



COVER STORY

# ARTIFICIAL INTELLIGENCE



## LEARNING ABOUT LEARNING MACHINES

by **Simon Purton**  
Section Head, Decision Support Division  
Headquarters Supreme Allied Commander Transformation





When I saw its moves, I wondered whether any of the moves I have ever known were the right ones. These were the words of Lee Sedol after his defeat while playing Go with Google's DeepMind in 2016. Mr Sedol is an 18-time international Go champion. When it comes to the game of Go, he is as near to superhuman as it gets. To give you an idea of how difficult it is to be good at Go, it is estimated that there are about 1 million trillion trillion trillion more possible configurations of a Go board than a chess one.<sup>1</sup> But a computer programme made Mr Sedol question everything he knew about the game.

So, if you thought Artificial Intelligence (AI) was the stuff of science fiction, then it might be time to re-evaluate your point of view.<sup>2</sup> Whether machines are *really* intelligent like people is debatable, but they are now "smart", and in cases much more so than us. Personally, I think it is important to avoid defining things, for example, whether or not computer programmes are sentient and what AI really is or is not; our inability to agree on such terms did not prevent these technologies from coming into existence and it is not going to prevent them from developing.<sup>3</sup> Bear in mind that Dr Raymond Kurzweil, a noted futurist and leading thinker in the area of AI, believes that the first fully sentient AI will exist by the year 2023, and that your desktop computer will be fully sentient by the end of the 2020s.<sup>4</sup>

Advances mean that learning computers are able to provide timely, accurate, data-based judgments that rival or exceed a human's ability. To make sense of what it is to be artificially "smart", I am splitting this article into three parts. **Part I** is about how people make decisions; because comparing how humans actually make decisions with how we could make decisions, helps define the space where computers are "smarter" than us. **Part II** covers why learning in machines is such a significant leap. And, **Part III** will wrap things up.

## PART I: Human decision-making

Current scientific theory proposes that people make decisions rationally. This idea has its roots in behavioural economic theory.<sup>5</sup> However, "rational" here is different from its more common usage in English; where it would mean "sane" as opposed to "insane". A rational decision-maker (also known as a rational actor) is assumed to rank possible choices by likely benefit and associated cost. They always choose the one that delivers the highest benefit for a cost. Within theory, it is accepted that decision-makers might not have all the information they need, or that all choices may not have been considered. This could result in bad decisions. But, a decision-maker is said to be "rational" in their approach if they always choose the most desirable option from their perspective.<sup>6</sup>

But, to be honest, you would not recognize many rational decision-makers in the real-world. Individuals rarely act in such a manner. We are often faced with time constraints, limited information, "cognitive limitations" (scientist code for our stupidity) and emotions, which means that frequently we do not solve problems optimally. We often rely on "rules

of thumb" to help out when confronted with complex decisions and such short-cuts enable us to make decisions efficiently and quickly. That we routinely take such short-cuts is really where machines can gain the advantage.<sup>7,8</sup>

As an Englishman it gives me no little pleasure to use an example of non-rational decision-making that illustrates the point and shows the frailty of American Football coaches. In a 16-game regular season, the average NFL<sup>9</sup> coach throws away 1.4 wins. Just think about that: there are 16 games, on average season a team only wins eight times. If the right decisions were taken the team would realise 1.4 more victories. So why doesn't this happen?

After spending 15 years in the Navy, Brian Burke founded Advanced Football Analytics as an outlet for his innovative ideas about analysis for the NFL. By 2009, Burke's analysis became a regular feature in the *New York Times*, particularly the "New York Times, 4th Down Bot".<sup>10</sup> This tool analyses what to do when a team reaches fourth down and faces a decision as to whether to punt, attempt a field goal, or "go for it" (for a touchdown or first down).<sup>11</sup> Analysis shows that a coach who followed the "4th Down Bot" would have on average won 1.4 more games a season than one that does not.<sup>12</sup> So the question is really, why do

**BELOW:** Go world champion Lee Sedol of South Korea seen ahead of the first game, the Google DeepMind Challenge Match, against Google's AlphaGo programme. Photo by @Stanislav Varivoda/TASS, Alamy



coaches ignore such advice? And the answer is pretty simple: because they are human.

