"Through the looking glass: envisioning new library technologies" understanding artificial intelligence

Peter Fernandez

As technology becomes more sophisticated, there is an increasing demand for flexible operating systems that can automatically respond to emerging challenges. The concept of artificial intelligence (AI) is, therefore, increasingly integral to many of the most exciting technological developments occurring today. Broadly defined, AI is the aspect of computing that enables the machine to take in information about its environment and make choices that help it succeed. This column will explore some of the key concepts behind AI and demonstrate how they apply to emerging technology.

Google and artificial intelligence

To get a sense of the importance of AI, we need look no further than Google. Although many companies are working on AI, Google is of particular interest to libraries because its mission to "organize the world's information and make it universally accessible and useful" shares many similarities with libraries. With that in mind, let us look at an incomplete list of ways that Google has made news with AI in recent months:

- Amit Singhal, the Senior Vice President in charge of search, has retired. His replacement, John Giannandrea, was previously in charge of AI for the company, and he intends to begin to merge the company's research in these two areas (Bergen, 2016).
- The Google-owned project DeepMind developed an AI called AlphaGo that, for the first time, beat a grandmaster at the game of Go (Metz, 2016).
- In the USA, the US National Highway Traffic Safety Administration has declared that AI can officially be considered as the driver of a car (Anthony, 2016).
- Utilizing AI originally intended to help identify photographs, an art show was held that displayed the work "dreams" created by Google's AI, as it subtly modified images to identify and understand underlying patterns (Brueck, 2016) (you can even upload images and experiment for yourself here: http://deepdreamgenerator.com/).
- In part to help with widespread misunderstandings, Google has launched a free online course on the topic of the machine learning that helps power many of its applications, including Google Voice Search, Inbox by Gmail, Google Photos and Google Translate (Novet, 2016).

This list gives a sense of the scope of potential applications for AI, which is increasingly integral to many new technologies, from mobile operating systems to efforts to make sense of big data. To appreciate how AI came to occupy its current role, it is helpful to first look at the history of AI to understand the core concepts that define it.

Understanding artificial intelligence

AI has been contemplated for as long as humans have considered the nature of their own intelligence and pondered how artificial beings might be able to accomplish similar tasks. The idea of AI can be seen as far back as the ancient Greek myth of Galatea. Philosophers and scientists from René Descarte to Alan Turing have posed questions that are still relevant to modern AI.

But in practical terms, AI as we think of it today is inexorably linked to modern computing. Once the first computers were built, programmers immediately set about developing ways to have the computer operate and solve problems independent of direct human input or programming. As early as the 1960s, researchers had introduced the concept of "fuzzy" logic, which enabled computers to go beyond binary logic and simulate degrees of truth to more closely mirror how humans think. However, it was in the 1980s and 1990s that AI truly began to take off. And the reasons why can be instructive as we attempt to understand the role of AI in our modern context.

First, it was in the 1980s that AI began to be used throughout industry to solve specific problems – for example, helping to fulfill orders by selecting products based on customer specifications. In these instances, the relatively difficult quest to develop human-like AI was usually set aside in favor of building programs that could help a human accomplish a task more quickly or perform a calculation that would have taken humans even longer to complete.

Even today, the idea of AI can bring to mind ideas about computers that operate like humans, and as we shall see, researchers continue to try and replicate the human brain in a variety of ways. Known as "strong" AI, it attempts to mimic what we know about the human brain. Efforts to create strong AI require not only that we develop more sophisticated computers but also that we understand human intelligence to better replicate it (Warwick, 2012). But these efforts are in their infancy compared to what has been accomplished by using computers to do things humans are not particularly good at and leveraging the unique abilities of AI. This "weak" AI is focused on solving immediate problems and leveraging the inherent strengths of computers.

These two frames through which to view AI can be helpful, as they reflect true divisions in how problems in AI have been approached. They serve to

highlight the degree to which our conceptual understanding of AI is dependent on our understanding of human intelligence. But it is worth noting that in recent years, many of the most powerful developments have come from someplace in-between these two ideals, combining our increased knowledge of how humans operate without trying to rely exclusively on these techniques.

