

Tutorial on Large Language Models (LLM)

LLM basics

1. LLM Slides by Dr. Kevin Duh:
<https://www.cs.jhu.edu/~kevinduh/a/llm-tutorial/2406-JSALT-LLM.pdf>
2. Additional materials:
 - a. Stanford course CS224N: Natural Language Processing with Deep Learning (Neural Network Basics): <https://web.stanford.edu/class/cs224n>

Post training (finetuning, RL, test-time scaling)

1. Overview of finetuning methods:
https://github.com/chrisvdweth/selene/blob/master/notebooks/llm_model_fine_tuning_overview.ipynb
2. Reinforcement Learning (RL) slides by Prof. Daniel Khashabi:
<https://self-supervised.cs.jhu.edu/sp2025/files/slides/14-15.alignment.pdf>
3. Test-time scaling slides by Prof. Daniel Khashabi:
 - a. <https://self-supervised.cs.jhu.edu/sp2025/files/slides/15-16.inference-scaling.pdf>
4. Additional materials:
 - a. Stanford course CS336: Language Modeling from Scratch (especially Lectures 15-17 Alignment - SFT/RLHF/RL): <https://stanford-cs336.github.io/spring2025/>
 - b. Stanford course CS 329X: Human-Centered LLMs:
<https://web.stanford.edu/class/cs329x/>
 - c. Hands-on examples:
 - i. LLM fine-tuning examples using the transformers and peft packages:
https://github.com/chrisvdweth/selene/blob/master/notebooks/llm_model_finetuning_lora_hf_kidsqa.ipynb
 - ii. Nanochat (a simple pipeline for training LLMs with ~\$100):
<https://github.com/karpathy/nanochat>
 - iii. Documentation for PEFT (Parameter-Efficient Fine-Tuning for LLMs):
<https://huggingface.co/docs/peft/en/index>
 - iv. Documentation for verl (an RL training framework for LLMs):
<https://verl.readthedocs.io/en/latest/>
 - v. A dataset preparation pipeline:
https://wandb.ai/capecape/alpaca_ft/reports/How-to-Fine-Tune-an-LLM-Part-1-Preparing-a-Dataset-for-Instruction-Tuning--Vmlldzo1NTcxNzE2

Tinker guide

1. Quick start
 - a. Tinker docs: <https://tinker-docs.thinkingmachines.ai/>
 - b. Tinker Cookbook: <https://github.com/thinking-machines-lab/tinker-cookbook>
2. A minimal testing example (with a very small cost)
 - a. SL: <https://tinker-docs.thinkingmachines.ai/supervised-learning/sl-basic>
 - b. RL: <https://tinker-docs.thinkingmachines.ai/rl/rl-basic>
3. Simple cost guidance: <https://www.datacamp.com/tutorial/tinker-tutorial>