

**DARPA-PA-25-07-02 Compositional Learning-And-Reasoning for AI Complex Systems  
Engineering (CLARA)  
Frequently Asked Questions (FAQs)  
as of 3/23/2026**

60Q: Given the 10-page limit, how should I include documentation on an existing technical project?

60A: External references may be included in a proposal via the bibliography, which can include further technical documentation. Proposers are encouraged to submit concise, but descriptive, proposals.

59Q: Is [a proposed area] a suitable application task domain for the program?

59A: The application task domains listed in the solicitation are only examples, proposals may include different application task domains. Any proposed application task domain must include the requirements laid out in the solicitation Section 1.D, such as the performance benchmark to compare against, laid out in the corresponding proposal.

58Q: Does the integration of AR methods into transformer decoding qualify as the kind of compositional AR-based ML that CLARA envisions?

58A: Most methodologies that integrate an automated reasoning method into the internal processing of a machine learning model would be considered as a valid form of AR-based ML.

-----↑↑↑New Q/A↑↑↑-----

57Q: Given the solicitation statement that “Assurance here means verifiability with strong explainability to humans, based on automated logical proofs and hierarchical, vetted logic building blocks,” we ask whether the referred “automated logic proofs and hierarchical, vetted logic building blocks,” would suffice to substantiate strong explainability to humans or it would be desirable to plan for IRB submissions for evaluating this aspect.

58A: DARPA does not expect non-IV&V performers to need to plan for IRB submissions.

56Q: Has the proposal template been released, or is there an expected date for when it will be made available?

56A: The majority of the proposal templates are provided in the Program Announcement (DARPA-PA-25-07): <https://sam.gov/opp/bae5fa32f4ee4cdb9a070a14309b2c66/view>.  
The “Streamlined Cost Buildup Workbook” (TA1 and TA2) are provided with the DO (DARPA-PA-25-07-02):  
<https://sam.gov/opp/0e76cc6102804ab68667605d9e28c900/view>.

55Q: I am confused on the \$1,940,000 vs \$2,000,000 price cap and how the Hackathon's are priced. Can we submit a price of \$2,000,000?

55A: Yes. The maximum proposed amount should not exceed \$2,000,000.

Per the solicitation: "The total award value for the combined Phase 1 base (Feasibility Study) and Phase 2 option (Proof of Concept) is limited to \$2,000,000."

Per the Solicitation: "Hackathon incentive payments will count towards the total \$2,000,000 funding cap allowed in the program, meaning that non-Hackathon milestone costs should total no more than \$1,940,000."

The total combined value of the 2 Hackathons (Phase 1 and Phase 2) will not exceed \$60,000. In their proposal, TA1 performers should assume that the submitting team will win each Hackathon (Phase 1 and Phase 2) and propose it as a milestone for each Phase.

54Q: When are proposals due?

54A: The proposal due date has been extended to April 17, 2026, at 4:00 p.m. ET, per Amendment 1 to the DO (<https://sam.gov/opp/0e76cc6102804ab68667605d9e28c900/view>).

53Q: I am hoping to add a collaborator to my proposed team who is a US citizen but works in [foreign country]. Would that be allowed?

53A: Yes, with potential risk mitigation.

52Q: Can proposals contain Controlled Unclassified Information (CUI)?

52A: CLARA proposals are not anticipated to contain CUI. Please see Section 8.2.2 of the Program Announcement (DARPA-PA-25-07) for more information on CUI.

51Q: A few kinds (e.g., Bayesian LP) are listed under both Automated Reasoning and Machine Learning. How does that reconcile with the multiplicity of kinds metric?

51A: Bayesian LP is a kind of AR, however there are Bayesian LP based forms of ML as well.

50Q: In the BAA, page 6, bullet 3 states: "A proposed unclassified application task domain (set of problems) for the approach." Does this mean the use case needs to be open-sourced? Are we allowed to use proprietary use cases? Are we required to share the use case with TA2 and other TA1 performers?

