

Hung Ba Huynh

+61 432 077 104 | [✉ hung.ba.huynh@proton.me](mailto:hung.ba.huynh@proton.me)
[🌐 hbhungg](#) | [📄 Ba Hung Huynh](#) | [📍 Melbourne, VIC](#)

Experience

Data Engineer

November 2023 - Present

Beyond Zero Technologies — Vermont, Melbourne, VIC

- Spearheaded the creation of our first data platform using Dagster and DBT to orchestrate SQL and Python jobs, deploying it on ECS for on-demand scaling
- Currently acting as the lead maintainer for both the data platform infrastructure and pipelines, collaborating with stakeholders and teams to develop new data pipelines and features.
- Developed an in-house programming language (similar to Excel, compiling to Polars and SQL) to enable non-technical stakeholders and teammates to quickly create high-level data points from low-level sensors and onboard new clients through our internal admin portal.
- Established standards for data analysis and data engineering workflows, onboarding and mentoring junior Data Analysts and Data Engineers to collaborate effectively and be productive.
- Planned and created an MVP using Iceberg and AWS EMR (Spark) to meet future scaling needs.

Cloud Infrastructure Engineer

April 2023 - Present

Beyond Zero Technologies — Vermont, Melbourne, VIC

- Led the project to design and migrate over 1TB of telemetry from AWS RDS (PostgreSQL) to TimescaleDB Cloud, which now serves as our main production database, achieving 5-10x performance on common telemetry queries.
- Spearheaded the IoT telemetry collection project, collaborating closely with on-site BMS engineers to transition on-premises computing resources to cloud infrastructure using AWS Greengrass and AWS IoT.
- Developed and deployed Python microservices on on-premises machines, enabling efficient communication with industrial sensors and streamlining data integration processes. Ensured seamless integration with customers' industrial sensors within BACnet and Modbus IoT networks.
- Worked with hardware team on research and development of custom IoT hardware solution using LoRaWAN, integrated with current AWS IoT stack.
- Implemented near real-time ETL pipelines using AWS MQTT and AWS API Gateway to transmit millions of IoT messages and handle thousands of API calls daily from edge machines, ensuring robust connectivity with cloud platforms
- Currently acting as the lead maintainer bridging the IoT collection stack and the data platform stack.

Software Engineer

February 2023 - April 2023

Beyond Zero Technologies — Vermont, Melbourne, VIC

- Developed internal monitoring tools to identify failures and bugs in Edge devices, comprising an agent written in Go sending telemetry to Grafana.
- Demonstrated expertise in backend development and API design using NestJS and TypeScript, contributing to the creation of diverse services.
- Maintain and improve Github CI/CD pipeline. Transition into cloud infrastructure team.

Data Scientist

Jan 2022 - Dec 2022

Monash Microgrid @ Monash University — Clayton, Melbourne, VIC

- Analyzed, processed, designed, and developed forecasting models using Monash's weather data, building energy usage, and solar panel energy production with tuned Random Forest and Gradient Boosting regression, achieving the highest prediction accuracy in the cohort.
- Applied mixed-integer linear programming to forecasted data to optimize building activity schedules (an NP-hard problem), aiming to minimize energy usage with Google's OR-Tools.
- Lead a team of 4 data and software engineers, under the supervision of Dr. Christoph Bergmeir. The project is open source under MIT license at [tinygrid](#).

Technical Skills

Programming Languages: Python, Typescript/Javascript, SQL, Go, C, Verilog, Bash

Data Engineering: DBT, Dagster, Polars/Pandas, PostgreSQL, TimescaleDB, NumPy, PyTorch, Apache Spark, Hive, Iceberg

Developer Tools: Linux/Unix, Git, Vim, Tmux, Terraform

AWS Cloud: EC2, IoT Core, Greengrass, Lambda, SQS, IAM, S3, RDS, ECS

Web Frameworks/Libraries: AWS GreengrassSDK, NestJS, Flask, FastAPI

Theoretical: Deep Learning, Statistical Analysis, Computer Theory, Hardware Design, System Architecture, Mathematic

Projects

[smalltensor](#) | Dec 2022 - Present

Python, Numpy

A basic Python Deep Learning library (tensor + autograd), similar to PyTorch, from the ground up.

[config](#) | Nov 2020 - Present

Bash, Neovim/Vim, Rmux, Git, Alacrity

Configuration files for a personal development environment.

[riscv-core](#) | January 2024 - Present

Python, Verilog

Simulate RISC-V 32-bit core in Python (and Verilog).

[DeepChess](#) | Nov 2021 - Present

PyTorch, TinyGrad, MLFlow

Recreating [DeepChess](#) paper.

Education

Monash University

Jan 2019 - Dec 2022

Bachelor of Computer Science — Clayton, Melbourne, VIC

Specialize in Data Science, minor in Computer Network and Security

Achievements & Activities

Advent of Code

Dec 2018 - Present

- Yearly participant in [Advent of Code](#).

UNIHack 2023 (Rank 3rd)

March 2023

- Worked in teams of 6 on Drowzalert, a solution to prevent drowsy driving and increase road safety
- Served as software engineer lead, responsible for webcam, facial detection and drowsiness prediction.
- Project is open source, technical write up at [Devpost](#)

MDSS & Servian Datathon Advance Stream (Rank 1st)

April 2022

- Fine-tuned pre-trained YOLOP models to classify images as drivable.

Google Hash Code (Rank 10th in Australia)

Feb 2022

- Compete in Monash competitive programming team to solve NP optimization problem in a 3-hour competition.

FIT3181 Kaggle's Animal image classification competition (Rank 1st)

Sept 2021

- Trained and fine-tuned an EfficientNetB0 model to classify low resolution images of different animals.

Future Goals & Interests

Join MLOps team.

Create Deep Learning hardware training/inference accelerator with RISC-V architecture.

References

References available upon request