

Defence Sub-Committee

Oral evidence: Developing AI capacity and expertise in UK Defence, HC 429

Tuesday 27 February 2024

Ordered by the House of Commons to be published on 27 February 2024.

[Watch the meeting](#)

Members present: Mrs Emma Lewell-Buck (Chair); Richard Drax; Mr Mark Francois; Jesse Norman; Sir Jeremy Quin; Gavin Robinson; Derek Twigg.

Questions 37-89

Witnesses

I: Neil Morphett, Chief Engineer, Advanced Technology, Lockheed Martin Rotary and Mission Systems UK, and Dr Simon Harwood, Director of Capability and Chief Technology Officer, Leonardo UK.

II: Richard Drake, General Manager, Anduril Industries; Phil Morris, Head of Defence AI (United Kingdom) Palantir Technologies; and Andrew van der Lem, Head of Defence, Faculty AI.

Examination of Witnesses

Witnesses: Neil Morphett and Dr Simon Harwood.

Q37 Chair: Good morning and welcome, everyone, to our second public session on developing AI capacity and expertise in defence. My name is Emma Lewell-Buck, and I will be chairing today's session. Can I ask our panellists to introduce themselves and say a little about what their organisations are involved in currently in defence and AI? Do you want to start, Neil?

Neil Morphett: Good morning, I am Neil Morphett. I am from Lockheed Martin UK in the Rotary and Mission Systems division. Lockheed Martin is obviously a worldwide defence and security organisation. The groups do everything from fighter aircraft to transport and space systems and rockets to munitions and defence systems.

The group that I am in is everything from helicopters to ship and submarine combat system sensors, training logistics systems and cyber and special effects. We also do some commercial work with civil organisations such as post offices. More AI is probably used in those than we have in defence, certainly in the UK, so far.

In the UK, programmes that we are involved with are things such as Merlin maritime helicopters. We do work for post offices around the world. We do training and logistics systems. We are partners with Babcock on the military flight training system, which has some AI in some of the training devices. We have some maritime navigation systems. We have cyber and electromagnetic support for various other Government Departments that have some emergent AI capabilities coming through.

I am sure there are things we do that I have forgotten, but we do things such as supporting F-35 in the UK and other Army programmes with logistic support systems that have some emergent AI capabilities. Thank you for the opportunity to be with you today.

Q38 Chair: Thank you, Neil. Dr Harwood?

Dr Simon Harwood: Thank you, first of all, for the opportunity to speak to you today. I have personally had a bit of an unusual career. It has always been in defence, but I have had 10 years in government, 10 years in industry and 10 years in academia; there is the way in which we talk about this triple helix. I was even, for my sins, a district councillor for a while—please excuse me for that—so I look at things in a slightly different way.

With Leonardo itself, it is important for the Committee to understand today that we are effectively data providers. We build sensors and systems that collect data. That data is then taken and used in certain ways in order to make decisions to generate knowledge to allow things to happen; you will all be familiar with the OODA loop in that sense. Leonardo itself is nowhere near as big as Lockheed Martin—I think we are about seventh in the world, with £13 billion to £15 billion.

Obviously, as the name suggests, we are an Italian-heritage company, but clearly we are a UK company. We have about 8,000 people onshore, supporting 10,000 people directly through the company across nine locations spread across Scotland and England in that sense. We have just opened a new defence artificial intelligence centre, for example, up in Newcastle in the north-east, and we are looking to do more in what we are doing there. It is about really coming back to look at the data lineage and taking that from a cyber or electromagnetic effect type of perspective.

Q39 Chair: Thank you both for that. As you will be aware, there are lots of documents, statements, strategies and units set up from a Government level around defence AI, but I am curious about how clearly the Government are communicating to industry what kind of AI capabilities they want to develop.

Neil Morphett: The defence AI strategy is actually a very clear document; it is very comprehensive and thorough, although there are probably a few areas where we would like to see more attention. The focus on early experimentation—on getting out and having a play with AI—is very thorough. What we are not seeing to the extent we would like is how it is going to make it into operational service in full capital projects, frontline platforms, frontline systems, and even to an extent the back-office enterprise systems. So yes, we are seeing a lot of experimentation right now.

The strategy recommends, or is considering mandating, “AI ready” in future capital projects. That might be a bit of a leap because no one is going to know what AI ready means, but having future programmes define what they are expecting to do with AI, and what provision they are making for future AI, is something that we would like to see in the road maps going forward so that we in industry can know what to prepare for and what we can engage the SME and academic community on.

At the moment, they can come in with a bright idea, but the gestation period for it to become a major programme and get delivered out is so long that they can't sustain the interest and need to move on. It is the same with people joining defence as a career. It is a long gestation period in defence, and we would like to see more of a road map for how these things are going to get into service so that we can prepare for it and aim to deliver them.

Dr Simon Harwood: Obviously, there are five or six main things that we are looking at. You feed down from the national AI strategy, to the defence AI strategy, to the playbook. DSTL's biscuit books are really useful, and there is the Army AI strategy and so forth.

Policy and strategy are just that: a destination—where you wish to get to. I don't think it is useful having documents that don't necessarily refer to one another when there is overlap and duplication in what they are saying. You can take any policy across Government and look at the coherence between different pieces of documentation and the references—the flow down—within them.

As my colleague said, it is about how we deploy the infrastructure and how the Government intend to engage with industry in order to make that vision a reality for what we are doing. In my expertise, we do an enormous amount of work in

AI, but it is about how you take disruptive technology and put it into the field—how you get that operationally through the system.

When you look at what other Governments are doing in terms of the levels of investment—hundreds of millions and billions of pounds—and the related technologies such as semiconductors' processing power, it is those things that the policy needs to focus on: what is the execution plan that the policy defines? So it is about deploying infrastructure to support AI.

The biggest barrier to innovation—I say this all the time—is the commercial models that the Government use in order to develop and deploy new technology. If we can overcome that, then we can look at the budget. There is no point taking a £1 million budget, splitting it 2 million ways and giving everybody 50p. You have to have a focused effort on what you want to do if you want to achieve the scale.

As I said, the playbook is really good in that it sort of says, "Here are some of the scenarios that we've got; these are the sort of things we want to achieve." I think the establishment of the Defence Artificial Intelligence Centre is an excellent initiative. If one thing stands above others, it is that that will add some clarity to the direction that we have got.

The last point I will make, and I will keep repeating this as we go forward, is that we can't just class industry as one thing. There seems to be a preoccupation in many instances of focusing on the funny acronym "SME", and I never know whether that means small to medium enterprise or subject matter expert. Industry is not industry. We are primes, and we are blessed to be in that position, and then there are small to medium enterprises.

There has never been a strategy about what is the relationship between industry primes and small to medium enterprises. That would be very helpful in exploiting this particular instance, because when I look at where we are spending our money on artificial intelligence, we are spending it with, interestingly, some of the people sitting behind us today, in terms of what you would class as small and medium enterprises in the artificial intelligence space. This is somewhere where you have to be agile. Typically, with companies like ours, for valid purpose, we are very regimented and structured in terms of governance, and so forth.

In summary, the documentation does a good job of saying, "What is the destination we want to achieve?", but a whole host of things are needed to deliver that aspiration, including funding, commercialisation, infrastructure and so forth to deliver it.

Q40 Gavin Robinson: That is a helpful opener from you both. Simon, you just mentioned the challenges and things that we don't have knowledge of or direction as to where we are going. Neil, you mentioned that the strategy was a comprehensive document. Last week, we took some evidence that highlighted two frailties. One was that the strategy talked about future AI 18 months ago—things that we have now seen deployed in Ukraine. We are seeing progress at an exponential rate. Secondly, we heard that it is devoid of an action plan and transparency around an action plan with matrices in which we mark or measure

progress.

At this stage, I would be interested to hear your reflections on whether you see that as a deficit in the strategy, and on the need for refresh in the strategy, or whether you think that the architecture is right and will allow the enablement for future development within the overall architecture, as it stands.

Neil Morphett: We can have a debate about which document it should be in. I echo Simon's point that more obvious linkage would be better.

I think the strategy probably does a good job, as it stands. Where it leaves it open is, "What are the individual programmes? What are the frontline commands? What is the major delivery that will give us proper operational capability? How is that going to happen?" We still see programmes coming out literally today—and if we went through a current major tender looking for the words "AI", "AI provision" or "AI road map today or future", it might not say it at all.

Some of these systems are going to be in service for 20, 30, 40 years. When we are doing a future action and we want to deploy what had been UORs and are now urgent capability requirements, you want, in an ideal world, to send out an AI model over Defence Digital and it will land in a platform and be flying the next morning, with that new capability. Right now, there is no road map to do that, and the platforms and systems are not being designed to enable that to happen. So in 10 years' time, when we are trying to do that, the answer is probably a plethora of laptops around a command centre, with everyone emailing data files around and everything taking a few days to work.

The strategy is good. The next demands are about what the individual programmes are doing to make that happen in their particular space, particular domain, particular programme and particular line of development. I will probably focus too much on the equipment line of development, but how will training work, and how will logistics work?

We have wargamed out how training will work in future. The Royal Navy today, when it is doing air defence training, goes down to the south-west approaches, and up come what used to be the Hawks out of Culdrose but are now some commercial business jets out of Bournemouth. They come in and are the aggressor. They pretend to be a missile and they have pods to look like an enemy missile.

