Silicon versus carbon: is bridge next?

Ke Jie, the world's undisputed best Go player, just lost 3-0 in his match against AlphaGo, the revolutionary new AI from Google's now world-famous Deep Mind team. What could this mean for bridge players? We try to draw a line from the past to the future…

The past: Deep Blue versus Garry Kasparov, 1997

May 11, 1997. Kasparov resigns the sixth game and loses the match.

People started thinking about chess computers very early on – Alan Turing, the father of computing, wrote a paper about the idea in 1953 – but it was a while before the computers started playing well. This resulted in some unfortunate predictions. The philosopher Hubert Dreyfus published a book in 1972 called *What Computers Can't Do*, where he argued that grandmaster-level chess was an example of a task that was impossible for a computer, even in principle. Dreyfus's reasoning wasn't as silly as it seems now. He looked at what human chess players did, and said that it wasn't a matter of following rules, but of using judgement and understanding: since computer programs can only follow rules and had no judgement or understanding, they can't play high level chess.

In fact, Dreyfus turned out to be wrong about chess requiring judgement and understanding. People play chess that way, because our brains are wired much better for judgement and understanding than for deep calculation, but as computer hardware got faster it turned out deep calculation was also a perfectly good way to play chess. The programmers found that every time the machine's hardware speed doubled, the program's Elo rating went up about 50 points. So for a while, they concentrated on building hardware that could calculate moves as quickly as possible. The culmination of this line of work was Deep Blue, the machine which won the famous match against Garry Kasparov in 1997. Some people said Deep Blue played "just by brute force", which wasn't really fair. The programmers put in some clever ideas. The "evaluation function" (the program's rough estimation of position quality in a position with no immediate tactics) was learned from looking at tens of thousands of master games. Also, when the machine was following its opening book, it thought about how good the players were whose moves it was copying, and was allowed to vary from them if it thought it had calculated a better line. But in the end, Deep Blue was the machine that first beat the World Chess Champion because it could calculate more quickly than any of the others, so "brute force" wasn't totally unreasonable either.

AlphaGo versus Lee Sedol and Ke Jie, 2016-2017

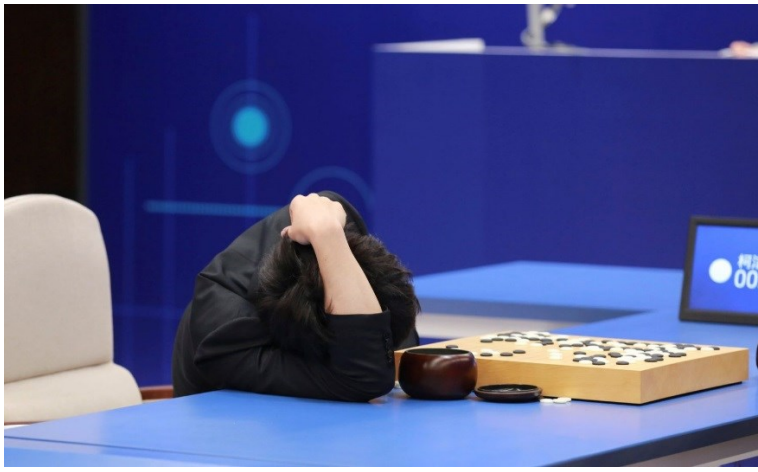

May 25, 2017. Ke Jie visibly shaken after losing the challenge match to AlphaGo.

Strategy game players hated to feel that they were inferior to computers, so they quickly switched over to explaining that chess was the wrong test. It just happened to be the case that chess was simple enough that it could be done by doing a large search. But other games, like Go, wouldn't crack that easily. Things happen much

more slowly in Go, there are far more possible moves at each turn, and the Dreyfus arguments about judgement and understanding got wheeled out again in a slightly different form. Here, Dreyfus seemed to have been right, at least up to a point: you couldn't play Go just by calculating, there was too much to calculate.

