McKinsey Digital

Teaching AI to sail like a world-class sailor and what it means for business

To help Emirates Team New Zealand defend their America's Cup title, a new sailor was introduced: a McKinsey-built AI bot taught using reinforcement learning.

March 2021

The America's Cup is the oldest trophy in international sport. Winning the America's Cup has always been as much about technology and innovation as it is about sailing. In this episode, we discuss the exciting and challenging work that has gone behind training a "digital twin"—an artificial intelligence (AI) bot within a simulator—using reinforcement learning to help Emirates Team New Zealand find the best hydrofoil design in their bid to defend their title. For more conversations on Future of Asia, subscribe to our podcast.

Oliver Tonby: You are listening to the Future of Asia Podcasts by McKinsey & Company. I am Oliver Tonby, your host and Chairman of McKinsey Asia. In this series, we feature leaders from across the region to discuss the forces, the opportunities, and the challenges that are shaping the Future of Asia.

Welcome everyone to the Future of Asia Podcast series. Today's topic is how McKinsey built a world-class AI sailor bot with Emirates team New Zealand. I am joined by three esteemed colleagues. I'm joined by Olie Fleming, he is the Chief Operating Officer (COO) of QuantumBlack in Australia. QuantumBlack is McKinsey's advanced analytics and machine learning firm. I am joined by Helen Mayhew—she is the COO of QuantumBlack in Europe. She's also a partner with the firm and she happens to be a world champion sailor as well. I am also joined by Nicolas Hohn, a senior expert with QuantumBlack out of Melbourne. Welcome all. Now I'm going to start straightaway, Olie, with a question to you, which is McKinsey and sailing. That doesn't seem like what we normally do. So why, why are we doing this in the first place?

Olie Fleming: Well, great question. Firstly. I mean I'm QuantumBlack, born and bred and QuantumBlack's heritage is actually in Formula 1 racing. So, there's actually quite a natural link between this work and QuantumBlack, but more broadly elite performance in sport is where new technology is proved out. And this opportunity allowed us to do highly relevant R&D in an application called deep reinforcement learning.

Oliver Tonby: Perfect. Thank you, Nicolas and Helen, let me ask each of you. Now, this must have been quite a bit different type of project to be doing with the firm. You know, just tell us some cool episodes that you've been part of over the last few months, Helen.

Helen Mayhew: I think one of the coolest episodes for me is when we built the bots, one of the unexpected consequences is that actually, the majority of the time, it outperformed the elite sailors. And the elite sailors were fascinated by this and it actually helped them improve their performance as they learned from the bot. So, I think a nice kind of synergy of human intelligence and machine learning coming together.

Oliver Tonby: Meaning as a former world champion and sailor, you're happy you got out of game in time. Is that what you're saying?

Helen Mayhew: I'm not sure I'd quite say this, but it's definitely a game changer when you can start to learn from technology in this way.

Oliver Tonby: Nicolas?

Nicolas Hohn: For me, one really cool moment was at the beginning of the project, when we actually got to go to New Zealand for the unveiling of the first boat. Those boats had been in the making for two years and no one knew whether those boats would actually sail and foil properly and being there for the first time when the boat actually foiled was, was absolutely magical.

Oliver Tonby: Excellent. Welcome, all three of you now let's dig in. But I do think it's useful for our listeners just to understand: who wants to explain what is America's Cup?

Helen Mayhew: I'll go a bit on what's the America's Cup. The America's Cup is the most prestigious race in sailing, and it's one of the

oldest sporting trophies in the world. It was first held in the Isle of Wight, just off the South coast of England in about 1851. And from the beginning was a real beacon of innovation. So, even in that very first series, the winner, which was a boat called "America," introduced what was quite radical innovation at the time: changing their sails from hemp to being made of cotton. And to this day, it's a hugely innovative competition where the design and the build of the boat is as critical as the sailing itself, much like Formula 1 where the race doesn't begin on race day. It begins with the development of the car. And in this instance, the parameters of the boat design, which are known as the "Class Rule," were first published in March 2018. And, from then until the racing actually starts on the water itself, the race is really one of innovation to build the fastest boat possible. And that's where we helped Emirates Team New Zealand.