Today, that works fine, because it is a blob on a radar and the crew says, "We think this is suspect", and they go and do their stuff. Imagine an AI world: the AI is plugged into Defence Digital and it says, "I can see the blob on the radar, but I have been tracking that thing since it left Bournemouth airport and it has been squawking commercial IFF codes for the past hour. The pilot's name is Bob. He was at Alton Towers last week according to Facebook. Yes, okay, for the last 20 minutes it has stopped squawking and it is pretending to be a missile, but it is friendly. Don't go and engage that thing."

So the world is going to change and some of those adjacent doctrines could do with thinking ahead to what the world will look like when that is a capability, because I think they are going to get surprised by it. That is without even

starting on the certification, accreditation and governance community, which is probably going to be very phobic.

Dr Simon Harwood: I spend my days talking about buzzwords, of which artificial intelligence is among the current list. We are doing data fusion, we are doing advanced machine learning, we are doing large data—we are doing a lot of things, so it depends what you mean.

The DSTL biscuit book explains this really well. We have been doing software development for years. The move into generative artificial intelligence is the next generation of—in simple terms—letting the computer write some code in order to make a decision. In order to do that, it is great to get some stories in the press about the impressive things that artificial intelligence can do, but if we are going to do this effectively, it is about doing the boring things really well: it is about the governance, the structure, the skills, the people and so on.

The one big question for me, post Levene, is who is in charge. Is it the Army, the Navy, the Air Force, StratCom, special forces or the Department for Science, Innovation and Technology? You do not buy AI; you buy a system that has AI within it, unless we are going to change the procurement system.

The good work that StratCom is doing under General Jim and others around the formation of the Integration Design Authority to set standards for what we are doing is a good first step down the pathway of getting the relevant ecosystem in place. As I have said, we need to work with the likes of DSTL, the future capability group in DE&S and others to set up that infrastructure.

We must not forget that the key thing in the move from machine learning to artificial intelligence is the ability to train the data. That is where the big risks come, if that data has bias within it or has not been developed in a similar way. Clearly, we work in a very classified world in many instances and the type of data that we are using is not widely available, or rather it has not been captured in the right format that we need to train the AI models. We then ask questions like, “What is the synthetic data we use? Is it good? Has it been validated?” That comes back to who is in charge.

In summary, it is about the Integration Design Authority in StratCom setting standards. It is a really good piece but, with the greatest respect, there is a little bit of incoherence between the different frontline commands on who is doing what. We would probably be better together than apart in that sense. Again, there is the question of what is a significant sum and how you contract for and buy this. I hope that that answers the question.

Q41 **Derek Twigg:** It is interesting that you have talked about money. Yesterday I had an answer to a parliamentary question saying that they do not have a clue what the spending on AI is. They are working on it at the moment to see whether they can better understand it, so maybe that says something.

Given your experience, the obvious question to you is whether the MoD is AI-ready, which is different from having a strategy and a policy. A sub-question, which fits in with what Dr Harwood has just mentioned, is what is happening in the MoD. You alluded to how the Army, the Navy and the Air Force are basically working in a silo when it comes to AI and would be a lot better working together.

If you take that as read, is that holding us back significantly from moving forward on AI? So the questions are “Are they AI-ready?” and “Is the fact that each service is apparently working in a silo holding us back from being AI-ready?”

Dr Simon Harwood: Are we AI-ready? That is a very difficult question to answer.

Derek Twigg: Are we readier than our potential adversaries, for instance?

Dr Simon Harwood: Are we, through the policy and doctrine and all the things we talked about, putting in place the building blocks to get to a point at which we can successfully roll out AI technology? Yes, I think we are heading in the right direction: we understand the policy, and the Bletchley declaration and elements like that are all good. But it is like asking, “Are you ready to run the London marathon?” Well, you might be the guy at the front who can nearly run it in two hours, or you might be the guy at the back they sweep up six hours later.

Where are we in that race, frankly, in terms of the investment and the level we are reaching? We are a good amateur, I would say; we are not in the professional pack at the front, but we are definitely a good amateur. That is to do with the scale and level of investment. Don’t get me wrong: some of these amateurs can perform really well, but it is about how serious we are about this. Are we spending enough on artificial intelligence? I come back to the fact that you do not buy AI; you buy a system that contains code and has artificial intelligence with it.

I do not want to be misinterpreted. The forces—the frontline commands—do work together, of course, but because you buy individual bits of equipment, there is not necessarily any core system. Look at Strike Network, Nexus or any of the software architectures across the individual commands. That is where interoperability comes in.

The other thing comes down to processing power—whether you are processing that on the jet, the plane, the ship or the submarine or doing that back at the other front end. We have got a huge number of—let me be careful with my terms here; we have a supply chain issues in relation to things such as semiconductors. We need those in order to have the sufficient levels of processing and compute power to effectively get artificial intelligence to work when deployed. Has my answer helped?

Derek Twigg: That’s fine.

Neil Morphett: I suspect the reason that the authorities do not know how much they spend on AI is the same reason they do not know how much they are spending on software. I do not think the authority captures the metrics to be able to say, “This fraction is on AI versus software versus any other part of the system.” To your wider question of whether the authority is AI-ready, I will echo Simon’s point: define ready. The answer is probably no. There are bits of it—it is quite sporadic—that are AI literate, have aspirations to get smart on it and are on the journey to do so. There are probably quite a lot of areas where, if it all

went away, they would be quite happy about it because then it would be a problem they did not have to deal with.

Getting the mindset from “This is why we can’t do this” to “This is how we will do this” is probably the key focus. It is important that individuals feel as though they are empowered to go and work out how we are going to certify an AI-ready platform and how the commercial infrastructure is going to work on something that is definitely going to be updated regularly. It is not a “fire and forget” project where you buy it and you are done. You are going to want to keep it going through life. You are going to want to engage with the supplier. The supplier is going to want to feed that product to other people. How does the IP work? It might have been their initial design, but it has had a load of MoD data into it. How is that going to work? Having people working out how they are going to solve that, rather than perhaps resisting it happening, is maybe the journey that MoD needs to go on.

Q42 Mr Francois: Because the MoD is so bureaucratic, if you want to get something changed, you need the strong support of either Ministers or at least a four-star officer, these days, to drive something. Who is in charge? When it comes to AI, is there actually a champion in the Department? Is it the head of StratCom? Is it the vice chief? Who, if anyone, has picked up this challenge and is driving it? Or, as I think you are suggesting, is the effort well meaning but highly atomised?

Dr Simon Harwood: Hmm.

Mr Francois: It is not a trick question.

Dr Simon Harwood: I just do not want to trip myself up. Each of the commanders in the frontline commands will take clear personal responsibility. Where does it roll up? Probably General Magowan through FMC. That would be the logical place. But then, obviously, we know about the Grand Canyon and valleys of death down at the sort of DE&S end. It comes back to what I said about development. We can do great research and development upstream. A contract is issued and, because that is time, cost and performance, there is not really time for innovation within that contract. Agile technology such as we are talking about here does not fit well in terms of development. I would have said that, overall, I think it is atomised—I like that word. There is not a clear lead within it. If you look at who is in the DAIC-X, it is future capability group, which is a fantastic model. What the team are doing down there at DE&S at FCG about prototyping the technology and getting it out pre-concept is fantastic. Defence Digital obviously is part of StratCom and then DSTL within that. But look at the three organisations within DAIC-X: StratCom, DSTL, and FCG, which is a component of DE&S. You have got to ask the question—where are the other commands?

Q43 Mr Francois: This is extremely helpful because this is sort of my point. To be fair to DE&S, they have got the Avon gorge down in Bristol, but I think the valley of death is probably a bit harsh. It is atomised, isn’t it? General Magowan has given evidence to this Committee multiple times. He is highly capable officer, but he has arguably got one of the busiest jobs in the entire Department, as FinMilCap. Even with the best will in the world, he probably does not have the

bandwidth to drive AI, given all the other plates he is trying to spin, to mix metaphors. The Committee is told that our potential enemies are exploiting this very heavily. If we are really going to make the most of this—if we are going to reply in kind, or better—do we not need someone at the centre who is empowered with a directive from the Secretary of State to really bang heads together and make something of this? If we do, who should that person be?

Dr Simon Harwood: Apologies, did you mean me to answer that question?

Mr Francois: Yes, but we would like Mr Morphett's view as well.

Dr Simon Harwood: General Magowan has taken an absolutely fantastic approach to collaboration with industry. I have been doing this for 25 or 30 years, and the openness with which General Magowan and the team share and collaborate with industry is fantastic, and a game changer. It is funny you should mention this—ironically, there is a gentleman here who was one of my first bosses in the MoD. Philip was in charge of DEC (CCII), the command, control and information infrastructure department for equipment capability at the time. That DEC structure, which sat under the predecessor of future military capability, was really good at taking control of bigger capabilities like this. There were two DECs—DEC (CCII) and DEC (ISTAR)—and nine directors of equipment capability on equipment procurement. The job of DEC (CCII) and DEC (ISTAR) was to—

Q44 **Mr Francois:** There is a risk of slightly drowning us in acronyms. Could you spell out what those organisations are?

