

Experience

Software Engineer Dune Jan 2021–Present

Remote

- Worked on the Data Ingest team, doing everything related to Spark ETL jobs streaming 100s of TBs of data. Designing new datasets, creating and operating Spark jobs.
- Designed, implemented and maintained the indexing job that powers search across 1.5m tables on dune.com.
- Wrote and open sourced [Harmonizer](#), a Python library for translating SQL queries from Spark and Postgres dialects into Dune's extended Trino dialect, by manipulating query ASTs. I used it to migrate 69 thousand lines of SQL code. Started as a full-stack feature I built which won 2 prizes in our 2023 hackathon. It quickly became an integral part of our strategy in Q2 of 2023 in helping users migrate their queries to our new query engine.
- Designed, implemented and maintained a table migration system for the data warehouse, with assisting command line and CI/CD tooling.
- In my first three months, I designed, implemented and maintained the initial version of the SQL query execution service. The architecture increased transparency and reliability into execution of queries, and reduced the number of technologies we depended on (got rid of Redis, RabbitMQ). This code was integral to the company and handled millions of query executions, and the caching scheme saved many millions of executions.
- Joined as second employee, currently at ~50 employees. Dune data is frequently cited in FT and other news outlets.
- Tech stack: Go, Kotlin, Rust, Python, Apache Spark, Apache Arrow, Delta Lake, AWS S3, Databricks, GitHub Actions, Postgres, Docker, Kubernetes.

Software Engineer II Cognite Jun 2019–Dec 2020

Oslo, Norway, and remote

- Built a document ingestion and search API used to search millions of documents. Wrote most of [its documentation](#).
- Wrote a high-throughput sync service for Postgres and Elasticsearch, using Postgres transaction IDs for atomicity.
- Drove Python developer experience in company: CI integrations, typing, linting, formatting.
- Held 100+ interviews and ran initial part of software engineer hiring pipeline; automated scheduling of interviews.
- Cognite grew from 150 to 400 employees while I was there.
- Tech stack: Python (with mypy), Postgres, Elasticsearch, Docker, Kubernetes, GCP, Jenkins.

Education

- **M.Sc. Statistics**, University of Oslo, Norway. GPA: A. 2017–2019
A paper summarizing my thesis work was [published](#) in 2022 as “A boosting first-hitting-time model for survival analysis in high-dimensional settings” in the [Lifetime Data Analysis](#) journal.
- **B.Sc. Mathematics and Computer Science**, University of Oslo, Norway. 2014–2017
Courses in applied math, computer science, and statistics. Completed courses for 350 credits, regular load is 300.

Personal Projects

- **Technical writing for LWN.net**: An article on [Go 1.18 for LWN.net](#), a magazine with 500k monthly visits.
- **Technical blogging** on vegardstikbakke.com: Two posts featured on the front page of Hacker News (HN), the biggest social news website for tech. [Problem solving with Unix commands \(HN\)](#) and [How do Unix pipes work? \(HN\)](#).
- [extract-table.com](#), a website and API for extracting tabular data from images. Has been used by thousands.