Artificial Intelligence and Manufacturing

DAVID J. CRANDALL, School of Informatics, Computing, and Engineering, Indiana University, USA

Artificial Intelligence technology is rapidly moving out of the research lab and into products, with the potential to fundamentally transform many facets of business and everyday life. This paper provides a brief overview of AI, including what it is and what it isn't, when it tends to work well, and when it tends to fail. We then specifically review how it could impact the manufacturing sector in particular.

1 INTRODUCTION

After six decades of research, Artificial Intelligence is finally moving out of the lab and into the real world. Computers now out-perform humans on a range of tasks, from everyday games like Chess and Jeopardy [14, 27] to advanced security systems that recognize faces [29] or read lips [11]. Some AI applications are already commonplace — e.g., smartphones that react to voice commands — while others loom large on the horizon — e.g., self-driving vehicles that could forever transform transportation. This excitement has come with hype and many mysteries: why can AI defeat every human Chessmaster that has ever lived, but a state-of-the-art AI-powered mall security robot can clumsily drown itself in a fountain because it didn't see it [15]?

2 WHAT IS AI?

Artificial Intelligence is surprisingly difficult: Dr. Herbert Simon, Nobel and Turing Prize laureate, is reported to have predicted in 1960 that "machines will be capable, within twenty years, of doing any work that a man can do" [19] — a goal that is still elusive some 60 years later. The original dream of AI was to replicate the *way that humans think*, so early AI researchers wrote programs that encoded rules for carrying out tasks and making decisions. This turned out to be impractical: even the simplest of everyday activities, like going for a walk around the mall, involves making innumerable choices, observations, and inferences that seem trivial to us but are extremely difficult to express algorithmically. (Even just deciding whether the path ahead is solid remains a challenge, as the aforementioned security robot learned). Today's work in AI aims for the more modest goal of creating systems that can *perform tasks that seem to require human-level intelligence*. Thus modern AI generally tries to reproduce *what* humans can do, not *how* they do it, through a variety of technologies and approaches.

3 MACHINE LEARNING

Instead of writing programs that explicitly instruct a computer how to carry out a task, many AI systems use *machine learning*. Although often described in grandiose terms of replicating humans' abilities to learn, in practice machine learning is simply about *finding patterns in data, and then using those patterns to make predictions about future data*. Consider the simple example of a robot learning to calculate the circumference of a circle from its radius. To do this, it collects "training data" by drawing circles of different radii, measuring the circumferences with a ruler, and then finding a mathematical relationship between the two. Many possible relationships fit the training data, as shown in Figure 1, each making different predictions about unseen data points — some even predict that the circumference can be negative! A human learner might

