

Essential Science Conversations

ARPA and You: Research Process and Funding Opportunities at Defense Advanced Research Project Agency (DARPA)

Tuesday March 7, 2023

Juliane Baron: Hello, good afternoon, good morning depending on where you are today. My name is Juliane Baron and I am the Executive Director of the Federation of Associations in Behavioral and Brain Sciences, otherwise known as FABBS. I'm delighted to welcome you to today's Essential Science conversation. ARPA and You: Research Process and Funding Opportunities at the Defense Advanced Research Projects Agency, otherwise known as DARPA. For those of you unfamiliar with FABBS, I just wanted to mention we are a coalition of 29 scientific societies with a shared goal to advance our sciences and provide a united voice for our disciplines here in Washington DC. I am honored to be joined by my co-host for this event, Dr. Adrienne Stith Butler, Deputy Chief Science Officer with the American Psychological Association.

We've also been putting together this series with our colleagues at the Association for Psychological Science. As many of you are aware, the federal government is in the process of standing up a new agency, the Advanced Research Projects Agency in Health or ARPA-H. FABBS and APA have been actively engaged in conversations about ARPA-H, both making the case for the importance of including behavioral and brain sciences in the fabric of ARPA-H, and also advocating for funding opportunities for our disciplines. As we thought about ways to engage our members of ARPA-H, we realized that our researchers tend to be more familiar with NIH and NSF funding mechanisms, and that it would be valuable to pull together a series of webinars on the ARPA model, even though they're different in a lot of ways, just to both increase familiarity with the ARPA model, but also to bring attention to existing opportunities that maybe our members are less aware of.

Today, we have the great opportunity to hear from three speakers. I will go ahead and give very brief introductions to each of those speakers. I really hope that while you listen to them you'll put some of your questions into the Q&A. First, we will hear from Dr. Bartlett Russell, the Deputy Director of Defense Sciences Office, where she has served since August 2022, having joined DARPA as a program manager in April of 2019. Her career has focused on the human domain and human effectiveness and military context. This includes understanding variability of human cognition and social behavior to enable the decision-maker to improve analytics. Next, we will have the opportunity to hear from Dr. Pirolli, a senior research scientist at the Institute for Human and Machine Cognition.

His research combines cognitive science, artificial intelligence, and human-computer interactions with applications on human-AI teaming, mobile health, sense-making, information, foraging, and measuring and enhancing cognitive capacities. I'll also mention he is named in 74 US patents. Then we will turn to Dr. Lebiere, who is research faculty in the psychology department at Carnegie Mellon University. His main research interests are cognitive architectures and their applications to psychology, artificial intelligence, human-computer interaction, decision-making, intelligent agents, network science, cognitive robotics, and human-machine teaming.

That might sound a little intimidating. Those are some pretty tough topics, but I'm really excited to hear from them today. I think we'll go ahead and start with you, Bart. You're on mute.

Dr. Bartlett Russell: All right. I'm going to share my screen. I do have some slides here. Thank you for inviting me. First of all, we're always happy to try to demystify DARPA and ARPAs more generally. Any way we can do that, we welcome because we are in the business of always getting the best ideas in the door that we can. I think we have a shared mutual interest there. Quick check that you can see my slides. Are they coming through?

Juliane Baron: Looks great.

Dr. Bartlett Russell: Excellent. All right, I'm going to provide a very high-level overview from DARPA in particular on how we use things that is, as you pointed out, a little different than maybe how NSF and NIH approach problem sets and approach research funding. I speak only for DARPA, of course, and particularly for DSO, which is our most basic science office within DARPA. As you said, with the ARPA model being exported to address other very critical problems across US industries, a lot of the fundamentals and focusing on here really will hold.

That said, one of the driving features of ARPA is that everything is program manager-driven and therefore there's a lot of variability. I'll say at the beginning, and I'll also say at the end, if you are interested in finding out more about how DARPA works and how to work with DARPA, the best and first place you should go is always to find a PM and go talk to them. It's taking a minute to advance. While it's doing, I'll talk ahead of that. It is not advancing.

Dr. Bartlett Russell: It's not advancing on my end, I suspect it is not on your end either. Let me try that again. Why don't I just talk to things while I try to troubleshoot at the same time? What I am going to start with is that one of the fundamental differences here between ARPA or between DARPA and NSF and NIH is that we are focused on 10x or more improvement on a given capability. We're not looking at incremental advances, which if you're doing an R01 grant, if you're doing an NSF grant, you're looking at what is the next right scientific step to go after, in order to advance the field.

In ARPA and DARPA, we like to think about what is a fundamentally different way to think about the problem that we can maybe disrupt the way the herd is thinking about it and have a 10x or more improvement. In most of our BAAs, you will see the language that says, 'specifically excluded is incremental advancement or solutions that will provide on incremental advances'. That's something that really is fundamental to what we do, and often it helps us answer the question of should we be in a space if we think it isn't incremental advancement, then we won't.

