *real* dataset, rather than being a generated sample. The discriminator’s objective will train it to distinguish between a sample x from training data \mathcal{X} , and a generated sample \hat{x} , in a binary classification setup. Specifically, we define the discriminator loss for x, \hat{x} as follows:

$$d(x, \hat{x}; \phi) = -\log(f_{\phi}(x)) - \log(1 - f_{\phi}(\hat{x})) \quad (2)$$

2.3 Learning

Generator parameters θ and discriminator parameters ϕ are trained together under following objective:

$$\min_{\theta} \left[\mathbb{E}_{x \in \mathcal{X}} \left[-\log p_{\theta}(x) + \lambda \max_{\phi} \mathbb{E}_{\hat{x} \sim p_{\theta}} [-d(x, \hat{x})] \right] \right] \quad (3)$$

Note, in addition to using a traditional adversarial objective, we also include a likelihood term to help stabilize the generator. λ is a hyperparameter which controls the relative weight of the two terms. Since sampling of \hat{x} from generator involves discrete choices, we use the REINFORCE (Williams, 1992) algorithm to train the generator using learning signal from the adversarial loss term. The generator simultaneously gets an exact gradient from the likelihood portion of the objective. We observe training is more stable when

| Model | Expected #Samples | |
|------------------|-------------------|------------|
| | SONNET | LIMERICK |
| DEEP-SPEARE | 153.8 | N/A |
| RHYME-LM | 169.5 | 500 |
| RHYME-GAN-NS | 4.8 | 26.6 |
| RHYME-GAN | 3.7 | 4.7 |

Table 1: Sampling efficiency: We obtain 10K samples of poetry without additional intervention during decoding, and report the expected samples as inverse of the fraction of samples satisfying valid rhyming patterns for the corresponding dataset. Lower values are better.

we pretrain the LSTM word encoder $g_{\phi_g}(\cdot)$ part of the discriminator, along with a separate LSTM decoder, using an auto-encoding objective on words in the vocabulary.

3 Experiments and Results

3.1 Datasets

We work with the Shakespeare SONNET dataset (Lau et al., 2018) and a new LIMERICK corpus. Each sonnet in the Sonnet dataset is made up of 3 quatrains of 4 lines each, and a couplet. The dataset consists of 2685 sonnets in train, and 335 each in validation and test splits. The quatrains typically have one of the following rhyming structures: AABB, ABAB, ABBA, though some deviations are observed in the dataset. This may be because rhyming patterns are not always strictly followed in writing quatrains, and there are possible inaccuracies in the word pronunciation dictionaries used (e.g. some words can have multiple different pronunciations based on context).

A limerick is a form of verse with five lines. Limericks typically follow a rhyming pattern of AABBA. We collect limericks from an online collection¹. Due to a large vocabulary in the full collection, we filter the dataset to retain only those limericks whose all the words are in a subset of 9K most frequent words. Our final dataset consists of 10, 400 limericks in train and 1300 each in validation and test splits. We train and evaluate the models separately on each corpus.

3.2 Poem Generator

Sampling efficiency We compute the expected number of samples needed before we sample a quatrain which satisfies one of the hand-defined rhyming patterns. Towards this end, we get 10K samples from each model without any constraints (except avoiding *UNK* - unknown tokens). Fol-

¹<http://hardsoft.us>. Accessed May 2019.

lowing prior work (Lau et al., 2018), words are sampled with a temperature value between 0.6 and 0.8. We use CMU dictionary (Weide, 1998) to look up the phonetic representation of a word, and extract the sequence of phonemes from the last stressed syllable onward. Two words are considered to be rhyming if their extracted sequences match (Parrish, 2015). We consider a generated quatrain to have an acceptable pattern if the four ending words follow one of the three rhyming patterns: AABB, ABBA, ABAB. Similarly for LIMERICK, we consider only those samples to be acceptable which have line endings of the rhyming form AABBA.

