

About me

I'm a staff-level software engineer, founder and consultant. I'm curious, product-minded and entrepreneurial, dedicated to- and motivated by solving complex problems that matter (what Seymour Papert called "hard fun"). I feel at home in teams working on or with large distributed systems, large-scale (scientific) computing and data processing infrastructures, cryptosystems, and even a little mechatronics. I'm currently available for full-time positions and contract work.

Employment history

Software Engineer / Consultant / TPM

Freelance

June 2006 – present

I intermittently do contracted assignments in a variety of roles. Some of my past clients include FEI Company (now [ThermoFisher](#)), [ING](#)'s Global CISO organisation, the Datalab of City of Amsterdam, the Department of Traffic Management of the City of The Hague, the [Information and Language Processing Systems \(ILPS\)](#) group at the University of Amsterdam, and others. See below for a selection of work.

- ▷ Cyber Security
- ▷ Big Data
- ▷ Hadoop
- ▷ Spark
- ▷ Scientific Computing

Engineering Manager

Uber

October 2021 – March 2024

I was manager of the Earnings Data Intelligence team, a team of talented backend engineers. The team built and ran the systems that powered several critical earnings-related features for Uber's earners, i.e. drivers and couriers. Important features included real-time trip earnings summaries and historical earnings reports, but also highly localized earnings forecasts generated using large-scale machine learning pipelines. At the core of the data was a complex taxonomy that models earnings specific to all of the 70+ markets that Uber still operates in, allowing Uber and its earners to comply with all rules and regulations specific to the earners locality. As engineering manager I was responsible for all aspects of this team's health, from the psychological safety and personal growth of its members to the team's vision, mission and execution, both tactical and strategic.

- ▷ Engineering management
- ▷ Complex taxonomies
- ▷ Mission critical real-time features
- ▷ Machine Learning
- ▷ Team health
- ▷ Vision, mission, execution

Software Engineer

Uber

July 2018 – October 2021

I was a senior / lead engineer on the Tokenizer team, a small team of 5 backend engineers and 4 devops engineers, on its own little hyper secure island. We designed, built and ran Tokenizer, Uber's bank card vault, which held 100s of millions of card numbers, bank account numbers and other tier 1 sensitive PII like SSNs. Tokenizer was (and to my knowledge, still is) on the critical path for both collections (from riders and eaters) and disbursements (to drivers and couriers). Given the very large scale of card processing at Uber, Tokenizer contributed huge savings by not depending on third party vaults and by allowing integrations with many different local card processors and banks around the world, without trading off stability, availability and security (Tokenizer is, of course, PCI compliant and as such is extensively audited every year). I worked on many features of the system over the years, both lower-level system features and higher-level user-facing features. Some example projects I led include the integration with Uber's service authentication and authorization, making it much easier for upstream systems to use Tokenizer; the migration of users' payment cards in their Postmates wallets after Uber acquired Postmates; and building a secure banking API for the Uber Bank project.

- ▷ PCI-DSS Compliance
- ▷ Cyber security
- ▷ Payments
- ▷ Fintech

Co-founder

Lucipher

May 2013 – February 2016

Inspired by [the Snowden leaks](#), Lucipher was a startup that addressed the need for end-to-end secure and confidential cloud-based file storage and collaboration. The company built and provided an online (NFS-based) file system that could be mounted locally, complete with identity management and exchange, all secured with a hybrid cryptosystem that was completely transparent and could be inspected on the wire. In other words: a verifiably secure, zero-trust system. Lucipher won an IBM innovation prize and was featured in a number of online and printed publications, but was not able to gain enough traction to survive. See (Dutch language) interviews in [Automatiseringsgids](#), [IBM Inspire](#) and [Computable](#).

- ▷ Startup
- ▷ Fundraising
- ▷ Cyber security
- ▷ File systems
- ▷ OpenPGP
- ▷ Big Data

Scientific consultant

SARA

October 2008 – March 2013

[SARA](#), now part of SURF, is the Dutch national centre for high-performance computing (HPC) and owns all national scientific computing infrastructure. At SARA I worked with academic researchers across different disciplines, including High Energy Physics, Astronomy, Bioinformatics, Information Retrieval, Ecology, and others. It was my job to help solve very large-scale computational problems on some of the largest distributed systems in the world. In 2010, with national funding, I created a new team to build what was then the largest Hadoop cluster in The Netherlands, to better serve researchers with relatively trivial (often machine learning heavy) processing of extremely large datasets, including genome sequences, web crawls (the complete [Common Crawl](#) and [ClueWeb'09](#) datasets), and other popular datasets. As part of my work at SARA I organized a number of training events in large-scale scientific computing for (post)grad students and professionals.

