

Master Thesis

Start date: As soon as possible

Topic: *Calibrated Decision Models for Onboard Spacecraft Autonomy*

An autonomous spacecraft has to decide, minute by minute, whether to charge, observe, process data, downlink or stay safe. Generative LLM agents can take such decisions from a textual state description [1], but each answer is slow, costly and carries no trustworthy statement of its own reliability. *System One* decision models instead answer typed questions in a single forward pass and return a calibrated probability for every option [2], [3]. Trustworthy confidence could decide when the spacecraft acts on its own and when it defers to a safer controller or to the ground [4]. Building on our work on autonomous satellite decision-making [5], the decision-model branch of our AUTOPS framework already connects the open-weight Kev model family (0.8B to 27B parameters) to an Orekit-based simulator of the EventSat 6U CubeSat [6], [7]. First zero-shot closed-loop runs fly complete simulated weeks and expose the open problems this thesis attacks: the energy cost of the decision cadence, a possible sensitivity to option order, confidence that is not yet calibrated for the mission, and a clear gap to a rule-based scheduler. The student works with the AUTOPS codebase, a data-centre GPU, an NVIDIA Jetson Orin platform and close supervision towards a joint publication.

Goals

- Turn a zero-shot System One model into a mission-specific decision model for EventSat by learning from privileged planners and from mission outcomes
- Make the model's confidence trustworthy: mission-specific calibration and robustness to option order
- Design a confidence-gated hybrid controller that acts autonomously above a threshold and defers to a fallback otherwise, and evaluate it in closed loop
- Quantify the onboard cost of System One decisions (latency, memory, energy) on flight-like embedded hardware and its trade-off with decision cadence

Task

Modeling and Learning:

- Reproduce the zero-shot System One baseline in AUTOPS and redesign the decision interface (question, option descriptions, decision cadence) under the same information budget as the LLM baselines
- Diagnose and mitigate option-order sensitivity of the hybrid recurrent/attention backbone (permutation tests, permutation averaging, permutation-consistency training)
- Fine-tune Kev models with LoRA by imitating a privileged analytical planner on the states the model itself visits (Dagger [8]); compare with counterfactual mission-outcome labels and analyse which privileged knowledge cannot be imitated from onboard information
- Calibrate the fine-tuned models on held-out scenarios (temperature scaling [2])

Simulation:

- Run week-long closed-loop episodes of the trained models on paired, held-out scenarios against rule-based, generative-LLM and latent world-model planners
- Implement a confidence-gated controller (act when the top probability exceeds a threshold τ , otherwise defer to a fallback) and sweep τ
- Measure latency, memory and power of small Kev models on an NVIDIA Jetson Orin, the payload-computer class intended for EventSat, and feed them into the simulator's energy model

Analysis:

- Evaluate calibration (reliability diagrams, ECE, Brier score) and coverage–risk curves, offline and in closed loop [4]
- Evaluate mission outcome with the AUTOPS metrics and seed-clustered bootstrap confidence intervals, and explain behaviour from replayed trajectories
- Map the trade-off between decision quality, model size, cadence and onboard energy, and position System One models against generative LLM agents

Documentation and Presentation:

- Contribute tested, documented code to the AUTOPS repository, with run manifests that make every reported number reproducible
- Write the master thesis and a manuscript in double-column conference format

Requirements

- Excellent academic record and the ambition to do publishable research
- Solid machine learning background; hands-on experience with PyTorch and transformer language models (e.g. LoRA fine-tuning) is a strong plus
- Very good Python and software engineering skills (Git, testing, Linux, remote GPU servers)
- Good command of probability and statistics (calibration, uncertainty, bootstrap)
- Basic knowledge of spacecraft operations (power, attitude, communication); experience with imitation learning, reinforcement learning or planning is a plus

Expected results:

- A mission-specific, calibrated System One decision model for EventSat
- A confidence-gated hybrid controller with a closed-loop evaluation against rule-based, LLM and world-model baselines
- Onboard cost measurements of small decision models on embedded hardware
- Tested, documented code and reproducible run manifests in the AUTOPS repository
- A master thesis and a manuscript intended for submission to a top-tier venue (e.g. AAAI, AAMAS or IAC)

- A short presentation of the results
- Following our internal evaluation, the candidate may be funded, on a merit basis, to present the results of the work at a world-class international conference.

References

- [1] V. Rodriguez-Fernandez *et al.*, "Language Models are Spacecraft Operators," arXiv:2404.00413, 2024.
- [2] C. Guo *et al.*, "On Calibration of Modern Neural Networks," in *Proc. 34th International Conference on Machine Learning (ICML)*, 2017.
- [3] Kev contributors, "Kev: Small Jev-like Decision Models," open-source software (Apache-2.0), 2026. [Online]. Available: <https://github.com/jaredpalmer/kev>
- [4] Y. Geifman and R. El-Yaniv, "Selective Classification for Deep Neural Networks," in *Advances in Neural Information Processing Systems 30 (NeurIPS)*, 2017.
- [5] C. J. Juan Oliver *et al.*, "Autonomous Decision-Making for Federated Satellite Systems: A Decentralized Multi-Agent Deep Reinforcement Learning Approach in Partially Observable Dynamic Environments," in *Proc. European Conference for Aeronautics and Space Sciences (EUCASS)*, 2025, doi: 10.13009/EUCASS2025-248.
- [6] L. Schuberth *et al.*, "Leveraging Event-Based Cameras for Enhanced Space Situational Awareness: A Nanosatellite Mission Architecture Study," *75th International Astronautical Congress (IAC)*, 2024. [Online]. Available: <https://mediatum.ub.tum.de/node?id=1759594>
- [7] S. Dolan *et al.*, "Design and Analysis of an Event Camera Payload for Space-Based Object Detection on the EventSat 6U CubeSat Mission," *15th IAA Symposium on Small Satellites for Earth System Observation*, 2025.
- [8] S. Ross, G. J. Gordon, and J. A. Bagnell, "A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning," in *Proc. 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.

Supervisor

Clemente J. Juan Oliver, M.Sc.

Phone: +49 89 289 – 55752

E-mail: clemente.juan@tum.de