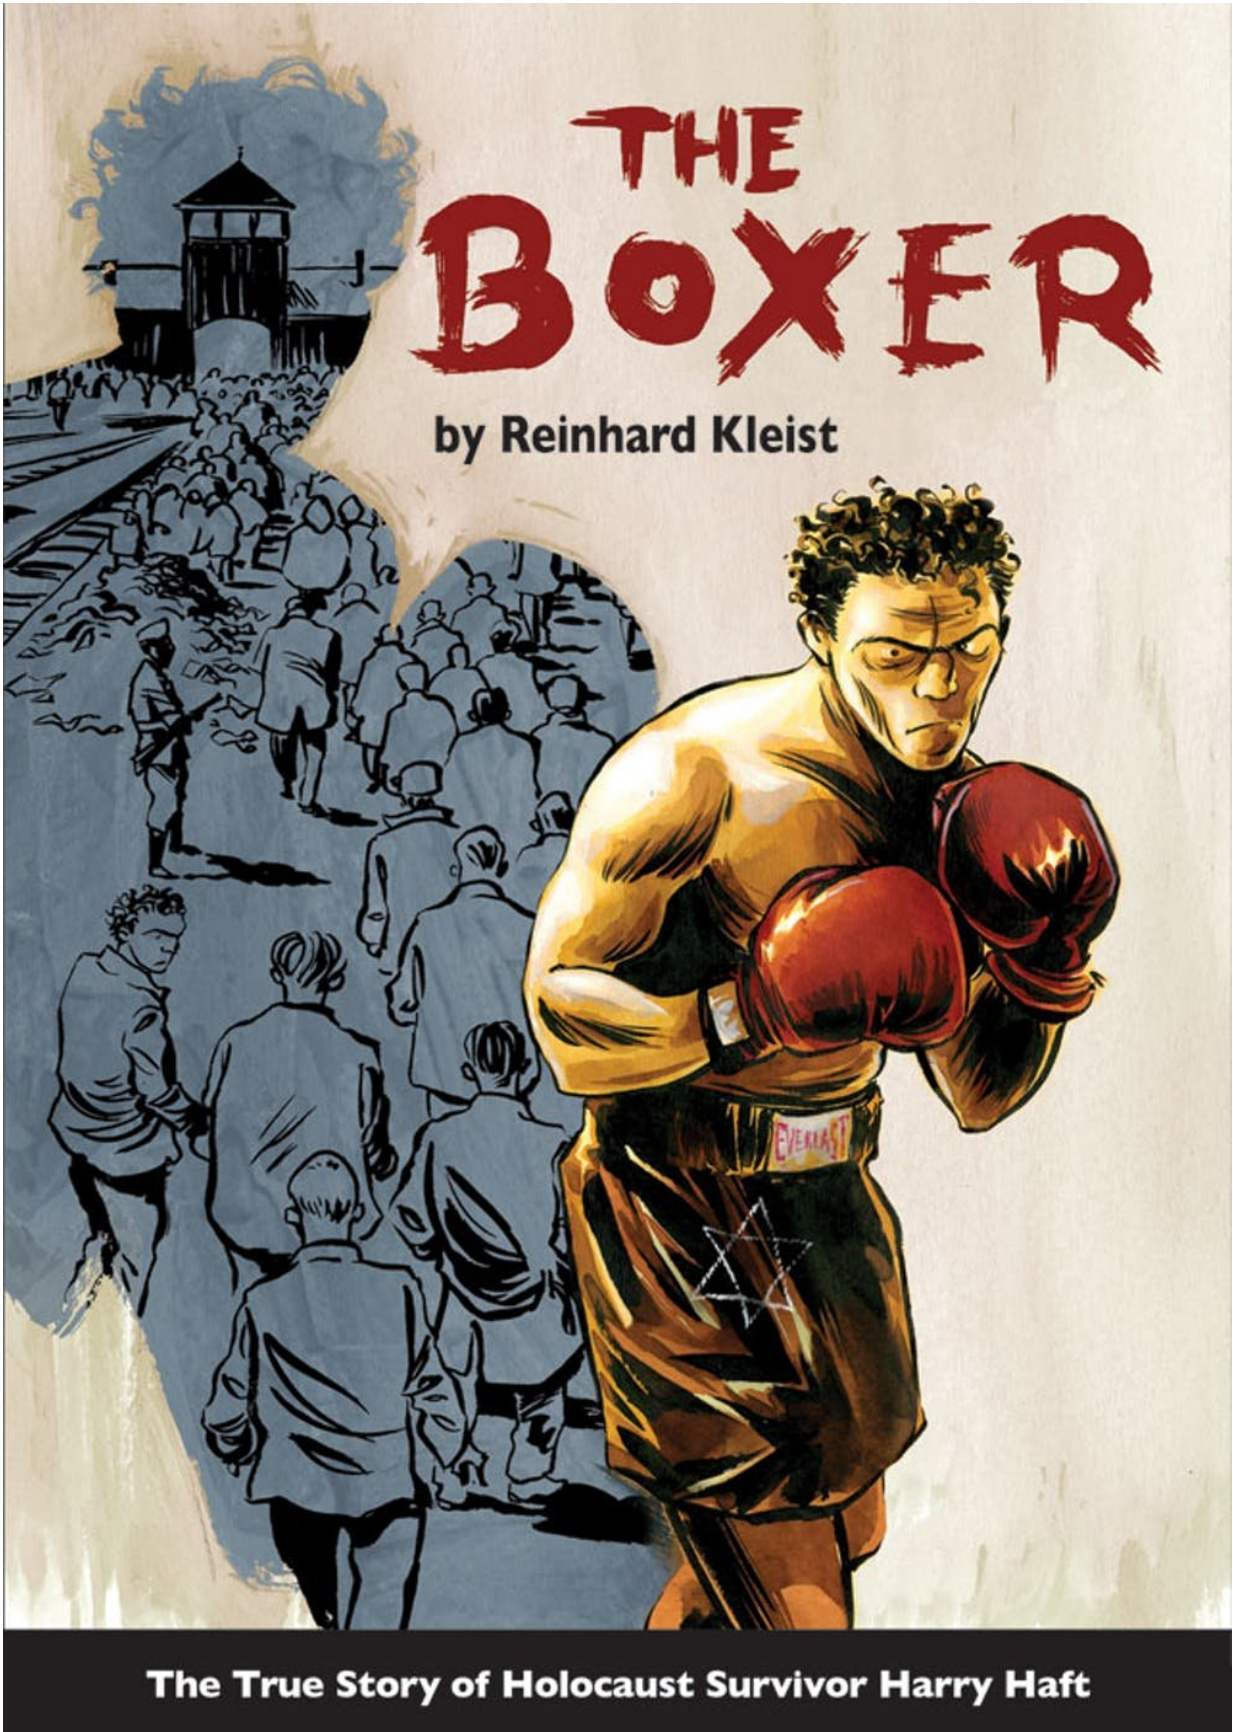# We asked artificial intelligence to analyze a graphic novel – and found both limits and new insights

