Writing Machines

EngCmp/CmpInf 1207

Fall 2024, University of Pittsburgh Prof. <u>Matthew Burton</u> and Prof. <u>Annette Vee</u>

Sponsored by a grant from the Humanities Center Co-Teaching Fellowship, and co-designed by Annette Vee and Matt Burton.

This course counts towards the Public and Professional Writing Major and the Digital Narrative and Interactive Design Major. It fills the Writing-intensive requirement for the Dietrich School of Arts and Sciences and the School of Computing and Information.

Course Description

Large Language Models (LLMs) like ChatGPT have dominated the news cycle in the last year, promising--or threatening--to upend everything we know about writing. Can machines really write?? It seems like the ability to use complex and sequential language, such as writing demonstrates, is something only humans can do! But, it turns out, people have pursued ways of automating writing for centuries—from mechanized automatons in the 18th century to ChatGPT in the 21st.

This class approaches the contemporary buzz about ChatGPT and Natural Language Generation (NLG) through a sociotechnical and historical lens, exploring automated writing through historical and contemporary writing machines, from 18th century automatons, Ouija boards (a Victorian writing machine), secretarial shorthand, combinatorial poetics, procedural algorithms, recurrent neural networks, and the transformer technology that drives cutting-edge LLMs like ChatGPT. The course is organized into five units focusing on different methods of automating writing: mechanical; spiritual; procedural; statistical; simulative. Each unit teaches students historical and technical contexts of its associated automation through texts and hands-on projects.

Students who take this class will learn to generate text using computational processes and critically engage with the text that they produce.

In so doing, students will engage with the political and ethical implications of large-scale data usage to train machine learning algorithms, including issues of underrepresented languages, bias, hate speech, further marginalizing communities, automating labor, environmental impacts, dirty data and privacy violations. We'll consider whose writing is being automated, and for what purposes as we trace the development and commercialization of NLG.

This course is Writing-Intensive and designed to fill the W requirement for the Dietrich School of Arts and Sciences and SCI, which means you'll do the equivalent of about 25 pages of writing. Writing in this course takes various forms: human-only composition; collaborations with text generation technologies such as LLMs through prompting, editing, and iterating; and coding or revising programs that generate text. You will construct and/or train Markov-style models, recurrent neural networks, and LLMs by choosing texts for training, iterating, and adjusting outputs from the models. Key to each of these assignments is revision of prompts, code, or training datasets as you work with text-gen tech to meet your composition goals. For each course project, you'll also compose a reflection (either alone or in collaboration with LLMs) about your process collaborating with various writing technologies. For each writing assignment, I'll ask that you disclose what portion or percentage of the work is human and which is technology-generated.

You can expect to spend 9 hours per week outside of class completing the work for this course.

Learning Objectives (developed with the help of Claude.ai)

Upon completion of this course, students will be able to:

- Generate original text using computational processes and tools alone and with refinement and collaboration with human methods of writing.
- Critically analyze, evaluate, and revise algorithmically generated text through prompting, editing, and reflection.
- Demonstrate understanding of the technical evolution of natural language generation systems and their social/historical contexts, from early writing machines to modern neural networks.
- Articulate the political, ethical, and social implications of large language models and their training data.
- Examine issues of underrepresented languages, bias, privacy, and labor in automated writing systems.
- Develop collaborative skills through group projects and shared critical reflection.

Assessment

Assessment for this course is based on *completion* of projects, posts, and participation. I will give you feedback on your work but will not grade it individually. If you turn in an assignment on time and it's not of B or above quality, then I will ask you to work more on it until it is. Once you have completed the assignment, then you've checked that box and you're done! For weekly Canvas posts, note that you will need to complete all of them to receive an A, and most of them to pass the course. Participation and attendance counts, too.

Late work is accepted with prior notice. Please let me know if you need an extension on any project.

Please see me at any time if you have any questions about your grade in the course. The grade matrix below is a guide that I will use for assigning a final grade. You can use it to self-evaluate during the term. A+ grades are rare, but I'm delighted to assign an A+ to someone who knocks it out of the park on every part of the course.

