The Path to Artificial General Intelligence

How to spot the right opportunities

Accenture Labs

Machine learning has come a long way

Artificial intelligence (AI) powered by machine learning (ML) is behind many of the incredible advances we've seen in areas like healthcare, retail, manufacturing, and more. From improved cancer screening to more reliable products, advances in ML are already improving many aspects of our lives.

Yet true artificial general intelligence is still some way off

Still, today's AI models are narrow, and typically fail if applied outside the domain for which they were trained. And despite exciting advances in language models from OpenAI's GPT-3 and other transformers, true artificial general intelligence (AGI)—which includes the ability to perform common-sense reasoning like a human—is still a distant prospect.

Advances toward AGI offer business value along the way

The good news: businesses don't need full AGI to see value from improved machine reasoning. Many practical business use cases—from "off script" chatbot conversations to automated regulatory monitoring and much else besides—can be significantly improved with even a light form of common-sense reasoning.

Companies should be looking beyond machine learning alone for the new possibilities opening up as AGI research progresses.

Copyright © 2021 Accenture. All rights reserved.

Think like a human: The promise of AGI

The original charter for artificial intelligence, defined in 1956, was as follows:

> **"Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves."**
>
> Cognitive scientist John McCarthy's original goal for AI defined at the 1956 Dartmouth Conference

Artificial general intelligence is a form of AI that aims to return to this emphasis on replicating broad human cognitive capabilities. Sometimes also called "strong AI", AGI aims to create machines capable of **general intelligence**—the kind we typically associate with broad competence such as common-sense reasoning.

Common-sense reasoning is something that we as humans do all the time. Unlike a conventional machine learning system, we don't need to have been "trained" by seeing 500 examples of large objects speeding toward us to know we need to get out of the way. We just do it because we can immediately infer the consequences of not doing so.

Copyright © 2021 Accenture. All rights reserved.

Achieving this same type of reasoning through AGI has long been a goal, but it's much easier said than done. Consider the image on the left.

If someone asked you whether the paper cup in the picture was garbage, how would you decide? Most likely you would want more information in order to make a decision: is there still coffee in the cup? Has it already been drunk from or not? Is the person in the photo (who the cup likely belongs to) done playing chess? You might not immediately know the answer, but you would know what additional information was needed to make a decision.

A current machine learning system, by comparison, might be able to easily identify the paper cup as a paper cup, but it would have no idea how to answer a question like "Is this paper cup garbage?" This type of problem is something that a true AGI system could handle, but today's machine learning systems aren't designed to address.

In fact, it's the very success of machine learning that has blinded us to the value of AGI. Since machine learning is designed to handle classification problems and does it so well, we've gotten used to construing every problem that way. We rarely stop to consider problems of a different kind altogether: those that require broad reasoning capabilities.

Almost 70 years after the charter was set for AI, it might seem disappointing to note that we're still relatively early on the journey toward true AGI, and that reaching it is likely a way off. But there's good news for business: there's value in the journey itself, and significant untapped potential waiting in the AI systems that are available today.

Figure 1. Easier for machines to win at chess than to decide if this coffee cup is garbage.

Copyright © 2021 Accenture. All rights reserved.

Language models now offer a new approach

In recent years, a series of “language models” developed using deep learning have been released that arguably represent a significant step on that journey. These models analyze vast amounts of text to enable them to “predict” what text is likely to follow a given prompt. This capability can be used for applications like question answering, semantic search, and text generation with results that sound surprisingly natural.

One of the most notable language models is OpenAI’s GPT-3, which is trained on text containing 300 billion words and employs 175 billion parameters (its predecessor, GPT2, had only 1.5 billion parameters). To process this amount of data, OpenAI’s system used 10,000 GPUs and 285,000 CPU cores.

The results can be very convincing. In our own research, for instance, we’ve seen how giving GPT-3 two sample press releases from a company can produce a surprisingly real-sounding “fake” press release. There are many such examples of GPT’s uncanny ability to produce human-like passages of text.

It’s easy to be seduced by GPT-3’s results. But it’s important to recognize it’s not actually communicating anything meaningful. It’s merely producing text that is statistically likely given what it has seen in the past.

