

How Machines Write Poetry

Written by Elizabeth Preston

December 15, 2015 // 09:00 AM EST

<http://motherboard.vice.com/read/how-machines-write-poetry>

As a teenager in Vermont, Sarah Harmon used Java to create a computer program that wrote poetry. She named it OGDEN. Then she submitted one of its poems in 2008 to her high-school literary magazine under the pen name Dan Goshen, an anagram of Ogden Nash.

“They accepted it,” Harmon laughs, “although they did say all these funny things about how abstract it was.”

OGDEN was nothing fancy, she says. It followed predefined rules of grammar and structure to compose poetical-sounding snippets. For example: *He was perfectly strange, / His world was shyly hopeless, / Then he tasted his dreams.*

These days, Harmon is a computer science PhD student at the University of California, Santa Cruz. She says it’s easy to write programs that generate poetry humans consider [publishable](#). But she’s interested in the next step: a program that’s actually creative.

Researchers can’t even agree on what creativity is. Yet Harmon and others are forging ahead, trying to write software that composes poetry in a more sophisticated way. In their successes and flops, they’re learning about what it means to be human.

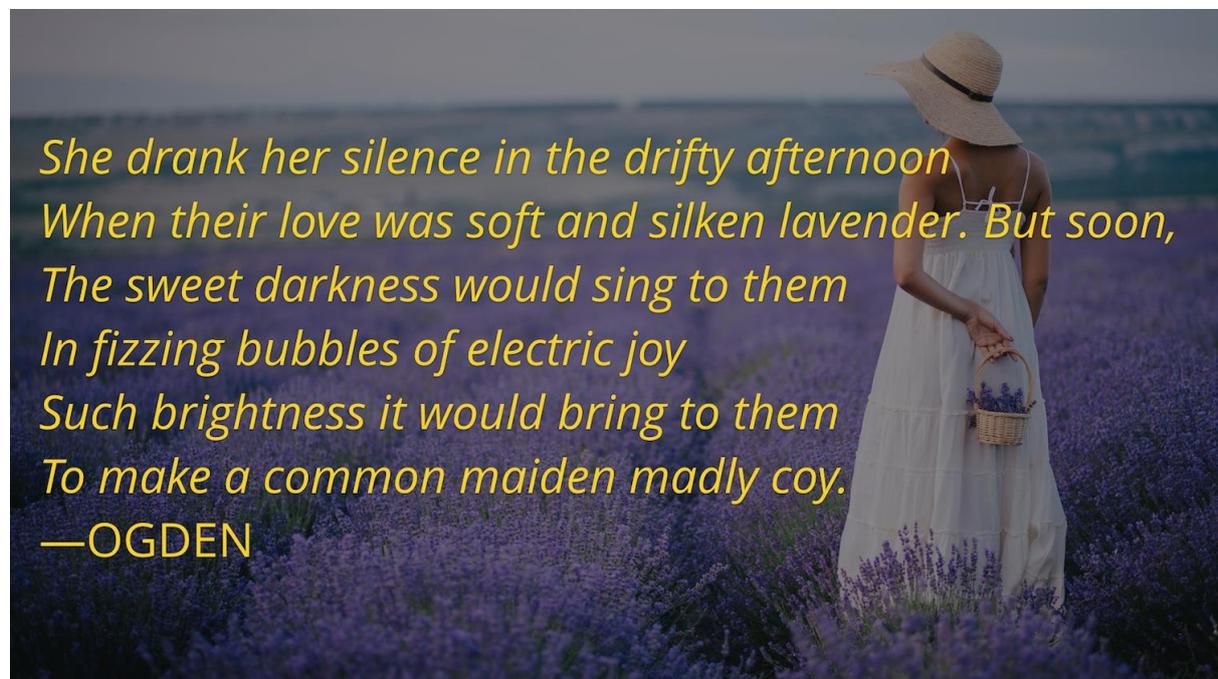


Image: Shutterstock

One of the first computer bards was named Racter. The program wrote [prose largely at random](#). Its most serendipitous bits of output became [a 1984 book of poetry](#) and surreal dialogues called *The Policeman's Beard Is Half Constructed*. But programs like this require cherry-picking by a person to make their work look good, Harmon says. She wanted to build something that could evaluate its own ideas and choose the best ones, like a human author.

Furthermore, she says, a poem written by OGDEN or Racter might be charming once or twice. But if you read enough of the program's outputs, you'll glimpse the patterns it's following. The poems won't surprise you.

"As system creators, we want to be surprised too," Harmon says. "For me that's one of the most important features of creative systems."

