

Automatic Generation of Poetry Inspired by Twitter Trends

Hugo Gonalo Oliveira^(✉)

CISUC, Department of Informatics Engineering, University of Coimbra,
Coimbra, Portugal
hroliv@dei.uc.pt

Abstract. This paper revisits PoeTryMe, a poetry generation platform, and presents its most recent instantiation for producing poetry inspired by trends in the Twitter social network. The presented system searches for tweets that mention a given topic, extracts the most frequent words in those tweets, and uses them as seeds for the generation of new poems. The set of seeds might still be expanded with semantically-relevant words. Generation is performed by the classic PoeTryMe system, based on a semantic network and a grammar, with a previously used generate & test strategy. Illustrative results are presented using different seed expansion settings. They show that the produced poems use semantically-coherent lines with words that, at the time of generation, were associated with the topic. Resulting poems are not really about the topic, but they are a way of expressing, poetically, what the system knows about the semantic domain set by the topic.

Keywords: Computational creativity · Creative systems · Poetry generation · Social media

1 Introduction

Creative systems are computer programs that exhibit behaviours that would be deemed as creative by unbiased observers [1]. Such behaviours are often rendered in the form of artefacts that go from visual art [2] to linguistic creativity including, but not limited to verbally-expressed humor [3], narratives [4], metaphors [5], neologisms [6], slogans [7] or poetry, one of the most popular tasks in this subfield. Poetry generation is a kind of natural language generation where the resulting text can be seen as a poem. This can be achieved by the presence of features, such as a regular metre, rhymes, or a figurative language.

PoeTryMe [8, 9] is a poetry generation platform that produces lines with the help of a grammar and a set of relation instances, and combines them according to a pre-defined strategy, towards the creation of a poem. PoeTryMe has a versatile architecture that provides a high level of customisation and can be the starting point for the development of different poetry generation systems. Several of its components can be changed: its semantic knowledge, the line templates, the

generation strategies and, of course, the poem configuration. The combination of all these components enables the generation of diverse poems, thus contributing to a positive perception of creativity.

This paper presents a new instantiation of PoeTryMe where the generation of a poem is inspired by information circulating in the Twitter¹ social network. The process starts with a given topic, which is used to retrieve associated words from Twitter, then used as seeds for poetry generation. Resulting poems are not clearly about the topic, but they are, at least, inspired by it, and an abstract connection is usually present. A bot was developed to publish the creations of the presented system in Twitter. Therefore, we see Twitter’s role in this process also as a way of continuously retrieving different seeds, and thus contributing to the generation of more diverse poems every time.

In the remaining of the paper, related work, mostly on poetry generation, is first reviewed. Then, a short description of PoeTryMe is provided. Before concluding and discussing cues for future work, the specificities of this instantiation are presented, together with some examples and a critical view.

2 Related Work

Computational Creativity is a multidisciplinary endeavour at the intersection of the fields of artificial intelligence, cognitive psychology, philosophy, and the arts². It is driven towards modelling, simulating or replicating creativity, using a computer, either to: (i) construct programs capable of human-level creativity; (ii) better understand human creativity and formulate an algorithmic perspective on creative behaviour in humans; (iii) design programs that can enhance human creativity without necessarily being creative themselves.

Poetry generation systems are artificial systems that produce text with poetic features and has thus creative value. The automatic generation of poetry is a complex task, as it involves several levels of language (e.g. phonetics, lexical choice, syntax and semantics) and usually demands a considerable amount of input knowledge. However, not all of those levels have to be strictly addressed. On the one hand, poetic text does not have to be extremely precise [10], as several rules, typically present in the production of natural language, need to (or should) be broken [11]. On the other hand, poetry involves a high occurrence of interdependent linguistic phenomena where rhythm, metre, rhyme and other features like alliteration, sentiment, or figurative language play an important role. For instance, it is sometimes enough to have a less clear message, in a trade-off for a pleasant sound given by a highly regular metre and rhymes.

¹ <https://twitter.com>.

² Check the website of the Association for Computational Creativity at <http://computationalcreativity.net/>.

Several poetry generation systems are based on poem or line templates, but most of them go further and combine the previous with other techniques (e.g. [12, 13]). Templates learned from human-created poetry are often sequences of words with gaps to be filled by the system, but they can also be sequences of parts-of-speech [14].

Produced word sequences usually evolve to meet the desired constraints, which often include a stress pattern (metre) and may additionally define the position of rhymes, syntactic rules, semantic constraints, and other features like the presence of alliteration or the use of figurative language. Evolution can be made through a generate-and-test approach [10, 15] or it can rely on evolutionary algorithms [11, 16]. Other approaches include case-based reasoning [17], probabilistic language models [18], constraint programming [13], or multi-agent systems [19].

