Deep Learning for Autonomous Vehicles—Perception

Artificial intelligence (AI) is taking the automotive industry by storm. Market research projects that by 2025 the automotive AI market will be valued at $11 billion and the installation rate of AI-based systems in new vehicles will rise by 109%. One of the most prominent ways that AI is revolutionizing the industry is through autonomous vehicles. Self-driving cars use camera-based machine vision systems and radar- and lidar-based detection units to perceive, understand, and safely navigate a nearly infinite range of driving scenarios.

In this Deep Learning Institute (DLI) workshop, developers will learn how to optimize performance for self-driving car perception applications such as lane navigation and pedestrian detection. They'll be able to build and train a semantic segmentation neural network to identify objects in videos such as roads, pedestrians, and other vehicles. Then they'll deploy the neural network on the NVIDIA DRIVE™ platform to power autonomous navigation of a vehicle. At the end of the workshop, developers will have the skills needed to build AI applications for a variety of autonomous driving scenarios, including traffic navigation, obstacle avoidance, and robotic ridesharing.

All workshop attendees get access to fully configured, GPU-accelerated servers in the cloud, guidance from a DLI-certified instructor, and the opportunity to network with other attendees. Attendees can earn a certificate to prove subject matter competency and support professional growth.

Prerequisites:	Familiarity with C++ and Python, experience with convolutional neural networks (CNNs). To learn the basics of C++, we suggest this C++ tutorial. To learn the basics of Python, we suggest this Python tutorial. To learn the basics of CNNs, we suggest reading this article on the NVIDIA Developer Zone.
Technologies:	TensorFlow, Keras, NVIDIA® TensorRT™, CUDA® C++, Python, DIGITS, semantic segmentation, deep learning
Price:	Contact us for pricing
Duration:	Approximately 8 hours

Learning Objectives

In this workshop, developers will learn how to:

- Run example code using different GPU memory configurations on the DRIVE platform to determine which is most effective in a variety of use cases
- Compare multiple methods of performance optimization using CUDA and timed-inference test cases on the DRIVE platform
- Create and train a semantic segmentation network to understand automotive scenes by combining a fully convolutional neural network (FCN) "head" with a MobileNets CNN "stem"
- Train a semantic segmentation model with the DIGITS tool using Cityscapes data to demonstrate pixel-wise semantic segmentation of scenes
- Convert Keras and Tensorflow semantic segmentation models into optimized TensorRT models that can be deployed for practical applications on the DRIVE platform
- Deploy and run an optimized TensorRT model on the DRIVE platform to demonstrate the development and deployment workflow for DRIVE applications

Workshop Outline

Introduction (15 mins)	
CUDA on DRIVE AGX (120 mins)	**Learn techniques to improve GPU performance for DRIVE applications through memory management and optimization techniques.** > Test and compare performance tradeoffs for conventional memory, pinned memory, and unified memory types. > Explore mixed-precision computation optimizations in inference. > Optimize performance using CUDA streams and load balancing techniques.
Lunch (60 mins)	
Training Semantic Segmentation for DRIVE (120 mins)	**Explore how to build and train a fully convolutional network (FCN) for semantic segmentation and deploy it to analyze automotive scenes.** > Build a semantic segmentation FCN based on a conventional CNN in TensorFlow. > Prepare a Cityscapes dataset to train an FCN using DIGITS. > Train an FCN in DIGITS and test it using inference to observe the resulting pixel-level semantic segmentation of the scenes.
Break (15 mins)	
Deployment of a Semantic Segmentation Network Using TensorRT (120 mins)	**Learn the TensorRT development workflow for an autonomous driving semantic segmentation use case.** > Optimize a pre-trained semantic segmentation model built with Keras to TensorRT for an embedded system. > Test and compare performance and accuracy across the Keras implementation, TensorRT FP32, and TensorRT INT8. > Build a calibration dataset and deploy the model to the embedded target system for additional performance comparisons.
Assessment and Q&A (15 mins)	

Why DLI Hands-On Training?

> Build deep learning, accelerated computing, and accelerated data science applications for industries such as autonomous vehicles, healthcare, manufacturing, media and entertainment, robotics, smart cities, and more.
> Gain real-world expertise through content designed in collaboration with industry leaders, such as Uber, the Children's Hospital of Los Angeles, Mayo Clinic, and PwC.
> Access course content anywhere, anytime with fully configured, GPU-accelerated servers in the cloud.
> Earn an NVIDIA DLI certificate to demonstrate subject matter competency and support career growth.
> Work with the most widely used, industry-standard software, tools, and frameworks.

Next Steps

Connect with your NVIDIA contact to schedule an onsite workshop for your team, or submit your request at www.nvidia.com/requestdli and the DLI team will be in touch.

Additional Resources

DLI offers other hands-on training and educational resources in AI for autonomous vehicles, including:

> Self-paced, online courses in computer vision, intelligent video analytics, and more at www.nvidia.com/dli
> Instructor-led workshops in multi-GPUs, robotics, intelligent video analytics, and more at www.nvidia.com/dli
> Blogs, webinars, and other resources on AI for autonomous vehicles at www.nvidia.com/drive