I've made several changes to my little chatbot over the past couple years, so I thought I'd give it a mention.

For one thing, I no longer train them like I used to. Instead, I use what is called "Retrieval Augmented Generation" (RAG) which essentially means you provide documents that you consider "truth" as training data. This adds to the pre-trained model.

Right now, it only has one little document for this purpose: A file that says, "Your name is Aliza" and then follows that with some drivel about the reason why I chose that name. I may add more stuff later, but I don't want to indoctrinate the thing.

Second big change is similar, in that Aliza can now search the internet for information. So it is not limited to the pre-training data, and can respond to what's new.

And lastly, it has a little bit of conversational memory, so it can "remember" the last few things it was asked. That helps it with context, and allows it to better follow a conversation. Sometimes it's better to open with some ideas and then refine as you go along. Without conversational memory, it can't do that. So it helps having some.

It does keep track of who's asking what, so you can't ask it what other people are saying.

If I were more ambitious, I'd add some genetic algorithm agents, so it could actually start to learn concepts. I think that's the next big step. And some shops have done it too - like DeepMind - with their AlphaGo, AlphaZero and AlphaFold systems, among others. They allow their systems to teach themselves very sophisticated concepts. But they're also very specialized. I'm thinking more along the lines of very simple concepts, like inside/outside, up/down, heavy/light and then build up from those.

But it's still incredibly useful, and a heck of a lot of fun. I've often said "LLMs are blind, deaf and dumb but they're very well read." That's not totally true, but closer to truth than false. And even still, they just have so much data to draw from that they can do some pretty incredible things.

Hope you all enjoy it!
Aliza