Dr Simon Harwood: The predecessor of the current future military capability—FMC—organisation was what we called the DECS: the directors of equipment capability. There were 11 DECS, which, as the name suggests, were responsible for capability areas. Procurement was within those capability areas, but there were two cross-cutting DECS: DEC (CCII)—command, control and information infrastructure—and DEC (ISTAR)—we know what that stands for. Their job was to look at the overall capability in the domain. With Levene and the move to the frontline commands, we do not have those capability functions sitting in the centre. It was an excellent model, and I think it is something that we have lost. Going back to something like that—he says humbly—in my personal opinion, that worked pretty well, and that was 20 or so years ago.

Q45 **Mr Francois:** That is very helpful. Mr Morphett?

Neil Morphett: There is an organisation called the Defence AI and Autonomy Unit, which owns the Defence AI Centre and the empowering activities such as the Digital Foundry. The expertise in there is the empowerment. Is the senior champion in there? I don't know—I suspect not. Is it a role that should be formed, or is there an existing role that should take it on? I don't think I am in a good position to comment. We tend to find that these things bend around the character and personality of the person put in that role. You get the real visionary in the role who corrals everyone and drives it forward.

Q46 **Mr Francois:** Sorry, I am being slow—remind me where the Defence AI Centre is located?

Neil Morphett: If it has a physical presence somewhere, I do not know where it is.

Mr Francois: It is a bit worrying if you don't know.

Neil Morphett: We have a Lockheed AI centre, but it is scattered all over the UK, in the cloud and in the US. The good news is that it doesn't—

Mr Francois: I don't need to know where yours is. I'd just like to know that you know where the MoD's is.

Neil Morphett: If there is a physical desk somewhere, I genuinely do not know where it is. Perhaps I should. But they are the people running organisations such as the Digital Foundry.

Q47 **Mr Francois:** What rank is the person in charge of the Defence AI Centre?

Neil Morphett: I would have to look it up, but Navy commodore.

Mr Francois: A one-star officer?

Neil Morphett: Yes.

Mr Francois: All right. I think the word "atomised" is perhaps not entirely inappropriate.

Q48 **Sir Jeremy Quin:** In response to Derek's questions regarding readiness for AI, both of you were very focused on the paper and the policy and the direction of travel as announced by the MoD. Simon, it was only towards the end of your answer that you referred to processing power. If there are two or three areas where the MoD needs to up its game, could you point in that direction? The policy could be perfect, but it is more than that. My questions are on the cloud and processing power. It is a collection of data. It is all wonderful saying, "We're going to move to AI", but lots of stuff across Government is not even done digitally, so how you do AI if you do not have the raw data is a problem. Is that something the Armed Forces are good at or bad at, and do they need to do more? The third thing is culture. Is there a risk that everyone across the Armed Forces is saying, "We have these brilliant people doing AI and they are in a box over there", when what we need to have for AI to be properly addressed is a culture throughout the services about how they will be collecting the data and seizing the opportunities?

Dr Simon Harwood: There are five things, if I may, rather than three, only two of which are technological. I hate to sound like a broken record, but the first thing is people and SQEP—suitably qualified and experienced persons. We are competing with every other industry that wants to implement and use AI and advanced machine learning technologies. That is something that there are a million and one working groups on, but we never seem to come up with a problem. The second is data, as you mentioned, and availability, primarily of the training data at this point. Again, there is a fine balance between understanding the system that you are trying to implement artificial intelligence on and being an expert in artificial intelligence. They are not necessarily the same thing. There

is having the people, then there is having the data with the people to interpret the data in the right way.

The third thing is digital technology, which I will generically refer to. That includes high performance computing and the ability to process on the edge, which links to things like semiconductors, communications, quantum and other buzzwords and things like that—free space optical communications. There is then a massive thing around liability and the commercial and legal repercussions of what we are doing here, as well as the ability and need to have that. Linked very closely to that is regulation and governance on explainable AI and verification and validation. There are the five things, only two of which are technology.

Q49 Jesse Norman: I want to go back to some of the things that you have said already and put in a little bit of context. Everyone in the general public is obsessed now with ChatGPT, and the natural reaction from the pros is to say, “Hold on a second, we have been doing machine learning all these years. Let’s not just put AI on top of things and pretend it’s anything enormously new.” But there obviously is something new, qualitatively. That is what we are wrestling with.

The Defence AI Strategy was published in June ’22, and ChatGPT did not come out until November ’22. There has been a massive change in the public’s sense of what the capabilities of AI, generative AI in particular, could be. We have heard evidence that they could be even closer than in some of the wildest estimates. It could be a relatively small number of years away before we start to see activity in that area. There is an enormous amount of public and private capital going into this.

One of you talked about incoherence in the MoD, and I want to explore what that incoherence might be. We are struggling to put aircraft carriers to sea at the moment. We are using very ancient technology. The idea that we have a Defence establishment that is AI-savvy or AI-ready seems risible to me. I would be very interested in your thoughts on that. I also want to explore the idea of a platform and the way in which you think AI ought to be procured. Can you talk a bit about the sources of incoherence and capability, in terms that can give it some flesh and detail for us?

Neil Morphett: Coherence needs to be designed in from the start. Every programme we are working on with the Navy or the Army or the Air Force individually is predominantly focused on how we interoperate with our own service. At some point someone will go, “Well, if I am the Fleet Air Arm, how do I interoperate with the Royal Air Force?” They may well say, “Actually, we interoperate with the US navy more than we interoperate with the Royal Air Force.” But it will get a little bit stuck because there is a problem with just focusing on what we are going to do with AI in, for example, the Fleet Air Arm and what that is going to look like. Before you go too far, what is the US navy doing and what is the Royal Air Force doing? They are two different things. We need to pick a path or try to jam the two together and hope NATO or whoever puts some standards out there.

Q50 Jesse Norman: We could theoretically end up with a situation where, might it be fair to say, we are the extension—possibly not the favoured

extension—of a US airborne AI strategy without necessarily having some of the controls that we would expect?

Neil Morphett: Potentially, or a European one. All these things come with their own balance of, “Who are we going to interoperate with?”. Yes, we can always do our own—in fact, we have a history of doing our own, and then realising a few years later that we have come up with yet a third standard, a third architecture, that does not work with anybody else’s. After a while, we are saying, “I have this platform that I developed to my standard, but it really also needs to work with the platforms of the US, NATO or the Europeans,” so you end up having two. Eventually, the UK one dies, and although it was a nice initiative, actually, it probably would have been better to sit down at the start and say, “Who are we going to work with? Let’s have a common architecture that we work with,” and get that started with them.

Q51 **Jesse Norman:** You are not sensing that that common architecture is being discussed adequately at the moment between NATO, US and European allies.

Neil Morphett: As a collection of them, I think not. There are lots of bilateral discussions going on. The AUKUS framework has vast potential to get those countries on a common framework. As a company, we work very closely with our US, Australian and Canadian colleagues. From an industry point of view, those four nations actually work very well, as do the Defence Departments. The initiative coming out of AUKUS is that we should have a common AI standard as part of AUKUS pillar 2, which says that everyone is going to subscribe to this standard, and, if little local variations are needed, those will be local variations, not completely opposing approaches. That would be a great thing to come out of AUKUS. At the moment, AUKUS comes across a bit like a larger version of the UK AI policy, which is to try to do everything. Let’s pick some things to get all the way through first, to exercise the whole system, before we try to boil the ocean.

Q52 **Jesse Norman:** Thinking of AI and airborne systems, when we get to Tempest, we will have to include Japan as well. The natural extension of this platform starts to become really quite global. Do you have a comment on this, Dr Harwood?

Dr Simon Harwood: I have a couple of points. As you understand, we have to be very careful when talking about AI. AI in large language models is somewhat different. We use AI in large language models and we use AI for data and sensor fusion as well. I guess that is the first thing to say.

To answer your question specifically, there are effectively two approaches to AI models: you either build your own, or you take somebody else’s. In large, complex procurement systems, if you take the scenario you have laid out there, if you are late to the party, you might just have to adopt somebody else’s. Of course, while that has many benefits, you do not know how it has been trained or how it has been deployed in that sense, unless you get into the real weeds of it, and even then it is difficult to do. Do we want to build our own, or do we want to deploy somebody else’s? In international collaborations, that may be—

Q53 **Jesse Norman:** Do you have a preference? Do you think we should be



HOUSE OF COMMONS

taking one approach or the other, and which one should we take?

Dr Simon Harwood: I'm a Brit; I would much rather we build our own.

Jesse Norman: Right, but that is precisely the point that Neil has made—we can't do that, really, based on history.

Dr Simon Harwood: We can do.

Neil Morphett: Take the middle ground. Before it is trained, the fundamental AI engine—the raw app, if you like—is something that can be exported around different countries, and you can train it locally using your own data. You have, if not the best of both worlds, a balance between the two worlds. In an ideal world, you would have it trained with all the data from all the allies, so that there is one best-of-breed system that has all the information in it, versus having a UK one that is really good at picking up this sort of thing, while the US one is really good at something else. Do you put both of them in the computer, with one seeing this and the other seeing something different, and between them they get the solution? That sounds inefficient.

