ARYAN GUPTA

(734) 819-3134 | guptaaryanr@gmail.com | LinkedIn | GitHub | Portfolio

EDUCATION

University of Michigan September 2019 – May 2023
Major: BSE Engineering Physics *Cum Laude*
Minors: Electrical Engineering, Computer Science, Entrepreneurship
Selected Coursework: Multivariable & Vector Calculus; Differential Equations; Matrix Algebra; Mathematical Methods in Physics; Honors Physics I-III; Advanced Mechanics, Electricity & Magnetism, Statistical & Thermal Physics; Data Structures & Algorithms; Logic Design; Signals & Systems; Embedded Systems (Intro & Advanced); Computer Architecture; Control Systems

WORK EXPERIENCE

Technip Energies – AI Solutions Initiative, Americas Dec 2025 – Present
Co-Founder, Co-Lead & Technical Lead, Internal AI Initiative

- Co-founded and co-lead an Americas-wide AI initiative across 4 operating centers and 8 engineering functions, scoping and piloting 5 AI workflows for document-heavy and data-heavy EPC tasks.
- Serve as technical lead for initiative architecture and project execution, translating high-value engineering pain points into reusable AI capabilities, MVP scopes, evaluation criteria, and implementation roadmaps.
- Defined a capability-first delivery strategy focused on reusable retrieval, extraction, document-intelligence, and structured-query components rather than one-off automations, enabling multiple EPC use cases to share common tooling, datasets, and evaluation patterns.
- Lead development of a Python LLM-assisted natural-language-to-SQL pipeline across 3 engineering datasets, translating engineering questions into executable structured queries for faster retrieval and analysis.
- Coordinated technical development of an LLM-assisted bid tabulation workflow that parses large unstructured vendor specification packages to extract, normalize, and compare requirement-level information, reducing manual review effort by ~80%.
- Build and coordinate a P&ID intelligence workflow for instrument, line, loop, label, connector, crossing, and cross-page relationship extraction using PDF/vector parsing, object/shape detection, OCR, and standards-aware validation.
- Established a repeatable AI delivery framework for EPC environments, incorporating agile sprints, design-verification-fitness checks, and Pugh decision matrices to prioritize, de-risk, and evaluate MVPs.
- Built and coordinated cross-disciplinary teams across multiple Americas operating centers, aligning parallel efforts to avoid duplication and enable shared tooling, datasets, and lessons learned.
- Partnered with engineering managers, discipline leads, and global digital stakeholders to select project portfolios, align with enterprise strategy, and define handoff paths for scalable MVPs.
- Drove internal visibility and adoption through presentations at internal events, town halls, and leadership forums, helping establish the initiative's credibility and project pipeline within months of launch.

Technip Energies Jul 2025 – Present
Specialist Engineer I – Instrumentation & Controls

- Standards & discipline onboarding: built working familiarity with ISA/IEC instrumentation and control standards, internal Technip specifications, and EPC deliverable workflows through training programs and document reviews.
- Blue hydrogen project support: reviewed vendor instrumentation documents for consistency against P&IDs, PFDs, and Cause & Effect diagrams; flagged discrepancies and supported senior engineers in document alignment and issue resolution.
- Controls documentation: assisted with verification and update of Cause & Effect diagrams by cross-referencing legacy logic, control narratives, and P&IDs to ensure functional consistency for brownfield scope changes.
- Carbon capture project exposure: contributing to early-stage I&C updates for pump addition scope, including review of control logic narratives and incremental updates to Cause & Effect documentation based on prior designs.
- Engineering rigor: applied version control, document traceability, and change-awareness practices while working within established management-of-change and QA processes.
- Used I&C project exposure to identify document-intensive engineering workflows suitable for AI-assisted retrieval, comparison, extraction, and validation, informing AISI project scoping and technical prioritization.

