

Towards a Research Community in Interpretable Reinforcement Learning: the InterpPol Workshop

Hector Kohler

CRISTAL, Université de Lille, Inria, CNRS, Centrale Lille, France

hector.kohler@inria.fr

Quentin Delfosse

AIML Lab, Computer Science Department, Darmstadt, Germany

quentin.delfosse@cs.tu-darmstadt.de

Paul Fester

UKRI Centre for Doctoral Training in AI for Healthcare, Imperial College London, UK

paul.fester19@imperial.ac.uk

Philippe Preux

CRISTAL, Université de Lille, Inria, CNRS, Centrale Lille, France

philippe.preux@inria.fr

Abstract

Embracing the pursuit of intrinsically explainable reinforcement learning raises crucial questions: what distinguishes explainability from interpretability? Should explainable and interpretable agents be developed outside of domains where transparency is imperative? What advantages do interpretable policies offer over neural networks? How can we rigorously define and measure interpretability in policies, without user studies? What reinforcement learning paradigms are the most suited to develop interpretable agents? Can Markov Decision Processes (MDPs) integrate interpretable state representations? In addition to motivate an Interpretable RL community centered around the aforementioned questions, we propose the first venue dedicated to Interpretable RL: the **InterpPol** Workshop¹.

1 Why care about interpretable RL?

Deep reinforcement learning (RL) agents are prone to suffer from a variety of issues that hinder them from learning optimal or generalizable policies. Prominent examples are *reward sparsity* (1) and *difficult credit assignment* (28; 39) or goal misalignment (20), but not at test time. Such misalignments can be difficult to identify (11). A recently popular approach for identifying such shortcut behaviours is via methods from the field of eXplainable AI (XAI), whose aim is to identify the reasons explaining decisions made using a certain model (22; 16; 32; 29; 30; 31). Although XAI methods have helped us better understand, detect and correct many misaligned behaviours, they suffer from issues such as a lack of faithfulness, i.e. the explanation does not faithfully present the model underlying decision process (18; 3; 10), and coarse, low-level explanation semantics. For example, importance-map explanations indicate the importance of an input element for a decision without indicating *why* this element is important (19; 34; 36). (11; 9) show that heatmaps can be particularly misleading when applied to decisions made by RL agents. They argue that, learning rules-based policies defined over extracted intermediate object-centric and relational representations of states, results in intrinsically explainable agents: this allows for detecting and for correcting the previously mentioned problems. Hence one needs to go *beyond explainability* and aim for interpretability (intrinsic explainability).

There have been mainly two approaches to learn interpretable policies (26). First, imitating neural network policies a posteriori with policies (24) such as decision trees or programs (15; 7; 17), which are interpretable, i.e. explainable in nature (14), allows for verifiability but often leads to complex policies (trees having a lot of nodes, hence hard to understand for a human being) as trees have to grow deep to represent neural policies (2; 38). These approaches rely on object extraction method such as FastSAM (41) or SPACE+MOC (8). Second,

¹<https://sites.google.com/view/interppol-workshop/home>




Figure 1: **We need XRL/IRL workshops.** The articles with "interpretable/explainable reinforcement learning" in the title against number of dedicated venues on XRL, using `google scholar`.

recent approaches have directly optimized interpretable policies with reinforcement learning leading to less complex solutions than imitated policies (33; 37) but suffer from stability issues, lack of theoretical grounding, or require inductive bias (21). Furthermore, independent of the approach to learn interpretable policies, these policies need to be defined over semantically meaningful state spaces, e.g not pixels but objects in images in the case of Atari games. This calls for, not only policy learning algorithms, but also for interpretable MDPs learning algorithms in the likes of AlignNet (5).

The biggest challenge in interpretable RL research remains the lack of definitions and a common paradigm. In the general machine learning (ML) literature, there have been attempts at quantifying how interpretable a model is either with metrics such as *simulatability* (time necessary for a human to compute the output of a model) (23), or with user studies (4). But there is no similar work specifically for RL tasks. There is also a lack of tools for comparing interpretable policies from different classes, e.g comparing a tree and a program.

