

06-002-107-3: Forschungsseminar B – Experimentelle Soziologie und Computational Social Science

Sascha Grehl, Wiebke Krügener, & Felix Lennert

Winter 2026/27

- E-mail: sascha.grehl@uni-leipzig.de und kruegener@studserv.uni-leipzig.de
- Course hours: Wednesdays and Thursdays, 11:15 – 12:45; for exact dates, see schedule
- Readings: see schedule
- Course materials: see website
- Student hours: are to be set up individually via email; Saschas office is H3 1.07; Wiebkes office is H2 1.13; there is a multitude of valid reasons why you should come see us – some are listed here:
 - things are unclear and you need help with the material
 - you want to discuss a research idea
 - you have come across a cool new paper that we might deem interesting
 - you have recommendations for us in terms of general course resources/references/structure/behavior
 - you want some career advice from someone roughly your age
 - you forgot your mensa card at home and want some coffee/tea
 - in case you need some free period supplies, no email is required, you can just get them at Leonie Steinbrinker’s office (we also have a key if she’s not there), H3 1.06

Course description

The Forschungsseminar in Computational Social Science (CSS) equips you with the tools to analyze human behavior, predict social trends, and tackle complex societal issues using cutting-edge data science techniques. From web scraping to AI-powered text analysis, you’ll learn to use your computer in new ways to gain insights into social phenomena.

The curriculum covers a range of topics including data management, web scraping, speech-to-text, and computational text analysis. Students will hone their R and develop skills in Python, applying these languages to real-world social science problems. The course progresses from fundamental concepts to advanced techniques, including the use of state-of-the-art AI models for text analysis.

The course structure consists of one lecture and one lab session per week, providing a balance of theoretical knowledge and practical application. Throughout the semester, students will benefit from hands-on coding exercises, one-on-one mentoring, and collaborative projects. The course culminates in a research paper, allowing students to apply their new skills to a topic of their choice.

This course is ideal for social scientists looking to enhance their computational skills. It is geared toward 2nd year Master’s students who are enrolled in the reformed Sociology Master’s program. Interested Bachelor’s students are also very welcome to attend, they shall send an email to the instructors no later than October 23, 2026. However, Bachelor’s students will not be able to earn credits with their attendance.

What to expect

This course is structured so that it provides you – the student – with theoretical **lectures** (on Thursdays) and hands-on coding labs the so called **practicals** (Wednesdays). The lectures aim to introduce you to how

you can use techniques to conduct empirical research and will mostly consist of the presentation of innovative and/or cutting-edge research that has harnessed digital data to solve exciting research puzzles. They also might contain some practical demonstrations already. The practicals then showcase how to perform similar things yourself and give you ample opportunities to perform the analyses on your own data. Exercises will be provided, but you are not expected to hand them in – we could not care less and not doing them is simply your loss. However, we are of course always available to help you with the exercises if you get stuck.

The main objective of this course is that students **perform their own research as the course unfolds**. To this end, we will talk about your interests. Study groups of matching research interests might be formed. Students can of course also work individually. Students are expected to continue working on their projects as the course unfolds.

A tried-and-tested strategy is to work on a topic you could later use for your Master’s thesis. However, this is not a requirement. In this case, this course would give you more time for the tedious data acquisition and cleaning phase, which is often underestimated.

At the very end of the course, it is time for a “peer reviewed presentations” of your project. One of your peers will be assigned your project and provide comments on it. This peer review ensures that you are right on track and that everyone has accomplished the course’s learning goals. After that, you are all good to go and can check your analyses and write up your resulting papers so that they resemble a proper empirical research paper (see more in “Expectations” below).

The actual “peer-review” will work as follows: first, you (the author of the project) present your work briefly (10 minutes). Then, the assigned reviewer group will provide their comments (5 minutes). To facilitate this, you must send presentations presenting the main idea, some theory (including hypotheses), a preliminary testing strategy, and the first results to us by **January 30, 2027**, so we can distribute them to your peers.

As a peer reviewer, you have 5 minutes to provide feedback on four components:

- (1) Research Question & Motivation: Is it clear and compelling? Are there ways to sharpen the focus?
- (2) Theory & Hypotheses – Is relevant literature cited? Are there gaps or alternative frameworks to consider?
- (3) Data & Methods – Is the approach appropriate? Are there concerns about data quality or better-suited methods?
- (4) Preliminary Results – Are visualizations clear? Do results address the hypotheses? What additional analyses might help?