I don't know if my slides are going to work. Why don't we just make them available afterwards if that works for you, because it has a lot of good information in there? Does that work, Juliane?

Juliane Baron: Yes. If you don't mind, maybe you can keep talking and, Shandol you can go ahead and share the slides. Should we do it that way?

Dr. Bartlett Russell: The PDF version? Yes, okay.

Juliane Baron: Let's try it.

Dr. Bartlett Russell: In part of that, we obviously in order to make those ones we have to take high-risk approaches, very high-risk approaches. Thank you, it's slide two. Perfect. The reason we can take high risk is because we're also willing to be our own harshest critics. We can start a program and if we have really good metrics salted in throughout the program, it means that we should be willing to cut them because if we're being high-risk enough, then not all of them will work. That is fundamental.

That's something also to know about DARPA projects, is that if you're involved in a DARPA project, and you're involved in a ARPA project, it is never a guarantee it will go through its entire four years, five years, however long the period a program is. We're constantly evaluating and making sure it still makes sense to push taxpayer money in this direction. Another point is that we are very problem oriented. Again, for an R01 grant, you might often say that this bit of research could eventually, may one day, really help us with something like suicide or something like depression.

We bring those used cases and those impacts much closer and say, "What is the thing we are going to do?" As a result, rather than have fund one or two researchers to work on separate versions of that problem, we'll often look at whole big teams that come together multidisciplinary approaches in one team to solve a problem. If you have not worked with DARPA before, think about how the problem can be solved not just from how you would do it, but also from what would be the teaming arrangement that would help solve the whole problem, not just one part of it.

All right, next slide, please. I started with know your PM and talk to your PM, go talk to your PM. It's also really useful to know what office things are coming out of. Different ARPAs are structured differently. This is the way that DARPA is structured. We have four what are known as tech offices from defense science office where I sit to BTO, that speaks for itself. It's a biotech office. I2O is where you'll find a lot of the computer science and AI microsystems technology, a lot of emphasis on microsystems these days, all very technically oriented towards the technical application space.

STO and TTO, where it's called our systems offices, those are much more applied. They create bigger programs, more applied. The 6.3 version instead of the 6.-1 where I sit, if you understand what that means. 6.1 is the most basic science, 6.3 is more applied solution. Those are the bigger teams. They want to go create things that actually do things in the world. They might build a satellite; they might build a plane. You've heard of DARPA's X-plane. That's the place where that happens.

Knowing your office is really going to help and knowing an office that a solicitation comes out of is really going to help you understand how to respond to that solicitation and how risky to be. Whereas in the defense sciences office, we get to be really out there and really think about how can we fundamentally change computing, for instance. STO and TTO is going to think about how can we take cutting-edge but existing capabilities and bake that into fieldable systems today.

All right next slide, please. When you go and have a talk with a PM because again, that is where the rubber meets the road. We don't impose a whole lot of top-down direction from an office level and from an agency level, we say, where are the opportunities and PMs come make them a reality. We empower the PMs, but all PMs are asked to answer these questions when they put together a program pitch.

When they go justify their program idea and ask for money to go pursue it, they have to answer these questions. It is a lingua franca within most of the ARPAs into how we think about problems and what rises to the level of an ARPA problem, a DARPA problem, as opposed to say again, the next incremental research approach. These are much harder to answer than they look like.

When you read them, they seem very simple, very straightforward. Let them not fool you. They're extremely difficult. Question three is one of the ones that we deal with most, that we work with most in DSO. They're all important, but it's the crux of often what makes or breaks a program. It's not just, why do we think we can be successful? It's what is new that other people are not doing?

Why is my unique approach going to succeed when so many others have agreed that this is a hard problem? Everybody wants to solve a problem, what is the thing that is fundamentally new that you are doing that nobody else is doing? That's what three is and that's the hardest problem. We're all very good at identifying a lot of the issues out there.

What is the new approach that's going to get us there that hadn't before is very important for DSO in particular, but really all of the offices. Another thing is we need to be able to measure it. Metrics. We know there's midterm and final exams. That's how we know if our high-risk bet is actually making dividends. Is it something that's paying off, or maybe it was a good try, but maybe we think about other things.

If you can measure and you have a good way to validate your system in the wild, these are really good ways to talk to a PM. When you go talk to a PM you say, "I've got to identify this PM, I think they're interested in my space. I'm going to find out what they're up to." This is a good way to structure what you bring to the table and what your research brings to the table.

Next slide. I'm going to move now and talk about some newer initiatives that DSO in particular is after, and then DARPA as well. I bring this up for two reasons. One, because I want to get them out there and I want you to know about them because they're hitting at some of the fundamental questions that I think you're interested in. How do we demystify DARPA? How do we make sure those good ideas get in the door?

Again, we're interested in the same thing. These are our mechanisms to try to get there. The other point though, the other reason for bringing this up is that even if you've worked with DARPA, say five years ago, even two years ago, it is the nature of a DARPA always to change. Nature of an ARPA always change. If you think you know how they're doing things, keep in mind that internal, we're always thinking about how to do things differently.

