

Deep Learning: Modern Research Topics

An advanced course on reading and evaluating modern deep-learning research, from efficient systems to reliable deployment.

COURSE CONTENT

Research practice and tooling	Fair baselines, ablations, uncertainty reporting, reproducibility, and PyTorch-centered framework concepts.
Training and inference efficiency	Compute and memory bottlenecks, profiling, quantization, pruning, distillation, and efficient serving.
Sequence models and LLMs	Efficient attention, state-space models, long-context modeling, grounding, and faithful generation.
Reliability and robustness	Calibration, selective prediction, distribution shift, OOD detection, adversarial robustness, and conformal prediction.
Causal reasoning	Causal graphs, confounding, interventions, identification, and counterfactuals.
Representation and multimodal learning	Contrastive and non-contrastive learning, transfer, parameter-efficient fine-tuning, cross-modal alignment, and generation.
Continual and test-time adaptation	Catastrophic forgetting, replay, constrained updates, and adaptation to changing deployment data.

LEARNING FOCUS

Read research claims critically and connect methods to their assumptions, evidence, computational costs, and failure modes.

Lectures and weekly research-paper presentations; midterm and final exams.