**Invariance, Causality, and Applications**
David M. Blei
Columbia University
Spring 2026

## Seminar overview

This seminar explores the idea of *invariance* and its role in causal inference and probabilistic modeling. We will study statistical methods for finding invariant relationships across multiple environments, how these ideas connect to causality and robustness, and how they extend to representation learning and empirical Bayes. Our goal is to explore research at the intersection of all of these ideas.

Our context for these subjects will be probabilistic modeling and machine learning. This is not a survey course or a gentle introduction. It is an exploratory seminar where we will think about and develop ideas together.

## Core themes

We will explore four core circles of ideas.

1. **Causal inference**
    - Prediction under intervention
    - Discovery of causal relationships
    - Counterfactual reasoning
    - Identification and estimation
2. **Invariance**
    - Identifying relationships that remain stable across environments
    - Connections to robustness and generalization
    - Invariance as a criterion for scientific explanation
3. **Causal representation learning**
    - High-dimensional observations with lower-dimensional causal structure
    - Inferring latent causal variables (e.g., health states from medical images)
    - The role and limits of "disentanglement"
4. **Empirical Bayes**
    - Bayesian models that adapt to populations
    - Learning priors from data
    - A blend of Bayesian and frequentist thinking

## Prerequisites

I expect you to be fluent in probabilistic modeling and machine learning.

This typically means you have taken *Probabilistic Models and Machine Learning*. You should be able to design probabilistic models and implement approximate inference.

Ideas from this field—such as deep generative models, amortized variational inference, stochastic optimization, graphical models, probabilistic diffusion models, and others—will be in the *background* of our discussion. I expect you to know most of them well.

## Seminar logistics

- Enrollment cap: 25 students
- Devices: No laptops, tablets, or phones during seminar meetings
- Format: Discussion-based, instructor-led
- Expectations: Reading, participation, progress on a doctoral-level research project

## Readings and reader reports

We will read a mix of papers and books. There will usually be an assigned reading.

Every week, you must post a reader report to Slack. It should be prepared in LaTeX and under one page. It can include ideas, questions, confusions, opinions.

Before each class, please read the other students' reader reports.

## Final project

You must complete a research project by the end of the semester. This project should pursue original work around the themes of the class. Projects should aim to develop new methods for solving important applied problems.

We will start projects early. Every two weeks, all students will report on their progress.

In developing your project, you should be able to answer the following questions:

1. What problem am I solving? State it clearly.
2. Why is this problem important? To whom?
3. What is my strategy for solving it?
4. Why is this solution good? Where does it fall short?
5. How can I demonstrate that the solution works?

We hope that your project will use real data and pursue real questions. We prefer large, non-benchmark datasets. Examples:

- Electronic health records
- Neuroscience recordings
- Climate measurements

- Educational assessments
- Political polls
- Genetic measurements

(We are less interested in finance applications.)

### Grading

- Final project: 75%
- Participation: 25%

Participation includes class discussion and reader reports.

### Schedule and Readings

The schedule is loose and will likely change as the semester evolves.

| Week | Topic | Readings |
|---|---|---|
| 1 | Orientation: What this class is about. | |
| 2 | Causality Foundations I | Pearl et al. (2016) |
| 3 | Causality Foundations II | Pearl et al. (2016) |
| 4 | Causality Foundations III | Peters et al. (2018) |
| 5 | Invariance and Causality | Peters et al. (2016); Bühlmann (2020) |
| 6 | Probabilistic Invariance | Wu et al. (2025b) |
| 7 | Invariant Risk Minimization | Arjovsky et al. (2019); Krueger et al. (2021) |
| 8 | Causal Representation Learning I | Schölkopf et al. (2021) |
| 9 | Causal Representation Learning II | Uhler and Zhang (2025) |
| 10 | Causal Representation Learning III | von Kügelgen et al. (2025) |
| 11 | Empirical Bayes I — Classical Foundations | Efron (2012) |
| 12 | Empirical Bayes II — Probabilistic Symmetries | Wu et al. (2025a) |
| 13 | Synthesis and discussion | |

## References


Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. (2019). Invariant risk minimization. *arXiv:1907.02893*.

Bühlmann, P. (2020). Invariance, causality and robustness. *Statistical Science*, 35(3):404–426. 2018 Neyman Lecture, Special Issue on Causal Inference.

Efron, B. (2012). *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*, volume 1. Cambridge University Press.

Krueger, D., Caballero, E., Jacobsen, J., Zhang, A., Binas, J., Zhang, D., Priol, R. L., and Courville, A. (2021). Out-of-distribution generalization via risk extrapolation. *arXiv:2003.00688*.

Pearl, J., Glymour, M., and Jewell, N. (2016). *Causal Inference in Statistics: A Primer*. John Wiley & Sons.

Peters, J., Bühlmann, P., and Meinshausen, N. (2016). Causal inference by using invariant prediction: Identification and confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):947–1012.

Peters, J., Janzing, D., and Schoelkopf, B. (2018). *Elements of Causal Inference : Foundations and Learning Algorithms*. The MIT Press.

Schölkopf, B., Locatello, F., Bauer, S., Ke, N., Kalchbrenner, N., Goyal, A., and Bengio, Y. (2021). Towards causal representation learning. *arXiv:2102.11107*.

Uhler, C. and Zhang, J. (2025). Causal structure and representation learning with biomedical applications. *arXiv 2511.04790*.

von Kügelgen, J., Ketterer, J., Shen, X., Meinshausen, N., and Peters, J. (2025). Representation learning for distributional perturbation extrapolation. *arXiv:2504.18522*.

Wu, B., Weinstein, E. N., and Blei, D. M. (2025a). Bayesian empirical bayes: Simultaneous inference from probabilistic symmetries. *arXiv 2512.16239*.

Wu, L., Yin, M., Wang, Y., Cunningham, J. P., and Blei, D. M. (2025b). Bayesian invariance modeling of multi-environment data. *arXiv 2506.22675*.