DARPA Information Innovation Office (I2O) Office Wide Proposers Day

Q&A

Question: What is DARPA's interface between traditional hardware and artificial intelligence?
Answer: Some of these programs are in fact already at the interface between software and hardware. There is another version of this slide that has MTO slides. Lok Yan is a PM in MTO who we work with on a regular basis. We certainly are interested in programs that span that just for the same reason that often vulnerabilities come at the boundaries. From a budget management standpoint, programs are only launched from single office, but the logical launching from two offices occurs frequently.

Question: Quantum computers are making progress; cybersecurity is getting into a new area of "quantum safe security." Is there any new plans or programs from I2O on cryptographic engineering modernization of cryptography for quantum-safe security?
Answer: We are not doing anything on quantum-safe security. We do have the QUANET program, which is using quantum on making networking more secure. DSO has a number of efforts on quantum.
Question Follow up: It is not related to quantum exactly. Quantum security that NIST is standardizing and the NSA is pushing it and endorsing algorithms deployment, migration, integration and evaluation is a large topic that the DoD should be interested in.
Follow Up Answer: If you think that what NIST and other organizations are doing is not sufficient, and you think I2O should be doing something that they are not, I encourage you to reach out to Allyson O'Brien who is the I2O program manager who does quantum and discuss this with her.

Question: How will the President's Executive Order on Safe, Secure, and Trustworthy AI affect DARPA?
Answer: I am not sure yet. There are restrictions on what it takes to work on a frontier model in the new document. They are particularly worried about who you can use a frontier model to generate bio weapons in cyber and how you can demonstrate how you are being appropriately careful with that. I have not yet had a chance to grasp what the implications would be. We intend to be able to do work on frontier models. We will get back with you.

Question: Looking over the past year, how does DARPA and the DARPA programs that pop up, how do they stay relevant with the fast-paced advancements in AI? How does DARPA maintain relevance when it is that fast paced?
Answer: One area is by program structure. The AI Cyber Challenge (AIxCC) is a competition where we partner with large language model (LLM) companies (Anthropic, Google, Microsoft, and OpenAI) to provide compute access to those in the competition.

As the capability advances so too will the performers using them be able to leverage the advanced capability at the same time. That is one model. Another piece is that we will be keeping an eye on what is happening if the capability that we are working on in the program becomes outmatched, we will stop the program and regenerate or do something else. Another thing is that not all the frontiers are advancing at the same pace. Reinforcement learning is not going as fast as the transformer model. The pace of the frontier models is slowing down a little bit. A lot of the results that we are seeing right now include understanding what they are doing and what they are not doing. They haven't released GPT5. They haven't really even started training GPT5 due to the slowdown in the release of the H100s due to the production problems at the Taiwan Semiconductor Manufacturing Company Limited (TSMC). So, we have a little bit of breathing space. The Gemini model, getting the planning piece integrated in the LLM, we are not sure, we lack full transparency, but there are large research problems that still need to be solved. Hearing people say we're "just a little bit away from full artificial general intelligence (AGI)" is a bit more optimistic than reality. There are things like the halting problem. We still have exponential things. We still need resources. I think there are still going to be super hard problems that are not going to be fixed by scaling.
Follow up Question: My question is, when you have, you might not have AGI, but you might have a system that helps humans and everyone in this room to advance so quickly that before AGI comes this apex where not an apex this asymptotic growth where we are dealing with that constantly?
Follow up Answer: We try very hard to not get in the way of what industry is going to do. We are trying to solve problems and work on problems that industry isn't going to do tomorrow. We aren't planning to work on multimodal large language modules because they are going to do that some time. We are not trying to work on incorporating new information into an LLM because they are going to do that as soon as they can. We are trying to work on things they won't work on right away. We haven't done this yet but we might do multi-level security because we think that is something the DoD might care more about than industry would. Maybe that is on industry's road map, but in a further future time frame. I don't know what the right answers are, but the question of "What are they going to do and in what time frame, and what should we do?" is something we talk about all the time. Do we have perfect answers? No, but do we ask that question constantly? Yes.

Question: You have been pointing out here today, code generated by AI systems is just going to increase in scope and scale in ways we can hardly imagine? How important is it to DARPA that that code get verified for correct functionality and security properties?
Answer: That's a thing that we have been noodling over quite a bit. Clearly companies are going to be working a lot on generating a lot of code. We are not so sure they are going to generate code that is high quality or care about generating code that is of high quality. Clearly generating proofs about code and generating specifications, specifications, codes, and proofs are all languages those are all in the wheel house of LLMs; tying them together could be hard. Definitely noodling over trying to generate

specifications, code, and proofs that are checked. I recommend you speak with Dr. Velasquez and Mr. Martin. There are tons of code on the web, a lot of it is not good code. I believe a study from five years ago from Stack Overflow there usually is a good security answer to a question but it is usually number 10. That means there are 9 bad answers before the good answer.

