



# Sandeep Malik

Principal / Staff AI Engineer

## CONTACT

### Email

samsandeepmalik@gmail.com

### Phone

(+1) 438-725-7870

### Location

Brossard (Greater Montreal), Canada

### Web

sandeep-malik.com

### LinkedIn

in/sandeep-malik-48a12a70

### GitHub

github.com/samsandeepmalik

## CORE SKILLS

### AGENTIC AI

Harness & runtime engineering, Agent & MCP design, BYO-agent/MCP platforms, orchestration, Multi-Agent, ReAct, Plan-and-Execute, Human-in-the-Loop

### GENAI & RAG

LLMs, RAG (multi-query, step-back, reranking, hybrid search, LLM-as-judge, evals), Prompt Engineering, Semantic Search, SBERT

### PLATFORMS

Azure OpenAI, AWS Bedrock (Guardrails), LangChain, LangGraph, Llamaindex, Google ADK, Langflow

### DATA & CLOUD

Snowflake, Databricks, Spark, Kafka, Airflow, AWS, Azure, Kubernetes, Terraform

### DATABASES

Pinecone, Elasticsearch/OpenSearch, PostgreSQL, MongoDB, Delta Lake

### LANGUAGES

Python, SQL, Java, Scala, Node.js

## CERTIFICATIONS

- ▶ Snowflake SnowPro Core
- ▶ AWS Solutions Architect – Associate
- ▶ AWS Cloud Practitioner
- ▶ Databricks Spark 3.0 – Associate
- ▶ CCA-175 (Spark & Hadoop)
- ▶ Microsoft AZ-900

## EDUCATION

### M.Sc. Computer Science

Manav Bharti University, 2013

## SUMMARY

AI Engineer & Data Architect with **13+ years** across 11 companies and 2 continents, now leading **agentic-AI platform & harness engineering**. A hands-on builder who ships **production-grade agentic systems** end-to-end — agents, MCP tools, RAG, and the discovery/orchestration layers that govern them — on a decade of data-platform and big-data architecture (handled up to **~2 PB** and **~200K events/sec**). Currently Principal Software Engineer & Technical Lead (Trust AI Engineering) at Autodesk, where a proactive agent I built **modernized 1,500+ internal apps** to GraphQL at **>85% success**.

## EXPERIENCE

### Principal Software Engineer / Technical Lead · Autodesk

Jul 2025 – Present | Montreal, Canada

- Technical lead of a central team building an **agentic platform** (BYO-agent / BYO-MCP) — the agent & MCP catalog, discovery & orchestration layers, and the **harness/standards** for building workers.
- Built a **proactive agent** that migrates legacy, vulnerability-prone APIs to **GraphQL** across **1,500+ internal applications** — handling direct and dependency-based usage in Python, Go, Java, Ruby on Rails, and Perl — at a **>85% success rate** (by code complexity), surfaced through a UI dashboard. Significant ROI.
- Shipped **3 production agents** (1 proactive, 2 reactive), **5 MCPs**, and **2 IDE plugins** used across Autodesk.
- Shift security **left** in the SDLC: MCP tools surface secure-coding & cloud standards inside AI IDEs (Cursor, Claude Code) so developers remediate in-cycle.

### Data Architect / Staff Engineer · Databook Labs (USA)

Dec 2022 – Jul 2025 | Montreal, Canada · Remote

- Built the data platform **from scratch** and its **first semantic layer (early 2023)** powering Databook AI, hardening data freshness, consistency & quality.
- Spearheaded a **GPT product for strategic sales** with advanced RAG (multi-query, step-back, reranking, LLM-as-judge) to lift engagement and personalization.
- Designed **agentic AI** workflows with LangGraph, Google ADK, Conductor & Langflow; semantic index over Pinecone & Elasticsearch on Snowflake/Databricks; led a team of data engineers.

### Principal Data Engineer · Sitecore

Mar 2022 – Dec 2022 | Montreal, Canada

- Led the core **Search** team handling **~200K events/sec**, enhancing AI-backed search and personalization for e-commerce.
- Architected a Data Lake & Warehouse for a global retailer, reducing storage costs; built Snowflake marts and automated monitoring. Stack: Elasticsearch, Spark, Kafka, Snowflake, AWS, Airflow, Docker.

### Solutions Architect · Cloudera

May 2021 – Mar 2022 | Montreal, Canada

- Professional Services expert onboarding customers to Cloudera Data Platform (public/private/hybrid); designed on-prem → AWS migrations and reduced cloud cost.

### Senior Data Engineer · Intact Insurance

Apr 2020 – May 2021 | Montreal, Canada

- Led a team of data engineers; built on-prem → AWS historical-data migration, refactored pipelines onto Databricks, authored Snowflake/Databricks/AWS best-practice guides.

**Earlier:** Data Engineer, Ericsson (2019–20 — 5G optimization on Spark/Airflow/Kubeflow/K8s, Azure) · Data Engineer, CGI (2018–19 — real-time prediction & rule engine for trade settlement) · Data Engineer, Synechron (2018, Pune) · Big Data Developer, Interglobe (2016–18, Dubai — batch + real-time pipelines over ~2 PB) · Software Engineer, MotherSumi/MIND (2014–16) · Software Developer, NIIT (2012–14).