



## Machine Learning Operations Engineer

We started out by wanting to make satellite imagery available to all. We've grown to understand it's even more powerful when you match it to real business needs.

Bird.i is a fast growing "big data from space" start-up backed by Hanley Wood | Meyers Research, the housing industry's foremost advisors and provider of rich data. We have unique partnerships with the world's leading satellite and airborne operators, giving us access to an unprecedented amount of Earth observation imagery.

We're determined to make it accessible, affordable and usable, putting the end user at the heart of the process, so that we provide a service that delivers smoothly, efficiently and effectively. Over and over again.

With the collaboration from our group companies, we process this imagery with machine learning techniques to monitor the construction of real estate projects and make these unique sources of insights available to builders, developers and investors.

As a Machine Learning Operations Engineer at Bird.i, a typical day might include the following:

- Celebrate with the team the latest improvement in the accuracy of our intelligence services
- Pour yourself a fresh cup of coffee, grab a biscuit, catch-up with the progress from our US colleagues, before participating in the daily team stand-up
- Review the previous day models performances and iron out a bug or two
- Enjoy your lunch while hearing from a fellow employee about a topic of interest
- Spend the afternoon peer programming with a team mate
- Finish the day with brainstorming on culture and discuss how we can make Bird.i an even more desirable employer

This job could be a good match if you:

- Take culture seriously and share our values: **respect, simplicity, communication** and always maintaining a **user-focus** in everything we do
- Can deliver consistent high quality, high-impact work in an agile and changing environment

- Understand that Machine Learning is primarily about alleviating customer problems as much as writing code
- Are keen to bear the full weight of your ML knowledge and DevOps skills to create world class ML Products
- Are happy to take risks, learn from failure and improve with constructive feedback
- Always keep customers in mind for every engineering decision you make

The techy bit:

You will be training Machine Learning models that will solve user problems.

You will be recording metrics of model performances and identifying models for production.

You will be creating and implementing data formats for input and output for the ML products.

You will be actively involved in collecting, preparing and curate data for training the models.

You will be creating and running infrastructure in AWS for development and production.

Our technology stack is mainly based on Python and PyTorch. We run our production workloads in AWS. You will extensively use Jupyter and Docker for experimenting and production respectively.

To land this gig, you must:

- Postgraduate degree in Computer Science, preferably with some majors in Big Data, Artificial Intelligence, Machine Learning or Data Science.
- Strong Machine Learning experience preferably using PyTorch/Tensorflow.
- Strong Mathematical foundations with a focus on Differential Calculus, Linear Algebra, Statistics and Probability
- Working experience in Cloud technologies, preferably in AWS
- Working knowledge of SQL
- Python

If you have the following skills, it's even better:

- Basics of GIS and Geography
- Linux
- Computer Vision knowledge as we work on images
- Agile Processes
- CI/CD and Source Control principles
- Experience in Docker; creation, deployment and management.

Think you have what it takes? Please send your cover letter and CV at [siva@hibirdi.com](mailto:siva@hibirdi.com)

**'NO RECRUITMENT AGENCIES PLEASE!**

**We want motivated candidates looking for their dream job by themselves!'**