



**HAL**  
open science

# LATENT PERSPECTIVE-TAKING VIA A SCHR ÖDINGER BRIDGE IN INFLUENCE-AUGMENTED LOCAL MODELS

Kevin Alcedo, Pedro U Lima, Rachid Alami

► **To cite this version:**

Kevin Alcedo, Pedro U Lima, Rachid Alami. LATENT PERSPECTIVE-TAKING VIA A SCHR ÖDINGER BRIDGE IN INFLUENCE-AUGMENTED LOCAL MODELS. World Modeling Workshop, Feb 2026, Montreal, Canada. ⟨hal-05562006⟩

**HAL Id: hal-05562006**

**<https://hal.science/hal-05562006v1>**

Submitted on 22 Mar 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

# LATENT PERSPECTIVE-TAKING VIA A SCHRÖDINGER BRIDGE IN INFLUENCE-AUGMENTED LOCAL MODELS

Kevin Alcedo<sup>1,2</sup>, Pedro U. Lima<sup>1,2</sup>, Rachid Alami<sup>3,4</sup>

<sup>1</sup>Institute for Systems and Robotics,

<sup>2</sup>Instituto Superior Técnico, Universidade de Lisboa (Lisbon ELLIS Unit),

<sup>3</sup>LAAS-CNRS,

<sup>4</sup>Artificial and Natural Intelligence Toulouse Institute (ANITI)

{kevin.alcedo, pedro.lima}@tecnico.ulisboa.pt, rachid.alami@laas.fr

## ABSTRACT

Operating in environments alongside humans requires robots to make decisions under uncertainty. In addition to exogenous dynamics, they must reason over others’ hidden *mental-models* and *mental-states*. While Interactive POMDPs and Bayesian Theory of Mind formulations are principled, exact nested-belief inference is intractable, and hand-specified models are brittle in open-world settings. We address both by learning structured mental-models and an estimator of others’ mental-states. Building on the Influence-Based Abstraction, we instantiate an Influence-Augmented Local Model to decompose socially-aware robot tasks into local dynamics, social influences, and exogenous factors. We propose (a) a neuro-symbolic world model instantiating a factored, discrete Dynamic Bayesian Network, and (b) a perspective-shift operator modeled as an amortized Schrödinger Bridge over the learned local dynamics that transports factored egocentric beliefs into other-centric beliefs. We show that this architecture enables agents to synthesize socially-aware policies in model-based reinforcement learning, via decision-time mental-state planning (a Schrödinger-Bridge in belief space), with preliminary results in a MiniGrid social navigation task.

## 1 INTRODUCTION

A socially-aware robot acts alongside humans and robots, pursuing its own goals while enabling others to achieve theirs and accommodating known or inferred needs. Hence, it must make decisions under partial observability: both of the environment and of others’ *mental-models* and *mental-states*. We adopt the lens of classical Interactive POMDPs (I-POMDPs) (Gmytrasiewicz & Doshi, 2004) and Bayesian Theory of Mind (BToM) (Baker et al., 2011): *mental-models* are (known or inferred) internal transition, observation, reward and action models, and *mental-states* are probability distributions over latent states (beliefs). The knowledge-representation community has formalized perspective-taking (Flavell, 1992) with ontologies/semantic memory (Milliez et al., 2014; Sarthou et al., 2019), epistemic human-aware task planners (Shekhar et al., 2025), and dynamic epistemic logic (Bolander & Andersen, 2011; Bolander et al., 2021), yielding interpretable and verifiable systems. Yet these symbolic models are brittle in open-world settings with stochastic dynamics, and exact nested-belief reasoning quickly becomes intractable; PSPACE at level-0, and EXPSPACE in the worst case with factored dynamics and exponential horizon (Mundhenk et al., 2000). In contrast, probabilistic sequential decision-making naturally handles uncertainty but typically assumes prescribed models and still suffers from the nested-belief explosion.

