Sukai Huang

Postdoctoral Research Fellow | Monash University | Melbourne, Australia | huangsukai1997@gmail.com

Professional Summary

Postdoctoral researcher at Monash University working on LLM/VLA-driven embodied agents and sequential decision-making. Publications at AAAI, ICAPS, AAMAS, ACL, EMNLP, ICLR, and NeurIPS on LLM planning, neuro-symbolic AI, and VLM/RL reward modeling; open-source VLA contributor (89.8% end-to-end success on LIBERO-Goal). Goal: verifiable, robust autonomous systems that couple neural perception and learning with symbolic planning.

Research Experience

Postdoctoral Research Fellow, **Monash University, Melbourne, Australia** **Jun 2025 – Present**

Supervised by A/Prof Hamid Rezatofighi (DARPA HARNESS PI; CVPR/NeurIPS/IROS Area Chair)

DARPA HARNESS — Neuro-Symbolic Planning for Autonomous Robotics

- Developed neuro-symbolic planning and instruction-following modules for autonomous robotic systems, integrating LLM-based mission interpretation with symbolic plan generation and reinforcement learning control.
- Integrated and deployed a VLA-based perception–planning system on physical UAV hardware, and conducted experiments as part of HARNESS engineering acceptance testing.
- Investigated the impact of instruction granularity on embodied agents' performance.
- Evaluated purely neural agents as part of the red team, assessing capabilities and limitations.

ONR SEA-AI — Multi-Agent Maritime Autonomy and Monitoring

- Designed and built a hierarchical multi-agent system coordinating LLM mission planning, maneuver control, and shared context for report validation, collision warning, anomaly inspection, and reconnaissance.
- Built a symbolic world model abstracting 3D scenes into 2.5D representations, with Bayesian belief updates, D* Lite navigation, visibility-based search, and evidence-triggered replanning under partial observability.
- Decision-core simulation benchmark (10 synthetic scenarios, ideal perception, no LLM calls): uncovered 189/210 corrupted vessel reports versus 93/210 for report-order patrol (+103%).
- Integrated MiniZinc/PDDL solvers and plan validation to produce executable statecharts, with priority-aware multi-mission scheduling, deadline and feasibility checks, and pre-execution operator approval.
- Developed an operator monitoring console with live mission status, evidence-grounded text/voice Q&A, human intervention controls, and auditable mission histories.

Student Supervision

- Co-supervised two honours/master students at Monash: PACE-Drone (preference-aware continual exploration with POMDP planning and Gaussian Process preference learning for drone missions) and WALR (width-aware language reward design for vision-language-action models that addresses the language ignoring problem in embodied AI).

PhD in Engineering and IT, **The University of Melbourne, Melbourne, Australia** **Jul 2021 – Oct 2025**

Thesis: Integrating Natural Language for Sequential Decision Problems; Supervised by A/Prof Nir Lipovetzky (ICAPS Executive Committee, Mentoring Chair) and Prof Trevor Cohn (Google Research)

- Investigated how natural language can be reliably integrated into sequential decision-making (SDM) systems across two paradigms: model-free reinforcement learning and model-based automated planning.
- RL line: exposed critical limitations of VLM-based reward models (susceptibility to false positive rewards) and proposed BiMI, a reward function that suppresses false positives with minimal trade-off on false negatives.
- Planning line: developed LLM-symbolic planning pipelines ("Planning in the Dark") that generate executable plans without expert-generated examples, and analysed robustness of end-to-end LLM plan generation strategies (ICAPS 2025 Oral).

Machine Learning Engineer (Part-time, Remote), **FINVISE London** **May 2022 – Nov 2023**

- Built an Automated Valuation Model (AVM) achieving $> 80\%$ accuracy; designed a multimodal data pipeline handling 20M+ property records.

Open-Source Contributions

StarVLA — Open-Source VLA Framework (Code Contributor, github.com/starVLA)

- Responded to Issue #59 in the Qwen VLA open-source repository, where the Adapter architecture's reproduction underperformed expectations: identified and refactored a subtle implementation deviation in the action-query embedding construction, aligning it with the official Qwen-OFT reference logic and fixing the training convergence anomaly.
- With a Qwen3-VL-2B backbone (language side frozen), completed 50k training steps on LIBERO-All and achieved an 89.8% end-to-end success rate on the 500-task LIBERO-Goal benchmark, reproducing the official baseline convergence trend; contributed fully reproducible training launch scripts.

