**INFO 290**

# Cultural Analytics

Fall 2025 · Kent Chang

## Frequently Asked Questions

(Last updated: May 22, 2025)

Thank you for your interest in this class! Below are answers to common questions. The most up-to-date version of the syllabus will be posted at:

`https://ca.kentkc.org`

**Q:** Who is this class for?

**A:** This class is designed for students interested in *connecting* interpretive approaches with computational tools, especially across the humanities and computer science, which can be roughly divided into two groups:

1. Students from the humanities who want to develop stronger intuitions for the algorithms in machine learning (ML) and natural language processing (NLP), and want to explore how they might apply them to their own work.
2. Students with some prior experience in ML and NLP who are interested in analyzing cultural and literary artifacts as real-world *data*—especially questions relevant to the humanities.

**Q:** Who can enroll in this class?

**A:** Graduate students can enroll directly through CalCentral. Undergraduate students must request a permission code (email Kent). The course is officially offered by the School of Information, but you do not need to be a student in the School to join.

**Q:** Can I use this class for the course requirement of my program?

**A:** This is a new class with a temporary course number (290's), so usually someone (me or you) has to petition or otherwise make a case for it to count toward the requirement. At the moment, this course has been **approved** as a **technical elective** for the Designated Emphasis (DE) in New Media. Reach out if you are interested and we'll see what we can do.

**Q:** What does it take to do well in this class?

**A:** The final project constitutes the most significant portion of your grade (see the grade components below). Given that cultural analytics encompasses a range of approaches, your project can lean more theoretical (or qualitative) or more computational, as long as it focuses on measuring some aspect of culture. The project is designed to allow you to draw upon your particular strengths—whether in coding, close reading, or both—as you apply the methods and frameworks explored in this course.

**Q:** Do I have to know _____ to take this class?

**A:** See the prerequisites below—but the course is designed to be interdisciplinary and largely self-contained, with built-in flexibility. If most students have strong close reading skills, for example, we may condense relevant lab sessions (below). If most are experienced with deep learning, we'll adjust discussion accordingly.

**Q:** My question isn't answered here?

**A:** Feel free to reach out to Kent at <`kentkchang@berkeley.edu`>.

**INFO 290**

# Cultural Analytics

## Machine Learning and Measurements of Culture, in Three Acts

Fall 2025 · 3 units · Kent Chang · *faculty advisor*: Prof. David Bamman ver. 2.135 (May 22, 2025)

---

While often defined as "the computational study of culture", cultural analytics might be best understood as a radical interdisciplinary experiment, one that seeks to understand cultures—socialties, histories, cognition, the literary—through empirical models and patterns, built on effective computational representations of relevant cultural constructs.[1] This experiment calls for a unique skillset: one needs to be familiar with approaches in the interpretive humanities and computer/information science; one also needs to cultivate an interdisciplinary mindset: recognize and appreciate the *affordances* and limitations of both qualitative and quantitative traditions. This class is imagined as a possible point of departure for those who are so inclined.

As such, this course is all about making *connections*: bridging the interpretive traditions of literary and cultural studies—which guide our critical engagement with texts and cultural artifacts—with computational methods, ranging from featurized classifiers to large language models. It pursues two complementary ends: students will develop interpretive strategies and critical vocabularies for cultural analysis, and, through hands-on practice grounded in engagement and experience with the text, they will learn to represent cultural data—whether text, image, audio, or video—and train machine learning models as algorithmic measuring devices to systematically characterize cultural phenomena of interest.

This class welcomes a range of inclinations: maybe you know how to implement an RNN or Transformer from scratch but are curious whether those models can be used to study literature, culture, or something beyond positive or negative sentiments. Or, maybe you've experimented with off-the-shelf topic models or word embeddings to explore humanities questions and want to see how far you can take them.

## LEARNING OBJECTIVES

- Students will have a solid grasp of theoretical and technical approaches to representation for cultural analytics.
- Students will further their understanding of deep learning methods and their affordance for cultural analytics.
- Students will be in a better position to navigate between disciplinary practices and traditions as they pertain to the computational study of culture.

## COURSE REQUIREMENTS

### PREREQUISITES

This course is self-contained and there will be no hard prerequisites. However, this class engages rigorously with both philosophical writings and technical papers, so previous exposure to concepts such as logistic regression and deep neural networks (from any of the following: DATA C200; INFO 251, 256, 159/259; CS 182/282; etc.) and theoretical writers (like Foucault, Deleuze) would be useful. During the first three weeks, resources will be provided to help you assess your preparation and fit for the course.