As a step toward unifying symbolic knowledge representation and probabilistic sequential decision-making, we propose a data-driven architecture that learns (i) factored, discrete mental-models and (ii) a tractable belief-to-belief map for perspective-taking in factored latent space. The key idea is to cast perspective-taking as belief transport via a Schrödinger Bridge (SB) (Fortet, 1940) defined over learned local reference dynamics, producing other-centric beliefs compatible with local dynamics. Building on the Influence-Based Abstraction (IBA) framework (Oliehoek et al., 2021),



Chetrite et al., 2021); the goal is to discover a new law describing the dynamics of the trajectory. SB can be viewed as a special case of Optimal Transport: *dynamic entropic optimal transport with reference dynamics* (Léonard, 2014); we minimize a pathwise KL divergence to the trajectory distribution induced by the reference transition model.

Given a reference local dynamics  $\dot{T}(x_{t+1} | x_t, a_t)$ , start/end marginals (beliefs)  $b_0$  and  $b_n$ , and conditioning on a fixed action sequence  $a_{0:n-1}$ , we seek a bridge  $\bar{T} = \{\bar{T}_t\}_{t=0}^{n-1}$  (a family of time-varying transitions) whose induced trajectory distribution is closest to that of the reference  $\dot{T}$ . Concretely, we compute  $\bar{T}$  such that its induced trajectory distribution  $\mathbb{P}_{\bar{T}}$  minimizes  $\text{KL}(\mathbb{P}_{\bar{T}} \| \mathbb{P}_{\dot{T}})$ , subject to the endpoint marginals  $(b_0, b_n)$ . This yields a first-order, time-inhomogeneous Markov chain. Here we focus on the space-time Sinkhorn method (Fortet, 1940), where forward and backward potentials  $\phi_t, \psi_t \in \mathbb{R}_{>0}^{|\mathcal{X}|}$  are propagated through the reference  $\dot{T}$  and produce the bridge. We parametrize the potentials with context  $c_t := D_{t+1} = d(h_t)$  with  $h_t = \langle x_1, a_1, \dots, a_{t-1}, x_t \rangle$ . The Schrödinger system becomes:

$$\begin{aligned} \psi_t(x_t; c_t) &= \sum_{x_{t+1}} \dot{T}(x_{t+1} | x_t, a_t) \psi_{t+1}(x_{t+1}; c_{t+1}), \\ \phi_{t+1}(x_{t+1}; c_{t+1}) &= \sum_{x_t} \phi_t(x_t; c_t) \dot{T}(x_{t+1} | x_t, a_t), \\ p_t(x_t) &= \phi_t(x_t; c_t) \psi_t(x_t; c_t), \quad p_0 = b_0, \quad p_n = b_n \end{aligned} \tag{2}$$

The steered transitions take the Doob  $h$ -transform form (Doob, 1957):

$$\bar{T}_t(x_{t+1} | \bar{s}_t, a_t) = \dot{T}(x_{t+1} | x_t, a_t) \frac{\psi_{t+1}(x_{t+1}; c_{t+1})}{\psi_t(x_t; c_t)}, \quad \text{where } \bar{s}_t := \langle x_t, D_{t+1} \rangle \tag{3}$$

The induced trajectory distribution, with initial state distribution  $\mu_0$ , is:

$$\mathbb{P}_{\bar{T}}(x_{0:n}) \propto \phi_0(x_0; c_0) \mu_0(x_0) \left[ \prod_{t=0}^{n-1} \dot{T}(x_{t+1} | x_t, a_t) \right] \psi_n(x_n; c_n) \tag{4}$$

One way to view this process is as control-as-inference (Todorov, 2006), targeting a minimum-effort steering of a noisy passive dynamic process. In Section 3 we instantiate an IALM as a neuro-symbolic world model that learns  $(\bar{\mathcal{S}}, \bar{T}, \mathcal{O}, \bar{O}, D)$ . We then realize  $\bar{T} = \{\bar{T}_t\}$  as an approximation of the ideal IALM transition in Eq. 1 via an amortized Schrödinger Bridge tilt of  $\dot{T}$  for perspective-taking. The potentials  $\phi$  and  $\psi$  are trained offline with epistemic counterfactuals.

