



Opening: Machine Learning Engineer

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 from Fidelity Investments Canada.

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 **Machine Learning Engineer** who is passionate about creating solutions with massive impact, solving challenging real-world problems, and working in unexplored areas, continuous improvement and Agile development. Applied machine learning experience in driving business value from large datasets in Enterprise environments is a must.

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

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 will you do?

- Take charge and manage state-of-the-art machine learning and deep learning projects with billions of data points from day one
- Collaborate with enterprise clients to customize machine learning solutions for their distinct business requirements
- Design and develop different machine learning systems, insights and algorithms, and then prototype, deploy to production, and test them to ensure all client's business needs are covered
- Manage overall software implementation life cycles of projects, ensuring compliance with existing and future enterprise implementation standards
- Facilitate the deployment and configuration of DeepLearn.ng software applications
- Support the deployment team's preparation of work plans, milestones, resource needs according to DeepLearn.ng's standard methods and practices
- Mentor junior machine learning engineers in developing their skills and knowledge
- Communicate ideas, updates, capabilities, and workflows to various business stakeholders as well as your fellow DeepLearn.ng team members

What do you need to succeed?

Must have:

- 3+ years of applied machine learning experience working with live datasets in enterprise environments solving real-world problems using different approaches and techniques
- Experience leading teams in deploying machine learning projects in enterprises
- Proven prior experience working with very large datasets and live production systems where your ML solutions have made a real impact with measurable implications
- An Engineering, Computer Science, Physics, Math or similar graduate degree
- Effective client-focused communication skills and comfort in interacting with diverse technical and business stakeholders at all levels
- Proven hands-on experience with Python Machine Learning libraries, such as TensorFlow, and PyTorch, distributed computing frameworks such as Spark, and NoSQL databases such as Cassandra and MongoDB
- Demonstrable understanding of and passion for statistics, probability theory, and data modeling, with a penchant for great algorithm design and strong CS fundamentals
- Ability to work on a diverse team or with a diverse range of people, and in a rapidly changing multi-disciplinary work environment

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.

Nice to have:

- Demonstrable understanding of deep learning
- Experience with semi-supervised, reinforcement learning, active learning, or natural language processing (NLP)
- Experience implementing high-performance Service Oriented Architectures
- Expertise with Docker or Kubernetes
- Contributions to open-source or personal projects, combined with an active voice on the AI/ML/DL scene such as record of published research papers, involvement at conferences/meetups/online teaching, etc.
- Experience working in the following sectors: banking, insurance, telecommunications

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
- Opportunity to attend exclusive seminars at the Vector Institute with world renowned researchers in deep learning and AI
- Split time between your work duties and pursuing personal machine learning projects
- 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.