SREE DHARSHAN G J

AI Research Engineer | Reinforcement Learning · Multi-Agent Systems · Memory-Augmented AI

Chennai, India | Email | GitHub | LinkedIn | Google Scholar | ResearchGate | ORCID

SUMMARY

Electronics and Computer Engineering graduate (B.Tech, 2026) specializing in Reinforcement Learning, Multi-Agent Systems, Memory-Augmented AI, and Deep Learning. IEEE Best Paper Award recipient and first author; Springer book chapter author. Builds and evaluates ML/DL systems in PyTorch and TensorFlow across computer vision, NLP, and multi-agent control, with hands-on experience in REST API development and a track record of rigorous, statistically validated experimentation.

AWARDS & PUBLICATIONS

IEEE Best Paper Award (2026), First Author — "AI-Based Event-Triggered Decentralized Swarm Coordination for Multi-Vehicle Collision Avoidance," IEEE RAEEUCCI 2026. D**OI:** 10.1109/RAEEUCCI67649.2026.11504821

Springer Book Chapter (Accepted, In Press) — "Testing and Using Benchmark Datasets for YOLOv7 Deepfake Detection," Springer LNEE Series

Best Project Award (2022–23) — Solar Power Bank and Wireless Charging, Techknow, SRMIST

EDUCATION

SRM Institute of Science and Technology, Chennai *Graduated May 2026*

B.Tech, Electronics and Computer Engineering — CGPA: 7.84/10.0 | Final Semester SGPA: 10.0/10.0

Relevant Coursework: Deep Learning (A+), Python Programming (A+), Machine Learning, Reinforcement Learning, Data Structures

TECHNICAL SKILLS

Languages: Python, C++, C, Java, SQL

ML & Deep Learning: PyTorch, TensorFlow, Transformers, Reinforcement Learning, Computer Vision, NLP, Multi-Agent Systems

Frameworks & Tools: FastAPI, Flask, React, OpenCV, NumPy, Pandas, Git

Core Concepts: Optimization, Distributed Systems, Control Systems, Data Structures, Algorithms

RESEARCH EXPERIENCE

AI-Based Event-Triggered Decentralized Swarm Coordination for Multi-Vehicle Collision Avoidance *Jan 2026 – May 2026*

- Reduced inter-agent communication overhead by 92.5% and computational cost by 84%, measured across simulation runs, by designing a novel event-triggered mechanism for decentralized multi-agent PSO-DMPC coordination.
- Achieved zero collisions across 10,000+ simulation timesteps, improving the MTTC safety metric by 90.5% and reducing speed variance by 81.9% over baseline methods.
- Validated robustness by testing the system across multiple traffic density configurations and parameter sensitivity scenarios.

Technologies: Python, SUMO, Particle Swarm Optimization, Distributed Model Predictive Control

YOLOv7 Deepfake Detection — Benchmark Dataset Analysis *Sep 2025 – Jan 2026*

- Achieved 92% accuracy and 0.958 AUC on FaceForensics++ and FakeAVCeleb by building a YOLOv7 + ResNet-50 + Vision Transformer pipeline with spatio-temporal attention modules.
- Reached 25 FPS real-time inference by optimizing the pipeline and reducing compute overhead, while evaluating generalization across multiple forgery datasets.

Technologies: PyTorch, Computer Vision, Transformers, Deep Learning

PROJECTS

Memory-Aware LLM Agents (CARM++) *2026*

- Engineered a rigorous multi-method evaluation protocol (warm-up phase, instrumentation probes, random-context control) to isolate genuine memory-retrieval gains from confounds like prompt length and token budget.
- Diagnosed and resolved subtle generation-truncation and prompt-leakage failure modes through systematic per-task error analysis, improving evaluation reliability across 7 experimental conditions.

Technologies: Python, PyTorch, FAISS, Sentence-Transformers, LLM Agents, Statistical Validation

Memory-Aware Mean Field Reinforcement Learning (MAMF-RL) *2026*

- Achieved statistically significant gains over mean-field baselines (MFAC, MFQ) across two multi-agent environments, with large effect sizes (Cohen's d up to 4.9, $p < 0.01$) validated via paired t-tests across 5 seeds.
- Conducted a generalization study across varying agent counts (N=2–16), showing memory-augmentation gains grow to +26% at smaller N, with effect sizes confirmed via Cohen's d and 95% CIs.
- Benchmarked computational overhead at scale (N=10–200), demonstrating sub-linear cost growth (O(N^0.33)) and a practical ~2–3x overhead versus non-memory baselines.
- Empirically validated a theoretical convergence proposition (O(1/√N) mean-field error bound), confirming the fitted scaling exponent and linear error contribution from approximate retrieval ($R^2 > 0.99$).

Technologies: Python, PyTorch, Reinforcement Learning, Multi-Agent Systems

Memory Utility Networks (MUN) *2026*

- Achieved Recall@5 = 0.829 and NDCG@5 = 0.914 with the label-aware utility model, significantly outperforming 7 baseline retrieval strategies (cosine similarity, recency, frequency, TF-IDF) with large effect size (Cohen's d ≈ 28.8, $p < 1e\text{-}10$).
- Conducted a rigorous ablation and ranking-loss study isolating the source of retrieval gains, quantifying a 34-percentage-point contribution from label access and documenting a negative result for label-free utility estimation to guide future research direction.

Technologies: Python, PyTorch, NLP, Information Retrieval