IF WE UNPACK what is happening, there are really several factors weighing on the coach's mind. What a coach does on 4th down reflects received wisdom and identified best practice—you might refer to this as "doctrine".<sup>13</sup>

In some respects, it is a hold-over from early days of the game, when scoring was less frequent, and punting was a way of ensuring that an opponent would struggle to score on their possession.<sup>14, 15</sup> Coaches also know that if they follow convention and punt on fourth, then any blame falls on the team and not them.

Even if it was a good decision, we live in a world where people who defy convention and fail take all the criticism. The idea about how risks have an impact on decision-making was neatly summarised by Daniel Kahneman when he stated that people generally consider multiple decisions in isolation, instead of aggregating them and looking at the collective outcome.<sup>16, 17</sup> So the coach rejects the 4th down gamble as risky, because people do not consider the collective benefit of a repeated gamble over time.

There is also the relative natural conservatism in people, which is explained by something called "Prospect Theory". Kahneman and Amos Tversky established that people fear a loss much more than they value an equivalent gain.<sup>18, 19</sup> This human built-in tendency toward risk aversion means that coaches are biased toward risk avoiding activities (kicking), rather than risk seeking activities (conversion attempts).

#### IDEA IN BRIEF

When making a decision you act on your preferences and try to make decisions that maximize benefits for you; however there are a number of identified biases that not only undermine your ability to make a good decision, they often end up encouraging you to make bad ones.



**ABOVE:** Captain Chesley "Sully" Sullenberger, best known for the U.S. Airways Flight 1549 emergency water landing on the Hudson River. Photo by @Stanislav Varivoda/TASS, Alamy

## PART II: Machines that learn

"So we hit harder than we would have had we been able to keep the nose up. That was a little-known part of the software that no airline operators or pilots knew about."<sup>20</sup> So said Captain Chesley "Sully" Sullenberger, referring to the software that prevented him from keeping the plane's nose higher during the last four seconds of Flight 1549 as he ditched into the Hudson River.

If you want to get a computer to do something, you have to programme it. Programming usually requires laying out in detail each single step you want the computer to do to achieve the goal, effectively encoding subject matter expertise into a format and language a computer understands.

If you want to programme a computer to fly an airplane you would have to develop a programme in which each relevant item

needs to be accounted for: gear, fuel pumps, fuel quantity, hydraulic pressure, intercoolers, brakes, flaps, throttle, propellers (say), generators, passengers, radio, flight controls. And now you are nearly ready to switch the aircraft on. Then you have to consider all of the possible configurations of the aircraft taking off, in-flight, landing. In different weather. For different terrain. Everything needs to be accounted for, and that covers all the situations or contingencies that may occur, with every combination of activity and environment. This is because you might be three minutes into a flight with 155 passengers on board when your aircraft ingests a flock of geese leaving you at 3,000 feet with no engines!

So, if programming a computer requires access to a great deal of subject matter expertise, what do you do if you want a computer to do something that you hadn't anticipated (such as landing an airliner in a river), or that you didn't know how to do yourself?

**"If you thought Artificial Intelligence was the stuff of science fiction then it might be time to re-evaluate your point of view."**



## IDEA IN BRIEF

Most computers "learn" by being programmed in minute detail. Any failure to "encode" an instruction means that the computer either doesn't know what to do, or it does something really dumb.

This was the challenge faced by a man called Arthur Samuel. In the 1950s, he wanted to get a computer to beat him at the board game checkers. But how do you write a programme to be better than yourself at a game? Samuel came up with an idea: he knew that there were many volumes of annotated checkers games with good moves distinguished from bad ones and so he was able to programme a computer to adjust its criteria for choosing moves so that it would choose better moves as often as possible. He allowed the programme to re-write itself so that it could learn from each game it played. And then, he let it play itself thousands and thousands of times. After a few years, this computer programme was able to beat the Number 4 ranked player in the United States.

