



Poetry at the first steps of Artificial Intelligence

Christina Linardaki

Abstract: This paper is about Artificial Intelligence (AI) attempts at writing poetry, usually referred to with the term “poetry generation”. Poetry generation started in the Digital Humanities, which developed out of humanities computing; nowadays, however, it is part of Computational Creativity, a field that tackles several areas of art and science.

No matter which field is undertaking the venture of poetry generation, it is certainly a fascinating process, as it essentially involves a reinvention of how the human brain works. Moreover, if the vision is to arrive at a novel form or idiom of poetry, the product can be altogether refreshing. But so far, AI attempts at poetry generation are not so different from equivalent human attempts based on algorithmic processes. This means that, in a way, AI attempts have dovetailed with poetry’s long tradition as an art form in evolution.

In the paper, it is first examined why poetry was chosen among other literary genres as a field for experimentation. Mention is made to the characteristics of poetry (namely arbitrariness and absurdity) that make it fertile ground for such endeavors and also to various text- and reader-centered literary approaches that favored experimentation even by human poets.

Then, a rough historic look at poetry generation is attempted. This is followed by a review of the methods employed — either for fun or as academic projects — along Lamb et al.’s taxonomy (“A Taxonomy of Generative Poetry Techniques”), which distinguishes between mere poetry generation and result enhancement. Another taxonomy by Gonçalo Oliveira (“A Survey on Intelligent Poetry Generation: Languages, Features, Techniques, Reutilisation and Evaluation”), dividing between form and content issues in poetry generation, is also briefly presented.

The results of poetry generators are evaluated against Manurung’s (*An evolutionary algorithm approach to poetry generation*) three preconditions for successful AI-generated poems: (i) grammaticality, (ii) meaningfulness, and (iii) poeticness. The reasons for any poorness occurring are also looked into: the inability of computers to understand words as signs with a signified, a lack of general intelligence, process- (rather than output-) driven attempts, etc.

Then, computer-like results (algorithmic processes) from a number of human poetic movements are also presented as a juxtaposition: Dada, stream of consciousness, OuLiPo, LangPo, Flarf, and blackout/erasure poetry. A remarkable similarity is found to surface between (i) human poets who are more concerned with experimentation than quality of results, and (ii) computer scientists who are process- rather than outcome-driven. This leads to a discussion of the characteristics of humanness, the possibility of granting future AI personhood, and the need to see our world in terms of a new and more refined ontology.

Introduction

Attempts at poetry generation using Artificial Intelligence (AI) started as early as 1959, but it was already in the 1940s that humanities computing had given rise to Digital humanities. Digital humanities (DH) was (and is) a field that applies automation processes to texts, and this included attempts at poetry generation, without however ever making this its main goal. It was only after the turn of the millennium that computer science started seeing poetry generation as fertile ground for AI experimentation within the scope of Computational Creativity, a field that tackles several areas of art (such as prose, music, painting, song lyrics, photobashing, etc.) and science.

If Digital Humanities' attempts at poetry generation were restricted to rearranging or otherwise mixing human-written poems to produce something different, the question now with computer science was how to arrive at "intelligent" ways of generating poems rather than ruminating pieces of existing ones. To achieve this, complex and often knowledge-intensive systems that involved several computational and linguistic layers were employed. Creators of poetry generators today are less focused on the output; they are rather focused on the process of generation *per se*, a process that involves the teaching of AI to perform desired tasks. Thus, the potential contribution of humanities was rather disregarded and jostled altogether.

Various attempts at poetry generation, either for fun or as academic projects, can be enlightening with regard to understanding the complex computational processes involved. But first, we will tackle the question of why poetry was chosen over other literary genres.

Why poetry in the first place?

Why did the creators of poetry generators choose poetry over another creative writing form, e.g. bonsai short stories (with a typical ceiling of around 600 words)? The answer lies in the characteristics of poetry relative to other literary genres. A short story, however small, must have a structure — characters, dialogues, plot, and so on. Poetry, on the other hand, incorporates a degree of arbitrariness since there are no strict or universal rules for what is acceptable or not. It also incorporates a degree of absurdity, as many a time it not only resists commonly acceptable meaning, but also reverses it. Despite being considered one of the most difficult literary genres to grasp due to its multitude of language levels (vocabulary, grammar, syntax, semantics, phonemics, metrics, figures of speech, etc.), poetry does not necessitate the use of all of them at the same time. In other words, it is not mandatory for all these elements to be present in a single poem, and this has certainly been noticed by computer scientists. Manurung (*An evolutionary algorithm approach to poetry generation*), for example, has noted that writing poetic text does not have to be extremely precise, as several rules of language that hold in the generation of natural language can be broken in poetry.

The above characteristics theoretically make poetry more fertile ground for experimentation than any other literary genre. This is not only the case with AI; personal experience in literary criticism has shown that many new and ambitious human writers start off with poetry because they consider it more convenient and easier, and as a consequence, they usually begin by writing mediocre poems. Why should we expect any different from AI?

Notwithstanding the above, there is also another reason to choose poetry as a field for experimentation: text- and reader-centered literary approaches have turned the headlights away from the author. The spotlight has shifted to the text first and the reader next — bringing about the "death of the author" (Barthes) and giving the reigns of determining the poem's meaning to the other two. Among text-centered approaches and schools

in literary theory, there are notably the following: Russian formalism (early 20th century), which looks for the meaning of a literary text in the mechanisms employed and its structural principles, i.e. the very qualities that make a text literary (Medvedev and Bakhtin); American New Criticism (1930s and on), which looks into the way that poetry interacts with the semantics of language (Richards; Eliot); and Structuralism (1970s), which focuses on the conventions that allow literature to exist (Todorov). Under such approaches, the literary text is seen as an autonomous entity, the understanding of which requires nothing apart from its very words.

Reader-centered approaches, on the other hand, include the following: Semiotics (late 19th century), which supports that the meaning of a literary text arises from the interaction of the sign, the interpreter, and the event, and can change over time (Culler); Psychoanalytic Criticism (1970s and on) which speaks of reading as a means of rebuilding the identity of the reader through a transactive relation that is established between him/her and the text (Holland); Reception Aesthetics (Rezeptionsästhetik, 1960s), which maintains that the past experiences embedded in a text merge with the interests, cultural stock, and personal experiences of the readers as they read (Jauss); and Reading Response Criticism (1970s), which talks of the ambiguity inherent in the text, creating gaps which the reader must bridge in order to arrive at an interpretation (Iser).

The autonomy of the text, coupled with the reign of the reader as the assigner of meaning to it, caused a multitude of experimentations in literature to emerge. And much like computer scientists who are fascinated by the process, several poetic movements favored experimentation over a result that would fit mainstream poetry standards, challenging even the idea of such standards. This point will be addressed further, after a review of AI poetry generation history, methods, and processes.

The baby steps of poetry generation

First attempts to develop poetry using computers are more than 60 years old. The first poetry generation program is considered to be Theo Lutz's "Stochastic Texts" in 1959. The program ambioned to write poem-like texts from scratch in the German language (Lutz). A sample of such a "poem" is the following:

NICHT JEDER BLICK IST NAH UND KEIN DORF IST SPÄT.
 EIN SCHLOSS IST FREI UND JEDER BAUER IST FERN.
 JEDER FREMDE IST FERN. EIN TAG IST SPÄT.
 JEDES HAUS IST DUNKEL. EIN AUGE IST TIEF.

NOT EVERY LOOK IS NEAR. NO VILLAGE IS LATE.
 A CASTLE IS FREE AND EVERY FARMER IS FAR.
 EVERY STRANGER IS FAR. A DAY IS LATE.
 EVERY HOUSE IS DARK. AN EYE IS DEEP.

(Translation by Helen MacCormac, 2005)

Same as with every other sample in the paper, this sample is seen through the lens of Manurung's (*An evolutionary algorithm approach to poetry generation*) three preconditions for successful AI-generated poetry; Manurung said that AI-generated poems should include language that ensures the three elements of (i) grammaticality, (ii) meaningfulness, and (iii) poeticness. According to these preconditions, it is obvious that the results of the "Stochastic Texts" may be grammatical, but they are otherwise poor in terms of meaning and even poorer in terms of poeticness. Of course, one could argue that it was the first such attempt, after all. And perhaps it was remarkable to arrive even at such a result at the time.

Digital Humanities vs Computational Creativity

Other early attempts at poetry generation mainly involved rearrangements or imaginative combinations of existing human-written poems. This was perhaps the result of the nature of the framework in which such poetry generation occurred: DH developed out of humanities computing (Henrickson), which has origins dating back to the 1940s and 50s vis-à-vis the pioneering work of scholars such as Roberto Busa and Josephine Miles (Stagnaro; Wimmer).

