universität wien

Research Group
Data Mining and Machine Learning

Sebastian Tschiatschek

July 5, 2023 // ASAI Networking Event 2023

6 working groups

Database Techniques for Data Mining

Christian Böhm

Machine Learning with Graphs

Nils Kriege

Data Mining

Claudia Plant

Natural Language Processing

Benjamin Roth

Probabilistic and Interactive Machine Learning

Sebastian Tschiatschek

Scalable Algorithms for Graph Mining

Yllka Velaj

Working group: Probabilistic and Interactive Machine Learning

3 key research "directions"

- Reinforcement Learning
  - Reward / constraint inference
  - Exploration & abstraction
- Interactive machine learning
- Probabilistic (Generative) Models

The group




Motivation

[C. Morrison et al., CHI'21]

Key goals and challenges

Intelligent agents enabling
efficient & seamless collaboration

Challenge: Collaboration in the face of

- (significant) mismatch in inputs,
- (initially) non-aligned goals and constraints,
- (complex) large state spaces

Key goals and challenges

Intelligent agents enabling
efficient & seamless collaboration

Challenge: Collaboration in the face of

- (significant) mismatch in inputs,
- (initially) non-aligned goals and constraints,
- (complex) large state spaces

Teaching and learning
desired and undesired behavior

Reinforcement learning in a nutshell

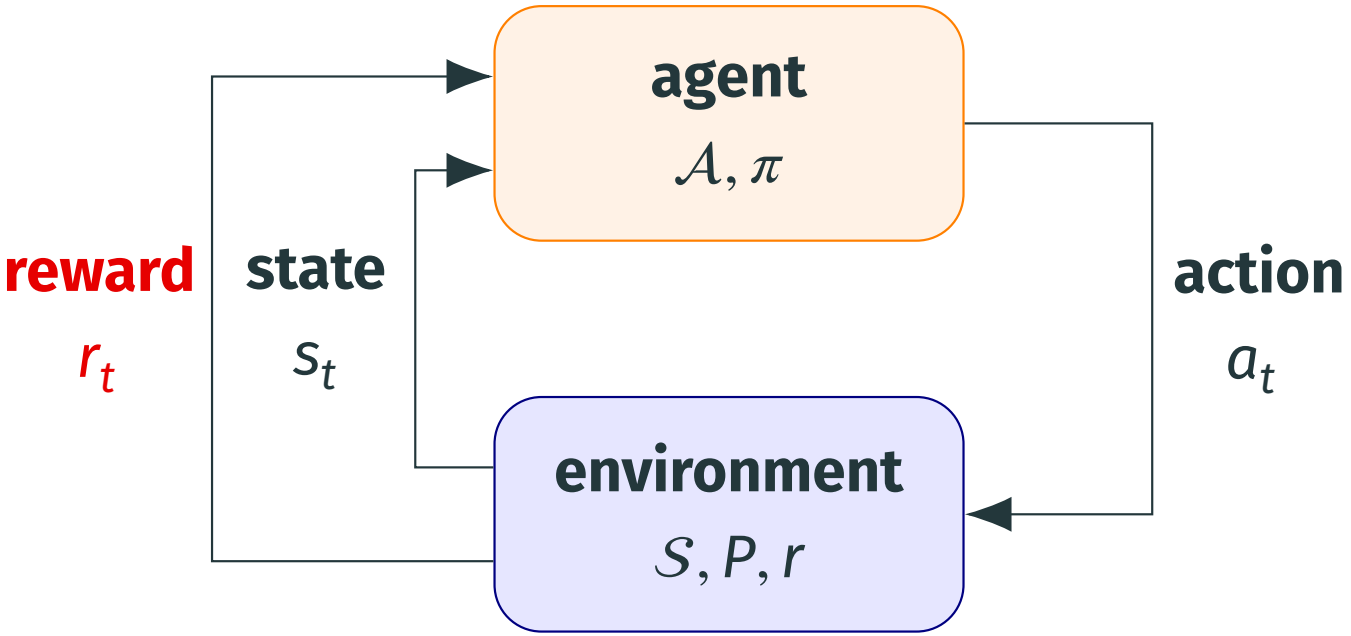

Past and future work

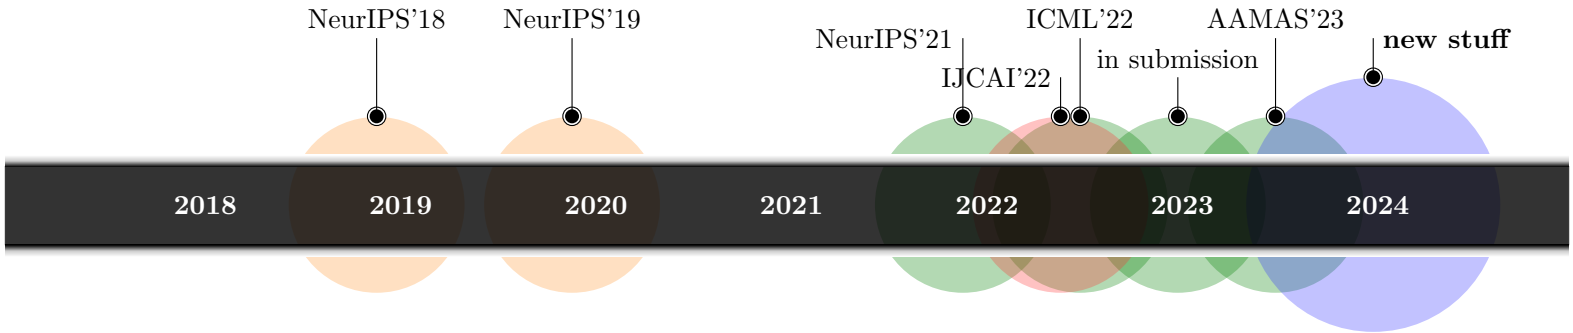

Teaching and learning under mismatch

- Reward linear in expert's features: $r.s/ = \langle \phi^{\mathsf{E}}, w^{\mathsf{i}} \rangle$
- *NeurIPS'18*: Learner's ϕ^{L} and expert's features ϕ^{E} differ
- *NeurIPS'19*: Constraints/prefs. on feature expectations

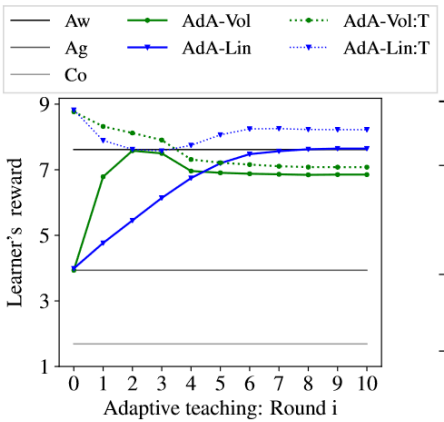

(a) Reward over teaching rounds

Teacher \ Env	10×10	15×15	20×20
AWARE-CMDP	7.62 ± 0.02	7.44 ± 0.04	7.19 ± 0.04
AGNOSTIC	3.94 ± 0.09	3.84 ± 0.06	3.95 ± 0.06
CONSERV	1.68 ± 0.01	1.67 ± 0.012	1.62 ± 0.02
ADAWARE-VOL (3^{rd})	7.50 ± 0.14	7.50 ± 0.04	7.29 ± 0.05
ADAWARE-VOL (end)	6.85 ± 0.33	7.06 ± 0.06	6.77 ± 0.08
ADAWARE-LIN (3^{rd})	6.14 ± 0.08	6.28 ± 0.10	6.37 ± 0.08
ADAWARE-LIN (end)	7.64 ± 0.02	7.53 ± 0.03	7.29 ± 0.06

(b) Varying grid-size

Past and future work

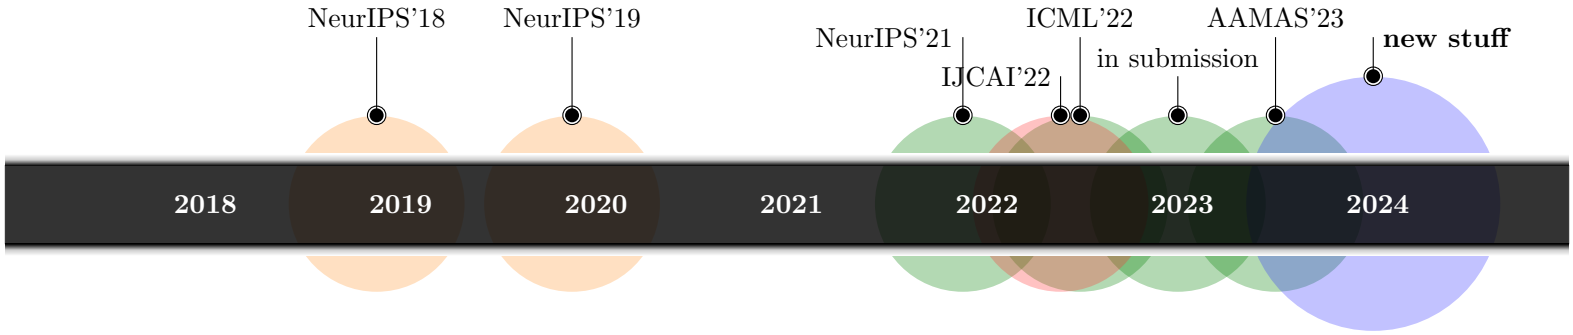

Effectively learning about rewards & constraints

- *NeurIPS'21/ICML'22/AAMAS'23/in submission*: Different FB types / constraints / information directed learning