Besides exhibiting poetic features, produced text should obey linguistic conventions and convey a conceptual message, meaningful under some interpretation [11]. The handling of linguistic rules is typically achieved with the help of natural language processing tools, such as morphological lexicons or grammars. On the other hand, meaningfulness is more subjective and difficult to achieve. Towards this goal, different systems have handled semantics differently. Some start generation from a textual document [12, 20, 21] or a set of seed words [22–24] to constrain the space of possible generations, in a way that the poem should use these exact words, or others semantically associated. The choice of relevant words may be achieved either by exploring models of semantic similarity, extracted from corpora [13, 22, 24], with the help of lexical-semantic knowledge bases [14, 23], or both [12].

Poetry generation has been mainly addressed for English, but there are attempts in other languages, including Spanish [10, 17], Basque [14], Finnish [13, 20], Chinese [24], Indonesian [25], or Bengali [26], among others. Our original effort targeted Portuguese [8], which is also the target language of the present work.

3 PoeTryMe

PoeTryMe [8, 9] is a poetry generation platform, on the top of which different strategies for poetry generation can be implemented. It relies on a modular architecture (see Fig. 1), which enables the independent development of each module and provides a high level of customisation, depending on the needs of the system and ideas of the user or developer. Among other parameters, users may define the semantic network to use, the rules of the generation grammar, the transmitted sentiment, the generation strategy and the structure of the poem. Developers may reimplement some of the modules and reuse the others.

A Generation Strategy organises lines, such that they suit, as much as possible, the structure of a poetic form and exhibit certain features. A structure file sets the number of stanzas, lines per stanza and of syllables in each line of the poem. An instantiation of the Generation Strategy does not generate the lines, but exploits the Sentence Generator module to retrieve natural language

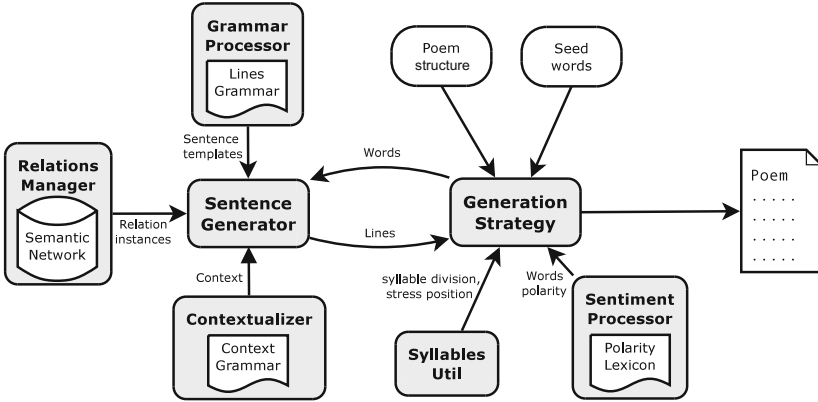


Fig. 1. The architecture of PoeTryMe.

fragments, which might be used as such. Each strategy may differ on the number of fragments requested from the Sentence Generator at any time, and how they are organised into the poem structure, considering for this purpose features like metre, rhyme, coherence between lines or others, depending on the desired goal.

Syllable-related features are evaluated with the help of the Syllable Utils. Given a word, this module may be used to divide it into syllables, to find its stress, or to extract its termination, useful to identify rhymes.

The Sentence Generator is a core module for PoeTryMe. It generates semantically-coherent natural language fragments, with the help of: (i) a semantic network, managed by the Relations Manager, that connects words according to relation predicates; and a generation grammar, processed by the Grammar Processor, with textual renderings for the generation of lines that express semantic relations. The generation of a line is a three-step interaction:

1. A random relation instance, in the form of a *triplet* = $\{word_1, predicate, word_2\}$, is retrieved from the semantic network. To narrow the space of possible generations, a set of seed words can be provided to the Relations Manager. This set defines the generation domain, represented by a subgraph of the main network that will contain all the triplets involving seed words. A surprise factor, ν , sets the probability of selecting also triplets involving nodes that are two levels far from the seeds.
2. A random rendering for the *triplet's* predicate is retrieved from the grammar. Grammar rules are natural language renderings of predefined semantic relations. So, there must be a direct mapping between the relation names, in the graph, and the rules' name, in the grammar. Besides terminal tokens, that will be present in the poem without change, rules have placeholders that indicate the position of the relation arguments (`< arg1 >` and `< arg2 >`). A simple example of a valid rule set, with three hypernymy patterns, is shown in Fig. 2.