Don't assume just because you worked with DARPA once, you always know how to work with DARPA. If we're doing our jobs correctly, we should be adjusting the way we do things to keep in time with all of the-- industry is changing. The triad between academia, industry, and government and pushing forward defense people, that's fundamentally changed. We have a much more BC-driven world we have to adapt as the world adapts around us.

These two particular efforts here are initiatives. These are experiments. They're experiments in engaging with the community are to lower the barrier to entry. The first is what we're calling advanced research concepts. The idea here is that often time when a PM is starting with a new concept, they will put a little seed funding, we call them seedlings out there to just answer a question. If this can work, then I should launch a bigger program. The idea, if I get this little bit of information, can it help me launch a program?

The problem is, is that really suffered from confirmation bias. Everybody's incentivized for its work, and then you don't really know, did I get the answer I wanted or is this motivated reasoning towards the answer I wanted because I wanted a program and because the performer wanted to give me a good answer that I liked? Instead, we're saying, we want seedling farms. We're going to pose a what-if question and solicitation was extremely fast, contract rates we're talking like within a month, and instead of funding 1Zs, 2Zs, we're going to fund maybe 30 and just see what happens. Oftentimes what you read in a proposal is not what ends up being the most effective thing.

We said, "We'll spend a little bit more money to get to see what works first and then up select from there." That's what the ARCs are coming up. We have our first one we launched this last year. We had our first topic posted in Quantum. That's on the street now. These are open also for six months. Oftentimes, people who are most familiar with DARPA know that you have to be on it. When a solicitation drops, you've got essentially 45 days to respond. We're trying to break that a little bit. Opening up, we've got six months to respond to this, about 300k for each seedling.

You have a year to do it, about a post-doc level of support over a year to get this done. Now, you might ask, "That's a whole lot of little seedlings to manage. How are you going to do all that?" Is the other side of it. We also launched an innovation fellowship. Our PMs are leaders in their field. Oftentimes, we have people come in very late in their career who really want to disrupt things. What we're trying to do here is bring people in earlier in their careers for two reasons. One is to demystify DARPA. These folks will be here for two years.

They'll understand how DARPA works from the inside, and then they will go out and work in industry, academia, wherever, and help again, bridge that gap between and demystify how DARPA works external to the building. Second one is also the cross-pollination. These fellows are working together not in isolation.

We have everything from environmental scientists, biologists, quantum scientists working together on same problems to bring multidisciplinary approaches. Last slide, please. This last one is a DARPA-wide initiative, and it goes to trying to overcome the challenge of sometimes, well, not sometimes, writing proposals in itself as a

learned art. We would really hate to leave great ideas on the table because people don't know how to write a DARPA proposal.

DARPA Connect again, just launched. This is very, very new. You can see that the email is there. The idea is to help non-traditional organizations, non-traditional performers learn how to read a BAA, learn how to read a proposal so they understand what's being asked for, and that they can better respond. We don't miss good ideas just because somebody doesn't understand how to write a proposal. There are more functions than that coming out of DARPA Connect.

There's going to be a series of events going on across the country with these discussions, but I wanted this community to know about it because I think it's targeted exactly where you're after here, which is again, demystifying. We don't want to leave those good ideas on the table. I'll leave these slides for you. I'm happy to answer any specific questions. While this is very DARPA specific, I will say that again, the fundamental principles about going to your PM, talking to the PM first, and thinking about how your questions in Heilmeier is going to be lingua franca across all the ARPA's for the most part.

Juliane Baron: Terrific. Thank you so much for that really useful overview and some very specific pieces of advice for our members. That was really terrific. Now we have the opportunity to hear from the researchers who have interacted with DARPA, and so they can speak a little bit more from that perspective. Peter, Dr. Pirolli, let's turn to you, please.

Dr. Peter Pirolli: Great, thank you. I have been lucky enough to get funding from a variety of agents including NSF and NIH, but I wanted to give you a ground-level view of what it's like to work on projects for DARPA and IARPA and other ARPA's, and a little bit about the uniqueness of writing proposals for DARPA.

The first thing to say is that DARPA and the other ARPA's typically support research that has at least implicitly a duality of being mission-oriented or impact-oriented, as well as supporting some real 10X fundamental or applied research advances. Christian and I have worked together many times and we're currently involved in programs or have been involved in programs that have done research on human information foraging and sense-making, more recently around human-AI interaction and teams.

I've worked on DARPA projects where we've developed new multisensor speech, EEG heart rate, aim-like measurement systems for really getting at novel assessments of cognition and resilience, non-invasive stimulation to improve learning. These are all things that I think are generally interesting to APA and APS. They're clearly of relevance to society, but there's really fundamental research to be done there.

What is it like to work on these projects? Typical DARPA projects are multi-organizational. They're multidisciplinary. The ones that I've been involved in typically involved working with computer scientists and other psychologists, and sometimes neuroscientists. I've worked on some where we've had geneticists and people doing human performance. As Bart mentioned, one aspect of these projects is that the programs have milestones, and they have metrics, and they're taken seriously.

