

In Brief:


At a time where amounts of money it's hard to believe actually exist are being spent trying to convince us that "AI" solutions are a panacea for every technical and social problem, it's hard to tell what these tools are actually useful for. This guide serves as an extremely brief primer on the most common kinds of AI currently being pushed, what the limitations of those technologies are, and what they might be useful for.

THIS MACHINE HAS NO GNOSIS

A WARNING FOR ASPIRANTS

This Zine is made available as a digital document which may be freely distributed as a file or in print.

The author reserves the right to sell printed copies for a nominal fee.



ARCANA LABS

Zachary J. Adam

Zine#

001


Issue

2026-A

Technology Survival Manual

THIS MACHINE HAS NO GNOSIS

A WARNING FOR ASPIRANTS



ARCANA LABS

ARTIFICIAL INTELLIGENCE DOESN'T EXIST YET

There's been a lot of talk in the last two years about "AI"-driven solutions to technological and sociological problems. So much so, that AI as a buzzword has become a sort of wunderwaffen against "problems". It's hard to touch anything with more than a few bits of computing power and not find an AI solution badly shoehorned into the side - a solution to a problem nobody had, solved in a way that nobody wants.

This manual is a short guide to the layman on why these products are not as desirable as they appear - and why that problem is baked into the technology itself, and not merely a question of "being ready for prime-time".

To begin with, I need to stress upon you the importance of this single idea: **Artificial Intelligence, as of the time of Writing, Does Not Exist**. The two major product classes we'll be discussing - Large Language Models and Generative Adversarial Networks - are not intelligent constructs. They are tools, which do have jobs, and we'll discuss that too, but it's not a question of whether or not they are "smart enough yet". The problem is that these are not machines for knowing things. They are not mechanisms with a comprehension system. *They have no Gnosis.*

This is an important concept that needs to be well understood: GANs and LLMs can both be used to create impressive output. Their outputs may even strongly resemble true outputs. An LLM can, more or less by accident, give the right answer to "What year is it", but as the internet will show you, accuracy is by no means guaranteed.

The distinction here is the difference between "Truth" and "Truth-Shaped". The AI summary on your search results may give the right answer to a question because right answers to the question are statistically possible, but it's equally likely to construct a grammatically-correct answer

that cites references which never exist. Those who say this is an issue with the original question don't understand the difference between Truth and Truth-Shaped.

To be short about it and spare you a full blown course on Theory of Knowledge, the distinction, for our purposes, is this:

- A statement is true when it both reflects the reality of a situation and can be demonstrated to be true, usually by evidence or measurement. A conclusion based on some source material can be said to be true if the conclusion represents the source material's thesis and the source material really exists.
- A statement is truth-shaped when it appears to be true but gets there based on spurious logic; a source that doesn't exist or a source which did not represent reality. Example, an AI responding to a question saying that a certain event happened when the source cited was actually a satirical article that doesn't accurately reflect reality, or a conclusion that was based on a paper that sounds like it existed, but didn't.

When Truth-Shaped answers are given, that is not an artifact of a "hallucination" or a deficiency in the Large Language Model's code: it is a consequence of Large Language Models not being premised on truth and falsehood, but statistics. Similarly, the tendency of GANs to create images that closely resemble the works of living artists but only poorly, or to struggle to accurately include text in rendered images, are not "issues" to be solved but a fundamental misunderstanding of what a GAN is.

WHAT AN LLM IS AND DOES

An LLM, or large language model – like Grok or ChatGPT – is exactly what it says on its tin: a large model of a natural language. Model in this sense is a statistical term. Large Language Models are trained by feeding large corpuses of written material into them. These texts are then analyzed statistically, with individual words and parts of speech taken as “tokens”. Vector mathematics is then used to describe the relationships between all these tokens. Given a large enough model, trained extensively enough, it is possible to do text generation via the large language model in response to a prompt.