3. The resulting fragment is returned after inserting the arguments of the *triplet* in the proper placeholders of the rule. For instance, the rules displayed in Fig. 2 could be used to generate the following fragments: *a tool like a hammer*, *mango is a delicious fruit*, *man before animal*.

```

HYPERNYM-OF ::= a <arg1> like a <arg2>
HYPERNYM-OF ::= <arg2> is a delicious <arg1>
HYPERNYM-OF ::= <arg2> before <arg1>

```

Fig. 2. Grammar example rule set.

In addition to the previous modules, the Contextualizer explains why certain words were selected and what is their connection to the seed words, as a list of triplets for each line. It can be used for debugging or evaluation purposes.

Besides the surprise factor, another way of increasing diversity is to expand the set of seeds with semantically-relevant words. For this purpose, before generation, a personalized version of the PageRank [27] algorithm is run in the full semantic network. Initial node weights are randomly distributed across the seeds, while the rest of the nodes have an initial weight of 0. After 30 iterations, nodes will be ranked according to their structural relevance to the seeds. The top- r ranked nodes are added to seed set.

The previous expansion feature can be biased to induce a target sentiment in the poem. For this purpose, the top- r nodes are previously filtered, in order to use only those with a target polarity. The typical polarity of words is obtained from the Sentiment Processor, an interface to a polarity lexicon that lists words and their typical polarities (positive, neutral or negative). For instance, suppose that the top-10 ranked words for the word “blue” are: *grim*, *blueness*, *gloomy*, *sexy*, *color*, *dark*, *dejected*, *low*, *dye*, *down*. When generating a negative poem with, say, the top-3 words, *grim*, *gloomy* and *dark* would be added to the seed set, together with *blue*. For a positive poem, the word *sexy* would be added, together with the next two positive words in the ranking.

A more detailed description of PoeTryMe’s architecture is available elsewhere [9]. Although PoeTryMe was originally developed to produce poetry in Portuguese, its flexible architecture enabled the adaptation to Spanish [28] and English. It has also been used to produce song lyrics for a given melody [29].

4 Poetry Inspired by Current Trends

This section describes a new instantiation of PoeTryMe where seed words are collected from Twitter. Feeding the system with words that, at a certain time,

are associated with a topic, enables the generation of different poems every time, with a shallow connection with present events, even if it is not immediately clear.

As the original PoeTryMe, the presented instantiation targets Portuguese although, given our recent adaptations, it could be adapted to Spanish and English with low effort. All the linguistic resources used were the same as those of previous instantiations for Portuguese [9], except for the generation grammar, which is described in the next section. After that, the current approach for producing poetry inspired by Twitter is described; a Twitterbot that publishes poems about trendy topics is introduced; the current setup of this system is detailed; and some illustrative examples are presented, followed by a critical view of our results.

4.1 Generation Grammar

The generation grammar used in this instantiation of PoeTryMe has two main updates: it is more strict and covers different kinds of text. The rules of the grammar are still automatically acquired from human-created poetry, by identifying lines where two words connected in the semantic network co-occur. Yet, current rules were only added to the grammar when the relation arguments matched the desired part-of-speech (POS). Previously, this did not always happen because, depending on the context, the same words might have different POS. For instance, most verbs can also be nouns (e.g. *break*, *cover*), or many nouns can behave as adjectives (e.g. *red*, *young*).

Another difference in the grammar is that, in addition to rules learned from human-created poetry, they were also acquired from proverbs and from Wikipedia sentences. As the lines of a poem are already kind of abstract or already involve figurative meanings, when PoeTryMe adds a new level of abstraction, the result can sometimes turn out to be more difficult, if not impossible, to interpret. On the other hand, Wikipedia text is not so creative but easily interpretable. Our intuition is that, combining both kinds of text, the previous issues will be more balanced, and the result may still slightly more clear, even when altered by PoeTryMe.

Being more strict when collecting rules resulted in a much smaller grammar. The current generation grammar, for Portuguese, covers about 1,500 renderings, which is substantially less than the previous 4,100 [9]. Hopefully, low quality grammar rules were left out.

4.2 Approach

The approach of the present system can be divided in three main steps, including the generation of a poem. Before generation, there is a seed acquisition stage, and a seed expansion stage. The seed acquisition stage goes as follows:

1. A topic t , in the form of a word or expression, is given as input.

2. Through the Twitter4J³ library, m tweets mentioning t are retrieved from Twitter.
3. Each tweet is processed and the top- f most frequent nouns, adjectives or verbs are collected and used as seed words.

