

SHASHANK DEV

Patna, Bihar, India | coderpersion043@gmail.com | linkedin.com/in/shashank-dev-sciencebanda09
github.com/sciencebanda09 | mindofshashank.vercel.app

SUMMARY

Independent AI/ML researcher focused on reinforcement learning and safe decision-making under uncertainty. Designed CCPL (Causal Consequence-Penalized Learning), an original RL framework that penalizes agent actions by downstream consequence severity rather than reward alone — currently under review at NeurIPS 2026. Conducted this and related applied ML/robotics research entirely independently, without institutional lab access, compute grants, or academic mentorship, including a published, peer-reviewed paper (IJNRD, 2023). Seeking a SPAR mentorship to bring this independent research background into structured, mentored alignment work.

TECHNICAL SKILLS

ML / RL Research: Reinforcement Learning (original algorithm design), Reward Shaping, Constrained Optimization, Kalman Filtering (EKF, IMM-EKF), Computer Vision, YOLOv8, Multi-Object Tracking, Sensor Fusion, Simulation Design, NumPy

Programming: Python, Julia, JavaScript (Electron), Flask, Git/GitHub, NixOS

Robotics & Systems: ROS2, PX4, Embedded Systems, Real-Time Systems, RTOS, Edge Deployment (Raspberry Pi), Arduino, Autonomous Systems

Research: Technical Paper Writing, Peer-Reviewed Publication, Field Data Collection, Algorithm Design, Experimental Validation

EXPERIENCE

Co-Founder & CTO- Syntheta AI

March 2026 – Present

- Co-founded and lead technical development of an industrial data infrastructure platform that converts real factory, robotic, drone, and egocentric captures into labeled, validated synthetic training data for physical AI.
- Architected a 7-stage data pipeline restoration (NeuroDepth-T4/EventCore), controlled domain-gap mutation (Real2Robust), neural scene synthesis (FunctaGen, implicit neural representations), and target-aware detection/validation — with full source provenance and quality gating on every sample.
- Directed R&D on emerging 3D pipeline: working structure-from-motion reconstruction with Gaussian-splat and SMPL-X human-compositing in progress.

Founder & CEO- DmechatronicX

Aug 2024 – Feb 2026

- Founded and lead a defense-tech venture building AEGIS-X, an offline-capable counter-UAS decision layer that fuses multi-sensor input, prioritizes aerial threats, assigns effectors, and enforces rules of engagement on customer-controlled infrastructure
- Directed development of a 5-layer architecture spanning physics/environment modeling, sensor fusion (RF, visual, LiDAR, radar/ADS-B), a decision engine (Hungarian-algorithm effector assignment, custom NumPy-based threat scoring), a multi-stage safety layer, and fault-tolerant watchdog systems.
- Built the system for air-gapped, sovereign deployment zero cloud dependency, indigenous decision framework, HMAC-signed licensing and command authentication.

KEY PROJECTS

Causal Consequence-Penalized Learning (CCPL)

2025 – 2026

- Designed and authored an original constrained reinforcement learning algorithm (submitted to NeurIPS 2026, under review) that corrects three provably incorrect assumptions in standard Lagrangian constrained RL: fixed-discount TD targets under unknown consequence delay, reward-correlated (rather than causal) safety penalties, and non-stationary Bellman targets from a single mixed Q-function.
- Derived a delay-corrected Bellman operator that learns the distribution of an unknown consequence delay and proved it remains a contraction mapping (with a machine-verified guarantee at every training run), and proved a state-conditioned Lagrange multiplier strictly dominates any scalar dual variable by a closed-form gap.
- Built an Interventional Consequence Net using Pearl's do-calculus to separate an action's true causal contribution to a bad outcome from consequences already caused by earlier actions, preventing the agent from over-penalizing all actions in a "dangerous-looking" state.
- Implemented the full system (delay net, causal attribution net, dual Q-functions) in pure NumPy and benchmarked against 8 baselines (PPO, CPO, RCPO, PID-Lagrangian, SAC-Lagrangian, DQN, etc.) across 6 environments; CCPL was the only method achieving both high reward and 100% constraint satisfaction simultaneously
- Validated zero-shot transfer to 3 held-out adversarial environments (incl. a reward-hacking benchmark), where causal attribution alone prevented reward hacking that a non-causal ablation fully exploited

