# Barış Deniz Sağlam

**PhD Researcher · continual learning, agents, reasoning**

Ankara, Türkiye · bdsaglam@gmail.com · +90 534 220 11 91 · barisdeniz.is-a.dev · github.com/bdsaglam · Google Scholar

I am a PhD researcher at METU working on continual learning, in particular language models that learn from experience at test time. My other interests are reinforcement learning, agents and reasoning. My open 27B document agent was a joint winner of the ICDAR 2026 DocVQA challenge, and I co-authored ArgGYM, a defeasible-reasoning benchmark. Before the PhD I worked as an engineer for 10+ years, most recently building LLM agents and their evaluation for production. I am open to research roles from internships to permanent positions, and fully available from 2027.

## RESEARCH

### Language of Learning · PhD research programme

2025 – present

In our position paper, we argue that continual learning from experience should be a capability of the model itself: learning signals should reach the model through language, the model should control its own learning through language, and what it learns should persist in its own memory, from context to weights.

We place this in a recurring pattern in AI, where hand-designed procedures become learned model capabilities, as search strategies moved into reasoning tokens, and propose learning tokens: generated language in which a model works out what an experience should change and directs updates to its own weights during inference, meta-trained so that lessons transfer to later, unseen situations.

### ArgGYM: a benchmark for defeasible reasoning · third author

2026

Our benchmark tests reasoning where new information can revise a conclusion; a symbolic engine verifies answers, and procedural generation gives fresh test sets and verifiable rewards for RL. I proposed turning the work into a benchmark, designed and built its evaluation scaffold, ran the evaluation experiments, and am running its RL training experiments.

### Perceive-Reason-Code: active perception for document VQA

Mar – Jun 2026

I built an agent in which a coding model in a persistent Python REPL directs a vision model to any region of any page, so it decides where to look instead of reading whole pages at a fixed resolution. When I started, code-executing agents were rarely used for multimodal tasks.

My open 27B agent was a joint winner of the 8–35B category of the ICDAR 2026 DocVQA challenge. On the held-out test set it beat the challenge's bare-model baselines, including Gemini 3 Pro and GPT-5.2.

My ablations show the gain needs both the REPL and the perception tool, while sub-agents, context compaction, OCR and search add nothing or cost accuracy. Scaling the reasoner moves accuracy about twice as much as scaling the vision model · blog · code

### Earlier projects

2023 – 2026

**epiq** (2026): we built an ARC-AGI agent that reasons by running code in a Python REPL and checks its program on a held-out training example · code

**RL for agentic multi-hop QA** (2025): I trained Llama-3.1-8B with GRPO and Dr. GRPO to plan retrieval across hops on MuSiQue, concurrently with Search-R1 · code

**Fine-tuning vs prompting** (2025): I compared LoRA fine-tuning with DSPy prompt optimisation for entity-relation extraction and found that string-match metrics and LLM-judge ratings rank the two differently · dataset · blog

**MSc thesis** (2024): I proposed Connect-the-Entities prompting, where a small LLM extracts entity–relation triplets before answering; it beat standard and chain-of-thought prompting on MuSiQue with fewer tokens. I also explored RL from task feedback with PPO for triplet extraction.

## PUBLICATIONS AND AWARDS

**Barış Deniz Sağlam**, Duygu Ataman, Altan Koçyiğit. *Language of Learning.* Position paper, SSRN preprint, 2026. doi:10.2139/ssrn.7552418

İbrahim Ethem Deveci, Funda Tan Çalık, **Barış Deniz Sağlam**, Duygu Ataman. *ArgGYM: A Procedural, Engine-Verified Benchmark for Structured Defeasible Reasoning.* Under review at ICLR 2027. arXiv:2609.38409

**Barış Deniz Sağlam**. *Knowledge Graph Augmented Multi-Hop Question Answering Using Large Language Models.* MSc thesis, METU, 2024. open.metu.edu.tr

ICDAR 2026 DocVQA challenge, joint winner of the 8–35B category, with Altan Koçyiğit as advisor. announcement

## EDUCATION

**PhD, Information Systems**, Middle East Technical University (METU) · 2025 – present · advisor Altan Koçyiğit · coursework completes by 2027

**MSc, Data Informatics**, METU · 2021 – 2024 · CGPA 3.93/4.00 · all AA in deep learning, reinforcement learning, sequence models, ML systems and GPU programming courses

**BSc, Mechanical Engineering**, METU · 2006 – 2012

## INDUSTRY EXPERIENCE

### Staff AI Engineer · Independent Contractor

Mar 2025 – present

At Bosch Legends Lab, built a neuro-symbolic research agent that finds and assesses experts across the web in runs of hours, which outperforms other deep-research tools by a large margin on our benchmarks, and set the architecture of its multi-agent platform. At Novee AI, built LLM agents for insurance documents with evaluation from day one, including a contract-comparison agent that caught material changes underwriters missed.

### Senior AI Software Engineer · Peppy Health

Mar 2023 – Feb 2025

Led the AI of an employee-benefits platform: an assistant agent, hybrid search and recommendations, each algorithm chosen by offline evaluation against expert labels. Launched a health chatbot with clinician evaluation and safety guardrails.

### Senior Software Engineer · Tractable

Aug 2021 – Feb 2023

Turned model predictions into cost-optimal repair decisions to flag overcharging, and built the human-in-the-loop platform that routes uncertain model decisions to expert reviewers.

### Software Development Engineer · Amazon Web Services

Oct 2020 – Jun 2021

High-scale, low-latency pricing and offer services in AWS Billing.

### Founding Engineer (ML and Mobile) · Mediate / Supersense

Nov 2018 – Mar 2020

Built on-device document-detection and intent-recognition models for an assistive app for blind users, end to end from data collection and labeling to training and deployment, with a synthetic-data pipeline for hard cases.

### Earlier

2013 – 2018

Co-founder and software engineer, Comind, speech-based English learning for children (2017–2018). R&D engineer in structural dynamics, Turkish Aerospace Industries, where I built flight-test modal analysis software (2013–2017).

## TEACHING AND OPEN SOURCE

Instructor, *Programming and Scientific Computations with Python*, two 8-week terms for engineers, UCTEA Chamber of Mechanical Engineers, Ankara, 2018 · materials

**fastmtl**, multi-task learning for fastai · code

**torch-scae**, PyTorch port of Stacked Capsule Autoencoders · code

## SKILLS

| | |
|---|---|
| **Research** | RL post-training (GRPO, Dr. GRPO), SFT and LoRA, prompt optimisation, agent design and ablations, benchmark and evaluation design, LLM-as-judge |
| **Tools** | Python, PyTorch, Hugging Face Transformers, TRL, verl, vLLM, DSPy, Weights & Biases |
| **Languages** | Turkish (native), English (professional) |