Now, they sometimes can change along the way, but they are things that drive the tempo and logic of the research that you're doing. Just to give you, for instance, Christian and I work on a program called ICARUS, where there were multiple teams. At various stages, we had to model experimental data and predict some number of biases. You start off, you model a few biases, and then a few months later, a few more, and finally, by the end of the program, you've modeled to some level of measurement about a dozen different cognitive biases like confirmation bias or anchoring.

That means that it really is a team scientific project, which requires organization and coordination. A typical good team will have a good leader, someone who's got some DARPA experience typically. You have to have lots of communication, so fairly regular meetings and clear roles and goals, and accountability. Often these projects have significant reporting in contracting and wrangling components to them that it's good to have someone who's essentially a project manager or administrator who deals with a lot of that stuff and leaves the science to the scientists.

In getting a team together, the issue of having common ground is an important issue because it facilitates communication and trust. Often, it's good to have people who are either colleagues or have worked together or at least have partial research overlaps with other members of the team. For example, people who do cognitive psychology but are also involved in machine learning might pair with someone who's within machine learning and also understands something about human cognition.

Developing the relationships within a team is not an insubstantial thing. I've been on some teams that did not work effectively, so I speak from experience. It's good to know who you're working with. That's what it's like to work on the projects. Then, the proposal process itself is different than NIH or NSF, or other agencies that people might be familiar with. I do think of it as a race. Typically, what happens is DARPA and the ARPA have various places where they will post "news" or The Broad Agency Announcements for novel programs.

Typically, there will be what's called a proposer's date where the program will be presented, and you can register just like any other conference, and that's where you'll hear about what the goals of the program are and all the various resources that you'll have access to. At that point, essentially the clock is ticking and at that point, you really should be thinking about the proposal and the team you want to put together, et cetera. One thing that's useful to know is that often at these proposer's days, there are organizations and companies that have been part of the DARPA ecosystem for some time.

Often, they have lots of experience, but they're not complete, and they're looking probably for people like you to partner with. You should be aware of that. Typically, there is a deadline for what's called an abstract, which is a five-page prospectus for what the proposal would be. Then there's a decision about that, that would tell you whether you're encouraged to submit or not. That's not a hard and fast decision, but typically it's a very good indicator that your idea is either good or not good. Then there's the proposal.

Within that, pretty much as soon as possible, as part of the process, you want to reach out to possible collaborators, try to assemble a team, identify people who you

think might lead various subparts. I can't emphasize enough that the Heilmeier catechism is a real crucial thing and whether you write it directly into your proposal or not, you should have gone through that and answered every question and iterate on it until you really feel confident that you've got good answers for all those questions. Basically, you want to get to a point where you have some secret sauce.

You're trying to tell people why it's innovative, and why it can be done now, and why it matters to the program, and why you think you're going to hit all those metrics. Also, as soon as possible, it's important to get a statement of work. What are the things that you really are going to do over the two, three, four period time and put names next to people? There's a variety of reasons for doing that. It really does get you down into the nitty gritty of what you're going to have to do.

It's also part of the proposal process in the sense that to put the proposal itself together requires a lot of financial and administrative documentation, and that statement of work drives what your admins and finance people are going to have to do. Very carefully read The Broad Agency Announcement in detail many times and get your collaborators to do it because it will tell you what they will reject. Identify all the statements that have the phrase, "Must have." Your proposals must have something, or your project must have something.

All the statements that say, "Should have," turn that into a checklist, and make sure that your proposal has it. Then, finally, there is a frequently asked question process where you can submit questions and get an answer that will go to whoever wants to look at it. You want to monitor that because often there's a lot of very good clarifications that go in that. DARPA will provide templates for the abstract and proposal itself, and within that, it'll say, "This section should describe X, Y, and Z, or state X, Y, and Z."

A thing that I typically see in a lot of proposals are typically a great unifying graphic that communicates like a marketing logo. What it is you're going to do. A very clear table of claims and significance to the program that is linked to the text sections because your reviewer is going to want to summarize it for other reviewers. That is the thing that gives them a mental model of your proposal. A table that describes the team and the expertise and how it's linked to various aspects of what needs to be done in a concise way.

Then some table that tells the reviewers, what the risks are to your research and your backup plans, the mitigation. Nothing goes as planned, and so you should always have a plan B. Anyway, that's like from the ground level, what it's like to work with DARPA. I look forward to any questions.

Juliane Baron: Terrific things and some very specific and useful advice, I think. [chuckles] Thank you. Christian?

Dr. Christian Lebiere: All right. Bart and Peter have already covered a lot of ground, so I'll try and be brief. First, I want to emphasize something that both of them mentioned the fact that DARPA programs are inherently multidisciplinary. I think that's a good thing in a natural fit for psychology and behavioral sciences, and cognitive sciences. If you look at the mass head of the Cognitive Science Journal, it says, "A multidisciplinary journal," with disciplines ranging from anthropology to

artificial intelligence, and presumably any field of human activity can benefit from a psychology or cognitive science contribution. Herb Simon himself pointed out that the likelihood for non-incremental advances which is a requirement for DARPA that Bart mentioned is likely to be found in that multi-disciplinary space rather than in well-trodden path within a single discipline. The challenge as Pete mentioned is that putting a team together and finding effective points of contact is very difficult. It can take an order of magnitude longer than you think it's going to take.

