
Python Software Developer

DESCRIPTION

We're transforming the semiconductor and electronics manufacturing industry with our AI platform that improves yield and lowers manufacturing costs at some of the largest chip makers in the world. We're looking for python software developers who are interested in implementing, improving, and scaling new and existing algorithms using open source machine learning libraries/frameworks or writing one from scratch. We look for some combination of: high creativity and productivity, depth and breadth of knowledge in the field, and the ability to implement and scale up state of the art and/or unproven algorithms. You'll work closely with other data scientists and software engineers to implement feature engineering, data transformations, and machine learning algorithms from scripting languages to high performance code that can scale. We look for people who are self-motivated and passionate about the transformative potential of AI. This position presents an exciting and unique opportunity to design/build the product and the platform from scratch utilizing high performance machine learning algorithms, data science and the associated frameworks.

Location: Santa Clara, CA

RESPONSIBILITIES

- Implementing scalable ML algorithms in production
- Improving computational performance of ML algorithms
- Interact with Software Engineers and Data Scientists to create clean API interfaces

REQUIREMENTS

- MS or PhD in Computer Science, OR, EE, Physics or related quantitative field
- 3+ years of experience in back-end development in python, ideally R as well.
- 3+ years of engineering experience with significant exposure to algorithm implementation/machine learning libraries/frameworks.
- Experience with relational and NoSQL databases such as PostgreSQL/Cassandra/Redis or similar.
- Experience with API calls.
- Hands-on experience with H2O framework source code/new algorithm additions etc.
- Ability to analyze the bottlenecks and create high performance code.
- Good understanding of CPU and memory issues in parallel computing

JOB DETAILS

This work involves building robust ML models in python and/or R to achieve desired accuracy on metrics of customer interest (e.g. yield, tolerance). Broken into steps, the work is roughly an equal mix of tweaking core ML algorithms for enhanced performance and creating deployment-ready software.

This role offers a unique opportunity to deploy non-trivial ML models into production and seeing the impact of those models in action at some of the world's largest chip manufacturers. This is **not** a Hadoop/Spark-focused big data role but rather an ML-focused role. Our deployment environment is markedly different from typical cloud deployment with constant access for updates or bug-fixes. Roughly stated, the production requirements are closer to a self-driving robot than a cloud-integrated image recognition service.

An ideal candidate has invested time in continuous learning and implementing algorithms either using an existing framework/library or from scratch and understands failure modes of various algorithms in production. Additionally, the ideal candidate has collaborated with Data Scientists to put ML code into production.

Please submit resumes to joblogic@pdf.com Subject header: **Python Software Developer**