Busa met Thomas J. Watson, the founder of IBM, in 1949 and convinced him to sponsor the *Index Thomasticus* project. The project involved a computer-generated concordance to Thomas Aquinas's writings, and many people today still consider it to mark the onset of the field of computational linguistics. Josephine Miles, on the other hand, a poet and professor of English at Berkeley, became project director of an abandoned index-card-based concordance to the poetical works of John Dryden in the early 1950s. Partnering with the Electrical Engineering department at Berkeley and using their IBM tabulation machine, Miles used automatic methods to complete the concordance. This was eventually published in 1957, around the time that Busa circulated early drafts of his own concordance to Aquinas's writings (Wimmer).

Following these pioneers, other scholars from the humanities began using mainframe computers to automate tasks like word-searching, sorting, and counting, which was much faster than processing information from texts manually. All of these tasks involved DH. DH also touched on attempts at poetry generation, although never as holistically and persistently as AI is doing today. A notable example of DH poetry generation was the French "Atelier of Literature Assisted by Maths and Computers" (ALAMO). ALAMO sprang in 1981 out of a poetic movement of the time, the OuLiPo, which was an acronym of *Ouvroir de Littérature Potentielle*, i.e. Workshop for Potential Literature (Oulipo; Berkman). An illustrious attempt of the ALAMO includes the *rimbaudelaires*, in which the altered structure of a sonnet by Rimbaud, acquired by cutting out nouns, verbs and adjectives, was filled in with words from Baudelaire's poems (Laws). *Rimbaudelaires* are a strange mixture, producing lines that echo 19th-century French poetry and rendering rather solid results in terms of both coherence and meaning. Seen against Manurung's (*An evolutionary algorithm approach to poetry generation*) preconditions (once more: grammaticality, meaningfulness, and poeticness), we may say that *rimbaudelaires* are good in terms of grammaticality, but they lack coherence (thus meaningfulness) in the succession of ideas, even though there is an allure of poeticness.

A good example of *rimbaudelaires* is the following (Laws):

Le Rêveur du bonheur
 C'est un lac de poitrine où passe une gamine
 Embrassant librement aux anges des sommeils
 D'argent; où le plaisir de la caresse fine
 Fuit: c'est un poudreux bonheur qui se rit de soleil

Which roughly translates into:

The Dreamer of happiness
 Is a bosom lake where a youngster passes
 Kissing freely the angels of sleeps
 Silver; where the pleasure of the fine caress
 Flees: it's a powdery happiness that laughs at the sun

After the turn of the millennium, computer science started seeing the field of poetry generation as fertile ground for AI experimentation. A sound result was suddenly secondary in importance. The process was what mattered the most; how could AI arrive at “intelligent” ways of generating satisfactory poems? Trying to answer these questions, AI experts put to work language generation systems. These were often knowledge-intensive and involving several layers of language, such as phonetics, diction, syntax and semantics (Gonçalo Oliveira, *Multilanguage extension and evaluation of a poetry generator*), and several computational methods. The sad thing was that the role of humanities in the process was completely upended. This is probably due to “the cultural and historical divisions between artists, humanists, engineers, and scientists” (Roque). It is sad, however, because humanists might tangibly help computer scientists keep a keener eye on the output of poetry generation as well as the input.

Online poetry generators for fun

There are numerous sites that produce poetry online mainly for recreation as a game. Some of them are Poem Generator (*Poem Generator*),¹ Language is a Virus (“Poetry”),² Fantasy Name Generators (for haiku) (“Haiku Generator”),³ Acrostic Poem Generator (for acrostic poems) (*Acrostic Poem Generator*),⁴ The Genuine Haiku Generator (“The Genuine Haiku Generator”),⁵ Poem Portraits (Delvin),⁶ and Verse by Verse (*Verse by Verse*).⁷ Other projects, such as Poem Creator, which mixes up Shakespearean poems with Hip Hop songs, are available only as an iPad app.

Let’s see an example of Google’s AI poetry:⁸

this was the only way.
it was the only way.

-
- ¹ This poem generator can generate many types of poems: free verse, haiku, rhyming couplets, sonnets, limericks etc. Poems are based on keywords that the user provides. See <https://www.poem-generator.org.uk/> All links in this paper were last accessed on 22.7.2021.
 - ² This site avails of a number of different poetry engines that include haiku and sonnets. It also avails of an interactive haiku generator that uses words which the user provides and are incorporated into templates created from famous haiku poems. See <http://www.languageisavirus.com/poetry.php>
 - ³ A random haiku generator that respects the standard (as perceived in the West) 5-7-5 syllable structure of the haiku. It also contains a disclaimer: “due to the randomness of a generator not all results will make perfect sense from a grammatical and/or strictly linguistic point of view”. See <https://www.fantasynamgenerators.com/haiku-generator.php>
 - ⁴ Generator of acrostic poems (i.e. poems that spell out a word down the left side of the poem, beginning a new line with each letter of the word). The generator uses the name that a user enters to generate the poem. See <http://www.acrosticpoem.org/>
 - ⁵ Random pseudo haiku generator. Using a Java Script application, this page automatically generates about $7.62e+39$ different haiku poems which, however, are admitted to be loaded with “too many incompatible images”. As for the generation mechanism, on the “What is this?” pad on the site, it is stated that “There are no predefined sentence structures; no pre-packaged lines. The words themselves decide which type of word will follow, as randomly as possible, while still maintaining some semblance of adherence to grammar and readability.” See <http://www.everypoet.com/haiku/>
 - ⁶ Google’s AI poetry generation project, undertaken with researchers at Stanford and the University of Massachusetts is “an experiment at the boundaries of AI and human collaboration”. Users are asked to donate a word and become part of an ever-evolving collective poem. If they also allow use of their camera, they will receive a photo portrait of theirs accompanied by a two-verse poem. See <https://artsexperiments.withgoogle.com/poemportraits>
 - ⁷ Google’s AI poems inspired by famous American poets of the 19th century. The user selects a poet from a list and then designs the desired poem structure. Then, s/he writes the first line and chooses the rest from a list of suggestions. See <https://sites.research.google/versebyverse/>
 - ⁸ More at: <https://www.writerswrite.com/google-ai-writes-poetry-51920161>

it was her turn to blink.
 it was hard to tell.
 it was time to move on.
 he had to do it again.
 they all looked at each other.
 they all turned to look back.
 they both turned to face him.
 they both turned and walked away.

Judged against Manurung's preconditions (*An evolutionary algorithm approach to poetry generation*), we see that Google is doing well with grammaticality, but not so much so with meaningfulness (there is little coherence and a significant amount of repetition that obscures meaning), let alone poeticness. Like this one, most results from such online poetry generators are generally clumsy. Additionally, they may require a generous donation of words or even chunks of speech from the user, according to what they would prefer to see in the generated poem. In the above example, "it was the only way" was the donated phrase.

Poetry generators as academic projects

Besides the fun of the absurdity (or the occasional random success) of online poetry generators like these previous ones, poetry generation has also been undertaken as part of meticulous academic projects. Such projects used a variety of methods, much like poetry generators for fun have used, only in the academic framework, they have become more transparent. Among early academic projects were PoeTryMe (Gonçalo Oliveira, "Po"TryMe"),⁹ Poevolve (Levy),¹⁰ Tra-La-Lyrics 2.0 (Gonçalo Oliveira, "Tra-la-Lyrics 2.0"; Gonçalo Oliveira, et al., "Text based on Rhythm"), and ASPERA (Gervás, "Generating poetry from a prose text"). A more up-to-date joint project would be Deep-Spear,¹¹ which produces lines like the following ones:

that is the world art we a lord of god
 give in him, evermore to teach my prayers
 and only to forget her for its sake
 it is not love, for all thy flesh and peace.

Against Manurung's preconditions (*An evolutionary algorithm approach to poetry generation*), it is evident that with Deep-Spear, grammaticality is achieved (albeit with a Middle English accent), but meaningfulness is not (certain phrases are simply nonsensical) and neither is poeticness. Thus, the output is again rather poor.

Lamb et al. ("A Taxonomy of Generative Poetry Techniques") have devised a taxonomy which separates between "mere" poetry generation and result enhancement. In mere generation, a computer produces poems based on existing human-written ones and an algorithm that prescribes how the former might be used. As stated in the paper, "the results of mere generation can appear nonsensical; this is not always a bad thing from an artistic perspective" (Lamb et al. "A Taxonomy of Generative Poetry Techniques"). This is an interesting point that is worth some pondering. Inasmuch as we are ready to accept that a human can produce nonsensical

⁹ Developed by the Center of Informatics and Systems of the University of Coimbra, Portugal. The generator comes with a diagram describing its operation. See: <https://poetryme.dei.uc.pt/>

¹⁰ See a detailed account by its creator at <https://github.com/rplevy/poevolve>

¹¹ Deep-Spear is a joint neural model of poetic language, meter and rhyme. See Lau et al. <https://aclanthology.org/P18-1181/> and <https://aclanthology.org/P18-1181.pdf>

poetry as a result of experimentation (because we know that they are capable of also producing poetry that makes perfect sense at the same time), we are less willing to accept that computers produce nonsensical poetry because it is their aim to do so. The primary reason for any nonsense is that AI uses words without understanding their meanings. AI may be able to tell that in the phrase “She has an apple”, “she” is a pronoun, “has” is a verb, and “apple” is a noun, or that “she” is the subject, “has” is the verb and “apple” is the object. But it cannot understand what this phrase is actually referring to in the real world nor draw a mental image, as the human brain does, of a woman holding an apple. Words for AI are just data sequences, referring to no real information (Harel and Rumpel). In Saussurian terms, we would say that words for AI are one-dimensional signs; they involve just the signifier, without having a clue on what the signified is.