Oliver Tonby: I understand. And, say a little bit more about the competition. So who's now in the final, how did they get there? Just explain a little bit to those of us that haven't been following this as closely.

Helen Mayhew: Absolutely. So the finals are a one-on-one race between Emirates Team New Zealand, who's the defender of the America's Cup, and the challenger, who has won the right to challenge them. And that is an Italian boat called Luna Rossa Prada Pirelli Team. In order to get to that position, the Italian boat had to beat off the British and the Americans. So there were four boats originally in the competition, and now, there are two that remain and they will be in the finals against one another.

Oliver Tonby: And for those of us that think of boats, as you know, small little dinghies that, you know, we rode. These are in a slightly different league given you're talking about four boats globally.

Helen Mayhew: Yeah. So these aren't typical boats, they travel upwards of a hundred kilometers an hour. They have hydrofoils, which means that they don't really sail on the water. They're in fact lifted about a meter or two above the water as they go. They have sophisticated onboard computing, and hundreds of Internet of Things (IoT) sensors. So, in some ways they're really like low-altitude aircraft, not boats at all. In the context of how important the America's Cup is to the sailors who are a part of it: the British team is headed by a guy called Sir Ben Ainslie, who is one of the most celebrated Olympians of all time. So, he has five Olympic medals, four of them gold. And he's gone on record on this America's Cup saying he would up all five of those metals if he could take the America's Cup back home to Britain. Which he won't be able to do this time, unfortunately, but just to give you a sense of what this means to the people who are competing.

Oliver Tonby: Superb. Olie, let me turn to you and ask you, can you explain a little bit what McKinsey actually did with Emirates Team New Zealand?

Olie Fleming: We've been their official technology partner. As part of that, we actually helped them focus on their two biggest areas of vulnerability that they saw coming into the campaign. The first was ensuring that these insanely complex 75-foot hydrofoiling monohulls were ready for sailing. They had as much uptime as possible. So, working on the maintenance and reliability. But the area we're talking about today was helping them optimize the design of their hydrofoils. As Helen said before, the "Class Rule," that, you know, the parameters of the boat design are fixed, but there are parts of the boat design, which there's quite a lot of space for the teams to be creative in the design of. And the hydrofoils are a particular area where the teams are able to become quite creative.

Oliver Tonby: So say a little bit more about this hydrofoil. So, what is it that the team would be looking at changing and how would this be done in the old way versus the new way?

Olie Fleming: One of the important rules that has changed in this America's Cup is that all of the design of the boats has to happen

in a digital simulator. Previously, teams could build mock-up boats, miniature boats, and pull them through tow tanks to understand their performance and help iterate through designs. They would also use wind tunnels. For this campaign, that's all been banned, and all teams need to use digital simulation. This is an area that Team New Zealand have pioneered and continue to lead. But there are some limitations to using the simulator. You need to take multiple Olympic athletes away from sailing on water and other forms of training, and actually get them to sit physically in the simulator and do laps.

Olie Fleming: So the design team can evaluate the relative performance of different designs. And the sailors themselves were also not completely consistent. So, they all control the boat slightly differently, which means they introduce noise into the data. And this means; one, there's a huge bottleneck for the design team to have to work through, to actually schedule the sailors, to get into the simulator, to do laps. But also, the marginal gains in performance that the team is hunting for, and in the case of the hydrofoil may be changing certain sections, the length, and the depth. It means some of that performance can actually be masked by the scatter of human testing. So, we set out on a highly ambitious and super exciting project to essentially take the humans out of the loop for a large part. And by deploying a nascent frontier technique called deep reinforcement learning, we were essentially able to build an autopilot, an AI agent, that was able to sail the optimal lap automatically in the simulator.

Oliver Tonby: What you're trying to do, if I understand, is you're trying to actually build a bot and train that bot to perform better than an Olympic sailor?

Olie Fleming: The actual criteria was: as good as an Olympic sailor. And the reason for that is that these marginal gains that the team are looking for—and at the end of the campaign, they're literally looking for tenths of a second in a 30-minute race—we needed to build a bot that could perform right at the limit so that they were able to compare two different designs at race pace.