Since the 1990s, the expansion of AI has been tied directly to the developments in robotics and computing. As technology of all kinds expanded, it created the new arenas, wherein AI could be used to solve problems. Additionally, as computing power expanded, many difficulties in developing AI were able to be solved through raw computing power. Instead of programming a more sophisticated AI, it became possible to have it run calculations multiple times until it found the correct answer.

Now, everything from search engines to fraud detection software to drones to bots in video games are powered by AI. Many major news sites even use AI to generate drafts of basic news summaries. Considering AI in this light yields insight about what kinds of things weak AI has been asked to accomplish. These tasks are often considered the 4D's: "Dangerous, Dirty, Dull, or Difficult". Each of these categories defines an area of work that humans may (or may not) be able to do but would rather not. AI can operate in conditions that humans would prefer not to go or perform calculations that would take humans a long time and be relatively boring.

Modern artificial intelligence

To understand the importance of AI to modern technology, we can look at recent editions of this very column. Cheap cloud-based computers, self-driving cars, the Internet of things and the quantified self are just a few of the concepts that have been featured in recent columns, and these concepts are profoundly affected by the quality of their AI. Perhaps more importantly, many of these technologies will not reach their full potential without even better AI.

For example, many libraries continue to provide readers with advisory services. In its most traditional form, advisory service occurs when library staffs take in information about a reader's preferences and desires, and then, based on the library staff's knowledge of the library's collection, and perhaps aided by some specialized software, they select and recommend titles that may also be of interest to the patron.

Very similar advisory tasks are being conducted by companies such as Netflix and Amazon using AI but with some key differences (Hammond, 2015). These interactions are happening online, and the scale and speed with which they are able to accomplish these tasks is radically advanced. Rather than relying on information solicited from the user, these online advisory services generally use data generated by the user's actions on their site. Finally, using the latest generation of AI systems, online sellers are beginning to develop the ability to make increasingly sophisticated recommendations by identifying key factors that tell them about the user's broader preferences. In so doing, they are relying on three important factors that affect modern AI.

Parallel computing

As far back as the 1980s, increased computing power has been used to help overcome many limitations of AI. In recent years, the widespread adoption of parallel processors has been transformative for the development of AI. This is notable because thinking, as humans do it, involves many simultaneous or near-simultaneous processes. When attempting to build AI applications that mimic how human neurons operate, it is extraordinarily advantageous to be able to model this simultaneousness, which was difficult, if not impossible, on traditional processors. For instance, researchers attempting to develop artificial neural networks require that millions of connections be made using cascading possibilities. Each decision (say the advisability of a single chess move) results in new possibilities that need to be analyzed (what subsequent chess moves may be the result). If the computer can track only one decision path at a time, then it takes much longer than if it can simultaneously analyze multiple possibilities at once. The processing time for this kind of calculation took weeks using traditional processors but has been reduced to a single day if parallel processors are used.

Big data

Advances in computing have also made it possible to generate and store massive data sets. This previously impossible task gives AI the advantage in many situations in which it can access and process more information at once than a human ever could. For the most part, this is how early chess AI was programmed (Warwick, 2012). It had some limited ability to make decisions, but its advantage came from being able to access and process possibilities using a huge number of historical chess games to give it an advantage over human players.

This advantage is magnified when AI has access to larger, more robust data sets. Not only does this give AI the advantage over humans when trying to solve the same problem, but more importantly, it makes solving certain kinds of problems possible because an individual human will be able to absorb and comprehend only a relatively small fraction of the amount of data a powerful computer can process. In fact, in the world of big data, it is often only by using algorithms, generated by humans or computers, that these massive data sets can be made at all useful.

