



Program Announcement (PA)  
Mathematics of Boosting Agentic Communication  
(MATHBAC)

DEFENSE SCIENCES OFFICE

DARPA-PA-26-05

April 7, 2026

**PROGRAM ANNOUNCEMENT OVERVIEW INFORMATION:**

- **Federal Agency Name:** Defense Advanced Research Projects Agency (DARPA), Defense Sciences Office
- **Funding Opportunity Title:** Mathematics of Boosting Agentic Communication (MATHBAC)
- **Solicitation Type:** Initial Solicitation
- **Funding Opportunity Number:** DARPA-PA-26-05
- **Assistance Listing Number:** Not applicable
- **Dates/Time - All Times are Eastern Time Zone (ET):**
  - Posting Date: April 7, 2026
  - Proposers Day: April 21, 2026
  - Proposal Abstract Due Date: April 30, at 4:00 p.m.
  - Question Submittal Closed: June 4, 2026, at 4:00 p.m.
  - Proposal Due Date: June 16, 2026, at 4:00 p.m.
  - Estimated Period of Performance Start: September 15, 2026
- **Anticipated individual awards:** Multiple awards are anticipated
- **Period of Performance:** The total period of performance is anticipated to be 34 months. The program will be broken into two phases. Phase I will be 16 months. Months 10-16 of Phase I will either involve wrapping up Phase I work, and/or beginning the transition from Phase I to the Phase II of 18 months. Months 17-34 will be the Phase II.
- **Types of instruments that may be awarded:** Other Transaction Agreements for Research
- **Anticipated Funding Available for Award:** Phase I awards will be limited to \$2,000,000 per award
- **NAICS Code:** 541715
- **Agency contact:**

The PA Coordinator for this effort may be reached at:

[MATHBAC@darpa.mil](mailto:MATHBAC@darpa.mil)  
DARPA/DSO  
ATTN: DARPA-PA-26-05  
675 North Randolph Street  
Arlington, VA 22203-2114

## SECTION I: FUNDING OPPORTUNITY DESCRIPTION

The Defense Sciences Office (DSO) at the Defense Advanced Research Projects Agency (DARPA) is soliciting innovative research proposals that advance the foundational mathematics, systems theory, and information theory required to enable/accelerate and understand science-discovery by autonomous agents and agent collectives. Proposed research should investigate innovative approaches that promote revolutionary advances in science, devices, or systems. Specifically excluded is research that primarily results in incremental improvements to the existing state of practice.

Mathematics of Boosting Agentic Communication (MATHBAC) will enable the next generation of Agentic Artificial Intelligence Platforms that systematically evolve to transform the nature of scientific discovery for defense. MATHBAC will develop the science of AI communication, leveraging/deriving mathematical theory, systems theory, and information theory to govern how AI agents collaborate and share information, as well as how they disentangle complex data streams to extract generalizable principles. MATHBAC enables agentic platforms to not only solve complex problems but also to understand the fundamentals behind their own functioning and solutions.

MATHBAC will understand/design/improve the communication protocols of interacting agents, discover new efficient protocols, and develop agentic coordination mechanisms (such as the distilling of shareable, generalizable principles from data). This will enable individual agents and agent collectives to effectively and efficiently address an array of tasks, from science-discovery to problem solving and instruction. It aspires to lower/remove barriers to understanding and help tackle challenges across domains ranging from the physical and mathematical sciences to military strategy, drastically accelerating and potentially reinventing their “ways of doing.” If successful, MATHBAC will accelerate science discovery and create tools to enable the systematic evolution/invention of new science itself.

In Phase I, MATHBAC will accelerate science by developing/deriving the mathematics for (1) understanding/designing agentic communication protocols and for (2) improving communication content; the latter will involve improving understanding by disentangling common from case-specific domain data; extracting generalizable domain principles from volumes of data that exceed human processing limits (i.e., from novel data streams pertaining to a science subdomain); and discovering novel generalizable principles, laws, or correlations. Successes in these steps will be used in Phase II, whose focus is to create tools that enable systematic evolution/invention of new science; this will be achieved by orchestrating agentic self-evolution towards improved science-discovery task-solving, distilling a curriculum for agent finetuning, and/or updating agent-shared core knowledge modules towards solving new tasks. This program will formulate and solve fundamental scientific and mathematical problems underpinning collective agentic intelligence; it is not aimed at building applications or specific products.

### 1. BACKGROUND

The cutting edge of machine learning (ML)-assisted scientific discovery continues to explode, driven by the ability of artificial intelligence (AI) agents to write quality code<sup>1</sup> and execute

---

<sup>1</sup> <https://www.anthropic.com/news/claude-opus-4-5>

workflows using scientific tools.<sup>2</sup> The explosion is fueled by agentic, interacting ML platforms that communicate, adapt, and are capable of rewriting their own code based on performance: Large Learning Models (LLM), Large Reasoning Models (LRM), Small Science Models (SSM), and smaller Foundation/Knowledge models (i.e., transformers).

Recent advances in AI for domain sciences like chemistry, for example, focus on autonomous systems capable of hypothesis generation, experimentation/computation, and refinement without human intervention.<sup>3</sup> Examples include SparksMatter<sup>4</sup> for materials discovery and Virtual Lab<sup>5</sup> for biomedical research. Relevant recent research in LLMs has evolved along four key directions: standalone LLMs fine-tuned for domain tasks,<sup>6</sup> tool-augmented LLMs integrating external resources for autonomous workflows,<sup>7</sup> multi-agent systems with specialized roles collaborating on complex tasks,<sup>8</sup> and multimodal LLMs combining text and visual data for enhanced experimental analysis.<sup>9</sup> The potential advantages of these advances led to a recent partnership between Google DeepMind and the UK government to support prosperity and security in the AI era<sup>10</sup> and the Genesis Mission,<sup>11</sup> a partnership between the Department of Energy and a swath of AI and AI-adjacent companies including Amazon, AMD, Anthropic, FutureHouse, NVIDIA, Google, OpenAI, and Scale AI, to name a few. Simultaneously, several Agent Protocols have been developed to establish communication standards for multi-agent systems, enabling them to collaborate on complex tasks. Protocols such as Agent2Agent (A2A),<sup>12</sup> Agent Communication Protocol (ACP),<sup>13</sup> and Model Context Protocol (MCP)<sup>14</sup> define the syntax, structure, and sequences of messages underpinning agentic cooperation.

Despite the impressive accomplishments of the agentic models and systems mentioned above, these approaches remain fundamentally heuristic, lacking a rigorous mathematical foundation. Current methodologies rely heavily on *ad hoc*, trial-and-error processes, which resemble an "Edisonian" approach to optimization, focusing on end-to-end performance evaluation. While this has led to notable successes, it highlights a critical limitation: the absence of principled frameworks for understanding, predicting, and systematically improving the various stages<sup>15</sup> of behavior of these systems. Without sound mathematical underpinning, the optimization of agentic communication and coordination will remain inefficient, inconsistent, and difficult to generalize across domains even when occasionally successful.

A prescient 2021 sentence in a review article by DARPA performers excelling in AI for chemical synthesis highlights this gap (paraphrased as): “We can find new molecules, and we can find new

---

<sup>2</sup> <https://www.elagente.ca/>

<sup>3</sup> <https://sakana.ai/ai-scientist-first-publication/>

<sup>4</sup> <https://arxiv.org/abs/2508.02956>

<sup>5</sup> <https://www.nature.com/articles/d41586-024-01684-3>

<sup>6</sup> <https://www.nature.com/articles/s42256-025-01066-y>

<sup>7</sup> <https://www.nature.com/articles/s42256-024-00832-8>

<sup>8</sup> <https://www.nature.com/articles/s41467-024-54457-x>

<sup>9</sup> <https://arxiv.org/abs/2408.07246>

<sup>10</sup> <https://deepmind.google/blog/strengthening-our-partnership-with-the-uk-government-to-support-prosperity-and-security-in-the-ai-era/>

<sup>11</sup> <https://genesis.energy.gov/>

<sup>12</sup> <https://github.com/a2aproject/A2A>

<sup>13</sup> <https://agentcommunicationprotocol.dev/introduction/welcome>

<sup>14</sup> <https://modelcontextprotocol.io/docs/getting-started/intro>

<sup>15</sup> <https://arxiv.org/pdf/2508.14111>

reactions, but we cannot find the periodic table, and we cannot find Bronsted-Lowry acid-base theory – because the space of solutions is different than the space of hypotheses.”<sup>16</sup> This underscores a fundamental challenge: while AI excels at navigating solution spaces, it struggles to systematically explore hypothesis spaces, which are essential for generating transformative *and generalizable* scientific insights. A recent survey on autonomous scientific discovery notes: “Effective collaboration enables multi-agent systems to address complex scientific problems that surpass the capabilities of any single agent.”<sup>17</sup> Protocols enabling such a collaborative approach are key to generating novel scientific hypotheses. It is precisely the rate of discovery of these important new hypotheses that MATHBAC aims to systematically accelerate by facilitating AI communication so as to enable breakthroughs in the efficiency of agentic scientific reasoning.