This basic idea—using feedback to re-write oneself—underpins the concept of an intelligent learning entity.<sup>21</sup> Ray Kurzweil believes that pattern recognition is not just for sensing the world, but for nearly all aspects of intelligent thought. He concludes that pattern recognition-based learning is critical to the idea of intelligence, and that an AI would only become so through repeated exposure to patterns that encourage it to self-organize, much like a child does.<sup>22</sup>

BY 1950s, COMPUTERS WERE beating us at Checkers. By the 1990s, computers were winning at Othello, and then Chess, when Kasparov lost to IBM's Deep Blue.<sup>23</sup> Earlier, I wrote about Lee Sedol's fate with DeepMind and Go in 2016. So there's an established track record that computers were getting better at playing more complicated games. And while predic-

tions as to when computers would beat humans at these games have not been very accurate, these games, even if they differ in orders of magnitude with respect to their complexity, are similar in type.<sup>24</sup> And they highlight one more human limitation—allow me to explain.

During the 2016 Chess World Championship, the finalists sat down to begin the third game. The game began innocuously enough. But things went weird on move 10. Instead of following a standard approach, World ranked Number 1 Magnus Carlsen made an unusual play leading to a board configuration seen only three times before in recorded history.<sup>25</sup> The game came to a grinding halt with the challenger, Sergey Karjakin, taking over 25 minutes to make his next move. And so the two found themselves playing a sequence of chess moves that no one had ever played before.

It seems strange to me that there are sequences in Chess that people have never played before.<sup>26</sup> Expertise in something comes partly from what is learned and partly from what is experienced. And what we learn is distilled best practice, which is basically other people's good and bad experiences. So it seems interesting that in something as well understood as Chess, there are things that are still new.

What anyone can experience is limited by how much time they have available to practice. Malcolm Gladwell believes that 10,000 hours of deliberate practice is sufficient to be considered to become world-class in any field.<sup>27,28</sup> Finding the time to spend 10,000 hours on a single activity is likely to take a human many years but isn't nearly as difficult for a computer.

Yet, fundamentally, when DeepMind plays Go, it makes moves that no one understands, including the team who made it.<sup>9</sup> It plays Go very differently than people. Which all begs the question: maybe we were actually not very good at playing Checkers, Chess, and Go; perhaps they were only good at playing other people? If that is the case, then that potentially has quite profound implications for us. What was hoped is that there were other types of game where humans might prevail.

I suspect you already know what I am going to tell you about that.

Poker is a dynamic card game in which players are locked in a complicated sequence of acting and reacting to escalating bets. Poker is about as complex as Go, but there is a fundamental difference about the information available to the players. Checkers, Chess and Go are known as "perfect information" games, which means that you get to see everything you need right in front of you when you make your decision. In poker there is hidden information only one player knows.<sup>30,31</sup>

In 2017, a programme called DeepStack defeated a group of professional human poker players in heads-up, no-limit Texas hold'em poker.<sup>32</sup> Before ever playing a real game of poker, DeepStack went through an intensive training period in which it played millions of randomly generated poker scenarios against itself and calculated how beneficial each play was. This reduced the problem space down to 10<sup>14</sup> abstract situations that DeepStack learnt during the period of playing against itself.<sup>33</sup> DeepStack then picks the best move from the reduced problem set, using what researchers compared to human intuition: "a gut feeling of the value of holding any possible private cards in any possible poker situation."<sup>34</sup> It solves all this in under five seconds. Gulp.

## PART III: The rise of the centaurs

Claude Shannon, described as the founder of information theory, did not believe that computers would play games in the same way that humans do. And he did not see computer victories over humans inevitable. He saw relative advantages and disadvantages for both. He contrasted the fact that machines were not emotional or lazy and always analyzed any position, against the fact that they do not have a capacity for imagination or an ability to learn.<sup>35</sup>

The famous Chess grandmaster, Gary Kasparov claimed to have visualized winning positions as far as fifteen moves in advance.<sup>36</sup>

**“Kasparov was able to balance tactical gains and losses against strategic gains; Deep Blue was not.”**





**ABOVE:** Arthur Samuel's Checkers programme, which was developed for play on the "IBM 701". It was demonstrated to the public on television on 24 February 1956. Photo Courtesy of International Business Machines Corporation (IBM), © International Business Machines Corporation.  
**BELOW:** Magnus Carlsen competes against Fabiano Caruana in the World Chess Championship, 28 November 2018. Photo by Bart Lenoir, Shutterstock



When he played IBM's Deep Blue, he was able to contrive trades which although they seemed beneficial for Deep Blue, such as sacrificing a rook for a bishop, it actually demonstrated the limits of the computer's thinking. Kasparov was able to balance tactical gains and losses against strategic gains; Deep Blue was not.

