

Computational Methods for Social Science Research

A Graduate Seminar on Nonparametric Modeling and Interpretable Machine Learning

Instructor Jessala A. Grijalva, Ph.D.

Credits 3 | **Format** Seminar + Workshop

Meeting Time TBD | **Location** TBD

Office Hours TBD (sign-up required)

COURSE DESCRIPTION

Social scientists routinely ask which features of a dataset are associated with an outcome, in what direction, and with what strength — yet the dominant tools impose functional form assumptions that are rarely tested and frequently violated in high-dimensional survey data. This course introduces a nonparametric framework for explanatory outcome modeling using Random Forest estimation and TreeSHAP decomposition. Organized as a complete research pipeline across sixteen weeks, students develop a research question, prepare survey data, fit and evaluate a model, interpret findings, and communicate results. All work is completed in R and rendered using Quarto. The course does not teach R as content — it teaches methodological reasoning: how to make defensible analytical decisions, how to let data characterize structure rather than imposing it, and how to report findings honestly.

LEARNING OBJECTIVES

By the end of this course, students will be able to:

- Distinguish between explanatory, predictive, and causal modeling goals and defend the analytical choices appropriate to each
- Prepare survey data for nonparametric analysis, documenting every preprocessing decision with sufficient detail for independent replication
- Fit, evaluate, and interpret a Random Forest model using TreeSHAP decomposition, reading results substantively rather than mechanically
- Produce a complete, reproducible analysis pipeline in R and Quarto and communicate findings in both a rendered technical document and a written scholarly paper

MATERIALS AND RESOURCES

All readings are posted as PDFs on the course management system or are freely available online. No textbook purchase is required. The primary methods reference is:

Molnar, Christoph. 2022. *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*. 2nd ed. Free at christophm.github.io/interpretable-ml-book.

Additional references appearing across multiple weeks:

Grimmer, Justin, Margaret E. Roberts, and Brandon M. Stewart. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton University Press. (Selected chapters.)

Shmueli, Galit. 2010. "To Explain or to Predict?" *Statistical Science* 25(3): 289–310.

Students need R and RStudio installed by the first day of class. All analytical work is completed in Quarto (.qmd) and rendered to PDF. Students new to R should begin working through the introductory materials provided before the semester starts.

GRADING

Component	Weight
Active Engagement (discussion, workshops, peer feedback)	25%
Data Preparation Plan	10%
Preliminary Analysis Checkpoint	15%
Final Presentation	20%
Final Deliverables (Rendered Analysis PDF + Written Paper)	30%

REQUIRED INDIVIDUAL MEETINGS

This course includes two required individual meetings built into the schedule.

Check-In 1 Week 4 — *Project Approval*

Before beginning data preparation, each student meets with the instructor to discuss their research question, proposed dataset, and planned outcome variable. The purpose is to confirm the project is feasible and that the dataset is appropriate for the analytical framework. No student begins the data preparation plan without an approved project. Sign-up opens at the end of Week 3.

Check-In 2 Week 13 — *Post-Analysis Review*

After completing the preliminary analysis and before the final presentation, each student meets with the instructor to review SHAP output, discuss interpretation, and identify any methodological issues. Students should bring a near-final rendered Quarto document, their beeswarm plot and at least two dependence plots, and their single most important unresolved interpretive question. This meeting is required. No student presents in Weeks 15 or 16 without completing it.

WEEKLY SCHEDULE

Each week assigns two priority readings to be completed before the first class meeting. One recommended reading provides additional depth and is not required. Workshop weeks have no new readings — both sessions are devoted to peer review and collaborative troubleshooting.

PHASE I — Research Design and Data Preparation

Week 1 — Explanatory Modeling and Research Design

What does it mean to characterize the associational structure between a feature space and an outcome? How is this goal different from causal inference and from prediction? This distinction is foundational — getting it wrong shapes every subsequent decision in the analysis.

Priority Reading

- Shmueli, Galit. 2010. "To Explain or to Predict?" *Statistical Science* 25(3): 289–310.
- Molnar (2022), Chapters 1–2 (Introduction; Interpretability).

Recommended

- Athey, Susan. 2017. "Beyond Prediction: Using Big Data for Policy Problems." *Science* 355(6324): 483–485.

Students identify research area by end of week.

Week 2 — Survey Data Architecture and the Assumptions We Impose

What do surveys actually contain, and what happens when we force that information through a linear model? This week examines the structure of survey data and the functional form assumptions that regression imposes — assumptions that are rarely tested and frequently violated.

Priority Reading

- Molnar (2022), Chapter 4, sections on linear regression and logistic regression.
- King, Keohane, and Verba. 1994. *Designing Social Inquiry*. Princeton University Press. Chapter 1.

Recommended

- Achen, Christopher H. 2005. "Let's Put Garbage-Can Regressions and Garbage-Can Probits Where They Belong." *Conflict Management and Peace Science* 22(4): 327–339.

