

Joshua Jimenez

jjime73@illinois.edu | [linkedin.com/in/jjime](https://www.linkedin.com/in/jjime) | github.com/cleanpxe | Chicago, IL

EDUCATION

University of Illinois Urbana-Champaign

BS in Systems Engineering and Design, Minor in Computer Science

Urbana, IL

Expected May 2028

- **Completed coursework:** Data Structures, Discrete Structures, Linear Algebra, Analysis of Data, Statics, Dynamics, Strength of Materials, Business Side of Engineering
- **In progress (Fall 2026):** SE 320 Control Systems; CS 341 System Programming; SE 310 Design of Mechanisms; SE 498 AI-Driven Intelligent Systems

RESEARCH INTERESTS

Field robotics and off-road autonomy; physics-grounded perception; synthetic data and sim-to-real transfer; multimodal sensing (RGB, near-infrared, lidar, acoustic); embedded inference under hard compute budgets

RESEARCH EXPERIENCE

Autonomous Systems Research Assistant

Center for Autonomous Construction and Manufacturing at Scale (CACMS)

Apr. 2026 – Present

Urbana, IL

Soil-Moisture Perception for an Unmanned Ground Vehicle | *Perception/ML lead*

Apr. 2026 – Present

Stack: Python, PyTorch, Blender (Cycles, bpy), OpenCV, ONNX, OpenVINO, ROS 2 Humble, Linux.

Collaborative project; embedded target is an UP 7000 (6 W quad-core x86 SBC).

System design and embedded constraints

- Designed a physics-grounded perception system that classifies terrain under a moving UGV as dry, semi-wet, or saturated soil, plus a no-verdict “not-dirt” class, fusing RGB, near-infrared, and a 64-beam Ouster OS1 lidar mounted on a 2.1 m forward-tilted mast
- Engineered against a hard embedded budget: a 6 W SBC that simultaneously decodes two 5 MP camera streams under ROS 2, with an end-to-end latency budget covering capture, decode, and inference on shared cores, not model inference alone
- Derived the imaging geometry from first principles (ground sample distance, motion blur in native sensor pixels, and per-frame observation counts across candidate speed bands), correcting a standing assumption that had overstated the temporal-voting budget by 2–5 $\times$ and setting camera-mode policy (5 MP @ 30 fps over VGA @ 112 fps)
- Built the ONNX→OpenVINO export path and an end-to-end latency benchmark harness (fp32/fp16/int8) for MobileNetV3-Large with a pre-approved MobileNetV3-Small fallback, designed to close the deployment gate on real hardware before any final model is trained

Synthetic data and sim-to-real

- Authored a physically based Blender/Cycles synthetic-data pipeline ($\sim$ 13k lines of Python, 201 passing tests, several of which encode experimental methodology rather than code behavior) producing 10,000+ labeled frames per batch with paired RGB/NIR imagery, per-pixel moisture, depth, segmentation, and simulated lidar point clouds
- Modeled moisture physically (darkening, texture loss under a water film, a specular coat above field capacity), with moisture–darkening curves calibrated to LUCAS soil spectra
- Targeted sim-to-real gaps systematically: a 12-texture photoscanned soil bank (CC0) with anti-tiling, a 336-HDRI sky bank with sun extraction and dome coupling, dusk/dawn and overcast regimes weighted to the deployment climate, relief-following litter and rocks, canopy shadow fields, and vehicle self-shadowing
- Enforced label-blind randomization: every nuisance factor (lighting, soil type, debris, sky) is drawn independently of the moisture label, and photometric augmentation is bounded by a measured brightness envelope so jitter cannot erase the class signal
- Measured the pipeline instead of trusting it: lighting randomization varied image brightness 3.55 $\times$ while the full dry-to-saturated class signal varied it 1.7 $\times$; modeled the real vehicle’s gray-card exposure lock to cut the nuisance to 1.48 $\times$ against a 1.68 $\times$ signal, and added a config validator that refuses any setup where that inequality inverts
- Calibrated the exposure model with an 84-render gray-card sweep of absolute scene irradiance: verified render-to-irradiance linearity to 0.01%, measured the sky-dome contribution at 7.8–15.1 W/m² (refuting the premise of my own recorded decision), and isolated the residual brightness confound to ground BRDF at the 60° deployment view zenith

