

# DHRUVA BISHT

Dublin, Ireland | +353 870376538 | dhruvabisht.18@gmail.com | github.com/dhruvabisht | linkedin.com/in/dhruvabisht

## PROFESSIONAL SUMMARY

---

Full-Stack Software Engineer with demonstrated expertise in building production-grade agentic AI pipelines, scalable backend systems and cloud-native infrastructure on AWS. Brings hands-on experience integrating large language models into real engineering workflows, designing multi-agent generative AI architectures and delivering GDPR-compliant automation at scale. Combines strong technical execution with product ownership experience across two commercial verticals, translating business requirements into sprint-ready specifications and driving cross-functional delivery. Holds an MSc in Computer Science from University College Dublin.

## TECHNICAL SKILLS

---

**Languages:** JavaScript, TypeScript, Python, Java, Swift, SQL

**Frontend and Backend:** React, SwiftUI, Node.js, Spring Boot, REST APIs, Microservices

**AI and LLM:** Claude API, OpenAI API, Gemini API, Agentic Workflows, Multi-agent Pipelines, NLP, scikit-learn

**Cloud and Infrastructure:** AWS (EC2, S3, Lambda, API Gateway), Docker, PostgreSQL, Supabase, Elasticsearch, Apache Kafka

**Tooling and Process:** Git, Jira, Agile Scrum, CI/CD, Code Reviews, AI-assisted Development (Claude Code)

## WORK EXPERIENCE

---

### Software Engineer | NomadEire | Dublin, Ireland | Oct 2025 – Present

- Built the entire payment transaction layer from scratch using JavaScript and Supabase, integrating Stripe, PayPal and a custom gateway to enable the platform to accept online payments end-to-end.
- Integrated a production-ready content moderation pipeline via the Claude API, applying agentic patterns and chain-of-thought reasoning to automate quality checks at scale.
- Optimised query performance across 18 PostgreSQL tables, restructuring indexing logic to bring average response times under 200 milliseconds.

### Software Development Intern | Advice Bytes | Dublin, Ireland | May 2025 – Aug 2025

- Trained an AI support chatbot on 20,000 historical customer interactions in Python, reducing repetitive inbound support volume by 30 percent.
- Engineered GDPR-compliant guardrails for AI response pipelines serving EU users, achieving 87 percent accuracy across automated compliance checks before production release.

### Packaged App Developer | Accenture | Jaipur, India | Oct 2023 – Aug 2024

- Automated a manual data ingestion process handling over 10,000 files, eliminating 15 hours of weekly manual processing for the engineering team.
- Resolved cross-system integration failures across three enterprise modules in .NET, unblocking user acceptance testing and reducing error rates across front-end to back-end communication.

## PROJECTS

---

### Sentinel-AI: Agentic Content Moderation System

*Python | Gemini API | OpenAI API | Streamlit | PostgreSQL | AWS Lambda | S3*

- Identified a real operational scaling problem — manual content moderation could not keep pace with multilingual, high-volume social media content — and designed an end-to-end AI solution to solve it autonomously.
- Architected a three-agent generative AI pipeline: a Policy Agent that validates content against community standards, a Hindi Cultural Agent that handles linguistic and regional context and an Auditor Agent that produces a final, explainable verdict with full chain-of-thought reasoning.
- Achieved 84 percent accuracy across five violation categories, reduced manual review time by 60 percent and deployed the full agentic system on AWS using Lambda for serverless inference and S3 for privacy-compliant data handling and audit log storage.

### Contextual Ad Targeting Engine

*Java | Spring Boot | Apache Kafka | Elasticsearch*

- Designed and built a real-time content classification pipeline using Spring Boot microservices to parse live web page content and categorise it into target advertising segments at high throughput.
- Streamed enriched classification events through Apache Kafka at 10,000 events per minute and indexed them in Elasticsearch for sub-100 millisecond query performance across the full classified corpus.

### Lynk: Custom macOS Desktop Browser

*Electron | React | TypeScript | Tailwind CSS | Node.js | better-sqlite3*

- Identified a gap in the browser market for a lightweight, developer-focused desktop browser with native tab management and local-first session persistence and built one from the ground up on Electron.
- Architected a multi-process browser engine model with a dedicated main process and isolated renderer contexts in React and TypeScript, resolving complex native module compatibility challenges on Apple Silicon (M1).
- Styled the full interface with a custom Tailwind CSS design system, delivering a production-quality UI that feels native to macOS.

### **Custom iOS Document Editor**

*Swift | SwiftUI | PDFKit | UIDocument | FileManager | UniformTypeIdentifiers*

- Identified a gap in the App Store — no lightweight, offline-capable app could natively open and edit both PDF and Word documents on iPhone without routing files to a cloud service — and built one end-to-end.
- Used PDFKit for complex file parsing, text rendering, annotation and fully on-device asset management; implemented Word support by parsing Open XML in Swift and writing changes back to valid .docx files via UIDocument.
- Optimised memory management for smooth scrolling and seamless document transitions across large files on older iPhone models, with no heavy third-party dependencies.

### **Personal iOS Music Player**

*Swift | SwiftUI | AVFoundation | MediaPlayer | Core Data*

- Built a fully native iOS music player with a SwiftUI now-playing screen, album art rendering, progress scrubbing and adaptive playback controls, prioritising a distraction-free listening experience.
- Integrated AVFoundation for local audio and MediaPlayer for lock-screen controls and background audio, handling call interruptions and route changes via AVAudioSession to match first-party Apple behaviour.
- Built a custom audio library indexer using Core Data for instant search and metadata retrieval across large local libraries, with persistent user preferences across sessions.

### **Twitter Sentiment Analysis**

*Python | NLTK | TF-IDF | scikit-learn | Tweepy | AWS EC2 | S3*

- Built a live sentiment classifier that ingests public Twitter data in real time via Tweepy, processes raw text through a full NLP pipeline (tokenisation, stopword removal, TF-IDF vectorisation) and surfaces signals in a graphical user interface.
- Trained a scikit-learn classifier achieving clean separation of positive and negative sentiment signals across diverse account types, making results readable and actionable for non-technical users.
- Deployed the full pipeline on AWS, hosting the application on EC2 and using S3 for scalable storage of ingested tweet data and model artefacts.

## **EDUCATION**

---

**Master of Science in Computer Science** | University College Dublin, Ireland | Sep 2024 – Sep 2025

**Bachelor of Technology in Computer Science Engineering** | Graphic Era University, India | Sep 2019 – Sep 2023

## **CERTIFICATIONS**

---

- AWS Certified Cloud Practitioner (CLF-C02)
- Microsoft: Azure AI Fundamentals
- Python for Data Science, AI and Development (IBM)
- Complete Web Development Bootcamp 2021