IBU'S RÉSUMÉ

with receipts

Mohamed Ibrahim A.

Software & AI engineer · Coimbatore, India

ibrahimanees.work@gmail.com · github.com/MdIbuA

linkedin.com/in/mohamedibrahimbinabdullah

main.d1n13ofdmf673t.amplifyapp.com

Printed 1 November 2026 · 14 claims, each with a receipt

VERIFIED every line has public proof
PARTIAL some lines are on my word
SELF-REPORTED no public record yet

PROFILE

Final-year B.Tech Information Technology student in Coimbatore, graduating in June 2027. I build AI agents and serverless apps on AWS, often in public and often against a weekend deadline.

My one rule: the model never invents a fact. This résumé keeps it too. Every claim links to its receipt, and every receipt prints its catch, the part that makes me look worse.

Ground the model. Facts come from data and tools, never from the model's memory.

Publish the number that hurts. An R² of 0.94 means little without the 47.93% beside it.

Own the loop. Fifty lines I can debug beat a framework I can't.

EXPERIENCE

#11 Three internships — PARTIAL

AI Intern · Infosys Springboard, Batch 5 — public
Backend AI Engineer Intern · FlyRank AI — no public receipt
MERN Stack Intern · AICTE program — no public receipt

The catch: Only the Infosys internship left public work behind.

Evidence: CryptoPulse README: Infosys Springboard, Batch 5

SELECTED WORK

#1 CryptoPulse — VERIFIED

R² 0.94 on Bitcoin close price · Infosys Springboard internship

R² · close · test set — 0.9402
Direction · close · test set — 47.93%

The catch: On close price the model calls up-or-down correctly 47.93% of the time.

Corrected: An earlier write-up of mine called this an 80/20 split. The LSTM cell in the notebook splits 70/30.

Evidence: Notebook: "Evaluating Model" cell

#2 Supermarket Ops Agent — VERIFIED

30 typed tools, one chat window · hiring take-home · 2026

Tools in ALL_TOOLS — 30
Where prices come from — Postgres, via tools

The catch: One store per bot.

Corrected: My README says 32 tools, and so did an earlier write-up. The registry holds 30. The code wins.

Evidence: registry.ts: the 30 tool schemas · 3 more on the receipt

#3 DocuMind AI — VERIFIED

Built in a weekend on Amazon Bedrock · AWS Builder Showcase Challenge · Aug 28–31, 2026

Model — Amazon Nova Pro
Backend — Lambda · Python 3.12

The catch: There is no retrieval.

Evidence: Published write-up on AWS Builder Center · 2 more on the receipt

#4 Tripwire — VERIFIED

A ~0.2 s reflex against prompt injection · Strands Agents + Amazon Bedrock · Oct 2026

Attacks that worked · Tripwire — 0 / 10
Median decision · guard — 233 ms

The catch: End to end, the median guard decision took 341 ms, not 0.2 s.

Evidence: README: results tables · 2 more on the receipt

#5 Ghost Probe — VERIFIED

A read-only S3 tool that found "deleted" data · AWS Builder Center "Give your Agent a Tool" challenge · Oct 2026

IAM actions granted — 11, all List/Get
Verdict on case_25 · before — COMPLETE
Verdict on case_25 · after — INCOMPLETE

The catch: That's one case from a fictional dataset, run against a demo bucket I seeded myself.

Evidence: Least-privilege IAM policy · 2 more on the receipt

#6 The Git Bard — VERIFIED

A bard that reads my commits every Friday · AWS "Set Your Creative App Free" weekend challenge

EventBridge schedule — cron(0 16 ? * FRI *)
Storyteller — Amazon Nova Micro

The catch: It only chronicles the repos I list (mine lists one), and in the SES sandbox it can only email addresses I've verified.

Evidence: template.yaml: the Friday schedule · 1 more on the receipt

#7 Voter Saathi — VERIFIED

An election guide in 12 Indian languages · civic tech · 2026

Languages — 12
Models tried, in order — 3 free OpenRouter models, then Gemini

The catch: Only English is hand-written.

Corrected: My README credits Gemini with the other 11 languages. The server tries three free OpenRouter models first (Gemma, Qwen and Llama) and calls Gemini only if they fail or no OpenRouter key is set.

Evidence: README: the 12 languages · 2 more on the receipt

#8 CelebrAI — VERIFIED

86 passing tests, zero runtime dependencies · event planner

Unit tests passing — 86 / 86
Runtime dependencies — 0

The catch: All 86 tests passed while a dialog was painted over the entire app.

Evidence: package.json: dependencies {} · 2 more on the receipt

#9 Tiny Tale Toy — VERIFIED

Every story branch generated live · Bedrock + Lambda + Amplify

Path — API Gateway → Lambda + Bedrock
Default model in code — Ministral 3 8B

The catch: Models don't always return clean JSON.

Corrected: The README still says Nova Lite in places. The code's default model is Mistral's Ministral 3 8B.

Evidence: lambda_function.py: the defensive parser · 1 more on the receipt

SKILLS

Agents that don't make things up
Amazon Bedrock · Strands Agents · Claude API · Gemini API · LangChain · RAG

Serverless on AWS
Lambda · API Gateway · EventBridge · S3 + CloudFront · SES · SAM · Amplify · IAM

Money-grade full-stack
TypeScript · React · Node + Express · FastAPI · PostgreSQL + Prisma · MongoDB · Oracle APEX

Honest ML
TensorFlow / Keras · scikit-learn · Computer vision · NLP

Also fluent in
Python · JavaScript · TypeScript · SQL & PL/SQL · Docker · GitHub Actions · Google Cloud Run · Firebase

Studied, not yet shipped
Java · C · C++ · Kubernetes · Azure · BigQuery · Snowflake

COMMUNITY

#10 Community roles — PARTIAL

Student programs

AWS Student Builder Campus Leader — no public receipt
Microsoft Learn Student Ambassador (Associate) — no public receipt
Author on AWS Builder Center — public

The catch: Two of these three roles have no public page I can link to, so they're on my word alone.

Evidence: AWS Builder Center profile · 1 more on the receipt

#12 HackerRank Orchestrate, Aug–Oct 2026 — PARTIAL

Monthly build challenges

October round: built on its dataset — public
Aug, Sep, Oct starter repos forked — public
My submissions and placings — no public receipt

The catch: The weakest receipt on this page.

Evidence: Ghost Probe audits the October round's dataset · 1 more on the receipt

#13 Ran an IBM Qiskit Fall Fest workshop — SELF-REPORTED

Off the clock · quantum computing

Workshop: quantum algorithms — self-reported

The catch: The only public record is my own profile README, so treat this one as self-reported.

Evidence: My GitHub profile README

EDUCATION

B.Tech, Information Technology · Coimbatore — June 2027

THIS PORTFOLIO

#14 This page, served by AWS Amplify Hosting — VERIFIED

The receipt for this receipt

Build and commit — stamped live by Amplify

The catch: I didn't type the build and commit: the Amplify build stamps them into the live page at deploy time.

Evidence: amplify.yml: type-check, test, compile, stamp · 2 more on the receipt

TOTALS

Claims on this résumé — 14
Backed by public proof — 10
Partly proven — 3
Self-reported — 1
Corrected by my own code — 4
Invented facts — 0

10 OF 14 VERIFIED

THANK YOU. CHECK AGAIN SOON.

Every # above is a live receipt you can re-check with AWS Lambda at main.d1n13ofdmf673t.amplifyapp.com