



Stefan Müller, PhD

Associate Professor

School of Politics and International Relations

University College Dublin

Belfield, Dublin 4, Ireland

✉ stefan.mueller@ucd.ie

🌐 <https://muellerstefan.net>

Level 4 Module; Autumn Trimester 2025

Applied Data Wrangling and Visualisation (POL42540)

Version: December 2, 2025

Latest version at: <https://muellerstefan.net/teaching/2025-autumn-adwv.pdf>

Time: Thursday, 09:00–10:50

Room: B002-CSI (→ [location](#))

Credits: 5.0

Format: Lecture and computer labs

Module Coordinator: Stefan Müller, PhD

stefan.mueller@ucd.ie | <https://muellerstefan.net>

Office: Newman Building, G312

Office hours: Thu, 11:30–12:30 ([sign up here](#))

TA: Mohamed Moheeb | mohamed.moheeb@ucdconnect.ie

Course Content

Welcome to Applied Data Wrangling and Visualisation! The module offers a comprehensive introduction to the essential techniques and tools required for effective data management and visualisation in R. Students will also learn how to use AI tools as coding assistants, manage projects and handle data in various file formats, ensuring a robust understanding of data cleaning, wrangling, and merging.

The course emphasises the fundamentals of data visualisation, moving from principles to applied data visualisation strategies for compelling data storytelling. Additionally, it delves into the use of relational databases with SQL and data collection through web scraping, enabling students to manage and analyse large datasets efficiently. Applied Data Wrangling and Visualisation prepares module participants for a range of data-intensive roles, equipping them with the knowledge to leverage data visualisation and management tools effectively in their future careers.

Module Instructors

This module is coordinated and taught by Stefan Müller. Jos-Dornschneider-Elkink will deliver the three classes on SQL.

Mohamed Moheeb will serve as a Teaching Assistant for this module. He is the first point of contact for you. The communication in this module will take place on Slack, usually in channels that everyone can contribute to. More details on Slack are provided below.

Course Structure

Week 1: Introduction to R and VSCode (11 September)	10
Week 2: Software, Project Management, and Replicability (18 September)	10
Week 3: Importing, Summarising, Transforming, and Merging Data (25 September)	10
Week 4: Using LLMs as Research Tools (2 October)	11
Week 5: Principles of Data Visualisation (9 October)	11
Week 6: Explorative Data Visualisation (16 October)	11
Week 7: Advanced, Intuitive, and Accessible Data Visualisation (23 October)	12
Week 8: No Class — Reading Break (30 October)	12
Week 9: SQL: Organising and Managing Data (6 November)	12
Week 10: SQL: Retrieving, Joining, and Summarising Data (13 November)	12
Week 11: Introduction to Web Scraping (20 November)	13
Week 12: Advanced Web Scraping (27 November)	13

Learning Outcomes

Upon successful completion of the course, students will be able to:

1. Manage and visualise data effectively in R, using VSCode and GitHub, enhancing your proficiency in handling data in diverse file formats.
2. Use relational databases (SQL) and conduct data collection through web scraping, enabling you to analyse large datasets for a range of data-intensive roles.
3. Develop skills in using AI tools as coding assistants, fostering your ability to manage projects efficiently and improve replicability in their work.
4. Robust understanding of data cleaning, wrangling, and merging techniques, preparing you for the challenges of managing large and complex datasets.
5. Solid foundation in the principles of data visualisation, learning to apply these strategies to create compelling data stories that are both intuitive and accessible.

General Readings

The seminar does not build on a single textbook, but relies on papers and book chapters. All readings used in this module are freely available online or accessible through the [UCD Library](#).

- H. Wickham, M. Çetinkaya-Runde, and G. Grolemund (2023). *R for Data Science: Import, Tidy, Transform, Visualize, and Model Data*. 2nd edition. Sebastopol: O'Reilly. URL: <https://r4ds.hadley.nz>
- N. B. Weidmann (2023). *Data Management for Social Sciences: From Files to Databases*. Cambridge: Cambridge University Press. URL: <https://cambridge.org/9781108845670>

- K. Healy (2019). *Data Visualization: A Practical Introduction*. Princeton: Princeton University Press. URL: <https://socviz.co>
- R. Alexander (2023). *Telling Stories with Data: With Applications in R*. New York: CRC Press. URL: <https://tellingstorieswithdata.com>
- C. O. Wilke (2019). *Fundamentals of Data Visualization: A Primer On Making Informative and Compelling Figures*. Sebastopol: O'Reilly. URL: <https://clauswilke.com/dataviz/>
- T. Monroe-White, D. Kozlowski, L. Domenech, C. Pradier, G. Rosati, and N. S. Shokida (2024). *Critical Computational Social Sciences*. URL: <https://critcss.github.io/CCSS/>
- S. Dattani (2025). *Saloni's Guide to Data Visualization: Why Data Visualization Matters, and How to Make Charts More Effective, Clear, Transparent, and Sometimes, Beautiful*. URL: <https://www.scientificdiscovery.dev/p/salonis-guide-to-data-visualization>.