Technip Energies Aug 2023 – Jul 2025,
Feb 2026 – Mar 2026
Technology Developer I (Applied Physics Lead, Novel Carbon Capture R&D Team)

- Established a first-principles molecular modeling capability for internal carbon capture technology development; reproducible MD/DFT pipelines (LAMMPS, NWChem) across HPC and cloud (GPU + parallel I/O) reduced simulation wall time by ~2/3 and meaningfully lowered cloud spend (verified via internal benchmarks).
- Bridged atomistic → continuum: injected MD/DFT outputs into diffusion and CFD models with a nucleation-delay term; sensitivities to critical nucleus size improved continuum accuracy and re-prioritized model assumptions; maintained diffusion-scale codebase with new verification tests.

- Physical fidelity: evaluated force fields beyond LJ (Morse, EAM, TraPPE, SPC/E; polarizable) and benchmarked interfacial energies, diffusion rates, crystallization barriers vs bench + literature with acceptable error bounds; designed and supervised crystal growth experiments.
- Automation at scale: parameter sweeps across composition-T-P grids; continuous ingest of experimental images + tables so analyses auto-updated.
- Image analytics: Python/C++ (OpenCV) pipelines with QA segmented terabyte-scale image sets; produced UQ-backed particle-size distributions (KS/AD) replacing multi-day manual workflows with a multi-hour pipeline.
- ML integration: engineered physics-informed features; PCA/UMAP/k-means for structure discovery; RF/linear surrogates identified crystallization-rate drivers later confirmed experimentally; deployed surrogate-assisted and active-learning workflows to adaptively select simulation grids, reducing simulation count while maintaining target error bounds.
- Rigor & reproducibility: multi-level UQ (sensitivities, CIs, error propagation); authored 15 internal technical notes (two-reviewer process) and a data processing code archive forming the team's reproducibility library.
- Collaboration & leadership: integrated atomistics with MATLAB/Fortran utilities for stress and thermowell analysis in large-scale experiments; mentored teammates and onboarded a new engineer to extend the pipelines.
- Strategic impact: distilled findings into decision-ready decks supporting multi-million-dollar internal funding decisions; authored a 70-page market analysis & business plan that informed the team roadmap and funding discussions.

PREPRINTS, MANUSCRIPTS & OPEN RESEARCH SOFTWARE

Costgate: CI Cost Regression Gate for LLM Inference Jan 2026 – Present

Manuscript in Preparation *Code DOI: 10.5281/zenodo.20148498*

- Developed a deterministic CI-native framework for detecting LLM cost regressions using structured policies, statistical checks, and reproducible experimental artifacts.
- Designed a configurable policy schema for cost, latency, success-rate, and token-usage gates, including regression thresholds, minimum absolute deltas, sample-size checks, and variance-aware comparisons.
- Implemented a comparison engine for evaluating baseline vs PR inference runs, with repeat-run aggregation, confidence interval analysis, cost per success metrics, and failure-mode diagnostics.
- Built a Python CLI and library for baseline generation, regression comparison, contract validation, integrity checks, experiment orchestration, and Markdown/JSON reporting for CI workflows.
- Emphasized determinism, interpretability, and engineering reliability, enabling transparent inference-cost validation without black-box heuristics or ad hoc manual review.

Trajectory-Only Structural Regime Detection in Deterministic Dynamical Systems Feb 2026 – May 2026

Preprint: 10.5281/zenodo.20278240 *Code DOI: 10.5281/zenodo.19465797*

- Developing a trajectory-only framework for identifying structural regime changes in deterministic nonlinear dynamical systems from simulated state trajectories without supervised labels.
- Built an indicator atlas and structural scoring pipeline for detecting qualitative changes in dynamics, extracting regime boundaries, and comparing structural transitions across parameter sweeps.
- Benchmarked the framework on FitzHugh-Nagumo and van der Pol systems, including boundary extraction, lead-distance analysis, robustness tests, specificity checks, and ablation studies.
- Implemented reproducible Python experiment suites with deterministic ODE simulation, contract checks, integrity checks, and publication-quality plots for manuscript preparation.
- Designed the project as a non-ML scientific software framework emphasizing interpretability, reproducibility, and trajectory-level dynamical evidence.

