

Alireza Salemi

Manning College of Information and Computer Sciences, University of Massachusetts, Amherst, MA, USA

+1 (413) 270-4879 | asalemi@cs.umass.edu | alirezasalemi7.github.io | github.com/alirezasalemi7 | 

Education

University of Massachusetts Amherst

PhD in Computer Science

- Received PhD Candidacy with **Distinction**
- Received the **Google PhD Fellowship 2025** in Natural Language Processing for *Collaborative Learning in Agentic AI*
- GPA: 3.923/4.0

Massachusetts, USA

Sep. '22 - Current

University of Massachusetts Amherst

MSc in Computer Science

- Courses: Information Retrieval, Computer Vision, Adv. Natural Language Processing, Neural Networks, Information Assurance, Algorithms.
- GPA: 3.9/4.0

Massachusetts, USA

Sep. '22 - Sep. '24

University of Tehran

BSc in Computer Engineering

- **Ranked 1st** among 103 Computer Engineering students
- GPA: 19.60/20.0 (4.0/4.0)

Tehran, Iran

Sep. '17 - Feb. '22

Interests

RAG

Enhancing LLMs with external knowledge by optimizing and accelerating retrieval pipelines, evaluating retrieval quality, and developing dense and multi-modal retrieval methods for knowledge-intensive applications.

Personalization

Tailoring Large Language Models for individual users through reinforcement learning, self-training, and reasoning-enhanced methods for personalized long-form text generation and question answering.

Alignment

Exploring different alignment strategies such as pluralistic alignment techniques to ensure generative models can effectively and safely cater to a diverse range of user needs, preferences, and perspectives.

Agents

Designing single and multi-agent systems for complex problem-solving, including blackboard architectures for collaborative information discovery and leveraging agentic workflows to improve RAG.

Reasoning

Developing complex reasoning in LLMs for open-ended tasks by exploring multi-directional thinking and designing distributed thinking frameworks, such as multi-agent blackboard systems, for collaborative problem-solving.

Publications [Link]

LLM-Based Multi-Agent Blackboard System for Information Discovery in Data Science

Under Review in ICLR'26

Alireza Salemi, Mihir Parmar, Palash Goyal, Yiwen Song, Jinsung Yoon, Hamed Zamani, Tomas Pfister, Hamid Palangi

Pathways of Thoughts: Multi-Directional Thinking for Personalized Question Answering

Under Review in WWW'26

Alireza Salemi, Cheng Li, Mingyang Zhang, Qiaozhu Mei, Zhuowan Li, Spurthi Amba Hombaiah, Weize Kong, Tao Chen, Hamed Zamani, Michael Bendersky

Learning from Natural Language Feedback for Personalized Question Answering

Under Review in WSDM '26

Alireza Salemi, Hamed Zamani

CIIR@LiveRAG 2025: Optimizing Multi-Agent Retrieval Augmented Generation through Self-Training

LiveRAG @ SIGIR '25

Alireza Salemi, Mukta Maddipatla, Hamed Zamani

LaMP-QA: A Benchmark for Personalized Long-form Question Answering

EMNLP '25

Alireza Salemi, Hamed Zamani

Yiyan Xu, Jinghao Zhang, **Alireza Salemi**, Xinting Hu, Wenjie Wang, Fuli Feng, Hamed Zamani, Xiangnan He, Tat-Seng Chua

Plan-and-Refine in RAG: Generating Diverse and Comprehensive Responses through Global Exploration and Local Exploitation

Under Review in WSDM '26

Alireza Salemi, Chris Samarinis, Hamed Zamani

ExPerT: Effective and Explainable Evaluation of Personalized Long-Form Text Generation

ACL '25 Findings

Alireza Salemi, Julian Killingback, Hamed Zamani

Reasoning-Enhanced Self-Training for Long-Form Personalized Text Generation

Under Review in WSDM '26

Alireza Salemi, Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, Tao Chen, Zhuowan Li, Michael Bendersky, Hamed Zamani

Beyond Factual Accuracy: Evaluating Coverage of Diverse Factual Information in Long-form Text Generation

ACL '25 Findings

Chris Samarinis, Alexander Krubner, **Alireza Salemi**, Youngwoo Kim, Hamed Zamani

