ANDREA SCHOLTZ

SENIOR MACHINE LEARNING ENGINEER

DETAILS

ADDRESS
Somerset West
Cape Town
South Africa

PHONE
+27 82 571 4559

EMAIL
andrea.scholtz1@gmail.com

DATE OF BIRTH
8 October 1986

NATIONALITY
South African

LANGUAGES

Afrikaans
●●●●●

English
●●●●●

LINKS

LinkedIn

SKILLS

MLOps

Python

CI/CD, Git

MLFlow

Pytorch, Keras, Sklearn

Kubernetes (EKS), Docker

Cloud computing (AWS, Azure)

Data-driven Product Development

Communication

Leadership

PROFILE

As a seasoned Machine Learning Engineer with a passion for MLOps Engineering and automation, I have spearheaded teams to revolutionise ML model productionisation. My leadership created a world-class model monitoring service, slashing model drift detection from weeks to hours. I pioneered an Automated Model Deployment Pipeline, a robust and efficient software project that streamlines model deployment. My passion for automation led to a service that exponentially boosted model training capacity, significantly accelerating model development timelines. I pride myself on crafting user-friendly interfaces that simplify complex processes, making model productionisation seamless for engineers. My work consistently delivers efficiency, productivity, and innovation, and I'm excited to bring these qualities to my next role.

EMPLOYMENT HISTORY

Senior Machine Learning Engineer (MLOps), Prediction and Portfolio Services, Jumo

Cape Town

Apr 2022 — Present

Role:

- Served as the **Product Owner** for the **Prediction Services Team**, leading **MLOps** and fostering collaboration with the Decision Science Team.
- Applied **deductive reasoning** to develop a set of **champion model architectures** that could be **generalised** across various modelling problems.
- Implemented initiatives to **enhance team services**, **improve efficiency** across teams, and **automate model deployment pipelines**.
- Took on a **mentorship** role, **supervising** and **guiding** junior software engineers.

Achievements:

- Implemented an **automated model deployment pipeline**, which **drastically reduced deployment time** from several days **to mere hours**.
- Developed a **scalable training service** for the modelling team, which **increased their productivity** from a few models per week to **up to 50 per day**.
- Established a **model monitoring pipeline** for timely **detection of drift** in predictions and features.

Projects:

- Led the development of the **Automated Model Training Pipeline v2.0** which **streamlined** the **model training process** within the Data Science and Analytics Team.
- Developed a **model monitoring solution** that enabled us to **detect** data anomalies in **minutes**, a **significant improvement** from the previous weeks-long detection period.
- Created a platform for the Decision Science Team to **train multiple ML models** at scale.

Machine Learning Engineer, Jumo

Cape Town

Jun 2021 — Mar 2022

INTERESTS

Continuous Learning

Reading current research

Life Sciences

Learning languages

AI Ethics

Music

Cooking

Role:

- Created **data and ML solutions** for **Mobile Credit**.
- Automated **scorecard deployment** and **ML pipelines**.
- Fostered **close collaboration** with decision scientists and **cross-functional teams**.
- Developed **production pipelines** for models in **cloud environments**.
- **Operationalised ML models** for client projects.

Skills obtained:

- Invented **novel ways** of doing complex tasks.
- Took **ownership and accountability**.
- Solved **complex problems** in an **ultra-fast-moving environment**.
- Rapidly prototyped solutions that added **significant value** to **scale and productivity**.

Projects:

- Integrated multiple systems to **deploy Machine Learning models to production** quickly and efficiently.
- Developed an efficient **production pipeline** for deploying mass-scale ML models for predicting the probability of default on different segments.

Machine Learning Engineer, Praelexis (Pty.) Ltd.

Stellenbosch

Oct 2016 — Jun 2021

Role:

- Developed **data and machine learning solutions** for **Retail banking (Capitec)** and **Private banking (Investec)**.
- Developed **production pipelines** for models in **cloud environments**.
- Delivered support for our **proprietary Machine Learning platform**.
- Designed **software and systems architecture** and **infrastructure**.
- **Conceptualised** and **implemented machine learning models** for client projects.
- Managed **data exploration**, management, and manipulation.

Skills obtained:

- Invented **novel ways** of doing complex tasks.
- Worked in a **multidisciplinary team**.
- Took **initiative** and managed projects holistically.
- Took **ownership and accountability**.

Projects:

- **Cognitive Banking platform for Investec** - Developed and productionised a **client-centric cognitive platform** for one of the leading private banks in South Africa.
- **Cognitive Banking platform for Capitec Bank** - Developed and productionised a **client-centric cognitive platform** for one of the leading retail banks in South Africa.
- **Statement Automation modelling for Capitec Bank** - Developed and productionised a **proprietary Machine Learning model** for one of the leading retail banks in South Africa.

Lecturer, North-West University

Potchefstroom

Jan 2014 — Sep 2016

About this role:

- Lectured various **Pure and Applied Mathematics** modules to undergraduate students.

Junior Lecturer, CTI Education Group

Potchefstroom

Aug 2011 — May 2014

About this role:

- Lectured **Statistics and Mathematics** for social sciences to undergraduate students.

PROJECTS

Automated Model Training Pipeline v1.0, Machine Learning Engineer, Jumo

Cape Town, South Africa

Oct 2021 — Dec 2021

In this project, I addressed a bottleneck in our data science workflow that was significantly limiting the data science team's productivity. Previously, our data scientists would manually launch an EC2 instance and set up a machine learning model for training. This was a time-consuming process - typically it took several hours for a data scientist to train a single model. Evaluating the results of this process usually led to a throughput of just one model per day per data scientist.

