Gaiku : Generating Haiku with Word Associations Norms

Yael Netzer[*] and **David Gabay** and **Yoav Goldberg**[†] and **Michael Elhadad**
Ben Gurion University of the Negev
Department of Computer Science
POB 653 Be'er Sheva, 84105, Israel
{yaeln,gabayd,yoavg,elhadad}@cs.bgu.ac.il

Abstract

creativeness / a pleasing field / of bloom

Word associations are an important element of linguistic creativity. Traditional lexical knowledge bases such as WordNet formalize a limited set of systematic relations among words, such as synonymy, polysemy and hypernymy. Such relations maintain their systematicity when composed into lexical chains. We claim that such relations cannot explain the type of lexical associations common in poetic text. We explore in this paper the usage of Word Association Norms (WANs) as an alternative lexical knowledge source to analyze linguistic computational creativity. We specifically investigate the Haiku poetic genre, which is characterized by heavy reliance on lexical associations. We first compare the density of WAN-based word associations in a corpus of English Haiku poems to that of WordNet-based associations as well as in other non-poetic genres. These experiments confirm our hypothesis that the non-systematic lexical associations captured in WANs play an important role in poetic text. We then present Gaiku, a system to automatically generate Haikus from a seed word and using WAN-associations. Human evaluation indicate that generated Haikus are of lesser quality than human Haikus, but a high proportion of generated Haikus can confuse human readers, and a few of them trigger intriguing reactions.

[*] Supported by Deutsche Telekom Laboratories at Ben-Gurion University of the Negev.

[†] Supported by the Lynn and William Frankel Center for Computer Sciences.

1 Introduction

Traditional lexical knowledge bases such as WordNet formalize a limited set of systematic relations that exist between words, such as synonymy, polysemy, hypernymy. When such relations are composed, they maintain their systematicity, and do not create surprising, unexpected word associations.

The human mind is not limited to such systematic relations, and people tend to associate words to each other with a rich set of relations, such as non systematic paradigmatic (*doctor-nurse*) and syntagmatic relations (*mash-potato*) as identified by Saussure (1949). Such associations rely on cultural (*mash-television*), emotional (*math - yuck*) and personal experience (*autumn - Canada*).

In linguistic creativity, such as prose or poetry writing, word associations play an important role and the ability to connect words into new, unexpected relations is one of the key mechanisms that triggers the reader involvement.

We explore in this paper the usage of Word Association Norms (WANs) as an alternative lexical knowledge source to analyze linguistic computational creativity. WANs have been developed in psychological research in the past 40 years. They record typical word associations evoked by people when they are submitted a trigger word. Such associations (*e.g.*, *table* to *chair* or *cloth*) are non-systematic, yet highly stable across people, time (over a period of 30 years) and languages. WANs have been compiled in various languages, and provide an interesting source to analyze word associations in creative writing.

We specifically investigate the Haiku poetic

Proceedings of the NAACL HLT Workshop on Computational Approaches to Linguistic Creativity, pages 32–39,
Boulder, Colorado, June 2009.

genre, which is characterized by heavy reliance on lexical associations. The hypothesis we investigate is that WANs play a role in computational creativity, and better explain the type of word associations observed in creative writing than the systematic relations found in thesauri such as WordNet.

In the rest of the paper, we refine our hypothesis and present observations on a dataset of English Haikus we collected. We find that the density of WAN-based word associations in Haikus is much higher than in other genres, and also much higher than the density of WordNet-based associations. We then present Gaiku, a system we developed to automatically generate Haikus from a seed word using word association norms. Evaluation we performed with a group of 60 human readers indicates that the generated Haikus exhibit interesting creative characteristics and sometimes receive intriguing acclaim.