Accelerating Retrieval-Augmented Generation

ASPLOS '25

Derrick Quinn, Mohammad Nouri, Neel Patel, John Salihu, **Alireza Salemi**, Sukhan Lee, Hamed Zamani, Mohammad Alian

Learning to Rank for Multiple Retrieval-Augmented Models through Iterative Utility Maximization

ICTIR '25

Alireza Salemi, Hamed Zamani

Comparing Retrieval-Augmentation and Parameter-Efficient Fine-Tuning for Privacy-Preserving Personalization of Large Language Models

ICTIR '25

Alireza Salemi, Hamed Zamani

Retrieval-Enhanced Machine Learning: Synthesis and Opportunities

Under Review in IRRJ

To Eun Kim, **Alireza Salemi**, Andrew Drozdov, Fernando Diaz, Hamed Zamani

LongLaMP: A Benchmark for Personalized Long-form Text Generation

Under Review in TKDD

Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, **Alireza Salemi**, Ryan Rossi, Franck Dernoncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, Nedim Lipka, Hamed Zamani

LaMP: When Large Language Models Meet Personalization

ACL '24

Alireza Salemi, Sheshera Mysore, Michael Bendersky, Hamed Zamani

Optimization Methods for Personalizing Large Language Models through Retrieval Augmentation

SIGIR '24

Alireza Salemi, Surya Kallumadi, Hamed Zamani

Towards a Search Engine for Machines: Unified Ranking for Multiple Retrieval-Augmented Large Language Models

SIGIR '24

Alireza Salemi, Hamed Zamani

Evaluating Retrieval Quality in Retrieval-Augmented Generation

SIGIR '24

Alireza Salemi, Hamed Zamani

Best Short Paper

A Symmetric Dual Encoding Dense Retrieval Framework for Knowledge-Intensive Visual Question Answering

PEACH: Pre-Training Sequence-to-Sequence Multilingual Models for Translation with Semi-Supervised Pseudo-Parallel Document Generation

ARMAN: Pre-training with Semantically Selecting and Reordering of Sentences for Persian Abstractive Summarization

UTNLP at SemEval-2021 Task 5: A Comparative Analysis of Toxic Span Detection using Attention-based, Named Entity Recognition, and Ensemble Models

Research Experience

Student Researcher at Google Cloud

- Working on blackboard multi-agent systems for data science.

Student Researcher at Google DeepMind

- Worked on Personalizing Large Language Models using reinforcement learning, self-training, inference-compute scaling, reasoning, and pluralistic alignment.

Data Science Intern at Lowe's

- Worked on personalizing large language models by optimizing retrievers with reinforcement learning from language model's feedback.

Research Assistant at the Center for Intelligent Information Retrieval

- Working on Retrieval-Enhanced Machine Learning (REML), Retrieval-Augmented Generation (RAG), Multi-modal Knowledge-intensive Tasks, and Personalizing Large Language Models.

Research Assistant at Information Retrieval Laboratory

- Worked on designing a pre-training objectives for multi-lingual and summarization language models.

Research Assistant at Data Analysis Laboratory

- Worked on designing a system for annotating toxic (hate speech) spans of a piece of text.

Professional Services

WWW'26

PC Member for THE ACM WEB CONFERENCE 2026.

SIGIR-AP'25

PC Member for The 3'rd International ACM SIGIR Conference on Information Retrieval in the Asia Pacific.

NuerIP'25

Reviewer for The Thirty-Ninth Annual Conference on Neural Information Processing Systems.

EMNLP'25

Reviewer for The 2025 Conference on Empirical Methods in Natural Language Processing.

SIGIR'25

Organized a Tutorial on Retrieval-Enhanced Machine Learning: Synthesis and Opportunities in The 48th International ACM SIGIR Conference on Research and Development in Information Retrieval.

ICTIR'25

PC Member for The 11th ACM SIGIR / The 15th International Conference on Innovative Concepts and Theories in Information Retrieval.

