

# Joshua Terranova

7748 Via Sorrento, Burbank, CA 91504

(818) 515-9043 | [jjt\\_373@usc.edu](mailto:jjt_373@usc.edu)

[github.com/thatrandomasiandev](https://github.com/thatrandomasiandev) | [linkedin.com/in/jterranovacs](https://linkedin.com/in/jterranovacs)

## CURRICULUM VITAE

June 12, 2026

---

### TABLE OF CONTENTS

<b>EDUCATION</b>	2	ML Research Benchmark Suite	6
B.Sc., Computer Science — USC	2	<b>METHODOLOGICAL CONTRIBUTIONS</b>	7
A.Sc., Mathematics — LAMC	2	Key algorithmic implementations	7
<b>TECHNICAL SKILLS</b>	2	Unifying research question	8
Languages	2	<b>SERVICE AND PEER REVIEW</b>	8
Machine Learning & Deep Learning	2	Peer review and proposal evaluation	8
ML Research (ML Benchmark Suite)	2	Mentorship and STEM outreach	8
Agentic AI & LLM Systems	2	<b>RESEARCH PROFILE</b>	8
Robotics & Autonomy	2	<b>RESEARCH INTERESTS</b>	9
Systems & DevOps	2	Primary research thrusts	9
Spacecraft & CAD	2	Causal Inference and Decision Theory	9
<b>PROFESSIONAL EXPERIENCE</b>	2	Machine Learning and AI	10
NASA Team 19: Ad Astra (LSE)	2	Uncertainty Quantification and Reliable ML	10
NASA Team 10: Forged in Orbit (LSE)	3	Robotics and Autonomy	10
<b>ADDITIONAL EXPERIENCE</b>	3	Spacecraft and Planetary Science	10
USC Field Robotics Lab	3	Privacy and Distributed ML	10
USC Rocket Propulsion Lab	4	Scientific ML	10
USC Racing (Formula SAE)	4	<b>RESEARCH PUBLICATIONS AND TECHNICAL REPORTS</b>	10
USC MESA	4	National Aeronautics and Space Administration (NASA), Team 19: Ad Astra, 2026	10
<b>CERTIFICATIONS</b>	5	National Aeronautics and Space Administration (NASA), Team 10, 2025	11
<b>PERSONAL PROJECTS</b>	5	Personal Projects	12
SpaAIder	5		
Pneumonia Detection (Chest X-Rays)	6		
LLM-V1.0	6		

## EDUCATION

---

### **B.Sc. in Computer Science**

*UNIVERSITY OF SOUTHERN CALIFORNIA, 2027*

*Relevant Coursework:* Data Structures and Object-Oriented Design, Introduction to Algorithms, Artificial Intelligence, Machine Learning, Computer Vision, Probability and Statistics for Electrical Engineering and Computer Science, Linear Algebra and Linear Differential Equations, Discrete Methods in Computer Science, Principles of Software Development, Introduction to Robotics

### **A.Sc. in Mathematics**

*LOS ANGELES MISSION COLLEGE, 2025*

*Relevant Coursework:* Calculus I–III, Linear Algebra, Differential Equations, Probability and Statistics, Discrete Mathematics

## TECHNICAL SKILLS

---

### **Languages**

- Python (primary), TypeScript, JavaScript, C++, MATLAB, SQL, Bash/Zsh,  $\LaTeX$

### **Machine Learning & Deep Learning**

- PyTorch, TensorFlow/Keras, scikit-learn, HuggingFace Transformers, PEFT (LoRA, adapters, prefix tuning), NumPy, SciPy, Pandas, Matplotlib, OpenCV, YOLO (ultralitics), EfficientNetV2

### **ML Research (ML Benchmark Suite)**

- Causal inference (meta-learners, NOTEARS, PCMCI, OPE), Bayesian deep learning (VI, MC Dropout, conformal prediction), generative models (DDPM, flow matching, diffusion posterior sampling), graph neural networks (GCN, GAT, GIN, GraphSAGE, temporal GNNs), self-supervised learning (SimCLR, BYOL, VICReg, MAE), federated learning (FedAvg, FedProx, DP-SGD, secure aggregation), scientific ML (PINNs, EGNN, scVAE, FNO)

### **Agentic AI & LLM Systems**

- Ollama, FastAPI, SQLite, PostgreSQL, server-sent events, retrieval-augmented generation, vector embeddings (TF-IDF, nomic-embed-text), episodic memory, multi-agent orchestration

### **Robotics & Autonomy**

- ROS 2 (Humble), Nav2, Gazebo, EKF sensor fusion, GNSS waypoint navigation, ArUco visual servoing, YOLO-based object detection, e-stop interlock systems

### **Systems & DevOps**

- Git, GitHub, Docker, Linux, pyenv, Google Colab, Jupyter Notebooks, pytest, YAML-driven experiment pipelines

### **Spacecraft & CAD**

- Autodesk Fusion 360, FEA (finite element analysis),  $N^2$  diagram, ICD matrix authoring, S-band link budget analysis, spacecraft subsystem trade studies, TRL assessment

## PROFESSIONAL EXPERIENCE

---

### **Lead Systems Engineer / Research Lead, NASA Team 19: *Ad Astra***

2026

*National Aeronautics and Space Administration*

- Sole Lead Systems Engineer on 18-member MCA team (PM Alejandro Gonzalez, Chief Scientist Hannah Kim, DPMR Alfredo Guerra); led Engineering subteam integration with Science and Programmatic leads across MCR–PDR.
- Authored spacecraft overview, subsystem summary table (mass/power/dimensions), and top-level subsystem

- requirements flow-down for a lunar PSR volatile-prospecting rover.
- Owned interface architecture: 19-pathway ICD matrix, N<sup>2</sup> diagram, and interface narrative linking payload instruments to vehicle subsystems.
- Co-led thermal subsystem (requirements, heat-flow maps, trade study) for 40–394 K PSR environments; supported CDH trade study and communications downlink analysis.
- Owned CDH subsystem through PDR: requirements, software architecture flowchart, TRL/sub-assembly baselines, recovery/redundancy plans, and verification matrix for science data acquisition and S-band downlink.
- Researched alternative mission concepts with science lead; developed science-to-engineering traceability from volatile-detection objectives through verification planning.
- Developed science-to-engineering traceability matrix linking volatile-detection science objectives (LiDAR, spectrometry, rotary-percussive sampling) through system-level KPPs to subsystem verification criteria.
- Supported communications downlink analysis for S-band relay architecture; evaluated link margin, data rate, and contact window constraints for Nobile Rim 2 orbital geometry.
- Coordinated Engineering subteam documentation for MCR–PDR milestone reviews; managed version control and cross-subsystem consistency across 19-interface ICD under PM review cadence.