Plagiarism

Although this should be obvious, plagiarism – copying someone else's text without acknowledgement or beyond 'fair use' quantities – is not allowed. Plagiarism is an issue we take very serious here in UCD. Please familiarise yourself with the definition of plagiarism on UCD's website¹ and make sure not to engage in it.

Module Policy on the Use of Artificial Intelligence (AI) Tools

I encourage the use of AI tools when completing the assignments for this module, but users of this technology must be aware of what it can and more importantly, what it cannot do well. It is crucial for you to exercise judgement when evaluating the quality and reliability of content generated through AI platforms. AI is not a panacea for all writing challenges; it will not automatically generate a flawless, logically coherent assignment. Instead, use AI as a tool to tackle specific issues such as brainstorming and idea formation, literature discovery, and text drafting issues.

View your preferred AI platform(s) as useful but imperfect tools that can offer inspiration, new perspectives, and supplementary areas for research for your own work. In-depth research on your part remains essential to ensure coherent, factual, and scientifically informed perspectives in your assignment. Always cross-reference the information AI offers against other independent and reliable sources.

Late Submission Policy

If a student or group submits an assignment late, the following penalties will be applied:

- Coursework received at any time within two weeks of the due date will be graded, but a penalty will apply.
 - Coursework submitted at any time up to one week after the due date will have the grade awarded reduced by two grade points (for example, from *B–* to *C*).
 - Coursework submitted more than one week but up to two weeks after the due date will have the grade reduced by four grade points (for example, from *B–* to *D+*). Where a student finds they have missed a deadline for submission, they should be advised that

¹<https://libguides.ucd.ie/academicintegrity>.

they may use the remainder of the week to improve their submission without additional penalty.

- Coursework received more than two weeks after the due date will not be accepted. Regulations regarding extenuating circumstances apply.

Release of Feedback and Grades

According to UCD's Academic Regulations, feedback on assessment must be provided no later than *20 working days* (roughly four weeks) after the deadline for submission of each piece of assessed work, except work submitted late or submitted as part of the final assessment component of the module. I will, of course, stick to these regulations.

Questions and Problems

In this module, we will discuss concepts, methods, and software you might not have heard of before. I am aware that parts of this module could be challenging, and I will assist you as best as I can.

We will use Slack for in this module.² Make sure to create a Slack account before the first seminar and join the Slack workspace. If you have a question that involve code or concepts, please share your question in `#coding`, `#homework`, or `#research-paper`.

If you struggle to solve problems relating to R, Python or SQL, please follow the steps outlined below before contacting your peers or me. It is very likely that at least one other person faced the same problem before or received the same error message.

1. You are welcome and encourage to use AI tools as coding assistants. Recall that [GitHub Copilot](#) is a powerful coding assistant.
2. Try to summarise the problem in your own words and then google this summary or use an AI tool, for example [ChatGPT](#). If the problem relates to R, add `rstats` to your search query. For example: `how to import csv file in rstats`. I am almost certain that you will find a solution to most of your questions.
3. If your R code returns an error, I would advise you to google the text of the error message. For example, you can google the error message `"Error: Can't subset columns that don't exist."`

→ If steps 1–3 still do not solve your problem or question, please ask your question in the Slack channel devoted to this module. Your peers and we will help you.

Contact and Office Hours

Mohamed Moheeb, our Teaching Assistant for this module, is your first point of contact. If you have questions about software, code, or assignments, please use Slack first (see more details below) or check if one of your peers already asked this question. We monitor Slack regularly and will respond to your question if other module participants cannot help you.

In addition, Stefan Müller also offers office hours, either in in person (Room G312, Newman Building) or online. Please sign up for a meeting at <https://calendly.com/mueller-ucd/office-hours>.

