

CSE 8803-SRM Fall 2026

Course Information

- **Instructor:** Bo Dai (bodai@cc.gatech.edu)
- **Course Prefix and Number:** CSE 8803 SRM
- **Term:** Fall 2026

Course Description

This course develops a principled, unified foundation for representation learning through the lens of spectral decomposition. Starting from the sufficiency of spectral representations for downstream tasks, we systematically build a framework that connects classical component analysis methods (PCA, CCA, Laplacian Embedding) to modern self-supervised learning algorithms (SimCLR, BYOL, CLIP, DINO) and extends to applications in reinforcement learning, causal inference, and controllable generation.

The course follows the theoretical backbone of "Spectral Ghost in Representation Learning" (Dai, Li & Schuurmans, 2026), supplemented by the RL-specific treatment in "Spectral Representation-based Reinforcement Learning" (Gao, Sun, Li, Schuurmans & Dai, 2026).

Prerequisites:

- solid knowledge of probability, statistics, calculus, and linear algebra;
- basic knowledge of machine learning;
- solid programming skills, preferably in Python.

Course Learning Outcomes

- Understand why spectral representations are sufficient and effective for a variety of downstream prediction tasks.
- Master the unified spectral framework that categorizes SSL algorithms by their parametrization: direct, energy-based, latent-variable, and nonlinear.
- Derive and analyze the learning objectives of major SSL methods (SimCLR, BYOL, Barlow Twins, VICReg, CLIP, DINO, etc.) from spectral principles.
- Extend spectral representations to multi-modal settings and understand modern vision-language models through this lens.
- Apply spectral representations to reinforcement learning: spectral decomposition of transition operators, Q-value parameterization, and provably efficient exploration.
- Connect modern deep SSL to classical manifold learning and component analysis.
- Critically read, present, and extend research papers in representation learning.

Required Course Materials

Primary References

[1] B. Dai, N. Li, D. Schuurmans. "Spectral Ghost in Representation Learning: from Component Analysis to Self-Supervised Learning." *arXiv:2601.20154*, 2026.

[2] C. Gao, H. Sun, N. Li, D. Schuurmans, B. Dai. "Spectral Representation-based Reinforcement Learning." *arXiv:2512.15036*, 2026.

Supplementary Reading by Topic

Contrastive & Non-Contrastive SSL

- Chen et al. "A Simple Framework for Contrastive Learning of Visual Representations (SimCLR)." *ICML 2020*.

- Grill et al. "Bootstrap Your Own Latent (BYOL)." *NeurIPS 2020*.

- Zbontar et al. "Barlow Twins." *ICML 2021*.

- Bardes et al. "VICReg." *ICLR 2022*.

- Chen & He. "Exploring Simple Siamese Representation Learning (SimSiam)." CVPR 2021.
- He et al. "Momentum Contrast for Unsupervised Visual Representation Learning (MoCo)." CVPR 2020.
- HaoChen et al. "Provable Guarantees for Self-Supervised Deep Learning with Spectral Contrastive Loss." NeurIPS 2021.
- Guo et al. "Representation Learning via Non-Contrastive Mutual Information (MINC)." 2025.

Latent-Variable & Clustering

- Caron et al. "Deep Clustering for Unsupervised Learning of Visual Features (DeepCluster)." ECCV 2018.
- Caron et al. "Emerging Properties in Self-Supervised Vision Transformers (DINO)." ICCV 2021.
- Caron et al. "Unsupervised Learning by Contrasting Cluster Assignments (SwAV)." NeurIPS 2020.
- Asano et al. "Self-Labeling via Simultaneous Clustering and Representation Learning (SeLa)." 2019.

Multi-Modal Representation

- Radford et al. "Learning Transferable Visual Models From Natural Language Supervision (CLIP)." ICML 2021.
- Zhai et al. "Sigmoid Loss for Language Image Pre-Training (SigLIP)." ICCV 2023.
- Sun et al. "AmorLIP." 2025.

Spectral RL & Control

- Ren et al. "Spectral Decomposition Representation for Reinforcement Learning (Speder)." ICLR 2023.
- Ren et al. "Latent Variable Representation for Reinforcement Learning (LV-Rep)." ICLR 2023.
- Shribak et al. "Diffusion Spectral Representation for Reinforcement Learning (Diff-SR)." 2024.
- Zhang et al. "Making Linear MDPs Practical via Contrastive Representation Learning (CTRL)." ICML 2022.
- Jin et al. "Provably Efficient Reinforcement Learning with Linear Function Approximation." COLT 2020.

Classical Methods & Theory

- *Belkin & Niyogi. "Laplacian Eigenmaps." NeurIPS 2001.*
- *Hotelling. "Relations Between Two Sets of Variates (CCA)." Biometrika 1936.*
- *Gutmann & Hyvärinen. "Noise-Contrastive Estimation." AISTATS 2010.*
- *Rahimi & Recht. "Random Features for Large-Scale Kernel Machines." NeurIPS 2007.*

Grading Policy

The grading breakdown is as follows:

- Paper Presentation — 40%

Present one research paper and analyze it through the spectral representation lens.

- Course Project — 40%

Original research: theoretical analysis, novel algorithm design, or empirical study extending spectral representation methods. Deliverables:

proposal, mid-term report, final report + presentation.

- Participation — 20%

In-class discussion and engagement during presentations.

Computing Resources: [Google Colaboratory](#) allows free access to run Jupyter Notebooks using GPU resources. The [Google Cloud Platform](#) and [AWS Educate](#) are also good resources. [The GitHub Student Developer Pack](#) also offers free Microsoft Azure and Digital Ocean credits. This semester, we are also offering [PACE ICE](#), Georgia Tech's in-home cluster to students.

Attendance Policy

All lectures will be in-person. The materials will be posted after class.

Attendance is not required but expected. The lectures are a chance to directly interact with your instructor and classmates.

Academic and Research Honesty/Integrity Statement

Georgia Tech aims to cultivate a community based on trust, academic integrity, and honor. Students are expected to act according to the highest ethical standards. Review the [Student Code of Conduct](#) and the [Academic Honor Code](#), especially [Appendix A: Graduate Addendum to the Academic Honor Code](#).

Students are expected to perform research in an ethical and responsible manner. All Doctoral and Master's Thesis students are required to take the [Responsible Conduct of Research training](#), and it is expected that students abide by the principles taught in that training while performing research for this thesis course.

Allegations of scientific or scholarly misconduct are handled in accordance with the procedures outlined by the [Policy for Responding to Allegations of Scientific or Other Scholarly Misconduct](#).

Accommodations for Students with Disabilities

If you are a student with learning needs that require special accommodation, [contact the Office of Disability Services](#) as soon as possible to make an appointment to discuss your special needs and to obtain an accommodations letter. Please also e-mail me as soon as possible in order to set up a time to discuss your learning needs.