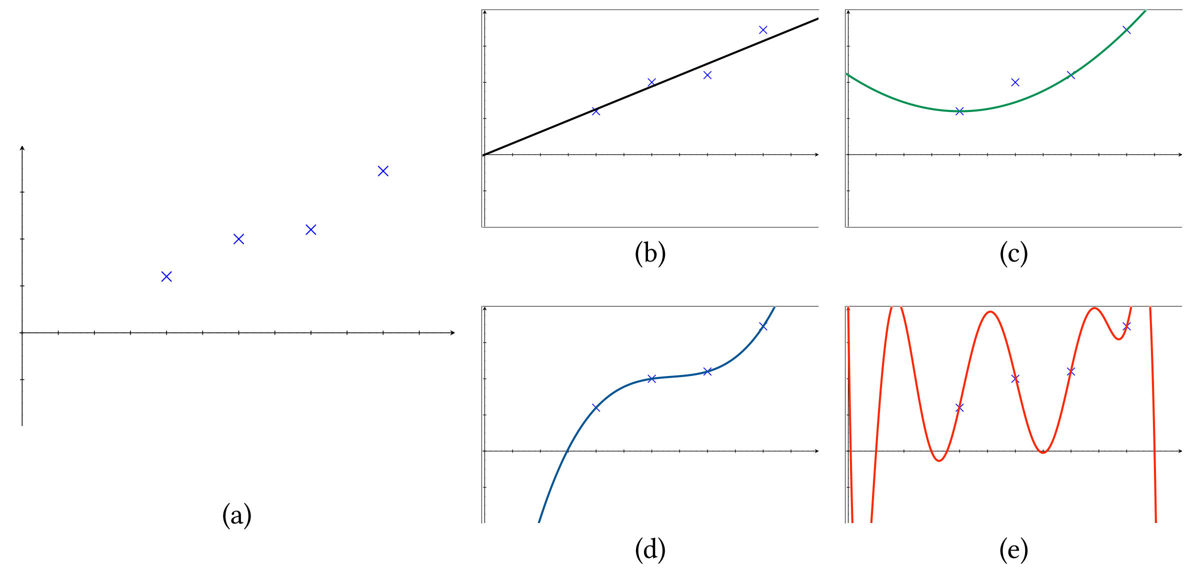

Fig. 1. A simple example of the challenges of fitting a mathematical model to observed training data, in order to describe the relationship between two variables (x and y axes). The same four data points in (a) can be fit with numerous possible models of different shapes and complexities, as shown in plots (b) through (e).

use intuition to choose between them (e.g., discarding models that predict negative circumferences), but algorithms do not have this "common sense." This means that our robot learner may perform very well on its own "training data," but fail spectacularly when it makes predictions about new circles.

Real applications of machine learning involve data that is much more complicated; for face recognition, for example, the data points are not single numbers but images encoded as vectors of millions of numbers. Nevertheless, the basic idea is the same — fitting models to training data. Just like the robot above, machine learning's success is at the mercy of its training data. (This is why face recognition works best on white men [7]: the face training datasets are often of AI researchers themselves, and thus reflect unfortunate biases of STEM demographics.) Ideally, the training set would be large enough to include every possible scenario. In most interesting real-world problems, however, observing all possible scenarios is difficult or impossible: in driving, for example, we *regularly* encounter events that are *individually rare:* a flooded roadway, a child chasing a ball into the street, a mattress flying off of a truck. People make reasonable (if not perfect) decisions even in scenarios they have never encountered before. A major remaining challenge for AI is to build systems that can similarly be trusted to "generalize" beyond the specific training examples that they have seen.

4 WHAT IS THE STATE OF THE ART?

Despite the limitations, AI is still a powerful tool because it turns out that pattern-finding on vast datasets can solve many problems that seem to require intelligence. For example, in 2016 a computer finally beat a human champion in Go, a board game so complex that it was thought to require human-level intelligence [26, 28]. But the algorithm solved it in a different way, by playing and finding patterns in nearly 20 million games — many thousands of times more than any human could play. AI's successes are not limited to games, of course. Figure 2 shows examples of AI recognizing specific species of animals in images, from among nearly

Fig. 2. Sample recognition results from [30], in which the system locates and identifies specific species of animals from among some 3,000 possible classes.

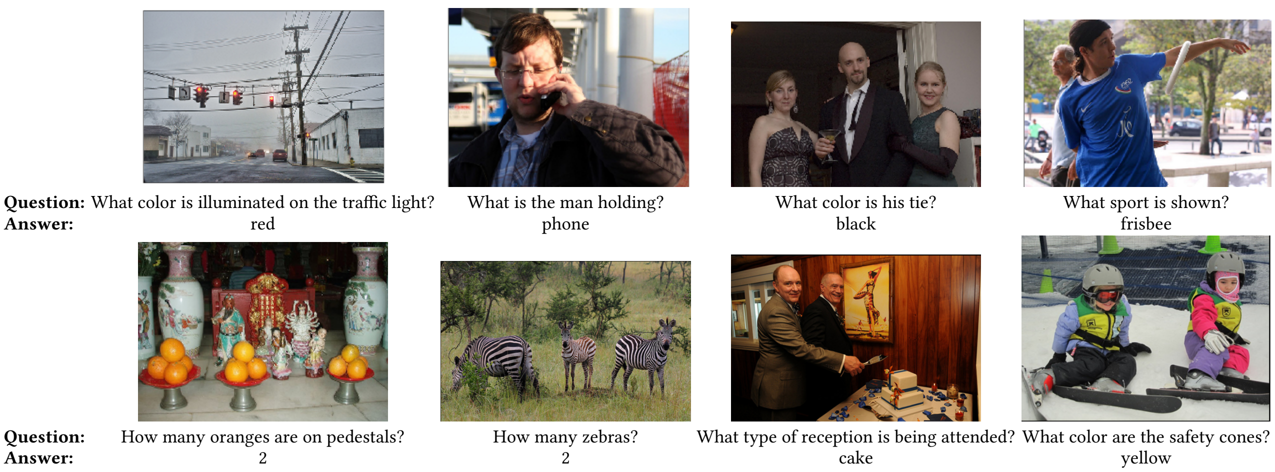

Fig. 3. Sample results of automatic question answering, from [3, 5].