Despite this advantage, he seemed to have been unsettled in the first game (which he won) by Deep Blue's penultimate move. Kasparov had determined that he would win the game, but Deep Blue's response was unusual. In trying to figure out the move, which the Deep Blue programmers later attributed to a glitch, Kasparov concluded that the computer must have been looking at least twenty moves out. Kasparov was not exactly sure what was going on, but he became convinced that Deep Blue was much more sophisticated and smarter than it was. He could not see that the computer had its own limitations and ultimately it affected how he approached playing.<sup>37</sup>

Some of you may have heard of the computer game World of Warcraft. It is a real-time strategy game that puts players in control of a group of units to achieve a variety of goals. Units have different capabilities and costs and are more or less effective against different adversary units. Essentially, it is a fantasy-based version of a strategic wargame. A free-to-play

**BELOW:** Pavilion on Gamescom. Gamescom is a trade fair for video games held annually at the Koelnmesse in Cologne, Germany. Photo by Pe3k, Shutterstock



**ABOVE:** World Champion Garry Kasparov during the third match with Deep Blue, the Chess computer programme, 1997. Photo by ITAR-TASS News Agency, Alamy.

version of this game is known as Defence of the Ancients (Dota).<sup>38</sup> On average half-a-million people are playing Dota 2 at any time, and it peaks over one million. If you were to guess what the prize was for the first placed team at the Dota 2 International Championship, would you have thought it was higher than \$1 million? Last year, "Team Liquid" took home the first prize of \$10.8 million.<sup>39</sup>

Daniil "Dendi" Ishutin is a professional Dota 2 player, winning the team event in 2011, and a runner-up in 2012, and 2013. Dendi was considered one of the world's best Dota 2 players. As part of an exhibition at the 2017 finals, Dendi played OpenAI's Dota 2 bot. The match occurred in front of a live audience and after a bit of back and forth between Dendi and the bot, Dendi went for the kill and lost.

The second game was worse.<sup>40</sup> Dendi described playing against it as "a little bit like a human, but at the same time, like something else." The OpenAI's Dota 2 bot only took two weeks to learn how to beat Dendi.<sup>41</sup> OpenAI's co-founder Greg Brockman said that, "OpenAI's bots play 100 human lifetimes of experience every single day."<sup>42</sup>

By now I understand that you may be a little tired of reading about humans losing to computers. Just bear with me a little longer. Immediately after Dendi's defeat, OpenAI offered a prize for the first 50 players to beat the bot.<sup>43</sup> All of the prizes were collected within 24 hours. Fifty people beat the unbeatable computer programme!<sup>44</sup> When beating OpenAI, players observed, experimented, collaborated, and then adapted.<sup>45</sup>

Obviously, the ability of machines to learn and process enormous amounts of data is incredibly impressive, we are still the apex



predator in the intellectual jungle. It was put recently, "if humans took thousands of years to learn how to play a single video game, we wouldn't be very far as a species."<sup>46</sup>

OpenAI has been testing its Dota 2 bot in team games. And what they have learned is, while the race between humans and machines ebbs and flows, once a human team gains an advantage the machines find it hard to recover. These machines have learned to grind out steady but predictable victories. Once a machine team has lost a lead, they are unable to identify more adventurous or creative ways to regain it. They learn very quickly but struggle to adapt. For humans, it is our creativity and confidence that allows us to see beyond conventional thinking. Yes, we have blind spots in thinking (as discussed in *Part I*). But computers have limitations as well, they can avoid only some of these failures by grinding through a larger number of possible moves.