1. Course website: https://ca.kentkc.org.

## COURSE LAYOUT

This course has a tripartite structure:

- **Act I**: theoretical and technical foundation
- **Act II**: connecting multimodal and multi-agent systems along with relevant critical theories
- **Act III**: three case studies in cultural analysis at scale, focusing on ideology, interaction, and context, particularly through sociological and sociolinguistic frameworks

We will learn by combining conceptual discussion with hands-on exploration:

- **Lecture**: Each week will open with a lecture on the week's topic. You are encouraged to do the readings listed for each lecture prior to attendance.
- **Lab**: For each week, there's an interactive portion of this class that we are going to call *labs*. We close-read a text ("annotation"), look at some code ("workshop"), and talk about some papers ("seminar") together. Participation is assessed by reflection entries (below).

## EVALUATION CRITERIA

Grade components:

| | |
|---|---|
| **Lab participation and presentation** | 20 pts |
| three reflection entries (15 pts, 5 pts each) | |
| one presentation (5 pts) | |
| **max(Quiz 1, Quiz 2)** | 20 pts |
| **Homework** | 20 pts |
| 1. Close reading paper (10 pts) | |
| 2. Workshop notebook (10 pts) | |
| **Project** | 40 pts |
| Proposal (5 pts) | |
| Mid-term report (10 pts) | |
| Presentation (5 pts) | |
| Final report (20 pts) | |

The course is not curved, and there's no *a priori* distribution. At the end of the semester, students will receive a letter grade based on the following conversion chart:

| LETTER | POINTS | LETTER | POINTS |
|---|---|---|---|
| A+ | 97–100 | C+ | 77–79 |
| A | 93–96 | C | 73–76 |
| A– | 90–92 | C– | 70–72 |
| B+ | 87–89 | D+ | 67–69 |
| B | 83–86 | D | 63–66 |
| B– | 80–82 | D– | 60–62 |
| S/P | 80–100 | F | 0–59 |

Additional comments on grading and expectations:

- **Lab participation**: Participation is graded on the completion of three reflection entries, each a **1-page write-up** to reflect on what you have learned in the relevant lab. Throughout the semester, our labs will engage with three different kinds of material: literary or cultural texts (Act I), code and algorithms (Act II), and technical papers (Act III), and you are required to submit only three reflection entries, each based on a lab session of your choosing. Additionally, you will do a presentation during the lab (which you're not going to reflect on) at least once; more details will be given in due course.

- **Close reading paper**: Weeks 2–5 introduce some approaches to cultural artifacts, and you will apply them to read one of your choosing closely. Ideally, you will identify a book, a TV series, or a movie at the beginning of the semester, and pick one passage or scene (suitable for a critical approach covered in class over those four weeks) to close read. On a higher level, you want to think thoroughly about what makes a passage or a scene interesting and articulate your judgment using the appropriate critical vocabulary. This homework should prepare you for the Workshop Notebook and Project (more below).

- **Workshop notebook:** Weeks 6–9 are focused on the implementation of models that work with data in multiple modalities. Your job is to study the code covered in the workshop and adapt them for *your data* and a toy task of your choosing, ideally building on what you close-read for Homework 1. The goal is to gain some practical experience of operationalizing a concept that's interesting at least to you using the methods we study.
- **Project:** Ideally you will build on Homework 1 and 2 to conduct a project that would be suitable for publication at your *dream* venue. It could take the form of a long analytical/argumentative essay (usually less than 30 pages) or a technical, workshop-style paper (usually less than 8 p.), depending on what you're more comfortable with. The main requirement here is a.) let your favorite cultural texts inspire and guide you, b.) carry out something empirical, adequately evaluated, and c.) experiment with different methods of representation. We will gain exposure to a variety of venues of interest to cultural analytics over the course of the semester.
- **Quiz:** Everyone has to take two quizzes in class during week 3 and 6. They are designed to ensure you have adequate background knowledge and literacy to read widely for your own project, and therefore happen relatively early in the semester (compared to typical midterm and final exams). Both quizzes will take a similar format (a mix of true or false, multiple choice, matching, short answer), focusing on the intuitions of key mathematical or algorithmic concepts and critical thought.

## POLICIES

### ACADEMIC INTEGRITY

All students will follow the UC Berkeley code of conduct. Attribute all ideas and code if they are not your original work. Every submitted lab or essay for this class must be your own work. While discussion with others is encouraged, you must write all your essays or solutions on your own, and acknowledge or credit resources you have consulted. When in doubt, ask me. For more information, see https://conduct.berkeley.edu/integrity/.

### AI ASSISTANTS & LARGE LANGUAGE MODELS

We will largely follow "NeurIPS 2025 Policy on the Use of Large Language Models":[2] you can (and probably *should*) use AI assistants like ChatGPT to boost your productivity—for brainstorming, editing, or handling routine tasks (e.g., formatting tables, which has little to do with the substance of your work). But you must fully understand everything you submit; if you use an AI tool to help with writing or coding, it's your responsibility to verify its output, correct mistakes, and, indeed, learn from the process. Again, all submitted work must be your own: if you can't explain it, you shouldn't turn it in.