### **Lead Systems Engineer / Research Lead, NASA Team 10: *Forged in Orbit***

2025

*National Aeronautics and Space Administration*

- Sole Lead Systems Engineer on 12-member NPWEE team (PI Joel Bhattarai, PM Alexis Gallardo, Chief Engineer Justin Schueler, Chief Scientist Riya Jain); integrated Engineering, Science, and Programmatic inputs for a \$10,000-class technology proposal.
- Researched thermodynamic, mechanical, and materials behavior of vacuum-environment cold welding; derived quantitative KPPs from ASTROBEAT ISS heritage, NTRS literature, and materials compatibility analysis.
- Authored cold-welding performance specifications in team Technical Research Memorandum; co-authored final proposal and quad-chart concept portfolio (cold-weld reclamation, vacuum dust ejection, triboelectric transport).
- Lead-authored L-MAP-FI: multi-agent lunar pathfinding and environmental forecasting framework (ROS2/Gazebo simulation, Fusion 360 FEA validation).
- Designed vacuum-test research protocol for weld-parameter characterization; developed alloy-specific vacuum-weld parameter matrix for austenitic stainless, aluminum, and copper alloy systems across –183 °C to +137 °C thermal envelope.
- Developed quantitative KPPs for vacuum cold-welding system: minimum contact pressure 100 MPa, actuation force 10 kN, target shear strength ≥80 MPa, derived from ASTROBEAT ISS heritage data and NTRS archival literature.
- Performed comparative cost and performance analysis across three ISRU technology concepts: cold-weld reclamation (94% cost reduction vs. TIG welding), ultrasonic vacuum dust ejection (1.5 g/min removal rate), and triboelectric autonomous transport.
- Evaluated technology readiness levels and NASA Technology Taxonomy alignment (TX4.3.3, TX4.6.1, TX12.4.1, TX12.4.6) for autonomous in-situ repair architectures.
- Lead Systems Engineer Intern, National Aeronautics and Space Administration (NASA), 2025–2026.
- Computer Engineer Intern, National Aeronautics and Space Administration (NASA), 2025 (multiple terms).

## **ADDITIONAL EXPERIENCE**

### **Software Engineer, USC Field Robotics Lab**

2026–Present

*University of Southern California*

- Developed full autonomy software stack for University Rover Challenge 2026 on ROS 2 (Humble); system integrated navigation, perception, manipulation, and science task modules.
- Implemented 7-target mission orchestrator managing state machine transitions across autonomous navigation, object retrieval, equipment servicing, and science site analysis tasks.
- GNSS waypoint navigation pipeline with coordinate-frame transforms (WGS-84 to local ENU), dead-reckoning fallback, and dynamic obstacle replanning via Nav2 costmaps.
- ArUco marker visual servoing for precision docking: camera intrinsic calibration, pose estimation via solvePnP, and proportional velocity controller with convergence tolerance.
- YOLO-based object detection with confidence gating; integrated detection results into Nav2 behavior tree for target approach and manipulation task triggering.
- EKF sensor fusion combining IMU (9-DOF), wheel odometry, and GNSS measurements; tuned process and measurement noise covariances for outdoor terrain traversal.
- E-stop interlock system with hardware and software layers; watchdog timers, heartbeat monitoring, and safe-state recovery procedures per URC safety requirements.
- Judge telemetry web dashboard: real-time ROS 2 topic visualization, rover state display, GPS track overlay, and task completion logging served over WebSocket.
- System integration testing on outdoor terrain; characterized localization drift, detection false-positive rates, and navigation replanning latency under competition conditions.
- Collaborated with 10+ team members across mechanical, electrical, and software subteams; maintained shared ROS 2 workspace, launch file architecture, and hardware abstraction layer.

### **Software Engineer, USC Rocket Propulsion Lab**

*2025–Present*

*University of Southern California*

- Developed software systems supporting experimental liquid and hybrid propulsion research: sensor data acquisition, test sequence automation, and post-test data reduction pipelines.
- Instrumentation software for pressure transducers, thermocouples, and load cells; real-time data logging with configurable sampling rates and alarm thresholds.
- Test operations tooling: automated pre-fire checklist verification, countdown sequencing, and abort logic with hardware interlock validation.
- Post-test analysis pipelines: chamber pressure trace parsing, specific impulse calculation, and thrust curve integration from raw force sensor data.
- Design analysis support: propellant mass flow rate modeling, nozzle expansion ratio trade studies, and combustion chamber thermal margin estimation.

### **Software Engineer, USC Racing (Formula SAE)**

*2025–2026*

*University of Southern California*

- Developed embedded software and data acquisition systems for Formula SAE vehicle; sensor integration for suspension, brake, and powertrain telemetry.
- Designed and implemented CAN bus data logging pipeline; parsed vehicle telemetry for real-time driver feedback display and post-session performance analysis.
- Contributed to vehicle dynamics simulation tooling for suspension geometry and lap time optimization.
- Supported systems integration testing and on-track data collection during development and competition events.