Like the original training data, the prompt is broken up into tokens so that the model can produce an analysis of the question. This analysis of the question is then used to produce an answer, by responding to the prompt with a sequence of tokens which is statistically likely to resemble a reply to the original prompt’s sequence of tokens.

This obtuse wording has been chosen for a reason. **Token** is an important term here. The word “fish” is likely to appear as a token in any english-trained LLM. A good model might have two tokens, one to represent the noun “fish” and one to represent the verb whose infinitive is “to fish”. This token is then very precisely related with vector expressions to other tokens in the model. However, this is just mathematics. The Large Language Model does not have a conception of what a fish is or what it means to fish for something. Instead, it simply knows what other words are closely associated with the word fish.

This is an important distinction. The word “fish” is itself a token for the idea of a fish. At risk of inspiring motion in the grave of René Magritte, the word “fish” is not directly equivalent to the idea of a fish or even the image conjured to mind by a fish. Without the capacity to hold the idea of a fish in mind and understand it, LLMs are unable to understand fish; they can merely give you statistical analyses of texts related to fishing.

This lack of understanding and basis in statistics has led to the rise of the term **Stochastic Parrot** to describe an LLM, and it's an apt comparison.

If you wouldn't take the word of a Parrot on a problem, you probably shouldn't take the word of an LLM. (It is worth noting that some species of parrots such as the African Grey Parrot may actually have language cognition, so if anything you should take their word over that of an LLM, which definitely is not cognizant.)

WHAT AN GAN IS AND DOES

One way in which a GAN is closer to “AI” than an LLM is that it is, arguably, a more advanced construct. This, however, is saying that a AA battery is closer to an automobile than a hammer is; it doesn’t mean much.

In a GAN, in general, you still have a very large statistical model built by consuming an enormous amount of training data. That training data is selected based on what you want the GAN to do. If you were creating an image generation GAN, you would want to feed it as many images as you can, tagged appropriately with descriptors of their content.

GANs are “adversarial” networks because they are agentic, and their agents are sorted into two broad categories:

- **Generator Agents**, which attempt to create an output that closely matches the statistical understanding of the original data set (bounded by some request), and;
- **Discriminator Agents**, which attempt to tell the difference between the output of the generator agents and the training data.

Basically, if you ask a GAN to draw you an image of your ideal solarpunk future, the generator agents are going to start creating image files aiming for statistical analysis results similar to the images tagged “ideal solarpunk future” in the training data set, and the discriminator agent is going to present you with the “most statistically significant” result.

This makes GANs slightly better than an LLM, in that they at least try to check their own work, but there are two major caveats there:

I. **Neither a GAN nor an LLM has gnosis.** They are both based on the same mathematics-based understandings of information theory and are reliant on the quality of their training data while also not having any concept of the truth or falsehood of their conclusions, only that their outputs “look right”, which of course is not the same thing.

II. **It is impossible for a GAN to create an original idea** because it is only ever trying to create conformity to its original training data, no matter how you word your creative prompt.

WHAT THEY CAN'T DO

The **inability to know, understand, and comprehend** severely limits these tools, in comparison to how they are sold to us. This lack of gnosis makes them profoundly inappropriate for what they're being pushed as solutions for:

1. Neither technology is appropriate to use to summarize search results, or even single texts, because they are attempting to summarize large volumes of data into a very small answer (itself a difficult problem with gnosis), without any conceptualization of the truth or falsehood of the source material. This is why the internet is rife with examples of the google or bing search AI summaries concluding that the year is in fact several years in the future, or that something sourced from singular Beaverton or Onion articles is the truth of the matter.

2. Neither technology is good for answering questions in general. While a GAN with a rigorous training set might have some uses in analyzing highly specific questions against highly constrained areas of expertise, the lack of gnosis means that any answer given needs to be taken with a huge grain of salt.