ACL'25	Reviewer for The 63rd Annual Meeting of the Association for Computational Linguistics.
SIGIR'25	PC Member for The 48'th International ACM SIGIR Conference on Research and Development in Information Retrieval.
SIGIR-AP'24	Organized a Tutorial on Retrieval-Enhanced Machine Learning: Synthesis and Opportunities in The 2'nd International ACM SIGIR Conference on Information Retrieval in the Asia Pacific.
SIGIR-AP'24	PC Member for The 2'nd International ACM SIGIR Conference on Information Retrieval in the Asia Pacific.
WWW'25	PC Member for THE ACM WEB CONFERENCE 2025.
ACL'24	Reviewer for The 62nd Annual Meeting of the Association for Computational Linguistics.
SIGIR'24	PC Member for The 47'th International ACM SIGIR Conference on Research and Development in Information Retrieval.
LLM-IGS'24	PC Member for The 1st and 2nd Workshop on Large Language Models for Individuals, Groups, and Society.
IMP'24	Mentoring MS Students in Industry Mentorship Program at UMass Amherst working on Personalizing LLMs in collaboration with industry mentors from Adobe.
ERSP'23-'24	Mentoring Undergrad Students in Early Research Scholar Program at UMass Amherst working on a research project about Unified Question Answering and Personalizing LLMs.
TOIS	Reviewer for ACM Transactions on Information Systems.
TKDE	Reviewer for IEEE Transactions on Knowledge and Data Engineering,

Work Experience

Sooton

Tehran, Iran

Natural Language Processing Data Scientist

Nov. '21 - Aug. '22

- Worked on an AutoML system designed specifically for text classification tasks, with a focus on sentiment analysis of customer's reviews and automatically acceptance or rejection of users' comments on products.

Achievements

'25-'27	Google Ph.D. Fellowship 2025 , This fellowship is awarded to exceptional PhD students who demonstrate outstanding research potential in computer science and related fields.	UMass Amherst
'25	PhD Candidacy with Distinction , This title is given to a PhD candidate that shows an extraordinary performance in the comprehensive exam.	UMass Amherst
'25	W. Bruce Croft Scholarship 2025 , The scholarship is given to students whose research area is about information retrieval and natural language processing.	UMass Amherst
'24	The Best Short Paper Award @ SIGIR 2024 , The award is given to the best short paper in the SIGIR 2024 conference.	UMass Amherst
'23	The Best Student Paper Award @ SIGIR 2023 , The award is given to the best paper mainly written by a student.	UMass Amherst
'23	W. Bruce Croft Scholarship 2023 , The scholarship is given to students whose research area is about information retrieval and natural language processing.	UMass Amherst
'22	UMass Amherst CICS Scholarship , The scholarship is given to some selected entering graduate students based on their application's merit.	UMass Amherst
'21	Best Thesis Award in Computer Engineering , BSc thesis was selected as the best thesis in the computer engineering department.	University of Tehran
'18-'20	University of Tehran Sponsors Foundation Top Student , This award is given to the top student at the computer engineering department.	University of Tehran
'17-'20	University of Tehran Scholarship , This scholarship is given to exceptional talents from the UT Sponsors Foundation.	University of Tehran
'18-'19	F.O.E (Faculty of Engineering) Award , The award is given to the top student among all of 103 Computer Engineering students.	University of Tehran

Invited Talks

May '25	Personalization in the LLMs Era , Talk about the state of personalization in generative AI system.	Google DeepMind
Apr. '25	Personalization in the LLMs Era , Talk about the state of personalization in generative AI system.	UMass Amherst
Dec. '23	Knowledge-Intensive VQA , Talk about Knowledge-Intensive Visual Question Answering.	QIAU
Oct. '23	When Large Language Models Meet Personalization , Talk about LaMP benchmark and personalizing LLMs.	UMass ML Club

Teaching Experience

'20-'21	Artificial Intelligence , Head TA, Under supervision of Yadollah Yaghoobzadeh and Hakimeh Fadaei.	<i>University of Tehran</i>
'20-'21	Programming Languages and Compiler Design , Head TA, Under supervision of Fatemeh Ghassemi.	<i>University of Tehran</i>
'20-'21	The Theory of Formal Languages and Automata , TA, Under supervision of Hossein Hojjat.	<i>University of Tehran</i>
'20	Database Design , TA, Under supervision of Azadeh Shakery.	<i>University of Tehran</i>
'20	Design and Analysis of Algorithms , TA, Under supervision of Hamid Mahini.	<i>University of Tehran</i>
'19	Engineering Mathematics , TA, Under supervision of Mahdi Tale Masouleh.	<i>University of Tehran</i>

Languages

English	Fluent (TOEFL score: 113/120, Reading: 30, Listening: 28, Speaking: 26, Writing: 29)
Persian	Native proficiency