PEDRO F. SILVESTRE

(+351) 916272368 ◇ p.silvestre21@imperial.ac.uk ◇ 30/06/1997 ◇ Portuguese
pfsilvestre.me ◇ PSilvestre ◇ in/pedro-silvestre ◇ pmfsilvestre

EDUCATION

PhD in Systems for Deep Reinforcement Learning *Feb 2021 - Present*
Imperial College London, London
Large Scale Data & Systems Group — Advisors: P. Pietzuch & H. Pirk

- Researching methods to accelerate, democratize and scale Deep Reinforcement Learning (DRL) research.
- Led the design and implementation of CORAL, a compiler for DRL algorithms which reduces memory use while accelerating training through a novel technique named Gradient Accumulation Through Time.

MSc in Computer Science *Sep 2018 - Dec 2020*
NOVA School of Science and Engineering, Lisbon — Grade Average: 18/20
Department of Informatics — Advisors: A. Katsifodimos & J. Leitão
Thesis: *Clonos: Consistent High-Availability for Distributed Stream Processing through Causal Logging*

Notable Project: (Concurrency and Parallelism) Developed and evaluated a C library providing work-efficient implementations of common parallel programming patterns using Cilk. Devised a variation of the Blelloch scan which accepts any input size. *Project Grade: 19/20*

Exchange Semester *Feb 2019 - Jul 2019*
Delft University of Technology, Delft — Grade Average: 9/10
Faculty of Electrical Engineering, Mathematics & Computer Science

Notable Project: (Deep Learning) Leveraged autoencoder architectures to study image dataset complexity and compared required latent space size to information-theoretic measures (e.g. entropy). Derived useful heuristics for choosing network size. *Project Grade: 8.5/10*

BSc in Computer Science *Sep 2015 - Jul 2018*
NOVA School of Science and Engineering, Lisbon — Grade Average: 17/20
Department of Informatics

Notable Project: (Distributed Systems) Built an HDFS clone with Namenodes and Datanodes. Ring replication was used for data fault-tolerance. Service discovery via Kafka or multicast communication. A functioning Map-Reduce engine was also implemented. *Project Grade: 20/20*

RESEARCH EXPERIENCE

Research Engineer *Jun 2019 - Nov 2020*
Delft University of Technology, Delft
Web Information Systems Group

- Led the design and implementation of Clonos (`delftdata.github.io/clonos-web`), a Stream Processing System using Causal Logging for consistent local recovery and high-availability.
- Developed automated distributed benchmarking infrastructure for Stream Processors by leveraging Kubernetes, capturing real-time end-to-end throughput, latency and recovery time with millisecond precision.
- Participated in the design, development and testing of rho (ρ), a stateful FaaS platform. Built tooling for the authoring and deployment of stateful functions.

Research Assistant *Sep 2018 - Dec 2018*
NOVA School of Science and Engineering, Lisbon
NOVA-LINCS Research Laboratory

- Implemented a middleware layer providing transparent δ-CRDT based state synchronization for wireless AdHoc sensor networks in C. A reliable message fragmentation protocol was also added.

PUBLICATIONS

SIGMOD'21
(ranked Core A*)

Clonos: Consistent Causal Recovery for Highly-Available Streaming Dataflows.
Silvestre, P. F., Fragkoulis, M., Spinellis, D., & Katsifodimos, A. (2021, June).
In Proceedings of the 2021 International Conference on Management of Data (pp. 1637-1650).

PROFESSIONAL EXPERIENCE

Big Data Software Engineering Internship *Jul 2018 - Sep 2018*
XPandIT, Lisbon

- Full-stack development of a web application for orchestrating Docker containers for data-science workloads, integrating with Kerberos for single sign-on into containers. Containers were automatically built from a web form describing the tools and resources the container should have.

(Academic) Software Engineering & Quality Assurance Internship *Mar 2018 - July 2018*
Feedzai, Lisbon Grade: 19/20

- Deployed Kubernetes in the on-premises cluster. Deployed a CI solution (Jenkins) with dynamic executor provisioning on Kubernetes cluster, improving CI resource usage by up to 30%.
- Achieved elasticity by joining AWS EC2 instances dynamically to Kubernetes automatically.
- Modified internal integration testing libraries to request resources from Kubernetes cluster.
- Presented the solution to over 100 colleagues during internal talks.

HONORS & ACHIEVEMENTS

Winner of the HackDelft 2019 Hackathon (40 teams)
Project: Built an early warning anomaly detection system for the Dutch railroad network which processed time series sensor data in real time. Warnings were presented in a web application which included automated visualization of abnormal sensor data.

Awarded 1^{st} prize in CLC Merit Scholarship (€5000)
Awarded the CM Azambuja Merit Scholarship (€1000) x4

SIDE PROJECTS

Process Controller Simulator: A highly flexible simulation framework implemented in Python. Able to concurrently simulate complex chemical processes, controllers (e.g. MPC) and more. Includes a web interface for creating and visualizing simulations. Done in collaboration with a chemical engineering PhD student.

JAX Visualizer: Created a tool which visualizes any JAX computational graph in an interactive and hierarchical way. Useful for debugging and understanding complex functions.

Raspberry Pi Cluster: Assembled a 4 node cluster with compact power and ethernet delivery. Runs Kubernetes and Slurm (for OpenMPI) on top of which I deploy services such as Jenkins and personal websites.

OTHER HIGHLIGHTS

SysML@ICL: Demonstrated leadership by creating the first interest group on Systems for Machine Learning at ICL. We have hosted 6 seminar sessions with prominent authors. (URL: `sysml.doc.ic.ac.uk`)

Teaching Assistant: Practiced public speaking skills by assisting in 4 different courses on diverse topics: System Performance Engineering, Operating Systems, Compiler Construction and Reinforcement Learning.

LANGUAGES

Portuguese	Native Proficiency
English	Full Professional Proficiency (IELTS: 8.5/9, CEFR level C2)
Spanish	Limited Working Proficiency

References - available upon request