

HR0011SB20254-10 Predictive Architectures for Decision-Making (PPADM)

Frequently Asked Questions

1. What is the role of In the Moment program referenced in the Topic?

A: ITM kicked off two years ago. The hypothesis is that alignable algorithms will increase a particular human's willingness to trust (instantiated through delegation) AI to make difficult decisions in high stakes domains where experts disagree. We need to understand what to align to. The SBIR Topic is intended to help us explore different approaches to human characterization and identifying key decision-making attributes that are predictive of trust. See the DARPA podcast on ITM for more information: <https://www.darpa.mil/news/podcast/when-should-machines-decide-83>

2. What is the critical unmet need that DARPA wants to fill with this SBIR?

A: This SBIR will produce research that supports DARPA's In the Moment program, particularly the effort to quantify the alignment of algorithmic decision-maker attributes with trusted humans in difficult domains such as medical triage. We are interested in difficult decisions where experts often disagree, not AI trained to conduct a task or perform competently with a clear right or wrong answer. We are open to solutions exploring different approaches to trust, levels of trust, and delegation.

3. What is a key decision-making attribute?

A: A key decision-making attribute is something that shapes an experts decision making outside of their competence and expertise when confronted with a difficult decision and/or ethical dilemma. Attributes include, as illustrative examples only, values, how much information experts need to decide, and how experts handle uncertainty.

4. How should we measure generalization?

A: This is a competition of ideas. Bring your best ideas that address the requirements.

5. Is there a canonical set of KDMA's?

A: No, the aim of this research is to surface KDMA's that generalize and are predictive of trust.

6. Is the focus on psycho-metrics?

A: A key interest is human characterization through human subjects research. Human characterization aims to identify, observe, quantify, and measure traits that define people. In this SBIR, we are interested in characterizing key decision making attributes that are predictive of trust. This is a competition of ideas. Bring your best ideas that address the requirements.

7. What are the metrics and targets?

A: This is a competition of ideas. Bring your best ideas and methods that address the requirements. Read the Topic and BAA carefully.

8. What are the run-time and compute requirements?

A. This is fundamental research and a competition of ideas. Bring your best ideas and methods that address the requirements.

9. Does DARPA expect the performer to integrate with ITM? What are expectations around integration?

A. DARPA does not have requirements around or an expectation for integration into any parts of ITM, though integration is a possibility. While ITM inspired this SBIR, DARPA encourages proposers to think about alternative domains, solutions, and approaches, rather than focusing on integration with ITM.

10. Are there any disallowed assumptions?

A. DARPA has no disallowed assumptions. This is a competition of ideas. Bring your best ideas and methods that address the requirements.

11. Does DARPA expect human subjects research?

A. DARPA encourages Human Subjects Research considering the nature of the requirement, but it is not required.

12. What are domains of interest?

A: To date, battlefield triage is a key research use case in ITM. We just added cyber defense and are exploring generality more systematically. Preferred domain criteria include high-stakes, limited information, fast-paced, high levels of uncertainty, and where experts often disagree. We are interested in a solution that generalizes. Don't plan around DARPA furnishing a domain, scenarios, or data.

13. What are the populations of interest?

A: Warfighters preferred, with an analogous civilian population acceptable if military recruitment a challenge. This is a competition of ideas. Bring your best ideas and methods that address the requirements.

14. Will you consider an engineering and behavioral science partnership?

A: We will consider an engineering and behavioral science partnership, but the focus of the SBIR is human characterization. This is a competition of ideas. Bring your best ideas that address the requirements.

15. Are you interested in software development?

A: Not necessarily. The SBIR is 6.1 fundamental research. This is a competition of ideas. Bring your best ideas that address the requirements.

16. What do a good proposal and team look like?

A: Please carefully review the evaluation criteria in the BAA.

17. Where do you see the program in 3-5 years?

A: It depends on the proposed solution. Please pay attention to the Commercialization Plan requirements in the BAA.

18. What do you need to qualify for Direct to Phase 2?

A: Proposals are evaluated based on the Evaluation criteria in the BAA. This topic is soliciting Direct to Phase II (DP2) proposals only. Prior research should have been conducted in the past five years. Reports which provide data, clearly present the analysis done and provide evidence of scholarly impact will be strongly preferred, but it is not necessary. Research which has been done on military decision-makers, especially those involved in making difficult decisions like those faced in combat, mass casualty, and triage events, will be strongly preferred. Proposers should already have a defensible framework and methodology to apply, and test based on their previous research. We will evaluate whether proposals are appropriate for Phase II based on the data, analysis, and evidence provided.

19. To qualify for Direct to Phase 2, do we need to have published widely?

A: Proposals are evaluated based on the Evaluation criteria in the BAA. This topic is soliciting Direct to Phase II (DP2) proposals only. Prior research should have been conducted in the past five years. Reports which provide data, clearly present the analysis done and provide evidence of scholarly impact will be strongly preferred, but not necessary. Research which has been done on military decision-makers, especially those involved in making difficult decisions like those faced in combat, mass casualty, and triage events, will be strongly preferred. Proposers should already have a defensible framework and methodology to apply, and test based on their previous research. We will evaluate whether proposals are appropriate for Phase II based on the data, analysis, and evidence provided.

