PEDRO F. SILVESTRE

Machine Learning Systems Researcher

(+44) 7599178925 ◇ pms20@ic.ac.uk ◇ 30/06/1997 ◇ Portuguese

https://www.doc.ic.ac.uk/~pms20/ ◇ PSilvestre ◇ in/pedro-silvestre ◇ pmfsilvestre

EDUCATION

PhD in Systems for Machine Learning *Feb 2021 - Present*

Imperial College London, London — Expected Graduation: Q1 2027

Large-Scale Data & Systems Group — Advisors: P. Pietzuch & H. Pirk

- Developed Tempo, a DL compiler that optimizes dynamic data dependencies in LLM programs to achieve automated KV-cache management, 7x lower latency than Torch/JAX and 32x lower peak memory.
- *Ongoing*: unified whole-program post-training parallelization; efficient fault-tolerance for MoE models.
- *Supervised 4 MSc students*: LLM parallelization on heterogeneous clusters; scalable off-policy reinforcement learning; code-generation for Tempo; and differentiable batch simulation rendering.

Integrated MSc in Computer Science *Sep 2015 - 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: (Distributed Systems) Built an HDFS and Map-Reduce clone complete with ring replication fault-tolerance and service discovery using Kafka. (*Project Grade: 20/20*)

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) Used autoencoders and information theory to study relationships between dataset complexity, over-fitting and model capacity using PyTorch. (*Project Grade: 9/10*)

RESEARCH EXPERIENCE

Quantitative Researcher (PhD Intern) *Jul 2024 - Sep 2024*

Citadel Securities, London — Alpha Research Team

- Built deep representation learning–based equity alpha signals, improving alpha–return correlation by 0.04.
- Contributed efficient data-loading and normalization to internal distributed training framework.

Research Engineer in RL Frameworks (PhD Intern) *May 2023 - Nov 2023*

InstaDeep Ltd., London — ML Frameworks Team

- Extended the distributed training strategies of the internal reinforcement learning framework.
- Developed a self-configuring strategy, saving 100s of engineer-hours optimizing systems parameters.
- Reduced costs by 2x while increasing throughput 8x with an efficient C++ batched inference server.

Research Engineer *Jun 2019 - Nov 2020*

Delft University of Technology, Delft — Web Information Systems Group

- Designed Clonos, a highly-available stream processor (SP) using causal logging for fault-tolerance. Clonos achieves 24x faster recovery, 35x lower latency for unaffected partitions, with low overhead.
- Participated in the development and testing of rho, a stateful serverless platform using SPs as executors.

Research Assistant *Sep 2018 - Dec 2018*

NOVA School of Science and Engineering, Lisbon — NOVA-LINCS Research Laboratory

- Implemented a δ-CRDT state synchronization protocol for wireless ad-hoc sensor networks in C++.

INDUSTRY EXPERIENCE

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

- Full-stack web app design for Docker container orchestration, image creation, deployment and log-in.

Software Engineering & Quality Assurance Internship *Mar 2018 - July 2018*
Feedzai, Lisbon

- Set-up on-premises Kubernetes, containerized Jenkins with dynamic executors, lowering CI costs by ∼30%.

PUBLICATIONS

OSDI'26
(under submission)

Facade: Balancing Homogeneous Parallelization Plans for Efficient Training on Heterogeneous GPU Clusters.
Yanda Tao, **Silvestre, Pedro F.**, Marcel Wagenlander, Peter Pietzuch.

SOSP'25
(ranked Core A*)

Tempo: Compiled Dynamic Deep Learning with Symbolic Dependence Graphs.
Silvestre, Pedro F., Peter Pietzuch (2025, October).
Proceedings of the ACM SIGOPS 31st Symposium on Operating Systems Principles.

EuroMLSys'25

Systems Opportunities for LLM Fine-Tuning using Reinforcement Learning.
Silvestre, Pedro F., Peter Pietzuch. (2025, March).
Proceedings of the 5th Workshop on Machine Learning and Systems.

SIGMOD'21
(ranked Core A*)

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

HONORS & ACHIEVEMENTS

- 3^{rd} place in the 2023 ICC Poster Competition **(∼25 posters)**
- 1^{st} place in the 2022 GResearch Quant Challenge Hackathon **(∼60 teams)**
- 1^{st} place in the HackDelft 2019 Hackathon **(∼40 teams)**
- Awarded 1^{st} prize in the 2018 CLC Merit Scholarship **(€5000)**
- Awarded the 2016-2020 CM Azambuja Merit Scholarship **(€1000) x4**

HIGHLIGHTS

SysML@ICL: Demonstrated leadership by creating the first interest group on Systems for Machine Learning at ICL. We have hosted 8 seminar sessions with prominent authors. (URL: `sysml.doc.ic.ac.uk`)

Teaching Assistant: Practiced communication skills assisting in 5 courses: Compiler Construction, Data Processing Systems, System Performance Engineering, Operating Systems and Reinforcement Learning.

Raspberry Pi Cluster: Assembled a 4 node cluster with compact power, cooling and ethernet delivery.

TOOLS & TECHNOLOGIES

DL Systems	PyTorch, JAX, XLA, Triton, Megatron, Deepspeed, W&B,
Programming Languages	Python, C++, C, Java, Bash
Data Analysis	Numpy, Pandas, Polars, Xarray, PyArrow/Parquet

LANGUAGES

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

UK IMMIGRATION STATUS

Pre-settled Status Valid until 6 December 2030

References - available upon request