### **Research Fellow, USC MESA**

*2022–Present*

*University of Southern California*

- Research mentor and STEM pathway advisor for underrepresented undergraduate students pursuing research readiness; conducted workshops on scientific communication, literature review, and research proposal writing.
- Developed curriculum materials covering research methodology, academic writing conventions, and

graduate school preparation strategies for STEM undergraduates.

- Facilitated peer mentorship cohorts; tracked mentee progress toward research placement, conference participation, and graduate program applications.
- Contributed to program assessment and reporting; collected outcome metrics on mentee research placement rates and academic progression.

## CERTIFICATIONS

---

- NASA L'SPACE Mission Concept Academy (MCA), Certificate of Completion, Team 19: *Ad Astra*, 2026.
- NASA L'SPACE NASA Proposal Writing and Evaluation Experience (NPWEE) Academy, Certificate of Completion, Team 10: *Forged in Orbit*, 2025.

## PERSONAL PROJECTS

---

### SpaAIder (Local-First Agentic Runtime)

2025–Present

- Foundational split: TypeScript control plane (`runOrchestrator`) owns user-facing planning and delegation; Python FastAPI substrate owns persistence, memory APIs, and service routes; Ollama supplies sovereign on-device inference across orchestrator, subagent, and embedding models.
- Memory substrate: `MemoryManager` delegates to contracted SQL tables (sessions, messages, knowledge, context, embeddings on SQLite/PostgreSQL) and an `EpisodicStore`; `/api/v1/memory/recall` fuses vector, episodic, and keyword channels before each plan.
- Episodic consolidation: session close triggers LLM synthesis of transcripts into episodes (summary, key facts, topics, tone) merged into durable user profiles.
- Context assembly: Python `ContextEngine` allocates a fixed token budget across system instructions, user profile, episodic hits, and working memory; an `EmbeddingBridge` (TF-IDF locally, Ollama `nomic-embed-text` when configured) scores similarity for recall and knowledge search.
- Ingress and observability: heterogeneous clients (web, mobile, Electron, CLI, Code-OSS IDE) POST to `/api/agent/stream`; an `AsyncEventQueue` multiplexes orchestration events—context fetch, plan emission, per-task status, subagent streams, final answer—over server-sent events.
- Planning contract: meta-orchestrator LLM receives recalled `UserContext` and must emit JSON `TaskPlan` only (intent, agent assignments, `dependsOn` DAG edges over a 135-agent type taxonomy) without executing tasks—pure delegation.
- Execution engine: dependency-aware parallel scheduler launches runnable tasks up to `maxParallelTasks`; domain subagents (code, filesystem, search, vision, memory, research, backend) run as async generators; `memory-agent` proxies read/write/recall to FastAPI; `AgentBus` provides in-process pub/sub and request-reply between agents.
- Quality loop and surfaces: critique and synthesis agents post-process subagent artifacts into one response; completed turns persist through the memory layer; IDE fork reuses the same agent stream with editor-grounded file, selection, and terminal context.
- Persistence architecture: PostgreSQL production / SQLite development with versioned SQL schema migrations; sessions, messages, knowledge, context, and embeddings tables with foreign-key integrity; `EpisodicStore` backed by same schema.
- IDE integration: Code-OSS fork with extension injecting active file path, text selection, and terminal stdout into the agent context window; same FastAPI backend serves both web client and IDE extension over identical SSE stream.
- Deployment targets: web (React), mobile (React Native), desktop (Electron), CLI (Node.js), and IDE extension; all share the TypeScript `runOrchestrator` and POST to the same `/api/agent/stream` endpoint.

## Pneumonia Detection from Chest X-Rays

2025–2026

- TensorFlow pipeline (`train.py`) for binary chest radiograph classification on 988 images (620 train / 368 validation).
- EfficientNetV2B0 transfer learning with focal loss, class balancing, oversampling, and two-phase fine-tuning.
- Test-time augmentation, threshold tuning, and evaluation via confusion matrices, ROC/AUC, and precision–recall curves.
- Achieved 83% validation accuracy; exported models and metrics to structured outputs/.
- Dataset: 988 chest radiographs (620 train / 368 validation) from Kaggle RSNA Pneumonia Detection Challenge subset; applied class-weighted oversampling to address 3:1 normal-to-pneumonia imbalance.
- Focal loss formulation:  $FL(p_t) = -\alpha_t(1 - p_t)^\gamma \log(p_t)$  with  $\gamma = 2.0$  and  $\alpha = 0.25$ ; suppresses easy-negative gradient contribution and focuses training on hard positives.
- Two-phase fine-tuning: phase 1 trains classification head with frozen EfficientNetV2 backbone (up to 25 epochs); phase 2 unfreezes top 40 backbone layers at 10× reduced learning rate ( $10^{-5}$ ) for up to 10 additional epochs.
- Evaluation suite: per-class precision/recall/F1, confusion matrix, ROC/AUC and precision–recall curves with average precision; threshold tuned on validation set to maximize F1 with test-time augmentation.

## Small-Scale Language Model (LLM-V1.0)

2025–2026

- Designed and trained a ~30M-parameter GPT-style transformer (8 layers, 8 heads, 512-dim embeddings, 8,192-token BPE vocabulary).
- PyTorch training, checkpointing, and text-generation pipelines for Colab and local workflows.
- Architecture: GPT-style decoder-only transformer; 8 layers, 8 attention heads, 512-dimensional embeddings, 512-token context window, feedforward dimension 2048 (4× expansion), ~30M total parameters.
- Tokenization: byte-pair encoding (BPE) with 8,192-token vocabulary trained on corpus via HuggingFace tokenizers; special tokens for BOS, EOS, and padding.
- Training: AdamW optimizer ( $\text{lr}=3 \times 10^{-4}$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ , weight decay=0.1); cosine learning rate schedule with warmup; gradient clipping at 1.0; mixed-precision (fp16) on CUDA; gradient accumulation for memory efficiency.
- Data pipeline: streaming dataset loader with dynamic batching; sequence packing to maximize GPU utilization; train/val split with per-epoch perplexity tracking.
- Checkpointing: saves model weights, optimizer state, and scheduler state every  $N$  steps; supports resume from checkpoint for Colab session continuity.
- Inference: top- $k$  and nucleus (top- $p$ ) sampling with temperature scaling; beam search option; streaming token generation with configurable `max_new_tokens`.
- Evaluation: perplexity on held-out validation set; qualitative generation samples logged at each checkpoint; attention weight visualization for interpretability analysis.
- Colab compatibility: optimized for T4/A100 free-tier constraints; automatic device detection (CUDA/MPS/CPU); fp16 auto-casting with fallback to fp32 for stability.

