Dr. Budesh Kanwar
Assistant Professor
Computer Science

Artificial Intelligence and its Future

Introduction

Artificial intelligence refers to the ability for computer technology to conduct a range of complex tasks for humans, it's rapidly becoming a part of everyday life. Currently, there are a variety of apps that incorporate at least some elements of AI technology, from speech recognition apps to apps that can solve complex mathematical equations. However, the complexity of AI apps is continuing to evolve, and the future of AI apps will reveal some remarkable advancements.

From SIRI to self-driving cars, artificial intelligence (AI) is progressing rapidly. While science fiction often portrays AI as robots with human-like characteristics, AI can encompass anything from Google's search algorithms to IBM's Watson to autonomous weapons.

Artificial intelligence today is properly known as narrow AI (or weak AI), in that it is designed to perform a narrow task (e.g. only facial recognition or only internet searches or only driving a car). However, the long-term goal of many researchers is to create general AI (AGI or strong AI). While narrow AI may outperform humans at whatever its specific task is, like playing chess or solving equations, AGI would outperform humans at nearly every cognitive task.

Recently, tech giants like Facebook and Google have created their own AI labs to produce robots capable of learning to play video games without any instructions.

Google now has their own form of AI called **Deep Learning**, which in 2015 doubled in the amount of projects using it. A huge focus right now for startups and giants alike is to learn how to use AI to help better automate a machine's ability to solve problems on the fly.

There's definite signs that machines with artificial intelligence will soon be taking over skilled manual work that now is typically handled by humans. Amazon recently showed interest in getting more robots working in its fulfillment centers by holding a contest called the Amazon Picking Challenge at a prominent robotics conference, which awarded $25,000 to the team that could design a robot to identify and grasp items from a storage shelf as quickly as possible.

While this may seem like a trivial task for us, grasping items haphazardly arranged on a shelf requires intelligence, expert coordination, and an understanding of one's environment — skills that up until this point have been quite difficult for robot kind.

Some of the examples:

Google Will Pave the Way

Google is one of the most visionary technology companies. They continue to unleash new and cutting-edge technologies that set them apart from competitors. "OK Google" is a voice recognition app that enables users to surf the web hands-free, simply by talking into their device. In the future, Google will continue to advance this technology. Google's search engine relies on a complex algorithm, and when searching via voice, it renders results based on specific keywords.

However, they are now beginning to use a technology called RankBrain, which will rely on identifying patterns within human speech to return results based around the user's specific intentions, rather than their words alone. This is useful for people who pose complex queries without specific keywords, as the speech recognition software will still be able to ascertain the types of search results the user really wants.

Improve Medical Care

Technology has long played a vital role in the advancement of medical care. AI has the potential to greatly expedite this process and guarantee a higher level of precision, which is particularly essential in the medical field. The Walter Reed National Military Medical Center has already begun implementing AI technology to their benefit by using it to construct drawings of decision support tools for surgical care. AI will be able to assist with operations, guiding surgeons through the surgery to improve success rates and ensure a higher degree of accuracy, whilst also assisting surgeons in finding the precise tools appropriate for each operation.

As robots enter the healthcare field, medical care will be more precise than ever.

Robots in Amazon

Amazon already employs the use of robots in its fulfillment centers according to the reports of those that have been given behind the scenes peeks. The word is that these centers are already highly automated and sophisticated, with robot workers that transport products between humans as endpoints in a finely tuned system.

Amazon is not the only organization backing up innovations in artificial intelligence. The DARPA Robotics Challenge, funded by the U.S. military, gave participants the lofty goal of designing humanoid robots capable of responding to highly dangerous disaster scenarios.

Photo Organisation

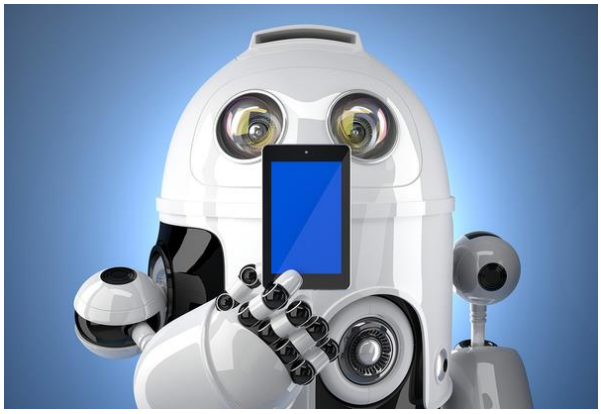

Though it may seem one of the more trivial applications, it can save professional and laymen photographers a lot of time. Facebook has already implemented a degree of AI technology in their photos, enabling facial recognition to automatically identify people in photos based on previous photos of the individual.

