HFT Review
High Frequency & Algorithmic Trading

Man Versus Machine, Part One

The first part of a two-part interview with Professor Dave Cliff
By Mike O'Hara, 1st December 2011

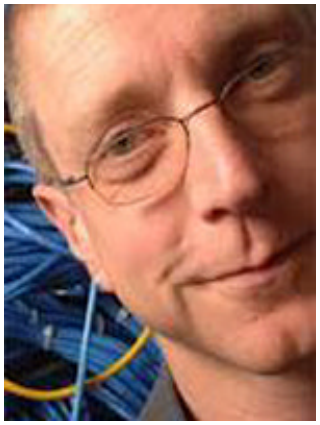

In this interview for the **High Frequency Trading Review**, Mike O'Hara talks to Professor Dave Cliff of the University of Bristol, Director of the UK Large-Scale Complex IT Systems Research Initiative, and a member of the Lead Expert Group of the UK Government's Foresight Project on The Future of Computer Trading in Financial Markets.

During 1998-2005 Professor Cliff worked in industry, as a Department Scientist for Hewlett-Packard Laboratories, and then as a Director/Trader in the Complex Risk Group at Deutsche Bank's London foreign exchange (FX) trading floor. In 1995 Professor Cliff invented an autonomous adaptive trading algorithm, which in 2001 was shown by IBM researchers to consistently beat human traders in experimental versions of financial-market auction systems, and which is now widely recognized as one of the two first autonomous adaptive algorithmic trading systems with real-world applicability.

This is Part One of a two-part interview.

HFT Review: *Can we start with some background on how you got involved in computer-based trading?*

Dave Cliff: The story starts with me doing an undergraduate degree in computer science in the mid-80's, when I got very interested in artificial intelligence (AI). I then went on to do a masters and a PhD in AI at the University of Sussex, when AI was going through something of a revolution.

Working as an academic at the University of Sussex in the mid-90's I didn't earn an awful lot of money but I did have a very fast internet connection in my office. This was the very early days of

the web when most people at home, if they wanted to access the internet, would have to dial up on a 14.4kbps modem, which would take half an hour to download a photo. And yet in my office I had essentially instant access to the web.

At that time, major exchanges like the London Stock Exchange and LIFFE were publishing data on the web for free. So I was day trading on the futures exchange from my office at Sussex, using my desk phone and mobile phone and the T1 line showing me the prices changing. I had no real background as a trader but I took five hundred quid out on my credit card and within about three months, I made so much money that if I drew a graph and extrapolated it out, I realized I could probably retire in about three or four years time. But then within another six weeks, I was back in the red [laughs].

I knew enough about artificial intelligence and about automating processes that a lot of what I was doing could in principle be replaced by machine. Ultimately, I wanted to just write a program that sucked the data off the webpage, traded automatically and made money for me while I was doing other things.

Then, out of the blue, I got a letter from Hewlett Packard Research Labs in Bristol asking if I wanted to spend a few months working there as a visiting academic. But the deal was that I had to choose a research topic that was totally different from anything I'd done before.

Up until that time, I'd been working on computational models using evolution to design new neural networks, which controlled autonomous mobile robots (autonomous in that they had to look after themselves, which meant that you couldn't pre-program them with an understanding of their environment). So I agreed to go to HP on the basis that I would work on an automated trader. In much the same way as my robots dealt with unpredictable environments and therefore had to learn from experience, I would create an artificially intelligent trader that observes what's happening in the market and learn from its experience in order to trade profitably.

HP was happy to let me do that and so I wrote this piece of software called ZIP, *Zero Intelligence Plus*. The intention was for it to be as minimal as possible, so it is a ridiculously simple algorithm, almost embarrassingly so. It's essentially some nested *if-then* rules, the kind of thing that you might type into an Excel spreadsheet macro. And this set of decisions

determines whether the trader should increase or decrease a margin. For each unit it trades, has some notion of the price below which it shouldn't sell or above which it shouldn't buy and that is its limit price. However, the price that it actually quotes into the market as a bid or an offer is different from the limit price because obviously, if you've been told you can buy something and spend no more than ten quid, you want to start low and you might be bidding just one or two pounds. Then gradually, you'll approach towards the ten quid point in order to get the deal, so with each quote you're reducing the margin on the trade. The key innovation I introduced in my ZIP algorithm was that it learned from its experience. So if it made a mistake, it would recognize that mistake and be better the next time it was in the same situation.

***HFTR**: When was this exactly?*

DC: I did the research in 1996 and HP published the results, and the ZIP program code, in 1997. I then went on to do some other things, like DJ-ing and producing algorithmic dance music (but that's another story!)