50A: See the solicitation section 1.H for details on expected licensing within the program. Software and algorithms developed over the course of the program should be open sourced. Our very strong preference is for work being performed on open non-

proprietary data due to the program plan including sharing and demonstrating results to other performers in events such as the Workshops and Hackathons.

49Q: Is the assumption that Automated Reasoning operates over a symbolic set of variables safe to make?

49A: The AI research community does not have a well agreed upon definition of “symbolic,” making it difficult to definitively answer this question one way or another. The solicitation refers to several acceptable kinds of automated reasoning in section 1.H.

48Q: Some Bayesian methods (e.g., exact inference in causal probabilistic programs) are auditable/interpretable/verifiable in ways similar to logic programs. Would it be OK to think of these methods as forming the Automated Reasoning (AR) component, even though they may not directly fit under any of the listed "AR kinds" (Logic Programs, Classical Logic, Answer Set Programs)?

48A: Some kinds of causal probabilistic programs can be viewed as a subkind of probabilistic logic programs. This will be evaluated on a case-by-case basis. Section 1.H of the solicitation outlines the evidence a proposal must include for kinds of AR that are not in the pre-identified list of allowable kinds of AR.

47Q: Is a performer allowed to participate on multiple teams within the same task area (TA)? I am intending to Prime a proposal, but a number of other teams have asked me to join in an ancillary way to help their project. If the scope of work is mutually exclusive, can one be a Prime on one proposal and be a sub on another?

47A: Potential awardees may be involved in multiple proposals. However, there is a potential conflict of interest if supporting multiple awards on the same program. You must provide a mitigation plan to demonstrate how you are going to deconflict the issue(s), protect stakeholders’ proprietary information, and provide/ensure sufficient support for the effort(s). Additionally, you will need to document how the work being assigned in one effort differs from the other. The Government will not pay for the same work performed on multiple awards.

46Q: Must the Prime organization be finalized by the time abstracts are submitted?

46A: Yes. A Prime organization needs to submit the abstract. If the question pertains to teaming, however, that does not need to be finalized at time of abstract submission.

45Q: If an organization is not the Prime, may it participate as a TA1 sub-awardee on more than one proposal submitted by different Primes? If permitted, are there any restrictions (e.g., limits on the number of proposals, required disclosure of overlapping roles/level-of-effort, or other conflict/commitment rules)?

45A: Subawardees may be involved in multiple proposals. However, there is a potential conflict of interest if supporting multiple awards on the same program. You must provide a mitigation plan to demonstrate how you are going to deconflict the issue(s), protect stakeholders' proprietary information, and provide/ensure sufficient support for the effort(s). Additionally, you will need to document how the work being assigned in one effort differs from the other. The Government will not pay for the same work performed on multiple awards.

44Q: Is there an order of magnitude approximation of the number of TA 1 performers? 4, 8, 16?

44A: The number of awards will depend on the quality of the proposals received and the availability of funds.

43Q: The documents of YFA 2026 precludes Chinese citizens (on H1B visa, without green card) from being eligible to serve as a Key person. Does this apply to CLARA too?

43A: No.

42Q: I did not see Foundational models, e.g. LLM, VLMs in the list of ML techniques. Are foundational models in scope?

42A: Please see 41Q and 41A.

41Q: Are Generative/Large Language Models in-scope for the potential ML "kinds" to wrap into an AR system? The list on page 15 of the solicitation does not include them.

41A: Large Language Models would be considered a subkind of Transformer based Neural Nets, which are explicitly listed in section I.H of the solicitation as an allowable kind of machine learning. Many types of generative models would fall under the Neural Net kind laid out in section I.H.

40Q: In the abstract template, it stated that the abstract should answer 3 questions from the Heilmeier Catechism. I took this to mean that the other 5 questions need not be addressed in the abstract. Is this correct?

40A: Yes.

39Q: Can you expand on the expectations on assurance in a probabilistic context? Reasoning about financial regulations can be done purely deductively, but most ML systems provide real-valued outputs. Reasoning with probabilistic semantics can only provide guarantees when the ML systems provide well-calibrated probabilities. Furthermore, providing clear decision support semantics in the context of probabilistic semantics seems difficult? In the example you mentioned, at what probability that terrorists are using a school should a school bus be allowed within an attack target?