- Replaced a derived sky model with a measured one, driving the sky dome to absolute diffuse irradiance coupled to sun geometry; in the process found two silent pipeline faults (the configured sky model does not exist in Blender 5.0, and a turbidity parameter draw had never reached a render)
- Analyzed the vehicle's own shadow as an in-frame nuisance (predicted for $\sim 1/3$ of headings with a $4\text{--}7\times$ brightness drop vs. a $\leq 2.15\times$ class signal; NIR/red ratio identified as the physical disambiguator), implemented an ego-shadow caster, and measured the actual confound at $1.38\text{--}1.9\times$
- Diagnosed a domain shift between dataset generations with strip-statistic PCA and random-forest probes, isolated a lamp color-temperature cause via same-seed re-renders, and fixed it with a white-balance-relative illuminant model; quantified a 44% drop in reliability of the blue-sky mirror cue and added luminance-preserving illuminant jitter to training

Experimental rigor and ML methodology

- Designed and ran a controlled A/B pilot (1,200 synthetic images, single variable: photometric jitter cap) showing the brightness cue was a trap: widening jitter improved macro-F1 by $+0.12$ with every class improving, eliminated all 21 saturated $\rightarrow$ dry confusions (saturated recall $0.66\rightarrow 0.90$), and revealed that the tight-cap model plateaued at epoch 4 on a learned shortcut; adopted into the training config as a documented decision
- Validated that the discriminative cue survives the classifier's own feature space with a frozen-feature probe (ImageNet-pretrained MobileNetV3 features, linear probe with controls), confirming $\sim 60\%$ three-class separability with the strongest margin on the saturated class before committing to a full training campaign
- Resolved a contradiction between two published studies of wet-soil reflectance as a viewing-geometry dependence (the brightness turnover exists only near the specular direction); because the deployment camera sits at 60° view zenith, the finding changed the lab protocol and the synthetic material model, and was later reproduced in the calibrated renderer
- Synthesized 26 primary sources into four structured research reviews, each ending in a recorded decision or an explicitly filed gap

Data collection and training infrastructure

- Ran controlled field and bench data collection (drying-tray soil series with mass tracking, lidar return modeling over trays, camera characterization) to anchor simulation parameters in measurements
- Built the PyTorch training and evaluation framework: multi-stream RGB / NIR / early-fusion arms, motion-blur and rolling-shutter augmentation derived from vehicle speed and exposure, and resumable multi-day render batches with manifest-based recovery and verification gates before each commit

Autonomous Navigation on a Clearpath Husky A200 | ROS 2, Nav2, C++, Python, Linux June – July 2026

- Own the ROS 2 software stack for the lab's Clearpath Husky A200, taking a 50 kg off-road UGV from bare hardware and tethered manual teleoperation to autonomous point-to-goal navigation with Nav2 on a SLAM-built map, operated over a robot-hosted wireless link
- Integrated a full perception payload (Ouster OS1 lidar, MultiSense S27 stereo camera, MicroStrain AHRS) into ROS 2 Humble
- Wrote an instrumented closed-loop controller (C++/Python) with configurable gains, a safety watchdog, and per-run metrics logging; authored a field runbook with an e-stop pre-flight checklist

Additional Lab Contributions and Engineering Process

Apr. 2026 – Present

- Annotated 1,000+ images in CVAT to train perception models for GPS-denied underwater autonomous vehicles
- Model autonomous-vehicle sensor platforms in Fusion 360 for simulation
- Maintain a version-controlled research wiki shared with the lab (200+ pages, 300+ commits) with per-dataset provenance, licensing, and reproducibility records, a milestone plan with exit criteria, and a live risk register
- Govern geometry, optics, materials, lighting, and deployment-envelope choices through 25+ architecture decision records with amendments, keeping measured facts, assumptions, and open questions separate; when the inference target changed four times in one day (Jetson $\rightarrow$ workstation GPU $\rightarrow$ SBC $\rightarrow$ back), amended a single record four times so voided and surviving analysis stayed auditable
- Enforced dataset provenance: registered datasets with provenance records and excluded one from all training after it failed a provenance check (confirmed AI-generated imagery)