Students identify dataset by end of week.

Week 3 — Data Preparation as Research Design

Data preparation is not a preliminary chore. It is a series of consequential research design decisions that shape everything downstream: how missing values are handled, how variables are encoded, which observations are included, and in what direction variables are coded.

Priority Reading

- Stekhoven, Daniel J., and Peter Böhmann. 2012. "MissForest." *Bioinformatics* 28(1): 112–118.
- Molnar (2022), Chapter 5, sections on Feature Importance and Permutation Feature Importance.

Recommended

- Lall, Ranjit. 2016. "How Multiple Imputation Makes a Difference." *Political Analysis* 24(4): 414–433.

Preprocessing checklist and dataset documentation due end of week for instructor review before the Week 4 check-in.

NO CLASS — PROJECT APPROVAL MEETINGS

Week 4 — Required Project Approval Meetings with Instructor

No class sessions. Each student meets with the instructor during a scheduled 25-minute appointment. Come prepared with a one-paragraph description of your research question, your dataset identified and accessible, your planned outcome variable and feature set, and your preliminary preprocessing checklist from Week 3. No student begins the data preparation plan without an approved project.

Week 5 — Workshop: Data Preparation

No new readings. Both sessions are devoted to peer review of data preparation plans and collaborative troubleshooting. Students bring their preprocessing checklists, memos, and data to class. Peer feedback is structured and evaluated as part of the Active Engagement grade.

Data Preparation Plan due Friday of Week 5.

PHASE II — Estimation and Interpretation

Week 6 — Nonparametric Estimation: Random Forests

Random Forests estimate the relationship between a feature space and an outcome without imposing functional form assumptions. This week covers the logic of ensemble methods, how decision trees aggregate, out-of-bag error, and why nonparametric estimation is suited to high-dimensional survey data.

Priority Reading

- Breiman, Leo. 2001. "Random Forests." *Machine Learning* 45(1): 5–32.
- Molnar (2022), sections on Decision Trees and Random Forests.

Recommended

- Muchlinski, David, et al. 2016. "Comparing Random Forest with Logistic Regression for Predicting Class-Imbalanced Civil War Onset Data." *Political Analysis* 24(1): 87–97.

BREAK

Week 7 — Spring / Fall Break

No class. No assignments due.

Week 8 — Model Fitting and Performance Evaluation

How do you train a model, evaluate whether it learned something real, and guard against overfitting? This week covers stratified train-test splitting, AUC, class imbalance handling, and hyperparameter tuning. Students fit their models this week.

Priority Reading

- Molnar (2022), sections on Model Performance and Measuring Performance.
- Chawla, Nitesh V., et al. 2002. "SMOTE: Synthetic Minority Over-Sampling Technique." *Journal of Artificial Intelligence Research* 16: 321–357.

Recommended

- Grimmer, Roberts, and Stewart (2022), Chapter 5, sections on evaluation and validation.

Week 9 — Workshop: Analysis

No new readings. Both sessions are devoted to peer review of preliminary analyses and collaborative interpretation of initial results. Students bring their model output, performance metrics, and preliminary SHAP visualizations to class.

Preliminary Analysis Checkpoint due Friday of Week 9.

Week 10 — Interpretation: TreeSHAP Decomposition

A learned forest cannot be read directly as a set of coefficients. TreeSHAP decomposes the model's output into additive feature contributions grounded in cooperative game theory. This week covers SHAP values, beeswarm plots, and dependence plots. The emphasis is on reading these outputs substantively — what do they tell us about the world, not just about the model?

Priority Reading

- Lundberg, Scott M., and Su-In Lee. 2017. "A Unified Approach to Interpreting Model Predictions." *Advances in Neural Information Processing Systems* 30: 4765–4774.
- Molnar (2022), sections on Shapley Values and SHAP.

Recommended

- Grijalva, Jessala A. [under review]. "Structure Without Assumptions: A Nonparametric Framework for Explanatory Modeling with Survey Data." Manuscript provided on course management system.

PHASE III — Validation, Robustness, and Communication

Week 11 — Validation, Robustness, and Limitations

Reporting that a feature ranks first in importance says nothing about whether that ranking would hold in a different sample. This week addresses why single-sample results require scrutiny, what validation approaches exist — including bootstrap rank stability, cross-validation, and sensitivity analysis — and how to discuss limitations honestly.

Priority Reading

- Grijalva [under review], sections on Bootstrap Rank Stability and the Pipeline (Steps 12–14).
- Molnar (2022), sections on Accumulated Local Effects and Feature Interaction.

Recommended

- Efron, Bradley. 1979. "Bootstrap Methods: Another Look at the Jackknife." *Annals of Statistics* 7(1): 1–26.
-

Week 12 — Communicating Findings: From Model Output to Scholarly Argument

A technically correct analysis can still fail to communicate. This week addresses how to translate model output into a scholarly argument — how to write about SHAP findings for an audience unfamiliar with the method, how to frame limitations honestly without undermining findings, and how to structure the written paper so that analysis and argument are integrated rather than parallel.

