



Opening: Site Reliability Engineer (SRE)

Who we are

DeepLearni.ng is a technology company transforming the way enterprises work with AI around the world. Based in Toronto, we work with enterprise leaders to collaboratively build and deploy customized AI systems that solve real business problems. DeepLearni.ng's creation of enterprise AI systems is powered by our end-to-end platform Frontiers, designed to optimize team collaboration on enterprise AI systems.

We are also an active player in the AI community, and have partnerships with industry leaders including NVIDIA, the Vector Institute and Amazon Web Services. In 2017, we were named AI Company of the Year by the DFI Institute and received Series A funding.

A few examples of our work:

- Reinforcement learning for the credit risk areas of retail banking
- Image classification for insurance adjudication and claims processing
- Self-optimizing environmental controls for vertical farming
- Machine vision for robot warehouse navigation

What is the opportunity?

We are looking for a **Site Reliability Engineer (SRE)** to join our high performing team. The SRE will be a foundational member creating the practice of stringently testing machine learning tools for enterprise, including an end-to-end machine learning platform. You will work with diverse enterprise teams across industry verticals, helping drive quality assurance, user acceptance testing, performance optimization, load testing and go-to-market for tools and metrics.

As a SRE, your main focus will be to maintain a high bar and rigorous standards in overall quality and user experience of new machine learning tools custom-built for enterprise. You will also

We are an equal opportunity employer and value diversity at our company. We do not discriminate on the basis of race, religion, color, national origin, gender, sexual orientation, age, marital status, veteran status, or disability status.

create new efficiencies for our Engineering team, driving programs designed to measure their efficacy and velocity, with KPIs reported to the CTO on a weekly basis.

Working with a diverse range of technologies and clients across industry verticals, you will have the opportunity to become a leading expert in delivering transformative machine learning solutions to industry designed for real-world results. As part of a vibrant and growing team, you will be at the frontier of the AI revolution.

What will you do?

- *Quality Assurance*: maintain, update and create strong QA methodology and testing of all our product lines to deliver bug-free code to our clients
- *Performance Testing*: ensure that our world class tools enable machine learning experts (and aficionados alike) to run experiments at break-neck speed (as fast or faster than any CLI)
- *Load Testing*: help us optimize and find the limitations of the system as a whole
- *Reliability and Usability Testing*: give the best of user experience in the paradigm of always up, 100% SLA world
- *User Experience Refinement*: empathy for the end user and ability to positively impact usability for a multitude of enterprise users
- *Foundational Leader*: set the process for the quality assurance practice and facilitate across team for our lines of products
- *Release Management*: produce deployment plans, run books and implementation plans; conduct release readiness reviews, milestone reviews, and go/no-go reviews
- *Project Management*: quality assurance, user acceptance testing to go-to-market strategies to optimize product launch success
- *Training & Collateral*: work with product marketing, product management and account management / professional services organizations to build collateral for training the wider organization, and to scale solutions globally
- *Enterprising*: collaborate on identification of new opportunities for DeepLearning to innovate on behalf of our customers in conjunction with Product Management, Business Development and Client input
- *Outreach*: active involvement in evangelism activities such as giving tech talks, running webinars, leading internal hack-a-thons and coordinating lunch & learns / DoJos

We are an equal opportunity employer and value diversity at our company. We do not discriminate on the basis of race, religion, color, national origin, gender, sexual orientation, age, marital status, veteran status, or disability status.

What do you need to succeed?

Must have:

- 3+ years of software development experience with hands-on experience in reliability testing, including experience testing low-level system concepts
- An Engineering, Computer Science, Physics, Math, or similar degree
- Demonstrable experience with Agile or similar methodologies designed for continuous product development
- Ability to root cause sources of instability in a large-scale distributed system
- Experience with configuration and troubleshooting of Linux, Java, Scala, Spark, PySpark, Hadoop, Kubernetes and other middleware technologies
- Strong understanding of large-scale complex systems from a reliability perspective
- Scripting abilities in scripting languages such as Python, Perl, Ruby, Bash, NodeJS or JVM-based languages
- Passion for resolving reliability issues and identify strategies to mitigate going forward
- Ability to dive extremely deep into code, network transports, REST APIs, websockets in a Linux environment

Nice to have:

- Excellent speaking, listening, and writing skills, attention to details, proactive self-starter
- Strong Linux system-level analysis capabilities
- Experience delivering solutions at scale based on a variety of cloud and on-premises services (in the datacenter, AWS, Google, Azure, Digital Ocean, etc.)
- Ensure community reach, having lead meetups, webinars, hacks, etc., including contributions to open-source or personal projects
- Experience working on a B2B application for enterprise clients
- Experience with or knowledge of machine learning and deep network analysis

What's in it for you?

- Work with purpose: contribute to deploying machine learning models that drive real business impact across industry verticals
- Opportunity to work with leading machine learning experts on cutting-edge problems
- Opportunities to develop relationships with a variety of industry decision-makers, subject-matter experts, software developers, and machine learning specialists

We are an equal opportunity employer and value diversity at our company. We do not discriminate on the basis of race, religion, color, national origin, gender, sexual orientation, age, marital status, veteran status, or disability status.

- Split time between your work duties and pursuing personal machine learning passion projects with mentoring from our machine learning experts
- Participate in feats of strength and lively company events that build your career toolkit while also enjoying a work-life balance
- Enjoy a comfortable company work environment located in downtown Toronto that encourages self-expression, creativity, honesty, and learning
- Amazing quality of life benefits: healthy snacks, gym memberships, travel accommodations, great health benefits, and more!
- Stock options possible after you have proven your impact on the job
- As needed: sponsorship for work visas

How to apply

We look forward to receiving your application! Send your resume and cover letter addressed to the Hiring Manager by email at joinus@deeplearni.ng. Please note that we will be in touch with applicants that make it to the next phase of the interview process directly.

By submitting your CV and application to this role, you authorize DeepLearnI.ng to store your information in our internal recruitment database and to reach out to you as necessary for the purpose of further evaluating your qualifications for this or other career opportunities at the company.

We are an equal opportunity employer and value diversity at our company. We do not discriminate on the basis of race, religion, color, national origin, gender, sexual orientation, age, marital status, veteran status, or disability status.