Yasemin Ozkut

MS Candidate | Computer Vision Engineer Intern | AI/ML New Grad (May 2026)

📍 Columbus, OH   ✉ ozkutyasemin@gmail.com   📞 +1 (614) 702-0362   🔗 yaseminozkut-portfolio.netlify.app/

in linkedin.com/in/yaseminozkut   ⌂ github.com/yaseminozkut   huggingface.co/yozkut

Education

MS	**The Ohio State University (OSU)**, Electrical and Computer Engineering	Aug 2024 – May 2026
	• **Focus:** Computer Vision, Machine Learning, Deep Learning	Columbus, OH, USA
BS	**Sabanci University**, Computer Science and Engineering	Sept 2019 – Feb 2024
	• **Minor:** Business Analytics, School of Management	Istanbul, Turkey

Skills

Languages: Python (3+ years hands-on, 1+ year professional), C++, SQL, JavaScript

Frameworks & Libraries: PyTorch, PyTorch Lightning, TensorRT, ONNX, YOLO, OpenCV, Unsloth, Hugging Face, LangChain, LangGraph, Pandas, NumPy, Scikit-learn, React, Hydra, Flask

Tools: Git, GitHub Actions (CI/CD), CUDA, SLURM, Linux, Claude Code, Cursor, Reachy Mini SDK, Docker, PostgreSQL

Experience

Ubihere | Computer Vision Engineer Intern (Full-Time) – Columbus, OH   Aug 2025 – Present
- Wrapped research-stage person re-identification modules into a GPU-accelerated Flask microservice with REST API, containerized with Docker and NVIDIA Container Toolkit; now integrated into core product.
- Reduced inference latency by 67% through GPU batching and parallelized image I/O.
- Currently developing an AI-powered admin dashboard with React that queries PostgreSQL for customer camera environments and delivers spatially aware tracking insights integrated with floor-plan data.

PCVLab, OSU | Graduate Research Assistant in ML – Columbus, OH   Aug 2024 – Aug 2025
- Benchmarked 8 CNN/Transformer architectures for ocular ultrasound classification across 5 binary tasks (40+ models) with selective top-k temporal pooling, using Lightning-Hydra and MONAI frameworks.
- Developed two-stage diagnostic pipeline for retinal detachment (97.4% accuracy) and macular status (88.2% accuracy) classification; published ERDES dataset and paper (10k+ Hugging Face downloads). ⌂
- Fine-tuned LLaMA-3.2-11B-Vision via 4-bit QLoRA (Unsloth/TRL) on curated 100K+ MIMIC-CXR chest X-rays for radiology report generation, on HPC (SLURM), reducing loss by 79%, and achieving BLEU-1: 28.1, ROUGE-1: 39.5. ⌂

DAI-Labor | AI/ML Researcher Intern – Berlin, Germany ⌂   July 2023 – Sept 2023
- Trained YOLOv8 on custom German supermarket products (87.2% accuracy) and built a real-time detection pipeline.
- Integrated Keras-OCR for German-language labels and benchmarked against YOLOv8 detection.

First-Author Publications

ERDES: A Benchmark Video Dataset for Retinal Detachment and Macular Status Classification in Ocular Ultrasound, *Nature Scientific Data* 🔗 Website | Hugging Face	In revisions July 2025
HOPE: A Benchmark Dataset for Pulmonary Embolism Classification in Spatiotemporal Heart Ultrasound, *Nature Scientific Data*	Under preparation Sept 2025 - Present

Projects and Open Source Contributions

Judgy Reachy, Reachy Mini Robot, Personal Project ⌂ \| Hugging Face   Jan 2026
- **NVIDIA GTC 2026 Golden Ticket Contest** – Top 5 out of 75 entries in Hugging Face x Pollen Robotics track.
- Built a real-time phone detection pipeline that delivers funny robot feedback using YOLO26m with TensorRT (2.69x speedup), Llama 3.1-8B (Groq API) with 9 personality profiles, ElevenLabs TTS, and expressive robot motion.
- Published as an open-source Hugging Face Space with CI/CD via GitHub Actions, supporting simulation, physical robot, and standalone WebGPU browser demo (ONNX Runtime + Transformers.js).

Reachy Mini SDK – Contributor ⌂ \| *Pollen Robotics (900+ stars)*   Jan 2026 – Present
- Fixed community app install bug by adding metadata persistence and Python package entry point matching (PR merged in v1.2.13). Multiple additional bug reports and open PRs.

Reachy Mini Conversation App – Contributor ⚙ | *Pollen Robotics (140+ stars)*   Jan 2026 – Present
  - Fixed missing auto-enable of Gradio UI flag in simulation mode, enabling the app to run without a physical robot (PR merged). Contributing feature enhancement for dynamic dance description loading (open PR).

Agentic Person Re-ID via LangGraph and VLM Reasoning, Personal Project ⚙   July 2025
  - Built a LangGraph pipeline integrating YOLOv8, Qwen2.5-VL-3B, and Qwen2.5-7B for person detection, description generation, and identity assignment across frames/videos.
  - Designed structured JSON prompts to extract rich semantic attributes (hair, clothing, accessories, face shape) and implemented persistent memory for global ID tracking.
  - Enabled explainable matching with logged reasoning, confidence scores, and memory updates for re-identification.

Serious Game for Children with Cerebral Palsy , Sabanci University ⚙   Feb 2023 – Feb 2024
  - Developed a serious game in Unity using C# for balance physiotherapy sessions in children with Cerebral Palsy.
  - Collected live data from Wii Balance Board for each game session and stored it into MongoDB.
  - Generated visual reports from the processed data using Python and Flask.

Spotify Music Artist Success Collaboration Network, Sabanci University ⚙   Feb 2023 – June 2023
  - Collected and preprocessed Spotify and Kaggle data to build a collaboration network of 8.6K artists and 13.3K connections using Python/NetworkX, incorporating genre, popularity, and centrality metrics.
  - Applied clustering and centrality algorithms to identify key influencers, and visualized the network with Gephi for success analysis.

Activities & Interests

Hiphop Dancer (2015–Present); **Board Member**, SuDance Club, Sabancı University (2019–2023); **Contestant**, Peak Games Unithon (2022); **Tutor**, Civic Involvement Project – 4th-grade students (2020); **Winter Sports Enthusiast** (Skiing, Snowboarding, Ice Skating)