DLI Generative AI Teaching Kit
Latest Release Note (Feb. 25th, 2025)

- Updated online syllabus with new content below

- Addition of module lab solutions for new modules below. **Note: for educators with an existing git clone obtained before Feb. 25th, 2025, we strongly recommend performing a fresh clone instead of a git pull to obtain the updated repo.**

- **New! Module 2: Word Embeddings, Tokens, and NLP**
    - **New!** Lecture Slide Deck 2.1 - Natural Language Processing and Language Models
    - **New!** Lecture Slide Deck 2.2 - Tokens and Tokenization
    - **New!** Lecture Slide Deck 2.3 - Word Embeddings
    - **New!** Knowledge Check 2.1 - Natural Language Processing and Language Models
    - **New!** Knowledge Check 2.2 - Tokens and Tokenization
    - **New!** Knowledge Check 2.3 - Word Embeddings
    - **New!** Demo 2.1 - Natural Language Processing and Language Models
    - **New!** Demo 2.2 - Tokens and Tokenization
    - **New!** Demo 2.3 - Word Embeddings
    - **New!** Lab 2 - Word Embeddings, Tokens, and NLP

- **New! Module 3: Large Language Models and the Transformer**
    - **New!** Lecture Slide Deck 3.1 - Language Models and Attention
    - **New!** Lecture Slide Deck 3.2 - The Transformer Architecture
    - **New!** Lecture Slide Deck 3.3 - Autoregressive Training Transformer LLMs
    - **New!** Knowledge Check 3.1 - Language Models and Attention
    - **New!** Knowledge Check 3.2 - The Transformer Architecture
    - **New!** Knowledge Check 3.3 - Autoregressive Training Transformer LLMs
    - **New!** Demo 3.1 - Language Models and Attention
    - **New!** Demo 3.2 - The Transformer Architecture
    - **New!** Demo 3.3 - Autoregressive Training Transformer LLMs
    - **New!** Lab 3 - Large Language Models and the Transformer
    - **New!** DLI Online Course: Rapid Application Development with Large Language Models (LLMs)

- **New! Module 4: LLM Scaling Laws and LLM Families**
    - **New!** Lecture Slide Deck 4.1 - LLM Architecture Variants
    - **New!** Lecture Slide Deck 4.2 - Scaling Laws in Training LLMs
    - **New!** Lecture Slide Deck 4.3 - Training Data for Larger LLMs
    - **New!** Knowledge Check 4.1 - LLM Architecture Variants
    - **New!** Knowledge Check 4.2 - Scaling Laws in Training LLMs
    - **New!** Knowledge Check 4.3 - Training Data for Larger LLMs
    - **New!** Demo 4.1 - LLM Architecture Variants

- • **New!** Demo 4.2 - Scaling Laws in Training LLMs
- • **New!** Demo 4.3 - Training Data for Larger LLMs
- • **New!** Lab 4 - LLM Scaling Laws and LLM Families

- - **New! Module 5: Multimodal Learning and its Applications**
  - • **New!** Lecture Slide Deck 5.1 - Introduction to Multimodal Models
  - • **New!** Lecture Slide Deck 5.2 - Vision Transformer and CLIP
  - • **New!** Lecture Slide Deck 5.3 - PerceiverIO and more generalized Multimodal Models
  - • **New!** Knowledge Check 4.1 - Introduction to Multimodal Models
  - • **New!** Knowledge Check 4.2 - Vision Transformer and CLIP
  - • **New!** Knowledge Check 4.3 - PerceiverIO and more generalized Multimodal Models
  - • **New!** Demo 5.1 - Introduction to Multimodal Models
  - • **New!** Demo 5.2 - Vision Transformer and CLIP
  - • **New!** Demo 5.3 - PerceiverIO and more generalized Multimodal Models
  - • **New!** Lab 5 - Multimodal Learning and its Applications
- -
- - **New! Module 7: Model Training (Pre-Training, Instruction Following, and PEFT)**
  - • **New!** Lecture Slide Deck 7.1 - Pre-training, continued pre-training, and task training
  - • **New!** Lecture Slide Deck 7.2 - Chat preparation: Reinforcement Learning with Human Feedback
  - • **New!** Lecture Slide Deck 7.3 - Parameter-Efficient Fine Tuning Methods
  - • **New!** Demo 7.1 - Pre-training, continued pre-training, and task training
  - • **New!** Demo 7.2 - Chat preparation: Reinforcement Learning with Human Feedback
  - • **New!** Demo 7.3 - Parameter-Efficient Fine Tuning Methods
  - • **New!** Lab 7 - Model Training (Pre-Training, Instruction Following, and PEFT)

- - **New! Module 9: Scaling Model Training to Distributed Workloads**
  - • **New!** Lecture Slide Deck 9.1 - Distributed Training with Data and Model Parallelism Strategies
  - • **New!** Lecture Slide Deck 9.2 - Challenges and Libraries for Distributed Training
  - • **New!** Lecture Slide Deck 9.3 - Scaling Inference and Deployment
  - • **New!** Demo 9.1 - Distributed Training with Data and Model Parallelism Strategies
  - • **New!** Demo 9.2 - Challenges and Libraries for Distributed Training
  - • **New!** Demo 9.3 - Scaling Inference and Deployment
  - • **New!** Lab 9 - Scaling Model Training to Distributed Workloads
  - • **New!** DLI Online Course: Introduction to NVIDIA NIM™ Microservices
  - • **New!** DLI Online Course: Sizing LLM Inference Systems