2 Background and Previous Work

2.1 Computational Creativity

Computational creativity in general and linguistic in particular, is a fascinating task. On the one hand, linguistic creativity goes beyond the general NLP tasks and requires understanding and modelling knowledge which, almost by definition, cannot be formalized (*i.e.*, terms like *beautiful, touching, funny* or *intriguing*). On the other hand, this vagueness itself may enable a less restrictive formalization and allow a variety of quality judgments. Such vague formalizations are naturally more useful when a computational creativity system does not attempt to model the creativity process itself, but instead focuses on 'creative products' such as poetry (see Section 2.3), prose and narrative (Montfort, 2006), cryptic crossword clues (Hardcastle, 2007) and many others. Some research focus on the creative process itself (see (Ritchie, 2006) for a comprehensive review of the field). We discuss in this paper what Boden (1998) calls *P-Creativity* (Psychological Creativity) which is defined relative to the initial state of knowledge, and H-Creativity (Historical Creativity) which is relative to a specific reference culture. Boden claims that, while hard to reproduce, *exploratory creativity* is most successful in computer models of creativity. This is because the other kinds of creativity are *even more elusive* due to *the difficulty of approaching the richness of human associative memory, and the difficulty of identifying our values and of expressing them in computational form.*

We investigate in our work one way of addressing this difficulty: we propose to use associative data as a knowledge source as a first approximation of human associative capabilities. While we do not explain such associations, we attempt to use them in a constructive manner as part of a simple combinational model of creativity in poetry.

2.2 Word Associations and Creativity

Associations and creativity are long known to be strongly connected. Mendick (Mendick, 1969) defines creative thinking as "*the forming of associative elements into new combinations which either meet specified requirements or are in some way useful*." The *usefulness* criterion distinguishes *original thinking* from *creative thinking*. A creative solution is reached through three main paths: *serendipity* (random stimuli evoke associative elements), *similarity* (stimuli and solution are found similar through an association) and *mediation* (both "problem" and "solution" can be associated to similar elements). In our work, we hypothesize that interesting Haiku poems exhibit creative word associations. We rely on this hypothesis to first generate candidate word associations starting from a seed word and following random walks through WANs, but also to rank candidate Haiku poems by measuring the density of WAN-based associations they exhibit.

2.3 Poetry Generation

Although several automatic and semi-automatic poetry generation systems were developed over the years, most of them did not rise above the level of "party tricks" (Manurung et al., 2000). In his thesis, (Manurung, 2003), defined a poem to be a text that meets three properties: meaningfulness, grammaticality and poeticness. Two of the few systems that attempt to explicitly represent all three properties are reported in (Gervas, 2001) and (Díaz-Agudo et al., 2002). Both systems take as input a prose message provided by the user, and translate it into formal Spanish poetry. The system proposed in (Manurung et al., 2000) is similar in that it focuses on the syntactic and phonetic patterns of the poem, putting less stress on the semantics. The sys-

tem starts with a simple seed and gradually develops a poem, by making small syntactic and semantic changes at every step.

Specifically in the subfield of Haiku generation, the Haiku generator presented in (Wong and Chun, 2008) produces candidate poems by combining lines taken from blogs. The system then ranks the candidates according to semantic similarity, which is computed using the results returned by a search engine when querying for words in each line. Hitch-Haiku (Tosa et al., 2008), another Haiku generation system, starts from two seed words given by the user. It retrieves two phrases containing these words from a corpus, and then adds a third phrase that connects both input words, using lexical resources.

In our work, we induce a statistical language model of the structure of Haikus from an analysis of a corpus of English Haikus, and explore ways to combine chains of lexical associations into the expected Haiku syntactic structure. The key issues we investigate are the importance of WAN-based associations in the Haiku generation process, and how a chain of words, linked through WAN-based associations, can be composed into a Haiku-like structure.

2.4 Haiku

Haiku is a form of poetry originated in Japan in the sixteenth century. The genre was adopted in Western languages in the 20^{th} Century. The original form of a poem is of three lines of five, seven and five syllables (although this constraint is loosened in non-Japanese versions of Haiku (Gilbert and Yoneoka, 2000)). Haiku, by its nature, aims to reflect or evoke emotion using an extremely economical linguistic form; most Haiku use present tense and use no judgmental words; in addition, functional or syntactic words may be dropped. Traditional Haiku involve reference to nature and seasons, but modern and western Haiku are not restricted to this theme[1].