Bostrom and Yudkowsky agree to that when they say that what AI lacks is generality. AI is trained to perform a single task. In this way, it resembles other-than-human biological life; in the example that the authors use, beavers can build dams and bees can build hives, but beavers cannot build hives and bees cannot build dams. A human watching beavers and bees, however, can build both, because human intelligence is applicable more generally, reflecting this generality. This has also been noted by the European Parliament Research Service, which has included in its report that a long-term goal of AI is the so-called Artificial General Intelligence (AGI).

Methods of mere poetry generation

Among the methods of mere generation that Lamb et al. (“A Taxonomy of Generative Poetry Techniques”) mention are the following:

(a) **Templates.** Template generation, also called slot-filling, has seen a wide application, as was the case with the *rimbaudelaires* we saw above. Templates can be extracted automatically from text, if lists of words or phrases pertaining to different categories (e.g. nouns or verbs) are drafted. Then, line templates with slots into which words from given lists can be inserted are created. Finally, words from the appropriate list are selected randomly to fill in the slots.

Otherwise, templates can be handcrafted. Colton et al., for example, have used

very short phrases (similes) mined from the internet, alongside the phrases of professional writers, namely journalists writing for the British *Guardian* newspaper. The former phrases fit into the long-standing tradition of using the words of the common man in poetry, and the latter reflects ‘the desire to increase quality while not appropriating text intended for poems.’ (97)

Theirs was a template-based approach which produced not only poetry but also brief stories and comments that accompanied each poem. This is because they identified the problem of computers having not been taken seriously as creative entities to be one of a lack of appropriate context. Thus, they saw that perhaps a story generator, able to produce “static framing information” and (in the future) to “reply with a story to any question asked of it in a dialogue situation” (101), would compensate for the lack of “human quality” in generating poems meant —after all— to be read by humans. The following is an example of a poem they generated, with the story first and then the poem entitled “Blue overalls,” following:

It was generally a good news day. I read a story in the *Guardian* culture section entitled: “South Africa's ANC celebrates centenary with moment in the sun”. It talked of south africans, interfaith prayers and monochrome photos. Apparently, “The heroic struggle against a racist regime was remembered:

those thousands who sacrificed their lives in a quest for human rights and democracy that took more than eight decades” and “At midnight he watched with amusement as Zuma lit the centenary flame, at the second attempt, with some help from a man in blue overalls marked ‘Explosives’”. I wanted to write something highly relevant to the original article. I wrote this poem. (101)

Blue overalls

the repetitive attention of some traditional african chants
a heroic struggle, like the personality of a soldier

an unbearable symbolic timing, like a scream
blue overalls, each like a blueberry
some presidential many selfless leaders

oh! such influential presidents
such great presidents
blueberry-blue overalls

lark-blue overalls
a knight-heroic struggle (*Ibidem*)

The story is better than the poem in terms of meaning, though the algorithm becomes transparent because of the irrelevance of the two quotes, which give an overall awkward impression. The poem is weaker because, judged against Manurung’s (*An evolutionary algorithm approach to poetry generation*) preconditions for sound AI-produced poetry, it has little meaningfulness (one of the main issues with template-based poetry generation). As Lamb et al. (“A Taxonomy of Generative Poetry Techniques”) point out in their discussion of templates, this poem too suffers from useless repetitiveness, but there is also little coherence in the example. And even though grammaticality is ensured, poeticness is present only in the form of the poem, i.e. the stanzas that create a poetic impression. Therefore, the problem is not one of context, as Colton et al. claim; the problem is one of meaning. Meaning, as we have seen from the examples so far, remains elusive for AI. As to the specific method, it should be added that the creativity of template-based generation is “open to discussion” (Gervás “Exploring Quantitative Evaluation of the Creativity of Automatic Poets”), as the input is pre-defined and the choices involved are narrow.

(b) **Markov chains.** A Markov chain is a statistical model applied to data in a series, calculating the probability distribution for each entry. The N entries used to make the probability prediction are each referred to as an *n-gram*. For poetry generation, entries would be words. Entries could also be letters but that would produce a multitude of non-words and neologisms, an effect which could be intentional but is not successful in terms of outcome. An example of characters used as entries comes from Roque; he introduces a poetry generator developed by R. W. “Bill” Gosper in the early 1970s at the Artificial Intelligence lab at the Massachusetts Institute of Technology, using character n-gram generation. Gosper approach was later called Dissociated Press. The following example cited by Rogue is a representative output of Dissociated Press on Lewis Carroll’s *Alice in Wonderland*.¹²

¹² A probability model can be generated from a specific work, as in this case, or from a broader corpus.

(...)
 Would guess, an't
 grom one foot to thistle,
 to keep and reachinah'll be she could not ever who had not atte-book
 hastily. Convers began to trings: into thing on wast the door, and
 the.

Involving an automatic shuffling of words, as it does, Markov chain poetry can be said to resemble magnetic poetry¹³ or even Dadaist “cut-ups.”¹⁴ A project that involves Markov chains is EMILY (Shihadeh and Ackerman). EMILY is a machine that aims to create original poems in the style of Emily Dickinson, and it was trained on publicly available Emily Dickinson poetry from the Gutenberg project. With the help of Markov chains, poems in the style of Emily Dickinson were attempted, relying on a dictionary that maps words found in Dickinson’s poems to the words that proceed them to capture the relationships between them. An example of a Dickinson-like poem produced by EMILY would be the following:

Poem 1

Some shook their yellow gown
 And certainly her eye, they
 Leap upon the rose smiling
 To die. The orchards Eternity!

Applying again Manurung’s (*An evolutionary algorithm approach to poetry generation*) preconditions for sound AI-produced poetry, we would say that the result is rather poor. Even though grammaticality is achieved (save for the last line), meaningfulness is not, and poeticness exists only as an impression that borders ridicule. A weak point of Markov chains is that even though they preserve many features of the input text, they cannot replicate grammar.

(c) **Found poetry.** This method involves skipping the generation process altogether and using a computer to harvest text written by humans. Found poetry preserves entire human-written sentences that are presented outside their original context. A splendid example of found poetry are the haikus generated by Hrešková and Machová.

Hrešková and Machová have produced haikus using two different approaches: an evolutionary algorithm and poem models. In the evolutionary algorithm approach,¹⁵ a corpus of haikus was built from two haiku portals containing human-written haikus. Cross-overs, i.e. exchanges, were attempted from two or three existing haikus randomly. This corpus was then evolved by human judges rating every attempt. Examples of such haikus are as follows:

¹³ Magnetic poetry, devised by Dave Kapell, is based on moving around on a metal board words that are written on magnets. See the official website of this poetic movement <https://magneticpoetry.com/>

¹⁴ At a surrealist rally in the 1920s Dada poet Tristan Tzara proposed to create a poem on the spot by pulling words out of a hat. Almost 40 years later, in the summer of 1959, painter and writer Brion Gysin cut newspaper articles into sections and rearranged the sections at random. “Minutes To Go” resulted from this initial cut up experiment. “Minutes To Go” contains unedited unchanged cut ups emerging as quite coherent and meaningful prose. See Burroughs, “The Cut Up Method.”

¹⁵ Evolutionary algorithms are just one of the various methods used in poetry generation. For other approaches, see Lamp et al., “A Taxonomy of Generative Poetry Techniques”; Gonçalo Oliveira, “PoeTryMe: a versatile platform for poetry generation”; Gervás, “Exploring Quantitative Evaluation of the Creativity of Automatic Poets.”

spring wind
 an old lady tries again
 wet with rain

aging beauty
 spring cascades into valleys
 as apple blossoms

The haikus achieved are robust and convey a true sense of “haikuness,” i.e. poeticness according to the norms of the specific genre. Both grammaticality and meaningfulness are also achieved. This is not to say that less meaningful haikus or haikus with confusing content were not created but that these were rather the minority. As enthusiastic as one can be with such a great result, the fact that whole lines of existing haikus were recombined bodes poor for originality or creativity.

But the success of the result makes us understand that, considering mere generation methods, the quality of input is paramount. Many generators take as input content from newspapers (such as Colton et al.), blogs (e.g. Wong and Chun), psychological tests (Netzer et al.), even the Twitter (e.g. the TwitSong, Lamb et al., “Can human assistance improve a computational poet?”). This does allow “using the words of the common man in poetry” (Colton et al. 2012), but it is far from a poetic result. The main reason is that not all words of common man have poetic value. By contrast, when the input is highly relevant, as with the haikus generated by Hrešková and Machová, the result is more in line with a sound poetic production, even though it is poorer in terms of creativity.