Despite research on interpretable RL being at an early stage, there have been encouraging applications of the latter in fields like healthcare (12; 40). Research is also carried to understand the link between XRL and the influence it has on the final decision maker (27). Interpretability has to be developed with a user-centric approach and raises questions of trust between users and AI. Benchmarks to evaluate interpretable methods are e.g GymDSSAT (13) and OCArari (6).

2 Venues for Interpretable RL

2.1 Existing venues

There is a need for crystallizing both a community and a set of research directions. As shown in Figure 1, there is a lack dedicated venues for people interested in explainable and interpretable RL: even though it is a growing field in terms of published papers, there has been only one dedicated venue ². **The community needs definitions and a formal research paradigm.**

2.2 InterpPol Workshop: an attempt to crystalize an Interpretable RL community

Motivated by the need for research in Interpretable RL described in section 1 and the growing publications on the topic shown in figure 1, we organize the first dedicated workshop on Interpretable RL: the Workshop

²<https://xai4dr1.github.io/>

on Interpretable Policies in Reinforcement Learning, **InterpPol**³. **InterPpol** will take place on August 9 at the Reinforcement Learning Conference in Amherst, MA ⁴.

The purposes of the workshop on Interpretable Policies in RL (**InterpPol**) are to create a community as well as better formalize and crystallize the problem of learning interpretable policies. The core topics of interest at **InterPpol** are (but not limited to):

- *why learn interpretable policies?* What are the interpretability requirements of sequential decision problems, e.g in the health sector? Are there practical advantages of learning interpretable policies such as decision trees rather than neural networks with respect to RL goal alignment? Why looking for intrinsically explainable –interpretable– policies rather than post-hoc explainability?
- *How to define interpretability?* What is an interpretable policy? How to compare different classes of interpretable policies such as trees or programs? How to quantify interpretability of policies with metrics or with (simulated) user studies?
- *How to learn interpretable policies?* What advantages does each learning paradigms (e.g imitation learning, direct reinforcement learning, or evolutionary methods) incorporate? When to favor learning interpretable policy rather than distilling? Do certain interpretable structures prove easier to learn than others, e.g trees against programs? How to link combinatorial optimization and interpretable RL?
- *What sequential problems can be solved with interpretable RL?* How can we refine MDPs to be most conducive to learning interpretable policies, especially in terms of interpretable state space representations? How to learn new models over interpretable state spaces? Do all interpretable MDPs have an interpretable solution? How to use reward shaping to trade-off between interpretability and performance of the learned policy?

We also call for connex topics that could lead to collaborations such as real-world data RL algorithms and benchmarks if they can be adapted to be interpretable.

Please submit your papers on <https://openreview.net/group?id=rl-conference.cc/RLC/2024/Workshop/InterpPol> and do not hesitate to contact us if you can help reviewing. **The submission deadline will be 26th of April and decisions will be released on 20th of May.** Authors from key papers of section 1 will also talk during the workshop and be present at poster sessions as shown in the **InterPpol** schedule ⁵.

2.3 Beyond the workshop

We aim to create an open community on Interpretable RL. We hope to draft a position paper during **InterPpol** that advocate for advancements and standards in the interpretable and explainable field, such as the one done by (35) on representation alignment, or by (25) on deployment. We will establish an open Google Group on Interpretable RL ⁶ to exchange ideas, research, and developments in the field of interpretable reinforcement learning. This group aims to foster collaboration and discussion among researchers, practitioners, and enthusiasts who are keen on advancing our understanding and application of interpretability in reinforcement learning algorithms. Finally, we hope to organize a series of online seminars on the topics of interest in the like of <https://sites.google.com/view/rltheoryseminars>.

If you are interested in Interpretable RL contact us at <mailto:interpPol.workshop@gmail.com>

