
ChatGPT vs. Hybrids: The Future Depends on Our Choices



By Philip M. Parker, INSEAD

Imposing the equivalent of a worldwide moratorium on generative artificial intelligence is absurd. But developers might want to pause and reflect on what they hope to achieve before releasing AI applications to the public.

On 29 March, some 1,000 AI funders, engineers and academics **issued an open letter** calling for an immediate pause on further development of certain forms of generative AI applications, especially those similar in scope to ChatGPT. With irony, some of the letter's signatories are engineers employed by the largest firms investing in artificial intelligence, including Amazon, DeepMind, Google, Meta and Microsoft.

This call for reflection came in the wake of overwhelming media coverage of ChatGPT since the chatbot's launch in November. Unfortunately, the media often fails to distinguish between generative methodologies and artificial intelligence in general. While various forms of AI can be dated back to before the Industrial Revolution, the decreasing cost of computing and the use of public domain software have given rise to more complex methods, especially neural network approaches. There are many challenges facing these

generative AI methodologies, but also many possibilities in terms of the way we (humans) work, learn or access information.

INSEAD's **TotoGEO** AI lab has been applying various generative methodologies to business research, scientific research, educational materials and online searches. Formats have included books, reports, poetry, videos, images, 3D games and fully scaled websites.

In this article and others to follow, I will explain various generative AI approaches, how they are useful, and where they can go wrong, using examples from diverse sectors such as education and business forecasting. Many of the examples are drawn from our lab's experience.

Decoding the AI jargon

AI jargon can be overwhelming, so bear with me. Wikipedia defines generative AI as “a type of AI system capable of generating text, images, or other media in response to prompts. Generative AI systems use generative models **such as** large language models to statistically sample new data based on the training data set that was used to create them.” Note the phrase “such as”. This is where things get messy.

AI methodologies can be divided into four broad categories: robotics, cognitive and symbolic, computational, and control. The first two are often referred to as “good old-fashioned AI”, or GOFAI for short. Cognitive – think decision-support systems used in dashboards, expert systems and other applications across many industries – and symbolic AI – such as WolframAlpha which can interpret mathematical symbols – is programmed to follow rules and mimic human experts.

Computational AI includes various machine learning methods including neural networks and a plethora of multivariate data analytic tools often used to “discover” the rules that expert systems fail to achieve. GPT and similar approaches leverage this relatively new form of AI, which gained momentum starting in the 1980s.

The fourth AI category, machine learning control, optimises objective functions (e.g., maximise fuel efficiency) within constraints (e.g., “never fly through a mountain”), often leveraging insights from the cognitive/symbolic

and computational branches of AI.

The graphic below, taken from a [video](#) summary of AI, highlights a few buzz words across the four branches. From the chart, it is interesting to note that all AI methodologies can and/or have already been used for generative purposes (e.g., a robot splashing paint on a canvass; an expert system analysing scanner data and automatically writing a memo to a marketing manager, etc.).

Chat was largely the domain of the rule-based cognitive and symbolic methods. Translating one language to another also was dominated by GOFAL. Then, researchers starting using computational machine learning (e.g., neural networks like OpenNMT) and achieved largely superior results.

My hypothesis about the future of chat is that it will be based on fast-evolving hybrid approaches; the mechanical robot will eventually use machine learning control, having been programmed to leverage both computational machine learning and a mix of cognitive/symbolic algorithms. Chatbots will also be built to purpose - a bot for poetry will likely not do well for welding.

The Four AI Approaches

	Automation & Robotics	Cognitive/Symbolic Computing	Computational Machine Learning	Machine Learning Control
Inputs	Engineering, industrial, design, physics, robotics, computer science	Linguistics, geography, maths, economics, core disciplines, computer science	Marketing science, data science, statistics, computer science	Operations research, industrial engineering, management science, computer science
Applications	ATMs, welders, toasters, dashboards	Expert systems, chatbots, poems, reports, stock picks	Forecasting models, classifiers, translators, image recognition	HVAC, autopilots, route planners, elevator controls, etc.
Logic	Kinematics, efficiency, likeability	Accepted rules, formulas, theories, logic, graphs	Data-driven, factor analysis, cluster analysis, econometrics, neural networks	Dynamic linear, nonlinear, stochastic, etc., programming
Tools	C++, Assembly, Python, Matlab, C, R, etc.	.net, C#, VBA, R, Python, Lua, Go, JavaScript, Java, etc.	Python, R, Julia, Java, C#/ML.net, SAS, SPSS, etc.	Julia, Python, Matlab, Wolfram, SAS/OR, COIN-OR, R, etc.