Please be constructive rather than critical, frame suggestions as questions (“Have you considered...?”), be specific with your feedback, and acknowledge that everyone is working within constraints. Aim to identify 2-3 main strengths and 2-3 areas for improvement, and suggest specific papers, methods, or framings where helpful.

The (tentative) deadline for the paper (“Forschungsbericht”) is **March 19, 2027**. Please send it – in PDF format – to Simone Müller (muellers@uni-leipzig.de) and CC Sascha (sascha.grehl@uni-leipzig.de). Code and data have to be sent to us via email (depending on data size, data can also be shared via Google Drive/Dropbox/Uni Leipzig Wolke). The code should run “out of the box” and contain everything necessary to replicate the results and graphs in the paper. Preferably, this is in chronological order and structured into sections with descriptive titles. We have no preferences regarding the programming language (R, Python), therefore your project can be in either language or a mix of both (e.g., using **reticulate** in R or **quarto** notebooks – the latter being our suggestion).

Extensions

Extensions can be granted for particular reasons. These involve, among others, internships and sickness. In the case of the former, please give us a quick heads-up so that we can arrange it (preferably with some sort of proof). If you need an extension for a different reason than the ones mentioned above, feel free to reach out anytime, and we will do our best to accommodate your needs.

Expectations

- The Basics
 - written in English (recommended) or German
 - file format: PDF
 - font size 12 pt, 1.5 line spacing
 - no typos, grammatical flaws, etc. (you are living in the age of helpers such as Grammarly, there are no more excuses)
 - length: 5000 words ($\pm 10\%$, approximately 15 pages, excluding the title page, references, and tables)
 - cite correctly and in a uniform manner. We suggest ASA as an citation style. In addition, it is strongly advised to use Zotero and Quarto/Overleaf; resources can be provided upon request.
- Structured resembles an empirical research paper:
 - the *introduction* contains an empirical social scientific research question that is theoretically and practically motivated (i.e., showing its scientific and real-world relevance)
 - the *theory section* provides a **brief** overview of relevant prior research; clearly testable hypotheses are derived from the literature/goals for exploratory analyses are formulated
 - in *data and methods*, the data (including acquisition strategy), as well as the analysis strategy, are described; in our case, the data consist of text, the analyses are related to the course content; data and methods need to enable valid results
 - *results* need to be visualized through tables and/or (gg)plots and described in the text; tables and visualizations need to be properly labeled so that they can “stand on their own”
 - *discussion* of the results is performed in the light of the theoretical foundations; potential shortcomings and reach of the paper are outlined
 - the *conclusion* circles back to the introduction and connects it to the results; it needs to clearly answer the research question

Basic rules of behavior

- If anything is unclear, ask us. This probably means that we have failed our job, and your question offers us a second chance to fix this.
- No discrimination. Never. If you witness any, tell us. We will find a way to deal with it.
- THIS IS IMPORTANT: If there are problems, reach out whenever. Do not let them become too big.
- Copy code from the internet – but you are responsible for the solution, so please make sure it works and solves your problem.
- Generative AI (i.e., ChatGPT et al.) is explicitly allowed. In our opinion, it is a tool that is here to stay, and you should use whatever resource you have to get the job(s) done. Plus, writing the right prompt is a skill in itself that you should definitely hone. If you use it for your writing, please make sure to proof-read everything properly, since you – and only you – will be held accountable for its content.
- Form groups with your peers for working on the material. Everything will be easier and more fun. Except for when you have free riders. Kick them out of your group.
- AGAIN: ask questions if needed. Anytime.

Schedule

As stated above, Wednesdays are practicals and Thursdays are lectures. Please bring a laptop to all of our meetings (if you don't have one, feel free to reach out and we will try our best to lend you one). It's also a good idea to charge it before, since outlets appear a bit scant.

Most of you will use RStudio. Sascha, for example, personally uses Positron, which is basically the same but more accommodating to Python. You can download it for free from posit.co. However, RStudio is perfectly fine, too.

Literature-wise, we will use a mix of textbooks and review papers to introduce theoretical concepts and

related studies to illustrate. In terms of programming, we will mostly rely on online resources. However, everything that is relevant in terms of R content (and more!) can be found in the R script.

Every reading will be either provided online or linked to in this syllabus (just click on “*online*”).

We do not expect you to read the literature and will do the theoretical sessions in a “lecture” style. This is because this is an applied course and not every piece of content has the same relevance for everyone. Having been students ourselves, we think that students should not be overwhelmed by having to read everything while working on their projects. We will upload additional readings on top of the ones this syllabus mentions in case you want to read more and need some inspiration. We also recommend looking at the references of the papers we read to find more relevant literature or just ask us anytime.

