Kamyar Zeinalipour

Postdoctoral researcher, DIISM, University of Siena

kamyar.zeinalipour2@unisi.it kamyarzeinalipour.github.io linkedin.com/in/kamyarzeinalipour github.com/KamyarZeinalipour
huggingface.co/Kamyar-zeinalipour ORCID 0009-0006-3014-2511 Google Scholar

Research interests

Language models for languages with little data; generating and evaluating educational material with language models; instruction data and fine-tuning for open European models; evaluation that checks meaning, not only form.

Positions

Aug 2026 to now
Postdoctoral researcher, University of Siena, DIISM
Natural language processing for education and for languages with few resources, at SAILab with Marco Gori and Marco Maggini.

Nov 2025 to now
AI engineer, Villanova.AI, Cagliari
Open language models for Europe's languages, part of the IPCEI-CIS European AI programme. Built Villanova-SFT-2603, the 1.71-million-conversation instruction dataset for the first-generation Villanova models, worked on Villanova-2B-2603, and was part of the team behind Villanova2.

Nov 2025 to now
Generative AI engineer, YUKAI!, Siena
Fine-tuned models, retrieval pipelines and agentic systems for production, with a focus on educational and creative applications.

Apr to Jun 2025
Visiting researcher, University of Gothenburg
A multilingual study of plural nouns that have no singular, which became an EACL 2026 paper.

Feb to Apr 2025
Visiting researcher, Inria, Nice
Built the Detoxification Fidelity Index and a reinforcement learning pipeline for detoxification that preserves meaning.

Sep 2024 to Jan 2025
Visiting researcher, University of Massachusetts Amherst
Built a 22,000-pair human-verified dataset for detoxification that does not destroy the original meaning.

Apr 2022 to May 2023
Machine learning specialist, University of Siena
Graph neural networks for ECG classification: signals turned into graphs and classified with a Graph Isomorphism Network, published at WIRN 2022.

Jan 2022 to Mar 2023
Data scientist, with Baker Hughes, University of Siena
Emission forecasting for gas turbines, deployed as a service, published at WIRN and AI-ML-Systems.

Education

Nov 2022 to Feb 2026
Ph.D. in Information Engineering and Science, University of Siena
Thesis: From Language to Learning, an LLM-driven framework for multilingual educational content. Supervised by Marco Gori, reviewed by Hava Siegelmann and Hanan Salam, defended 10 February 2026 with the grade eccellente.

2017 to 2021
M.Sc. in Applied Mathematics (Data Science), University of Siena
Thesis on classifying ECG signals by turning time series into graphs, which became the ts2g package.

Sep 2013 to Jan 2017
B.Sc. in Applied Statistics, University of Hormozgan, Iran
Alongside it, statistician at Gilda Research and data analyst at the State Management and Planning Organization in Bandar Abbas.

Awards

Sep 2026
Best Fooling AI Agent, Turing Hotel Italy, CLiC-it 2026, Palermo
With team YUKAI!, for the AI agent that was most often mistaken for a human, especially in longer conversations. Awarded at the live Turing test of the Twelfth Italian Conference on Computational Linguistics, supported by Villanova.AI and AILC.

Publications (31, 19 as first author)

2026

31. **K. Zeinalipour**, A. Sadeghi, G. Angelini, L. Rigutini, M. Gori, M. Maggini. MultiHint-9K: Orchestrating Human and AI Agency for Scalable Pedagogical Hint Generation. *AIED 2026 workshops, CEUR 4248*.
30. **K. Zeinalipour**, A. Zugarini, A. Zanollo, L. Rigutini. PharmaQA.IT: an Italian dataset for Q&A in the pharmaceutical domain. *EACL 2026, industry track*.
29. M. Mehak, **K. Zeinalipour**, P. Soomro, C. Chesi, M. Gori, M. Maggini. Enabling Structured Reasoning in Sindhi with Culturally Grounded Instruction Tuning. *LoResLM 2026*.
28. **K. Zeinalipour**, S. Severini, A. Borghini, S. Cardarelli, M. Maggini, M. Gori. From Graph to Text and Back: Semantic Fidelity in Automated Industrial Knowledge Graphs. *ACL 2026, industry track*.
27. A. Bendjeddou, **K. Zeinalipour**, D. Bardou, M. Maggini, F. Scarselli, M. Bianchini. PharMistral: Drug generation for novel protein targets by Mistral Large Language Model. *Computer Methods and Programs in Biomedicine Update 10*.
26. N. Jamshidi, **K. Zeinalipour**, F. Akbari, M. Bianchini, M. Maggini, M. Gori. TalkFa: A Unified Benchmark for Farsi Dialogue Generation and Understanding. *Preprint, arXiv:2609.01810*.
25. D. Baiamonte, **K. Zeinalipour**, E. Fano, S. Simonazzi, A. Zugarini. Multi-SFT: Permissively-Licensed Multilingual Instruction-Following. *Preprint, arXiv:2610.22086*.
24. E. Spaziani, **K. Zeinalipour**, P. Cassotti, N. Tahmasebi. Elections go bananas: A First Large-scale Multilingual Study of Pluralia Tantum using LLMs. *EACL 2026*.
23. M. Nissim, D. Croce, V. Patti, P. Basile and the CALAMITA community. Challenging the Abilities of Large Language Models in Italian: a Community Initiative. *Italian Journal of Computational Linguistics*.
22. S. Fioravanti, M. Zavatteri, R. Confalonieri, **K. Zeinalipour**, P. Frazzetto and others. Iterative In-Context Learning to Enhance LLMs Abstract Reasoning: The Case-Study of Algebraic Tasks. *ACM SAC 2026*.
21. **K. Zeinalipour**, L. Barriere, D. Ghelardi, M. Gori. Application of Machine Learning Models for Carbon Monoxide and Nitrogen Oxides Emission Prediction in Gas Turbines. *Neural Networks: Overview of Current Theories and Applications, Springer*.

