

AI in the Real World

Data, Power, & Society - Day 06

Predicting the Next Word

CMSE 492 (aka CMSE 101) • Fall 2026

Prof. Danny Caballero

Updates to Calendar

Completing DTPA investigations

- This week (Week 3): Practices (Generative AI)
- Next week (Week 4): Actions (Using Generative AI in Decision Making)

Two-Week Case Studies

Starting week 5, you may choose your new groups. Moved Case Design Scaffold I to Oct 9th.

- **Week 5+6:** Healthcare, Art & Creative Industries, Education, or Athletics & Sports
- **Week 7+8:** Law & Criminal Justice, Finance & Banking, Journalism & Media, Hiring & Employment, or Military & Defense
- **Week 9+10:** Agriculture, Transportation, Retail & Consumer, Entertainment & Gaming, or Environmental & Climate Science

Full topic breakdown and turn-in schedule: [Two Week Case Analyses Explainer](#)

Predicting the Next Word

Given some text, an LLM produces a **probability distribution over what comes next**, then picks from that distribution, then does it again.

That's all it does.

The essay it produced, the code you asked for, the email you edited, and the fabricated citations are that loop running over and over.

Today we build toy models of this loop by hand, with dice.

Activity: Building a Story Word by Word

In your groups, start with the word "**Once**" and for 2 minutes rotate adding a word to build a story.

- Does it make sense?
- Are there multiple sentences?

Again, in your groups, start with the phrase "**We took the dog for a walk and...**" For 2 more minutes rotate adding a word to build a story.

- Does it make sense?
- Are there multiple sentences?

Discussion Questions

1. What differences did you notice about the story you built from a single word compared to a phrase?
2. How does this story-building exercise connect to Generative AI like ChatGPT?

Activity: How do LLMs work anyhow?

Rolling a Sentence with One Die

Designed by Vashti Sawtelle and Danny Caballero with the help of Mad Libs (TM)

You have each been given dice to roll. For the narratives below, **each of you will roll dice individually** and compare the resulting narratives (highlight or mark each roll).

Work through Narrative 1 and Narrative 2 as a group, answer the reflection questions, and we'll share out — then do the same after Narrative 3.

Narrative 1

Once upon a time, there lived a **[word 1]** in a **[word 2]** town who was known to be exceptionally **[word 3]**. One morning they found **[word 4]** box on their front doorstep. A note attached to the top of the box contained a **[word 5]** message. Knowing what they had to do they prepared for a **[word 6]** journey.

Equal Probabilities

Roll	Word 1	Word 2	Word 3	Word 4	Word 5	Word 6
1	boy	tiny	generous	heavy	cryptic	perilous
2	girl	forgotten	cunning	black	threatening	thrilling
3	queen	hidden	brave	dusty	urgent	tedious
4	prince	magical	impulsive	polished	exciting	solitary
5	witch	prosperous	stubborn	wooden	demanding	long
6	wizard	gloomy	unlucky	crystal	unexpected	unexpected

Narrative 1 – Weighted Probabilities

Once upon a time, there lived a **[word 1]** in a **[word 2]** town who was known to be exceptionally **[word 3]**. One morning they found **[word 4]** box on their front doorstep. A note attached to the top of the box contained a **[word 5]** message. Knowing what they had to do they prepared for a **[word 6]** journey.

Roll	Word 1	Word 2	Word 3	Word 4	Word 5	Word 6
1	boy	forgotten	brave	polished	cryptic	perilous
2, 3	queen	hidden	generous	dusty	threatening	thrilling
4, 5, 6	witch	gloomy	cunning	wooden	unexpected	long

Narrative 2

The professor **[word 1]** across the room and handed me an envelope. "I've reviewed your project," she said, her voice sounding surprisingly **[word 2]**. The letter inside was incredibly **[word 3]**. Leaving the classroom, I felt a deep sense of **[word 4]** about the **[word 5]** work ahead.

Equal Probabilities

Roll	Word 1	Word 2	Word 3	Word 4	Word 5
1	marched	gentle	detailed	worry	unpredictable
2	strolled	exhausted	encouraging	relief	exhausting
3	danced	hesitant	brief	disappointment	annoying
4	rushed	enthusiastic	formal	dread	easy
5	strode	excited	blunt	determination	exciting
6	plodded	stern	confusing	gratitude	challenging

Narrative 2 – Weighted Probabilities

The professor **[word 1]** across the room and handed me an envelope. "I've reviewed your project," she said, her voice sounding surprisingly **[word 2]**. The letter inside was incredibly **[word 3]**. Leaving the classroom, I felt a deep sense of **[word 4]** about the **[word 5]** work ahead.

