

Data Scientist

About Peritus.ai

Peritus enables self-healing autonomous datacenters with automated, cognitive support for infrastructure software and hardware. It is a funded startup co-created at The Hive in Palo Alto, CA that delivers artificial intelligence based virtual support expert systems for datacenter service fulfillment and incident resolution.

As datacenter vendors move from on premise to the cloud their existing support system lacks the agility and cost-effectiveness for the cloud. Peritus significantly enhances operational efficiencies of existing support services, and enables managed service providers & system vendors to offer new business continuity entitlements. Peritus assists & automates a wide spectrum of decisions in system support including incident classification, routing, contract coverage, incident resolution recipes and orchestration of incident management between subject matter experts (SMEs).

Peritus' unique vectorization of system log data drives predictive modeling with highly granular feature extraction for early detection of system events. The platform's advanced natural language processing (NLP) capabilities drive Peritus' incident modeling and predictive capabilities. The core service fulfillment engine uses a combination of supervised and unsupervised methods to predict incident features from system log data. Peritus delivers automated orchestration of incident resolution through its close integration with existing incident management platforms.

About the role

As part of the engineering team, drive and execute the Machine Learning/Artificial Intelligence technology roadmap to solve the challenges of dealing with structured, semi-structured and unstructured data in the context of a datacenter. In short, it's about machines learning about machines, i.e., Systems ML.

Responsibilities

- Advance the Machine Learning capabilities of the product.
- Analyze, select and implement Machine Learning/AI algorithms that improve the relevance of related incidents, accuracy of resolutions and routing of incidents.
- Experiment with multiple algorithms simultaneously to identify the optimal algorithm.
- Enhance the experimentation platform by incorporating implicit and explicit feedback from users.
- Understand the data by interacting with the Subject Matter Experts.

About you

- You have a Master's / PhD in Computer Science, Machine Learning or equivalent field.
- You have experience in building products from scratch using machine learning.
- You have proven experience in developing algorithms using Python or C++ or Scala in one or more of the following fields: machine learning, data analytics, information retrieval.
- You have experience applying statistics, probability, and machine learning to real business problems.
- You have experience in Natural Language Processing such as topic extraction, building knowledge graph (ontology).
- You have experience in building machine learning models for time series data.
- You have experience with Deep Learning frameworks (TensorFlow), applications, or services.
- You have excellent knowledge in optimal computational environment required for data processing, mining and machine learning
- Your knowledge in one or more of the following data center supportability technologies will be a big plus - networking, storage, supportability (e.g. incident management, systems diagnostics).



PERITUS

720 University Avenue Suite 200
Palo Alto CA 94301 United States
Ph: (650) 234-0400

<http://peritus.ai>
Twitter: @Peritus_AI
Email: info@peritus.ai

- You are smart, get things done, have great energy and thrive in a fast paced early-stage startup environment.
- You have passion for creating customer value by applying cutting edge innovations in data science & technology.

Please send your resume to jobs@peritus.ai.