So, while AGI research is well ahead of where it was a decade ago, we’re still a long way from seeing full-blown common-sense reasoning. What GPT-3 shows, however, is that language models are producing interim capabilities that are useful not only in advancing AGI research, but also in addressing specific business problems.

Copyright © 2021 Accenture. All rights reserved.

Three example business problems language models might help to solve

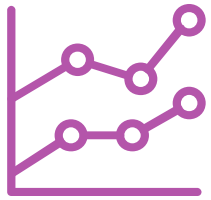

Populating knowledge graphs

Knowledge graphs—which represent collections of interlinked concepts, people, objects, and events—are increasingly important in integrating, analyzing, and sharing data. These graphs form the base of "knowledge" that many automated reasoning systems rely on. Google's knowledge graph, for example, is used to enhance its search results with information about the term searched for.

GPT-3 can be used to overcome one of the big challenges with knowledge graphs: populating them. Consider the problem of contact tracing during a pandemic. The number of contacts a person has is clearly central to this challenge. But there is no readily available resource that can provide that information. However, one data point that is often present is a person's occupation, and this can act as a predictor of their likely number of contacts.

For example, in some preliminary tests we found GPT-3 was able to make qualitative assessments of a person's number of contacts based on their profession (security guard: small; salesperson: large; etc). Of course, these results are not precise. But they could still help prioritize the order in which contact tracing is done. Imperfect but broad results from reasoning applications can complement the deep but narrow results of traditional machine learning applications.

Copyright © 2021 Accenture. All rights reserved.

Localization

One frequent marketing challenge is localization—tailoring content for different geographies. A line like "Stroll down Michigan Avenue then stop for a deep dish pizza in this casual and trendy jacket" would need to be tailored to different locations to be relevant. Among the many things that GPT-3 has learned is the association between locations, typical foods, and landmarks. Given only three examples of the above, GPT-3 is able to start generating localized versions for different cities:

Paris:

Stroll down the Champs-Élysées then stop for a macaroon in this casual and trendy jacket.[1]

Buenos Aires:

Stroll down the Paseo de la Recoleta then stop for a dulce de leche ice cream in this casual and trendy jacket.

Tokyo:

Stroll down a Tokyo backstreet then stop for a bowl of ramen in this casual and trendy jacket.

Milan:

Stroll down the Via Montenapoleone then stop for a cappuccino in this casual and trendy jacket

[1] Note that GPT-3 confused the French macaron with a macaroon here – a nice example of getting us close to a good answer, but perhaps requiring a human to provide a final review and approval.

Copyright © 2021 Accenture. All rights reserved.

Enriching metadata

In practice, important data about a product, transaction, or event is often incomplete or exists only in an unstructured form. This can make it difficult for a machine to process, even if there's enough information present for a human being to make the necessary inferences.

GPT-3 could potentially help here too by using its knowledge to infer and fill in missing metadata. For example, given the following unstructured description of a clothing item, GPT-3 was able to label its function.

Prompt:

> "For sophisticated style with superior comfort, choose the refined modern fit and unrestrictive TH Flex stretch design of these suit separates from Tommy Hilfiger."

GPT-3 label: Business formal

Copyright © 2021 Accenture. All rights reserved.

Do we even need AGI?

There's a lot of untapped potential in today's machine learning approaches. And common-sense reasoning – the hallmark of what we'd call AGI – is already present in every employee and every customer. So do businesses even need machines with common sense? Wouldn't it be common sense to rely on people for common sense?

The answer is that there are many situations where even a small degree of common-sense reasoning can make a big difference to machine operations. As with other areas of automation, the best results are often achieved when humans and machines work together. Let's look at some possible – and practical – situations.

Copyright © 2021 Accenture. All rights reserved.

Better customer service through scriptless conversations

Chatbots are great at answering direct questions from customers. But if the conversation doesn't stay within predefined parameters, even the smartest chatbots will struggle. They also usually fail to understand the broader context of a customer interaction in a way that would be obvious to a human customer service agent.