If successful, MATHBAC will fundamentally change the “ways of doing,” whether for scientific discovery or for instruction, by exploiting/synthesizing both established and novel tools/results from mathematics and systems/information theory. This synthesis will provide a new vantage point for observing data and information propagated within and communicated across AI agents – not based on the beginning-to-end performance success but rather based on the nature and content of the interactions. Under current state of the art (SOA), the agents interact with each other, their environment, and tools under *ad hoc* protocols.

One way of leveraging mathematics and systems/information theory is by considering, for example, an ML agent as a mathematical operator, operating on inputs (e.g., token strings, number strings, also possibly multimodal input including images); processing these inputs in its internal latent space (e.g., through the attention heads, routing heads, and various architectural components of a transformer); and producing outputs – new token strings, number strings, or images. This echoes the “input-state-output” view of a system in control theory and opens the door to tools and techniques developed over decades to identify the nature of the operator.

Such an approach empowers the evaluation and scheduling of interactions between heterogeneous operators and codes, enabling the efficient integration of different fidelity models, scales, partial system components, and operational stages. Planning, monitoring, manipulating, and understanding these interactions will allow for more intrinsic comparisons of communication protocols beyond merely evaluating final results. This will require developing and using advanced algorithms for system (e.g., agent) identification and interaction planning. It will also require exploiting analogies with applied mathematics algorithms harnessing the communication between heterogeneous operators/codes, such as (1) preconditioning operators with other operators; (2) operator splitting; (3) domain decomposition; and/or (4) multifidelity, multiscale, or mixed precision computations. These types of numerical algorithms enable powerful tools to quantify communication dynamics and optimize the associated protocols.

Combining two operators—such as a low-fidelity operator that preconditions a high-fidelity one—creates a new composite operator whose spectrum reflects the synergy of the constituent operators. Analyzing this new operator reveals how the strengths of the individual components are combined to achieve faster convergence and improved solutions, demonstrating the transformative potential of applied mathematics in advancing computational efficiency and precision. Understanding and exploiting such identification technology/numerical technology holds the promise of better-

---

<sup>16</sup> <https://onlinelibrary.wiley.com/doi/abs/10.1002/anie.201909987>

<sup>17</sup> <https://arxiv.org/pdf/2508.14111>

informed design and control of agentic communication to accelerate discovery without relying exclusively on trial-and-error experimentation or “black box” reinforcement learning to discover optimal communication protocols (and eventually, optimal agents).

The overall boosting of agentic communication task is facilitated by both the broad use of natural language as a common communication tool and multimodal inputs to a transformer embedded as numbers in its internal latent space, so we can use data analysis techniques to understand the evolution of the problem-solving campaign. We will take advantage of this by studying the role of different communication architectures and protocols *on what is exchanged* between agents, especially as a function of the length of exchanged prompts. In Phase I MATHBAC will start with *fixed* pretrained agents (with no finetuning) because we believe that current knowledge and its intelligent communication, even with fixed agents, can already lead to nontrivial successes; the significantly larger number of abstractions an agent can simultaneously buffer compared to the corresponding physiological limits of an expert human<sup>18</sup> supports this conviction. In Phase II, MATHBAC will explore how *adaptively evolving* agents may specialize or even decide to disaggregate into multiple agents with distinct knowledge domains/skills.

In exploring agentic communications, we clearly see an expansion beyond traditional Shannon information (and so-called “syntactic” communication protocols), where we are interested in the accuracy of the transmission of bit strings to the two higher levels of information – semantic information/communication and communication effectiveness.<sup>19,20</sup> Currently, agents interact with each other, their environment, and tools under *ad hoc* protocols, prescribed by the platform designer. Agents need to coordinate their understanding of the communication signals, of the requests/goals of the communication they undertake with each other (and the overarching problem target), and of their shared knowledge bases—the shared long-term “memory” in memory networks.<sup>21,22</sup> These all relate to semantic communication and to communication effectiveness. Currently, this coordination is all hand-curated.

Beyond evaluating and optimizing interaction protocols, MATHBAC will also focus on features of the *content* of agentic communication: in particular, it will focus on the discovery of “principles” (laws, correlations) from data—the extraction of compact, generalizable “nuggets” that should become part of the common knowledge module (“memory”) of cooperating agents. This is analogous to common core knowledge between cooperating humans and the ways we adjust it (e.g., through targeted supplementary training) towards joint success.

An example: when reaction progression data from different acid-base chemical experiments become embedded in a transformer latent space, we can study using new data science methods what is *common* across all studied reactions and what pertains to each particular experiment separately. The low-dimensional manifold pertaining to commonality across the cases pertains then to the acid-base nature of the reactions; parametrizing this manifold should “rediscover” the fundamental similarity of acid-base reactions and embody its acid-baseness. Data-driven studies

---

<sup>18</sup> Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97. <https://doi.org/10.1037/h0043158>

<sup>19</sup> <https://arxiv.org/pdf/2212.01485>

<sup>20</sup> Shannon, C. E., & Weaver, W. (1949). *The mathematical theory of communication*. University of Illinois Press.

<sup>21</sup> <https://arxiv.org/abs/1410.3916>

<sup>22</sup> <https://arxiv.org/pdf/2405.06394>

can help us reversely articulate such “principles” in natural language,<sup>23</sup> possibly inventing new terms (like click chemistry, high entropy alloys, or even reticular chemistry—to include this year’s Nobel prize) when these terms have not yet been coined.

In the case of generalizable principle discovery from data, we propose to focus on (1) tracking the exchange streams of information between agents, especially as a function of the length of the prompts (the exchanged token streams) and (2) exploring what is “the part of the communication that matters” in a domain context, i.e., what is semantically shared in the communications that arise in processing similar phenomena/queries versus what is case-dependent. The fact that multimodal information (from sentences to experimental data) after tokenization becomes jointly embedded in a transformer’s latent space provides an avenue towards the systematic data-driven study *in the digital latent space of the transformer*.

Improving coordination between different cooperating agents includes possibly developing a completely new *domain language*. Chemists use words and equations but also diagrams; architects and construction engineers use words and blueprints; the military uses doctrine in manuals as well as laminated instruction cards for operating drones. There exist beautiful literature cases illustrating the evolution of a new language (both from Facebook<sup>24</sup> in the case of chatbots working to equitably share a treasure and in the Google Neural Machine Translation Program<sup>25</sup> that came up with its own Interlingua) pointing towards what may well develop into the “Neuralese” of AI2027.<sup>26</sup> The development of useful domain languages is part of semantic communication, which will be a component of the second phase of MATHBAC. Phase I will focus on natural language and floating-point numbers as the communication substrates.

Fifty years ago, DARPA<sup>27</sup> guided the development of the TCP/IP protocol for transmitting Shannon information between computers and adopted it forty years ago in ARPANET. The future of MATHBAC is to move beyond the already ambitious semantic extension of TCP/IP to the case where agents can adapt themselves *and their knowledge* through experience, new data, communication, and shared knowledge/memory. Phase II will focus on agent evolution as part of improving agentic communication. Up to now the evolutionary pressure on AI and agentic platforms comes from evolutionary pressures on the humans involved in their development. MATHBAC wants to move this (not Darwinian, but Lamarckian) pressure onto the agents themselves and their cooperation skills. MATHBAC will systematically explore, understand, and design the best protocols, content, and possibly even language of collaborative AI agent communication.

