

Bart de Goede | Resume

New York NY, United States

✉://bart@degoe.de in://bart-de-goede 🌐://bartdegoede https://bart.degoe.de

Summary

I am a staff software engineer with more than 12 years experience. I co-founded a web development agency, worked as a senior software engineer on the search and data teams of one of the fastest growing unicorns in Europe, on the search team at a digital library, e-book and audiobook subscription service, the growth team at a North American grocery delivery service, ran search at OnFrontiers, and help LLMs answer questions at Perplexity AI. On the side I'm I love building scalable systems, designing sensible APIs, working on products that users can't live without, have a soft spot for search and recommendation systems, writing blog posts, helping others succeed, and applying machine learning in such a way it feels like magic.

Skills

Natural languages English (fluent), Dutch (native)

Programming languages Python, Ruby, SQL, Javascript, Go, Scala

Frameworks Django, Flask, FastAPI, Rails, Grape, D3, React

Databases MySQL, PostgreSQL, Elasticsearch/OpenSearch, Solr, MongoDB, Redis, DynamoDB, Vespa

Tools and versioning Git, Mercurial, Virtualbox, Vagrant, Ansible, Fabric, Chef, AWS, Terraform

Other keywords Agile software development (Scrum, Kanban), user-focused, AB-testing, Eppo, Nginx, ELK, Resque, Grafana, Datadog, (micro)services, architecture, (RESTful) API design, NLTK, Gensim, scikit-learn, Spark, Airflow, Scrapy, asyncio, Linux (server) environments, bash scripting, Unicorn, uWSGI, RabbitMQ, Kafka, React

Experience

Perplexity AI

Member of Technical Staff

April 2024 – Present

Role: I work on the answer engine at Perplexity, where we turn LLMs and search into the most powerful knowledge tool on the web.

Interesting challenges: Optimizing the file upload flow, helping the LLMs understand PDFs better.

Skills used: Python, FastAPI, AWS, (Postgre)SQL, Redis, Vespa, Snowflake

OnFrontiers

Head of Search Engineering, Search

April 2023 – April 2024

Role: I lead search engineering at OnFrontiers, where we help organizations gain insight into the knowledge and experience of their employees. We leverage LLM-generated knowledge and descriptions in a custom taxonomy to do semantic search and expansion on user knowledge profiles.

Interesting challenges: Architected, designed and built the search APIs from scratch, reducing indexing time from over an hour to less than 10 seconds, and lowering request latency 8x. I maintain the Opensearch clusters, search ETL pipelines and tune the Opensearch queries.

Skills used: Python, FastAPI, OpenSearch, AWS, (Postgre)SQL, Redis, React/Javascript, OpenAI

Instacart

Senior Software Engineer II, Growth

October 2021 – April 2023

Role: Senior software engineer on the Growth team, sending 75+ million relevant messages to our users. This includes emails, push notifications and in-app merchandising.

Interesting challenges: Built and scaled a system for personalized email content with DynamoDB and Sidekiq. This system lets us send over 10 million emails with personalized groceries within an hour to our entire user base.

Designed a system that automates email and push notification campaigns for dozens of retailers on the Instacart Platform.

Skills used: Ruby, Python, RSpec, Airflow, (Postgre)SQL, Redis, Memcached, React/Javascript, Braze, Snowflake

Scribd

Senior Software Engineer, Search and Web Infrastructure

October 2018 – October 2021

Role: Senior software engineer on the search and web infrastructure teams, making sure our 250 million monthly users can find the content they're looking for in Scribd's vast catalog. This includes developing search and ranking features, as well as being the team subject matter expert for Elasticsearch.

Interesting challenges: Migrated Elasticsearch and search applications to AWS in less than 6 months. This increased the resiliency, reliability and scalability of the Elasticsearch clusters, reduced downtime risks, and unlocked testing different analysis pipelines and relevance tuning experiments at scale. Rebuilt controller and frontend architecture for content preview pages in 8 weeks, allowing for rollout of other content types in 2 days, while improving core web vital metrics.

Skills used: Ruby, Python, RSpec, Elasticsearch, (My)SQL, Redis, Spark, Chef, Go, Terraform, React/Javascript

Catawiki

Senior Software Engineer, Search & Discovery and Data Science teams

February 2016 – September 2018

Role: Senior developer responsible for site search, auction and category pages, and the integration of data science models with the Catawiki platform. This includes technical responsibilities for ETL, query building and performance optimization, product responsibilities such as integrating business requirements with search rankings and team responsibilities such as cross-team mentoring, pair programming and code reviews.

Interesting challenges: Migrated the site search from the legacy monolith application with Solr to a separate microservice using Elasticsearch. The migration was required to have no downtime, and have no negative impact on the business; resulted in increase of 47% in visitor-to-bidder ratio in sessions using search.

Company activities: Chief editor for engineering content and board member of The Catawikian, the internal magazine.

Skills used: Ruby, RSpec, Grape, Elasticsearch, Solr, MySQL, Redis, Python, Ansible, ELK

Dispectu

Co-founder

September 2012 – November 2015

Dispectu was a full-stack software development agency that specialized in web applications, search engine implementations and data visualization. We built search engines and interactive data visualization applications for a variety of customers. Examples include building an open source API for publicly available cultural open data, an [ngram-viewer](#) for 150 years worth of OCR-ed Dutch newspapers and natural language processing APIs.

University of Amsterdam

Research assistant

2010–2012

Role: Developing software for the PoliticalMashup research project, where we built a database of Dutch political data (parliamentary proceedings, biographical data, laws, etcetera).

Skills used: Python, eXist-db, XML, Xpath, Solr

Additional activities

Elasticsearch core training

2014

E&Y NextGen Academy

2011

Education

University of Amsterdam

Master of Information Science

2011–2012

Thesis: [Sustainable Questions](#)

Thesis published at SIGIR 2012 Workshop on Time-aware Information Access

University of Amsterdam

Bachelor of Information Science

2008–2011

Thesis: [Political slant in Dutch public broadcasting](#)

Student member of the Department Committee

2008–2011

The Department Committee is responsible for safeguarding and advising the department board and dean on the quality of teaching, organisation and the course curriculum.

Attended ESSIR in Koblenz (European Summer School in Information Retrieval)

2011

Publications

2013 Huurnink, B., Bronner, A., Bron, M., Van Gorp, J., de Goede, B. & van Wees, J. AVResearcher: Exploring Audiovisual Metadata. In *CEUR Workshop Proceedings* (2013), vol. 986, pages 1–2.

2013 De Goede, B., van Wees, J., & Marx, M. Politicalmashup ngramviewer. In *Proceedings of the 13th Dutch-Belgian Workshop on Information Retrieval* (April 2013), vol. 986, pages 54–55. Winner of the conference *Best demonstrator award*.

2012 Migut, M., van Wees, J., Bakker, D., de Goede, B., Steltenpohl, H., Oude Lenferink, N., & Worring, M.. VAST challenge 2012: Interactively finding anomalies in geo-temporal multivariate data. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, Seattle, WA, October 2012.

2012 Goede, B. de, Schuth, A. & Rijke, M. de (2012). Sustainable Questions: Determining the Expiration Date of Answers. In *SIGIR 2012 Workshop on Time-aware Information Access: #TAIA2012*, Portland, OR, August 2012. Accepted papers. Microsoft Research.

2011 De Goede, B., Marx, M., Nusselder, A., & van Wees, J. (2011). Succinct summaries of narrative events using social networks. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*, pages 299–304.