

DPA26TZ05-DV003
SPEED DIAL
Frequently Asked Questions (FAQs)

September 08, 2026

Q1: For the Small Business Technology Transfer team, how should responsibilities be divided between the research institution's discovery-engine research and the small business's production platform, workflow integration, and commercialization activities?

A1: Proposers must define the division of responsibilities while strictly adhering to standard STTR allocation requirements (e.g., minimum 40% small business, minimum 30% research institution).

Q2: What verification, reproducibility, version-control, and benchmarking infrastructure is required before an automatically discovered algorithm can enter a trusted engineering workflow?

A2: Proposers must identify, propose, and justify the necessary infrastructure required to ensure trust, verification, and successful transition into operational engineering workflows.

Q3: Which engineering environments must be supported first (MATLAB/Simulink, COMSOL, Python, C++, other), and what level of auto-code generation/compilation/deployment is expected?

A3: Target environments and the extent of auto-compilation are dependent on the proposer's selected domain. Proposers must justify their choices based on their proposed domains.

Q4: What is the desired relationship between the discovered-algorithm library and online/on-demand discovery: retrieve-and-adapt, warm-start new searches, or fully rediscover per engineering problem?

A4: There is no mandated relationship. Proposers should define a viable concept of operations that optimally supports their target domains and workflow integration.

Q5: How should physical constraints, engineering requirements, hardware constraints, and certification requirements be encoded into the algorithm-discovery search space?

A5: Encoding methods are domain-dependent. Proposers must define and justify their approach to integrating physical, engineering, and hardware constraints into the discovery process.

Q6: What discovery-engine paradigms must the platform support - transformers, reinforcement learning, genetic programming, program synthesis, differentiable programming, or an extensible plug-in architecture?

A6: No specific paradigms are recommended. Proposers are responsible for proposing and justifying an architecture that best meets the topic objectives.

Q7: How will DARPA evaluate human interpretability of discovered algorithms, and what representation or proof/verification artifacts are expected?

A7: Interpretability will be evaluated based on the algorithm's composition of transparent, mathematically provable building blocks that domain experts can understand, verify, and certify.

Q8: What quantitative performance evidence is sufficient to establish Phase-I-equivalent feasibility, and must the prior discovered algorithms be in the same application domains proposed for Phase II?

A8: Proposers must provide written evidence (e.g., independent technical reports, peer-reviewed publications, or independent technical reviews) of their capability to discover novel algorithms. Prior discovered algorithms are not required to be in the same application domain proposed for Phase II, provided the proposer can substantiate that the capability translates effectively.

Q9: What constitutes a “novel algorithm” for Direct-to-Phase-II feasibility - a new symbolic/mathematical structure, a learned algorithmic policy, an automatically composed solver, or any method that demonstrably outperforms known baselines?

A9: A novel algorithm is a new method that systematically solves a family of problems, supported by rigorous mathematical backing. Methods based on simple heuristics, rigid conditional logic (e.g., standard if/then statements), or "black-box" approaches without a mathematical explanation are not of interest.

Q10: What classes of scientific/engineering algorithms should Phase II prioritize for discovery and deployment (optimization, control, estimation, signal processing, partial differential equation solvers, linear algebra, or other domains)?

A10: No specific classes of algorithms are prioritized. Proposers must identify the classes of algorithms relevant to their selected application domain and justify their dual-use potential.

Q11: Does prior work using Symbolic Regression and Genetic Programming to discover interpretable partial differential equation (PDE) meta-solvers satisfy the Phase I feasibility requirements for Direct-to-Phase-II entry?

A11: It may, provided the proposer supplies sufficient evidence, as noted in A8, that the prior work demonstrates discovery of novel, interpretable algorithms with quantitative performance gains.

Q12: Must Phase I-equivalent feasibility evidence derive from prior Department of War- or DARPA-funded work specifically?

A12: No. Evidence may come from work supported by any funding source, *except prior SBIR/STTR efforts*.

Q13: May a feasibility case be assembled from several prior efforts funded from different sources?

A13: Yes. Proposers may draw on multiple prior efforts, provided the evidence is clearly attributed and collectively demonstrates the required capability as noted in A8.

Q14: For an STTR team, may prior work performed by the research institution partner be used as feasibility evidence, or must the evidence be attributable only to the small business concern?

A14: Prior work performed by the research institution partner is acceptable, provided the proposal clearly attributes the work and explains its relevance to the team’s capability (also see A8 and A13).

Q15: Would automated algorithm discovery applied to non-linear PDEs in extreme shock physics, including prediction of dynamic metal failure and additive manufacturing processes, qualify as a suitable Department of War use case?

A15: Potentially yes. Proposers should justify the relevance of the selected use case and its dual-use potential.

Q16: Does an auto-compilation workflow concept, in which an AI engine compiles bespoke simulation modules directly into production C++ and commercial software environments, align with the required in-context integration toolchain?

A16: Based on the provided question, it is unclear how this concept addresses automated algorithm discovery. Proposers should justify their target environments, integration approach, and extent of auto-compilation based on the selected application domains.

Q17: Does “novel algorithm discovery” require a general-purpose symbolic or mathematical discovery capability, or can discovery of novel AI/ML algorithms and model architectures also satisfy the feasibility requirement?

A17: What matters is the demonstrated capability to discover novel algorithms that systematically solve a family of problems, are supported by rigorous mathematical backing, and are not black-box approaches without mathematical explanation.

Q18: Would bilevel or co-evolutionary structural search over interpretable algorithmic primitives, producing a new executable algorithm rather than only tuned parameters, fall within SPEED DIAL’s intended definition of automated algorithm discovery?

A18: It may, provided the result is a novel algorithm rather than parameter tuning or a black-box heuristic.

Q19: For the Phase I feasibility requirement, would prior results on a surrogate-assisted bilevel search engine, together with preliminary structural-discovery results, be sufficient if sparse recovery is used as the initial benchmark and Phase II demonstrates transfer to a second distinct Department of War-relevant problem?

A19: Prior discovered algorithms need not be in the same application domain proposed for Phase II, provided the proposer substantiates that the capability translates effectively.

Q20: Would a detailed self-authored technical report, together with a public benchmark leaderboard result and a publicly available code repository, be sufficient written evidence? What form of written evidence is sufficient for feasibility?

A20: No. Self-authored materials, public repositories, or leaderboard results alone are not sufficient. As noted in A8, documentation supporting the Phase I Feasibility claims includes technical reports, peer-reviewed publications, and demonstration results—all verified and validated by independent parties.

Q21: Should feasibility documentation be contained within the 10-page feasibility section of the Direct-to-Phase-II Technical Volume or may supporting technical reports and publications be

submitted as attachments or appendices outside the page limit?

A21: Proposers should provide sufficient evidence within the 10-page feasibility section, summarizing the relevant points from independent technical reports and publications and, when available, providing public links to the documentation.

Q22: Should a SPEED DIAL framework include verification processes, formal methods, or other trust infrastructure before discovered algorithms are used in standard workflows?

A22: Proposers should identify and justify the verification, reproducibility, benchmarking, and related trust infrastructure needed for their selected domains and workflows. (see A2)

Q23: Are blockchain or distributed-ledger elements required for recording verification evidence or providing non-repudiation?

A23: No specific infrastructure is required. Proposers may include such elements if justified by the selected application domains and workflow concept.

Q24: Is causal discovery, analysis, and reasoning in scope for SPEED DIAL?

A24: No specific discovery paradigms are prescribed. Proposers may include causal discovery, analysis, or reasoning if they justify its relevance to the topic objectives and proposed application domains.
