Research Statement

Cracking the Curse: Scalable Learning for High-Dimensional Multi-Agent Systems

My research develops the theoretical foundations and practical algorithms that enable large-scale autonomous systems to coordinate intelligently in real-world environments. Its potential applications range from warehouse robot fleets numbering in the thousands to complex articulated systems with dozens of degrees of freedom. The central challenge unifying this work is the “curse of dimensionality”: as the number of agents or joints increases, the joint state-action space grows exponentially, rendering centralized planning computationally intractable. My approach breaks this complexity through hierarchical, decentralized, and distributed control frameworks mainly powered by deep reinforcement learning (RL) and imitation learning (IL), enabling individual agents to learn local decision-making policies that collectively yield emergent, large-scale coordination.

Since joining NUS as an Assistant Professor in July 2019, I have established a distinctive and internationally recognized research program, spanning from theoretical to experimental research with a consistent focus **on theoretic grounding and robotic validation**. This program is entirely independent of my doctoral and postdoctoral advisors (Prof. Max-Olivier Hongler, EPFL; Prof. Howie Choset, CMU). Since my work lies at the **interface of robotics and artificial intelligence**, my scientific output has balanced conference and journal publications, as is common in these rapidly evolving fields. My scholarly output during this period includes 53 peer-reviewed conference papers, of which 41 as sole/first/last/corresponding (SFLC) author, and 20 journal articles, of which 14 as SFLC author, including 10 in Scopus Top-10% venues. This resulted in a cumulative citation count exceeding 2,600 (h-index: 24) on Google Scholar.

These numbers are further contextualized by roboranking.org, a newly established, community-adopted benchmark for robotics research productivity inspired by the widely used csrankings.org for computer science, which ranks me **3rd among all NUS robotics faculty** for the 2019–2026 period. This standing places me immediately behind two senior, internationally established full professors (David Hsu (School of Computing) and Marcelo Ang (Mechanical Engineering)) when the comparison is appropriately restricted to Roboranking's core robotics venues (i.e., excluding computer vision conferences such as CVPR, ICCV, and ECCV, which fall outside my research scope and carry substantially higher citation volumes that would otherwise skew cross-faculty comparisons). At the institutional level, my output contributes to placing NUS 16th worldwide in the Roboranking standings (as of 16 March 2026), a strong result given that the score naturally favors larger departments with more faculty. Taken together, these indicators confirm that my scholarly contributions are internationally competitive in absolute terms, while remaining remarkable given the early stage of my independent research career.

Consistently, my work has been recognized both institutionally and internationally through multiple awards. At the institutional level, I received the **NUS College of Design and Engineering (CDE) Outstanding Early Career Award (2023)**, a faculty-wide distinction recognizing exceptional contributions across research, teaching, and service, for which I was one of only three recipients and the sole awardee from Mechanical Engineering. Internationally, an **Amazon Research Award (2022)**, a highly competitive grant recognizing exceptional early-career contributions (~$80k USD + $20k AWS credits, awarded globally to top researchers across academia), specifically recognized my research direction on distributed learning for human-aware multi-agent pathfinding. Most visibly, my work received the **Best Student Paper Award** and **Best Paper Award on Multi-Robot Systems at IEEE ICRA 2025**, the flagship robotics conference, where only four Best Student Papers were selected from 1,606 accepted papers (top 0.25%), as well as **Best Paper Awards at DARS 2021 and DARS 2022**.

My research leadership is further evidenced by sustained editorial responsibilities at the field's top venues: I serve as Associate Editor for the *International Journal of Robotics Research* (IJRR, Scopus Top-1%, IF 9.3), one the most prestigious journal in robotics, and *IEEE Robotics and Automation Letters*

(RA-L, Scopus Top-10%, IF 5.3). I have been appointed Area Chair for *Robotics: Science and Systems (RSS) 2026*, arguably the most selective robotics conference, and have served as Associate Editor for ICRA in multiple areas since 2022. My expertise is regularly sought through keynote invitations at leading venues, including CVPR 2025, AAMAS 2024, IROS 2024, and AAAI 2023, as well as invited seminars at top institutions including the University of Pennsylvania's GRASP Lab, Stanford's Multi-Robot Systems Lab, CMU's Robotics Institute, Caltech, MIT, Cambridge, EPFL, and ETH Zürich.

