

Rohit Sharma

Copenhagen, Denmark
+45 50128907

rs@rohit.sh
rohit.sh
github.com/r0hitsharma
linkedin.com/in/r0hit

Software engineer working across the stack on data-heavy products — the inspection and visualisation interfaces, the services that feed them, and the platform underneath.

Experience

Archon Tech / Senior Engineer

Nov 2025 – Present · Copenhagen, Denmark

Joined as a frontend specialist and grew into full-stack and platform ownership across a research-led product group building on-chain financial-risk infrastructure. Work spans four distinct areas, all shipped as production software rather than prototypes.

Data platform and risk modelling. Contributed across a Go/Python/TypeScript monorepo that indexes Ethereum and L2 state in real time, tracks DeFi lending and DEX positions, and computes required risk capital for institutional allocators. Built a second independent data path that reconciles the platform's own on-chain figures against a counterparty's self-reported balance sheet, replacing per-request live HTTP integrations with a properly indexed pipeline and eliminating a class of silent-degradation bugs. Rolled a Monte-Carlo risk model into production, including the capacity study — peak resident memory measured across nine markets — that set its pod sizing. Root-caused and fixed two production incidents others could not explain: a TimescaleDB hypertable recompressed without segment keys, defeating chunk pruning and pegging CPU on every query; and a connection pooler OOMing mid-transaction, leaving database connections desynced in a state standard health-checking could not detect.

Frontend architecture and a shared design system. Designed and built a TypeScript/React toolkit published as fifteen npm packages — a Panda CSS design system on headless primitives with a semantic token contract and AA-verified dark mode, a visx-based charting layer with cross-chart cursor synchronisation and large-series downsampling, a virtualised data grid, a schema-driven dashboard engine, and HTTP client and mock packages whose types are both derived from the same generated OpenAPI contract so mocks cannot drift from the client. Ran repeated bundle-size forensics on the result, measuring every module in isolation to find that one library barrel import cost 82 kB minified on its own, and restructuring the package into lazy tiers — so that a consuming application no longer had to hand-maintain its own chunk map to get its entry chunk down from 1,004 kB to 80 kB.

Infrastructure, observability and reliability. Built an AWS observability and paging stack end to end for a production conversational-AI service: CloudWatch alarms and metric filters derived from both request metrics and structured application logs, a Slack incident-paging pipeline on Lambda and Step Functions with re-notification and thread grouping, synthetic uptime canaries, and a self-hosted Grafana platform migrated off a managed offering once its sharing and version constraints blocked a business need. Tuned alarm thresholds against the observed log-normal latency distribution rather than naive percentiles. On the shared Kubernetes/GitOps platform, provisioned durable storage and workload identity for new tenants, sized autoscaling capacity with worked arithmetic, and governed alert noise by root-causing the underlying deploy behaviour. Repeatedly found alerting that looked correct but structurally could not fire — in one case quantifying that a staleness alert had been silently disarmed for weeks before rewriting it.

Language and systems research. Contributed to the core engine of a from-scratch logic and term-rewriting language — pattern matching, unification, non-deterministic rewriting under a typed discipline — including new language-level semantics for the boundary between data and executable code. Separately designed and built the compiler front end for a domain-specific financial-specification language: a parser compiling a restricted subset of a general-purpose language into symbolic-math expression graphs, a code generator emitting a typed importable client, and a graph-database-backed API and web UI for browsing the resulting knowledge graph with branch history, diffing and dependency visualisation. Built and operated an event-sourced simulation platform to stress-test that engine at production-like scale, including checkpointed boot, idempotent replay and a real-time analytics dashboard whose incremental fold cut per-event cost by roughly five orders of magnitude.

Engineering practice. A consistent method runs through all of it: measure before and after and publish the numbers; reproduce a bug before fixing it; sabotage-test a fix by reintroducing each historical defect to confirm exactly the intended test goes red; and treat tooling — build reproducibility, CI gates, dev environments, AI coding-agent configuration — as infrastructure held to the same correctness bar as the product.

Most work at Archon Tech is under NDA and is described here at the level of the problem rather than the product; the projects below are the parts that are public.

SELECTED PROJECTS

UI Kit — A shared front-end toolkit for a company whose products all show the same kind of thing — positions, balances, risk, movement over time — ...

Sentinel Verify — A dashboard over a platform that watches Ethereum and its layer twos in real time, indexes what lending protocols are holding, and works...

Chatbot for sky.money — An AI assistant that shipped inside a DeFi application, answering questions about a user's positions and walking them through actions.

Edge Delta / Senior Front-end Engineer

Apr 2023 – Oct 2025 · Copenhagen, Denmark

Led the transition of the core product from Angular to React, and the adoption of a shared component library to go with it.

Extended Metrics Explorer — a Grafana-like query and charting surface — with support for multiple queries and formulas, and redesigned it around the new component system.

Built **Service Map** and **Metrics Inventory** for introspection into deployment architecture and data cardinality, and worked on the **Monitors** and **Composite Monitors** overhauls.

Delivered **Kubernetes Metrics** monitoring, shipped with a walkthrough demo.