Week 1: Kick Off

Practical (Wed, 14 October 2026): Welcome

- Let’s get to know each other!

Lecture (Thu, 15 October 2026): Thinking about your projects

Week 2

Practical (Wed, 21 October 2026): Setting up your workstation

- Preparing everything for Scientific Computing
 - Acquire access to *sc.uni-leipzig.de*
 - Get a Github account if you don’t have one yet
 - Connect your PC and the SC Account with your Github account
- R recap (the corresponding chapters can be found in the R4DS book – *online*)
 - RMarkdown/Quarto – chapters 28 & 29
 - dplyr – chapter 4
 - tidyr – chapter 6
 - ggplot2 – chapters 2 & 10 & 11 & 12
 - purrr & loops in different flavors – chapter 27
 - functional programming – chapter 26
 - set up **reticulate** in RStudio – *online*

Lecture (Thu, 22 October 2026): Python & Regexes

- Brief intro to Python
 - *reticulate* – how to run Python in R studio
 - data types
 - loops
 - functions
 - pandas

Week 3

Practical (Wed, 28 October 2026): Regular Expressions

- Toolbox CSS Website: *Chapter 6*
- **stringr** & Regular Expressions – R4DS book, *online*, chapters 15 & 16

Lecture (Thu, 29 October 2026): How the Web is Written and Ethics

- Stoltz and Taylor (2024) – chapter 5
- Blog post on CSS selectors – *online*
- Blog posts on API calls – *online*

Week 4

Practical (Wed, 04 November 2026): `rvest`

- Toolbox CSS Website: *Chapter 7*
- `rvest` Web scraping 101 – *online*

Lecture (Thu, 05 November 2026): Dynamic Pages and Forms

- TBA

Week 5

Practical (Wed, 11 November 2026): APIs

- Toolbox CSS Website: *Chapter 8*
 - `selenium` documentation – *online*
- Toolbox CSS Website: *Chapter 9*
 - `httr2` documentation – *online*

Lecture (Thu, 12 November 2026): OCR and Audio Transcription

- Stoltz and Taylor (2024) – chapter 5
- Tesseract documentation – *online*

Week 6

Practical (Wed, 18 November 2026): Audio Transcription in Practice

- Stoltz and Taylor (2024) – chapter 5
- OpenAI Whisper Python package documentation – *online*

Lecture (Thu, 19 November 2026): Text Preprocessing and Featurization

- TBA

Week 7: Introduction to Text as Data

Practical (Wed, 25 November 2026): Text Preprocessing and Featurization

- CSS Handbook Chapter 12

Lecture (Thu, 26 November 2026): Bag of Words

- Evans and Aceves (2016)
- Grimmer, Roberts, and Stewart (2022), chapters 3–5, 11, & 15
- Stoltz and Taylor (2024), chapters 4–9

Week 8

Practical (Wed, 02 December 2026): Sentiment Analysis, TF-IDF, and NER/POS

- CSS Handbook Chapter 1e
- Grimmer et al. (2022), chapter 11
- Jurafsky and Martin (n.d.), chapter 21 – *online*
- Silge and Robinson (2017) – *online*, chapters 2 & 3

Lecture (Thu, 03 December 2026): Supervised Machine Learning

- Barberá et al. (2021)
- Grimmer et al. (2022), chapters 17–20
- Stoltz and Taylor (2024), chapters 9 & 12

Week 9

Practical (Wed, 09 December 2026): Supervised Machine Learning in Practice

- Hvitfeldt and Silge (2022) – *online*, chapters 6 & 7
- Silge and Hvitfeldt (2019) – *online*

Lecture (Thu, 10 December 2026): Unsupervised ML

- Blei (2012)
- DiMaggio, Nag, and Blei (2013)
- Grimmer et al. (2022), chapters 10, 12–3
- Stoltz and Taylor (2024), chapters 10 & 11
- Silge and Robinson (2017) – *online*, chapter 6

Week 10

Practical (Wed, 16 December 2026): Unsupervised ML in Practice

- ToolBox CSS Website: *Chapter 15*

Lecture (Thu, 17 December 2026): Word Embeddings

- Jurafsky and Martin (n.d.), chapter 6 – *online*
- Stoltz and Taylor (2021)

— CHRISTMAS BREAK —

Week 11

Practical (Wed, 06 January 2027): Word Embeddings in Practice

- Hvitfeldt and Silge (2022) – *online*, chapter 5
- Stoltz and Taylor (2024), chapter 11
- `text2map`: R Tools for Text Matrices – *online*

