Has the evolution of artificial intelligence reached its limits?

But the peculiar thing about deep learning is just how old its key ideas are. Hinton's breakthrough paper…was published in 1986. The paper elaborated on a technique called backpropagation, or backprop for short. Backprop…is "what all of deep learning is based on—literally everything."

…

[M]aybe we're not actually at the beginning of a revolution. Maybe we're at the end of one.
…

[T]hese "deep learning" systems are still pretty dumb, in spite of how smart they sometimes seem. A computer that sees a picture of a pile of doughnuts piled up on a table and captions it, automatically, as "a pile of doughnuts piled on a table" seems to understand the world; but when that same program sees a picture of a girl brushing her teeth and says "The boy is holding a baseball bat," you realize how thin that understanding really is, if ever it was there at all.
…

[T]he latest sweep of progress in AI has been less science than engineering, even tinkering. And though we've started to get a better handle on what kinds of changes will improve deep-learning systems, we're still largely in the dark about how those systems work, or whether they could ever add up to something as powerful as the human mind.

It's worth asking whether we've wrung nearly all we can out of backprop.

The GLP aggregated and excerpted this blog/article to reflect the diversity of news, opinion, and analysis. Read full, original post: Is AI Riding a One-Trick Pony?