Grade matrix

	Major projects (5)	Weekly posts or comments on canvas	Participation and attendance (including being prepared for class)
A (95%)	All 5 projects completed on time and of good quality	All posts completed on time and of good quality, following prompt, engaging with peers and texts from class	Excellent attendance; always prepared for class (readings and activities completed); engaged and supportive of peers in discussion and group work.
A- (90%)	All 5 projects completed on time and of good quality	Almost all posts completed on time and of good quality, following prompt, engaging with peers and texts from class	Excellent attendance (no more than 1 class missed); prepared for class (readings and activities completed); engaged and supportive of peers in discussion and group work.
В (85%)	All 5 projects completed and of good quality	Almost all posts completed and of good quality, following prompt	Good attendance (no more than 2 classes missed); usually prepared for class (readings and activities completed); supportive of peers in discussion and group work.
B- (80%)	All 5 projects completed	Most posts completed on time and decently	Good attendance (no more than 2 classes missed); often prepared for class (readings and activities completed); supportive of peers in discussion and group work.
C (75%)	At least 4 projects completed, including the final	Most posts completed	Decent attendance (no more than 3 classes missed); most of the time prepared for class (readings and activities completed); supportive of peers in discussion and group work.
C- (70%)	At least 3 projects completed,	At least half of posts completed	Attended most classes (no more than 4 classes missed); sometimes prepared for class (readings and activities completed); supportive

	including the final		of peers in discussion and group work.
D (65%)	At least 3 projects completed	Some posts completed	Attended most classes (no more than 4 classes missed)

Course Policies

[omitted]

Course Schedule

Most weeks will follow a rough pattern of Monday texts and discussions and Wednesday hands-on activities or "lab."

Unit 0: Automating Writing

Week 1-2

This unit is a general introduction to the theme of the course: how and why we automate writing, and what it means for human intelligence, and the history of artificial intelligence. We begin with a canonical text in computing (Turing) that lays out the early conversation about language generation and AI, and several contemporary texts describing the current state and problems with Large Language Models (LLMs). Activities are focused on general understanding of statistical models of language and computing, and Shane provides a good introduction to (a circa 2018) machine learning as a technique, plus the foundation for an activity that will give you hands-on experience with this ML technique, plus data collection and augmentation.

Week 1 - Mon, 8/26/24

Survey on AI interest and use Madlibs in class
Course introduction:
Preview homework for Weds

- ⇒ Homework for Weds 8/28:
 - Read: Lee, Timothy and Sean Trott, "Large language models, explained with a minimum of math and jargon," *Understanding AI blog*, July 27, 2023. https://www.understandingai.org/p/large-language-models-explained-with
 - > Write an introduction to you by collaborating with an AI such as ChatGPT or Claude (no time limit, but aim for 15min)

Week 1 - Weds, 8/28/24

Discuss survey results, Lee and Trott

Watch Joy Buolamwini together: Buolamwini, Joy. "AI, Ain't I a Woman?"

https://www.youtube.com/watch?v=HZxV9w2o0FM [about Buolamwini and her PhD project, Gender Shades, and the ways that AI facial recognition works differently for different skin color.]

MadLibs replit activity

- ⇒ Homework for Weds 9/4:
 - > Start a Replit account and fork the MadLibs version here: https://replit.com/@mcburton1/WM-Python-Mad-Libs If you're not familiar with Python, or code in general, just take a look at the code there and see if you can tell what it's doing. If you are comfortable with the code, you can start playing around with it in preparation for the MadLibs assignment.
 - > Look at/begin reading: Turing, Alan. "Computing Machinery and Intelligence." Mind 59, no. 236 (October 1950): 433–60. This is an early discussion of the use of computers to reproduce language, and the connection to the history of AI, plus a canonical text in the history of computing. Figure out: what is he arguing here? What are some strange points, or points you don't get what he's saying? Annotate a copy of the reading, or take notes on it.
 - > Post to Canvas a quote of the most interesting sentence that you see so far. Explain why it's interesting in one sentence beginning with something other than "This is interesting because..." Then give the sentence to an AI platform of your choice and ask it to tell you why it's interesting. Include your AI sentence in your post, along with which platform/version you used.
 - > Read Lee and Trott if you didn't read it for Weds, 8/28.

