

Applications of AI for Predictive Maintenance

According to the International Society of Automation, \$647 billion is lost globally each year due to downtime from machine failure. Organizations across manufacturing, aerospace, energy, and other industrial sectors are overhauling maintenance processes to minimize costs and improve efficiency. With artificial intelligence (AI) and machine learning, organizations can apply predictive maintenance to their operation, processing huge amounts of sensor data to detect equipment failure before it happens. Compared to routine-based or time-based preventative maintenance, predictive maintenance gets ahead of the problem and can save a business from costly downtime.

In this Deep Learning Institute (DLI) workshop, developers will learn how to identify anomalies and failures in time-series data, estimate the remaining useful life of the corresponding parts, and map anomalies to failure conditions. They'll learn how to prepare time-series data for AI model training, develop an XGBoost ensemble tree model, build a deep learning model using a long short-term memory (LSTM) network, and create an autoencoder that detects anomalies for predictive maintenance. At the end of the workshop, developers will be able to use AI to estimate the condition of equipment and predict when maintenance should be performed.

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 earn a certificate to prove subject matter competency and support professional growth.

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 for Computer Vision .
Technologies:	Python, TensorFlow, Keras, XGBoost, NVIDIA RAPIDS™, cuDF, LSTM, autoencoders, artificial intelligence, deep learning
Price:	Contact us for pricing
Duration:	Approximately 8 hours

Learning Objectives

In this workshop, developers will learn how to:

- > Use AI-based predictive maintenance to prevent failures and unplanned downtimes
- > Identify key challenges around detecting anomalies that can lead to costly breakdowns
- > Use time-series data to predict outcomes with XGBoost-based machine learning classification models
- > Use an LSTM-based model to predict equipment failure
- > Use anomaly detection with time-series autoencoders to predict failures when limited failure-example data is available

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)	
Training XGBoost Models with RAPIDS for Time Series (120 mins)	<p>Learn how to predict part failures using XGBoost classification on GPUs with cuDF.</p> <ul style="list-style-type: none"> > Prepare real data for efficient GPU ingestion with RAPIDS cuDF. > Train a classification model using GPU-accelerated XGBoost and CPU-only XGBoost. > Compare and discuss performance and accuracy results for XGBoost using CPUs, GPUs, and GPUs with cuDF.
Lunch (60 mins)	
Training LSTM Models Using Keras and TensorFlow for Time Series (120 mins)	<p>Learn how to predict part failures using a deep learning LSTM model with time-series data.</p> <ul style="list-style-type: none"> > Prepare sequenced data for time-series model training. > Build and train a deep learning model with LSTM layers using Keras. > Evaluate the accuracy of the model.
Break (15 mins)	
Training Autoencoders for Anomaly Detection (120 mins)	<p>Learn how to predict part failures using anomaly detection with autoencoders.</p> <ul style="list-style-type: none"> > Build and train an LSTM autoencoder. > Develop and train a 1D convolutional autoencoder. > Experiment with hyperparameters and compare the results of the models.
Assessment and Q&A (15 mins)	

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.

Related Training

If your organization is interested in using AI to detect defects in manufacturing equipment, we recommend the instructor-led workshop [Applications of AI for Anomaly Detection](#). Your team will learn how to apply multiple techniques across accelerated XGBoost, deep learning-based autoencoders, and generative adversarial networks (GANs) to detect anomalies, specifically to identify network intrusions.

Additional Resources

DLI offers other hands-on training and educational resources on AI for manufacturing, including:

- > Self-paced, online courses on deep learning, accelerated computing, and accelerated data science at www.nvidia.com/dli
- > Instructor-led workshops on deep learning for anomaly detection, industrial inspection, multi-GPUs, and more at www.nvidia.com/dli
- > Blogs, webinars, and other resources on AI for manufacturing at www.nvidia.com/industrial