While I conceive and lead my research program independently, I have also cultivated strategic international collaborations that amplify its reach and impact. These partnerships were initiated by me, involved students I supervised visiting partner labs on competitive NUS fellowships (ORIA awards), and resulted in work where I served as co-senior or senior author. Recent partnerships with Prof. Auke Ijspeert (EPFL) on bio-inspired locomotion, Prof. Mac Schwager (Stanford) on collaborative manipulation, and Prof. Jiaoyang Li (CMU) on scalable pathfinding have produced publications at RSS, CoRL (oral talk, ~5% acceptance), and ICRA (dual Best Paper Awards). These successes demonstrate the scientific community's appreciation of my intellectual leadership and contributions to the field.

My research has attracted over **S$8 million (USD ~6.3M)** in funding as PI (excluding Co-PI and Co-I roles), spanning grants from Singapore's Ministry of Education (MOE Tier 1), the National Research Foundation (NRF), Temasek Laboratories, Singapore's Ministry of Defense, and Singapore's Defense Science Organization (DSO), as well as industry partners including ST Engineering, Cisco Systems, Amazon, and the Singapore Maritime Institute. This funding has enabled me to build the **Multi-Agent Robotic Motion (MARMot) Lab** into a leading research group, graduating 7 PhD and 2 MEng as well as exceptional undergraduate students, now placed at MIT (PhD), Stanford (postdoc), TikTok, Alibaba, Genesis, and other top career destinations. Within MARMot Lab, I currently supervise a team of 15+ graduate students and postdoctoral researchers.

The following sections detail my four core research thrusts: multi-agent pathfinding, informative path planning, articulated locomotion and manipulation, and intelligent transportation. Each represents a coherent body of work with demonstrated international impact. I conclude with my vision of future research directions advancing toward generalizable robot intelligence through foundation models, test-time adaptation, and collaborative embodied AI.

1. Scalable Multi-Agent Pathfinding (MAPF)

Multi-agent pathfinding, i.e., the coordination of hundreds or thousands of agents to reach their destinations without collisions, is a cornerstone problem in robotics and AI, with direct applications in warehouse automation, autonomous vehicle coordination, and logistics optimization. The problem is NP-hard in its optimal form, and classical bounded-optimal search-based planners have hit fundamental scalability limits: state-of-the-art methods like Priority-Based Search and Conflict-Based Search struggle beyond 256–512 agents, far short of the thousands envisioned by leading logistics companies. My research has pioneered a paradigm shift toward **decentralized, learning-based MAPF**, demonstrating that neural policies can not only match, but more and more exceed highly optimized classical solvers while scaling to unprecedented team sizes.

Foundational Work: PRIMAL and PRIMAL2

My seminal work, **PRIMAL** (RA-L, 2019; 614 citations as of 11 March 2026), was the first to demonstrate that combining reinforcement learning with imitation learning could enable real-time coordination of over 1,000 agents, an order-of-magnitude improvement over existing methods. This work, conducted during my postdoctoral research at CMU, established the conceptual foundation I have since developed into a comprehensive research program at NUS.

Upon joining NUS, I extended this framework to address **lifelong MAPF**, the warehouse problem where agents continuously receive a new goal upon reaching their previous one. **PRIMAL2** (RA-L, 2021; 263 citations as of 11 March 2026), developed entirely during my time at NUS with undergraduate and master's students I supervised (Mehul Damani, Zhiyao Luo, Emerson Wenzel), introduced **learned**

conventions: implicit behavioral norms enabling agents to navigate high-traffic corridors and resolve deadlocks without explicit communication. The key insight was that certain coordination strategies, e.g., “never enter a corridor with opposing traffic,” “never reverse unless deadlocked”, cannot be discovered through pure RL but can be instilled through selective imitation learning. PRIMAL2 scaled to 2,000+ agents while matching centralized planners, and convention learning has since become a standard technique in multi-agent coordination research.

Addressing Fundamental Challenges

Building on PRIMAL2, my lab systematically addressed the limitations constraining learning-based MAPF. **SCRIMP** (IROS 2023; 100 citations as of 11 March 2026) introduced a scalable Transformer-based communication mechanism enabling effective coordination with fields of view as small as 3×3 cells (a 10× reduction compared to prior works) by learning precisely what information to share and heed. **ALPHA** (ICRA 2024) addressed long-horizon reasoning by fusing ground-truth proximal information with “fuzzy” distal guidance, achieving 15% improvement over PRIMAL2 in high-density warehouses. **SIGMA** (ICRA 2025) introduced **sheaf theory** from algebraic topology to provide formal structure for decentralized cooperation, enabling agents to learn geometric cross-dependencies with theoretical guarantees. **SYLPH** (Artificial Intelligence, 2025; IF 4.6) addressed symmetry-induced deadlocks through learned social behaviors, such as dynamic shifting between selfish and pro-social roles and hence breaking symmetric configurations.