### ACCESSIBILITY & ACCOMMODATION

It is important to me that this course be accessible to you. If you have been issued a letter of accommodation from the Disabled Students Program (DSP), please see me to work out the necessary arrangements. If you need an accommodation and have not yet seen a Disability Specialist at the DSP, please do so when you can.

### LATE POLICY

All assignments will be due 5 pm on the due date, with a 2-hour grace period. All students will have 3 slip days, each extending the deadline by 24 hours. Otherwise, a 10% penalty applies for each late day. Please feel free to talk to me if you anticipate difficulties in meeting any of the deadlines for any reason, and we can work together to figure out a solution.

2. https://neurips.cc/Conferences/2025/LLM.

## SCHEDULE

- The list is subject to change.
- Announcements and readings can be found on bCourses.
- *Murphy* = Kevin Murphy's *Probabilistic Machine Learning*, book I (2022) and II (2023)

### PROLOGUE ■ WEEK 1. Representing, Operationalizing, . . .

*Lecture.* In the introductory lecture, we focus on the fundamental ideas underlying the computational study of culture: representation and operationalization. How do we translate what might be of interest into something computationally viable?

**Keywords.** operationalization, featurized models
- Quist, "Laurelled Lives" (2017)
- So, *Redlining Culture* (2020): chap. 3 (Recognition: Literary Distinction and Blackness)
- Bengio, Ducharme, and Vincent, "A Neural Probabilistic Language Model" (2000)

(Optional: math self-assessment; cf. Kun, *A Programmer's Introduction to Mathematics*)

### ACT 1 ■ Termination and Indetermination

#### WEEK 2. Death of Theory

*Lecture.* To unpack the concept of a *learnable representation*, we look at *learning* and *representation* from two perspectives: neural and literary-critical.

**Keywords.** formalism and New Criticism; word embeddings
- Eliot, "Tradition and Individual Talent—I" (1919)
- *Murphy* II. 32.1–4 (Representation learning)
- Mikolov et al., "Distributed Representations of Words and Phrases and their Compositionality" (2013)

**Extension.** Radford et al., "Language Models Are Unsupervised Multitask Learners" (2019); A. Liu et al., "SuperBPE: Space Travel for Language Models" (2025)

*Annotation.* Benj Pasek and Justin Paul, *Dear Evan Hansen* (2017)

(Optional: Python self-assessment; cf. Walsh, *Introduction to Cultural Analytics & Python*)

#### WEEK 3. Death of Subjectivity

*Lecture.* Continuing our discussion on learning and representation, we examine the idea of aesthetic autonomy as it intersects with machine-learning techniques. Given an abundance of data, what and how might we estimate with probabilistic models?

**Keywords.** formalism and narrative structure; generative models
- Wimsatt and Beardsley, "The Intentional Fallacy" (1946), up to §II; Piper, So, and Bamman, "Narrative Theory for Computational Narrative Understanding" (2021)
- *Murphy* II. 20.1–20.4.1 (Generative models: an overview)
- Vafa, Naidu, and Blei, "Text-Based Ideal Points" (2020)

**Extension.** Seaver, *Computing Taste* (2022): chap. 1 (Introduction); Gopalan, Hofman, and Blei, "Scalable Recommendation with Hierarchical Poisson Factorization" (2015)

**QUIZ 1** (1 hr, closed-book, in class)

### WEEK 4. Death of the Author

*Lecture.* This week, we consider signifying structures through French structuralism, focusing on Saussurean synchrony and diachrony via Roman Jakobson. From there, we explore how sequence models formalize dependencies across ordered data—whether the order encodes time, syntax, or positional dependencies.

**Keywords.** French structuralism; sequence models
- Jakobson, "Closing Statement: Linguistics and Poetics" (1960)
- Weatherby and Justie, "Indexical AI" (2022)
- *Murphy* I. 15.1–3 (Neural sequence modeling)

**Extension.** *Murphy* II. 29.1–4 (State-space models, esp. HMMs); Geoghegan, "From Information Theory to French Theory: Jakobson, Lévi-Strauss, and the Cybernetic Apparatus" (2011)

*Annotation.* David Angell, Peter Casey, and David Lee, *Frasier*: I.i (The Good Son), I.xxiv (My Coffee with Niles), VI.xvii (The Dinner Party); cf. Sedgwick and Frank, "Shame in the Cybernetic Fold" (1995)

### WEEK 5. Death of the Man

*Lecture.* Continuing our discussion on structuralism, we now study deconstruction at work and reflect on the *network* concept/metaphor, which can be anti-humanist and anthropological at the same time.

