

Muhammad Khalifa

khalifam@umich.edu

<https://mukhal.github.io>

[Google Scholar](#)

Phone: (+1) 734-510-2583

--

Researcher in Machine Learning and NLP focusing on LLM reasoning, post-training, model merging, and web agents. My work develops reward models for test-time scaling. My work is published at NeurIPS, ICLR, ACL, and EMNLP. I aim to build reliable, scalable, and safe reasoning agents.

Education

- 09/2021 - Now *University of Michigan*
PhD. in Computer Science and Engineering.
Advisors: Honglak Lee and Lu Wang.
Focus: Reasoning with LLMs, knowledge attribution and test-time scaling.
- 02/2018 - 08/2021 *Cairo University*
MSc. in Computer Science.
GPA: 90.1% (A+)
Thesis: Transfer Learning for Low-resource Arabic Natural Language Processing.
- 09/2011 - 06/2016 *Mansoura University*
BSc. in Computer Engineering.
GPA: 87% (Excellent with Honors).

Employment

- 05/2025 - 11/2025 *LG AI Research, Research Intern*
 - Training and analysis of LLM-as-a-Judge for web agents
 - Test-time scaling of web agents with outcome and process verifiers
- 06/2024 - 11/2024 *Cohere Inc., Research Intern at CodeGen team*
 - Developed scalable model merging methods ('model soups') to recycle 100B+ LLM checkpoints, mitigating tradeoffs across code, math, and instruction-following tasks.
 - My work was incorporated into frontier model training e.g., [Command A](#) model.
- 05/2023 - 09/2023 *Allen Institute for AI, Research intern at AllenNLP team*
 - Introduced *intrinsic source citation*, where LLMs are required to provide evidence for their responses.
 - Proposed a two-stage recipe for training LLMs to cite their pretraining data and support their parametric knowledge.
 - Inspired future work inside the OLMo team such as [OLMoTrace](#).
- 10/2020 - 07/2021 *Amazon AWS, Applied Scientist Intern*
 - Adapting LLMs to dialogue summarization: Aspects investigated included reasoning, handling informal language and multi-speaker dialogues.
 - Pretraining semi-structured language models for document image understanding; Trained a BERT-style model that was used in production.
 - Developed CLIP-style loss for zero-shot classification of visually rich documents.

- 05/2020 - 09/2020 *NAVER Labs Europe, Research Intern*
- Developed a principle energy-based alternative to LLM alignment under KL-penalties.
 - Proposed a general framework supporting instance-level and distributional constraints for text generation
- 04/2018 – 02/2020 *Sypron Solutions, Machine Learning Engineer*
- Developed anomaly detection system based on VAEs for predictive maintenance.
 - Built a pipeline for big data streaming of sensor data using Apache Spark and Kafka.

Awards

- 2022 Rackham Professional Development Grant.
- 2022 Rackham Fellowship for Outstanding International Students. Award amount is \$10,000.
- 2020 Best Presentation Award, Intern Day at NAVER Labs Europe, against 8 other presenters.
- 2019 1st Place at IDAT@FIRE 2019 shared task on *Irony Detection in Arabic Tweets*, against 18 systems.
- 2016 1st Place in the Journey to Mars Track in the NASA’s Space Apps Challenge in Cairo, Egypt.

