Affordances of AI

Adam Vosburgh

GSAPP AI Retreat

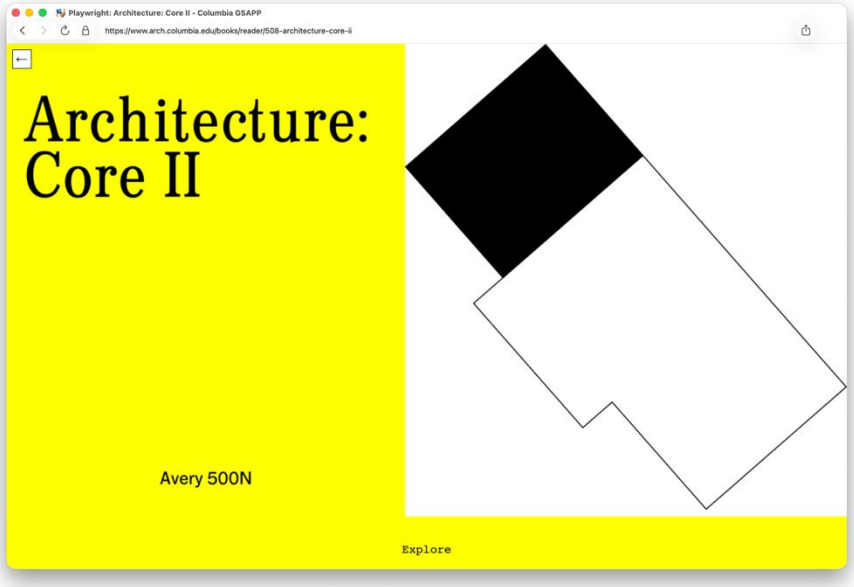
September 18th, 2026

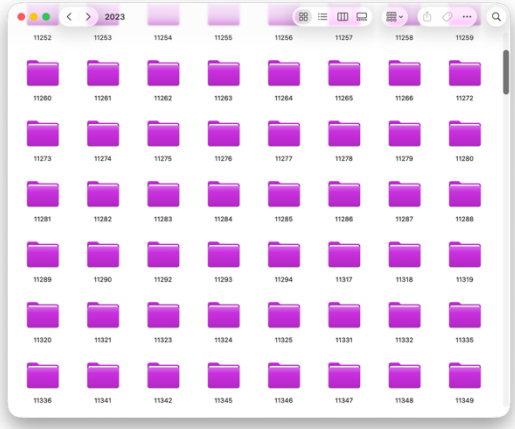
presentation overview

- "Artificial Intelligence"
- How an AI model is built
- What an AI model does
- Affordances of AI
- The view from MS CDP
- Questions?

References, visualizations, slide pdf and further reading at ***retreat.adamvosburgh.com***
password: ***gsapp-ai***

presentation overview





References, visualizations, slide pdf and further reading at ***retreat.adamvosburgh.com***
password: ***gsapp-ai***

"Artificial Intelligence"

"Artificial Intelligence"

machine learning — a computational method that is a subfield of artificial intelligence and that enables a computer to learn to perform tasks by analyzing a large dataset without being explicitly programmed (Merriam-Webster)

deep learning — a form of machine learning in which the computer network rapidly teaches itself to understand a concept without human intervention by performing a large number of iterative calculations on an extremely large dataset (Merriam-Webster)

natural language processing (NLP) — the branch of computer science that involves giving computers the ability to interpret and generate human language (Cambridge Advanced Learner's Dictionary)

large language model (LLM) — a language model that utilizes deep methods on an extremely large data set as a basis for predicting and constructing natural-sounding text (Merriam-Webster)

computer vision — a robot analogue of human vision in which information about the environment is received by one or more video cameras and processed by computer (Dictionary)

diffusion model — a type of machine learning algorithm trained by having the algorithm add noise to images and then remove it from them: used especially for generating pictures from text prompts or for repairing, transforming, or expanding existing images (Dictionary)

"Artificial Intelligence"

machine learning — a computational method that is a subfield of artificial intelligence and that enables a computer to learn to perform tasks by analyzing a large dataset without being explicitly programmed (Merriam-Webster)

deep learning — a form of machine learning in which the computer network rapidly teaches itself to understand a concept without human intervention by performing a large number of iterative calculations on an extremely large dataset (Merriam-Webster)

natural language processing (NLP) — the branch of computer science that involves giving computers the ability to interpret and generate human language (Cambridge Advanced Learner's Dictionary)

large language model (LLM) — a language model that utilizes deep methods on an extremely large data set as a basis for predicting and constructing natural-sounding text (Merriam-Webster)

computer vision — a robot analogue of human vision in which information about the environment is received by one or more video cameras and processed by computer (Dictionary)

diffusion model — a type of machine learning algorithm trained by having the algorithm add noise to images and then remove it from them: used especially for generating pictures from text prompts or for repairing, transforming, or expanding existing images (Dictionary)

When we refer to "AI", we are almost always referring to a mix of these things

"Artificial Intelligence"

machine learning — a computational method that is a subfield of artificial intelligence and that enables a computer to learn to perform tasks by analyzing a large dataset without being explicitly programmed (Merriam-Webster)

deep learning — a form of machine learning in which the computer network rapidly teaches itself to understand a concept without human intervention by performing a large number of iterative calculations on an extremely large dataset (Merriam-Webster)

natural language processing (NLP) — the branch of computer science that involves giving computers the ability to interpret and generate human language (Cambridge Advanced Learner's Dictionary)

large language model (LLM) — a language model that utilizes deep methods on an extremely large data set as a basis for predicting and constructing natural-sounding text (Merriam-Webster)

computer vision — a robot analogue of human vision in which information about the environment is received by one or more video cameras and processed by computer (Dictionary)

diffusion model — a type of machine learning algorithm trained by having the algorithm add noise to images and then remove it from them: used especially for generating pictures from text prompts or for repairing, transforming, or expanding existing images (Dictionary)

When we refer to "AI", we are almost always referring to a mix of these things.

"Artificial Intelligence"

machine learning — a computational method that is a subfield of artificial intelligence and that enables a computer to learn to perform tasks by analyzing a large dataset without being explicitly programmed (Merriam-Webster)

deep learning — a form of machine learning in which the computer network rapidly teaches itself to understand a concept without human intervention by performing a large number of iterative calculations on an extremely large dataset (Merriam-Webster)

natural language processing (NLP) — the branch of computer science that involves giving computers the ability to interpret and generate human language (Cambridge Advanced Learner's Dictionary)

large language model (LLM) — a language model that utilizes deep methods on an extremely large data set as a basis for predicting and constructing natural-sounding text (Merriam-Webster)

computer vision — s received by one or more video cameras and

diffusion model — to images and then remove it from them: used es , or expanding existing images (Dictionary)

It gets kind of complicated, but generally:

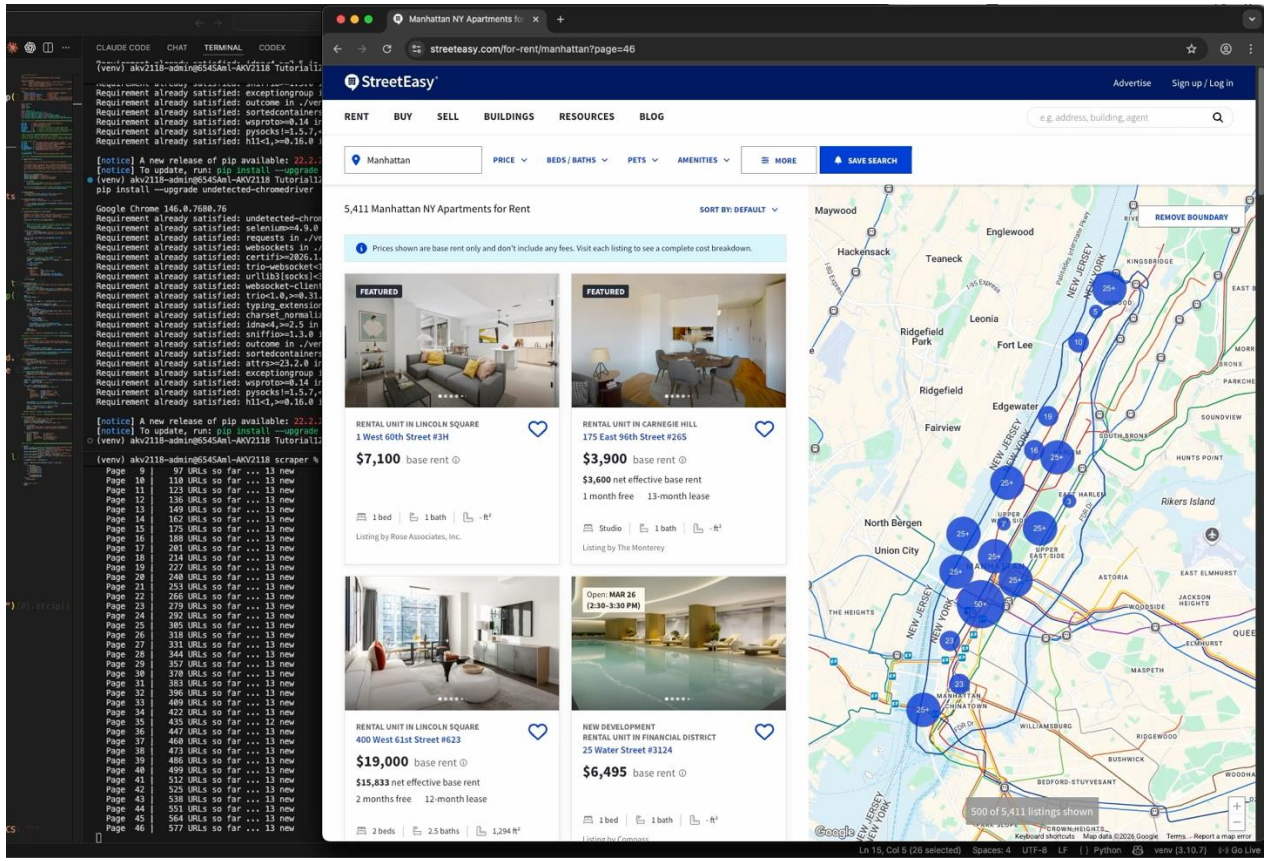
LLM's generate text (writing, programming code, tabular data, etc)