Genesis An artificial-life agent driven by pure intrinsic curiosity (<https://github.com/sciencebanda09/Genesis>)

- Designed and implemented a GRU-based reinforcement learning agent trained entirely on intrinsic reward (Random Network Distillation) with zero task reward, using a delay-corrected Bellman operator for stable value learning across 14,700+ updates with no divergence
- Built a battery of falsifiable verification scripts (pre-registered pass/fail bars) to test claims about world-model accuracy, concept formation, and exploration including a forward world model that predicted true transitions with 29–32× lower error than shuffled controls
- Implemented a meta-learning "self-editing" layer (SEAL, per Zweiger et al. NeurIPS 2025) using rejection-sampling RL to adapt training hyperparameters online, improving world-model accuracy by 23.7% and exploration coverage by 8.7%

Turbojet Digital Twin (<https://github.com/sciencebanda09/Turbojet-Digital-Twin>)

- Built a production-oriented digital twin for a four-stage turbojet engine, fusing a physics-based Brayton-cycle thermodynamic model with a learned ML surrogate through Bayesian state estimation (EKF/UKF/particle filter) to produce calibrated health trajectories.
- Implemented uncertainty-aware prognostics (conformal prediction, MC-dropout, ensemble calibration) to generate remaining-useful-life (RUL) quantiles, failure-probability estimates, and thrust/fuel-efficiency forecasts from streaming sensor data.
- Designed a fault-injection and root-cause-analysis pipeline (compressor fouling, turbine erosion, bearing wear, sensor drift/bias) that propagates injected faults through the physics model, estimator, and health/RUL pipeline to validate the twin's diagnostic sensitivity.

SHWFS Adaptive Optics Pipeline(<https://github.com/sciencebanda09/shwfs-ao-pipeline>)

- Built an end-to-end adaptive optics pipeline (Shack-Hartmann wavefront sensing → reconstruction → deformable-mirror control) implementing classical (SVD, MMSE/Bayesian) and CNN/U-Net-based wavefront reconstruction, achieving Strehl ratio 0.9986 with 2× lower RMS wavefront error (37.1 nm vs. 75.1 nm) than classical SVD reconstruction.
- Implemented atmospheric turbulence profiling (SLODAR $C_n^2(h)$ profiling, Fried parameter r_0 and coherence time τ_0 estimation via Von Kármán PSD fitting) and an LQG (Kalman filter + LQR) controller alongside an LSTM-based temporal predictor for closed-loop correction.
- Modeled the physical system end-to-end: Zernike polynomial wavefront decomposition (Noll ordering), Kolmogorov/Von Kármán turbulence statistics, Fried-geometry actuator placement, and Gaussian deformable-mirror influence functions.

EDUCATION

Himalayan Public School

2024-2026

Class XI-XII (Secondary Education) Patna, Bihar. Building and shipping production robotics and AI systems concurrently with coursework since 2021.

Gyan Niketan

2012-2024

UKG to X Patna, Bihar. Building and shipping production robotics and AI systems concurrently with coursework since 2021.

PUBLICATIONS

- “Smart Safety Protector: A Comprehensive Approach to Enhancing Personal Safety for Women” International Journal of Novel Research and Development (IJNRD), Dec 2023.
- “Constrained RL with Causal Consequence-Penalized Learning (CCPL)” Submitted to NeurIPS 2026, under review.

HONORS & AWARDS

- World Champion, Drone Rescue Technoxian 2024 (world's largest robotics championship, New Delhi); Bihar's first world-stage robotics title.
- 2nd Place, Entrepreneurship Advanced Robotex India 2025 (100+ teams); qualified for Robotex Global 2025, Estonia.
- 2nd Rank, TECHNOXIAN Innovation Challenge (2025)
- Best Durability Award, Blixathon, Techfest IIT Bombay (2024).
- India Book of Records, INSPIRE Award, IGNITE Award for Smart Safety Protector, 2023.