³<https://sites.google.com/view/interpPol-workshop/home>

⁴<https://rl-conference.cc/index.html>

⁵<https://sites.google.com/view/interpPol-workshop/program?authuser=0>

⁶<https://groups.google.com/g/interpretable1>

References

- [1] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. *Advances in neural information processing systems*, 2017.
- [2] Osbert Bastani, Yewen Pu, and Armando Solar-Lezama. Verifiable reinforcement learning via policy extraction. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2018.
- [3] Chun Sik Chan, Huanqi Kong, and Guanqing Liang. A comparative study of faithfulness metrics for model interpretability methods. In *Conference of the Association for Computational Linguistics (ACL)*, 2022.
- [4] Valerie Chen, Nari Johnson, Nicholay Topin, Gregory Plumb, and Ameeet Talwalkar. Use-case-grounded simulations for explanation evaluation. *Advances in Neural Information Processing Systems*, 2022.
- [5] Antonia Creswell, Kyriacos Nikiforou, Oriol Vinyals, Andre Saraiva, Rishabh Kabra, Loic Matthey, Chris Burgess, Malcolm Reynolds, Richard Tanburn, Marta Garnelo, and Murray Shanahan. Alignnet: Unsupervised entity alignment, 2020.
- [6] Quentin Delfosse, Jannis Blüml, Bjarne Gregori, Sebastian Sztwiertnia, and Kristian Kersting. Ocatari: Object-centric atari 2600 reinforcement learning environments, 2023.
- [7] Quentin Delfosse, Hikaru Shindo, Devendra Singh Dhami, and Kristian Kersting. Interpretable and explainable logical policies via neurally guided symbolic abstraction. *Advances in Neural Information Processing (NeurIPS)*, 2023.
- [8] Quentin Delfosse, Wolfgang Stammer, Thomas Rothenbacher, Dwarak Vittal, and Kristian Kersting. Boosting object representation learning via motion and object continuity. In Danai Koutra, Claudia Plant, Manuel Gomez Rodriguez, Elena Baralis, and Francesco Bonchi (eds.), *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML)*, 2023.
- [9] Quentin Delfosse, Sebastian Sztwiertnia, Wolfgang Stammer, Mark Rothermel, and Kristian Kersting. Interpretable concept bottlenecks to align reinforcement learning agents. *arXiv*, 2024.
- [10] Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. ERASER: A benchmark to evaluate rationalized NLP models. In *Conference of the Association for Computational Linguistics (ACL)*, 2020.
- [11] Lauro Langosco di Langosco, Jack Koch, Lee D. Sharkey, Jacob Pfau, and David Krueger. Goal misgeneralization in deep reinforcement learning. In *International Conference on Machine Learning ICML*, 2022.
- [12] Paul Festor, Giulia Luise, Matthieu Komorowski, and A Aldo Faisal. Enabling risk-aware reinforcement learning for medical interventions through uncertainty decomposition. *arXiv*, 2021.
- [13] Romain Gautron, Emilio J. Padrón, Philippe Preux, Julien Bigot, Odalric-Ambrym Maillard, and David Emukpere. gym-dssat: a crop model turned into a reinforcement learning environment, 2022.
- [14] Claire Glanois, Paul Weng, Matthieu Zimmer, Dong Li, Tianpei Yang, Jianye Hao, and Wulong Liu. A survey on interpretable reinforcement learning, 2022.
- [15] Peter Graf and Patrick Emami. Three pathways to neurosymbolic reinforcement learning with interpretable model and policy networks. *arXiv*, 2024.
- [16] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM Computing Surveys*, 2019.