Generative AI: A brief history

If AI is defined as any non-human method that mimics human intelligence, then the 1967 Cal-Tech handheld calculator must be seen as the ultimate generative AI. You prompt it with, say $2*2$, and it gives you a human readable answer “= 4”. This is called rule-based AI. It is not rule-based machine learning because the calculator is not learning but following a rule from expert intelligence. In fact, you do not want to program a pocket calculator with a learning algorithm or user feedback as this may lead to incorrect results. Each generative algorithm has a unique purpose and should be held up to the standards of the end user, not the programmer.

Rule-based machine learning, as the phrase implies, ingests a large quantity of data that allows the algorithm to better learn or infer the plethora of rules that are not easily documented. For example, we can replace each word in the sentence “George Washington was the first President of the United States,” with position-of-speech tags (e.g., proper noun = NP0, VVB = Verb, ADJ = Adjective, etc.) = “NP0 NPO VBD ATO AJ0 NN1 PRF ATO AJ0 NN2.”

This so-called natural language processing (NLP) lets us infer grammatical patterns likely to be understood by humans – the more frequent the pattern, the more likely it will be understood by an end user. NLP also allows us to divide sentences at the verb (VBD = “was”) – as humans do – to extract the “right-hand” and “left-hand” knowledge, mirroring English grammatical rules that distinguish between subject and object. IBM Watson, introduced in 2011, is an example of an AI system that uses NLP.

Scraping billions of sentences from digitised sources, this method gives a simple, non-reasoning, knowledge-oriented chatbot, with one sentence yielding two chunks of knowledge:

Q: Who was the first president of the United States? (searches right side)

A: George Washington (returns the left side)

and

Q: Who was George Washington? (searches left side)

A: The first president of the United States (returns right side)

Note, in both cases, the text exists in nature and is plagiarised from some source, though only fragments are returned. No need for a neural network for simple knowledge bots – old-school NLP can do the trick.

Now imagine we replace the text with “synonyms” that can be used for speculation or scientific inference:

Aspirin inhibits headaches in chimpanzees.

becomes

<molecule> inhibits <symptom> in <genomic sequence or species>.

One sentence alone can represent the crystallisation of years of research and millions of dollars (e.g., “the cure for malaria is XYZ”). Since the knowledge is in the public domain, the computer can free ride on billions of dollars of sunk investments reflected in petabytes of text. It can now apply reasoning and forecast what a scientist might speculate. Cool stuff.

The problem with modern chatbots

Neural networks, of which generative pre-trained transformers or the now-famous GPT is one class, are useful when the exact rules are not documented enough for programmers to get the desired results. They are also useful when the rules or strategies to implement the rules are simply not well understood or too costly to program compared to cognitive approaches. Deep generative models and the underlying algorithms confidently make what appear to be educated guesses. But they may not follow known rules even in cases where the rules are fully understood. That is a problem.

If users know the rules, and the AI fails to follow them, the output can be construed as “disinformation”, “deception”, “lying”, “propaganda” or worse, “social manipulation”. If programmers **add rule-based logic on top of the “black box”** then concerns over manipulation are amplified. Bad advice or instructions rendered without attribution or caveats might kill someone or cause a riot.

Let's compare two didactic or instructional limericks generated from different methodologies, but with the same prompt: "love" (e.g., write a didactic limerick about love). One uses black-box deep learning, the other uses rules. The imposition of "didactic" is to emphasise that the poems should teach the reader the meaning or the definition of love. A limerick is a five-line poem, usually humorous and frequently rude or vulgar, in predominantly **anapestic trimeter**. It has a strict rhyme scheme of AABBA, in which the first, second and fifth line rhyme, while the third and fourth lines are shorter and share a different rhyme.

Poem #1: ChatGPT

There once was a love so fine
It shone like a bright sunbeam
It brought joy to every day
And made everything gleam
It was a love that was meant to be, not a scheme

Poem #2: TotoGEO

There was a young girl from the Dove,
And she wanted a word for love.
Like flesh injection?
No, it's affection!
That literate gal from The Dove.

Which poem is better? This may depend on your taste. Which poem is a better limerick? The answer is clear: Poem #2. Why? Because it is a limerick and defines love as affection. The problem of the first generative result, from ChatGPT, is that it does not inform the reader that it does not know how to write limericks. A young reader might infer that Poem #1 is a limerick, and forever be misinformed. The second poem, which follows the known rules, can never be "wrong" as long as the rules themselves are known and programmed, as cognitive AI approaches typically do.

