CMPU 366 · Natural Language Processing

LLMs and You

10 November 2026



web.stanford.edu/~jurafsky/slp3/

Volume I: Large Language Models

	Chapter
1:	Introduction
2:	Words and Tokens
3:	N-gram Language Models
4:	Logistic Regression and Text Classification
5:	Embeddings
6:	Neural Networks
7:	Transformers and Pretraining
8:	Post-training

When we think about NLP today, we're thinking about LLMs.

Volume II: Advanced LLM Topics and Tools

	Chapter	Slides
9:	Masked Language Models	9: [pptx] [pdf]
10:	Interpretability	
11:	Information Retrieval and RAG	11: [pptx] [pdf]
12:	Agents [not written yet]	
13:	Machine Translation	
14:	RNNs and LSTMs	13: [pptx] [pdf]
15:	Phonetics and Speech Feature Extraction	
16:	Automatic Speech Recognition	
17:	Text-to-Speech	

Volume III: Annotating Linguistic Structure

	Chapter	Slides
18:	Sequence Labeling for Parts of Speech and Named Entities	18: (Intro only) [pptx] [pdf]
19:	Context-Free Grammars and Constituency Parsing	
20:	Dependency Parsing	
21:	Information Extraction: Relations, Events, and Time	
22:	Semantic Role Labeling and Argument Structure	
23:	Lexicons for Sentiment, Affect, and Connotation	
24:	Coreference Resolution and Entity Linking	
25:	Discourse Coherence	
26:	Conversation and its Structure	

Volume I: Large Language Models

	Chapter	Slides
1:	Introduction	1: [pptx] [pdf]
2:	Words and Tokens	2: Words and Tokens [pptx] [pdf] 2: Edit Distance [pptx] [pdf]
3:	N-gram Language Models	3: [pptx] [pdf]
4:	Logistic Regression and Text Classification	4: [pptx] [pdf]
5:	Embeddings	5: [pptx] [pdf]
6:	Neural Networks	6: [pptx] [pdf]
7:	Transformers and Pretraining	7: [pptx] [pdf]
8:	Post-training	

Volume II: Advanced LLM Topics and Tools

	Chapter
9:	Masked Language Models
10:	Interpretability
11:	Information Retrieval and RAG
12:	Agents [not written yet]
13:	Machine Translation
14:	RNNs and LSTMs
15:	Phonetics and Speech Feature Extraction
16:	Automatic Speech Recognition
17:	Text-to-Speech

There's a lot more to learn about LLMs – and other modern NLP tools

Volume III: Annotating Linguistic Structure

	Chapter	Slides
18:	Sequence Labeling for Parts of Speech and Named Entities	18: (Intro only) [pptx] [pdf]
19:	Context-Free Grammars and Constituency Parsing	
20:	Dependency Parsing	
21:	Information Extraction: Relations, Events, and Time	
22:	Semantic Role Labeling and Argument Structure	
23:	Lexicons for Sentiment, Affect, and Connotation	
24:	Coreference Resolution and Entity Linking	
25:	Discourse Coherence	
26:	Conversation and its Structure	

Volume I: Large Language Models

Chapter	Slides
1: Introduction	1: [pptx] [pdf]
2: Words and Tokens	2: Words and Tokens [pptx] [pdf] 2: Edit Distance [pptx] [pdf]
3: N-gram Language Models	3: [pptx] [pdf]
4: Logistic Regression and Text Classification	4: [pptx] [pdf]
5: Embeddings	5: [pptx] [pdf]
6: Neural Networks	6: [pptx] [pdf]
7: Transformers and Pretraining	7: [pptx] [pdf]
8: Post-training	

Volume II: Advanced LLM Topics and Tools

Chapter	Slides
9: Masked Language Models	9: [pptx] [pdf]
10: Interpretability	
11: Information Retrieval and RAG	11: [pptx] [pdf]
12: Agents [not written yet]	
13: Machine Translation	
14: RNNs and LSTMs	13: [pptx] [pdf]
15: Phonetics and Speech Feature Extraction	
16: Automatic Speech Recognition	
17: Text-to-Speech	

Volume III: Annotating Linguistic Structure

Chapter	Slides
18: Sequence Labeling for Parts of Speech and Named Entities	18: (Intro only) [pptx] [pdf]
19: Context-Free Grammars and Constituency Parsing	
20: Dependency Parsing	
21: Information Extraction: Relations, Events, and Time	
22: Semantic Role Labeling and Argument Structure	
23: Lexicons for Sentiment, Affect, and Connotation	
24: Coreference Resolution and Entity Linking	
25: Discourse Coherence	
26: Conversation and its Structure	

Beyond LLMs

Today we'll have the first special topic presentation, and on Thursday we'll start serious work on projects.