If there is one, the main sentiment about the topic can also be estimated by counting the total number of smileys and emojis in the retrieved tweets. For each happy face (positive), a counter c is incremented by 1 ($c = c + 1$), and for each sad or crying face (negative) it is decremented ($c = c - 1$). The estimated polarity depends of the value of c . If $c > \theta$, it is positive, and if $c < -\theta$, it is negative, where θ is a predefined threshold.

Additional seeds can be obtained through the seed expansion procedure described earlier, which can be biased towards the polarity estimated in the previous stage. Alternatively, if there is a Wikipedia article about topic t , open words from its first sentence can also be used as additional seeds. This is an attempt to mix long-term data about the topic, in Wikipedia, with fresh information, from Twitter, and has similarities with what Toivanen et al. [20] do with Wikipedia and recent news.

Generation is the final stage and starts by feeding PoeTryMe with the set of seed words. As most of its previous instantiations (e.g. [28,29]), a generate & test strategy at the line level is followed. This means that, for each line in the target poem structure, text fragments are successively generated and scored against the target metre and presence of rhymes, while the best scoring are kept. The generation of each line stops either after a predefined number of generated candidates (n), or when a candidate line has precisely the target number of syllables and target rhyme, if there is one.

4.3 Twitterbot

Twitter is increasingly becoming a popular tool in the Computational Creativity community, not only as a source of information, but also as a platform for exhibiting the results of creative systems. While Twitterbots are autonomous systems, connected to Twitter that, from time to time, post messages, for creative Twitterbots, messages have a creative value. This includes the production of novel metaphors [30], riddles [31], or Internet memes [32], among others. Poetry has also been produced from the re-organisation of tweets [33].

Following the previous trend, the *@poetartificial*⁴ Twitterbot was developed for tweeting poems inspired by the current trends. Every hour, it reads the Twitter trends for Portugal, selects one of the three top trends, and runs the previously described approach for producing a set of poems inspired by the selected trend. The best-scoring poem is selected. Given the size limitations for tweets (140 characters), generated poems are currently blocks of four 10-syllable lines.

³ <http://twitter4j.org/>.

⁴ In Portuguese, *poeta artificial* means artificial poet.

4.4 Setup

PoeTryMe has several customisable parameters, most of them mentioned earlier. For the presented system, including the Twitterbot, the following parameters were set:

- Poem structure: block of four lines, each with 10 syllables;
- Relevant seeds obtained from expansion, $r = 5$;
- Surprise factor, $\nu = 0$;
- Strategy: generate & test
 - Maximum generations / line, $n = 2,500$
 - Increasing factor $\sigma = 0.8$; (to increase the probability of rhymes, $n = n + i * \sigma * n$, where i is the position of the line in the stanza)
- Score:
 - Each syllable missing / out of metre: +1 penalty point
 - Each rhyme: -2 penalty points

It should be highlighted that the surprise factor was set to 0 because we believe that there is already enough richness in the Twitter seeds, especially if they are expanded. Apart from that, there was not a big difference from previous instantiations of the system. The following additional parameters, specific of this instantiation, were set:

- #Retrieved tweets, $m = 200$;
- #Frequent words, $f = 5$;
- Polarity threshold, $\theta = m/20$ (5 % of the retrieved tweets);
- #Wikipedia words, $w = 5$.

The Twitterbot is currently generating 25 poems each hour, but publishing only the one with the best overall generate &test score, out of those that fit in a tweet (140 characters).

4.5 Examples

The new instantiation of PoeTryMe is illustrated by examples displayed in this section. All of them were produced in 19th January, 2016, using, as an example topic, “David Bowie”, a well-known musician whose death, 9 days earlier, was still echoed throughout Twitter.

The examples are presented in their original form, in Portuguese, together with a rough English translation. Of course that, due to the vagueness of language and to the inherent abstraction in a poem, the resulting translations were hard to reach, and often resulted in odd constructions. Examples were produced with different sets of seeds words, obtained with different seed expansion settings, enumerated in Table 1, together with the resulting seed words. The examples of Fig. 3 only use the top-5 frequent words in the retrieved tweets, while the others use five additional seeds, obtained by different means. In the poems of Fig. 4, the regular expansion procedure was applied to retrieve relevant words from the

semantic network. The poems of Figs. 5 and 6 used the same expansion algorithm but the former is biased towards positive seeds, and the second towards negative seeds. Finally, the poems in Fig. 7 use five additional seeds extracted from the abstract of David Bowie’s article in the Portuguese Wikipedia.

