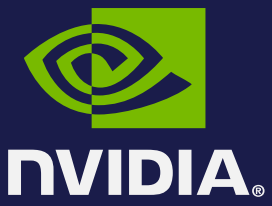

סדנת
NVIDIA AI

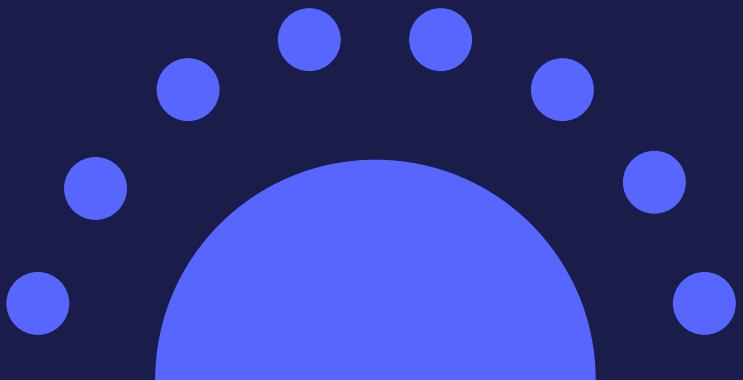

אטלס
האקדמיה הטכנולוגית
של מערכת הביטחון

NVIDIA

על הקורס

סדנה מעמיקה ומעשירה, שנועדה להעניק כלים וידע פורצי דרך בתחום הבינה המלאכותית. הסדנה כוללת חמישה קורסים ייחודיים המתמקדים בתחומים מבוקשים ומתפתחים, ומעניקים הסמכות בינלאומיות מטעם: NVIDIA.

על התוכנית

היכרות עם למידה עמוקה (Getting Started with Deep Learning)
עסקים ברחבי העולם משתמשים בבינה מלאכותית כדי להתמודד עם אתגרים מרכזיים: אנשי רפואה משתמשים בה לאבחון מהיר ומדויק יותר; קמעונאים מציעים חוויות קנייה מותאמות אישית; ויצרני רכב משפרים את הבטיחות והיעילות בתחבורה. ללמידת עומק, שהיא שיטה מתקדמת של בינה מלאכותית המשתמשת ברשתות נוירונים בשכבות, מאפיינת ביצועים מתקדמים במשימות כמו זיהוי אובייקטים, זיהוי דיבור ותרגום שפה. גישה זו מאפשרת למחשבים לזהות דפוסים מורכבים שאינם ניתנים לזיהוי באמצעות תוכנה מסורתית.

מודלים מבוססי ראייה ממוחשבת עם נתונים סינתטיים
(Bootstrapping Computer Vision Models With Synthetic Data)
בקורס זה, נשתמש ב־NVIDIA Omniverse Replicator ובתוסף Omniverse Defect Extension ליצירת נתונים סינתטיים. לאחר מכן, נעבור תהליך של חזרה על מערך הנתונים כדי לאמן רשת נוירונים עמוקה (DNN) לאיתור אובייקטים מטרה (כגון שריטות) בצבעה.

בניית מערכות AI מבוססות לאבטחת סייבר
(Building AI-Based Cybersecurity Pipelines)
שיטות אבטחת סייבר מסורתיות כוללות יצירת חומות הגנה מסביב לתשתיות כדי להגן עליהן מפני פולשים. עם זאת, ככל שחברות מאמצות את תהליך הטרנספורמציה הדיגיטלית שלהן, הן מתמודדות עם ריבוי מכשירים, תוקפים מתוחכמים יותר ורשת נתונים עצומה שיש להגן עליה – דבר שמצריך גישה חלופית. גישה חלופית היא להתייחס לאבטחת סייבר כמערכת מידע נתונים: המטרה היא לזהות ולנטר את כל המשתמשים והפעולות ברשת, כדי שניתן יהיה לזהות אילו עסקאות ואילו פעולות עשויות להיות זדוניות.

