

KAUSHIK KUMAR



kaushikkumar.208@gmail.com

[Portfolio](#)

[LinkedIn](#)

(520) 609-4957

[GitHub](#)

Dear Hiring Committee,

Patterns don't just live in the data, they hide in assumptions, in workflows, in what's not tracked yet. That's why the best data scientists aren't just modelers; they're translators between ambiguity and structure, context and code. That's how I approach the work: not just as a technician, but as someone who builds trust in data by treating it with precision and respect.

I'm **Kaushik Kumar**, a Master's student in Data Science at the University of Arizona. My training is in machine learning, AI systems, and large-scale data engineering, but my perspective was shaped by a personal experience, losing my father in a car accident caused by a faulty vehicle during my undergrad. It taught me that the systems we build have real consequences. Why a predictive maintenance model saves costs in theory but fails under noisy sensor data. Why a reinforcement learning agent succeeds in simulation but crashes when safety constraints shift. These questions led me deeper into reproducibility, fault tolerance, and the human context of AI systems.

Some of the most critical lessons I've learned came not from models that worked, but from models that shouldn't have. At **Johnson Electric**, I designed AI-powered quality assurance pipelines, reducing test times by 72% while learning how fragile assumptions could derail factory floor deployment. At **Selector AI**, I replaced legacy models with sparse retrieval algorithms (BM25S), discovering how scalability demands force different tradeoffs than research benchmarks. At the **ACT Lab**, I built **ShieldNN-AM**, a provably safe reinforcement learning framework that achieved zero collisions in CARLA simulations, showing me how theory and safety guarantees must align in practice. And as **CTO of Kamuit**, I architected a detour-aware ridesharing platform, where real-time reliability and user trust mattered as much as algorithmic performance.

Generalization has never just meant cross-validation to me, it's meant understanding whether this model makes sense when the baseline shifts, the incentives change, or the edge cases become the new norm. That's why my work always begins with a question: not just can we predict this, but should we? So I've built anomaly detection systems that flag blind spots, reinforcement learning pipelines stress-tested against safety margins, and data platforms designed with monitoring, retraining, and feedback loops baked in from the start.

What I prioritize isn't just model performance at the moment, but the staying power of the entire system, building AI products that remain understandable, safe, and genuinely useful long after deployment, even as the context shifts and new hands take over.

Thank you for considering my application. I'd be thrilled to bring this mindset to a team that values thoughtful modeling, practical impact, and technical elegance.

Sincerely,
Kaushik Kumar