Week 2 - Mon, 9/2: NO CLASS (UNIVERSITY HOLIDAY)

Week 2 Weds 9/4

Discuss Turing in class, and play with a Turing GPT.

- \Rightarrow Homework for Mon, 9/9:
 - > Work on Assignment 0: MadLibs in Python
 - > Continue reading Turing.

Unit 1: Automating Writing Mechanically

Week 3-4

This unit dives into early types of mechanical AI. In a mechanistic worldview that dominated in 18th century Europe, the creation of automata--especially automata that performed human activities such as writing--the recreation of human activities was tantamount to recreating

human intelligence. This very early AI history is also connected to the history of calculation, as Daston points out. Our activities connect mechanical calculation in the early 20th century with cryptography, and involve visiting one of our archival holdings at Pitt: a cryptographic machine Leibniz designed in the 17th century but was not built until Pitt Philosophy Professor N. Rescher worked with an engineer to do so. This machine was the most sophisticated in mechanical cryptographic design until the 20th century.

Week 3 - Mon, 9/9

Workshop on MadLibs in Replit

- \Rightarrow Homework for 9/11:
 - > Madlibs assignment

Week 3 - Weds 9/11: MEET IN HILLMAN OPEN LAB

Turn in: Assignment 0: MadLibs in Python, plus reflection on Canvas (1159pm) Visit Open Lab in Hillman, and do the class in the Open Lab. Cipher discs lab https://exhibits.library.cmu.edu/cipherdiscs/

- \Rightarrow Homework for 9/16:
 - > Read: Daston, Lorraine. "Calculation and the Division of Labor: 1750-1950." 32nd Annual Lecture of the German Historical Institute, Washington, DC, Nov 9, 2017.

Week 4 - Mon 9/16

Visit Hillman Special Collections to see the Leibniz machine Homework for 9/18:

> Cipher project and reflection

Week 4 - Weds 9/18

Turn in Cipher Project and reflection (1159pm)
Discuss Daston and connections between calculation and AI
Watch Enigma machine animation

- \Rightarrow Homework for 9/25:
 - Read McRobbie, Linda Rodriguez, "The Strange and Mysterious History of the Ouija Board," Smithsonian Magazine, Oct 27, 2013.
 https://www.smithsonianmag.com/history/the-strange-and-mysterious-history-of-the-ouija-board-5860627/

Unit 2: Automating Writing Spiritually

Week 5-6

While mechanical automata put humans in the role of God, designing machines that performed human activities, "automatic writing," was a process that lent human bodies to spirits to write

with. Through seances, trances, or instruments like the Ouija board, people summoned spirits, who then communicated with them through writing. This spiritual attempt at automating writing is echoed in present services that resurrect people who have passed through mining their texts, writings, and other text they leave behind.

**Note that discussions around the Ouija board and spiritualism might be sensitive for folks, and that this whole unit deals with death and loss, plus complicated questions about consent from those passed.

Week 5 - Mon, 9/23

Discuss Spiritualism and automatic writing and McRobbie essay on Ouija.

Automatic writing exercise (10min, close your eyes, keep writing, using the prompt "I remember..." if you get stuck)

Introduce Ghostwriting project (Assignment #2)

- ⇒ Homework for Weds, 9/25:
 - > Read Vara, Vauhini. "Ghosts" Believer. 9 August 2021. https://www.thebeliever.net/ghosts/
 - > Read Fagone, Jason. "The Jessica Simulation: Love and Loss in the Age of AI." San Francisco Chronicle. 23 July 2021. https://www.sfchronicle.com/projects/2021/jessica-simulation-artificial-intelligence/

Week 5 - Weds, 9/25 Automatic writing exercise Discuss Vara and Fagone

- \Rightarrow Homework for 9/30:
 - > Canvas post on Ghostwriting project

Week 6 - Mon, 9/30 (Prof. Vee is out of town and Prof. Burton will run class)

Listen to Ghostwriter excerpts: "Ghostwriter," from Episode 757, "The Ghost in the Machine," This American Life, December 2021 (23min)