Education

Doctor of Philosophy (PhD) – Engineering and IT **Jul 2021 – Oct 2025**
The University of Melbourne | Melbourne Research Scholarship
Bachelor of Advanced Computing (Honours, First Class) **Jul 2017 – Jul 2021**
Australian National University (ANU) | GPA: 6.84/7.00; Specialization: Machine Learning Systems;
ANU Dean's Award 2019

Publications

1. Sukai Huang, Chenyuan Zhang, Fucai Ke, Zhixi Cai, Gholamreza Haffari, Lizhen Qu, and Hamid Rezatofighi. "Mini-BEHAVIOR-Gran: Revealing U-Shaped Effects of Instruction Granularity on Language-Guided Embodied Agents." International Conference on Empirical Methods in Natural Language Processing (EMNLP), 2026. [CORE-A*] — First author.
2. Sukai Huang, Chenyuan Zhang, Hamid Rezatofighi, Mor Vered, and Buser Say. "Neurosymbolic Active Goal Recognition in Partially Observable Environments." International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2026. [CORE-A*] — First author.
3. Chunhua Liu, Kabir Manandhar Shrestha, and Sukai Huang. "ALIGN: Word Association Learning for Cultural Alignment in Large Language Models." Annual Meeting of the Association for Computational Linguistics (ACL), 2026. [CORE-A*] — Joint First author.
4. Sukai Huang, Nir Lipovetzky, and Trevor Cohn. "Chasing Progress, Not Perfection: Revisiting Strategies for End-to-End LLM Plan Generation." International Conference on Automated Planning and Scheduling (ICAPS), 2025. Oral Presentation (top 5%). [CORE-A*] — First author.
5. Sukai Huang, Nir Lipovetzky, and Trevor Cohn. "Planning in the Dark: LLM-Symbolic Planning Pipeline without Experts." AAAI Conference on Artificial Intelligence (AAAI), 2025. [CORE-A*] — First author.
6. Sukai Huang, Shu-Wei Liu, Nir Lipovetzky, and Trevor Cohn. "The Dark Side of Rich Rewards: Understanding and Mitigating Noise in VLM Rewards." ICAPS 2025 Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL). — First author.
7. Zhixi Cai, Fucai Ke, Sukai Huang, Maria Garcia de la Banda, Peter J. Stuckey, Gholamreza Haffari, and Hamid Rezatofighi. "JRDB-AVR: An Active Visual Reasoning Benchmark for Real-World Embodied Environments." Conference on Neural Information Processing Systems (NeurIPS), 2026. [CORE-A*] — Co-author.
8. Zhixi Cai, Fucai Ke, Kevin Leo, Sukai Huang, Maria Garcia de la Banda, Peter J. Stuckey, and Hamid Rezatofighi. "MATA: A Trainable Hierarchical Automaton System for Multi-Agent Visual Reasoning." International Conference on Learning Representations (ICLR), 2026. [CORE-A*] — Co-author.
9. Simon De Deyne, Sukai Huang, Lea Frermann, and Chunhua Liu. "Fine-tuned Large Language Models Predict Human Word Associations and Improve Performance across Lexical and Semantic Processing Tasks." Annual Meeting of the Cognitive Science Society (CogSci), 2026. — Co-author.

Awards & Scholarships

- **AAAI 2025 Scholarship**, Association for the Advancement of Artificial Intelligence
- **ICAPS 2025 Scholarship**, International Conference on Automated Planning and Scheduling
- **Invited Speaker**, Kingston AI Symposium 2026, Adelaide, Australia
- **IEEE COG 2020 Champion**, Angry Birds Level Generation Competition, Aesthetic Track
- **ANU Dean's Award**, 2019; Chancellor's Letters of Commendation, 2018 & 2020

Academic Service

- Reviewer: ICAPS, ACL Rolling Review (ARR), Artificial Intelligence journal.

Technical Skills

- **Programming & Tools**: PyTorch, LangChain, Mattcopock Skills, Codex, SQL, C++, Haskell; Git, Docker, W&B
- **Domain Expertise**: Vision-Language-Action (VLA) models, embodied agents, LLM reasoning & planning, LLM agent systems, reinforcement learning (RL), neuro-symbolic AI, automated planning (PDDL), sequential decision-making (POMDP), multimodal LLMs / VLMs, LLM evaluation & benchmark design.
- **Languages**: Chinese (native), English (professional)