## 3 METHOD

### 3.1 NEURO-SYMBOLIC WORLD MODEL

We instantiate the IALM as a neuro-symbolic world model, composed of a hierarchy of discrete variational-autoencoders (dVAEs) with Gumbel-Softmax (Jang et al., 2017; Maddison et al., 2017) trained end-to-end on transitions  $(o_t, a_t, o_{t+1})$ , explicitly modeling the structural components of a 2DBN (See Fig. 1). Raw observations are encoded into a latent of shape  $(C, F)$ , where  $C$  is the number of discrete observation factors and each has  $F$  categories. These are then mapped to factored state latents  $(N, K)$ , with  $N$  state factors each with  $K$  categories. A local, influence-naïve transition model  $\dot{T}$  maps the current factored state and a one-hot action to the next factored latent state. The transition model uses  $\alpha$ -entmax (Peters et al., 2019) cross-attention over the input state factors to induce exact-zero, data-driven sparsity, exploiting latent parent-child structure.

### 3.2 PERSPECTIVE-SHIFT OPERATOR

Extending Alcedo et al. (2025), we estimate others’ beliefs via an SB-based perspective-shift operator:  $\hat{b}_i = \Phi_{\text{SB}}(b_0, \hat{x}_i, a_{0:n-1}, \bar{T})$ , where  $b_0$  is the egocentric belief,  $\hat{x}_i$  is the locally observed relative state of another agent  $i$ ,  $\bar{T}$  is the learned local influence-naïve transition model (the reference dynamics),  $a_{0:n-1}$  is a proposed action sequence hypothesized for agent  $i$  generated via an upstream

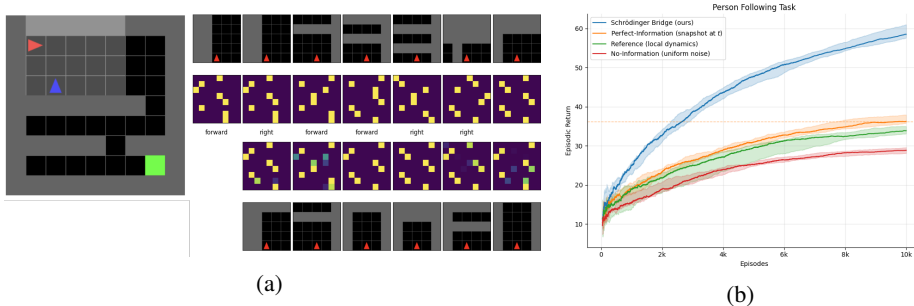


Figure 2: (a) Perspective-shift representations. Left: global state. Right: ground-truth observations (top); factored, discrete latent states (second); actions (third); next-step latent states (fourth); reconstructed observation (last). (b) Median episodic returns over 5 seeds, shaded band denotes 95% bootstrap CI of the median; dashed line: Perfect-Information baseline final median.

local planner, and  $\hat{b}_i$  is the estimated other-centric belief. We implement  $\Phi_{SB}$  by learning the potentials  $(\hat{\phi}_t, \hat{\psi}_t)$  using a bi-directional sequence model (BiGRU) with  $\alpha$ -entmax attention on the action-local-state history  $h_t = \langle x_1, a_1, \dots, a_{t-1}, x_t \rangle$  generated from a roll-out using the reference model  $\hat{T}$ . The sparsity induced by  $\alpha$ -entmax on  $h_t$  serves as the deterministic update to compute the approximate minimal d-set from the action-local-state history  $D_{t+1} = d(h_t)$ . At training time, we sample “epistemic counterfactuals” (other-centric beliefs)  $b_n$  in a single-agent task from the environment to amortize the potentials that realize  $\hat{T}$  via the Doob  $h$ -transform of  $\hat{T}$ , minimizing  $\text{KL}(\mathbb{P}_{\hat{T}} \parallel \mathbb{P}_{\hat{T}})$  subject to the endpoints. At inference, we tilt  $\hat{T}$  with the amortized potentials  $(\hat{\phi}_t, \hat{\psi}_t)$  parametrized by  $c_t := D_{t+1} = d(h_t)$  to reach an estimate for  $\hat{b}_i$ . The learning objective is done at the per-factor level, enabling  $d(h_t)$  to contextually manage the tilt by masking factors. The objective becomes:  $\mathcal{L} = \mathcal{L}_{\hat{O}} + \mathcal{L}_{\hat{T}} + \mathcal{L}_{\Phi}$ . We employ training stabilization techniques that will be discussed and ablated in the full manuscript; these include: Gumbel-softmax  $\tau$  annealing, knowledge distillation confidence-gating, and socially-weighted epistemic counterfactuals objectives at two levels of reasoning.