Discuss ghostwriting posts (due at midnight) and begin work on Ghostwriting project

- \Rightarrow Homework for 10/2:
 - > First iteration of Ghostwriting project

Week 6 - Weds, 10/2

Work on Ghostwriting project and discuss first iteration Check out preserving human sites like Hereafter.ai https://www.hereafter.ai/

- ⇒ Homework:
 - Read Roberts, Siobhan. "Christopher Strachey's Nineteen Fifties Love Machine," The New Yorker, Feb 14, 2017, https://www.newyorker.com/tech/annals-of-technology/christopher-stracheys-nineteen-fifties-love-machine

o (Click here for the free version on the Wayback Machine)

Unit 3: Automating Writing Procedurally

Week 7

Procedural methods for automating writing have been used for centuries, beginning as early as the I, Ching. In the early 20th century, Oulipo poets came up with parlor games like Exquisite Corpse and N+7 to create combinatorial work. In the 1950s, Christopher Strachey (a contemporary of Alan Turing) created a program for the Manchester Computer that wrote love letters. These combinatorial methods also draw on more quotidian template and job printing. We will play with some of these combinatorial methods and create our own new poetry or texts with them.

Week 7 - Mon, 10/7

Discuss Strachey and Eliza
Play with the Strachey Love Letter Generator on JS Fiddle version:
https://jsfiddle.net/mcburton11/nz1L9dft/ (Javascript, modifiable)

- \Rightarrow Homework for 10/9:
 - > Ghostwriting project due on Weds
 - > Play with an Eliza bot online, such as this one: https://www.masswerk.at/elizabot/

Week 7 - Weds, 10/9

Ghostwriting project due with reflection Introduce Oulipo writing exercises Discussion - What is personal Style? Lab activities for the day:

Oulipo processes, working from <u>Exercises in Style</u> by Raymond Queneau (<u>text is here</u>) or Crista Lopes's <u>Exercises in Programming Style</u>

(based on) Nick Montfort's "Through the Park":

https://replit.com/@annettevee/Through-The-Park#main.py (run link here)

Nick Montfort's Taroko Gorge

Example procedural text: "Sermon on the Warpland," Lillian-Yvonne Bertram https://a-new-sermon-on-the-warpland.glitch.me/ (Gwendolyn Brooks' Sermon on the Warpland as reference)

N+7 game

- → Homework for 10/16:
 - > Read Roald Dahl, "The Great Automatic Grammatizator," 1954.

> Read Cal Newport, "What kind of Writer is ChatGPT" *The New Yorker*. Oct 3 2024.

https://web.archive.org/web/20241003145439/https://www.newyorker.com/culture/annals-of-inquiry/what-kind-of-writer-is-chatgpt

Unit 4: Automating Writing Logically

Week 8-9

Midcentury computing tried to crack the code of logical reasoning, and computational AI--as demonstrated by Newell, Shaw and Simon's Logic Theorist, a reasoning program presented at the Dartmouth Summer Session on Artificial Intelligence in 1956, the birth of AI. At this time, there was fear of "giant brains" or computers replacing the work of writers (Dahl), along with an origin of the computer as a metaphor for the mind. Margaret Masterson offered a somewhat opposing view, asking researchers to consider how AI might be a tool for new insights, rather than a more efficient, faster way of calculating.

Week 8 - Mon, 10/14 Fall break (no class)

Week 8 - Weds, 10/16 (Prof. Vee is out of town and Prof. Burton will run class)
Discuss Markov Chains, Introduce Markov Chains
Introduce Assignment 3: Text Generation with Markov Chains or RNNs

- ⇒ Homework:
 - > Find a text dataset to work with on for (think about: song lyrics, quotes from characters, stories on Project Gutenberg, poetry, etc)
 - Read the Training Teardops on my Markov Notebook
 https://colab.research.google.com/drive/1TzLLCbQDjSdnJulzWYmOINXEO59090
 md?usp=sharing
 - > If you haven't done so, read Dahl and Newport!