Priority Reading

- Grimmer, Roberts, and Stewart (2022), Chapter 2, sections on social science goals and validation.
- Molnar (2022), Chapter 9 (Explaining Individual Predictions).

Recommended

- Healy, Kieran. 2018. *Data Visualization: A Practical Introduction*. Chapter 1. Free at socviz.co.

Students should have complete SHAP output and a draft interpretation by end of week.

NO CLASS — POST-ANALYSIS CHECK-IN MEETINGS

Week 13 — Required Post-Analysis Check-In Meetings with Instructor

No class sessions. Each student meets with the instructor during a scheduled 25-minute appointment. Come prepared with your near-final rendered Quarto document, your beeswarm plot and at least two dependence plots, a one-paragraph plain-language description of what your model found, and your single most important unresolved interpretive question. This meeting is required. No student presents in Weeks 15 or 16 without completing it.

Week 14 — Workshop: Final Presentation Preparation

No new readings. Both sessions are devoted to rehearsal, peer feedback on presentation drafts, and collaborative troubleshooting of final analyses. Students present a 5-minute version of their talk to a small group and receive structured feedback.

Final presentation slides or materials due to instructor by end of Week 14.

Week 15 — Final Presentations, Part I

No new readings. The first half of the seminar presents during Week 15. Each student presents their complete analysis for approximately 10 minutes, followed by 10 minutes of questions and structured peer feedback.

Presentation schedule posted by end of Week 13.

Week 16 — Final Presentations, Part II and Course Conclusion

No new readings. The second half of the seminar presents during Week 16. The final class meeting includes a brief synthesis discussion: what have we learned about what machine learning can and cannot do for social science, and what does letting the data speak actually mean?

Rendered Analysis PDF and Final Paper due during finals week.

ASSIGNMENT DESCRIPTIONS

Data Preparation Plan (Due Week 5)

Two components. First, a preprocessing checklist documenting how the student will handle special values (refusals, don't-know responses, not-applicable codes), any subsetting decisions, directional recoding plans, and encoding strategy for each variable type. The checklist should be specific enough that the instructor can identify any issues before the student proceeds. Second, a brief memo (1–2 pages) describing the research question, why it matters, and what the student expects to learn.

Preliminary Analysis Checkpoint (Due Week 9)

A report (3–5 pages) including the train-test split specification, hyperparameter tuning results, AUC on the held-out test set, initial SHAP output (beeswarm plot and at least two dependence plots for the highest-ranked features), and a preliminary interpretation of what the model has found. The interpretation should engage substantively with the findings, not simply describe the output.

Final Presentation (Weeks 15–16)

A formal presentation covering the research question and its motivation, the analytical approach, and the findings. Approximately 10 minutes, followed by 10 minutes of questions and discussion. Students must complete the Week 13 check-in meeting before presenting.

Final Deliverables (Finals Week)

Two submissions. First, a Quarto document (.qmd) rendered to PDF containing the full reproducible analysis — all code, output, and SHAP visualizations. Students submit both the rendered PDF and the source .qmd file. Second, a written paper (4–7 pages) presenting the research question, summarizing the findings, and interpreting them substantively. The paper should read as a self-contained document and incorporate feedback from the Week 13 check-in and the final presentation. No code in this document.

COURSE POLICIES

Attendance and Participation

Students are expected to attend all class sessions. Workshop weeks in particular depend on everyone being present — your absence affects your peers' learning, not just your own. If you need to miss a class, please email in advance. Frequent absences will affect your engagement grade.

Late Work

If you need an extension, ask before the due date and we will work out a plan. Late work submitted without advance notice loses 10% of the total assignment points per day. The final deliverables deadline is firm.

Academic Integrity

All work submitted must be your own. Code and analysis must be independently produced. Collaboration on conceptual questions during workshops is encouraged, but each student's analysis, code, and written work must be their own.

Use of AI Tools

Large language models and AI-assisted coding tools may be used as aids for debugging code, brainstorming, and revising drafts. They may not be used to generate analysis, produce SHAP interpretations, or write substantive sections of assignments. If you use AI tools in any part of your workflow, include a brief statement describing how. The analytical reasoning in this course must be yours.

R and Quarto

This course uses R as the computational environment and Quarto for all rendered documents. The course does not teach R as content, and class time will not be spent on R instruction. Students are responsible for developing and maintaining their R skills using the resources provided and available workshops. Students with limited R experience should begin working through the introductory materials before the semester starts.

Accessibility

This course is committed to being accessible to all students. If there are circumstances that make any aspect of this course difficult for you, please reach out so we can discuss accommodations. For formal accommodations, please contact your institution's disability services office.

This syllabus is subject to change. Any revisions will be announced in class and posted to the course management system.