#### 4 PRELIMINARY RESULTS, LIMITATIONS & NEXT STEPS

We evaluate our approach for influence-aware policy learning in a multi-agent model-based reinforcement learning setting. **Environment:** partially-observable Minigrid (Chevalier-Boisvert et al., 2023), with discrete observations and actions. **Person-following Task:** follow a fixed-policy agent to an unknown target while satisfying social norms (Alcedo et al., 2025). The environment rewards proximity and visibility, gives a large terminal reward for reaching the goal, penalizes collisions, and truncates episodes when the follower stays too far away for too long. **Influence-aware Policy:**  $a_t = \pi_{\text{DQN}}(b_t^0, \hat{b}_t^i)$ . **Baselines:** (i) Perfect-information, true current belief of the other agent with no context/history:  $\hat{b}_t^i = b_t^i$ , (ii) No-information:  $\hat{b}_t^i \sim \text{Uniform}(\Delta_X)$ , and (iii) Reference (local dynamics) multi-step rollout:  $\hat{b}_t^i = \prod^{n-1} \hat{T}(\cdot)$ . **Results:** SB learns faster and achieves higher returns than all baselines (See Fig. 2), highlighting the benefit of decision-time, context-aware mental-state planning over context-free estimates. Currently, our approach assumes homogeneity in the perception and action models of other agents, we aim to address this in future work by adapting the belief-to-belief transform objective via heterogeneous epistemic counterfactuals.

#### 5 RELATED WORK

Latent perspective-taking has been approached via meta-learned agent embeddings (Rabinowitz et al., 2018), latent state-space world models that generate other-centric observations (Chen et al., 2023), and latent opponent/teammate modeling under partial observability (Labash et al., 2020; Papoudakis et al., 2021). In contrast, we cast perspective-taking as belief-to-belief transport, implementing an amortized SB within a learned IALM to produce other-centric beliefs for decision-time mental-state planning; we extend the local-model-only approach of Alcedo et al. (2025). Suau et al.

(2022) instantiate IALMs at scale in multi-agent reinforcement learning settings, with prescribed influence sources and learned predictors, and Suau et al. (2025) learn influence-aware memory and a d-set selector in pixel space (ours operates in latent space). Finally, Wang et al. (2023) proposed the generalized belief transport method for teacher-student learning settings modeled via a static unbalanced entropic optimal transport, whereas we use a Schrödinger Bridge.

#### ACKNOWLEDGMENTS

This work was supported by LARSyS FCT funding (DOI: 10.54499/LA/P/0083/2020, 10.54499/UIDP/50009/2020, and 10.54499/UIDB/50009/2020) and the Portuguese Recovery and Resilience Plan (PRR) through project C645008882-00000055, Center for Responsible AI; and partially supported by ANR (Agence Nationale de la Recherche) AAPG- 2024/Ostensive project; and the project European Learning and Intelligent Systems Excellence (ELIAS, Grant Agreement No. 101120237).

