

Oktay Ozturk

Location: Phoenix, Arizona | P: 316-553-9272 | E: oktay.ozturk010@gmail.com | LinkedIn: linkedin.com/in/ozturkoktay

SUMMARY

Lead Machine Learning Engineer with proven experience designing, building, and deploying production-grade ML systems across computer vision, NLP, and generative AI. Delivered \$1M+ in revenue impact through ML-driven products, reduced model security vulnerabilities by 60%, and improved user engagement by up to 35%. Strong background in end-to-end ML systems, including model development, deployment, MLOps, and scalable data pipelines. Experienced in bridging research and real-world systems, with publications submitted to top-tier AI conferences and a track record of leading high-impact technical initiatives.

TECHNICAL SKILLS

Programming & Core ML: Python, PyTorch, TensorFlow, NumPy, Scikit-learn

Machine Learning: Computer Vision, NLP, Transformer Models, LLMs, Generative Models, A/B Testing

ML Systems & MLOps: Model Deployment, Training Pipelines, Model Optimization, Edge ML, Jenkins CI/CD

Data & Cloud Infrastructure: SQL, Apache Airflow, Spark, Docker, Redis, Elasticsearch, Cloud Platforms (AWS, Azure)

APIs & Tools: FastAPI, REST APIs, Microservices, Git/GitHub, Shell Scripting, Streamlit

EXPERIENCE

Lead Machine Learning Engineer

Jun. 2025 – Present

AcSYS Automation

Gilbert, Arizona

- Robotic Perception Systems:** Led development of an industrial vision pipeline integrating Basler cameras and FANUC robots, achieving 97% detection accuracy and reducing manual sorting effort by 90% across production lines.
- Generative AI Platforms:** Architected the AcSYS LLMs REST API for intelligent document processing, automating document workflows and reducing data retrieval time by 60% for engineering and operations teams.
- Motion Planning & Digital Twins:** Designed a digital twin system with IK solvers and collision-aware path planning, cutting robot commissioning and validation time by 25% while improving deployment safety.
- Edge MLOps Leadership:** Built an edge AI training and deployment framework that standardized model lifecycle management and reduced new SKU onboarding time by 80% (weeks to hours).

Senior AI Researcher

Feb. 2022 – Dec. 2024

University of Tennessee

Knoxville, TN, USA

- Research Excellence:** Led cutting-edge AI research on deep learning security, authoring a paper submitted to CVPR with potential to impact AI professionals, strengthening brand authority in AI security.
- Innovation & Implementation:** Designed and implemented a novel defense mechanism using PyTorch and diffusion autoencoders, securing deep learning models against data exploitation while processing 100,000+ data points, significantly reducing business risks tied to adversarial threats.
- Security Enhancement:** Reduced model vulnerability by 60%, fortifying AI systems against adversarial attacks and unauthorized data extraction, enhancing the security and reliability of enterprise-level AI solutions.
- Leadership:** Provided strategic leadership by mentoring a team of 5 junior researchers, optimizing code efficiency, and ensuring 100% on-time delivery of high-impact research, reinforcing operational excellence and development.

Senior Machine Learning Engineer

Oct. 2020 – Dec. 2021

ZackAI

Remote

- Model Optimization:** Engineered intent classification models (LSTM, CNN, TensorFlow), achieving a 25% drop in misclassifications, significantly enhancing chatbot accuracy, boosting customer satisfaction, and improving operational efficiency. Served the developed deep learning models using microservice architecture.
- Revenue Growth:** Designed and deployed a Transformer-based recommendation system (BERT4Rec), increasing chatbot engagement by 35%, lifting conversions by 20%, and driving \$1M+ in direct revenue impact.
- Customer Experience:** Delivered a highly accurate (93%) sentiment analysis model (BERT, PyTorch), analyzing 50K+ customer interactions, reducing complaint calls by 20%, and unlocking \$100K in additional revenue through streamlined customer support.
- Operational Efficiency:** Automated customer complaint data collection from 10+ sources via Python web scraping (Selenium, BeautifulSoup), reducing manual effort by 40%, accelerating resolution times by 20%, and significantly cutting operational costs.