SELECTED PROJECTS

Service Map — A live node-link graph of service-to-service traffic, drawn from kernel-level observation rather than instrumentation, with error rate and...

Grammars & Language Support — Two ANTLR grammars and the editor built on top of them: a templating language for monitor notifications where the conditionals decide who...

Monitors — An overhaul of how a monitor collection is managed, splitting one overloaded screen into surfaces for status, management and event history...

Metrics Explorer — A query and charting surface over every metric Edge Delta ingests or derives, rebuilt on the component system that came out of the React...

Churney / Lead Front-end Developer

Oct 2022 – Feb 2023 · Copenhagen, Denmark

Worked directly with the founders to establish a web platform and dashboard giving clients self-service access to their models and insights.

Built embeddable interactive visualisations used for marketing and outreach.

DataRobot / Senior Front-end Developer

May 2019 – Sept 2022 · Copenhagen, Denmark

Joined the Geospatial Domain at an early stage to develop the inspection and visualisation capabilities that were integrated into the wider application.

Later joined the Batch Predictions team to build a job designer and scheduler for recurring prediction jobs, integrating with every supported data source and providing build-time validation and run-time error tracing.

SELECTED PROJECTS

Drift Drill Down — Drill Down chart visualizes the difference in distribution over time between the training dataset of the deployed model and the datasets...

Job Definitions UI — Job definitions allow users to create a template for prediction jobs which can be invoked via a trigger or run on a schedule.

Geospatial Domain — We built capability into the DataRobot AutoML platform to help users explore patterns in their data using dynamic map visualizations.

Leapcraft / Data Visualisation & Full-stack Development

Jan 2015 – Apr 2019 · Copenhagen, Denmark

Built bespoke visualisations for consulting projects with large client data sets, generally using D3, Leaflet and Highcharts.

Grew into leading development and architecture of the cloud infrastructure for data ingestion, model execution and insight delivery on AWS — consuming real-time data from in-house sensor platforms and running ML models for interpretation and calibration.

SELECTED PROJECTS

Ambinode — AmbiNode is an automated sensing as a service solution which helps you monitor and manage your indoor ambience, integrating a wide range of...

Flight Sense — An analytics service to help optimize corporate travel.

Bike Maps — The city of Copenhagen conducted an experiment to detect the flow of bicycle traffic along some important corridors in the city.

Creative Clusters Map — Based on the Danish Central Business Register records (CVR) we build a map that explored the geographic clustering of companies from the...

Freelance / Creative Technologist

Jan 2013 – Dec 2014 · Copenhagen, Denmark

Consulted on several projects specialising in data visualisation and user experience, for clients including Leapcraft, KL7 and Kirt Thomsen — the first of which became a full-time role in 2015.

Also worked on Appeer, which received project funding through LaunchPad Denmark.

SELECTED PROJECTS

Wind Power In Denmark 1977-2020 — WIND ELECTRICITY IN DENMARK 1977-2020 shows patterns of wind turbines installed – and taken down – in Denmark over the years in combination...

Vores Fart (Our Speed) — We were given the task to build a data visualization based on three speed measurement spots around Denmark that would engage the media and...

HITLabNZ / Research Assistant

Sept 2011 – Dec 2011 · Christchurch, New Zealand

Worked on augmented reality projects on the Android platform, alongside visual design for the lab's output.

SELECTED PROJECTS

CityViewAR — CityViewAR is a mobile AR (augmented reality) application that allows people to see how the city was before the earthquakes and building...

Purple Leaf Design / UX Designer & Developer

Nov 2009 – Dec 2009 · Guwahati, India

Worked on branding and identity alongside web development for the studio's clients, with competitive research feeding both.

Zensar Technologies / Usability Engineering Intern

May 2009 – Jul 2009 · Pune, India

Interned with the Usability Engineering Group, producing wireframes, use cases and information architecture for enterprise software, and contributing to heuristic evaluations, personas and user interviews.

Tools & Frameworks

Languages — TypeScript, Python, Go, SQL

Front-end — React ecosystem, TanStack Query & Router, Panda CSS, Ladle, D3 / visx / Deck.gl

Back-end & data — Node.js, FastAPI, PostgreSQL, TimescaleDB, OpenAPI contracts

AI & agent tooling — MCP / WebMCP, Claude Code / Codex / Copilot CLI, Agent skills & routing, LangChain / LangSmith, LLM tracing & safety metrics

Platform & practice — AWS, Kubernetes, Terraform / OpenTofu, Grafana, Playwright, Vitest

Education

CIID / Interaction Design Programme

Jan 2012 – Dec 2012 · Copenhagen, Denmark

Interaction and Service Design

Department of Design at the Indian Institute of Technology Guwahati / Bachelor of Design

Jul 2007 – May 2011 · Guwahati, India

Industrial and Interaction Design, Usability and Ergonomics Research

Awards & Features

Finalist at IXDA Interaction Awards 2013

LaunchPad Denmark 2013 Startup Capital Competition Winner
Published and showcased at IndiaHCI 2011