Week 9 - Mon, 10/21

Discuss Newport and Dahl, transition between template style and current statistical models Early attempts to automate writing Newell and Simon's Logic Theorist, and Masterson's ideas Homework

→ What were three moments in the Dahl story that struck you as related to our conversations in class? (add your answer to the Canvas discussion thread) ⇒ By 10/28, Read Shane, Janelle. <u>You Look Like a Thing and I Love You</u>. Voracious, 2019. (Chap 1-3, which is up to p 108.)

Week 9 - Weds, 10/23

Work through Text generation RNN tutorial on Colab

- ⇒ Homework (due by 10/28 classtime):
 - > By 10/28, Read Shane, Janelle. *You Look Like a Thing and I Love You*. Voracious, 2019. (Chap 1-3, which is up to p 108.)
 - > Optionally, for a more technical explanation of RNNs, you may want to read Andrej Karpathy's post on "The Unreasonable Effectiveness of RNNs," which the tutorial for 10/23 is based on.

Unit 5: Automating Writing with Simulation

Week 10-13

This week takes us into contemporary machine learning methods of automating writing, from recurrent neural networks to transformers and LLMs like ChatGPT. We'll read a few of the most circulated papers on the topic, some explainers on how LLMs work, discussions of how humans are involved in training, the dangers of LLMs, and a murder mystery written almost entirely by an LLM, but prompted and arranged by a human. You will begin work on your final project in this unit: making a writing machine.

Week 10 - Mon, 10/28

Discuss: Janelle Shane

- \Rightarrow Homework for 10/30:
 - > Read OpenAI, "Better Language Models and Their Implications." 14 Feb. 2019, https://openai.com/blog/better-language-models/. (The blog post that launched GPT-2, and really the first impressive model using transformers as a technology.)

Week 10 - Weds, 10/30

Talk about GPT-2

Text generation actively (Markov, RNN, GPT-2)

Introduce: Assignment 4: Prompting Game Design, "Connections"

- ⇒ Homework for 11/04
 - > Read OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic by Billy Perrigo for Time Magazine
 - https://time.com/6247678/openai-chatgpt-kenya-workers/

- > DELVE into How cheap, outsourced labour in Africa is shaping AI English by Alex Hern for The Guardian
 - https://www.theguardian.com/technology/2024/apr/16/techscape-ai-ga dgest-humane-ai-pin-chatgpt
- Read Baack, Stephan. "The Human Decisions that Shape Generative AI and Who is Accountable for What," Mozilla Foundation, Aug 2, 2023.
 https://foundation.mozilla.org/en/blog/the-human-decisions-that-shape-generative-ai-who-is-accountable-for-what/
- > Play "Connections" in prep for Assignment #4 (due 11/19)
 - https://www.nytimes.com/games/connections

Week 11 - Mon, 11/4

Play with Sudowrite or AI Dungeon or other creative LLM platform, discuss creative work with LLMs

Introduce Writing Machines final project

Listen to Death of an Author, by Aiden Marchine [Al driven by Stephen Marche]. Read by Edoardo Ballerini. Pushkin Industries. https://www.pushkin.fm/audiobooks/death-of-an-author (Murder mystery written 95% by Al LLMs, prompted and edited by Stephen Marche, a regular New Yorker contributor. Pay attention especially to Marche's afterward reflecting on the composition process. Note: \$5 cost to purchase this audiobook, which takes 2.5h to listen to.

Don't forget to vote on 11/5!

Week 11 - Weds, 11/6

Discuss Death of an Author in class

Watch Sunspring (2016) in class: https://www.youtube.com/watch?v=LY7x2Ihqjmc
Due Friday Nov 8: Assignment 3: Text Generation with Markov Chains or RNNs

Discuss Death of an Author and other creative work with LLMs

- \Rightarrow Homework for 11/11:
 - > Keep playing Connections https://connections.swellgarfo.com/archive (this is the archive that the tech staff of NYTimes has suggested you use while they are on strike.)