## ML Research Benchmark Suite (10 modules)

2025–2026

- Designed and implemented ten independent reproducible research frameworks spanning the major subfields of modern machine learning; each module follows a shared architecture of synthetic ground-truth DGPs, YAML-configured benchmark runners, pytest suites, and structured JSON/Markdown evaluation reports.
- Module 01—Causal ML for Decision-Making: five meta-learners (S/T/X/R/Causal Forest), doubly robust and TMLE estimators, NOTEARS linear and nonlinear DAG structure learning, PCML time-series causal discovery, IPS/SNIPS/DR off-policy evaluation, conformal CATE intervals, E-value sensitivity analysis; 47 source modules, 12 test files, 71 passing tests.
- Module 02—Foundation Model Adaptation: LoRA rank decomposition, adapter bottlenecks, prefix tuning, full fine-tune baselines; LiSSA influence function attribution; TracIn checkpoint-based data attribution;

- ECE/MCE calibration; risk-coverage curves for selective prediction; bootstrap confidence intervals.
- Module 03—Reinforcement Learning for Real Environments: Conservative Q-Learning (CQL) and Implicit Q-Learning (IQL) offline RL; CMDP Lagrangian safe RL with primal-dual lambda updates; ensemble world models with epistemic/aleatoric uncertainty decomposition; CEM planning over model rollouts.
  - Module 04—Probabilistic and Bayesian Deep Learning: mean-field variational inference with Bayesian-Linear layers and closed-form KL; MC Dropout uncertainty; deep ensembles with variance decomposition; temperature scaling and vector scaling calibration; split conformal prediction with RAPS adaptive score function; reliability diagrams.
  - Module 05—Generative Modeling for Science: DDPM with cosine and linear schedules; DDIM deterministic sampling; conditional score network with classifier-free guidance; flow matching with straight-path conditional vector fields; Diffusion Posterior Sampling (DPS) for inverse problems; Bayesian molecular optimization with EI/UCB/PI acquisition.
  - Module 06—Graph ML and Relational Systems: GCN (spectral), GAT (multi-head attention with stored attention weights), GIN (learnable epsilon, BatchNorm), GraphSAGE (mean/max/LSTM aggregators), DiffPool differentiable pooling; temporal GNN with GRU memory and negative sampling; link prediction with AUC/AP/MRR evaluation.
  - Module 07—Self-Supervised Representation Learning: SimCLR (NT-Xent loss), MoCo (momentum queue), BYOL (EMA target network, asymmetric predictor), VICReg (invariance + variance + covariance), Barlow Twins (cross-correlation regularization), masked autoencoder (MAE) with 75% masking ratio; linear probe and  $k$ -NN evaluation; feature alignment/uniformity metrics.
  - Module 08—Neurosymbolic Reasoning: differentiable forward chaining with product/Gödel/Łukasiewicz  $t$ -norms and fixed-point iteration; neural theorem prover with embedding-space unification similarity; DSL with type system for program synthesis; beam search over typed program space; fuzzy logic circuit with learnable gates.
  - Module 09—Federated and Privacy-Preserving ML: FedAvg, FedProx (proximal regularization), FedNova (normalized updates), FedBN (local BN statistics); DP-SGD with per-sample gradient clipping and Gaussian noise injection; Rényi DP accountant with  $(\epsilon, \delta)$ -DP conversion; pFedMe Moreau envelope personalization; MAML-Federated; pairwise masking secure aggregation; membership inference attack evaluation.
  - Module 10—ML for Science: Physics-Informed Neural Networks (PINNs) with autograd PDE residuals for heat and Burgers equations; E(3)-equivariant GNN (EGNN) for molecular property prediction; scVAE with negative binomial likelihood for single-cell RNA-seq; ComBat and Harmony batch correction; Bayesian optimization for materials discovery (EI/UCB/PI over GP surrogate); Fourier Neural Operator (FNO) for PDE solution operators; protein contact map prediction.
  - Shared infrastructure: YAML experiment configs, seed-controlled DGPs with oracle ground truth (PEHE, calibration error, graph SHD, equivariance residual), JSON result serialization, Markdown benchmark reports; ~33,000 lines of Python source, ~5,300 lines of tests, 310 pytest cases across 271 modules.

## METHODOLOGICAL CONTRIBUTIONS

---

### Key algorithmic implementations (research-grade, from scratch)

- NOTEARS nonlinear DAG structure learning (Module 01): augmented Lagrangian optimizer with matrix-exponential acyclicity constraint  $h(W) = \text{tr}(e^{W \odot W}) - d = 0$  and per-variable MLP structural equations; gradient  $\nabla_W h = 2W \odot e^{W \odot W}$  derived analytically and verified against finite-difference Jacobian.
- TMLE targeted minimum loss-based estimation (Module 01): clever covariate  $H(A, W) = A/\hat{e}(W) - (1 - A)/(1 - \hat{e}(W))$ , epsilon fluctuation via logistic sub-model, one-step bias correction; doubly robust influence-function standard error with bootstrap overlap and E-value sensitivity.