The Roll is a newer app that expands on that idea. It can scan an entire photo library and automatically identify similarities between the photos, and then group them into their appropriate albums without any human intervention. It also can apply scores to each photo based on lighting and overall picture quality, which can help professional photographers to continually improve upon their art.

AI is already here, with a plethora of apps designed to simplify everyday tasks and save people time. The technology is rapidly advancing, and it will soon play a significant role in the way apps are developed, with more and more apps containing voice recognition features and other AI technologies.

AI Safety Research

In the long term, an important question is what will happen if the quest for strong AI succeeds and an AI system becomes better than humans at all cognitive tasks. As pointed out by I.J. Good in 1965, designing smarter AI systems is itself a cognitive task. Such a system could potentially undergo recursive self-improvement, triggering an intelligence explosion leaving human intellect far behind. By inventing revolutionary new technologies, such a super intelligence might help us eradicate war, disease, and poverty, and so the creation of strong AI might be the biggest event in human history. Some experts have expressed concern, though, that it might also be the last, unless we learn to align the goals of the AI with ours before it becomes super intelligent.

There are some who question whether strong AI will ever be achieved, and others who insist that the creation of super intelligent AI is guaranteed to be beneficial. Beside both of these possibilities, we should also recognize the potential for an artificial intelligence system to intentionally or unintentionally cause great harm. Research today will help us better prepare for and prevent such potentially negative consequences in the future, thus enjoying the benefits of AI while avoiding pitfalls.

How AI is Dangerous:

Most researchers agree that a super intelligent AI is unlikely to exhibit human emotions like love or hate, and that there is no reason to expect AI to become intentionally benevolent or malevolent. Instead, when considering how AI might become a risk, experts think two scenarios most likely:

1. **The AI is programmed to do something devastating:** Autonomous weapons are artificial intelligence systems that are programmed to kill. In the hands of the wrong person, these weapons could easily cause mass casualties. Moreover, an AI arms race could inadvertently lead to an AI war that also results in mass casualties. To avoid being thwarted by the enemy, these weapons would be designed to be extremely difficult to simply "turn off," so humans could plausibly lose control of such a situation. This risk is one that's present even with narrow AI, but grows as levels of AI intelligence and autonomy increase.
2. **The AI is programmed to do something beneficial, but it develops a destructive method for achieving its goal:** This can happen whenever we fail to fully align the AI's goals with ours, which is strikingly difficult. If you ask an obedient intelligent car to take you to the airport as fast as possible, it might get you there chased by helicopters and covered in vomit, doing not what you wanted but literally what you asked for. If a super intelligent system is tasked with a ambitious geoengineering project, it might wreak havoc with our ecosystem as a side effect, and view human attempts to stop it as a threat to be met.

Stephen Hawking, Elon Musk, Steve Wozniak, Bill Gates, and many other big names in science and technology have recently expressed concern in the media and via open letters about the risks posed by AI, joined by many leading AI researchers.

The idea that the quest for strong AI would ultimately succeed was long thought of as science fiction, centuries or more away. However, thanks to recent breakthroughs, many AI milestones, which experts viewed as decades away merely five years ago, have now been reached, making many experts take seriously the possibility of super intelligence in our lifetime. While some experts still guess that human-level AI is centuries away, most AI researches at the 2015 Puerto Rico Conference guessed that it would happen before 2060. Since it may take decades to complete the required safety research, it is prudent to start it now.

Conclusion:

Because AI has the potential to become more intelligent than any human, we have no surefire way of predicting how it will behave. We can't use past technological developments as much of a basis because we've never created anything that has the ability to, wittingly or unwittingly, outsmart us. The best example of what we could face may be our own evolution. People now control the planet, not because we're the strongest, fastest or biggest, but because we're the smartest. If we're no longer the smartest, are we assured to remain in control?

All above examples illustrate, the concern about advanced AI isn't malevolence but competence. A super-intelligent AI will be extremely good at accomplishing its goals, and if those goals aren't aligned with ours, we have a problem. **You're probably not an evil ant-hater who steps on ants out of malice, but if you're in charge of a hydroelectric green energy project and there's an anthill in the region to be flooded, too bad for the ants. A key goal of AI safety research is to never place humanity in the position of those ants.**

Future of Life Institute's view is that our civilization will flourish as long as we win the race between the growing power of technology and the wisdom with which we manage it. In the case of AI technology, FLI's position is that the best way to win that race is not to impede the former, but to accelerate the latter, by supporting AI safety research.