COURSE AND TEAM PROJECTS

Audio Rate – Passive Acoustic Exercise Exertion Estimation

Sep. – Dec. 2026

Team lead and principal designer, SE 498: AI-Driven Intelligent Systems

UIUC

- Lead a three-person team building a smartphone system that estimates exercise exertion (%HRR per 10 s window) from passive breathing audio, with no wearable on the user; own the model architecture, signal-processing design, and deliverable coordination

- Designed the signal-processing front end: unprocessed 16 kHz capture with OS audio effects disabled and acoustically verified, calibration-tone normalization, 10 s windows with a 5 s hop, and a frozen 64-band log-mel feature extractor
- Architected the model and inference stack: a compact per-window CNN feeding a GRU over 30–60 s of context, Huber-loss %HRR regression, and int8 TFLite / Core ML export for on-device inference
- Defined the evaluation protocol: chest-strap RR intervals converted to per-participant %HRR labels; leave-one-participant-out and leave-one-device-out splits across 3+ phone models; an incremental-validity analysis over heart rate using Borg and differentiated RPE
- Built an iOS capture app (48 kHz audio, motion logging, BLE Polar H10 client, event marks, session bundling) with a capture-qualification sequence (calibration tone, room noise, sync clap) that verifies OS processing has not removed the signal; confirmed clock sync to within ~6 ms
- Built the Polar H10 reference pipeline (130 Hz ECG, RR intervals over BLE, –7 ppm clock drift) for ground-truth labels, and a two-repo workflow (public release repo for reproducibility, private development repo with CI-ready scaffold and docs)
- Positioned the project’s novelty with an AI-assisted PRISMA-style review of 100+ records on smartphone and earable respiration and exertion sensing (closest precedent: Ren 2021; also BreathPro, VitalEar): consumer air-coupled mics, a continuous participant-normalized target, and explicit cross-person and cross-device generalization
- Owned the human-subjects data plan under the course’s ethical-conduct policy: informed consent, de-identified participant IDs, raw audio stored securely outside version control, and on-device inference

CATrack Inspect AI – HackIllinois Caterpillar Challenge

Python, FastAPI, Modal, Whisper, Claude Vision

Feb. 2026

Urbana, IL

- Built a multimodal AI pipeline fusing Whisper speech-to-text and Claude Vision to detect equipment anomalies from technician voice notes and photos of a Cat D6N dozer, returning severity-graded results
- Designed the FastAPI backend on Modal (REST endpoints and JSON contracts integrating two cross-team services) and authored prompt guardrails encoding known LLM failure modes
- Ran sprint-based development with acceptance criteria

LEADERSHIP

Product Lead, Product Space UIUC

Product Fellowship (6% acceptance rate); student product-consulting organization

Feb. 2026 – Present

Urbana, IL

Client Engagement: Brownlee Lighting | Co-lead, six-person consulting team

Sep. – Dec. 2026

Signed 12-week statement of work evaluating additive manufacturing as a scalable way to design and produce decorative shades for the client’s existing LED light engines, ending in a production-feasibility and insourcing recommendation.

Leadership and client management

- Co-lead the team across three phases (discovery, design and prototype, validation and feasibility); own the feasibility track end to end while my co-lead owns design; delivered Phase 1 on schedule (landscape report, initial concepts, design workflow, phased project plan)
- Run the weekly call with the client’s Head of Product Strategy and Lighting Engineer, converting each call into decisions, open items, and next-week deliverables within 48 hours; drove the scope change from pellet extrusion to filament FFF when the client dropped its planned machine
- Enforce a three-hour weekly budget per consultant: every deliverable sheet carries per-item minute estimates, overflow moves to the next week by rule, and workloads are rebalanced when a consultant returns from leave
- Coordinated the engagement’s first vendor interaction: a quote request to the client’s contract printer across two nozzle sizes and four volume bands, the key input to make-versus-buy

