Applied Scientist

HyperCube develops a high-performance enterprise platform for deploying machine learning and artificial intelligence applications. The nature of our product is such that our technical staff must be technically outstanding, scientifically grounded, and customer outcome driven.

Applied Scientists at HyperCube build components of high performance data systems, design and implement high-performance algorithms, follow proper development processes, generate insights from data, perform feature engineering and model training, benchmark and experiment different solutions, communicate with the open source community, produce and publish new research, present at technical and academic conferences, and communicate with customers.

Required skills and qualifications

- At least 5 years of experience working as a data scientist, software engineer, research engineer, architect, scientist, or product manager. Relevant graduate degrees are counted as years of experience.
- A Masters or PhD Degree in a technical domain such as Computer Science, Mathematics, Physics, Software Engineering, Statistics, and Operations Research.
- Proven ability to effectively communicate priorities, opportunities, challenges, plans, solutions, decisions, goals, and priorities.
- Being able to understand, communicate, and implement technical white-papers and academic publications in machine learning, algorithms, or distributed systems.

Preferred skills and qualifications

- Experience with proper development and deployment processes. These include scrum or other agile methodology, working in continuous deployment mode, performing/passing code reviews, operating and monitoring remote machines and distributed systems.
- Thorough understanding and experience with internal workings, codebase, and algorithms of different deep learning and shallow learning frameworks.
- Experience with performing novel research in algorithms, machine learning, data science, or high performance computing. Experience publishing results in academic venues.
- Experience building end-to-end machine learning applications including problems definition, building data pipelines, feature engineering, model training, and deployment into production systems.