Finding collaborators that you've worked with before that at least understand your work is invaluable. The payoff of that is that you get to test the assumptions and limitations of cognitive theory and application domain that use-inspired basic research past the squadron. It's often something that Peter mentioned confirmation bias. It's often a confirmation bias of just you have a theory or a tool or an architecture and you're just applying to the things that it does well.

DAPA will force you out of your comfort zone, and that can be really invaluable. Speaking of being out of the comfort zone, one key aspect of DARPA program is that they move at a breakneck pace. The advantage of that is that again you will fail quickly. You will fail very quickly, test and refine your approach and figure out what doesn't work, and then go again. The cycles of the inner loop is as fast as six months, and the outer loop is maybe a year or a year and a half.

That's the speed at which you have to operate, and that can be challenging. Sometimes you don't always have the time to go and cross all the Ts, dot all the Is and carefully publish your research at your leisure. You have to plan for that and figure out how to overlap that breakneck pace with the traditional research and scholarly activities. Then the last thing I'll mention is for cognitive scientists and psychologists, being part of a DARPA program can involve different kinds of contributions.

A lot of it that I'm a computational modeler so a lot of the contributions take the form of software. Some kind of software there that embodies the cognitive theory. That's one way to do that. You can take a more experimental or theoretical approach, and that's another contribution to the team. It's important to define those contributions upfront and figure out where you are on the critical path of meeting those metrics that Pete mentioned and how because that will determine your work cycle there in terms of what you can do and how you should do it. That's it.

Juliane Baron: Terrific. Thank you so much. You have me thinking about what APA and FABBS can do together to be helpful too. Before I turn it over to Adrienne for the first question, Bart, I was just going to ask if you had any responses or reactions, either about what Christian or Peter had to say.

Dr. Bartlett Russell: I don't disagree with anything they say. I think what they're saying from the respondent side is very similar to what I'm saying. I used to write proposals on the other side too, so what they're saying I've been there. The only other thing I would say is that sometimes, and forgive me if this is repeating, I've been answering some questions in the Q&A. If a BAA drops and it's the first time you've heard of it, just be aware that that might be a little too late and that you should be part of maybe somebody else's team.

The best time to learn about an effort is when the PM is thinking about it, when they are conceiving of what that program would like. Oftentimes, we will get people that see a program associated with my name, for instance, and they'll say, "Oh, I'm interested in your Habitat program. I'm interested in this other program. What can I do to get funded?"

What they don't realize is that once the BAA is out and once we've solicited proposals and funded them, the money is basically spent. We've already placed our bets on the table and we're just running those programs and managing those programs forward. That's usually not the time to go after a particular program. That might still be a time to figure out what the PM is interested in and go talk to the PM and see if they're doing anything else in that space. That's an important distinction. It is not enough just to wait for the solicitations to drop in most cases, and most beneficial if you have a sense of what's coming before it's coming.

Juliane Baron: Can I just ask? How do you do that? Where in that process is a proposal day because I was thinking, "Oh, we need to definitely get some of our members there." How do you just proactively-

Dr. Bartlett Russell: Go on the website and look at different PMs. We have 100-some PMs. You look at what they're into, you look at what their interest areas are and you use the-- Most often we see a website, it has a generic email. You're like, "That's never going to get answered." Ours does. They get shunted to specific PMs. I will tell you I'm answering those on a daily basis and it's a priority for us to answer them.

Even if it is a polite, "Maybe this doesn't really fit," just getting in the door. "Can I have a half an hour meeting to talk to you about this idea?" That is the best way to do it. It all really begins and ends with just going to talk to PM.

Juliane Baron: Excellent advice. Thank you. Adrienne, can I turn it over to you to kick off the questions?

Dr. Adrienne Stith Butler: Certainly. Thank so much, Juliane, and thank you all for such great presentations. This was incredibly helpful and useful very practical information. It seems like one of the biggest challenges and opportunities is finding partners. I think, Peter, you mentioned that there are sometimes companies or other entities that have holes that they're looking for partners. I'm wondering if you can talk a little bit about the process of finding partners and ways to go about doing that.

Dr. Peter Pirolli: Often, you can look at people who have participated in other programs, and I'm guessing that if you talk to a program manager, they will say, "Oh, yes. This company or that company was participating in some other programs and projects." Then at the Proposers Day, often you can actually see the list of people who are attending. Like I said, some of these companies are very large and some of them are small. Often, they have a particular focus, and it may not include having psychologists on the team.

You want to look for some of these where you may actually know some people that are in that organization and often, they're more than happy to talk to you because they're looking in the other direction. They're saying, "Look, we don't have everybody

internally." When a new program comes out, that might be about suicidality or something like that, we need to get people who know about suicidal ideation or about FMRI or whatever.

