[Fall 2025] ECE 5290/7290 and ORIE 5290
Distributed Optimization for Machine Learning and AI

Course Project Description

Objective of This Project

The course project is a cornerstone of this class, designed to extend your understanding of distributed optimization through hands-on research and critical analysis. The project encourages both creativity and rigor, allowing you to apply course concepts to real-world or theoretical problems.

Key learning goals:

- Deepen your understanding of distributed optimization methods for large-scale machine learning.
- Develop the ability to read, critique, and synthesize modern research papers.
- Strengthen your skills in clear technical writing and professional presentation.
- Gain experience conducting original research or educational synthesis.

Project Options and Expectations

You may choose between two equally rigorous but differently oriented project formats:

(1) Educational Project (recommended for 5290 students). Educational projects are ideal for students who enjoy explaining, organizing, and teaching technical concepts. You will act as a "mini instructor," preparing detailed and self-contained lecture notes that explain a specific topic related to (but not directly covered in) this course.

Deliverables and expectations:

- Develop a comprehensive set of lecture slides (15–20 pages) written in a clear, pedagogical style using LaTeXand the template provided by the instructor.
- Include at least one empirical example or small experiment that demonstrates a key algorithmic idea.
- Design a short problem set (1–2 problems) with solutions to reinforce understanding.
- Prepare a 3-minute highlight presentation assuming that you want to attract future students to learn the topic or future employers to adopt this technology.

Examples of educational topics:

- Decentralized optimization and consensus algorithms
- Federated learning and communication-efficient SGD

- Primal-dual methods and ADMM in large-scale learning
- Low-rank adaptation (LoRA) for large models
- Memory and latency analysis of LLM inference and fine-tuning

What makes a strong educational project? A high-quality educational project will clearly explain the mathematical intuition, use figures or examples to illustrate key ideas, and design a problem set that deepens conceptual understanding. Clarity, structure, and accessibility are valued as much as technical depth.

(2) Research Project (recommended for advanced 7290 students). Research projects allow students to engage more deeply with cutting-edge problems in distributed optimization, federated learning, LLM inference and fine-tuning or scalable AI systems. Projects may explore theoretical questions, propose new algorithms, or adapt existing methods in innovative ways.

Deliverables and expectations:

- Conduct a small but meaningful piece of research (theoretical or empirical).
- Write a concise workshop-style report (max 5 pages excluding references and appendix); see NeurIPS 2025 LaTeX style file.
- Develop reproducible code and experiments shared via a GitHub repository.
- Present findings in a 20-minute "NeurIPS oral"-style talk.

Examples of research directions:

- New communication-efficient distributed optimization algorithms
- Empirical analysis of fairness or robustness in federated learning
- Theoretical study of convergence trade-offs in local SGD or ADMM
- Parameter-efficient optimization for large foundation models
- Multi-objective optimization or preference-guided learning

What makes a strong research project? A strong project demonstrates curiosity, originality, and rigor—showing that you not only understand existing methods but also attempt to extend, compare, or critique them. Novelty may come from a new theoretical insight, a practical improvement, or a well-designed empirical study.

Deliverables and Deadlines

(1) Project Proposal (5%)

- **Due: November 1 (via Gradescope)** Submit a concise (max one page) project proposal.
  - For research projects: briefly describe the problem, method, and motivation.
  - For educational projects: include a detailed lecture outline and an example problem draft.

(2) Final Deliverables (25%)

- **Due: December 13 (via Gradescope submission link)** Each project must include:
  (a) A GitHub repository containing reproducible code, scripts, or notebooks (if applicable).
  (b) A presentation slide deck for an educational project (in both PDF and LaTeX formats).
  (c) A polished, workshop-style report for an educational project (in PDF format, up to 6 pages).

(3) In-Class Presentations

- **ECE/ORIE 5290 (Master's):** Deliver a 3-minute "NeurIPS Spotlight"-style presentation on **December 1 or 3**, followed by a short Q&A. Presentations will be evaluated by peer review for clarity, depth, and delivery.
- **ECE 7290 (Ph.D.):** Deliver a 20-minute "NeurIPS Oral"-style presentation on **December 3 or 8**, followed by Q&A and peer review. Students are encouraged to present theoretical insights, experimental validation, or ablation analyses.

Suggested Project Topics

The topics below are examples intended to spark inspiration; you are encouraged to propose your own ideas, provided they align with distributed optimization, federated learning, or scalable AI systems.

