Charles Yee, Ph.D.

Director of AI, Machine Learning & NLP | Healthcare

Nashville, TN • charles.wj.yee@gmail.com • (720) 771-1828

Profile Summary

• Director of **AI**, **Machine Learning**, and **NLP** with over **10 years** of experience delivering end-to-end **healthcare and biomedical informatics systems**, including vaccine safety surveillance, medication adherence prediction engines, and clinical trial matching platforms for **Fortune 500** companies and **VC-backed startups**.

• Comprehensive technical expertise building scalable **AI** solutions using machine learning frameworks (**TensorFlow**, **PyTorch**, **XGBoost**), **NLP** libraries (**spaCy**, **NLTK**, **Stanford CoreNLP**, **Spark NLP**), and **AWS** healthcare services (**SageMaker**, Comprehend Medical, **Bedrock**). Collaborated with **MLOps** teams to deploy models via **CI/CD** pipelines and monitor performance with **CloudWatch.**

• Led cross-functional programs across **Product, Engineering, Clinical, and Regulatory** teams in **Agile** and **SAFe** environments, aligning roadmaps with technical feasibility and business impact.

• Experienced **people manager** and **mentor** leading multidisciplinary teams through research, product development, and delivery, contributing to peer-reviewed publications, filing **multiple patents** in **biomedical NLP**, and representing organizations as a conference speaker and technical evangelist.

• Native bilingual in **English** and **Mandarin** with proficiency in **German**, and **co-founder** of multiple **AI-driven healthcare** ventures with successful Series A/B outcomes.

Technical Skills

Programming Languages: Python, Java, R
Libraries & Frameworks: Pandas, Numpy, SciPy, Matplotlib, TensorFlow Keras, PyTorch, Flask, Spring Boot, JUnit, Mockito, RDKit, FastAPI
Databases: PostgreSQL, Snowflake SQL, MongoDB
Cloud & DevOps: AWS, Azure, Docker, Kubernetes, Terraform, Jenkins, AWS Lambda, Step Function, Cloud Development Kit (CDK)
API & Interface Tools: GraphQL, Swagger, Postman APIs
Data Management & Processing: ETL, Tableau, Apache

Machine Learning Tools & Techniques: XGBoost, Support Vector Machine (SVM), Random Forest, gradient boosting, Latent Dirichlet Allocation (LDA), ELMo, BERT, LSTM, Retrieval Augmented Generation (RAG), Foundation Models, Scikit-learn, LLM, Generative Pre-trained Transformers (GPT), Model Context Protocol (MCP), Agentic AI
Natural Language Processing: spaCy, Stanford CoreNLP, Elasticsearch, OpenNLP, Spark NLP, Word Embeddings, LangChain, NLTK

Work Experience

Inovalon

Bowie, MD

AI, Machine Learning Director

Apr. 2023 - Present

• Directed translation of complex healthcare use cases into machine learning solutions across provider, payer, and specialty pharmacy workflows, supporting over **80 product lines,** including medication adherence prediction, eligibility verification, and domain-specific chatbots.

• Built and deployed predictive models using **XGBoost** and **Gradient Boosting** on **AWS SageMaker**, achieving **95% accuracy** in identifying non-adherent patients and reducing eligibility errors via **EC2**, **S3**, and **feature optimization** techniques.

• Implemented a **Retrieval-Augmented Generation (RAG)** chatbot using **Claude Sonnet 3.7**, **AWS Kendra**, and **OpenSearch** to improve internal triage and installation workflows within the ScriptMed Cloud platform, enhancing response accuracy through a **Modular Task-Oriented Dialogue System**.

• Engineered scalable and reproducible model training pipelines, collaborating with **MLOps teams** to deploy via **AWS-native CI/CD tools** (CodePipeline, CodeDeploy, Lambda), and monitor performance and drift using **CloudWatch**.

• Developed a rule-based optimization layer to refine **AWS Comprehend Medical** outputs, eliminating false positives in diagnosis extraction, accelerating coder workflows, and improving **Risk Adjustment** model accuracy and inference latency.

• Worked cross-functionally with the **VP of Engineering**, **Product Owner**, **Project Manager**, and **Clinical team**, applying Agile practices such as daily standups and weekly sprints to deliver under tight deadlines while adapting communication to both technical and non-technical audiences.

• Represented Inovalon at **Customer Congress 2023**, delivering a panel presentation and co-hosting a podcast on healthcare AI, strengthening the company's thought leadership and client engagement.

AstraZeneca

Gaithersburg, MD

Principal Scientist

Feb. 2021 - Mar. 2023

• Led **R&D** initiatives on **AZD1222 COVID-19 vaccine** adverse event detection and prevention using pharmacovigilance methodologies, managed safety reporting focused on myocarditis, coagulopathy, and thrombocytopenia, and delivered analytical findings to corporate executives, the **Chief Medical Officer,** and **regulatory bodies,** including the **FDA** and **ECDC.**

• Built statistical models and pipelines to detect abnormal rates of adverse events linked to specific vaccine manufacturing lots, using population-level baselines and custom spelling correction via Levenshtein distance and n-gram frequency ranking.

• Constructed biomedical text mining pipelines for the ***IQ DILI Consortium*** to identify cases of drug-induced liver injury (**DILI**) from scientific publications and pharmaceutical reports, using **PubMedBERT, ClinicalBERT**, and **SciSpacy** for named entity recognition, supported by custom RegEx patterns and manual validation to ensure clinical relevance.

iQuartic

Boston, MA

Director of Data Science

Oct. 2019 - Oct. 2020

• Developed the company's front-end and back-end microservice architecture supporting health insurance risk adjustment and **AI**-assisted **ICD-10** coding, and served as product owner for all offerings related to **NLP**, **OCR,** and **Machine Learning**.

