

WHAT WE TALK ABOUT WHEN WE TALK ABOUT AI (PART THREE)

Proteins, Factories, and Wicked Solutions

Part 3- But What is AI *Good* For?

**(Go to Part Two)(Go to Part
four)**

There are many frames and metaphors to help us understand our AI age. But one comes up especially often, because it is useful, and perhaps even a bit true: the Golem of Jewish lore. The most famous golem is the Golem of Prague, a magical defender of Jews in the hostile world of European gentiles. But that was far from the only golem in Jewish legends.



T
h
e
g
o
l
e
m
w
a
s
o
f
t
e
n
a
t
r

ustworthy and powerful servant in traditional Jewish stories – but only when in the hands of a wise rabbi. To create a golem proved a rabbi's mastery over Kabbalah, a mystical interpretation of Jewish tradition. It was esoteric and complex to create this magical servant of mud and stone. It was brought to life with sacred words on a scroll pressed into the mud of its forehead. With that, the inanimate mud became person-like and powerful. That it echoed life being granted to the clay Adam was no coincidence. These were deep and even dangerous powers to call on, even for a wise rabbi.

You're probably seeing where this is going.

Mostly a golem was created to do difficult or dangerous tasks, the kind of things we fleshy humans aren't good at. This is because we are too weak, too slow, or inconveniently too mortal for the work at hand.

The rabbi activated the golem to protect the Jewish community in times of danger, or use it when a task was simply too onerous for people to do. The golem could solve problems that were not, per se, impossible to solve without supernatural help, but were a lot easier with a

giant clay dude who could do the literal and figurative heavy lifting. They could redirect rivers, move great stones with ease. They were both more and less than the humans who created and controlled them, able to do amazing things, but also tricky to manage.

When a golem wasn't needed, the rabbi put it to rest, which was the fate of the Golem of Prague. The rabbi switched off his creation by removing the magic scroll he had pressed into the forehead of the clay man.

Our Servants, Ourselves


The parallels with our AIs are not subtle.

If the golem was not well managed, it could become a violent horror, ripping up anything in its path mindlessly. The metaphors for technology aside, what makes the golem itself such a useful idea for talking about AI is how human shaped it is. Both literally, and in its design as the ultimate desirable servant. The people of Prague mistook the golem for a man, and we mistake AI for a human-like mind.

Eventually, the rabbis put the golems away forever, but they had managed to do useful things that made life easier for the community. At least, sometimes. Sometimes, the golems got out of hand.

It is unlikely that we're going to put our new AI golem away any time soon, but it seems possible that after this historical moment of collective madness, we will find a good niche for it. Because our AI golems are very good at doing some important things humans are naturally bad at, and don't enjoy doing anyway.

Folding Proteins for Fun and Profit



Alphafold 3 logo

Perhaps the originally famous example of our AI golem surpassing our human abilities is AlphaFold, Google's protein folding AI. After throwing many technological tools at the problem of predicting how proteins would shape themselves in many circumstances, Google's specialist AI was able to predict folding patterns in roughly 2/3rds of cases. AlphaFold is very cool, and could be an amazing tool for technology and health one day. Understanding protein folding has implications in understanding disease processes and drug discovery, among other things.

If this seems like a hand-wavy explanation of AlphaFold, it's because I'm waving my hands wildly. I don't understand that much about AlphaFold – which is also my point. Good and useful AI tends to be specialized to do things humans are bad at. We are not good golems, either in terms of being able to do very difficult and heavy tasks, or paying complete attention to complex (and boringly repetitive) systems. That's just not how human attention works.

One of our best Golem-shaped jobs is dealing with turbulence. If you've dealt with physics in a practical application, anything from weather prediction to precision manufacturing, you know that turbulence is a terrible and crafty enemy. It can become impossible to calculate or predict. Often by the time turbulent problems are perceivable by humans or even normal control systems, you're already in trouble. But an application-specific AI, in, for instance, a factory, can detect the beginning of a component

failure below both human and even normal programatic detection.

A well-trained bespoke AI can catch the whine of trouble before the human-detectable turbulence starts. This is because it has essentially “listened” to how the system works so deeply over time. That’s its whole existence. It’s a golem doing the one or two tasks for which it has been “brought to life” to do. It thrives with the huge data sets that defeat human attention. Instead of a man shaped magical mud being, it’s a listener, shaped by data, tirelessly listening for the whine of trouble in the system that is its whole universe.



