DARPA Mapping Machine Learning to Physics (ML2P) Program - Q&A

This document summarizes questions and answers related to the DARPA ML2P program, based on the DARPA-PS-25-32 solicitation posted to SAM.gov on September 23, 2025.

Scope and Alignment with Program Objectives

1. **How would you react to an abstract that does not address the full scope of the call, but focuses on what could be a component in a larger project?**

   A: In accordance with Attachment A1, proposers should include all information in the template to constitute a fully conforming abstract submission. Abstracts that fail to address the content requirements of Attachment A1 may be found less favorable (weaknesses, significant weaknesses) during the evaluation process described in 6.3 of the solicitation.

2. **Assuming sufficient proof of concept, is the following out of scope for ML2P: Developing new AI/ML technology that natively runs on very low power devices so that technology can also perform regression, classification, clustering, etc.; and can be benchmarked/mapped in.**

   A: The ML2P program aims to improve power consumption and performance of machine learning (ML) models by preserving local energy semantics and tuning energy-performance objective functions. ML2P will explore a set of multi-objective functions, balancing the trade-off of joules with performance, covering a range of common ML tasks. Technical solutions should design experiments and collect the energy semantics of machine learning (ES-ML) to discover interactions between upstream optimizations and their downstream effects. New AI/ML technology other than hardware is welcomed. Note that the ML2P program objective is producing code, algorithms and documentation; so new hardware is not in scope.

3. **Is developing novel low-power ML models in scope, or is ML2P singly focused on developing a rigorous power benchmarking/mapping technology for existing ML**

   A: ML2P is focused on the entire model pipeline: not just the model itself but data ingest and preparation, model building and evaluation, and inference. In particular, ML2P is not concerned with the power used by a single component, but rather the power used as a result of interactions between each step or component of the process.

4. **Is online learning in scope of the program or are only train-then-inference models of interest?**

A: Development teams may use any ML model or algorithm they choose, as long as they are able to capture the energy semantics of the entire ML pipeline (not just the energy used by the model).

5. **Regarding discovery of new energy efficient algorithms, are distributed multi-node algorithms (like branch-train-merge mentioned at proposers' day) within scope?**

   A: Development teams may use any machine learning model or algorithm they choose, as long as they are able to capture the energy semantics of the entire ML pipeline (not just the energy used by the model).

6. **My team was wondering if large language models were within scope for ML2P or whether other types of models (e.g., strictly classification or clustering models) were meant to be the focus?**

   A: Development teams may use any machine learning model or algorithm they choose, as long as they are able to capture the energy semantics of the entire ML pipeline (not just the energy used by the model).

8. We would like clarification on which software and hardware design/configuration variables fall within the scope of this program. Specifically, we are wondering whether different learning paradigms, such as distributed learning, federated learning, fine-tuning of pre-trained models, incremental learning, and prototype learning, are regarded as design variables subject to optimization. Given that these approaches exhibit distinct trade-offs between accuracy and energy consumption, it seems natural to consider and compare them within the ES-ML context. Furthermore, in the case of distributed or federated learning, should the energy costs associated with communication be explicitly included as part of the overall energy dissipation?

A: The ML2P program seeks to optimize energy usage across the entire ML lifecycle, making all design decisions within the ML pipeline in scope. This includes software, hardware, training paradigms, and activities with inherent tradeoffs within the ML pipeline. Specific learning paradigms like distributed, federated, fine-tuning, incremental, and prototype learning can be considered as design variables, considering the trade-offs between accuracy and energy consumption. In distributed or federated learning, the energy costs associated with network communication must be explicitly included in the overall energy dissipation, reflecting the program's focus on energy efficiency throughout the ML lifecycle and across optimization considerations.

TECHNICAL REQUIREMENTS AND APPROACH

1. **Can broad leeway be assumed when converting a trained model to specific hardware e.g., we can change the model architecture itself (or the model itself)?**

   A: Technical solutions should focus on developing algorithms and software to convert trained or partially trained ML models for use on the different hardware types identified within the technical solution. The goal of ML2P is to create a way for ML model builders and users to choose from a list of ML pipeline options based on the energy resources they have available. This coupled with the key transition objective of ML2P – to make ML2P software the gold standard for ML construction and simulation of power usage and trade-offs – means that ML2P is seeking solutions that can be widely adopted and easily implemented. Any combination of model and physical hardware is within scope.