- Mean-field variational inference BNN (Module 04): BayesianLinear with softplus reparameterization  $\sigma = \log(1 + \exp(\rho))$ , local reparameterization trick  $w = \mu + \sigma\varepsilon$ , and closed-form KL divergence  $\sum_i [\log(\sigma_p/\sigma_q) + (\sigma_q^2 + \mu_q^2)/(2\sigma_p^2) - 1/2]$ ; ELBO loss with Monte Carlo likelihood averaging.
- DDPM + DDIM sampling with classifier-free guidance (Module 05): forward process  $q(x_t|x_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)I)$ ; DDIM deterministic reverse  $x_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\hat{x}_0 + \sqrt{1 - \bar{\alpha}_{t-1}}\varepsilon_\theta$ ; guidance  $\varepsilon_{\text{guided}} = \varepsilon_{\text{uncond}} + w(\varepsilon_{\text{cond}} - \varepsilon_{\text{uncond}})$  with configurable scale  $w$ .
- Multi-head Graph Attention Network and GIN (Module 06): per-head attention  $e_{ij} = \text{LeakyReLU}(a^T [Wh_i || Wh_j])$  with masked softmax over adjacency support; stored attention weights ( $K \times N \times N$ ) for interpretability; GIN update  $h_v^{(k)} = \text{MLP}^{(k)}((1 + \varepsilon^{(k)})h_v^{(k-1)} + \sum_{u \in \mathcal{N}(v)} h_u^{(k-1)})$  with learnable  $\varepsilon$ .
- DP-SGD with Rényi DP accounting (Module 09): per-sample gradient clipping via vmap, Gaussian noise at calibrated scale  $\sigma$ ; RDP of Gaussian mechanism  $\varepsilon_\alpha = \alpha/(2\sigma^2)$ ; composition and conversion to  $(\varepsilon, \delta)$ -DP via  $\varepsilon(\delta) = \min_\alpha [\varepsilon_\alpha + \log(1 - 1/\alpha)/(\alpha - 1) - \log(\delta)/(\alpha - 1)]$ .
- FedProx / pFedMe personalization (Module 09): FedProx local objective  $F_i(w) + (\mu/2)\|w - w_{\text{global}}\|^2$ ; FedNova normalized update aggregation for heterogeneous local step counts; pFedMe Moreau envelope  $w_i^* = \arg \min_w F_i(w) + (\lambda/2)\|w - \theta\|^2$ .
- Physics-Informed Neural Networks (Module 10): autograd PDE residual  $\mathcal{L}_{\text{PDE}} = \|\mathcal{N}[u_\theta](x, t)\|^2$  for heat equation  $u_t = \alpha u_{xx}$  and Burgers equation  $u_t + uu_x = \nu u_{xx}$ ; E(3)-equivariant message passing satisfying  $f(Rx + t) = Rf(x)$  for SO(3) rotations  $R$ .

### Unifying research question across the suite

- All ten modules share a common experimental protocol: synthetic data-generating processes with tunable assumption-violation parameters, oracle ground truth, and metrics that directly measure failure under those violations (PEHE for CATE estimators, SHD for graph recovery, ECE for calibration, equivariance residual for symmetry-constrained models). The overarching question is: *how and at what rate do modern ML estimators degrade as their identifying assumptions are systematically violated?* This adversarial framing produces findings generalizable beyond individual architectures.

## SERVICE AND PEER REVIEW

---

### Peer review and proposal evaluation

- Formal reviewer on NASA Review Panel #6 (2025): independently evaluated complete technology proposals from Teams 16, 17, and 18 using NASA-structured scoring methodology covering technical merit, feasibility, and programmatic risk.
- Cross-team technical reviewer for Teams 19 and 10 internal milestone documents; provided written feedback on CDH requirements narratives, ICD completeness, and subsystem requirement traceability prior to formal NASA review.
- Open-source research software stewardship: ML Research Benchmark Suite released under structured module architecture with reproducibility guarantees (seed-controlled DGPs, pinned dependency environments, complete test suites); designed to lower the barrier for replication studies.

### Mentorship and STEM outreach

- USC MESA Research Fellow (2022–Present): mentored cohorts of underrepresented undergraduates through research placement, graduate school preparation, and scientific communication; developed workshop curriculum and facilitated peer mentorship cohorts. (See Additional Experience for full details.)

## RESEARCH PROFILE

---

My central research question is how to build learning systems that remain reliable under distribution shift, sparse data, and deployment constraints—from observational causal inference and calibrated uncertainty

quantification to autonomous robots operating in unstructured environments. I pursue this through two complementary threads: rigorous ML methodology (controlled benchmarks with known ground truth) and systems engineering (spacecraft, rovers, and agentic runtimes where failure has real cost). At NASA L'SPACE I translated science objectives into verifiable subsystem architectures; in parallel I built SpaAIder, a retrieval-augmented multi-agent runtime with tiered episodic–semantic–vector memory, and a ten-module ML research benchmark suite that studies estimator and policy failure modes rather than reporting best-case metrics alone.

Independently architected a ten-module ML research benchmark suite spanning causal inference through scientific machine learning. Module 01 studies heterogeneous treatment effects and offline policy evaluation under confounding; Module 04 quantifies predictive uncertainty via variational Bayes, deep ensembles, and conformal prediction; Module 03 addresses offline and safe reinforcement learning under model epistemic error; Module 10 integrates physics-informed neural networks, equivariant molecular models, and neural operators for scientific computing. Each module implements production-quality research software: synthetic DGPs with oracle ground truth, YAML-configured experiment sweeps, pytest suites, and structured evaluation reports—the infrastructure required to systematically study when modern ML methods fail, not only when they succeed.

In robotics, developed the full autonomy stack for USC's University Rover Challenge 2026 entry on ROS 2 (Humble): 7-target mission orchestrator, EKF-fused localization, GNSS waypoint navigation, ArUco visual servoing, and YOLO-based detection with behavior-tree integration. This systems engineering work complements spacecraft experience at NASA, where two consecutive L'SPACE assignments as sole Lead Systems Engineer produced formal reviews (MCR through PDR), subsystem ownership (CDH, thermal, ICD), and co-authorship on twelve technical reports.