SIMILARLY, THE UNITED STATES' concept of the "Third-Offset Strategy" postulated that alone machines would always be better than people, but that machines collaborating with people would always beat machines.<sup>47</sup> In his book "Smarter Than You Think", Clive Thompson explores the history of human-machines in Chess.<sup>48</sup> In 2005, there was a Chess tournament in which a team could consist of any number of humans or computers, in any combination. While there were teams consisting of only grand masters and only supercomputers, the winning team comprised of relative amateurs with their computers. He attributes to Kasparov the idea that a new form of Chess intelligence was emerging: a Chess grand master was good; a Chess grand master playing with a laptop was better; but they could all be bested if you could integrate the human with the machine. Kasparov concluded that "human strategic guidance combined with the tactical acuity of a computer was overwhelming." This type of player would later be called a centaur—a hybrid beast endowed with the strengths of each.

Paul Scharre, Senior Fellow and Director of the Technology and National Security Program at the Center for a New American Security, makes the point that the debate in the military tends to presume a choice between human or machines.<sup>49</sup> Whereas in reality in most situations, human-machine teaming in decisions will not only be possible but prefer-

## "We need to defend ourselves in the digital age, and in the age of Artificial Intelligence."

*An extract from the speech by NATO Deputy Secretary General Rose Gottemoeller at the 8th Xiangshan Forum in Beijing on 25 October 2018*



"Today, machines or algorithms can far outperform humans in certain respects—this has become a 'new normal' (...) but while we marvel, we must also learn—we must adapt our thinking and our planning to the realities we face; augmented or otherwise.

"NATO's Strategic Command located in Norfolk Virginia, Allied Command Transformation, is leading NATO's work on innovation and disruptive technologies—big data, Artificial Intelligence, and robotics. This work is key to ensuring that we are ready for the future.

"NATO's member countries are also leading innovation in Artificial Intelligence research and

development. The United States is charting the course—the Defence Advanced Research Projects Agency has pioneered ground-breaking research and development in Artificial Intelligence for more than five decades. France and the United Kingdom allocate significant resources to develop Artificial Intelligence in the defence sector. And a number of other Allies are making advances in this field.

"We know that we cannot fight tomorrow's threats with today's tools. Defending ourselves is no longer about just looking at a map and deciding where to position troops. We need to defend ourselves in the digital age, and in the age of Artificial Intelligence."

able. He says, "hybrid human-machine cognitive architectures will be able to leverage the precision and reliability of autonomy without sacrificing the robustness and flexibility of human intelligence."<sup>50</sup>

Human-machine teaming and in particular creating human-machine trust is going to be key to marrying the tactical superiority of a computer with the strategic innovation of humans to make the best decisions. To paraphrase Clive Thompson: "Which is better: humans or computers? The answer is neither. Instead in the future it will be the two, fused together, like the mythological Centaur."<sup>51</sup>

So, what is being done about all this? Well, finally, NATO is beginning to consider

**ABOVE:** The full transcript of the speech given by Rose Gottemoeller can be found at [https://www.nato.int/cps/en/natohq/opinions\\_160121.htm](https://www.nato.int/cps/en/natohq/opinions_160121.htm)

seriously how to integrate these types of technologies into the enterprise. NATO Deputy Secretary General Rose Gottemoeller noted that, "we need to defend ourselves in the digital age, and in the age of Artificial Intelligence (...) we are thinking about it in terms of both policy and practice."<sup>52</sup>

In NATO, however, we have to overcome many obstacles on the way to developing technologies enabled by AI and machine learning. Some foundational activities are currently underway with efforts coordinated by Allied Command Transformation together with the



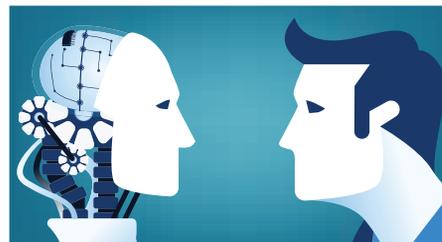
NATO Industry Advisory Group (NIAG) and the Science and Technology Organization, as well as through activities, including TIDE Sprint and Coalition Warrior Interoperability eExercise (CWIX), and Exercise TRIDENT JUNCTURE 2018.<sup>53</sup> There are also more thematic projects based on problems that are benefiting from the use of learning machines.<sup>54</sup>

## Conclusion

I hope this article has given you an insight into what AI might be and why these things are often, but not always, *smarter* than we are. I wrote about why people sometimes don't make the best decisions, and I have previously written on the ACT blog on military decision-making and problem-solving, and what the implications are for NATO Transformation.<sup>55</sup>

I have also written about how learning machines differ from conventional computer programming and how this is being used.

