

Research Scientist – Imitation Learning

The Role

We are looking for a Research Scientist to join our Imitation Learning team on a full-time basis, reporting into our CTO. You will manage 1 or more machine learning engineers, contribute to the development of state-of-the-art *Imitation Learning* (IL) and *reinforcement learning* (RL) algorithms, and oversee their application to learn control policies for simulated traffic agents from real-world data.

What we'll ask from you:

- Passion for the state-of-the-art in IL, RL, and the wider field of machine learning, and a drive to keep up-to-date the latest academic literature;
- The ability to lead collaborative research efforts, through planning and managing the activities of a small team, to ensure that the R&D objectives of the company are met in a timely manner;
- Commitment to delivering high quality software at the output of R&D activities, promoting and implementing software engineering best practices such as code reviews and thorough documentation;
- A willingness to get things working, flexibly helping more junior team members with development efforts where and when needed;
- Investment in developing strong working relationships, and to mentoring and training less experienced team members;
- Active engagement with the commercial team to understand the needs of our customers and how we can best address these in the future R&D roadmap.

We offer an open-minded, diverse and non-hierarchical culture where your ideas are valued and you can work alongside leading experts, in a fast-paced start-up environment where you are close to the action and can make a real impact on our culture, strategy and IP.

We may offer share options so you can share the future success of our growing company.

Your Experience

We are looking for candidates with a PhD in Computer Sciences or Machine Learning (or a related field) and who:

- Have published academic work on the topics of machine learning or decision theory;
- Have implemented deep learning methods in practice and can demonstrate contributions;
- Are familiar with the theoretical basis for reinforcement learning and/or decision-making under uncertainty;
- Are confident with Python and have an understanding of deep learning software libraries such as Tensorflow or Pytorch;
- Have some experience using Linux for software development;

- (Optionally) have had some experience of collaborative software development;
- (Optionally) have worked with the state-of-the-art in reinforcement learning, ideally with some experience in imitation learning

Direct experience in self-driving cars is not required. We have a diverse team and international candidates are welcomed.

Your Values

Successful candidates will share our values and help to role-model them. This means we look for candidates who are:

Honest and open: you prefer straight-forward communications. You believe that open but empathetic communication is the key to building trust and successful relationships. You like the idea of working in and promoting a culture of continuous 360-degree feedback so you and others keep learning and developing skills

Collaborative and practical: you want to have an impact on the world through translating state-of-the-art research to software that addresses real-world challenges. You like solving problems in teams, and are happy to help with practical, short-term tasks to get the job done. You want to build high-quality code that lasts, so you appreciate the importance of software engineering best practices and welcome supervision from our senior software developers to that end.

Pioneering and intellectually-curious: our mission is to extend the state-of-the-art in our fields of computer vision and imitation learning – which means coming up with new and evolving solutions to hard problems that no-one's answered before. You take structured approaches to open research questions and understand the uncertain nature of scientific research.

Open-minded: you try to take an objective approach to cracking real-world problems. You believe good ideas can come from anywhere and want to create an environment where less experienced members of the team feel free to suggest their own ideas and think for themselves. You take time to understand where other people are coming from, and believe you can learn from them, even if you don't immediately agree with them.

Start Date: January 2019 (or earlier by agreement). **Location:** Oxford, U.K.