The top frequent words in the retrieved tweets include the word ‘music’, because Bowie was a musician. The word ‘homage’ is present because several tweets refer planned homages, especially from Lady Gaga, another musician. The verb ‘to make’ is often used with ‘homage’ – *to make an homage* – and was thus also frequent. The other two words, ‘partnership’ and ‘good’, were used several times to mention that Bowie refused to make a partnership with the band Coldplay, claiming that their music was not good.

Table 1. Seeds collected from Twitter and additional seeds obtained from different expansion settings.

Seeds	<i>homenagem, msica, fazer, parceria, bom</i> (homage, music, make, partnership, good)
Expansion	<i>associao, sociedade, lugar, comemorao, promessas</i> (association, society, place, commemoration, promises)
Expansion+	<i>virtuoso, harmonia, glria, alegria, carola</i> (virtuous, harmony, glory, joy, prayer)
Expansion-	<i>vo, lbia, treta, vassalas, partida</i> (vain, wordy, bullshit, vassals, departure)
Wikipedia	<i>musical, artistico, inglts, nome, ator</i> (musical, artistic, english, name, actor)

Besides the original seeds, the poems in Fig. 4 use directly related words, obtained from the semantic network. The exact connection of these words with the domain can be confirmed with the help of the Contextualizer. Related words include synonyms of homage (glory, proof), music (harmony) or of ‘to make’ (invent, practise, charge), as well as a hypernym of partnership (association). Some of these relations are not held by the same sense of the seed words,

<i>glria da homenagem em cruz</i>	glory of homage in the cross
<i>homenagem a prova de luz</i>	homage is a proof of illumination
<i>de associaes de parceria</i>	of partnership associations
<i>sem achar msica, nem harmonia</i>	without finding music, nor harmony
<i>fazer em frente, ir a inventar</i>	make ahead, go and invent
<i>fazer em frente, ir a praticar</i>	make ahead, go and practice
<i>fazer em frente, ir a facturar</i>	make ahead, go and charge
<i>mais vale fazer que inventar</i>	it is better to do than to invent

Fig. 3. Poems inspired by the topic David Bowie, without seed expansion.

and they are often not the sense we would first thought of. Although the system does not do it intentionally, but because word senses are not handled, we like to see this as an open door to the presence of figurative language.

The poems in Fig. 4 mix the utilisation of the original seeds with other semantically-relevant, together with words directly-related to one of the previous. This includes synonyms of ‘to make’ (proceed, conclude) or of ‘to collect’ (gather), a hyponym of recordings (record), also words like ‘albums’ that may be collected, or ‘promises’ that should be devoted.

<i>coleccionar sem tocar em lbuns</i>	collect without touching in albums
<i>coligir para coleccionar</i>	gather to collect
<i>melhor fazer que representar</i>	it is better to do than to act
<i>um registo de gravaes doiradas</i>	a record of golden recordings
<i>sempre a fazer e a proceder</i>	always making and proceeding
<i>sempre a concluir e a fazer</i>	always completing and making
<i>no o queira depois o lugar dar</i>	do not want then to give your place
<i>ter promessas sem saber consagrar</i>	to have promises without knowing to devote

Fig. 4. Poems inspired by the topic David Bowie, with expanded seeds.

The poems in Fig. 5 use the original seeds and other semantically-relevant words with a positive polarity. Words directly related to the previous include synonyms of harmony (communion) and joy (satisfaction), hypernyms of partnership (society) and harmony (art), a hyponym of homage (proof) and an action that causes joy (to rejoice).

<i>alegrar resulta em alegria</i>	rejoice results in joy
<i>de sociedades de parceria</i>	of partnership societies
<i>alegrar resulta em alegria</i>	rejoice results in joy
<i>a arte principal a harmonia</i>	the main art is harmony
<i>homenagem a prova de luz</i>	homage is a proof of illumination
<i>comunho da harmonia em cruz</i>	harmony communion in the cross
<i>no h satisfao sem alegria</i>	there is not satisfaction without joy
<i>de sociedades de parceria</i>	of partnership companies

Fig. 5. Poem inspired by the topic David Bowie, with expanded positive seeds.

The poems in Fig. 6 use the original seeds and other semantically-relevant words with a negative polarity. Words directly related to the previous include a synonym of vain (lied), a hypernym of bullshit (verbiage), a hyponym of music (cheap music) and an association with tax (efficient), among others.