Diffusion models generate images, videos

When we refer to "AI", we are almost always referring to a mix of these things.

How an AI model is built

How an AI model is built



Scraping StreetEasy, and making spatial data from plain-text, Methods in Spatial Research Spring 2026

How an AI model is built

"...in this book I argue that AI is neither *artificial* nor *intelligent*. Rather, artificial intelligence is both embodied and material, made from natural resources, fuel, human labor, infrastructures, logistics, histories, and classifications."

— Kate Crawford, Atlas of AI (2021), p. 8

How an AI model is built

Overview: Data, training, tuning, weights

LLM:

specifically Meta's LLaMA (2023):

1 DATA	roughly 1.4 trillion tokens of text - web crawl, books, code, Wikipedia, papers. For this example, somewhere inside it there are a few million apartment listings
2 TRAIN	show the model text with the next word hidden, compare its guess to the real word, nudge the weights, repeat "Sunny 1BR with a private rooftop ___" -> terrace billions of times, for weeks, on thousands of GPUs
3 TUNE	people write model answers and rank the model's attempts, so it learns what a useful reply looks like and in what register
4 WEIGHTS	what is left is about 65 billion numbers. The listings are not in there; what remains is the pattern they made

Diffusion model:

specifically Stable Diffusion v1 (2022):

1 DATA	billions of web images with their alt text. For this example, somewhere inside it there are a few million listing photos
2 TRAIN	add noise to a photo, have the model guess the noise from the noisy photo and its caption, compare, nudge the weights, repeat photo + "Sunny 1BR with a private rooftop terrace" -> noise about 2.4 billion times, on 256 GPUs
3 TUNE	keep training only on images a scoring model rates as good-looking, which pulls results toward that taste
4 WEIGHTS	what is left is about 1 billion numbers. Most photos are not in there, but some can be pulled back out almost exactly

How an AI model is built:

1: Pre-Training Data

LLM:

Pre-training mix for Meta's LLaMA (2023)

CommonCrawl	67.0%	filtered web crawl, 2017–2020 snapshots
C4	15.0%	a second, differently filtered web crawl
GitHub	4.5%	public code repositories
Wikipedia	4.5%	articles in 20 languages
Books	4.5%	Project Gutenberg + pirated book collections
arXiv	2.5%	scientific papers, LaTeX source
StackExchange	2.0%	Q&A across 28 sites

total: about 1.4 trillion tokens

Source: Touvron et al., "LLaMA: Open and Efficient Foundation Language Models," Meta AI, 2023, Table 1.

Diffusion model:

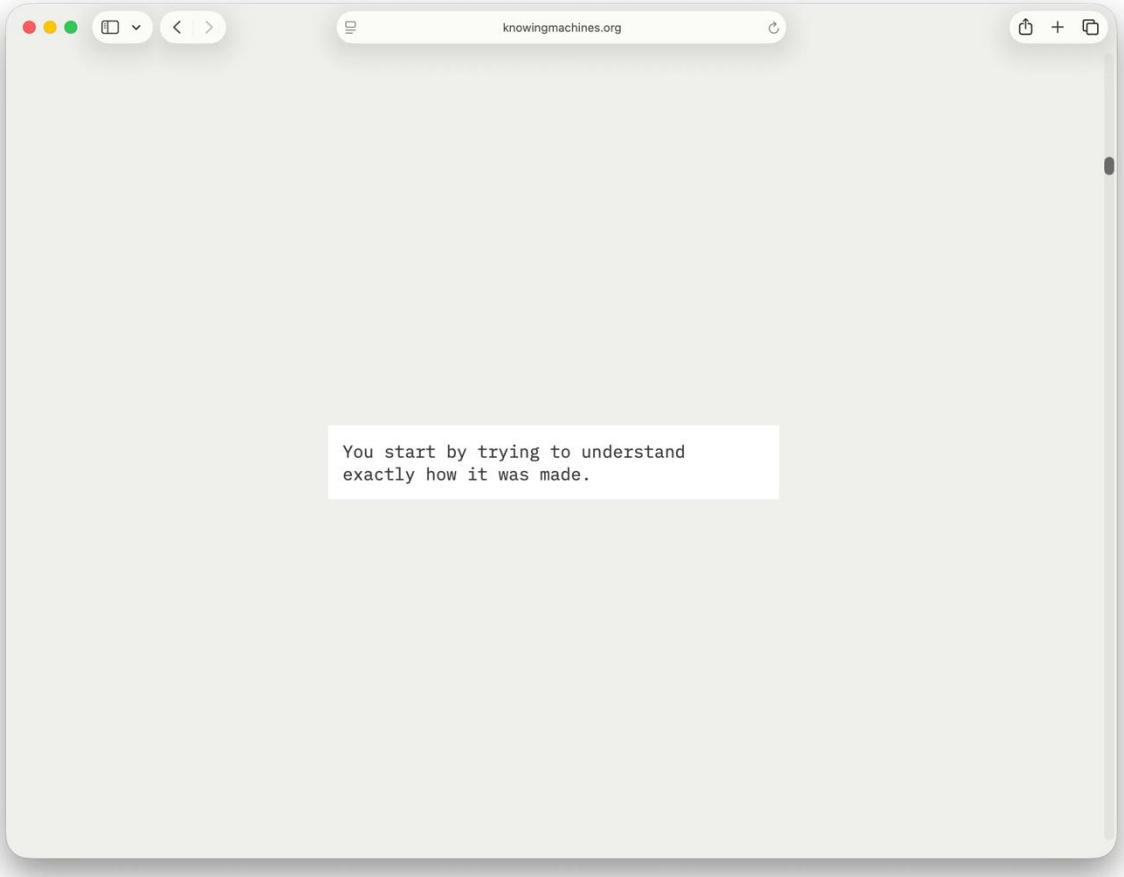
Training data for Stable Diffusion v1 (2022)

LAION-5B: 5.85 billion images paired with their alt text, collected from web pages saved by Common Crawl

laion2B-en	39.7%	2.32 billion pairs, English text
laion2B-multi	38.6%	2.26 billion pairs, 100+ other languages
laion1B-nolang	21.7%	1.27 billion pairs, language unclear

Source: CompVis, "Stable Diffusion v1-4" model card, 2022; Schuhmann et al., "LAION-5B," NeurIPS 2022; LAION, "Releasing Re-LAION-5B," 2024.

How an AI model is built:



Christo Buschek & Jer Thorp. "Models All the Way Down". A visualization of LAION-5B.

