

JAI PARAKH

Software Developer | Backend Developer | Full Stack Developer | AI / ML Engineer

contact.jaiparakh@gmail.com | parakhjs.me | linkedin.com/in/jai-parakh-5626b4178 | github.com/JaiParakh

PROFESSIONAL SUMMARY

Software and ML engineer with 4+ years building production systems across AI, full-stack, and data domains in healthcare, fintech, and marketing. Recently completed an MSc in Artificial Intelligence and Machine Learning at the University of Limerick, researching autoencoder based anomaly detection for wind turbine fleets. Built backend for a high throughput AI radiology platform, owning micro-services and real-time APIs at clinical scale. Shipped production AI/ML pipelines and LLM powered features for an AI marketing platform processing 10M+ daily events, and built ML pipelines from raw data to SageMaker inference alongside full-stack web applications end to end. Comfortable with ambiguity and high ownership. Based in Dublin, Ireland. Available immediately.

TECHNICAL SKILLS

Languages: Python, JavaScript (ES6+), TypeScript, SQL, C#

Backend: Node.js, Express, Flask, REST APIs, WebSockets, microservices, async job patterns, Puppeteer

Databases: PostgreSQL, MongoDB, Redis, DynamoDB, Firebase

Frontend: React, Next.js, HTML5, CSS3, Tailwind CSS, Material UI

Cloud & Infra: AWS (SageMaker, Lambda, SQS, S3, EC2, Cognito, DynamoDB, CloudWatch, X-Ray, IAM), Vercel, GCP, Docker, Kubernetes, GitHub Actions, Terraform

ML / AI: PyTorch, Pandas, NumPy, scikit-learn, CNNs, Autoencoders, MoE, Optuna, DEAP, LLM prompt engineering, batch inference

Test Frameworks: Jest, MochaJS, Selenium, Cypress

EXPERIENCE

AI/ML Research Intern - BDS Group (Lero), University of Limerick

Jan 2026 - Aug 2026

- Developed a generalized CNN autoencoder framework for industrial SCADA time-series anomaly detection, removing the need for a separate deep model per turbine. Applied Condition-Binned Standardisation and decoder fine-tuning to raise fleet-average AUC from 0.498 (near-random cross-turbine baseline) to 0.735, with per-turbine calibration needing only about 60 normal operating windows.
- Evaluated a Mixture-of-Experts routing extension against the CBS plus fine-tuning result using a matched-pairs statistical test, found no significant improvement ($p = 0.56$), and recommended the simpler model as the defensible final approach.

SDE 2 - Blkbox.ai (AI Marketing Platform)

Nov 2022 - Jun 2025

- Shipped a multi-modal AI pipeline end to end (data model, REST APIs, PostgreSQL) turning ad video into structured, schema-validated LLM output, processing hundreds of videos per batch run and integrating OpenAI APIs.
- Built and scaled a production TypeScript backend handling 10M+ daily events from the Facebook and TikTok APIs at 99.9% uptime.
- Reduced infrastructure costs by 30% and churn by 3% by automating a previously manual client reporting process.
- Led and mentored 3 engineers through architecture discussions and code review, improving delivery quality and cutting review cycle time.

Backend Developer - 5C Network (AI Radiology Platform)

Jun 2021 - Oct 2022

- Built the backend for one of India's largest teleradiology platforms using REST APIs, WebSockets, and microservices on AWS, supporting real-time report delivery to clinicians.
- Designed an intelligent auto-assignment algorithm that routes radiology orders by specialist availability and modality, cutting clinical turnaround time by 25% in production.
- Reduced report load times by 96% (10s to 0.36s) through schema redesign, indexing, and query optimization.

PROJECTS

CocktailAI - GA + SageMaker Async Pipeline (PyTorch, DEAP, AWS, Next.js)

2025 - 2026

- Trained 9 binary taste classifiers on a scraped 4,500-recipe dataset to drive a genetic algorithm; batch inference cut SageMaker calls by 99.9% (150K to 150 per run).
- Deployed on a full async AWS stack: S3, SageMaker endpoint, SQS with a dead letter queue, Lambda, DynamoDB, and CloudWatch/X-Ray for monitoring. (parakhjs.me/projects/cocktail-ai)

ConvNeXt Sketch Classifier - Ablation Study (PyTorch, TorchVision)

2026

- Ran 9 systematic ablation experiments classifying 300K Google QuickDraw sketches across 30 classes, reaching 87.88% accuracy pretrained and 83.84% from scratch with zero overfitting.
- Traced the overfitting to a distribution shift between ImageNet pretraining and sketch data, and resolved it by training from scratch with cosine annealing. (<https://github.com/JaiParakh/ConvNext-Quickdraw>)

Genetic Programming Classifier - Adult Income Dataset (Python, DEAP)

2025

- Built a genetic programming classifier from scratch with DEAP for binary income classification, reaching 78%+ accuracy and a top 4 finish on the Kaggle leaderboard. (<https://github.com/JaiParakh/GP-Classifer-Adult-Income-Dataset>)

RapidMock - API Mocking SaaS (Next.js, React, Node.js, Express, MongoDB)

2023

- Designed, built, and deployed solo, a full-stack SaaS tool letting developers mock API endpoints with configurable, realistic responses, across frontend, backend, auth, and infrastructure. (<https://rapidmock.vercel.app/>)

EDUCATION

MSc, Artificial Intelligence and Machine Learning - University of Limerick, Ireland

Sep 2025 - Aug 2026

Research: autoencoder anomaly detection for wind turbine fleets, BDS Group (Lero) | Coursework: Deep Learning, Cloud AI Ecosystems, Data Engineering, Evolutionary Algorithms