Publications and Preprints

- Arxiv '25 *Process Reward Models That Think* [[paper](#), [code](#), [models](#), ~2k downloads/month]
Muhammad Khalifa, Rishabh Agarwal, Lajanugen Logeswaran et al.
- NeurIPS '25 *MLRC-Bench: Can Language Agents Solve Machine Learning Research Challenges?* [[paper](#)]
Yunxiang Zhang, Muhammad Khalifa, et al.
- Arxiv '24 *Optimizing Model Merging at Scale Mitigates Performance Tradeoffs* [[paper](#)]
Muhammad Khalifa, Yi-Chern Tan, Arash Ahmadian et al.
- NeurIPS '24 *Learning to Reason via Program Generation, Emulation, and Search* [[paper](#)]
Nathaniel Weir*, Muhammad Khalifa*, Linlu Qiu, Orion Weller, Peter Clark.
- ACL '24 Findings *Small Language Models Need Strong Verifiers to Self-Correct Reasoning* [[paper](#), [code](#)]
Yunxiang Zhang, Muhammad Khalifa, Lajanugen Logeswaran et al.
- COLM '24 *Source-Aware Training Enables Knowledge Attribution in Language Models* [[paper](#), [code](#)]
Muhammad Khalifa, David Wadden, Emma Strubell et al.
- ICLR '24 *Lightweight Calibration of LMs for Open-ended QA over Varied-Length Answers* [[paper](#), [code](#)]
Xin Liu, Muhammad Khalifa, Lu Wang.
- EMNLP '23 Findings *GRACE: Discriminator-Guided Chain-of-Thought Reasoning* [[paper](#), [code](#), [LLM Reasoners](#)]
Muhammad Khalifa, Lajanugen Logeswaran et al.
- EMNLP '23 *Merging Generated and Retrieved Knowledge for Open-domain QA* [[paper](#), [code](#)]
Yunxiang Zhang, Muhammad Khalifa, Lajanugen Logeswaran et al.
- ACL '23 Findings *Contrastive Training Improves Zero-Shot Classification of Semi-structured Documents* [[paper](#)]
Muhammad Khalifa, Yogarshi Vyas, Shuai Wang et al.
- ACL '23 *Few-shot Reranking for Multi-hop QA via Language Model Prompting* [[paper](#), [code](#)]
Muhammad Khalifa, Lajanugen Logeswaran et al.
- ACL '23 *BOLT: Fast Energy-based Controlled Text Generation with Tunable Biases* [[paper](#), [code](#)]
Xin Liu, Muhammad Khalifa, Lu Wang.
- Me-FoMo ICLR '23 *Exploring Demonstration Ensembling for In-context Learning* [[paper](#), [code](#)]
Muhammad Khalifa, Lajanugen Logeswaran et al.
- EMNLP '21 *A Bag of Tricks for Dialogue Summarization* [[paper](#)]
Muhammad Khalifa, Miguel Ballesteros, and Kathleen McKeown.

ICLR '21	<i>A Distributional Approach to Controlled Text Generation (Oral, top 2.1%)</i> [paper , code] <u>Muhammad Khalifa</u> *, Hady Elsahar*, and Marc Dymetman*.
EACL '21	<i>Self-Training Pretrained LMs for Zero-and Few-Shot Multi-Dialectal Arabic Sequence Labeling</i> [paper] <u>Muhammad Khalifa</u> , Muhammad Abdul-Mageed, and Khaled Shaalan.
CSL '19	<i>Character convolutions for Arabic Named Entity Recognition with LSTM Networks</i> [paper] <u>Muhammad Khalifa</u> and Khaled Shaalan.

Skills

Python, C++, PyTorch, DeepSpeed, Jax, Ray, VeRL, vLLM, SGLang, Apache Spark, Kafka, Optuna.

Invited Talks

2025	"Verify Reasoning with Reasoning: Process Reward Models That Think" - Tsinghua University, June 2025.
2025	"Optimized Model Merging for Recycling Generalist LLMs" - ML Collective (Deep Learning Classics and Trends), April 2025.
2024	"Discriminator-Guided Chain-of-Thought Reasoning" - ML Collective (Deep Learning Classics and Trends), March 2024.
2023	"Guided Language Models Reasoning" - NAVER Labs Europe, November 2023.
2021	"Controlling Stochastic Parrots" - reciTAL, February 2021.
2021	"A Distributional Approach to Controlled Text Generation" - ML Collective, March 2021 (With Hady Elsahar and Marc Dymetman).
2020	"Self-training Language Models for Arabic Sequence Labeling" - Virtual Workshop on Information Extraction from the Web using Machine Learning Techniques, November 2020.