

Multimodal Algorithmic Reasoning Workshop (MAR 2025)

June 11 or 12th, 2025, Nashville, TN

Held in conjunction with CVPR 2025

<https://marworkshop.github.io/cvpr25/>

CALL FOR CONTRIBUTIONS

Deep learning–powered AI systems have rapidly advanced in their data modeling capabilities, yielding compelling applications that often seem to rival human intelligence. Despite these impressive achievements, questions remain about whether these systems possess the foundational elements of general intelligence, or whether they simply excel at task-specific computations without human-like understanding. Addressing these questions calls for new methods of both developing and assessing such models.

In this workshop, we aim to bring together researchers working in neural algorithmic learning, multimodal reasoning, and cognitive models of intelligence to showcase cutting-edge research, tackle current challenges, and highlight critical yet underexplored problems in perception and language modeling—issues at the core of achieving true artificial general intelligence. A key focus is on the emerging field of multimodal algorithmic reasoning, which explores neural representations of algorithms to devise novel solutions for real-world tasks. These span a wide range of areas, including multimodal learning, algorithms over foundational models for solving problems related to analysis, synthesis, or planning, mathematical problem-solving, procedural learning in robotic manipulation, and more.

Our goal is to delve deeply into this exciting intersection of multimodal algorithmic learning and cognitive science, reflecting on the current progress in machine intelligence while examining the gaps that distinguish it from human cognition. Through talks by leading researchers and faculty, we aim to inspire participants to explore the "missing rungs" on the ladder to true intelligence.

We invite you to submit high-quality papers to the workshop that propose innovative approaches, theoretical insights, or practical applications towards advancing this exciting field, as well as foster meaningful discussions and collaborations.

IMPORTANT DATES & DETAILS

Submission deadline: *March 19, 2025*** (11:59 PM PDT)**

Paper decisions to authors: April 3, 2025

Camera-ready deadline: April 7, 2025

TOPICS

We invite submissions of high-quality research papers in the topics related to multimodal algorithmic reasoning. The topics for MAR 2025 include, but are not limited to:

- * Multimodal machine reasoning
- * Algorithmic reasoning in vision, including program synthesis, planning, and procedural learning
- * Neural architectures and approaches for mathematical reasoning
- * Architectures for aligning/integrating multimodal foundation models, including vision, language, audio, and 3D content.
- * Architectures for solving abstract multimodal reasoning/language-based IQ puzzles, e.g., using sketches, diagrams, audio-visual clips, etc.
- * New tasks, datasets, benchmarks, and models for multimodal reasoning including algorithmic reasoning, neuro-symbolic reasoning, abstract reasoning, mathematical reasoning, etc.
- * Extreme generalization to new tasks and few-shot concept induction
- * Synthetic data and automatic verification for reasoning
- * Multimodal agents including programmable agent, tool-use agent, etc., for reasoning tasks
- * Position papers on novel perspectives to understand AI and human problem solving
- * Studies comparing AI and human problem solving skills, including but not limited to: i) Perspectives from psychology, neuroscience, and educational science, ii) Children's cognitive development, and iii) Limitations of large vision-and-language models
- * Vision-and-language applications.

SUBMISSION INSTRUCTIONS FOR PAPER TRACK

We have two tracks for paper submissions:

1. Papers with IEEE/CVF workshop proceedings (≤ 8 pages)
2. Papers without workshop proceedings (≤ 8 pages)

For track 1, we are inviting only original, previously unpublished papers, and dual submissions are not allowed. The page limits described above are excluding the references. Papers accepted to track 2 will not be included in the proceedings, however will be publicly shared on the workshop website. The submissions to this track can be novel/ongoing work (limited to 4 pages) or accepted/previously published papers (limited to 8 pages), both excluding references. Please see the workshop website for more details.

* All submissions are handled via the workshop's CMT website:

<https://cmt3.research.microsoft.com/MAR2025/>.

* Submissions should be made in PDF format and should follow the official CVPR 2025 template and guidelines.

- * All submissions should maintain author anonymity and should abide by the CVPR conference guidelines for double-blind review.
- * Accepted papers will be presented as either an oral, spotlight, or poster presentation. At least one author of each accepted submission must present the paper at the workshop.
- * Presentation of accepted papers at our workshop will follow the same policy as that for accepted papers at the CVPR main conference
- * Papers accepted in track 1 will be part of the CVPR 2025 workshop proceedings.
- * Authors may optionally upload supplementary materials, the deadline for which is the same as that of the main paper and should be submitted separately.

WORKSHOP ORGANIZERS

[Anoop Cherian](#), Mitsubishi Electric Research Laboratories
[Kuan-Chuan Peng](#), Mitsubishi Electric Research Laboratories
[Suhaz Lohit](#), Mitsubishi Electric Research Laboratories
[Honglu Zhou](#), Salesforce AI Research
[Kevin A. Smith](#), Massachusetts Institute of Technology
[Tim K. Marks](#), Mitsubishi Electric Research Laboratories
[Joshua B. Tenenbaum](#), Massachusetts Institute of Technology

CONTACT

Email: smart101@googlegroups.com

Website: <https://marworkshop.github.io/cvpr25/>