How an AI model is built:

2: Pre-Training

LLM:

```
Training loop (simplified):

  input:    "Sunny 1BR with a private rooftop"

  predict:  terrace 34%   deck 21%   garden 17%   pool 11%   ...

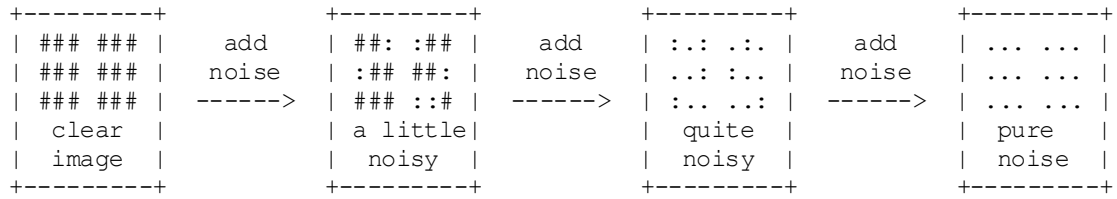
  actual:   "terrace"

  repeat:   once per token in the corpus — about 1.4 trillion times
```

Diffusion model:

```
Training loop (simplified):

  input:    photo + alt text "Sunny 1BR with a private rooftop terrace"
```




```
  predict:  given the noisy photo and the text, what noise was added?

  actual:   the noise that was really added; compare, nudge the weights

  repeat:   about 2.4 billion times (1.17 million steps x 2,048 images)
```

How an AI model is built:

3: Annotation and Tuning — supervised fine-tuning

```
Annotator writes example conversations:

  Prompt:    "Summarize this listing in one sentence for a first-time renter."
  Response:  "A small but bright one-bedroom on a quiet block, about a
              ten-minute walk to the 7 train."

  Prompt:    "Write a listing for a 2BR in Jackson Heights with a balcony."
  Response:  "Sun-filled two-bedroom on a tree-lined street, ..."

  × tens of thousands of examples

          ↓

  Model is tuned to imitate these: it learns the tone and format
  a person expects in a reply
```

RLHF

```
Prompt: "What's a good neighborhood in Manhattan for a GSAPP student?"

Model writes 4 responses; an annotator ranks them:

  Response A  ==================   1st
  Response B  ================     2nd
  Response C  ==============       3rd
  Response D  ==========           4th

          ↓
  A "reward model" is trained on thousands of these rankings
          ↓
  The main model is tuned to produce answers the reward model scores highly
          ↓
  Result: the model learns which answers people consider "better"
```

What an AI model does

(I'll do an LLM first, then diffusion)

What an AI model does

Tokenize, embed, predict

```
TOKENIZE   "Charming pre-war 1BR"  →  [Charm] [ing] [ pre] [-war] [ 1] [BR]
           pieces of words — roughly 50–150k pieces

EMBED      each token → a list of numbers: a position in "meaning space"
           tokens used in similar ways end up near each other
             · cozy · charming · intimate          · doorman · concierge · elevator
             · loft · studio · 1BR                 · Astoria · Sunnyside · Woodside

PREDICT    given "The apartment has a rooftop", score every token in the vocabulary:
             terrace   34%  ████████
             deck      21%  ██████
             garden    17%  ████
           pick one, append it, and repeat for the next token
```

What an AI model does: Embedding

Tokens → vectors

```
Input text:

  "Charming pre-war 1BR in Morningside Heights"

Tokens:

  ["Charm", "ing", " pre", "-war", " 1", "BR", " in", " Morning", "side", " Heights"]

Token IDs (integers):

  [ 8580,  278,  1248,  3162,  16,  1647,  287,  22770,  1735,  21879 ]

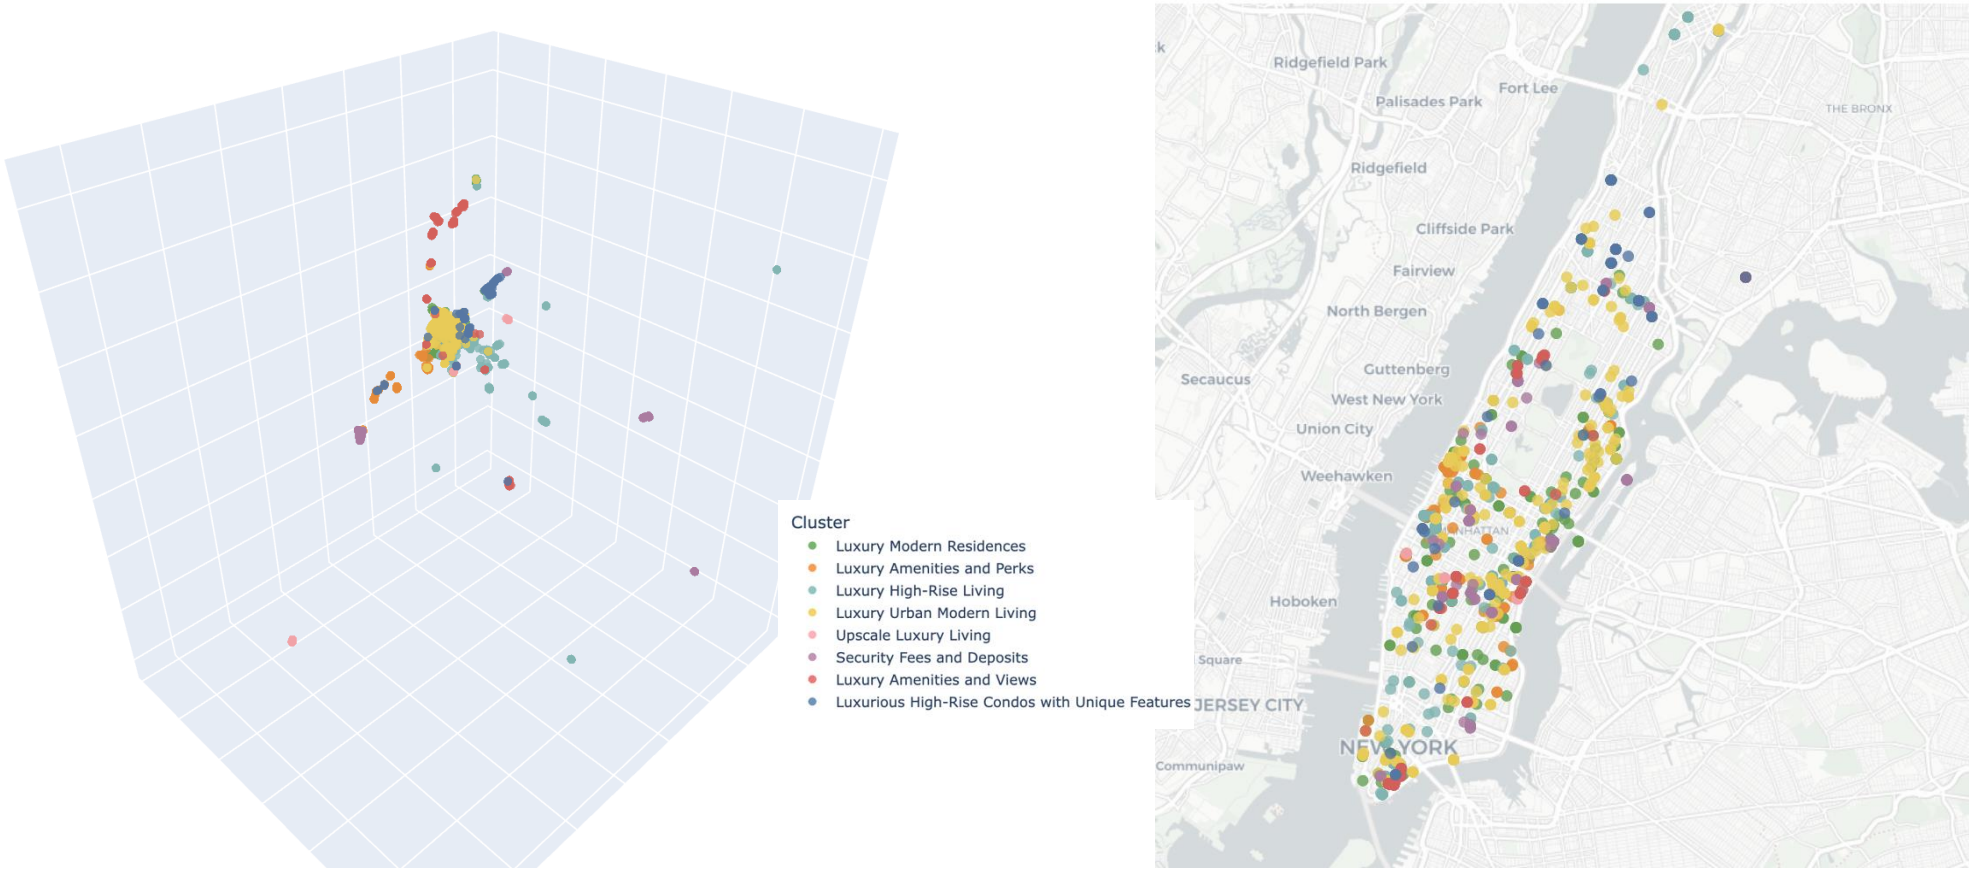
One token as a vector:

  Token:    " Morning"
  ID:       22770
  Vector:   [ 0.21, -0.54,  0.88,  0.03, -0.77,  0.41, ... ]   (768 numbers with this small
                                                                 model; thousands in a large one)

Vectors encode meaning – each number is an axis in an embedding space.

Tokens used in similar ways end up near each other
   · cozy · charming · intimate          · doorman · concierge · elevator
   · loft · studio · 1BR                 · Astoria · Sunnyside · Woodside
```

What an AI model does: Embedding



Sentence embeddings of ~500 StreetEasy listings, Methods in Spatial Research Spring 2026