Culmination: SILLM and 10,000-Agent Coordination

This line of research culminated in **SILLM** (ICRA 2025), which received both a **Best Student Paper Award** and **Best Paper Award on Multi-Robot Systems**, a dual recognition at the flagship robotics conference (4 Best Student Papers from 1,606 accepted, top 0.25%). SILLM, a collaboration with Prof. Jiaoyang Li at CMU, demonstrated lifelong coordination of **up to 10,000 agents**, outperforming state-of-the-art search-based baselines and international competition winners.

SILLM's innovations synthesized insights from my lab's prior work: **spatially sensitive communication** that explicitly preserves spatial relationships during information aggregation (extended from SCRIMP); **multi-modal global guidance** selecting optimal combinations of shortest-path, highway, and congestion-avoidance signals per map type (extended from ALPHA); **Collision Shield PIBT** converting neural outputs to guaranteed collision-free actions (extended from SCRIMP); and **scalable imitation** from the competition-winning WPPL solver. My lab contributed critical hardware validation: we implemented SILLM on 10 physical robots and 100 simulated robots in a mock warehouse at NUS, demonstrating zero-shot sim-to-real transfer.

Impact and Recognition

My MAPF research has achieved substantial academic impact, with core papers accumulating over **1,000 citations** (PRIMAL: 614; PRIMAL2: 263; SCRIMP: 100; ALPHA: 19; SILLM: 16 in <1 year; SYLPH: 10 in < 1y, as of 11 March 2026). Beyond citations, this work has influenced the research community through:

- **Competition success**: I led an NUS team to **First Place (Round 1)** and **Fourth Place Overall** in the Reinforcement Learning track of the NeurIPS 2020 “Flatland” Competition, organized by the Swiss and German national railways (SBB/DB) to advance large-scale train scheduling.
- **Keynote invitations**: I have delivered keynotes on MAPF at AAAI 2023 (Workshop on Multi-Agent Pathfinding), IROS 2024 (Workshop on Multi-Robot Path Planning), ICRA 2025 (Workshop on Multi-Agent Robotic Construction), and CVPR 2025 (Workshop on Multi-Agent Embodied Intelligent Systems).
- **Industrial relevance**: My Amazon Research Award (2022, ~$80k USD + $20k AWS credits) specifically recognized this research direction, funding work on “Distributed Learning for Human-Aware Multi-Agent Pathfinding.”

This body of work establishes learning-based MAPF as a viable, and often superior alternative to classical planning, with direct implications for autonomous warehouse operations at scale. Building on these foundations, my lab is currently developing **PRIMAL3**, targeting extreme scalability and robustness through novel representations, training strategies, and multi-agent reasoning techniques informed by recent advances in foundation models.

2. Informative Path Planning: Autonomous Search and Exploration

Beyond navigating to known destinations, autonomous robots must often explore unknown environments, search for targets of interest, and maintain persistent awareness of dynamic scenes; and achieving all of this while managing constraints like limited battery life, communication range, and sensor coverage. These challenges fall under **informative path planning (IPP)**: selecting trajectories that maximize information gain under uncertainty. My research has developed a comprehensive suite of learning-based methods for exploration, adversarial search, and persistent surveillance, capabilities essential for search-and-rescue, environmental monitoring, and security applications.

Foundational Work: Non-Myopic Exploration with ARiADNE

Classical exploration approaches employ greedy heuristics, such as visiting the nearest frontier, that fail to account for how current decisions affect future opportunities. Prior deep RL methods achieved only marginal improvements because fixed grid representations cannot capture the multi-scale spatial reasoning required for efficient exploration.

ARiADNE (ICRA 2023; 59 citations), led by my PhD student Yuhong Cao (graduated 2024), fundamentally shifted learning-based exploration by introducing an **attention-based architecture** over dynamically constructed graphs. Rather than processing fixed grids, ARiADNE represents the environment as a collision-free graph where stacked attention layers capture both local constraints and global topology through different attention heads. A soft-actor-critic-trained critic implicitly predicts long-term information gains, providing **temporal non-myopicity** without explicit trajectory optimization.

ARiADNE was the first learning method to outperform TARE, the leading hierarchical planner, achieving 5% higher efficiency while planning in real-time (sub-0.1s versus TARE's multi-second cycles). We validated these results on physical robots with 3D LiDARs. ARiADNE established attention-based architectures as the foundation for modern exploration planning, with its **interpretable attention weights** revealing how different heads specialize on local versus global reasoning.