What should be preferred, then? An intensely human-driven poetic result, which is less creative in terms of the auxiliary role of AI, or a bad poetic result arrived at by AI alone? It could be argued that poetry generators should be about producing a sound poetic output (otherwise, why *poetry* generators?). But the truth is that it should not be poetic at all costs. The real issue of AI poetry generation is that it should be intelligent; in other words, the computer should be able to actually do it on its own. Thus, the process by which a result is achieved is as important as the result itself. Of course, the result is ultimately the measure by which the success of the attempt is rated. And the result often is, as we have seen, poor, especially so in terms of meaningfulness.

Enhancement of poetry generation results

Moving away from mere generation, Lamb et al. (“A Taxonomy of Generative Poetry Techniques”) next tackle the issue of enhancement with regard to poetic results. Enhancement may be conducted either by a person or by the computer itself.

Human enhancement

As is obvious, a human can enhance the output of poetry generators, editing the poems themselves. Although this might ultimately render the generator useless in the first place, it is undeniably an established practice. It works marvelously for those who see computational text generation as nothing more than a good starting point for writing poetry, equal to the provision of raw material.

The practice of human enhancement is evident in Gnoetry (*Gnoetry Daily*), an application for interactive text generation in which some decisions are made by the computer and others by a human user. Generated poetry is based on n-grams, but it is the user that provides the corpus. Thus, a dialogue is established between the human and the computer (Lamb et al. “A Taxonomy of Generative Poetry Techniques”). An example from Gnoetry (*Gnoetry Daily*) is found below:¹⁶

The day stirs the bare limbs above me,
dry breath from the frozen fields.

Inside,

charred, still smoldering, I sit patiently

with my despair, having done too little,
for too long, to fan these flames, having been

content enough, secure enough in my same

fearful ways, narrow, shut in, obscure
to myself.

Computer enhancement

Poetry generation results can also be enhanced by the computer itself. As pointed out by Lamb et al. (“A Taxonomy of Generative Poetry Techniques”), “this set of methods comes not from the humanities but from scientists in the discipline of computational creativity” (198) and involves two main purposes. One is “optimization of the system’s output on some metric,” and the other is connection to underlying knowledge about the real world.

Beginning with the latter, Lamb et al. (“A Taxonomy of Generative Poetry Techniques”) acknowledge that the nonsensical results of mere poetry generation are the “result of the computer’s lack of real-world experience” (198). It is the problem with the phrase “She has an apple,” as described above. The computer may be able to say *which word is what kind of thing* in terms of grammar or its place in a syntactic tree, but it cannot understand what this phrase is really about.

This is a serious issue and the real reason behind the production of nonsense, yet this can be solved to some degree by the provision of knowledge bases. In poetry generation, a knowledge base programmed into the generation process may guide its output. This is done in much the same way as e.g. in Gnoetry (*Gnoetry Daily*), where human input inside the generation process optimizes results. In particular, a knowledge base built on the calculation of the co-occurrence of different words in a source text can give the computer a sense of which words go better together or are more relevant to a topic. Knowledge bases may also be extracted from dictionaries and cover relations between words, such as junction, composition, opposition, causation, etc. (Gonçalo Oliveira, “PoeTryMe: a versatile platform for poetry generation”). Netzer et al. have produced quite good haikus from a list of word associations that are combined on the basis of syntactic templates:

¹⁶ More can be retrieved at <https://gnoetrydaily.wordpress.com/>

early dew
 the water contains
 teaspoons of honey

spring bloom
 showing
 the sun's pyre

Regarding optimization, a computer may be fed with some formal definition of the desired characteristics in a poem and can thus test different possibilities in order to choose the ones that best fit such characteristics. These are baby steps to critical thinking, but it is a start. Lamb et al. (“Can human assistance improve a computational poet?”) list several optimization methods, some of which are as follows: stochastic hill-climbing search, generate-and-test, case-based reasoning and recurrent neural networks. Many such methods concentrate on issues of poetic form (e.g. rhyme and meter), with fewer on meaning, style, or emotions. TwitSong (Lamb et al. (“Can human assistance improve a computational poet?”) browses tweets to choose those that are relevant to a topic and produce sentiment. The process generates poems that are composed of tweets which, although written by different people at different occasions, make up something reasonably meaningful, even though its poetic value could be strongly challenged:

Hey Nashville... 2014 is pretty awesome!
 Happy 2014 friends! Be safe out there!!
 Had a great New Year's Eve at Magic Kingdom
 We started off 2014 with a prayer.

Lamb et al. (“Can human assistance improve a computational poet?”) conclude that poetry generation is still at an infancy stage and that there is a lot to be expected. Hopefully, that will be the case. Why hopefully? Because, it could lead to a new style or a new content of poetry, which would be simply refreshing.

Poetry generation in terms of target features

Lamb et al. (“A Taxonomy of Generative Poetry Techniques”) are not the only ones with a taxonomy of poetry generation systems; other scholars have too provided their own taxonomies. For example, Gervás (“Exploring Quantitative Evaluation of the Creativity of Automatic Poets”) has presented four types of AI techniques employed in the generation of poetry and Gonçalo Oliveira (“A Survey on Intelligent Poetry Generation: Languages, Features, Techniques, Reutilisation and Evaluation”) has tackled the point from a different perspective, allowing a closer glimpse at the workings of generators. More particularly, Gonçalo Oliveira makes a useful distinction between methods used to address the form features considered by poetry generators and methods used to address content features, both of which the following section will explore.

Form features

Form features play a prominent role in poetry generation, as they allow recognition at first glance that the output is, in fact, a poem. Meter and rhyme are such form features that can be quite competently handled by computer programs. Meter involves the number of syllables that each line of the poem must have, taking into consideration stress patterns. Haikus, for example, have three lines that include 5-7-5 syllables, respectively (as

is perceived in the West, at least), and limericks typically have 5 lines. Moreover, sometimes stress patterns are also important.

Rhyme is the repetition of the same sound in different lines, according to the appropriate scheme, e.g. AAAA (monorhyme), AABBC (coupled rhyme), ABAB (alternate rhyme), ABBA (enclosed rhyme), AABBA (limerick), or ABA BCB CDC etc. (*terza rima*, or chain rhyme). Some systems deal with other kinds of rhyming apart from end-rhyme, e.g. assonance or alliteration. These involve the repetition of the same vowel or consonant sound, respectively, in some parts of the poem or throughout it.

According to Gonalo Oliveira, for “less phonetic languages such as Portuguese or Spanish [...] it is often enough to design a set of orthography-based rules to handle meter and rhyme” (Gonalo Oliveira et al., “A Survey on Intelligent Poetry Generation”; Gervás, “An expert system for the composition of formal Spanish poetry”). For English, poetry generators typically resort to pronunciation dictionaries, often online ones (see e.g. Manurung, “A chart generator for rhythm patterned text”; Tobing and Manurung, “A chart generation system for topical metrical poetry”).

Content features

Even if form features are more easily tackled with AI, the issue of content remains paramount. No text can be considered a poem if it just follows a recognizable poetry form. It must contain language that, according to Manurung (*An evolutionary algorithm approach to poetry generation*), ensures the central three elements we have already discussed: (i) grammaticality, (ii) meaningfulness and (iii) poeticness.

Out of these three, grammaticality seems to be the most easily achievable one; after all, natural language generation does involve syntactic rules, and lexical-syntactic patterns can also be achieved. Meaningfulness, on the other hand, is more elusive. Meaning is only loosely achieved in AI by using, for instance, words that belong to the same semantic domain (e.g. fruit-apple-orange-mango and other hypernymic-hyponymic relations, or synonyms/antonyms etc.) or by creating semantic clusters of words that are connected to each other in various ways (for instance, successful combinations of adjective and noun, verb and noun). Such techniques may allow for the generation of other natural language texts, such as weather forecasts. For a weather forecast to be generated, there must be a computer that is able to recognize weather phenomena, e.g. storms, precipitation, hurricanes, and their movements. Putting all these into a forecast would then require the computer to be able to name the weather phenomena, select interesting pieces of information while rejecting other not-so-interesting ones, and decide on a good textual structure, among other things (Wright and Purver).

In poetry generation, meaning is attempted by exploiting a semantics model, a semantic knowledge base, or a statistical model of distributional semantics (Gonalo Oliveira “Multilanguage extension and evaluation”). Semantics models are based on labelled relations between words and reflect a more theoretical view of linguistic knowledge. Knowledge bases we have already discussed above in “Computer enhancement.” Distributional models of semantics consider how language is actually used, e.g. in a collection of texts, and assume that words occurring in similar contexts have similar meanings. In other systems, text is generated according to a grammar that handles syntax (Manurung *An evolutionary algorithm approach to poetry generation*) or a grammar that is tightly related to semantics (Gonalo Oliveira, “PoeTryMe: a versatile platform for poetry generation”).