What an AI model does: Embedding

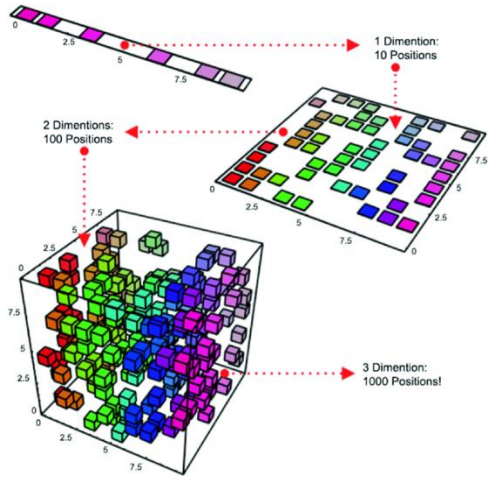
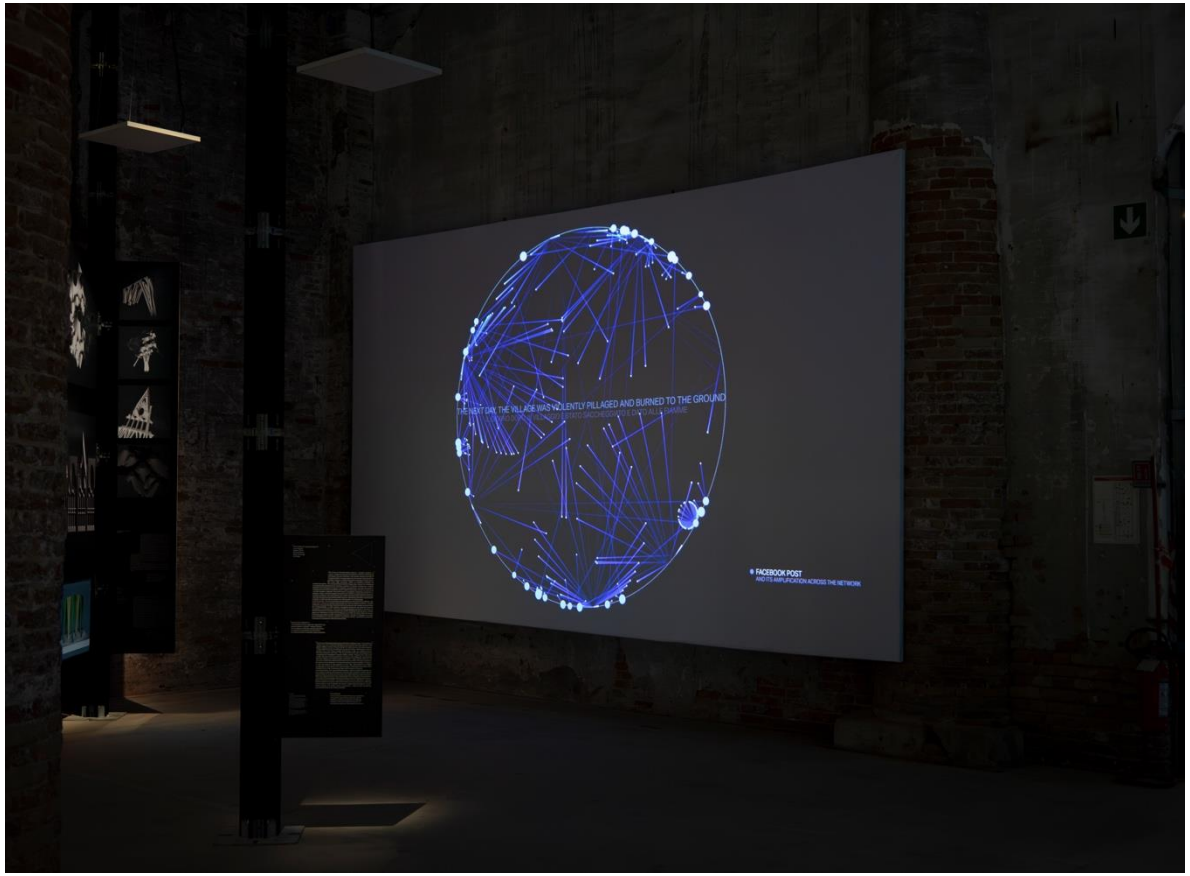


Diagram of "The Curse of Dimensionality"



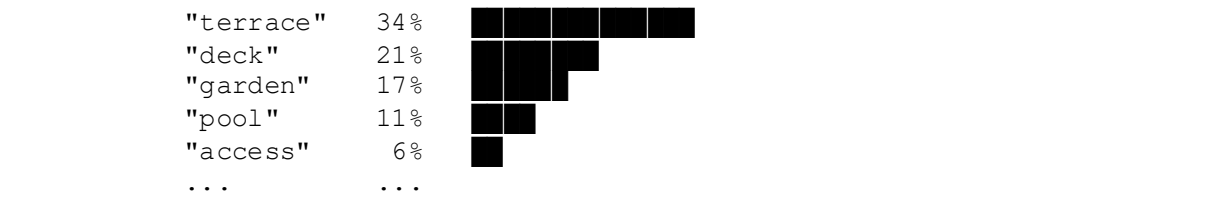
Chum, Krisch, Kurgan, Vosburgh, Zhang. "The Curse of Dimensionality". Venice Architecture Biennale 2025.

What an AI model does: Prediction in LLMs

Predicting the next token

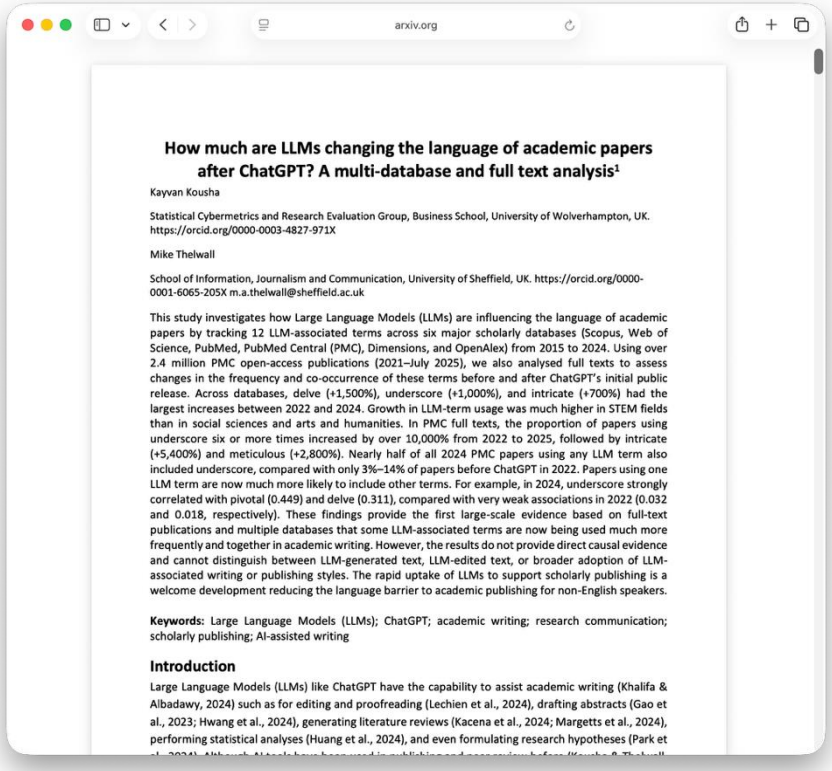
Context: "The apartment has a rooftop ___"

Model output: a probability for every one of its ~50k+ tokens



One token is chosen (usually, but not always, the most likely one), added to the context, and the process repeats.

Affordances of AI: Prediction



How much are LLMs changing the language of academic papers after ChatGPT? A multi-database and full text analysis[1]

Kayvan Kousha

Statistical Cybermetrics and Research Evaluation Group, Business School, University of Wolverhampton, UK. https://orcid.org/0000-0003-4827-971X

Mike Thelwall

School of Information, Journalism and Communication, University of Sheffield, UK. https://orcid.org/0000-0001-6065-205X m.a.thelwall@sheffield.ac.uk

This study investigates how Large Language Models (LLMs) are influencing the language of academic papers by tracking 12 LLM-associated terms across six major scholarly databases (Scopus, Web of Science, PubMed, PubMed Central (PMC), Dimensions, and OpenAlex) from 2015 to 2024. Using over 2.4 million PMC open-access publications (2021–July 2025), we also analysed full texts to assess changes in the frequency and co-occurrence of these terms before and after ChatGPT's initial public release. Across databases, delve (+1,500%), underscore (+1,000%), and intricate (+700%) had the largest increases between 2022 and 2024. Growth in LLM-term usage was much higher in STEM fields than in social sciences and arts and humanities. In PMC full texts, the proportion of papers using underscore six or more times increased by over 10,000% from 2022 to 2025, followed by intricate (+5,400%) and meticulous (+2,800%). Nearly half of all 2024 PMC papers using any LLM term also included underscore, compared with only 3%–14% of papers before ChatGPT in 2022. Papers using one LLM term are now much more likely to include other terms. For example, in 2024, underscore strongly correlated with pivotal (0.449) and delve (0.311), compared with very weak associations in 2022 (0.032 and 0.018, respectively). These findings provide the first large-scale evidence based on full-text publications and multiple databases that some LLM-associated terms are now being used much more frequently and together in academic writing. However, the results do not provide direct causal evidence and cannot distinguish between LLM-generated text, LLM-edited text, or broader adoption of LLM-associated writing or publishing styles. The rapid uptake of LLMs to support scholarly publishing is a welcome development reducing the language barrier to academic publishing for non-English speakers.

