Research Statement

Hoong Chuin Lau
School of Computing and Information Systems, Singapore Management University

Vision

I build decision systems that help organisations act well when the future is uncertain and resources are scarce. My work spans the intersection of Artificial Intelligence and Operations Research (OR), as well as tests where quantum and quantum-inspired computation can provide computational improvements.

The problems I study — routing and scheduling fleets, deploying police and emergency responders, managing maritime traffic, securing supply chains — are dynamic, constrained and only partly observed. Prediction alone does not tell an operator what to do; an optimisation model that assumes a static, fully known world gives answers that break on contact with reality. My central method is therefore to put learning and optimisation in one loop: anticipate what may happen, decide what to do, and revise the plan as conditions change.

Three commitments run through everything I do:

- **Use-inspired problems.** I start from operational problems posed by partners in logistics, public safety, maritime and healthcare, and judge methods by whether they improve real decisions, not only benchmark scores.
- **Rigour against strong baselines.** New tools, deep learning, large language models, quantum circuits, earn a place only when they beat well-engineered classical methods on fair comparisons.
- **Translation to working systems.** Over three decades I have taken research into deployed software, patents, a corporate laboratory and a start-up, so that ideas are tested against real operating conditions.

My research is organised in three thrusts:

Thrust 1: Human-centred planning for resilience

Question. *How should scarce resources be planned and re-planned when demand shifts, information is incomplete, and the decisions fall on people?*

I combine prediction, simulation and optimisation in a single operational loop. In public safety, machine-learning-informed optimisation produced deployment plans that cut manpower while holding response-time standards (IJCAI 2019), and became the GRAND-VISION and OFFICERS systems (ICAPS 2021, 2022). Multi-agent reinforcement learning then let us dispatch and reschedule police patrols as incidents unfold (ICAPS 2022, IJCAI 2023). In logistics and maritime operations, the same principle drives work on coordinated mall deliveries with limited docks (*Transportation Science* 2022), risk-aware procurement under COVID-19 disruption (*TRE* 2025), tugboat scheduling that trades fuel against punctuality (*TRE* 2025), and storage and corridor design for energy security under chokepoint disruption (*TRE* 2026).

Two lessons from this work shape what comes next. First, robustness must be specified against the disruptions operators actually fear; my current position work for the planning community questions robustness notions that hedge against mathematically convenient but operationally irrelevant

uncertainty. Second, an efficient plan is not enough: it must respect service quality, fairness and workforce limits, and be explainable to the people who carry it out.

Thrust 2: Learning to search, and explaining the result

Question. *Can optimisation algorithms learn how to search well across changing instances, while staying transparent to the people who use them?*

No single algorithm dominates across the instances and time budgets that operations present. My earliest work here, on automated parameter tuning and algorithm portfolios (AAAI 2015; LION 2015 best-paper runner-up), treated algorithm design itself as a learning problem. Today I use reinforcement learning and neural models to control and augment classical search rather than replace it: deep RL that adapts Adaptive Large Neighbourhood Search online (ICAPS 2024), autoregressive RL for mixed-integer programs (LION 2025), one-shot learned insertion for dynamic routing (*IEEE TKDE* 2025), search-trajectory networks for multi-objective algorithm configuration (ECAI 2025), and diffusion-enhanced generative flow networks for vehicle routing (NeurIPS 2026).

I am equally interested in what these models learn. Probing studies reveal how neural combinatorial-optimisation models encode problem structure (NeurIPS 2025), and graph-based methods produce step-wise explanations for why a constraint problem has the solution it has (CP 2026). A related line discovers process routings and decision drivers directly from enterprise databases (BPM 2026), turning operational records into models that can be optimised.

The goal is algorithm engineering that is systematic rather than hand-tuned: systems that diagnose an instance, choose and configure a search strategy, and can justify their answer to a planner.

Thrust 3: Quantum and hybrid methods for hard decisions

Question. *Where, if anywhere, can quantum and quantum-inspired computation improve real decisions — and how do we tell?*

My group came to quantum computing from practice: tuning and enhancing industrial quantum-inspired QUBO solvers for routing and permutation problems (GECCO and IJCAI workshops 2022; IEEE QCE 2021). That experience made the hard questions concrete — how to encode constraints, preserve feasibility, exploit structure, and combine quantum subroutines with classical search and learning. Our answers include a feasibility-preserving quantum approximate solver for capacitated vehicle routing (*Quantum Information Processing* 2024), slack-free penalty functions and variational imaginary-time evolution for constrained problems (IEEE QCE 2025), light-cone cancellation for noisy Max-Cut (*Scientific Reports* 2026), and LLM-guided closed-loop search over adaptive variational policies (IEEE QCE 2026).