**Keywords.** deconstruction; transformer models
- Barthes, *S / Z* (1974), up to §XV
- Hayles, "Information or Noise? Economy of Explanation in Barthes's *S/Z* and Shannon's Information Theory" (1987)
- Tay, Luu, and Hui, "Compare, Compress and Propagate" (2018)

**Extension.** Derrida, "Structure, Sign, and Play in the Discourse of the Human Sciences" (1979); Li et al., "Optimus: Organizing Sentences via Pre-trained Modeling of a Latent Space" (2020)

*Annotation.* Oscar Wilde, *The Importance of Being Earnest* (1895); cf. Craft, "Alias Bunbury" (1990) and Sedgwick, *Tendencies* (1994): "Tales of the Avunculate"

**HW 1** (Close reading, 5–10 pp. double-spaced) due on Friday at 5 pm. Only `pdf` files are accepted

## ACT 2 ■ Being and Vectorizing

### WEEK 6. Imagetext

*Lecture.* Equipped with modular, probabilistic, and deconstructionist thinking, we now start from scratch and revisit foundational ideas underlying interpretation and representation. What constitutes a cultural artifact, and how can we represent it computationally?

**Keywords.** limits of critique; diffusion models
- Mitchell, *Picture Theory* (1995): chap. 1 (The pictorial turn)
- Vilnis and McCallum, "Word Representations via Gaussian Embedding" (2015)
- Ruhe et al., "Rolling Diffusion Models" (2024)

**Extension.** Abrams, "The Deconstructive Angel" (1977); Taylor, *The Archive and the Repertoire* (2003): "Acts of Transfer"; Crawford and Paglen, "Excavating AI" (2021)

**QUIZ 2** (cumulative; 1 hr, closed-book, in class)

**PROJECT PROPOSAL** (maximum 1 p.) due on Friday at 5 pm

### WEEK 7. Time-Image

*Lecture*. This week we consider videos as time-images and introduce the challenging task of long-form video understanding. How do we model and reason over a sequence of images with long-range dependency?

**Keywords.** time-image; long-form video understanding

- Deleuze, *Cinema 2* (1989): chap. 2 (Recapitulation of images and signs)
- Korbar, Huh, and Zisserman, "Look, Listen and Recognise" (2024)
- Sun et al., "video-SALMONN: Speech-enhanced Audio-Visual Large Language Models" (2024)

*Workshop*. Embeddings for prediction and clustering

- Zhang et al., *Dive Into Deep Learning*: skim through chaps. 13–16
- *Murphy* I. 21.3 ($K$-means clustering); 21.4 (Clustering using mixture models)
- *Murphy* II. 3.10 (Hypothesis testing)

### WEEK 8. Assemblage

*Lecture*. To further explore theory post structuralism, we look at the concept of *assemblage*, which we juxtapose with modular systems that are distinct from typical end-to-end networks.

**Keywords.** assemblage, poststructuralism; mixture-of-experts, multi-agent systems

- Deleuze and Guattari, *A Thousand Plateaus* (1987): Introduction; chap. 6 (How Do You Make Yourself a Body Without Organs?)
- Tsing, *The Mushroom at the End of the World* (2021): Introduction; chap. 4 (Agential Realism)
- Fedus, Zoph, and Shazeer, "Switch Transformers" (2022)

*Workshop*. Inference and sampling

- Wortsman et al., "Model Soups" (2022)
- Bertsch et al., "It's MBR All the Way Down" (2023)
- Qineng Wang et al., "Rethinking the Bounds of LLM Reasoning" (2024)

### WEEK 9. Relationality

*Lecture*. To conclude our foray into entangling encounters, we turn to *relationality*: how do relations constitute entities, rather than merely linking them? How might computational models formalize social, political, and cultural relations beyond language itself and embed complex structures of connection into measurable *form*?

**Keywords.** relational forms, intersubjectivity; graphs, connotation frames

- Simondon, *Individuation in Light of Notions of Form and Information* (1964 [2021]): "Form and Matter"
- Ferreira Da Silva, *Unpayable Debt* (2019): chap. 2 (the most perfect hallucination)
- Bordes et al., "Translating Embeddings for Modeling Multi-relational Data" (2013)

*Workshop*. Wrap up previous workshops; plus: writing and presentation

- Sims, Park, and Bamman, "Literary Event Detection" (2019)
- Bode, "Why You Can't Model Away Bias" (2020)
- Shechtman, "Command of Media's Metaphors" (2021)

## ACT 3 ■ Everywhere and Nowhere

### WEEK 10. Form and Scale

*Lecture*. What's the implication of *scale* in the context of both cultural analysis and language models? Does scale lead to *knowledge*?

