

RAMCHAND KUMARESAN

AI Researcher & Program Leader · LLM Systems · On-Device Inference · Enterprise-Scale Delivery

PMP, PgMP | Liberty Hill, TX | mechramc@gmail.com | linkedin.com/in/ramchandeaswar | murailabs.com

PROFESSIONAL SUMMARY

AI researcher and program leader who pairs hands-on machine-learning research with a 16-year record of enterprise-scale delivery. Author of five research papers (arXiv, NeurIPS 2026, TMLR) on large-language-model systems — spanning direct Apple Neural Engine programming, cooperative model training, multi-model orchestration, and evolutionary fine-tuning — with headline results including a **+21.76% cross-lingual model-fusion gain** and an **8.5× Apple Neural Engine compilation speedup**, every result reproducible from committed code. Builder of production AI systems, from a live personalized-storytelling platform to a clinical-triage assistant that runs a 5GB language model fully offline on a phone. This depth sits on a foundation of dual PMI certifications (PMP, PgMP) and \$25M+ multi-state technology programs delivered on time across cloud, SaaS, and government domains.

EXPERIENCE

Founder & AI Researcher — Murai Labs

2025 – Present · Austin, TX

- Authored five AI research papers (arXiv, NeurIPS 2026, TMLR) on large-language-model systems — spanning direct Apple Neural Engine programming, cooperative model training, multi-model orchestration, and evolutionary mixture-of-LoRA fine-tuning.
- Shipped production AI products end-to-end: **StoryForge Kids**, a live iOS + web personalized-storytelling platform, and **Marunthagam**, a Tamil-first clinical-triage assistant running a 5GB language model fully offline on an Android phone (0/12 missed emergencies, 100% safety-refusal on adversarial prompts).
- Built **PINNAL**, a genomic-foundation-model RNA-therapeutic candidate ranker (Evo 2, Borzoi, Enformer, ViennaRNA) — in a pre-registered SMN2 / ISS-N1 falsification test it surfaced all 3 clinically-validated ISS-N1 tiles in its top 10 (enrichment@10 4.10 vs 0.00).
- Delivered headline research results — +21.76% cross-lingual model-fusion gain (KALAVAI) and 8.5× Apple Neural Engine compilation speedup (Orion) — each reproducible from committed code and audit-tracked artifacts.
- Established disciplined research methodology — pre-registered hypotheses, frozen-hash provenance, and cheap-baseline falsification — publishing negative results with the same rigor as positive ones.

Enterprise Project Manager — Procore Technologies

Dec 2025 – Present · Austin, TX (Remote)

- Manage enterprise project delivery for public sector and government clients; lead cross-functional coordination across engineering, product, customer success, and compliance.
- Drive change management using SDLC methodology with zero disruption to government operations; partner with the Change Advisory Board for FedRAMP / SOC 2 alignment.
- Facilitate Agile sprint planning for distributed teams — 15% reduction in release cycle time.

Senior Project Manager, PMO — Aspira

2017 – 2025 · Dallas, TX

- Led multi-year, multi-state modernization programs (Utah, Texas, Kansas, NJ, Alberta, Maryland) valued over \$25M, delivering SaaS platforms for park reservations, licensing, and customer engagement.
- Directed enterprise API integrations connecting DMV, DCF, banking, and payment systems, enabling real-time data exchange and supporting \$40M+ revenue in year one.
- Led the enterprise migration to AWS (EC2, S3, Redshift) — 30% cost reduction, 99.99% uptime; migrated 600K+ customer records across 28 parks with zero data loss.
- Deployed AI-driven conversational chatbots — 35% reduction in call volume, \$250K+ annual support savings; established the PMO and standardized program management practices.

Senior Project Manager — Doodleblue Innovations

2014 – 2017 · Chennai, India

- Delivered mobile-first iOS/Android solutions for finance, retail, and services; founded an in-house product division and mentored 10+ project managers.

Enterprise Release Project Manager — Cognizant

2012 – 2014 · Hartford, CT

- Supervised enterprise application releases (Real-Time Eligibility, Claims Adjudication) via SDLC; drove continuous delivery for pharmacy claims and directed a team of 3 release managers.

AI RESEARCH & PUBLICATIONS

Orion — arXiv 2603.06728

Programming Apple's Neural Engine: the first open end-to-end system for direct ANE inference and training. → **8.5× ANE compilation speedup; bypasses CoreML entirely.**

KALAVAI — NeurIPS 2026 (under review) · arXiv 2603.22755

Predicting when independent specialist fusion works in cooperative LLM training. → **+21.76% cross-lingual fusion gain; gain predicted before compute is spent.**

ACAR — arXiv 2602.21231

Adaptive complexity routing for multi-model ensembles via self-consistency variance. → **1,510 tasks across 4 benchmarks, 7,550+ auditable runs.**

UYIR — TMLR 2026 (under review) · arXiv 2605.11153

Decomposing evolutionary mixture-of-LoRA architectures with a controlled factorial. → **Rigorous attribution and honest-negative-result methodology.**

Rigel — arXiv 2606.12765

Reverse-engineering the Apple M4 Max Metal 4.1 tensor compute path. → **11 hidden hardware specs recovered; +12.9% hand-fused GEMM kernel.**

KEY PROJECTS

StoryForge Kids — Live personalized children's storytelling platform (iOS + web) with a multi-pass story engine and persistent character memory. **Shipped end-to-end with real users.**

PINNAL — RNA-therapeutic candidate ranker prioritizing antisense-oligo interventions for splice-altering variants via an ensemble of genomic foundation models (Evo 2, Borzoi, Enformer). **Surfaced all 3 ISS-N1 tiles in its top 10 (enrichment@10 4.10 vs 0.00).**

Marunthagam — Tamil-first clinical-triage assistant running a 5GB LLM fully offline on an Android phone for rural community-health workers. **0/12 missed emergencies; 100% safety-refusal.**

KALAVAI — Cooperative LLM-training protocol that fuses independently trained specialists and predicts the fusion gain before any compute is spent. **+21.76% cross-lingual gain; NeurIPS 2026.**

Multi-State SaaS Platform — Enterprise reservation/licensing system across six jurisdictions. **\$25M delivered on time; zero government-ops disruption.**

SKILLS & EXPERTISE

AI / ML: LLM Fine-Tuning (LoRA/QLoRA) · PyTorch · HuggingFace Transformers · Mixture-of-Experts Routing · Model Quantization · On-Device / Edge Inference · Apple Neural Engine · RAG · Model Evaluation · Prompt Engineering · AI Product Strategy · MLOps & Reproducibility

Cloud & Data: AWS · SQL · PLSQL · API Integration · RDBMS

Program & Delivery: Agile · Scrum · SAFe · SDLC · Six Sigma · JIRA · Confluence · Microsoft Project · PMO · Change Management · Stakeholder Management · Risk Management

CERTIFICATIONS & EDUCATION

- Project Management Professional (PMP) — PMI
- Program Management Professional (PgMP) — PMI
- Business English Certification — Cambridge (Class A)
- MS, Industrial Engineering & Management — Florida A&M, 2010
- BE, Mechanical Engineering — Anna University, 2006