**Global order and the (mis)perception of powerful AI**

**M.L. Cummings**

There is currently significant worldwide concern that advances in artificial intelligence (AI) could significantly shift the center of military power from the United States (US) to other countries with less democratic principles like China or Russia. A fundamental issue with such discussions is the assumption that AI is actually advanced to the point of dramatically changing how militaries operate, which it may not be. However, it is not clear whether the achievement of such advances is actually important, as it may be to a country's advantage to act as if it has advanced AI capabilities, which is relatively easy to do. Such a pretense could then cause other countries to attempt to emulate potentially unachievable capabilities, at great effort and expense. Thus, the perception of AI prowess may be just as important as having such capabilities.

To further examine such issues, the first question that must be asked is, "What successes AI has achieved, both in commercial terms as well as for militaries worldwide?" To answer this question, AI must be defined more precisely. For the purposes of this effort, AI is defined per the Oxford definition as "the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages."

Given this definition, in the commercial world there has been qualified success across many of these elements. AI has dramatically improved voice recognition, which is now a mainstay of commercial businesses (Brynjolfsson and McAfee 2017). Other successes are more muted. While it is true that computer vision has improved over the past 10 years, particularly for object recognition, the brittleness of underlying machine learning approaches has also become more evident over time.

Figure 1 illustrates the brittleness of computer vision using a deep learning AI algorithm. In three examples, a typical road vehicle (school bus, motor scooter, firetruck) is shown in a normal pose, with 3 other unusual poses, along with the probabilistic estimates of what the underlying computer vision algorithms sees. These results demonstrate that this form of AI is unable to cope with different presentations of the same object, and this is well-known problem in driverless cars. Computer vision problems have been cited as contributing factors in many fatal Tesla crashes (Crowe 2016, Risen 2016) and the death of a pedestrian by an Uber self-driving car (Griggs and Wakabayashi 2018).

The other major area where AI has been heralded as amazingly successful in commercial settings is in game playing, specifically the TV game show Jeopardy and the board games of Alpha Go and Chess (Reddy 2017). While this may seem to be a breakthrough for AI-enabled decision making, the reality is more mundane. Such successes were achieved because the domains of games are deterministic, which means that the number of moves or the number of choices that can be made are *known*, albeit numerous. Computers excel over humans when searching a large space of known options. Where AI is decidedly much less capable is in the presence of uncertainty or in drawing abstract conclusions that require judgment under uncertainty (Cummings 2014). Indeed, Watson, the decision-making engine behind the Jeopardy AI success has been deemed a general failure when it was extended to medical applications (Strickland 2019). Alphabet's DeepMind medical AI is also facing similar questions (Lu 2019).

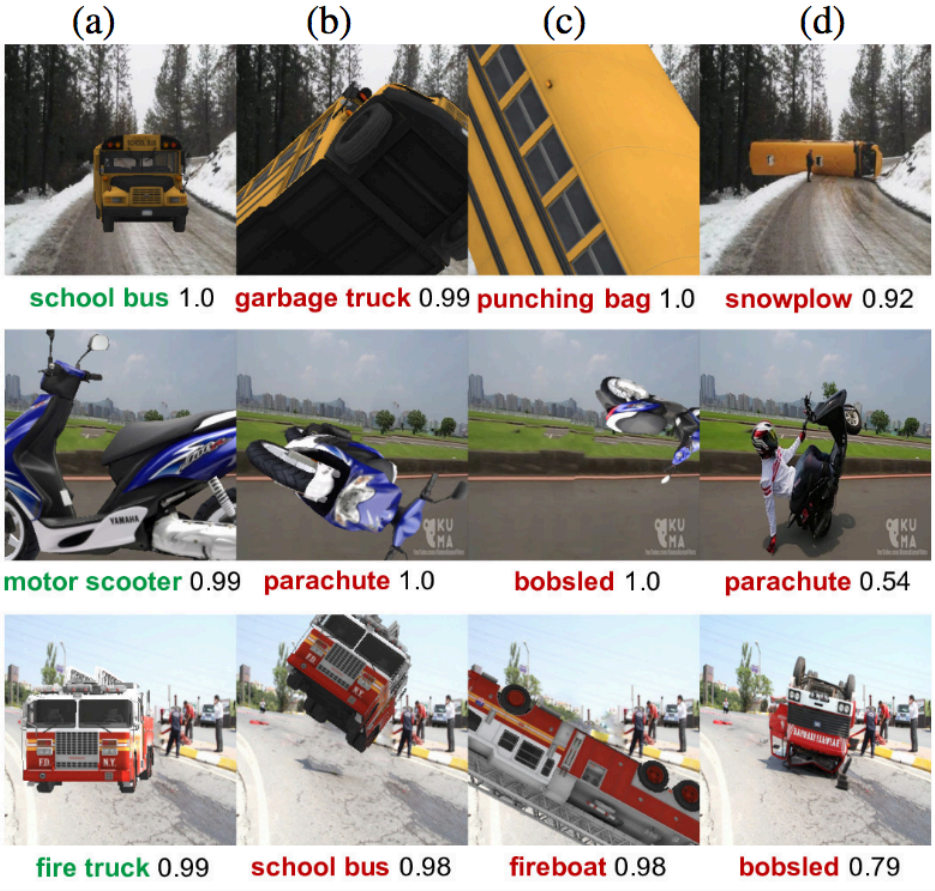


**Figure 1: A deep learning algorithm prediction for typical road vehicle poses in a 3D simulator (a) and for unusual poses (b-d). The computer's estimate of its probability of correctness follows the algorithm's label of what it thinks the object is (Alcorn, Li et al. 2018).**