Keywords: Large Language Models (LLMs); ChatGPT; academic writing; research communication; scholarly publishing; AI-assisted writing

Introduction

Large Language Models (LLMs) like ChatGPT have the capability to assist academic writing (Khalifa & Albadawy, 2024) such as for editing and proofreading (Lechien et al., 2024), drafting abstracts (Gao et al., 2023; Hwang et al., 2024), generating literature reviews (Kacena et al., 2024; Margetts et al., 2024), performing statistical analyses (Huang et al., 2024), and even formulating research hypotheses (Park et

Kousha & Thelwall, "How much are LLMs changing the language of academic papers after ChatGPT?", Scientometrics (2026; preprint 2025)

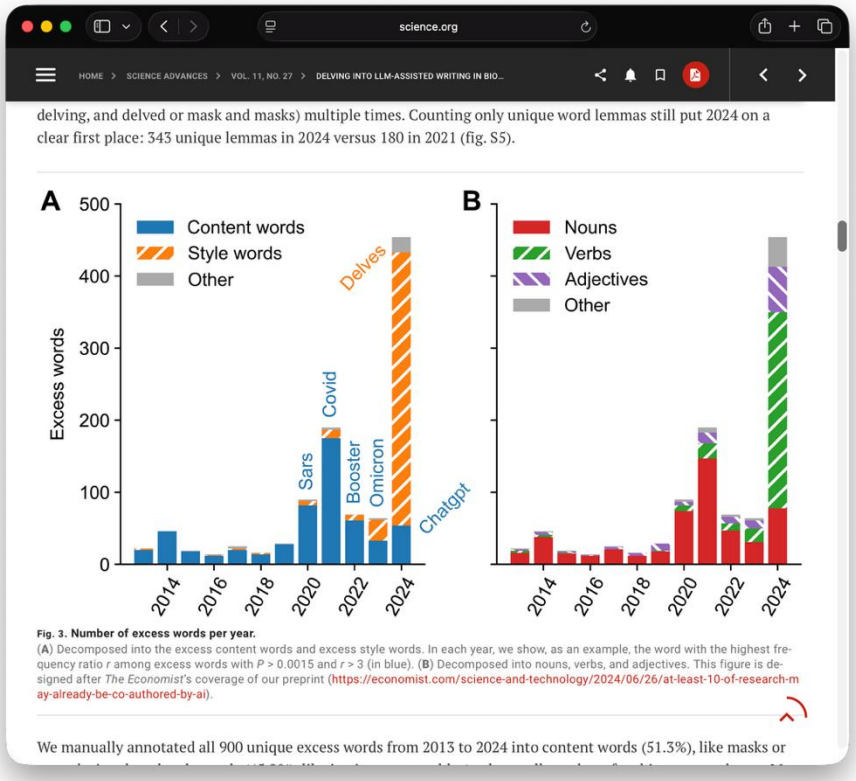
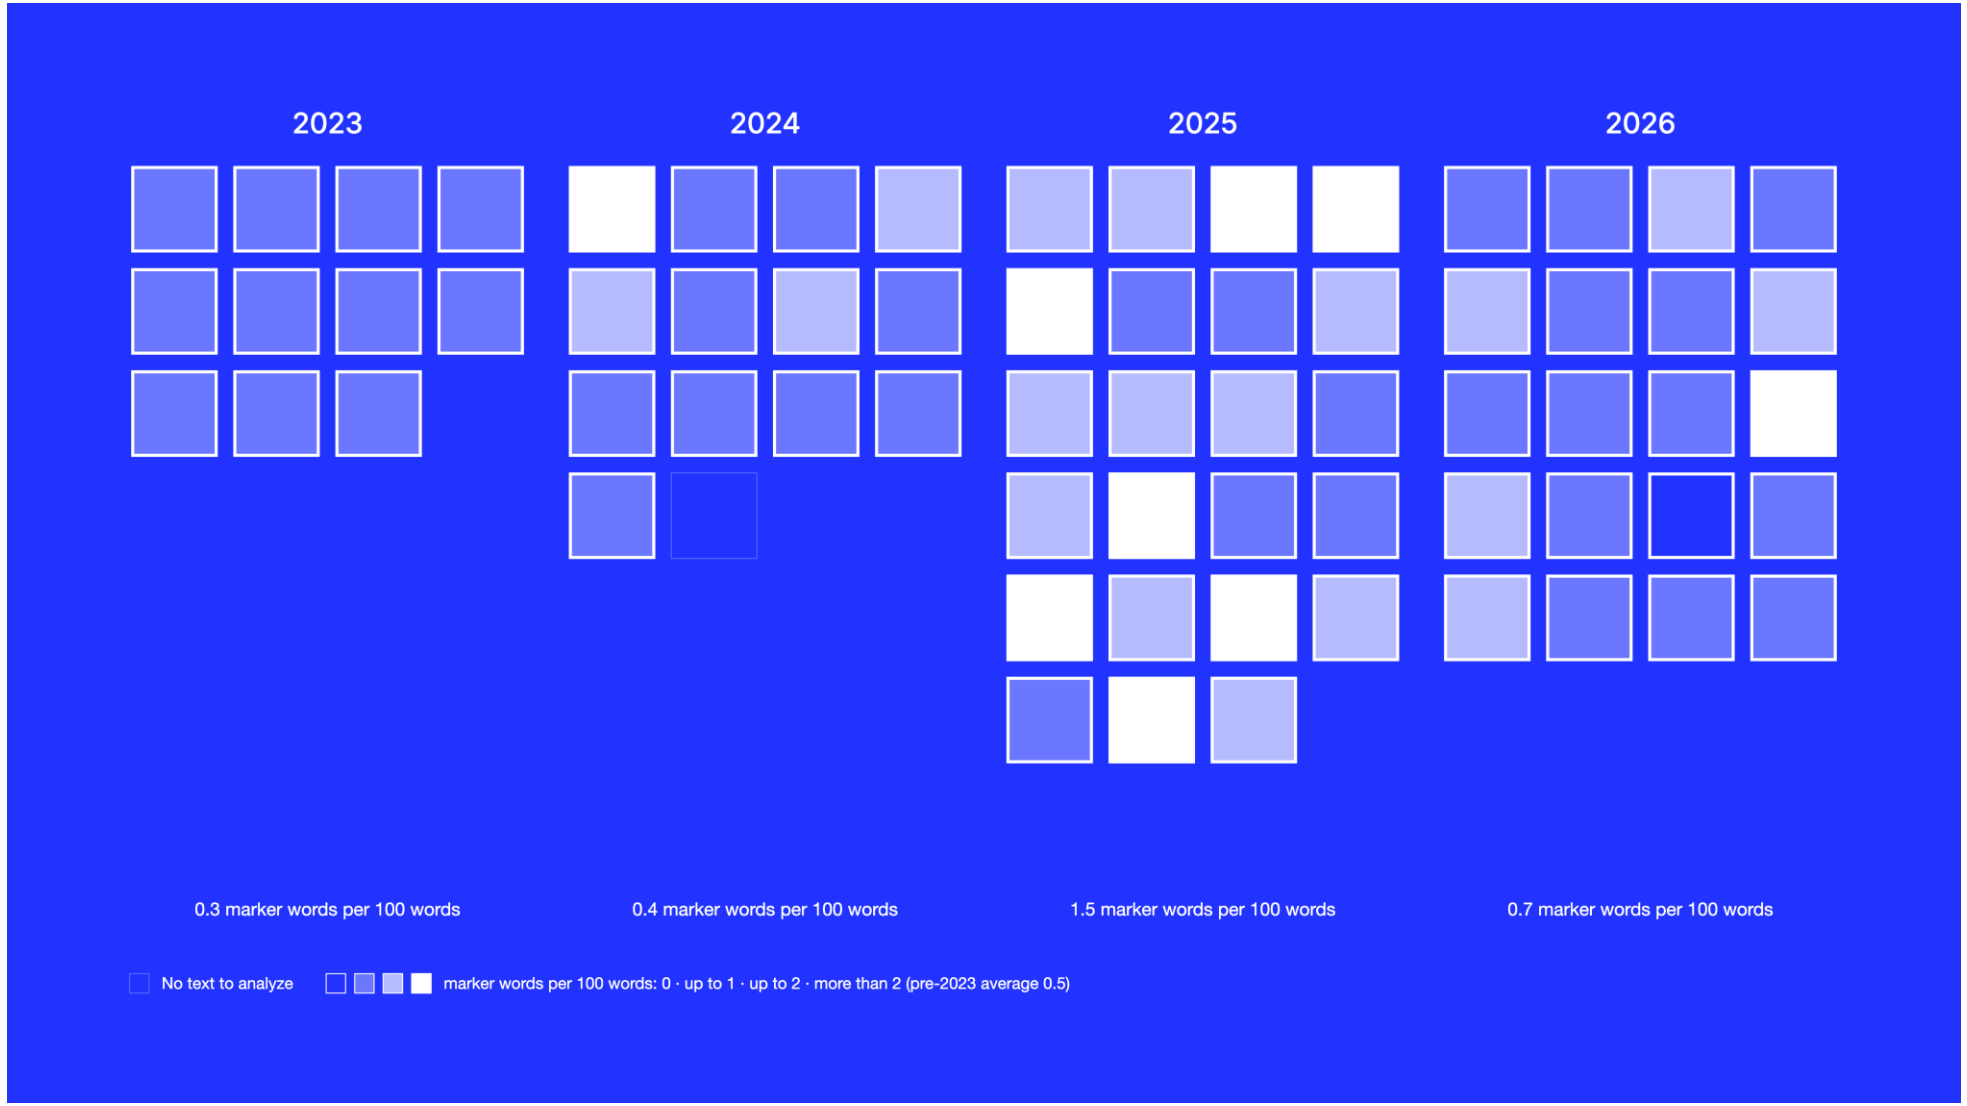