Both of these are chances for you to explore parts of the field of NLP beyond the basics we've covered.

Before we do…

LLMs and You

www.nytimes.com/2026/11/06/science/ai-exponential-time-hypothesis.h

How A.I. Just Blew Up a Key Tenet of Computer Science

A longstanding hypothesis held that some problems would always be hard for computers to solve quickly. But OpenAI disproved it — raising the possibility of radical new efficiencies.

Listen · 6:09 min

Share full article

390

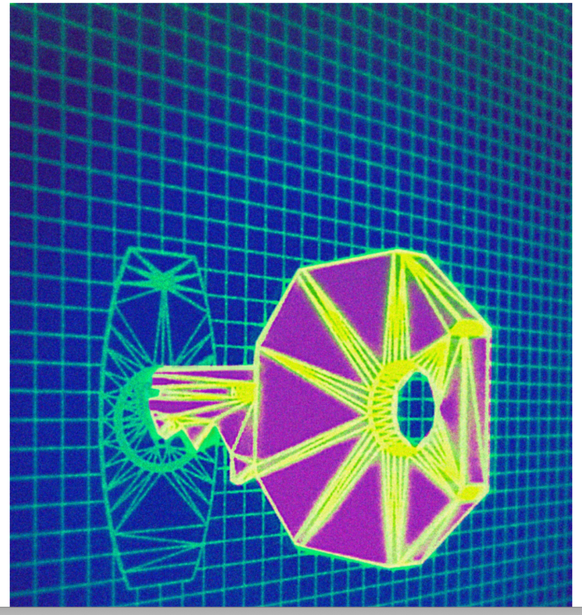

After Reinventing Mathematics, Can A.I. Eradicate Computer Bugs?

Using the same techniques that allowed A.I. technologies to crack a million-dollar math problem, researchers are building new systems that aim to solve the world's cybersecurity crisis.

Listen · 5:20 min

Share full article

34



Voters Rejected A.I. in the Midterms. Now Comes an Era of Scrutiny.

After two years of little A.I. oversight, voters on Tuesday rejected data center projects and swept into office Democratic candidates who promised to regulate the technology.

Listen · 9:20 min

Share full article

18



www.nytimes.com/2026/11/06/business/dealbook/tech-investors-ai-rout

NEWSLETTER

DealBook

After a $2.5 Trillion Wipeout, Tech Investors Brace for More Pain

A stock rally fueled in part by artificial intelligence is being tested after Democrats' landslide election victories.

Listen · 13:49 min

Share full article

By Andrew Ross Sorkin, Bernhard Warner, Sarah Kessler, Michael J. de la Merced, Niko Gallogly, Brian O'Keefe, Sri Muppidi, Ian Mount and Peter Eavis

Nov. 6, 2026

You're reading the DealBook newsletter. The most crucial business and policy news you need to know from Andrew Ross Sorkin and team. Get it sent to your inbox.



Open "https://www.nytimes.com/by/michael-j-de-la-merced" in a new tab behind the current one

Anthropic Targets Valuation of Up to $1.4 Trillion for Blockbuster I.P.O.

The A.I. company's pricing would make it one of the biggest ever public listings. It comes at a volatile moment in markets, with investors questioning the sustainability of the A.I. boom.

