



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>

Quantitative Text Analysis (POL42050)

Level 4 Module; Spring Trimester 2026

Version: April 29, 2026

Latest version at: <https://muellerstefan.net/teaching/2026-spring-qta.pdf>

Time: Wed, 09:00–10:50

Location: L3.15-LEA (Centre for Fut. Learning)

Credits: 10.0

Format: Lecture and computer labs

Module coordinator: Stefan Müller

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

Office: Newman Building, G312

Office hours: Wed, 13:00–13:45 ([sign up here](#))

Course Content

Computational text analysis has become increasingly popular in political science in recent years. With the vast availability of text data on the web, political scientists increasingly view quantitative text analysis (or “text as data”) as a valuable approach for studying various forms of social and political behaviour.

This module introduces social science students to the quantitative analysis of textual data.¹ The course is structured in two complementary parts, each addressing different approaches to text analysis with distinct methodological properties.

The *first part* covers classical quantitative text analysis approaches that prioritise interpretability and transparency. It covers the theoretical foundations, practical applications, and technical implementations of these text-as-data methods using the R statistical programming language, and how to validate such approaches. Classical methods follow a three-step framework: identifying texts and units of analysis; extracting measurable features from these texts and converting them into a quantitative feature matrix; and analysing this matrix using statistical techniques. These include dictionary-based approaches, supervised document classification, scaling models, and topic modelling. These approaches offer clear, interpretable results that can be explained and validated at each stage of the analysis, making them particularly valuable for rigorous social science research.

The *second part* introduces more powerful, state-of-the-art techniques that leverage large language models. Students will gain hands-on experience with word embeddings, transformer models, and generative AI approaches. Whilst these methods often deliver superior predictive performance and

¹**Note to module instructors:** You are welcome to reuse and build upon this document for your teaching. If you design a module based on this document, I would greatly appreciate it if you could cite and include a link to this syllabus. I would also be grateful if you could contact me to let me know how you have used or adapted these materials.

capture nuanced linguistic patterns, they frequently operate as “black boxes” with limited interpretability. The module incorporates the Hugging Face Python infrastructure, a leading resource for implementing transformer models and other state-of-the-art natural language processing tools. Students will learn how to fine-tune pre-trained models and leverage generative AI for text analysis, whilst understanding the important trade-offs between performance and interpretability.

Each session combines lectures with practical exercises, allowing students to apply these methods to political texts. These exercises address real-world challenges at each stage of the research process. By engaging with both classical and modern approaches, students will gain insights into the full spectrum of text analysis methods available to political science researchers and will be equipped to select appropriate techniques for their own research questions.

Learning Outcomes

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

1. Understand fundamental issues in quantitative text analysis, such as inter-coder agreement, reliability, validation, accuracy, and precision.
2. Master classical text-as-data approaches by converting texts into quantitative matrices of features and analysing them using statistical methods, scaling models, and topic modelling.
3. Understand the strengths and limitations of classical approaches, particularly their interpretability, and recognise contexts where they are most appropriately applied.
4. Apply and fine-tune modern neural network approaches, including transformer models, to text analysis tasks.
5. Understand the capabilities and limitations of generative AI for text analysis, including the trade-off between performance and interpretability.
6. Use human coding of texts to train and evaluate both supervised classifiers and fine-tuned transformer models.
7. Select and justify appropriate text analysis techniques (classical or modern) for their own research questions and text corpora.
8. Critically evaluate social science research that uses text analysis methods, assessing methodological choices and the appropriateness of different techniques.

Prerequisites

Prior familiarity with the statistical programming language R (or Python) is essential for this course. The module uses R extensively for implementing classical text analysis methods and Python for modern neural approaches. Without prior programming experience, students will find it difficult to follow the practical exercises, complete assignments, and work with text corpora in the research paper component.

If you have already completed a module on computational text analysis in the past, I recommend taking [AI and Large Language Models \(POL42560\)](#) instead.

General Readings

The seminar does not build on a single textbook, but relies on papers and book chapters. The following books and articles are recommended for a general overview of quantitative text analysis, natural language processing, and computational social science.

- K. L. Nielbo, F. Karsdorp, M. Wevers, A. Lassche, B. R. B., M. Kestemont, and N. Tahmasebi (2024). “Quantitative Text Analysis”. *Nature Reviews Methods Primers* 2 (24).
- K. Benoit (2020). “Text as Data: An Overview”. *Handbook of Research Methods in Political Science and International Relations*. Ed. by L. Curini and R. Franzese. Thousand Oaks: Sage: 461–497.
- K. Watanabe and S. Müller (2023). *Quanteda Tutorials*. URL: <https://tutorials.quanteda.io>.
- D. S. Stoltz and M. A. Taylor (2024). *Mapping Texts: Computational Text Analysis for the Social Sciences*. Oxford: Oxford University Press.
- E. Hvitfeldt and J. Silge (2021). *Supervised Machine Learning For Text Analysis in R*. Boca Raton: CRC Press.
- J. Grimmer, M. E. Roberts, and B. M. Stewart (2022). *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press.
- D. Jurafsky and J. H. Martin (2025). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models*. 3rd edition.

Technical Background

The following books are helpful to refresh and extend knowledge of R, Python, Quarto, data visualisation, and transformer-based NLP, or to move beyond the content covered in the module. Most books listed here are published in print and also freely available online.

R, Quarto, and Data Science

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

Statistical Analysis and Causal Inference

- N. Huntington-Klein (2025). *The Effect: An Introduction to Research Design and Causality*. 2nd edition. Boca Raton: CRC Press.
- S. Cunningham (2021). *Causal Inference: The Mixtape*. New Haven: Yale University Press.

Python for Social Scientists

- A. Turrell (2024). *Coding for Economists*. URL: <https://aeturrell.github.io/coding-for-economists>.

Data Visualisation

- K. Healy (2019). *Data Visualization: A Practical Introduction*. Princeton: Princeton University Press.

NLP with Transformers

- L. Tunstall, L. von Werra, and T. Wolf (2022). *Natural Language Processing with Transformers: Building Language Applications with Hugging Face*. Beijing: O'Reilly.

Software and Packages

In this module, we will work with [R](#) and [Python](#) code. You are welcome to use any integrated development environment (IDE) of your choice, including [VSCode](#), [RStudio](#), or [Positron](#). Each IDE has distinct advantages: VSCode is lightweight and flexible for multiple languages, RStudio excels at R-focused work, and Positron is a newer IDE that combines the strengths of both for data science.

The lectures will demonstrate code using [Positron](#) because it integrates both R and Python seamlessly in a single environment, reducing the friction of switching between languages and IDEs. This integration is particularly valuable when working with both classical text analysis in R and modern neural approaches in Python.

Brightspace contains detailed guidelines for installing R, Python, and your preferred IDE, including setup instructions specifically for VSCode and Positron (see the [Installation Instructions](#) page under [My Learning](#)). We will also use [GitHub Copilot](#) (free to use for students) to assist with code generation and debugging.

Important: Make sure to install and set up R, Python, your preferred IDE, and GitHub Copilot before our first lecture in Week 1. If you have any question related to installing the software, please post it in our Slack workspace.

We will use the following packages in this module:

- Quantitative text analysis: [quanteda](#), [quanteda.textstats](#), [quanteda.textmodels](#)
- Importing and converting text files: [readtext](#); [markitdown](#)
- Topic models: [keyATM](#), [BERTopic](#)
- Word embeddings: [wordvector](#)
- Data wrangling and visualisation: [tidyverse](#) (esp. [dplyr](#), [tidyr](#), [lubridate](#), and [ggplot2](#))
- Creating documents and reports: [Quarto](#)
- Transformer models: [HuggingFace](#) and [transformers](#) library (Python)
- R Packages for Classification using Large Language Models: [rollama](#), [quallmer](#)

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 familiarize yourself with the definition of plagiarism on UCD’s website² and make sure not to engage in it.

