Title: Deep Learning for Computer Vision – A Cloud-based GPU Deployment and Testing

Abstract

Businesses worldwide are using artificial intelligence to solve their greatest challenges. For example, healthcare professionals use AI to enable more accurate, faster patient diagnoses. Retail companies use it to offer personalized customer shopping experiences. Automakers use it to make personal vehicles, shared mobility, and delivery services safer and more efficient. Deep learning is a powerful AI approach that uses multi-layered artificial neural networks to deliver state-of-the-art accuracy in object detection, speech recognition, and language translation tasks.

This tutorial aims to demonstrate deep learning techniques for various computer vision tasks through a series of hands-on exercises. The participants will work with widely used deep learning tools, frameworks, and workflows to train and deploy neural network models on a fully configured, GPU-accelerated workstation in the cloud. They will train deep learning models from scratch, learning tools and tricks to achieve highly accurate results. They also learn to leverage freely available, state-of-the-art pre-trained models to save time and get their deep learning application up and running quickly. At the end of the tutorial, the participants will have access to additional resources to create new deep-learning applications independently. This tutorial is a collaboration between the NVIDIA Deep Learning Institute and the University of Regina.

Objectives:

1. Learn the fundamental techniques and tools required to train a deep-learning model
2. Gain experience with common deep-learning data types and model architectures
3. Enhance datasets through data augmentation to improve model accuracy
4. Leverage transfer learning between models to achieve efficient results with fewer data and computation
5. Build confidence to take on your project with a modern deep-learning framework

Goals:

The major objective of this tutorial is to provide an opportunity for researchers to deploy their deep learning algorithms on cloud-based GPUs and train and test deep neural networks.

Audience:

Research students, academics, and developers from industry.

Presenter:

Abdul Bais is an Associate Professor in the Faculty of Engineering and Applied Science at the University of Regina. His research interests include deep learning/machine learning, image processing, and computer vision. Abdul Bais's research is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC Alliance and Discovery programs), Mitacs (www.mitacs.ca), the Saskatchewan Ministry of Agriculture, and the AgTech industry. He is a certified instructor with the NVIDIA Deep Learning Institute, a senior member of the IEEE and a Licensed Professional Engineer in Saskatchewan, Canada.