For her thesis, Harmon is working on another challenge in the field of [computational creativity](#), as it's called. She's trying to use software to adapt narrative prose into simpler language, such as for readers with limited English. Her work on computer poetry is just "for fun," she says.

Her latest program is called FIGURE8. Instead of composing entire poems, Harmon wanted to focus on one crucial element of poetry: figurative language. (Hence "figure-ate.") The program's output would be single similes. "I wanted beautiful content that was understandable and meaningful and surprising," Harmon says.

She preferred 'The snow falls like a dead cat', which got only middling ratings from humans. "They might have been cat lovers," she says.

First, Harmon dug through psychology and linguistics research to learn what makes a good simile. She found that clichés are bad, of course. A good simile is something new. A good simile is also understandable—the reader should see how the two things are "like" each other. And it should be apt, but unexpected. Harmon found that the cleverest-seeming similes link things that share some characteristics, but not too many. Comparing a strawberry to a cherry, for example, won't impress anybody.

FIGURE8 is what's called a "case-based reasoning" system. When Harmon asks it to describe a certain noun, the program checks its internal library to see how other authors have described the same thing. What words did they use? What properties does the thing have? What actions can it take? Harmon has loaded up the program's library with all the public-domain stories she could find. The program also has access to the internet.

Unlike most poetry programs, which stick to templates, FIGURE8 can learn new ways of composing a sentence from other writers. It can also infer the "unwritten rules" of language, Harmon says. For example, it's learned that you can put two or three adjectives in a row, or add a sentence fragment to the beginning of a simple sentence. One of its first constructions was *Like a pale moon, the garden lit up in front of him*.

Like a human author brainstorming and revising, FIGURE8 generates many possible similes, then goes back to analyze what it's written. It ranks all of its similes according to Harmon's criteria of clarity, novelty, aptness and surprise. For example, if a web search reveals that no one has associated these two things before, FIGURE8 guesses that its simile might be

unclear. If the two things share a major category—like the cherry and strawberry, both fruits—the simile gets a low score for surprise.

Some other researchers are also using case-based reasoning systems to generate poetry. [ASPERA](#), created by Pablo Gervás of the Universidad Complutense de Madrid, can compose Spanish poetry in three different traditional forms.



Image: Shutterstock

When Harmon saw the first outputs from FIGURE8, “I had no idea where it was getting a lot of these connections and even sentence structures from,” she says. “That was really exciting to me.”

For an objective assessment of her program’s work, Harmon recruited subjects from Amazon Mechanical Turk. In a pilot study, people read FIGURE8’s more complex sentences. But the task seemed to wear them out. “It became more of an English assignment,” she says. So for the [study](#) she presented at this year’s International Conference on Computational Creativity, Harmon restricted her program to simple sentences that fit certain patterns. It wasn’t FIGURE8’s most thrilling writing, but these shorter phrases would be easier for subjects to compare.

Each person read sets of five FIGURE8 sentences at a time. People ranked the similes for their clarity, as well as how enjoyable they were to read. Harmon then compared those ratings to scores the computer program had given itself for clarity and overall quality. Harmon found that her subjects’ ratings were pretty close to how FIGURE8 rated itself. For example, humans and computer agreed that *The queen stands like a strong castle* was a good, clear

simile. But they also agreed that *It was her pearl, crumpling like a drowned corpse* didn't make much sense.

Harmon would love to have writers or other experts judge FIGURE8's work too. Her online subjects tended to rate the similes better if they were obvious. *The snow continued like a heavy rain* got high scores, for example, even though Harmon thought this was quite a bad effort on FIGURE8's part. She preferred *The snow falls like a dead cat*, which got only middling ratings from humans. "They might have been cat lovers," she says.

Despite a few clunkers, Harmon calls FIGURE8 "much more successful than I initially expected."