3,000 possibilities; its accuracy of about 60%, while not perfect, is probably far above that of the average person [30]. Figure 3 shows sample outputs from a system that can answer questions about photos [3, 5, 32], which requires AI to solve multiple problems including understanding a question, recognizing photo content, and producing a correct answer.

These examples all use machine learning, and thus depend on fitting mathematical models in millions of images provided as training data. But because these models are not perfect and do not operate the same way that people do, the results they produce can be perplexing or even nonsensical. The second row of Figure 3, for example, shows visual questions that were incorrectly — and largely inexplicably — answered by the algorithm. Figure 4 shows some examples [20] of pairs of images that look nearly identical but are classified very differently by the computer: just like how some of the circle models would predict negative circumferences despite fitting the training data, the model found here correctly classifies many images but fails on images that are only slightly different. Unfortunately, these misclassifications mean that AI systems can be easily fooled: pranksters could modify the appearance of road signs to make them invisible to autonomous vehicles (Figure 5), for example.

Fig. 4. Sample adversarial examples from [20]. The first row shows images and the object that the algorithm recognizes. The second row shows the same images corrupted with noise that is nearly imperceptible, but nevertheless causes the classifier to recognize the wrong object.

Fig. 5. Graffiti patterns that confuse an autonomous car's sign classification system, causing it to recognize the wrong type of sign, from [12].

5 WHEN DOES AI WORK WELL?

Given its imperfections and limitations, deploying AI in the real world must be done carefully. AI works well in applications with very large amounts of high-quality data, and in applications that are resilient to potential errors, typically because (1) a human is "in the loop" to double-check the AI's decisions and intervene if needed, (2) the context or environment is constrained, and/or (3) the consequences of failure are minimal. Board games like Go are perfect for AI because large training datasets are available (the computer can generate many games by just playing against itself), consequences of failure are low, and the environment is constrained by the rules of the game. Autonomous driving, on the other hand, is much more difficult, which explains why many companies are targeting limited scenarios such as requiring a human to have their hands on the wheel to override the system if needed, or working only in controlled scenarios such as Interstate highways.

While computers cannot yet surpass human intelligence, they *can* outperform humans in sheer speed of calculations and ability to search vast amounts of information. They can also be programmed to perform the same task over and over again, impartially and without fatigue. These properties enable new capabilities that would simply not be possible for humans to do alone. For example, AI technology can be used to communicate with hundreds or even thousands of sensors and other devices at a time, collecting data and making decisions in real-time. Such a network of small devices, or "Internet of Things," can range in scale from dozens of sensors in an automobile, to thousands of sensors in a manufacturing plant, to millions of senors in consumer products scattered around the world. These sensors can be collecting many different types of information — video, audio, sensor readings, text, etc. Meanwhile, computers' objective and fast calculations let them quickly make quantifications that would not be possible by a human. For example, instead of simply predicting that an important piece of equiment may soon fail based on sensor readings and other data, AI algorithms could predict the *probability* of failure, and compare the expected cost of repairing it once it fails versus the downtime cost of taking it off-line for predictive maintenance.

6 HOW CAN AI BE APPLIED IN MANUFACTURING?

Many manufacturing applications are well suited to these advantages of AI. For example:

