Vasilis Siomos

siomvas@protonmail.com
vsiomos.me
London, UK

Summary

I specialize in Federated Learning and Privacy-Preserving ML for medical imaging, with first-author publications at **MICCAI** (Spotlight, top 11%), **TMLR**, and **ISBI** achieving up to **4% Dice improvement** over prior SOTA. I build distributed training systems, collaborate with clinicians on problem formulation, mentor junior researchers, and review for **ICLR**, **NeurIPS**, and **MICCAI**.

Education

PhD in Computer Science **2022 – Expected Mar 2026**
City St George's, University of London *London, UK*

- Thesis: Multi-task Federated Learning for Medical Imaging
- Advisor: Dr. Giacomo Tarroni (joint supervision with Imperial College London)

MSc in Computing (AI & Machine Learning) **2020 – 2021**
Imperial College London — *Distinction* *London, UK*

- Bodossaki Foundation Scholar (full merit-based funding, <5% acceptance rate)
- Thesis: Towards Federated Reinforcement Learning

Engineer's Degree in Electrical & Computer Engineering **2013 – 2019**
Aristotle University of Thessaloniki — GPA: 9.0/10 (Top 5%) *Thessaloniki, Greece*

- Thesis: Learning Under Label Noise via Ensemble Head Fine-Tuning
- Proposed RANSAC-inspired method matching SOTA accuracy with significantly less compute on CIFAR-10 and Fashion-MNIST

Experience

Doctoral Researcher **2022 – Present**
City St George's, University of London *London, UK*

Federated Dataset Simulation (in progress)

- Leading research on VLM-guided partitioning with an undergraduate research engineer; creates non-IID splits with 2× higher heterogeneity than Dirichlet sampling while preserving semantic coherence (publication in progress)

FedCLAM — MICCAI 2025 (Spotlight, top 11%)

- Proposed client-adaptive momentum aggregation with foreground intensity matching; achieved **+3.8% Dice** on cardiac MRI and **+2.1% Dice** on abdominal CT vs. prior SOTA
- Open-sourced: github.com/siomvas/FedCLAM

ANFR — TMLR 2025

- Designed first FL-native architecture combining scaled weight standardization with channel attention; reduces non-IID accuracy degradation by **40%** across 4 datasets and 3 aggregation methods
- Maintains favorable privacy–utility trade-offs under differential privacy (ϵ=1)
- Open-sourced: github.com/siomvas/ANFR

ARIA — IEEE ISBI 2024

- First systematic benchmark of architecture–aggregation interactions in FL across 5 medical imaging datasets (12 sites, 50K+ images); identified batch normalization failure modes under distribution shift
- Open-sourced with NVIDIA FLARE: github.com/siomvas/ARIA

Teaching & Infrastructure

- Graduate TA for Computer Vision, Neural Computing, and Python Programming (200+ students/year); designed exercises adopted in subsequent course iterations
- Onboarded 4 PhD students to departmental HPC: environment setup, SLURM orchestration, distributed training debugging
- Mentored 1 MSc student; currently supervising undergraduate RE on publication-track research

Applied Research Engineer **2022 – 2023**
Equideum Health *London, UK*

- Benchmarked FL contribution methods (Shapley value, Least Core) on 3 healthcare datasets; validated MSc research (NeurIPS FL Workshop 2021), achieving 15% improvement in fair reward allocation

MSc Thesis Researcher **2021**
Imperial College London *London, UK*

- Built PyTorch framework bridging Federated Learning with Reinforcement Learning; achieved **8% improvement** in early sepsis detection on MIMIC-III (40K patient episodes) vs. centralized baselines

Research Intern **Winter 2019**
Centre for Research and Technology Hellas (CERTH) *Thessaloniki, Greece*

- Delivered TensorFlow modules for Horizon 2020 nutritional-coaching robot; built 5K-image food dataset programmatically using OpenCV

Selected Publications

- **Siomos, V.**, Passerat-Palmbach, J., Tarroni, G. FedCLAM: Client Adaptive Momentum with Foreground Intensity Matching for Federated Medical Image Segmentation. *MICCAI 2025* **(Spotlight)**.
- **Siomos, V.**, Passerat-Palmbach, J., Tarroni, G. Addressing Data Heterogeneity in Federated Learning with Adaptive Normalization-Free Feature Recalibration. *TMLR 2025*.
- Zenk, M et al. (incl. **Siomos, V.**). Towards fair decentralized benchmarking of healthcare AI with the Federated Tumor Segmentation challenge. *Nature Communications 2025*.
- **Siomos, V.**, Naval-Marimont, S., Passerat-Palmbach, J., Tarroni, G. ARIA: On the Interaction Between Architectures, Initialization and Aggregation Methods for Federated Visual Classification. *IEEE ISBI 2024*.
- Naval-Marimont, S., **Siomos, V.**, et al. Ensembled Cold-Diffusion Restorations for Unsupervised Anomaly Detection. *MICCAI 2024*.
- Naval-Marimont, S., **Siomos, V.**, Tarroni, G. MIM-OOD: Generative Masked Image Modelling for Out-of-Distribution Detection. *MICCAI 2023*.
- **Siomos, V.**, Passerat-Palmbach, J. Contribution Evaluation in Federated Learning: Examining Current Approaches. *NeurIPS FL Workshop 2021*.

Full list: scholar.google.com/citations?user=zUdoBe0AAAAJ

Service & Presentations

- **Peer Review:** ICLR, NeurIPS, MICCAI (2023–2025)
- **Workshop Presentations:** MICCAI DeCaF Workshop (2025), EurIPS Medical Imaging Workshop (2025)
- **Invited Seminars:** UCL Centre for Medical Image Computing, University of Bamberg (scheduled)

Technical Skills

ML & Distributed Systems: PyTorch, TensorFlow, NVIDIA FLARE, DeepSpeed, Differential Privacy
Infrastructure: Docker, SLURM, Git, Linux, Multi-GPU/Multi-Node Training
Languages: Python (expert), C++, Java, MATLAB
Spoken: English (fluent), Spanish (fluent), Greek (native), Italian (basic)

Honors

Bodossaki Foundation Scholar, Imperial College London · Selected for Advanced Data Science Team, Imperial College London · National Mathematics Competition Finalist · National Debate Championship Finals