The GPT approach is broad in its abilities but expensive to implement and open to guesswork. A recent media [report](#), for example, concluded that “ChatGPT Is Pretty Bad At Poetry, According To Poets”.

In contrast, the rule-oriented approach is built on a large and general linguistic graph applied to a narrow formulaic writing style. The algorithms for limericks are not the same ones that write sonnets even though they share the same back-end data graph and other code. The method for Poem #2 in the above example is far less expensive to implement than GPT and, like a calculator, is never wrong in what might be expected as an outcome. Visit totopoetry.com and type in a word to get a few examples.

The rules and back-end methodology for the algorithms collectively called “TotoPoetry” (referenced above) have achieved a few milestones: the world’s longest poem; the first unabridged dictionary consisting of definitions written in verse across some 20 genres of poetry; and the largest collection of poetry written by the same source.

Since the data input into TotoPoetry are themselves partially generated from neural networks, the graph-generated text can be seen as a hybrid generative methodology. TotoPoetry creates millions of poems so a second algorithm acts as an editor to choose the best ones to post.

The real battle

What is really going on here is the expectation that output generated by AI should be, well, correct, reliable, unbiased or at least not completely wrong. An autopilot should not try to fly an airplane through a mountain. Likewise, a generative AI platform should not mislead. At the minimum, it should indicate that it might be clueless on a topic it was not trained on. When I asked ChatGPT to write a Wikipedia article about Tinwarrole, here is what I got (note the bot’s assertive tone – likely from rules imposed on the language model):

ChatGPT: Tinwarrole is a fictional substance with no real-world basis. It does not exist in reality and therefore does not have any known properties or uses. It is not a recognised term in any scientific or medical fields, and there is no information available about it. As a result, it is not possible to write a Wikipedia-style article about Tinwarrole, as there is no reliable, verifiable information to include in such an article.

Our lab, leveraging exacting rules, gives a factual (non-fictional) answer, derived from an ASCII file [indexed by Google](#) some 10 years ago:

TotoGEO: Tinwarrole was the name of an African slave aboard the NS de Montserrat. Tinwarrole was registered as a male of 44 years and enslaved in 1818; the NS de Montserrat was a Spanish schooner captained by Idriso de Pasqual. This vessel had a recorded mortality rate of 1 percent. Descendants of Tinwarrole, if any, were freed from slavery in the United States in 1863 following the Emancipation Proclamation.

Generative AI is in its infancy. Things will improve over time. Asking to pause research in this area appears naïve in its lack of specificity. What needs to happen is that firms selling this technology must be cautious in releasing it too early or too widely. Minimal quality standards should be encouraged. Many of the platforms developed at the INSEAD TotoGEO AI Lab are in the hands of selected curators verifying veracity and provenance. Our work has already included the use of generative AI to exacting standards in agriculture and research.

Soon we can expect to see more hybrid AI methodologies that blend deep learning with cognitive, symbolic, control and/or various rule-based approaches. In all cases, engineers can and have introduced strategic biases into the systems produced, leading to calls for transparency, provenance and accountability. If a plane crashes because of a bad autopilot system, the company creating and/or knowingly using such a faulty system may face liabilities.

In the articles to follow in this series, I will explain how these methodologies can take business consulting, education, journalism and search engines to a whole new level.

Find article at

<https://knowledge.insead.edu/strategy/chatgpt-vs-hybrids-future-depends-our-choices>

About the author(s)

Philip M. Parker is a Professor of Marketing at INSEAD and the INSEAD Chaired Professor of Management Science.

About the research

The INSEAD AI Lab was founded on the Fontainebleau campus in 2000. The lab's overarching project is called **TotoGEO**: toto meaning everything in Latin and GEO, an acronym for global education and outreach.

TotoGEO was originally started to create personalised materials for use in executive education, but its projects now span virtually all subject domains, including education, science, agriculture, finance and marketing. It has since created research studies, poetry, crossword books, 3D games, short videos and mobile applications. More recently the lab is generating local newspapers (across "news deserts"), developing first drafts of encyclopedic articles, and fully operational search engines for underserved industries, languages or geographies.

Now based out of INSEAD's Singapore campus, the lab comprises 18 engineers and two editors, and has partnerships with non-profits and private sector investors across three continents. These have included the Bill and Melinda Gates Foundation, the Grameem Foundation, Farmer Voice Radio, and various universities and companies involved in content distribution.

About the series

AI: Disruption and Adaptation

Delve deeper into developments in artificial intelligence, especially the disruptions across value chains. This series examines AI's impact on a range of sectors, including business consulting, education and the media. It also sizes up the regulatory and ethical questions tied to this game-changing technology.