

Ryan Inghilterra

Senior Data Scientist II, CrowdStrike | Applied AI, Production ML & Football Analytics

San Diego, CA, USA · Open to relocating to London; available for international travel
ringhilterra.github.io | linkedin.com/in/ryan-inghamterra-017 | ringhilterra17@gmail.com

Senior Data Scientist II with 10+ years of experience building production machine learning and applied AI systems in high-growth technology environments, including customer-facing capabilities for CrowdStrike's Falcon platform. Experienced across the full modelling lifecycle, from problem framing and feature engineering through validation, explainability, monitoring, stakeholder adoption, and production delivery. Applies tracking and event data to football analytics research, building reproducible models, metrics, and analyst-facing tools that turn player, team, and academy questions into practical decision-support insight.

PROFESSIONAL EXPERIENCE

CrowdStrike – Senior Data Scientist II

Aug 2021 – Present

Promoted from Data Scientist III to Senior Data Scientist I to Senior Data Scientist II.

- Lead applied ML/AI research and production delivery across high-scale Falcon product environments, partnering with data science, data engineering, product, and business stakeholders to turn ambiguous problems into shipped capabilities used by thousands of global customers.
- Data science lead for [Falcon Next-Gen SIEM's first ML-powered product feature](#), owning model development, evaluation, explainability, monitoring design, and delivery for security-analyst triage.
- Built first-generation production ML for CrowdStrike Data Protection, including USB exfiltration detection and a [TensorFlow source-code classifier](#) that reduced training time from 227 hours to 6 hours.
- Researching NVIDIA Nemotron models for agentic triage use cases, fine-tuning with PyTorch on GCP GPU infrastructure, and evaluating outputs for security-analysis workflows.
- Led cloud workload anomaly detection research using unsupervised learning over billions of container events, resulting in a published CrowdStrike research article; mentored data science interns on reproducible, version-controlled delivery.
- Earlier CrowdStrike work included scoping cross-functional data science initiatives and building a calibrated XGBoost churn-risk model with SQL ETL and SHAP explainability to identify high-risk customers for proactive retention outreach.

Viasat – Senior Data Scientist

Sep 2017 – Aug 2021

Promoted from Data Scientist to Senior Data Scientist.

- Built large-scale geospatial, optimisation, and predictive modelling systems for network-planning decisions, including a Global Demand Model identifying 2M target communities from 1B+ data points and a linear-programming model projecting an 18% increase in target market opportunity.
- Developed SQL/Postgres ETL workflows, data architecture, Python applications, XGBoost models, and distributed simulation frameworks, including a speed-test simulation framework with 100x throughput improvement.

Teradata – Software Engineer

Jan 2015 – Sep 2017

- Built SQL developer tools, big-data platform features, and UI components for a next-generation business intelligence product, using Java, TypeScript, Node.js, and distributed-systems concepts to help users query, manage, and interact with large-scale data environments; won 1st Prize at the Teradata Analytics Hackathon in 2017.

FOOTBALL ANALYTICS & RESEARCH

Pitch to the Pros 3 – Finalist

2026

ringhilterra.github.io/projects/pitch-to-pros-3

- Selected as one of four finalists for Training Ground Guru's Pitch to the Pros 3 and presented live to football analytics professionals, including David Sumpter.
- Built "Was the Best Pass Really Playable?", a tracking-data decision-window framework that adds execution context to pass value by identifying when high-value options are available, constrained, or no longer playable.
- Analysed 10 A-League 2024/25 matches: 9.6M player-frame tracking samples, 47.9K dynamic events, and 24.4K linked passing options.
- Created practitioner-facing visuals and an analyst workflow for reviewing ranked alternatives, decision timing, playability constraints, and visual evidence.

Training Ground Guru – Football Analytics Consultant

Contract | 2026 – Present

- Contracted by Training Ground Guru to rebuild the Academy Productivity Rankings analysis and co-author the annual article.
- Building automated scraping, ingestion, and integration workflows to analyse player and academy histories across FBref, Transfermarkt, and other public sources.
- Developing an updated scoring methodology, reproducible ranking pipeline, QA process, and source-tracking workflow to evaluate player-development contribution across the English academy system.

Support Creation Dynamics in Football – Research Collaborator

2026 – Present

ringhilterra.github.io/projects/support-creation-dynamics

- Collaborating with Rui Marcelino, Professor at the University of Maia and sports scientist with the Portuguese Football Federation men's beach soccer team, on off-ball support creation research.
- Developing a temporal Support Opportunity Quality framework to quantify how off-ball movement creates, improves, weakens, or removes support options during possession.
- Using PFF FC 2022 FIFA World Cup tracking and event data to evaluate whether support creation dynamics relate to possession retention, territorial progression, and threat generation.

Active Support Index – SkillCorner Analytics Cup Project

2025

ringhilterra.github.io/projects/asi

- Created the Active Support Index, a metric framework for quantifying off-ball teammate support during pressure events using SkillCorner tracking and dynamic event data.
- Analysed 10 A-League 2024/25 matches, 7,063 pressure events, and 10fps tracking data; validated ASI against season-level physical aggregates from 175 matches, showing a strong movement-availability signal ($r = 0.74$, $p < 0.001$, $n = 167$).
- Built reusable SkillCorner tracking/event processing and visualisation code that became the foundation for the Pitch to the Pros 3 project.

Possession Foul-Win Probability Model – StatsBomb Event Data Project

2024

ringhilterra.github.io/projects/foul-win-probability

- Built an end-to-end predictive modelling project estimating the probability that a player wins a foul after receiving or recovering possession using 2015/16 English Premier League event data.
- Framed the problem, extracted possession-level examples, engineered spatial/contextual features, and trained an XGBoost classifier with time-based validation and cross-validation.
- Evaluated model quality through calibration, error analysis, spatial foul-rate visualisations, and SHAP interpretation; delivered a reproducible Python/Jupyter notebook and HTML report.

PUBLICATIONS & RECOGNITION

- Published CrowdStrike technical articles on [source code classification training optimisation](#) and [cloud workload anomaly detection research](#).
- 7 pending patents in applied machine learning and cybersecurity, including published applications in source-code classification, behavior-based asset classification, anomaly detection, graph ML, and interpretable alert grouping; selected filings at ringhilterra.github.io/patents-publications.
- XLR8 "Fanatical About the Customer" Award, CrowdStrike company-wide Innovation Week competition.

CORE TECHNICAL STRENGTHS

Predictive Modelling & Applied AI: supervised and unsupervised learning, classification, gradient boosting, clustering, anomaly detection, calibration, explainability, scikit-learn, XGBoost, PyTorch, TensorFlow

Full Modelling Lifecycle: problem framing, feature engineering, model validation, reproducible pipelines, automation, monitoring, stakeholder adoption, production delivery

Football Analytics & Research: tracking/event data, spatiotemporal analysis, action valuation, pass value and decision context, off-ball support dynamics, player movement modelling, academy player development analysis, analyst-facing tools

Data, Cloud & Engineering: Python, SQL, Pandas/Polars, DuckDB, PySpark, Docker, Git, CI/CD, AWS/GCP, Airflow

Decision Support & Technical Leadership: research-to-product translation, stakeholder engagement, visual explanation, analyst workflow design, mentoring, coding standards, technical and non-technical communication

EDUCATION

- University of California, San Diego — Master's degree, Data Science and Engineering, GPA 4.0, 2017–2019.
- Stanford University — Graduate Professional Program, Artificial Intelligence, 2020–2022.
- University of California, San Diego — B.S. Computer Science, Minor in Economics, 2010–2014.