December 5 2018, by Leonie Hintze And Arend Hintze

THE BOXER

by Reinhard Kleist

The True Story of Holocaust Survivor Harry Haft

A graphic novel examined by artificial intelligence. Credit: [Reinhard Kleist/Self Made Hero](#)

With one spouse studying the [evolution of artificial and natural intelligence](#) and the other researching the [language, culture and history of Germany](#), imagine the discussions at our dinner table. We often experience the stereotypical clash in views between the quantifiable, measurement-based approach of natural science and the more qualitative approach of the humanities, where what matters most is how people feel something, or how they experience or interpret it.

We decided to take a break from that pattern, to see how much each approach could help the other. Specifically, we wanted to see if aspects of artificial intelligence could turn up new ways to interpret a nonfiction graphic novel about the Holocaust. We ended up finding that some AI technologies are not yet advanced and robust enough to deliver useful insights – but simpler methods resulted in quantifiable measurements that showed a new opportunity for interpretation.

Choosing a text

There is plenty of research available that [analyzes large bodies of text](#), so we chose something more complex for our AI analysis: Reinhard Kleist's "[The Boxer](#)," a graphic novel based on the true story of how Hertzko "Harry" Haft survived the Nazi death camps. We wanted to identify emotions in the facial expressions of the main character displayed in the book's illustrations, to find out if that would give us a new lens for understanding the story.

In this black-and-white cartoon, Haft tells his horrific story, in which he and other concentration camp inmates were made to box each other to the death. The story is written from Haft's perspective; interspersed throughout the narrative are panels of flashbacks depicting Haft's memories of important personal events.

The humanities approach would be to analyze and contextualize elements of the story, or the tale as a whole. Kleist's graphic novel is a reinterpretation of a 2009 biographical novel [by Haft's son Allan](link), based on what Allan knew about his father's experiences. Analyzing this complex set of authors' interpretations and understandings might serve only to add another subjective layer on top of the existing ones.