פרטים ליצירת קשר | office@atlastech.co.il

אטלס
האקדמיה הטכנולוגית
של מערכת הביטחון

NVIDIA

על התוכנית

פיתוח מהיר של יישומים עם מודלי שפה גדולים
(Rapid Application Development with Large Language Models – LLMs)
בקורס זה נלמד לפתח יישומים מבוססי LLM על ידי חקר קהילת הקוד הפתוח, הכוללת LLMs מאומנים מראש.

בניית סוכני RAG עם מודלי שפה
(Building RAG Agents with LLMs)
סוכנים המופעלים על ידי מודלים שפתיים גדולים (LLMs) הוכיחו יכולות אחזור מרשימות לשימוש בכלים, בחינת מסמכים ותכנון גישות פעולה. קורס זה ילמד כיצד לפרוס מערכת סוכנים בפועל, עם גמישות להרחיב את המערכת כך שתתאים לדרישות של משתמשים ולקוחות.

קהל יעד

הסדנה מתאימה למגוון משתתפים ותספק ידע, כלים, וניסיון מעשי בעולמות הבינה המלאכותית, הראייה הממוחשבת, הסייבר, ופיתוח יישומים מבוססי מודלים מתקדמים.

היקף הקורס

40 שעות
בסיום ההכשרה תינתן תעודת הסמכה בינלאומית של אינבידה

תנאי סף

- **פייתון כללי, בדגש על סינטקס, פונקציות בסיסיות, מערכים ומילונים.**
- **דוקרים ברמת היכרות של הכלי**

האקדמיה הטכנולוגית
של מערכת הביטחון

Day 1 ◗ Fundamentals of Deep Learning

Course Details

Duration: 09:00 – 17:00

Level: Technical – Beginner

Subject: Deep Learning

Language: English

Course Prerequisites: An understanding of fundamental programming concepts in Python 3, such as functions, loops, dictionaries, and arrays; familiarity with Pandas data structures; and an understanding of how to compute a egression line.

Suggested materials to satisfy prerequisites: Python Beginner's Guide.

Technologies: PyTorch, Pandas

Assessment Type: Skills–based coding assessments evaluate students' ability to train a deep learning model to high accuracy.

Certificate: Upon successful completion of the assessment, participants will receive an NVIDIA DLI certificate to recognize their subject matter competency and support professional career growth.

Hardware Requirements: Desktop or laptop computer capable of running the latest version of Chrome or Firefox. Each participant will be provided with dedicated access to a fully configured, GPU–accelerated server in the cloud.

האקדמיה הטכנולוגית
של מערכת הביטחון

About this Course

Businesses worldwide are using artificial intelligence to solve their greatest challenges. Healthcare professionals use AI to enable more accurate, faster diagnoses in patients. Retail businesses 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 tasks such as object detection, speech recognition, and language translation. Using deep learning, computers can learn and recognize patterns from data that are considered too complex or subtle for expert-written software

Learning Objectives

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

Topics Covered

- PyTorch
- Convolutional Neural Networks (CNNS)
- Data Augmentation
- Transfer Learning
- Natural Language Processing

Course Outline

The below is a suggested timeline for the course. Please work with the instructor to find the best timeline for your session.

Introduction	• Meet the instructor. • Create an account at courses.nvidia.com/join
The Mechanics of Deep Learning	**Explore the fundamental mechanics and tools involved in successfully training deep neural networks:** • Train your first computer vision model to learn the process of training. • Introduce convolutional neural networks to improve accuracy of predictions in vision applications. • Apply data augmentation to enhance a dataset and improve model generalization.

האקדמיה הטכנולוגית
של מערכת הביטחון

Course Outline

Pre-trained Models and Large Language Models	**Leverage pre-trained models to solve deep learning challenges quickly. Train recurrent neural networks on sequential data:** • Integrate a pre-trained image classification model to create an automatic doggy door. • Leverage transfer learning to create a personalized doggy door that only lets in your dog. • Use a Large Language Model (LLM) to answer questions based on provided text.
Final Project: Object Classification	**Apply computer vision to create a model that distinguishes between fresh and rotten fruit:** • Create and train a model that interprets color images. • Build a data generator to make the most out of small datasets. • Improve training speed by combining transfer learning and feature extraction.
Final Project: Object Classification	• Discuss advanced neural network architectures and recent areas of research where students can further improve their skills.

