

Florian Strub

Senior Research Scientist

180 rue de Vaugirard
Paris, 75015 - France
✉ florian.strub@gmail.com
🌐 [fstrub95](#)

Last update: July 2023

Summary

I am a Senior Research Scientist in Machine Learning at Google DeepMind, specializing in studying large language model training with RL and multimodal techniques. Throughout my academic journey, I have worked in many renowned AI laboratories worldwide (UK, Canada, France, and South Korea), publishing influential papers and accumulating thousands of citations. I now hold a pivotal role at Google DeepMind, where I effectively coordinate multiple scientists, leveraging my expertise to drive impactful research initiatives. In addition to my research proficiency, I have gained valuable product experience through previous roles in the industrial sector. I am recognized for being proactive and open-minded, with a versatile skill set and strong interpersonal abilities. I am adept at rigorously undertaking a wide range of new research projects and contributing to developing innovative products.

Professional Experiences

2019-2023 **Senior Research Scientist**, *DeepMind*, [Paris & London](#)

I coordinate a team of six scientists developing state-of-the-art algorithms for training large-language models with multi-agent RL techniques and vision capabilities. My specialty is to design self-improving methods to foster network capacities without human intervention. My research journey has been focused on three main applications:

1. Enhancing language agent skills with multi-agent and RL methods by interacting together language models:
 - Scaling up RL(HF) methods for NLP and developing methods to counter language drift.
 - Analyzing the emergence of communication protocols among conversational agents during co-training.
 - Training populations of communicative language agents through self-play.
2. Advancing multimodal models through self-supervised pretraining:
 - Co-inventing the BYOL algorithm, which pretrains visual backbone networks solely with raw images.
 - Extending BYOL to incorporate multimodal data such as video and audio or leveraging scarce labeled data.
3. Scaling up multi-agent and game-theory methods with deep learning:
 - Co-developed DeepNash, the first RL agent to achieve professional-level mastery in Stratego, a turn-based imperfect information game, later published in Science.
 - Contributing to the AlphaStar project, the pioneering RL agent that achieved professional-level mastery in Starcraft 2. (Listed in acknowledgment section due to interning at DM.)

2016-2019 **Ph.D. in Deep Learning**, *Univ. Lille*, [Lille](#)

I studied deep learning method required to learn consistent multimodal representations and finetune them through interactions with reinforcement learning. This Ph.D. was supervised by Prof. Olivier Pietquin, Dr. Jeremy Marie.

- Collected one of the largest visually grounded dialogue dataset GuessWhat?! containing 150k samples
- Precursive exploration of Deep Reinforcement Learning methods to finetune generative language models.
- Co-designed neural architectures to fuse vision and language modalities, e.g., Conditional Batch Norm, FiLM.
- Published six papers in A* conferences (NeurIPS, CVPR, IJCAI, AISTATS) and four workshop papers.
- All papers and experiments are open-sourced with multiple independent reproductions.
- The Ph.D. defense was reviewed by leading ML researchers, e.g., Aaron Courville, Patrick Gallinari.

2017 **Research Internship**, *Mila*, [Montreal - Canada](#)

Research collaboration on embodiment, and release of the multimodal and interactive 3D HOME environment. Three co-authored papers, including the widespread FiLM layer for multimodal network architectures.

2015-2016 **Data Scientist**, *Inria - SequeL*, [Lille](#)

I explored online recommendation systems and collaborative filtering with neural networks. This project aimed to predict future user preferences by leveraging their past ratings and meta-information.

2013-2015 **C++ Software developer in Finance, Front Office, Société Générale, Paris**

- I implemented a high-frequency trading API for low latency automatons. It connected trading automatons to worldwide markets and verified billion of trades (Stock, Future, Option) a day while forwarding them to the market in a few microseconds. As a technical leader, I was in charge of developing, testing, and releasing the API.
- Designed and developed from scratch an API for High-Frequency trading following business requirements
 - Provided technical support for IT Quants and Traders: Support Level 2 and 3
 - Maintained and refactored a 300k line codebase inside a team of 30 developers.
 - Developed a platform of tests (Test Driven Development) to sustain continuous deliveries.
 - Integrated sensitive features to the API dealing with official market regulations, trade monitoring, KPI, etc.

March 2012 – **Research Assistant (Master thesis), Imperial College, London - UK**

August 2012 I enhanced a C++ expert system that aids doctors in delivering precise breast-cancer diagnoses. The input data consisted of pre-processed scientific texts formatted as logic rules based on Argumentation Based Assumption (ABA) theory developed by Francesca Toni and Robert Craven.

May 2011 – **Research Assistant (Internship), KAIST – RCV Laboratory, Daejeon – Rep. of Korea**

Sept. 2011 I developed a web platform aimed at promoting researchers' works and facilitating the testing of cutting-edge algorithms in computer vision.

Academic Research Activities

Impact 7800 cites with a h-index of 20 over 40 papers (July. 2023).

- Selected Publications
- J. Perolat, B. De Vylder, B., D. Hennes, E. Tarassov, **F. Strub**, et al. *Mastering the game of Stratego with model-free multiagent reinforcement learning*. Science (2022).
 - JB. Grill*, **F. Strub***, F. Altché*, C. Tallec*, P. Richemond*, E. Buchatskaya, et al. *Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning*. In Proc. NeurIPS (2020) - Oral.
 - E. Perez, **F. Strub**, H. Vries, V. Dumoulin, A. Courville *FiLM: Visual Reasoning with a General Conditioning Layer*. In Proc. AAAI (2018) - Oral.
 - H. de Vries*, **F. Strub***, J. Mary, H. Larochelle, O. Pietquin, A. Courville *Modulating Early visual Processing by language*. In Proc. of NIPS (2017) - Spotlight.
 - **F. Strub**, H. de Vries, J. Mary, B. Piot, A. Courville, O. Pietquin *End-to-end Optimization of Goal-Driven and Visually Grounded Dialogue Systems*. In Proc. of IJCAI (2017).
 - H. de Vries, **F. Strub**, S. Chandar, O. Pietquin, H. Larochelle, A. Courville, *GuessWhat?! Visual Object Discovery through Multi-Modal Dialogue*. In Proc. of CVPR (2017) - Spotlight.
 - **F. Strub**, J. Mary, R. Gaudel, *Hybrid Recommender System based on Autoencoders*. In Proc. of Recsys workshop DLRS (2017).

Review Regular reviewer at NeurIPS, ICML, ICLR, CVPR with two outstanding reviewer awards.

Organizer Co-organized the ViGIL workshop (NeurIPS 2017, 2018, 2019 & NAACL 2021), European Workshop of RL (2018), the RL Summer School (2019), and internal Google DM workshops (2022 and 2023)

Supervision I am supervising two Ph.D. students, and already supervised two Ph.D students and five interns.

Education

2011–2012 **MSc in Advanced Computing, Imperial College of London, Distinction - Top 10%**

2009–2011 **Joint Master in Engineering and Management, École Nationale Supérieure des Mines de Saint-Étienne, First class honor - Top 10%**

Double Diploma: Research Master in Web Intelligence from University of Jean Monnet

*Equal Contribution