20. Is there a plan for continued work after Phase 2?

A: Please refer to the Topic. There is a Phase 3 Option.

21. Do we need to provide a budget?

A: Yes, but note DARPA tried to streamline the SBIR budget process. It is a different process than standard DARPA BAA requirements. Please carefully review the requirements.

22. What is the deliverable?

A. Provide architecture capable of modeling KDMAs for human decision-makers. The focus is on human characterization. Human characterization aims to identify, observe, quantify, and measure traits that define people. In this SBIR, we are interested in characterizing key decision making attributes that are predictive of trust.

23. Does DARPA have resources to support small businesses who want to work with DARPA?

A: DARPA Connect helps to broaden DARPA's reach and stimulate growth and collaboration with small businesses and education institutions new to the national security space. <https://www.darpa.mil/about/darpaconnect>

28 August

1. Could DARPA clarify the type of leadership you are interested in?

A: The SBIR research is intended to support US warfighters. This is a competition of ideas. Bring your best ideas that address the requirements. DARPA is primarily interested in decisions where experts disagree. This could include leaders or technical experts. Novices are often put into highly complex decisions where experts could disagree when a leader or technical expert is not available on the battlefield. DARPA is interested in diverse and novel approaches to this problem.

2. The topic description specifically identifies a specific type of decision – the decision to delegate difficult decisions (to either an HDM or ADM) – but also mentions decision making more broadly. Could DARPA clarify if you are specifically requesting research on KDMAs around delegating decisions, or KDMAs regarding decision making as a whole / more broadly?

A: DARPA is interested in KDMAs or alternative architectures that are predictive of trust in and delegation of difficult decisions to AI in high stakes domains where experts disagree.

2 SEPT

1. Is there a canonical KDMA set?

A: No. DARPA aims to generate a set that is predictive of trust in high-stakes domains from this study.

2. What are the target cohorts?

A: DARPA does not specify cohorts, but they should all within high-stakes domains and that make difficult decisions where experts disagree.

3. What are required metrics?

A: Please read the BAA criteria carefully. A key SBIR evaluation criteria is whether the performer's architecture is predictive of trust.

4. Are there any run-time and compute parameters?

A: DARPA does not specify specific run-time and computer parameters. Please let us know what you expect. If you need additional compute, DARPA can help provide additional GPUs through the HPC program, if the performer and program qualify.

5. Is there expected integration with ITM? Will there be any government furnished datasets available? What is the demo cadence?

A: DARPA does not require integration with ITM, but this is possible if the use cases and research are compatible. They don't have to be to answer the SBIR requirements.

A: Don't rely on government furnished datasets because DARPA does not specify a domain or research use case that performers must utilize. DARPA doesn't know a priori what you would need.

A: Demo cadence to be determined in coordination with the Government. Please propose your solution and associated schedule.

6. Are there any disallowed assumptions?

A: No disallowed assumptions.

10 September

1. Will an existing neurocognitive biomarker platform be relevant to this Topic?
 - a. In this SBIR, we are interested in characterizing key decision making attributes that are predictive of trust.
 - b. DARPA is interested in KDMAs or alternative architectures that are predictive of trust in and delegation of difficult decisions to AI in high stakes domains where experts disagree.
 - c. This is a competition of ideas. Bring your best ideas that address the requirements.

11 September

1. Could you provide a description for what is meant by "architecture"? Is this topic soliciting a technology stack comprising sensors, processing, storage, and a user interface? Or is it soliciting a concept for collecting data along with modeling approaches?
 - a. The aim is to provide an architecture capable of modeling KDMAs for human decision-makers. Architecture in the context of this Topic is flexible.
 - b. DARPA is interested in KDMAs or alternative architectures that are predictive of trust in and delegation of difficult decisions to AI in high stakes domains where experts disagree.
 - c. DARPA is not necessarily asking for a technology stack. The SBIR is 6.1 fundamental research. This is a competition of ideas. Bring your best ideas that address the requirements.
 - d. We will consider an engineering and behavioral science solution, but the focus of the SBIR is human characterization.

15 September

1. Does Darpa wish to receive proposal to conduct (1) more research on identifying even more human psychological traits that will affect, alter, upgrade or down grade, speed up or slow down - human decision making capabilities -especially under stress? Or will Darpa accept a new research proposal that seeks to identify the building blocks for a complete Decision Making Architecture by which humans (if applied correctly) can significantly increase the probability of making the right decision - in time - more often - with higher probabilities of mission success; even under high stress combat situations ?
 - a. DARPA is interested in KDMAs or alternative architectures that are predictive of trust in and delegation of difficult decisions to AI in high stakes domains where experts disagree.
 - b. This is a competition of ideas. Bring your best ideas that address the requirements.

18 September

1. We see the PPADM page states that recent prior research within the past five years is strongly preferred, but we did not see a statement that such results are strictly required for responsiveness.
 - a. Proposers should already have a defensible framework and methodology to apply, and test based on their previous research. Empirical research involving military decision-makers is strongly preferred, but not required.