As a US parent company, we have certainly looked at it. The US is pouring a huge amount of money into AI. How do we in the UK get access to it? We do not want to reinvent our own, because the company will not pay for the same investment twice. But we can get the basic engines from the US. I can literally go on my phone right now and get the basic engines that we have a copy of on our server farm—I think it is in New Jersey, but I genuinely don't know where it is physically. We have access to that. We cannot get access to the training data and the finished model, but I can take it to the Wintermute exercises the British Army runs, and go and train it with pictures of Russian tanks banging around Salisbury plain and train it to recognise tanks.

Q54 **Jesse Norman:** I will come back to that in a second. Talking about platforms, there seems to be a kind of model evolving within defence AI circles that when we talk about systems, sitting underneath them, ideally, would be a platform that unites the uses of different systems together in some way, ideally with a single engine and consolidated datasets of the kind you described. It is kind of an iPhone view of the world. You then sporadically patch it from time to time and there are very high levels of interoperability, people can write different programmes and submit them to some quality assessment process before uploading into the system or into the network. Is that how you are broadly thinking?

Dr Simon Harwood: Yes, and the iPhone analogy is an interesting one. It is more iOS than it is iPhone. iPhone is the thing and iOS is the operating system that sits within. That analogy has many weaknesses, but the principle of what you are trying to say—

Jesse Norman: In a public forum, we are trying to get our hands round something people can get their heads round.

Dr Simon Harwood: Just to reassure you, you mentioned Tempest before, but look at the bigger programme: GCAP, the global combat air programme. Please be reassured. We have a thing called ISANKE, the integrated sensing and non-

kinetic effects element, which is what the UK is investing in for GCAP. We are doing a lot of work—you may have seen publicly that we did the Leonardo combat air AI challenge recently, where we are looking at deep reinforcement learning and integrated sensor fusion. Some of the companies, again, and other small to medium enterprises fed into that. We are doing a significant amount of work on what you would clarify as the platform, in order to be world leaders in that.

I just keep coming back to the question: what is enough? What is enough to be spending? At what time period and at what epoch do we want to deploy that capability? We are doing lots of amazing and fantastic work and I do not want the Committee to leave thinking that we are not doing a lot of cutting-edge stuff. It is just that from a personal perspective we could do a lot more a lot more quickly if there was more time, money, investment and commercial elements.

Q55 Jesse Norman: The point that Neil made about strategic risk of incoherence and lack of interoperability would remain. There is a governance and a driving of change in procurement challenge that is being articulated there as well.

Neil Morphett: There is a need for a coherent architecture, almost a system-design approach for the whole thing, even if individual services are only living in part of it. A platform that can run multiple AI systems from multiple places is not necessarily a bad thing. I mentioned earlier that we do commercial work for the postal services. We recognise envelopes: we read the address and the stamp. We actually have banks of different AI engines in there; some are ours, some are from third-party vendors and some are from the customer. We throw the image at all of them, and we know over time that these ones are really good at handwriting and those ones are really good at text. You can translate that into a defence environment and say, “We know these ones are really good at picking out tanks. Those are really good at picking out small surface contacts in a maritime environment. These others are really good at picking up RF signals in the ether.” Provided they are all sitting there on the same platform and you can deploy them all, versus—let us jump analogies for a moment—this one needs to be running on Android and this one needs to be running on Apple, if they are all sitting there on the same infrastructure, in defence digital there is no reason it cannot deliver that.

So next time we are somewhere like a Camp Bastion and someone literally says, “I have a new threat out there”—I do not want to make one up—someone back at base can come up with a new engine that can detect it, it goes out overnight, it is on the system, it is running the next day and it already has access to all the data. You can train it overnight on this wealth of data that we already have because it is the same format as we trained everything else on. It may do well on some things and weaker on others, but all the hard stuff—the infrastructure, the connectivity to the sensors and all that legacy data that is currently sitting in disk drives in cupboards in bits of MoD—all that is out in the open. We can throw it all at the new models and train them very quickly, getting a lot of, as Simon says, the boring stuff out of the way. The AI bit is sometimes the easy bit. You have done the best bit in two weeks. The gestation to get it through into service and get it connected to what it needs to do to be useful is the hard bit.

Q56 Richard Drax: Good morning, gentlemen. Can I go further on the point you touched on, Dr Harwood, which is the challenge of recruiting and retaining people who have or do not have these skills? In the evidence we have heard to date, the picture is a bit bleak. For example, the lack of sufficient STEM graduates has been listed as one. A highly competitive world, both in this country and internationally, means keeping people is very difficult. At the Defence AI Centre, a recent job as deputy head of centre had a salary of approximately £60,000. Professor Payne said that salary was not competitive. We have also heard that many of the young to mid-career experts in AI are Chinese. That is a bit controversial and difficult where defence is concerned. On developing defence AI and working with the MOD, contributions to the inquiry suggested that AI expertise within MOD “lags significantly behind academia and the broader UK AI industry”.

We heard some very good evidence from retired Air Marshal Ed Stringer, who suggested that one of the ways of resolving this issue was to look at careers, and maybe move people from civilian jobs into the military, and perhaps vice versa, to retain the skills, rather than keep someone in the defence business or armed services for all their career. I will start with you, Neil. What is your strategy to attract these talented minds? What specific skills are most difficult to find? Is it true that there is a specific challenge in finding UK nationals to fill these posts?

Neil Morphett: Are we experiencing challenges? Yes. Is it a massive problem today? No. It is not a massive problem because we do not have huge programmes of record that are demanding AI. We are still very much in that experimentation phase of, “Let’s go to this exercise,” with Army warfighting experiments and those sorts of things. We have ample capacity to deal with those. If the customer were to turn around and say, “I want to productionise that, deliver into service and turn it into a massive AI programme,” then yes, we absolutely would.

We know where we get good AI graduates and apprentices from—don’t underestimate the apprentice scheme. It is a domain where people can get stuck in very quickly. You do not need to have a graduate degree, though there are some good MScs around in AI. There are all sorts of sources of people. The disproportionate overseas student issue you mentioned is a problem. Clearly, we will need residency requirements and UK nationals for a lot of the programme, which brings us down to a smallish proportion of that base. It needs to be grown. Of what talent there is, where is it going when it leaves education? Is it coming to defence? Some of it is, and some of it is going into entertainment, fintech, health and all sorts of other places.

I mentioned earlier whether defence is an attractive place for an AI practitioner to come. The people we are really interested in are the AI architects—people who understand how it works, how to get it trained with the right sort of data, and understand the weaknesses. You can throw data at an AI and it will learn stuff, but it might learn the wrong stuff, and you will not find out until later. If, for example, every time you train it on a tank, the tank is in a field and pointing right—because that is the way the Army parks them—then it will think it is a tank if it is pointing right. We are really interested in people who understand how this works.

There is also the issue of the intellectual payback period for them. If they come and say, "I've got this really bright idea. I have done something exciting I would like to show you. I did it over the weekend at home with a bit of scripting, and it is brilliant. I threw some data at it I got off the internet," then that is fantastic. We might throw some resource at it to get it useful and demonstrable. We can take it to an Army warfighting experiment, which is great, but when we say, "What are you bringing us next year?" the bright ideas guy or girl is going, "Well, what's the point of me pumping these bright ideas in there and spending my weekend doing this cool stuff if the gestation period to get it out into the hands of the real user is two years?" If you have a live programme, it is probably two years by the time you get through all the testing, certification, safety—doing it properly. If you have to go to the customer and say, "You have a brilliant new idea here. Create a business case, get it through scrutiny, and get a programme established," then that is five years. That is an awfully long return time for someone at the start of their career who is full to the brim with vigour and has some great ideas. They will be thinking, "I can go to Amazon or Google or the film industry with this great thing, and it will be in usage."

Q57 Richard Drax: Can I rudely interrupt you? If I may paraphrase what you are saying, people may come up with brilliant ideas outside the MoD and give the MoD a fantastic idea, but by the time it is developed, it is old hat. It is gone, because AI is moving so fast. How do you keep those skills in the MoD if these individual contributors are no longer wanted, having given an idea that is no longer developed because it is taking too long? Do you follow my thinking?

Neil Morphett: Yes, I do. I would rephrase it and say that they are not no longer wanted; we would love to keep them, but they may feel like their ideas are not going anywhere. They may say, "I can go and feel like I am getting more intellectual reward out of some other industry." How do you keep them going? That is a good question, and the answer is to keep them doing the stuff that they like doing.

There is a volume of work going on. That could be seconding someone into a command and saying, "Go and be the AI person within this project team, Territorial regiment or experimental group. Go and take those skills in and use them there. Then come back to us and maybe second somewhere else and bounce around." That is a useful way of keeping that energy and interest up. Otherwise, you may have to say, "It is great that you have come in as an AI architect. It has all gone a bit quiet after that demonstration last week. Would you mind going and doing something completely different for a while, until we have more AI work?" That is a pretty depressing place for someone who is full of brewing vigour and has lots of ideas to bring to the party. They atrophy and say, "Well, I'll go and work in fintech in the City."

Q58 Richard Drax: Does that come back to my colleague Mr Francois's point that you need someone to have more control over the direction of travel? If each individual silo is getting people full of ideas but nothing is happening because they all take so long, and you are not necessarily getting or retaining the skills that you need, does it need a more co-ordinated approach where defence is concerned? It seems to not be working very well.