We subsequently extended this foundation through **graph rarefaction with privileged learning** (RA-L, 2024), enabling models trained in small scenarios to scale to environments 10 times larger, and **HDPlanner** (RA-L, 2025), a hierarchical framework that decomposes exploration into subgoal selection, trajectory planning, and low-level control for joint exploration-navigation in search-and-rescue settings.

Most recently, **HEADER** (submitted to the International Journal of Robotics Research (IJRR), 2026) marks what we believe is a milestone for the field: the first learning-based planner to surpass state-of-the-art conventional planners in both exploration efficiency and scalability. HEADER extends ARiADNE with a community detection-based global graph that adaptively partitions the robot's belief with linear complexity and no environment-specific tuning, guidepost features that inject global path information as a soft reference into local decision-making, and a parameter-free privileged expert reward replacing handcrafted shaping with a theoretically grounded surrogate of the true exploration objective. In benchmarks up to 330m × 250m, HEADER achieves on average 20% shorter travel distance than TARE, the strongest conventional baseline, and was validated on hardware in a 300m × 230m outdoor campus area. To our knowledge, this is the largest-scale deployment of any learning-based exploration planner to date.

Multi-Robot Coordination Under Real-World Constraints

Extending to multi-robot exploration introduces coordination challenges when communication is unreliable or sensors have limited coverage. **IR2** (IROS 2024) addresses **sparse, intermittent**

connectivity by learning implicit rendezvous strategies: rather than predetermined meeting schedules, agents learn to position themselves such that connectivity emerges naturally when valuable information needs exchange. **MARVEL** (ICRA 2025) tackles exploration with **constrained directional fields of view** (e.g., cameras or directional LiDARs) by learning to selectively prune the large joint action space (position × heading per robot), bridging theoretical algorithms assuming omnidirectional sensing with practical camera-equipped deployments.

Adversarial Search: Pursuit-Evasion

Many scenarios involve targets that actively avoid detection. **ViPER** (CoRL 2025), led by my PhD student Yizhuo Wang, introduces **visibility-based reinforcement learning** for pursuit-evasion in unknown environments. ViPER addresses the "art gallery problem" where pursuers must ensure no evader enters cleared regions undetected, learning policies that jointly optimize exploration and visibility maintenance. ViPER has been validated through simulations with agent attrition and on physical robots.

This research direction has been supported by continuous collaboration with Singapore's **Defence Science Organisation (DSO)** through Temasek Laboratories at NUS, spanning three completed projects and two ongoing initiatives totalling approximately **S$1.25M as PI**. Current efforts include a project on generative and foundation model approaches for multi-robot exploration, and my team's selection for **DSO's National Drone Challenge**, a competitive program focusing on deep RL-based multi-UAV target search in urban environments.

Persistent Surveillance and Emerging Directions

Beyond one-time exploration, persistent surveillance requires continuous monitoring of dynamic targets. **STAMP** (IROS 2023) introduced spatio-temporal attention for predicting target motion and coordinating robot trajectories. **COMPASS** (MRS 2025, **Best Paper Award Finalist**) extends this to multi-agent surveillance of evasive targets, learning intelligent handoff strategies for smooth coverage transitions as targets move between drones' regions.

Most recently, **CogniPlan** (CoRL 2025) proposed to integrate **foundation models** for semantic exploration, using generative layout prediction to guide robots toward regions likely to contain task-relevant objects. This constitutes a novel shift from purely geometric exploration toward semantically-informed search enabled by pretrained world knowledge.

Impact Summary

Compared to my work on MAPF, informative path planning represents a **research direction I initiated entirely at NUS**. Starting from ARiADNE in 2023, I have built a comprehensive portfolio spanning single-robot exploration, multi-robot coordination under connectivity constraints, adversarial pursuit-evasion, and persistent surveillance, culminating in foundation model integration with CogniPlan and, on the exploration front, **HEADER**, which is the first learning-based planner to outperform state-of-the-art conventional methods in both efficiency and scalability. This body of work has achieved recognition through publication at **CoRL** (ViPER, CogniPlan), the premier robot learning venue; **Best Paper Award Finalist at MRS 2025**; and sustained DSO/Temasek **funding totaling over S$1.25M** across five projects, ongoing for the past six years. This progression demonstrates both research independence and a coherent trajectory addressing increasingly sophisticated real-world challenges in autonomous exploration.