Mere text generation, however, simply does not suffice. The poem has to be meaningful, i.e., it must be about something, otherwise it must at least be able to respond to various stimuli ((Gonalo Oliveira, “A Survey on Intelligent Poetry Generation”). Such stimuli could be any of the following: a list of semantic predicates, e.g.

love (John, Mary) (Manurung *An evolutionary algorithm approach to poetry generation*); one or more keywords that will set a semantic domain (e.g. Netzer et al.; Gonçalo Oliveira, “PoeTryMe: a versatile platform for poetry generation”); a text which can be a single sentence with a message (Gervás “An expert system for the composition of formal Spanish poetry”) such as a chunk taken from a newspaper (Colton et al.).

As for poeticness, figurative language is often implicitly present as a consequence of using human-produced poetry as input. We have seen the apt example of Hřešková and Machová’s haikus. There is, however, ongoing research in the field of figure-of-speech generation, e.g. regarding the generation of similes (*as-as* or *like*) mined by using Google n-grams (Veale) or the dynamic generation of similes from scratch (Chakrabarty, Muresan & Peng).¹⁷ Such similes are then used for the generation of other similes or metaphors or yet other figures of speech. But in real poetry, using figures of speech alone does not suffice; a supporting poetic context should also be in place.

Human or Not

Poetry generators are not the only “poets” that produce strange results. Like the computer scientists discussed earlier who were more interested in process than product, several modern poetic movements have favored experimentation over a result that would fit mainstream poetry standards, challenging even the idea of them. This certainly makes poetry generation look like another link in the long chain of poetic tradition, as will become clearer at the end of this section. Some of the most striking such movements, involving algorithmic or algorithmic-like processes, are mentioned below.

Dada

As early as in 1916, Dada made its appearance as a poetic movement of the European avant-garde. Coming into being around of World War I, its founders tried to “embody a simultaneously playful and nihilistic spirit alive among European visual artists and writers” (“Dada”). They attempted to salvage a sense of freedom from the instability that characterized European society in the aftermath of the war, embracing both “everything and nothing” in their desire to “sweep, sweep clean,” as Tristan Tzara wrote in the Dada Manifesto in 1918. These Dada poems were made of cut-ups, combined in the form of collage to produce a syntactic result that encouraged readers to question what they were reading and make up their meanings for the poems’ words.

The following is a representative Dada poem entitled “Assessment” from Three Poems by Tristan Tzara:

¹⁷ The achievement of Chakrabarty et al. makes us realize that the real benefit lies in seeing through human brain processes and trying to reproduce them, as if reinventing them. This would necessarily also make AI another way of understanding human intelligence better.

ASSESSMENT

voltic arc of these two nerves that don't touch
 near the heart
 we note the black shivers under a lens
 is this feeling this white spouting
 and methodical love
 splits my body into rags
 toothpaste pastry
transatlantic
tickets
the crowds crash the column conched in wind
range of rockets
 on my head
the bloody revenge of the liberated two-step
directory of determinations at prix fixe
 folly at 3:20 am
 or **5.50 francs**
 cocaine slowly gnaws the walls for its pleasure
satanic horoscope dilates under your vigor
 VIRGIL'S VIGILANCE VERIFIES THE VIRAL WIND
eyes through once more

 **TRISTAN TZARA**

Considering that the origins of this poem are from different sources, as is evident by the different styles of the cut-ups, one could argue that — but for the stylistics — a poem generator that uses human-written texts as input seems to walk on a very similar path to Dada.

Stream of consciousness

Stream of consciousness is another notable example of literature written without much concern for the outcome. The term “stream of consciousness” was coined by the psychologist William James in *The Principles of Psychology* (1890). This style intended to render “the flow of impressions — visual, auditory, physical, associative, and subliminal — that impinge on the consciousness of an individual and form part of his awareness along with the trend of his rational thoughts.”¹⁸ This technique was adopted by many poets, one of which is notably Gertrude Stein. Here follow two poems of hers:

Red Faces

Red flags the reason for pretty flags.
 And ribbons.
 Ribbons of flags
 And wearing material
 Reason for wearing material.
 Give pleasure.
 Can you give me the regions.
 The regions and the land.
 The regions and wheels.
 All wheels are perfect.
 Enthusiasm.¹⁹

America

Once in English they said America. Was it English to them.

¹⁸ <https://www.britannica.com/art/stream-of-consciousness>

¹⁹ <https://allpoetry.com/Red-Faces>

Once they said Belgian.
 We like a fog.
 Do you for weather.
 Are we brave.
 Are we true.
 Have we the national colour.
 Can we stand ditches.
 Can we mean well.
 Do we talk together.²⁰

Both poems are asyntactic to some extent and contain a flow of ideas that is not easy to follow. They resemble very much the current output of poetry generators.

OuLiPo

OuLiPo, another poetic movement mentioned in “Digital Humanities versus Computational Creativity,” emphasized systematic, self-restricting processes of making texts that resemble algorithms. For example, the technique known as n+7 replaced every noun in an existing text with the noun that followed seven entries after it in the dictionary. By applying the n+7 rule to Wallace Stevens’s poem “(Stephens), for example, one gets a completely new poem, “The Soap Mandible” (“A Brief Guide to OULIPO”):

The Soap Mandible

One must have a miniature of wisdom
 To regard the fruit and the boulders
 Of the pinions crusted with soap;

And have been colic a long time
 To behold the junkyards shagged with Idaho,
 The spun-yarn rough in the distant gloom

Of January surgery; and not to think
 Of any mishap in the south of the winter,
 In the south of a few lectures,

Which is the south of the language
 Full of the same winter
 That is blowing in the same bare plague

For the lithographer, who listens in the soap,
 And, now himself, beholds
 Now that is not thermal and the now that is.

²⁰ <https://allpoetry.com/poem/8529995-America-by-Gertrude-Stein>

Another OuLiPo exercise uses the “snowball” technique, where the first line is one letter long, the second line has two letters, and so on. In the 15-line snowball poem that follows by Harry Mathews, each line is $n+1$ character longer than the previous:

I
 am
 the
 text
 which
 begins
 sparely,
 assuming
 magnitude
 constantly,
 perceptibly
 proportional,
 incorporating
 unquestionable
 incrementations

As becomes evident from the above samples, Oulipians were more interested in experimentation than in the outcome, as is often the case with art. Why should AI generated poetry be any different?

Language Poetry

Language Poetry, or LangPo, is yet another example of poetry that could have been written by a poetry generator. Among the key beliefs of LangPo lies the idea that language is what dictates meaning. LangPo poets also intend to make the reader actively involved in the text: “by breaking up poetic language, the poet[s are] requiring the reader to find a new way to approach the text” (“A Brief Guide to Language Poetry”). An example of a LangPo poem is the following, an excerpt from “Letter 7” written by Michael Palmer (Palmer):

Letter 7 (excerpt)

But the buried walls and our mouths of fragments,
no us but the snow staring at us . . .

And you Mr. Ground-of_what, Mr. Text, Mr. Is-Was,
 can you calculate the ratio between wire and window,

between tone and row, copula and carnival
 and can you reassemble light from the future-past

in its parabolic nest
 or recite an entire winter’s words,

its liberties and psuedo-elegies,

the shell of a street-car in mid-turn

or scattered fires in the great hall
 I would say not-I here I'd say *The Book of Knots*

I'd say undertows and currents and waterspouts,
 streaks of phosphorus and rivervine winds

Dear Z, I'd say it's time, it's nearly time, it's almost, it's
 just about, it's long
 past time now time now for the vex- for the vox- for the
 voices of shadows,

time for the prism letters, trinkets and shrouds,
 for a whirl in gauzy scarves around the wrecked piazza

Flarf

A 21st century poetry movement reflecting similar themes is Flarf, which thrives on language from the web to arrive at odd juxtapositions and grammatical inaccuracies. Flarf also celebrates deliberately bad or “incorrect” poetry by forcing clichés, swear words, onomatopoeia, and other linguistic aberrations into poetic shape (“Flarf”). Flarf poets mine such odd terms from Google to arrive at often hilarious and sometimes disturbing poems, plays, and other texts (a technique called “Google sculpting”). The extensive use of internet resources is much akin to most poetry generators. An example of Flarf poetry is the following poem by Rodney Koenke (Koenke):

bug's psalm

The bug's psalm: don't get crushed.
 Afterlives feel meaningless
 but spring will come,
 push out the nubs
 the kids braid into pallets.
 Take up your pallet
 from lawns noon's hardly touched.
 The small think gods
 just loll on clouds.
 Bugs think gods just crush.


The incoherence of this poem makes one think that it could actually have been written by a computer, but it is again an example of human-written poetry for the sake of experimentation.