Neil Morphett: It is an absolutely valid point. If that road map is clearer, someone out there can say, "If I develop a bright idea here, I can see it going through on to these programmes and this new initiative that the MoD is doing. I can see it ending up on Tempest. I can see it going into UK exports elsewhere. Great, I can see where my ideas are going. I will keep pushing them into the funnel and following them through. If I get to bounce around and work with the Navy for a little bit, and I can see how it is being used and help put it through, that can only be a good thing." At the moment, there are too many ways of it stalling during that process.

Dr Simon Harwood: In my previous job, I was the academic adviser to the Security Minister, and I ran a university department for eight years. I feel quite passionate about this particular question. Let's start with STEM. Clearly, you can always have more people doing maths, physics and chemistry, and everything comes back to maths. You can incentivise that system at the bottom end. The point I am trying to make really early on here is that a lot of this is out of the control of the Ministry of Defence; other Departments of State and Government need to incentivise things. You have these come-to-school days where parents come and talk about their career. It is quite funny; if you have kids at school, which I do, and the school appeals for parents to come in and talk about their careers, try telling them you work in the defence industry. You will not hear back from the school again—I'll tell you that very quickly. Talking about working in the defence industry is not PC. Schools at that STEM age do not like that. I am being very generic there, and I am sure that there are schools that do like it, but that has been my personal experience.

We could go on about UK nationals all day. Clearly, we need them and we need more of them, but they need to be trained in order to do that. Obviously, we will not discuss the clearance system today, but it needs to be much snappier in turning around the security grade and classifications only going one way, and we need to do that.

We can talk about Gen X, Gen Y and all these different generations of people. I think they are probably all the same, but there are some idiosyncrasies to the different generations. In defence, we do not help ourselves because we do not tell people what we are actually doing. When you look at pharma, automotive or something else, you see these really exciting careers. You may say, "Oh look! I can do Formula 1. Oh look! I can go and put a spaceship on the moon or something like that." People ask us, "What do you do?", and we have to say, "Sorry, I can't tell you. It is classified." That does not help us in terms of that attraction. When we get people, we tend to put them on the same project for many years. We lock them in a room with no windows, we do not let them out very often and we cannot tell you what they are doing. The point I am trying to make there is that I do not know how you overcome that problem, because it is inherent in the system that we have. All I can say in public is, "My goodness me, when you go inside the locker and learn about what we are doing, you'll see it is absolutely cutting edge. But we cannot tell anybody that." That is the big problem we have.

With respect to your second question about specific skills, I do not get particularly wrapped up in any particular skills because there are not enough

people coming out of universities with specialist degrees. We have to look at what the related skills are, and generally those include data specialists and system engineering. In order to apply AI, you need to understand how to apply it, but you also need to understand the systems you are applying it to.

You referred to “moving between”, and I think Ed Stringer mentioned it; we would refer to that as a zig-zag career, or career passporting. We have had many attempts at career passporting and zig-zag careers, but we have never really managed to pull it off. The industry, through the Defence Suppliers Forum and other groups, is very keen on it, so everybody is leaning forward, but for a number of reasons we have not managed to execute that finally.

I will finish with a couple of things. On pay and compensation—yeah, we are not competitive. Why are we not competitive? Because the Government, coming back to my commercial point, runs a scheme called MEAT—most economically advantageous tender—which means the cheapest person gets the procurement. When you are pricing like that, obviously you have to drive your costs down. The costs refer ultimately to salary and labour costs; therefore, you are not competitive. I do not think the cheapest is the best a lot of the time if you want to achieve what you are implying.

I will finish by coming back to where I started in terms of STEM. I absolutely love the apprenticeship schemes and the apprenticeship levy. From a university perspective, classical education is great, but this ability to learn on the job is very important. We have to go back and look at what happened to the apprenticeship levy and apprenticeship scheme. We have to look at why it is not following through, why the industry is not spending as much money, and what the level 6, 7 and 8 apprenticeship schemes that we have involve. We need to avoid the silly headlines that you see in some of the newspapers about fat cats getting paid to do apprenticeships. That is not just true, it is about on-the-job experience at that top point of the pyramid for learning.

Q59 Richard Drax: Very quickly, people join the Armed Forces to fly, to be a soldier, to drive a tank. What if the young were to ask to join the AI department, which would be made up specifically of apprentices, as you say, and would perhaps be less military-minded and more AI-minded? What if we breathed them through the system that way?

Dr Simon Harwood: I would challenge whether they need to join the military. We need to be more collaborative between Government, industry and academia. You need places where you can co-develop—the term in the integrated review is “co-create”. When you talk about procurement you buy a ship, you buy a tank, you buy a submarine; going back to what we were saying earlier, you do not buy AI.

What we should be doing to incentivise and encourage people is developing prototype systems. They should be “agile”—that is a software term—rather than “waterfall.” That is so you do not start at one end and drop off the other, as in the traditional “waterfall” type of approach to procurement. We should follow the agile method of do a bit, test a bit, build a bit, deploy a bit, test a bit. Those are the big things we can do, and that is what gets people excited. That is the visible stuff, and that is what James Gavin and the Future Capability Group are doing.

We just need to ramp that up on steroids and look at co-investment. UKRI spends £6 billion a year, and Defence is spending about £600 million to £1 billion a year on overall research and development—there is coherence to be sought there.

There are many solutions, and this is a real opportunity. You can come from these things and be really negative about it, but we are making a difference; we could just do even more.

Chair: We are going to have to leave it there. Thank you very much, Dr Harwood and Neil Morphett. We will pause the Committee for a few moments while we swap over to our next panel.

Examination of Witnesses

Witnesses: Richard Drake, Phil Morris and Andrew van der Lem.

Q60 **Chair:** Good morning and welcome. You have already met the members of the Committee—I think you were sitting in on the previous panel. Are you happy to introduce yourselves and say a little about what your company does?

Andrew van der Lem: My name is Andrew van der Lem. I lead the defence team in Faculty AI. We are a UK-owned and UK-based AI company. We are slightly bigger than an SME—we have about 350 employees—and specialise in trying to solve hard problems using AI. About a quarter of what we do is in the defence space.

Phil Morris: My name is Phil Morris. I am head of defence AI at Palantir in the UK. Thank you very much for giving me the opportunity to speak here this morning.

Palantir is a software company that delivers AI-enabled capabilities to our users to allow them to make decisions with their data. We have proudly delivered those capabilities to the MoD for the last 15 years, as well as most of the UK's western allies, including the United States and the Armed Forces of Ukraine. The capability we deliver that enables us to bring these cutting-edge technologies to users can largely be thought of as plumbing, connecting that data that users have access to and enabling them to make decisions from it. There is a plethora of applications to this type of software—not least those that some of the previous witnesses mentioned around enabling a “best of breed” type of approach. My background is in science and engineering, and it is that standpoint that I will bring to answering the Committee's questions this morning.

Richard Drake: Good morning. My name is Richard Drake. I am the GM for Anduril Industries in the UK. Anduril Industries is a defence technology company based on the thesis of bringing the innovation that silicon valley can bring into the defence world. We use combinations of artificial intelligence, computer vision and machine learning and fuse them with hardware to create solutions for our allies and ourselves.

Chair: Welcome.

Q61 Sir Jeremy Quin: I have a very broad question to open with—so broad that you could each speak on it for an hour. Please don't; it will get me and you into trouble. What do you see as the most likely important advances in defence AI over the next five to 10 years? By all means give one, two or three points, but can we cap it there? Start anywhere you like. Phil, do you want to start?

Phil Morris: Thank you for the question. I will avoid future prediction, because we all wanted flying cars and ended up with 140 characters as a text limit on Twitter, so perhaps I will avoid that type of domain. I would say that very specifically, AI will bring the ability to make critical decisions more effectively and much faster, using the breadth of the data that the analysts, officers and senior leaders should have access to.

One analogy that I would give you there is to think of anyone with a decision-making capacity in the MoD, from a new analyst all the way through to a senior leader or even the Chief of the Defence Staff. You can think of the capability that AI could bring them as each one of those people having an army of 1,000 interns. If you imagine every single decision maker in the MoD having that army of 1,000 interns, their ability to move through even simple procedural and process-based work would significantly improve. I say "interns" versus experts or subject matter expertise, as the capabilities are currently in a state whereby you would want to check that work. Those interns will do diligent work for you, but you are unlikely to pass it out unsighted and make a critical decision off the back of it before you have had an eye over it or existed as a human in the loop to determine if that is an output you wish to deliver.

However, we can then take that a step forward and think about the amount of procedural and process-based work that armed services personnel conduct every single day—the amount of sitreps and assessreps, the amount of engineering plans or build plans, whatever it may be. Whether it is in the main building down the road, in Rosyth shipyard or a more sensitive intelligence hub, the ability of those individuals to be powered by advanced capability is obviously a benefit not just for the MoD and its outputs, but for the quality of life of those individuals. Therefore it is about the speed at which they are able to get through their work and spend their time doing analysis, rather than processing. That opportunity for every single decision maker in the MoD is immense.