Fig. 3. Number of excess words per year.
(**A**) Decomposed into the excess content words and excess style words. In each year, we show, as an example, the word with the highest frequency ratio *r* among excess words with $P > 0.0015$ and $r > 3$ (in blue). (**B**) Decomposed into nouns, verbs, and adjectives. This figure is designed after *The Economist*'s coverage of our preprint (https://economist.com/science-and-technology/2024/06/26/at-least-10-of-research-may-already-be-co-authored-by-ai).

We manually annotated all 900 unique excess words from 2013 to 2024 into content words (51.3%), like masks or

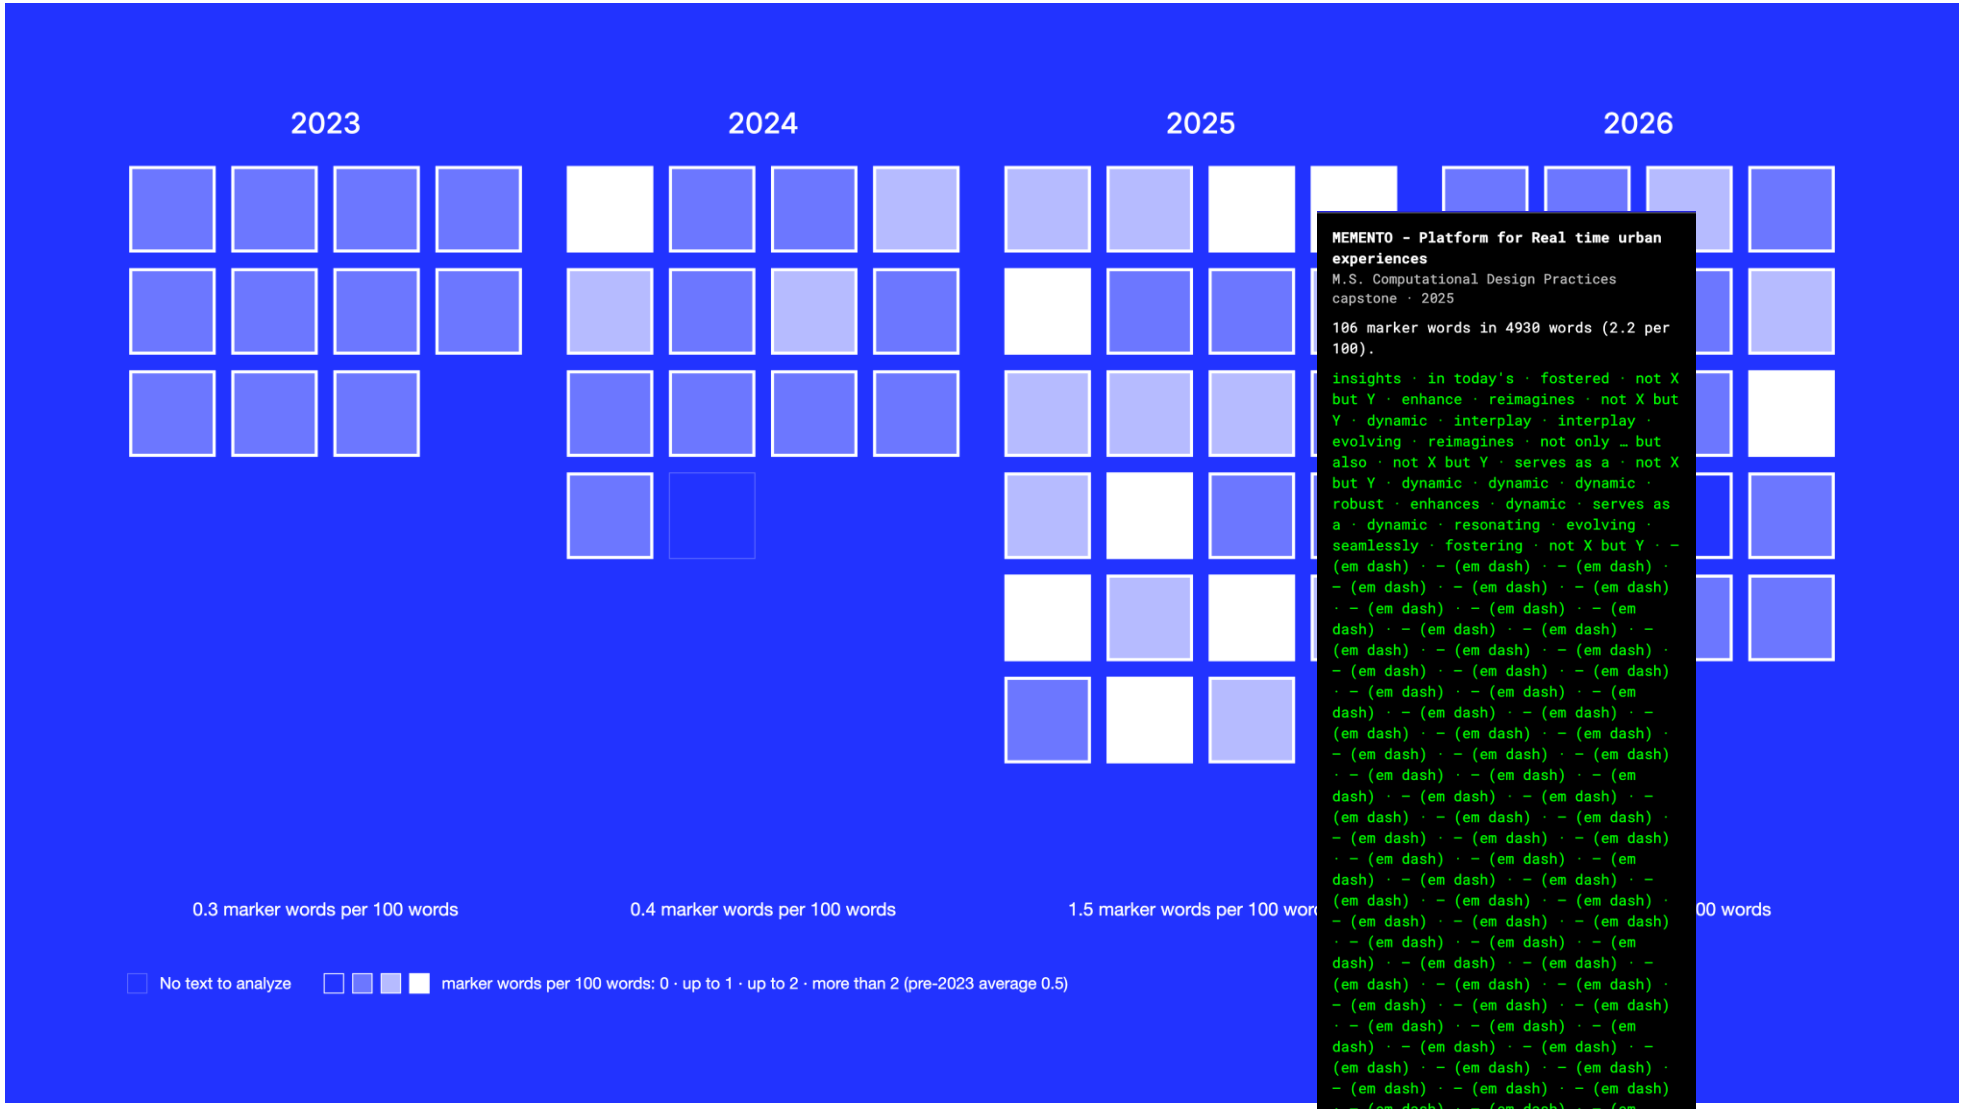
Kobak et al., "Delving into LLM-assisted writing in biomedical publications through excess vocabulary", Science Advances (2025)