We adopt the less "constraining" definition of the author Jack Kerouac (2004) for a Haiku "I propose that the "Western Haiku" simply say a lot in three short lines in any Western language. Above all, a Haiku must be very simple and free of all poetic trickery and make a little picture and yet be as airy and graceful as a Vivaldi Pastorella." (pp. *x-xi*). In addition, we are guided by the saying " *The best haiku should leave the reader wondering* " (Quoted in (Blasko and Merski, 1998))

[1] *Senryu* poetry, similar in form to Haiku, is the Japanese genre of poems that relate to human and relationships, and may be humorous. Hereafter, we use Haiku for both the original definition and the Senryu as well.

2.5 Word Association Norms

The interest in word associations is common to many fields. Idiosyncrasy of associations was used as a diagnostic tool at the beginning of the 20^{th} century, but nowadays the majority of approaches deal less with particular associations and more with general patterns in order to study the structure of the mental lexicon and of semantic memory (Rubinsten et al., 2005).

Word Association Norms (WAN) are a collection of cue words and the set of free associations that were given as responses to the cue, accompanied with quantitative and statistical measures. Subjects are given a word and asked to respond immediately with the first word that comes to their mind. The largest WAN we know for English is the University of South Florida Free Association Norms (Nelson et al., 1998).

Word Association Norms and Thesauri in NLP Sinopalnikova and Smrz (2004) have shown that when building and extending semantic networks, WANs have advantages over corpus-based methods. They found that WANs cover semantic relations that are difficult to acquire from a corpus: 42% of the non-idiosyncratic cue-target pairs in an English WAN never co-appeared in a 10 words window in a large balanced text corpus. From the point of view of computational creativity, this is encouraging, since it suggests that association-based content generation can lead to texts that are both sensible and novel. (Duch and Pilichowski, 2007)'s work, from a neuro-cognitive perspective, generates neologisms based, among other data, on word association. (Duch and Pilichowski, 2007) sums "creativity requires prior knowledge, imagination and filtering of the results."

3 WordNet vs. Associations

Word association norms add an insight on language that is not found in WordNet or are hard to acquire from corpora, and therefore can be used as an additional tool in NLP applications and computational

creativity.

We choose the Haiku generation task using word associations, since this genre of poetry encapsulates meaning in a special way. Haiku tend to use words which are connected through associative or phonological connections (very often ambiguous).

We hypothesize that word-associations are good catalyzers for creativity, and use them as a building block in the creative process of Haiku generation. We first test this hypothesis by analyzing a corpus of existing Haiku poems.

3.1 Analyzing existing text

Can the creativity of text as reflected in word associations be quantified? Are Haiku poems indeed more associative than newswire text or prose? If this is the case, we expect Haiku to have more associative relations, which cannot be easily recovered by WordNet than other type of text. We view the WAN as an undirected graph in which the nodes are stemmed words, and two nodes are connected *iff* one of them is a cue for the other. We take the *associative distance* between two words to be the number of edges in the shortest path between the words in the associations-graph. Interestingly, almost any word pair in the association graph is connected with a path of at most 3 edges. Thus, we take two words to be associatively related if their associative distance is 1 or 2. Similarly, we define the *WordNet distance* between two stemmed words to be the number of edges in the shortest path between any synset of one word to any synset of the other word[2]. Two words are WordNet-related if their WordNet distance is less than 4 (this is consistent with works on lexical-cohesion, (Morris and Hirst, 1991)).

We take the *associativity* of a piece of text to be the number of associated word pairs in the text, normalized by the number of word pairs in the text of which both words are in the WAN.[3] We take the *WordNet-relations level* of a piece of text to be the number of WordNet-related word pairs in the text.

[2]This is the inverse of the path-similarity measure of (Pedersen et al., 2004).

[3]This normalization is performed to account for the limited lexical coverage of the WAN. We don't want words that appear in a text, but are not covered by the WAN, to affect the associativity level of the text.

SOURCE	AVG. ASSOC RELATIONS (<3)	AVG. WORDNET RELATIONS (<4)
News	0.26	2.02
Prose	0.22	1.4
Haiku	0.32	1.38

Table 1: Associative and WordNet relations in various text genres

We measure the average associativity and WordNet levels of 200 of the Haiku in our Haiku Corpus (Section 4.1), as well as of random 12-word sequences from Project Gutenberg and from the NANC newswire corpus.

