

Karthik Rajeev

Nottingham Two, Midland Way, NG7 3LP

+44 77 99 25 8301 | karthikpr10@gmail.com

Website: <http://www.karthikrajeev.me/> | LinkedIn:

<https://www.linkedin.com/in/karthik-p-r-728843194/>

Professional Summary

I deploy architectures into environments where theoretical models collide with operational reality. In high-stakes operations, a product is a mediating tool designed to drive outcomes through human-machine teaming. True system resilience requires mutual alignment: the operator must understand the exact capabilities and limitations of the tool, while the system architecture operates as a cognitive safety net. It must bypass the risk of human blindsight errors, while actively unlocking and augmenting the domain expertise of the user in the loop.

Education

University of Nottingham, MSc Human-Computer Interaction | 2025 - 2026 (Present)

- **Core Focus:** Simulation Optimisation for Decision Support, Cognitive Ergonomics in Design, Design Ethnography, Advanced Human-Machine Interfaces, Human Factors & Performance and Statistical Analysis encompassing ANOVA, t-tests, Mann-Whitney U, Wilcoxon Signed-Rank, Friedman, Kruskal-Wallis, Chi-square, and further more.

National Institute of Fashion Technology, Bachelor of Design | 2019 - 2023

Technical Implementation & Tools

- **Data Analysis & Simulation:** Data Engineering & Analytics (R studio), Discrete Event Simulation, Agent-Based Modelling, Optimisation (Anylogic), Statistics (IBM SPSS).
- **Research & Evaluation:** Design Ethnography (Workflow Immersion), Mixed-Methods Research (Quantitative & Qualitative), Post Mortem Data Triangulation, Heuristic Evaluations, Contextual Inquiry, Stakeholder Analysis, Operations Research.
- **Product Telemetry & Workflows:** : Figma, Jira, Atlassian Suite, Microsoft Clarity, Google Analytics.

Experience

Makerble | UX Design Intern | 2024-2025 | 8 Months

- Implemented UI/UX fixes and interface improvements within an Agile/Jira sprint workflow, translating Product Manager requirements into shipped design changes.

Applied Research & Case Studies

Clinical EPR Decision Infrastructure | User Research & Systems Design

- Diagnosed a structural flaw in clinical medication software: patient data was presented as static, retrospective ledgers instead of interdependent, time-relevant information - root cause of recurring workflow and safety failures.
- Designed and built a prototype restructuring medication history around an "atomic data model" - a spatial past/present/future timeline exposing dependencies between data points instead of isolating them, cutting forced context switches from 7 to 0 per workflow.
- Ran a counterbalanced within-subjects study (n=12 Clinical domain experts) against the incumbent EPR interface, applying McNemar's exact test and Wilcoxon signed-rank tests across 3 objective and 4 subjective measures - total task time dropped 61% (352s → 136s) with error rates reduced from (66.7% → 8.3% max), and NASA-TLX cognitive load fell 56% (52.9 → 22.9).
- Identified "Semantic Abstraction" in legacy systems - patients reduced to mechanistic data points - and designed a "Narrative Anchor" intervention (synthetic photo + psychosocial tag) to test whether contextual data could close the empathy gap; 25% of clinicians adopted it to actively personalise care, while qualitative probing revealed the remaining 75% were blocked not by the product feature but by resource-constrained, throughput-optimised ward environments that penalise time spent on bedside rapport.
- Triangulated objective and subjective data to isolate a single-participant outlier - a case where performance and preference sharply diverged, undetectable in aggregate results - and used it as a targeted probe to surface a broader "Preference vs. Performance Paradox" hypothesis for future validation, without overgeneralising from a single case.
- Authored the production architecture spec: PostgreSQL/JSONB data model with a rule-based "Action Engine" middleware to auto-propagate clinical dependencies (e.g., flagging a dependent dose when a parent dose is missed), including contraindication guardrails, OCC-based concurrency control, and FHIR-compliant external alerting. (Design proposal - not built.)

Sparse Network Anomaly Detection | Data Engineering & Visualisation Project

- Built a fault-tolerant ETL pipeline in R to clean and structure a highly sparse knowledge graph - 75,000+ identifiers in inconsistent formats and 37,857 untyped edges - using regex-based parsing and boolean normalisation to produce a stable, joinable schema before analysis.
- Built an anomaly-scoring model using glmnet as the underlying regression framework, oversampling rare cases to correct for an 11,360:1 class imbalance without distorting the overall signal.

- Engineered composite features for "popularity" and "structural influence" from raw edge data, evaluating 13,000+ entity relationships via custom heuristics to convert sparse connections into actionable, exportable scores.
- Built a D3.js frontend to make the scores usable: a time-series view comparing 40-year career trajectories, and a log-scale dot plot isolating the top 5 statistical outliers.

Resource Load Equilibrium Simulation & Optimisation | Operations Research

- Built a hybrid agent-based/discrete-event simulation (AnyLogic) of a two-node service network where fixed staffing (40 total) was structurally mismatched to a 60/40 spatial demand bias, with the majority of capacity positioned furthest from where demand actually originated.
- Held staffing constant and varied only the adoption rate of a routing-information layer (0%, 50%, 80%) across 1,000 replications per phase, isolating whether the failure was caused by information flow rather than physical capacity constraints.
- Found a non-linear relationship between adoption and stability: 0% adoption produced 140.07 mean SLA breaches; 50% adoption cut breaches 82.9% (23.91) but retained high variance, indicating the system was still cascading under stochastic pressure; 80% adoption crossed a stability threshold, cutting breaches 94.8% (7.31 mean) and reaching load equilibrium with zero added staff or restructuring.
- Identified residual brittleness at the 80% stable state - variance ($\sigma=12.68$) exceeded the mean, revealing a bimodal outcome where stochastic shocks (e.g. sudden weather events) forced agents back into default proximity-based behaviour, temporarily destabilising the network. Concluded this required a structurally different resilience mechanism, not a resourcing fix.
- Ran a constrained optimisation (OptQuest, 500 iterations $\times$ 100 replications) at 90% adoption to find the cost-minimal staffing configuration - minimising wages plus a £50-per-SLA-breach penalty, subject to capacity caps and hard wellbeing constraints (frustration/burnout capped at 0.3, preventing cost savings through staff exhaustion) - arriving at a feasible optimum of £5,778/shift.