

## 2019-01287 - Post-Doctoral Research Visit F/M Postdoc: High Performance Deep Reinforcement Learning

**Type de contrat :** CDD de la fonction publique  
**Niveau de diplôme exigé :** Thèse ou équivalent  
**Fonction :** Post-Doctorant

### A propos du centre ou de la direction fonctionnelle

Inria the French national institute for research in computer science and control, is dedicated to fundamental and applied research in information and communication science and technology (ICST). Inria has a workforce of 3,800 people working throughout its eight research centers established in seven regions of France.

Grenoble is the capital city of the French Alps. Combining the urban life-style of southern France with a unique mountain setting, it is ideally situated for outdoor activities. The Grenoble area is today an important centre of industry and science (second largest in France). Dedicated to an ambitious policy in the arts, the city is host to numerous cultural institutions. With 60,000 students (including 6,000 foreign students), Grenoble is the third largest student area in France.

### Mission confiée

Reinforcement learning goal is to self-learn a task trying to maximize a reward (a game score for instance) interacting with simulations.

Recently, researchers have successfully introduced deep neural networks enabling to address more complex problems. This is often referred as

Deep Reinforcement Learning (DRL). DRL managed for instance to play many ATARI games. The most visible success of

DRL is probably AlphaGo Zero that outperformed the best human players (and itself) after being trained without using data from human games but solely through reinforcement learning. The process requires an advanced infrastructure for the training phase. For instance AlphaGo Zero trained during more than 70 hours using 64 GPU workers and 19 CPU parameter servers for playing 4.9 million games of generated self-play, using 1,600 simulations for each Monte Carlo Tree Search.

The general workflow is the following. To speed up the learning process and enable a wide but thorough exploration of the parameter space, the learning neural network interacts in parallel with several instances of actors, each one consisting of a simulation of the task being learned and a neural network interacting with this simulation through the best winning strategy it knows. Periodically the actor neural networks are being updated by the learned neural network.

This workflow has evolved through various research works combining parallelization, asynchronism and novel learning strategies (GORILA, A3C, IMPALA,...).

The goal of this postdoc is to push forward the scalability of these approaches, and to proposing novel learning strategies to learn more rapidly and more complex tasks (multiple heterogeneous tasks at once, non deterministic games, simulations of complex industrial or living systems).

This work will be performed in close collaboration in between the Sequel INRIA team specialized in DRL (<https://team.inria.fr/sequel/>) and the DataMove team specialized in HPC (<https://team.inria.fr/datamove/>).

Datamove has developed the Melissa (<https://melissa-sa.github.io/>) solution to manage large ensembles of parallel simulations and aggregate their data on-line in a parallel server. Melissa enabled to run thousands of simulation on up to 30 000 cores. So far Melissa was used to compute advanced statistics. But we expect this framework to be a sound base for a DRL workflow. The Sequel team has strong activities in reinforcement learning, either deep or not, ranging from theoretical aspects to applications. Among other projects, Sequel has collaborated with Mila (Montréal) to design and develop the Guesswhat?! experiment (<https://guesswhat.ai/>). As early as 2006, Sequel worked on go and designed the first go program (Crazy Stone) able to challenge a human expert player.

### References

- AlphaGoZero: <https://deepmind.com/blog/alphago-zero-learning-scratch/>
- TensorFlow: <https://www.tensorflow.org/>
- Gorila <https://arxiv.org/pdf/1507.04296>
- A3C <https://arxiv.org/abs/1602.01783>
- Rainbow <https://arxiv.org/abs/1710.02298>
- Impala <https://arxiv.org/abs/1802.01561>
- Elf: <https://arxiv.org/abs/1707.01067>
- RAY/RLlib: <https://ray.readthedocs.io/en/latest/rllib.html>
- Melissa: "<https://hal.inria.fr/hal-01607479v1>

### Principales activités

- Requirement: PhD in computer Science
- Location: Grenoble or Lille
- Hosting Teams:
  - Sequel (INRIA Lille): <https://team.inria.fr/sequel/>
  - DataMove (INRIA Grenoble): <https://team.inria.fr/datamove/>

### Informations générales

- **Thème/Domaine :** Calcul distribué et à haute performance  
Statistiques (Big data) (BAP E)
- **Ville :** Montbonnot
- **Centre Inria :** CRI Grenoble - Rhône-Alpes
- **Date de prise de fonction souhaitée :** 2019-09-02
- **Durée de contrat :** 2 ans
- **Date limite pour postuler :** 2019-08-31

### Contacts

- **Equipe Inria :** DATAMOVE
- **Recruteur :**  
Raffin Bruno / [bruno.raffin@inria.fr](mailto:bruno.raffin@inria.fr)

### A propos d'Inria

Inria, l'institut national de recherche dédié aux sciences du numérique, promeut l'excellence scientifique et le transfert pour avoir le plus grand impact. Il emploie 2400 personnes. Ses 200 équipes-projets agiles, en général communes avec des partenaires académiques, impliquent plus de 3000 scientifiques pour relever les défis des sciences informatiques et mathématiques, souvent à l'interface d'autres disciplines. Inria travaille avec de nombreuses entreprises et a accompagné la création de plus de 160 start-up. L'institut s'efforce ainsi de répondre aux enjeux de la transformation numérique de la science, de la société et de l'économie.

### Consignes pour postuler

#### Sécurité défense :

Ce poste est susceptible d'être affecté dans une zone à régime restrictif (ZRR), telle que définie dans le décret n°2011-1425 relatif à la protection du potentiel scientifique et technique de la nation (PPST). L'autorisation d'accès à une zone est délivrée par le chef d'établissement, après avis ministériel favorable, tel que défini dans l'arrêté du 03 juillet 2012, relatif à la PPST. Un avis ministériel défavorable pour un poste affecté dans une ZRR aurait pour conséquence l'annulation du recrutement.

#### Politique de recrutement :

Dans le cadre de sa politique diversifiée, tous les postes Inria sont accessibles aux personnes en situation de handicap.

**Attention:** Les candidatures doivent être déposées en ligne sur le site Inria. Le traitement des candidatures adressées par d'autres canaux n'est pas garanti.

- Contact: [Bruno.Raffin@inria.fr](mailto:Bruno.Raffin@inria.fr) and [Philippe.Preux@inria.fr](mailto:Philippe.Preux@inria.fr)
- Period: to start somewhere in 2019
- Duration: 24 months

We are looking for a candidate with a PhD either in deep learning, reinforcement learning or high performance computing (a combination of these expertise would be ideal) for a 24 month contract at INRIA. The candidate will have the possibility to join either the Sequel team at Lille or the Grenoble Team at Grenoble.

The postdoc will have access to large supercomputers equipped with multiple GPUs for experiments. We expect this work to lead to international publications sustained by advanced software prototypes.

## Avantages

- Subsidized meals
- Partial reimbursement of public transport costs
- Leave: 7 weeks of annual leave + 10 extra days off due to RTT (statutory reduction in working hours) + possibility of exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking (after 6 months of employment) and flexible organization of working hours
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Social, cultural and sports events and activities
- Access to vocational training
- Social security coverage

## Rémunération

Salary: 2 653 € gross/month.

Monthly salary after taxes : around 2 136,39 € (medical insurance included, income tax excluded).