But I left out a lot too, such as details on how the mathematics of AI works, because there are better people to tell you about this. If you are interested, I would encourage you to read Ray Kurzweil's book, a copy is on my office book shelf, but then again so is Kahneman's and Klein's, and since these will help you to understand how you make decisions, they may be a better place to start. But overall, it was never my intention to cover everything, that would have taken a few decently sized books. For example, I intentionally left out ethical, legal, and moral issues. Although these are important, I think that they tap into more personal perspectives, where you would end up reading about what *I* think instead. †



## About the author

**SIMON PURTON** is no expert in AI or learning machines, although he is an enthusiastic amateur. By day, he is an operational analyst working in HQ SACT. He has a Bachelor's degree in Mathematics, is a Fellow of the Institute of Mathematics, a chartered Mathematician, and an Associate Fellow of the Operations Research Society, all of which may account for his poor communication skills. He compiled this article by plagiarism, "re-search" (the process of "Googling stuff"), and by standing on the shoulders of his brighter colleagues. He also read a few books and spoke to some people.

**“Machines learn very quickly, but struggle to adapt. For humans, it is our creativity and confidence that allows us to see beyond conventional thinking.”**



TRIDENT JUNCTURE 2018 LIVEX. A Danish tactical air picture operator during the "War At Sea" Exercise. Photo by Fran C. Valverde, NATO