But in fact, Dreyfus was wrong, and in a much more interesting and fundamental way than had first appeared. As neural net programming made progress and the Deep Learning revolution started in the early years of the 21st century, it turned out that computers *could* in fact develop judgement and understanding. They didn't need to have their rules painstakingly programmed by human beings: a deep learning neural net in effect develops its own set of rules, by looking at data. The AlphaGo program learned to play by watching master games, then by playing more games against itself and learning from them. Deep Blue had only learned to optimize an evaluation function which had been constructed by people, but AlphaGo went much further. It learned both the "evaluation function", this time more or less on its own, and also the "move generation function", the rules it uses to choose the next move. So AlphaGo thinks in quite a human-like way. It looks at the position, and using its accumulated experience (you might as well call it "judgement") it immediately sees that one of a small number of moves is likely to be the right one. It calculates out a small number of continuations, and at the end of each one it uses its evaluation function (more "judgement") to estimate how well it will be doing if it reaches that position. Then based on those carefully focussed calculations, it makes its choice.

The hard part is putting in the "judgement", which comes from the deep learning. The Deep Mind team, who are world leaders in neural net technology, were able to solve this extremely difficult problem and advance computer Go from weak master level to World Champion level. Last year, AlphaGo beat Lee Sedol, one of the world's three top players, by a decisive 4-1 score. It has just beaten Ke Jie, the undisputed top player, by a clean sweep.

?? versus ??, 20??

Well: is bridge next? We can of course take the Dreyfus line and say that bridge is different because, unlike chess and Go, it requires *real* human judgement and understanding. But given what's happened so far, this seems optimistic. From the point of view of the AI engineer, the thing that makes Bridge hard is that each player has only partial information, so the search space includes all the possible distributions of the unknown cards. That means a lot more to think about. But as we saw with Go, a very large search space doesn't mean that machines can't do it.

There have been a couple of false starts. GIB, which every bridge player knows, was supposed to become the world's best bridge player a little after the Deep Blue breakthrough. GIB can basically do double-dummy analysis perfectly. It handles partial information by generating a hundred or so layouts that fit what it already knows, doing double-dummy on all of them, and then picking the choice which works in the largest number of layouts. It does bidding by using rules that tell it what the allowed bids are in a given situation, generating layouts that fit the bidding, then again making the choice that works in most layouts.

As GIB's inventor Matthew Ginsberg discovered, this doesn't give you more than a strong amateur player. But if you applied deep learning methods and the same kind of hardware as AlphaGo uses (it runs on a network containing hundreds of processors), I think you would see a huge increase in strength. There are plenty of online hand records to train the neural nets. The "move generation function" would be one net, which looks at the current situation and gives you the plausible candidates for next bid or play. The "evaluation function" would be another net, which looks at a layout and estimates how likely each contract is with single-dummy play – basing everything on artificial double-dummy play is one of the reasons why GIB's judgement has never been that great. If you have enough

processors to use, you wouldn't just be limited to creating a hundred layouts to model what you don't know. You could create more layouts to model the other player' uncertainties too, and in effect think about what they are thinking.

Of course, this sketch is simplistic. Building a world-class bridge AI would probably be a big software project that required dozens of person-years of expert effort. But all the pieces now seem to be there. It took 54 years to get from Turing's initial paper on computer chess to Deep Blue, and it took another 20 years to get from Deep Blue to AlphaGo. My guess is that it will take significantly less than 20 years to get to the point where a deep learning system will beat the best human bridge players. It's mainly a question of finding someone who has a strong enough desire to make it happen and enough money to pay for the work. Well: it isn't hard to think of a person who's very rich, has access to hundreds of highly talented AI experts, and likes bridge. I'm starting to wonder why this hasn't already happened.

