Modern Deep Learning with PyTorch

07-10, 13:30–17:30 (America/Chicago), Classroom 106

Schedule

1. Introduction to Deep Learning (1:30 - 2:00 pm)
2. Understanding the PyTorch API (2:00 - 2:30 pm)
3. Training Deep Neural Networks (2:30 - 3:00 pm)

10 Min Break

4. Accelerating PyTorch Model Training (3:10 - 3:45 pm)
5. Organizing PyTorch Code (3:45 - 4:15 pm)
6. More Tips and Techniques (4:15 - 4:45 pm)

10 Min Break

7. Finetuning LLMs (4:55 - 5:25 pm)
8. Wrap Up & Questions (5:25 - 5:30 pm)

Workshop Goal:

Understand and Use PyTorch Confidently,
Including Advanced Features

More LLMs?

Finetuning Large Language Models

An introduction to the core ideas and approaches

https://magazine.sebastianraschka.com/p/finetuning-large-language-models

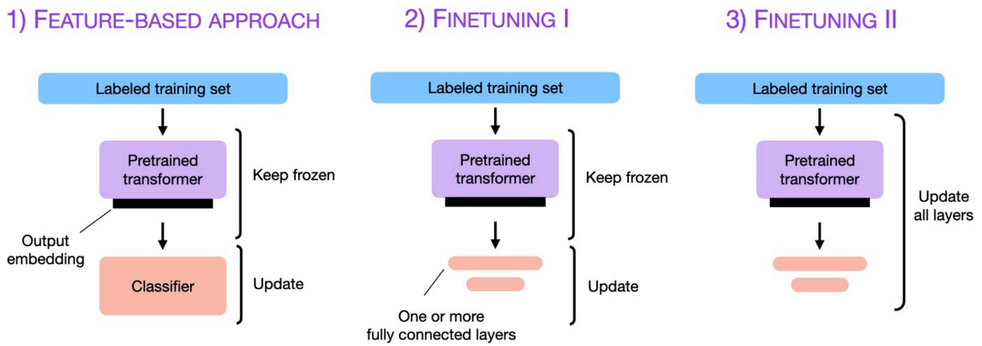

Understanding Encoder And Decoder LLMs

https://magazine.sebastianraschka.com/p/understanding-encoder-and-decoder

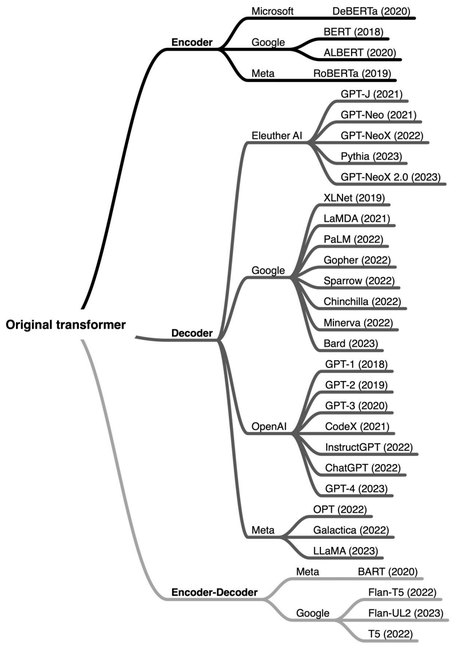

More LLMs?

1. Download the data and generate an instruction tuning dataset:

```
python scripts/prepare_alpaca.py
```

2. Run the finetuning script

For example, you can either use

Adapter (Zhang et al. 2023):

```
python finetune/adapter.py
```

or Adapter v2 (Gao et al. 2023):

```
python finetune/adapter_v2.py
```

or LoRA (Hu et al. 2021):

```
python finetune/lora.py
```

https://github.com/Lightning-AI/lit-gpt

About

Hackable implementation of state-of-the-art open-source LLMs based on nanoGPT. Supports flash attention, Int8 and GPTQ 4bit quantization, LoRA and LLaMA-Adapter fine-tuning, pre-training. Apache 2.0-licensed.

Readme

Apache-2.0 license

Computer Vision?

Ahead of AI #10: State of Computer Vision 2023

Large language model development (LLM) development is still happening at a rapid pace. At the same time, leaving AI regulation...

