

PROFILE

A full-stack data scientist with **9 years of experience** across enterprise technology, FMCG, the internet industry, data architectures, and AI agents. Confident in data wrangling, data visualization, machine learning, and building data products end-to-end.

EXPERIENCE

2025 – Present

Virtuals Protocol

Head of DevRel & Forward Deployed Engineering

- Lead a technical team building and maintaining AI products — building agents, mentoring junior devs / DevRel's, building agent frameworks, and running developer relations.
- Built multiple agents to seed growth in the agent marketplace — companion agents, evaluator agents, adtech agents.
- Demonstrating thought leadership on the agentic economy by speaking at technical panels including ETH Denver, OpenClaw events, and the Coinbase hackathon.

2024 – 2025

Oliver Wyman

Lead Data Scientist / Engineer

- Designed the data model for Thailand's largest bank, optimised for a lakehouse architecture.
- Databricks partner champion.

2022 – 2024

Sourced Group

Consultant — Data Analytics & Machine Learning

- Designed and built an Azure data mesh platform for a major Japanese telco — data pipelines with monitoring (Databricks, Azure Data Factory), data quality checks (Great Expectations), data catalogue (Purview).
- Built an ML model to predict overwork in an HR analytics project, registered and served via Databricks model registry & model serving.
- Designed and built an AWS data platform for a major Malaysian FSI — Spark pipelines and data governance (Glue, Lake Formation, QuickSight).
- Earned 11 certifications across cloud & data (5x Databricks, 3x Azure, 3x AWS).
- Additional roles: tech interviewer, pre-sales, project lead.

2019 – 2022

Grab Holdings

Senior Data Scientist / Engineer

- Built data pipelines and datamarts for 5+ SuperApp categories, integrating with 9+ external partners within 10 months.
- Built causal inference models and ran A/B tests to quantify product feature impact (Presto, PySpark, Databricks, Tableau).
- Built and maintained an internal tool to democratise causal inference methods across the org.

SKILLS

Python

SQL

R

Spark / PySpark

Databricks

AWS

Azure

Tableau

MongoDB

QuickSight

Agent Frameworks

LLMs

Causal Inference

ML Engineering

Lakehouse

Data Mesh

PROJECTS & WORK SAMPLES

DevRel Content & Events

- Product recap — x.com/celesteanglm/status/2005609882293272959
- Short-form interview — x.com/BaseinSG/status/2041858032553095419
- Live talk — x.com/virtuals_io/status/2031611132243488961
- Panel — x.com/virtuals_io/status/2026865549230354679
- Article — x.com/celesteanglm/status/2024646955352416524

Intelligent Detection of Street Numbers with CNNs

- Scored Distinction in Georgia Tech's Computer Vision course.
- Preprocessing with image pyramids and Maximally Stable Extremal Regions (MSER).
- Predicting digit sequences with Convolutional Neural Networks.

Ocubulum

- Finalist for the NUS School of Computing Innovation Prize (2016).
- Automated data collection and dashboards to revolutionise researcher performance analysis at NUS.
- Built with Django, Python, Vue.js, and D3.

EDUCATION

MSc, Computer Science

Georgia Institute of Technology

2018 – 2020

Specialization in Machine Learning.

BSc, Business Analytics

National University of Singapore

2014 – 2017

NUS Overseas College Scholar — Stanford University exchange.

ACHIEVEMENTS

- 2019** 2nd Place, Rolls-Royce Data Challenge — Analysed factors for flight delay using ML and Monte Carlo simulation, with dashboard visualizations.
- 2017** Certificate of Distinction, Analytics Techniques — Awarded by NUS Computing to high-performing students.
- 2017** Certificate of Merit, Data Science — Awarded by NUS Computing to high-performing students.
- 2016** NUS Overseas College Scholarship, Silicon Valley — Year-long internship in Silicon Valley startups with coursework at Stanford University.