

Machine Learning Scientist

About Us

Latent Logic (www.latentlogic.com) is a dynamic deep tech start-up spun out from Oxford University. Our mission is to use state-of-the-art A.I. to learn complex human behaviours from real demonstrations, such as video data, so that we can design autonomous systems which can work collaboratively alongside humans. Our first product is building the virtual humans for the simulations used to test and train safer self-driving cars.

We closed our seed round in February 2018 and are looking for bright, motivated people to join our growing team and help us achieve our mission, and help make the world a safer place.

We are currently recruiting for a Machine Learning Scientist to join the Imitation Learning team. **To apply for this position**, please send your CV and a brief covering email to Applications@LatentLogic.com

The Role

You will report into our CTO and manage 2 or more machine learning engineers. Working alongside another similar-sized team, you will lead the development of state-of-the-art *learning from demonstration* (LfD) and *reinforcement learning* (RL) algorithms, and oversee their application in a practical setting to learn control policies for simulated traffic agents from real-world data.

What we look for:

- The ability to lead collaborative development efforts, through planning and managing the activities of a small team, to ensure that the product roadmap of the company is delivered in a timely manner;
- Commitment to delivering high quality software at the output of development 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 building 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 product development roadmap;
- Passion for the state-of-the-art in LfD, RL, and the wider field of machine learning, and a drive to keep up-to-date the latest academic literature

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 join in 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 application of research in a commercial context or 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 can sponsor Tier 2 visas.

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 coaching 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 are listened to.

We believe hard problems like ours can only be cracked by people from different perspectives and backgrounds coming together with a shared mission, which is why we encourage women, minorities and members of under-represented groups to apply to join us, and why we are committed to improving the diversity and inclusiveness of our industry.

Location: Oxford, U.K.

Start Date: Spring / Summer 2019