Multimodal Algorithmic Reasoning Workshop (MAR-CVPR 2026)

June 3rd / 4th, 2026, Denver

Held in conjunction with CVPR 2026

https://marworkshop.github.io/cvpr26/

CALL FOR CONTRIBUTIONS

Large AI frameworks have been rapidly increasing their data modeling capabilities in recent years, with compelling applications emerging frequently, some of which may even appear to challenge human intelligence. Yet, despite this impressive performance, there remain open questions about whether these models possess the foundations of general intelligence, or whether they succeed without human-like understanding. This motivates the development of better tools for assessing such models, alongside continued advances in model design.

This workshop focuses on multimodal algorithmic reasoning, where an agent must assimilate information from multiple modalities for complex problem solving. Real-world examples of such problems include: (i) chain-of-thought reasoning across modalities, (ii) vision-and-language problem solving, (iii) agentic reasoning and tool use, and (iv) reasoning under physical constraints, among others. Over the past year, we have seen rapid advances in AI that more effectively bridge modalities, inspiring both optimism about superhuman capabilities and skepticism about the limits of current approaches. This is an opportune moment to explore critical challenges, including new architectures for visual and physical reasoning, data generation via simulators, and the theoretical limits of reasoning in large models.

Through talks by outstanding researchers and faculty, we aim to delve deeply into this topic at the intersection of multimodality, algorithmic foundations, and cognitive science, to better understand what has been achieved in machine intelligence and what remains missing relative to human cognition, as we seek the next rungs on the ladder toward advancing AI to the next frontier.

IMPORTANT DATES & DETAILS

Submission deadline: March 5 (Anywhere on Earth) (earlier Feb 27th)

Final decision notification: March 20

Camera Ready: April 10

TOPICS

We invite submissions of original and high-quality research papers in the topics related to multimodal algorithmic reasoning. The topics for MAR-CVPR 2026 include, but are not limited to:

* Multimodal structured and multi-step reasoning across vision, language, audio, and other modalities, including compositional and programmatic inference.
* Multimodal foundation models and world models for reasoning, planning, and decision-making, and their connections to general intelligence.
* Reasoning under physical, geometric, and causal constraints, including embodied agents, simulators, and digital twins.
* Multi-agent reasoning and collaboration, including debate, coordination, mixture-of-experts, and reward- or critique-based aggregation.
* Extreme generalization and concept learning, including few-shot, zero-shot, and out-of-distribution multimodal reasoning.
* Scaling laws, efficiency, and test-time reasoning, including inference-time optimization, self-refinement, and tool-augmented reasoning.
* Benchmarks, datasets, diagnostics, and evaluation, including synthetic data, interpretability, and systematic analysis of shortcomings and failure modes in multimodal AI models.
* Theoretical and cognitive perspectives on multimodal reasoning, including limits of current models and insights from human cognition.
* Human–AI reasoning comparisons and foundations, including perspectives from psychology, neuroscience, and child development; theoretical limits of reasoning in large models; and position papers on how current multimodal AI reasoning differs from human cognition.

SUBMISSION INSTRUCTIONS

We are inviting submissions of both original and previously published works.

* All submissions are handled via the workshop's OpenReview website: https://openreview.net/group?id=thecvf.com/CVPR/2026/Workshop/MAR.

* Submissions should be made in PDF format and must follow the CVPR 2026 submission style provided here: https://github.com/cvpr-org/author-kit/archive/refs/tags/CVPR2026-v1(latex).zip.

* We allow three types of submissions:

1. Original and unpublished papers of up to 8 pages, which will be published as part of the CVPR 2026 workshop proceedings and will be released on the workshop website upon acceptance

2. Original and unpublished papers of up to 4 pages, which will not be included in the CVPR workshop proceedings and will be released only on the workshop website upon acceptance; and

3. Previously accepted or published papers of up to 8 pages, which will be released only on the workshop website upon acceptance to our workshop.

* All the page limits above are excluding references, acknowledgements, and other non-technical content (e.g., scope, limitations, impact statement).
* Authors may upload an optional Appendix, containing additional details, proofs, images, etc. as part of the submission pdf (after the references) or in a separate zip file (with a max of 50MB in size). The deadline for submitting these supplementary materials is the same as that for the main paper.
* All submissions should maintain author anonymity and should abide by the CVPR 2026 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 **in-person**.
* Presentation of accepted papers at our workshop will follow the same policy as that for accepted papers at the CVPR 2026 main conference.
* Accepted papers will be made publicly accessible on the workshop website shortly after the camera-ready deadline.
* The submitting authors are expected to also be reviewers for the workshop.

WORKSHOP ORGANIZERS

Anoop Cherian, Mitsubishi Electric Research Laboratories
Suhas Lohit, Mitsubishi Electric Research Laboratories
Kuan-Chuan Peng, Mitsubishi Electric Research Laboratories
Honglu Zhou, Salesforce AI Research
Kevin A. Smith, Massachusetts Institute of Technology
Joshua B. Tenenbaum, Massachusetts Institute of Technology

CONTACT

Email: smart101@googlegroups.com
Website: https://marworkshop.github.io/cvpr26/