Week 12 - Mon, 11/11

Discuss Death of an Author and artistic approaches to generation

- ⇒ Homework for 11/18:
 - > Read CARE: Structure for Crafting Prompts
 - o https://www.nngroup.com/articles/careful-prompts/
 - > 11/19: Connections project due

Week 12 - Weds, 11/13

Final project discussion/workshop on Prompt Engineering

- ⇒ Homework for next week:
 - > Write a proposal for your Writing Machines final project (discuss in conference 11/8-11/20)
 - > 11/19: Connections project due
- \Rightarrow Homework for after break (12/2):
 - > Bender, Gebru, McMillan-Major, Mitchell. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. March 2021. 610–623. https://doi.org/10.1145/3442188.3445922 [open access]
 - This paper was at the center of Gebru and Mitchell's firings from Google. See this story for context: https://www.wired.com/story/google-timnit-gebru-ai-what-really-happened

Week 13 - Mon, 11/18 Conferences on final project in lieu of class.

Week 13 - Weds, 11/20

Due on 11/19, 1159pm: Assignment 4: Prompting Game Design, "Connections" Play everyone's "Connections" game and discuss prompt engineering best practices. Cautionary tales about LLMs, including the LLAMA model release, ethical dilemmas.

BREAK

Week 14 - Mon, 12/2

Read Stochastic Parrots with AI (you should have at least skimmed the article before this)

- > Bender, Gebru, McMillan-Major, Mitchell. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. March 2021. 610–623. https://doi.org/10.1145/3442188.3445922 [open access]
 - This paper was at the center of Gebru and Mitchell's firings from Google.
 See this story for context:
 https://www.wired.com/story/google-timnit-gebru-ai-what-really-happened

Week 14 - Weds, 12/4

Wrap UP
Student evaluations
final project presentations

Week 15 - Mon, 12/9 - final project presentation (last day of class)

Course Bibliography

Unit 0: Intro to Writing Machines

- Turing, Alan. "Computing Machinery and Intelligence." Mind 59, no. 236 (October 1950):
 433–60. [early discussion of the use of computers to reproduce language, and the connection to the history of Al. A canonical text in the history of computing.]
- Chiang, Ted. "ChatGPT is a blurry JPEG of the Web." The New Yorker, Feb 9, 2023.
 https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web
- Buolamwini, Joy. "AI, Ain't I a Woman?"
 https://www.youtube.com/watch?v=HZxV9w2o0FM
 [about Buolamwini and her PhD project, Gender Shades, and the ways that AI facial recognition works differently for different skin color.]
 - > Shane, Janelle. You Look Like a Thing and I Love You. Voracious, 2019. (Read at least Chap 1-3, which is up to p 108. Please read through the whole book if you can! It's a fast and hilarious read.)
 - > Lee, Timothy and Sean Trott, "Large language models, explained with a minimum of math and jargon," *Understanding AI blog*, July 27, 2023. https://www.understandingai.org/p/large-language-models-explained-with

Unit 1: Automating Writing Mechanically

- Daston, Lorraine. "Calculation and the Division of Labor: 1750-1950." 32st Annual Lecture of the German Historical Institute, Washington, DC, Nov 9, 2017.
- Hugo. Paramount Pictures, 2011, Dir by Martin Scorsese. Based on the book by Brian Selznick, The Invention of Hugo Cabret.
 - Trailer: https://www.youtube.com/watch?v=FG9rvNhdOX8
- Rescher, Nicholas. "Leibniz's Machina Deciphratoria: A Seventeenth-Century Proto-Enigma," Cryptologia Vol. 38 Iss. 2, 2014, 103-115.

Unit 2: Automating Writing Spritually

- Fagone, Jason. "The Jessica Simulation: Love and Loss in the Age of AI." San Francisco Chronicle. 23 July
 2021. https://www.sfchronicle.com/projects/2021/jessica-simulation-artificial-intelligence/
- Hendrickson, Leah. "Chatting with the dead: The hermeneutics of thanabots," Media, Culture and Society 45.5, Jan 2023. https://journals.sagepub.com/doi/10.1177/01634437221147626
- Vara, Vauhini. "Ghosts" Believer. 9 August 2021. https://believermag.com/ghosts/

McRobbie, Linda Rodriguez, "The Strange and Mysterious History of the Ouija Board,"
 Smithsonian Magazine, Oct 27, 2013.

 https://www.smithsonianmag.com/history/the-strange-and-mysterious-history-of-the-ouija-board-5860627/