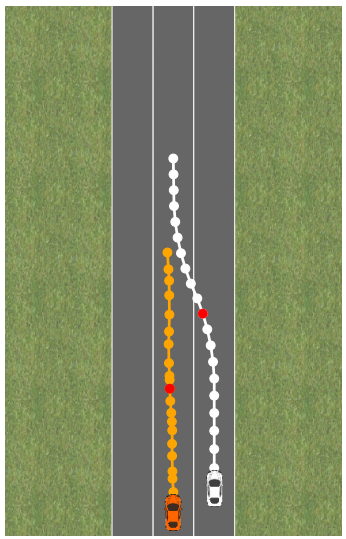

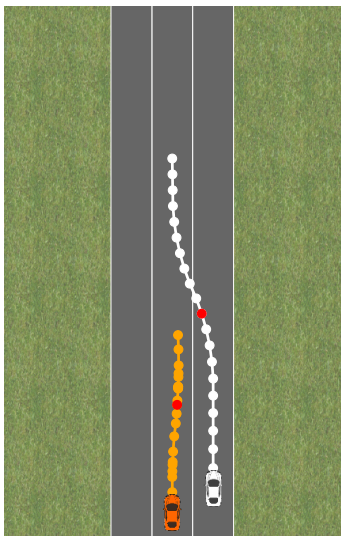

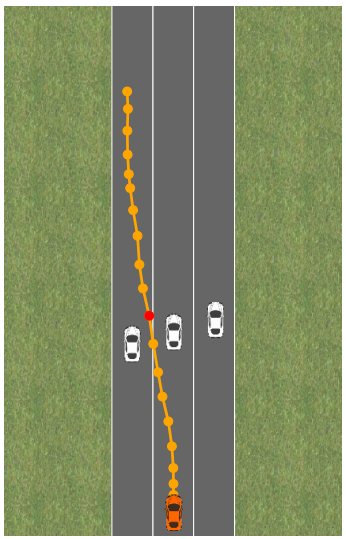

Past and future work

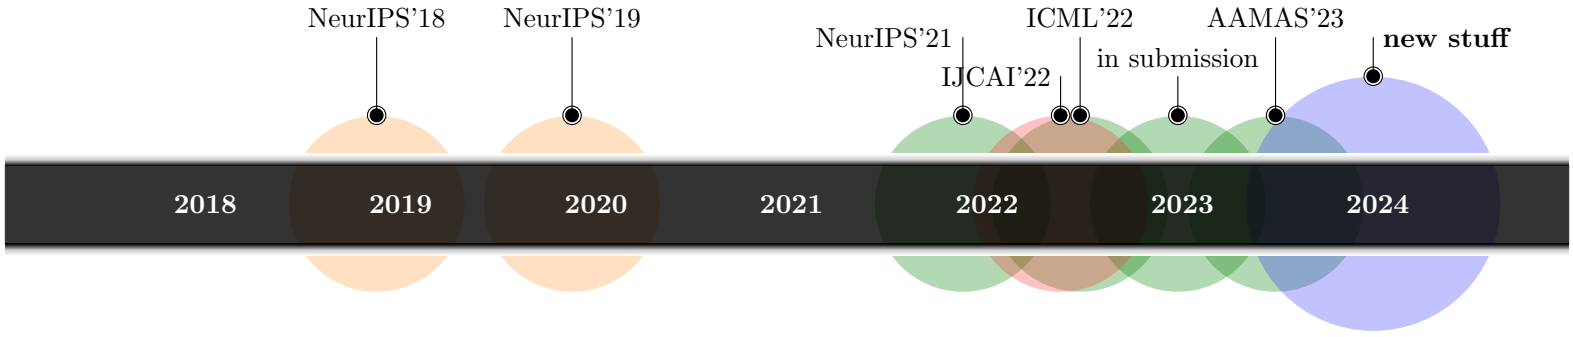

Understanding large state spaces

- *IJCAI'22 / ongoing*: Abstractions of large state spaces

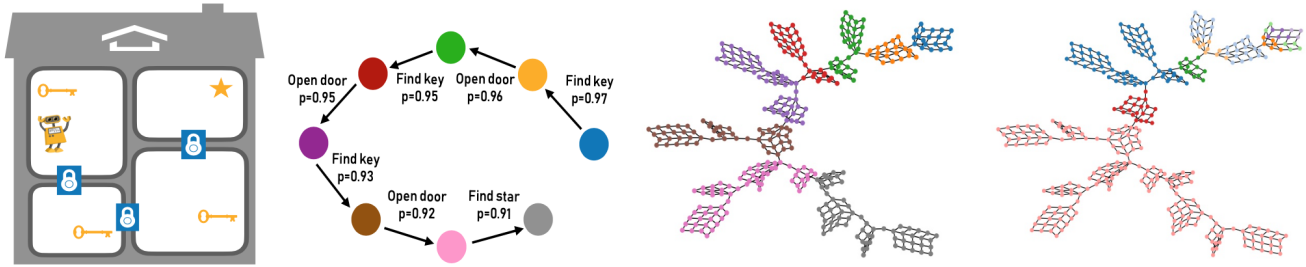

(a) Object-Rooms (b) Abstract ψ-SMDP (c) Successor Homomorphism (d) Q (all): Q^*-irrelevance Abstraction

