

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 been increasing in their data modeling abilities at an ever more vigor in recent times, with compelling applications emerging frequently, many of which may even appear to challenge human intelligence. Yet despite such impressive performance, there remain open questions about whether these models include the foundations of general intelligence, or whether they perform these tasks without human-like understanding. This necessitates development of better tools for assessing these models in tandem with developing the models themselves.

In this workshop, we plan to gather researchers working in neural algorithmic learning, multimodal reasoning, and cognitive models of intelligence to showcase their cutting-edge research, discuss the latest challenges, as well as bring to the forefront problems in perception and language modeling that are often overlooked but are pivotal in achieving true artificial general intelligence. An important goal is to bring to the forefront problems in perception, language modeling, and cognition that are often overlooked in state-of-the-art research and that are important for making true progress in AI. An emphasis of this workshop is on the emerging topic of multimodal algorithmic reasoning, where a reasoning agent is required to automatically deduce new algorithms/procedures for solving real-world tasks, e.g., algorithms that use multimodal foundational models for analysis, synthesis, and planning, new approaches towards solving challenging vision-and-language mathematical reasoning problems, deriving winning strategies in multimodal games, procedures for using tools in robotic manipulation, etc. We hope to deep dive into this exciting topic at the intersection of multimodal learning and cognitive science to understand what we have achieved thus far in machine intelligence and what we are lacking in relation to the human way of thinking -- through talks from outstanding researchers and faculty that could inspire the audience to search for the missing rungs on the ladder to true intelligence.

IMPORTANT DATES & DETAILS

Submission deadline: *March 12, 2025*** (11:59 PM PDT)**

Rebuttal: March 25-26, 2025

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
[Suhas 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/>