Stefano Alberto Russo - Curriculum Vitae

stefano.russo@gmail.com http://sarusso.github.io

SUMMARY

- **Technical leader** with an interdisciplinary background spanning **research**, **industry**, and **entrepreneurship**, focused on **robust data** and **machine-learning** systems operating under **real-world constraints**.
- Proven **track record** of leading small, **high-performance teams** to build **reliable**, **production-ready** systems emphasizing **CI/CD**, **testing**, and **observability**.
- Proficient in **Python**, with broad hands-on experience across **ML/DL frameworks**, **containerization** and **orchestration**, **HPC** and **distributed systems**, **web applications**, and **databases**.
- Comfortable operating in **noisy**, **ambiguous** environments, combining **technical depth** with strong **problem-framing** skills to balance r**apid prototyping** and long-term **engineering sustainability**.
- Engages **technical and non-technical audiences** through clear, structured, and empathetic communication.

EXPERIENCE

May 2020 - Dec. 2025

Senior Research Fellow / Co-Lead

Italian National Center For HPC, Big Data and Quantum Computing & INAF (Italian National Institute for Astrophysics)

• Co-lead of a **7-person innovation grant** with Intesa Sanpaolo, **Italy's largest banking group**, translating **anomaly-detection** methodologies from astrophysics to financial systems (from 2024).

• Led **applied research** on **robust, real-time time-series anomaly** detection under strict **real-world constraints**, focusing on reliability, uncertainty quantification, and minimizing false positives in **critical** an **noisy** environments.

• Designed and implemented a **container-based scientific computing platform** for resource-intensive, interactive workloads, enabling **reproducible** and **seamless execution** across **cloud** and **HPC** systems; later adopted as an **official institutional service** to **simplify access** to computing resources.

• **Mentored** junior team members; **taught Python** for AI and statistics (~100 students/year); developed a web-based **autograding platform** delivering real-time feedback to improve engagement and reinforce **clean**, **test-driven** engineering practices.

• Participated in **technology transfer** and **spin-off** evaluation, supporting early-stage **deep-tech initiatives** within the institution.

Jul. 2017 - May 2020

Co-Founder

Sharpsense AI, London (UK)

• Co-founded a startup developing **real-time monitoring solutions** for infrastructural and environmental engineering, driving **product**, **technology**, and early **fundraising**.

• Built a small, **high-performance team** to develop an **IoT monitoring platform** supporting **large-scale** field deployments, involving **4,000+ underground sensors** and **multi-million-euro** client projects.

• Provided **hands-on technical leadership** and owned the **design** and **implementation** of the platform's **architecture**, real-time data **pipelines**, and **production reliability**, operating **systems at scale**.

• Raised **pre-seed investment** and secured **strategic partnerships** supporting early commercial validation across Europe, North America, and Asia.

Jan. 2017 - Jul. 2017

Founder in Residence

Entrepreneur First, London (UK)

• Selected for EF's **highly competitive talent investment** program.

• Engaged in an intensive **startup-building process**, exploring and validating **deep-tech venture** ideas.

• Collaborated with **top-tier technical and commercial talent** to identify co-founding opportunities.

• Developed expertise in **lean startup** methodologies, **iterative ideation**, **business model** development, **fundraising**, and **market validation** through mentorship from experienced entrepreneurs and investors.

Mar. 2014 - Jan. 2017

Founding Data Engineer / Data Lead

Early-stage energy startup, Trieste (IT)

• Designed and built an **IoT platform** for **smart meter data** collection and analysis.

• Led **data infrastructure** development, including **real-time data processing** and analytics pipelines.

• Worked **cross-functionally** on **smart meter** hardware **design**, **testing**, and pre-production **planning** to ensure **seamless integration** with the software platform.

• Defined technical **requirements** in alignment with **business goals**.

Feb. 2013 - Feb 2014

Openlab Fellow

CERN, Geneva (CH)

• Part of **CERN Openlab**, a **public-private partnership** collaborating with **Intel, Google, IBM, Oracle, Siemens**, and other major tech companies.

• Explored novel methodologies for **big data analysis**, with a focus on **machine learning** for time series.