Oliver Tonby: I would suggest the answer to the question is: yes.

Helen Mayhew: Well to perform more consistently than an Olympic sailor, but also to perform endlessly. So, you know, Olympic sailors have to eat and sleep and go to the gym, and do all sorts of other inconvenient things. So, the ability to take that out of the equation is critical.

Oliver Tonby: Wow. Got it. I understand. Olie you, you said the words "deep reinforcement learning;" and Nicolas, can you explain a little bit about what is that?

Nicolas Hohn: So, when you think about AI in general and the different types of models that our data scientists would use to solve a problem, you have more traditional models like predictive analytics and descriptive analytics. And reinforcement learning is a new type of model, which was very much developed and proved to work in the last maybe five years, mostly in video games, because video games offer a very good training ground for academics to try themes and see what works in a very controlled environment. And, we decided to use the same technique here. One good analogy is to think of a system that really learns with no prior knowledge, just with trial and errors. So, imagine us as humans, when we were toddlers, trying to learn to walk. You stand up, you take one step, and you fall because you haven't understood gravity, whatsoever. And then you try again, and you fall a little bit less, and slowly but surely, you manage to learn to walk and then you can run. And so here, the reinforcement learning algorithm does the same thing in sailing where starting from knowing absolutely nothing about sailing, slowly, but surely it learns what works, it learns what doesn't work, and tends

towards that optimal control that enables it to be as good as the sailors.

Oliver Tonby: And Nicolas, in this case, how long did it take the bot to learn to sail first, say, like a normal human being, and then sail like an Olympic sailor?

Nicolas Hohn: So all in all, it took us about a year to build that system that could ultimately be as performant as the Olympic sailors. In practice what happens is you start by teaching the bot to do very simple things like just sailing the straight line with no wind and no waves. Try to do that well. Once you master that, then you add the waves and the wind, and slowly, then you add maybe maneuvers like attack, or a jibe and made sure the bot is still very performant. For a long time, I would say we were maybe, 98 percent as good as the sailors. And it took us maybe another month or two at the end to just really break that barrier and get to 100 percent or maybe a bit more. All-in-all, like once all that development is done, the bot itself, when it's given a new design and trying to learn how to sail it, that just happens overnight. But it can happen over a short period of time because it's doing that in parallel, maybe over a thousand computers in the cloud. And so, it's managing all those different experiments together to get all the right sailing experiences, to ultimately gain all the knowledge it needs over a short period of time to be as good as the sailors.

Oliver Tonby: So Helen, we heard it took Nicolas and team a year to build an Olympic-level sailor. How long did it take you to win the world championship?

Helen Mayhew: I think I won my first when I was 16.

Oliver Tonby: There you go. That's pretty early. That's pretty good by the way. Asia standing in the world has changed. And it's clear that where the focus once was on how quickly the region would rise, the reality is now all about how Asia will lead. Keep listening to the Future of Asia Podcast.

Oliver Tonby: Listen, joking aside, I want to get back to Olie and ask you, what are some of the challenges of building this AI simulator bot?

Olie Fleming: Goodness me, I mean, there's a huge infrastructure orchestration component of it. This couldn't be done without the cloud. Nick just mentioned we needed to take Team New Zealand's simulator and run it in parallel up to a thousand different times in the cloud to generate enough learning for the bot. So, huge technical challenge to achieve this. I was privileged to work alongside Nick and a large global QuantumBlack team that managed to stand it up. So that would be the first one: getting the plumbing in place, I think, was huge. The second, not so much of a challenge; it was an exciting piece of really needing to work super closely and integrately with Team New Zealand, both in terms of their designers, you know, the hydrodynamicists, their own software engineers, but equally, the sailors. So that we could build trust in a really new area from the outset.

Olie Fleming: You know, at the beginning, when the sailors used to come down to the simulator, they used to love beating the computer. But over time as the performance improved, they'd also built their trust along the way. I thought that was a very exciting element of it. And then another one is not just ourselves on this project when COVID-19 hit, we went from a team that was broadly co-located in Sydney with a client in New Zealand to almost overnight have that large technical team being fragmented across London, Germany, Montreal, Melbourne, Sydney, and then our client in Auckland. I know it's not unique to this specific project, but it was a challenge that we had to overcome.