With respect to some of the questions in the previous session, we should really be thinking about how we best silo or centralise, or whether we should AI—thinking of AI in that sense, around everyone having an army of 1,000 interns and its being an embedded technology across the full spectrum of effects that the MoD provides, from warfighting in the operational space all the way through to legal and policy, in the J8 and J9-type space. I am conscious of time now. As you said, there is hours' worth of stuff to talk about.

Q62 Sir Jeremy Quin: The bottom line of your answer is that this is the most extraordinary enabler. We gave you the opportunity to gaze into a crystal ball and say, "It can do this, this and this," and instead, perfectly reasonably, you said, "For anything involving process, it is an enabler that will speed up and make more efficient all of those processes." Is there anything that Andrew or



HOUSE OF COMMONS

Richard wish to add to that response? In particular, if you want to say, “And what about this?” you are welcome to do so.

Andrew van der Lem: I agree with the premise. You are talking about a technology that is a huge driver of productivity. “Interns” is correct, but maybe slightly narrow. The consequence of the technology, even in the quite short term for the Ministry of Defence—but also for warfighting capabilities more generally—is going to be revolutionary, because those magical interns are not necessarily going to be in the back office; they can be on weapons systems. You are automating human-like processing and putting it into lots of different places.

I also think that the economics of it will mean that the current capabilities of the UK or other Armed Forces will change. The relative importance—the relative capability that we have—will be revolutionised. We could talk about drones, for example, which could be controlled and co-ordinated much more accurately than they can at the moment, and of course that is going to change things in the way that warfighting is changing in Ukraine. Therefore, for the UK, it is really important that we think hard about the shift in the potential balance of power.

Something that I was struck by when listening to this is that a slightly depressing conversation has been going on about how difficult it is and how pessimistic we all are, but actually, it is a hugely disruptive technology, so why couldn’t the UK take advantage of it? Everyone is starting, more or less, from a standing start with this, so there are some great opportunities as well as risks in what is happening.

Richard Drake: For me, it is about doing more with less. Whichever way we look at it—the 1,000 interns or the drones, or whatever it is—the speed of iteration will increase, and continue to increase. We are seeing that in Ukraine around electronic warfare. Things are changing on a weekly basis where, in the past, they may have changed in timescales of months and years. We need to be able to keep up with that. The speed of decision making will increase.

For where we are with our Armed Forces, we are not as big as we used to be, but AI could bring us more mass—it could bring us more force power—for less. We can take away the dull, dirty and dangerous jobs from the humans, and AI allows us to give them to other people, keeping our people safe and increasing their ability to fight.

Q63 **Sir Jeremy Quin:** Andrew said, “Why couldn’t the UK take advantage of this?” That is the next question, so thank you for teeing it up. Compared with your relationships with other allied nations and commercial customers, what is it like working with the MoD, and do you sense that it is set up to succeed in this, or are there major flaws that need to be addressed? That is to any one of you. Richard, do you want to start?

Richard Drake: I don’t necessarily believe that there are flaws. I think that the MoD is more set up to deal with older-style procurement and delivery, and the commercial models will need revisiting. Considering that we are a software-driven business—we all are here—change is measured in weeks. Larger businesses tend to charge Government for a software update, whereas businesses like ours will drop a software update every two weeks. That requires a different approach—a subscription model almost—and the iPhone parallel is

pretty similar here. That is where change needs to occur at speed for businesses like ours to flourish.

Q64 Chair: Have any of you had any involvement with Commercial X? The MoD have told us that that is supposed to be more agile and less risk-averse in terms of defence procurement. Have you had any involvement with them, and is it working? Is it less risk-averse and more agile?

Phil Morris: Yes, they are less risk-averse, and yes, they are more agile. I would not understate the level of challenge that it takes to come from a hardware-based procurement system that takes decades to deliver something grey and made out of metal, and turn into something that was out of date 10 seconds ago. The scale of the challenge that Commercial X faces there is immense, and we at Palantir have certainly seen significant improvement in the iterative development of those procedures. Yes, of course we would always like to see things move faster, but we are very pleased with the speed of those developments.

In response to your original question, I believe the MoD faces a set of interlocking obstacles to deploying AI in an operationally relevant sense. You might expect me to say this, working for a software and technology company, but we believe that those obstacles largely, at their root, are technically based. Those interlocking obstacles then interweave technically into commercial and capital.

You mentioned the performance of allies, and how the MoD has performed with respect to them. If we look at the Department of Defence, they ran a number of programmes in the late noughties that worked out how to effectively acquire, accredit, assure, evaluate, deploy, monitor and improve AI capabilities, and in doing so essentially overcame their technical obstacles. Richard mentioned the speed in Ukraine at which we are seeing the development of capabilities and the compression of that OODA loop. A number of those programmes in the late noughties were based around the rapid deployment of AI capabilities to the warfighter.

Through those programmes, the Department of Defence essentially removed some of the commercial obstacles that these companies face, because the evaluation of those capabilities was taking place over days or weeks versus months and years. Therefore, the speed at which the Department of Defence could look at quantitative measures of evaluation to say that this or that model or capability may be better than the baseline, or may not be, was rapidly compressed. As a result, more and more enterprises, whether they be primes or medium-sized or SMEs, were able to enter into this type of programme. Removing the technical and commercial barriers in itself allowed a significant amount of private capital to flow into the defence industry, and specifically the defence AI industry, in the United States. Since the late noughties—around the time of 2017-18—that has dramatically improved the level of funding available to those companies, and thus their ability to cross the valley of death.

Just to come back to your original question on what it is like working with the MoD and how that Department compares to its allies, there is huge opportunity to allow Commercial X to take that same approach to address those technical,

root-cause obstacles. That in itself should allow for the lowering of the commercial and capital obstacles that are also in the way.

Sir Jeremy Quin: Don't feel obliged to add anything, Andrew. Are you happy not to—?

Andrew van der Lem: It has been a relatively positive experience—that is what I would say. And I have not disagreed with anything that my fellow panellists have said, nor do I disagree with what the first panel said about some of the problems that exist. But increasing professionalisation of procurement processes and understanding what software is are really important in all of this. The bits of MoD that embrace that are very positive for the advance of AI within the UK, or within the UK defence sector.

Q65 **Sir Jeremy Quin:** I have one last question and it is just for Phil. There is always an issue when an AI provider comes in to set up a service, which is that it depends on the data from the client. The data from the client goes into the AI provider and it is just—it is the same with the MoD as with any other part of Government, and I suspect it is the same with other Governments.

You referred to—the exact phrase was “allowing the clients to take decisions with their data.” Does the way you operate enable your clients to have flexibility about how their data is used and operated, or is there a risk that data, which is important and an incredibly valuable commodity, gets into your black box and they can have one line of answers from it, but if they cannot tweak or play with that data they get locked in, and it can be a very expensive process? I do not want to misrepresent you, but that is a concern, I presume?

Phil Morris: It is a common concern across the software industry, particularly in respect of working with Government and in defence. The approach that I just mentioned is that if the MoD was to focus its considerable talent, opportunity and resources on looking to remove those technical barriers—the platform-type approach that was mentioned earlier—that would give it the ability to quantitatively evaluate the performance of models, or even greater technologies, systems or platforms.

I say “quantitatively”, so that rather than just being dependent on how good the sales pitch was, or on whatever the particular urgent operational requirement might be, there could be testing of the particular model not just in one or two scenarios but in thousands of scenarios across millions of pieces of data. That would allow the MoD to create a best-of-breed approach, ensuring that there is never a lock-in to whichever AI capability or software capability was chosen. In fact, the capabilities that are utilised are those that are most appropriate and effective at the task at hand.

It might be that an in-house model built in conjunction with the Defence AI Centre and one of the frontline commands is used for ISTAR-type activities, perhaps by object detection or recognition through full-motion video. Then, it might be that a third-party model is used to understand patterns of signals that might be emanating from electronic devices on the battlefield in Ukraine. And perhaps then that could be combined with a model from a partner—an ally like the United States—to understand sentiment across social media-type posts, or whatever it may be, in that regard.

Taking each one of those capabilities and allowing them to be compared against essentially their competitors or other providers in that space would allow the MoD to break through the concept of black boxes, without the manufacturer needing to disclose their source code—their intellectual property—or the weights and biases of their models. It would allow the MoD to work things out and to say, “Ah yes, this one is very good in this setting, but not in desert terrain when the altitude is of x thousand feet for the drone,” or whatever it may be. So there are certainly routes to break through the perceived lock-in and black box of AI technologies that stem from the effective implementation of technical infrastructure.

Q66 Derek Twigg: Can I put a different slant on Jeremy’s question? You keep repeating the word “speed”. It seems apparent to me from the hearings that we have had so far that one of the issues—going back to your point about opportunities—is how the MoD keeps on top of the speed of change, given its record on conventional weapons, with some of them being outdated by the time they have been completed, and technology having moved on some of it even becoming obsolete. You have answered Jeremy’s question in a way by saying, “Yes, the MoD’s doing really well, it’s working hard, it’s changing things.” How does it keep on top of the sheer scale and speed of change, in terms of software, data and artificial intelligence? Is it doing that?