3. Neither are any good for replacing knowledge or skill on the part of human users in the pursuit of an activity, in any of the applications where this has been suggested. They are especially not appropriate for either accuracy-bound tasks (like code generation) or creativity-bound tasks (like producing creative writing or artwork), **because their lack of understanding needs to be offset by skilled professionals who understand their outputs.**

An example I like to use, because it suits my expertise, is the inclusion of Microsoft Copilot AI in Visual Studio Code, and its willingness to generate code for a user based on a description of what they want the code to do. The ability to type code into a text file is roughly 10% of the skillset required to develop software. The vast majority is actually

in understanding the program architecture, the way design decisions will impact the ability to scale the software's function up or down in accordance with the need, and of course the accurate, secure, or correct nature of the code. The short version is, to be any good at coding, you need to be very good at spotting and solving bugs (and the opportunity for bugs).

Because an LLM-based code generator will be very good at creating code that "looks" correct, the threshold for spotting these bugs is actually raised. This is not a problem, necessarily, for a senior and well-seasoned coder, familiar with the type of application being built and the languages being used. It is, alternatively, an enormous problem for the novice, or for someone using a language they are less familiar with, or solving a problem that is novel to them.

A special caution is deserved for any application where it's encouraged that you offload some part of cognition to the AI. Having AI write (or grade!) assignments for your classes or reports for your employer being prime examples.

WHAT THEY MIGHT HELP WITH

There are, however, cases where the LLM or GAN might have some use. Accessibility tooling comes to mind straight away. Having an LLM that is skilled at spoken language transcription is profoundly useful for helping to make media, especially older media, which is not adequately captioned more accessible. The ability to use such technology for real-time captioning could make a status quo where most livestreams and meetings have closed captioning realistic. Another possibility is producing alt text for images so that screen reader users have a better-than-nothing fallback, though they'd never replace human-generated alt-text.

Image-parsing AI might be useful for finding aberrations in medical scans or other diagnostic imaging that, supported by human expertise, might be helpful in diagnosing medical or engineering issues.

There are also interesting applications for these tools in areas where accuracy and integrity might not matter as much. A few popular RPGs have mods available that use GPT-3 (a related technology) to make it possible for the player to have characters respond to any prompt provided. More than one TTRPG company has suggested the possibility of using an LLM as a portable GameMaster, making it possible to play such games alone.

EVEN STILL, THERE ARE OTHER PROBLEMS

This little manual has been a warning to you about one specific problem the sudden mass adoption of these “AI” technologies: the fact that they aren’t Intelligent systems even in the most generous possible consideration, and the problem of their inability to understand, and how that relates to their utility.

We haven’t discussed two other major problems that face the technology, both of which are based more on ethics than fitness for service:

1. **There are serious ethical concerns about how the data to train the models was collected.** In many cases, AI companies have collected their training data by simply scraping the web as widely as they could and consuming things regardless of any attached licensing or real or implied permissions. This is such a problem that people who self-host their websites now have to include special systems for tar-pitting the AI scrapers, in order to conserve bandwidth. Essentially, these models were trained by taking photos without permission inside art galleries. The people who generated the original training data have been compensated in no way, which is an issue especially prominent for writers and artists whose work is now being (poorly) replicated by these systems.

2. **The energy requirements to train these models** (and in some cases, to host the running models at scale which matches their usage) **have serious environmental impacts.** We’re building tools that use way more power than they are worth and we’re burning that carbon and consuming that water (for data centre coolant) in generally-impooverished areas. Other tools may solve the same problems more efficiently.

So, even if you think you have a reasonable application for these technologies, these ethical problems need to be solved in your implementation, and I’m leaving that as an exercise for you, dear reader.

About This Manual

The Technology Survival Manual series is created and maintained by Arcana Labs to share knowledge we find to be important to foster an attitude and skillset of self-sufficiency, freedom, and cooperation with regard to technology.

To check for newer versions, visit:

<https://www.arcanalabs.ca/writing/zines.html>