**Keywords.** literary sociology; scaling laws

- English and Underwood, "Shifting Scales" (2016)
- Klein, "Dimensions of Scale" (2020)
- Brown et al., "Language Models are Few-Shot Learners" (2020); skim: J. Kaplan et al., "Scaling Laws for Neural Language Models" (2020)

*Seminar* (TBD—check bCourses for updates).

- Lester, Al-Rfou, and Constant, "The Power of Scale for Parameter-Efficient Prompt Tuning" (2021)
- Levy, Jacoby, and Goldberg, "Same Task, More Tokens: the Impact of Input Length on the Reasoning Performance of Large Language Models" (2024)
- Muennighoff et al., "s1: Simple Test-Time Scaling" (2025)

**HW 2** (Workshop notebook) due on Friday at 5 pm. Only `ipynb` files are accepted

### WEEK 11. Form and Ideology

*Lecture*. How might language models be used to make judgments about society and culture? How does scale shape our understanding of culture and society?

**Keywords.** ideology, affect; domain adaptation

- Jameson, *The Political Unconscious* (1994): chap. 1 (On Interpretation)
- Clough, *The User Unconscious* (2018): "The Datalogical Turn"
- Ziems et al., "Can Large Language Models Transform Computational Social Science?" (2024)

*Seminar* (TBD—check bCourses for updates).

- Gururangan et al., "Don't Stop Pretraining: Adapt Language Models to Domains and Tasks" (2020)
- Arora and Goyal, "A Theory for Emergence of Complex Skills in Language Models" (2023)
- Qunbo Wang et al., "Soft Knowledge Prompt" (2024)

**PROJECT MID-TERM REPORT** (min. 4 pp./paragraph outline) due on Friday at 5 pm

### WEEK 12. Form and Interaction

*Lecture*. LLMs and cultures emerge out of interaction. But for any given interaction, what's involved, and how do we model it?

**Keywords.** recognition, sociality; preference alignment

- Hegel, *Phenomenology of Spirit* (1807 [2018]): section A. "Self-Sufficiency and Non-Self-Sufficiency of Self-Consciousness; Mastery and Servitude"
- Goffman, *Forms of Talk* (1981): chap. 2 (Replies and Responses)
- Stiennon et al., "Learning to Summarize with Human Feedback" (2020)

*Seminar* (TBD—check bCourses for updates).

- Rafailov et al., "Direct Preference Optimization" (2023)
- Dingemanse and Enfield, "Interactive Repair and the Foundations of Language" (2024)
- Z. Liu et al., "A Dynamic LLM-Powered Agent Network for Task-Oriented Agent Collaboration" (2024)

### WEEK 13. Form and Context

*Lecture.* If all interactions are contextual, how might language models make use of such information beyond the document level? This week we think about knowledge (situational and social) in terms of its representation and retrieval for inference.

**Keywords.** situated knowledge; long-context models, knowledge retrieval

- Silverstein, *Language in Culture* (2022): Introduction; chap. 8 (Knowledge)
- Fish, "Is There a Text in This Class?" (1995)
- Jurgens et al., "Your Spouse Needs Professional Help: Determining the Contextual Appropriateness of Messages through Modeling Social Relationships" (2023)

*Seminar* (TBD—check bCourses for updates).

- Chevalier et al., "Adapting Language Models to Compress Contexts" (2023)
- Khanuja et al., "An Image Speaks a Thousand Words, but Can Everyone Listen? On Image Transcreation for Cultural Relevance" (2024)
- Edge et al., "From Local to Global: A Graph RAG Approach to Query-Focused Summarization" (2024)

### (WEEK 14. No class; Thanksgiving)

### EPILOGUE ■ WEEK 15. What Is Cultural Analytics?

To wrap up the class, we reflect on what it means to study culture vis-a-vis computation.

*Conversation starters*—**Keywords.** modernity, validity.

- Foucault, "What Is Enlightenment" (1983)
- Mbembe, "The Universal Right to Breathe" (2021)

**PROJECT PRESENTATION** (in class)

### (WEEK 16. RRR)

**PROJECT FINAL REPORT** due

## INDICATIVE LIST OF READINGS


Quist, J (2017). "Laurelled Lives". In: *New Left review*.

So, Richard Jean (2020). *Redlining Culture*. New York: Columbia University Press.

Bengio, Yoshua, Réjean Ducharme, and Pascal Vincent (2000). "A Neural Probabilistic Language Model". In: *Advances in neural information processing systems* 13. URL: https://proceedings.neurips.cc/paper_files/paper/2000/hash/728f206c2a01bf572b5940d7d9a8fa4c-Abstract.html.

Eliot, T S (Sept. 1919). "Tradition and Individual Talent—I". In: *Egoist* 6.4, pp. 54–55.

