

# Resume

## Martin Krasser

I'm a freelance machine learning engineer and software developer. I support customers in all phases of their machine learning projects, from initial research, prototyping, iterative solution development up to serving and maintaining machine learning solutions in production. My current focus is on neural search engines and multimodal representation learning. I'm also interested in Bayesian methods for machine learning. My experience as distributed systems engineer comes from the development of numerous large-scale event-driven systems, including globally distributed systems that provide causal consistency guarantees. I'm a continuous learner and contributor to several open source projects.

<https://martin-krasser.com>

<https://twitter.com/mrt1nz>

<https://github.com/krasserm>

## Specialties

Neural search engines, multimodal representation learning, deep learning, machine learning, machine learning operations (MLOps), data-intensive applications, distributed systems, event-driven systems, event sourcing and CQRS, system integration, functional programming, object-oriented programming, probabilistic programming.

## Work experience

### Since 08.2023: Senior director machine learning

Canto, <https://www.canto.com/>, Freelance

- Responsible for AI/ML development of Canto's digital asset management system.
- ML research and engineering team lead.
- Core ML software developer.

#### **04.2018 - 08.2023: Lead machine learning engineer**

MerlinOne, <https://merlione.com>, Freelance

- ML research and engineering team lead, core ML software developer.
- Responsible for all ML research and development efforts at MerlinOne.
- Development of the Merlin Accelerated Intelligence (AI) Suite:
  - AI-powered, metadata-free multimodal search engine (NOMAD).
  - Low search latency under high load, scalability to very large asset datasets.
  - Large-scale training of asset encoding models on in-house news datasets.
  - Used as backend for AP’s AI-powered search since May 2023.
  - Visual similarity search based on global and local image features.
  - Face recognition in images and videos.
  - Image aesthetics assessment with models trained on custom aesthetics datasets.
  - Merlin AI Suite is used at many customer sites in production at large scale.

#### **Since 04.2011: Software engineer, Machine learning engineer**

<https://martin-krasser.com>, Freelance

- Research and development for industry machine learning projects (mainly deep learning).
- Software development services with a focus on backend software and distributed systems.
- See other positions since 2011 for further details.

#### **05.2017 to 04.2018: Machine learning sabbatical**

Sabbatical year

- Deep dive into mathematics, statistics, “traditional” machine learning and deep learning.
- Certifications received in online courses. Self-study of numerous books and papers.
- Publication of some exercise work in articles and open source projects.
- Equal focus on scientific and engineering aspects in machine learning and deep learning.

#### **02.2017 - 04.2017: Distributed systems engineer**

fine.lines.software GmbH, <https://bullet.io>, Freelance

- Design of a distributed workflow management system that also includes mobile devices.

- Event-sourced application with explicit causality for distributed data consistency.
- Consulting activities and development of prototypes.

**09.2014 - 08.2017: Distributed systems engineer**

Red Bull Media House, <http://www.redbullmediahouse.com>, Freelance

- Global distribution of the in-house digital asset management platform. Responsible for architecture, design and implementation.
- Conception and development of Eventuate, a toolkit for event sourcing and event collaboration at global scale with causal consistency.
- Development of a service framework for operation-based CRDTs.

**10.2015 - 07.2017: Software engineer and architect**

agido GmbH, <http://www.agido.com>, Freelance

- Development of an online sports betting web application.
- Architecture based on event sourcing and event stream processing.
- Usage of Eventuate as technical basis for the web application backend.
- Development of streaming data analytics extensions for Eventuate with Apache Spark.

**10.2013 – 02.2014: Software engineer and architect**

Lightbend, <http://www.lightbend.com/>, Freelance

- Conception and development of Akka Persistence.

**03.2012 – 10.2013: Software engineer and architect**

Eligotech BV, <http://www.eligotech.com>, Freelance

- Development of an E-Wallet management web service for customers in the online gambling industry.
- Conception and development of Eventsourced, an open source project for building reliable, scalable and distributed event-sourced applications in Scala.

**02.2011 – 04.2011: Senior software architect and engineer**

Talend, <http://www.talend.com>

- Development of scalability and high-availability concepts and prototypes for the Talend Integration Factory.

**01.2005 – 01.2011: Lead software architect and engineer**

InterComponentWare AG, <https://icw-global.com/>

- Conception and development of the Open eHealth Integration Platform (IPF), an eHealth integration platform on top of Apache Camel. Development and deployment of IPF applications in several international customer projects.
- Development of LifeSensor, a web-based personal health record.

**09.2000 – 12.2004: Senior software architect and engineer**

LION bioscience AG, <http://www.lionbioscience.com>

- Development and optimization of hierarchical clustering algorithms for the drug discovery pipeline of an international pharmaceutical company.
- Development of an integration platform for globally distributed chemical compute services.
- Development of an activity tracking web application for pharmaceutical research projects.

**02.1999 – 08.2000: Software developer, research assistant**

CAME, University of Salzburg, <http://uni-salzburg.at>

- Development of algorithms and software for ab-initio protein structure prediction.
- Application of statistical mechanics for structure optimization and evaluation.