Question: What is DARPA specifically interested in related to protecting electrical power systems and their industrial control systems?
Answer: DARPA did the RADICS program a while ago about exactly this problem. Some of the results from that were really quite interesting. The initial response from the power industry was like “We completely know how to cold start a power plant. This is in our wheel house we do this all the time from hurricanes and a natural disaster.” The part that was not so much in their wheel house was how do you do that when your sensors are lying to you. Which of course is completely in the wheel house of attackers who take over the output of sensors. So, I think the program was a success from the point of view of opening the eyes of the power industry of what a cyberattack would look like. It transitioned a bunch of tools to power grid operators around the country. So, we don’t have any current efforts. There is some stuff going on in that space still. The current approach or current effort in power is part of the larger effort of cyber infrastructure or infrastructure in general which is the AIxCC effort which got launched at Black Hat in Las Vegas which is: “Can we use AI based tools to help automatically find and suggest repairs to open-source software?” There was a paper that came out a few months ago that showed ChatGPT out of the box was roughly as good as bespoke tools for finding and suggesting fixes to software but that in addition, its most common answer was “I need more information.” With ChatGPT, you can ask “What more information would you like?” in human native form and then you can have a conversation with it and in the process of the conversation it was able to find and fix substantial and additionally more information. Leveraging that insight, we launched the AI Cyber Challenge, which is focused on open-source software and we are partnered with the Open Source Software Foundation to help guide the competition to the kinds of software that is typical in these power plants and other kinds of infrastructure software to be appropriate for exactly those types of challenges and then be able to feed into the acceptance processes so that humans could then vet the suggested changes and get the software released. In some ways it is in response to the Avril Haines testimony like what if we had to find and fix bugs at scale really really fast.
Sergey Bratus Answer: So, we do have engagement with National Renewable Energy Laboratory (NREL). They have the problem of meeting the Net 0 by 2050 vision. What this means for them is that they need microgrids that communicate reliably and securely on the scale that no one has achieved before. They are looking into formal methods and quick ways of repairing legacy software along those lines.

Question: How seriously does DARPA consider the possibility of software being developed by AI?
Answer: DARPA has a position on the topic. My opinion is that it will be a tool that will help people write software faster. Particularly boring boiler plate software faster but it will not automate the process. I don't think that people who write good code will be out of a job anytime in the foreseeable future. Maybe I am overly optimistic but that seems inconceivable. I think a lot of the boiler plate software like coming in frameworks or something like that, the code everyone hates to write, I think AI will write anyway in the near future.

Question: Can you give the office view of a minimum viable program (MVP) and how you think it is going to affect program size complexity and funding?
Answer: MVP is the way that agency leadership is looking at program structures. What is the MVP, you might think about that as inspired by startup culture. It is driven by what is the core hypothesis for a program. Likely that is a technical hypothesis but it doesn't have to be a technical hypothesis. What is the disruptive affect you want to achieve, how do you think you are going to achieve it, what is the smallest, most efficient program structure to validate or refute the core program hypothesis? That is what agency leadership is looking for in terms of the initial program structure. That doesn't mean it will be the only thing we will do. If we validate the hypothesis, what the agency wants to do is double down and expand that. Sometimes we might not know at program formation time what is the right thing to be doing longer term until we have figured out the core technical approach or that core disruptive program formation approach. We are still figuring out within the agency what that means. I think you will see an evolution over time, I am not certain what that evolution will look like. Some of you asked about the impact on program length and I think this goes back to one of the questions that came up for Kathleen in terms of how does DARPA still be relevant in the face of a super-fast moving field like AI. I think Kathleen gave you a number of very good answers including working on problems that we think are long term and important and that industry is not going after. But part of that is much more agile program structures. Now that does tend to mean shorter programs for those MVPs and then maybe bigger investments on the back end. There is appetite to take on more risk. That sounds odd in the context of MVP. There is appetite in taking more risk in the definition of that hypothesis and what we are trying to prove and then what we do on the back end if we have proved out that hypothesis.

Question: Which PMs would be interested in ideas on computer vision?
Answer: I don't know that we have program managers specifically targeting that part of AI. But obviously it is a part of a number of problems. If you need to have autonomous systems operating in the real world, they need to be able to perceive. Obviously as Kathleen related, some of these foundational models are going to have multi-modality capabilities. I would encourage broad engagement with the PMs working in the AI space – Dr. Shafto, Dr. Marge, Dr. Corvey, Dr. Velasquez.