Fast-forward to 2001, when I started to get a bunch of calls because a team at IBM's Research Labs in the US had just completed the first ever systematic experimental tests of human traders competing against automated, adaptive trading systems. Although IBM had developed their own algorithm called MGD, (Modified Gjerstad Dickhaut), it did the same kind of thing as my ZIP algorithm, using different methods. They had tested out both their MGD and my ZIP against human traders under rigorous experimental conditions and found that both algorithms consistently beat humans, regardless of whether the humans or robots were buyers or sellers. The robots always out-performed the humans.

IBM published their findings at the 2001 IJCAI conference (the International Joint Conference on AI) and although IBM are a pretty conservative company, in the opening paragraphs of this paper they said that this was a result that could have financial implications measured in billions of dollars. I think that implicitly what they were saying was there will always be financial markets and there will always be the institutions (i.e. hedge funds, pension management funds, banks, etc). But the traders that do the business on behalf of those institutions would cease to be human at some point in the future and start to be machines. Now when IBM says something like that, it gets a lot of attention. So IBMs result, and my work, was suddenly appearing in New

Scientist, in The Economist, it was even in the UK's Daily Mail, which is not a newspaper you associate with in-depth scientific reporting!

***HFTR**: Where did your research go from there?*

DC: In 1998 I'd been working at the MIT AI Lab and I'd had the idea of allowing evolution to alter the details of the design of the market mechanism, so the way in which buyers and sellers were matched in the electronic market place. But it was only after the IBM results in 2001 that I re-activated that line of research.

So I ran some experiments that proved the computer could not only design traders better than I could using an evolutionary process, but it could also design market places that were more efficient than the conventional markets you see offered by major exchanges, screen-based markets, or indeed any dark pool.

Essentially the question I asked was what happens if you know that all the traders are robots, pieces of software? Maybe there are better ways of organizing the markets than what is basically an electronic version of some mechanism that was dreamed up in a coffee house off of Lombard Street about 250 years ago.

So that got fairly well known too and on the back of the press coverage, HP's major corporate clients, including major stock exchanges, ECNs and investment banks were talking to us. So I suddenly went from being a jeans and t-shirt guy in HP's research labs to spending an awful lot of time in a suit wandering around the City of London - or indeed New York or Boston – talking to HP's clients in the financial world. I ended up giving the academic keynotes at TradeTech in 2004, which was the very early days of what has since turned into this huge wave of automated trading and HFT

Looking back, perhaps I could have been more diplomatic because I basically stood up and said, "I don't know you. I don't work in a bank or an exchange. I come from a technology company. But banking and exchanges are becoming technology industries. And here is the technology that makes traders redundant." As you can imagine, this got people's attention.

A short time after that, Deutsche Bank London were looking for someone to join their team working on what we would now called high frequency trading in foreign exchange. When I arrived, it wasn't entirely easy for me socially in the sense that I was interacting with the spot FX dealers, safe in the knowledge that I was developing technology that would probably make them redundant.

This was 2005, and the widely quoted figure at the time was that global FX transactions were about $3 trillion a day. London was about a third of that volume. Most of the FX transactions in London were run either by UBS or by Deutsche. And about 70% of Deutsche's spot transactions were run by this one piece of software.

The conversation that I had with my manager at the time, he told me that over the past three years, their number of FX spot traders had has dropped 40% and he expected the number to drop by another 40% in the next three years. If you just do a sales trader's job where an order comes in and you have to work that order, the bottom line is usually there's no reason whatsoever why a machine can't do that.

***HFTR**: But isn't there still a requirement for human intelligence in there? What's your perspective on what can and can't be automated and how do you expect in the future that computers and humans will interact in the markets?*

DC: Well it turns out that at Trade Tech 2004, one of things they did is they videoed all the presentations and all the participants got a DVD. Not long ago, I looked at the video of my talk because I just wanted to check what I said seven or eight years ago. And the example I came up with then I think still holds true, which is what has happened to pilots.

When I was growing up, to be a pilot was a perfectly admirable career choice because you might fly commercial jets or you might fly for the military. Whereas now, maybe it doesn't look such a good career choice. Certainly in the military domain, everyone in aerospace engineering admits that the current generation of combat aircraft are the last that will be carrying human flesh. The reason for that is you have to engineer it very differently if you know there is a human being on board that you have to keep alive, breathing and shielded from weapons. And that human being can't take very heavy G-forces, so you have to moderate the design for that too.

Air combat has always been a job requiring very skilled, very brave human beings. But those people will essentially be redundant within the next 20 to 30 years because the job is now devolved to teams of air force officers who are essentially flying drones by remote control. Most drone strikes in Afghanistan and Iraq are executed by people flying the planes from Nevada. So they're not even ground troops in the field, they're just essentially playing a video game, but it's for real.