39A: The focus is on providing inferences that have probabilities and also meta-knowledge about those probabilities, e.g., interval bounds, probability distributions over probability distributions, approximation characteristics.

38Q: For TA1 performers utilizing hardware-in-the-loop architectures for deterministic verification, what is the anticipated flexibility within the TA2 library to support the ultra-low-latency hardware interfaces required for real-time mission-critical applications like Kill Web or active Wargaming?

38A: CLARA's focus is on software rather than (directly) hardware aspects.

37Q: Exploiting additive or multiplicative structure can reduce probability distribution size dramatically — but that assumes that the model structure is known upfront. In real-world ML deployment, structure is often unknown or changes with data distribution. How do performers handle the case where the probabilistic structure must itself be learned, and does that re-introduce the intractability you're trying to escape?

37A: DARPA expects learning probabilistic structure to be one of the technical challenges of the program; training of models is discussed throughout the solicitation. DARPA does not agree with the assumption that model structure must be known upfront to be sufficiently exploited, and that the process of learning the structure would necessarily introduce tractability constraints of the kind the program is seeking to avoid.

36Q: Would you please confirm that there is no cost share requirement?

36A: Cost sharing may be required in certain situations for OTs for prototype projects awarded under the authority of 10 U.S.C. § 4022. Please see Appendix A of DARPA-PA-25-07 for more information.

35Q: For this proposal will the actual model we are building a product for be a model that DARPA owns or is to be built on a certain open source model or a foundation model?

35A: DARPA will not be providing a model it owns nor recommending a particular open source or foundation model. Performers are anticipated to be constructing their own AR-based-ML models.

34Q: If we are building off a foundation model will we focus on using API calls or the General Model that has memory, etc., and if a general model is used will we have access to the foundation models settings that are normally not visible, to avoid contradicting commands?

34A: CLARA's focus is on open IP (see section I.H. of DARPA-PA-25-07-02). DARPA is not expecting to provide foundation models nor access to them to performers.

33Q: Are computational resources, e.g. for LLM access, going to be made available to performers, or part of their budget?

33A: Computation resources would be part of the performer's budget.

32Q: The way the problem is posed, it appears that the AR is used for learning the models. A decade or so back, an area called Lifted Inference was popular for using KR during inference. Is this interesting as well? Or should we only focus on learning?

32A: AR (a.k.a. KR or KRR) has been used for ML inference in numerous previous approaches. CLARA is aiming for development of approaches that can also do ML training as well as ML inferencing, as presented in the solicitation's "Multiplicity of AI Kinds in Composition" metric in section I.E of the DO. ML training notably requires ML inferencing.

31Q: If I wished to provide limited componentry potentially useful by proposers, will I have the opportunity to do so?

31A: DARPA would encourage the submission of a teaming profile for this and similar cases.

30Q: Is the white paper really due March 2nd, which is just 11 days away for the Information Session? If so, would you accept a late submission?

30A: Abstracts are limited to a single page of technical content and should be submitted no later than the date and time listed in the Solicitation (March 2nd, 4pm ET).

29Q: The solicitation mentions a \$2M cap for the whole 2-year program. What do you anticipate being individual proposal awards?

29A: Per the solicitation: "The total award value for the combined Phase 1 base (Feasibility Study) and Phase 2 option (Proof of Concept) is limited to \$2,000,000. This total award value includes Government funding and performer cost share if required or proposed." This amount is per proposer that receives an award.

28Q: Can you address the appropriateness for work involving development of new theoretical Machine Learning or Automated Reasoning kinds (frameworks, structures, etc.) not mentioned in the technical overview?

28A: The development of new kinds of Machine Learning or Automated Reasoning is appropriate insofar as it contributes to the success of CLARA's overall objectives, metrics, etc.

27Q: Are UARCs invited to submit proposals?

27A: See section III of the DO for eligibility information.

26Q: When are abstracts due?

26A: No later than March 2, 2026, at 4:00 p.m. ET.