## 2. PROGRAM DESCRIPTION AND SCOPE

MATHBAC is founded on the observation that, in current practice, autonomous scientific discovery is mostly treated as an end-to-end black-box performance optimization problem. By

---

<sup>23</sup> “the interaction of RL and LLMs is not merely additive, it is multiplicative: RL equips LLMs with the ability to learn from interaction and feedback, while LLMs equip RL with contextual awareness, explainability, accessibility, and code generation capability” - Bertsekas, D., A Course in Reinforcement Learning, 2nd-edition, Athena Scientific, 2025, page 136.

<sup>24</sup> <https://engineering.fb.com/2017/06/14/ml-applications/deal-or-no-deal-training-ai-bots-to-negotiate/>

<sup>25</sup> <https://research.google/pubs/googles-multilingual-neural-machine-translation-system-enabling-zero-shot-translation/>

<sup>26</sup> <https://ai-2027.com/>

<sup>27</sup> <https://www.darpa.mil/about/innovation-timeline>

focusing on distinct components of agentic communication, MATHBAC will make the overall task tractable and efficient by first individually and then collectively optimizing the components to clarify when and how the current process fails. In Phase I, the program focuses on optimizing distinct components of the communication process (communication protocols, communication content) *for a fixed set of agents*. Knowing how to routinely optimize these components in the short time scale should allow the agents in Phase II to gradually, slowly (adiabatically!) self-evolve their capabilities more efficiently than a simultaneous multi-objective Reinforcement Learning optimization.

The MATHBAC program in Phase I requires advances in two technical areas (TAs):

TA1: The Mathematics of Agentic Communication Protocols

TA2: Geometry and Evolution of Shared Knowledge - Generalizable Principles from Data

In Phase I, TA1 and TA2 development will be separate. In Phase II the two TAs will be combined into “Evolution Teams” (defined below) via a Phased Acquisition Approach (details below). As such, proposals must address TA1 **or** TA2 in Phase I; they should also briefly describe how they anticipate capabilities developed by both TA1 and TA2 researchers in Phase I could be orchestrated to address Phase II in a draft Phase II Statement of Work. Proposals that do not describe how they could address Phase II by leveraging anticipated Phase I capabilities may be considered non-conforming and removed from consideration. Proposers may propose to both TAs, but they must be submitted as separate proposals.

MATHBAC will also fund a Federally Funded Research and Development Center (FFRDC), University-Affiliated Research Center (UARC), or Government entity as the Independent Verification & Validation (IV&V) team (not solicited under this Program Announcement).

As noted above, MATHBAC is dedicated to advancing the foundational mathematics, systems/control theory, and information theory required to enable/accelerate (and understand) science-discovery by autonomous agents and agent collectives. In order to make this effort tractable, the MATHBAC program will restrict its focus to science domains since (1) science domains have a vast, structured curated knowledge corpora; (2) groundbreaking discoveries often require interdisciplinary synthesis and communication; (3) discoveries require hypothesis generation via analogical reasoning; and (4) science domains provide a clear method for validation and measuring success.

TA1 and TA2 proposers are expected to identify *their selected* science subdomain(s) and two families of scientific tasks in their subdomain(s) they plan to address—for example, for a chemistry subdomain, two such task families might be (1) Structure Elucidation (e.g., predict spectra from structures, identify structures consistent with spectra, deconvolute, and analyze complex mixtures); and (2) Reaction Development (e.g., predict new products of known reactions); TA2 performers must identify at least one well-established scientific principle, law, or correlation *to be rediscovered* in their domain(s) (e.g., rediscovering the Karplus equation or coupling rules in Nuclear Magnetic Resonance spectra); and their curated science corpora. TA2 proposers must also convincingly demonstrate that these corpora do not explicitly contain nor purposefully seed the targeted principle(s), law(s), or correlation(s). TA1 and TA2 proposers must (i) specify the best performance they have achieved with their current methods, as a baseline and (ii) identify potential additional metrics to be adopted by the program. For each potential additional metric, the proposer must provide the following:

- A description of what it measures
- The method for its calculation
- A description of how it measures progress against program goals
- An explanation of how it supplements current program metrics
- An argument for why the metric should be adopted (e.g., what insights it will provide)
- Literature sources providing metric provenance and application (if available)

Lastly, proposers must identify and provide rationale about other scientific domain(s)/subdomain(s) to which they believe their current approach can generalize.

In TA1, Phase I approaches will be compared to a SOA model, e.g., a Mixture of Experts model.

In TA2, Phase I will include two evaluations. In the first evaluation, proposers will rediscover at least one known generalizable principle, law, or correlation with agents trained on domain literature *chronologically prior* to the original discovery. In the second evaluation, agents will be required to discover at least one (but ideally, more than one ) previously unknown principle, law, or correlation from novel data streams (a DARPA-hard goal, for instance, would start with the data-driven Mendeleev-level rediscovery of the periodic table for atoms and proceed to a “multi-dimensional analog” of a periodic table for *molecules*).

Toward the end of Phase I, DARPA will host a PI meeting for performers to demonstrate their current capabilities and share their test results to facilitate potential Phase II collaborations between TA1 and TA2 performers. Specifically, performers are highly encouraged to collaborate in what will be termed an “*Evolution Team*” for their Phase II proposals. An Evolution Team is defined as:

*A single performer or a team of performers that addresses both TA1 and TA2 in Phase II; no performer can be part of more than a single Evolution Team.*

In Phase II, Evolution teams’ approaches will be compared to (and expected to significantly outperform) a collective of domain experts with hand-crafted protocols and access to SOA models.

While performers will have the option not to collaborate with other performers in Phase II, collaboration is strongly encouraged. It should be noted that the decision to work independently or collaborate with others will not impact Phase II determinations.

DARPA will use a Phased-Acquisition Approach for the MATHBAC program. Proposers should provide a fully-definitized proposal for Phase I and a draft Statement of Work (SOW) and Rough Order of Magnitude (ROM) cost for Phase II. Proposed price for Phase I must not exceed \$2,000,000. DARPA anticipates releasing proposal instructions for Phase II prior to month 10 of Phase I; Phase II awards are anticipated by month 16. Phase II proposals will require the Evolution team to have one performer acting as the prime performer and the rest as subperformers. It is incumbent upon the members of the Evolution Teams to choose and negotiate their own prime-sub relationships. Phase II awards are subject to availability of funds, assessed Phase I performance, and the evaluation of Phase II proposals. Evaluation criteria for Phase II proposals will be in the Phase II Proposal Instructions and will be consistent with the criteria in this solicitation. DARPA may invite one or more Evolution Teams to continue to Phase II of the program. Phase II may be negotiated as a modification to an existing Phase I award or may result in a new award.

## Phase II Agentic Evolution towards New Science

This section describing MATHBAC Phase II is provided for planning purposes only. Proposers should use this information to justify how their Phase I effort will support Phase II and how they plan to integrate into an “Evolution Team” for Phase II. Initial proposals should describe potential Phase II approaches, with supporting technical details, assuming all of the expected outcomes from Phase I TA1/TA2 are successful. During Phase I, DARPA will issue proposal requests to Phase I performers detailing the content and evaluation criteria DARPA will use to determine Phase II awardees.

The expertise developed in both TA1 and TA2 of Phase I will be on *fixed* agents, i.e., agents that do not change/adapt during (or after) a campaign. The TA2 Orchestrator tool is the precursor of what will be the main goal of Phase II: allowing the agents to determine *what, when, how, and where* to evolve the agents. Currently, evolutionary pressure for agent platform design is mainly applied on the platform creators (individuals or companies). Agent self-evolution (rewriting the agent code based on performance, e.g., fine-tuning, updating the agent common knowledge modules, updating the agent communication protocols, disaggregating existing agents to subspecialized ones, and/or creating new agents and placing them in the protocol) may occur either in a punctuated (inter-campaign) or continuous (intra-campaign) manner. Performers will implement agentic self-evolution via optimization (reward based, possibly using reinforcement learning) that may be general or focus on a particular science domain, driven by agent performance (time to task completion, communication token bandwidth, energy requirements, etc.). The proposers must identify the rewards they anticipate to leverage in Phase II and why.

