Curriculum Vitae - Jesse van Remmerden

PhD Candidate, Eindhoven University of Technology (TU/e)
j.v.remmerden@tue.nl | jessevanremmerden@gmail.com
Website | Google Scholar | ORCID | GitHub | LinkedIn

Research Profile

I develop learning-based methods for **combinatorial optimization**: reinforcement learning, discrete diffusion models and language models for job-shop scheduling, vehicle routing and infrastructure maintenance planning. One question runs through my work: *what can learned planners do without a simulator, without expert demonstrations, and without a hand-specified objective?* My results show that none of these is a fundamental requirement. I built the first fully offline RL approach to job-shop scheduling (*Machine Learning*, 2025) and its successor, which performs best when trained on *random* schedules (*TMLR*, 2026). I built the first neural routing solver that follows constraints stated in plain English, satisfying them in 98–99% of cases where existing neural solvers manage ~32% (under review). I also built multi-objective RL that optimizes the risk measure the City of Amsterdam's asset managers use in practice (*Neural Computing and Applications*, 2025). For every project I built the training data, the automatic verifier and the benchmark myself, and all published work comes with open-source code.

Research interests: neural combinatorial optimization; offline, distributional and multi-objective reinforcement learning; discrete diffusion and flow models; language-conditioned optimization; learning for scheduling and vehicle routing.

Education

PhD Candidate, Artificial Intelligence for Planning and Scheduling Jan. 2023 - Jan. 2027 (exp.)
Eindhoven University of Technology (TU/e), Information Systems group, Dept. of Industrial Engineering & Innovation Sciences

- Promotor: Dr. Yingqian Zhang; daily supervisor: Dr. Zaharah A. Bukhsh.
- Five first-author papers (three journal, one workshop, one under review); three open-source codebases and public datasets.

MSc Artificial Intelligence, Utrecht University Feb. 2020 - Jun. 2022

- *Cum Laude* (8.2/10). Thesis (8.6/10): Utilising Reinforcement Learning for the Diversified Top-k Clique Search Problem.
- Honours Young Innovators programme, Utrecht Honours College.

BSc Computer Science, University of Amsterdam Sept. 2016 - Oct. 2019

- Thesis (9/10): detecting dolphins in video footage with deep learning (Faster R-CNN).

Publications

All first-authored; author name in bold. Full list on Google Scholar.

Journal Articles

[J3] **J. van Remmerden**, Z. Bukhsh, Y. Zhang. Generalizing Beyond Suboptimality: Offline Reinforcement Learning Learns Effective Scheduling through Random Solutions. *Transactions on Machine Learning Research (TMLR)*, 2026. Paper · arXiv · Code · Data · Project page

[J2] **J. van Remmerden**, Z. Bukhsh, Y. Zhang. Offline reinforcement learning for learning to dispatch for job shop scheduling. *Machine Learning*, 114(8), 191, 2025. Presented in the ECML PKDD 2025 Journal Track. DOI · arXiv · Code

[J1] **J. van Remmerden**, M. Kenter, D. M. Roijers, C. Andriotis, Y. Zhang, Z. Bukhsh. Deep multi-objective reinforcement learning for utility-based infrastructural maintenance optimization. *Neural Computing and Applications*, 37, 24719–24742, 2025. DOI · arXiv · Code

Under Review

[P1] **J. van Remmerden**, Y. Wu, Y. Zhang, Z. Bukhsh. DROVER: Diffusion-Based Routing with Verbal Preferences. Under review, 2026.

Workshop Papers

[W1] **J. van Remmerden**, M. Kenter, D. M. Roijers, Y. Zhang, C. Andriotis, Z. Bukhsh. A Deep Multi-Objective Reinforcement Learning Approach for Infrastructural Maintenance Planning with Non-Linear Utility Functions. *Multi-Objective Decision Making Workshop (MODeM) at ECAI 2023*, Kraków, Poland, 2023. Paper

Research Experience and Collaborations

Visiting Researcher, Technical University of Munich (TUM) Nov. 2025 - Jan. 2026
Munich, Germany; research visit during the PhD

STABILITY consortium, NWO Dutch Research Agenda (NWA.1431.20.004) 2023 - 2027
Sustainable Circular Life Extension Strategies for Inner-City Bridges and Quay Walls; funded by NWO and the City of Amsterdam

- Worked directly with the City of Amsterdam's asset managers to formulate 50-year maintenance planning for the city's historic quay walls as a multi-objective RL problem, encoding their FMECA risk-scoring method as the objective ([J1], [W1]).

Selected Research

DROVER: vehicle routing from constraints written in plain English 2025 - Present
First author; under review [P1]; presented at DSO @ IJCAI-ECAI 2026. Method: language-conditioned discrete diffusion

- Neural routing solvers optimize a fixed scalar objective and cannot take requests such as “visit 7 before 4” or “keep these three stops on one truck”. I built the first benchmark for this setting: 26+ request types, ~3×10^9 possible phrasings, reference solutions from established solvers, and an automatic verifier that checks whether a tour obeys the request.
- Designed a discrete diffusion model that reads the request through a frozen language model and steers its denoising toward satisfying it with a new guidance technique.
- Obeys the request in 98–99% of cases on 50 to 1,000 stops, versus ~32% for the strongest existing neural solvers, and 87–90% on request types never seen in training (existing solvers: 0%). Without any request, it remains competitive with the best neural solvers on tour length (best at 100 stops).