The inability of AI to handle uncertainty raises serious questions about its success in military settings. The fog of war is the definition of uncertainty and any AI-based system that has to reason about dynamic and uncertain environments is likely to be extremely unreliable, especially in situations never before encountered. Unfortunately, this is exactly the nature of warfare.

So, given the degrees of commercial success of AI, how has AI fared in military settings? There are very few actively deployed military systems that rely on AI. Drones, aka unmanned aerial vehicles, have advanced automated flight control systems but rely on rules-based coding to operate with no AI as defined above. The Tomahawk missile system, which is over 30 years old, uses primitive AI to match digital camera scenes from its onboard camera to images in its database as it flies close to the earth (Larson 1990). While it is highly accurate, it cannot respond to dynamic scene changes and cannot cope with uncertainty.

Automated target recognition such as that in the Tomahawk missile is an area that the US military is keen to use AI to improve. Such a capability would allow weapons systems to detect and potentially destroy targets on their own in real time. While no military publishes exact statistics about such weapons systems, current reports suggest that little progress has been made in this area (Ratches 2011, Boulanin and Verbruggen 2017), undoubtedly due to issues with computer vision as illustrated in Figure 1.

Some consulting groups like Deloitte suggest that the best use of military AI is in readiness development due to known weaknesses in AI. Indeed, using AI to analyze intelligence like satellite images, acoustic data, and logistics information may be able to improve the planning process to better prepare troops for warfare (Strickland, Mariani et al. 2019). This kind of AI-as-a-support-tool approach is quite different than the AI-driven weapon-toting killer robots envisioned by some (Human Rights Watch 2012, Future of Life Institute 2015).

Despite the fact that AI has not been as successful in either military or commercial settings as many people think, it is entirely possible that the perception of having all-powerful AI may be just as important as the existence of all-powerful AI. A major factor driving the perception of who has the most advanced AI is who spends the most on AI. Alphabet has spent more than $2B on DeepMind, which has a reputation as one of the most advanced AI companies in the world. However, DeepMind has produced very little in terms of revenue, and beyond successes in its playing of deterministic games like Alpha Go, DeepMind's supposed successes have been questioned (Marcus 2019, Powles 2019).

The questionable successes of AI matter in the international arms race because there is significant concern that China is outpacing the US in AI applications. But given the significant weaknesses of current AI development, the question must be asked, "Is China really ahead of the US in AI development or has the AI overhype and well-placed demonstrations made us perceive that China is ahead, and what are the ramifications of such a misperception?"

The practice of claiming to have all powerful AI without having actual AI-driven systems is currently an issue in the commercial world vis-à-vis driverless cars. Companies developing driverless cars must rely on humans to significantly augment the computer vision systems through data labelling, which means humans must tell the car what it is seeing (road, bush, pedestrian etc.), with the hope that after enough examples, the car will "learn" these relationships on its own. Figure 1 illustrates just how problematic this approach is and as a result, companies have missed all promised self-driving car deployments (Dennis 2019). To date, no company has demonstrated the ability for sustained driving operations without a safety driver behind the wheel.

This practice of "fake-it-'til-you-make-it" is well known in Silicon Valley and has shown up in other commercial settings like smart email, where hired employees have edited peoples' actual emails to make them think AI accomplished the task, and in voice-to-text translation where call center employees acted as transcription AI (Solon 2018).

The ramifications of the "fake-it-'till-you-make-it" culture in driverless cars has led to inflated and unrealistic expectations that are driving a hypercompetitive first-to-market race, which is quickly becoming prohibitively expensive. More than $100 billion has been spent on driverless car development (Eisenstein 2018), with no end in sight due to the significant problems with computer vision as illustrated in Figure 1. Because of spiraling costs, there have been several company consolidations and partnerships in recent months and there is speculation that many will not survive (Masters 2019). The automotive industry's top investor at Soft Bank has stated "The risks are so big and opportunities so massive that there will be few players that have intellectual capital and financial capital (Weinberg, Tilley et al. 2019)…"

Investments in military AI are escalating just as they are in commercial applications, fueling the concern that China may be outpacing the US in military AI prowess. Indeed, just as the US countered the Soviet Union's conventional military through outspending, particularly in terms of technological advancements, China may be doing the very same to the US through AI investments (Work and Gran 2019).