2025

20. **K. Zeinalipour**, N. Jamshidi, S. Hejazi, M. Maggini, M. Bianchini, S. Paoletti, M. Gori. FarSense: A Comprehensive Commonsense Benchmark and Evaluation Framework for the Farsi Language. *IJCNLP-AACL 2025*.
19. **K. Zeinalipour**, M. Z. Saad, O. Attafi, M. Maggini, M. Gori. Shawarma Chats: A Benchmark Dialogue and Evaluation Platter in Egyptian, Maghrebi and Modern Standard Arabic — A Triple-Dialect Feast for Hungry Language Models. *ArabicNLP 2025*.
18. S. Lamsiyah, **K. Zeinalipour**, S. El-amrany, M. Brust, M. Maggini, P. Bouvry and others. ArabicSense: A Benchmark for Evaluating Commonsense Reasoning in Arabic with Large Language Models. *WACL-4 2025*.
17. **K. Zeinalipour**, N. Jamshidi, F. Akbari, M. Maggini, M. Bianchini, M. Gori. PersianMCQ-Instruct: A Comprehensive Resource for Generating Multiple-Choice Questions in Persian. *LoResLM 2025*.
16. **K. Zeinalipour**, M. Z. Saad, M. Maggini, M. Gori. From Arabic Text to Puzzles: LLM-Driven Development of Arabic Educational Crosswords. *LoResLM 2025*.
15. T. Iaquinta, A. Zanollo, A. Fusco, **K. Zeinalipour**, C. Chesi. Surprisal and Crossword Clues Difficulty: Evaluating Linguistic Processing between LLMs and Humans. *CLiC-it 2025*.
14. **K. Zeinalipour**, T. Iaquinta, A. Zanollo, G. Angelini, L. Rigutini, M. Maggini, M. Gori. Italian Crossword Generator: An In-depth Linguistic Analysis in Educational Word Puzzles. *Italian Journal of Computational Linguistics 11*.
13. **K. Zeinalipour**, L. Barriere, D. Ghelardi, M. Gori. Adaptive NOx Emissions Forecasting for Gas Turbines via Dynamic PCA-Driven Clustering. *AI-ML-Systems 2025*.

2024

12. A. Zugarini, **K. Zeinalipour**, S. S. Kadali, M. Maggini, M. Gori, L. Rigutini. Clue-Instruct: Text-Based Clue Generation for Educational Crossword Puzzles. *LREC-COLING 2024*.
11. **K. Zeinalipour**, Y. G. Keptiğ, M. Maggini, M. Gori. Automating Turkish Educational Quiz Generation Using Large Language Models. *ISPR 2024*.
10. **K. Zeinalipour**, Y. G. Keptiğ, M. Maggini, L. Rigutini, M. Gori. A Turkish Educational Crossword Puzzle Generator. *AIED 2024*.
9. **K. Zeinalipour**, A. Fusco, A. Zanollo, M. Maggini, M. Gori. Harnessing LLMs for Educational Content-Driven Italian Crossword Generation. *CLiC-it 2024*.
8. **K. Zeinalipour**, N. Jamshidi, M. Bianchini, M. Maggini, M. Gori. Design Proteins Using Large Language Models: Enhancements and Comparative Analyses. *Language + Molecules @ ACL 2024*.
7. **K. Zeinalipour**, M. Mehak, F. Parsamotamed, M. Maggini, M. Gori. Advancing Student Writing Through Automated Syntax Feedback. *AIEER 2024*.
6. A. Zugarini, **K. Zeinalipour**, A. Fusco, A. Zanollo. ECWCA, Educational Crossword Clues Answering: A CALAMITA Challenge. *CLiC-it 2024*.