Roll	Word 1	Word 2	Word 3	Word 4	Word 5
1	marched	exhausted	detailed	worry	exhausting
2, 3, 4	strode	stern	brief	dread	annoying
5, 6	rushed	excited	encouraging	determination	challenging

Narrative 1 and 2 Reflection Questions

1. Which narratives seemed more coherent or consistent? Was this true across your individual rolls?
2. Which narratives seemed less understandable or inconsistent? Was this true across your individual rolls?
3. What is the difference between Narrative 1 and Narrative 2 in terms of consistency or coherence? Was one narrative more likely to be coherent?
4. What is the difference between equal and weighted rolls? Was one more likely to be coherent?

Activity: How do LLMs work anyhow?

Rolling a Sentence with Two Dice

Each of you roll **two dice and add them together**. Unlike Narratives 1 and 2, there is no "equal probability" version of this table — some sums are simply more common than others, no weighting required.

Write the narrative that your rolls produce.

Narrative 3

The forecaster stood before the map and declared that today would be **[word 1]**. The chance of rain, she explained, was **[word 2]**. Looking out the window, everything outside appeared **[word 3]**. The crowd's mood shifted to **[word 4]** as they braced for what was shaping up to be a **[word 5]** afternoon.

Narrative 3

Sum	Ways to roll it	Word 1	Word 2	Word 3	Word 4	Word 5
2	1/36	apocalyptic	certain	terrifying	panic	disastrous
3	2/36	catastrophic	highly likely	ominous	dread	dreadful
4	3/36	severe	probable	gray	worry	difficult
5	4/36	unsettled	possible	hazy	unease	uneventful
6	5/36	cloudy	uncertain	overcast	indifference	routine
7	6/36	mild	fifty-fifty	ordinary	calm	ordinary
8	5/36	pleasant	unlikely	bright	relief	pleasant
9	4/36	sunny	doubtful	clear	cheer	enjoyable
10	3/36	gorgeous	improbable	radiant	joy	wonderful
11	2/36	perfect	remote	dazzling	excitement	magical
12	1/36	miraculous	impossible	otherworldly	euphoria	legendary

Reflection Questions for Narrative 3

1. Roll the dice **20 times individually** and record each sum. Count each roll in the table above.
2. Compare your 20 rolls with your group mates. Are your distributions identical? Should they be?
3. Add all your counts together and sketch the distribution of your group's rolls.
4. How is this different from the "weighted" tables in Narratives 1 and 2? There, *we* decided which word was more likely. Here, who or what decided it?
5. What might happen to the narrative if the number of dice was increased to 1,000? Or 10,000? Notice which words would fall near the middle of any distribution compared to those near the ends (e.g., for two dice: 6, 7, 8 vs. 2 and 12).

One thing this activity gets wrong on purpose

Dice tables are a **lookup**. A transformer is not.

The rolling shows you what the model does *with* the probabilities — the last step of a long process. The transformer is how those probabilities get **computed in the first place**, using everything in the context at once.

That's the piece the dice cannot show you.

Activity: Simulations — Watermarking AI-Produced Text

Our dice activities produced toy models of LLMs — but the text an LLM produces is random, *contextually* random.

The [European Union has passed a regulation](#) requiring generative AI outputs to be detectable as such.

Anthropic has developed a [watermarking technique](#) that biases the selection of the next word using the prior context (i.e., what was written before).

Go to deClaude.org/watermarking and work through the simulations.

Reflection Questions for Watermarking

1. How does an LLM typically produce the next word in a sentence?
2. How is Claude changing that with a "secret key"? What is it changing about the selection of the next word?
3. How does this secret key approach allow generated text to be detected?
4. What does this biasing of text generation using a secret key say about how these GenAI companies think about the task of writing?

Reminders

Exit Tickets, D2L, and the Course Website

- Don't forget to complete today's Exit Ticket by 5pm (on D2L)
- Case Study 3 due 11:59pm Sunday
- Individual and Metacognitive Reflection 3 due 11:59pm Sunday
- Spend time before Sunday (~90 minutes) going over Week 4's preparatory materials
 - Forum Post due 11:59pm **Sunday**

Your instructors will give detailed feedback on Week 2's assignments by Wednesday.