

Gabriel Paludo Licks

Ph.D. student

Education

- 2020/2–current **Ph.D. in Engineering in Computer Science**, *Sapienza University of Rome*, Italy.
2018/2–2020/2 **M.Sc. in Computer Science**, *PUCRS*, Porto Alegre, Brazil.
2014/1–2018/2 **B.Sc. in Computer Science**, *UPF*, Passo Fundo, Brazil.

Ph.D. thesis (*proposal*)

- Title *Sample-efficient approximation of average-reward reinforcement learning*
Advisor Prof. Dr. Giuseppe De Giacomo
Description Sample-efficient learning and planning in the average-reward reinforcement learning setting, optimizing agent behavior in infinite-horizon tasks.

M.Sc. thesis (*with honors*)

- Title *Automated database indexing using model-free reinforcement learning*
Advisor Prof. Dr. Felipe Meneguzzi
Description Learning the value of indexes in a database and policies that adjust the index configuration to maintain optimal performance in dynamic workloads.

Academic experience

- 2019/1–2020/2 **Fellowship project with SAP (during M.Sc.)**, *PUCRS/SAP*, Porto Alegre, Brazil.
Automated Planning for Optimizing the Deployment of Data Pipelines.
2018/2 **Fellowship project with SAP (during M.Sc.)**, *PUCRS/SAP*, Porto Alegre, Brazil.
Automated Database Indexing for Dynamic Workloads using Reinforcement Learning.
2019/2 **AI teaching assistant (during M.Sc.)**, *PUCRS*, Porto Alegre, Brazil.
Artificial Intelligence (Undergraduate course), Prof. Dr. Felipe Meneguzzi.
2016/2–2017/1 **B.Sc. exchange student in Computer Science**, *UHasselt*, Hasselt, Belgium.
Taking courses as an exchange student via university bilateral agreement.
2014/2–2018/1 **Undergraduate research assistant**, *UPF*, Passo Fundo, Brazil.
Working on collecting data and assisting M.Sc. and Ph.D. students in their research. The papers I participated as an undergrad research fellow are written in Portuguese, thus I do not list them in this CV.

Other academic activities

- 2016 **Student representation.**
(during B.Sc.) *Vice-president of the Informatics Academic Directory at the University of Passo Fundo.*

Publications

Conference papers

- 2020 **Using Self-Attention LSTMs to Enhance Observations in Goal Recognition.** Leonardo Amado, Gabriel Paludo Licks, Matheus Marcon, Ramon Fraga Pereira, and Felipe Meneguzzi. *The International Joint Conference on Neural Networks (IJCNN 2020)*

Workshops

- 2020 **Automated Database Indexing Using Model-Free Reinforcement Learning.** Gabriel Paludo Licks and Felipe Meneguzzi. *The ICAPS Scheduling and Planning Applications workShop (ICAPS SPARK 2020)*

Demos

- 2020 **LatRec+: Learning-based Goal Recognition in Latent Space.** Leonardo Rosa Amado, João Paulo Aires, Ramon Fraga Pereira, Maurício Magnaguagno, Roger Granada, Gabriel Paludo Licks, Matheus Marcon, and Felipe Meneguzzi. *The AAAI Workshop on Plan, Activity, and Intent Recognition (AAAI PAIR 2020)*
- 2019 **LatRec: Recognizing Goals in Latent Space.** Leonardo Rosa Amado, Ramon Fraga Pereira, João Paulo Aires, Maurício Magnaguagno, Roger Granada, Gabriel Paludo Licks, and Felipe Meneguzzi. *The 29th International Conference on Planning and Scheduling (ICAPS 2019)*

Journal papers

- 2019 **SmartIX: A Database Indexing Agent Based on Reinforcement Learning.** Gabriel Paludo Licks, Julia Colleoni Couto, Priscilla de Fátima Míche, Renata de Paris, Duncan Dubugras Ruiz, and Felipe Meneguzzi. *The International Journal of Research on Intelligent Systems for Real Life Complex Problems (Applied Intelligence – APIN)*

Research interests

My general research interests are in learning and planning for sequential decision making. Currently, I'm focusing on average-reward reinforcement learning, which is a problem formulation suitable for continuing tasks, i.e. infinite-horizon MDPs. Specifically, during my Ph.D., my interests are in integrating learning and planning in the average-reward setting, especially in non-Markovian environments. This relates to efficiently learning models online from samples in a replay memory, and extending algorithms such as Dyna-Q and n-step bootstrapping techniques. Such techniques, inspired by Monte Carlo methods, are able to plan and compute estimates a few steps ahead before making a decision, and hypothetically should be able to outperform one-step lookahead techniques in non-Markovian environments.

Skills

- Preprocessing data and training ML models (NNs and RNNs in PyTorch and Keras).
- Implementing (deep) RL agents and modeling environments and reward functions.
- Writing planning domains (PDDL).