

---

Subject: Re: Would you pay to keep this woman in your house?

Posted by [Wayne Parham](#) on Fri, 19 Jan 2024 18:07:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

LOL! :lol:

Yes, we're all used to these chatbots being free, and also the voice assistants like Alexa. So nobody is going to want to pay for them.

And even more an issue - in my opinion - is the accuracy problem with neural networks. Assistants like Alexa tend to be narrow AI applications that have pretty tight "guardrails." They only can respond to a fairly narrow set of specific things that they've been trained on. This makes them act a little more like rules-based systems, which don't stray off their knowledgebases.

General-purpose AI is exposed to a lot more training data, so it can respond to a lot more topics. But that also means it can make a lot more inferences that don't track with reality. So general-purpose AI has a problem with accuracy.

But if you think about it, accuracy is even more of a problem with human neural networks, especially politicians. :d

Jokes aside, the whole mechanism of neural networks prevents them from acting in the ways we normally associate as "computer behavior," which is rules-based and deterministic. There is no guarantee of accuracy from any neural network. It is literally a statistical inference engine, and it gains its "understanding" from training data that is inaccurate and incomplete.

Hopefully its training data is mostly accurate, but even if its perfect, it isn't complete. Nothing ever can be. The whole underpinnings of machine intelligence tell us that, right from the start, when guys like Alan Turing and John von Neumann studied Kurt Gödel's incompleteness theorem and started making thinking machines using that understanding.

There are always holes in our understanding, and in those holes, inference fills in the blank spots. Sometimes its with useful stuff, but as often as not, it's nonsense. A good machine then studies what it puts in those holes, and attempts to test it to see if its useful or not. If useful, that's good knowledge. If not, its nonsense that should be discarded.

But our machine learning systems haven't quite gotten there yet. We're moving our systems in that direction, but they're not there yet. Neither are the meat-suit artificial intelligences that walk this planet. :)

---