What might happen to the bridge world, if a world-class bridge AI emerges? Looking at what's happened in chess, it probably would be more good than bad. Since everyone who can afford a basic laptop now has access to a world-class chess player, chess has taken off in many countries where the game was hardly played before. All grandmaster chess tournaments are now broadcast online with reliable real-time computer commentary, so amateurs can follow what's going on. And, a development that might interest bridge players, chess AIs are good at unmasking cheats. Since the machines know what the right move is in most positions, they can spot when someone is playing too well and give statistically significant evidence that something funny is going on. The US chess master and computer expert Ken Regan has been a pioneer in this field.

In fact, when you think more about it, a strong AI might be exactly what bridge needs…

More about neural nets and deep learning

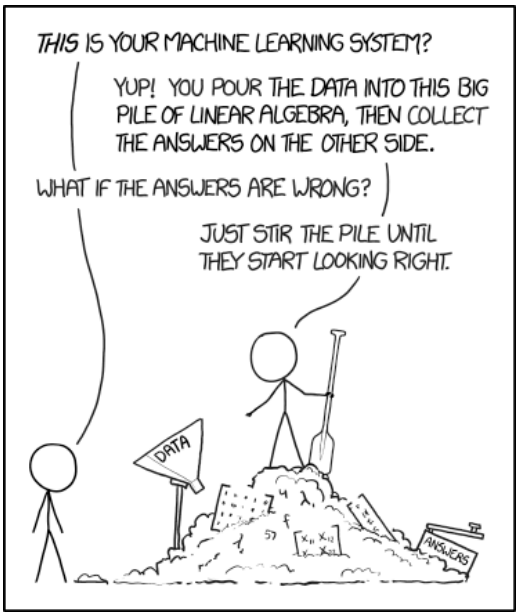

We were going to say that neural nets and deep learning are a highly technical subject that's impossible to explain in a few sentences. Luckily we have Randall Munroe's xkcd strip. He's pretty much nailed it.

Neural nets have been around for a long time. An early success, back in the 90s, was Neurogammon, a neural net backgammon program that became a world-class player. But other things didn't immediately work so well, and neural nets went out of fashion for a while, except in the field of image analysis. About fifteen years ago, people found solutions to some technical problems that had been holding up progress, and then things really took off. Google played a large part in making this happen, and they now use neural nets with deep learning for most of their core business. In particular, it's given them very strong performance in machine translation (Google Translate), speech recognition (voice search) and image recognition (Google inverse image search). Al-

phaGo shows that they're still just starting to exploit this new technology. Another recent success is lip reading. A few years ago, most AI experts were saying that the famous scene from *2001* would be science-fiction for the foreseeable future. It turned out they were too pessimistic.

2023 61st

The Numbers Game - obvious trends

	Pairs					Teams			
	Open	Seniors	Intermediate Intermediate Restricted Novice	Swiss Pairs Swiss Pairs	Total	Open	Seniors	Under 1500 Intermediate Restricted Novice	Total
2006	324	98	156		**578**	244	56	86	**386**
2007	297	91	180		**568**	243	45	103	**391**
2008	332	104	162		**598**	274	48	96	**418**
2009	284	106	180		**570**	246	46	96	**388**
2010	314	102	204		**620**	214	44	138	**396**
2011	304	98	292		**694**	240	46	150	**436**
2012	264	94	282		**640**	220	46	138	**404**
2013	194	88	308	242	**832**	200	52	174	**426**
2014	198	70	312	296	**876**	196	46	184	**426**
2015	198	52	328	320	**898**	208	36	192	**436**
2016	184	52	350	350	**936**	206	34	216	**456**
2017	186	48	326	470	**1030**	214	40	228	**482**
2018	190	56	352	448	**1046**	226	36	202	**464**
2019	201	35	378	406	**1020**	212	20	234	**466**
2020	228	0	396	408	**1032**	236	0	232	**468**
2022	108	0	192	162	**462**	70	0	138	**208**
2023	194	0	268	310	**772**	144	0	192	**336**