Byte

Machine Learning Customer Engineer

Description

(This role can be remote or hybrid role based in New York City or Redwood City or Redwood City + San Francisco) As an ML Customer Engineer, you are integral to the post-sales journey for our enterprise customers. In this role, you will do more than manage issues and SLAs, you will help solve complex customer problems, collaborate cross-functionally with field and engineering resources, and serve as a trusted advisor. The MLCE role extends to educating customers on machine learning concepts, proposing creative solutions, actively contributing to the company's growth and helping shape our product. If growing your skills and solving some of the most challenging real world machine learning problems excites you, continue reading. Main Responsibilities Partner with Snorkel Flow users to design, build, troubleshoot and deploy AI applications. Lead the resolution of critical technical issues, providing prompt and complete resolution to technical challenges and business issues. Perform live working sessions to analyze and address customer reported issues. Prioritize, document and coordinate customer issues with account assigned ML Success Managers and the Snorkel engineering team. Contribute to internal and external guides and docs, improving our self-service support materials. Become an expert in the Snorkel Flow platform and assist our customers do the same. Drive improvements in issue triage, reporting, and analysis to better understand customer pain points. Minimum Qualifications 2+ years experience working in a technical customer-facing role, e.g., technical support or account management, management consultant or customer success. B.S. degree in a quantitative field such as Computer Science, Engineering, or comparable degree/experience. Previous experience working on machine learning projects or industry knowledge of standard technologies in the machine learning space. Proficiency in writing Python for either Data Science, Machine Learning, or other distributed systems workloads. Outstanding organizational skills and ability to multitask in order to effectively prioritize and manage customer requests. Experience with common support software like Zendesk, Jira, and Slack. Preferred Qualifications Expertise in modern ML frameworks and libraries, e.g., PyTorch, Scikit-learn, Numpy, SciPy, and HuggingFace. Track record of collaboration across field and engineering teams to manage support issues and resolution within accounts. Previous experience with cloud infrastructure providers such as Amazon Web Services, Microsoft Azure, or Google Cloud Platform. Experience operating Kubernetes orchestration tools in a production setting. The salary range for our Tier 1 locations of San Francisco, Seattle, Los Angeles & New York is $120,000 – $160,000. All offers include equity compensation in the form of employee stock options.   #LI-SH1Please mention the word **RESTFUL** and tag RMTA3LjE3OC4yMzIuMjQy when applying to show you read the job post completely (#RMTA3LjE3OC4yMzIuMjQy). This is a beta feature to avoid spam applicants. Companies can search these words to find applicants that read this and see they're human.

Contacts

Job listing via RemoteOK.com

Hiring organization
Snorkel AI

Job Location
Redwood City, California, United States

Base Salary
$ 60000 - $ 110000

Date posted
June 6, 2024

Apply Now