To optimize this process, I developed an API using FastAPI with Python. The API is integrated with the Kubernetes API to automatically launch an EKS job for training a specified model pipeline.

This innovative approach allowed us to initiate multiple model training runs via the API. The model pipeline to be trained could be easily specified through distinct configuration files, enabling us to train a variety of models simultaneously. As a result, we significantly increased our model training capacity and reduced the time spent per model, with the only limitation being the resource constraints of our Kubernetes cluster. We could launch at most 50 training runs on a test dataset.

Through this project, we were able to greatly enhance our data science team's productivity and efficiency.

Model monitoring solution, Senior Machine Learning Engineer, Jumo

Cape Town, South Africa

Aug 2022 — Nov 2022

In this project, I spearheaded the development of a model monitoring solution which drastically reduced data anomaly detection time from weeks to mere minutes. This was a significant enhancement for our operational efficiency.

This solution automated the entire batch scoring process using Airflow, which orchestrated Amazon EKS Jobs for each segment of our system, including feature materialisation, model scoring, and monitoring subprocesses.

To ensure comprehensive and reliable data monitoring, I utilized the Great Expectations library along with a custom API that I developed. This combination was used to validate datasets, as well as to compute the Population Stability Index (PSI) for model scores and the Characteristic Stability Index (CSI) for features with every daily run.

In addition, the project involved creating an automated mechanism for our Airflow DAGs. Whenever a new model was deployed to our proprietary registry, a corresponding DAG was automatically created and scheduled to run the next day.

Lastly, to provide a comprehensive view of our model's performance, all metrics were logged to a DataDog dashboard. This enabled us to monitor model performance consistently and make data-driven decisions for improvements.

Automated Model Deployment Pipeline v1.0, Senior Machine Learning Engineer, Jumo

Cape Town, South Africa

Oct 2022 — Dec 2022

In this project, I significantly enhanced our model deployment process, which had been a time-consuming and labor-intensive task requiring about a week of work per engineer for each model deployment.

Initially, the task involved establishing a formalized Model Life Cycle process. Subsequently, I designed and implemented an Automated Model Deployment pipeline that leveraged our team-developed, proprietary Artifact Registry.

Our data scientists would train models and log them to MLFlow. Once logged, my automated process would retrieve the model and associated feature set artifacts from the underlying S3 location corresponding to the experiment on the MLFlow server. The artifacts would then be converted into packages compatible with our Artifact Registry. Following the conversion, the artifacts would be uploaded to the registry and initial tests would be conducted in the Development environment.

Only after all checks, balances, and model validation steps were successfully passed, the model would then be deployed to the Staging environment. Here, our decision science team could validate the model scores.

Upon successful verification of all tests and checks in the Staging environment, the model would be deployed to Production. Once in Production, a script on the Airflow server would identify and pick up the new model, decommission the old model, and automatically generate a DAG scheduled to run the next day.

This newly streamlined process, which previously took up to a week, was now completed within hours. Furthermore, this process is parallelizable, allowing for the simultaneous deployment of multiple models. This represented a significant boost in efficiency and productivity for our team.

Automated Model Deployment Pipeline v2.0, Senior Machine Learning Engineer, Jumo Cape Town, South Africa

Apr 2023 — Present

In the original project scenario, our data scientists had to undertake a laborious and time-consuming multi-step process before they could commence actual model training. Firstly, they needed to create features from relatively large data sets for specific time windows. Our supportive engineering team helped to implement a process to manage this task utilizing EMR on EKS.

Once the features were prepared, the data scientists then needed to train various models and select the optimal one for deployment. Upon experimentation, I noticed that certain model architectures consistently performed well with our data and seemed to generalize well across different scenarios, so I constructed some "champion" model architectures using deductive reasoning.

Recognizing the tediousness and inefficiency of the existing process, I decided to introduce automation into the system to streamline it. Although there is no frontend at the time, the provided logs illustrate the operations of the service.

The automated service triggers the data/feature creation process. Once the Metaflow DAG completes and the data has been prepared, the experiment is automatically initiated in MLFlow. This process triggers multiple model training runs, each with distinct model configurations based on the successful architectures I'd identified earlier.

The result is a significantly accelerated process, with candidate models available for review and deployment within hours instead of days or weeks. This automation will greatly improve our productivity and efficiency in training and deploying machine learning models.

EDUCATION

PhD, Bioinformatics and Computational Biology, Stellenbosch University Stellenbosch, South Africa

Jan 2023 — Present

I am currently enrolled in a postgraduate programme in Bioinformatics and Computational Biology, with the goal of completing a PhD.

Under supervision of Prof. Hugh Patterton.

M.Sc., Applied Mathematics (cum laude), North-West University Potchefstroom, South Africa

Jan 2013 — Oct 2015

Specialization in Numerical Analysis, Control Theory, Optimization and Mathematical Epidemiology.

Under supervision of Prof Ilse Schoeman and Prof Japie Spoelstra.

Thesis available on request.

B.Sc.Hons., Applied Mathematics (cum laude), North-West University Potchefstroom, South Africa

Jan 2012 — Dec 2012

B.Sc., Computer and Mathematical Science (cum laude), North-West University Potchefstroom, South Africa

Jan 2007 — Dec 2012

REFERENCES

References available upon request