²I have had very positive experiences with Slack in my modules. Müller (2023) discusses both the advantages and shortcomings of Slack for teaching and learning.

Proper communication practices are crucial for formal email exchanges. Please familiarise yourself with [this summary](#)³ on professional, efficient, and respectful communication.

Software

In this module, we will run **R** and **SQL** code in the code editor **VSCode**. We will also use **GitHub Copilot** (free to use for students).

On Brightspace, we have uploaded a detailed guidelines on installing the relevant software and extensions (see the **Installation Instructions** page under **My Learning**).

Important: Make sure to install and set up **R**, **VSCode**, and **GitHub Copilot** *before* our first lecture in Week 1. If you have any question related to (installing) the software, please ask this Question in our Slack workspace.

Syllabus Modification Rights

I reserve the right to reasonably alter the elements of the syllabus at any time by adjusting the reading list to keep pace with the course schedule. Moreover, I may change the content of specific sessions, depending on the participants' prior knowledge and research interests. If I make adjustments, I will email all seminar participants and upload the revised syllabus to Brightspace.

Dignity and Respect

UCD is committed to the promotion of an environment for work and study which upholds the dignity and respect of all members of the UCD community and which supports your right to study and/or work in an environment which is free of any form of bullying, harassment or sexual misconduct (including sexual harassment and sexual violence).

There are a number of supports in place if you are experiencing bullying, harassment or sexual misconduct and you are strongly encouraged to come forward to seek confidential support and guidance on the range of informal options and formal options for resolving issues as appropriate. Reports of bullying, harassment or sexual misconduct can also be made anonymously through UCD's Report and Support tool.

UCD is actively promoting a culture where bullying, harassment and sexual misconduct is not tolerated, where everyone is respected and feels valued, included and that they belong in UCD.

You can find more details on UCD's Dignity and Respect Website at: <https://www.ucd.ie/equality/support/dignityrespect/>.

³<https://www.beyondbera.org/blog/2022/05/03/email-etiquette-for-students/>.

Expectations and Grading

The assignments consist of two Multiple Choice Questionnaires and a group project. Table 1 lists the deadlines and weight of each assignment for the final grade.

Assignments 1 and 2: Multiple Choice Questionnaires

Tables 2 and 3 show the grade scales used for the two MCQ assignments. More details on the MCQ assignments will be provided in class.

Table 1: Overview of Assignments and Deadlines

Date	Assignment
Week 7 (23 October)	Multiple Choice Questionnaire #1: 25%
Week 11 (20 November)	Multiple Choice Questionnaire #2: 25%
5 December	Data Report: Descriptive analysis and visual representation of insights derived from a large dataset: 50%

Table 2: Alternative Linear Conversion Grade Scale for first MCQ Assignment (with 15 questions)

Grades	Lower %	Upper %
A+	≥95	100
A	≥90	<95
A-	≥85	<90
B+	≥80	<85
B	≥75	<80
B-	≥70	<75
C+	≥65	<70
C	≥60	<65
C-	≥55	<60
D+	≥50	<55
D	≥45	<50
D—	≥40	<45
E+	≥35	<40
E	≥30	<35
E-	≥25	<30
F+	≥20	<25
F	≥15	<20
F-	≥10	<15
G+	≥5	<10
G	≥0.02	<5
G-	≥0.01	<0.02
NM	0	<0.01

ABS: No work was submitted
by the student or
the student was absent from assessment

Assignment 3: Data Report (Group Project)

In this final project, your group will craft a compelling data story based on a dataset that you find online (or optionally need to collect/scrape yourself). You will use descriptive analysis, visualisation techniques, and critical thinking to explore and explain trends, patterns, or insights from the data. The submission deadline for this assignment is **Friday, 5 December, 8pm**.

Table 3: Grade Scale Mapped to Percentage Bands and Correct Answers for Second MCQ Assignment (with 12 questions)

Grades	Lower %	Upper %	Correct
A+	≥ 95	100	12
A	≥ 90	< 95	11
B+	≥ 80	< 85	10
B	≥ 75	< 80	9
C+	≥ 65	< 70	8
C-	≥ 55	< 60	7
D+	≥ 50	< 55	6
D-	≥ 45	< 50	5
E	≥ 30	< 35	4
E-	≥ 25	< 30	3
F	≥ 15	< 20	2
G+	≥ 5	< 10	1
NM	0	< 0.01	0
ABS: No work was submitted by the student or the student was absent from assessment			

Groups will be allocated as follows: students registered for the Connected_Politics module will work in the same group as in Connected_Politics. The remaining students will be randomly assigned to a group. We try to ensure that all group members are registered for a degree in the same School. Further information will be provided in class.