From the perspective of science philosophy, that level of analysis would only make things more complicated. Scholars might have differing interpretations, but even if they all agreed, they would still not know if their insight was objectively true or if everyone suffered from the same illusion. Resolving the dilemma would require an experiment aimed at generating a measurement others could reproduce independently.

Reproducible interpretation of images?

Rather than interpreting the images ourselves, subjecting them to our own biases and preconceptions, we hoped that AI could bring a more objective view. We started by scanning all the panels in the book. Then we ran [Google's vision AI](link) and [Microsoft AZURE's face recognition and emotional character annotation](link) as well.

The algorithms we used to analyze "The Boxer" were previously trained by Google or Microsoft on hundreds of thousands of images already labeled with [descriptions of what they depict](link). In this [training phase](link), the AI systems were asked to identify what the images showed, and those answers were compared with the existing descriptions to see if the

system being trained was right or wrong. The training system strengthened the elements of the underlying deep neural networks that produced correct answers, and weakened the parts that contributed to wrong answers. Both the method and the training materials – the images and annotations – are crucial to the system's performance.

Then, we turned the AI loose on the book's images. Just like on "Family Feud," where the show's producers ask 100 strangers a question and count up how many choose each potential answer, our method asks an AI to determine what emotion a face is showing. This approach adds one key element often missing when subjectively interpreting content: reproducibility. Any researcher who wants to check can run the algorithm again and get the same results we did.

Unfortunately, we found that these AI tools are optimized for digital photographs, not scans of black-and-white drawings. That meant we did not get much reliable data about the emotions in the pictures. We were also disturbed to find that none of the algorithms identified any of the images as relating to the Holocaust or concentration camps – though human viewers would readily identify those themes. Hopefully, that is because the AIs had problems with the black-and-white images themselves, and not because of negligence or bias in their training sets or annotations.

Bias is a well-known phenomenon in machine learning, which can have really offensive results. An analysis of these images based solely on the data we got would not have discussed or acknowledged the Holocaust, an omission that is against the law in Germany, among other countries. These flaws highlight the importance of critically evaluating new technologies before using them more widely.

Finding other reproducible results

Determined to find an alternative way for quantitative approaches to help the humanities, we ended up analyzing the brightness of the pictures, comparing flashback scenes to other moments in Haft's life. To that end, we quantified the brightness of the scanned images using [image analysis software](link).

We found that throughout the book, emotionally happy and light phases like his prison escape or Haft's postwar life in the U.S. are shown using bright images. Traumatizing and sad phases, such as his concentration camp experiences, are shown as dark images. This aligns with [color psychology](link) identifications of white as a pure and happy tone, and black as symbolizing sadness and grief.

Having established a general understanding of how brightness is used in the book's images, we looked more closely at the flashback scenes. All of them depicted emotionally intense events, and some of them were dark, such as recollections of cremating other concentration camp inmates and leaving the love of his life.

We were surprised, however, to find that the flashbacks showing Haft about to punch opponents to death were [bright and clear](link) – suggesting he is having a positive emotion about the upcoming fatal encounter. That's the exact opposite of what readers like us probably feel as they follow the story, perhaps seeing Haft's opponent as weak and realizing that he is about to be killed. When the reader feels pity and empathy, why is Haft feeling positive?

The middle image in this sequence shows an example of a bright flashback. Credit: Reinhard Kleist/Self Made Hero

This contradiction, found by measuring the brightness of pictures, may reveal a deeper insight into how the Nazi death camps affected Haft emotionally. For us, right now, it is unimaginable how the outlook of beating someone else to death in a boxing match would be positive. But perhaps Haft was in such a desperate situation that he saw hope for survival when facing off against an opponent who was even more starved than he was.

Using AI tools to analyze this piece of literature shed new light on key elements of emotion and memory in the book – but they did not replace the skills of an expert or scholar at interpreting texts or pictures. As a result of our experiment, we think that AI and other computational methods present an interesting opportunity with the potential for more quantifiable, reproducible and maybe objective research in the humanities.

It will be challenging to find ways to use AI appropriately in the humanities – and all the more so because current AI systems are not yet sophisticated enough to work reliably in all contexts. Scholars should also be alert to potential biases in these tools. If the ultimate goal of AI research is to develop machines that rival human cognition, artificial intelligence systems may need not only to behave like people, but understand and interpret feelings like people, too.

This article is republished from The Conversation under a Creative Commons license. Read the original article.

Provided by The Conversation

Citation: We asked artificial intelligence to analyze a graphic novel – and found both limits and new insights (2018, December 5) retrieved 16 July 2024 from https://phys.org/news/2018-12-artificial-intelligence-graphic-limits-insights.html