

# Report on the outcomes of a Short-Term Scientific Mission<sup>1</sup>

Action number: CA22145

Grantee name: Achille MORENVILLE

## **Details of the STSM**

Title: **Exploring Approaches in Reinforcement Learning for Imperfect-Information Games**

Start and end date: 8/04/2024 to 31/05/2024

## **Description of the work carried out during the STSM**

The primary goal of the STSM was to visit the Department of Advanced Computing Sciences at Maastricht University, and engage in a scientific collaboration on Game AI and Reinforcement Learning (RL) with Dr. Dennis Soemers and Prof. Mark Winands.

The collaboration focused on two key objectives:

1. Developing a new representation specifically designed for imperfect-information games within the Ludii framework was a key focus of the STSM. The current extensive-form representation lacks functionalities such as public and private information sets, and cannot encode the temporal aspect of observation in games, making it impossible to differentiate between different histories that lead to the same state. To overcome these limitations, we investigated the feasibility of incorporating public and private information sets and belief states into the Ludii framework, while ensuring that the representation remains free of human knowledge. In addition, we explored alternative formalisms such as partially observable stochastic games, which prioritize observations over states, to determine whether they are suitable for representing imperfect-information games within Ludii.

Our research has shown that a hybrid approach of search and learning is likely to be essential for strong performance in imperfect information games. While search-only methods have proven successful in perfect information games, the added complexity of imperfect information significantly reduces their effectiveness. To facilitate our research and to better determine the necessary modifications to Ludii, we decided to start by developing a prototype.

---

<sup>1</sup>This report is submitted by the grantee to the Action MC for approval and for claiming payment of the awarded grant. The Grant Awarding Coordinator coordinates the evaluation of this report on behalf of the Action MC and instructs the GH for payment of the Grant.

This prototype is designed to encompass games in general, including imperfect-information games, and will allow for easy integration of research done on agents for imperfect-information games. This will help us identify areas in Ludii that need to be reworked.

2. Enhancing the dataset used in a future competition focused on characterizing the effectiveness of Monte-Carlo Tree Search (MCTS) variants in various game types. The collaboration with Dr. Soemers allowed to expand the dataset and ensure it accurately reflects diverse gameplay scenarios. As the competition is still under development and won't be open for another few months, I can't go into more detail.

Furthermore, I took Dr. Soemers' Reinforcement Learning (RL) course to deepen my understanding of reinforcement learning algorithms and concepts, which will be a crucial part of my future research. The course provided an introduction to RL, combining lectures with practical assignments that covered a broad spectrum of RL topics. These assignments provided hands-on experience that is directly applicable to my research.

The course also covered advanced topics essential to my research, most notably Multi-Agent RL, Transfer Learning in RL, and Deep RL. Multi-Agent RL involves techniques and strategies for training agents in multi-agent environments, which are essential for games where multiple players interact. Transfer Learning in RL includes methods for transferring knowledge from one domain to another, crucial for developing agents capable of adapting to different games. Deep RL applies neural networks to reinforcement learning problems to handle high-dimensional state spaces.

Specifically, the insights gained on Multi-Agent RL will aid in designing agents that can handle the complexities of imperfect-information games. Additionally, the transfer learning techniques will help in creating a versatile agent capable of playing a variety of games without requiring extensive retraining. The knowledge gained from this course will directly benefit my research and will undoubtedly accelerate my work on games with imperfect information.

### **Description of the STSM main achievements and planned follow-up activities**

The STSM achieved its expected outcomes, and we made significant progress in both key objectives:

1. We initiated the development of a prototype for a new representation within Ludii that aims to incorporate public and private information sets and belief states. This will enable the analysis of a broader range of games within Ludii and significantly contribute to research in the field of imperfect-information games. It will also help to define the requirements necessary to incorporate learning in the model. We plan to continue refining and testing the prototype, with the ultimate goal of fully integrating it into Ludii.
2. The dataset for the competition was substantially enriched, providing a valuable future resource for the AI community and fostering further research. We plan to launch the competition in a few months and analyze the results to gain further insights into the effectiveness of MCTS variants.

In addition, the collaborative environment during the STSM laid the groundwork for a research paper that will showcase the exploration of novel game representations and their integration within Ludii. We plan to continue our collaboration and submit the paper as soon as possible, with the aim of contributing to knowledge dissemination and inspiring further research in the field.

The STSM also established valuable connections for the grantee within the game AI community and paved the way for further collaborative research. We plan to continue our collaboration, with the aim of further advancing the state of the art in imperfect-information games and reinforcement learning.