Consider, for example, a passenger asking an airline's chatbot whether they can bring their crutches on board. Today, a bot would be able to answer that question directly, but no more. By contrast, a human customer service agent would use this as an opportunity to ensure the passenger had the best possible experience given their mobility needs. They'd ask questions like "do you need help with your luggage?", "would you like a wheelchair?", "will you need assistance on board from flight attendants?", and "what about arrangements at the destination airport?"

Accenture Labs has been prototyping a chatbot capable of scriptless interaction. It involves modelling the domain with enough granularity to start capturing fine distinctions. This is what allows the bot to draw the right number of relevant inferences (and ignore the irrelevant) so that it can solve problems rather than just answer narrow questions.

The interaction thus takes a single datapoint ("has crutches") from a single question ("can I take crutches on board?") and quickly moves beyond a simple yes-or-no response to a conversation about enhancing the whole flight experience. There are countless such situations in customer service that we don't stop and think about because common sense is so second nature to humans. But if we can add more of that reasoning to chatbot solutions, they can address a larger proportion of the calls received without having to anticipate every question.

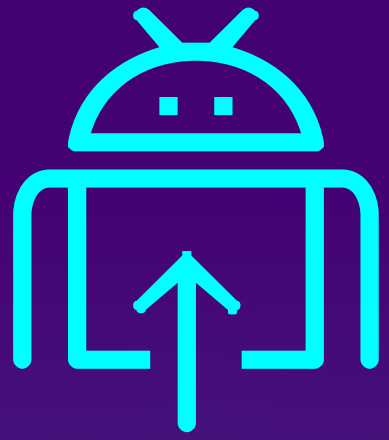

Simple reasoning about a patient's context to inform their diagnosis and treatment

Spicy foods aggravate conditions such as irritable bowel syndrome. Any doctor who is told that a patient suffering from IBS went to a Thai restaurant may reasonably suspect that could be a source of the patient's discomfort – because the doctor knows that Thai food is often spicy. How can a machine make such connections? We could, in principle, have a database of spicy foods. But would we do that for everything that could aggravate such a condition? That's clearly not practical. As discussed earlier, much of this kind of knowledge is implicitly present. By making such knowledge explicit, nobody has to sit down and enumerate a list of all spicy foods. We can try to extract such knowledge from language models for things like spiciness – and other such common attributes. This way we can plug the informal knowledge gaps upon which so much common sense depends.

Copyright © 2021 Accenture. All rights reserved.

Better compliance through continuous monitoring

Many business processes depend on monitoring the news in a particular domain, recognizing relevant classes of events, and making inferences that trigger action.

For example, pharmaceutical companies need to monitor news stories and scientific publications for reports of adverse drug events in order to meet their pharmacovigilance requirements. Manufacturing companies need to stay abreast of potential disruptions to their complex supply chains. And every business in a regulated industry needs to check regularly for changes in the legal and regulatory environment.

The difficulties of getting enough human resources to do this kind of monitoring quickly and comprehensively at scale make it an ideal problem for an AI solution. Accenture Labs is working on a broad framework to allow the development of structured domain-specific event monitoring applications that support a limited amount of reasoning.

In regulatory compliance, for example, we're building a system capable of monitoring relevant news sources and detecting a defined collection of regulatory events (a new reporting threshold, a change in the date a rule becomes effective, and so on). Once detected, these events are integrated into a knowledge graph. That allows us to infer, for example, who needs to be notified.

This is an emerging space, and there's potential to add more common-sense reasoning to the system for use in other business domains. We're confident this kind of intermediate approach will provide significant business value without requiring full AGI.

In all of these examples, the ability to do a light form of common-sense reasoning about the immediate situation could have a major impact on the efficiency and effectiveness of the task at hand.

In other words: you don't need to wait for all-powerful general artificial intelligence to develop significantly more powerful business applications today. You can achieve a great deal with systems that go beyond traditional machine learning and begin to use a little knowledge and reasoning to handle a broader range of use cases.

In fact, most of the time full-blown AGI would be overkill. Think of a restaurant serving its customers. An AGI-powered digital waiter might be able to converse with diners on any topic. But in practice, the vast majority of customers will only want to talk about the menu and the wine list. As with this example, most customer and business needs can be met within a reasonably tightly defined domain. No need for an all-knowing AGI waiter; just an AGI that knows the restaurant and its customers' needs well.