2023

5. **K. Zeinalipour**, M. Z. Saad, M. Maggini, M. Gori. ArabIcros: AI-Powered Arabic Crossword Puzzle Generation for Educational Applications. *ArabicNLP @ EMNLP 2023*.
4. **K. Zeinalipour**, T. Iaquinta, A. Zanollo, G. Angelini, L. Rigutini, M. Maggini, M. Gori. Italian Crossword Generator: Enhancing Education through Interactive Word Puzzles. *CLiC-it 2023*.
3. **K. Zeinalipour**, T. Iaquinta, G. Angelini, L. Rigutini, M. Maggini, M. Gori. Building Bridges of Knowledge: Innovating Education with Automated Crossword Generation. *ICMLA 2023*.
2. **K. Zeinalipour**, M. Gori. Graph Neural Networks for Topological Feature Extraction in ECG Classification. *WIRN 2022, published 2023*.
1. G. Angelini, M. Ernandes, T. Iaquinta, C. Stehlé, F. Simões, **K. Zeinalipour**, A. Zugarini, M. Gori. The WebCrow French Crossword Solver. *EAI INTETAIN 2023*.

Open research releases

Each with its code on GitHub and its models and data on Hugging Face.

- **FarSense**: Farsi commonsense benchmark, 6 fine-tuned adapters
- **Shawarma Chats**: Dialogue benchmark in three Arabic varieties
- **Sindhi Reasoning**: Culturally grounded instruction tuning for Sindhi
- **Arabic crosswords**: Arabic clue generation, annotation interface
- **TalkFa**: Farsi dialogue generation and understanding
- **ArabicSense**: Arabic commonsense benchmark and models
- **PersianMCQ-Instruct**: Persian multiple-choice question generation
- **Round-trip KG**: Text to knowledge graph and back, with fidelity scores
- **Turkish quizzes**: Quiz generation from Turkish educational text
- **Turkish crosswords**: Turkish clue generation models
- **Italian crosswords**: Italian clue generation across three papers
- **Essay syntax feedback**: Automated syntax feedback for student writing
- **MultiHint-9K**: Pedagogical hints, 9,000 of them, with five models
- **Protein design**: Protein sequence generation with language models
- **PharmaQA.IT**: Italian question answering over regulatory documents
- **Crossword Surprisal**: Clue difficulty predicted from token probabilities
- **PharMistral**: Code and models for the CMPB Update 2026 paper

Teaching and supervision

July 2026	**Lecturer, Large Language Models for Developers,** Zucchetti A 20-hour course in English for software developers: ten two-hour lectures from the transformer to production systems, through prompt engineering, retrieval-augmented generation, graph RAG, agents, evaluation, LoRA fine-tuning and multimodal models. Every lecture has a hands-on lab; slides, labs and solutions are public.
Oct 2022 to Feb 2026	**Teaching assistant, Machine Learning,** University of Siena Ran lab sessions and office hours for the Machine Learning course, built and revised exercises, marked assignments and gave feedback on student projects.
2023 to 2026	**Master's thesis supervision,** University of Siena Guided ten master's students through their theses, from framing a question to running the experiments and writing them up. Several of those theses became papers in the list above.

Service

since 2024	**Area chair, then reviewer,** ACL Rolling Review Area chair in the March cycle; a reviewer since then.
Sep 2026	**Co-organiser, Turing Hotel,** CLiC-it 2026, Palermo A live, multi-party Turing test in Italian, built on UNaIVERSE: people and artificial agents at one table, each trying to tell the other apart.
Oct 2025	**Publicity co-chair,** AIML Systems 2025, Bengaluru
2024 and 2025	**Program committee,** BEA, the workshop on NLP for educational applications
2024 and 2025	**Reviewer,** AAAI
2024	**Publicity and communication chair,** CoLLAs 2024, Pisa
2024	**Challenge co-proposer, ECWCA,** CALAMITA at CLiC-it 2024

Talks

Sep 2026	**Villanova-2B, poster with the Villanova.AI team,** CLiC-it 2026, Palermo
Jun 2026	**From Graph to Text and Back,** ACL 2026, industry track
Dec 2025	**FarSense,** IJCNLP-AACL 2025
Feb 2025	**Speaker,** LT4All 2025, UNESCO Headquarters, Paris
Nov 2023	**Applications of large language models in education,** SAILab seminar, University of Siena
Feb 2023	**Research directions, with Andrea Zugarini,** TAILOR WebCrow workshop, CINI and Expert.ai

Certificates

2021	**Mathematics for Machine Learning, specialization,** Imperial College London
2021	**Deep Learning, specialization,** DeepLearning.AI
2020	**Machine Learning,** Stanford University, Andrew Ng
2019	**Python for Everybody, specialization,** University of Michigan
2024	**Graph Neural Networks, Ph.D. course with exam,** University of Siena, Franco Scarselli

Skills and languages

Language models	fine-tuning (LoRA, QLoRA, PEFT), RLHF, instruction-data curation, retrieval-augmented generation and graph RAG, agents, evaluation
Frameworks	PyTorch, Hugging Face Transformers and PEFT, vLLM, DeepSpeed, LangChain, scikit-learn, XGBoost
Infrastructure	Linux, Docker, Git, CUDA, Slurm, Weights & Biases, Azure ML, Google Cloud
Programming	Python, SQL, R, Bash
Languages	Farsi (native), English (C1), Italian (B1)

References available on request. The current version of this CV is at kamyarzeinalipour.github.io/cv-full.html.