ThermoBench-Consist (v1.0) Sep 2025 – Oct 2025

Preprint: 10.5281/zenodo.17489426 *Code DOI: 10.5281/zenodo.17330440*

- Designed a CPU-only diagnostic/benchmark to vet ML equations-of-state and VLE surrogates before they are plugged into CFD/combustion/process simulators.
- Formalizes four physics checks (C1-C4): ρ-p monotonicity, $\kappa_t>0$ stability, Clapeyron slope, and speed of sound (a^2) with configurable tolerances, critical-region exclusion, and near-spinodal flagging.
- Provides reference grids for CO_2 and N_2, deterministic sampling (and seeded random grids), and guardrails (phase filtering, critical-band avoidance) to reduce false positives in fragile regimes.
- Emits one-page MD/HTML reports with plots, severity badges (info/warn/fail), and a machine-readable JSON score; runs CLI <30 s and Colab notebook <60 s on a laptop (CPU-only).
- Ships a clean adapter API (finite-difference fallbacks, unit handling via pint) and a deliberately inconsistent surrogate to illustrate failure modes and expected report signatures.
- Includes tests/CI, tiny datasets, and reproducible example outputs to enable drop-in evaluation of third-party surrogates across labs.

REB-1: Robot Energy Benchmark (v0.1) Aug 2025 – Aug 2025

Code DOI: 10.5281/zenodo.17204853

- Built a WSL-friendly CLI micro-benchmark that logs power via nvidia-smi (or a deterministic demo source) and writes tidy CSV traces with reproducible 60-second workloads.

- Provides an analysis notebook that integrates power to Wh, converts to gCO_2 with a user-set grid-intensity factor, and exports bar charts and a short demo GIF for side-by-side comparisons.
- Focuses on power/impact, deliberately complementing existing ROS tools centered on latency/throughput; enables fast A/B of autonomy workloads on commodity laptops.
- Includes schema-checked outputs, lightweight tests/CI, and turn-key examples so other groups can replicate figures exactly and contribute logs.

Assessing the Limits of Graph Neural Networks for Vapor-Liquid Equilibrium May 2025 – Sep 2025
Preprint: 10.48550/arXiv.2509.10565
- Negative-results study on cryogenic mixtures: documents phase-dependent errors and liquid-phase inconsistency that prevent trained GNN surrogates from supporting VLE solvers.
- Traces failure to derivative pathologies (e.g., density/enthalpy behaviors across phases), demonstrating why seemingly accurate pointwise fits can violate global thermodynamic identities.
- Provides a diagnostic workflow (finite-difference checks, phase-aware tests, hybrid fallback with classical property libraries) and released failure cases to help other teams avoid silent breaks.
- Built phase-aware evaluation routines to test whether learned thermodynamic surrogates remained usable inside downstream VLE and solver workflows rather than only matching pointwise training targets.
- Reframed the project from prediction accuracy toward solver-safety and physical consistency, showing why apparently acceptable aggregate errors can still produce unusable scientific ML models.

Energy-Efficient Robotics Software (2020-2024): Systematic Literature Review Jan 2025 – Aug 2025
Preprint: 10.48550/arXiv.2508.12170 Code DOI: 10.5281/zenodo.16907564
- 79-study synthesis of post-2020 work on energy in autonomy stacks (planning, perception, middleware, hardware acceleration) across mobile and manipulator platforms.
- Proposes a taxonomy of energy-aware techniques (e.g., scheduling/DVFS, algorithmic refactoring, compute off-loading, task-level policy design) and a reporting checklist for fair energy claims.
- Quantifies coverage gaps (e.g., Wh/mission, energy-latency trade-offs, standardized workloads) and surfaces under-measured components (sensing, comms, OS services).
- Releases a full replication package (screening spreadsheet, code to regenerate figures, inclusion/exclusion rationale) to support future surveys and benchmark design in the community.
- Provides the analytical springboard for later software artifacts (e.g., REB-1's Wh→gCO_2 framing and quick-run workloads).

RESEARCH EXPERIENCE

BIRDS Lab, UMich Jan 2021 – Apr 2023
Research Assistant
- Built adaptable quadcopter HW/SW stack; implemented Wi-Fi ↔ UART bridges (ESP32/FTDI) for real-time control/telemetry; redesigned power boards; contributed to motor control reliability in a 5-person team.
- Supported integration and debugging across embedded electronics, flight-control software, telemetry links, and mechanical constraints for experimental aerial robotics platforms.