Andrew van der Lem: There was a reference to agile ways of working earlier and thinking, “What does this mean for -software?”. I think it is true that people are in the mindset of “this is a 20-year programme” and therefore it can all get a little bit flabby or difficult to track the individual inputs. Part of it is making sure that your procurement is set up so that you are regularly going to the market, and that you are interested in getting incremental change.

The other thing is linking up with the users better. Because we are used to these large hardware programmes in the MoD, you are often very far away from the sailors, the soldiers, the pilots or the analysts. These are people who love technology—they use it all the time at home—so these are people who are very informed about what is going on and demand and would like new features. I understand that it is difficult within the secure environment. We have been working with the MoD for only two years, and we are really proud of all the things we have done and it has gone really well, but the one thing that we often stumble over is actually talking about our software solutions with the people who are going to use them. If you harness that, I think that really accelerates change and builds up demand for the absolute state-of-the-art solution for the particular problem that they want to solve.

Q67 Derek Twigg: That is a very important point, but just to pin you down—anyone can answer the question—how can we as a Committee be assured, or have the confidence, that the MoD is on top of this? What is it that we need to be looking at? What will give us confidence that that is happening?

Richard Drake: I am probably not going to answer your question directly, so I will apologise immediately. The way that it is done now gives you that confidence, because someone will specify something down to the paint colour, and then there is a whole industry—I used to work for a major UK prime—

around answering those questions to prove to the MoD that we have given them exactly what they wanted.

Q68 **Derek Twigg:** There is also over-specification.

Richard Drake: Yes, atomising everything. Where we need to be is in an 80:20, because 80% of the work is done in the last 20%, and co-developing with the end user allows us to do that. For you to see whether the MoD is on top of it is more difficult in that in that regard, but that is the way that success will be achieved.

Phil Morris: I think you will know that we have started to flip the balance of risk from commercial and financial and into the potential of operational failure of not deploying AI. When you start to hear more about these programmes that have come into difficulty—I don't want to say failed—in developing the software, the only way that software companies like ours are able to make progress in the field of software is to try something and find out that it doesn't work, and to do that again and again. I am not suggesting that that would need to take place at technology readiness level 9—on the battlefield. You would want to ensure that that was the case much further down the maturity spectrum of these technologies.

To your original question as to how you would have the confidence in ensuring that the MoD was on top of this, one marker of success or metric for that progress will be in seeing what things didn't work. Coming back to some previous evidence, which I believe was raised by Mr Black from RAND last week, service personnel are not necessarily empowered to fail in these regards, and yet, if you spoke to any UK, US or western technology company, every one of their employees is almost certainly allowed to fail, because that is how software is developed—it is software development.

There are certainly markers around allowing these programmes to utilise capital funds in a way that means that they are not necessarily throwing it all away, but a way that allows them to explore areas that may never have been thought of, because it is from there that we will find the once-in-a-lifetime type discoveries. I think there is a point there around seeing the flipping of commercial and financial risk to that of operational failure. If we do not do this, we could fail on these operations. Therefore, some of the increased risk tolerance in that regard may mean that we see that some of these programmes perhaps are not quite as successful.

Chair: We have just had news that there is an urgent question in the Chamber at 12.30 pm about defence equipment to Ukraine, so we might lose a lot of the Committee at about 12.25 pm. Can we keep that in mind as I hand over to Gavin Robinson?

Q69 **Gavin Robinson:** Thank you, Chair. I was not here for the start of this session, but I think you have all highlighted your experience with Commercial X and the important relationships that you have with the Ministry of Defence. Turning to the Defence AI Centre, one of its functions or roles and its focus should be to encourage relationships between defence and the technology sector. Has it been effective? I don't mind who wishes to take that up first, but

have you found that that role of the Defence AI Centre is working, or could it do better?

Andrew van der Lem: Over the last two years we have done quite a lot directly for the Defence AI Centre and, with others, helped with the process of setting it up. In those two years we saw a lot of progress. We have seen a lot of the building blocks that the other panel referred to being developed and tested and new ways of working. We have seen a very positive attitude from people and a different attitude towards accessing data, training data, and working at different levels of classification. Is it enough? Is it fast enough? Absolutely not, because the scale of the threats and the opportunities is so huge, as we have seen in Ukraine and elsewhere, but we found it positive. It really helped us as a British company to develop capabilities that we can now use in other parts of the MoD, so the balance is definitely more pro than con.

Phil Morris: I agree with Andrew entirely. I believe the Defence AI Centre has a significant part to play in the delivery of AI capabilities to the personnel on the frontlines. It is important that we don't lose sight of the fact that AI is a general purpose embedded technology, so centralising to a specific Department has its risks. Last week it was noted, "Who should be the SRO for Defence?" and it is the Chief of the Defence Staff. It is that significant to the entirety of the Department. While it has shown significant progress in its nascent life, AI cannot get to the point of essentially being outsourced. We would not outsource, or I think it would be very unlikely that in governmental spheres we would think about outsourcing cloud computing, semiconductors or general purpose technologies.

The challenge in the focus of this centralisation is going to be scale. We see quite commonly across the UK AI companies with which we partner and which leverage our technology to bring their solutions to MoD and Government that they feel as though there is not the sharing of challenge from the MoD. While there are multiple innovation frameworks and procedures that allow these companies to have access to MoD-type problems, those challenges can often get stuck at technology readiness levels that essentially end at demonstrators, perhaps on Salisbury plain or inside a laboratory-type setting.

The challenge that the MoD faces—the DoD faces it, too; the world faces it—is taking these innovations and allowing them to be scaled to the extent that they could be utilised by the entire fighting force of the MoD or of the UK. It is very clear to us in our conversations with venture capitalists, the private capital class or even just the firms themselves that there is enormous appetite for these companies and for the capital to be deployed in these types of problem sets, but they do not feel as though the scale of challenge that the UK faces is open, accessible and available to them to try to put their technology on to help against it.

Richard Drake: I have nothing more to add. Scaling is the thing for us.

Q70 **Richard Drax:** Good morning, gentlemen. Our evidence says that the UK does not have a defence ecosystem that encourages risk taking or innovation in AI. How can the MoD better incentivise innovation and encourage investment in the sector? You touched, I think, on this, so the second part of the question is

about how defence can bridge the gap between experimentation and fielding capability—the so-called valley of death, which sounds very dramatic. Andrew, would you like to go first?

Andrew van der Lem: Trying to apply AI to real-world problems, as the panel have said, is part of this. The valley of death exists exactly because things have stayed in the lab. There is a disjoint—again, I come back to the end users, the actual military users—between the users and the drivers of innovation in defence.

Q71 **Richard Drax:** Are you saying that the people who make all this up—the backroom boys and girls, call them what you will—are not talking to the potential users, so the users do not really know what is being done? Is that what you are saying?

Andrew van der Lem: Crudely, yes. There is not enough conversation and not enough—

Richard Drax: So, better co-ordination?

Andrew van der Lem: Better co-ordination, or just valuing the frontline users more highly and involving them in the setting of the questions. Not doing it as an academic exercise, to ask whether something is interesting or not, but applying it to real-world problems is, I would say, one of the key things in this. I will have to disagree a little with what was said. I am not completely convinced that the answer is necessarily to bring in a lot of— Let me put it another way. It may be that you could get innovation from defence companies in other countries, but our experience is as a company that has worked in other sectors outside defence and has come into defence.

There are huge amounts of learnings from healthcare, retail or wherever else that you can bring in, and one of the interesting things about AI is that it is being driven by consumers—everyone knows about AI, everyone has it in their phone. The innovation is coming from the consumer sector and that is what is exciting and can be applied in the military and defence sphere. Part of it is not just opening up to more foreign companies and foreign capital—that has got to be part of the innovation—but opening up to non-traditional defence suppliers.

Phil Morris: I think you will find that the senior leaders of the MoD are wholeheartedly on board with the idea of AI being an embedded technology that horizontally scales across their force. If you were then to go and meet operators and service personnel on the frontline, so to speak, you would find a willing and driven set of individuals who also want to use those technologies. It is the enabling functions in the middle, to go back to one of my earlier points about the technical foundations we are lacking in the UK, that would allow the rapid deployment of those capabilities, with the sponsorship of the senior leaders, into the hands of the end users on those frontlines. If we can get through those technical barriers, that breaks down some of the commercial and capital barriers.

In terms of bridging the gap and incentivisation, as well as the investment of capital, that comes back to a previous answer about allowing these companies to see that their technology could be used on something beyond an innovation

sandbox. If it was proved to be successful in an innovation sandbox, it could then be used on the frontlines within six months. Every technology company we work in concert with at Palantir would say that they would desperately want their technology to be deployed in a way that helped the UK prosper, whether that is in a military sense, a health setting, commercially or whatever. To incentivise that means shortening the time it takes for a three-person spin-out from a university, or however a company may come about, to go from saying, "Look, I've got a great idea," to being allowed to take a technology that is then tested, evaluated and able to be deployed to make a difference.

If we can put those opportunities in front of the plethora of companies that exist in the UK, we will rapidly bridge that valley of death, and attract that level of capital into those companies to allow them to expand their offerings and allow the UK and the MoD to succeed.

Q72 Richard Drax: Just for the record, the MoD is not doing that; it is not getting to the frontline. An idea, an innovation or something that may be slightly more risky—whatever it is—is not being experimented with by those on the frontline. Is that what you are saying?