- ▷ High Performance Computing (HPC)
- ▷ Scientific computing
- ▷ Hadoop
- ▷ Spark
- ▷ Big Data
- ▷ Teaching

Selected projects

AuthZ

Amsterdam Datalab

July 2017 – June 2018

An OAuth 2.0 compatible authorization service written in Go (see the [website](#)). This service:

- 1: Implements the implicit flow as described in RFC6749 section 4.2,
- 2: Uses JSON Web Keys (RFC7517) for key management,
- 3: Creates JSON Web Tokens (RFC7519) using HMAC or ECDSA (HS256, HS384, HS512, ES256, ES384, ES512), and
- 4: Provides interfaces for identity providers, state storage (with implementations for single node in-memory storage and Redis), authorization providers (to map users to scopes) and client registries.

- ▷ OAuth 2.0
- ▷ Authorization
- ▷ Microservices
- ▷ Open Gov Data
- ▷ Cyber Security

Core Intel

ING Global CCERT

October 2014 – March 2016

An intrusion detection system that detects possible malicious activity on the internal ING network using both open- and closed source threat intelligence, and that analyzes traffic patterns of strategically collected, large volumes of Netflow (V5).

- ▷ Cyber Security
- ▷ Netflow
- ▷ Threat detection
- ▷ Hadoop

Semanticizer

University of Amsterdam, ILPS

November 2012 – November 2013

A service for real-time semantic linking. It is able to take in a text stream, and return links to online resources (such as Wikipedia) that provide background information on the subject(s) of the input, including heuristic measures that estimate the likelihood of a link being correct. See more background in "[Feeding the Second Screen: Semantic Linking based on Subtitles](#)" (D. Odijk, E. Meij, M. de Rijke, University of Amsterdam, 2013).

- ▷ Named Entity Recognition (NER)
- ▷ Information Retrieval (IR)
- ▷ Natural Language Processing (NLP)
- ▷ Machine learning

LifeWatch.eu

University of Amsterdam, IBED

March 2011 – March 2013

[LifeWatch](#) is an [ESFRI](#) project to build a modern European digital infrastructure for biodiversity research in the EU. I did technical management of the project, including significant fundraising for the infrastructure in The Netherlands.

- ▷ Research Infrastructure
- ▷ Fundraising
- ▷ Biodiversity

SCULPTOR

MTA SZTAKI, CIM

February 2006 – February 2007

The SCULPTOR project researched a dynamic system for sheet metal forming. Instead of relying on the traditional two-stamp process, SCULPTOR proposes a vertical arrangement of two robotarms each equipped with tools to massage the sheet into a desired form. In this setup, one tool shapes the metal from above while the other provides necessary support from below. (See [this presentation](#) for a short intro.) MTA SZTAKI is the research institute of the Hungarian Academy of Sciences (MTA) focused on computer science and control systems.

- ▷ Research
- ▷ Robotics
- ▷ Mechatronics
- ▷ RT-Linux

Outreach

A curated archive of my engagements in teaching, conferences, interviews, meetups, and other related events (see more details on [my website](#)):

- ▷ Cyber Security Sessions, 2015 (Meetup organizer)
- ▷ Norvig Web Data Science Award, 2014 edition (Review committee)
- ▷ Hadoop Summit Europe 2014 (Chair of the Future of Hadoop track)
- ▷ Norvig Web Data Science Award, 2013 edition (Organiser, review committee)
- ▷ Hadoop Summit Europe 2013 (Chair of the Operating Hadoop track)
- ▷ Apache Drill development workshop, 2013 (Organiser)
- ▷ NL-HUG, Netherlands Hadoop User Group, 2010-2013 (Founder, organizer)
- ▷ HPC Hadoop course, 2012 (lecturer)
- ▷ IIR Big Data Analytics, Hadoop lectures, 2012 (lecturer)
- ▷ SIKS Big Data course for PhD students, 2011 (Organizer, lecturer)
- ▷ Several LHC grid tutorials, 2009-2013 (Lecturer)