The methodological thread across all projects is a commitment to measuring what fails rather than what succeeds. In causal inference this means evaluating CATE estimators under violations of overlap and unconfoundedness, not just on benign DGPs. In Bayesian deep learning it means measuring calibration degradation as dataset shift increases, not just reporting ECE on in-distribution data. In federated learning it means quantifying privacy–utility tradeoffs as a function of noise multiplier and clipping threshold under heterogeneous data partitions. This failure-mode orientation shapes both what I build and how I evaluate it, and it reflects a research philosophy I intend to carry into graduate study: rigorous empirical methodology demands adversarial evaluation, not optimistic reporting.

I aim to continue this trajectory in doctoral study—connecting rigorous ML methodology with embodied autonomy and aerospace systems—where interdisciplinary programs can bridge the gap between theoretical guarantees and deployment-constrained systems. My goal is to produce research that is simultaneously mathematically precise, empirically rigorous, and grounded in the failure modes that matter in practice: distribution shift, model misspecification, privacy constraints, and physical symmetry requirements. The ten-module benchmark suite represents a working prototype of this research methodology at undergraduate scale; doctoral study would allow me to pursue these questions at the depth required to produce publishable theoretical and empirical contributions.

## RESEARCH INTERESTS

---

### Primary research thrusts

- Trustworthy machine learning: heterogeneous treatment effect estimation, uncertainty quantification, calibration, and conformal prediction under distribution shift.
- Embodied AI and autonomous systems: ROS 2 autonomy, sensor fusion, visual servoing, and multi-agent orchestration for deployment-constrained environments.

### Causal Inference and Decision Theory

- Heterogeneous treatment effect estimation under unmeasured confounding; doubly robust and semipara-

metric efficient estimators; CATE sensitivity analysis (E-values, Rosenbaum bounds); off-policy evaluation and counterfactual policy optimization; structural causal models and interventional distributions; time-series causal discovery (PCMCI, Granger, VAR); causal representation learning and identifiability theory.

### **Machine Learning and AI**

– Generative models (diffusion, flow matching, VAEs) for scientific applications; self-supervised and contrastive representation learning; graph neural networks and temporal relational learning; neurosymbolic reasoning and differentiable program synthesis; retrieval-augmented generation and agentic AI systems; local LLM inference and transformer architecture optimization.

### **Uncertainty Quantification and Reliable ML**

– Bayesian deep learning (variational inference, MC Dropout, deep ensembles); conformal prediction and distribution-free coverage guarantees; calibration theory and recalibration methods; selective prediction and risk-coverage tradeoffs; epistemic vs. aleatoric uncertainty decomposition; out-of-distribution detection; uncertainty-aware decision making under deployment constraints.

### **Robotics and Autonomy**

– Autonomous mobile robotics and ROS 2 systems; EKF sensor fusion and state estimation; visual servoing and marker-based localization; behavior tree mission orchestration; multi-agent pathfinding and swarm coordination.

### **Spacecraft and Planetary Science**

– Spacecraft systems architecting and subsystem integration; lunar PSR volatile characterization; in-situ resource utilization; fault-tolerant command and data handling for extra-terrestrial environments; subsystem-level interdependency modeling.

### **Privacy and Distributed ML**

– Federated learning under data heterogeneity; differential privacy and the Rényi DP composition framework; secure aggregation protocols; personalized federated learning; membership inference attack evaluation; privacy–utility tradeoff characterization under varying noise multiplier and clipping threshold.

### **Scientific ML**

– Physics-informed neural networks; equivariant architectures for molecular and materials property prediction; neural operators for PDE solution; single-cell genomics and batch effect correction; Bayesian optimization over GP surrogates for materials discovery; symmetry-constrained learning and physical law enforcement via loss design.

## **RESEARCH PUBLICATIONS AND TECHNICAL REPORTS**

---

### *National Aeronautics and Space Administration (NASA), Team 19: Ad Astra, 2026*

*18-member L'SPACE MCA team. Lead Systems Engineer / Research Lead (sole LSE). Reports to PM Alejandro Gonzalez; integrates with Chief Scientist Hannah Kim and DPMR Alfredo Guerra.*

1. Gonzalez, A., Guerra, A., Chang, A., Dixon, D., Zhou, E., Senteza, J., Mendez, J., Kumar, K., Richmond, R., **Terranova, J.**, Kim, H., French, S., Noor, M., Saad, S., Durkin, J., Raihan, F., Gutierrez, N., & Saxena, H., Preliminary engineering design, verification matrix, and integrated vehicle baseline for an autonomous lunar volatile-prospecting rover at Nobile Rim 2 (Preliminary Design Review, Team 19), National Aeronautics and Space Administration (NASA), 2026.

– Lead Systems Engineer contribution: owned CDH subsystem through PDR—requirements narrative, software architecture flowchart, TRL/sub-assembly baselines, recovery/redundancy plans, and verification matrix for science data acquisition, processing, and S-band downlink; coordinated Engineering subteam inputs under PM review.

2. Gonzalez, A., Guerra, A., Chang, A., Dixon, D., Zhou, E., Senteza, J., Mendez, J., Kumar, K., Richmond,

R., **Terranova, J.**, Kim, H., French, S., Noor, M., Saad, S., Durkin, J., Raihan, F., Gutierrez, N., & Saxena, H., Definitive mission architecture, operational paradigm, and programmatic baseline for in-situ lunar volatile characterization (Mission Definition Review, Team 19), National Aeronautics and Space Administration (NASA), 2026.

– Lead Systems Engineer contribution: researched mission architecture and operational paradigms for in-situ volatile characterization at Nobile Rim 2; linked Artemis-class PSR science objectives to traverse, communications, and programmatic constraints with Science and Programmatic subteams.