25Q: I sense a connection between the CLARA and CODORD program at DARPA. Where possible, can you express your thoughts on any links between CLARA and CODORD?

25A: ErgoAI, cited in the CLARA solicitation, is used in CODORD as an important kind of Automated Reasoning in that program. CODORD's and CLARA's success are potentially complementary, but neither is dependent on the other.

24Q: The Hackathon evaluation process was described as a "winner take all" payoff. To support program goals for using each other's work, might an alternative reward structure, like crediting what components were used in a winning synthesis system?

24A: The reward structure for the Hackathon performance may involve criteria such as crediting other performers' work, but the emphasis will be on simplicity.

23Q: Can universities participate in OT?

23A: Yes.

22Q: Are Machine Learning (ML) systems anticipated to produce symbolically ground outputs for Automated Reasoning (AR), or will the grounding between ML outputs and AR symbols be part of the infrastructure created by the TA 2 performer?

22A: Per the solicitation section I.C, the focus of TA2 is to provide low-lift integration especially pluggability in a composition library of TA1 techniques' software.

"Grounding" could be addressed by the techniques developed in TA1.

21Q: In distinguishing approaches where AR is tacked onto ML versus approaches that achieve AR-driven ML, what is the most important difference?

21A: AR-based ML has AR involved in the guts of the ML, i.e., the ML model does ML inferencing and training that is based on AR. AR that is tacked onto ML will often look like ML models sending requests to external AR systems.

20Q: Could you please clarify if the goal is to do this compositional approach during training of the AI component, or training and inference, or only training or only inference?

20A: The overall program goal is to use the compositional approach during training and inferencing. Per the solicitation, Phase 1 metrics only include inferencing while Phase 2 metrics include both inferencing and training.

19Q: Which part of the Feb 19 webinar will be posted online, and when?

19A: The CLARA overview portion of the webinar can be found on the CLARA program page: <https://www.darpa.mil/research/programs/clara>.

18Q: Are software, minimal hardware, materials and travel allowed in the budget proposal?

18A: They are not precluded from being included in the cost proposal. Please include all cost elements as appropriate per the solicitation.

17Q: Are there any preferred scenarios for Phase 1?

17A: No. In the solicitation section I.D, several potential performer level application task domains are specified. In the solicitation section I.C, a single example Hackathon Scenario is presented. These are intended for illustrative purposes; they are not preferred scenarios. IV&V Hackathon scenario development will depend on what approaches TA1 performers take and will be created in consultation with the TA1 and TA2 performers.

16Q: Will data be available for any preferred scenarios for Phase 1?

16A: Scenario data will not be available during the period of preparing your proposals.

15Q: What is the maximum allowed awards ceiling for Phase I & 2 budgeting purposes?

15A: Per the Solicitation: "Hackathon incentive payments will count towards the total \$2,000,000 funding cap allowed in the program, meaning that non-Hackathon milestone costs should total no more than \$1,940,000. Further information on Milestones can be found in section I.F." These amounts are per competitive performer team.

14Q: IV&V will develop the composition patterns used in hackathons. Given that the performers need to optimize for those patterns, isn't there a risk that the TA1 approaches will overfit to IV&V's chosen scenarios rather than generalizing to novel compositions? How does CLARA guard against that?

14A: Hackathon performance is only a portion of one of the six metrics in the program, most other metrics, in particular "Verifiability without loss of performance" and "Composed Task Reliability of ML+AR system," are relative to the performer's own proposed application domains. IV&V scenario development is expected to involve consultation with both TA1 and TA2 performers, in particular during the workshops, and is not expected to begin until after the performers have already proposed their own application task domains.

13Q: Where does reinforcement learning (RL) fit into your categorization of machine learning kinds?

13A: RL has its own category as it is somewhat in between the other machine learning techniques, for instance it may utilize different kinds of ML.

12Q: Is CLARA interested both in approaches that would train new ML models and in approaches that would compose/integrate existing models?

12A: Yes.

11Q: Can performers assume that logic programs for automated reasoning (e.g., constraints from financial regulations) are given in our application domains rather than learned from data. Or, succinctly, is structure learning out of scope?