- **Quality inspection:** In the restricted environment of a manufacturing plant, computer vision can perform many inspection tasks more quickly, accurately, and efficiently than a human. For example, an aircraft engine manufacturer recently began applying computer vision to inspect turbofan blades in 3d with micrometer precision [9]. The system checks several hundred properties of a blade in just 15 seconds, which has allowed the manufacturer to inspect *every* blade it manufactures instead of just a random sample. Moreover, the system applies a consistent standard, eliminating variations across different human inspectors. Automated inspection may also significantly improve efficiency of consumer product manufacturing: an automated system adopted by a hot sauce maker checks the placement of labels at a rate of over 1,000 per minute [10]. Many of these systems are custom-designed for one particular inspection task and (unlike a human) are unable to be easily retrained. Machine learning-based approaches may change this; machine learning pioneer Andrew Ng recently announced landing.ai, a start-up which promises more flexible inspection systems.
- **Optimizing supply chains:** AI can be used to collect and monitor fine-grained data along the supply chain, and then manage inventory, predict future demand, spot inefficiencies, etc. For example, Walmart is testing indoor drones to monitor its warehouse inventories [2]. It also uses machine learning to forecast product demand based on local weather, for example, and has discovered subtle patterns that may not have occurred to a human forecaster (e.g., that steaks sell better than ground beef when it is cloudy and windy) [21]. The algorithms are not able to explain *why* these patterns occur, or even if they are reliable patterns or simply coincidences, but this is acceptable in this application: the consequences of a few incorrect predictions are minor as long as the system improves efficiency overall. (This is unlike, say, autonomous weapons where explaining why the system chose a particular target would be crucial.)
- **Fine-grained equipment monitoring and predictive maintenance:** AI can monitor manufacturing equipment at a very fine grain through hundreds of networked sensors, picking up on subtle changes — e.g. greater than usual vibration, or slight changes in machine noise — that may indicate impending failure. Mueller Industries, a manufacturer of industrial products, is testing such a system, and already identified a problem with bearings on one of its machines that could have caused significant downtime if it had not been discovered and repaired [8]. This technology has the potential to move from "preventative maintenance" to "predictive maintenance," avoiding machine downtime both from machine failure and unnecessary preventative maintenance.

- **Advanced robotics:** Robots have long been used in manufacturing, but typically must be custom-built for one particular task, cannot be easily "retrained," and are typically blind to their surroudings, simply performing the same task over and over regardless of what (or who) might be in the way. New technology is starting to allow robots to perceive human activities and safely collaborate with them [18]. Other research is investigating robots that can automatically learn by imitating human actions — which could dramatically reduce development costs — or that can learn on their own by simply "practicing" a task over and over again until they succeed [13]. Most of this work is still in the proof-of-concept stage, but the technology is advancing quickly.
- **Generative design:** AI can be used to simulate how a design would perform in the real world, without physically building it, and then automatically "evolve" modifications until an optimal design is reached. As just one example, Airbus reportedly used generative techniques to create aircraft parts that are significantly lighter than those designed by humans [4].
- **Augmenting human capabilities:** Collaborating humans and AI can potentially perform better and more efficiently than either individually. For example, Augmented Reality (AR) can enhance efficiency by showing workers important information as they perform a task, and allowing them to see views that would not otherwise be possible (e.g., infrared imaging to see in low light). One study reported a 34% improvement in productivity for a worker performing a wiring task when AR glasses were used to guide the process [6].
- **Transportation:** Autonomous vehicles have the potential to revolutionize the world's transportation systems *eventually*, but numerous technical, social, legal, and ethical problems remain before they will likely see widespread consumer use [17]. But autonomous vehicles in more restricted settings, such as manufacturing floors, are already being deployed. Amazon reportedly uses tens of thousands of robots to automatically move products in its warehouses [25], for example. And autonomous long-haul trucking may arrive much sooner than autonomous consumer vehicles, since navigating the restricted setting of Interstate highways is significantly easier than handling all possible roadways [24]. Semi-autonomous trucks with assistive safety features are already becoming commonplace.

7 WHAT IS THE ROAD AHEAD?

Although AI can outdo humans on some very specific tasks, humans still dramatically outperform in practically all real-world tasks requiring intelligence [23]. Moreover, machine learning algorithms require huge training sets, whereas humans can learn with very limited experience. Finally, while humans can offer reasoning to support their conclusions, machine learning is a "black box" that typically cannot explain or defend its answers. It may just be a matter of time before these limitations are solved, or they may be more fundamental. Some believe that human learning is nothing more than a sophisticated version of model fitting [1], while others believe that current AI techniques are inherently "wrong" and could never mimic the complex reasoning that people do [16, 22, 31]. Regardless, AI is advancing rapidly, having achieved milestones that seemed unreachable even a few years ago. Current AI technology can already be usefully applied in many applications, particularly in the manufacturing sector, where data is copious, operating environments are restricted, and trained humans can oversee the automatic systems.

