

Palantir Technologies – London – Software Engineer (Machine Learning)

We are looking for talented and bold engineers who love to code, get their hands dirty with raw data, and derive meaningful and actionable insights.

Our technology enables customers to recognize and understand complex relationships within heterogeneous data. Our products facilitate the discovery, iterative refinement, and testing of domain-specific hypotheses, and present results in an understandable form for end-users.

The datasets we work with are as varied as our customers (e.g. government agencies, financial institutions, media companies, disaster relief), who trust us to help them use their data to solve their most important problems. Often, these problems are not clearly defined – they are discovered and refined through hard work, rapid iteration, and perceptiveness.

Ideal candidates learn and adapt quickly, and will be able to use every tool at their disposal – software, algorithms, statistical models, and beyond – to understand and effectively tackle hard problems. They appreciate the difference between explaining and fitting statistical models, the importance of good metrics, and the tradeoff between exploration and exploitation. They can perceive common structure between superficially unrelated problems, and can use this to build tools, algorithms, and products with superlinear value.

RESPONSIBILITIES

Work directly with customer data to derive actionable insights

Develop statistical or machine-assisted approaches to problems at massive scale

Build out tools and infrastructure for data analysis

REQUIREMENTS

BS/MS in Computer Science, Statistics, Mathematics, or related field (started or completed a PhD a plus)

Substantial experience developing scalable machine learning or quantitative analysis software in an industry or research environment

Ability to understand and deliver in the presence of rapidly evolving product, customer, and business needs

Proficiency in at least one compiled language (e.g. C, C++, Java) and one scripting language (e.g. Python, R, MATLAB)

A desire to transform dirty/noisy signals within customer data into valuable results through code and data grooming

Ability to travel preferred

DESIRABLE

Experience developing software within a distributed computation framework (e.g. Hadoop, Spark, Storm, GraphLab)

Experience developing distributed systems, data visualization or enterprise software systems

Working familiarity with SQL

Prior work in natural language processing