-
- [17] Jiaming Guo, Rui Zhang, Shaohui Peng, Qi Yi, Xing Hu, Ruizhi Chen, Zidong Du, Ling Li, Qi Guo, Yunji Chen, et al. Efficient symbolic policy learning with differentiable symbolic expression. *Advances in Neural Information Processing Systems*, 2024.
 - [18] Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. A benchmark for interpretability methods in deep neural networks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
 - [19] Subbarao Kambhampati, Sarath Sreedharan, Mudit Verma, Yantian Zha, and L. Guan. Symbols as a lingua franca for bridging human-ai chasm for explainable and advisable ai systems. In *AAAI Conference on Artificial Intelligence*, 2021.
 - [20] Jack Koch, Lauro Langosco, Jacob Pfau, James Le, and Lee Sharkey. Objective robustness in deep reinforcement learning, 2021.
 - [21] Hector Kohler, Riad Akrou, and Philippe Preux. Limits of actor-critic algorithms for decision tree policies learning in ibmdps, 2024.
 - [22] Sebastian Lapuschkin, Stephan Wäldchen, Alexander Binder, Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. Unmasking clever hans predictors and assessing what machines really learn. *Nature communications*, 2019.
 - [23] Zachary Chase Lipton. The mythos of model interpretability. *ArXiv*, 2016.
 - [24] Lirui Luo, Guoxi Zhang, Hongming Xu, Yaodong Yang, Cong Fang, and Qing Li. Insight: End-to-end neuro-symbolic visual reinforcement learning with language explanations. *arXiv*, 2024.
 - [25] Shie Mannor and Aviv Tamar. Towards deployable rl—what’s broken with rl research and a potential fix. *arXiv*, 2023.
 - [26] Stephanie Milani, Nicholay Topin, Manuela Veloso, and Fei Fang. Explainable reinforcement learning: A survey and comparative review. *ACM Computing Surveys*, 2023.
 - [27] Myura Nagendran, Paul Fester, Matthieu Komorowski, Anthony C Gordon, and Aldo A Faisal. Quantifying the impact of ai recommendations with explanations on prescription decision making. *NPJ Digital Medicine*, 2023.
 - [28] David Raposo, Samuel Ritter, Adam Santoro, Greg Wayne, Théophane Weber, Matthew M. Botvinick, H. V. Hasselt, and Francis Song. Synthetic returns for long-term credit assignment. *ArXiv*, 2021.
 - [29] Gabrielle Ras, Ning Xie, Marcel van Gerven, and Derek Doran. Explainable deep learning: A field guide for the uninitiated. *Journal of Artificial Intelligence Research*, 2022.
 - [30] Nicholas A. Roy, Junkyung Kim, and Neil C. Rabinowitz. Explainability via causal self-talk. 2022.
 - [31] Waddah Saeed and Christian W. Omlin. Explainable AI (XAI): A systematic meta-survey of current challenges and future opportunities. *Knowledge-Based Systems*, 2023.
 - [32] Patrick Schramowski, Wolfgang Stammer, Stefano Teso, Anna Brugger, Franziska Herbert, Xiaoting Shao, Hans-Georg Luigs, Anne-Katrin Mahlein, and Kristian Kersting. Making deep neural networks right for the right scientific reasons by interacting with their explanations. *Nature Machine Intelligence*, 2020.
 - [33] Andrew Silva, Matthew Gombolay, Taylor Killian, Ivan Jimenez, and Sung-Hyun Son. Optimization methods for interpretable differentiable decision trees applied to reinforcement learning. In Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*. PMLR, 2020.
 - [34] Wolfgang Stammer, Patrick Schramowski, and Kristian Kersting. Right for the right concept: Revising neuro-symbolic concepts by interacting with their explanations. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2021.

-
- [35] Ilya Sucholutsky, Lukas Muttenthaler, Adrian Weller, Andi Peng, Andreea Bobu, Been Kim, Bradley C Love, Erin Grant, Jascha Achterberg, Joshua B Tenenbaum, et al. Getting aligned on representational alignment. *arXiv*, 2023.
 - [36] Stefano Teso, Öznur Alkan, Wolfgang Stammer, and Elizabeth Daly. Leveraging explanations in interactive machine learning: An overview. *Frontiers in Artificial Intelligence*, 2023.
 - [37] Nicholay Topin, Stephanie Milani, Fei Fang, and Manuela Veloso. Iterative bounding mdps: Learning interpretable policies via non-interpretable methods. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
 - [38] Abhinav Verma, Vijayaraghavan Murali, Rishabh Singh, Pushmeet Kohli, and Swarat Chaudhuri. Programmatically interpretable reinforcement learning. In *International Conference on Machine Learning*. PMLR, 2018.
 - [39] Yue Wu, Yewen Fan, Paul Pu Liang, Amos Azaria, Yuanzhi Li, and Tom M Mitchell. Read and reap the rewards: Learning to play atari with the help of instruction manuals. *Advances in Neural Information Processing Systems*, 2024.
 - [40] Kristine Zhang, Henry Wang, Jianzhun Du, Brian Chu, Aldo Robles Arévalo, Ryan Kindle, Leo Anthony Celi, and Finale Doshi-Velez. An interpretable rl framework for pre-deployment modeling in icu hypotension management. *npj Digital Medicine*, 2022.
 - [41] Xu Zhao, Wenchao Ding, Yongqi An, Yinglong Du, Tao Yu, Min Li, Ming Tang, and Jinqiao Wang. Fast segment anything, 2023.