We need to be able to rapidly assemble those people to go forward, and they will contact you if you've made contact with them. Again, it's like a lot of other relationship-building, where you might want to find colleagues in other departments. You just have to get an awareness of who participates in these programs, who might be going after projects in the past that you think are interesting.

When you see these names on a Proposers Day, they've probably already got a good idea about what they're going to do because they've been talking to people like Bart and so know what the program is going to be about. I think that's one way that you can get your foot in the door and start participating on some teams.

Dr. Christian Lebiere: If I may add something. Pete mentioned Industry Day, which is the official unveiling of the program either before or after the BAA is out. I was part of one. I attended one last week of an IARPA program. The first part of the Industry Day was the program manager detailing what the program was about and some of the contractual specifics and the other nitty-gritty.

The second part was quick lightning, five-minute presentations by researchers or teams interested in the program, and typically that took the form of, "Here's what I do. Here's how I think it would be relevant, and I would be interested in teaming with someone who can provide this or that piece of research or capability." I think that's another mechanism there for matching the pieces and finding people who are complementary to your knowledge and specialty and putting a team together.

Dr. Adrienne Stith Butler: Thank you.

Juliane Baron: I had some questions too thinking about our members who are typically faculty and not thinking about proposals or research in this way at all. I'm just curious, is there a moment where they get the bug, or is there a moment where they want to take that chance? I think our researchers, basically the whole point of their career, is to be accurate about something. Being wrong seems a little bit scary. You know what I mean?

I just wonder if you've seen some transformation or some ways to help some of our researchers who might not think of themselves as this high risk, or on these very quick timelines, I don't know go with a hunch and go all in. I'm just curious if you see that it seems a little bit of a cultural difference, just trying to figure it out.

Dr. Bartlett Russell: Maybe but, isn't it really the scientific drive to try to disprove hypotheses constantly, not necessarily prove them? By definition, we've gotten into this kind of, you're right. Culturally, we don't want to be wrong, but is that really how it should be? I'll just offer that. Most often, it's not though people are going out on a limb, sorry. In most cases, people have-- and any researcher, I know I've talked to, they always have-- they're working on 1 thing, or 2 things, or 3 things, but they have 10 other ideas knocking around in the back of their head. They're always thinking about different ways to apply it.

Most often we find the right-- when it really works you get the right alignment of that idea that they have good foundations to bring forward to say, "I've got this hunch and haven't had the ability to work on it yet, but what if we could do this, and let's give it a shot." To go back to some of the things that Peter was saying about looking in the BAA and finding out what is a must-have versus a what would be nice to have.

Another thing that is really important is if you're going to make a claim, it's much more effective if it is justified. Unsubstantiated claims get flagged when we're reviewing proposals, and we're like, "That sounds like a tall order, but they've provided no substance as to how they're going to get there." Being high-risk doesn't mean being void of justification in science. It means that you just extend a little bit. I'll just offer that, we, and especially in DSO, we're working on really fundamental science, we're not looking to break the laws of physics. We understand what the theoretical limits are, we're trying to push as far up against them as we can.

Juliane Baron: Thank you. I also had some questions about PMs, if I may, two questions, I guess. One, how do you become one? I think about our sciences, both as horizontal and vertical, it's woven throughout, as Dr. Panch at NSF likes to say, and also standalone, some fundamentally behavioral science questions. This probably is more for ARPA-H that we're looking for, but just questions about how you become it sounds like quite a powerful PM.

Then two, are there structures around the PM to help them? Nobody can anticipate everything or know what they don't know, but that feedback of the back and forth of, "Oh, and have you thought about this or explored this or don't miss that?" I'm just curious about those two pieces around PMs.

Dr. Bartlett Russell: I'm not sure if everybody knows this. I know that a PM is only here for four to five years. It is part of the model that we have very high turnover on purpose to keep the ideas fresh. That means that we're always looking for people, and the best way is to email me or another office director, or office DD to say, "Hey, I'm interested in PM, and here are my ideas. Here are the things that I want to work on." That is key.

We don't hire people just based on their resume. They could have an amazing resume, but if they don't have any new fresh ideas on how they're going to change the world, then maybe it's not a good fit. We want people who come in and say-- our job talks from seminars, "Here's the list of things I've done in the past." It's usually a slide or two on who I am, and the rest of it is, "Here's what I want to do." We asked usually for one pretty good-baked program idea and then a couple of others to back it up with if the first one doesn't work out. Your second question I wasn't as clear on. Can you just restate that one?

Juliane Baron: Sure. It just sounds like the PM has so much power and influence, and so I'm just wondering are their structures to support the PM? You've mentioned metrics, but it just feels like that's a lot of diverse expertise that one person would have.