Design and manufacturing engineering

- Defined the design-for-additive workflow for the shade family (light engine modeled as padded keep-out cylinders, fit ring before full shade, one change per print, no print without a log entry) plus a 13-rule FFF printability checklist tied to UL and clearance constraints
- Led development of a parametric shade-family generator in Python with a Fusion 360 API front end: one spreadsheet row (shape × size × texture × nozzle) yields the solid, a printability, clearance, and glare check, a true-scale silhouette sheet, and sliceable STLs with a fit-test ring; runs headless on Linux for screening and inside Fusion for the build
- Scoped production as parts per eight-hour shift on filament FFF; mapped campus and commercial printers by build volume, nozzle size, and reliability; designed a nozzle test that separates wall thickness from print time

- Direct a materials screen across three PETG candidates against the client’s gates (UL 94 HB, relative thermal index at the fixture contact, clear and frosted optics, recycled content), and sourced a second recycled supplier to remove single-source risk
Feasibility and economics
- Own the business case: a cost-per-shade model with every assumption and source beside its input, a crossover-volume analysis for in-house versus outsourced printing, and a risk register with dated triggers reviewed weekly
- Reframed the recommendation from “insource or not” to a three-way comparison (in-house, outsourced, split by shade size) after establishing that outsourced lead time fits inside the client’s quoted customer lead time
Knowledge system and tooling
- Designed and operate the team’s knowledge system: a git-tracked Obsidian vault with an embedded AI agent (Claude Code) governed by a written operating manual that draws a hard line between what the agent may draft and what the team must write; YAML frontmatter on every note drives Dataview dashboards, Kanban boards, a design- and print-keyed iteration log, a materials matrix, and the cost model as live queries
- Automated the weekly operating loop: an idempotent ingest command that routes raw inputs into linked notes, decision records, and task cards (client calls transcribed locally with Whisper); a planner that re-derives the week from deadlines and the risk register; a scaffolding script that carries unfinished items forward; and a WeasyPrint publishing chain for client- and team-facing briefs
- Enforce integrity and confidentiality by script: a vault linter (broken links, duplicate names, template syntax, unfilled placeholders), a [VERIFY] convention that keeps unsourced specs, prices, and print times out of client documents, a gitignore backstop for CAD and mesh formats, and client files kept outside version control; team updates pushed to Notion via its API only on explicit approval

Fellowship Project: International Games Company

- Developed a U.S. go-to-market strategy through market research, competitive analysis, and user segmentation; trained in product discovery, metrics, and roadmapping under FAANG+ PM mentors

Philippine Student Association
University of Illinois Urbana-Champaign

May 2025 – Present
Urbana, IL

Treasurer

May 2026 – Present

- Serve as the lead financial advisor for the Philippine Student Association (200+ members)
- Manage financial operations for the largest student-led Filipinx-American conference in the Midwest
- Promote Filipinx culture by advising funding for award-winning cultural groups Harana and Barkada

Philanthropy Chair

May 2025 – May 2026

- Managed fundraising operations, establishing partnerships with Memorial Stadium
- Developed a volunteer management system to track assignments for 40+ volunteers
- Streamlined the organization’s processes and documentation management

TECHNICAL SKILLS

Robotics and sensing: ROS 2 (Humble), Nav2, SLAM, Ouster lidar, MultiSense stereo, MicroStrain AHRS, RGB/NIR cameras, V4L2, Ubuntu Linux

Perception and ML: PyTorch, OpenCV, NumPy/SciPy, scikit-learn, ONNX/OpenVINO, Blender/Cycles (bpy automation, procedural materials, geometry nodes), domain randomization, sim-to-real

Programming: Python, C++, C, Swift/SwiftUI, MATLAB, SQL (Postgres, SQLite), FastAPI, Pandas

Tools and practice: Git, Fusion 360 (incl. API), CVAT, Obsidian (Dataview, Templater), Notion API, Claude Code, WeasyPrint, Whisper, Excel, architecture decision records

ADDITIONAL INFORMATION

Languages: English (fluent), Japanese (intermediate)

Interests: Sports psychology, philosophy, classical literature, modern art, techno music