Unit 3: Automating Writing Procedurally

- Brown, James, Jr. "The Machine that Therefore I Am," Philosophy and Rhetoric 47.4, 2014. (available on PittCat under the journal archives, or in our files on canvas; the article is playing on Derrida and Erasmus, on robot rhetorics and the procedures of writing)
- Roberts, Siobhan. "Christopher Strachey's Nineteen Fifties Love Machine," The New Yorker, Feb 14, 2017, https://www.newyorker.com/tech/annals-of-technology/christopher-stracheys-nineteen-fifties-love-machine
- Gitelman, Lisa. "A Short History of ______" Chapter 1 in Paper Knowledge: Toward a
 Media History of Documents. Duke University Press, 2014, p 21-52. (Paper Knowledge is
 on PittCat; this chapter is on job printing, template documents, and forms, with super
 cool pics and historical court cases) [optional: read this alongside excerpt from They Say,
 I Say, composition textbook by Graff & Birkenstein]
- Brock, Kevin. "One Hundred Thousand Billion Processes: Oulipan Computation and the Composition of Digital Cybertexts." Technoculture, vol. 2, 2012, https://tcjournal.org/vol2/brock.
- Borges, Jorge Luis. "The Garden of Forking Paths," Ficciones, 1941. online here; read more at the Pitt Borges Center
- Calvino, Italo. "Cybernetics and Ghosts." The Uses of Literature. Translated by Patrick Creagh, New York, NY: Harcourt Brace, 1982, pp. 3-27. Available here and <a href=here.

Unit 4: Automating Writing Logically

- Gigerenzer and Goldstein, "Mind as Computer: Birth of a Metaphor," *Creativity Research Journal*.
- Simon and Newell, <u>The Logic Theorist</u> [early AI, presented at the Dartmouth Summer Session on Artificial Intelligence in 1956, often considered the birth of AI, and the origin of the term "artificial intelligence."]
- Dahl, Roald. "The Great Automatic Grammatizator," 1954. [Short story of an engineer becoming a successful writer by automating language.]
- Masterson, Margaret. "The Intellect's New Eye," Times Literary Supplement, 1962.
 [Masterson, a computational linguist, proposes the computer as something like a telescope—a tool that gives us new insights—rather than the prevailing idea of it to automate boring tasks or calculation.]
 https://aclanthology.org/www.mt-archive.info/50/TLS-1962-Masterman.pdf

Unit 5: Automating Writing with Simulation

- "Better Language Models and Their Implications." OpenAI, 14 Feb.
 2019, https://openai.com/blog/better-language-models/. (The blog post that launched GPT-2, and really the first impressive model using transformers as a technology.)
- Aiden Marchine [Al driven by Stephen Marche]. Death of an Author. Read by Edoardo Ballerini. Pushkin Industries. https://www.pushkin.fm/audiobooks/death-of-an-author (Murder mystery written 95% by Al LLMs, prompted and edited by Stephen Marche, a regular New Yorker contributor. Pay attention especially to Marche's afterward reflecting on the composition process.)
- Bender, Gebru, McMillan-Major, Mitchell. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. March 2021. 610–623. https://doi.org/10.1145/3442188.3445922 [open access] This paper was at the center of Gebru and Mitchell's firings from Google. See this story for context: https://www.wired.com/story/google-timnit-gebru-ai-what-really-happened
- Vee, Annette. "Automated Trolling: The Case of GPT-4Chan." Interfaces: Essays and Reviews in Computing and Culture, Charles Babbage Institute, University of Minnesota. Dec, 2022. https://cse.umn.edu/cbi/interfaces#edgelordy
- Baack, Stephan. "The Human Decisions that Shape Generative AI and Who is Accountable for What," Mozilla Foundation, Aug 2, 2023. https://foundation.mozilla.org/en/blog/the-human-decisions-that-shape-generative-ai-who-is-accountable-for-what/

Assignments

Projects

Detailed assignment sheets will be posted on Canvas, and all assignments will be discussed in class. Due dates are always 11:59pm on the date listed. Please see me in office hours or by appointment or email me if you have questions.