

AI Engineer Job Description:

OVERVIEW

The use of machine learning to reimagine software applications and service development is exploding. Companies from every corner of the industry -- the biggest cloud service providers to corporate industrials to financial services to healthcare to retailers -- are exploring new ways of building products and services using data-centric learning models in place of traditional explicit programming. The drive to deliver more timely and more accurate results is compelling an ever greater need for specialized computing power. GPUs have been hailed as the solution to those computing needs but the industry is actively searching for a better, more efficient solution. Graphcore has that solution.

Graphcore's Intelligence Processing Units, or IPUs, are specifically designed for artificial intelligence and compute-dense graph applications. These are not GPUs, Graphics Processing Units, but rather graph processors especially adept at the kinds of computations used in understanding relationships within a sea of data. Architecturally, IPUs look nothing like GPUs. They offer performance, latency, and power efficiency advantages that uniquely distinguish them from GPU, CPU, and FPGA alternatives. Graphcore offers support for popular industry ML frameworks and a full tool suite for developers to innovate both within and outside those frameworks.

As an AI Engineer, you will be working closely with Graphcore's customers and partners to help them in understanding and getting the most from our Intelligence Processing Unit (IPU) technology. In your work you will support some of the world's top machine learning innovators at deep learning research groups, at academic institutions, at innovative machine learning start-ups, at leading automotive companies, and at some of the world's largest cloud and internet companies.

You will need to develop a deep understanding of the state-of-the-art in artificial intelligence & machine learning domains and work with our customers to develop new techniques which exploit the unique features of our IPU architecture. We want you to become an industry thought leader on Graphcore technology and machine learning applications in the cloud, in automotive and in embedded applications. You should be interested and keen to present at industry conferences and will be able to back this up with written blogs and compelling content.

RESPONSIBILITIES

- Develop strong technical relationships with researchers and engineers at our customers and partners and help them to develop new algorithms

using Graphcore's IPU technology and achieve breakthroughs in artificial intelligence

- Support Sales teams as they engage new customers, understanding their AI challenges and describing how Graphcore's technology can help
- Become a thought leader on machine learning and advocate for Graphcore's IPU technology
- Work with the Product Management and Engineering to ensure a good flow of customer and market feedback that can be incorporated into future products

REQUIREMENTS

- 4+ years of related experience ideally in Machine Learning, HPC or a math's intensive engineering field
- Bachelors in Engineering, Computer Science, Mathematics, Physics or similar field
- Experience working with modern deep learning software architecture and frameworks including: Tensorflow, MxNet, Caffe, Caffe2, Torch, and/or PyTorch, or a strong mathematics background and a keen interest to learn about these exciting new techniques.
- Comfortable in a customer-facing environment and experience working with developers or researchers
- Ability to multitask effectively in a fast paced environment
- Action-oriented with strong analytical and problem solving skills
- Strong time-management and organization skills for coordinating and prioritizing multiple projects and initiatives
- Strong written and oral communications skills with the ability to effectively interface with management and engineering
- Strong team-working and excellent interpersonal skills

DIFFERENTIATORS

- Masters or PhD in related computationally intensive science or engineering field
- Experience with C/C++, parallel programming and knowledge of computer architectures
- Experience working with PCIe form-factor accelerators such as GPUs, DSPs or FPGAs
- Experience in the cloud or automotive space

This position will be based at our office in Palo Alto and could require frequent travel to customer & partner sites.

Graphcore is an equal opportunity workplace and we would like to hear from all qualified applicants. If you have a disability or any special needs that we might need to accommodate, please just let us know.