JOSHUA **SYLVESTER**

Room 239, Kennedy Building, University of Kent, Canterbury, Kent, CT2 7FS
Email: jrs71@kent.ac.uk
Website: jsylvester.com GitHub: github.com/flaxa

I am a 3rd year PhD researcher at the University of Kent specialising in reinforcement learning for autonomous cyber defence, with a focus on building agents that detect and respond to adversarial drift through novelty identification in simulated network environments. I have five publications across IEEE CSR 2024, ESORICS 2025, EDCC 2026, and arXiv covering RL algorithms, LLMs, knowledge retention, and adversarial drift. With two of them winning the best paper award. I recently completed a research placement working along side the Alan Turing Institute on LLM-based agents evaluated against the CAGE Challenge 2 benchmark, with prior industry experience as a Security Researcher at Orange Cyberdefense applying NLP to threat intelligence. I am an active participant in reverse engineering CTF competitions. My motivation is to increase the robustness and resilience of advanced AI systems, particularly at the intersection of machine learning and security.

EXPERIENCE

AUGUST 2025 – MARCH 2026

VISITING RESEARCHER, ALAN TURING INSTITUTE

Investigated the use of large language models as reinforcement learning agents on the CAGE Challenge 2 cyber defence benchmark. Conducted user research with SOC analysts to evaluate agent behaviour and incorporate practitioner feedback into the research, and presented findings to technical stakeholders.

JANUARY 2023 – PRESENT

GRADUATE TEACHING ASSISTANT, UNIVERSITY OF KENT

Taught across 14 modules spanning undergraduate and postgraduate levels in computer science, including Deep Learning, Computer Security, AI Systems Implementation, Blockchain Systems, and Introduction to Digital Forensics. Invited to deliver guest lectures on three occasions in *Advanced Topics in Cyber Security* (PG) and *Artificial Intelligence and Cyber Security* (PG). Responsibilities include leading classes, designing and delivering content, grading assignments, providing feedback, and offering academic guidance.

JUNE 2022 – DECEMBER 2022

SECURITY RESEARCH INTERN, ORANGE CYBERDEFENSE

Applied machine learning to a variety of problems such as classification of text and Named-Entity Recognition (NER). Created custom datasets using web scraping and automated labelling, normalised and enriched ransomware data, created custom Docker containers for production deployment within our team, produced a monthly report as a team which involved authoring articles, and produced a final year report where we evaluated the security landscape and gave insight into what we had seen develop over the past year and what we expected in the future.

EDUCATION

JANUARY 2023 - PRESENT

PHD, UNIVERSITY OF KENT

Research focuses on reinforcement learning for autonomous cyber defence,

including reward shaping, knowledge retention, and multi-attack environments. Contributed to the Dstl-funded ARCD (Autonomous Resilient Cyber Defence) project on "Reward shaping for network defence based on reinforcement learning."

SEPTEMBER 2019 – JUNE 2022

BACHELOR'S DEGREE (HONS) IN COMPUTER SCIENCE, UNIVERSITY OF KENT

First Class Degree

Certificates awarded for:

- Outstanding overall performance in an undergraduate programme in the School of Computing
- The best mark in the final year computing project in the School of Computing

SEPTEMBER 2017 – JUNE 2019

A-LEVELS, EAST NORFOLK SIXTH FORM

Computer Science (B), Chemistry (A), Mathematics (B)

PUBLICATIONS

Hicks, C., Bates, E., McFadden, S., Thompson, I.S., Foley, M., Chapman, E., Dice, N.E., Samaddar, A., **Sylvester, J.**, Neema, H. and Butts, N., 2026. Building Better Environments for Autonomous Cyber Defence. Accepted at AI4CNI 2026. *arXiv:2604.08805*.

ElZemity, A., **Sylvester, J.**, Arief, B. and De Lemos, R., 2026. Agentic Knowledge Distillation: Autonomous Training of Small Language Models for SMS Threat Detection. To appear in EDCC 2026. arXiv:2602.10869.

Sylvester, J. and de Lemos, R., 2025, September. Automated Cyber Defence with Reinforcement Learning in Multi-attack Environments. In European Symposium on Research in Computer Security (pp. 57-74). Cham: Springer Nature Switzerland.

Sylvester, J. and de Lemos, R., 2025, September. Knowledge Retention for Generic Reinforcement Learning Policies in Autonomous Cyber Defence. In European Symposium on Research in Computer Security (pp. 39-56). Cham: Springer Nature Switzerland.

Sylvester, J. and de Lemos, R., 2024, September. Identifying novelty in network traffic. In 2024 IEEE International Conference on Cyber Security and Resilience (CSR) (pp. 506-511). IEEE.

Full list available on Google Scholar.

TECHNICAL SKILLS

- **Languages:** Python, C, SQL
- **ML Frameworks:** PyTorch, TensorFlow, Hugging Face Transformers
- **Reinforcement Learning:** Stable-Baselines3, Gymnasium, CAGE Challenge 2
- **LLMs:** LoRA, vLLM, LangChain, RAG
- **Infrastructure:** Docker, Linux, Git
- **Cloud / HPC:** AWS, SLURM
- **Research Areas:** Autonomous Cyber Defence, Novelty Detection
- **Reverse Engineering:** Ghidra, gdb, Wireshark