## ENDNOTES

- 1 In Go, there are over 32,000 possible first moves, around 1,000 of these are considered to be good first moves compared with Chess, which has 400 possible of which around 60 are identified good first moves, around 98% of Chess games start with one of four moves.
- 2 Intelligence is the ability to acquire and apply knowledge and skills.
- 3 From Wikipedia: "Sentience is the capacity to feel, perceive or experience subjectively... [it] is necessary for the ability to suffer, and thus is held to confer certain rights."
- 4 "The Singularity Is Near: When Humans Transcend Biology", 2005, Ray Kurzweil.
- 5 "Rationality", 2007, Blume and Easley.
- 6 Proponents of "rational actor" theory argue that it does not always predict the way individuals make decisions, but it does predict the collective outcome of human decision-making, and so it is "as if" individuals were making decisions in this manner.
- 7 If you want to read more about the way that people make decisions, I would recommend to you Daniel Kahneman's book "Thinking, Fast and Slow". Kahneman differentiates between two types of thinking: fast, which he called System 1 and deliberate, slow thinking, which he called System 2. For more on fast thinking, typified by people who have to make split-second decisions (fire fighters, medics, military, etc.), then I should recommend Gary Klein's book "Sources of Power", in which Klein describes in detail "Recognition Primed Decision-Making".
- 8 Simon note: in fairness, how people make decisions is based on some fairly sound evolutionary reasons. Making decisions quickly, often when things were far from certain was probably a massive survival advantage for our ancestors. Knowing what to eat or what might want to eat us was pretty important and waiting was impossible, dangerous, or not good for our long-term health, either because the "food" would run away or we might become the "food".
- 9 The National Football League (NFL) is the professional American football league and consists of 32 teams. The style of the game, with many breaks and planned plays, has allowed a large body of statistical studies to be developed.
- 10 <https://www.nytimes.com/2014/09/05/upshot/4th-down-when-to-go-for-it-and-why.html>
- 11 In American football, a team has four attempts (known as downs) to move the ball at least ten yards down the pitch. If successful, the team is allowed four more attempts until they score or fail to move the ball ten yards, in either case they give the ball to the other team. What a team should do when it has used three downs and is on its last attempt is the subject of the 4th Down Bot. The options are punting the ball (kicking it a long way into the opponent's territory, usually the safest option), attempting a field goal (if you are within range), or "go for it" (either going for first down or a touchdown). If you fail to convert a field goal or a first down, then the opponent gets the ball at the position of play.
- 12 <http://archive.advancedfootballanalytics.com/2009/08/decision-theory-in-football.html>
- 13 Simon note: I have mixed feelings about the compatibility of doctrinal and innovative thinking. Famously Max Planck said that science advances one funeral at a time: "A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents die." Working within NATO it seems that we innovative at the rate that proponents of old thinking rotate.
- 14 "The Hidden Game of Football", 1988, Carroll, Palmer, and Thorn.
- 15 It is also possible that the value of "going for it" is the unexpected nature of the play, the benefit of which clearly could diminish with continued use.
- 16 "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking", 1993, Kahneman and Lovallo.
- 17 Although what he actually wrote was: "(...) evaluations of single risky prospects neglect the possibilities of pooling risks and are therefore overly timid."
- 18 "Prospect Theory: An Analysis of Decision Under Risk", 1979, Kahneman and Tversky.
- 19 Simon note: It might have been easier for me and better for you to ask Kahneman to write this article himself, but he was too busy collecting Nobel Prizes. This aversion has been calculated to be is about 2:1; so people are generally not interested in taking a gamble until the payoff is approximately double that of any potential loss.
- 20 "Who's Really Flying the Plane?" 26th March 2012, Thom Patterson.
- 21 "How to Create a Mind: The Secret of Human Thought Revealed" 2012, Ray Kurzweil.
- 22 The Blue Brain Project is attempting to create a full brain simulation by 2023. For more information access webpage: <https://bluebrain.epfl.ch/page-56882-en.html>
- 23 Othello: Logistello beat world champion Takeshi Murakami. Chess: Deep Blue beat Gary Kasparov. For the story behind Kasparov and Deep Blue you should read the fantastic article by Nate Silver, "Rage against the Machines" <https://fivethirtyeight.com/features/rage-against-the-machines/>.
- 24 It was thought that computers would not beat a human at Go for another 10 years or more. Read this article from 2006 "Computers just can't seem to get past Go" which tries very hard to convince us that humans will not be bested. <https://www.theguardian.com/technology/2006/aug/03/inside-it.guardianweeklytechnologysection>.
- 25 ChessBase GmbH is a German company that maintains and sells massive databases, containing most historic Chess games. Its annotated database contains more than 7 million games, with over 70,000 with expert commentary.
- 26 Good opening moves are well studied and more familiar to players (and computers) than board positions later on. However, unique board positions are more common (but not common) than you might expect. In the first game between Kasparov and Deep Blue, after he had made three moves and Deep Blue two, the board position had only been recorded once before. This seemed to be a deliberate attempt by Kasparov to confuse Deep Blue.
- 27 "Outliers: The Story of Success", 2008, Gladwell
- 28 Simon note: this number is debated; ten years of full-time effort is another estimate. The exact contribution of "deliberate practice" can also vary by activity see "Deliberate Practice and Performance in Music, Games, Sports, Education, and Professions: A Meta-Analysis", 2014, McNamara, Hambrick, and Oswald.
- 29 And this is in stark contrast with Deep Blue. Deep Blue was written more conventionally, although enormously complex, its programmers could inspect the code to see why Deep Blue had made a specific decision. Modern learning programmes employ Neural Networks and Hidden Markov models. It is much more difficult for programmers to peer into these, while the outcome is visible, the state of the programme at any time is not.
- 30 In an imperfect information game, players are not aware of the actions of other players, although they understand who the players are and what strategies or choices they have available to them.
- 31 Simon note: If you are really interested in imperfect information games and mental torture, try googling a form of Chess called Kriegspiel.
- 32 "DeepStack: Expert-Level Artificial Intelligence in No-Limit Poker," 2017, Moravcik, et al.