Two features set this work apart. First, it is benchmarked on real hardware against strong classical baselines, including head-to-head comparisons of standard and qubit-efficient formulations (*Quantum Science and Technology* 2026). Second, it reaches beyond combinatorial search to decisions under risk: quantum amplitude estimation for simulation-based optimisation (IEEE QCE 2024), with manuscripts in preparation on risk-averse inventory decisions and on quantum-assisted scenario exploration for pandemic digital twins — the latter a direct bridge to HOPE. I recently wrote a survey, *Beyond QAOA*, maps where AI and quantum methods for combinatorial optimisation meet.

My stance is deliberately sceptical: identify where a quantum or hybrid approach might help, test it fairly, and report plainly when it does not.

Looking Forward

The three thrusts converge on one question: *how do we build decision technology that works in the world, not just on a benchmark?* Over the next five years I will pursue it along four lines.

1. **Integrated crisis planning.** Through HOPE, connect supply-chain, emergency-response and healthcare models into digital twins that agencies can use to rehearse and coordinate decisions across sectors.
2. **Robustness that matches reality.** Develop planning methods whose uncertainty sets come from the disruptions operators actually face, and that adapt from operational feedback rather than fixed assumptions.
3. **Human-centred and explainable optimisation.** Build fairness, workforce sustainability and step-wise explanations into planning models, and study how planners work alongside AI, including large language models, without ceding judgement to them.
4. **Honest assessment of quantum advantage.** Test quantum and hybrid methods, especially amplitude estimation for risk and scenario analysis, on decision problems that matter, against the best classical methods, on real hardware.

Research Leadership

My research culminates in **HOPE (Human-Centered AI in Operational Planning and Resilience)**, a five-year, Ministry of Education–funded research programme I lead at SMU starting July 2026. HOPE asks how Singapore can plan and coordinate national resources before, during and after public-health crises and other large-scale emergencies.

HOPE aims to tackle problems at national scale through three connected pillars — resilient supply chains and logistics, resilient emergency response, and resilient healthcare operations. Methods for these domains are usually built in isolation; HOPE's aim is shared decision infrastructure, so that agencies coordinate rather than improvise when a crisis arrives. The ambition is to develop a science for operational resilience that help organisations and communities meet uncertainty and crisis through principled, responsible and human-centered planning.

Aligning with and extending from this effort, I am starting a research center called HOPE (Human-Centered AI for Operational Planning and Resilience) which will give SMU a distinctive institutional platform for research at the intersection of artificial intelligence, operations research, social science and public health. The Centre for HOPE will tackle the central question of how organizations and intelligent systems can learn, decide, coordinate and act safely under uncertainty, especially in healthcare, emergencies, supply chains and other mission-critical settings. The center will be organized into three pillars:

Pillar	Core Scientific Scope	Role in Operational Loop
P1: Human-Centred Decision Intelligence	AI-OR integration; optimisation, planning and scheduling; forecasting; robust and stochastic decisions; causal and behavioural modeling; explainability and human-AI trust; digital twins; fairness, equality and responsible AI.	Understand the situation and determine desirable actions under operational, ethical and human constraints.
P2: Agentic AI for Planning & Coordination	Goal-directed, tool-using agents (incl. LLM-based reasoning); multi-agent planning and negotiation; human-agent teaming; solver and simulator integration; verification, guardrails and override.	Decompose goals, formulate plans, coordinate participants and adapt as circumstances change.
P3: Collaborative AI for Human-Machine Execution	Human-robot teaming; shared situational awareness; mixed-fleet coordination (drones, ambulances, police, trucks); shared autonomy and safe handoff; simulator-to-real transfer; field safety, verification and human override.	Execute, monitor and adapt selected actions in the physical world.
Unifying Loop	Understand and decide → plan and coordinate → act and learn. Human goals, values, supervision, and authority remain present throughout the loop.	

Translation and impact (since 2020)

Projects	Partners	Outcomes
AI supply-chain resiliency	IBM	US$35M in reported savings; 2026 SICC Award, Most Scalable Collaboration
Collaborative Urban Delivery Optimisation (CUDO)	Keppel Logistics, Commonwealth Kokubu, Y3/YCH, Gurusoft, uParcel; AI Singapore	20% efficiency gain; Singapore patent granted 10202101416T (granted 2026) and licensed to logistics operators
GRAND-VISION deployment and scheduling	MHA, SPF and SCDF	Licensed to Fujitsu Ltd for deployment in police departments across Japan
Crowd management system	Suntec Convention Center	Singapore patent 10201509975U (granted 2017)
Fujitsu–SMU Urban Computing and Engineering Corp Lab (UniCEN)	Fujitsu; A*STAR IMDA, LTA, MPA Jurong Port	Director, 2014–2021; Maritime traffic, port, crowd management and logistics systems
horaion.com (startup)	TTSH Various hotels	Commercial hybrid OR–AI manpower planning and scheduling
HOPE (Human-Centered AI for Operational Planning and Resilience) Center	MOH, MHA, CDA, HTX SGH, Singhealth, NHG Health CNRS@Create, NUS	Director, 2026-2031; Ongoing