Phil Morris: It most certainly does happen. It happens through the number of innovation frameworks and pathways that exist to allow private companies to bring their technologies to the MoD. However, if we look to the application of these technologies in Ukraine, that OODA loop takes hours and days, rather than the months and years it might take in a post-war-type state, versus the pre war we are in now.

Q73 Richard Drax: War tends to speed things up, doesn't it? Richard.

Richard Drake: For me, it is around business models. We are currently set up for a business model that encourages slow speed, because the margins we are allowing companies to make on large-scale defence projects are quite small in percentage. People want to make money from this; we are companies, we are a business.

If the business model is different, where smaller businesses like ours are happy to take more risks than the larger businesses, it will encourage investment, new brains and new skills if people can see a way of growing their business and becoming successful. Getting the right business model that allows us to move from being the businesses that do the clever things in the bottom left-hand corner of scale in business to that larger scale of business makes a big difference.

Q74 Jesse Norman: Gentlemen, we have heard about the importance of the right platform or platforms, and we have explored the analogy with iOS or the equivalent. What state is the MoD's platform for AI in at the moment? When will it be ready and who is doing it?

Richard Drake: It is hard to generalise, because it is atomised—in a good way, because certain teams have certain aspects for what they need to be able to say—

Q75 Jesse Norman: Give me an example of a good one and a less good one—one that is, as it were, in development, just so we can be a bit more concrete.

Otherwise, we are swimming around in generalisations.

Richard Drake: I agree, but I am going to struggle to find a not-good one, because I can only talk about the ones we have worked on. With a lot of the ISR platforms, a lot of the calculation that works behind that is very good, because they are need-driven. In areas where the need is not quite so great, it is not quite so developed.

Q76 **Jesse Norman:** You will have developed the platform in that area, or others will have, to a commissioned set, put by MoD under advice.

Richard Drake: Yes, I believe it is not uniform, because the need is not uniform; the need right now and the need next week are different in different areas. I hope that helped.

Q77 **Jesse Norman:** We are about 5% towards where I hoped we would get to. Phil or Andrew, do you want to have a shot?

Phil Morris: Sure. It is really a question of what we define as doing well. Is that taking the capability and delivering it to the war fighter? Is it exploring new state of the art? Is it winning on the battlefield? To your original question, I would say the state for the UK is nascent. It reflects the fact that the DAIC has made significant strides over its first couple of years of existence, but is still in a nascent state.

As mentioned already, there is some challenge in the post-Levene world as to where the ownership lies in driving that forward. We would certainly like to see much more embedding of AI across the services, versus in a specific area. Therefore, we would like to see a little bit more around the idea that there is a place in which these capabilities—let's call them models—are developed, tested, evaluated and monitored. It is most likely that that will take place in the DAIC, but due to its nascent state that is still being set up in that regard.

Q78 **Jesse Norman:** You are hinting at the incoherence arguments addressed by Neil earlier on?

Phil Morris: Correct, yes.

Q79 **Jesse Norman:** That is helpful. Does Palantir do a country-by-country scoring of how well advanced their AI is? You must know what good and bad looks like across Europe, for instance.

Phil Morris: I am sorry to disappoint you. We do not have scoring per se, but we do have a good understanding of what good looks like.

Q80 **Jesse Norman:** But you will be thinking, "Are these folks ready for me to sell into, as I and my colleague have done successfully in another country?" That will be the kind of situation you will be thinking about.

Phil Morris: Yes, and the MoD is absolutely ready to be able to buy that type of capability.

Q81 **Jesse Norman:** Notwithstanding its nascent capability, would you regard the MoD at the moment as being at the top end?



Phil Morris: Notwithstanding its nascent capability, I would regard it as able to buy these capabilities and put them on to the frontline, mostly certainly.

Q82 **Jesse Norman:** Whereas Estonia, to pick a random name as an example, might be further advanced?

Phil Morris: Because of some very specific reasons as to how they are digitally developed. The very obvious comparison for the United Kingdom is the United States, and my American colleagues would say that, while significant advancements have been made in the Department of Defense—referring back to some of the programmes in the late noughties that I mentioned previously—there is still a significant amount of money that should be spent but is not being spent.

Q83 **Jesse Norman:** Within DoD?

Phil Morris: Within the Department of Defense, yes, to uplift their software spend to at least 1% of their budget, which it is not currently at.

Andrew van der Lem: I do not quite agree with the paradigm or the similes and metaphors we have been using about creating this platform. We as a company build the models and the applications that sit on the platforms. Sometimes we have to help out on the platform or work with companies such as Palantir in order that the platform exists or is stable in order for our work to happen. That is where we come from in it.

Part of the way that you get the infrastructure in place is by building the models and using the models. That creates demand. Our experience has been that if you try to solve every single problem sequentially, it will take forever. Part of the solution is putting end-to-end AI solutions in place and showing that they work. That then creates the demand for the infrastructure and for the platforms to follow and be built up as well. I do not disagree with what others have said, but—

Q84 **Jesse Norman:** You are throwing the ingredients into the pot, seeing how good they are looking and then tweaking and developing and, if necessary, replacing.

Andrew van der Lem: If you bake some great cakes, someone will pay for a new oven so you can bake more cakes.

Jesse Norman: We are now in the world of food—a bit of an overstretch I think, but I am getting the point.

Q85 **Derek Twigg:** Phil, you referred a few times to some non-traditional companies and their access. What barriers are holding back these non-traditional companies that could be doing things in defence?

Phil Morris: There are a number—a couple have been mentioned already in this inquiry—including clearances, access to secure spaces and access to data. As I have mentioned, if the MoD is able to address these technical infrastructure pieces to allow a company to very rapidly get close to those opportunities, perhaps with a middle level-type clearance versus the highest levels that are required, that would give a much greater appetite for the company to say, “Yes,

we are going to double down on this particular line of development, in which we have something new and state of the art.” The challenges are certainly around the infrastructure, both technical and non-technical, and the clearances and access to secure facilities.

Q86 Derek Twigg: Do you think that that is a mentality problem, in that we are potentially on a war footing because of what is going on in the world, but the MoD’s psyche hasn’t got to that stage? It is still working on a pre-Ukraine situation. Therefore, we would probably move more quickly where we are today, compared with two years ago.

Phil Morris: We no longer have the luxury of time in a pre-war state. I agree with the statements made previously, in that the strategies are well developed in themselves, but most of those capabilities that were “AI next”—I can’t remember them all specifically off the top of my head—are commercially available off the shelf today. We have mentioned the speed at which large language models have been developed. Last year, text to video created laughably bad capabilities that you could clearly tell were not real, but just a couple of weeks ago, a leading AI firm released a model that makes completely human-realistic capability.

If you combine the speed with which the technology is developing and the speed at which the observe, orient, decide, act loops are needing to take place on the battlefield in Ukraine, there is somewhat of a disconnect with the pace at which the MoD must move to best take advantage of these capabilities, bring in the significant expertise that exists within the UK market and industry, and deploy them on to operationally critical and relevant problem sets.

Q87 Derek Twigg: Is there anything you want to add to that?

Andrew van der Lem: Just on one facet of it. On security clearance for individuals, if that cannot be done quickly, you cannot have a non-traditional supplier provide a solution. It is as simple as that. It must be a solvable problem—to change how we do vetting to allow people to come in with, I don’t know, a few hours of work to get their clearances, rather than weeks and months.

Derek Twigg: In what could be critical developments in terms of moving forward.

Andrew van der Lem: Yes, they could be massively critical.

Q88 Chair: Just one final question from myself. How can the MoD turn AUKUS into a success story for UK AI, or can it not?

Phil Morris: I would like to draw on the movement away from the previous accreditation scheme towards secure by design. I think that there is significant opportunity to exploit the secure-by-design framework to allow, essentially, a “build one, ship many” type of approach. If some of the small to medium enterprises that we have been talking about, which have perhaps been struggling to gain access to these problems due to clearances or facility access, could utilise secure by design and essentially be accredited in the UK once, and therefore applicable in the United States and in Australia on very serious critical operational defence problems, that is a huge opportunity that a whole number of

these very innovative, novel, nimble, agile companies would want their technology to be deployed upon. The opportunity is immense. The danger is that we take some of the failings or issues with commercial, procedural and financial risk that exist in each of those nations and combine them.

Richard Drake: The worst of all worlds.

Phil Morris: Exactly. There would be nervousness around us, rather than making an agile technology decision, making a tripartite decision for technology across the board and therefore losing the balance of risk against operations.

Q89 **Chair:** Andrew or Richard, did you want to add anything to that?

Richard Drake: I think the same could be said of export control and export rules. We have started to see movement, certainly between the States and Australia, around ITAR, for example. Any level of acceleration there to allow us—or any business, really—to do once, use many, with less paperwork and fewer barriers, while still remaining secure, is important.

Andrew van der Lem: Yes, it is also about showcasing the good stuff that has happened in the UK to the Australians and listening to what they have got—the very simple stuff of being proud of what we have got and allowing them to reuse it quickly and cheaply. That would put benefit in the hands of the Australians, which again would be a positive thing to do.

Chair: Thank you all very much. Thank you to our witnesses, Committee and staff. I will end the session now.