3. Articulated Locomotion and Manipulation

The principles underlying multi-agent coordination extend naturally to the "internal agents" of a single robot and its degrees of freedom (DoFs), joints, or limbs. A hexapod's six legs, a quadruped's four limbs, or a humanoid's dozens of joints can be viewed as cooperative agents that must coordinate to produce coherent whole-body behavior. My research applies distributed learning architectures to articulated systems, enabling complex locomotion and manipulation while improving robustness to disturbances and adaptability across tasks and morphologies.

Bio-Inspired Foundations: From CPGs to Hierarchical Control

During my postdoctoral work at CMU, I established the conceptual foundation for this research direction. My work on **distributed learning of decentralized control policies** (IEEE Transactions on Robotics, 2019) demonstrated that treating each joint as an independent learning agent, i.e., one that shares only local information with neighbors, could produce coordinated whole-body behaviors while naturally providing robustness to actuator failures and morphological variations.

Concurrently, my work on **Central Pattern Generators (CPGs)** (ICRA 2018) established bio-inspired mechanisms for locomotion stability, drawing on how vertebrate spinal circuits generate rhythmic movement patterns. At NUS, I extended this foundation through collaboration with Prof. Auke Ijspeert at EPFL, a world leader in bio-inspired robotics. Our joint work on **hierarchical control emulating the central nervous system** (IROS 2024), led by my PhD student Ge Sun (graduated 2024) during his ORIA-funded visit to EPFL, demonstrated how high-level commands can modulate low-level limb coordination for diverse gaits, mirroring the interplay between brain, spinal CPGs, and local reflexes in vertebrates.

SATA: Torque-Based Control for Safe, Adaptive Locomotion

Prior to my work, the legged robotics community had largely abandoned torque-based control (directly outputting motor torques rather than position commands) due to severe challenges: highly nonlinear dynamics, inefficient exploration, and premature convergence to unnatural gaits. Position-based policies, while easier to learn, produce stiff behaviors unsuitable for safe human-robot interaction.

SATA (RSS 2025; 7 citations in <1 year), led by my PhD student Peizhuo Li with collaborators at EPFL and ETHZ, **reopened this research direction** by demonstrating that torque-based policies can achieve superior compliance, safety, and adaptability. SATA introduced two bio-inspired innovations. First, a **biomechanical muscle model** where neural networks output muscle activation signals passing through biologically-inspired filters, such as activation dynamics, force-velocity relationships, and fatigue tracking, that naturally constrain exploration toward safe, efficient behaviors. Second, a **developmental curriculum** mimicking how animal infants gradually acquire motor skills: training begins with limited torque output and slow control speeds, progressively unlocking capabilities while shifting rewards from basic forward motion to precise velocity tracking.

SATA achieved 100% success across four challenging scenarios (sideways pushing, soft terrain, tunnel traversal, vertical stomping) where position-based baselines including state-of-the-art methods achieved only 0–80%. Our compliance demonstrations are striking: the robot can be strongly pushed or have its legs safely manipulated with minimal resistance, performing corrective steps only when balance is genuinely at risk. A 1.2km autonomous walk across diverse terrains without human correction demonstrated unprecedented robustness. SATA's acceptance at **RSS**, arguably robotics' most prestigious conference, signals strong community recognition.

Building on SATA, we recently developed **APEX** (under review), in collaboration with ETHZ, which integrates expert demonstrations via decaying action priors to bias early exploration toward natural movement styles. APEX enables legged robots to acquire complex dynamic skills within minutes and distill them into deployable, multi-gait policies with no sim-to-real gap.

Collaborative Manipulation: Theory of Mind for Multi-Robot Coordination

Bridging high-DoF control and multi-agent coordination, my recent work extends from single-robot control to **collaborative manipulation**, where multiple robots must jointly manipulate objects. The core challenge is inferring partner intent without explicit communication: each robot observes only its own sensors but must anticipate how partners will act.

LatentToM (CoRL 2025, Oral presentation; ~5% acceptance rate), developed in collaboration with Prof. Mac Schwager at Stanford and led by my PhD student Chengyang He (graduated 2025), introduces a **theory-of-mind** approach to decentralized manipulation. Each robot maintains a latent embedding of partner states inferred through a learned decoder, enabling tight coordination in complex bimanual tasks without explicit communication. The oral presentation granted at CoRL, the top venue for robot learning, shows recognition of this work as an important contribution to the field.

Extending further, **FALCON** (submitted to IEEE Transactions on Robotics (T-RO), 2025), also led by Chengyang He, addresses **loco-manipulation** for mobile manipulators by decoupling locomotion and manipulation into specialized visuomotor diffusion policies coordinated by a central vision-language foundation model. Our new architecture enables coherent whole-body behaviors for a wide range of tasks defined in natural language, despite fundamentally different observation and action spaces for the base and arm subsystems.