The results are presented in Table 1.

Perhaps surprisingly, the numbers for the Gutenberg texts are lower on all measures. This is attributed to the fact that Gutenberg texts have many more pronouns and non-content words than the Haiku and newswire text. Haiku text appears to be more associative than newswire text. Moreover, newswire documents have many more WordNet-relations than the Haiku poems – whenever words are related in Haiku, this relatedness tends to be captured via the association network rather than via the WordNet relations. The same trend is apparent also when considering the Gutenberg numbers: they have about 15% less associations than newswire text, but about 30% less WordNet-relations. This supports the claim that associative information which is not readily available in WordNet is a good indicator of creative content.

3.2 Generating creative content

We now investigate how word-associations can help in the process of generating Haikus. We define a 5 stage generative process: **theme selection** in which the general theme of the Haiku is decided, **syntactic planning**, which sets the Haiku form and syntactic constraints, **content selection / semantic planning** which combines syntactic and aesthetic constraints with the theme selected in the previous stages to form good building blocks, **filtered overgeneration** of many Haiku based on these selected building blocks, and finally **re-ranking** of the generated Haiku based on external criteria.

The details of the generation algorithm are presented in Section 4.2. Here we focus on the creative aspect of this process – theme selection. Our main claim is that WANs are a good source for interest-

ing themes. Specifically, interesting themes can be obtained by performing a short *random walk* on the association graph induced by the WAN network.

Table 2 presents the results of several random walks of 3 steps starting from the seed words "Dog", "Winter", "Nature" and "Obsession". For comparison, we also present the results of random walks over WordNet glosses for the same seeds.

We observe that the association network is better for our needs than WordNet. Random walks in WordNet are more likely to stay too close to the seed word, limiting the poetic options, or to get too far and produce almost random connections.

4 Algorithm for generating Haiku

4.1 Dataset

We used the Word Association Norms (WAN) of the University of South Florida [4] (Nelson et al., 1998) for discovering associations of words. The dataset (Appendix A, there) includes 5,019 cue words and 10,469 additional target that were collected with more than 6,000 participants since 1973.

We have compiled a Haiku Corpus, which includes approximately 3,577 Haiku in English of various sources (amateurish sites, children's writings, translations of classic Japanese Haiku of Bashu and others, and 'official' sites of Haiku Associations (*e.g., Haiku Path - Haiku Society of America*).

For the content selection part of the algorithms, we experimented with two data sources: a corpus of 1TB web-based N-grams supplied by Google, and the complete text of Project Gutenberg. The Gutenberg data has the advantage of being easier to POS-tag and contains less restricted-content, while the Google Web data is somewhat more diverse.

4.2 Algorithm Details

Our Haiku generation algorithm includes 5 stages: theme selection, syntactic planning, content selection, filtered over generation, and ranking.

The **Theme Selection** stage is in charge of dictating the overall theme of our Haiku. We start with a user-supplied seed word (*e.g.* `WINTER`). We then consult the Association database in order to enrich the seed word with various associations. Ideally, we would like these associations to be close enough to the seed word to be understandable, yet far enough away from it as to be interesting. After some experimenting, we came up with the following heuristic, which we found to provide adequate results. We start with the seed word, and conduct a short random walk on the associations graph. Each random step is comprised of choosing a random direction (either "Cue" or "Target") using a uniform distribution, and then a random neighbor according to its relative frequency. We conduct several (8) such walks, each with 3 steps, and keep all the resulting words. This gives us mostly close, probable associations, as well as some less probable, further away from the seed.

The **syntactic planning** stage determines the form of the generated Haiku, setting syntactic and aesthetic constraints for the generative process. This is done in a data-driven way by considering common line patterns from our Haiku corpus. In a training stage, we POS-tagged each of the Haiku, and then extracted a pattern from each of the Haiku lines. A line-pattern is a sequence of POS-tags, in which the most common words are lexicalized to include the word-form in addition to the POS-tag. An example for such a line pattern might be `DT_the JJ NN`. We kept the top-40 frequent patterns for each of the Haiku lines, overall 120 patterns. When generating a new Haiku, we choose a random pattern for the first line, then choose the second line pattern conditioned on the first, and the third line pattern conditioned on the second. The line patterns are chosen with a probability proportional to their relative frequencies in the training corpus. For the second and third lines we use the conditional probabilities of a pattern appearing after the previous line pattern. The result of this stage is a 3-line Haiku skeleton, dictating the number of words on each line, their POS-tags, and the placement of specific function words.