Deliverable

A written report in the form of a blog post-style data story. The story should present a clear narrative that takes the reader through the data analysis process, supported by descriptive graphs and thoughtful commentary. A good example of the end goal of the project are the [Data Stories](#) blog from the Swiss National Science Foundation, for example [this blog post on funding for early-career researchers](#). We will also upload excellent examples from previous years to Brightspace to give you a better idea of the scope and structure.

Individual Accountability and Group Contribution Assessment

Group work is a core part of this module because it cultivates essential skills such as collaboration, problem-solving, and communication, which are vital for professional success. The policy on individual accountability within groups is designed to prevent issues related to uneven contributions in group assignments. Each group member is required to submit a document along with their group project that clearly specifies their individual contributions. This document must transparently detail each student's roles, responsibilities, and actual input. While the aim is to award the same grade to all members of the group, divergent contributions will result in varied grades to ensure fairness. This policy not only promotes equitable grading but also deters disparities in the distribution of workload among team members. Students are encouraged to report any issues with group dynamics promptly to facilitate swift intervention and support from the course instructor.

Steps to Follow

1. Select a Dataset

- You may find a dataset online or collect your own data (this is optional). The dataset should *not* be one we have used in class. You can find links to large data repositories in the “Helpful Links” section on page 9 of this document.
- Ensure the dataset is complex enough to support meaningful analysis but manageable within the scope of the project.
- Depending on your research question, you might also collect and merge two or more datasets.
- Once you have identified a dataset, contact Mohamed Moheeb [via Email](#) or Slack for quick approval before proceeding with the analysis. You must contact Sarah three weeks before the submission deadline (i.e., no later than 15 November).

2. Define Your Research Question(s)

- Formulate 1–2 research questions that will guide your analysis. Your questions should help frame the narrative of your data story and be answerable using the data.
- Example(s): *“What are the key factors influencing student performance in higher education across different countries?”* *“What are the voting patterns in urban vs. rural areas over the last three U.S. presidential elections?”* *“How has government spending on social welfare changed in response to economic recessions?”*

3. Explore the Data

- Conduct an initial exploration of the dataset to understand its structure, key variables, and any missing data.
- Summarise your findings in a brief introductory section, giving readers context for what the dataset contains.

4. Analyse the Data

- Perform descriptive statistical analyses.
- Identify any notable relationships, patterns, or trends that help answer your research questions.

5. Create Visualisations

- Construct 3–5 clear and informative graphs that help illustrate your findings.
- Each graph should be accompanied by a brief description that explains what the graph shows and how it relates to your research question.
- Follow best practices for data visualization (clear labels, appropriate scales, and use of color) discussed in class. The readability and interpretability of the graphs will be a crucial component of this assignment.

6. Tell the Story

- Organise your findings into a coherent narrative. Your story should walk the reader through the data, describing your research question and how the analysis addresses it.
- Highlight interesting insights or unexpected findings. Where applicable, connect your findings to broader academic or policy debates.

- Keep your audience in mind – assume they have a basic understanding of the topic but not expert knowledge.

7. Conclusion

- Summarise your key takeaways. What did the data reveal? What questions remain unanswered?
- Discuss any limitations of your analysis and suggest potential avenues for future research.

8. Submission Requirements

- Word count: 1,000–1,300 words (excluding graphs and references).
- Include 3–5 descriptive graphs.
- Submit the assignment as a PDF document. Word documents will not be accepted. Make sure to save all graphs in a vector graphic format (ideally, PDF or EPS) and include these graphs in high-quality resolution.
- Submit the files containing the data and code to complete the project in a reproducible format. This means that you need to upload the raw dataset(s), and your R scripts to wrangle, (merge), and visualise the data. The code might be checked during the grading process if questions emerge.
- Include a reference section for any external data sources or literature cited.

Grading Criteria

- **Data Exploration & Analysis (40%)**: Depth of analysis and insight into the dataset.
- **Visualisations (35%)**: Quality, clarity, and appropriateness of graphs.
- **Narrative & Storytelling (25%)**: Coherence and structure of the data story. Make sure to follow [this guide](#) on how to structure and write paragraphs.