Mikolov, Tomas et al. (2013). "Distributed Representations of Words and Phrases and their Compositionality". In: *Advances in Neural Information Processing Systems*. Ed. by C J Burges et al. Vol. 26. Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf.

Radford, Alec et al. (2019). "Language Models Are Unsupervised Multitask Learners". In: *OpenAI* 1.8, p. 9. URL: https://life-extension.github.io/2020/05/27/GPT%E6%8A%80%E6%9C%AF%E5%88%9D%E6%8E%A2/language-models.pdf.

Liu, Alisa et al. (Mar. 17, 2025). "SuperBPE: Space Travel for Language Models". In: *arXiv [cs.CL]*.

Wimsatt, W K and M C Beardsley (1946). "The Intentional Fallacy". In: *The Sewanee Review* 54.3, pp. 468–488. URL: http://www.jstor.org/stable/27537676.

Piper, Andrew, Richard Jean So, and David Bamman (Nov. 2021). "Narrative Theory for Computational Narrative Understanding". In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, pp. 298–311. URL: https://aclanthology.org/2021.emnlp-main.26.

Vafa, Keyon, Suresh Naidu, and David Blei (July 2020). "Text-Based Ideal Points". In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Ed. by Dan Jurafsky et al. Online: Association for Computational Linguistics, pp. 5345–5357. URL: https://aclanthology.org/2020.acl-main.475.

Seaver, Nick (Dec. 2022). *Computing Taste*. en. 1st ed. Chicago, IL: University of Chicago Press.

Gopalan, Prem, Jake M Hofman, and David Blei (2015). "Scalable Recommendation with Hierarchical Poisson Factorization". In: *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*. UAI'15. Amsterdam, Netherlands: AUAI Press, pp. 326–335.

Jakobson, Roman (1960). "Closing Statement: Linguistics and Poetics". In: *Style in Language*. Ed. by Thomas Albert Sebeok. MIT Press Paperback Series. Cambridge, Mass: MIT Press, pp. 350–377.

Weatherby, Leif and Brian Justie (Jan. 2022). "Indexical AI". en. In: *Critical inquiry* 48.2, pp. 381–415. URL: https://www.journals.uchicago.edu/doi/10.1086/717312.

Geoghegan, Bernard Dionysius (Sept. 2011). "From Information Theory to French Theory: Jakobson, Lévi-Strauss, and the Cybernetic Apparatus". In: *Critical inquiry* 38.1, pp. 96–126. URL: http://dx.doi.org/10.1086/661645.

Sedgwick, Eve Kosofsky and Adam Frank (1995). "Shame in the Cybernetic Fold". In: *Critical Inquiry* 21.2, pp. 496–522.

Barthes, Roland (1974). *S / Z*. United Kingdom: Blackwell.

Hayles, N Katherine (1987). "Information or Noise? Economy of Explanation in Barthes's *S/Z* and Shannon's Information Theory". In: *One Culture: Essays in Science and Literature*. Ed. by George Lewis Levine and Alan Rauch. Science and Literature. Madison, Wis: University of Wisconsin Press, pp. 119–142.

Tay, Yi, Anh Tuan Luu, and Siu Cheung Hui (2018). "Compare, Compress and Propagate". In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Ed. by Ellen Riloff et al. Brussels, Belgium: Association for Computational Linguistics, pp. 1565–1575. URL: https://aclanthology.org/D18-1185.

Derrida, Jacques (1979). "Structure, Sign, and Play in the Discourse of the Human Sciences". In: *The Structuralist Controversy: The Languages of Criticism and the Sciences of Man*. Ed. by Richard Macksey and Eugenio Donato. Fourth printing. Baltimore: Johns Hopkins Press.

Li, Chunyuan et al. (Nov. 2020). "Optimus: Organizing Sentences via Pre-trained Modeling of a Latent Space". In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Ed. by Bonnie Webber et al. Online: Association for Computational Linguistics, pp. 4678–4699. URL: https://aclanthology.org/2020.emnlp-main.378.

Craft, Christopher (July 1990). "Alias Bunbury". In: *Representations* 31, pp. 19–46. URL: https://online.ucpress.edu/representations/article-pdf/doi/10.2307/2928398/238708/2928398.pdf.

Sedgwick, Eve Kosofsky (1994). *Tendencies*. London: Routledge.

Mitchell, W J T (Sept. 1995). *Picture Theory*. en. University of Chicago Press.

Vilnis, Luke and Andrew McCallum (2015). "Word Representations via Gaussian Embedding". In: *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Ruhe, David et al. (2024). "Rolling Diffusion Models". In: *Proceedings of the 41st International Conference on Machine Learning*. Ed. by Ruslan Salakhutdinov et al.