Oliver Tonby: Thank you Olie. I want to go to Helen because you know, you've been talking a lot as a sailor, but you are also a partner and the COO of QuantumBlack in Europe. So, you do know a little bit about the AI and the machine learning behind this as well. I would assume that, you know, for a company or a

team that's trying to build this, you need a certain type of skill set at your disposal, so to speak. So, tell us a little bit about what kind of skills did we have inside of the team here, the QuantumBlack team?

Helen Mayhew: So, within QuantumBlack, we've got a deep bench of technical people. So that ranges from data architects who will set up data architecture for us; data engineers, who wrangle that data; data scientists who work on that data to create insights all the way through to designers who help people visualize what the insight is that's coming out of the models. We also have a big range of people whom we call translators, who help speak between the technical teams and the business teams—the teams who need to implement the insights on the ground and in reality. So, at QuantumBlack, we have quite a large technical bench and Olie, I don't know if you just want to talk specifically about the type of technical talent we used in Team New Zealand?

Olie Fleming: Absolutely. I'm going to start with our translator. We were fortunate enough to have a guy called Henry who was not only a semi-professional sailor, but he also had a PhD in optimization. So, sort of a rare unicorn that was absolutely perfect for this. But we had a team of data scientists, cloud and data architects to help set up the infrastructure that we needed. And another really important skill set that we had was a software engineer. And we were delivering an operational testing platform based in the cloud for Team New Zealand. So, it was critical that the solution that we developed was production-ready.

Oliver Tonby: So listening to all of this, can you just say a little bit about what is the outcome? What's the impact of all of this? Did you achieve that the bots were better than the sailors?

Olie Fleming: We did. And in fact, in two-thirds of the situations, our bot was actually better than the sailors were able to achieve in the simulator. So, this was quite a remarkable achievement and it was a wonderful moment. And it was, it was a moment when the sailors became even more interested in the work that we were doing as it helped them eke out even more performance for a given design. Almost like AI-assisted coaching. For the designers, by automating the simulator, were able to deliver a 10X uplift in the number of designs that could be tested. So, we were essentially reducing the cost of testing to zero to allow the design team, to explore much more of the design space in their search for the optimal hydrofoil design.

Nicolas Hohn: One thing to add here is that compared to maybe a traditional industrial problem, Team New Zealand here are really running against the clock. And so, they are innovating at maximum speed, but with absolutely no room to delay the release of their product. And so being able to increase the rate of testing by 10X means that we really increase the chances of them finding the best design within the time window that they have.

Oliver Tonby: Now, ladies and gentlemen, let's shift topics and just spend a few minutes on what is the relevance of this beyond sports? Helen?

Helen Mayhew: Yeah, I'm happy to start. So, I think it's actually interesting to reflect first on why elite sports organizations are such early adopters of technology. I mean, we don't just see it in sailing with boat design, but in Formula 1, with race strategy and Premier League football teams, predicting injuries of players. They are always at the leading edge of adopting new technology. And in this instance, machine learning technology. It's because they are after that marginal performance gain and they're really willing to innovate to get there. And so, what can like mainstream organizations learn from how they do that? I think the first thing that mainstream

organizations can learn is to adopt AI well, you really have to choose something that's mission critical to you. So that may be how fast you innovate your product; what your product innovation looks like. It may be how highly performant your factories are.

Helen Mayhew: It may be how well you engage your customers. So, what's the thing, that is mission critical in your organization and that people are going to rally around in order to make the technology successful? And that will lead to the biggest performance gain. I think the second thing is then once you have found that thing that's mission critical, how do you express that in a clear metric? Because ultimately models end up training themselves around metrics. So, in sailing, the metric is Velocity Made Good or VMG, which basically means how fast are you going in the right direction? If you are interested in product innovation, it might be how quickly is your R&D cycle working? In productivity in a factory, you might be measuring, how can you reduce the downtime in your equipment? So, what's the metric that really matters to you? And I think the third thing to say is, how do you foster that real sense of human-machine engagement? So, you have augmented learning between the two. In essence, how do you make your workforce as curious about the technology as the sailors were about the bot?