Helpful Links

Here are some recommended repositories where you can find datasets for your project:

- [PolData](#): PolData provides a comprehensive overview of political datasets.
- [Ireland's Open Data Portal](#): Irish Public Sector data in open, free and reusable formats.
- [Harvard Dataverse](#): One of the world's largest open-access repositories, with data from various research disciplines including social sciences, law, and health.
- [Google Dataset Search](#): Google's search engine for publicly available datasets across a wide range of fields.
- [World Bank Open Data](#): Global development data on economics, demographics, and more.
- [UCI Machine Learning Repository](#): A large repository with datasets in health, economics, and more, often used for machine learning tasks.
- [United Nations Data](#): Global statistics covering population, education, employment, and development indicators.
- [European Union Open Data Portal](#): Public data on various sectors such as economics, transportation, and environment.

- [Pew Research Center](#): Survey data covering politics, social issues, media, and more.
 - [Open Data UK](#): Datasets provided by UK government departments, covering various topics including health, crime, and transportation.
 - [OECD Data](#): Data on global economic and social issues, including education and health.
-

Week 1: Introduction to R and VSCode (11 September)

Instructor: Stefan Müller

Mandatory Readings:

- H. Wickham, M. Çetinkaya-Runde, and G. Grolemund (2023). *R for Data Science: Import, Tidy, Transform, Visualize, and Model Data*. 2nd edition. Sebastopol: O'Reilly: chapters 2 and 27.
- C. Ismay and A. Y. Kim (2025). *Statistical Inference via Data Science: A Modern Dive into R and the tidyverse*. 2nd edition. Boca Raton: CRC Press: chapter 1.

Important: Please follow the installation guide (see the [Installation Instructions](#) page under [My Learning](#)), available on Brightspace, in order to set up R, VSCode, and GitHub Copilot. Make sure to install the software *before* our first lecture. If you have any question related to (installing) the software, please ask this question through our Slack workspace.

Week 2: Software, Project Management, and Replicability (18 September)

Instructor: Stefan Müller

Mandatory Readings:

- H. Wickham, M. Çetinkaya-Runde, and G. Grolemund (2023). *R for Data Science: Import, Tidy, Transform, Visualize, and Model Data*. 2nd edition. Sebastopol: O'Reilly: chapters 3 and 5.
- Posit Team (2025). *Tutorial: Hello, Quarto*. 2025. URL: <https://quarto.org/docs/get-started/hello/positron.html>: select your preferred tool.

Optional Readings:

- J. Bryan (2020). *Happy Git and GitHub for the useR*. URL: <https://happygitwithr.com> (skim if you are interested in version control).
- W. G. Jacoby and R. N. Lupton (2016). *American Journal of Political Science: Guidelines for Preparing Replication Files. Version 2.1*.

Week 3: Importing, Summarising, Transforming, and Merging Data (25 September)

Instructor: Stefan Müller

Mandatory Readings:

- H. Wickham, M. Çetinkaya-Runde, and G. Grolemund (2023). *R for Data Science: Import, Tidy, Transform, Visualize, and Model Data*. 2nd edition. Sebastopol: O'Reilly: chapters 7 and 8.

Week 4: Using LLMs as Research Tools (2 October)

Instructor: Stefan Müller

Note: please read Wratil and Hobolt (2019) before class to get a general understanding of the dataset we will analyse. Afterwards, check out the codebook and class guidelines (uploaded on Brightspace in the folder for week 4) and make sure you have a free account for ChatGPT, Claude, or another generative AI chatbot (e.g., Mistral, Gemini, Microsoft Copilot).

Mandatory Readings:

- C. Wratil and S. B. Hobolt (2019). “Public Deliberations in the Council of the European Union: Introducing and Validating DICEU”. *European Union Politics* 3 (511-531).
- S. Verdi (2024). *How to Use AI Coding Tools to Learn a New Programming Language*. URL: <https://github.blog/developer-skills/programming-languages-and-frameworks/how-to-use-ai-coding-tools-to-learn-a-new-programming-language/>.

Optional Readings:

- S. Arnett (2025). *In a First, Google has Released Data on How Much Energy an AI Prompt Uses*. MIT Technology Review. URL: <https://technologyreview.com/2025/08/21/1122288/google-gemini-ai-energy/>.
- GitHub Docs (2025). *Asking GitHub Copilot Questions in your IDE*. URL: <https://docs.github.com/en/copilot/using-github-copilot/asking-github-copilot-questions-in-your-ide>.
- D. Autor and J. Manyika (2025). *A Better Way to Think About AI: Artificial Intelligence is Ready to Collaborate. Why Fixate on Automation?*. The Atlantic. URL: <https://www.theatlantic.com/technology/archive/2025/08/ai-job-loss-human-enhancement-google/683963/>.