Industrial Translation: Collaborative Tugboating

The capabilities detailed above are being translated to real-world applications through a new collaboration with **ST Engineering** (2025–2027, S$620k) on **autonomous multi-agent vessel tugboating**. This project applies multi-robot coordination and collaborative manipulation principles to coordinate multiple autonomous tugboats maneuvering large vessels—a safety-critical application requiring tight coordination under physical coupling constraints.

Impact and Trajectory

This line of research has achieved recognition through publication at **RSS 2025** (SATA) and **CoRL 2025 Oral** (LatentToM), representing acceptance at the field's most selective venues. The progression from bio-inspired CPGs through torque-based locomotion to collaborative manipulation demonstrates a coherent trajectory toward robots that can safely and adaptively interact with humans and each other. My lab is currently extending this work to **humanoid robots** and **heterogeneous multi-robot teams** combining legged platforms with aerial systems for collaborative tasks in unstructured environments.

4. Intelligent Transportation and Smart Cities

Urban traffic control presents a large-scale coordination problem that mirrors the challenges in multi-robot systems: decentralized decision-makers (intersections) must adapt to local conditions while achieving network-wide efficiency. Traditional fixed-time or rule-based signal control struggles with dynamic traffic patterns, contributing to billions in annual economic losses. For instance, in New York City alone, drivers lost 101 hours/year amounting to $9.1B due to congestion, according to the 2023 INRIX report. Aiming to optimize traffic control, my research translates multi-agent coordination principles to urban transportation, developing learning-based methods that can deploy across heterogeneous city networks while coordinating with emerging autonomous vehicle fleets.

Foundations: Cooperative Learning for Traffic Signal Control

I embarked in this new line of research at NUS in 2020 through a collaboration with **ST Engineering** (2020–2022, ~S$780k) on city-wide traffic optimization. This partnership produced **SocialLight** (AAMAS 2023), led by my students Harsh Goel (undergraduate) and Yifeng Zhang (PhD, graduated 2025), which introduced **counterfactual reasoning** to multi-agent traffic signal control. In SocialLight, each intersection estimates its marginal contribution to network-wide flow by comparing actual outcomes against counterfactual scenarios where it acted differently, enabling scalable cooperation without centralized coordination. SocialLight demonstrated that such principles from cooperative game theory could be effectively integrated into distributed reinforcement learning for traffic systems. Extending this work, **CoordLight** (published in 2025 in IEEE Transactions on Intelligent Transportation Systems (T-ITS); Scopus Top-10%; IF 8.4) strengthened decentralized coordination through neighbor-aware policy optimization, learning to anticipate and respond to adjacent intersections' decisions.

Traffic Signal Control Across Heterogeneous Networks

Building on this foundation, we addressed a critical limitation: real city networks contain diverse intersection topologies (3-way, 4-way, complex geometries) and traffic patterns that homogeneous policies cannot handle. **HeteroLight** (IROS 2024), led by Yifeng Zhang, first introduced attention-based representations that adapt to arbitrary intersection structures. This research trajectory recently culminated in **Unicorn** (T-ITS, 2026). Unicorn addresses the dual heterogeneity that prevents real-world deployment of learning-based control: internal heterogeneity (varying intersection topologies and demands) and external heterogeneity (diverse network interconnection patterns).

Unicorn introduces four key innovations. First, a **unified traffic movement representation** using fine-grained connections between incoming and outgoing lanes as a “common language” that links intersection states with signal phases across arbitrary topologies. Second, a **Universal Traffic Representation (UTR)** module employing decoder-only architecture with cross-attention to process variable-length phase and state vectors. Third, an **Intersection Specifics Representation (ISR)** module using a Variational Autoencoder with contrastive learning to distinguish intersection-specific features while maintaining shared decision principles. Fourth, **attention-based collaborative learning** that models state-action dependencies across neighboring intersections for efficient coordination.

Unicorn demonstrated improved performance across **eight benchmark datasets including five real-world city networks** (Ingolstadt, Cologne, Fenglin, Nanshan, and Monaco), representing the first MARL framework successfully integrating universal policy learning with multi-agent collaboration for heterogeneous networks. A single trained model can deploy across entire city networks regardless of intersection diversity, hence dramatically reducing deployment complexity compared to approaches requiring per-intersection tuning.