Beyond the technical challenges, AI also raises important legal, ethical, and public policy questions. AI will have to make potentially life-or-death decisions — how should a self-driving car choose between crashing itself and potentially killing its passengers, versus striking a child who has run into the road? To what extent should the algorithms that make such choices be subject to government oversight? How do we assign liability for when AI makes mistakes? How do we safeguard AI systems, to protect both the data they collect and decisions they make from hackers and other adversaries? In general, what protections are needed, if any, to ensure that AI does more good than harm?

REFERENCES

[1] Why humans learn faster than AI — for now. *MIT Technology Review*, March 7, 2018.

[2] Rachel Abrams. Walmart looks to drones to speed distribution. *The New York Times*, June 2, 2016.

[3] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.

[4] Ravi Akella. What generative design is and why it's the future of manufacturing. *New Equipment Digest*, March 16, 2018.

[5] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and VQA. *arXiv:1707.07998*, 2017.

[6] Magid AbrahamMarco Annunziata. Augmented reality is already improving worker performance. *Harvard Business Review*, MARCH 13, 2017.

[7] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency, FAT 2018*, pages 77–91, 2018.

[8] Tim Caldwell. Case study: Mueller moves from preventive to predictive maintenance. *Control Design*, Apr 6, 2017.

[9] Jim Camillo. Robots and Machine Vision Automate Inspection of Jet Engine Parts. *Assembly Magazine*, September 3, 2015.

[10] James Carroll. Machine vision system inspects Tabasco hot sauce products. *Vision Systems Design*, October 10, 2013.

[11] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Lip reading sentences in the wild. *arXiv:1611.05358*, 2016.

[12] Ivan Evtimov, Kevin Eykholt, Earlence Fernandes, Tadayoshi Kohno, Bo Li, Atul Prakash, Amir Rahmati, and Dawn Song. Robust physical-world attacks on machine learning models. *arXiv preprint arXiv:1707.08945*, 2017.

[13] Daniel Faggella. Machine learning in robotics – 5 modern applications. *Techemergence.com*, December 5, 2017.

[14] David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A. Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John Prager, Nico Schlaefer, and Chris Welty. Building watson: An overview of the deepqa project. *AI Magazine*, 31(3), 2010.

[15] Natt Garun. DC security robot quits job by drowning itself in a fountain. *The Verge*, July 17, 2017.

[16] Douglas Hofstadter. The shallowness of google translate. *The Atlantic*, Jan 30, 2018.

[17] Nidhi Kalra and Susan M. Paddock. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A: Policy and Practice*, 94:182 – 193, 2016.

[18] Will Knight. This company tames killer robots. *MIT Technology Review*, June 15, 2018.

[19] Raymond Kurzweil. What is Artificial Intelligence Anyway? *American Scientist*, 73(3), 1985.

[20] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. *arXiv preprint arXiv:1610.08401*, 2016.

[21] Jack Neff. Cloudy with a chance of meatballs: How weather forecast predicts Walmart's sales outlook. *AdAge*, October 27, 2014.

[22] Judea Pearl. Theoretical impediments to machine learning with seven sparks from the causal revolution. *arXiv:1801.04016*, 2018.

[23] P. Jonathon Phillips, Matthew Q. Hill, Jake A. Swindle, and Alice J. O'Toole. Human and algorithm performance on the pasc face recognition challenge. In *IEEE International Conference on Biometrics Theory, Applications and Systems*, 2015.

[24] PwC. *Industrial mobility: How autonomous vehicles can change manufacturing*. 2018.

[25] Sam Shead. Amazon now has 45,000 robots in its warehouses. *Business Insider*, January 3, 2017.

[26] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

[27] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy P. Lillicrap, Karen Simonyan, and Demis Hassabis. Mastering Chess and Shogi by self-play with a general reinforcement learning algorithm. *arXiv 1712.01815*, 2017.

[28] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.

[29] Yaniv Taigman, Ming Yang, Marc'Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1701–1708, 2014.

[30] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The iNaturalist Species Classification and Detection Dataset. *arXiv:1707.06642*, 2018.

[31] James Vincent. Facebook's head of AI wants us to stop using the Terminator to talk about AI: Yann LeCun chats about super-intelligent AI and the future of virtual assistants. *The Verge*, Oct 26, 2017.

[32] Qi Wu, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Ask me anything: Free-form visual question answering based on knowledge from external sources. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4622–4630, 2016.