Week 5: Principles of Data Visualisation (9 October)

Instructor: Stefan Müller

Mandatory Reading:

- C. O. Wilke (2019). *Fundamentals of Data Visualization: A Primer On Making Informative and Compelling Figures*. Sebastopol: O'Reilly: chapters 2, 17, 19, and 29.

Optional Reading:

- C. O. Wilke (2019). *Fundamentals of Data Visualization: A Primer On Making Informative and Compelling Figures*. Sebastopol: O'Reilly: chapters 1, 4, 18, 26.

Week 6: Explorative Data Visualisation (16 October)

Instructor: Stefan Müller

Mandatory Readings:

- K. Healy (2019). *Data Visualization: A Practical Introduction*. Princeton: Princeton University Press: chapters 1, 3, and 4.

Week 7: Advanced, Intuitive, and Accessible Data Visualisation (23 October)

Instructor: Stefan Müller

Mandatory Readings:

- K. Healy (2019). *Data Visualization: A Practical Introduction*. Princeton: Princeton University Press: chapters 5 and 8.
- J. P. Kastellec (2025). *Practical Advice for Producing Better Graphs*. URL: https://jkastellec.scholar.princeton.edu/sites/g/files/toruqf3871/files/documents/kastellec_graphs_practical_tips_v2.pdf.

Optional Readings:

- S. Dattani (2025). *Saloni's Guide to Data Visualization: Why Data Visualization Matters, and How to Make Charts More Effective, Clear, Transparent, and Sometimes, Beautiful*. URL: <https://www.scientificdiscovery.dev/p/salonis-guide-to-data-visualization>.
- J. L. Steenwyk and A. Rokas (2021). “ggpubfigs: Colorblind-Friendly Color Palettes and ggplot2 Graphic System Extensions for Publication-Quality Scientific Figures”. *Microbiology Resource Announcements* 10 (44): e00871–21

Week 8: No Class — Reading Break (30 October)

Week 9: SQL: Organising and Managing Data (6 November)

Instructor: Jos Dornschneider-Elkink

Mandatory Readings:

- N. B. Weidmann (2023). *Data Management for Social Sciences: From Files to Databases*. Cambridge: Cambridge University Press: chapter 8.
- A. Beaulieu (2020). *Learning SQL: Generate, Manipulate, and Retrieve Data*. 3rd edition. Sebastopol: O'Reilly Media: 1–50.

Note: Beaulieu (2020) is accessible online through the UCD Library.

Week 10: SQL: Retrieving, Joining, and Summarising Data (13 November)

Instructor: Jos Dornschneider-Elkink

Mandatory Readings:

- N. B. Weidmann (2023). *Data Management for Social Sciences: From Files to Databases*. Cambridge: Cambridge University Press: chapter 9.
- A. Beaulieu (2020). *Learning SQL: Generate, Manipulate, and Retrieve Data*. 3rd edition. Sebastopol: O'Reilly Media: 51–100; 147–160.

Note: Beaulieu (2020) is accessible online through the UCD Library.

Week 11: Introduction to Web Scraping (20 November)

Instructor: Stefan Müller

Mandatory Readings:

- R. Alexander (2023). *Telling Stories with Data: With Applications in R*. New York: CRC Press: chapter 7.

Optional Reading:

- H. Mühleisen and K. Müller (2025). *duckplyr: A ‘DuckDB’-Backed Version of ‘dplyr’*. R package version 1.1.1. URL: <https://CRAN.R-project.org/package=duckplyr>.

Week 12: Advanced Web Scraping (27 November)

Instructor: Stefan Müller

Mandatory Reading:

- H. Wickham (2022). *rvest: Selector Gadget*. URL: <https://rvest.tidyverse.org/articles/selectorgadget.html>.

Optional Reading:

- M. A. Brown, A. Gruen, G. Maldoff, S. Messing, Z. Sanderson, and M. Zimmer (2025). “Web Scraping for Research: Legal, Ethical, Institutional, and Scientific Considerations”. *Big Data & Society* (4): 1–17.