Figure 1: MDP abstraction in the Object-Rooms domain. (b) Abstract ψ-SMDP induced by our successor homomorphism from the ground MDP as shown in (c), the abstract states in (b) correspond to aggregated ground states of the same colour in (c). (d) Abstraction induced by approximate Q^*-irrelevance abstraction (cf. Appendix A.4) for *find key*; the abstraction does not carry temporal semantics, and is not reusable for other tasks e.g., *find star*.. Another example can be found in Figure 10 in Appendix A.6 and more details are in the experiments section.

Past and future work

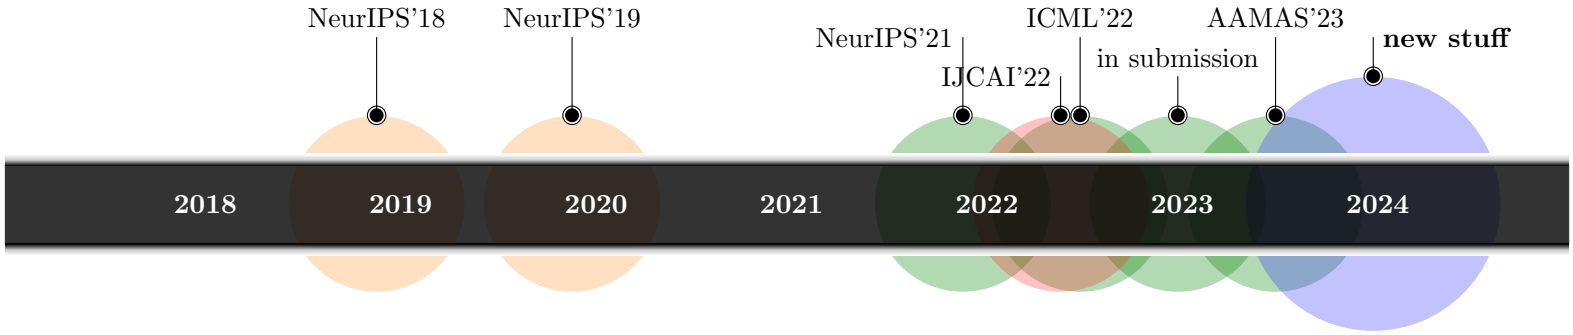

What's next/ongoing—collaborations welcome

- Active 3rd person imititation learning
- Deeper understanding about learning from stop-feedback
- Abstractions for efficient exploration
- Policy design for continual reinforcement learning

[Photo credit: vog.photo]

References

- Luis Haug, Sebastian Tschiatschek, Adish Singla, *Teaching inverse reinforcement learners via features and demonstrations*, NeurIPS'18
- Sebastian Tschiatschek, Ahana Gosh, Luis Haug, Rati Devidze, Adish Singla, *Learner-aware Teaching: Inverse Reinforcement Learning with Preferences and Constraints*, NeurIPS'19
- Cecily Morrison, Edward Cutrell, Martin Grayson, Anja Thieme, Alex Taylor, Geert Roumen, Camilla Longden, Sebastian Tschiatschek, Rita Faia Marques, Abigail Sellen, *Social Sensemaking with AI: Designing an Open-ended AI experience with a Blind Child*, CHI'21
- David Lindner, Matteo Turchetta, Sebastian Tschiatschek, Kamil Ciosek, Andreas Krause, *Information Directed Reward Learning for Reinforcement Learning*, NeurIPS'21
- Han Dongge, Sebastian Tschiatschek, *Option Transfer and SMDP Abstraction with Successor Features*, IJCAI'22
- David Linder, Sebastian Tschiatschek, Katja Hofmann, Andreas Krause *Interactively Learning Preference Constraints in Linear Bandits*, ICML'22
- Sebastian Tschiatschek, Maria Knobelsdorf, Adish Singla, *Equity and Fairness of Bayesian Knowledge Tracing*, EDM'22
- Silvia Poletti, Alberto Testolin, Sebastian Tschiatschek, *Learning Constraints From Human Stop-Feedback in Reinforcement Learning*, AAMAS'23