So there's a situation where something that certainly 25 or 30 years ago you would have thought would be a job for life is a job that soon will simply no longer exist. In civil aviation it's a little different because anyone who looks at the technology knows that actually, if you sit in a jumbo jet and fly across the Atlantic, you are basically in the belly of a flying robot. As long as a tractor driver can tow the jumbo to the start of the runway, all it really needs is for you or me to press the green "go" button and the plane can take off, navigate to its destination and land. And then all it needs is another tractor driver to tow it to the stand.

In that sense the pilot is redundant, but there are a couple of other senses in which the pilot is certainly not redundant. One being that passengers would be quite nervous knowing that they're trusting their life to a flying machine. So in a way, the pilot exists to keep the customers happy. In fact, I've heard anecdotally that the reason why when you board a plane, the pilot door is open is so that you can see that there are indeed real people in there with lots of gold braid on their cuffs, who look qualified and have authority.

The other thing is that there are still things that machines can't do. Every now and again, a situation will occur where there is no way or it's almost impossible to believe that engineers could have expected such a situation to arise, and to have developed and tested a control system which could deal with that situation safely.

The example I used a lot (because it's so dramatic) is that double bird strike on climb out from JFK a few years ago. A twin engine passenger jet, flying at its most vulnerable point having just left the runway, hits a flock of geese, loses both engines and the pilot who is a Vietnam vet just calmly says, "I've lost both engines and I'm ditching in the Hudson". Then he turns around, lands his plane in the river, nobody dies and it's a great success story.

It's interesting that war is often characterized as 99% boredom and 1% sheer terror. And there's a sense in which you can characterize trading as something similar to that. The tricky bit for technologists interested in automating trading will always be the terror. But the question is, what can we do with currently available technology? I think we can do an awful lot of what a human does during the routine 99% of trades, but the 1%, the terror situations, are still deeply problematic.

***HFTR**: So would you say that humans are very good at the black swan type events whereas machines are very good at handling the more predictable stuff? When things go wrong, when things outside the ordinary happen, humans can gather all of the data at their disposal whereas machines generally will be working on a fixed set of data?*

DC: Yes. The other thing to say is that machines are very, very good at doing things that humans are not good at, such as assimilating and integrating large quantities of numeric data, doing heavyweight calculations on that data and comparing them to some statistical model or prediction of what the number should have been. That's the kind of thing computers are built to do. But people are really good at watching a news screen on a trading floor and thinking, "Oh Lord, that's just happened and because X has just happened, that means Y will then happen and so I should sell A and buy B". That causal logical reasoning about abstract things that aren't easily representable as numbers is something that humans are way, way better at than machines. And there are good reasons to suspect that actually, making those trading decisions is something that humans will always be good at.

I've never argued that the City of London will become a ghost town with just three employees because there will still be a need for humans to make investment and trading decisions. But at the point of execution, you can expect to see many fewer people. In fact already we see far fewer people in that role than we had 10 or 15 years ago. The depopulation of the trading desks probably isn't fully played out yet, though.

But now we get into realms of what's concerned me over the last few years. When I was at Deutsche, I'd go for a beer in the evenings, not only with my colleagues but with colleagues who did similar roles in other banks. And there was a real sense of camaraderie and of people sharing information. So it was common knowledge that every now and again, someone's

automated trading system would go horribly wrong and would set a price that was a giveaway price. And at that point, everyone else's automated trading system would hit that broken robot because it was basically giving away free money and quite often what would happen is that so many traders would start trading at the wrong price that actually it became the market price. And this was for major currency pairs.

In my time at Deutsche, I watched it happen on the news wires, on the data feed for GBP/USD. It would happen for about 30 minutes and the company concerned would lose a lot of money. Then someone would pull the plug out, they'd reprogram it and then hope it wouldn't happen again. But that sort of thing isn't supposed to happen at all. And in particular, global exchange rates aren't supposed to be set by one machine going wrong.

***HFTR**: So now you are working in the field of large scale complex IT systems at the University of Bristol. How did that come about and what does your work there involve?*

DC: In 2005, I decided I was tired of working in industry and wanted to go back into academia. So I quit Deutsche and went to become a professor at Southampton and just a couple of weeks after that I was appointed to lead this national research and training initiative in the science and engineering of ultra large scale complex IT systems. And the world is full of these: national-scale health services, defence systems, critical infrastructure, and financial systems, are all massively dependent on networks involving hundreds or thousands of computers, and the users that interact with them. It's a big change that's happened in the last 10 or 15 years as everything has become computerized and as every computer can talk to any other computer. Suddenly, in principle, an error or a failure in one system, that would have been an isolated event, can have negative effects that ripple out in a chain reaction over a whole network.

***HFTR**: And then you have the whole "butterfly effect"?*

DC: Yes exactly. And one of the things that we have focused on in that project for the last five years is the extent on which the global financial markets are now essentially a single, planetary-wide, ultra-large scale complex IT system. And the extent to which there are failure modes like those I saw in FX back in 2005 might, in principle, ripple out over the entire system and cause big problems.