Blackout or Erasure Poetry

Last but not least, blackout or erasure poetry can be said to resemble the workings of poetry generators, erasing as it does all the words in a given text save for those that it chooses to create a poem. It is therefore a

process of narrowing down the words from a given text (newspaper articles are used very often; however, most any text will do) to a few ones according to what the poet sees fit (Glatch). An example of blackout poetry — rather bad, but helpful for the sake of the argument — would be the following:

great success
 against time is
 You look
 on the outside
 like you feel



Miller, E. Ce. “Blackout Poetry Is A Fascinating Art Form You Can Try At Home Right Now.” August 28, 2017. <https://www.bustle.com/p/what-is-blackout-poetry-these-fascinating-poems-are-created-from-existing-art-78781>

All the foregone examples help to show that sometimes people, too, write poetry as if they were computer generators, not minding incoherence, elliptical meaning, or even nonsense being produced. This fact highlights the elusiveness of poetry and creativity; not only it cannot be solidly defined, but it also means different things to different people or groups of people over time. What’s more, AI attempts at poetry generation are not so different from similar human attempts. This means that, in a way, AI attempts are actually worth a place in the long tradition of poetry as it evolves.

Human or Bot

Having said all of the above, is it even possible to understand nowadays if a poem has been written by a human being or a computer? The answer is that sometimes, indeed, it is not. This is because some AI attempts have, in fact, proven to be of very good quality — and because human beings have written poems that can closely mimic computer attempts. So, if there are computers that write as computers, computers that write as human beings, human beings that write as human beings, and human beings that write as computers, “what on earth is going on”? This question was posed in a TedX speech by Oscar Schwartz, a writer, poet, and creator

of an online site on AI versus human-written poetry (Schwartz).²¹ Schwartz further talks about the capacity of being human and what can define it. Is it something with which we are born, or something that we are trained to become? And how do we define who — or what — is part of it? Generalizing therefore on the concept of AI and not restricting it to poetry, let us consider it, for a while, in relation to humanness.

In 2017, six scientists wrote and signed a manifesto in which they also wonder: Is the capacity of being human something that can be lost or won? (Barzov, “Manifesto: Time for Humanness Learning”). Are those who torture, kill, humiliate human existence still considered human? Is a person born in a western country more human than an indigenous in, say, New Zealand? In the end, is human nature a social, thus changing, construct? And they end up saying: “It’s time for us to come out and start building a community of AI scientists, neuroscientists, philologists, game developers, historians of literature, movie producers, writers, entrepreneurs — all those who are united by the passion to discover the first principle of humanity and —furthermore— by the genuine desire to provide humans and machines with compelling and efficient tools of learning and relearning to be human” (Barzov, “Manifesto: Time for Humanness Learning”).

But how can computers and AI “learn to be human” if they are not? The idea of granting “personhood” to AI has already become a point of scrutiny. To begin with, it has been the case for non-humans to be considered persons in certain countries. In the USA, corporations like the South Pacific Railway have been granted the rights of free speech and freedom of religion since 1881 (Winkler). In other countries, like New Zealand or Mexico, natural sites such as the Te Urewera Forest or Mount Taylor, respectively, have person-like rights (Colwell). On the other hand, there are law scholars today who have shown that anyone can assign legal personhood to a computer by having it run a limited liability corporation in the USA. In this way, it could own property, hire lawyers, sue, and enjoy protections provided under law as if it were a regular human being. And what about the robot Sophia to which Saudi Arabia granted citizenship, causing objections from Saudi women who noted that the robot was given more rights than them? (Yampolskiy)?

Granting personhood to AI is an issue that has occupied a number of law scholars, also involving scholars from other fields like ethics (Garcia). Gunkel and Wales have recently looked at it from two different standpoints, one involving philosophy and law (Gunkel) and another involving theology and psychology (Wales). Gunkel (Gunkel and Wales) draw a useful division between natural and legal persons; they state that natural persons are those who possess the capability of speech, have consciousness, and are able to experience pleasure and pain, whereas legal persons are entities defined as such not due to their innate qualities but rather following a legal decision. Then, Gunkel cites “human beings of diminished intellectual capacity or in a persistent non-responsive state” (475) who are not full legal persons. This makes us understand that, although there are non-human entities that are considered legal persons as we saw above, there are also humans who, by contrast, are not considered to be such. One can also note at this point the suggestion of the European Parliament for AIs to be considered “electronic persons,” which was reflexively followed by an open letter signed by 250 scientists, engineers, and AI professionals opposing the suggestion and asserting that robots and AI are nothing more than tools (Open Letter).

Wales (Gunkel and Wales), on the other hand, tackles the issue of the interior life of human beings, who engage in “voluntary self-gift” as they meet other human beings’ interiorities in a fusion of minds through

²¹ That site was Bot or Not (botpoet.com). The user had to decide whether the poem that he saw on his screen was written by a computer or a human. It operated, in other words, on the basis of “Turing tests.” A Turing test, named after the computer scientist Alan Turing, is intended to show whether a machine can convince someone that it is human. Sadly, this site is no longer live.

empathy and conscious understanding. Artificial neural networks, as he points out, are just biologically inspired simulations without any such characteristics or potential. AI, Wales argues, “is a sedimentary reflection of our own personhood. It is artificial in the original sense of that word – an artifact, a work of skill that we have brought forth by gazing into a computational pool of Narcissus” (480). In other words, even if we create computers with human form or capacities, they will remain mirrors of human beings, who — among other things — feature self-awareness and a sense of identity. These are qualities that something that a machine, which is programmed just to execute orders, cannot ever possess. Transhumanist thinkers have long pondered such issues (Bostrom, “A history of transhumanist thought”). Is it, then, ever imaginable for a mirror to reach or overcome that which it mirrors? This remains to be seen, literally.

All these discussions and ideas are, of course, far-fetched when it comes to poetry generators that employ human-driven imaginative combinations of human-created examples to come to a result. Computational Creativity, at present, remains useful in helping humans with their own creative endeavors, even though its attempts at imitating human creativity have yielded interesting and at the same time ineffectual results. It does not possess imagination, however, nor inspiration of its own, and any talks about personhood in this case would only seem irrelevant (Barzov, “Human Intelligence (HI) is underrated”). The real value of Computational Creativity perhaps rests in what it can reveal about human creativity and cognition (Carleton University).

But what about the future? According to Dormehl, “Computational Creativity has been ignored. Right now, either fondly or maliciously, it is being laughed at. Next it will start fighting our preconceptions.” It is true that when AI creates something, we have the tendency to credit its programmer alone. After all, machine creativity today typically involves humans making at least some of the decisions. However, the boundary between programmer and computer is becoming increasingly fluid and indistinguishable, as many of the computer methods of today produce results with production details that the programmer has no clue how to explain. The scholar Matthias corroborates this when he says that the programmer of a neural network, for instance, has little control past the point of coding. Soon, perhaps, the credit for any AI achievement will be split between humans and machines in much the same way as with films written by two scriptwriters, even if this will, in fact, mean that AI has to acquire legal personhood.

Such a development will seem less and less bizarre as AI grows more and more capable of performing the duties assigned to it. After all, philosopher Immanuel Kant held that moral concern should extend equally to all rational beings, including ones that may appear from other planets. According to Boddington, Kant might also have added that it could apply to AI, if AI had had then the widespread presence it has today (13).

It is also possible that the train of thoughts we currently follow is distracting us from perceiving the real opportunities and challenges which confront us in the age of AI (Gunkel and Wales). Maybe what we need is a more refined ontology that does not see the world in terms of black and white. And maybe it is Computational Creativity that will show us the way to it, which would make it truly invaluable.

Works Cited

“A Brief Guide to Language Poetry.” *Poets.org*, Academy of American Poets, 2004, <https://poets.org/text/brief-guide-language-poetry>. Accessed 22 July, 2021.

“A Brief Guide to OULIPO.” *Poets.org*, Academy of American Poets, 2004, <https://poets.org/text/brief-guide-ouliipo>. Accessed 22 July, 2021.