Vol. 235. Proceedings of Machine Learning Research. PMLR, pp. 42818–42835.
Abrams, M H (1977). "The Deconstructive Angel". In: *Critical inquiry* 3.3, pp. 425–438.
Taylor, Diana (Sept. 2003). *The Archive and the Repertoire*. a John Hope Franklin Center Book. Durham, NC: Duke University Press.
Crawford, Kate and Trevor Paglen (June 2021). "Excavating AI". en. In: *AI & society*.
Deleuze, Gilles (1989). *Cinema 2*. University of Minnesota Press.
Korbar, Bruno, Jaesung Huh, and Andrew Zisserman (Apr. 2024). "Look, Listen and Recognise". In: *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 2975–2979.
Sun, Guangzhi et al. (2024). "video-SALMONN: Speech-enhanced Audio-Visual Large Language Models". In: *Proceedings of the 41st International Conference on Machine Learning*. ICML'24. Vienna, Austria: JMLR.org.
Deleuze, Gilles and Félix Guattari (1987). *A Thousand Plateaus*. Minneapolis: University of Minnesota Press.
Tsing, Anna Lowenhaupt (June 2021). *The Mushroom at the End of the World*. en. Princeton, NJ: Princeton University Press.
Fedus, William, Barret Zoph, and Noam Shazeer (2022). "Switch Transformers". In: *Journal of Machine Learning Research* 23.120, pp. 1–39.
Wortsman, Mitchell et al. (2022). "Model Soups". In: *International conference on machine learning*. PMLR, pp. 23965–23998.
Bertsch, Amanda et al. (Dec. 2023). "It's MBR All the Way Down". In: *Proceedings of the Big Picture Workshop*. Ed. by Yanai Elazar et al. Singapore: Association for Computational Linguistics, pp. 108–122. DOI: 10.18653/v1/2023.bigpicture-1.9. URL: https://aclanthology.org/2023.bigpicture-1.9.
Wang, Qineng et al. (Aug. 2024). "Rethinking the Bounds of LLM Reasoning". In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Ed. by Lun-Wei Ku, Andre Martins, and Vivek Srikumar. Bangkok, Thailand: Association for Computational Linguistics, pp. 6106–6131. DOI: 10.18653/v1/2024.acl-long.331. URL: https://aclanthology.org/2024.acl-long.331.
Simondon, Gilbert (Jan. 1964 [2021]). *Individuation in Light of Notions of Form and Information*. en. Posthumanities. Minneapolis, MN: University of Minnesota Press.
Ferreira Da Silva, Denise (Dec. 2019). *Unpayable Debt*. en. Sternberg Press / The Antipolitical. Germany: Sternberg Press.
Bordes, Antoine et al. (2013). "Translating Embeddings for Modeling Multi-relational Data". In: *Advances in Neural Information Processing Systems*. Ed. by C J Burges et al. Vol. 26. Curran Associates, Inc.
Sims, Matthew, Jong Ho Park, and David Bamman (2019). "Literary Event Detection". In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, pp. 3623–3634. URL: http://dx.doi.org/10.18653/v1/P19-1353.
Bode, Katherine (Mar. 2020). "Why You Can't Model Away Bias". en. In: *Modern Language Quarterly* 81.1, pp. 95–124. URL: https://read.dukeupress.edu/modern-language-quarterly/article/81/1/95/152221/Why-You-Cant-Model-Away-Bias.
Shechtman, Anna (June 2021). "Command of Media's Metaphors". en. In: *Critical inquiry* 47.4, pp. 644–674. URL: https://www.journals.uchicago.edu/doi/abs/10.1086/714512.
English, James F and Ted Underwood (Sept. 2016). "Shifting Scales". In: *Modern Language Quarterly* 77.3, pp. 277–295. URL: http://dx.doi.org/10.1215/00267929-3570612.
Klein, Lauren F (Jan. 2020). "Dimensions of Scale". In: *PMLA* 135.1, pp. 23–39. URL: http://dx.doi.org/10.1632/pmla.2020.135.1.23.
Brown, Tom et al. (2020). "Language Models are Few-Shot Learners". In: *Advances in neural information processing systems* 33, pp. 1877–1901. URL: https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html.
Kaplan, Jared et al. (Jan. 22, 2020). "Scaling Laws for Neural Language Models". In: *arXiv [cs.LG]*.
Lester, Brian, Rami Al-Rfou, and Noah Constant (Nov. 2021). "The Power of Scale for Parameter-Efficient Prompt Tuning". In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Ed. by Marie-Francine Moens et al. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, pp. 3045–3059. URL: https://aclanthology.org/2021.emnlp-main.243.
Levy, Mosh, Alon Jacoby, and Yoav Goldberg (Aug. 2024). "Same Task, More Tokens: the Impact of Input Length on the Reasoning Performance of Large Language Models". In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Ed. by Lun-Wei Ku, Andre Martins, and Vivek Srikumar. Bangkok, Thailand: Association for Computational Linguistics, pp. 15339–15353. DOI: 10.18653/v1/2024.acl-long.818. URL: https://aclanthology.org/2024.acl-long.818.
Muennighoff, Niklas et al. (Jan. 2025). "s1: Simple Test-Time Scaling". In: *arXiv [cs.CL]*.
Jameson, Fredric (1994). *The Political Unconscious*. 1. publ., 7. pr. Ithaca: Cornell University Press.
Clough, Patricia Ticineto (Mar. 2018). *The User Unconscious*. Minneapolis, MN: University of Minnesota Press.