**07.1998 – 01.1999: Software developer**

Austrian Red Cross, <http://www.rotekreuz.at/>

- Development of a patient transportation management system.

**Tools experience (excerpt)**

- Programming languages: Python, Scala, Java, C/C++, Haskell
- Machine learning: PyTorch, PyTorch Lightning, Hugging Face, LangChain, Tensorflow, JAX, scikit-learn, PyMC3, Spark MLlib
- Machine learning operations: AWS SageMaker, Weights & Biases, MLFlow
- Data engineering: Apache Spark, Dask
- Stream processing: Spark Streaming, Akka Streams, Kafka Streams
- Messaging: Apache Kafka, Amazon SQS, Apache ActiveMQ, ZeroMQ
- Databases: MySQL, PostgreSQL, Oracle Database, Apache Cassandra, Apache HBase, Amazon DynamoDB, MongoDB
- Cloud providers: Amazon Web Services (AWS), Google Cloud Platform (GCP)

- Java SE/EE stack plus many other Java libraries and frameworks
- Lightbend Platform plus many other Scala libraries and frameworks
- Confluent Platform and Alpakka Kafka.

## Open source projects

The following is an incomplete list of open source projects I'm currently working on or worked on in the past. See also [Stories](#) for some details and my [GitHub](#) account for a complete list.

### Machine learning

- Perceiver IO. A PyTorch implementation of Perceiver, Perceiver IO and Perceiver AR.
- Bayesian machine learning. Collection of notebooks about Bayesian machine learning.
- Super resolution. Tensorflow 2.x based implementation of WDSR, EDSR and SRGAN for single image super-resolution.
- Machine learning notebooks. Stanford machine learning course exercises re-written in Python and scikit-learn.

### Distributed systems and stream processing

- Eventuate. A toolkit for event sourcing and event collaboration at global scale with causal consistency. Founder and lead developer during my work at the Red Bull Media House.
- Streamz. A combinator library for integrating Functional Streams for Scala (FS2), Akka Streams and Apache Camel. A functional programming API for Apache Camel.
- Akka. A toolkit for building highly concurrent, distributed, and resilient message-driven applications for Java and Scala. Contributed Akka Persistence and Akka Camel.
- akka-persistence-cassandra. Storage plugin for Akka Persistence. Started as personal project and now officially maintained by Lightbend.
- akka-stream-eventsourcing. A project that brings to Akka Streams what Akka Persistence brings to Akka Actors: persistence via event sourcing.

### System integration

- Open eHealth Integration Platform. A platform for developing eHealth integration solutions. Founder and lead developer during my work at ICW.
- Apache Camel. An integration framework based on Enterprise Integration Patterns. Contributed several integration components as well as core functionality during my work on the Open eHealth Integration Platform.

## Certifications

- 05.2021, Full Stack Deep Learning
- Advanced Machine Learning Specialization.
  - 03.2018 Bayesian Methods for Machine Learning.
- 02.2018, Deep Learning Specialization.
  - 02.2018, Sequence Models.
  - 12.2017, Convolutional Neural Networks.
  - 12.2017, Structuring Machine Learning Projects.
  - 12.2017, Improving Deep Neural Networks: Hyperparameter tuning, Regularization and Optimization.
  - 12.2017, Neural Networks and Deep Learning.
- 11.2017, Bayesian Statistics.
- 11.2017, Inferential statistics.
- 04.2014, Machine Learning.
- 01.2014, Principles of Reactive Programming.
- 05.2012, Functional Programming Principles in Scala.

## Articles

- Blog posts
  - at [krasserm.github.io](http://krasserm.github.io) since 01.2015,
  - at [krasserm.blogspot.com](http://krasserm.blogspot.com) until 12.2014.
- Introduction to the Open eHealth Integration Platform. DZone, 05.2011.
- Akka and Camel, Camel in Action, Manning, 12.2010.
- Geschützte Objekte - Instanzbasierte Zugriffskontrolle: Architekturkonzepte und Implementierung. Java Magazin, 07.2007, 42-49.

## Talks

- 11.2015, Reactive Systems Hamburg - Event Sourcing and CQRS with Akka Persistence and Eventuate.
- 10.2015, Scala User Group Munich - Event Sourcing and CQRS with Akka Persistence and Eventuate.
- 06.2011, Scala eXchange 2011 - Connecting Akka to the Rest of the World.
- 11.2009, W-JAX, Munich - Die Open eHealth Integration Platform.
- 05.2009, ICW Developer Conference, Walldorf - The Open eHealth Integration Platform.
- 05.2008, ICW Developer Conference, Walldorf - Instance-based Security with the Security Annotation Framework.
- 11.2007, SOACon, Munich - Implementierung von Enterprise Integration Patterns auf einem JBI ESB.

## Education

- 1997, MS in Chemistry (Mag. rer. nat.) with distinction, Karl-Franzens Universität Graz.

- 1996, Various courses in computer science, Karl-Franzens Universität Graz.