Multimodal Algorithmic Reasoning (MAR 2024)
Workshop and Challenge
June, 2024, Seattle

Held in conjunction with CVPR 2024

https://marworkshop.github.io/cvpr24/

CALL FOR CONTRIBUTIONS

The focus of this workshop is to bring together researchers in multimodal reasoning and cognitive models of intelligence, towards contextualizing current advancements in AI within the broader aspiration of attaining machine intelligence that parallels human capabilities. While the emergence of large language and multimodal models has garnered considerable acclaim for their prowess, their capabilities are not flawless. Empirical evaluations have exposed a pronounced shortfall in their ability to engage in sophisticated reasoning and exhibit the intelligence necessary for effortless adaptation and generalization in novel and complex scenarios. A central aim of this workshop is to bring to the forefront problems in perception, language modeling, cognition, embodiment, and multimodal learning, all of which are widely acknowledged as essential elements for further advancing AI capabilities. One such problem is multimodal algorithmic reasoning, where a reasoning agent needs to derive or innovate algorithms using multimodal data for solving problems, e.g., algorithms over perceptual foundational models across various modalities like vision, language, audio, 3D, tactile input, and more. In this workshop, we plan to gather outstanding faculty/researchers working in multimodal learning and cognition to provide their opinions on the recent breakthroughs, discuss critical yet often overlooked steps in the recipe for artificial general intelligence, as well as showcase their cutting-edge research on the aforementioned topics that could inspire the audience to search for the missing pieces in the puzzle of true artificial intelligence.

IMPORTANT DATES

* Paper Track
  - **Submission deadline: ***March 11, 2024*** (11:59PM EDT)
  - Paper decisions to authors: April 5, 2024
  - Camera-ready deadline: April 10, 2024

* SMART-101 Challenge Track
  - Challenge open: March 21, 2024
  - Submission deadline: ***May 31, 2024***
  - Arxiv paper deadline to be considered for awards: June 5, 2024
  - Public winner announcement: June 17, 2024

TOPICS FOR PAPER TRACK

We invite submissions of original and high-quality research papers in the topics related to multimodal algorithmic reasoning. The topics for MAR 2024 include, but are not limited to:

* Multimodal machine cognition and learning
* Large language models, vision, and cognition including children's cognition
* Foundation models of intelligence, including vision, language, and other modalities
* Artificial general intelligence / general-purpose problem solving architectures
* Neural architectures for solving vision & language or language-based IQ puzzles
* Embodiment and AI
* Large language models, neuroscience, and vision
* Functional and algorithmic / procedural learning in vision
* Abstract multimodal reasoning, e.g., using sketches, diagrams, etc.
* Perceptual reasoning and decision making
* New vision-and-language abstract reasoning tasks and datasets
* Vision-and-language applications

SUBMISSION INSTRUCTIONS FOR PAPER TRACK

We have four tracks for paper submissions:

1. Short papers with IEEE/CVF workshop proceedings (≤ 4 pages)
2. Long papers with IEEE/CVF workshop proceedings (≤ 8 pages)
3. Papers without proceedings (≤ 8 pages), and
4. Previously published papers (≤ 8 pages).

For tracks 1–3, we are inviting only original and previously unpublished papers and dual submissions are not allowed. The page limits described above are excluding the references. We plan to accept only a very limited number of previously accepted papers in track 4 if our final program schedule permits. Please see the workshop website for more details.

* All submissions are handled via the workshop's CMT Website: https://cmt3.research.microsoft.com/MAR2024.
* Submissions should be made in PDF format and should follow the official CVPR 2024 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 paper tracks 1–2 will be part of the CVPR 2024 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.

INSTRUCTIONS FOR PARTICIPATING IN THE SMART-101 CHALLENGE TRACK

As part of MAR 2024, we are hosting a challenge based on the Simple Multimodal Algorithmic Reasoning Task – SMART-101 – dataset, which is available for download here: https://smartdataset.github.io/smart101/. The accompanying CVPR 2023 paper "Are Deep Neural Networks SMARTer than Second Graders" is available here: https://arxiv.org/abs/2212.09993.

* Details regarding the challenge will be announced soon.
* The challenge participants are required to make arXiv submissions detailing their approach. These are only used to judge the competition, and **will not be reviewed** and will not be part of workshop proceedings.
* Winners of the challenge are determined both by performance on the leaderboard over a private test set as well as the novelty of the proposed method (as detailed in the arXiv submission). Details are made available on the challenge website.
* Prizes will be awarded on the day of the workshop.

WORKSHOP ORGANIZERS

Anoop Cherian, Mitsubishi Electric Research Laboratories
Suhas Lohit, Mitsubishi Electric Research Laboratories
Honglu Zhou, Salesforce Research
Moitreya Chatterjee, Mitsubishi Electric Research Laboratories
Kuan-Chuan Peng, Mitsubishi Electric Research Laboratories
Kevin A. Smith, Massachusetts Institute of Technology
Tim K. Marks, Mitsubishi Electric Research Laboratories
Joanna Matthiesen, Math Kangaroo USA
Joshua B. Tenenbaum, Massachusetts Institute of Technology

CONTACT

Email: smart101@googlegroups.com
SMART-101 project: https://smartdataset.github.io/smart101/
Website: https://marworkshop.github.io/cvpr24/