In the **Content Selection** stage, we look for possible Haiku lines, based on our selected theme and syntactic structure. We go over our candidate lines[5], and extract lines which match the syntactic patterns and contain a stemmed appearance of one of the stemmed theme words. In our current implementation, we require the first line to contain the seed word, and the second and third line to contain any of

[4] http://w3.usf.edu/FreeAssociation/

[5] These are POS-tagged n-grams extracted from a large text corpora: the Google T1 dataset or Project Gutenberg

SEED	WAN	WORDNET
Dog	puppy adorable cute	heel villain villainess
Dog	cat curious george	hound scoundrel villainess
Winter	summer heat microwave	wintertime solstice equinox
Winter	chill cold alergy	midwinter wintertime season
Nature	animals instinct animals	world body crotch
Nature	natural environment surrounding	complexion archaism octoroon
Obsession	cologne perfume smell	fixation preoccupation thought
Obsession	compulsion feeling symptom	compulsion onomatomania compulsion

Table 2: Some random walks on the WordNet and WAN induced graphs

the theme words. Other variations, such as choosing a different word set for each line, are of course possible.

The **over generation** stage involves creating many possible Haiku candidates by randomly matching lines collected in the content selection stage. We filter away Haiku candidates which have an undesired properties, such as repeating the same content-word in two different lines.

All of the generated Haiku obey the syntactic and semantic constraints, but not all of them are interesting. Thus, we **rank** the Haiku in order to weed out the better ones. The top-ranking Haiku is the output of our system. Our current heuristic prefers highly associative Haikus. This is done by counting the number of 1st and 2nd degree associations in each Haiku, while giving more weight to 2nd degree associations in order to encourage "surprises". While all the candidate Haiku were generated based on a common theme of *intended* associative connections, the content selection and adherence to syntactic constraints introduce additional content words and with them some new, *unintended* associative connections. Our re-ranking approach tries to maximize the number of such connections.[6]

5 Evaluation

The ultimate goal of a poetry generation system is to produce poems that will be considered good if written by a human poet. It is difficult to evaluate to what extent a poetry generation system can meet this goal (Ritchie, 2001; Manurung et al., 2000). Difficulties arise from two major sources: first, since a creative work should be novel, it cannot be directly evaluated by comparison to some gold standard. Second, it is hard for people to objectively evaluate the quality of poetry. Even determining whether a text is a poem or not is not an easy task, as readers expect poetry to require creative reading, and tolerate, to some extent, ungrammatical structures or cryptic meaning.

5.1 "Turing Test" Experiment

To evaluate the quality of Gaiku, we asked a group of volunteers to read a set of Haiku, indicate how much they liked each one (on a scale of 1-5), and classify each Haiku as written by a human or by a computer.

We compiled two sets of Haiku. The first set (AUTO) contained 25 Haiku. 10 Haiku chosen at random from our Haiku corpus, and 15 computer generated ones. The computer generated Haiku were created by identifying the main word in the first line of each human-written Haiku, and passing it as a seed word to the Haiku generation algorithm (in case a first line in human-written Haiku contained two main words, two Haiku were generated). We included the top-ranking Haiku returning from a single run of the system for each seed word. The only human judgement in compiling this set was in the identification of the main words of the human Haiku.

The second set (SEL) was compiled of 9 haiku poems that won awards[7], and 17 computer Haiku that were selected by us, after several runs of the automatic process. (Again, each poem in the automatic poems set shared at least one word with some poem in the human Haiku set).

The subjects were not given any information about the number of computer-generated poems in the sets.

[6] While this heuristic works well, it leaves a lot to be desired. It considers only the quantity of the associations, and not their quality. Indeed, when looking at the Haiku candidates produced in the generation stage, one can find many interesting pieces, where some of the lower ranking ones are far better than the top ranking.

