

TAMAS MADARASZ

JPMorgan Chase

@ tamasmadarasz1@gmail.com
in linkedin.com/in/tamas-madarasz-040ba91a5

+44-740-3553-719

London, UK., EU citizen

www.tasmadarasz.com

EXPERIENCE

AI/ML Lead (Vice President)

JPMorgan Chase

October 2022 – ongoing

I oversee the International Consumer Bank's **NLP** initiatives as team lead, training and deploying language models for **customer experience**, **voice of customer**, and **business analytics**. The team has deployed and is developing models for **topic and sentiment analysis**, **FAQs**, **digital assistants**, **quality assurance and summarisation** by **fine-tuning LLMs** and building custom architectures and pipelines.

I also work with leadership to support business-critical decisions such as setting the bank's interest and reward offerings with the help of **forecasting**, **optimization** and **reinforcement learning** models of customer and competitor behaviours.

Driven by these use cases, I also developed an agenda for **fundamental research** with the team, using a combination of **multi-agent reinforcement learning** and **NLP** for improved veracity of language models, as well as privacy in LLMs.

Staff Research Scientist/Deep Learning Researcher

MediaTek Research

December 2021 – October 2022

At MediaTek I led applied research projects focused on **reinforcement learning for chip design**, and fundamental research in the domains of responsible decision making and continual reinforcement learning.

- Used hierarchical RL and quality diversity to develop a chip design paradigm that can interact with human designers to enable better designs.
- Introduced a novel method for efficient planning and transferring knowledge between tasks by learning the positional invariances in event sequences[1].

Machine Learning Researcher

Huawei R&D UK

December 2020 – December 2021

At Huawei R&D I was part of the flagship AI Theory team in Noah's Ark lab, working on fundamental research as well as applications in areas like compiler optimization and self-driving technologies.

- Developed solutions for transfer learning problems in the domains of **combinatorial optimization** and **planning**.
 - Participated in developing a new compositional transfer learning algorithm for **computer vision** [2].
-

AI/ML fellow

GSK

June 2020 – December 2020

During my fellowship, I worked on reinforcement learning approaches to assist scientists in the optimization of new medicines, by learning to extract and supplement human expertise using databases from past **drug discovery** cycles.

Postdoctoral Researcher

University of Oxford

March 2018 – May 2020

Developed sample-efficient deep reinforcement learning algorithms that quickly adapt to solve new tasks in a **continual** and **multi-task learning** setting[4], and by decomposing complex task structures into subtasks[3].

Postdoctoral Researcher

University of Geneva

October 2015 – December 2017

Researched reinforcement learning and planning algorithms for partially observable decision problems and developed a dynamical systems model of representation learning for sensory systems in the brain[6].

EDUCATION

Ph.D. in Neuroscience

New York University

2015

B.A. in Mathematics

University of Cambridge

Studies and work experience as a classical cellist

M.A. @ Robert-Schumann-Academy, Dusseldorf & Ecole Normale de Musique de Paris

BACKGROUND AND SKILLS

- Programming: Python Pandas Matlab C++
- DL frameworks: Tensorflow Pytorch Keras Theano
- ML areas: Reinforcement Learning and Planning NLP Deep Learning Probabilistic Modelling
Bayesian Statistics Computer Vision , Causal Inference

SELECTED RESEARCH PROJECTS

- Some of the open challenges preventing RL agents from autonomously and safely navigating real-world problems are the understanding of the nonstationarities and feedback loops their decisions can create, the ability to break down complex tasks into manageable subtasks, and the potential to transfer knowledge from one task/sub-task to another. In a series of papers and ongoing work I have contributed new algorithmic tools to decompose and reason about the structure of complex, non-Markovian reinforcement learning problems, while also allowing efficient planning in new settings. By representing temporal and causal invariances in the reward functions of such tasks, and a savvy re-evaluation of past experience given the current problem setting, this approach allows fast generalization of to new tasks within a task family.
- Published at the Workshop on Responsible Decision Making in Dynamic Environments at **ICML 2022** and at the Causal Learning and Decision Making Workshop at **ICLR 2020**.

-
- An important problem when transferring knowledge in the RL setting is to know how to act when our goals or the available rewards in the environment change. But how do we know when these changes happen, which policies to reuse, how to quickly adapt them, and how to discover newly available rewards? The Bayesian Successor Representation algorithm gives some answers to these questions by continually evaluating task similarity online, clustering similar tasks, and quickly approximating optimal policies for newly emerging reward functions. By combining nonparametric clustering methods with flexible deep RL, it provides an efficient approach to quickly solve new tasks as rewards change, and outperforms competing approaches on continual learning problems.
 - Awarded a **spotlight talk** at **NeurIPS 2019**.