#### REFERENCES

- Kevin Alcedo, Pedro U. Lima, and Rachid Alami. Perspective-shifted neuro-symbolic world models: A framework for socially-aware robot navigation. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, Eindhoven, Netherlands, 2025. Accepted for publication.
- Chris Baker, Rebecca Saxe, and Joshua Tenenbaum. Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the annual meeting of the cognitive science society*, volume 33, 2011.
- Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21:34 – 9, 2011.
- Thomas Bolander, Lasse Dissing, and Nicolai Herrmann. DEL-based Epistemic Planning for Human-Robot Collaboration: Theory and Implementation. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning*, pp. 120–129, 11 2021. doi: 10.24963/kr.2021/12. URL <https://doi.org/10.24963/kr.2021/12>.
- Craig Boutilier, Thomas Dean, and Steve Hanks. Decision-theoretic planning: structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11(1):1–94, July 1999. ISSN 1076-9757.
- Kaiqi Chen, Jing Yu Lim, Kingsley Kuan, and Harold Soh. Latent emission-augmented perspective-taking (leapt) for human-robot interaction. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8006–8013, 2023. doi: 10.1109/IROS55552.2023.10342126.
- Raphaël Chetrite, Paolo Muratore-Ginanneschi, and Kay Schwieger. E. schrödinger’s 1931 paper “on the reversal of the laws of nature”. *European Physical Journal H*, 46:28, 2021. doi: 10.1140/epjh/s13129-021-00032-7.
- Maxime Chevalier-Boisvert, Bolun Dai, Mark Towers, Rodrigo Perez-Vicente, Lucas Willems, Salem Lahlou, Suman Pal, Pablo Samuel Castro, and J Terry. Minigrad & mineworld: Modular & customizable reinforcement learning environments for goal-oriented tasks. In *Advances in Neural Information Processing Systems*, volume 36, pp. 73383–73394. Curran Associates, Inc., 2023.
- J.L. Doob. Conditional brownian motion and the boundary limits of harmonic functions. *Bulletin de la Société Mathématique de France*, 85:431–458, 1957. doi: 10.24033/bsmf.1494. URL <https://www.numdam.org/articles/10.24033/bsmf.1494/>.
- John H. Flavell. *Perspectives on perspective taking.*, pp. 107–139. The Jean Piaget symposium series. Lawrence Erlbaum Associates, Inc, Hillsdale, NJ, US, 1992. ISBN 0-8058-1050-1 (Hardcover).

- Robert Fortet. Résolution d'un système d'équations de M. Schrödinger. *Journal de Mathématiques Pures et Appliquées*, 9e série, 19(1-4):83–105, 1940. URL [https://www.numdam.org/item/JMPA\\_1940\\_9\\_19\\_1-4\\_83\\_0/](https://www.numdam.org/item/JMPA_1940_9_19_1-4_83_0/).
- P.J. Gmytrasiewicz and P. Doshi. Interactive pomdps: properties and preliminary results. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004.*, pp. 1374–1375, 2004.
- Alvin I. Goldman. In defense of the simulation theory. *Mind & Language*, 7:104–119, 1992. doi: <https://doi.org/10.1111/j.1468-0017.1992.tb00200.x>.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, Toulon, France, 2017.
- Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009. ISBN 0262013193.
- Aqeel Labash, Jaan Aru, Tambet Matiisen, Ardi Tampuu, and Raul Vicente. Perspective taking in deep reinforcement learning agents. *Frontiers in Computational Neuroscience*, Volume 14 - 2020, 2020. ISSN 1662-5188. doi: 10.3389/fncom.2020.00069. URL <https://www.frontiersin.org/journals/computational-neuroscience/articles/10.3389/fncom.2020.00069>.
- Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems*, 34(4):1533–1574, 2014. ISSN 1078-0947. doi: 10.3934/dcds.2014.34.1533. URL <https://www.aims sciences.org/article/id/d5bcf817-901d-4104-b7da-eade7847c53e>.
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, Toulon, France, 2017.
- Grégoire Milliez, Matthieu Warnier, Aurélie Clodic, and Rachid Alami. A framework for endowing an interactive robot with reasoning capabilities about perspective-taking and belief management. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication, IEEE RO-MAN 2014, Edinburgh, UK, August 25-29, 2014*, pp. 1103–1109. IEEE, 2014. doi: 10.1109/ROMAN.2014.6926399.
- Martin Mundhenk, Judy Goldsmith, Christopher Lusena, and Eric Allender. Complexity of finite-horizon markov decision process problems. *Journal of the ACM*, 47(4):681–720, July 2000. ISSN 0004-5411. doi: 10.1145/347476.347480. URL <https://doi.org/10.1145/347476.347480>.
- Frans Oliehoek, Stefan Witwicki, and Leslie Kaelbling. A sufficient statistic for influence in structured multiagent environments. *J. Artif. Int. Res.*, 70:789–870, May 2021. ISSN 1076-9757. doi: 10.1613/jair.1.12136.
- Georgios Papoudakis, Filippos Christianos, and Stefano Albrecht. Agent modelling under partial observability for deep reinforcement learning. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 19210–19222. Curran Associates, Inc., 2021. URL [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/a03caec56cd82478bf197475b48c05f9-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/a03caec56cd82478bf197475b48c05f9-Paper.pdf).
- Ben Peters, Vlad Niculae, and André F. T. Martins. Sparse sequence-to-sequence models. In Anna Korhonen, David Traum, and Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1504–1519, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1146. URL <https://aclanthology.org/P19-1146/>.
- Pascal Poupart. *Exploiting structure to efficiently solve large scale partially observable markov decision processes*. PhD thesis, University of Toronto, CAN, 2005. AAINR02727.