<i>andar ao palavreado da treta</i>	wandering to the verbiage of bullshit
<i>com msica santa e musiqueta</i>	with holy music and cheap music
<i>eficientes, imposto de mo</i>	efficient, hand tax by heaven's decree, lied and in
<i>por decreto do cu, mentido e vo</i>	vain
<i>tudo comeou a um vo embalde</i>	it all started to one useless vain
<i>tudo comeou a um vo debalde</i>	it all started to one worthless vain
<i>olha os guelfos da nossa partida!</i>	looks the Guelphs of our departure!
<i>de que turma fazer a minha vida</i>	of what class to make my life

Fig. 6. Poem inspired by the topic David Bowie, with expanded negative seeds.

The poems in Fig. 7 use the original seeds and other collected from Wikipedia. As expected, the words collected from Wikipedia are more stable associations with David Bowie: he was an English singer, musical producer, and also an actor. Semantically-related words include synonyms of producer (creator) or name (power), a hypernym of singer (artist) and of musical (movie).

<i>de seu filme, lugar e musical</i>	of your movie, place and music
<i>seu filme de musical palatal</i>	his musical film is palatal
<i>, portanto, um nome de poder</i>	it is therefore a name of power
<i>fazer para chegar a fazer</i>	make to get to make
<i>de criador e de produtor</i>	it is of a creator and a producer
<i>que alma, que produtor, que criador</i>	what a soul, what a producer, what a creator
<i>as responsveis no tm produtor</i>	the responsible do not have a producer
<i>o artista e depois o cantor</i>	the artist and then the singer

Fig. 7. Poems inspired by the topic David Bowie, with additional seeds from Wikipedia.

Additional real-time examples can be checked in the Twitter feed of the user @poetartificial, where the bot operates in real time.

4.6 Critical View

In the displayed examples, features like the regular metre and the frequent presence of rhymes arise. With the current settings, these features are often met. In fact, with the current linguistic resources, meeting them is mostly a matter of increasing the number of generations per line. Grammatical constraints are also frequently satisfied, especially now, with more strict grammar on the POS of the arguments. Though smaller than previous grammars, we can say that the richness of the underlying resources still enables the generation of poems with an interesting degree of variation, which contributes to a positive perception of creativity.

Each line is semantically-coherent, because semantically-related words are basically put in the position of other words that hold the same relation, in what

can be seen as a shallow exploitation of analogy. Stranger situations might occur from fixed words in the rule’s body, too specific and with a strong connection with the original words, but not so much with the replacements, at least as long as the semantic network is well-structured. At the same time, this is where the poem may become interesting, or when it may fail. But this situation is not that frequent in 10-syllable lines, typically restricted to functional words (e.g. determiners and prepositions), besides the pair of related words.

In the displayed examples, when the previous situations happened, we can say it went quite well in lines such as: *homage is a proof of illumination*, or *harmony communion in the cross*, where the words ‘illumination’ and ‘cross’ were fixed. On the other hand, the line *his musical film is palatal*, where the word ‘palatal’ is fixed, is odder. As most of the more specific words come from the rules learned from Wikipedia, its usage for this purpose should be rethought.

Still on semantics, although they are generated independently, lines end up having some unity, together and with the topic. This happens because they are based on the same semantic domain, set by the topic. Though not frequent, issues might arise from an odd order of the lines. This happens because the system does not have any concern on showing some kind of evolution from the beginning to the end of the poem.

Finally, the connection of the poem with the topic is sometimes too tenuous. Using associated words is not always enough for this purpose. Moreover, there might be strongly-associated words that are not in the semantic network, and will thus never be used. This is why we say that the poem is inspired by the topic, in a sense that it uses related words in semantically-coherent sentences, but we do not claim that it is about the topic.

5 Concluding Remarks

A new instantiation of PoeTryMe, a poetry generation platform, was presented. The singular feature of the presented system is that its poetry is inspired by Twitter trends, more precisely, words that are associated with those.

Presented examples illustrate the potential of this system, but also its limitations. Despite the presence of a regular metre, rhymes, grammatical sentences and semantically-related words, the connection with the topic is not always very clear. If it is not known, it is often hard to identify the original topic. Especially when fresh associations, possibly valid just for a short period of time, are used.

This is why the poems are not about a topic, but inspired by it – given a topic, currently associated words are extracted and fed to PoeTryMe, which expresses “what it knows” about those words, poetically. Writing about the topic would require deeper linguistic processing of the tweets and possibly other sources of knowledge. On this problem, Tobin and Manurung [21] extract predicate-argument structures from an input article and try to keep the same structure during generation. They admit, however, that considering this together with other features results in too much complexity and a long time for producing a poem.