Similarly, the giant datasets of NOAA and NASA could take a thousand human life years to comb through to find everything you need to accurately predict a hurricane season, or the transit of the distant exoplanet in front of its sun.

The trajectories and interactions of the space junk enveloping Earth are dangerously out of reach of human calculation – but maybe not with AI. The thousands of cycles of an Amazon cloud server hosting a learning model that gets just close enough to modeling how the stochastic processes of weather and space are likely to work will never be human readable.

That third-of-the-time-wrong quality of AlphaFold is kind of emblematic of how AI is mostly, statistically right, in specific applications with a lot of data. But it’s no divine oracle, fated to always tell the truth. If you know that, it’s often close enough for engineering, or figuring out what part of a system to concentrate human resources next. AI is not smart or creative (in human terms), but it also doesn’t quit until it gets turned off.

Skynet, But for Outsourcing

AI can help us a lot with doing things that humans aren't good at. At times a person can pair up with an AI application and fill in each other's weaknesses – the AI can deliver options, the human can pick the good one. Sometimes an AI will offer you something no person could have thought of, sometimes that solution or data is a perfect fit; the intractable, unexplainable, wicked solution. But the AI doesn't know it has done that, because an AI doesn't know in the way we think of as knowing.

There's a form of chess that emerged out of computers becoming better than humans at this cerebral hobby, like IBM's Deep Blue. It's called cyborg, or centaur, chess, in which both players are using a chess AI as well as their own brains to play. The contention of this part of the chess world is that if a chess computer is good, a chess computer plus human player is even better. The computer can compute the board, the human can face off with the other player.

This isn't a bad way of looking at how AI can be good for us; doing the bits of a task we're not good at, and handing back control for the parts we are good at, like forming goals in a specific context. Context is still and will likely always be the realm of humans; the best chess computer there is still won't know when it's a good idea to just lose a game to your nephew at Christmas.

Faced with complex and even wicked problems, humans and machines working together closely have a chance to solve problems that are intractable now. We see this in the study and prediction of natural systems, like climate interacting with our human systems, creating Climate Change.


Working with big datasets lets us predict, and sometimes even ameliorate, the effects of climate on both human built systems and natural systems. That can be anything from predicting

weather, to predicting effective locations to build artificial reefs where they are most likely to revitalize ocean life.

It's worthwhile to note that few, or maybe even none, of the powerful goods that can come from AI are consumer facing. None of them are the LLMs (Large Language Models) and image generators that we've come to know as AI. The benefits come from technical datasets paired with specialized AIs. Bespoke AIs can be good for a certain class of wicked problems- problems that are connected to large systems, where data is abundant and hard to understand, with dependancies that are even harder.

But Can Your God Count Fish All Day

Bespoke AIs are good for Gordian knots where the rope might as well be steel cord. In fact, undoing a complex knot is a lot like guessing how protein folding will work. Even if you enjoyed that kind of puzzle solving, you simply aren't as good at it as an AI is. These are the good tasks for a golem, and it's an exciting time to have these golems emerging, with the possibility of detecting faults in bridges, or factories, or any of our many bits of strong-then-suddenly-fragile infrastructure.



Students in Hawaii worked on AI projects during the pandemic, and all of them were pretty cool

Industrial and large data AI has the chance to

change society for the better. They are systems that detect fish and let them swim upstream to spawn. They are NOAA storm predictions, and agricultural data that models a field of wheat down to the scale of centimeters. These are AI projects that could help us handle climate change, fresh water resources, farm to table logistics, or the chemical research we need to do to learn how get the poisons we already dumped into our environment back out.

AI, well used, could help us preserve and rehabilitate nature, improve medicine, and even make economies work better for people without wasting scarce resources. But those are all hard problems, harder to build than just letting an LLM lose to train on Reddit. They are also not as legible for most funders, because the best uses of AI, the uses that will survive this most venal of ages, are infrastructural, technical, specialized, and boring.

The AIs we will build to help humanity won't be fun or interesting to most people. They will be part of the under-appreciated work of infrastructure, not the flashy consumer facing chatbots most people think is all that AI is. Cleaning up waterways, predicting drug forms, and making factories more efficient is never going to get the trillion dollars of VC money AI chatbots are supposed to somehow 10x their investments on. And so, most people seeing mainly LLMs, we ask if AI is good or bad, without knowing to ask what kind of artificial intelligence we're talking about.

(Go to Part Two)