Ziems, Caleb et al. (Mar. 2024). "Can Large Language Models Transform Computational Social Science?" In: *Computational Linguistics* 50.1, pp. 237–291. DOI: 10.1162/coli_a_00502. URL: https://aclanthology.org/2024.cl-1.8/.

Gururangan, Suchin et al. (July 2020). "Don't Stop Pretraining: Adapt Language Models to Domains and Tasks". In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Ed. by Dan Jurafsky et al. Online: Association for Computational Linguistics, pp. 8342–8360.

Arora, Sanjeev and Anirudh Goyal (July 2023). "A Theory for Emergence of Complex Skills in Language Models". In: *arXiv [cs.LG]*.

Wang, Qunbo et al. (Aug. 2024). "Soft Knowledge Prompt". In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Ed. by Lun-Wei Ku, Andre Martins, and Vivek Srikumar. Bangkok, Thailand: Association for Computational Linguistics, pp. 6132–6143. DOI: 10.18653/v1/2024.acl-long.332. URL: https://aclanthology.org/2024.acl-long.332.

Hegel, Georg Wilhelm Fredrich (Feb. 1807 [2018]). *Phenomenology of Spirit*. en. Ed. by Terry Pinkard and Michael Baur. Cambridge Hegel translations. Cambridge, England: Cambridge University Press.

Goffman, Erving (1981). *Forms of Talk*. en. University of Pennsylvania Press.

Stiennon, Nisan et al. (2020). "Learning to Summarize with Human Feedback". In: *Advances in Neural Information Processing Systems*. Ed. by H Larochelle et al. Vol. 33. Curran Associates, Inc., pp. 3008–3021.

Rafailov, Rafael et al. (2023). "Direct Preference Optimization". In: *Thirty-seventh Conference on Neural Information Processing Systems*. URL: https://openreview.net/forum?id=HPuSIXJaa9.

Dingemanse, Mark and N J Enfield (Jan. 2024). "Interactive Repair and the Foundations of Language". en. In: *Trends in cognitive sciences* 28.1, pp. 30–42. URL: http://dx.doi.org/10.1016/j.tics.2023.09.003.

Liu, Zijun et al. (2024). "A Dynamic LLM-Powered Agent Network for Task-Oriented Agent Collaboration". In: *First Conference on Language Modeling*. URL: https://openreview.net/forum?id=XII0Wp1XA9.

Silverstein, Michael (Dec. 2022). *Language in Culture*. en. Cambridge University Press.

Fish, Stanley (1995). "Is There a Text in This Class?" In: *Campus wars*. Routledge, pp. 49–56.

Jurgens, David et al. (July 2023). "Your Spouse Needs Professional Help: Determining the Contextual Appropriateness of Messages through Modeling Social Relationships". In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Ed. by Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki. Toronto, Canada: Association for Computational Linguistics, pp. 10994–11013. URL: https://aclanthology.org/2023.acl-long.616.

Chevalier, Alexis et al. (Dec. 2023). "Adapting Language Models to Compress Contexts". In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Ed. by Houda Bouamor, Juan Pino, and Kalika Bali. Singapore: Association for Computational Linguistics, pp. 3829–3846. URL: https://aclanthology.org/2023.emnlp-main.232.

Khanuja, Simran et al. (Nov. 2024). "An Image Speaks a Thousand Words, but Can Everyone Listen? On Image Transcreation for Cultural Relevance". In: *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Ed. by Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen. Miami, Florida, USA: Association for Computational Linguistics, pp. 10258–10279. URL: https://aclanthology.org/2024.emnlp-main.573.

Edge, Darren et al. (2024). "From Local to Global: A Graph RAG Approach to Query-Focused Summarization". In: *arXiv*. eprint: 2404.16130 (cs.CL). URL: https://arxiv.org/abs/2404.16130.

Foucault, Michel (Nov. 1983). "What Is Enlightenment". en. In: *The Foucault Reader*. Ed. by Peter Rainbow. Pantheon Books.

Mbembe, Achille (2021). "The Universal Right to Breathe". In: *Critical Inquiry* 47.S2, S58–S62.