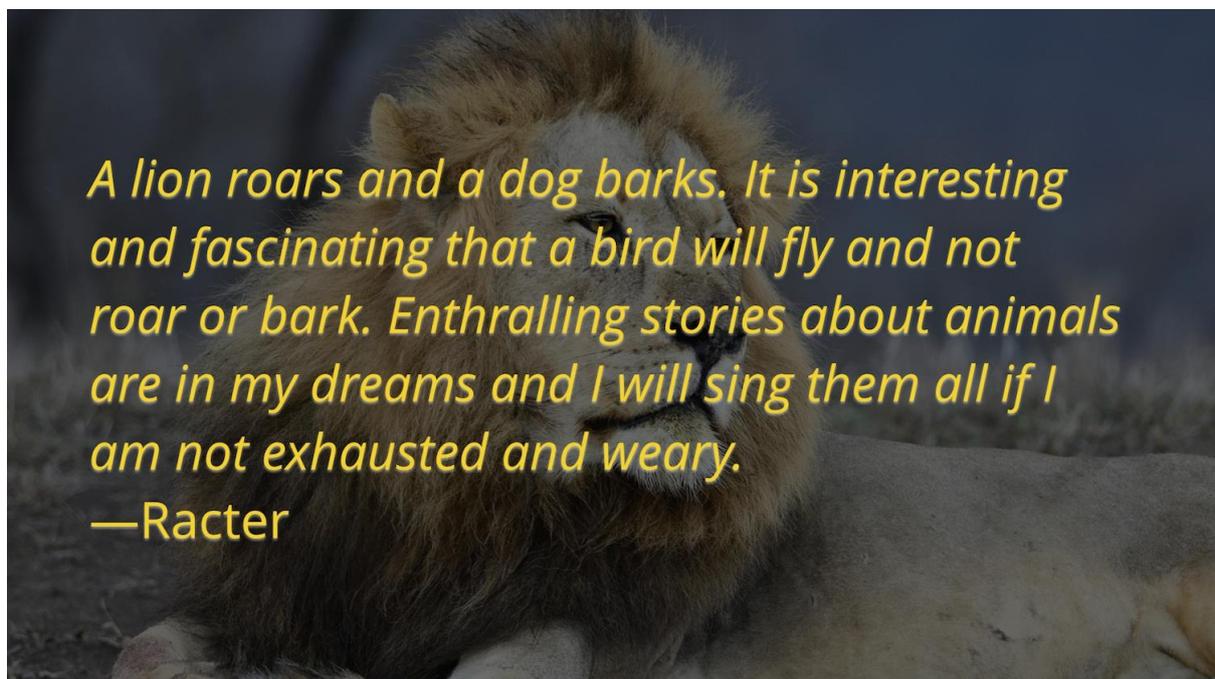


Image: Shutterstock

No matter how good the output, it's hard to argue that FIGURE8 is writing poetic sentences because it wants to. Can the program be creative without intent? "We can make up all kinds of stories in our heads for what a poem means," Harmon says. "It doesn't necessarily mean the system itself was creative." She thinks someday, computers may write poems out of their own intention. But we're not there yet.

University of London computer scientist Geraint Wiggins has [defined](#) computational creativity as the study of "behavior... which would be deemed creative if exhibited by humans." Are behavior and output the same thing? Researchers are still wrestling with questions like this.

Artificial poets don't always need creativity to please humans. Randomness and serendipity can be even better. "Some bots achieve surprising effects with the simplest language tools,"

computer scientist Tony Veale [wrote](#) in 2015. He pointed to [@Pentametron](#), a charming Twitter bot created by artist Ranjit Bhatnagar. The bot searches for tweets that happen to be in iambic pentameter, then retweets them in rhyming couplets: *Is this a pimple or mosquito bite / I need a dancing partner for tonight.*

Veale leads a computational creativity group at University College Dublin whose projects include the Twitter bot [@MetaphorMagnet](#) and an ironic simile generator called [Sardonicus](#). The group is also one of five research teams collaborating on the [What-If Machine](#) (WHIM). This software generates fictional scenarios, such as *What if a vacuum cleaner appeared in a backyard, and suddenly became a cat that was able to swim?* (It's a work in progress.)

The leader of the WHIM project is Simon Colton, who heads a computational creativity group at the University of London. Colton describes his own views on computer-generated poetry as "hardcore." He says Sarah Harmon's FIGURE8 is a step forward in producing poems that look like a human's work. But he thinks Harmon, like most computational creativity researchers, is ignoring a crucial problem. He calls it the humanity gap.

"Poetry is a way by which people connect," Colton says. Knowledge about a poem's author is part of our experience of reading it. We may read the same poem very differently depending on whether we believe the author to be a man or a woman, for example.

Without this human connection, Colton says, "We have to ask ourselves: what is the point of computer-generated poetry?"

Colton thinks computer poetry should be treated as its own genre, separate from human poetry. And rather than judging an artificially intelligent poet only on its output, Colton wants to address the "creative" process itself. His own program, called Full-FACE, writes poems based on newspaper articles. Along with each poem, it writes a summary of its process. *I read a story in the Guardian culture section entitled: "South Africa's ANC celebrates centenary with moment in the sun,"* Full-FACE [wrote in part](#) to explain one poem. The program had focused on a sentence about a man at the celebration who was wearing 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

Stephen McGregor, a PhD student in computational creativity at Queen Mary University of London, agrees with Colton that Turing-style tests of computer poetry are less interesting than exploring the creative process. He calls this a “holistic” approach to computational creativity.