- Acrostic Poem Generator*, <http://www.acrosticpoem.org/>. Accessed 22 July, 2021.
- Barthes, Roland. "The Death of the Author." Translated by Richard Howard, *Aspen*, no. 5-6, 1967, www.ubu.com/aspen/aspen5and6/threeEssays.html. Accessed 22 July, 2021.
- Barzov, Yuri. "Human Intelligence (HI) is underrated." *Medium*, May 25, 2018, medium.com/@yuribarzov/humanness-is-underrated-d2b745124dd4. Accessed 22 July, 2021.
- Barzov, Yuri, et al. "Manifesto: Time for Humanness Learning." *Medium*, March 22, 2017, towardsdatascience.com/manifeto-time-for-humanness-learning-51051befa5f4. Accessed 22 July, 2021.
- Berkman, Natalie. "Digital Oulipo: Programming Potential Literature." *Digital Humanities Quarterly* vol. 11, no. 3, 2017. <http://www.digitalhumanities.org/dhq/vol/11/3/000325/000325.html>. Accessed 22 July, 2021.
- Boddington, Paula. *Towards a Code of Ethics for Artificial Intelligence*. Springer International Publishing, 2017. Print.
- Bostrom, Nick and Yudkowsky, Eliezer. "The ethics of Artificial Intelligence." *Cambridge Handbook of Artificial Intelligence*. Eds. William Ramsey and Keith Frankish. Cambridge University Press, 2011. Print.
- Bostrom, Nick. "A history of transhumanist thought." *Journal of Evolution & Technology*, vol. 14, no. 1, 2005. www.nickbostrom.com/papers/history.pdf. Accessed 22 July, 2021.
- Burroughs, William. "The Cut Up Method." *The Moderns: An Anthology of New Writing in America*. Ed. Leroi Jones. New York: Corinth Books, 1963. <https://www.writing.upenn.edu/~afilreis/88v/burroughs-cutup.html>. Accessed 22 July, 2021.
- Carleton University, DIGH500 cohort. *Digital Humanities: A Primer*, chapter 10, 2018-2019, carletonu.pressbooks.pub/digh5000/. Accessed 22 July, 2021.
- Chakrabarty, Tuhin, et al. "Generating similes effortlessly *like a Pro*: A Style Transfer Approach for Simile Generation." *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*. Eds. Bonnie Webber, et al. Association for Computational Linguistics, 2020. 6455-6469. Accessed 22 July 2021.
- Colton, Simon, et al. "Full-FACE poetry generation." *Proceedings of 3rd International Conference on Computational Creativity (ICCC 2012)*, Dublin, Ireland, University College Dublin, 2012. 95-102. <https://www.computationalcreativity.net/proceedings/ICCC-2012-Proceedings.pdf>. Accessed 22 July 2021.
- Colwell, Chip. "What If Nature, like Corporations, Had the Rights and Protections of a Person?" *The Conversation*, 8 Oct. 2021, <https://theconversation.com/what-if-nature-like-corporations-had-the-rights-and-protections-of-a-person-64947>. Accessed 22 July 2021.
- Culler, Jonathan. "Semiotics as a Theory of Reading." *The Pursuit of Signs: Semiotics, Literature, Deconstruction*, Cornell University Press, 1981. 47-51. Print.
- "Dada." *Poetry Foundation*, Poetry Foundation, <https://www.poetryfoundation.org/learn/glossary-terms/dada>. Accessed 22 July 2021.
- Delvin, Es. *Poem Portraits*, Google Arts & Culture, <https://artsexperiments.withgoogle.com/poemportraits>. Accessed 22 July 2021.

- Dormehl, Luke. “Truly creative A.I. is just around the corner. Here’s why that’s a big deal.” *Digital Trends Media Group*, January 5, 2018, www.digitaltrends.com/cool-tech/artificial-intelligence-creativity-future/. Accessed 22 July 2021.
- Eliot, T. S. (Thomas Stearns). *The Sacred Wood: Essays on Poetry and Criticism*, London: Methuen, 1928. Print.
- European Parliament Research Service. *The ethics of artificial intelligence: Issues and initiatives*. European Parliament, 2020, [https://www.europarl.europa.eu/RegData/etudes/STUD/2020/634452/EPRS_STU\(2020\)634452_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2020/634452/EPRS_STU(2020)634452_EN.pdf). Accessed 22 July 2021.
- “Flarf.” *Poetry Foundation*, Poetry Foundation, <https://www.poetryfoundation.org/learn/glossary-terms/flarf>. Accessed 22 July 2021.
- Garcia, Robert. “Artificial Intelligence and Personhood.” *Cutting Edge Bioethics: A Christian Exploration of Technology and Trends*. Eds. John Kilner et al. Grand Rapids, MI: W.B. Eerdmans, 2002. Print.
- Gervás, Pablo. “Exploring Quantitative Evaluation of the Creativity of Automatic Poets.” *Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems*. Eds. Tony Veale et al. Cham: Springer, 2019. 275-304. Internet resource.
- Gervás, Pablo. “Generating poetry from a prose text. Creativity versus faithfulness.” 2001a, www.researchgate.net/publication/2925641_Generating_Poetry_from_a_Prose_Text_Creativity_versus_Faithfulness. Accessed 22 July 2021.
- Gervás, Pablo. “An expert system for the composition of formal Spanish poetry.” *Journal of Knowledge-Based Systems*, vol. 14, no. 3-4, 2001. 181-188. <https://www.sciencedirect.com/science/article/abs/pii/S0950705101000958>. Accessed 22 July 2021.
- Glatch, Sean, et al. “What Is Blackout Poetry? Examples and Inspiration.” *Writers.com*, 11 June 2021, <https://writers.com/what-is-blackout-poetry-examples-and-inspiration>. Accessed 22 July 2021.
- Gnoetry Daily*, <https://gnoetrydaily.wordpress.com/>. Accessed 22 July 2021.
- Gonçalo Oliveira, Hugo. “A Survey on Intelligent Poetry Generation: Languages, Features, Techniques, Reutilisation and Evaluation.” *Proceedings of the 10th International Natural Language Generation conference, INLG 201* (Santiago de Compostela, Spain), Association for Computational Linguistics, 2017. 11-20. <https://aclanthology.org/W17-3502/> Accessed 22 July 2021.
- Gonçalo Oliveira, Hugo, et al. “Multilanguage extension and evaluation of a poetry generator.” *Natural Language Engineering*, 2017, nil.fdi.ucm.es/sites/default/files/multilingual_poetry.pdf. Accessed 22 July 2021.
- Gonçalo Oliveira, Hugo. “Tra-la-Lyrics 2.0: Automatic Generation of Song Lyrics on a Semantic Domain.” *Journal of Artificial General Intelligence*, vol. 6, no. 1, 2015. https://www.researchgate.net/publication/291386243_Tra-la-Lyrics_20_Automatic_Generation_of_Song_Lyrics_on_a_Semantic_Domain. Accessed 22 July 2021.
- Gonçalo Oliveira, Hugo. “Poetryme.” *PoeTryMe*, Center of Informatics and Systems of the University of Coimbra, 2015, <https://poetryme.dei.uc.pt/>. Accessed 22 July 2021.

- Gonçalo Oliveira, Hugo. "PoeTryMe: a versatile platform for poetry generation." *Proceedings of ECAI 2012 Workshop on Computational Creativity, Concept Invention, and General Intelligence*, Montpellier, France, C3GI 2012, 2012. https://eden.dei.uc.pt/~hroliv/pubs/GoncaloOliveira2012_c3gi_CRC.pdf. Accessed 22 July 2021.
- Gonçalo Oliveira, Hugo, et al. "An Approach to Generate Text based on Rhythm." *New Trends in Artificial Intelligence*. Eds. José Neves et al. Guimarães, Portugal: Associação Portuguesa para a Inteligência Artificial, 2007. 57-68.
- "Google's AI Writes Poetry after Reading Thousands of Romance Novels." *Writers Write*, Writers Write, Inc., 19 May 2016, <https://www.writerswrite.com/google-ai-writes-poetry-51920161>. Accessed 22 July 2021.
- Gosper, R. W. "Item 176." *HAKMEM*, Memo 239. Eds. Michael Beeler et al. Cambridge: Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1972. Internet resource.
- Gunkel, David and Wales, Jordan Joseph. "Debate: What is personhood in the age of AI?" *AI & Society*, vol. 36, 2021: 473-486. Print.
- "Haiku Generator." *FantasyNameGenerators.com*, <https://www.fantasynamegenerators.com/haiku-generator.php>. Accessed 22 July 2021.
- Harel, David and Rumpe, Bernhard. "Modelling Languages: Syntax, Semantics and All That Stuff – Part I: The Basic Stuff." 2003, https://www.researchgate.net/publication/2877092_Modeling_Languages_Syntax_Semantics_and_All_That_Stuff_Part_I_The_Basic_Stuff. Accessed 22 July 2021.
- Holland, Norman N. "Reading and Identity: A Psychoanalytic Revolution." *Academy Forum*, American Academy of Psychoanalysis, vol. 23, 1979: 7-9. Print.
- Hrešková, Miroslava and Machová, Kristína. "Haiku poetry generation using interactive evolution vs. poem models." *Acta Electrotechnica et Inforomatica*, vol. 17, no. 1, 2017: 10-16. DOI:10.15546/aeeci-2017-0002
- Henrickson, Leah. *humanities computing, digital humanities, and computational humanities: what's in a name*. 3:AM Magazine, October 24, 2019. <https://www.3ammagazine.com/3am/humanities-computing-digital-humanities-and-computational-humanities-whats-in-a-name/>. Accessed 22 July 2021.
- Iser, Wolfgang. "Indeterminacy and the Reader's Response in Prose Fiction." *Aspects of Narrative* (English Institute Essays). Ed. Joseph Hillis Miller. New York : Columbia University Press, 1971. 1-45. Print.
- James, William. *The Principles of Psychology*. New York: Henry Holt and Company, 1890. Print.
- Janes, Tyler L. "Legal personhood for Artificial Intelligence: Citizenship as the exception to the rule." *AI and Society*, vol. 35, no. 2, 2020: 343-354. Print.
- Jauss, Hans R. "Literary History as a Challenge to Literary Theory." *Toward an Aesthetic of Reception*. Transl. Timoty Bahti. Minneapolis: University of Minnesota Press, 1967/1982. 3-45. Print.
- Kant, I. *The Moral Law, translation of groundwork for the metaphysics of morals*. London : Hutchinson University Library, 1972. Print.
- Kapell, Dave. Magnetic Poetry: Magnets that Reveal the Poet Within. <https://magneticpoetry.com/>. Accessed 22 July 2021.

