Moshe Mash, Ph.D.

AI Researcher | NLP & LLMs | Multi-Agent Systems | Human-Centric AI

moshikmash@gmail.com | +972-55-259-0991 | WhatsApp +1-412-626-1676 | moshemash.com | LinkedIn

SUMMARY

AI researcher with a Ph.D. in Software and Information Systems Engineering (Ben-Gurion University) and a postdoc at the Carnegie Mellon Robotics Institute. 8+ years taking AI from research to production: fine-tuned LLMs, text embeddings and deep-learning models, predictive ML systems in industry (Amdocs, Diagnostic Robotics), and multi-agent and human–agent decision-making research published at top venues (Best Paper Award, ACM EC; JACM; IJCAI; CHI).

PROFESSIONAL EXPERIENCE

AI Researcher (NLP & LLM research lead) | DICTA – The Israel Center for Text Analysis *2022–present*

- Lead the center's NLP and LLM research on large Hebrew and rabbinic text collections: set the research agenda, design the methods, and take systems from prototype to tools in active use.
- Fine-tuned LLMs and custom text embeddings trained with contrastive learning, powering information extraction, semantic search, similarity, clustering and question answering over long, complex documents.
- Built the end-to-end training and inference pipeline, with tooling to analyze model outcomes, and deployed document-understanding systems to production.
- Authorship attribution models that resolve long-debated questions of disputed authorship; citation detection and cross-referencing across a large corpus.

Principal AI Researcher & Founder | MashInnovateAI *2024–present*

- Independent AI R&D consultancy: identify where AI can solve a company's core challenge, then design and deliver the solution end to end, from proof of concept to production, for tech companies, startups and academic groups.

Postdoctoral Researcher | Carnegie Mellon University, Robotics Institute, School of Computer Science *2019–2022*

- With Prof. Reid Simmons and Dr. Stephanie Rosenthal. Developed DSWorkFlow, a framework that captures data scientists' workflows from Jupyter notebooks (CHI 2021).
- Built a model that predicts when ML developers get stuck from real-time signals, and an algorithm that times interventions to its onset (VL/HCC 2022). Funded by J.P. Morgan, renewed for a second year based on results.

AI Researcher | Diagnostic Robotics *2019*

- ML models for healthcare risk prediction (ER diagnoses, hospitalization forecasting), developed with clinicians.

AI Researcher | Amdocs *2018–2019*

- Designed, built and deployed from scratch a churn-prediction system for maintenance-contract renewals, saving millions of dollars; presented results to senior leadership.

EDUCATION

Ph.D., Software and Information Systems Engineering | Ben-Gurion University *2014–2018*

- Advisor: Prof. Kobi Gal. Fair agreements between people and computer agents: polynomial-time algorithms for envy-free, fair rent division, deployed on Spliddit since 2015 (Best Paper Award, ACM EC 2016; JACM 2017; CACM Research Highlight 2018); ML + cooperative game theory to predict coalition formation, and a negotiating agent that outperformed human players (IJCAI 2017; ACM TIST 2020).

M.Sc., Computer Science (AI) | Bar-Ilan University *2011–2013*

- Advisor: Prof. David Sarne. Joint exploration with self-interested agents; equilibrium analysis of multi-agent search (AIJ 2014; IAT 2012).

B.Sc., Software Engineering | Shenkar College of Engineering and Design *2007–2011*

AWARDS

- Best Paper Award, ACM Conference on Economics and Computation (EC 2016); Communications of the ACM Research Highlight (2018).
- Excellence in Teaching Award, Ben-Gurion University (2016, 2017); Ph.D. Scholarship for Academic Excellence (2014).

SELECTED PUBLICATIONS

- Mash, Oryol, Rosenthal & Simmons. Predicting Data Scientist Stuckness During the Development of ML Classifiers. VL/HCC 2022.
- Mash, Rosenthal & Simmons. DSWorkFlow: A Framework for Capturing Data Scientists' Workflows. CHI 2021.
- Mash, Fairstein, Bachrach, Gal & Zick. Human-Computer Coalition Formation in Weighted Voting Games. ACM TIST 2020.
- Gal, Mash, Procaccia & Zick. Which Is the Fairest (Rent Division) of Them All? JACM 2017; EC 2016 (Best Paper).
- Mash, Bachrach, Gal & Zick. How to Form Winning Coalitions in Mixed Human-Computer Settings. IJCAI 2017.
- Rochlin, Sarne & Mash. Joint Search with Self-Interested Agents and the Failure of Cooperation Enhancers. AIJ 2014.
- Also: Science Translational Medicine (2025); Northwestern J. of Technology and IP (2026); AAMAS 2014. Full list: Google Scholar.

EXPERTISE & SERVICE

- **AI/ML:** LLM fine-tuning, embeddings & contrastive learning, NLP, deep learning, predictive modeling, multi-agent systems, game theory, human–AI interaction, user studies, end-to-end training/inference pipelines and production deployment.
- **Teaching & service:** Lecturer and TA in ML, game theory, databases and programming at four Israeli universities; mentored M.Sc. students; reviewer for AAAI, IJCAI, AAMAS, KDD, CHI, EC, ECAI.