There is one critical difference in this historical comparison in that physical military systems are tangible illustrations of advancements whereas claims of advanced AI are much harder to verify. For example, in the Cold War, the US Navy was building up to 600 ships (compared to ~400 ships today) and the presence of such ships in ports around the world communicated progress towards this goal. Claims of superior AI are much harder to verify since

they are software-based and as discussed previously, it is not obvious in any AI demonstration whether the results are real, or that team of humans are actually driving the success of an AI system in a Wizard-of-Oz fashion.

Going forward, it is imperative for governments to monitor developments in military-related artificial intelligence, especially for weapons systems and in cybersecurity. However, it is equally important that they arm themselves with the capabilities to detect inflated or faked claims, so as not to invest large sums of money developing a counter-capability to a non-existent threat. Just as in ballistic missile defense where the Chinese use balloons to look like incoming threats to draw scarce counter-missile resources, humans must be able to detect when AI is a balloon or an actual threat in order to determine the most timely and cost-effective response.

**References**


Boyle v. United Technologies Corp., 487 U.S. 588, U.S. Supreme Court 1988.

Alcorn, M. A., Q. Li, Z. Gong, C. Wang, L. Mair, W. S. Ku and A. Nguyen (2018). "Strike (with) a Pose: Neural Networks Are Easily Fooled by Strange Poses of Familiar Objects." arXiv: 1811.11553.

Boulanin, V. and M. Verbruggen (2017). Mapping the Development of Autonomy in Weapon Systems. Solna, Sweden, Stockholm International Peace Research Institute.

Brynjolfsson, E. and A. McAfee. (2017). "The Business of Artificial Intelligence." Retrieved 20 August, 2019, from https://hbr.org/cover-story/2017/07/the-business-of-artificial-intelligence.

Crowe, S. (2016). Tesla Autopilot Causes 2 More Accidents, Robotics Trends.

Cummings, M. L. (2014). "Man vs. Machine or Man + Machine?" IEEE Intelligent Systems **29**(5): 62-69.

Dennis, E. P. (2019). Announced Deployment Timeline. The Center for Automotive Research.

Eisenstein, P. A. (2018). "Not everyone is ready to ride as autonomous vehicles take to the road in ever-increasing numbers." Retrieved 23 August, 2019, from https://www.cnbc.com/2018/10/14/self-driving-cars-take-to-the-road-but-not-everyone-is-ready-to-ride.html.

Future of Life Institute. (2015). "Autonomous Weapons: an Open Letter from AI & Robotics Researchers." 2016, from http://futureoflife.org/open-letter-autonomous-weapons/

Griggs, T. and D. Wakabayashi (2018). How a Self-Driving Uber Killed a Pedestrian in Arizona. New York Times. NY, NY, The New York Times Company.

Human Rights Watch (2012). Arms: New Campaign to Stop Killer Robots, Human Rights Watch **2013**.

Larson, E. V. (1990). Technological Risk: The Case of the Tomahawk Cruise Missile. Santa Monica, RAND.

Lu, D. (2019). It's too soon to tell if DeepMind's medical AI will save any lives. New Scientist. London.

Marcus, G. (2019). DeepMind's Losses and the Future of Artificial Intelligence. Wired. NY, NY, Conde Nast.

Masters, B. (2019). "Self-driving car companies find that going it alone is difficult." Retrieved 23 August, 2019, from https://www.ft.com/content/39c01b56-9be5-11e9-9c06-a4640c9feebb.

Powles, J. (2019). "DeepMind's Latest A.I. Health Breakthrough Has Some Problems." Retrieved 21 August, 2019, from https://onezero.medium.com/deepminds-latest-a-i-health-breakthrough-has-some-problems-5cd14e2c77ef.

Ratches, J. A. (2011). "Review of current aided/automatic target acquisition technology for military target acquisition tasks." Optical Engineering **50**(7).

Reddy, T. (2017). "Why it matters that AI is better than humans at games like Jeopardy." Retrieved 20 August, 2019, from https://www.ibm.com/blogs/watson/2017/06/why-it-matters-that-ai-is-better-than-humans-at-their-own-games/.

Risen, T. (2016). Tesla Updates Radar in Wake of Autonomous Car Crashes, US News & World Report.

Solon, O. (2018). The rise of 'pseudo-AI': how tech firms quietly use humans to do bots' work The Guardian.

Strickland, E. (2019). IBM Watson, Heal Thyself. IEEE Spectrum. Piscataway, NJ, IEEE.

Strickland, F., J. Mariani and I. Jenkins (2019). Military readiness through AI: How technology advances help speed up our defense readiness. Deloitte Center for Government Insights.

Weinberg, C., A. Tilley and K. McLaughlin. (2019). " SoftBank's Ronen Says Self-Driving Market "Big Boys' Game"." The Information Retrieved August 22, 2019, from https://www.theinformation.com/articles/softbanks-ronen-says-self-driving-market-big-boys-game.

Work, R. O. and G. Gran (2019). Beating the Americans at their Own Game: An Offset Strategy with Chinese Characteristics. Washington DC, Center for New American Security.