Toward Mixed-Autonomy Traffic

Looking toward the future of urban mobility, where human-driven and autonomous vehicles will share roads, my lab is developing coordination frameworks for mixed-autonomy traffic. **MACH3** (under review), a collaboration with Prof. Lorenzo Sabattini at UNIMORE (Italy) co-led by his visiting PhD student Alessandro Bonetti and my PhD students Tanishq Duhan and Yifeng Zhang, introduces a hierarchical architecture decomposing navigation into global cooperation, trajectory planning, and motion control, thus enabling autonomous vehicles to coordinate with both infrastructure and each other. **COIN** (under review) identifies critical local vehicle interactions through a collaborative interaction-aware critic, learning which nearby vehicles most influence safe navigation. **V2XFormer** (ITSC 2025) exploits vehicle-to-everything (V2X) communication via a multi-stage transformer to fuse heterogeneous infrastructure and vehicle data, demonstrating how connected autonomous vehicles and intelligent infrastructure can improve both traffic flow and safety. As these new approaches are currently being further extended and refined, we are currently looking at ways to pilot-test these approaches in controlled urban scenarios, to qualitatively and quantitatively assess their positive impact on traffic under real-life conditions.

Industrial Partnership and Impact

This research direction has been sustained through continuous industrial collaboration. Following the initial ST Engineering partnership, I joined the **Cisco-NUS Accelerated Digital Economy Corporate Laboratory** (2021–present, funded by the National Research Foundation Singapore (NRF), ~S$650k for my project) as PI for Work Package 3, developing decentralized coordination for autonomous vehicle routing. These partnerships provide access to real traffic data and deployment pathways while grounding research in practical constraints.

The progression from SocialLight's counterfactual cooperation all the way to Unicorn's universal framework demonstrates a coherent trajectory addressing increasingly realistic deployment challenges. Publication in **IEEE T-ITS**, the leading journal in transportation systems research, ensures broad dissemination to both robotics and transportation engineering communities. This body of work positions my lab to address the emerging challenges of coordinating autonomous vehicles with smart infrastructure as cities transition toward mixed-autonomy mobility, with potential foreseeable next steps including real-life deployment.

Future Directions: Toward Generalizable Robot Intelligence

My research is progressively shifting from task-specific coordination toward **generalizable robot intelligence**—systems that can operate seamlessly in unstructured human environments, adapt to novel situations without retraining, and collaborate fluidly with humans and other robots. Three interconnected thrusts define this vision: foundation models for embodied AI, strategy learning for multi-agent teams, and collaborative heterogeneous systems.

Foundation Models and Test-Time Adaptation

The emergence of large-scale pretrained models, namely visual-language models (VLMs), large language models (LLMs), and visual-language-action (VLA) models, offers unprecedented opportunities to imbue robots with common-sense reasoning and world knowledge. My lab is pioneering the integration of these models into robotic systems while addressing their key limitation: brittleness when deployed in environments that differ from training distributions.

Search-TTA (CoRL 2025), led by my PhD student Derek Tan, introduces a **multimodal test-time adaptation framework** for visual search in the wild. When robots use satellite imagery to search for targets that cannot be directly seen from above, such as wildlife hidden under vegetation, initial predictions from vision-language models like CLIP can be inaccurate due to domain mismatch. Search-TTA addresses this by dynamically refining these probability maps on-the-fly during search using uncertainty-weighted gradient updates inspired by Spatial Poisson Point Processes. The framework accepts queries across multiple modalities (text, images, even sound) through alignment with a shared representation space, achieving zero-shot generalization without modality-specific fine-tuning. This work demonstrates that foundation models can serve as flexible backbones for robotic search when combined with principled online adaptation mechanisms.

Building on this, my lab is developing approaches that leverage foundation models for **semantic reasoning during exploration**. Our recent work, CogniPlan (CoRL 2025), uses generative layout prediction to guide exploration toward task-relevant regions. We are now extending this to multi-agent settings where teams of robots share semantic understanding besides geometric maps, enabling coordination based on object-level and task-level abstractions rather than raw sensor data.

A major ongoing effort integrates VLA models with multi-agent coordination. Current VLA architectures assume single-robot settings; extending them to teams requires addressing how language instructions decompose across agents, how visual observations should be shared or fused, and how action generation should account for coordination constraints. Through my collaborative project with TL@NUS (**funded by Singapore's Ministry of Defense**) on **Cooperative Multi-Agent Learning Powered by Generative and Foundation Models** (2025–2028, S$500k), I am leading the development of architectures that leverage foundation models to enable instruction-following and adaptive coordination at the team level, bridging the gap between single-robot VLA capabilities and multi-agent deployment.