-
- Given limited data, we can construct many different predictive models of our environments, with varying statistical complexity. Some of these will focus on a small number of strong predictive relationships, while others will encompass a broader range of associations. During my PhD, I formulated a computational framework explaining how biological learning deals with ambiguity in the predictive relationships in the environment, through weighing different plausible explanatory models.
 - Published as an article in **Nature Neuroscience**.

AWARDS

- Samuel J. and Joan B. Williamson Dissertation Fellowship
- NYU Dean's Dissertation Fellowship
- MacCracken Graduate Fellowship
- Cambridge Overseas Trust and Trinity College full undergraduate scholarship
- NeurIPS, RLDM and Cosyne travel awards

SERVICE

- Ad hoc reviewer: Science
Nature Neuroscience Biological Cybernetics
IBM Journal of Research and Development
- **Mentorship:**
 - David Ireland, PhD Intern, MediaTek Research
 - Antonin Vidon, Intern, Huawei R&D UK
 - Frank Catuela, Omar Akhand, Samit Roy Undergraduate Researchers, NYU.

PUBLICATIONS

Liu R*, Madarasz TJ* (2023)

Planning into the fog of uncertainty: when to observe in partially observable planning tasks.

Under review

[1] Madarasz TJ (2022)

LPI: Learned Positional Invariances for Transfer of Task Structure and Zero-shot Planning.

ICML, 39th International Conference on Machine Learning, Workshop on Responsible Decision Making in Dynamic Environments.

[2] Parisot S, Esperanca PM, McDonagh S, Madarasz TJ, Yang Y, Li Z (2022)

Long-tail Recognition via Compositional Knowledge Transfer.

CVPR, 2022 IEEE Conference on Computer Vision and Pattern Recognition .

[3] Madarasz TJ, Behrens TEJ (2020)

Learning transferable task schemas by representing causal invariances.

ICLR, Eighth International Conference on Learning Representations, Causal learning for decision making workshop.

[4] Madarasz TJ, Behrens TEJ (2019) Better transfer learning with inferred successor maps.

NeurIPS, 33rd Conference on Neural Information Processing Systems, Vancouver, Canada.

Spotlight oral presentation (<3% of submissions).

[5] Madarasz TJ, Behrens TEJ (2019) Inferred predictive maps in the hippocampus for better transfer learning. **RLDM,**

Multidisciplinary Conference on Reinforcement Learning and Decision Making, Montreal.

[6] Yamada Y*, Bhaukaurally K*, Madarasz TJ, Pouget A, Rodriguez I, Carleton A (2017)

Context- and output layer-dependent long-term ensemble plasticity in a sensory circuit. *Neuron, Volume 93, Issue 5,*

1198 - 1212.

[7] Madarasz TJ, Diaz-Mataix L, Akhand O, Ycu EA, LeDoux, JE, Johansen JP (2016)

Evaluation of ambiguous associations in the amygdala by learning the structure of the environment. *Nature Neuroscience*

19, 965-972.

[8] Madarasz TJ, LeDoux JE, Johansen JP (2015) Evaluating predictive variables by a dual system of structure and parameter learning. **RLDM,**

Multidisciplinary Conference on Reinforcement Learning and Decision Making, Edmonton.

CONFERENCE PRESENTATIONS

Madarasz, TJ, Behrens TEJ (2019) Flickering hope? Inferred hippocampal maps and splitter cells support multi-task learning COSYNE: *Computational and Systems Neuroscience, Lisbon.*

Fink AE, Madarasz TJ, LeDoux JE (2015) Short-term plasticity as a homeostatic mechanism in the lateral amygdala. *Society for Neuroscience.*

Madarasz TJ, Diaz-Mataix L, Akhand O, LeDoux JE, Johansen JP (2015) Evaluating ambiguous associations in the amygdala by learning the structure of the environment. COSYNE: *Computational and Systems Neuroscience, Salt Lake City, Utah.*

Madarasz TJ, Johansen JP, LeDoux JE (2013) Causality and its neural underpinnings in active and passive aversive learning. *Society for Neuroscience.*

Madarasz TJ, Diaz-Mataix L, Boyden SE, LeDoux JE, Johansen JP (2012) Temporally specific optogenetic inactivation of lateral amygdala pyramidal neurons reverses the effects of contingency degradation on fear learning. *Society for Neuroscience.*

Madarasz TJ, Roy SS, Boyden ES, LeDoux JE, Johansen JP (2011) Making predictions in a complex world: mechanisms of contingency degradation in fear conditioning. *Society for Neuroscience.*

Gervan P, Berencsi A, Madarasz TJ, Kovacs I (2010) Development and plasticity of primary visual and motor function in humans. *II. Dubrovnik Conference on Cognitive Science. x*

INTERNSHIPS

2012, 2013 RIKEN Brain Science Institute Summer Intern