At a recent poetry festival in London, McGregor demonstrated the poetry program he’s working on. The program uses a language model that makes words into points in a many-dimensional space. How these points relate in space imitates how the words are connected in English, and the program uses these relationships to compose poetry.

“The space itself is where the making of new meaning happens,” McGregor says—which is how he thinks of creativity.

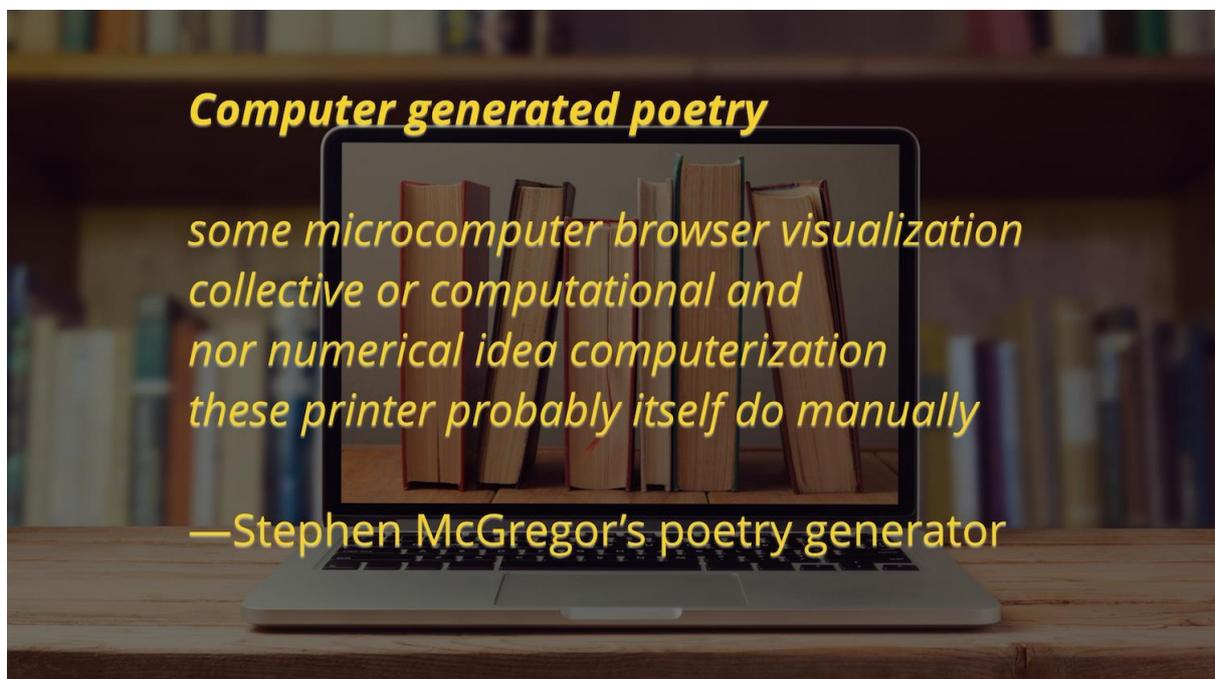


Image: Shutterstock

Even before computer scientists perfect their poetry programs, teaching computers about creativity may help them work better with humans.

That’s the goal of [Poetry for Robots](#), a project started by [Corey Pressman](#), who’s a partner at the digital agency Neologic. On the website, humans can enter lines of their own poetry in response to photos. This input is becoming a database of human associations and emotions that go with various images. Pressman hopes that eventually, this database will improve search algorithms. After all, implicit associations with objects or events—what we might call poetry—are what computers call metadata.

The experiment has been a success so far, Pressman says. It's even getting some attention from large search companies. But he thinks true poetry by robots would look quite different.

"They might write poems about how their algorithms are feeling slippy, or 'my chips feel wet today,'" he says. "Other robots would understand that, but we wouldn't. I think that would be a real success."

Sarah Harmon also wants to help computers and humans work together. Ultimately, her goal with FIGURE8 isn't a program that can pen great poetry, but an aid to human creativity.

"This is a discovery process for helping us figure out how humans are creative," she says. "How can [computers] support that process to help us be the best we can be?"

If computers can simulate how a human generates ideas, maybe they can help us with all kinds of creative problem solving, Harmon says. That might mean writing music or coming up with recipes. It could also mean inventing new techniques for processes like building bridges.

And although these partnerships could take humans to new places creatively, we may never see a true computer poet.

"If a computer were to somehow have its own idea about how to write poetry, the computer would probably write about what it's like to be a computer," Stephen McGregor says. However, he adds, "I very much doubt there's anything it's like to be a computer."

Topics: [creativity](#), [robots](#), [machines](#), [poetry](#), [algorithms](#), [features](#)