האקדמיה הטכנולוגית
של מערכת הביטחון

Course Outline

Final Project:

Object Classification

Final Review

- Integrate a pre-trained image classification model to create an automatic doggy door.
- Leverage transfer learning to create a personalized doggy door that only lets in your dog.
- Use a Large Language Model (LLM) to answer questions based on provided text.

האקדמיה הטכנולוגית
של מערכת הביטחון

Day 2 Computer Vision for Industrial Inspection

Course Details

Duration: 09:00 –17:00

Level: Technical – Intermediate

Subject: Deep Learning

Language: English

About this Course

Whether companies are manufacturing semiconductor chips, airplanes, automobiles, smartphones, or food or 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 resultset to train a DNN to find target objects (scratches) in a scene.

Learning Objectives

- Extract meaningful insights from the provided data set using Pandas DataFrame.

האקדמיה הטכנולוגית
של מערכת הביטחון

Learning Objectives

- 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.

Topics Covered

- 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 Getting Started with Deep Learning.
- **Technologies:** Python, Pandas, DALI, NVIDIA TAO Toolkit, NVIDIA TensorRT™, and NVIDIA Triton™ Inference Server
- **Certificate:** Upon successful completion of the assessment, participants will receive an NVIDIA DLI certificate to recognize their subject matter competency and support professional career growth.

האקדמיה הטכנולוגית
של מערכת הביטחון

Course Outline

Introduction	• Meet the instructor. • Create an account at courses.nvidia.com/join
Data Exploration and Pre-Processing with DALI	**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
Efficient Model Training with TAO Toolkit	**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
Model Deployment for Inference	**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

האקדמיה הטכנולוגית
של מערכת הביטחון

Day 3 ◗ Building AI-Based Cybersecurity Pipelines

Course Details

Duration: 09:00 – 17:00

Subject: Deep Learning

Language: English

Technologies: NVIDIA Morpheus, NVIDIA Triton Inference Server, RAPIDS™, CLX, Helm, Kubernetes

Assessment Type: Skills-based coding assessments evaluate students' ability to build end-to-end Morpheus cybersecurity pipelines. Multiple-choice questions test students' understanding of the Morpheus-related concepts presented in the workshop.

Certificate: Upon successful completion of the assessment, participants will receive an NVIDIA DLI certificate to recognize their subject matter competency and support professional career growth.

Hardware Requirements: Desktop or laptop computer capable of running the latest version of Chrome or Firefox. Each participant will be provided with dedicated access to a fully configured, GPU-accelerated workstation in the cloud.

האקדמיה הטכנולוגית
של מערכת הביטחון

About this Course

Traditional cybersecurity methods include creating barriers around your infrastructure to protect it from intruders. However, as enterprises continue to digitally transform, they're faced with a proliferation of devices, more sophisticated cybersecurity attacks, and an incredibly vast network of data to protect—which means new cybersecurity methodologies must be explored. An alternative approach is to address cybersecurity as a data science problem: Aim to better understand all the users and activities across your network so that you can identify which transactions are typical and which are potentially nefarious.

Learning Objectives

- Build Morpheus pipelines to process and perform AI–based inference on massive amounts of data for cybersecurity use cases in real time
- Utilize several AI models with a variety of data input types for tasks like sensitive information detection, anomalous behavior profiling, and digital fingerprinting
- Leverage key components of the Morpheus AI framework, including the Morpheus SDK and command–line interface (CLI), and NVIDIA Triton™ Inference