References

- Alexander, R. (2023). *Telling Stories with Data: With Applications in R*. New York: CRC Press.
- Arnett, S. (2025). *In a First, Google has Released Data on How Much Energy an AI Prompt Uses*. MIT Technology Review. URL: <https://technologyreview.com/2025/08/21/1122288/google-gemini-ai-energy/>.
- Autor, D. and J. Manyika (2025). *A Better Way to Think About AI: Artificial Intelligence is Ready to Collaborate. Why Fixate on Automation?*. The Atlantic. URL: <https://www.theatlantic.com/technology/archive/2025/08/ai-job-loss-human-enhancement-google/683963/>.
- Beaulieu, A. (2020). *Learning SQL: Generate, Manipulate, and Retrieve Data*. 3rd edition. Sebastopol: O’Reilly Media.
- Brown, M. A., A. Gruen, G. Maldoff, S. Messing, Z. Sanderson, and M. Zimmer (2025). “Web Scraping for Research: Legal, Ethical, Institutional, and Scientific Considerations”. *Big Data & Society* (4): 1–17.
- Bryan, J. (2020). *Happy Git and GitHub for the useR*. URL: <https://happygitwithr.com>.
- Dattani, S. (2025). *Saloni’s Guide to Data Visualization: Why Data Visualization Matters, and How to Make Charts More Effective, Clear, Transparent, and Sometimes, Beautiful*. URL: <https://www.scientificdiscovery.dev/p/salonis-guide-to-data-visualization>.
- GitHub Docs (2025). *Asking GitHub Copilot Questions in your IDE*. URL: <https://docs.github.com/en/copilot/using-github-copilot/asking-github-copilot-questions-in-your-ide>.
- Healy, K. (2019). *Data Visualization: A Practical Introduction*. Princeton: Princeton University Press.
- Ismay, C. and A. Y. Kim (2025). *Statistical Inference via Data Science: A Modern Dive into R and the tidyverse*. 2nd edition. Boca Raton: CRC Press.
- Jacoby, W. G. and R. N. Lupton (2016). *American Journal of Political Science: Guidelines for Preparing Replication Files. Version 2.1*.

- Kastellec, J. P. (2025). *Practical Advice for Producing Better Graphs*. URL: https://jkastellec.scholar.princeton.edu/sites/g/files/toruqf3871/files/documents/kastellec_graphs_practical_tips_v2.pdf.
- Monroe-White, T., D. Kozlowski, L. Domenech, C. Pradier, G. Rosati, and N. S. Shokida (2024). *Critical Computational Social Sciences*.
- Mühleisen, H. and K. Müller (2025). *duckplyr: A ‘DuckDB’-Backed Version of ‘dplyr’*. R package version 1.1.1. URL: <https://CRAN.R-project.org/package=duckplyr>.
- Müller, S. (2023). “How Slack Facilitates Communication and Collaboration in Seminars and Project-Based Courses”. *Journal of Educational Technology Systems* 51 (3): 303–316.
- Posit Team (2025). *Tutorial: Hello, Quarto*. 2025. URL: <https://quarto.org/docs/get-started/hello/positron.html>.
- Steenwyk, J. L. and A. Rokas (2021). “ggpubfigs: Colorblind-Friendly Color Palettes and ggplot2 Graphic System Extensions for Publication-Quality Scientific Figures”. *Microbiology Resource Announcements* 10 (44): e00871–21.
- Verdi, S. (2024). *How to Use AI Coding Tools to Learn a New Programming Language*. URL: <https://github.blog/developer-skills/programming-languages-and-frameworks/how-to-use-ai-coding-tools-to-learn-a-new-programming-language/>.
- Weidmann, N. B. (2023). *Data Management for Social Sciences: From Files to Databases*. Cambridge: Cambridge University Press.
- Wickham, H. (2022). *rvest: Selector Gadget*. URL: <https://rvest.tidyverse.org/articles/selectorgadget.html>.
- Wickham, H., M. Çetinkaya-Runde, and G. Grolemund (2023). *R for Data Science: Import, Tidy, Transform, Visualize, and Model Data*. 2nd edition. Sebastopol: O’Reilly.
- Wilke, C. O. (2019). *Fundamentals of Data Visualization: A Primer On Making Informative and Compelling Figures*. Sebastopol: O’Reilly.
- Wratil, C. and S. B. Hobolt (2019). “Public Deliberations in the Council of the European Union: Introducing and Validating DICEU”. *European Union Politics* 3 (511–531).