The 6th May Flash Crash was the first real sign that actually our concern was justified, that events could happen at an unprecedented scale, in terms of the magnitude of the drop and the speed at which it happened. The market is not supposed to crash while you go out and make a cup of tea.

***HFTR**: Yes, and then recover again while you drink it!*

DC: Exactly. So it definitely feels like we're not in Kansas anymore, Toto! The concern I have at the moment as a researcher and a scientist is to what extent are the dynamics and failure modes of these systems understood and to what extent are those bits we don't understand, risky or dangerous in the sense they might give rise to major market fluctuations or crashes that we really don't want to see?

***HFTR**: Given the fact that this global network has become more and more complex and there are more and more dependencies and interactions between systems, do you think it's even possible to map that complexity?*

DC: Well, that is a really good question. In 2009, Andy Haldane at the Bank of England published the text of a speech he gave called "Rethinking the Financial Network". He's a very, very bright guy and he writes brilliantly. You can download the PDF of that speech from the Bank of England website.

In there he basically argues that we need to first map the financial system because no one's got a map, no one really knows what the network of interacting entities is. And only once we've got that map do we stand a chance of managing what we've got. Only then, once we're able to manage things, do we have the chance to perhaps modify that network of interactions in order to make it more stable or less likely to suffer problems we don't want to see. But although the notion of mapping and then managing and modifying is something that Andy Haldane came up within 2009, it's still very much a research issue. Some people question whether it is indeed possible in principle. Maybe it's just not possible.

***HFTR**: And the fact that it's constantly evolving, must make it even more difficult.*

DC: Yes, but then it becomes more of a philosophical question. What do you do? Do you just give up and say it's dark and unknowable, it's just this mystery and we should just leave it be? Live with the fact that things will blow up every now and again? Or do you say that, as scientists and engineers, we should be able to do something?

After all, there are other complicated socio-economically critical systems that we have a good understanding of. And at least in part that's because of accidents that have happened in the past, which have forced the people who run those systems to reassess whether they want any more accidents like that, say in aerospace or in nuclear power engineering. For a start there's regulation but also there's a history of engineering for safety critical systems that is really quite different from the traditional engineering methods and practices applied in the construction of automated trading systems, or the electronic plumbing that glues together the financial market.

***HFTR**: So are we now facing a whole new set of risks we've never faced before because of these interconnected networks? Does current risk management culture need to change and if so, how?*

DC: Well, it can be informative to compare what goes on in other fields, other systems where risky technologies are being deployed on the front line and where the consequences of those technologies going wrong are high. The financial markets have become heavily dependent on technology and there are very high consequences if they go wrong.

And yet, to the best of my knowledge, they're not engineered with the same attention to risk assessment as you see in defense and aerospace and nuclear and critical national infrastructure. In part maybe that's just an accident of history. Plus, regulators have tended to not introduce as much legislation governing the design and construction and operation of automated trading systems as they have for airplanes or pharmaceuticals. In defense & aerospace for example, there is a set of techniques known as probabilistic risk assessment where essentially it's like a calculus, it allows you to look at the probability of individual events occurring but then the calculus allows you to compute probabilities over the entire system by appropriately combining those individual probabilities in various ways. At its simplest, you can think of it like a network of logic gates. There are some events where you need multiple failures

to occur simultaneously, like having an AND gate and others where you only need one or perhaps some subset of the systems to fail before things ripple forward, which is like having an OR gate.

The other issue you have to look at is incentives. Let's just say I decide tomorrow to set up a hedge fund. So I go out and I find a bunch of people who are prepared to invest and I have a pot of money and then I have my algorithm, which I think is going to make money. The thing that I'm massively incentivised by is my algorithm making money. So I will test this thing until I'm blue in the face to make sure it's not going to lose any money and indeed to make sure it's going to make as much money as possible. Then at that point, when I'm sure that I've done all the testing I can and I know that it will make money, the point where I release it into the market, I have no real control over what actually happens in the market thereafter. The only control I have is over my own automated trading system, I have no control over the other systems that it interacts with.

That is where events like May the 6th become an interesting warning because that was a system level failure. So I have my algorithm and I've tested it and I believe it's going to work and it's fine in my testing rig. And you have your algorithm, where the same is true. You've tested it, you don't want to lose money, and you believe it's going to work.

But what you don't know, in fact what neither you nor I know, is what will happen when your algorithm meets my algorithm. Because when I was testing mine I didn't know what yours was about to do and the same applies to you. So some of the interesting problems are understanding the system-level dynamics and that's a very big challenge, that is something about which very little has been done this far.

***HFTR**: I look forward to talking to you about that, and more, in part two of this interview. Thanks Dave.*