CDQAC & Offline-LD: learning to schedule from data, not simulators 2023 - 2026
First author; Machine Learning (2025) [J2] and TMLR (2026) [J3]. Method: offline and distributional RL for (flexible) job-shop scheduling

- Offline-LD was the first fully end-to-end offline RL approach for job-shop scheduling: maskable Quantile-Regression DQN and discrete SAC trained with Conservative Q-Learning on a fixed set of past schedules. From only 100 constraint-programming solutions it matches or beats simulator-trained baselines that need millions of interactions.
- CDQAC (Conservative Discrete Quantile Actor-Critic) matches or beats state-of-the-art simulator-trained methods on four benchmark suites. It generalizes to instances with ten times as many jobs as seen in training, reaches competitive performance from only 10–25 training instances, and performs *best* when trained on schedules from a random heuristic.
- Showed that offline performance is governed by how much of the decision space the data covers, not by solution quality (rank correlation −0.95). This gives a practical rule for building datasets for learned decision-making.

MO-DCMAC: 50-year maintenance planning for Amsterdam's historic quay walls 2023 - 2025
First author, with the City of Amsterdam; Neural Computing and Applications (2025) [J1]. Method: multi-objective RL

- Developed a deep multi-objective actor-critic that optimizes non-linear utility functions directly instead of a shaped scalar reward, with FMECA, the risk-scoring method the city's asset managers use, as the objective. It found the best plan in 5 of 6 settings, up to 46% better than the rule-based strategies in use, and is the first deep multi-objective RL method of its kind applied to a real-world problem.

Talks, Tutorials and Posters

Tutorial

Aug. 2026 **Deep Reinforcement Learning for Combinatorial Optimization.** Half-day tutorial at IJCAI-ECAI 2026, Bremen, Germany; one of five speakers (with Z. Bukhsh, Y. Wu, Y. Zhang, Y. Yu), including a hands-on part with code walkthroughs for job-shop scheduling and CVRP solvers. Materials

Conference and Workshop Talks

Aug. 2026 **DROVER: Diffusion-Based Routing with Verbal Preferences.** Data Science Meets Optimisation (DSO) Workshop at IJCAI-ECAI 2026, EURO working group, Bremen, Germany.

Apr. 2026 **Generalizing Beyond Suboptimality: Offline RL Learns Effective Scheduling through Random Solutions.** NWO ICT.OPEN 2026, Smart Industry track, Rotterdam, the Netherlands.

Sept. 2025 **Offline Reinforcement Learning for Learning to Dispatch for Job Shop Scheduling.** ECML PKDD 2025, Journal Track, Porto, Portugal.

Apr. 2024 **Utility-based multi-objective RL for infrastructural maintenance planning (MO-DCMAC).** Oral presentation, NWO ICT.OPEN 2024, Utrecht, the Netherlands.

Oct. 2023 **A Deep Multi-Objective RL Approach for Infrastructural Maintenance Planning with Non-Linear Utility Functions.** MODeM Workshop at ECAI 2023, Kraków, Poland.

Posters

Jun. 2026 **Generalizing Beyond Suboptimality: Offline RL Learns Effective Scheduling through Random Solutions (CDQAC).** 20th Learning and Intelligent Optimization Conference (LION20), Milan, Italy.

Jul. 2024 **Multi-objective RL for infrastructural maintenance planning (MO-DCMAC).** DLRL Summer School 2024 (CIFAR, Amii, Mila, Vector Institute), Toronto, Canada. Poster

Teaching and Supervision

- **Thesis supervision:** supervised thesis students in the Information Systems group, TU/e.
- **Teaching:** teaching-assistant work at TU/e; co-presented the IJCAI-ECAI 2026 tutorial on deep RL for combinatorial optimization (see above).

Service and Leadership

- **Member, PhD Council**, Eindhoven University of Technology (2024 - 2025).
- **Co-author**, 2022 municipal election programme, Volt Amsterdam (2021).
- **Head of the Health, Welfare and Sports Committee**, Jonge Democraten Amsterdam (2016 - 2017).

Awards, Honours and Schools

- **Cum Laude**, MSc Artificial Intelligence, Utrecht University (2022).
- **Honours Young Innovators programme**, Utrecht Honours College (MSc).
- **Participant**, Deep Learning and Reinforcement Learning (DLRL) Summer School, CIFAR, Amii, Mila and Vector Institute, Toronto (2024).
- **First prize**, Odyssey Hackathon 2019 (Nature 2.0 track): a federated-learning, blockchain and evolutionary-computing proof-of-concept built in one weekend.

Open-Source Software and Data

- **CDQAC** [J3]: github.com/jesserem/CDQAC_scheduling; offline RL scheduling datasets and checkpoints on Hugging Face.
- **Offline-LD** [J2]: github.com/jesserem/Offline-LD.
- **MO-DCMAC** [J1]: github.com/jesserem/MODCMAC.

Other Professional Experience

Blockchain Engineer, Kryha, Amsterdam Aug. 2018 - Jan. 2020

- Built proofs-of-concept with smart contracts, tokenization, cryptography and federated learning, connecting new technologies to real-world use cases.

Technical Skills

Programming: Python (8+ yrs.; PyTorch, PyTorch Lightning, JAX), C/C++, Rust, SQL, LaTeX.
Machine learning: offline and online RL (CQL, QR-DQN, SAC, multi-objective RL), discrete diffusion and classifier-free guidance, graph transformers, conditioning on frozen LLM encoders, verifier and benchmark design, synthetic data generation.
Optimization and compute: OR-Tools (CP-SAT), Gurobi, LKH; Linux/HPC (SLURM, Apptainer, Kubernetes/Run:AI), multi-GPU training on A100/H100/B200 (Snellius), Docker, Git, Weights & Biases.
Languages: Dutch (native), English (fluent, C1).

References

Dr. Yingqian Zhang, PhD promotor, Information Systems, Eindhoven University of Technology, yqzhang@tue.nl
Dr. Zaharah A. Bukhsh, daily supervisor, Information Systems, Eindhoven University of Technology, z.bukhsh@tue.nl