Late Submission Policy

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

- Coursework submitted up to 10 working days past the due date will be graded but a grade penalty will be applied.
 - Coursework submitted at any time up to and including 5 working days after the due date will have the grade awarded reduced by one grade point (for example, from *B–* to *C+*).
 - Coursework submitted more than 5 working days but up to and including 10 working days after the due date will have the grade reduced by two grade points (for example, from *B–* to *C*).
- Coursework received more than 10 working days after the due date will not be accepted or graded. Regulations regarding extenuating circumstances apply.

Office Hours

My office hours take place on Wednesday from 13:00–13:45, either in person (Room G312, Newman Building) or online. Please sign up for a meeting at <https://calendly.com/mueller-ucd/office-hours>.

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 in this module.³ Make sure to create a Slack account before the first seminar and join the Slack workspace. If you have a question that involves code or concepts, please share your question on Slack.

If you struggle to solve problems relating to R, Python, or concepts, 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. Try to summarise the problem in your own words and then search online. If the problem relates to R, add `rstats` to your search query. For example: `how to import csv file in rstats`. Similarly, for Python questions, search for relevant documentation or Stack Overflow discussions.
2. If your R or Python code returns an error, copy the error message and search for it online. For example, searching “`Error: Can’t subset columns that don’t exist`” will often lead directly to solutions on Stack Overflow or package documentation.
3. Query a large language model (LLM) such as Claude, ChatGPT, GitHub Copilot, or Gemini with your problem. You can paste your code and error message and ask for debugging help. Be specific about what you are trying to do and what error you received. You can also use

²<https://libguides.ucd.ie/academicintegrity>.

³I have had very positive experiences with Slack in my modules. Müller (2023) discusses both the advantages and shortcomings of Slack for teaching and learning.

local models if you prefer. LLMs are particularly effective at debugging code, explaining error messages, and suggesting alternative approaches.

→ 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 I will help you.

Use of Generative Artificial Intelligence (AI) Tools

I encourage the use of generative AI tools when completing the assignments for this module but all work relying on AI-generated content must adhere to the highest academic standards. 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, and factually correct 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.

Documenting AI Use (Mandatory)

Since generative AI is still a relatively novel tool in an academic context, its capabilities are not yet fully understood and are evolving rapidly. What is infeasible today may become trivial tomorrow, subject to the academic standards outlined above. To address this uncertainty, you must provide a transparent account of any generative AI tools used and how they were used in a mandatory appendix to your assignment. This appendix forms part of the assessment, with marks awarded for clear and accurate communication of the methods used to generate content.

For each instance in which a generative AI tool is used, you are required to provide:

1. An in-text citation or footnote, for example:
 - “Some AI-generated text (GPT-4o 2025)”
 - “Some AI-generated text⁴”
2. A bibliographic reference to the tool used, including the date of access.
3. A corresponding entry in the mandatory AI appendix detailing the use of the tool (see Table 1).

The mandatory AI appendix must include:

- An AI usage statement describing the tool used, the prompts provided, and how the tool contributed.
- Any retained AI-generated text clearly marked with green highlighting or included as an appendix.
- Complete and verifiable sources for all claims, with no fabricated citations.

⁴Text based on content generated by OpenAI’s GPT-5.2 on 5 January 2026. See Appendix 1 for details.

Table 1: Example table demonstrating how generative AI was used to complete the research paper

AI Tool	Explanation	Prompt used
ChatGPT	Topic brainstorming	“Provide an overview of potential political science research questions that could be answered using quantitative text analysis methods.”
Elicit	Literature search	“Compile a list of academic publications detailing advantages and shortcomings of topic models.”
Notebook LM	Synthesising research papers	“Create a study guide summarising the key findings from these papers on sentiment analysis validation.”
GitHub Copilot	Debugging transformer fine-tuning	“I’m getting RuntimeError: Expected all tensors to be on the same device when fine-tuning my model. Can you help me debug this?”

Generative AI Tools

Below I have listed some AI tools that might help you drafting your research paper:

- Brainstorming and Coding
 - ChatGPT: <https://chat.openai.com>
 - Claude: <https://claude.ai>
 - Gemini: <https://gemini.google.com/app>
- Literature Discovery
 - Elicit: <https://elicit.org>
 - Notebook LM: <https://notebooklm.google>
- Structure and Drafting
 - Grammarly: <https://www.grammarly.com>
 - Quillbot: <https://quillbot.com/>
 - Jenni: <https://jenni.ai>

Alignment with UCD’s AI Guidelines

This module use falls under UCD’s *green* category for the usage of generative AI. More specifically, you may use AI tools like the ones above for:

- Brainstorming topics, outlining, and planning study schedules.
- Clarifying definitions, rephrasing your own text, grammar or surface edits.
- Editing code.
- Generating practice questions, summarising given readings (which cannot replace in-depth preparation).

As outlined on page 6 above, you **must** disclose your use of AI in the homework assignment and research paper.

More information on the AI tools supported by UCD:

- <https://www.ucd.ie/itservices/ourservices/communicationcollaboration/officeproductivity/aiservices/>
- <https://www.ucd.ie/teaching/resources/generativeai/>

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/>.

Module Structure

Week 1: Introduction to Quantitative Text Analysis (21 January 2026)	14
Week 2: Workflow, R and Quarto (28 January 2026)	15
Week 3: Tokenisation and Document-Feature Matrix (4 February 2026)	16
Week 4: Dictionaries and Sentiment: Old and New Approaches (11 February 2026)	16
Week 5: Describing and Comparing Texts (18 February 2026)	17
Week 6: Supervised, Unsupervised, and Semi-Supervised Scaling (25 February 2026)	18
Week 7: Topic Models (4 March 2026)	19
Week 8: Retrieving, Loading, and Wrangling Text Corpora (25 March 2026)	19
Week 9: Fundamentals of Supervised Document Classification (1 April)	20
Week 10: Transformers: Fundamentals (8 April 2026)	21
Week 11: Fine-Tuning Transformer Models for Classification (15 April 2026)	22
Week 12: Large Language Models: Capabilities and Implications (22 April 2026)	22

Expectations and Grading

- Students are expected to read all papers or chapters assigned under **Mandatory Readings**. These readings serve as the basis for in-class discussions about the advantages, disadvantages, and applicability of the various approaches to social science questions. For each session, I also assign a variety of optional readings. I strongly encourage students to (at least) skim these readings. Both the required and the optional readings consist of technical readings and at least one practical application of the respective method.
- The first assessment is an in-class **Multiple Choice Questionnaire**, administered in Week 5 of class (18 February). Students will answer questions drawn from the lecture slides and lab exercises. Use of external notes, reference materials, or AI tools are prohibited. The assignment counts towards 15% of the final grade.
- Students submit one **Homework assignment**, which counts 25% of the final grade. The assignments will be distributed as a [Quarto](#) file 14 days before the submission deadline. Students fill in the answers and solutions in the same [Quarto](#) file, render it as an `html` file, and submit it via Brightspace. Only rendered `html` files will be accepted! The Homework will be submitted by 2 April 2026 (Week 9). More details on the homework will be provided in the first session(s) of the course.

Table 2: ‘Alternative Linear Conversion Grade Scale’ used for multiple choice questionnaire

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		

- Students also submit a short **Research Paper** of 3,000 words (excluding references and appendices). Please keep footnotes to a minimum. They are intended for supplementary information only and should not be used to extend the word count or to present material that belongs in the main text. The research paper counts towards 60% of the final grade and must be submitted by 4 May 2026. During the module, students will be required to briefly outline the research question they want to test in the research paper and describe which textual data they will use for testing this question.