We consider a baseline **RHYME-LM**, which has the same generator architecture as RHYME-GAN but is trained without the discriminator. We also compare with **RHYME-GAN-NS** which uses a simpler non-structured discriminator. Specifically, it uses a discriminator which first runs a character-level encoder for each ending word - similar to RHYME-GAN - but then instead of computing pairwise similarity matrix, it utilizes a LSTM on the sequence of the computed representations.

As can be observed from Table 1, RHYME-GAN needs fewer samples than other methods to produce an acceptable quatrain or a limerick, indicating that it has learned natural rhyming structures more effectively from data. Note we do not report DEEP-SPEARE on Limerick due to their SONNET specific assumption that for a given end-of-line word there is exactly one more rhyming word among other end-of-line words. Additionally, RHYME-GAN-NS performs worse than RHYME-GAN, and the difference in performance is more prominent in LIMERICK – demonstrating that the proposed structure in the discriminator provided useful inductive bias. Note that compared to 4 line quatrains in SONNET, LIMERICK has 5 line poems and has arguably more complex rhyming pattern constraints.

Likelihood on held out data We report negative log likelihood (NLL) on test splits (Table 2). For SONNET, RHYME-GAN achieves a test set NLL of 3.98. Our model without adversarial learning i.e. RHYME-LM, achieves a test set NLL of 3.97. DEEP-SPEARE reports a test set NLL of 4.38. Note that our language model is hierarchical while DEEP-SPEARE has a linear model. The NLL for RHYME-LM and RHYME-GAN are very similar,

| Model | NLL | |
|-----------------|-------------|-------------|
| | SONNET | LIMERICK |
| DEEP-SPEARE | 4.38 | N/A |
| RHYME-LM | 3.97 | 3.48 |
| RHYME-GAN | 3.98 | 3.49 |

Table 2: Held out negative log likelihood per token for poems in test split.

though RHYME-GAN gets much better sampling efficiency scores than RHYME-LM.

Our generator implementation is largely based on that of Lau et al. (2018). The main difference is that we first generate all the line-ending words and then condition on them to generate the remaining words. The change was made to make it more amenable to our proposed discriminator. However, our hierarchical language model (RHYME-LM) performs worse than DEEP-SPEARE as per sampling efficiency. Therefore, structured discriminator is the driving factor behind the observed improvement with RHYME-GAN. However, committing to the ending words of all lines before completing preceding lines can be a limitation, and addressing it is a possible future direction.

3.3 Analyzing Learned Discriminator

We probe the the word representations $g(\cdot)$ to check if rhyming words are close-by in the learned manifold. We consider all pairs of words among the ending words in a quatrain/limerick, and label each pair to be rhyming or non-rhyming based on previously stated definition of rhyming. If the cosine similarity score between the representations of pairs of words is above a certain threshold, we predict that word pair as rhyming, else it is predicted as non-rhyming. We report F1 scores for the binary classification setup of predicting word-pairs to be rhyming or not. We consider some additional baselines: **RHYM-EM** (Reddy and Knight, 2011) uses latent variables to model rhyming schemes, and train parameters using EM. **GRAPHEME-K** baselines predict a word pair as rhyming only if the last $K = \{1, 2, 3\}$ characters of the two words are same.

For SONNET data, we observe that RHYME-GAN obtains a F1 score of 0.90 (Table 3) on the test split (threshold chosen to maximize f1 on dev split). We repeat the above analysis on the LIMERICK dataset and observe an F1 of 0.92 for RHYME-GAN. DEEP-SPEARE model reports F1

| Model | SONNET | LIMERICK |
|------------------------|-------------|-------------|
| GRAPHEME-1 | 0.71 | 0.79 |
| GRAPHEME-2 | 0.78 | 0.79 |
| GRAPHEME-3 | 0.69 | 0.67 |
| RHYM-EM | 0.71 | 0.73 |
| DEEP-SPEARE/MAX-MARGIN | 0.91 | 0.81 |
| RHYME-GAN-NS | 0.85 | 0.87 |
| RHYME-GAN | 0.90 | 0.92 |