Strategy Learning: From Coordination to Collective Intelligence

Current multi-agent learning methods excel at specific coordination tasks but struggle to transfer skills across scenarios or discover novel strategies to tackle new tasks. The next frontier is strategy learning: enabling agents to acquire generalizable coordination primitives from diverse experiences and transpose them for unseen challenges.

My lab is pursuing this through two complementary approaches. First, **offline-online hybrid multi-agent learning** that extracts coordination strategies from datasets of demonstrations while refining them through targeted environment interactions. Our recent work, **HyGen** (NeurIPS 2025 Workshop on MAS), introduces a novel framework combining offline skill discovery with hybrid policy learning. HyGen uses a global trajectory encoder to extract general skills common across diverse tasks from multi-task offline datasets, then employs a linearly decreasing hybrid ratio that initially leverages offline data efficiency while progressively incorporating online exploration diversity. This approach achieved state-of-the-art performance on the StarCraft Multi-Agent Challenge, demonstrating remarkable zero-shot generalization to unseen tasks by discovering and reusing transferable coordination primitives. Second, **hierarchical strategy representations** that separate what to achieve (strategic objectives) from how to achieve it (tactical execution), enabling transfer across scenarios that share strategic structure but differ in low-level details.

This direction broadly builds upon my lab's MAPF work, where we demonstrated that learned conventions (PRIMAL2), communication protocols (SCRIMP), and social behaviors (SYLPH) could transfer across map types and team sizes. The goal is to systematize this transfer capability: rather than hand-designing which abstractions should generalize, agents should discover reusable strategic primitives through exposure to diverse coordination challenges. This is precisely the capability HyGen's skill discovery mechanism provides.

My new collaboration with **ST Engineering** and **HTX** on **hierarchical cooperation and competition learning** (2026–2029, S$1.5M), in collaboration with colleagues at CMU and UPenn, specifically targets scaling these ideas to swarm scenarios. We are developing architectures that combine the planning capabilities of foundation models with the reactive control of learned policies. These architectures will enable teams to reason about and sequence high-level, task-agnostic strategies while executing robust low-level coordination in complex multi-agent environments involving both cooperative teammates and adversarial opponents.

Collaborative Loco-Manipulation and Heterogeneous Teams

My work on articulated locomotion (SATA, APEX) and collaborative manipulation (LatentToM, FALCON) is converging toward **integrated loco-manipulation** for mobile manipulators and humanoids. The core challenge is coordinating fundamentally different subsystems, such as legs optimized for stability and traversal, arms optimized for dexterity and manipulation, under a unified control architecture that can reason about whole-body behavior.

Current efforts extend FALCON's decoupled visuomotor policies to more complex platforms, including quadrupeds with mounted arms and early-stage humanoid systems. The vision-language coordination layer that proved effective for single-robot loco-manipulation naturally extends to **multi-robot collaborative loco-manipulation**, where heterogeneous robots (e.g., a legged platform and a drone) must coordinate to accomplish tasks neither could achieve alone, such as a quadruped manipulating objects while a drone provides aerial observation and guidance.

The **ST Engineering collaboration** on **autonomous multi-agent vessel tugboating** (2025–2027, S$620k) serves as a testbed for these ideas: multiple autonomous tugboats must tightly coordinate their forces to maneuver large vessels, requiring both individual platform control and team-level cooperation under physical coupling constraints. Success here will demonstrate that principles developed for legged robots and manipulators are transferable to maritime systems, thus validating the generality of our coordination frameworks.

Convergent Vision

My lab's directions converge toward robots that exhibit **everyday competence**: the ability to understand natural language instructions, perceive and reason about complex environments, coordinate with humans and other robots, and adapt when situations deviate from expectations. The progression from my foundational work on multi-agent pathfinding and articulated control, which established that decentralized learning could achieve scalable coordination, all the way to current efforts integrating foundation models and strategy learning, traces a coherent trajectory toward this vision.

My sustained funding pipeline (>S$3M in active grants as PI), established international collaborations (EPFL, Stanford, CMU, ETHZ), and growing research group (15+ PhD students and postdocs) provide the resources to pursue this ambitious agenda. The recognition my work has received (ICRA 2025 Best Paper Awards, publications at RSS and CoRL, Amazon Research Award, editorial roles at IJRR and RA-L) reflects the community's confidence that this research direction addresses fundamental challenges in robotics and AI. I am committed to advancing the field toward robot systems that can operate as capable, adaptive partners in human environments.