Table 3: Grade conversion scale for homework assignment

Homeworks	UCD Grade
97–100 points	A+
94–96 points	A
91–93 points	A–
88–90 points	B+
85–87 points	B
83–84 points	B
80–82 points	C+
77–79 points	C
74–76 points	C–
71–73 points	D+
68–70 points	D
65–67 points	E+
54–64 points	E
44–53 points	E–
0–32 points	F

In the research paper, the students should succinctly but clearly write up the results of a small research project using quantitative text analysis methods discussed in this module. Students are free to collect their own data or use existing data. Creativity is encouraged. Students are free to answer questions from all fields of political science or related disciplines but must justify their choice and the relevance of the question. This paper should contain the following elements:

1. Introduction and research question: introduction to the topic, research question, and relevance.
2. Expectations: a concise overview of the theoretical expectation(s) that will be tested in the results section.
3. Data and methods: description of the data sources as well as the methods employed.
4. Results: a discussion (with figures and tables) of the results of the analysis. This section forms the bulk of the paper.
5. Conclusion: a brief evaluation of the results and steps to push the research forward.

Replicability and Transparency: As part of your submission, you must upload all data and code (R scripts, Python notebooks, or Quarto files) used to create every table and figure in your paper. This requirement reflects best practices in computational social science and ensures that your work is transparent and replicable. For further guidance on best practices for reproducible research, see Chapter 3 of R. Alexander (2023). *Telling Stories with Data: With Applications in R*. New York: CRC Press.

Table 4: Overview of assignments

Date	Assignment
Wednesday, 18 February 2026 (Week 5)	In-class Multiple Choice Questionnaire (15%)
Thursday, 2 April 2026 (Week 9)	Homework (25%)
Monday, 4 May 2026	Research Paper (60%)

Grading Criteria for Research Paper

In essence, markers assess four crucial elements in any answer:

- Analysis/understanding
- Extent and use of reading
- Organisation/structure
- Writing proficiency

The various grades/classifications listed below reflect the extent to which an answer displays essential features of each of these elements (and their relative weighting). At its simplest: the better the analysis, the wider the range of appropriate sources consulted, the greater the understanding of the materials read, the clearer the writing style, and the more structured the argument, the higher will be the mark.

The following provides an indicative outline of the criteria used by markers to award a particular grade/classification. If you are in any confusion about how to correctly approach referencing and bibliography issues, please refer to the following guidelines: APSA Committee on Publications (2018). *Style Manual for Political Science (Revised 2018 Version)*. URL: <https://connect.apsanet.org/stylemanual/>.

Proper referencing is ESSENTIAL in a good assignment. I strongly recommend using bibliography management software such as [Zotero](#) or maintaining a `.bib` file when writing your research paper in Quarto or L^AT_EX. These tools ensure consistent citation formatting across your document, reduce referencing errors, and are valuable skills for academic and professional writing.

Grade: A (Excellent Performance)

A deep and systematic engagement with the assessment task, with consistently impressive demonstration of a comprehensive mastery of the subject matter, reflecting:

- A deep and broad knowledge and critical insight as well as extensive reading
- A critical and comprehensive appreciation of the relevant literature or theoretical, technical or professional framework
- An exceptional ability to organise, analyse and present arguments fluently and lucidly with a high level of critical analysis, amply supported by evidence, citation or quotation;
- A highly-developed capacity for original, creative and logical thinking
- An extensive and detailed knowledge of the subject matter
- A highly-developed ability to apply this knowledge to the task set
- Evidence of extensive background reading
- Clear, fluent, stimulating and original expression
- Excellent presentation (spelling, grammar, graphical) with minimal or no presentation errors
- Referencing style consistently executed in recognised style

Grade: B (Very Good Performance)

A thorough and well organised response to the assessment task, demonstrating:

- A thorough familiarity with the relevant literature or theoretical, technical or professional framework
- Well-developed capacity to analyse issues, organise material, present arguments clearly and cogently well supported by evidence, citation or quotation;

- Some original insights and capacity for creative and logical thinking
- A broad knowledge of the subject matter
- Considerable strength in applying that knowledge to the task set
- Evidence of substantial background reading
- Clear and fluent expression
- Quality presentation with few presentation errors
- Referencing style for the most part consistently executed in recognised style

Grade: C (Good Performance)

An intellectually competent and factually sound answer, marked by:

- Evidence of a reasonable familiarity with the relevant literature or theoretical, technical or professional framework
- Good developed arguments, but more statements of ideas
- Arguments or statements adequately but not well supported by evidence, citation or quotation
- Some critical awareness and analytical qualities
- Some evidence of capacity for original and logical thinking
- Adequate but not complete knowledge of the subject matter
- Omission of some important subject matter or the appearance of several minor errors
- Capacity to apply knowledge appropriately to the task albeit with some errors
- Evidence of some background reading
- Clear expression with few areas of confusion
- Writing of sufficient quality to convey meaning but some lack of fluency and command of suitable vocabulary
- Good presentation with some presentation errors
- Referencing style executed in recognised style, but with some errors

Grade: D (Satisfactory Performance)

An acceptable level of intellectual engagement with the assessment task showing:

- Some familiarity with the relevant literature or theoretical, technical or professional framework
- Mostly statements of ideas, with limited development of argument
- Limited use of evidence, citation or quotation
- Limited critical awareness displayed
- Limited evidence of capacity for original and logical thinking
- Basic grasp of subject matter, but somewhat lacking in focus and structure
- Main points covered but insufficient detail
- Some effort to apply knowledge to the task but only a basic capacity or understanding displayed
- Little or no evidence of background reading
- Several minor errors or one major error
- Satisfactory presentation with an acceptable level of presentation errors
- Referencing style inconsistent

Grade: D– (Acceptable)

The minimum acceptable of intellectual engagement with the assessment task which:

- The minimum acceptable appreciation of the relevant literature or theoretical, technical or professional framework
- Ideas largely expressed as statements, with little or no developed or structured argument
- Minimum acceptable use of evidence, citation or quotation
- Little or no analysis or critical awareness displayed or is only partially successful
- Little or no demonstrated capacity for original and logical thinking
- Shows a basic grasp of subject matter but may be poorly focused or badly structured or contain irrelevant material
- Has one major error and some minor errors
- Demonstrates the capacity to complete only moderately difficult tasks related to the subject material
- No evidence of background reading
- Displays the minimum acceptable standard of presentation (spelling, grammar, graphical)
- Referencing inconsistent with major errors

Grade: E (Fail [marginal])

A factually sound answer with a partially successful, but not entirely acceptable, attempt to:

- Integrate factual knowledge into a broader literature or theoretical, technical or professional framework develop arguments
- Support ideas or arguments with evidence, citation or quotation
- Engages with the subject matter or problem set, despite major deficiencies in structure, relevance or focus
- Has two major error and some minor errors
- Demonstrates the capacity to complete only part of, or the simpler elements of, the task
- An incomplete or rushed answer (e.g. the use of bullet points through part/all of answer)
- Little or no referencing style evident

Grade: F (Fail [unacceptable])

An unacceptable level of intellectual engagement with the assessment task, with:

- No appreciation of the relevant literature or theoretical, technical or professional framework
- No developed or structured argument
- No use of evidence, citation or quotation
- No analysis or critical awareness displayed or is only partially successful
- No demonstrated capacity for original and logical thinking
- A failure to address the question resulting in a largely irrelevant answer or material of marginal relevance predominating
- A display of some knowledge of material relative to the question posed, but with very serious omissions / errors and/or major inaccuracies included in answer
- Solutions offered to a very limited portion of the problem set
- An answer unacceptably incomplete (e.g. for lack of time)
- A random and undisciplined development, layout or presentation
- Unacceptable standards of presentation, such as grammar, spelling or graphical presentation

- Evidence of substantial plagiarism
- No referencing style evident

Grade: G (Fail [wholly unacceptable])

No intellectual engagement with the assessment task

- Complete failure to address the question resulting in an entirely irrelevant answer
- Little or no knowledge displayed relative to the question posed
- Little or no solution offered for the problem set
- Evidence of extensive plagiarism
- No referencing style evident

Grade: NG (No Grade)

No work was submitted by the student or student was absent from the assessment, or work submitted did not merit a grade.

