

# SPECIES Scholarship Host Application

Dennis Wilson  
dennis.wilson@isae-supero.fr

## Topic

Evolutionary algorithms have recently demonstrated impressive capabilities as optimization methods for finding policies in reinforcement learning (RL). As opposed to gradient-based methods, these approaches simulate natural selection on a population of competing policies to optimize an agent's behavior. This can be performed with various policy representations, such as function graphs [1] and deep neural networks [2]. The results from evolutionary methods are competitive with state-of-the-art RL methods such as Deep Q-Networks [3] and outperform these methods on many benchmark tasks.

There are many open questions for evolutionary RL methods. While gradient-based methods are required to optimize differentiable functions, evolutionary methods can optimize over any function, or multiple functions, and the development of optimal objective functions for evolutionary RL is an open question. Evolutionary methods also maintain a population of different policies and optimal population management strategies, such as island-based models or diversity metrics, are also open for exploration in the context of RL. Finally, many ideas from traditional reinforcement learning, such as exploration-exploitation trade-off, have yet to be fully considered in the evolutionary context. These could greatly improve sample efficiency, one of the main drawbacks of current evolutionary RL.

The goal of this project is to improve the state-of-the-art of evolutionary reinforcement learning. As there are many possible directions to explore in improving evolutionary RL, the precise topic definition will be determined in discussion with the candidate. The ideal candidate would have experience with evolutionary computation and/or RL and would be motivated to advance this topic; experience with the Julia programming language is a bonus.

## References

- [1] Wilson, Dennis G., et al. "Evolving simple programs for playing Atari games." Proceedings of the Genetic and Evolutionary Computation Conference. ACM, 2018.
- [2] Such, Felipe Petroski, et al. "Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning." arXiv preprint arXiv:1712.06567 (2017).
- [3] Mnih, Volodymyr, et al. "Human-level control through deep reinforcement learning." Nature 518.7540 (2015): 529.

## **Research Group and Host Institution**

This internship will take place at ISAE-SUPAERO in Toulouse, France, in the SuReLI research team inside the Department of Complex Systems Engineering. The ISAE-SUPAERO Reinforcement Learning Initiative (SuReLI) is a vibrant group of researchers thriving to design next generation AI. We maintain a constructive, lively environment in a human-sized team that range from undergrad students to permanent academic staff, focussed on understanding AI and developing solutions that contribute positive impact on science and society. We work on state of the art algorithms, study their properties mathematically and empirically, and apply them to challenging problems to bridge a gap between theory and real applications. We focus on Reinforcement Learning for sequential decision making problems and its links with other disciplines. Our applications include: autonomous vehicles, control of optimization processes, rehabilitation exoskeleton control, Atari games, aircraft landing scheduling, satellite resources planning, sailboat and UAV planning and control.

ISAE-SUPAERO is one of the top engineering schools in France and a worldwide leader in aerospace engineering. It is a part of the University of Toulouse and there will be a strong link in this internship with other research groups in Toulouse, notably the REVA Artificial Life team at IRIT. ISAE-SUPAERO is also member of the new Artificial and Natural Intelligence Toulouse Institute. ANITI has been selected as one of four institutes spearheading research on AI in France and aims to make Toulouse one of the world leaders in Artificial Intelligence in research, education, innovation and economic development. ISAE-SUPAERO has recently advanced its research and educational focuses in Artificial Intelligence with the growth of the Decision Systems research group and Science of Decision and Data Master's program.

Over the course of this internship, the student will be able to profit from the lively environment surrounding AI in Toulouse, attending SuReLI, ANITI, REVA, and TidDL (the Toulouse Interdisciplinary Deep Learning Group) events. These range from formal workshops and seminars to informal events like the regular TidDL DeepBeers. Depending on the internship timing, the student could attend the upcoming Reinforcement Learning Summer School in Toulouse and Julia4AI Workshop. The candidate will be surrounded by an international and talented group of students and researchers working on various aspects of evolutionary computation, reinforcement learning, and AI in general.

The student will also be encouraged to explore Toulouse and its surroundings. Once the capital of the Visigoths and now the capital of the air and space industry, Toulouse offers an excellent way to experience France. It is home to and surrounded by UNESCO World Heritage Sites such as the Basilica of St. Sernin, Carcassonne, and Albi. It is frequently ranked as one of the best student cities and best cities for foreigners. For this internship, we can offer housing on the campus of ISAE-SUPAERO, which is along the beautiful Canal du Midi and easily accessible by bike, metro, or bus.

## **Advisor**

Dennis G. Wilson is an Associate Professor of AI and Data Science at ISAE-SUPAERO in Toulouse, France. He obtained his PhD at the Institut de Recherche en Informatique de Toulouse (IRIT) on the evolution of design principles for artificial neural networks. Prior to

that, he worked in the Anyscale Learning For All group in CSAIL, MIT, applying evolutionary strategies and developmental models to the problem of wind farm layout optimization. He is a former chair of the Emergent Researchers in Artificial Life group and is currently co-chair of the Complex Systems track at GECCO.

His research focuses on genetic programming, neural networks, and the evolution of learning. He has published in multiple A rank conferences and has co-authored two book chapters on developmental neural networks. He previously organized a competition on wind farm layout optimization at GECCO, which cumulated in a journal publication on state-of-the-art wind farm optimization methods. He currently organizes the Developmental Neural Networks Workshop, which will take place at ALIFE 2020 for its 3<sup>rd</sup> edition. He is passionate about environmental applications and is an active member of the Climate Change AI group and the co-founder of Nautilia Computing, which provides AI for understanding ocean ecosystems.