

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 environments. Current focus is on deep learning for computer vision and natural language processing, mainly in context of multimodal neural search engines. I also have a special interest 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

Distributed systems, deep learning, machine learning, machine learning operations (MLOps), data-intensive applications, event sourcing and CQRS, event stream processing, event-driven architectures, message-driven architectures, system integration, resilience, scalability, high availability, functional programming, object-oriented programming, probabilistic programming.

Work experience

See also Stories for further information.

Since 04.2018: Machine learning engineer

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

- Development of a multimodal neural search engine for the MerlinOne digital asset management system.
 - Support for text-text, text-image and image-image search modes, optionally using personal information.
 - Training and fine-tuning of neural models on public and high-quality in-house datasets.
 - Low search latency and scalability to very large datasets.
- Image aesthetics assessment with criteria learned from custom in-house datasets.
- Single image super-resolution to improve quality of low-resolution images.
- Production deployments to several customer sites (hosted and on-premise).
- Responsible for scientific and engineering work in all machine learning projects.

Since 04.2011: Freelance software engineer, machine learning engineer

<https://martin-krasser.com>

- Research and development for industry machine learning and deep learning projects.
- 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 theoretical 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 and implementation of a distributed system architecture based on event sourcing with explicit causality that also includes mobile devices.
- 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.

10.2015 - 07.2017: Software engineer and architect

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

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

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 to evaluate generated protein structures.

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: Scala, Java, Python, C/C++, Haskell.
- Machine learning: scikit-learn, Tensorflow/Keras, PyTorch, JAX, PyMC3, Spark MLlib, ...
- Databases: MySQL, PostgreSQL, Oracle, Cassandra, HBase, DynamoDB, MongoDB, LMDB, LevelDB, ...
- Stream processing: Spark Streaming, Akka Streams, Kafka Streams.
- Cloud computing: Amazon Web Services and Google Cloud Platform.
- Java SE and Java EE stack plus many other Java libraries and frameworks.
- Lightbend Reactive Platform plus many other Scala libraries and frameworks.
- Confluent Platform (Apache Kafka and its ecosystem) and Reactive 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

- Bayesian machine learning. Collection of notebooks related to Bayesian machine learning. Extended when I'm writing new articles on that topic.
- Super resolution. Tensorflow 2.x based implementation of WDSR, EDSR and SRGAN for single image super-resolution.
- Image captioning transformer. Transformer-based image captioning. An extension of the pytorch/fairseq sequence modeling toolkit.
- Face recognition. Deep face recognition with Keras, Dlib and OpenCV.
- 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. Personal project that I'm still actively developing. Functional programming API on top of Apache Camel and successor of Akka 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. Personal project that brings to Akka Streams what Akka Persistence brings to Akka Actors: persistence via event sourcing. This is experimental work, inactive at the moment.
- Eventsourced. A library for building reliable, scalable and distributed event-sourced applications in Scala. Founder and lead developer during my work at Eligotech. Evolved into Akka Persistence and is not maintained any more.

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 known 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 since 01.2015,
 - at 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.