11A: No, performers should not assume logic programs are given in the application domains rather than learned from data. Structure learning is in scope.

10Q: How important is it for performers to have deep domains (e.g., not MNIST addition) versus being able to quickly adapt to IV&V provided domains?

10A: The depth of performer domains is up to the performers to decide. See the solicitation section I.D for some example task domains. As outlined in the solicitation section I.D, TA1 proposals should include “A proposed unclassified application task domain (set of problems) for the approach, with explicit state-of-the-art (SOA) performance benchmarks to compare against and corresponding train/test corpuses. The proposed domain should be relevant to the DARPA mission, with several potential options laid out at the end of this section. At least one SOA benchmark should have an associated, identified pair or larger set of previously existing ML and AR models. The SOA performance should have an associated sample complexity.” IV&V will be responsible for verifying and validating performer system’s performance relative to the performer proposed domains, with performer proposed SOA comparison points, and performer proposed datasets. Hackathon scenarios, developed by IV&V, are in addition to performer proposed domains.

9Q: Can you explain the award structure of the hackathons? For budgeting, should we have a milestone with an expected value of \$30k which is payable on placing first in the hackathon for each of 2 hackathons?

9A: Per the solicitation: "Hackathon attendance is included in the milestone schedule. In addition, each Hackathon will have a milestone and corresponding payment associated with it, which only the top performing TA1 Hackathon team may qualify for. This milestone payment will be significantly smaller than all other milestone payments and is intended to incentivize Hackathon performance. Hackathon incentive payments will count towards the total \$2,000,000 funding cap allowed in the program, meaning that non-Hackathon milestone costs should total no more than \$1,940,000. Further information on Milestones can be found in section I.F."

8Q: Are temporal logics in scope?

8A: Logics that treat temporality are potentially in scope. This will depend on whether they have other characteristics that might make them suitable - versus not - for CLARA. See the solicitation for details on expected characteristics, in particular sections I.A, I.E, and I.H.

7Q: Is the following within CLARA's scope: transforming current Neural Net based ML architectures so they are significantly more compute/flops (and power) efficient; more human understandable/explainable (and therefore more predictable, trustworthy) but preserve performance?

7A: Modifying existing ML and AR techniques, as well as inventing/developing quite new ones, is in scope.

6Q: Are teaming profiles required?

6A: No. Potential proposers may choose to share their information with other registrants via the Proposer Profile. For further information please see DARPA-SN-26-28.

5Q: How can we submit teaming information?

5A: All profiles must be emailed to [CLARA@darpa.mil](mailto:CLARA@darpa.mil) no later than 4:00 p.m. on February 23, 2026.

4Q: Can you clarify what is required for the “Proposer profile for teaming” due on Feb 23?

4A: Information session registrants interested in discovering collaboration opportunities for this announcement may submit a one-page profile. The proposer’s profile should include their contact information (name, organization, email, telephone number, mailing address, and, if applicable, organization website), a brief description of their technical competencies, and, if applicable, their desired expertise from other teams/organizations. All profiles must be emailed to [CLARA@darpa.mil](mailto:CLARA@darpa.mil) no later than 4:00 p.m. on February 23, 2026.

3Q: Will the CLARA program consider proposals with a domain focus on Physical AI, that is AR+ML hierarchical architecture for robots interacting with the physical world?

3A: Physical AI is a fine application task domain area; similar examples are laid out in the solicitation section 1.D.

2Q: Are there specific mechanisms for Service-Disabled Veteran-Owned Small Business firms to partner with academic institutions while maintaining prime contractor status?

2A: There are no such mechanisms specifically outlined under this PA/DO.

1Q: For Phase 1 feasibility studies, would a focused demonstration on a single ML/AR integration be considered an acceptable scope?

1A: Per Section 1.E of the Solicitation, Phase 1 metrics will require a minimum of a single ML and a single AR to be integrated. The integration of such a pair would be an acceptable scope for meeting Phase 1 metrics, though insufficient to meet Phase 2 metrics.