- 33 Simon note: it's all relative, I guess!
- 34 "Programming a Computer for Playing Chess" Shannon, Philosophical Magazine, Ser.7, Vol. 41, No. 314 - March 1950.
- 35 Simon note: Here, I'm going to take a liberty. I'm going to say that Shannon did not mean that computers could never learn, but they could not adapt during execution. More on this point later.
- 36 [https://www.huffingtonpost.com/2010/01/22/gary-kasparov-on-chess-me\\_n\\_432043.html](https://www.huffingtonpost.com/2010/01/22/gary-kasparov-on-chess-me_n_432043.html).
- 37 In the second game, Deep Blue was left with a choice of moving its queen or exchanging pawns. Convention would have Deep Blue moving its queen, but it went for the pawn exchange. Kasparov was initially relieved, but the more he thought the more he became convinced that Deep Blue had elected to not move its queen because it had calculated out a better choice and he simply could not figure out what that was. The computer could not have miscalculated. He resigned shortly afterwards "knowing" that Deep Blue was going to win anyway. Post-game analysis showed that Kasparov could have forced a draw. It was later described as an "embarrassing, unprecedented mistake". Grandmasters covering the match "couldn't recall the last time a champion made such an error".
- 38 Dota 2 a computer game developed by Valve Corporation. It was released in July 2013. It is the successor to the very popular Warcraft 3 mod, Defense of the Ancients.
- 39 <http://dota2.prizetrac.kr/international2018>.
- 40 You can watch the games here: <https://www.youtube.com/watch?v=wiOopO9jTZw>; a more edited video can be found here: <https://www.youtube.com/watch?v=5zaJ58q9vuI>.
- 41 Instead of programming the rules of Dota 2, the OpenAI bot is left to figure things out for itself. OpenAI's engineers reward the bot for completing certain tasks but nothing more than that. In the area of machine learning, computers can operate under what is called unsupervised learning. Supervised learning is essentially when we assist a computer programme to develop the correct answers. This is where we codify expert intuition. It can seem a little counter-intuitive that a computer could be left to learn in an unsupervised manner. However, we have all experienced unsupervised learning. No one taught you how to see or how to hear. You learned these skills. For sight, you learned to distinguish edges, lines, colours, contrast, and depth. Given sufficient inputs computers using unsupervised learning can visually recognize objects. They can do so having never been "told" how to see.
- 42 <https://www.theverge.com/2018/6/25/17492918/openai-dota-2-bot-ai-five-5v5-matches>.
- 43 The prize was a rare in-game item.
- 44 Enjoy this description of a winning strategy from a Reddit thread "(...) win by claiming first tower. At 0:00, you aggro the enemy creep wave so that they start following you. Walk around in a circle around and the enemy wave will start to form a [conga] line that will follow you around, so that on the next wave spawn, you can aggro the wave again and continue to walk around in circles. The AI will burn glyph when your creep wave hits the tower, and for some reason it can't really decide between chasing you or defending the town, after about 5 minutes, your creep waves will eventually destroy the tower and you win the 1v1... [figuring this out]... took a lot of experimenting and observation."
- 45 Does that seem vaguely familiar to anyone? In a nutshell, is this not the process of Transformation?
- 46 <https://www.theverge.com/2018/6/25/17492918/openai-dota-2-bot-ai-five-5v5-matches>.
- 47 Memorandum "The Defense Innovation Initiative", November 2014, The Honorable Chuck Hagel.
- 48 "Smarter Than You Think: How Technology is Changing Our Minds for the Better", 2013, Clive Thompson
- 49 "Centaur Warfighting: The False Choice of Humans vs. Automation," 2016, Paul Scharre.
- 50 Ibid.
- 51 Something to consider here is that effective teaming will likely require us to be adaptive when it comes to role-taking. This would require "give and take" between person and machine, not just one-way. Another issue is that we tend to think about encoding human-like judgement into machines, which would be a mistake. We should prefer to ask machines to do things that we were not good at.
- 52 [https://www.nato.int/cps/en/natohq/news\\_160122.htm](https://www.nato.int/cps/en/natohq/news_160122.htm).
- 53 TIDE Sprint is the "Think-Tank for Information and Decision Execution Superiority" (TIDE) and is part of continuum of related events focused on interoperability.
- 54 Information Environment Assessment, Strategic Foresight Analysis, Protection of Civilians, Lessons Learned Modernisation, and Defence Planning are good examples of existing efforts where Artificial Intelligence and machine learning are being employed to solve Alliance problems.
- 55 <https://blog.act.nato.int/wordpress/2016/08/toto-ive-a-feeling-were-not-in-kansas-anymore-a-guide-to-tackling-transformational-problems/>



The Data Centre at Joint Warfare Centre, which is one of the most advanced data centres within NATO. Photo by JWC PAO