

Whether companies are manufacturing semiconductor chips, airplanes, automobiles, smartphones, or food and beverages, quality and throughput are key benefits of optimization. Poor quality and throughput can result in significant operational, financial, and reputational costs. Deep learning-based computer vision technology enables manufacturers to perform automated visual inspection. Compared to traditional visual inspection processes—which are often manual and rules-based—visual inspection AI can improve efficiency, reduce operating costs, and deliver more consistent results.

In this Deep Learning Institute (DLI) workshop, developers will learn how to create an end-to-end hardware-accelerated industrial inspection pipeline to automate defect detection. Using NVIDIA's own real production data set as an example, we'll illustrate how the application can be easily applied to a variety of manufacturing use cases. Developers will also learn to identify and mitigate common pitfalls in deep learning-based computer vision tasks, and be able to deploy and measure the effectiveness of their AI solution.

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 developers, data scientists, and researchers attending the workshop. Attendees can also earn a certificate to prove subject matter competency and support professional growth.

Duration:	Approximately 8 hours
Price:	Contact us for pricing
Prerequisites:	> Experience with Python; basic understanding of data processing and deep learning. > To gain experience with Python, we suggest this Python tutorial. > To get a basic understanding of data processing and deep learning, we suggest DLI's Fundamentals of Deep Learning.
Tools, libraries, and frameworks:	Python, Pandas, DALI, NVIDIA TAO Toolkit, NVIDIA TensorRT™, and NVIDIA Triton™ Inference Server

Learning Objectives

In this workshop, developers will learn how to:

- > Extract meaningful insights from the provided data set using Pandas DataFrame.
- > Apply transfer-learning to a deep learning classification model.
- > Fine-tune the deep learning model and set up evaluation metrics.
- > Deploy and measure model performance.
- > Experiment with various inference configurations to optimize model performance.

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 the Children's Hospital of Los Angeles, Mayo Clinic, PwC, and Uber.
- > Access content anywhere, anytime with a fully configured, GPU-accelerated server 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.

Workshop Outline

Introduction (15 mins)	Meet the instructor > Create an account at courses.nvidia.com/join
Data Exploration and Pre-Processing with DALI (120 mins)	Learn how to extract valuable insights from a data set and pre-process image data for deep learning model consumption > Explore data set with Pandas > Pre-process data with DALI > Assess scope for feasibility testing
Lunch (60 mins)	
Efficient Model Training with TAO Toolkit (120 mins)	Learn how to efficiently train a classification model for the purpose of defect detection using transfer learning techniques > Train a deep learning model with TAO Toolkit > Evaluate the accuracy of the model > Iterate model training to improve accuracy
Break (15 mins)	
Model Deployment for Inference (120 mins)	Learn how to deploy and measure the performance of a deep learning model > Optimize deep learning models with TensorRT > Deploy model with Triton Inference Server > Explore and assess the impact of various inference configurations
Assessment and Q&A (15 mins)	

For the latest DLI workshops and trainings, visit www.nvidia.com/dli

For questions, contact us at nvdl@nvidia.com

© 2022 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, TensorRT, and Triton are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. All other trademarks and copyrights are the property of their respective owners. JUN22