Course Outline

Introduction

- Meet the instructor.
- Create an account at courses.nvidia.com/join

האקדמיה הטכנולוגית של מערכת הביטחון

Course Outline

An Overview of the NVIDIA Morpheus AI Framework	**Explore the fundamental mechanics and tools involved in successfully training deep neural networks:** • Understand the need for AI-based cybersecurity. • Learn about the components of the Morpheus framework. • Discover how institutions are building solutions with Morpheus.
Morpheus Pipeline Construction	• Get an overview of the Morpheus SDK and CLI. • Learn about pipeline types and commands. • Learn about data input/output (IO) and processing.
Inference in Morpheus Pipelines	• Get an overview of NVIDIA Triton Inference Server. • Understand how models are deployed. • Explore a sensitive-information-detection pipeline.
Case Study: AI-Based Machine Logs Parsing at Splunk	• Apply your understanding to a real-world example.
Digital Fingerprinting Pipeline	• Use the Morpheus autoencoder pipeline. • Discover compromised credentials.

האקדמיה הטכנולוגית
של מערכת הביטחון

Course Outline

Time Series Analysis	• Apply time series analysis within a Morpheus pipeline. • Combine time series analysis with digital fingerprinting.
Case Study: Cybersecurity Flyaway Kit at Booz Allen Hamilton	• Apply your understanding to a real-world example.
Assessment 1: Test Your Understanding	• Assess your conceptual understanding of the topics covered.
Assessment 2: Practical Demonstration	• Build an end-to-end Morpheus pipeline to identify a cybersecurity breach.
Wrap Up	• Get resources for further development with Morpheus. • Provide feedback on the workshop.

האקדמיה הטכנולוגית
של מערכת הביטחון

Day 4 Rapid Application Development with Large Language Models (LLMs)

Course Details

Duration: 09:00 – 17:00

Level: Technical – Beginner

Subject: Generative AI/LLM

Language: English

Course Prerequisites:

· Introductory deep learning, with comfort with PyTorch and transfer learning preferred. Content covered by DLI's Getting Started with Deep Learning or Fundamentals of Deep Learning courses, or similar experience is sufficient.

· Intermediate Python experience, including object–oriented programming and libraries. Content covered by Python Tutorial (w3schools.com) or similar experience is sufficient.

Tools, libraries, frameworks used: Python, PyTorch, HuggingFace, Transformers, LangChain, and LangGraph

About this Course

Recent advancements in both the techniques and accessibility of large language models (LLMs) have opened up unprecedented opportunities to help businesses streamline their operations, decrease expenses, and increase productivity at scale.

About this Course

Additionally, enterprises can use LLM-powered apps to provide innovative and improved services to clients or strengthen customer relationships. For example, enterprises could provide customer support via AI companions or use sentiment analysis apps to extract valuable customer insights. In this course you will gain a strong understanding and practical knowledge of LLM application development by exploring the open-sourced ecosystem including pretrained LLMs, enabling you to get started quickly in developing LLM-based applications.

Learning Objectives

By participating in this workshop, you will:

- Find, pull in, and experiment with the HuggingFace model repository and Transformers API.
- Use encoder models for tasks like semantic analysis, embedding, question-answering, and zero-shot classification.
- Work with conditioned decoder-style models to take in and generate interesting data formats, styles, and modalities.
- Kickstart and guide generative AI solutions for safe, effective, and scalable natural data tasks.
- Explore the use of LangChain and LangGraph for orchestrating data pipelines and environment-enabled agents.

האקדמיה הטכנולוגית
של מערכת הביטחון

Topics Covered

The workshop covers large language models from beginning to end, starting with fundamentals of transformers, progression into foundational large language models, and finishing in model/agentic orchestration. Each of these sections is designed to equip participants with the knowledge and skills necessary to progress further in developing useful LLM-powered applications.

Course Outline

The table below is a suggested timeline for the course. Please coordinate with the instructor for the best timeline for your session.

Course Introduction	• Overview of workshop topics and schedule. • Introduction to HuggingFace and Transformers. • Discuss how LLMs can enhance enterprise applications.
Transformers and LLMs	• Introduce and motivate the transformer-style architecture from deep learning first principles. • Understand input-output processing with tokenizers, embeddings, and attention mechanisms.
Task-Specific Pipelines	• Profile encoder models for different NLP tasks where they are most useful.