- Neil Rabinowitz, Frank Perbet, Francis Song, Chiyuan Zhang, S. M. Ali Eslami, and Matthew Botvinick. Machine theory of mind. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4218–4227. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/rabinowitz18a.html>.
- Guillaume Sarthou, Aurélie Clodic, and Rachid Alami. Ontologenius: A long-term semantic memory for robotic agents. In *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pp. 1–8, 2019. doi: 10.1109/RO-MAN46459.2019.8956305.
- Erwin Schrödinger. Über die umkehrung der naturgesetze. *Sitzungsberichte der Preußischen Akademie der Wissenschaften, Phys.-Math. Klasse*, pp. 144–153, 1931.
- Shashank Shekhar, Anthony Favier, and Rachid Alami. An epistemic human-aware task planner which anticipates human beliefs and decisions. In *Social Robotics: 16th International Conference, ICSR + AI 2024, Odense, Denmark, October 23–26, 2024, Proceedings, Part II*, pp. 208–222, Berlin, Heidelberg, 2025. Springer-Verlag. ISBN 978-981-96-3518-4. doi: 10.1007/978-981-96-3519-1\_20. URL [https://doi.org/10.1007/978-981-96-3519-1\\_20](https://doi.org/10.1007/978-981-96-3519-1_20).
- Miguel Suau, Jinke He, Matthijs T. J. Spaan, and Frans Oliehoek. Influence-augmented local simulators: a scalable solution for fast deep RL in large networked systems. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 20604–20624. PMLR, 17–23 Jul 2022.
- Miguel Suau, Jinke He, Elena Congeduti, Rolf A. N. Starre, Aleksander Czechowski, and Frans A. Oliehoek. Influence-aware memory architectures for deep reinforcement learning in pomdps. *Neural Computing and Applications*, 37(19):13145–13161, Jul 2025. ISSN 1433-3058. doi: 10.1007/s00521-022-07691-7. URL <https://doi.org/10.1007/s00521-022-07691-7>.
- Emanuel Todorov. Linearly-solvable markov decision problems. In B. Schölkopf, J. Platt, and T. Hoffman (eds.), *Advances in Neural Information Processing Systems*, volume 19. MIT Press, 2006. URL [https://proceedings.neurips.cc/paper\\_files/paper/2006/file/d806ca13ca3449af72a1ea5aedbed26a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2006/file/d806ca13ca3449af72a1ea5aedbed26a-Paper.pdf).
- Junqi Wang, Pei Wang, and Patrick Shafto. Generalized belief transport. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA, 2023. Curran Associates Inc.