Phase II focuses on techniques that allow the agentic interactions—and even the agents themselves—to adapt and evolve, possibly differentiating or even disaggregating themselves as the collaboration campaigns progress.

Proposers are expected to identify the mathematics/systems and information theory algorithms and tools that will be leveraged and advanced to allow the protocols and the agents/agent collectives themselves to adapt and evolve. Proposals must describe the (1) methods that will be used to (a) identify and understand evolutionary pressures from monitoring evolving systems, (b) create evolutionary pressure mechanisms for directed evolution, (c) measure and monitor agents’ knowledge (and latent space if accessible), and (d) evolve and measure efficiency of the agents’ selected language; (2) the experiments to be used to provide baseline improvement comparison; (3) how the team will implement an oracle (what information is needed, how much, what type, etc.) to confirm successful reception/understanding of a semantic communication by the receiving agent; (4) how TA1 and TA2 will be extended to address Phase II; and (5) how the proposed approach will achieve the Phase II notional outcomes, see Table 1: Program Metrics. It is expected that proposals will include a grounded example in their selected subdomain, demonstrating the expected capabilities of their approach.

It is assumed that in Phase II performers will need to develop new methods that will be used to

- implement shared latent space representation across agents (or identify/infer it, when not accessible) for a given task,
- understand and engineer the form (language) of the content for particular pairs in a collective, and
- design agentic pressures that force agent collectives to converge to optimal representations.

### 3. PROGRAM STRUCTURE

The MATHBAC program is a 34-month, two phase effort with two TAs that will run concurrently in Phase I and become combined in Phase II. Phase I is 16 months (Base), and Phase II is planned to be 18 months. Proposals must address a single TA in Phase I and describe how the results of both TA1 and TA2 combined will address the agentic evolution towards new science in Phase II.

Any changes to the Phase II focus will be delineated in the Phase II Proposal Instructions.

- TA1 –The Mathematics of Agentic Communication Protocols
- TA2 – Geometry and Evolution of Shared Knowledge—Generalizable Principles from Data

Phase I will hold agents static and not allow them to evolve, and the two TAs may be completed separately. In Phase II the agents are planned to evolve, and the TAs are expected to be unified within an Evolution team.

In Phase II performers are expected to work on both TAs as part of a unified Evolution Team that need not be defined until performers submit proposals for Phase II. These Evolution Teams should possess expertise in a range of topics including mathematics, information theory, proposer selected science subdomain, and science-discovery agents. The performers should also have access to agentic platforms allowing access to latent spaces and communication logs.

#### Phase I Accelerating Agentic Communication towards Scientific Discovery

##### Technical Area 1: The Mathematics of Agentic Communication Protocols

- Goal: To establish a quantitative mathematical framework for designing, optimizing, and establishing guarantees for the performance of communication protocols within a collective of *fixed* pretrained AI agents, moving beyond current *ad-hoc* methods.

This TA focuses on methods to understand/design/improve the communication protocols of interacting agents in agent platforms. The protocols will include how many and what type of agents are needed on the team and which agents “talk to” which agents at each stage of the science-discovery campaign.

Proposers are expected to identify the mathematics, systems theory, and information theory methods that will be leveraged and advanced to understand/design and improve the communication protocols of interacting agents. These may include input-output attempts to create for each agent approximate “internal models” of the agents with which it interacts (a sort of “theory of mind” of agents). Proposers must state:

- (1) the methods that will be used to
  - (a) evaluate and schedule interactions between heterogeneous agents
  - (b) understand and enable more efficient (e.g., lower communication prompt length/faster completion) agent interactions
    - (b1) without access to any latent spaces
    - (b2) with access to transformer latent spaces
    - (b3) with mixed access agents (with/without accessible latent spaces)

(c) characterize and quantify the progress of communication dynamics throughout the campaign

(d) optimize communication protocols.

(2) the experiments and baseline protocol(s) (which can be handcrafted) to be used for comparison and to quantify improvement.

(3) which “small science” foundation models will be used; these must be science models whose latent spaces can be accessed. Proposers must describe how they will access these latent spaces and how they will use the information they access.

(4) the method in which data from agentic communication stream and agent latent spaces during a campaign will be captured, processed to improve communication protocols, and how they will be shared with IV&V.

(5) how the proposed approach will achieve and/or exceed the TA1 metrics (see Table 1: Program Metrics for a description).

Proposals must include a clearly documented example in their selected science subdomain, demonstrating the expected capabilities of their approach as well as two families of scientific tasks in their subdomain(s) they plan to address. Additionally, as noted above, proposers must identify and provide rationale about other scientific domain(s) / subdomain(s) to which they believe their current approach can generalize.

TA1 posits that by systematically designing agent communication protocols we can markedly improve speedup and percentage success in agentic task completion.

- Proposals must state how they will address the following anticipated (and possibly additional expected) mathematical challenges:
  - Agents as Operators/Systems: Develop mathematical techniques, such as system identification and input-output based reduced-order modeling, to locally approximate/represent individual AI agents as (1) mathematical operators operating on token strings (or more generally, multimodal input streams) and/or (2) systems with inputs, internal states (in the agent latent space), and outputs.
  - Communication Protocols in Operator Networks/System of Systems: Adapt and extend the fidelity (e.g., via hierarchical spectral augmentation of the rank) of the identified models of the operators/systems as well as graph theory tools to analyze, evaluate the performance of, and design efficient complex networked interactions (including internal cascades and recurrences) toward the desired goals.
  - Performance Prediction: Formulate discrete mathematical models to estimate/analyze how different successive combinations/permutations of agent (operator, system) actions systematically affect the collective's efficiency, stability, and convergence over broad classes of related domain tasks.
  - Protocol Optimization: Formulate and solve discrete/continuous optimization problems (with defined cost functions (e.g., token usage, solution time) and performance metrics) to minimize end-to-end overall campaign cost while maximizing task success frequency.

- Construction of an Oracle: What information is needed, how much, what type, etc. to confirm semantic communication between agents was understood.
- Deliverables:
  - A Quantitative Mathematical Framework: A theory to describe, analyze, and evaluate the design and implementation of multi-agent communication.
  - A Computational Design Tool: A software suite implementing the above framework in specific science-discovery subdomains to allow researchers (and even, in Phase II, **the agents** themselves) to implement and validate optimal protocols without *ad hoc* large-scale experimentation.
  - A Catalog of Protocol Design Principles and of Domain-Specific Optimal Protocols: A set of demonstrably superior communication protocols for specific tasks (the program focuses on scientific discovery), verifiably proven through the mathematical framework. Additionally, a distillation of the interpretable guiding principles leading to optimal protocol design is highly desirable.

**Technical Area 2: Geometry and Evolution of Shared Knowledge—Generalizable Principles from Data.**

- Goal: To create a mathematical foundation based on disentanglement for discovering and incrementally sharing new, abstract, generalizable principles, laws, and correlations at meaningful spatiotemporal scales (in contrast to principles discovered through wisdom/serendipity).

This TA focuses on methods to monitor and usefully analyze multiple data streams (e.g., experimental data streams or streams recording *the content* of agentic communication), quantify its efficiency, discover new generalizable principles and abstractions, and use inductive and/or abductive reasoning in science-discovery.

Proposers are expected to identify the mathematics, information theory, and data science methods that they intend to leverage and advance towards discovering new generalizable principles by monitoring and analyzing multiple data streams. Proposals must describe:

- (1) Methods the proposer will use to
  - (a) enable combination/comparisons of data streams across modalities and time scales.
  - (b) identify common and uncommon/particular features of the data streams, e.g., experimental data streams or streams arising from agentic interactions.
  - (c) extract generalizable principles, laws, or correlations; explicitly describe cost function/rewards used to drive principle/law/correlation extraction.
  - (d) enable interpretability of the extracted principles, laws, or correlations.
  - (e) generate and validate a minimum curriculum (content) for succeeding in a given task/science domain, if possible.
- (2) How the proposed approach will achieve the TA2 metrics (see Table 1: Program Metrics for a description).