- **Revenue Optimization:** Engineered an SQL-driven recommendation system optimizing customer SMS and data usage, boosting engagement by 35%, conversions by 20%, and contributing \$1M+ annually in additional revenue.
- **Database Excellence:** Led optimization and maintenance of a critical Oracle SQL database (3M+ records), achieving 99.9% uptime and automating manual tasks, saving 1,000+ labor hours and approximately \$250K annually.
- **Data Engineering:** Designed and deployed a high-performance ETL (Airflow, Python), enhancing data processing efficiency by 30%, enabling daily management of 1TB+ data, and reducing infrastructure costs by \$200K per year.
- **DevOps Efficiency:** Streamlined CI/CD workflows using Git, Jenkins, and Docker, reducing deployment time by 20% and minimizing downtime by 30%, driving higher operational efficiency and saving approximately \$300K annually.

PROJECTS

Insurance LLM Framework | *Python, PyTorch, HuggingFace, Streamlit* github.com/ozturkotay/insurance-llm-framework

- Developed a comprehensive LLM-powered prompt engineering and evaluation framework for insurance applications, including policy summarization, claim drafting, risk assessment, and regulatory compliance.
- Implemented an evaluation dashboard with automated metrics (ROUGE, BLEU, BERTScore), enabling detailed quality assessments and comparative performance analysis across multiple open-source models.
- Optimized model inference for both CPU and GPU environments, incorporating quantization methods and achieving performance improvements of up to 10x in inference speed for resource-constrained deployments.
- Designed a customizable prompt engineering library supporting advanced prompting strategies (Zero-shot, Few-shot, Chain-of-Thought), improving accuracy and quality of generated insurance-related content.

Image Segmentation API | *Python, OpenAI, Flask, Firebase, Docker* github.com/ozturkotay/image-segmentation-api

- Built endpoints leveraging OpenAI's LLMs for extracting structured information and generating summaries from complex technical documents.
- Developed a Flask-based API integrating AI models (YOLO, UNet, Roboflow) to analyze and segment architectural floor plans from PDFs.
- Implemented real-time PDF document handling with storage, retrieval, and Base64 encoding using Firebase.
- Created modular services for automated legend mapping, room segmentation, and electrical symbol detection, enhancing accuracy in architectural plan analysis.

EDUCATION

University of Tennessee Knoxville, TN

M.Sc., Computer Engineering

Related Coursework: *Machine Learning, Natural Language Processing, Data Mining and Analytics, Artificial Intelligence for Robotics, Computer Vision, Reinforcement Learning*

SELECTED RESEARCH PAPERS

- [1] A. S. Sener, I. F. Ince, H. B. Baydargil, I. Garip, and O. Ozturk, “Deep learning based automatic vertical height adjustment of incorrectly fastened seat belts for driver and passenger safety in fleet vehicles,” *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 2021. DOI: [10.1177/09544070211025338](https://doi.org/10.1177/09544070211025338). eprint: <https://doi.org/10.1177/09544070211025338>.
- [2] O. Ozturk and A. Ozcan, “Ideology detection using transformer-based machine learning models,” pp. 30–52, 2021. DOI: [10.13140/RG.2.2.12303.51362](https://doi.org/10.13140/RG.2.2.12303.51362).
- [3] M. Shoaeinaeini, O. Ozturk, and D. Gupta, “Twitter-informed prediction for urban traffic flow using machine learning,” in *2022 6th International Conference on Universal Village (UV)*, 2023. DOI: [10.1109/UV56588.2022.10185516](https://doi.org/10.1109/UV56588.2022.10185516).
- [4] O. Ozturk and A. Ozcan, “Sentiment analysis in turkish using transformer-based deep learning models,” in *4th International Conference on Artificial Intelligence and Applied Mathematics in Engineering*, Cham: Springer International Publishing, 2023, pp. 1–15, ISBN: 978-3-031-31956-3.