Table 3: Rhyming probe: We use the cosine similarity score of the learned representations to predict a word pair as rhyming or not, and report F1 score for this classification task. RHYM-EM is an unsupervised rhyming pattern discovery method. GRAPHEME-K baselines predict based on exact match of last k characters.

of 0.91 on SONNET. As stated earlier, DEEP-SPEARE’S model is not amenable to LIMERICK - we do compare though with the max-margin classifier in DEEP-SPEARE model trained on LIMERICK which gets F1 score of 0.81. The scores are understandably lower since the AABBA pattern in limericks is not amenable to SONNET specific assumptions made in DEEP-SPEARE model. On the other hand, RHYME-GAN achieves high F1 scores for both the datasets without incorporating any domain specific rhyming pattern information.

RHYME-GAN performs much better than RHYM-EM and GRAPHEME-K baselines. RHYM-EM does not perform well - probably because it operates at word-level and fails to generalize. Note that **RHYME-GAN-NS** gets F1 score of 0.85 in case of SONNET dataset and 0.87 for LIMERICK. These values are lower than corresponding scores for RHYME-GAN, demonstrating that the proposed structure in the discriminator was useful in learning the notion of rhyming.

3.4 Human Evaluations

Following prior work (Lau et al., 2018), we requested human annotators to identify the human-written poem when presented with two samples at a time - a quatrain from the Sonnet corpus and a machine-generated quatrain, and report the annotator accuracy on this task. Note that a lower accuracy value is favorable as it signifies higher quality of machine-generated samples. Using 150 valid samples (i.e. samples belonging to one of the allowed rhyming patterns), we observe 56.0% annotator accuracy for RHYME-GAN, and 53.3% for DEEP-SPEARE – indicating that the post-rejection sampling outputs from the two methods are of comparable quality (the difference in annotator accuracy is not statistically significant as per McNemar’s test under $p < 0.05$). If we use pre-rejection

samples, we observe 60.0% annotator accuracy for RHYME-GAN, and 81.3% for DEEP-SPEARE (the difference being statistically significant as per McNemar’s test under $p < 0.05$) – indicating that unfiltered samples from RHYME-GAN are of higher quality compared to DEEP-SPEARE.

4 Related Work

Early works on poetry generation mostly used rule based methods (Gervás, 2000; Wu et al., 2009; Oliveira, 2017). More recently, neural models for poetry generation have been proposed (Zhang and Lapata, 2014; Ghazvininejad et al., 2016, 2017; Hopkins and Kiela, 2017; Lau et al., 2018; Liu et al., 2019). Yan et al. (2013) retrieve high ranking sentences for a given user query, and repeatedly swap words to satisfy poetry constraints. Ghazvininejad et al. (2018) worked on poetry translation using an unconstrained machine translation model and separately learned Finite State Automata for enforcing rhythm and rhyme. Similar to rhyming and rhythm patterns in poetry, certain types of musical compositions showcase rhythm and repetition patterns, and some prior works model such patterns in music generation (Walder and Kim, 2018; Jhamtani and Berg-Kirkpatrick, 2019). Generative adversarial learning (Goodfellow et al., 2014) for text generation has been used in prior works (Fedus et al., 2018; Wang et al., 2018, 2019; Rao and Daumé III, 2019), though has not been explored with regard to the similarity structure proposed in this paper.

5 Conclusions

In this paper we have proposed a novel structured discriminator to learn a poem generator. The generator learned utilizing the structured adversary is able to identify rhyming structure patterns present in data, as demonstrated through the improved sampling efficiency. Through the rhyming classification probe, we demonstrate that the proposed discriminator is better at learning the notion of rhyming compared to baselines.

Acknowledgements

We are thankful to anonymous EMNLP conference reviewers for providing valuable feedback. We thank Graham Neubig for useful discussions on Limerick generation. This project is funded in part by the NSF under grant 1618044 and by the NEH under grant HAA-256044-17.