- (3) The computations/experiments to validate the generalizability of the discovered principle(s), law(s), or correlation(s).

Proposals must include a clearly documented example in their selected domain, demonstrating the expected capabilities of their approach.

TA2 posits that agentic interactions improve reasoning and thinking capabilities<sup>28</sup> that should make common concepts “mineable” across a collective of data streams/agentic communication streams/agentic latent spaces. To test this, the TA2 metric will involve “understanding” that measures the ability to extract, distill, and share generalizable principles (laws or correlations).

- Proposals must state how they will address the following anticipated (and possibly additional expected) mathematical challenges:
  - Develop algorithms for learning to identify consistent, low-dimensional structures shared across multiple, distinct latent space point clouds pertaining to the same knowledge subdomain, separating (disentangling) them from case-dependent or random features. One way of approaching this is through joint manifold learning techniques. The data come from (1) prior domain knowledge, (2) data collected by the collective’s tools during a campaign, and/or (3) data from observation of the agentic communication streams and agent latent spaces during the campaign.
  - Geometric and Topological Characterization of these common data manifolds, using tools from computational differential geometry/harmonic analysis as well as Topological Data Analysis (TDA) to compute meaningful distances in the “common” submanifolds and in their complements.
  - Interpretability as an Inverse Problem: Develop mathematical methods to solve the inverse problem of projecting the discovered “common” geometric object (in transformer latent space) back into an interpretable (e.g., natural language) space.
- Deliverables:
  - A “Generalizable Principle Discovery Tool from Observations”: An algorithm suite capable of extracting the shared principles, laws, and correlations that connect a family of problems
  - A “Principle Speaker”: A tool that meaningfully projects the geometrically discovered principles into human-interpretable expressions
  - An “Orchestrator”: A tool that selects either (1) the fine-tuning curriculum of agents so that they embody the new discoveries or (2) the update of the agent shared core memory modules in view of these discoveries.

As noted above, MATHBAC will have an IV&V team who will prescribe and administer baseline comparisons for team performance. The IV&V team will by month 7 of Phase I (1) design bespoke challenge problems for each performer team in their scientific subdomain and (2) create or identify a current SOA scientific discovery agentic platform as the basis of comparison. This SOA platform will provide the basis of a “living metric”; IV&V will assess performer progress against the best and most recent version of the platform available at the time of comparison. The bespoke challenge

---

<sup>28</sup> <https://arxiv.org/pdf/2601.10825>

problems will be released at month 7 of Phase I and can be completed on a rolling basis. The IV&V team will adopt a rubric, in consultation with DARPA, to assess the performance for TA1 and TA2. The final rubrics will be shared with the performers by Month 4 of Phase 1. TA1 and TA2 performers are expected to report on their progress using the IV&V rubric.

**The TA1 rubric shall include, at a minimum,**

**(1) % success rate improvement for agentic campaigns:**

Performers will need to demonstrate that their multi-agent communication design and optimization tool(s) will improve the end-to-end success rate for a campaign, which starts with a human posing an *unsolved* scientific task/problem and ends when the agentic platform returns one or more *confirmable (computationally or experimentally)* solutions. This will be compared to the TA1 baseline protocols until the IV&V living metric is created/identified.

**(2) speedup of task completion(s):**

Performers will need to demonstrate that their multi-agent communication design and optimization tool(s) will speed up the completion of their scientific subdomain tasks (e.g., characterization of the natural transition orbitals of 4-(dimethylamino) benzonitrile). This will be compared to the TA1 baseline protocols until the IV&V living metric is created/identified.

**(3) adaptability/generalizability to related scientific-discovery tasks:**

Performers will need to demonstrate that their multi-agent communication design and optimization tool(s) will adapt/generalize to nearby tasks. Until the IV&V living metric and bespoke challenge problems are available, performers will demonstrate this capability using tasks from their two proposed families.

**The TA2 rubric shall include, at a minimum,**

**(1) # of data streams**

Performers will need to demonstrate that their data fusion and analysis platform can successfully integrate and process multiple concurrent data streams. These streams could include diverse data types such as sensor outputs, experimental results, and computational models.

**(2) amount of data, # of samples, and richness**

Performers will need to demonstrate that their data processing and analysis tool(s) can effectively handle and extract meaningful information from large and complex datasets. This includes processing a high volume of data (amount), a large number of individual data points (samples), and data with high dimensionality or multiple modalities (richness).


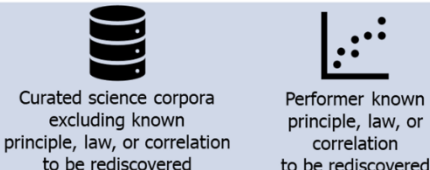

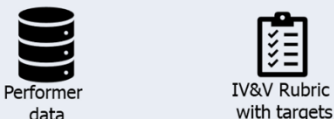
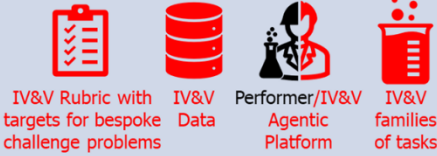
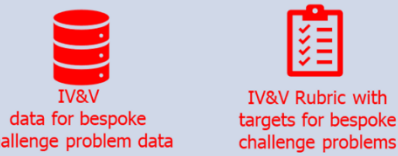
**(3) time to disentangle**


Performers will need to demonstrate that their causal inference and data disentanglement tool(s) will reduce the time required to separate and identify the underlying common factors or signals within a complex, multimodal dataset.

**(4) Interpretability of results**

Performers will need to demonstrate that their data analysis and modeling platform produces outputs that are interpretable and explainable to a human scientific expert. The goal is to ensure that the disentangled factors, causal relationships, or generated models are not merely predictive but also provide clear insights into the underlying scientific phenomena.

IV&V will further develop the rubrics to appropriately challenge each team according to the chosen subdomain.

Months	TA1 Phase 1 Baseline	TA2 Phase 1 Baseline
0 – 4 months	 <p>Performer Data Performer Agentic Platform Performer families of tasks</p>	 <p>Curated science corpora excluding known principle, law, or correlation to be rediscovered Performer known principle, law, or correlation to be rediscovered</p>
5 – 7 months	 <p>IV&amp;V Rubric with targets for performer's families of tasks</p>	 <p>Performer data IV&amp;V Rubric with targets</p>
8 – end of Phase I	 <p>IV&amp;V Rubric with targets for bespoke challenge problems IV&amp;V Data Performer/IV&amp;V Agentic Platform IV&amp;V families of tasks</p>	 <p>IV&amp;V data for bespoke challenge problem data IV&amp;V Rubric with targets for bespoke challenge problems</p>



Evaluated by  
IV&V & DARPA

*Phase I baselines: All results will be evaluated by IV&V and DARPA. (On the left) In months 0-4, TA1 performers will use their own data, agentic platform, and families of tasks as their baseline; in months 5-7 TA1 performers will have to achieve the IV&V rubric targets using their own data, agentic platform, and families of tasks; in months 8 – the end of Phase I, performers will compare to the living metric, that is IV&V data, IV&V agentic platform, IV&V families of tasks, and the IV&V rubric targets for the bespoke challenge problem. (On the right) In months 0-4, TA2 teams will use their own selected correlation, principle, or law, and their own curated science corpora, that excludes their target known correlation, principle, or law, to be rediscovered; in months 5-7, TA2 performers will discover a new correlation, principle, or law, to be rediscovered; in months 8 – the end of Phase I, performers will use the IV&V data for their bespoke challenge problem and achieve the IV&V rubric targets.*