Expectations for Data Interpretation, Visualisation, and Reporting

For assignments employing quantitative text analysis methods, outstanding work (Grade A) demonstrates the following characteristics:

Data Interpretation: Deep understanding of what the chosen text analysis method (dictionary-based, supervised classification, scaling, topic modelling, transformers) actually measures and its interpretive implications. Results are contextualised within the broader text analysis literature with critical reflection on generalisability and limitations. Complex methodological trade-offs are identified, such as relationships between feature selection and interpretability or the influence of training data characteristics on classifier performance.

Visualisation: Multiple visualisation types employed strategically to reveal different data dimensions. Figures are exceptionally clear and self-explanatory with minimal reliance on surrounding text. Professional standards are consistently applied: appropriate colour schemes, proper labelling, legends, and axis descriptions. Where applicable, uncertainty estimates are integrated (confidence intervals, robustness checks). Visualisations are publication-ready.

Reporting: Narrative seamlessly integrates quantitative findings with textual evidence. Summary statistics, classifications, or topic descriptions are accompanied by representative document excerpts or characteristic word lists that make the analysis concrete and interpretable. Technical decisions are clearly communicated with discussion of their implications (e.g., number of topics chosen, dictionary refinement process, classification thresholds). Results are reported with appropriate precision, and limitations are acknowledged and discussed in relation to substantive conclusions.

Week 1: Introduction to Quantitative Text Analysis (21 January 2026)

- What are quantitative text analysis and natural language processing?
- What is the structure of the module, and what are the expectations?

- *Application*: installing packages and setting up a project in Positron (VSCode/RStudio)

Mandatory Readings

- K. Benoit (2020). “Text as Data: An Overview”. *Handbook of Research Methods in Political Science and International Relations*. Ed. by L. Curini and R. Franzese. Thousand Oaks: Sage: 461–497.
- J. Grimmer and B. M. Stewart (2013). “Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts”. *Political Analysis* 21 (3): 267–297.

Optional

- K. L. Nielbo, F. Karsdorp, M. Wevers, A. Lassche, B. R. B., M. Kestemont, and N. Tahmasebi (2024). “Quantitative Text Analysis”. *Nature Reviews Methods Primers* 2 (24)
- J. Wilkerson and A. Casas (2017). “Large-Scale Computerized Text Analysis in Political Science: Opportunities and Challenges”. *Annual Review of Political Science* 20: 529–544.
- M. Gentzkow, B. T. Kelly, and M. Taddy (2019). “Text as Data”. *Journal of Economic Literature* 57 (3): 535–574.

Week 2: Workflow, R and Quarto (28 January 2026)

- What are the underlying assumptions of text-as-data approaches?
- How to set up and use R?
- *Application*: overview of important R functions; structure the workflow for a quantitative research project?

Mandatory Readings

- K. Watanabe and S. Müller (2023). *Quanteda Tutorials*. URL: <https://tutorials.quanteda.io>: chapter 1.
- J. Grimmer, M. E. Roberts, and B. M. Stewart (2022). *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press: chapter 4.
- Posit Team (2025). *Tutorial: Hello, Quarto*. 2025. URL: <https://quarto.org/docs/get-started/hello/positron.html>.
- S. Couch and S. Altman (2025). *Privacy and AI Assistant*. Posit Blog. URL: <https://posit.co/blog/trust-llm-tools/>.

Optional

- R. Alexander (2023). *Telling Stories with Data: With Applications in R*. New York: CRC Press: chapter 3.
- E. Ash and S. Hansen (2023). “Text Algorithms in Economics”. *Annual Review of Economics* 15 (659-688).

- K. Benoit, K. Watanabe, H. Wang, P. Nulty, A. Obeng, S. Müller, and A. Matsuo (2018). “[quanteda: An R Package for the Quantitative Analysis of Textual Data](#)”. *The Journal of Open Source Software* 3 (30): 774.

Week 3: Tokenisation and Document-Feature Matrix (4 February 2026)

- What are tokens, types, and features?
- What is the difference between stemming and lemmatisation?
- What information can we extract from a document-feature matrix?
- *Application*: tokenising texts, removing features, and creating a document-feature matrix

Mandatory Readings

- J. Grimmer, M. E. Roberts, and B. M. Stewart (2022). *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press: chapter 5.
- K. Watanabe and S. Müller (2023). *Quanteda Tutorials*. URL: <https://tutorials.quanteda.io>: chapter 2–3.

Optional

- M. W. Denny and A. Spirling (2018). “Text Preprocessing For Unsupervised Learning: Why It Matters, When It Misleads, And What To Do About It”. *Political Analysis* 26 (2): 168–189.
- K. Welbers, W. Van Atteveldt, and K. Benoit (2017). “Text Analysis in R”. *Communication Methods and Measures* 11 (4): 245–265.

Week 4: Dictionaries and Sentiment: Old and New Approaches (11 February 2026)

- Classical dictionary approaches: how to create, validate, refine, and apply automated dictionaries for sentiment and content analysis
- The limitations of dictionary methods and when they remain most useful for rigorous research
- Modern alternatives: can large language models replace dictionaries, and what are the trade-offs between interpretability and performance?
- *Application*: creating multi-word expressions and applying dictionaries to tokens objects and document-feature matrices

Mandatory Readings

- S.-O. Proksch, W. Lowe, J. Wäckerle, and S. N. Soroka (2019). “Multilingual Sentiment Analysis: A New Approach to Measuring Conflict in Legislative Speeches”. *Legislative Studies Quarterly* 44 (1): 97–131.

- M. Foramitti, U. M. Nater, C. Lamm, and M. Martins (2025). “Societal Crises Disrupt Long-Term Increases in Stress, Negativity, and Simplicity in US Billboard Song Lyrics from 1973 to 2023”. *Scientific Reports* 15: 41733.
- S. Rathje, D.-M. Mirea, I. Sucholutsky, R. Marjeh, C. E. Robertson, and J. J. Van Bavel (2024). “GPT is an Effective Tool for Multilingual Psychological Text Analysis”. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 121 (34): e2308950121.

Optional

- A. Muddiman, S. C. McGregor, and N. J. Stroud (2019). “(Re)Claiming Our Expertise: Parsing Large Text Corpora With Manually Validated and Organic Dictionaries”. *Political Communication* 36 (2): 214–226.
- S. Müller (2020). “Media Coverage of Campaign Promises Throughout the Electoral Cycle”. *Political Communication* 37 (5): 696–718.
- D. Jurafsky and J. H. Martin (2025). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models*. 3rd edition: ch. 22.

Applications in the News

- The Economist (2025a). “Scraping a Leader’s Mind”. 457 (9472): 55.

Week 5: Describing and Comparing Texts (18 February 2026)

- How do texts differ in their ‘readability’ and complexity? What are measures to estimate the similarity and distance between texts?
- How can we identify distinct features in texts?
- What are word embeddings and how can we use them for social science research?
- *Application*: estimating readability, similarity, and “keyness”; train Word2Vec model and identify similarities between word vectors

Mandatory Readings

- K. Benoit, K. Munger, and A. Spirling (2019). “Measuring and Explaining Political Sophistication Through Textual Complexity”. *American Journal of Political Science* 63 (2): 491–508.
- P. L. Rodriguez and A. Spirling (2022). “Word Embeddings: What Works, What Doesn’t, and How to Tell the Difference for Applied Research”. *The Journal of Politics* 84 (1): 101–115.

Optional

- E. Hengel (2022). “Publishing While Female: Are Women Held to Higher Standards? Evidence from Peer Review”. *The Economic Journal* 132 (648): 2951–2991.
- H. C. Shulman, D. M. Markowitz, and T. Rogers (2024). “Reading Dies in Complexity: Online News Consumers Prefer Simple Writing”. *Science Advances* 10 (23): eadn2555.