- Kapell, Dave, and Sally Steenland. *The Magnetic Poetry Book of Poetry*. New York: Workman Pub, 1997. Print.
- Koeneker, Rodney. "bug's psalm." <https://www.poetryoutloud.org/poem/bugs-psalm/>. Accessed 22 July 2021.
- Lamb, Carolyn, et al. "A Taxonomy of Generative Poetry Techniques." In *BRIDGES Finland Conference Proceedings*, 2017. 195-202. <https://archive.bridgesmathart.org/2016/bridges2016-195.pdf>. Accessed 22 July 2021.
- Lamb, Carolyn, et al. "Can human assistance improve a computational poet?" In *BRIDGES Finland Conference Proceedings*, 2015. 37-44. <http://t.archive.bridgesmathart.org/2015/bridges2015-37.pdf>. Accessed 22 July 2021.
- Lau, Jey Han et al. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers). Melbourne: Association for Computational Linguistics, 2018. <https://aclanthology.org/P18-1181> and <https://aclanthology.org/P18-1181.pdf>. Accessed 22 July 2021.
- Laws, Amy E. *Rimbaudelaire*. ELMCIP, <https://elmcip.net/creative-work/rimbaudelaire>. Accessed 22 July 2021.
- Levy, Robert P. "Rplevy/Poevolve: A Naive Model of Human Creativity / Dadaistic Song-Writing Code." *POEVOLVE*, GitHub, <https://github.com/rplevy/poevolve>. Accessed 22 July 2021.
- Lutz, Theo. "Stochastische Texte." *augenblick* 4 (1959), H. 1, S. 3-9. https://www.stuttgarter-schule.de/lutz_schule_en.htm. Accessed 22 July 2021.
- Mathews, Harry. 15-line snowball poem. <https://poetrywithmathematics.blogspot.com/2010/08/snowballs-growingshrinking-lines.html>. Accessed 22 July 2021.
- Manurung, Hisar. *An evolutionary algorithm approach to poetry generation*. Ph.D. Dissertation. University of Edinburgh, 2004. Computer file.
- Manurung, Hisar. "A chart generator for rhythm patterned text." *Proceedings of the First International Workshop on Literature in Cognition and Computer (Tokyo)*, 1999. Internet resource.
- Matthias, Andreas. "The responsibility gap: Ascribing responsibility for the actions of learning automata." *Ethics and Information Technology*, vol. 6, no. 3, 2004: 175-183. <https://doi.org.libproxy.uoregon.edu/10.1007/s10676-004-3422-1>. Accessed 22 July 2021.
- Medvedev, Pavel Nikolaevich and Bakhtin, Mikhail Mikhailovich. "The Object, Tasks and Methods of Literary History." *The Formal Method in Literary Scholarship: A Critical Introduction to Sociological Poetics*. Transl. A.J. Wehrle. Baltimore: Johns Hopkins University Press, 1978. 26-37. Print.
- Miller, E. Ce. "Blackout Poetry Is A Fascinating Art Form You Can Try At Home Right Now." August 28, 2017, <https://www.bustle.com/p/what-is-blackout-poetry-these-fascinating-poems-are-created-from-existing-art-78781>. Accessed 22 July 2021.
- Netzer, Yael, et al. "Gaiku: Generating haiku with word associations norms." *Proceedings of the Workshop on Computational Approaches to Linguistic Creativity*, Association of Computational Linguistics, 2009. 32-39. Internet resource.
- Open Letter to the European Commission Artificial Intelligence and Robotics, <http://www.robotics-openletter.eu/>. Accessed 22 July 2021.

- Oulipo (Ouvroir de Littérature Potentielle). *Atlas de littérature potentielle*, vol. 1 in *Collection Idées*. Paris: Gallimard, 1981. Print.
- Palmer, Michael. "Letter 7." *Poetry Foundation*, Poetry Foundation, <https://www.poetryfoundation.org/poems/51141/letter-7>. Accessed 22 July 2021.
- Poem Generator*, Powered by Aardgo, <https://www.poem-generator.org.uk/>. Accessed 22 July 2021.
- "Poetry." *Language Is a Virus*, <http://www.languageisavirus.com/poetry/index.php>. Accessed 22 July 2021.
- Richards, Ivor Armstrong. *Poetics and Sciences: A Reissue of "Science and Poetry" (1926, 1935) with Commentary*. London: Routledge and Kegan Paul, 1970. Print.
- Roque, Antonio. "Language Technology Enables a Poetics of Interactive Generation." *Journal of Electronic Publishing*, vol. 14, no. 2, 2011. <https://doi.org/10.3998/3336451.0014.209>. Accessed 22 July 2021.
- Schwartz, Oscar. "Can a Computer Write Poetry?" TEDx. TEDxYouth@Sydney, Sydney, Australia. https://www.ted.com/talks/oscar_schwartz_can_a_computer_write_poetry. Accessed 22 July 2021.
- Shihadeh, Juliana and Ackerman, Margaret. "EMILY: An Emily Dickinson machine." *Proceedings of the 11th International Conference on Computational Creativity (ICCC 2020)*, 2020. Internet resource.
- Stagnaro, Angelo. *The Italian Jesuit Who Taught Computers to Talk to Us*. January 28, 2017. <https://www.ncregister.com/blog/the-italian-jesuit-who-taught-computers-to-talk-to-us>. Accessed 22 July 2021.
- Stein, Gertrude. "Red Faces"; "America." <https://allpoetry.com/Gertrude-Stein>. Accessed 22 July 2021.
- Stephens, Wallace. "The Snow Man by Wallace Stevens." *Poets.org*, Academy of American Poets, <https://poets.org/poem/snow-man>. Accessed 22 July 2021.
- "The Genuine Haiku Generator." *Everypoet.com*, Selendy Communications, <http://www.everypoet.com/haiku/>. Accessed 22 July 2021.
- Tobing, Berty Chrismartin Lumban and Manurung, Ruli. "A chart generation system for topical metrical poetry." *Proceedings of the 6th International Conference on Computational Creativity (ICCC 2015)*, Park City, Utah, 2015. Internet resource.
- Todorov, Tzvetan. *Introduction to Poetics*. Transl. Richard Howard. Brighton: Harvester Press, 1968/1981. Print.
- Tzara, Tristan, Dada Manifesto, 1918. https://writing.upenn.edu/library/Tzara_Dada-Manifesto_1918.pdf. Accessed 22 July 2021.
- Tzara, Tristan. *Three Poems by Tristan Tzara*. Translate by Forrest Pelsue. <https://exchanges.uiowa.edu/issues/cae-sura/three-dada-poems/>. Accessed 22 July 2021.
- Veale, Tony. "Less rhyme, more reason: Knowledge-based poetry generation with feeling, insight and wit." *Proceedings of the 4th International Conference on Computational Creativity (ICC 2013)*, Sydney, Australia, 2013. 152-159. Internet resource.
- Verse by Verse*, Google AI Semantic Experiences, <https://sites.research.google/versebyverse/>. Accessed 22 July 2021.

- Winkler, Adam. “‘Corporations Are People’ Is Built on an Incredible 19th-Century Lie.” *The Atlantic*, Atlantic Media Company, 21 Apr. 2020, <https://www.theatlantic.com/business/archive/2018/03/corporations-people-adam-winkler/554852/>. Accessed 22 July 2021.
- Wimmer, Mario. “Josephine Miles (1911–1985): Doing Digital Humanism with and without Machines.” *History of Humanities*, vol. 4, no. 2, 2019, www.journals.uchicago.edu/doi/abs/10.1086/704850?journalCode=hoh. Accessed 22 July 2021.
- Wong, Martin Tsan and Chun, Andy Hon Wai. “Automatic haiku generation using VSM.” *Proceedings of 7th WSEAS Int. Conf. on Applied Computer & Applied Computational Science (ACACOS '08)*, Hangzhou, China, 2008. Internet resource.
- Wright, George A. and Purver, Matthew. “Creative language generation in a society of engagement and reflection.” *Proceedings of the 11th International Conference on Computational Creativity (ICC 2020)*, Coimbra, Portugal, 2020. 169-172. Internet resource.
- Yampolskiy, Roman V. “Could an Artificial Intelligence Be Considered a Person under the Law?” *PBS*, Public Broadcasting Service, 7 Oct. 2018, <https://www.pbs.org/newshour/science/could-an-artificial-intelligence-be-considered-a-person-under-the-law>. Accessed 22 July 2021.