To return to the example of reader (or media) advisory service, companies like Amazon have a massive database of information about a given individual user's preferences based on books they have looked at on the site, as well as those they have purchased. If a user has a Kindle e-reader, then Amazon might even know which books the user has read to completion. Additionally, Amazon has data about what other similar users have done in the past. So when Amazon's AI makes a recommendation that you might like a particular author based on the last book you purchased, it is drawn not only on knowledge about you but also from information mined from the millions of other users. Moreover, as it makes recommendations, the AI can learn and make better recommendations going forward. Any individual librarian is going to work with a relatively limited set of users and have limited opportunities to learn and improve. Depending on the circumstance, the resulting recommendation is not necessarily better, human intelligence still has many advantages, particularly in the realm of emotional intelligence. But it is indicative of the ways that AI when paired

with data can allow for increasingly sophisticated analysis that can assist a wider audience than any single human can interact with.

Open source

Google, Facebook, Yahoo and Microsoft, among others, have recently made large parts of their AI available open source, so that it can be downloaded and used by anyone (Shafto, 2016). While this sharing may seem strange at first, it helps reveal some fundamental facts about the future of AI. First and foremost, it demonstrates that AI will increasingly be a platform that connects a whole host of services. Companies are competing to be the platform that others use. Therefore, the algorithms in specific applications built on these open-source versions will only enhance the underlying AI.

Partially, open source AI is business philosophy. But it is also a direct implication of the technology itself. The AI systems that are being released in this way are usually strong AI. That is, they attempt to mimic the human brain and increasingly engage in what is known as "deep learning" or "machine learning". Rather than depending on custom algorithms written entirely by humans and modified by the data the AI has access to, machine learning attempts to give the AI the tools to teach itself.

Take, for example, the recent victory of AlphaGo over a grandmaster Go player. Go is a game with orders of magnitude more complex then chess. In each turn of chess, a player has only about 20 possible moves, whereas in Go, players have approximately 200 possible moves. Previous estimates suggested that it would take another 10 years before AI would be able play at a high enough level to defeat a grandmaster. By using machine learning, however, AlphaGo was able to learn to play Go from basic rules that were not custom-built for the game of Go. Rather, the rules were designed to teach the computer how to learn more generally, and then, it was given access to information about the specific game of Go (Heath, 2016). This resulted in a more complete understanding of the game than was previously possible.

As these kinds of powerful algorithms becoming increasingly ubiquitous, they transform what is possible to accomplish with AI. They require huge amounts of data and powerful processing, but the results are equally impressive. The nature of this type of programming also helps explain why companies are willing to give up the potential advantages of proprietary AI to make rapid advancements. Eventually, this kind of programming could lead to general-purpose AI that does everything from powering your cell phone's personal assistant to driving your car to recommending the next book for you to check out from the local library. Once that starts to happen, the company with the advantage will occupy an important space in the new technology frontier.

Concerns

Given the power of these systems, Elon Musk, Bill Gates and Stephan Hawking, among many others, have begun to urge that we develop checks to ensure that this technology is used and developed in only an ethical manner (Sainato, 2015). As Hawking put it, "A superintelligent AI will be extremely good at accomplishing its goals, and if those goals aren't aligned with ours, we're in trouble" (King, 2015). As communities grapple with these ethical dilemmas, they will want to consider, for instance, the privacy implications of an AI that can make connections that were never before possible. Libraries can serve as a resource.

One particularly interesting solution to the problem of how to train ethical AI whose goals align with the general welfare of humans comes from the world of storytelling. Just as AlphaGo is programmed to work from general rules on how to make decisions, researchers at the School of Interactive Computing at the Georgia Institute of Technology believe that they can use stories to help define those rules in the broadest sense (Wagner, 2016). As AI is increasingly able to interpret and understand natural human texts, incorporating our deepest values may be accomplished by imbuing it with our values through stories, rather than distinct rules.

Applications

These developments have particular importance to those of us who care about the role of information in the society. Personal assistants are one of the most ubiquitous applications that currently use primarily weak AI but may be transformed as AI advances. These programs are often part of the operating systems in phones (think Android's Google Now, Siri for iPhones and Cortana for Microsoft) and use the information in the phone to anticipate user needs. They can help us respond to routine e-mails, suggest appointments for your calendar, make note of your location to provide customized search results and even integrate with wearables, such as activity trackers, to provide you with customized information about how you should behave to achieve defined goals.