- D. Bischof and R. Senninger (2018). “Simple Politics for the People? Complexity in Campaign Messages and Political Knowledge”. *European Journal of Political Research* 57 (2): 473–495.
- E. M. Wirsching, P. L. Rodriguez, A. Spirling, and B. M. Stewart (2025). “Multilanguage Word Embeddings for Social Scientists: Estimation, Inference, and Validation Resources for 157 Languages”. *Political Analysis* 33 (2): 156–163.
- M. R. Holman, R. Johnson, and T. Simko (2025). “Measuring Conflict in Local Politics”. *Urban Affairs Review* published ahead of print (doi: 10.1177/10780874251355893).
- J. Blumenau (2021). “The Effects of Female Leadership on Women’s Voice in Political Debate”. *British Journal of Political Science* 51 (2): 750–771.

Applications in the News

- The Economist (2025b). “The Perils of Book-Spurning”. 456 (9464): 75–76.

Week 6: Supervised, Unsupervised, and Semi-Supervised Scaling (25 February 2026)

- What are the assumptions, advantages, and problems of supervised and unsupervised scaling?
- How can we use supervised scaling to answer substantive questions?
- *Application*: Wordscores, Wordfish, and Latent Semantic Scaling

Mandatory Readings

- M. Laver, J. Garry, and K. Benoit (2003). “Extracting Policy Positions from Political Texts Using Words as Data”. *American Political Science Review* 97 (2): 311–331.
- J. B. Slapin and S.-O. Proksch (2008). “A Scaling Model for Estimating Time-Series Party Positions from Texts”. *American Journal of Political Science* 52 (3): 705–722.
- K. Watanabe (2021). “Latent Semantic Scaling: A Semisupervised Text Analysis Technique for New Domains and Languages”. *Communication Methods and Measures* 14 (2): 81–102.

Optional

- P. Parschan and C. Jakob (2025). “Computational Measurement of Political Positions: A Review of Text-based Ideal Point Estimation Algorithms”. *Quality & Quantity* published ahead of print (doi: 10.1007/s11135-025-02500-4).
- G. Le Mens and A. Gallego (2025). “Positioning Political Texts with Large Language Models by Asking and Averaging”. *Political Analysis* 33 (3): 274–282.
- T. O’Grady (2019). “Careerists Versus Coal-Miners: Welfare Reforms and the Substantive Representation of Social Groups in the British Labour Party”. *Comparative Political Studies* 52 (4): 544–578.
- S. Müller, S. Brazys, and A. Dukalskis (2024). “Discourse Wars and ‘Mask Diplomacy’: China’s Global Image Management in Times of Crisis”. *Political Research Exchange* 6 (1): 2337632.

Week 7: Topic Models (4 March 2026)

- How does unsupervised document classification work? What are the assumptions, advantages, and caveats of topic models?
- *Application*: Compare unseeded and keyword-assisted topic models

Mandatory Readings

- T. Gessler (2022). “[Topic Models](#)”. *Elgar Encyclopedia of Technology and Politics*. Ed. by A. Ceron. Cheltenham: Edward Elgar Publishing: 108–111.
- S. Eshima, K. Imai, and T. Sasaki (2024). “[Keyword-Assisted Topic Models](#)”. *American Journal of Political Science* 68 (2): 730–750.

Optional

- M. E. Roberts, B. M. Stewart, D. Tingley, C. Lucas, J. Leder-Luis, S. K. Gadarian, B. Albertson, and D. G. Rand (2014). “[Structural Topic Models for Open-Ended Survey Responses](#)”. *American Journal of Political Science* 58 (4): 1064–1082.
- R. Parthasarathy, V. Rao, and N. Palaniswamy (2019). “[Deliberative Democracy in an Unequal World: A Text-As-Data Study of South India’s Village Assemblies](#)”. *American Political Science Review* 113 (3): 623–640.
- S. Müller, G. Kennedy, and T. Maher (2023). “[Reactions to Experts in Deliberative Democracy: The 2016–2018 Irish Citizens’ Assembly](#)”. *Irish Political Studies* 38 (4): 467–488.
- A. Catalinac (2016). “[From Pork to Policy: The Rise of Programmatic Campaigning in Japanese Elections](#)”. *The Journal of Politics* 78 (1): 1–18.

Week 8: Retrieving, Loading, and Wrangling Text Corpora (25 March 2026)

- What are typical text corpora you can use for your final research paper?
- What are APIs and how can we use them to retrieve data?
- How can we load various types of text corpora and transform them into a quantified corpus object?
- What are legal and ethical requirements and challenges when working with social media data?
- *Application*: Manifesto Corpus, UN General Debate Corpus, Guardian API, Twitter API

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: ch. 5–7.
- P. C. Bauer and C. Landesvatter, eds. (2024). *APIs for Social Scientists: A Collaborative Review*: skim potentially relevant chapters.

Optional

- N. Merz, S. Regel, and J. Lewandowski (2016). “The Manifesto Corpus: A New Resource for Research on Political Parties and Quantitative Text Analysis”. *Research & Politics* 3 (2): 1–8.
- S. Barari and T. Simko (2023). “LocalView, a Database of Public Meetings for the Study of Local Politics and Policy-Making in the United States”. *Scientific Data* 10: 135.
- M. J. Salganik (2017). *Bit by Bit: Social Research in the Digital Age*. Princeton: Princeton University Press: ch. 6.
- A. Fourney (2024). *MarkItDown: Python Tool for Converting Files and Office Documents to Markdown*. Version 0.0.1a3. URL: <https://pypi.org/project/markitdown/>.

News Article

- J. Burn-Murdoch and S. O’Connor (2026). *The AI Shift: Agentic AI is Coming for Quantitative Research*. The Economist. URL: <https://www.ft.com/content/91836acb-ab0c-4680-9dcf-c4de9e9917c5>.

Week 9: Fundamentals of Supervised Document Classification (1 April)

- How do we create training and test sets for document classification? How do we evaluate classification performance using metrics such as precision, recall, and F1 scores?
- Understanding bag-of-words classifiers and when they perform well
- *Application*: supervised machine learning using `quanteda` with human-coded training data

Mandatory Readings

- S. Müller (2022). “The Temporal Focus of Campaign Communication”. *The Journal of Politics* 84 (1): 585–590.
- S. Lacy, B. R. Watson, D. Riffe, and J. Lovejoy (2015). “Issues and Best Practices in Content Analysis”. *Journalism & Mass Communication Quarterly* 92 (4): 791–811.
- F. Gilardi, M. Alizadeh, and M. Kubli (2023). “ChatGPT Outperforms Crowd-Workers for Text-Annotation Tasks”. *Proceedings of the National Academy of Sciences of the United States of America* 120 (3): e2305016120.

Optional

- D. Jurafsky and J. H. Martin (2025). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models*. 3rd edition: Appendix B (Naïve Bayes).
- C. Hanretty (2023). *Replicating Mueller, ‘The Temporal Focus of Campaign Communication’*. URL: <https://www.chrishanretty.co.uk/posts/finetuning/>.
- S. Kapoor and A. Narayanan (2023). “Leakage and the Reproducibility Crisis in Machine-Learning-Based Science”. *Patterns* 4: 100804.

- A. Peterson and A. Spirling (2018). “Classification Accuracy as a Substantive Quantity of Interest: Measuring Polarization in Westminster Systems”. *Political Analysis* 26 (1): 120–128.
- B. Castanho Silva, D. Pullan, and J. Wäckerle (2025). “Blending In or Standing Out? Gendered Political Communication in 24 Democracies”. *American Journal of Political Science* 69 (2): 653–668.
- K. Benoit, D. Conway, B. E. Lauderdale, M. Laver, and S. Mikhaylov (2016). “Crowd-Sourced Text Analysis: Reproducible and Agile Production of Political Data”. *American Political Science Review* 110 (2): 278–295.
- L. Birkenmaier, C. M. Lechner, and C. Wagner (2024). “The Search for Solid Ground in Text as Data: A Systematic Review of Validation Practices and Practical Recommendations for Validation”. *Communication Methods and Measures* 18 (3): 249–277.