Oliver Tonby: Thank you, Helen. Olie, would you care to add?

Olie Fleming: Yeah, it's worth saying that at the heart of this problem, it was, it was a complex control problem. We were getting a bot to control a boat in a simulator. And the characteristics of that, you know, it was a complex dynamic, and high-frequency system. It was a complex environment. There were 14 different controls that the sailors needed to orchestrate at the same time. The optimization goals for this were quite loosely defined, you know, velocity made good speed in the direction that you should be going. And so, there are other situations where we think problems can be tackled that have these characteristics; high throughput manufacturing environments, production environments such as energy generation and 5G network management. And indeed, how would you rapidly respond to changes in customer preferences and tastes? So, we successfully applied deep reinforcement learning on top of Team New Zealand's digital twin. Digital twins are not just found in a sailing environment. They are found in many traditional industrial settings, such as high throughput manufacturing, other advanced industries, automotive, aerospace, and applications such as logistics. We are excited about applying deep reinforcement learning on top of digital twins in those more traditional settings to optimize performance.

Oliver Tonby: Let me go to Nicolas and continue this train of thought. So, Nicolas and your experiences, can you give us a couple of other examples where you've seen AI or machine learning applied in the business world?

Nicolas Hohn: Sure. I think there are more and more examples. A lot of organizations are now quite good at maybe setting up one or two of those use cases. Examples could be "next best action"—when you're trying to serve your customer in the best possible way. So, based on what they've done before, you can recommend a new product or a new service. And a lot of that is now handled quite effectively with AI. Another one is definitely on quantifying the risk of a customer. For instance, if you're a bank, having very sophisticated AI models means that you can really make sure that you are proposing the right product within the right risk tolerance to your customer. A lot of the issues we see with AI currently are more about how you scale and how, instead of just doing one or two use cases, you really want to try to do a hundred of them. All the effort that you put in the first one, you can't put the same effort in the next one. And that's where automation and things like MLOps and DevOps techniques to really

streamline the way you build and deploy those use cases becomes absolutely essential.

Oliver Tonby: Thank you, all three of you. Now, we're going to wrap up and I want to ask you one question, same question to each one of you. If you put yourselves in the shoes of a senior executive, what is your advice to that senior executive that wants to get value from AI or machine learning? Nicolas.

Nicolas Hohn: For me, my one sentence answer would be making sure that you solve the right problem with the right people and the right technology. It's too easy to be carried away with very fancy technology that doesn't necessarily apply. And here we've talked about extremely fancy technology, but I would advise anyone to actually start simple things first.

Olie Fleming: Mine would be, be humble and upskill yourself so that you can reimagine your business. So, raise the ambition level and then surround yourself with the best technical talent you can find.

Helen Mayhew: In my view, the barriers to value in this area are actually often not technical; they are more cultural. They are about adoption and how you drive frontline impact. So, I would say for every dollar or hour you spend developing the technology, spend that dollar or hour again on embedding it into organization and really driving the frontline adoption and value.

Oliver Tonby: Perfect. Helen, Nicholas, Olie: thank you so much. This has been a fascinating conversation. I must admit, I've learned a ton about sailing and a ton about AI as well in this last half hour. Thank you so much for joining us and to everybody listening. Thank you and stay to take care of everyone. You have been listening to the Future of Asia Podcast by McKinsey & Company, to learn more about McKinsey, our people, our latest thinking, visit us at mckinsey.com/futureofAsia, or find us on LinkedIn, Twitter and Facebook.

Oliver Fleming is the chief operating officer of QuantumBlack (a McKinsey company) in Australia, and an expert associate partner in McKinsey's Sydney office; **Nicolas Hohn** is a senior expert in the Melbourne office; **Helen Mayhew** is chief operating officer of QuantumBlack in Europe and a partner in McKinsey's London office; and **Oliver Tonby** is the chairman of McKinsey Asia in the Singapore office.

Copyright © 2021 McKinsey & Company. All rights reserved.