

EXPERIENCE

Employer: AMAZON WEB SERVICES

Hank, 12 W 39th St, New York, NY 10018

SENIOR SPECIALIST SOLUTIONS ARCHITECT — GENERATIVE AI, WWSO (L6), AWS

01/10/2025 – Present

RESPONSIBILITIES:

- Co-authored enterprise GenAI architectural guidance — including the AWS Generative AI Lens multi-tenant scenario, APG platform maturity guidance, and an AWS ML Blog — adopted across multiple customer engagements and account teams.
- Authored model customization readiness frameworks and delivered training to 25+ stakeholders, enabling AWS service teams to scale customer advisory for fine-tuning engagements.
- Co-led the Model Customization Champions community within the ML Technical Field Community, growing membership from 2 to ~50 specialists, and served as an AI League TFC Champion, upskilling practitioners to deliver events independently.
- Secured a seven-figure, multi-year competitive migration win against Azure by leading the technical evaluation and providing deep architectural, cost, and scalability guidance, establishing AWS as the customer's strategic AI partner.
- Designed and delivered a Model Customization Platform Accelerator — a reusable LLMOps fine-tuning platform on SageMaker AI — published to AWS Samples with a workshop offering. Presented at the AWS North America Tech Summit (CSAT 4.8) and accepted for a re:Invent 2025 Builder's Session.
- Served as a key scientific advisor for SageMaker's agentic model customization launch (Sherpa), contributing RLVR workflow design, training recipes, and a cross-org preview testing program with ~30 participants that directly shaped product readiness.
- Played a central role in the AWS AI League go-to-market, contributing toward 11,326 registrations and ~30 Lighthouse enterprise customers. Built the Nova Customization Workshop and POC for the re:Invent finale, and supported a 16-company hackathon across 5 time zones.
- Delivered talks at re:Inforce 2025 (CSAT 4.84), re:Invent 2025, the London Summit, and internal summits. Produced multiple AWS ML blogs on VLM fine-tuning, multi-provider GenAI gateways, and Bedrock model customization.

SENIOR MACHINE LEARNING ENGINEER (L6), AWS Professional Services

10/01/2023 – 01/19/2025

RESPONSIBILITIES:

- Led the delivery of an LLMOps platform for the end-to-end fine-tuning, evaluation, load-testing & deployment of LLM's for Booking.com. Working with AWS Trainium and Inferentia service teams on customer feature requests & future product roadmap.
- Leading the development of Booking's AI Trip Planner on Amazon SageMaker alongside the AWS GenAI Innovation Centre, generating destination recommendations using fine-tuned, open-source LLMs for Booking.com customers on the mobile app.
- Collaborating with SageMaker & Bedrock Service Teams to shape and refine the roadmap for our GenAI services by leveraging customer feedback and market trends. Focus areas include enhancing governance (Responsible AI) features and beta-testing preview features such as the Bedrock custom model import.
- GenAI Technical lead for AWS ProServe Energy UKIR, helping define the go-to-market strategy and scoping ProServe engagements for energy customers such as Total Energies, ENI & bpx Energy.
- Actively mentoring and training customers through AWS GenAI Deep Dive Days (D3) (e.g. running Bedrock agent workshops) and providing 1-on-1 hands-on advisory to C-level executive customers as part of the [UKIR GenAI Solutions Launchpad](#).
- Sole technical advisor and ML engineer for an Aviva GenAI Hackathon 2023 team, developing AutoClaim, an automated insurance claim bot powered by Amazon Bedrock & Amazon Kendra. [Nick's hackathon team won the hackathon](#).
- Active member of the GenAI speaking community. Notably, Nick led and co-presented "Scaling LLM Inference using Amazon SageMaker" at the [AWS 2024 AI & Data Conference Ireland](#) and presented "Building Generative AI Responsibly with Amazon Bedrock" externally at the [re:Invent 2023 re:Cap Manchester](#).

MACHINE LEARNING ENGINEER (L5), AWS Professional Services

07/01/2022 - 10/01/2023

RESPONSIBILITIES:

- Implemented an [MLOps framework for BPX Energy](#) on Amazon SageMaker that reduced model deployment time by 30%, improved model accuracy by 60% and provided a 40% ROI increase potential through model monitoring and re-training.
- Led the delivery of an MLOps platform on Amazon SageMaker for Ericsson, reducing onboarding time for 4 ML use cases by an average of 83%, saving \$453K in deployment costs.
- Delivered external presentations at [AWS Manchester User Group \("Machine Learning & Blockchain, The Perfect Marriage?"\)](#), [UCISA23 \("Responsible use of AI in Higher Education"\)](#), led EMEA AWS Partner ML Immersion days and organized two GenAI Hackathons for AWS partner Rackspace.
- Played a key role in strategic, pre-sales Professional Services discussions with C-level executives, influencing outcomes that resulted in 23 unique ProServe opportunities worth \$31M in bookings.

- Acted as the lead AWS advisor in designing and developing a new [CI/CD pattern for AWS partner ML6](#) that promotes ML code between environments rather than ML models, resulting in €200K savings across 10 ML6 engagements.
- Led the creation and delivery of the “AI Air Hockey” session at the Builder’s Fair at re:Invent 2022, attended by 250 participants.