Dr. Bartlett Russell: It's too much to ask of any one person. There's a lot of support, the whole organization is designed around supporting the PM. Everything from what is-- with high turnover, you have very short institutional memory. We have technical,

SETA, who provide a lot of that. Any one PM will have a SETA team who can be a mixture of disciplines. I always look for people on my SETA team who were not in my same discipline or complimented my capability. When it came to multidisciplinary approaches, I would have one area covered and my SETA would have the other areas covered. When you're building the program, that's really important.

We also have a lot of cross-pollination across the industry. One of the best parts of being a PM is getting to work with other PMs. You're very much on an island and working on your own programs. The interaction with other PMs and seeing what they're doing and getting their advice is critical to making sure you don't make naive mistakes in an area that you're not familiar with.

There's a lot of cross-pollination across the offices to helping each other say, "We have a PM that wants to get into this," and he touches on all this throughout Quantum since it's not relevant here. They say, "Oh, we've got somebody in Quantum, maybe they can help you," and we put them together our network. It's that thing, but it is very much oriented towards supporting the PM.

Juliane Baron: Terrific, thank you. I've been hogging the question, so I'm going to try and pull some from the audience. Someone writes "I love the point about failure, how you fail quickly, how have Peter and Christian use failure if they have as starting point for something else?" Peter or Christian, do you want to respond to have you failed?

Dr. Christian Lebiere: Well, we fail many times, and obviously, you can't fail too much because you may not continue, but it's important to learn from your failure. In my case, I spent the last few decades developing a cognitive architecture that's meant to represent a fairly comprehensive picture of human cognitive abilities. Every new program is just an opportunity to challenge that and learn what's wrong or what's not there and try to add a new piece to that and grow, as Bob said, extend a little bit.

Necessity is the mother of invention. You're not going to push in that direction unless you have some forcing function to get out of your comfort zone and then do that. The nice aspect of my field of cognitive architecture is that there is a mechanism there in the form of the cognitive architecture to accumulate progress in that particular form.

Juliane Baron: Thank you. Peter, did you want to add anything to that?

Dr. Peter Pirolli: I was trying to think of something fairly specific. Let me just mention one from, there was a recent DARPA program called XAI, Explainable AI. I don't think this was just a failure of our project, but I think what our going-in preconceptions were and where we ended up. One going in pre-conception was the program was about the fact that AI, like ChatGPT or black boxes, nobody really understands what's going on inside of them, but it does all this amazing stuff.

I think there was this naive view at the very beginning, it's like, "Well, if we can just generate an explanation for everything that's going on inside of this black box, people would trust it more, and they'd be able to explain to themselves what the thinking was." Well, it turns out, there's all sorts of subtleties about that. If you

explain something that people already understand, they will immediately discount you in the future for ever being credible.

The technical explanations that you might give are not the explanations that you would want. I could probably go through a list of things that violated the hypotheses that we had right away, but as Christian said, science is about finding out that you're wrong, and you get to a deeper understanding. Those are just fairly concrete examples, I think. Ultimately, I think the program did very well, but we had to rapidly course-correct what we thought were the right solutions to those kinds of problems.

[silence]

Dr. Adrienne Stith Butler: Juliane, can I ask you one other question? Did you have one in the Q&A? I don't want to jump ahead.

Juliane Baron: No, go ahead.

Dr. Adrienne Stith Butler: It's a question about, this is for Bartlett, how are proposals put together within DARPA? How are funding priorities determined, and who makes those decisions in that way?

Dr. Bartlett Russell: Sure, we call them programs. Proposals are what people send us, but it takes a-- I was actually just reading that question in the Q&A. It takes a PM anywhere from I'd say, six weeks to a year to put together a program. The major task is answering the Heilmeier questions. Remember, earlier I said, these seem simple, they're not. If it takes a PM six months of dedicated work just to answer those questions correctly to scope a program, that gives you a sense of how deviling they can be.

Most often it's about creating the right sandbox to drive the outcome you want without specifying the solution space, without specifying the solution. You want to create solution space so you can invite innovation without over-specifying how it's done. That's the magic of creating a program. As far as priorities, that comes partially in an office level, partially at an agency level, and it really depends on the office. From the system's office perspective, it's going to be much more, what do we need when we're thinking about disrupting advanced capabilities, what do we really need no kidding capability in downrange? For basic science questions, it's a little bit more opportunity driven. What is right for disruption?

Regardless of how we come about making those sense of priorities, a good place to see where our priorities are is that each office publishes and puts on an open office BAA on a yearly basis. It's after we go through our strategy effort. In those BAAs, there are general topics, areas that we are interested in. DSO has one. We have six different thrust areas, areas that we are particularly interested in hearing ideas about. That's a really good place to start to get a sense of where we're thinking about going in the future.

Dr. Adrienne Stith Butler: How are those disseminated? How do people go about finding that?

Dr. Bartlett Russell: SAM.gov is a really good place to find them. You can come to our website, and it can point you towards the SAM.gov portal. That's where all of our solicitations are posted. You can get access to them there. It's the same way that any other BAA or OT is posted. It's going to be on that same platform. The DARPA website tries to make it as accessible as possible; you can click a link and go right there. You might have to create a SAM.gov account in order to get access. That's open to anybody. There's a lot of questions about who can attend Proposers Day. For basic science, the idea is to be as open as possible. It's to learn about the potential opportunity.