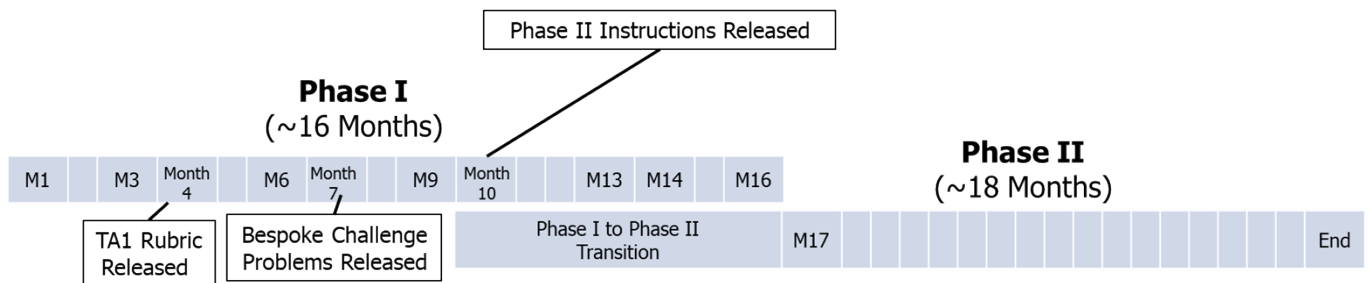
In addition to meeting and or exceeding the rubric, performers must explain why their approach worked (or, equally importantly, why it did not work) and the IV&V team must validate, through implementation in collaboration with the team, the proposed methodology. In the Table below, an example of a family of tasks could be the characterization of the natural transition orbitals of various organic compounds.

In Phase II, we anticipate that the rubric will have to be updated based on program results and progress published/shared by the broader community. Instead of metrics, we provide here a notional idea of expected outcomes for Phase II. An updated rubric will be included in the Phase II Proposal Instructions.

Phase I Metrics	Phase II Notional Outcomes
<b>Accelerating Agentic Communication towards Scientific Discovery</b>	<b>Agentic Evolution towards New Science</b>
For TA1: Meet or exceed IV&V rubric and mathematically explain/rationalize why approach worked (or did not)	For Combined TA1 and TA2: Distill curriculum for a family of new tasks  Demonstrate effectiveness of new agent collective, capable of solving tasks from unknown (to the agent) problems
For TA 2: Demonstrate ability to identify common and uncommon features across novel data streams pertaining to a science subdomain  Discover at least one new generalizable principle, law or unknown surprising correlation  Meet or exceed IV&V rubric and mathematically explain/rationalize why approach worked (or did not)	

Table 1: Program Metrics

**Program Schedule**



**Milestones and Deliverables**

Proposers must address the following fixed payable milestones in their proposals. Proposers must complete the “Schedule of Milestones and Payments” tab of the DARPA Streamlined Cost Buildup Workbook Excel Attachment F provided with this solicitation. If selected for award negotiation, the fixed payable milestones provided will be directly incorporated as an Attachment to the OT agreement (“Schedule of Milestones and Payments”). Proposers must use the Task Description Document (TDD) template, provided as Attachment H to this solicitation. The final TDD will be incorporated as Attachment to the OT agreement.

MATHBAC is adopting a living metric, as the goal of the program is to develop and understand the mathematics/systems theory/information theory basis for improving the agentic protocols and

the discovery of principles/laws/correlations. MATHBAC is not interested in performers “chasing” metrics; it is expected that every milestone, deliverable, and improvement achieved on the program will be accompanied by reporting/describing the mathematical and/or systems/information theory tools underpinning the success of the effort—or the reasons of failure. This mathematical justification of success (or explanation of failure) will be evaluated by the IV&V team and DARPA.

Phase I fixed milestones and deliverables, which are broken out by TA1 and TA2, must include, at a minimum, the following:

TA1: The Mathematics of Agentic Communication Protocols

- Month 1:
  - Attend Phase I project kick-off meeting to include presentation describing selected mathematical/information/systems theories to be leveraged, the rationale for their use, and a complete description of the chosen foundation models, including how their latent spaces will be accessed and leveraged. Presentation should include description of their baseline protocol(s) and a quantitative analysis of their current work evaluated using the initial TA1 rubric on their chosen scientific discovery tasks. All supporting positions identified in the proposal are assigned to personnel and names are provided to the Government.
- Month 3:
  - Report on initial successes and/or failures in effort for target scientific discovery task and initial description of quantitative mathematical framework to describe, analyze, and evaluate the design and implementation of multi-agent communication. Confirm all proposed personnel are working on the project at the planned level of effort.
- Month 6:
  - Attend Phase I PI meeting to include thorough description of quantitative mathematical framework and progress towards Phase I metrics. Demonstrate current capabilities and test results in collaboration with IV&V. Demonstrate ability to capture and process data from agentic communication streams and agent latent spaces during a campaign; share with IV&V.
- Month 9:
  - Report describing initial software suite implementing the quantitative mathematical framework in specific science-discovery subdomains to allow researchers to design, simulate, and validate optimal protocols without ad hoc large-scale experimentation. Report describing initial software suite implementing the quantitative mathematical framework in specific science-discovery subdomains to allow researchers to design, simulate, and validate optimal protocols without ad hoc large-scale experimentation and progress against IV&V provided bespoke challenge problem(s). Software should demonstrate successful creation of reduced-order models for selected agents. The models should accurately predict agent output based on a given input string with a specified degree of fidelity.

- Month 13:
  - Attend Phase I PI meeting to include thorough description of software suite, progress towards Phase I metrics, progress against IV&V provided bespoke challenge problem(s), and initial description of a set of demonstrably superior communication protocols for specific tasks, verifiably proven through the mathematical framework. The report must include a side-by-side comparison against the baseline, verifying success in the IV&V rubric explained through the mathematical framework.
- Month 14:
  - Demonstrate current capabilities of the computational design tool and test results in collaboration with IV&V.
- Month 16
  - Delivery to DARPA of final report that includes the mathematical theory, information theory, and systems theory developments and implementations used in their Phase I accomplishments; and plans for Phase II

## TA1 Deliverables:

- Month 12:
  - Quantitative Mathematical Framework: A theory to describe, analyze, and evaluate the design and implementation of multi-agent communication.
- Month 14:
  - A Computational Design Tool: A software suite implementing the above framework in specific science-discovery subdomains to allow researchers (and even, in Phase II, **the agents** themselves) to implement and validate optimal protocols without *ad hoc* large-scale experimentation.
- Month 16:
  - A Catalog of Protocol Design Principles and of Domain-Specific Optimal Protocols: A set of demonstrably superior communication protocols for specific tasks, (the program focuses on scientific discovery), verifiably proven through the mathematical framework. Additionally, the distillation of the interpretable guiding principles leading to optimal protocol design is highly desirable.

## TA2: Geometry and Evolution of Shared Knowledge—Generalizable Principles from Data.

- Month 1:
  - Attend Phase I project kick-off meeting to include presentation describing their selected science subdomain(s); at least one well-established scientific principle, law, or correlation to be rediscovered in their subdomain(s); their curated science corpora; overview of TA 2 approach; ) and a quantitative analysis of their current work evaluated using the initial TA2 rubric. All supporting positions identified in the proposal are assigned to personnel and names are provided to the Government.
- Month 3:

- Report on initial successes and/or failures in effort for target domains and initial description algorithm suite capable of extracting the shared principles, laws, and correlations that connect a family of problems. Confirm all proposed personnel are working on the project at the planned level of effort.
- Month 6:
  - Attend Phase I PI meeting to include thorough description of algorithm suite and progress towards Phase I metrics. Demonstrate current capabilities and test results in collaboration with IV&V.
- Month 9:
  - Report describing initial tool that meaningfully projects the geometrically discovered principles into human-interpretable expressions and progress against IV&V provided bespoke challenge problem(s).
- Month 13:
  - Attend Phase I PI meeting to include thorough description of principle projection tool, progress towards Phase I metrics, progress against IV&V provided bespoke challenge problem(s), and progress on disentangling data from agentic communication streams and agent latent spaces during a campaign.
- Month 14:
  - Demonstrate current capabilities of the principle discovery tool and test results in collaboration with IV&V. Initial description of orchestration tool that selects either (1) the fine-tuning curriculum of agents so that they embody the new discoveries or (2) the update of the agent shared core memory modules in view of these discoveries.
- Month 16
  - Delivery to DARPA of final report that includes the mathematical theory, information theory, and data science developments and implementations used in their Phase I accomplishments; and plans for Phase II.