Copyright © 2021 Accenture. All rights reserved.

Decision points:

What are the situations where just a little knowledge would let a system do a lot more?
As the airline chatbot example shows, there's a lot of value to be gained by adding even limited AI knowledge to a system. By reducing the number of common-sense questions that go to humans, like the question about crutches on the plane, companies can free up their human workers to focus on the trickier cases that aren't yet within AI's grasp. That gives workers more time to ensure that customers' needs are fully addressed, through a smooth, comprehensive customer experience. You don't need to automate a process entirely to make an AI investment worthwhile. Identify the situations where a just a little bit of reasoning could free humans up to focus on harder problems.

Where are you limited by a lack of breadth of knowledge, and can you address it using the knowledge implicit in language models?
Think back to the contact tracing example or the spicy food example. We're not going to go out and get databases of spicy food or the number of contacts by profession. But if we could infer that information – albeit imperfectly – and improve our processes, that's worth a look. Imperfect common-sense reasoning is a way that humans deal with uncertainty and novelty. Providing even a little bit of this ability to machines in a responsible way allows us to consider a set of applications and capabilities that have until now been beyond our reach.

Copyright © 2021 Accenture. All rights reserved.

Time to explore the expanding AI frontier

Machine learning and narrow AI solutions have been behind so many of the extraordinary recent advances we've seen in healthcare, customer experience, predictive maintenance, and elsewhere. And they will continue to be a central driver of business value into the future.

As we progress toward the long-term goal of common-sense machine reasoning, AI solutions will gradually expand the breadth of situations they can handle. Emerging language models will improve, as will our ability to extract knowledge from them and apply that knowledge in productive business applications.

Even as true AGI remains a futuristic goal, we will see a host of incremental opportunities for creating value and driving new growth. That's why it's time to embrace the expanding frontier of AI. Even before machines can truly "think like humans," businesses can harness recent advancements toward AGI to transform areas like customer service, regulatory compliance, supply chain management, and address other significant challenges today.

Copyright © 2021 Accenture. All rights reserved.

Authors

Andrew Fano, Ph.D
Managing Director,
Global Director of AI Research
and Development, Accenture Labs
andrew.e.fano@accenture.com

Jean-Luc Chatelain
CTO, Applied Intelligence,
Accenture
jean-luc.chatelain@accenture.com

Shubhashis Sengupta, Ph.D
Managing Director and
Accenture Labs Fellow
shubhashis.sengupta@accenture.com

Vivek A. Khetan
Applied Research Scientist,
Artificial Intelligence,
Accenture Labs
vivek.a.khetan@accenture.com

Copyright © 2021 Accenture. All rights reserved.

About Accenture

Accenture is a global professional services company with leading capabilities in digital, cloud and security. Combining unmatched experience and specialized skills across more than 40 industries, we offer Strategy and Consulting, Interactive, Technology and Operations services—all powered by the world's largest network of Advanced Technology and Intelligent Operations centers. Our 537,000 people deliver on the promise of technology and human ingenuity every day, serving clients in more than 120 countries. We embrace the power of change to create value and shared success for our clients, people, shareholders, partners and communities. Visit us at www.accenture.com.

About Accenture Labs

Accenture Labs incubates and prototypes new concepts through applied R&D projects that are expected to have a significant impact on business and society. Our dedicated team of technologists and researchers work with leaders across the company and external partners to imagine and invent the future. Accenture Labs is located in seven key research hubs around the world: San Francisco, CA; Washington, D.C.; Dublin, Ireland; Sophia Antipolis, France; Herzliya, Israel; Bangalore, India; Shenzhen, China and Nano Labs across the globe. The Labs collaborates extensively with Accenture's network of nearly 400 innovation centers, studios and centers of excellence to deliver cutting edge research, insights and solutions to clients where they operate and live. For more information, please visit www.accenture.com/labs

This content is provided for general information purposes and is not intended to be used in place of consultation with our professional advisors.

Copyright © 2021 Accenture.
All rights reserved. Accenture and its logo are registered trademarks of Accenture.