Dr. Adrienne Stith Butler: Great. Thank you.

Juliane Baron: Terrific. Thanks. I did want to mention the slides and recording will be available after the webinar. Somebody had asked. Just wanted to reassure you that will be coming in the following weeks. Actually, I also had a question for you, just like, what could be helpful? You mentioned that you do want to disseminate and take away some of the mystique. Is there anything that APA and FABBs could do to be helpful to-- We'll certainly be pointing people to this, the tremendous presentations here today. Is there anything else we can be helpful to familiarize or to spread the word?

Dr. Bartlett Russell: I appreciate that. Thank you. Our new initiatives would be great if opportunities for fellowships. I think about this, not as a short term, but this is a long play to make sure we remain accessible and to foster and grow the community. We don't want to be constantly fishing in the same pond. That's not good. That's not part of the art of the model. We constantly want to be good. The more that you have up-and-coming fresh PhD students who are considering a fellowship, maybe they're not sure academia is right for them, and they want to get in the business of seeing how they can push science elsewhere, that's a great opportunity.

Also, keeping an eye out for some of these efforts that are more geared, or these initiatives are more geared towards people who aren't as used to the big, huge, prime, different multi-performer proposals that we've been talking a lot about today. We're trying to get those initiatives out there. The more you can disseminate that, the better.

Juliane Baron: Absolutely. There was a question about eligibility, if working professionals can also apply and is it country specific?

Dr. Bartlett Russell: Apply to which? Be a performer or a PM? I can answer the second question.

Juliane Baron: Yes, it doesn't specify. [laughs] Is it country-specific?

Dr. Bartlett Russell: We are open. You have to be a citizen to be a PM and to be a fellow. Also, that is true. You don't necessarily have to be a performer. Our 6.1 work, it's fundamental research. It is open. We have international performers. We like to get good ideas from anywhere. Obviously, some of what we do is not so open, and that stuff is going to be protected accordingly. Each solicitation should make that clear.

Juliane Baron: Great. Actually, Dr. Mitch Prinstein and I were visiting with the chairs of departments of psychology recently and talking about these opportunities. There were some concerns about if you're leaving a faculty position for three years and just practical logistical.

Dr. Bartlett Russell: We do have IPAs where we essentially borrow people from universities for a while. They maintain their role at the university. They come, spend a little while with us, and then they go back. That's a PM role. What we like to do is we are interested in ideas about how to make that an easier transition. That's something we do very commonly. Usually, it may be associated with a sabbatical, that thing. It's usually longer than a typical year-long sabbatical, but we do that a lot because we are concerned about what happens when you go back. We do that for National Labs. That happens often.

Juliane Baron: I think maybe the chairs were a little bit nervous they wouldn't be able to fill it in the interim, they had a unique set of worries. Christian and Peter, is that what you did? You weren't PM, sorry. This was just part of your research and faculty roles. Okay, sorry. [laughs] Gosh, there are still a few more. I like this question about any current favorite exciting programs, recognizing we have three minutes left, but that seems a nice one to end us on. [laughs]

Dr. Bartlett Russell: I'm going to give you a terrible answer. I'm always most excited about the ones that I can't talk about because they're about to come out.

Dr. Adrienne Stith Butler: [laughs]

Dr. Bartlett Russell: Sorry.

Juliane Baron: Fair. [laughs] All right. Well, then we'll just have to stay tuned and we'll just have to be watching for that. Then just before we let you go, there is a specific question about military sexual trauma. I don't know, maybe we can respond to that afterwards just because it's particularly specific.

Dr. Bartlett Russell: Well, I saw that one. There are a lot of hard problems, and they're very important problems. Not all of them have a technological solution. The nature of the question, I wasn't sure. I wasn't clear if the person asking was interested in a study, that's the thing that maybe a think tank would do that DARPA tends not to do studies. We do develop technology. If there is a technological solution to something like that, then we would put forward the technological solution with maybe that as a particular problem to be solved, the H1, the Heilmeier Question 1.

I think that the bird in that question, the first thing to ask is, is this a technological solution, or maybe is there another solution like organization or policy? We don't do policy and organizational-type stuff. We do more the technology solution. I worked up a program that I was very interested in. It wasn't to do with that particular topic, but I realized ultimately as I was working, it wasn't a technology solution. It was a data access problem. That was fundamentally on DARPA.

Juliane Baron: Terrific. Well, I realized we're at 1:59, I think that I should probably, despite my temptation to try and fit in a few more questions, I just want to say thank

you so much. This was incredibly informative and really helpful. You've given me a number of ideas of how FABBs might be able to be helpful to bring attention to some of the opportunities. Just really valuable to hear from our colleagues in psychology as well. Thank you, Adrienne, too, for co-hosting. I also did want to thank Pat Kobor and Corbin Evans and Shandol for helping make this all go so smoothly.