Affordances of AI: Prediction

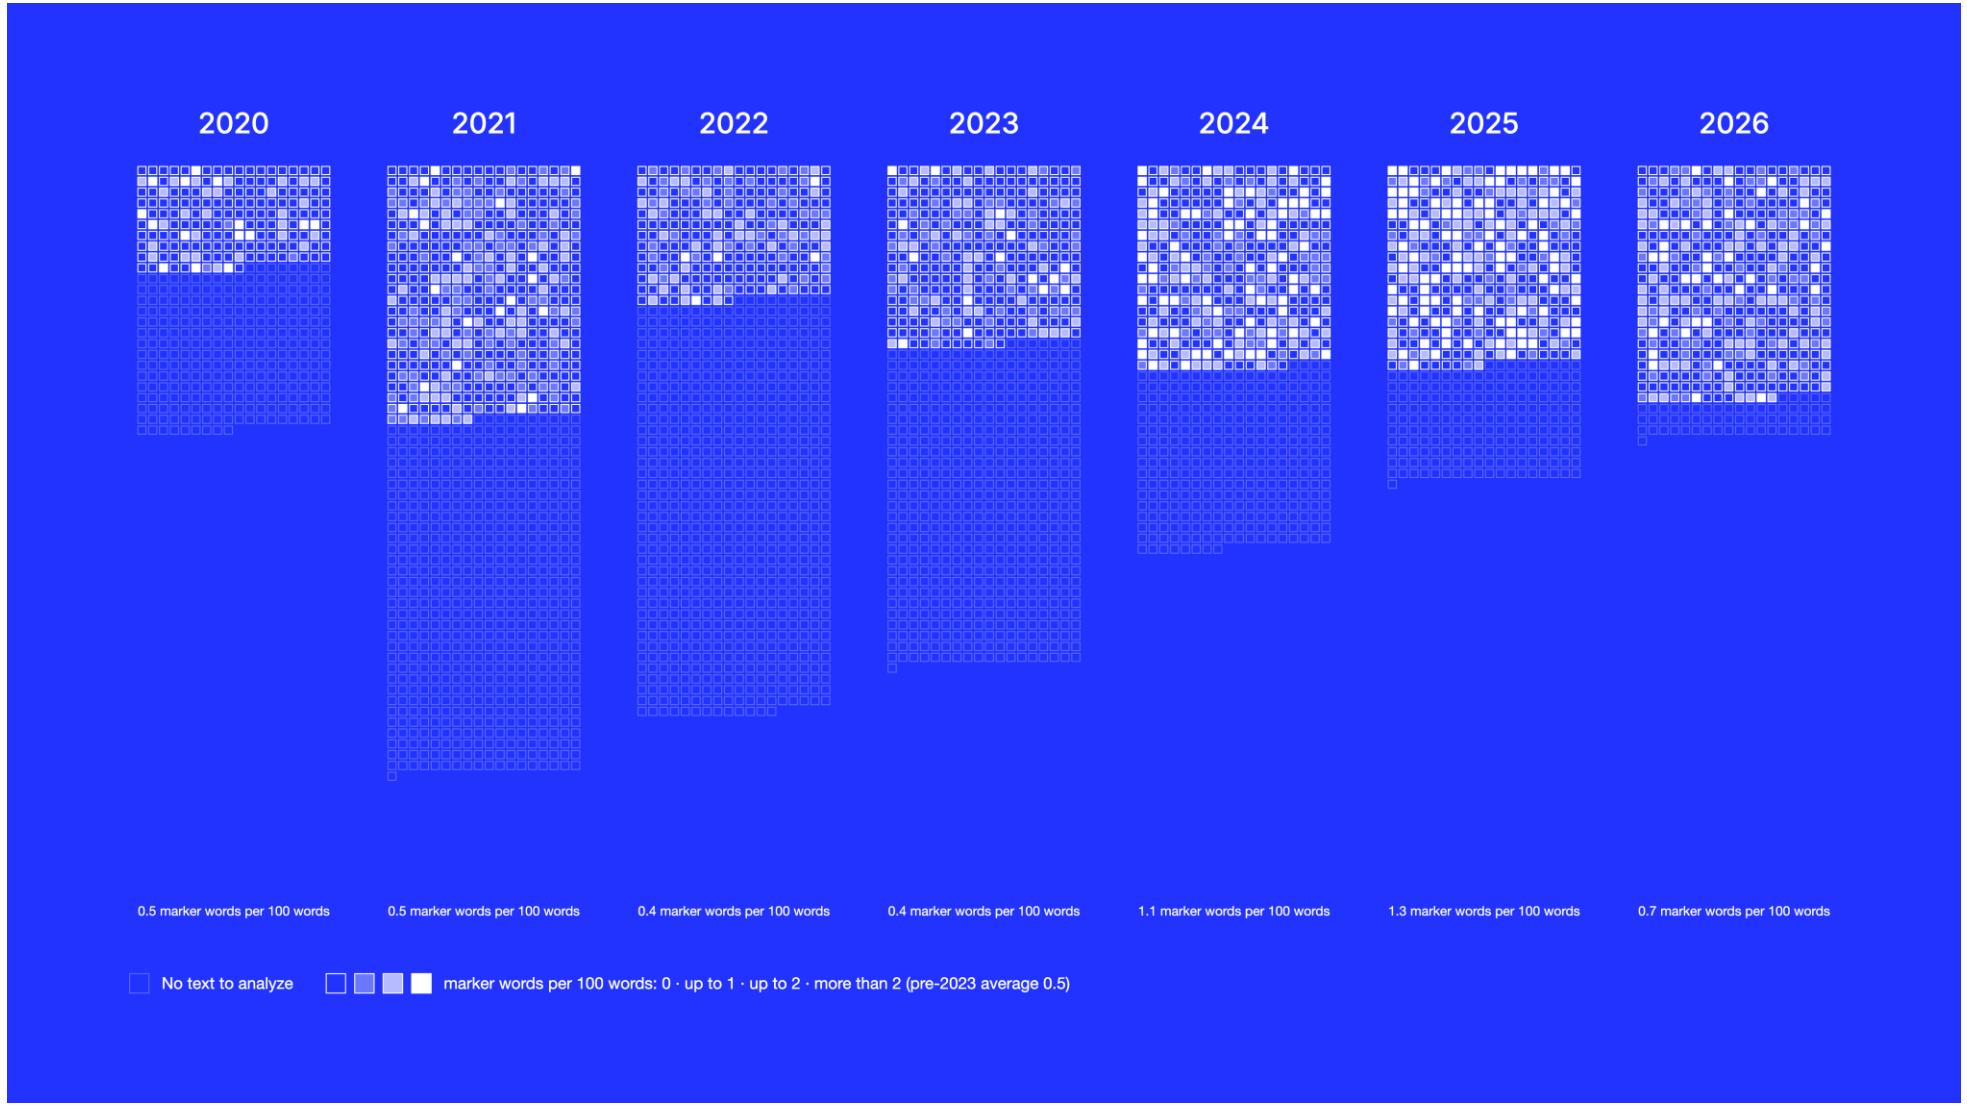


AI marker words in GSAPP CDP work – 2023 to 2026

Affordances of AI: Prediction

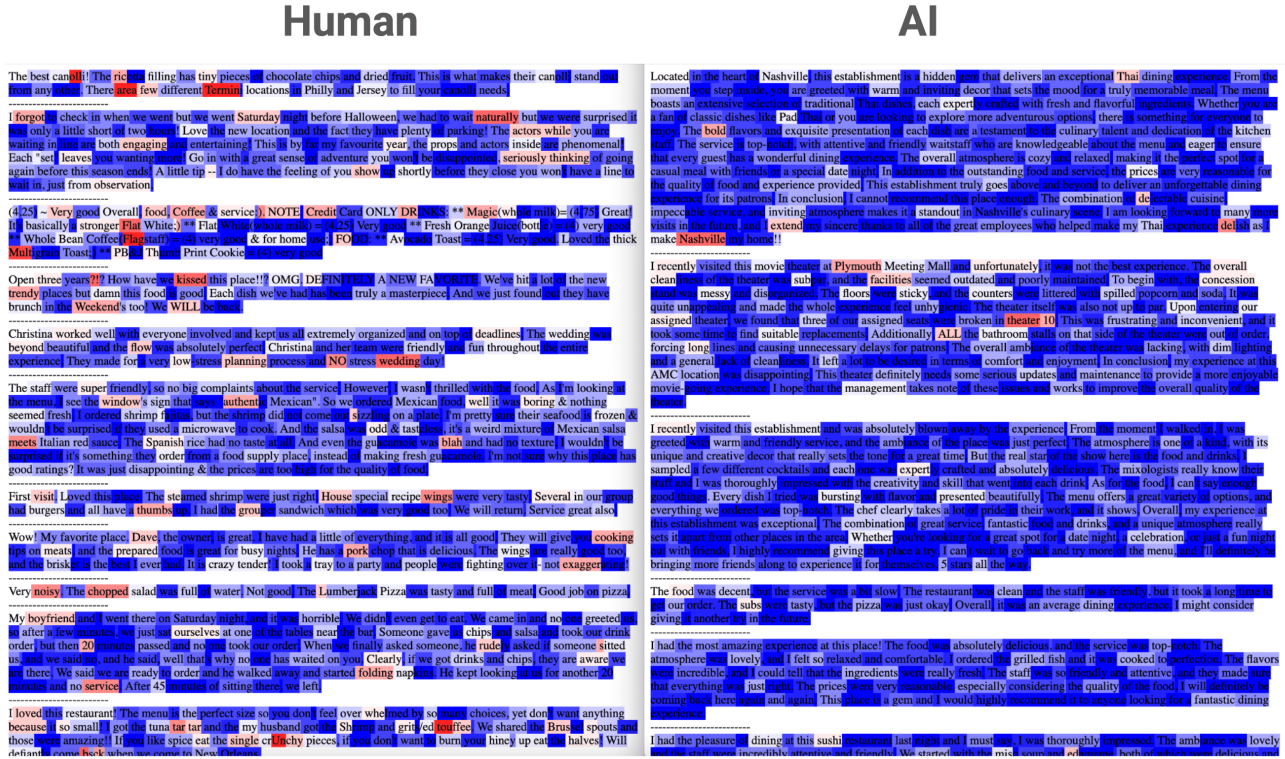

2023
2024
2025
2026
0.3 marker words per 100 words
0.4 marker words per 100 words
1.5 marker words per 100 wor
00 words
No text to analyze
marker words per 100 words: 0 · up to 1 · up to 2 · more than 2 (pre-2023 average 0.5)
MEMENTO - Platform for Real time urban experiences
M.S. Computational Design Practices capstone · 2025
106 marker words in 4930 words (2.2 per 100).