[7] Gerald Brady Memorial Award Collection http://www.hsa-haiku.org/bradyawards/brady.htm 2006-2007

The AUTO questionnaire was answered by 40 subjects and the SEL one by 22. (Altogether, 52 different people took part in the experiment, as some subjects answered both versions). The subjects were all adults (age 18 to 74), some were native English speakers and others were fully fluent in English. Except a few, they did not have academic background in literature.

5.2 Results and Discussion

Results are presented in Table 3 and Figure 1.

Overall, subjects were correct in 66.7% of their judgements in AUTO and 61.4% in SEL. The average grade that a poem - human or machine-made - received correlates with the percentage of subjects who classified it as human. The average grade and rate of acceptance as written by human were significantly higher for the Haiku written by people. However, some computer Haiku rivaled the average human poem in both measures. This is true even for AUTO, in which both the generation and the selection processes were completely automatic. The best computer Haiku of SEL scored better than most human Haiku in both measures.

The best computer poem in SEL was:

early dew / the water contains / teaspoons of honey

which got an average grade of 3.09 and was classified as human by 77.2% of the subjects.

At the other extreme, the computer poem (SEL):

space journey / musical instruments mythology / of similar drugs

was classified as human by only 9% of the subjects, and got an average grade of 2.04.

The best Haiku in the AUTO set was:

cherry tree / poisonous flowers lie / blooming

which was classified as human by 72.2% of the subjects and got an average grade of 2.75.

The second human-like computer generated Haiku in each set were:

spring bloom / showing / the sun's pyre

(AUTO, 63.8% human) and:

blind snakes / on the wet grass / tombstoned terror

(SEL, 77.2% human).

There were, expectedly, lots of disagreements. Poetry reading and evaluation is subjective and by itself (in particular for Haiku) a creative task. In addition, people have very different ideas in mind as to a computer's ability to do things. (One subject said, for example, that the computer generated

holy cow / a carton of milk / seeking a church

is *too stupid* to be written by a computer; however, content is very strongly connected and does not seem random). On the other end, subjects often remarked that some of the human-authored Haiku contained metaphors which were *too obvious* to be written by a human.

		Human Poems	**Gaiku**
AUTO	avg. % classified as Human	72.5%	37.2%
	avg. grade	2.86	2.11
SEL	avg. % classified as Human	71.7%	44.1%
	avg. grade	2.84	2.32

Table 3: Turing-test experiment results

Every subject was wrong at least 3 times (at least once in every direction); every poem was wrongly-classified at least once. Some really bad auto-poems got a good grade here and there, while even the most popular human poems got a low grade sometimes.

6 Discussion and Future Work

Word association norms were shown to be a useful tool for a computational creativity task, aiding in the creation of an automatic Haiku-generation software, which is able to produce "human-like" Haiku. However, associations can be used for many other tasks.

In the last decade, *lexical chains* are often used in various NLP tasks such as text summarization or text categorization; WordNet is the main resource for detecting the cohesive relationships between words and their relevance to a given chain (Morris and Hirst, 1991). We believe that using word association norms can enrich the information found in WordNet and enable the detection of more relevant words.

Another possible application is for assisting word-finding problem of children with specific language impairments (SLI). A useful tactic practiced as an assistance to retrieve a forgotten word is by saying all words that come to mind. The NLP task, therefore, is for a set of a given associations, reconstruct the targeted word.