References

- Antoine Bosselut, Asli Celikyilmaz, Xiaodong He, Jianfeng Gao, Po-Sen Huang, and Yejin Choi. 2018. Discourse-aware neural rewards for coherent text generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 173–184.
- William Fedus, Ian Goodfellow, and Andrew M Dai. 2018. Maskgan: Better text generation via filling in the... *arXiv preprint arXiv:1801.07736*.
- Pablo Gervás. 2000. Wasp: Evaluation of different strategies for the automatic generation of spanish verse. In *Proceedings of the AISB-00 symposium on creative & cultural aspects of AI*, pages 93–100.
- Marjan Ghazvininejad, Yejin Choi, and Kevin Knight. 2018. Neural poetry translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, volume 2, pages 67–71.
- Marjan Ghazvininejad, Xing Shi, Yejin Choi, and Kevin Knight. 2016. Generating topical poetry. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1183–1191.
- Marjan Ghazvininejad, Xing Shi, Jay Priyadarshi, and Kevin Knight. 2017. Hafez: an interactive poetry generation system. *Proceedings of ACL 2017, System Demonstrations*, pages 43–48.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ari Holtzman, Jan Buys, Maxwell Forbes, Antoine Bosselut, David Golub, and Yejin Choi. 2018. Learning to write with cooperative discriminators. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1638–1649.
- Jack Hopkins and Douwe Kiela. 2017. Automatically generating rhythmic verse with neural networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 168–178.
- Harsh Jhamtani and Taylor Berg-Kirkpatrick. 2019. Modeling self-repetition in music generation using generative adversarial networks. *Machine Learning for Music Discovery Workshop, ICML 2019*.
- Jey Han Lau, Trevor Cohn, Timothy Baldwin, Julian Brooke, and Adam Hammond. 2018. Deep-speare: A joint neural model of poetic language, meter and rhyme. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1948–1958.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Zhiqiang Liu, Zuohui Fu, Jie Cao, Gerard de Melo, Yik-Cheung Tam, Cheng Niu, and Jie Zhou. 2019. Rhetorically controlled encoder-decoder for modern chinese poetry generation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 1992–2001.
- Hugo Gonçalo Oliveira. 2017. A survey on intelligent poetry generation: Languages, features, techniques, reutilisation and evaluation. In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 11–20.
- Allison Parrish. 2015. [Pronouncing: Interface for the cmu pronouncing dictionary](#). [Online; accessed May 2019].
- Sudha Rao and Hal Daumé III. 2019. Answer-based adversarial training for generating clarification questions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 143–155.
- Sravana Reddy and Kevin Knight. 2011. Unsupervised discovery of rhyme schemes. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 77–82.
- Christian Walder and Dongwoo Kim. 2018. Neural dynamic programming for musical self similarity. In *International Conference on Machine Learning*, pages 5092–5100.
- William Yang Wang, Sameer Singh, and Jiwei Li. 2019. Deep adversarial learning for nlp. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials*, pages 1–5.
- Xin Wang, Wenhui Chen, Yuan-Fang Wang, and William Yang Wang. 2018. No metrics are perfect: Adversarial reward learning for visual storytelling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 899–909.
- R Weide. 1998. The CMU pronunciation dictionary, release 0.6. *Carnegie Mellon University*.

- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Xiaofeng Wu, Naoko Tosa, and Ryohei Nakatsu. 2009. New hitch haiku: An interactive renku poem composition supporting tool applied for sightseeing navigation system. In *International Conference on Entertainment Computing*, pages 191–196. Springer.
- Rui Yan, Han Jiang, Mirella Lapata, Shou-De Lin, Xueqiang Lv, and Xiaoming Li. 2013. i, poet: Automatic chinese poetry composition through a generative summarization framework under constrained optimization. In *IJCAI*, pages 2197–2203.
- Xingxing Zhang and Mirella Lapata. 2014. Chinese poetry generation with recurrent neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 670–680.
- Xingxing Zhang and Mirella Lapata. 2017. Sentence simplification with deep reinforcement learning. *arXiv preprint arXiv:1703.10931*.