3. Gonzalez, A., Guerra, A., Chang, A., Dixon, D., Zhou, E., Senteza, J., Mendez, J., Kumar, K., Richmond, R., **Terranova, J.**, Kim, H., French, S., Noor, M., Saad, S., Durkin, J., Raihan, F., Gutierrez, N., & Saxena, H., System requirements baseline, subsystem allocation framework, and interface control architecture for a lunar PSR exploration rover (System Requirements Review, Team 19), National Aeronautics and Space Administration (NASA), 2026.

– Lead Systems Engineer contribution: authored spacecraft overview, subsystem summary table (mass/power/dimensions), 19-interface ICD matrix, N<sup>2</sup> diagram, and interface narrative; co-led thermal heat-flow analysis and CDH/thermal trade studies for 40–394 K PSR operating environments.

4. Gonzalez, A., Guerra, A., Chang, A., Dixon, D., Zhou, E., Senteza, J., Mendez, J., Kumar, K., Richmond, R., **Terranova, J.**, Kim, H., French, S., Noor, M., Saad, S., Durkin, J., Raihan, F., Gutierrez, N., & Saxena, H., Conceptual elaboration of an autonomous lunar volatile-reconnaissance architecture for permanently shadowed region prospecting (Mission Concept Review, Team 19), National Aeronautics and Space Administration (NASA), 2026.

– Lead Systems Engineer contribution: co-authored alternative mission concept comparison with Chief Scientist Hannah Kim; developed conceptual architecture integrating LiDAR navigation, rotary-percussive regolith sampling, and onboard spectrometric volatile analysis for PSR reconnaissance.

5. Gonzalez, A., Guerra, A., Chang, A., Dixon, D., Zhou, E., Senteza, J., Mendez, J., Kumar, K., Richmond, R., **Terranova, J.**, Kim, H., French, S., Noor, M., Saad, S., Durkin, J., Raihan, F., Gutierrez, N., & Saxena, H., Command and data handling subsystem interface specification: message protocols, timing constraints, and fault-response contracts for a lunar volatile-prospecting rover (CDH Interface Document, Team 19), National Aeronautics and Space Administration (NASA), 2026.

– Lead Systems Engineer contribution: authored CDH interface specification documenting all command and telemetry interfaces between the CDH subsystem and payload instruments, power subsystem, communications transponder, and mobility electronics; established message formats, timing constraints, and fault-response interface contracts for PDR.

6. Gonzalez, A., Guerra, A., Chang, A., Dixon, D., Zhou, E., Senteza, J., Mendez, J., Kumar, K., Richmond, R., **Terranova, J.**, Kim, H., French, S., Noor, M., Saad, S., Durkin, J., Raihan, F., Gutierrez, N., & Saxena, H., Thermal subsystem requirements, heat-flow analysis, and MLI vs. active-heating trade study for a permanently shadowed region rover (Thermal Subsystem Report, Team 19), National Aeronautics and Space Administration (NASA), 2026.

– Lead Systems Engineer contribution: co-authored thermal subsystem analysis establishing heat-flow maps for PSR orbital and surface phases; performed comparative trade study of multi-layer insulation vs. active heating architectures for 40–394 K operational envelope; linked thermal KPPs to CDH safe-mode entry thresholds.

### *National Aeronautics and Space Administration (NASA), Team 10, 2025*

*12-member L'SPACE NPWEE team. Lead Systems Engineer / Research Lead (sole LSE). PI Joel Bhattarai; PM Alexis Gallardo; Chief Engineer Justin Schueler; Chief Scientist Riya Jain.*

7. Bhattarai, J., Schueler, J., **Terranova, J.**, Rojas, J., Gallardo, A., Ek, S., Jain, R., Concepcion, W., Jonson, M., Li, Q., Vasquez, G., & Atwell, J., Solid-state in-situ metallurgical reconstitution via autonomous cold-welding attachment for lunar infrastructure salvage (Final Proposal, Team 10: Forged in Orbit), National Aeronautics and Space Administration (NASA), 2025.

– Lead Systems Engineer contribution: co-authored \$10,000-class final proposal; synthesized literature-backed cold-welding KPPs (100 MPa pressure, 10 kN force,  $\geq 80$  MPa shear strength) and autonomous in-situ repair architecture aligned to NASA Technology Taxonomy (TX4.3.3, TX4.6.1, TX12.4.1, TX12.4.6).

8. Bhattarai, J., Schueler, J., **Terranova, J.**, Rojas, J., Gallardo, A., Ek, S., Jain, R., Concepcion, W., Jonson, M., Li, Q., & Vasquez, G., Tripartite technology concept portfolio for autonomous lunar surface operations and in-situ resource utilization (Quad Charts, Team 10), National Aeronautics and Space Administration (NASA), 2025.

– Lead Systems Engineer contribution: researched tripartite lunar ISRU and mobility concept portfolio—oxide-scrubbing cold-weld reclamation, ultrasonic vacuum dust ejection, triboelectric autonomous transport—with comparative cost and performance analysis (94% cost reduction vs. TIG welding; 1.5 g/min vacuum dust removal).

9. **Terranova, J.**, & Vasquez, G., L-MAP-FI: A multi-agent pathfinding and predictive environmental forecasting interface for autonomous lunar reconnaissance swarms (Technical Presentation, Team 10), National Aeronautics and Space Administration (NASA), 2025.

– Lead author: designed L-MAP-FI multi-agent pathfinding and predictive environmental forecasting framework for autonomous lunar reconnaissance swarms; validated through ROS2/Gazebo simulation and Fusion 360 FEA (70× coverage, 160× cost efficiency vs. state-of-the-art).

10. Schueler, J., **Terranova, J.**, Bhattarai, J., & Atwell, J., Thermodynamic, mechanical, and materials characterization of vacuum-environment cold-welding parameters for robotic lunar repair systems (Technical Research Memorandum, Team 10), National Aeronautics and Space Administration (NASA), 2025.