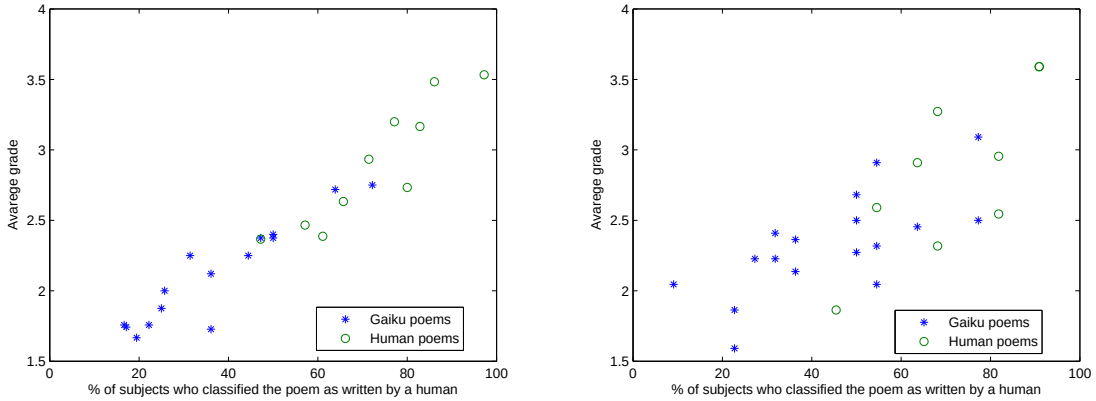

Figure 1: Average grades and percentages of subjects who classified poems as written by humans, for AUTO (left) and SEL. Circles represent Haiku written by people, and stars represent machine-made Haiku

References

D.G. Blasko and D.W. Merski. 1998. Haiku poetry and metaphorical thought: An invention to interdisciplinary study. *Creativity Research Journal*, 11.

M.A. Boden. 1998. Creativity and artificial intelligence. *Artificial Intelligence*, 103(1–2).

F. de Saussure, C. Bally, A. Riedlinger, and A. Sechehaye. 1949. *Cours de linguistique generale*. Payot, Paris.

B. Díaz-Agudo, P. Gervás, and P. A. González-Calero. 2002. Poetry generation in COLIBRI. In *Proc. of EC-CBR*.

W. Duch and M. Pilichowski. 2007. Experiments with computational creativity. *Neural Information Processing, Letters and Reviews*, 11(3).

P. Gervas. 2001. An expert system for the composition of formal Spanish poetry. *Journal of Knowledge-Based Systems*, 14.

R. Gilbert and J. Yoneoka. 2000. From 5-7-5 to 8-8-8: An investigation of Japanese Haiku metrics and implications for English Haiku. *Language Issues: Journal of the Foreign Language Education Center*.

D. Hardcastle. 2007. Cryptic crossword clues: Generating text with a hidden meaning BBKCS-07-04. Technical report, Birkbeck College, London.

J. Kerouac. 2004. *Book of Haikus*. Enitharmon Press.

H.M. Manurung, G. Ritchie, and H. Thompson. 2000. Towards a computational model of poetry generation. In *Proc. of the AISB'00*.

H.M. Manurung. 2003. *An evolutionary algorithm approach to poetry generation*. Ph.D. thesis, University of Edinburgh.

S.A. Mendick. 1969. The associative basis of the creative process. *Psychological Review*.

N. Montfort. 2006. Natural language generation and narrative variation in interactive fiction. In *Proc. of Computational Aesthetics Workshop at AAAI 2006*, Boston.

J. Morris and G. Hirst. 1991. Lexical cohesion computed by thesaural relations as an indicator of the structure of text. *Computational Linguistics*, 17.

D.L. Nelson, C.L. Mcevoy, and T.A. Schreiber. 1998. The University of South Florida Word Association, Rhyme, and Word Fragment Norms. http://www.usf.edu/FreeAssociation/.

T. Pedersen, S. Patwardhan, and J. Michelizzi. 2004. Wordnet::similarity - measuring the relatedness of concepts. In *HLT-NAACL 2004: Demonstrations*.

G. Ritchie. 2001. Assessing creativity. In *Proc. of AISB'01 Symposium*.

G. Ritchie. 2006. The transformational creativity hypothesis. *New Generation Computing*, 24.

O. Rubinsten, D. Anaki, A. Henik, S. Drori, and Y. Faran. 2005. Free association norms in the Hebrew language. *Word Norms in Hebrew*. (In Hebrew).

A. Sinopalnikova and P. Smrz. 2004. Word association thesaurus as a resource for extending semantic networks. In *Communications in Computing*.

N. Tosa, H. Obara, and M. Minoh. 2008. Hitch haiku: An interactive supporting system for composing haiku poem. In *Proc. of the 7th International Conference on Entertainment Computing*.

M. Tsan Wong and A. Hon Wai Chun. 2008. Automatic Haiku generation using vsm. In *Proc. of ACACOS'08*, April.