2. **During the proposer's day, first order logic was mentioned as the formalism for ES-ML energy semantics representation. Are approaches not using logic to represent ES-ML considered out of scope?**

   A: Approaches not using logic to represent ES-ML are not necessarily out of scope. The use of first-order predicates is strongly encouraged. ES-ML may be extensions of existing logic, calculus, or language (e.g., linear temporal logic, modal μ-calculus, systems modeling language). Technical solutions should develop a formal representation of energy semantics for ML designed for machine readability; the use of first-order predicates is strongly encouraged.

3. **Can emulators be used as potential hardware choices? Built hardware may not exist for advanced compute fabrics, but emulation of individual operations exists. This may enable better hardware design using ML2P capability.**

   A: The PS does not explicitly address the use of emulators as potential hardware choices. However, given the program's focus on accurate prediction of power and performance of future ML models and providing a foundation for simulation research in hardware design, the use of emulators could be a valid approach.

   The hardware used for this program is expected to be physical hardware for the following reasons: a) The solicitation states that development performers should plan to deliver one unit of each hardware component used for processing and power measurement to the T&E Team for independent verification and validation; and b) The code and algorithms being produced by ML2P developers is intended to be widely adopted and used quickly, thus the hardware should already exist.

OPEN SOURCE, LICENSING, AND DELIVERABLES

1. **Must all software be open sourced? What are the circumstances in which open sourcing the software will not be required?**

A: Yes, open-source publication is expected. DARPA strongly discourages the submission of technical solutions that offer restrictive licensing, and per section 3.3.1 of the solicitation "Restrictive licensing within proposed technical solutions will be found to be a significant weakness during the scientific review process." The key transition objective of the program is to make ML2P software the gold standard for ML construction and simulation of power usage and trade-offs. Performers will be required to publish documentation, algorithms, code, and tutorials they will develop and generate under ML2P awards, as open-source (e.g., the MIT license is strongly preferred) to existing ML repository sites (e.g., scikit-learn) and, when available, the forthcoming DARPA GitHub page, in addition to publishing in conferences (e.g., NeurIPS) and peer-reviewed journals (e.g., IEEE).

2. **If we are proposing an approach that runs on an edge device (such as Xilinx ZCU 102 FPGA) with training on a separate workstation/server (such as NVIDIA 8* H200), are we expected to send both the edge device and workstation/server to the T&E team, or just the workstation/server?**

   A: The T&E team's role is to replicate and independently validate *all* results from each Development team, including the lower bound for each performer. Therefore, the delivery of all hardware components used for processing and power measurement is essential for the T&E team to perform accurate and comprehensive evaluations.

ELIGIBILITY, SUMBMISSIONS, AND AWARD INFORMATION

1. **Do you know if this opportunity (or potential Opp) is solely targeting academia?**

   A: This solicitation encourages submissions from all responsible sources capable of satisfying the Government's needs, including large and small businesses, nontraditional defense contractors as defined in 10 U.S.C. § 3014, and research institutions as defined in 15 U.S.C. § 632.

2. **Are we eligible to join as a performer given that we are a UARC.**

   A: UARCs are highly discouraged from proposing. UARCs interested in this solicitation, either as a prime or a subcontractor, should contact the Agency POC listed in the Overview section prior to the proposal (or abstract) due date to discuss potential participation as part of the government team or eligibility as a technical performer.

3. **The RFI does not mention any possibility of potential resulting RFPs. Is this true?**

   A: The ML2P Program Solicitation (DARPA-PS-25-32) is a formal request for submissions and has been published to SAM.gov. Please review the solicitation for abstract and proposal submission information.

4. **Q: How are abstracts to be submitted?**

   A: Attachment A1 (ML2P Abstract Template), published with the solicitation, describes the page limit, format, content and submission requirements for abstracts.