האקדמיה הטכנולוגית של מערכת הביטחון

Course Outline

Task-Specific Pipelines	• Investigate the use of lightweight models for natural language embedding, classification, subsetting, and zero-shot prediction.
Seq2Seq with Decoders	• Introduce GPT-style decoder models for sequence generation and autoregressive tasks. • Apply encoder-decoder architectures for applications like machine translation and few-shot task completion.
Multimodal Architectures	• Integrate different data modalities (text, images, audio) into LLM workflows. • Explore multimodal models like CLIP for cross-modal learning, visual language models for image question-answering, and diffusion models for text-guided image generation.
Scaling Text Generation	• Explore LLM inference challenges and deployment strategies, including optimized server deployments. • Incorporate LLMs into interesting applications that can scale to larger repositories and user bases.

האקדמיה הטכנולוגית
של מערכת הביטחון

Course Outline

Orchestration and Agentics	• Introduce LangChain for LLM orchestration and agentic workflows. • Investigate use of agentics and tool-calling for integrating natural language with standard applications and data.
Final Assessment	• Build an LLM-based application integrating text generation, multimodal learning, and agentic orchestration.
Review and Wrap-up	• Review key learnings and answer final questions. • Earn a certificate upon successful completion. • Complete the workshop survey.

האקדמיה הטכנולוגית
של מערכת הביטחון

Day 5 Building RAG Agents with LLMs

Course Details

Duration: 09:00 – 17:00

Level: Technical – Intermediate

Subject: Generative AI/LLM

Language: English

Course Prerequisites:

· Introductory deep learning knowledge,
with comfort with PyTorch and transfer learning preferred.

· Intermediate Python experience,
including object–oriented programming and libraries.

About this Course

Agents powered by large language models (LLMs) are quickly gaining popularity as people are finding new capabilities and opportunities to greatly improve their productivity. An especially powerful recent development has been the popularization of retrieval–based LLM systems that can hold informed conversations by using tools, looking at documents, and planning their approaches. These systems are fun to experiment with and offer unprecedented opportunities to make life easier, but they also require many queries to large deep learning models and need to be implemented efficiently.

האקדמיה הטכנולוגית
של מערכת הביטחון

About this Course

This course will observe how you can deploy an agent system in practice and scale up your system to meet the demands of users and customers. Along the way, you'll learn advanced LLM orchestration techniques for internal reasoning, dialog management, tooling, and retrieval.

Learning Objectives

This course will observe how you can deploy an agent system in practice and scale up your system to meet the demands of users and customers. Along the way, you'll learn advanced LLM orchestration techniques for internal reasoning, dialog management, tooling, and retrieval.

- Compose an LLM system that can interact predictably with a user by leveraging internal and external reasoning components.
- Design a dialog management and document reasoning system that maintains state and coerces information into structured formats.
- Leverage embedding models for efficient similarity queries for content retrieval and dialog guardrailing.
- Implement, modularize, and evaluate a RAG agent that can answer questions about the research papers in its dataset without any fine-tuning.

Topics Covered

The workshop includes topics such as LLM Inference Interfaces, Pipeline Design with LangChain, Gradio, and LangServe, Dialog Management with Running States,

האקדמיה הטכנולוגית
של מערכת הביטחון

Topics Covered

Working with Documents, Embeddings for Semantic Similarity and Guardrailing, and Vector Stores for RAG Agents. Each of these sections is designed to equip participants with the knowledge and skills necessary to develop and deploy advanced LLM systems effectively.

Course Outline

- Introduction to the workshop and setting up the environment.
- Exploration of LLM inference interfaces and microservices.
- Designing LLM pipelines using LangChain, Gradio, and LangServe.
- Managing dialog states and integrating knowledge extraction.
- Strategies for working with long-form documents.
- Utilizing embeddings for semantic similarity and guardrailing.
- Implementing vector stores for efficient document retrieval.
- Evaluation, assessment, and certification.