Ryan Witzman

Westfield, IN | ryan@rwitz.com | rwitz.com
linkedin.com/in/ryanwitzman | huggingface.co/rwitz | github.com/ryanwitzman

EDUCATION

Indiana University Indianapolis, Luddy School of Informatics, Computing, and Engineering — Indianapolis, IN
B.S. in Artificial Intelligence, Minor in Mathematics · **GPA 4.00 / 4.00** — Expected Dec 2027

EXPERIENCE

Data Management Intern — Jan 2026 – Present
Indiana Center for Musculoskeletal Health, IU School of Medicine — *Indianapolis, IN*

- Built an automated Snakemake pipeline that runs DeepLabCut SuperAnimal pose estimation as Slurm GPU jobs on IU's Quartz HPC cluster, taking open-field mouse behavior videos from SharePoint ingest to filtered keypoint data and QC reports; **481 videos processed**.
- Validated half-resolution preprocessing against full-resolution output before adopting it; added duplicate detection, staging validation, a pytest suite and HPC documentation so lab members can run the pipeline independently.
- Trained an nnU-Net model to segment fracture callus and cortical bone in microCT scans, reaching **0.95 Dice** on held-out specimens across three healing timepoints with callus volume error within 5%.

AI Engineer Intern — May 2024 – Jul 2024
Triangle Labs — *Los Angeles, CA*

- Fine-tuned reranking models in the RAG pipeline of an AI study platform used by **hundreds of thousands of students**, improving document retrieval performance by **20%**.
- Built and integrated backend APIs connecting the platform's software components.

PROJECTS

NASA Technical Reports OCR Dataset · *Hugging Face* — Nov 2026

- Re-transcribed **9,000 NASA technical reports (498,206 pages, 1960–1975)** with the LightOnOCR-3 vision-language model, producing Markdown text with HTML tables and LaTeX equations.
- Wrote a crawler with an automated rights filter (public, U.S.-government works only) and per-page quality flags; released under CC0 with 160+ downloads.

Pinball Score Reader · *live demo on Hugging Face Spaces* — 2024 – 2026

- Trained YOLOv11 detectors on 1,100 hand-labeled arcade photos (**0.986 mAP@50**) and paired them with a vision-language OCR model to read scores from a single photo; deployed as a Gradio app on ZeroGPU.

SenseNova-Loop-9B Visual Puzzle Model · *Hugging Face* — Nov 2026

- LoRA fine-tuned a 9B vision-language model on image-grid puzzles using a self-collected dataset, then merged and quantized it to GGUF (Q8_0, Q4_K_M) for local inference; 500+ downloads across versions.

Trinity Metals Website · *trinitymetals.com* — 2026

- Designed and built the production website for an Indianapolis critical-metals recycler using a custom WordPress theme.

GolfCart YOLO11s · *live demo on Hugging Face Spaces* — Fall 2025

- Fine-tuned a YOLO11s golf-cart detector on a hand-annotated drone dataset (**0.93 mAP@50**, 0.978 precision, 14 ms inference, 22 MB) and deployed it as an interactive Gradio demo.

HaikuGPT-17M · *Hugging Face* — Fall 2025

- Trained a 17M-parameter decoder-only transformer from scratch on a curated haiku corpus; it learns 5-7-5 structure from data alone, with no templates or constrained decoding.

SKILLS

Languages: Python, JavaScript, PHP, SQL
ML & AI: PyTorch, Hugging Face Transformers/PEFT, LoRA fine-tuning, RAG, rerankers, quantization (GGUF, bitsandbytes)
Computer vision: YOLO (Ultralytics), nnU-Net, DeepLabCut, vision-language models, OpenCV, FFmpeg
Infrastructure: Slurm HPC, Snakemake, Git, pytest, conda, REST APIs, Gradio, Hugging Face Spaces