ASSOCIATE MACHINE LEARNING ENGINEER (L4), AWS Professional Services

09/01/2020 – 07/01/2022

RESPONSIBILITIES:

- For Deutsche Fußball Liga (DFL), designed and implemented Amazon SageMaker Clarify experiments to explain in-production (expected goals) XGB Boost models. Resulted in an AWS ML [blog post](#), a publication in [German DevOps magazine](#) and the submission of a Product Feature Request for a missing SHAP KernelExplainer explainability feature.
- For AstraZeneca's (AZ) Center for Genomics Research (CGR), designed and implemented a new deployment workflow reducing deployment time by 87% and a Genomics Data Lake saving 120 hours per annum for scientists, resulting in annual cost savings of \$490K.
- Designed and deployed an MLOPs platform (AI Bench) for AZ on Amazon SageMaker along with a portfolio of 5 SageMaker Pipeline Service Catalog products, enabling 300+ researchers to manage over 1,200 ML experiments annually, accelerating medical innovation on their AI platform.
- Leading a team of 4 backend engineers, developed [AWS Dash](#), an internal KPI tracker web application using API Gateway, Lambda, Glue and RDS. Used by over 3000 consultants and managers.

Employer: DECISION LAB
DATA SCIENTIST

120 Main Street, Newton Linford, Leicestershire, UK, LE6 0AF
11/04/2019 – 08/28/2020

RESPONSIBILITIES:

- Led a project with Frazer-Nash Consultancy for the Defence Science and Technology Laboratory (DSTL) to develop a state-of-the-art reinforcement learning (RL) system for devising firefighting strategies in Royal Navy submarines. Created a simulation of emergency firefighting scenarios, utilizing advanced RL algorithms to enhance mission systems software and improve autonomous defense technology.
- Collaborated with a team of 3 data scientists on a project for the UK Space Agency to develop a machine learning system using Convolutional Neural Networks (CNNs) for image segmentation and spatial data analysis. Leveraged TensorFlow and remote sensing data to produce prospectivity maps for battery metals, improving exploration efficiency.

Employer: XAI ASSET MANAGEMENT
DATA SCIENTIST

30 Churchill Place, London, UK, E14 5EU
06/03/2019 – 09/27/2019

RESPONSIBILITIES:

- Integrated three state-of-the-art, continuous, model-free Reinforcement Learning (RL) algorithms ([PPO](#), [DDPG](#), [SAC](#)) into xAI’s proprietary trading platform, mapping raw investment signals to target weights of a financial portfolio.
- Designed and implemented a number of novel RL trading environments that provided superior out-of-sample (cumulative) returns relative to xAI’s legacy long-only trading strategies, subsequently explaining them to non-technical senior investors.
- Performed significant data preprocessing and analysis on vast quantities of macro data to acquire an informative agent state - a single time series for a given asset. Researched and tested multiple custom neural network architectures and action distributions.

EDUCATION

University College London (UCL)

Machine Learning MSc, Distinction (77%).

London, UK • Sept 2018 – August 2019

Thesis: “Deep Reinforcement Learning for Portfolio Construction & Optimisation” – in collaboration with xAI Asset Management.
Grade: Distinction.

Physics BSc, First Class Honours

London, UK • Sept 2014 – June 2017

Thesis: “Investigation of the hyperfine structures of Rubidium using Doppler-free, frequency-modulated saturated absorption spectroscopy”.
Grade: First class.

PUBLIC ARTIFACTS

1. ["Reinforcement fine-tuning on Amazon Bedrock with OpenAI-Compatible APIs: a technical walkthrough"](#), 25th March 2026
2. ["Enhanced performance for Amazon Bedrock Custom Model Import"](#), 26th November 2025
3. ["Streamline AI operations with the Multi-Provider Generative AI Gateway reference architecture"](#), 21st November 2025
4. ["Fine-tune VLMs for multipage document-to-JSON with SageMaker AI and SWIFT"](#), 10th November 2025
5. ["Architect a mature generative AI foundation on AWS"](#), 30th May 2025
6. ["Time series forecasting with Amazon SageMaker AutoML"](#), 8th October 2024

7. ["Prepare training and validation dataset for facies classification using a Snowflake OAuth connection and Amazon SageMaker Canvas"](#), 17th May 2023
8. ["BMW Cloud Efficiency Analytics powered by Amazon QuickSight and Amazon Athena"](#), 15th November 2023
9. ["Explaining Bundesliga Match Facts xGoals using Amazon SageMaker Clarify"](#), 18th March 2021

PUBLICATIONS

1. **"SentiMATE: Learning to play Chess through Natural Language Processing"**, accepted for oral presentation at the AIIDE 2019 conference for Artificial Intelligence in Strategy Games – [arXiv link](#), [MIT Tech Review article](#).
2. **"Amazon SageMaker Clarify: Machine Learning Bias Detection and Explainability in the Cloud"**, Amazon Science paper, September 2021 – [arXiv link](#), [Amazon Science](#).

AWARDS & CERTIFICATIONS

GenAI with Large Language Models	Apr. 2024
Hashicorp Certified - Terraform Associate	May 2023
Certified AWS Senior Speaker	Apr. 2023
Scrum Alliance Certified ScrumMaster (CSM)	Dec. 2022
AWS Certified Data Analytics	Aug. 2022
AWS ProServe Worldwide Learn & Be Curious Award	Mar. 2021