Week 10: Transformers: Fundamentals (8 April 2026)

- Basic knowledge of the inner workings of transformers
- *Application*: Exploring the `transformers` Python library and HuggingFace infrastructure; using existing models for classification
- Explore Google Colab for running Python code in the cloud without local installation

Mandatory Readings

- A. Kroon, K. Welbers, D. Trilling, and W. van Atteveldt (2024). “Advancing Automated Content Analysis for a New Era of Media Effects Research: The Key Role of Transfer Learning”. *Communication Methods and Measures* 18 (2): 142–162.
- J. C. Timoneda and S. Vallejo Vera (2025). “BERT, RoBERTa, or DeBERTa? Comparing Performance Across Transformers Models in Political Science Text”. *The Journal of Politics* 87 (1).

Optional

- D. Jurafsky and J. H. Martin (2025). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models*. 3rd edition: ch. 8 (Transformers).
- S. Wankmüller (2024). “Introduction to Neural Transfer Learning with Transformers for Social Science Text Analysis”. *Sociological Methods & Research* 53 (4): 1676–1752.
- L. Tunstall, L. von Werra, and T. Wolf (2022). *Natural Language Processing with Transformers: Building Language Applications with Hugging Face*. Beijing: O’Reilly.
- M. Laurer, W. van Atteveldt, A. Casas, and K. Welbers (2025). “On Measurement Validity and Language Models: Increasing Validity and Decreasing Bias with Instructions”. *Communication Methods and Measures* 19 (1): 46–62.
- G. Okasa, A. de León, M. Strinzel, A. Jorstad, K. Milzow, M. Egger, and S. Müller (2025). “A Supervised Machine Learning Approach for Assessing Grant Peer Review Reports”. *Quantitative Science Studies* 6: 1189–1214.

News Article

- Financial Times Visual Storytelling Team (2023). *Generative AI Exists Because of the Transformer*. URL: <https://ig.ft.com/generative-ai/>.

Week 11: Fine-Tuning Transformer Models for Classification (15 April 2026)

- Fine-tune a pre-trained transformer model for a domain-specific classification task
- *Application*: Using the `transformers` Python library to fine-tune transformer for domain-specific task

Mandatory Readings

- HuggingFace (2025). *Transformers: State-of-the-art Machine Learning for PyTorch, TensorFlow, and JAX*. V4.53.3. URL: <https://huggingface.co/docs/transformers/>: skim tutorials for a basic understanding of the transformers library.
- M. J. J. Bucher and M. Martini (2024). *Fine-Tuned ‘Small’ LLMs (Still) Significantly Outperform Zero-Shot Generative AI Models in Text Classification*. arXiv PrePrint. URL: <https://arxiv.org/abs/2406.08660>.

Optional

- S. Müller and S.-O. Proksch (2024). “Nostalgia in European Party Politics: A Text-Based Measurement Approach”. *British Journal of Political Science* 54 (3): 993–1005.
- S. Müller and N. Fujimura (2025). “Campaign Communication and Legislative Leadership”. *Political Science Research and Methods* 13 (3): 545–566.
- B. Warner, A. Chaffin, B. Clavié, O. Weller, O. Hallström, S. Taghadouini, A. Gallagher, R. Biswas, F. Ladhak, T. Aarsen, N. Cooper, G. Adams, J. Howard, and I. Poli (2024). *Smarter, Better, Faster, Longer: A Modern Bidirectional Encoder for Fast, Memory Efficient, and Long Context Finetuning and Inference*. arXiv PrePrint. URL: <https://arxiv.org/abs/2412.13663>.
- Z. P. Dickson and S. B. Hobolt (2025). “Going Against the Grain: Climate Change as a Wedge Issue for the Radical Right”. *Comparative Political Studies* 58 (8): 1733–1759.

Week 12: Large Language Models: Capabilities and Implications (22 April 2026)

- What can LLMs do well for text analysis? Zero-shot and few-shot classification and summarisation
- Understanding the trade-offs: when LLMs excel and when classical approaches are preferable
- Ensuring robust and reproducible research with LLMs: open science and replicability
- *Application*: Running LLMs, such as GPT and Llama, through APIs and locally; comparing LLM outputs with classical approaches

Mandatory Readings

- K. Benoit, S. De Marchi, C. Laver, M. Laver, and J. Ma (Forthcoming). “Using Large Language Models to Analyze Political Texts Through Natural Language Understanding”. *American Journal of Political Science*.
- C. Barrie, A. Palmer, and A. Spirling (2025). *Replication for Language Models: Problems, Principles, and Best Practices for Political Science*. URL: https://arthurspirling.org/documents/BarriePalmerSpirling_TrustMeBro.pdf.
- S. F. Maerz and K. Benoit (2026). *quallmer: Qualitative Text Analysis with Large Language Models*. R Package Version 0.3.0 (read tutorials).
- J. T. Ornstein, E. N. Blasingame, and J. S. Truscott (2025). “How to Train Your Stochastic Parrot: Large Language Models for Political Texts”. *Political Science Research and Methods* 13 (2): 264–281.

Optional

- A. Spirling (2023). “Open Generative AI Models are a Way Forward for Science”. *Nature* 616: 413.
- J. Cova and L. Schmitz (2024). *A Primer for the Use of Classifier and Generative Large Language Models in Social Science Research*. OSF PrePrint. URL: <https://doi.org/10.31219/osf.io/r3qng>.
- T. Hu, Y. Kyrychenko, S. Rathje, N. Collier, S. van der Linden, and J. Roozenbeek (2025). “Generative Language Models Exhibit Social Identity Biases”. *Nature Computational Science* 5: 65–75.
- D. Jurafsky and J. H. Martin (2025). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models*. 3rd edition: ch. 7 (Large Language Models).
- I. Haaland, C. Roth, S. Stantcheva, and J. Wohlfart (2025). “Understanding Economic Behavior Using Open-ended Survey Data”. *Journal of Economic Literature* 63 (4): 1244–1280.
- S. J. Westwood (2025). “The Potential Existential Threat of Large Language Models to Online Survey Research”. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 122 (47): e2518075122.

News Articles

- K. Cunningham (2025). *Who are the People who Spoiled their Ballot and Why? An Initial Analysis*. URL: <https://kevcunningham.substack.com/p/who-are-the-people-who-spoiled-their>.
- J. Burn-Murdoch and S. O’Connor (2025). *The AI Shift: Is AI About to Break Polling?*. Financial Times, 27 November 2025. URL: <https://www.ft.com/content/1298a2cd-5623-480c-b30e-ff81fc5c788d>.

References

- Alexander, R. (2023). *Telling Stories with Data: With Applications in R*. New York: CRC Press.
- APSA Committee on Publications (2018). *Style Manual for Political Science (Revised 2018 Version)*. URL: <https://connect.apsanet.org/stylemanual/>.
- Ash, E. and S. Hansen (2023). “Text Algorithms in Economics”. *Annual Review of Economics* 15 (659-688).