Question: Are there opportunities at DARPA for small business to network with other small businesses?
Answer: That is a place where we need to do a little bit more work. I don't know if we have formalized structures to enable that but maybe we should, and I can have some conversations with our Small Business Programs Office to enable that. Having these sort of outreach events, these Proposers' Days, on a regular basis is one mechanism. That is part of the reason things like lanyards, to help people network over technical areas. Kathleen and I talked about more events similar to AI Forward. We found that a very useful way to engage with a portion of the performer base. We are thinking about how to take that mechanism and process to have broader base engagement with our performers. I don't know if I have a concrete answer on enabling networking across the small business but I will take that back, that it is great feedback for us.

Question: Has DARPA run numbers to converting SBIRs to commercially viable products ; i.e., percentage of efforts that go and become commercial products and come back around? As a follow-up to that, what does success look like for the Embedded Entrepreneurship Initiative (EEI)? Do they raise their percentage of programs that go from being seedlings to SIBRS or reduce that timeline to do that?
Answer: I will direct you to our Small Business Programs Office and commercialization team for the exact numbers. I would not be surprised if we are paying attention in terms of quantitative performance. I think there is also a number of dimensions for how that EEI can be successful. That was focused on helping create companies, helping create industry, helping turn small businesses into larger business and as our Commercial Strategy team pointed out earlier, sometimes there is a stumbling block that is not technical and that is what the commercialization team and that effort is designed to help with. If you have a concrete issue that you feel is not being addressed, happy to hear that, if others in the audience have feedback, we are happy to take it.

Question: With the arrival of the MVP construct, are DARPA seedlings no longer available to your DARPA program managers?
Answer: The answer to that is absolutely no, they are still available. There is still a blurred line between a new start and a seedling. One of the dynamics that we don't love is "new PM has brilliant idea, they are channeled down this seedling path, they study it for a year, then they come to the new start." We would much rather, in many cases, lower the bar of the new start. We want messier, earlier, half-baked new starts so we can find the right ones and have the higher velocity and higher throughput. In many cases, that seedling isn't necessary. For the MVP we don't need that level of evidence, we want the really disruptive concept and we aren't violating laws of physics, and it has some massive national security impact that we care about. If it has those ingredients, it is going to be a go. You don't need to study it quite has hard before you present the new start. All that being said, again, there is no one size fits all. Certainly, program

managers will still be executing seedlings to go explore ideas before they even get it to the point where they know they can conceive of what the right MVP might look like.

Question: Regarding technology transition, can you talk to some of the options that you are exploring to facilitate transition of the science to programs or platforms?
Answer: So, we have a program called Constellation that is an I2O program that is sort of a joint venture between DARPA and CYBERCOM. It is constructed in a way that, especially in this domain, that the velocity of change and when there is operational value for a technical insight from a DARPA program and the traditional mindset of a DARPA program over a structured number of years with an output and a transition is becoming nonsensical. Constellation is an experiment for how we might think about that type of activity differently. In essence what it does is constantly survey the I2O portfolio and when there is an operationally relevant insight or exploit or capability, it immediately starts to bundle that into packages within the Constellation structure. Which then CYBERCOM is reaching into and we have IDIQ over a set of performers that are entities that speak R&D and are also integrated into operational tools and they're the halfway house that bundles it up hardens it puts the documentation on it and puts it onto the operational floor. That isn't transition of any one DARPA program or transition of any particular program of record on the CYBERCOM side; it is a different way of thinking about continuous transition. That is one example. There are other models like product centers. There's something in the DOD called the Tactical Assault Kit (TAK) Product Center. Services pay into a central entity that manages development and integrates new code that can run on this platform and then serves it out to that entire community. The services don't each have a program of record that is funding an R&D and integrating their own thing. It is a community joint activity that sort of gets the capabilities out distributed out very quickly and efficiently. We see examples like that, some that we are doing, some that we see are observing out in the world and we like those features. Constellation is still a very early stage I2O experiment, but if it works it is something we are certainly watching to expand across the rest of the DARPA portfolio.
Follow up Question: As a small business is that a PM referral or how do you connect with this?
Follow up Answer: There is a process with DARPA and CYBERCOM and the PMs will pull things into it as appropriate. On the small business side, my perspective is that the best thing DARPA can do for a small business is link it into our full programs. We have this ecosystem; all of our energy and activity is into making these programs successful and transition them and we want our small businesses in the mix with that exposing their capabilities to a broader set of performers. It creates this opportunity for the magic to happen for whatever the next step is. SBIRs are a really great tool for PMs to do in the margins and I don't want to get rid of that, but I want to drive the program more towards in the end state every BAA has a small business topic released on top of it. And so, we would be co-funding SBIRs with a program. You would go to the program reviews; you'd be integrated into that ecosystem. Not all of that is done yet, I am not sure all of that is legal yet but that is the direction we are exploring.