Already, Amazon has released a personal assistant for home use, titled Amazon Echo. Among other things (including support for third party applications), users can search the Internet using voice or simply receive answers to basic questions (Harding, 2016). Similarly, Samsung is working to integrate Bluetooth headsets with its phones, so users can have a more seamless experience using the phone's personal assistant in a wider variety of situations.

Libraries

For libraries, the question is not so much what technology will be affected, but rather what technology, if any, will remain unaffected by AI. As the technology already in existence filters down and becomes more accessible to libraries, it is easy to see the outlines of how AI could have a transformative effect on our work. Online discovery tools and databases may begin to work in a fundamentally different way if AI assistance becomes commonplace. Is not that the necessity for specialized expertise in searching will necessarily, or even likely, go away. Rather, it is that the kind of inputs necessary to produce the best results may be entirely different in a world of computer-assisted AI and huge data sets.

These same principles will likely apply across the library organization, from training staff to storing materials. Powerful AI will make many routine decisions easier, freeing information professionals to focus on new challenges.

REFERENCES

Anthony, S. (2016), "NHTSA rules that AI can be sole driver of Google's self-driving cars", available at: http://arstechnica.com/cars/2016/02/googles-self-driving-car-ai-

can-be-the-vehicles-legal-driver-us-government-says/ (accessed 4 March 2016).

Bergen, M. (2016), "Google search SVP Amit Singhal retiring, research head taking role", available at: https://recode.net/2016/02/03/google-search-king-amit-singhal-retiring-artificial-intelligence-research-head-taking-over/ (accessed 23 February 2016).

Brueck, H. (2016), "Google's DeepDream computers make thousands as artists", *Fortune*, available at: http://fortune.com/2016/03/01/google-deepdream-art/ (accessed 8 March 2016).

Hammond, K. (2015), "The A.I. ecosystem", available at: www.computerworld.com/article/2918161/emerging-technology/the-ai-ecosystem.html (accessed 7 March 2016).

Harding, X. (2016), "Amazon echo family grows with 2 new, A.I. devices: dot and tap", available at: www.popsci.com/amazon-echo-family-adds-dot-and-tap (accessed 4 March 2016).

Heath, N. (2016), "How Google's AI breakthroughs are putting us on a path to narrow AI", available at: www.techrepublic.com/article/how-googles-ai-breakthroughs-are-putting-us-on-a-path-to-narrow-ai/ (accessed 4 March 2016).

King, H. (2015), "Stephen Hawking answers questions on Reddit AMA", available at: http://money.cnn.com/2015/10/08/technology/stephen-hawking-reddit-ama-answers/index.html (accessed 8 March 2016).

Metz, C. (2016), "In major AI breakthrough, Google system secretly beats top player at the ancient game of go", available at: www.wired.com/2016/01/in-a-huge-breakthrough-googles-ai-beats-a-top-player-at-the-game-of-go/ (accessed 23 February 2016).

Novet, J. (2016), "Google launches a deep learning course on Udacity", available at: http://venturebeat.com/2016/01/21/google-launches-a-deep-learning-course-on-udacity/ (accessed 23 February 2016).

Sainato, M. (2015), "Stephen Hawking, Elon Musk, and Bill Gates warn about artificial intelligence", available at: http://observer.com/2015/08/stephen-hawking-elon-musk-and-bill-gates-warn-about-artificial-intelligence/

Shafto, P. (2016), "Why big tech companies are open-sourcing their AI systems", http://phys.org/news/2016/02-big-tech-companies-open-sourcing-ai.html (accessed 4 March 2016).

Wagner, D. (2016), "How storytelling makes robots, AI more human", *Information Week*, available at: www.informationweek.com/it-life/how-storytelling-makes-robots-ai-more-human/d/d-id/1324384 (accessed 4 March 2016).

Warwick, K. (2012), *Artificial Intelligence: The Basics*, Routledge, New York, NY.

Peter Fernandez (*pfernand@utk.edu*) is based at Research Services Agricultural Sciences and Natural Resources, University of Tennessee, Knoxville, Tennessee, USA.