Selected publications (for other publications not listed, please refer to my homepage http://www.mysmu.edu/faculty/hclau/)

Human-centred planning, logistics and resilience

- Sharma M. & Lau H. C. (2026). Securing the Flow: Strategic Storage and Corridor Design for Maritime Energy Security under Chokepoint Disruptions. Transportation Research Part E: Logistics and Transportation Review, Forthcoming.
- Case, J., Lau, H. C., Yang, J., & Liu, L. (2025). Multi-period risk-aware procurement optimization under COVID-19 disruption. Transportation Research Part E: Logistics and Transportation Review, 202, 104272.
- Wei, X., Lau, H. C., Xiao, Z., Fu, X., Zhang, X., & Qin, Z. (2025). Bi-objective dynamic tugboat scheduling with speed optimization under stochastic and time-varying service demands. Transportation Research Part E: Logistics and Transportation Review, 193, 103876.
- Chen, D., Huang, C., Fan, T., Lau, H. C., & Yan, X. (2025). Predictive modeling for vessel traffic flow: A comprehensive survey from statistics to AI. Transportation Safety and Environment, 7(3).
- Yang, J., Lau, H. C., & Wang, H. (2024). Optimization of customer service and driver dispatch areas for on-demand food delivery. Transportation Research Part C, 165, 104653.
- Joe, W., & Lau, H. C. (2023). Learning to send reinforcements: Coordinating multi-agent dynamic police patrol dispatching and rescheduling via reinforcement learning. Proceedings of IJCAI 2023.
- Song, R., Lau, H. C., Luo, X., & Zhao, L. (2022). Coordinated delivery to shopping malls with limited docking capacity. Transportation Science, 56(2).
- Chase, J., Nguyen, D. T., Sun, H., & Lau, H. C. (2019). Improving law enforcement daily deployment through machine learning-informed optimization under uncertainty. Proceedings of IJCAI 2019.

Learning-enhanced algorithm engineering

- Zhang, N., Zhang Z., Pan L., Lau H. C., Cao, Z (2026). Diffusion-Enhanced GFlowNet for Solving Vehicle Routing Problems. Proceedings of Neural Information Processing Systems (NeuRIPS), Sydney, Australia, 2026.
- Zhang, Z., Cao, Z., & Lau, H. C. (2025). Neuro-Ins: A learning-based one-shot node insertion for dynamic routing problems. IEEE Transactions on Knowledge and Data Engineering, 37(9), 5495–5507.
- Reijnen, R., Bukhsh, Z. A., Lau, H. C., Wu, Y., & Zhang, Y. (2025). Search trajectory network-enhanced multi-objective dynamic algorithm configuration. Proceedings of ECAI 2025.
- Zhang, Z., Ma, Y., Cao, Z., & Lau, H. C. (2025). Probing neural combinatorial optimization models. Proceedings of NeurIPS 2025.
- Reijnen, R., Zhang, Y., Lau, H. C., & Bukhsh, Z. (2024). Online control of Adaptive Large Neighborhood Search using deep reinforcement learning. Proceedings of ICAPS 2024, 475–483.
- Wang, S., Zhang, X., Li, J., Wei, X., Lau, H. C., et al. (2024). Fuel-saving route planning with data-driven and learning-based approaches: A systematic solution for harbor tugs. Proceedings of IJCAI 2024, 7483–7490.

Quantum and hybrid optimisation

- Sharma M. & Lau, H. C. (2026). From Circuits to Hardware: Benchmarking Standard and Qubit-Efficient Quantum Optimization on Real Hardware. *Quantum Science and Technology* , Volume 11, no. 4.
- Lee, X. W., & Lau, H. C. (2026). Light cone cancellation for variational quantum eigensolver in solving noisy Max-Cut. Scientific Reports, 16, 9597. https://doi.org/10.1038/s41598-025-31798-1
- Lee, X. W., & Lau, H. C. (2025). Implementing slack-free custom penalty function for QUBO on gate-based quantum computers. Proceedings of IEEE Quantum Computing and Engineering Conference (QCE 2025). https://arxiv.org/abs/2504.12611
- Lee, X. W., & Lau, H. C. (2025). Solving constrained combinatorial optimization problems with variational quantum imaginary time evolution. Proceedings of IEEE QCE 2025. https://arxiv.org/abs/2504.12607
- Xie, N., Lee, X. W., Cai, D., Saito, Y., Asai, N., & Lau, H. C. (2024). A feasibility-preserved quantum approximate solver for the capacitated vehicle routing problem. Quantum Information Processing, 23, 291. https://arxiv.org/abs/2308.08785
- Sharma, M., Lau, H. C., & Raymond, R. (2024). Quantum enhanced-simulation based optimization for newsvendor problems. Proceedings of IEEE QCE 2024. https://arxiv.org/abs/2403.17389
- Sharma, M., Yan, J., Lau, H. C., & Raymond, R. (2024). Quantum relaxation for solving multiple knapsack problems. Proceedings of IEEE QCE 2024.