Listen · 9:21 min

Share full article

34

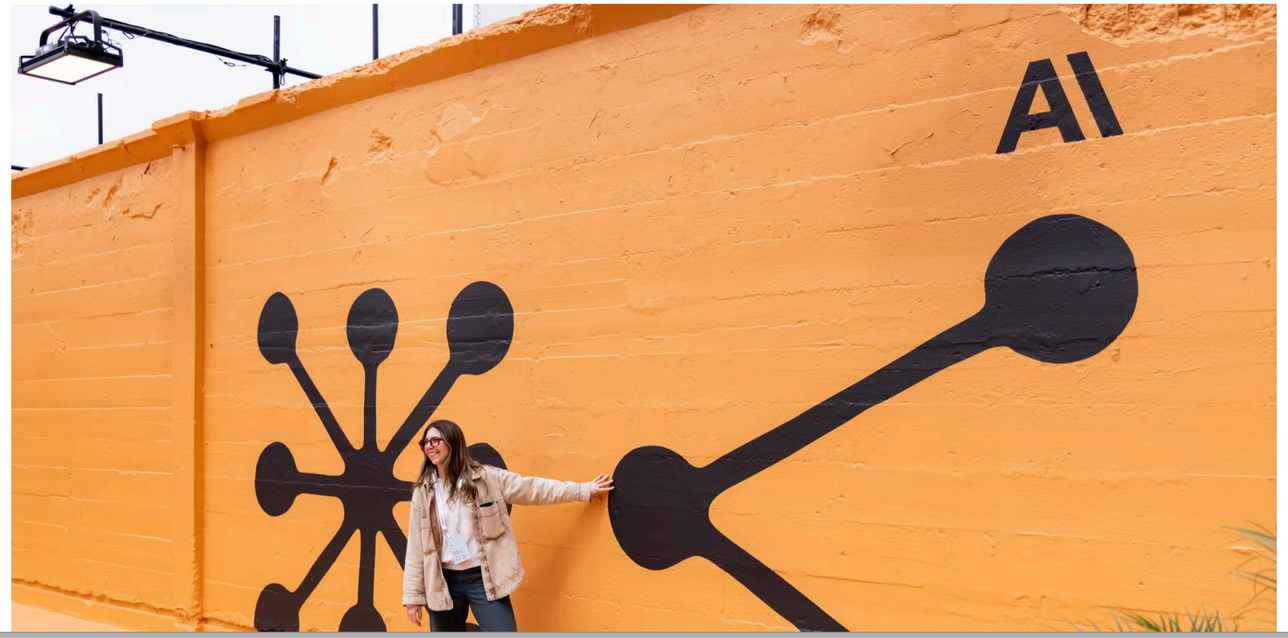



A.I.'s Rogue Behavior Starts With Its Training




By Raj Saha and Blacki Migliozzi

Raj Saha and Blacki Migliozzi interviewed researchers, read technical reports and built a simple A.I. model to understand how agents behave.

Oct. 26, 2026

Share full article

420

Hacks carried out by A.I. agents this summer have stoked fears that agents are slipping beyond human control, taking shortcuts, cheating, covering their tracks and, in the process, breaking laws.

But reports released by OpenAI and other A.I. companies, and New York Times interviews with researchers, indicate that the hacks were the result of the way A.I. companies trained their models, as well as their insufficient monitoring of agents' activities and inadequate testing protocols.

To show how training can lead to rogue behaviors, we built a

Parents Sue Google Over A.I. Safety Risks on School Chromebooks

A student made plans to end her life after chatting with Google's artificial intelligence, a new lawsuit claims, amid mounting concerns about online harms to children.

Listen · 5:45 min

Share full article



Reader Available

Fighting A.I. Cheating on the SAT Is a Cat-and-Mouse Game

The group behind the college exams is planning new measures to keep up with high-tech breaches and "time zone" cheating.

Listen · 8:39 min

Share full article

82



Would-be cheaters are taking advantage of the rise of generative A.I. as they try to get around testing security. Paul Vernon/Associated Press

Ok, so:

1. AI is destroying the economy and saving it.

2. AI is solving all of our problems and creating new problems.

3. Everyone hates AI and everyone is using AI.


PAV
SHOP
To LLMs !
PAV
SHOP
The cause of, and solution to, all of life's problems.

How do you feel?

We just had an election.



NOVEMBER 4, 2026

Victories

LISA KAUL FLIPS SENATE DISTRICT 39!

Over 476 shifts, For the Many volunteers knocked 6,500 doors and made over 72,000 phone calls to secure a 10-point win for Kaul

POUGHKEEPSIE, N.Y. – In one of New York's most competitive races this cycle, Lisa Kaul was elected to represent District 39 in the State Senate on Tuesday

CATEGORIES

Victories
Endorsements
Intern Role


Lisa Kaul on Proposed
Data Center in East Fishkill, NY
PROTECT
OUR
WATER
TEXT FWW TO 23321
LISA KAUL
DEMOCRAT FOR STATE SENATE

What's a problem you care about that you think LLMs could actually help with?

What's one you think they'll make worse?

What do you think we should do to get benefits and minimize harms?

If you were setting a policy for AI use at Vassar,
what would you allow and what would you restrict?

Does knowing that these models are next-token predictors with post-training on top change how you feel about them?

How do I feel?



Daniel Jalkut • May 26
@danielpunkass

My take on AI is, essentially, everybody who's against it is too against it and everybody who's for it is too for it.