• Developed multiple **proof-of-concept** evaluating potential use cases for **advanced data analysis**.

• Organized and participated in **cross-organization workshops**, contributing to discussions on project evolution and future development.

Feb. 2011 - Feb 2013

Technical Student and Graduate Researcher

CERN, Geneva (CH)

• Designed, developed, and deployed a **real-time analytics engine** for monitoring the performance of the LHC's **1,500+ node storage cluster**, using Hadoop MapReduce.

• Conducted a feasibility study and built a proof of concept to adapt **Hadoop MapReduce** for **High Energy Physics** analyses, within the **ATLAS experiment**.

• **Published and presented** findings at CHEP 2013 (Computing in High Energy Physics).

EXTRA

• Built Pythings.io, an **IoT platform** and **microcontroller OS** enabling browser-based programming, remote updates, and orchestration of connected devices for **rapid prototyping** and **deployment**.

• Developed Timeseria, an **open-source Python library** for **time series** data, designed to handle **real-world challenges** such as missing values, irregular sampling, time zones etc., enabling **robust forecasting**, **reconstruction**, and **anomaly detection**.

• Serving as **board member** and **advisor** to an **early-stage startup** in the energy monitoring sector, helping define the **technological roadmap** and align product development with market needs.

• Worked as a **contractor** on projects at the intersection of **data**, **machine learning** and **computer science** for clients including EUMETSAT, Sky UK, eXact Lab, Generali Group, and more.

• Delivering **talks**, **lectures**, and **seminars** on best practices in modern scientific computing, software containerization, and software development in general.

• Acting, skiing, sailing, and focusing on personal growth in my free time.

EDUCATION

Dec. 2020 - Mar. 2025

PhD in Computer Science

University of Trieste, Italy.

Thesis: **"Robust anomaly detection for time series data in sensor-based critical systems"**, carried out in collaboration with an industrial partner in water infrastructure monitoring.

Oct. 2009 - Mar. 2013

MSc in Computational Physics

University of Udine, Italy.

Thesis: **"Using the Hadoop MapReduce approach for monitoring the CERN storage system and improving the ATLAS computing model"**, carried out at CERN.

Oct. 2005 - Mar. 2009

BSc in Computer Science

University of Trieste, Italy.

Thesis: **"Benchmarking of scientific applications on High Performance and Grid computing infrastructures**", carried out at SISSA. Graduated magna cum laude.

LANGUAGES

	UNDERSTANDING		SPEAKING		WRITING
	Listening	Reading	Spoken interaction	Spoken production	
English	Fluent	Fluent	Fluent	Fluent	Fluent
French	Moderate	Moderate	Basic	Basic	Basic
Italian	Native	Native	Native	Native	Native

SELECTED TECHNLOGIES

- **Languages & frameworks**: Python, Django, ML/DL frameworks.
- **Systems**: Backend services, REST APIs, data pipelines, web applications.
- **Infra & DevOps**: Docker, Singularity, Kubernetes, Slurm, AWS, ArgoCD.
- **Research**: time series modeling, anomaly detection, uncertainty quantification, conformal prediction.

SELECTED PUBLICATIONS AND PRESENTATIONS

- **"Timeseria: An object-oriented time series processing library"**, **SoftwareX**, Volume 29, 2025**.**
- **"Containerization for scientific computing and infrastructure"**, **INAF** Workshop, Turin (Italy), 2025.
- **"Computing infrastructure"**, in Data-Intensive Radio Astronomy, **Springer Nature** 2024.
- **"Rosetta: A container-centric science platform for resource-intensive, interactive data analysis"**, Astronomy and Computing, 2022.
- **"Modern software development for scientific computing"**, Seminars, University of Trieste, 2022.
- **"A microservice-oriented science platform architecture"**, Proceedings of Astronomical Data Analysis Software and Systems (ADASS), 2020.

Earlier work includes publications on high-performance and data-intensive computing (CERN, ICTP, SISSA, 2009–2015), including one of the first systematic tests of NVIDIA Tesla GPUs and CUDA with real scientific codes.