## TA2 Deliverables:

- Month 11:
  - A “Generalizable Principle Discovery Tool from Observations”: An algorithm suite capable of extracting the shared principles, laws, and correlations that connect a family of problems.
- Month 12:
  - A “Principle Speaker”: A tool that meaningfully projects the geometrically discovered principles into human-interpretable expressions.
- Month 16:
  - An “Orchestrator”: A tool that selects either (1) the fine-tuning curriculum of agents so that they embody the new discoveries or (2) the update of the agent shared core memory modules in view of these discoveries.

All proposals must include the following meetings and travel in the proposed schedule and costs for Phase I:

- For budgeting purposes, plan for 3 three-day meetings over the course of 16 months: two meetings in the Washington, D.C. area and one meeting in the San Francisco, CA area.
- Regular teleconference meetings will be scheduled with the Government team for progress reporting as well as problem identification and mitigation. Proposers should anticipate at least one site visit per phase by the DARPA Program Manager during which they will have the opportunity to demonstrate progress towards agreed-upon milestones.

Performers will be expected to provide at a minimum the following deliverables:

- Comprehensive quarterly technical reports due within ten days of the end of the given quarter, describing progress made on the specific milestones as laid out in the SOW.
- A phase completion report submitted within 30 days of the end of each phase, summarizing the research done.
- Other negotiated deliverables specific to the objectives of the individual efforts. These may include registered reports, experimental protocols, corpora, demonstrations, prompts, publications, intermediate and final versions of software libraries, small science models, code, and APIs, including documentation and user manuals.
- Working system capable of being evaluated by the Test & Evaluation team in Phase I.
- Performers will brief DARPA on evaluation results and also submit a written summary of results due within two weeks of completion of each evaluation.

Performers should identify (1) how much compute they expect to use in Phase I for the proposed project broken out by month; (2) what resources they have available at their organization; and (3) what additional resources they will request. Resources include compute, compute time, and models.

It is encouraged that all noncommercial software (including source code), software documentation, and technical data generated by the program be provided as deliverables to the Government with a minimum of Government Purpose Rights (GPR) as lesser rights may adversely impact the lifecycle costs of affected items, components, or processes. Any commercial software that is a necessary component of a performer's approach must be identified and include a description of the license rights to be provided to the Government and how the license rights provided will enable use in any follow-on Government research programs.

## SECTION II: EVALUATION CRITERIA

Proposals will be evaluated using the following criteria listed in **descending order of importance**. Overall Scientific and Technical Merit; Potential Contribution and Relevance to the DARPA Mission; Price; and Proposer's Capabilities or Related Experience.

- **Overall Scientific and Technical Merit:** The proposed technical approach is innovative, feasible, achievable, and complete. Detailed technical rationale is provided delineating why the proposed approach can achieve the program goals and metrics. Task descriptions and associated technical elements provided are complete and logically sequenced with all proposed deliverables clearly defined so the final outcome of the award's work achieves the program's goal. The proposal identifies major technical risks and planned mitigation efforts that are clearly defined and feasible.
- **Potential Contribution and Relevance to the DARPA Mission:** The potential contributions of the proposed effort bolster the national security technology base and support DARPA's mission to make pivotal early technology investments that create or prevent technological surprise.
- **Price:** The proposed price represents a practical understanding of the effort and is based on meeting the key technical challenges in an efficient manner. A price evaluation will be conducted to ensure that the proposed price is reasonable. Satisfactory price competition is anticipated to establish price reasonableness. If needed, the Government may use various price evaluation techniques and methodologies to ensure the proposed price is reasonable. A price should be sufficiently detailed to demonstrate its reasonableness. The burden of demonstrating price reasonableness rests with the proposer. An assessment that the proposal price is not reasonable may result in the proposal not being selected for award.
- **Proposer's Capabilities or Related Experience:** The proposer provides evidence that the team and organization is capable of realizing the proposed research effort. In particular, the proposer team demonstrates capabilities and or related experience with (1) agentic AI; (2) small science models; and (3) the proposed mathematics/information/systems/data science theory and algorithms to be deployed. Similar efforts completed or ongoing by the proposer in this area are fully described including identification of other Government sponsors.

Unless otherwise specified in this announcement, for additional information on how DARPA reviews and evaluates proposals through the Scientific Review Process, please visit: [Proposer Instructions: General Terms and Conditions](#).

## SECTION III: SUBMISSION GUIDELINES

- **MODEL OTHER TRANSACTION FOR RESEARCH:** This solicitation allows for Other Transaction for Research award type only. Proposers must complete and submit the Model Research Other Transaction (OT) provided as **Attachment G**. DARPA has provided the model OT to expediate the negotiation and award process. The model Research OT for is representative of the terms and conditions that DARPA intends to award for the MATCHBAC Program. Proposers must edit the blue text/complete/redline the text within the model OT to the best of their knowledge to reflect their negotiation position.

### Proposers must review the following links:

- **Proposer Instructions: General Terms and Conditions:** <https://www.darpa.mil/about/offices/contracts-management/proposer-general-terms>
- **Other Transaction Agreements:** <https://www.darpa.mil/about/offices/contracts-management/proposer-transactions>
- **ABSTRACT:** This solicitation contains an abstract phase. Abstracts are strongly encouraged but not required. DARPA will review the abstracts and provide feedback as to whether a full proposal submission is encouraged. Regardless of DARPA’s response to an abstract, proposers may submit a full proposal. DARPA will review all conforming full proposals using the published evaluation criteria and without regard to any comments resulting from the review of an abstract. Proposers should note that a favorable response to an abstract is not a guarantee that a proposal based on the abstract will ultimately be selected for award negotiation. It is DARPA policy to attempt to reply to abstracts within thirty calendar days. These official notifications will be sent via email to the Technical POC and/or Administrative POC identified on the abstract coversheet. Abstracts are due no later than the date and time stated in the Overview section. Additional instructions for abstract submission are contained within **Attachments A and B**.
- **PROPOSALS:** Full proposals are due no later than the date and time stated in the Overview section. **Attachments C, D, E, F, G, H, I, and X** contain specific instructions and templates and constitute a full proposal submission. Please visit [Proposer Instructions: General Terms and Conditions](#) for specific information regarding submission methods through the Broad Agency Announcement Tool (BAAT). for specific information regarding submission methods through the Broad Agency Announcement Tool (BAAT).
- **MILESTONES:** Fixed payable milestones are payments based on successful completion of the milestone accomplishments agreed to within the milestone plan. Proposals may suggest modifications or additions to the Schedule of Milestones and Payments; **please note that any suggested edits may not be accepted by DARPA**. A Schedule of Milestones and Payments is included as a tab in **Attachment F**. Payments are triggered by completed performance of observable technical events (milestones).
- **Representations and Certifications:** All proposers are required to submit the DARPA-specific representations and certifications for Research OT awards (**Attachment I**) to be eligible to receive an OT award. See <https://www.darpa.mil/research/opportunities/rebs-certs> for further information on required representations and certifications for Research OT awards.