JUL 6 • SEBASTIAN RASCHKA, PHD

https://magazine.sebastianraschka.com/p/ahead-of-ai-10-state-of-computer

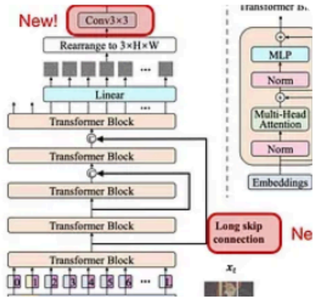

Optimizing Memory Usage for Training LLMs and Vision Transformers in PyTorch

Read More

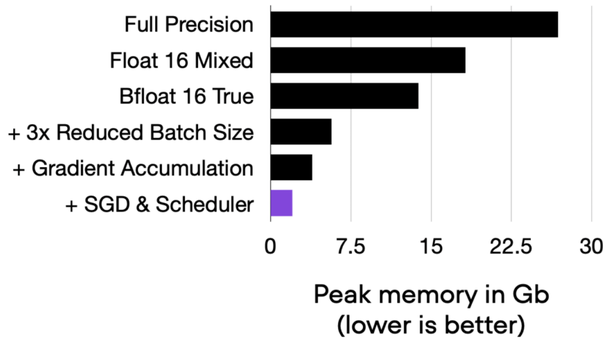

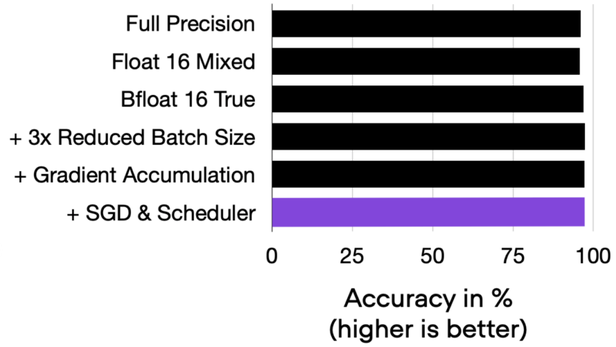

https://lightning.ai/pages/community/tutorial/pytorch-memory-vit-llm/

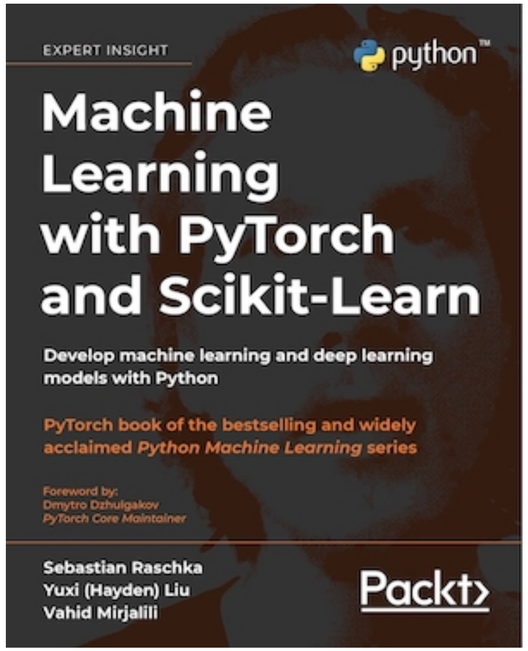

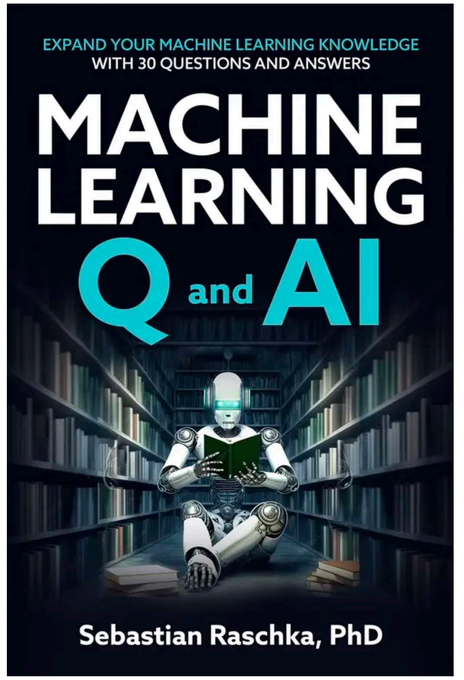

https://sebastianraschka.com/books/

Contact

@rasbt sebastianraschka

sebastian@lightning.ai

https://sebastianraschka.com

https://lightning.ai

Code & slides

https://github.com/scipy2023-deeplearning