In order to improve the connection with the topic, we are devising an alternative approach for setting the generation domain, but simpler than the previous. In the preprocessing step, whenever a tweet uses two words that, according to our semantic network, are related, this relation will be added to the set of identified relations. Then, instead of using seed words for setting the generation domain, this domain should consist of a graph with the identified relations, which is possible in the PoeTryMe architecture. If the graph is too small, additional relations can be used, based on their distance to the domain, or the current strategy can still be used as a fallback. The issue of only using words from the semantic network remains, though.

Anyway, trends are always changing and people say different things about them. So, independently of writing about the topic or not, the connection to Twitter enables the continuous generation of different poems every time. This can also be seen as a test to the limits of the system, and will certainly give hints for further improvements. Moreover, the Twitterbot will hopefully shake a little bit the (still small) Portuguese community of Twitter users.

Acknowledgements. This work was developed in the scope ConCreTe – *Concept Creation Technology*. The project ConCreTe acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 611733.

References

1. Colton, S., Wiggins, G.A.: Computational creativity: the final frontier? In: Proceedings of 20th European Conference on Artificial Intelligence (ECAI 2012). Frontiers in Artificial Intelligence and Applications, Montpellier, France, vol. 242, pp. 21–26. IOS Press (2012)
2. Machado, P., Cardoso, A.: All the truth about NEvAr. Appl. Intell. Spec. Issue Creat. Syst. **16**, 101–119 (2002)
3. Binsted, K., Ritchie, G.: An implemented model of punning riddles. In: Proceedings of 12th National Conference on Artificial Intelligence, AAAI 1994, vol. 1, pp. 633–638. AAAI Press, Menlo Park (1994)
4. Gervás, P., Díaz-Agudo, B., Peinado, F., Hervás, R.: Story plot generation based on CBR. Knowl.-Based Syst. **18**, 235–242 (2005)
5. Veale, T., Hao, Y.: A fluid knowledge representation for understanding and generating creative metaphors. In: Proceedings of 22nd International Conference on Computational Linguistics, COLING 2008, vol. 1, pp. 945–952. ACL Press, Manchester (2008)
6. Smith, M.R., Hintze, R.S., Ventura, D.: Nehovah: a neologism creator nomen ipsum. In: Proceedings of 5th International Conference on Computational Creativity, Ljubljana, Slovenia, ICC3 2014 (2014)
7. Tomašić, P., Znidaršić, M., Papa, G.: Implementation of a slogan generator. In: Proceedings of 5th International Conference on Computational Creativity, Ljubljana, Slovenia. ICC3 2014, pp. 340–343 (2014)

8. Gonalo Oliveira, H.: PoeTryMe: a versatile platform for poetry generation. In: Proceedings of ECAI 2012 Workshop on Computational Creativity, Concept Invention, and General Intelligence, Montpellier, France, C3GI 2012 (2012)
9. Gonalo Oliveira, H., Cardoso, A.: Poetry generation with PoeTryMe. In: Besold, T.R., Schorlemmer, M., Smaill, A. (eds.) *Computational Creativity Research: Towards Creative Machines*. Atlantis Thinking Machines, pp. 243–266. Atlantis/Springer, Lynnwood/Heidelberg (2015)
10. Gervás, P.: WASP: evaluation of different strategies for the automatic generation of Spanish verse. In: Proceedings of AISB 2000 Symposium on Creative and Cultural Aspects and Applications of AI and Cognitive Science, Birmingham, UK, pp. 93–100 (2000)
11. Manurung, H.: An evolutionary algorithm approach to poetry generation. Ph.D. thesis, University of Edinburgh (2003)
12. Colton, S., Goodwin, J., Veale, T.: Full FACE poetry generation. In: Proceedings of 3rd International Conference on Computational Creativity, Dublin, Ireland, ICC3 2012, Dublin, Ireland, pp. 95–102 (2012)
13. Toivanen, J.M., Järvisalo, M., Toivonen, H.: Harnessing constraint programming for poetry composition. In: Proceedings of the 4th International Conference on Computational Creativity, ICC3 2013, pp. 160–167. The University of Sydney, Sydney (2013)
14. Agirrezabal, M., Arrieta, B., Astigarraga, A., Hulden, M.: POS-tag based poetry generation with wordnet. In: Proceedings of the 14th European Workshop on Natural Language Generation, pp. 162–166. ACL Press, Sofia (2013)
15. Manurung, H.: A chart generator for rhythm patterned text. In: Proceedings of 1st International Workshop on Literature in Cognition and Computer (1999)
16. Levy, R.P.: A computational model of poetic creativity with neural network as measure of adaptive fitness. In: Proceedings of the ICCBR-01 Workshop on Creative Systems (2001)
17. Gervás, P.: An expert system for the composition of formal Spanish poetry. *J. Knowl.-Based Syst.* **14**, 200–201 (2001)
18. Kurzweil, R.: Cybernetic poet. <http://www.kurzweilcyberart.com/poetry/rkcp-overview.php>
19. Misztal, J., Indurkha, B.: Poetry generation system with an emotional personality. In: Proceedings of 5th International Conference on Computational Creativity, ICC3 2014, Ljubljana, Slovenia (2014)
20. Toivanen, J.M., Gross, O., Toivonen, H.: The officer is taller than you, who race yourself! Using document specific word associations in poetry generation. In: Proceedings of 5th International Conference on Computational Creativity, Ljubljana, Slovenia, ICC3 2014 (2014)
21. Tobing, B.C.L., Manurung, R.: A chart generation system for topical metrical poetry. In: Proceedings of the 6th International Conference on Computational Creativity, ICC3 2015, Park City, Utah, USA (2015)
22. Wong, M.T., Chun, A.H.W.: Automatic haiku generation using VSM. In: Proceedings of 7th WSEAS International Conference on Applied Computer and Applied Computational Science (ACACOS 2008), Hangzhou, China (2008)
23. Netzer, Y., Gabay, D., Goldberg, Y., Elhadad, M.: Gaiku: generating haiku with word associations norms. In: Proceedings of the NAACL 2009 Workshop on Computational Approaches to Linguistic Creativity, CALC 2009, Boulder, Colorado, pp. 32–39. ACL Press (2009)