- **Educational Topics:**
  - Federated optimization and privacy-preserving distributed learning
  - Physics-informed or graph-based neural networks
  - Low-rank adaptation (LoRA) and parameter-efficient fine-tuning
  - Distribution shifts and domain adaptation in distributed learning
  - Hyperparameter optimization for large-scale training
- **Research Topics:**
  - Extending a recent NeurIPS, ICML, or ICLR paper
  - Designing new methods for decentralized or asynchronous optimization
  - Developing algorithms for low-rank or sparse tensor decomposition
  - Exploring fairness and robustness in federated learning
  - Investigating adversarial or trustworthy optimization in large models
  - AI alignment and preference-guided optimization in LLMs

Collaboration policy: Projects must be done individually, not in groups.

Tip: Choose a topic that genuinely excites you. A well-executed small project that demonstrates understanding and creativity will be valued more than an overly ambitious but incomplete one.

Course Reading List

A. Resource Efficiency

- Mishchenko, K., Malinovsky, G., Stich, S., and Richtárik, P. *ProxSkip: Yes! Local Gradient Steps Provably Lead to Communication Acceleration! Finally!* In *ICML*, 2022. [PMLR link]
- Karimireddy, S. P., Rebjock, Q., Stich, S., and Jaggi, M. *Error Feedback Fixes SignSGD and Other Gradient Compression Schemes.* In *ICML*, 2019. [PMLR link]
- Demidovich, Y., Malinovsky, G., Shulgin, E., and Richtárik, P. *MAST: Model-Agnostic Sparsified Training.* In *ICLR*, 2025. arXiv preprint.

B. Label Efficiency

- Wang, L., Zhang, K., Li, Y., Tian, Y., and Tedrake, R. *Does Learning from Decentralized Non-IID Unlabeled Data Benefit from Self-Supervision?* In *ICLR*, 2022. [OpenReview link]
- Lu, N., Wang, Z., Li, X., Niu, G., Dou, Q., and Sugiyama, M. *Federated Learning from Only Unlabeled Data with Class-Conditional Sharing Clients.* In *ICLR*, 2021. [OpenReview link]
- Alam, S., Liu, L., Yan, M., and Zhang, M. *FedRolex: Model-Heterogeneous Federated Learning with Rolling Sub-Model Extraction.* In *NeurIPS*, 2022. [NeurIPS link]

C. Efficient Analog Training

- Gokmen, T., and Haensch, W. *Algorithm for Training Neural Networks on Resistive Device Arrays. Frontiers in Neuroscience*, vol. 14, 2020. [Frontiers link]
- Wright, L. G., et al. *Deep Physical Neural Networks Trained with Backpropagation. Nature*, vol. 601, no. 7894, pp. 549–555, 2022. [Nature link]
- Wu, Z., et al. *Analog In-memory Training on General Non-ideal Resistive Elements: The Impact of Response Functions.* In *NeurIPS*, 2025. [arXiv]

D. Generalization and Fairness

- Zhu, T., He, F., Zhang, L., et al. *Topology-Aware Generalization of Decentralized SGD.* In *ICML*, 2022. [PMLR link]
- Ray Chaudhury, B., Li, L., Kang, M., Li, B., and Mehta, R. *Fairness in Federated Learning via Core-Stability.* In *NeurIPS*, 2022. [NeurIPS link]
- Lei, Y., Sun, T., and Liu, M. *Stability and Generalization for Minibatch SGD and Local SGD.* arXiv preprint, 2023.
- Chang, H., and Shokri, R. *Bias Propagation in Federated Learning.* In *ICLR*, 2022. [OpenReview link]

E. Distributed inference and training with LLMs

- Borzunov, A., Ryabinin, M., Chumachenko, A., et al. *Distributed Inference and Fine-Tuning of Large Language Models Over The Internet.* In *NeurIPS*, 2023. arXiv preprint

- Liu, Z., Wang, J., Dao, T., Zhou, T., Yuan, B., Song, Z., Shrivastava, A., Zhang, C., Tian, Y., Ré, C., and Chen, B. *Deja Vu: Contextual Sparsity for Efficient LLMs at Inference Time.* In *ICML*, 2023. [PMLR link]

- C. Hou, et al. *Pre-text: Training language models on private federated data in the age of LLMs* In *ICML*, 2023. arXiv.

- Wang, B., Zhang, Y. J., Cao, Y., Li, B., McMahan, H. B., Oh, S., and Zaheer, M. *Can Public Large Language Models Help Private Cross-Device Federated Learning?* In *ICML workshop*, 2023. arXiv