- Barari, S. and T. Simko (2023). “LocalView, a Database of Public Meetings for the Study of Local Politics and Policy-Making in the United States”. *Scientific Data* 10: 135.
- Barrie, C., A. Palmer, and A. Spirling (2025). *Replication for Language Models: Problems, Principles, and Best Practices for Political Science*. URL: https://arthurspirling.org/documents/BarriePalmerSpirling_TrustMeBro.pdf.
- Bauer, P. C. and C. Landesvatter, eds. (2024). *APIs for Social Scientists: A Collaborative Review*.
- Benoit, K. (2020). “Text as Data: An Overview”. *Handbook of Research Methods in Political Science and International Relations*. Ed. by L. Curini and R. Franzese. Thousand Oaks: Sage: 461–497.
- Benoit, K., D. Conway, B. E. Lauderdale, M. Laver, and S. Mikhaylov (2016). “Crowd-Sourced Text Analysis: Reproducible and Agile Production of Political Data”. *American Political Science Review* 110 (2): 278–295.
- Benoit, K., S. De Marchi, C. Laver, M. Laver, and J. Ma (Forthcoming). “Using Large Language Models to Analyze Political Texts Through Natural Language Understanding”. *American Journal of Political Science*.
- Benoit, K., K. Munger, and A. Spirling (2019). “Measuring and Explaining Political Sophistication Through Textual Complexity”. *American Journal of Political Science* 63 (2): 491–508.
- Benoit, K., K. Watanabe, H. Wang, P. Nulty, A. Obeng, S. Müller, and A. Matsuo (2018). “quanteda: An R Package for the Quantitative Analysis of Textual Data”. *The Journal of Open Source Software* 3 (30): 774.
- Birkenmaier, L., C. M. Lechner, and C. Wagner (2024). “The Search for Solid Ground in Text as Data: A Systematic Review of Validation Practices and Practical Recommendations for Validation”. *Communication Methods and Measures* 18 (3): 249–277.
- Bischof, D. and R. Senninger (2018). “Simple Politics for the People? Complexity in Campaign Messages and Political Knowledge”. *European Journal of Political Research* 57 (2): 473–495.
- Blumenau, J. (2021). “The Effects of Female Leadership on Women’s Voice in Political Debate”. *British Journal of Political Science* 51 (2): 750–771.
- Bucher, M. J. J. and M. Martini (2024). *Fine-Tuned ‘Small’ LLMs (Still) Significantly Outperform Zero-Shot Generative AI Models in Text Classification*. arXiv PrePrint. URL: <https://arxiv.org/abs/2406.08660>.
- Burn-Murdoch, J. and S. O’Connor (2025). *The AI Shift: Is AI About to Break Polling?*. Financial Times, 27 November 2025. URL: <https://www.ft.com/content/1298a2cd-5623-480c-b30e-ff81fc5c788d>.
- Burn-Murdoch, J. and S. O’Connor (2026). *The AI Shift: Agentic AI is Coming for Quantitative Research*. The Economist. URL: <https://www.ft.com/content/91836acb-ab0c-4680-9dcf-c4de9e9917c5>.
- Castanho Silva, B., D. Pullan, and J. Wäckerle (2025). “Blending In or Standing Out? Gendered Political Communication in 24 Democracies”. *American Journal of Political Science* 69 (2): 653–668.
- Catalinac, A. (2016). “From Pork to Policy: The Rise of Programmatic Campaigning in Japanese Elections”. *The Journal of Politics* 78 (1): 1–18.
- Couch, S. and S. Altman (2025). *Privacy and AI Assistant*. Posit Blog. URL: <https://posit.co/blog/trust-llm-tools/>.
- Cova, J. and L. Schmitz (2024). *A Primer for the Use of Classifier and Generative Large Language Models in Social Science Research*. OSF PrePrint. URL: <https://doi.org/10.31219/osf.io/r3qng>.
- Cunningham, K. (2025). *Who are the People who Spoiled their Ballot and Why? An Initial Analysis*. URL: <https://kevcunningham.substack.com/p/who-are-the-people-who-spoiled-their>.
- Cunningham, S. (2021). *Causal Inference: The Mixtape*. New Haven: Yale University Press.
- Denny, M. W. and A. Spirling (2018). “Text Preprocessing For Unsupervised Learning: Why It Matters, When It Misleads, And What To Do About It”. *Political Analysis* 26 (2): 168–189.
- Dickson, Z. P. and S. B. Hobolt (2025). “Going Against the Grain: Climate Change as a Wedge Issue for the Radical Right”. *Comparative Political Studies* 58 (8): 1733–1759.
- Eshima, S., K. Imai, and T. Sasaki (2024). “Keyword-Assisted Topic Models”. *American Journal of Political Science* 68 (2): 730–750.
- Financial Times Visual Storytelling Team (2023). *Generative AI Exists Because of the Transformer*. URL: <https://ig.ft.com/generative-ai/>.
- Foramitti, M., U. M. Nater, C. Lamm, and M. Martins (2025). “Societal Crises Disrupt Long-Term Increases in Stress, Negativity, and Simplicity in US Billboard Song Lyrics from 1973 to 2023”. *Scientific Reports* 15: 41733.
- Fourney, A. (2024). *MarkItDown: Python Tool for Converting Files and Office Documents to Markdown*. Version 0.0.1a3. URL: <https://pypi.org/project/markitdown/>.

- Gentzkow, M., B. T. Kelly, and M. Taddy (2019). “Text as Data”. *Journal of Economic Literature* 57 (3): 535–574.
- Gessler, T. (2022). “Topic Models”. *Elgar Encyclopedia of Technology and Politics*. Ed. by A. Ceron. Cheltenham: Edward Elgar Publishing: 108–111.
- Gilardi, F., M. Alizadeh, and M. Kubli (2023). “ChatGPT Outperforms Crowd-Workers for Text-Annotation Tasks”. *Proceedings of the National Academy of Sciences of the United States of America* 120 (3): e2305016120.
- Grimmer, J., M. E. Roberts, and B. M. Stewart (2022). *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press.
- Grimmer, J. and B. M. Stewart (2013). “Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts”. *Political Analysis* 21 (3): 267–297.
- Haaland, I., C. Roth, S. Stantcheva, and J. Wohlfart (2025). “Understanding Economic Behavior Using Open-ended Survey Data”. *Journal of Economic Literature* 63 (4): 1244–1280.
- Hanretty, C. (2023). *Replicating Mueller, ‘The Temporal Focus of Campaign Communication’*. URL: <https://www.chrishanretty.co.uk/posts/finetuning/>.
- Healy, K. (2019). *Data Visualization: A Practical Introduction*. Princeton: Princeton University Press.
- Hengel, E. (2022). “Publishing While Female: Are Women Held to Higher Standards? Evidence from Peer Review”. *The Economic Journal* 132 (648): 2951–2991.
- Holman, M. R., R. Johnson, and T. Simko (2025). “Measuring Conflict in Local Politics”. *Urban Affairs Review* published ahead of print (doi: 10.1177/10780874251355893).
- Hu, T., Y. Kyrychenko, S. Rathje, N. Collier, S. van der Linden, and J. Roozenbeek (2025). “Generative Language Models Exhibit Social Identity Biases”. *Nature Computational Science* 5: 65–75.
- HuggingFace (2025). *Transformers: State-of-the-art Machine Learning for PyTorch, TensorFlow, and JAX*. V4.53.3. URL: <https://huggingface.co/docs/transformers/>.
- Huntington-Klein, N. (2025). *The Effect: An Introduction to Research Design and Causality*. 2nd edition. Boca Raton: CRC Press.
- Hvitfeldt, E. and J. Silge (2021). *Supervised Machine Learning For Text Analysis in R*. Boca Raton: CRC Press.
- Jurafsky, D. and J. H. Martin (2025). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models*. 3rd edition.
- Kapoor, S. and A. Narayanan (2023). “Leakage and the Reproducibility Crisis in Machine-Learning-Based Science”. *Patterns* 4: 100804.
- Kroon, A., K. Welbers, D. Trilling, and W. van Atteveldt (2024). “Advancing Automated Content Analysis for a New Era of Media Effects Research: The Key Role of Transfer Learning”. *Communication Methods and Measures* 18 (2): 142–162.
- Lacy, S., B. R. Watson, D. Riffe, and J. Lovejoy (2015). “Issues and Best Practices in Content Analysis”. *Journalism & Mass Communication Quarterly* 92 (4): 791–811.
- Laurer, M., W. van Atteveldt, A. Casas, and K. Welbers (2025). “On Measurement Validity and Language Models: Increasing Validity and Decreasing Bias with Instructions”. *Communication Methods and Measures* 19 (1): 46–62.
- Laver, M., J. Garry, and K. Benoit (2003). “Extracting Policy Positions from Political Texts Using Words as Data”. *American Political Science Review* 97 (2): 311–331.
- Le Mens, G. and A. Gallego (2025). “Positioning Political Texts with Large Language Models by Asking and Averaging”. *Political Analysis* 33 (3): 274–282.
- Maerz, S. F. and K. Benoit (2026). *quallmer: Qualitative Text Analysis with Large Language Models*. R Package Version 0.3.0.
- Merz, N., S. Regel, and J. Lewandowski (2016). “The Manifesto Corpus: A New Resource for Research on Political Parties and Quantitative Text Analysis”. *Research & Politics* 3 (2): 1–8.
- Muddiman, A., S. C. McGregor, and N. J. Stroud (2019). “(Re)Claiming Our Expertise: Parsing Large Text Corpora With Manually Validated and Organic Dictionaries”. *Political Communication* 36 (2): 214–226.
- Müller, S. (2020). “Media Coverage of Campaign Promises Throughout the Electoral Cycle”. *Political Communication* 37 (5): 696–718.
- Müller, S. (2022). “The Temporal Focus of Campaign Communication”. *The Journal of Politics* 84 (1): 585–590.

