

Jose Sanchez Gonzalez

AI / Machine Learning Engineer

(503) 660-9763 | joseszgz2021@gmail.com | github.com/joses2017smjh
linkedin.com/in/jose-j-sanchez-gonzalez-84a800257/ | jose-sanchez-portfolio-com.vercel.app

Summary

AI/ML engineer with an M.S. in Artificial Intelligence. Builds computer vision, reinforcement learning, and retrieval systems in Python and PyTorch, with reproducible evaluation, inference services, and documented failure analysis.

Technical Skills

Languages: Python, C++, C, SQL

Machine Learning: PyTorch, scikit-learn, XGBoost, computer vision, Transformers, reinforcement learning

Retrieval and Agents: Retrieval-Augmented Generation (RAG), BM25, pgvector, LangGraph, MCP

Systems: FastAPI, Docker, ONNX Runtime, PostgreSQL, Linux, Git, Slurm, Isaac Sim, MuJoCo

Projects

SPUR: Metric Depth for Robotic Pruning | PyTorch, DINOv2, FastAPI

- Built Blender synthetic-data pipelines with RGB, depth, segmentation masks, and camera calibration; evaluated monocular and multi-view models under a fixed 10-tree, 5-seed validation protocol.
- Reduced synthetic validation depth RMSE from 0.0550 m to 0.0445 m (19%) using DINOv2 RGB-D refinement across three stereo pairs; reported seed variation of 0.0057 m.
- Packaged inference in a Docker/FastAPI service with split ONNX export; benchmarked the PyTorch refiner at 156 ms p50 on a V100 (FP16, six precomputed-depth views).

Humanoid Robustness Ladder | Isaac Lab, MuJoCo, PPO, ONNX

- Built a headless simulation-to-simulation evaluation harness using the upstream deployment controller; automated policy export, domain-randomization sweeps, and MuJoCo rollout metrics.
- Evaluated five randomization settings over 450 MuJoCo episodes; measured 23.3% falls without randomization versus 0% for the default setting across 90 episodes per setting.

MetaNaviT: Research File Retrieval | Python, PostgreSQL, pgvector, MCP

- Implemented hybrid BM25/dense retrieval, reciprocal rank fusion, staleness filtering, and typed graph traversal; added query routing and approval-gated file operations through Model Context Protocol (MCP).
- Measured Recall@50 of 0.938 and nDCG@10 of 0.493 on 136 labeled queries over a 61-file fixture corpus, using hash embeddings and overlap reranking.

Agentic Match Forecasting | XGBoost, LangGraph, FastAPI

- Built a forecasting prototype with XGBoost, isotonic calibration, and conformal prediction sets; connected three MCP servers to a LangGraph orchestrator with explicit approval for staked recommendations.
- Reported 62.7% accuracy versus a 47.1% prior baseline and 0.9013 log loss on 102 World Cup 2026 matches in a historical evaluation with pre-tournament training.

Experience

Graduate Teaching Assistant, Oregon State University | Sep 2025 – Jun 2026

- Led 2-hour studio sessions for 25+ students in Operating Systems I (C++) and Computational Methods (COMSOL); supported debugging, concurrency, memory management, and numerical simulation.

Math and Computer Science Tutor, Oregon State University | Nov 2022 – Jun 2025

- Tutored 50+ students in mathematics, programming, and computer science; adapted explanations and debugging guidance to individual learning needs.

Education

Oregon State University, Corvallis, OR

Master of Science in Artificial Intelligence | Jun 2026

Bachelor of Science in Computer Science (Applied: Data Science); Mathematics Minor | Jun 2025

Additional Experience

Berkeley Humanoid Lite hardware build (in progress): Soldered 22 magnetic encoders to motor controllers, 3D-printed structural parts, and began actuator calibration and limb assembly.