

# Mohammad Atif Quamar

✉ [mohammad.atif@mbzuai.ac.ae](mailto:mohammad.atif@mbzuai.ac.ae)  
🌐 [atifquamar.com](http://atifquamar.com)  
🔗 [atifquamar07](https://github.com/atifquamar07)  
🌐 [atifquamar](https://www.linkedin.com/in/atifquamar)  
📄 [Google Scholar](https://scholar.google.com/citations?user=atifquamar)

Visiting Researcher at MBZUAI

## Research Interests

Reinforcement Learning, Efficient Training, Multi-modal Reasoning, Multi-Agentive Systems, AI Alignment and Safety, Natural Language Processing, Computer Vision

## Education



**M.Sc. in Machine Learning** 2026 – 2028

**Mohamed bin Zayed University of Artificial Intelligence (MBZUAI)**, Abu Dhabi, UAE

• Upcoming: Fall 2026



**B.Tech. in Computer Science and Biosciences** 2020 – 2024

**Indraprastha Institute of Information Technology (IIIT)**, Delhi, India

• GPA: 8.24/10.00

## Publications

\* *Equal contribution*

- [P1] **Adaptive Blockwise Search: Inference-Time Alignment for Large Language Models**  
**Mohammad Atif Quamar**<sup>\*</sup>, *Mohammad Areeb*<sup>\*</sup>, *Nishant Sharma*, *Ananth Shree Kumar*,  
*Jonathan Rosenthal*, *Muslum Ozgur Ozmen*, *Mikhail Kuznetsov* and *Z. Berkay Celik*  
arXiv, 2025. *Under Review*. [[arXiv](#)]
- [P2] **Reliable Chain-of-Thought via Prefix Consistency**  
*Naoto Iwase*, *Yuki Ichihara*, **Mohammad Atif Quamar** and *Junpei Komiyama*  
arXiv, 2026. *Under Review*. [[arXiv](#)]
- [P3] **Learning Modal-Mixed Chain-of-Thought Reasoning with Latent Embeddings**  
*Yifei Shao*, *Kun Zhou*, *Ziming Xu*, **Mohammad Atif Quamar**, *Shibo Hao*, *Zhen Wang*,  
*Zhiting Hu* and *Biwei Huang*  
arXiv, 2026. *Under Review*. [[arXiv](#)]
- [C1] **STARS: Synchronous Token Alignment for Robust Supervision in LLMs**  
**Mohammad Atif Quamar**<sup>\*</sup>, *Mohammad Areeb*<sup>\*</sup>, *Mikhail Kuznetsov*, *Muslum Ozgur Ozmen*,  
and *Z. Berkay Celik*  
Structured Probabilistic Inference & Generative Modeling Workshop at ICML 2026. [[arXiv](#)]
- [C2] **Logit–Entropy Adaptive Stopping Heuristic for Efficient Chain-of-Thought Reasoning**  
**Mohammad Atif Quamar** and *Mohammad Areeb*  
Efficient Reasoning Workshop at NeurIPS 2025. [[arXiv](#)] [[OpenReview](#)]
- [C3] **Decoding Histone Modification Signatures of Non-Coding RNAs via FMs**  
*Nishant Sharma*, **Mohammad Atif Quamar** and *Pengtao Xie*  
Workshop on Multi-modal Foundation Models and Large Language Models for Life Sciences  
at NeurIPS 2025. [[OpenReview](#)]

---

## Research Experience



### Visiting Researcher

Feb 2026 – Present

**MBZUAI**, Abu Dhabi, UAE

Advisor: **Prof. Junpei Komiyama**

- Working on methods to improve reasoning in language models through reinforcement learning, self-consistency, and test-time inference strategies, to enhance mathematical reasoning and coding performance while reducing unreliable reasoning paths. Publication - [P2].



### Research Intern

Sep 2025 – Feb 2026

**University of Virginia**, Virginia, United States

Advisor: **Prof. Ferdinando Fioretto**

- Developed a drift-resilient memory framework for code-execution agents using KL-constrained adapter updates, mitigating embedding distribution shifts during online learning to reduce unsafe code generation without compromising task success rates.



### Research Intern

Aug 2025 – Sep 2025

**UC San Diego**, California, United States

Advisor: **Prof. Biwei Huang**

- Proposed a multimodal Chain-of-Thought framework that interleaves text with latent visual representations in VLMs, with a two-stage SFT + RL training setup for modality switching, achieving strong gains across 11 multimodal reasoning benchmarks. Publication - [P3].



### Research Intern

Mar 2025 – Jul 2025

**Purdue University**, Indiana, United States

Advisor: **Prof. Z. Berkay Celik**

- Studied Bayesian sampling and inference-time alignment methods for language models, improving response quality and steering outputs toward harmlessness and positive sentiment. Publications - [C1] and [P1].



### Research Intern

Jan 2025 – Feb 2025

**UC San Diego**, California, United States

Advisor: **Prof. Pengtao Xie**

- Built a deep learning model predicting 50 histone modification–ncRNA interactions from transcript and genomic context. Publication - [C3].

---

## Industry Experience

### Co-Founder

Sep 2023 – Jan 2025

**Insituate**, New Delhi, India

- Founded Insituate and grew it to \$250K ARR, building on-premise AI agents using open-source LLMs and RAG systems for enterprise use cases; deployments included Singapore Judiciary, Indian High Courts, Mizuho Bank, and PNC Bank.

---

## Awards & Achievements

2026	<b>MBZUAI Scholarship.</b> 100% tuition scholarship for M.Sc. in ML; \$100K scholarship package.
2024	<b>NASSCOM GenAI Cohort.</b> Selected among the top 15 AI startups in India.
2024	<b>NVIDIA Inception.</b> Secured \$150K worth of AWS credits.
2024	<b>Microsoft Founder's Hub.</b> Received \$150K worth of Azure credits.
2020	<b>JEE Main.</b> Ranked in the top 3% among 1.1 million candidates.
Present	<b>Competitive Programming.</b> 1800+ rated on CodeChef and LeetCode; world top 3%.

---

## Academic Service

2025	<b>Program Committee.</b> AAI AI Alignment Track.
2025	<b>Program Committee.</b> Efficient Reasoning Workshop at NeurIPS.
2024	<b>Teaching Assistant.</b> ECO332: Valuation and Portfolio Management at IIITD.
2023	<b>Teaching Assistant.</b> DES102: Introduction to HCI at IIITD.