Lecture (Thu, 07 January 2027): Supervised Learning on Steroids: BERT

- Do, Ollion, and Shen (2022)
- Laurer et al. (2024)
- Törnberg (2023)
- Wankmüller (2022)

Week 12

Practical (Wed, 13 January 2027): Active Learning with BERT

- Set up environments
- Augmented Social Scientist tutorial – *online*
- BERTopic – *online*

Lecture (Thu, 14 January 2027): LLMs for information extraction

- Stuhler, Ton, and Ollion (2025)

Week 13

Practical (Wed, 20 January 2027): Local LLMs – a primer

- `ellmer` documentation – *online*
- Tutorial from IC2S2 by Etienne Ollion, Emilien Schultz, Julien Boelaert – *online*

Lecture (Thu, 21 January 2027): Project Work Session

- Attendance is optional. Work on your presentations and projects. We will be available for questions and guidance.

Week 14: Presentation Preparation Week

Practical (Wed, 27 January 2027): Project Work Session

- Attendance is optional. Work on your presentations and projects. We will be available for questions and guidance.

Lecture (Thu, 28 January 2027): Project Work Session

- Attendance is optional. Work on your presentations and projects. We will be available for questions and guidance.

IMPORTANT: Deadline for sending presentations: January 29, 6PM.

Week 15: Presentation & Wrap Up Week

Practical (Wed, 03 February 2027): Presentation Day I

- Presentations

Lecture (Thu, 04 February 2027): Presentation Day II

- More Presentations & Wrap-up

Deadline Forschungsbericht

March 19, 2027 (tentative).

References

- Barberá, Pablo, Amber E. Boydstun, Suzanna Linn, Ryan McMahon, and Jonathan Nagler. 2021. “Automated Text Classification of News Articles: A Practical Guide.” *Political Analysis* 29(1):19–42.
- Blei, David. 2012. “Probabilistic Topic Models.” *Communications of the ACM* 55(4):77–84.
- DiMaggio, Paul, Manish Nag, and David Blei. 2013. “Exploiting Affinities Between Topic Modeling and the Sociological Perspective on Culture: Application to Newspaper Coverage of U.S. Government Arts Funding.” *Poetics* 41(6):570–606.
- Do, Salomé, Étienne Ollion, and Rubing Shen. 2022. “The Augmented Social Scientist: Using Sequential Transfer Learning to Annotate Millions of Texts with Human-Level Accuracy.” *Sociological Methods & Research* 004912412211345.
- Evans, James A. and Pedro Aceves. 2016. “Machine Translation: Mining Text for Social Theory.” *Annual Review of Sociology* 42(1):21–50.
- Grimmer, Justin, Margaret Roberts, and Brandon Stewart. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press.
- Hvitfeldt, Emil and Julia Silge. 2022. *Supervised Machine Learning for Text Analysis in R*. First edition. Boca Raton London New York: CRC Press, Taylor & Francis Group.
- Jurafsky, Dan and James Martin. n.d. “Speech and Language Processing.”
- Laurer, Moritz, Wouter Van Atteveldt, Andreu Casas, and Kasper Welbers. 2024. “Less Annotating, More Classifying: Addressing the Data Scarcity Issue of Supervised Machine Learning with Deep Transfer Learning and BERT-NLI.” *Political Analysis* 32(1):84–100.
- Silge, Julia and Emil Hvitfeldt. 2019. “Predictive Modeling with Text Using Tidy Data Principles.” in *useR2020*.
- Silge, Julia and David Robinson. 2017. *Text Mining with R: A Tidy Approach*. First edition. Beijing ; Boston: O’Reilly.
- Stoltz, Dustin S. and Marshall A. Taylor. 2021. “Cultural Cartography with Word Embeddings.” *Poetics* 88:101567.
- Stoltz, Dustin S. and Marshall A. Taylor. 2024. *Mapping Texts: Computational Text Analysis for the Social Sciences*. New York, NY: Oxford University Press.
- Stuhler, Oscar, Cat Dang Ton, and Etienne Ollion. 2025. “From Codebooks to Promptbooks: Extracting Information from Text with Generative Large Language Models.” *Sociological Methods & Research* 54(3):794–848.
- Törnberg, Petter. 2023. “How to Use LLMs for Text Analysis.”
- Wankmüller, Sandra. 2022. “Introduction to Neural Transfer Learning With Transformers for Social Science Text Analysis.” *Sociological Methods & Research* 004912412211345.