- All technical, contractual, and administrative questions regarding this notice must be emailed to [MATHBAC@darpa.mil](mailto:MATHBAC@darpa.mil). Emails sent directly to the Program Manager, or any other address, may result in a delayed or no response. DARPA will attempt to answer all questions in a timely manner and post a “Frequently Asked Questions” document on the DARPA website. This will be updated on an ongoing basis until the closing date listed in the Overview section.
- **PS Attachments:**
  - **(required if submitting an abstract) Attachment A:** Abstract Summary Slide Template
  - **(required if submitting an abstract) Attachment B:** Abstract Instructions and Template
  - **(required) Attachment C:** Proposal Summary Slide Template
  - **(required) Attachment D:** Proposal Instructions and Volume I Template (Technical and Management)
  - **(required) Attachment E:** Proposal Instructions and Volume II Template (Price)
  - **(required) Attachment F:** Streamlined Cost Buildup Workbook
  - **(required) Attachment G:** Model Other Transaction for Research
  - **(required) Attachment H:** Task Description Document (TDD) Template
  - **(required) Attachment I:** Other Transaction Certification Template
  - **(required) Attachment X:** Proposal Overview & Proposed Metrics

## SECTION IV: SPECIAL CONSIDERATIONS

- This solicitation, stated attachments, and websites incorporated by reference constitute the entire solicitation. In the event of a discrepancy between the announcement, attachments, or websites, the solicitation takes precedence.
- All responsible sources capable of satisfying the Government's needs, including both U.S. and non-U.S. sources, may submit a proposal that shall be considered by DARPA. Historically Black Colleges and Universities, Small Businesses, Small Disadvantaged Businesses and Minority Institutions are encouraged to submit proposals and join others in submitting proposals; however, no portion of this announcement will be set aside for these organizations' participation due to the impracticality of reserving discrete or severable areas of this research for exclusive competition among these entities. Non-U.S. organizations and/or individuals may participate to the extent that such participants comply with any necessary nondisclosure agreements, security regulations, export control laws, and other governing statutes applicable under the circumstances.
- As of the time of publication of this announcement, all proposal submissions are anticipated to be unclassified.
- Federally Funded Research and Development Centers (FFRDCs), University Affiliated Research Centers (UARCs), and Government Entities to include National Laboratories are not eligible to propose to this solicitation.
- Organizational Conflicts of Interest (OCI): An organization cannot simultaneously provide scientific, engineering, technical assistance (SETA), advisory and assistance services (A&AS), or similar support to DARPA and also be a performer on a DARPA research program. If a prospective proposer believes a conflict of interest exists or may exist (whether organizational or otherwise) or has questions on what constitutes a conflict of interest, the proposer must send their contact information and a summary of the potential conflict to the specific email address identified in this PA before time and effort are expended in preparing any submission documentation.
- Department of War (DoW) policy states that the publication of products of fundamental research will remain unrestricted to the maximum extent possible. National Security Decision Directive (NSDD) 189 defines fundamental research as follows:
  - “Fundamental research” means basic and applied research in science and engineering, the results of which ordinarily are published and shared broadly within the scientific community, as distinguished from proprietary research and from industrial development, design, production, and product utilization, the results of which ordinarily are restricted for proprietary or national security reasons.
- As of the date of publication of this solicitation, the Government expects that program goals as described herein may be met by proposers intending to perform fundamental research and does not anticipate applying publication restrictions of any kind to individual awards for fundamental research that may result from this solicitation. Notwithstanding this statement of expectation, the Government is not prohibited from considering and selecting research proposals that, while perhaps not qualifying as fundamental research under the foregoing definition, still meet the solicitation criteria for submissions. If proposals are selected for

award that offer other than a fundamental research solution, the Government will either work with the proposer to modify the proposed statement of work to bring the research back into line with fundamental research or else the proposer will agree to restrictions in order to receive an award. For additional information on fundamental research, please visit [Proposer Instructions: General Terms and Conditions](#).

- DARPA’s Fundamental Research Risk-Based Security Review Process (FRRBS) is an adaptive risk management security program designed to help protect the critical technology and performer intellectual property associated with DARPA’s research projects by identifying the possible vectors of undue foreign influence. DARPA will create risk assessments of all proposed senior/key personnel selected for negotiation of a fundamental research OT award. The DARPA risk assessment process will be conducted separately from the DARPA scientific review process and adjudicated prior to final award. Proposer instructions have been updated on the DARPA landing page and can be found here: [Proposer Instructions: Other Transactions | DARPA](#). Anyone submitting a proposal requesting a fundamental research OT must comply with the Fundamental Research Risk-Based Security Reviews (FRRBS) requirements as outlined.

All submissions proposing to this solicitation must complete the Common Disclosure Forms, biosketch, and current/pending support forms. Forms can be found here and submitted to us:

- <https://www.nsf.gov/policies/nspm-33/common-form-biosketch>
- [https://nsf-gov-resources.nsf.gov/files/commonform-cps-r.pdf?VersionId=jCesM9lvHh51pYgd\\_ffVAdSm4f62YQn2](https://nsf-gov-resources.nsf.gov/files/commonform-cps-r.pdf?VersionId=jCesM9lvHh51pYgd_ffVAdSm4f62YQn2)
- Proposers should indicate in their proposal whether they believe the scope of the research included in their proposal is fundamental or not. While proposers should clearly explain the intended results of their research, the Government shall have sole discretion to determine whether the proposed research shall be considered fundamental and to select the award instrument type. Appropriate language will be included in resultant awards for non-fundamental research to prescribe publication requirements and other restrictions, as appropriate. This language can be found at [Proposer Instructions: General Terms and Conditions](#).
- For certain research projects, it may be possible that although the research to be performed by a potential awardee is non-fundamental research, its proposed sub-awardee’s effort may be fundamental research. It is also possible that the research performed by a potential awardee is fundamental research while its proposed sub-awardee’s effort may be non-fundamental research. In all cases, it is the potential awardee’s responsibility to explain in its proposal which proposed efforts are fundamental research and why the proposed efforts should be considered fundamental research.
- The APEX Accelerators program, formerly known as the Procurement Technical Assistance Program (PTAP), focuses on building strong, sustainable, and resilient U.S. supply chains by assisting a wide range of businesses that pursue and perform under contracts with the DoD, other federal agencies, state and local governments, and government prime contractors. See [www.apexaccelerators.us/](http://www.apexaccelerators.us/) for more information.

APEX Accelerators helps businesses:

- o Complete registration with a wide range of databases necessary for them to participate in the government marketplace (e.g., SAM).
  - o Identify which agencies and offices may need their products or services and how to connect with buying agencies and offices.
  - o Determine whether they are ready for government opportunities and how to position themselves to succeed.
  - o Navigate solicitations and potential funding opportunities.
  - o Receive notifications of government contract opportunities on a regular basis.
  - o Network with buying officers, prime contractors, and other businesses.
  - o Resolve performance issues and prepare for audit, only if the service is needed, after receiving an award.
- Project Spectrum is a nonprofit effort funded by the DoD Office of Small Business Programs to help educate the Defense Industrial Base (DIB) on compliance. Project Spectrum is vendor-neutral and available to assist businesses with their cybersecurity and compliance needs. Their mission is to improve cybersecurity readiness, resilience, and compliance for small/medium-sized businesses and the federal manufacturing supply chain. Project Spectrum events and programs will enhance awareness of cybersecurity threats within the manufacturing, research and development, and knowledge-based services sectors of the industrial base. Project Spectrum will leverage strategic partnerships within and outside of the DoD to accelerate the overall cybersecurity compliance of the DIB.  
[www.projectspectrum.io](http://www.projectspectrum.io) is a web portal that will provide resources such as individualized dashboards, a marketplace, and Pilot Program to help accelerate cybersecurity compliance.
  - DARPAConnect offers free resources to potential performers to help them navigate DARPA, including “Understanding DARPA Award Vehicles and Solicitations”, “Making the Most of Proposers Days”, and “Tips for DARPA Proposal Success”. Join DARPAConnect at [www.DARPAConnect.us](http://www.DARPAConnect.us) to leverage on-demand learning and networking resources.

## **RESEARCH OTHER TRANSACTIONS (OT) AUTHORITY**

Research OT projects are awarded under the authority of 10 U.S.C. § 4021. OTs are award instruments that are generally not subject to the Federal laws and regulations governing procurement contracts. OTs are not subject to the Federal Acquisition Regulations (FAR), the Defense Federal Acquisition Regulations Supplement (DFARS), or any regulations governing grants or cooperative agreements. OT awardees are also not subject to the FAR/DFARS cost accounting standards. As a result, the use of OTs should streamline the award process, making it faster than the traditional Government contracting process and more akin to a commercial contract award timeline. For additional information on OTs and how they are used in the Department of War, please go to [acquisitioninnovation.darpa.mil](http://acquisitioninnovation.darpa.mil).