MiTEE (Miniature Tether Electrodynamics Experiment) Jan 2020 – Dec 2020
Plasma Team; Liaison to Structures
- Studied electrodynamic tether propulsion feasibility for pico/nanosats; supported thermodynamic analysis & electronics prototyping; coordinated interfaces with Structures team.

PROJECTS

Electric Longboard Remote Controller Sep 2022 – Dec 2022
- Designed a modular remote on STM32L010C6T6 with UART/LPUART smartphone link; added speed/turn indicators, music control, battery telemetry, and display updates. Modified VESC firmware (GPIO/PWM) and produced a 3D-printed enclosure; delivered under 3.5-month / $1k constraints.

Out-of-Order Processor Architecture Feb 2022 – Apr 2022
- Implemented a MIPS R10K-style OOO core in System Verilog with I/D caches, LSQ, and branch predictor; 2-way superscalar issue and precise exceptions; added baseline SMT (round-robin), cache optimizations, and a visual debugger for bring-up.

Low-Autonomy Follower Robot Feb 2022 – Apr 2022
- Built a vision + ultrasonic follower platform (Arduino/Nucleo, PID control) using LED-encoded commands; warehouse POC for leader-follower flows.

TEACHING EXPERIENCE

UMich College of Engineering Sep 2022 – Apr 2023
EECS 314 (Circuit Design and Analysis) Instructional Aide
- Led discussion groups for ~80 students; ran lab of 20 on DC/AC circuits, RLC transients, op-amps; shaped course structure for 160+ students and supported bench instrumentation (Keysight/HPE, scopes, DMMs).

ENTREPRENEURIAL EXPERIENCE

HelixCases Feb 2022 – Aug 2022

Co-founder & Software Developer

- Secured a $10k Lyft contract, designed a coding competition platform reaching 550+ students at 20 universities, built UI/backend systems, and hosted custom competitions.
- Led early customer discovery, product scoping, and competition operations for university-facing technical recruiting events.

LEADERSHIP & SERVICE

Citizens for Animal Protection (CAP) Dec 2023 – Nov 2024

Volunteer

- Assisted with dog walking, kennel cleaning, and shelter operations in Houston community.

StartUM Jan 2021 – Apr 2022

Member & Mentor

- Guided a pre-seed medical-volunteer startup through ideation, validation, and re-targeting of its core technology to a new market.

Project RISHI Sep 2020 – Apr 2021

Fundraising Co-Chair

- Filtered/applied to 25+ grants; secured $16k toward a $35k budget; organized campus fundraising events (raffles, trivia nights, crowdfunding, game nights) with 100+ participants.

SKILLS

- **Programming & Scientific Computing:** Python, C/C++, MATLAB/Simulink, Linux/Unix, Git, Docker, Conda, CI/GitHub Actions, Pandas, NumPy, SciPy, scientific software development
- **Machine Learning, Optimization & Evaluation:** PyTorch, scikit-learn, OpenCV, surrogate modeling, active learning, uncertainty quantification, numerical optimization, model calibration, design of experiments, benchmarking, evaluation harnesses
- **LLM, Document AI & Data Systems:** LLM/VLM workflows, natural-language-to-SQL, OCR, PDF/vector parsing, document intelligence, structured extraction, data validation, PostgreSQL, SQLAlchemy, large-scale data pipelines
- **Cloud, HPC & Reproducibility:** GCP, Vertex AI, MLflow, Kubernetes, Terraform, SLURM, GPU/cloud pipelines, profiling/optimization, reproducible cloud and cluster workflows
- **Scientific Modeling & Simulation:** LAMMPS, NWChem, MD/DFT, atomistic-to-continuum modeling, diffusion/CFD coupling, computational physics, thermodynamic consistency testing
- **Embedded, Robotics & Control:** STM32/ARM microcontrollers, ESP32/FTDI, UART, FreeRTOS, embedded prototyping, embedded systems integration, classical control/PID
- **Engineering Practice:** technical documentation, requirements analysis, engineering QA, version control, traceability, management-of-change awareness, cross-functional technical coordination

HONORS

- University Honors: 2019, 2020
- Dean’s List: 2019, 2020