

A good read: AI evaluates quality of short stories

August 21 2017

The idea that artificial intelligence will someday be able to understand and even generate narratives has inspired and motivated researchers for years. A question inextricably bound to both lines of research remains unresolved, however: Can AI recognize a good story if it sees one?

In the first large-scale attempt to tackle this challenging problem, scientists at Disney Research and the University of Massachusetts Boston have developed neural networks that can evaluate short narratives. These AIs don't try to mimic the professional judgment of literary critics, but attempt to predict which stories will appeal to a large population.

"Our neural networks had some success in predicting the popularity of stories," said Boyang "Albert" Li, a research scientist at Disney Research. "You can't yet use them to pick out winners for your local writing competition, but they can be used to guide future research."

The team will present its findings Aug. 23 at the International Joint Conference on Artificial Intelligence in Melbourne, Australia.

"The ability to predict narrative quality impacts on both story creation and story understanding," said Markus Gross, vice president at Disney Research. "To evaluate quality, the AI needs some level of understanding of the text. And if AIs are to create narratives, they need to be able to judge the quality of what they are producing."

One of the challenges in developing automated evaluation of story quality is the lack of large databases of stories that have been evaluated by humans and can be used to train AIs, noted Tong Wang, a computer science Ph.D. student at UMass Boston and a lab associate at Disney Research. The scientists found the question and answering website Quora, to be a good data source, as many of its answers are in the form of stories, he added. Reader upvotes, likewise, measure a story's popularity and serve as a proxy for narrative quality.

After gathering almost 55,000 answers, the researchers developed an algorithm to classify them as either stories or non-stories, yielding more than 28,000 stories with an average of 369 words.

To understand the complex semantics of the stories, the research team looked for ways to represent the influence of story structures in [neural networks](#) because a sequence of events can interact to reveal character intentions.

The researchers developed an AI that evaluated separate regions of each story—including the question that prompted each. They created a network that looked at the regions interdependently and another network that took a holistic view, looking for how the meaning of the events and story regions emerged from the entire story. In each case, the AIs made predictions of which texts would prove most popular with readers.

In experiments, the neural nets showed improvement over a baseline text evaluation system, with the holistic neural network registering an 18 percent improvement.

In addition to Li and Wang, Ping Chen, Wang's PhD adviser and associate professor of computer science and engineering at UMass Boston, contributed to this research.

Combining creativity and innovation, this research continues Disney's rich legacy of leveraging technology to enhance the tools and systems of tomorrow.

Provided by Disney Research

Citation: A good read: AI evaluates quality of short stories (2017, August 21) retrieved 19 April 2024 from <https://phys.org/news/2017-08-good-ai-quality-short-stories.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.