- Müller, S. (2023). “How Slack Facilitates Communication and Collaboration in Seminars and Project-Based Courses”. *Journal of Educational Technology Systems* 51 (3): 303–316.
- Müller, S., S. Brazys, and A. Dukalskis (2024). “Discourse Wars and ‘Mask Diplomacy’: China’s Global Image Management in Times of Crisis”. *Political Research Exchange* 6 (1): 2337632.
- Müller, S. and N. Fujimura (2025). “Campaign Communication and Legislative Leadership”. *Political Science Research and Methods* 13 (3): 545–566.
- Müller, S., G. Kennedy, and T. Maher (2023). “Reactions to Experts in Deliberative Democracy: The 2016–2018 Irish Citizens’ Assembly”. *Irish Political Studies* 38 (4): 467–488.
- Müller, S. and S.-O. Proksch (2024). “Nostalgia in European Party Politics: A Text-Based Measurement Approach”. *British Journal of Political Science* 54 (3): 993–1005.
- Nielbo, K. L., F. Karsdorp, M. Wevers, A. Lassche, B. R. B., M. Kestemont, and N. Tahmasebi (2024). “Quantitative Text Analysis”. *Nature Reviews Methods Primers* 2 (24).
- O’Grady, T. (2019). “Careerists Versus Coal-Miners: Welfare Reforms and the Substantive Representation of Social Groups in the British Labour Party”. *Comparative Political Studies* 52 (4): 544–578.
- Okasa, G., A. de León, M. Strinzel, A. Jorstad, K. Milzow, M. Egger, and S. Müller (2025). “A Supervised Machine Learning Approach for Assessing Grant Peer Review Reports”. *Quantitative Science Studies* 6: 1189–1214.
- Ornstein, J. T., E. N. Blasingame, and J. S. Truscott (2025). “How to Train Your Stochastic Parrot: Large Language Models for Political Texts”. *Political Science Research and Methods* 13 (2): 264–281.
- Parschan, P. and C. Jakob (2025). “Computational Measurement of Political Positions: A Review of Text-based Ideal Point Estimation Algorithms”. *Quality & Quantity* published ahead of print (doi: 10.1007/s11135-025-02500-4).
- Parthasarathy, R., V. Rao, and N. Palaniswamy (2019). “Deliberative Democracy in an Unequal World: A Text-As-Data Study of South India’s Village Assemblies”. *American Political Science Review* 113 (3): 623–640.
- Peterson, A. and A. Spirling (2018). “Classification Accuracy as a Substantive Quantity of Interest: Measuring Polarization in Westminster Systems”. *Political Analysis* 26 (1): 120–128.
- Posit Team (2025). *Tutorial: Hello, Quarto*. 2025. URL: <https://quarto.org/docs/get-started/hello/positron.html>.
- Proksch, S.-O., W. Lowe, J. Wäckerle, and S. N. Soroka (2019). “Multilingual Sentiment Analysis: A New Approach to Measuring Conflict in Legislative Speeches”. *Legislative Studies Quarterly* 44 (1): 97–131.
- Rathje, S., D.-M. Mirea, I. Sucholutsky, R. Marjeh, C. E. Robertson, and J. J. Van Bavel (2024). “GPT is an Effective Tool for Multilingual Psychological Text Analysis”. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 121 (34): e2308950121.
- Roberts, M. E., B. M. Stewart, D. Tingley, C. Lucas, J. Leder-Luis, S. K. Gadarian, B. Albertson, and D. G. Rand (2014). “Structural Topic Models for Open-Ended Survey Responses”. *American Journal of Political Science* 58 (4): 1064–1082.
- Rodriguez, P. L. and A. Spirling (2022). “Word Embeddings: What Works, What Doesn’t, and How to Tell the Difference for Applied Research”. *The Journal of Politics* 84 (1): 101–115.
- Salganik, M. J. (2017). *Bit by Bit: Social Research in the Digital Age*. Princeton: Princeton University Press.
- Shulman, H. C., D. M. Markowitz, and T. Rogers (2024). “Reading Dies in Complexity: Online News Consumers Prefer Simple Writing”. *Science Advances* 10 (23): eadn2555.
- Slapin, J. B. and S.-O. Proksch (2008). “A Scaling Model for Estimating Time-Series Party Positions from Texts”. *American Journal of Political Science* 52 (3): 705–722.
- Spirling, A. (2023). “Open Generative AI Models are a Way Forward for Science”. *Nature* 616: 413.
- Stoltz, D. S. and M. A. Taylor (2024). *Mapping Texts: Computational Text Analysis for the Social Sciences*. Oxford: Oxford University Press.
- The Economist (2025a). “Scraping a Leader’s Mind”. 457 (9472): 55.
- The Economist (2025b). “The Perils of Book-Spurning”. 456 (9464): 75–76.
- Timoneda, J. C. and S. Vallejo Vera (2025). “BERT, RoBERTa, or DeBERTa? Comparing Performance Across Transformers Models in Political Science Text”. *The Journal of Politics* 87 (1).
- Tunstall, L., L. von Werra, and T. Wolf (2022). *Natural Language Processing with Transformers: Building Language Applications with Hugging Face*. Beijing: O’Reilly.
- Turrell, A. (2024). *Coding for Economists*. URL: <https://aeturrell.github.io/coding-for-economists>.

- Wankmüller, S. (2024). “[Introduction to Neural Transfer Learning with Transformers for Social Science Text Analysis](#)”. *Sociological Methods & Research* 53 (4): 1676–1752.
- Warner, B., A. Chaffin, B. Clavié, O. Weller, O. Hallström, S. Taghadouini, A. Gallagher, R. Biswas, F. Ladhak, T. Aarsen, N. Cooper, G. Adams, J. Howard, and I. Poli (2024). [Smarter, Better, Faster, Longer: A Modern Bidirectional Encoder for Fast, Memory Efficient, and Long Context Finetuning and Inference](#). arXiv PrePrint. URL: <https://arxiv.org/abs/2412.13663>.
- Watanabe, K. (2021). “[Latent Semantic Scaling: A Semisupervised Text Analysis Technique for New Domains and Languages](#)”. *Communication Methods and Measures* 14 (2): 81–102.
- Watanabe, K. and S. Müller (2023). *Quanteda Tutorials*. URL: <https://tutorials.quanteda.io>.
- Welbers, K., W. Van Atteveldt, and K. Benoit (2017). “[Text Analysis in R](#)”. *Communication Methods and Measures* 11 (4): 245–265.
- Westwood, S. J. (2025). “[The Potential Existential Threat of Large Language Models to Online Survey Research](#)”. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 122 (47): e2518075122.
- 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.
- Wilkerson, J. and A. Casas (2017). “[Large-Scale Computerized Text Analysis in Political Science: Opportunities and Challenges](#)”. *Annual Review of Political Science* 20: 529–544.
- Wirsching, E. M., P. L. Rodriguez, A. Spirling, and B. M. Stewart (2025). “[Multilanguage Word Embeddings for Social Scientists: Estimation, Inference, and Validation Resources for 157 Languages](#)”. *Political Analysis* 33 (2): 156–163.