24. Yan, R., Jiang, H., Lapata, M., Lin, S.D., Lv, X., Li, X.: I, poet: automatic Chinese poetry composition through a generative summarization framework under constrained optimization. In: Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI 2013, pp. 2197–2203. AAAI Press (2013)
25. Rashel, F., Manurung, R.: Pemuisi: a constraint satisfaction-based generator of topical Indonesian poetry. In: Proceedings of 5th International Conference on Computational Creativity, ICCC 2014, Ljubljana, Slovenia (2014)
26. Das, A., Gambäck, B.: Poetic machine: computational creativity for automatic poetry generation in Bengali. In: Proceedings of 5th International Conference on Computational Creativity, ICCC 2014, Ljubljana, Slovenia (2014)
27. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. *Comput. Netw.* **30**, 107–117 (1998)
28. Gonçalves Oliveira, H., Hervs, R., Daz, A., Gervs, P.: Adapting a generic platform for poetry generation to produce Spanish poems. In: Proceedings of 5th International Conference on Computational Creativity, ICCC 2014, Ljubljana, Slovenia (2014)
29. Gonçalves Oliveira, H.: Tra-la-lyrics 2.0: automatic generation of song lyrics on a semantic domain. *J. Artif. Gen. Intell.* **6**, 87–110 (2015). Special Issue: Computational Creativity, Concept Invention, and General Intelligence
30. Veale, T., Valitutti, A., Li, G.: Twitter: the best of bot worlds for automated wit. In: Streitz, N., Markopoulos, P. (eds.) DAPI 2015. LNCS, vol. 9189, pp. 689–699. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-20804-6_63](https://doi.org/10.1007/978-3-319-20804-6_63)
31. Guerrero, I., Verhoeven, B., Barbieri, F., Martins, P. Perez y Perez, R.: TheRid-dlerBot: a next step on the ladder towards creative Twitter bots. In: Proceedings of the 6th International Conference on Computational Creativity, ICCC 2015, pp. 315–322. Brigham Young University, Park City (2015)
32. Costa, D., Gonçalves Oliveira, H., Pinto, A.: “In reality there are as many religions as there are papers” - first steps towards the generation of Internet memes. In: Proceedings of the 6th International Conference on Computational Creativity, ICCC 2015, Park City, UT, USA, pp. 300–307 (2015)
33. Charnley, J., Colton, S., Llano, M.T.: The FloWr framework: automated flowchart construction, optimisation and alteration for creative systems. In: 5th International Conference on Computational Creativity, ICCC 2014, Ljubljana, Slovenia (2014)

Knowledge Discovery, Knowledge Engineering and
Knowledge Management

7th International Joint Conference, IC3K 2015, Lisbon,
Portugal, November 12-14, 2015, Revised Selected
Papers

Fred, A.; Dietz, J.; Aveiro, D.; Liu, K.; Filipe, J. (Eds.)

2016, XIX, 534 p. 187 illus., Softcover

ISBN: 978-3-319-52757-4