– Primary research author: derived cold-welding performance specifications from ASTROBEAT ISS in-orbit repair experiments, NTRS heritage data, and austenitic stainless/aluminum/copper alloy compatibility analysis across a  $-183$  °C to  $+137$  °C thermal envelope.

11. Bhattarai, J., Gallardo, A., Ek, S., Schueler, J., **Terranova, J.**, Rojas, J., Jain, R., Concepcion, W., Jonson, M., Li, Q., & Vasquez, G., Interdisciplinary workforce development and competency matrix for NASA-aligned technology infusion proposals (Teaming & Workforce Development Section, Team 10), National Aeronautics and Space Administration (NASA), 2025.

12. **Terranova, J.**, Formal peer evaluation of competing NASA technology proposals via NASA-structured scoring methodology (Review Panel #6, 2025). Evaluated proposals from Teams 16, 17, and 18.

– Lead Systems Engineer: conducted formal peer evaluation of competing NASA technology proposals (Teams 16–18) using NASA-structured scoring methodology on Review Panel #6.

### *Personal Projects*

13. **Terranova, J.** (2025–Present). SpaAIder: A local-first multi-agent agentic runtime with tiered episodic-semantic-vector memory, dependency-aware parallel subagent scheduling, and heterogeneous client support (web, mobile, Electron, CLI, Code-OSS IDE) (Personal research software).

– Lead architect and sole developer: designed TypeScript orchestration plane, Python FastAPI memory substrate, 135-agent taxonomy, AsyncEventQueue SSE multiplexer, and Code-OSS IDE fork.

14. **Terranova, J.** (2025–Present). SpaAIder System Architecture Specification: orchestration plane, memory substrate, agent taxonomy, and multi-client SSE protocol for a local-first agentic runtime (Technical design document, personal research software).

– Sole developer: formal architecture specification documenting TypeScript/Python split-plane design, 135-agent taxonomy, AsyncEventQueue SSE protocol, EpisodicStore consolidation algorithm, ContextEngine token-budget allocation, and PostgreSQL/SQLite schema.

15. **Terranova, J.** (2025–2026). EfficientNetV2B0 transfer learning for chest radiograph pneumonia classification: focal loss, class-balanced oversampling, two-phase fine-tuning, and test-time augmentation (Personal project).

- Sole developer: designed training pipeline achieving 83% validation accuracy on 368-image held-out set with TTA and tuned threshold; exported model artifacts and structured evaluation reports.
16. **Terranova, J.** (2025–2026). LLM-V1.0: A ~30M-parameter GPT-style decoder-only transformer with BPE tokenization, AdamW optimization, and nucleus sampling for local and Colab training (Personal project).
    - Sole developer: implemented full training pipeline including mixed-precision training, gradient accumulation, cosine LR schedule, checkpoint resumption, and streaming text generation.
  17. **Terranova, J.** (2026). URC-CV: ROS 2 autonomy software stack for University Rover Challenge 2026 including 7-target mission orchestrator, EKF sensor fusion, GNSS waypoint navigation, ArUco visual servoing, YOLO detection pipeline, and judge telemetry dashboard (USC Field Robotics Lab research software).
    - Software lead: 7-target mission orchestrator, EKF sensor fusion, GNSS waypoint navigation, ArUco visual servoing, YOLO detection pipeline, and judge telemetry dashboard.
  18. **Terranova, J.** (2025–2026). ML Research Benchmark Suite: Ten reproducible research frameworks for causal inference, foundation-model adaptation, offline and safe reinforcement learning, probabilistic and Bayesian deep learning, scientific generative modeling, graph and temporal neural networks, self-supervised representation learning, neurosymbolic reasoning, federated and privacy-preserving ML, and scientific machine learning; ~33,000 source lines, ~5,300 test lines, 310 pytest cases (Personal research software).
    - Sole architect: ten reproducible frameworks with synthetic ground-truth DGPs, YAML benchmark runners, and pytest evaluation harnesses across causal inference, adaptation, RL, uncertainty, generative modeling, graph ML, SSL, neurosymbolic reasoning, federated privacy, and scientific ML.
  19. **Terranova, J.** (2025–2026). Causal ML for decision-making: meta-learners, doubly robust estimators, NOTEARS structure learning, PCMCI, and off-policy evaluation under synthetic confounding (Module 01 Technical Report, ML Research Benchmark Suite).
    - Module 01 technical memo: from-scratch implementations of NOTEARS linear and nonlinear structure learning (matrix-exponential acyclicity constraint), PCMCI time-series causal discovery, five meta-learners (S/T/X/R/Causal Forest), TMLE targeted learning, IPS/SNIPS/DR off-policy evaluation, and E-value sensitivity analysis; 71 passing pytest cases.
  20. **Terranova, J.** (2025–2026). Probabilistic and Bayesian deep learning: variational inference, conformal prediction, and calibration under distribution shift (Module 04 Technical Report, ML Research Benchmark Suite).
    - Module 04 technical memo: mean-field variational inference BNN with closed-form KL divergence, MC Dropout, deep ensembles, temperature and vector scaling calibration, split conformal prediction with RAPS, and reliability diagrams; systematic calibration degradation study under dataset shift.
  21. **Terranova, J.** (2025–2026). Federated and privacy-preserving ML: DP-SGD, Rényi differential privacy accounting, FedProx, and personalized federated learning under heterogeneous data partitions (Module 09 Technical Report, ML Research Benchmark Suite).
    - Module 09 technical memo: DP-SGD per-sample gradient clipping with vmap, Gaussian noise injection, Rényi DP accounting, FedAvg/FedProx/FedNova/FedBN, pFedMe Moreau envelope personalization, MAML-Federated, pairwise masking secure aggregation, and membership inference attack evaluation.