• Recruited, mentored, and managed a team o**f 7 Data Scientists**. Conducted **1** on **1** meetings, led code and performance reviews, and collaborated on employee personal development plans. Created a psychologically safe environment to ensure well-being and maximize productivity.

• Oversaw the development of a **React.js**–based front-end microservice architecture built on **Flask**, leading a team of two full-stack developers and ensuring timely delivery of performant, modular web interfaces aligned with project goals.

• Implemented a backend microservice architecture using **Python** for **API** development and integrated persistent data storage with **MongoDB**, applying the **API Gateway** Pattern to support efficient **API** response times.

• Enhanced tagging of over **10**,**000** daily **Electronic Medical Record** pages with encounter and **ICD** codes using **ClinicalBERT**, **SciSpacy**, and a custom **LSTM** classifier trained on IOB-formatted data annotated by internal medical coders, optimizing downstream processing by enabling entity-level validation to filter false positives.

• Executed evaluation and hyperparameter tuning of deep learning models using **TensorFlow**, the **Keras LSTM** layer, and **Optuna**, and refined the models through **Knowledge Distillation** and **Model Pruning** to reduce model inference latency.

• Built **CI/CD** workflows using **Jenkins** and integrated automated test coverage and linting, achieving a deployment time reduction to **5** minutes.

• Performed end-to-end **User Acceptance Testing** and **Operational Acceptance Testing** to improve maintainability and reduce test execution time.

Mixfit Inc.

Boston, MA

Director of Data Science

Aug. 2018 - Oct. 2019

• Implemented a software project for a personalized nutrition recommendation and dispensing system powered by scientific guidance, utilizing optimization techniques to improve nutritional dispensing and business cost efficiency, resulting in two intellectual property filings.

• Authored invention disclosures that resulted in **2** intellectual properties representing the company's core value propositions and managed one vendor and a development team of **6** engineers.

• Designed and implemented a **Ruby on Rails** microservice using dynamic programming optimization to generate personalized nutritional recipes from wearable, dietary, biometric, genetic, and demographic data, improving computation efficiency across user profiles.

• Introduced backend logic to transform personalized nutrition plans into device-level instructions using **MQTT** for lightweight messaging, enabling integration with **Raspberry Pi**–based dispensers managed by hardware engineers.

• Implemented dynamic programming with **Memoization (top-down)** techniques and integrated with **Microservice** and **Event-Driven Architecture,** improving dispensation efficiency and reducing shipping costs.

Philips North America
Boston, MA

Senior Biomedical Informatics Scientist
Mar. 2016 - Aug. 2018

• Collaborated with the **development team, internal business unit,** and **senior executives** to propose new product features and managed validation and precision-recall improvements of the final solution through close coordination with **pathologists** and **oncologists**.

• Launched research exhibits at ***MD Anderson***, ***Dana-Farber***, and ***Westchester Medical Center*** by coordinating with clinical staff to set up product demonstrations and gather feedback on performance and usability, aiming to benchmark solutions against real-world clinical needs.

• Built **NLP** prototypes using **TensorFlow Keras** with **RNNs** (**LSTM**) and **CNN**s, fine-tuned for **NER** and classification tasks in oncology clinical trial matching applications.

• Designed a **Clinical Phenotype NER** system using **LSTM-CRF** architecture in **TensorFlow Keras**, trained on IOB-formatted data annotated manually to extract relevant entities from unstructured text.

• Created a **genome spelling correction algorithm** using custom gene dictionaries and n-gram co-occurrence frequency from biomedical texts, applying Levenshtein distance and probabilistic ranking to resolve ambiguous gene mentions.

• Integrated **NegEx** for precise interpretation of negations in clinical narratives and executed concept extraction for patient profiling using **CTakes, SciSpacy**, and **RegEx**, employing **Rule-Based Pattern**.

Optum Analytics, UnitedHealth Group
Boston, MA

NLP Engineer
Mar. 2014 - Mar 2016

• Engineered ontology-aligned **NLP pipelines** to extract complex medical concepts from unstructured **EMR/EHR** data across **130M+** patients, utilizing **SVM**-based classification to address linguistic challenges and improve concept recognition.

ACT
Denver, CO

Research Scientist, Computational Linguist
Aug. 2011 - Nov. 2013

• Designed and implemented an automated essay scoring engine for **K–12 standardized testing** using **Random Forest**, **Gradient Boosting**, and **LDA**, achieving human-level accuracy and securing top placement in the **Kaggle Essay Scoring** competition.

Education

University of Stuttgart Ph.D. in Computational Linguistics, NLP — Stuttgart, Germany Jun.2004 - Nov.2010

King's College London M.Sc in Computational Linguistics, NLP — London, UK Sep. 2001 - Oct. 2003

Ohio University B.Sc in Computer Science — Athens, OH Jul. 1997 - Jul. 2001

Patents

US PCT/EP2019/071747, Nomenclature-Informed Gene Name Recognition, Filed Aug 13, 2019
US PCT/EP2018/075263, Natural Language Processing using Ontology Mapping, Filed Sep 19, 2018
US PCT/EP2019/066322, A Method and Apparatus for Genome Spelling Correction and Acronym Standardization, Filed Jun 20, 2019
US WO/2018/060838, A Method and System for Matching Subjects to Clinical Trials, Issued May 4, 2018
US PCT/EP2019/072816, Method and System for Cancer Stage Annotation within a Medical Text, Issued Mar 5, 2020
US WO/2019/068870, Methods and Systems for Healthcare Clinical Trials, Issued Apr 11, 2019