insights · in today's · fostered · not X but Y · enhance · reimagines · not X but Y · dynamic · interplay · interplay · evolving · reimagines · not only … but also · not X but Y · serves as a · not X but Y · dynamic · dynamic · dynamic · robust · enhances · dynamic · serves as a · dynamic · resonating · evolving · seamlessly · fostering · not X but Y · – (em dash) · – (em dash) · – (em dash)

Affordances of AI: Prediction



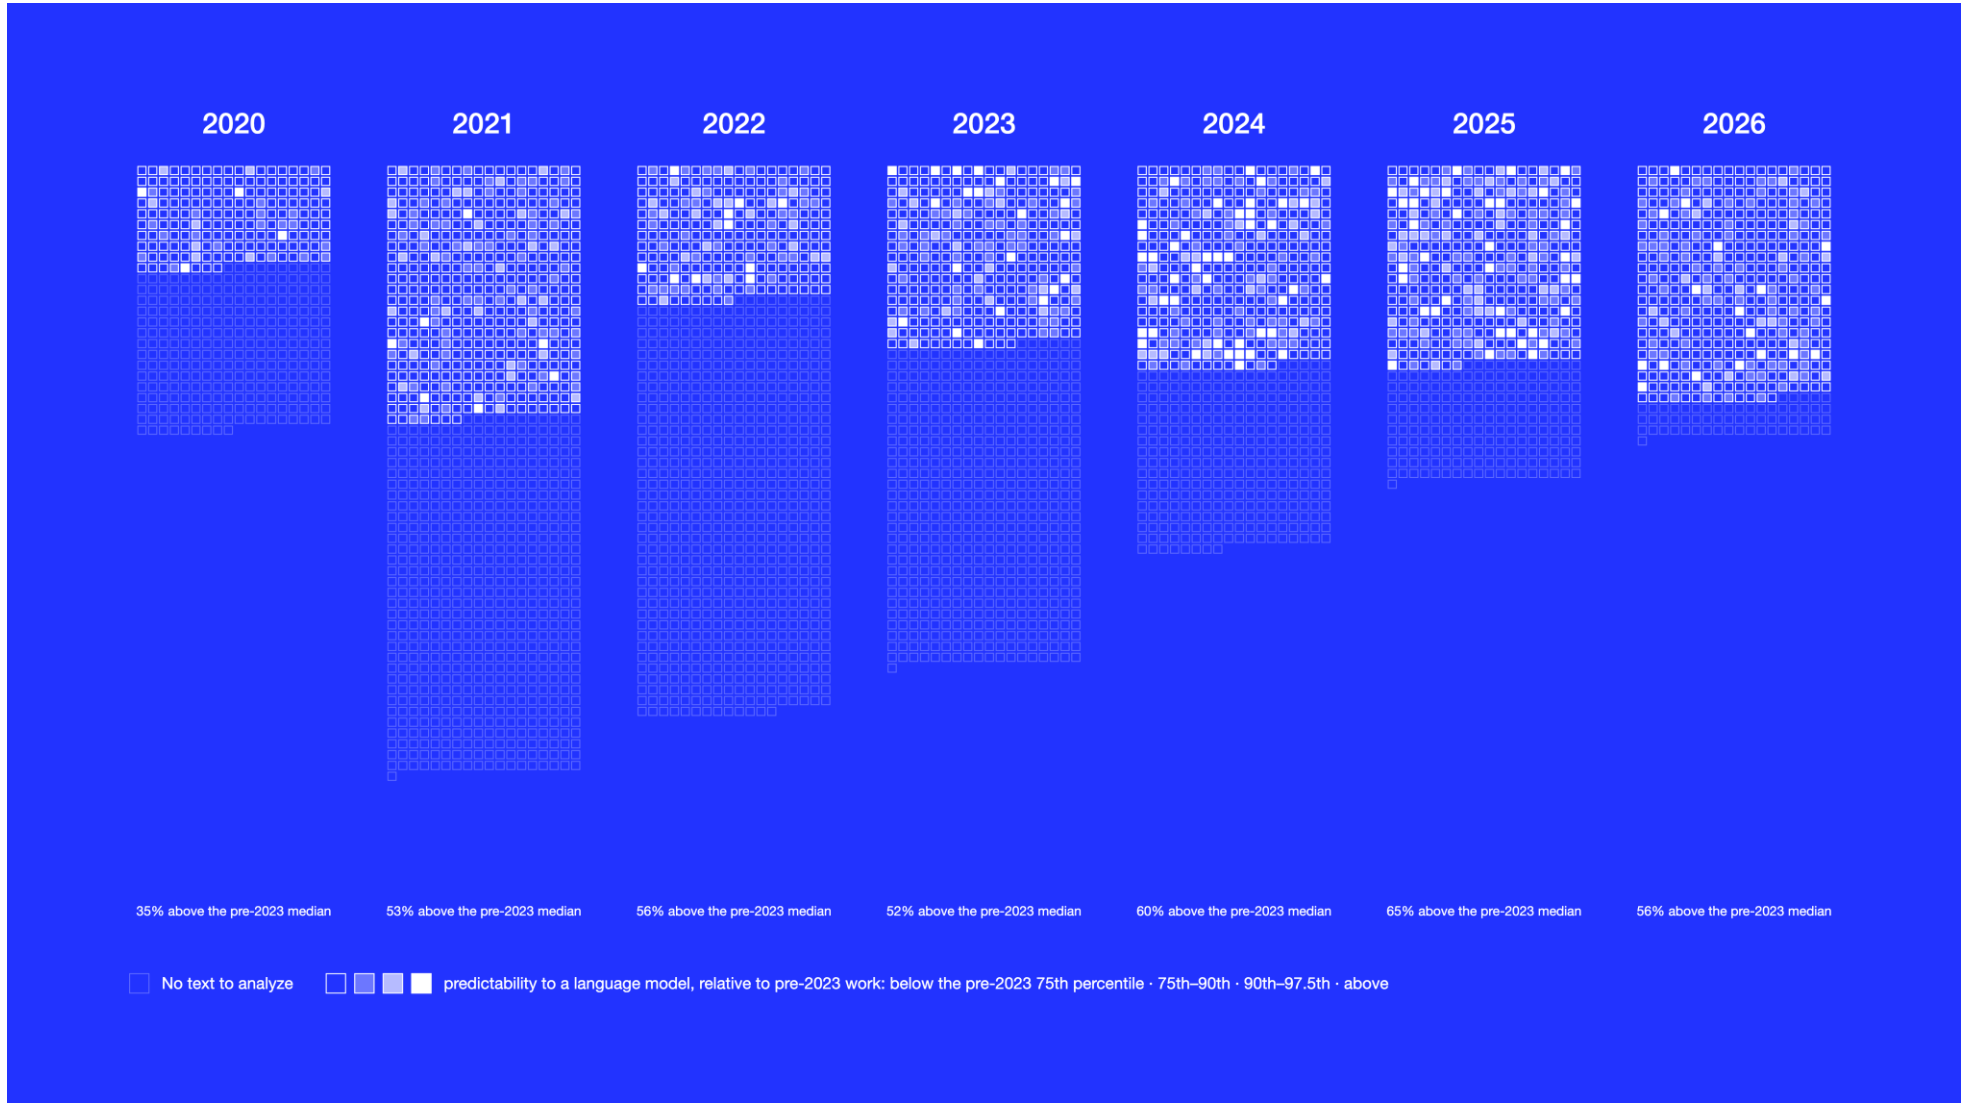
AI marker words in GSAPP EOYS work – 2020 to 2026

Affordances of AI: Perplexity



"perplexity" score of word-choice in a given text

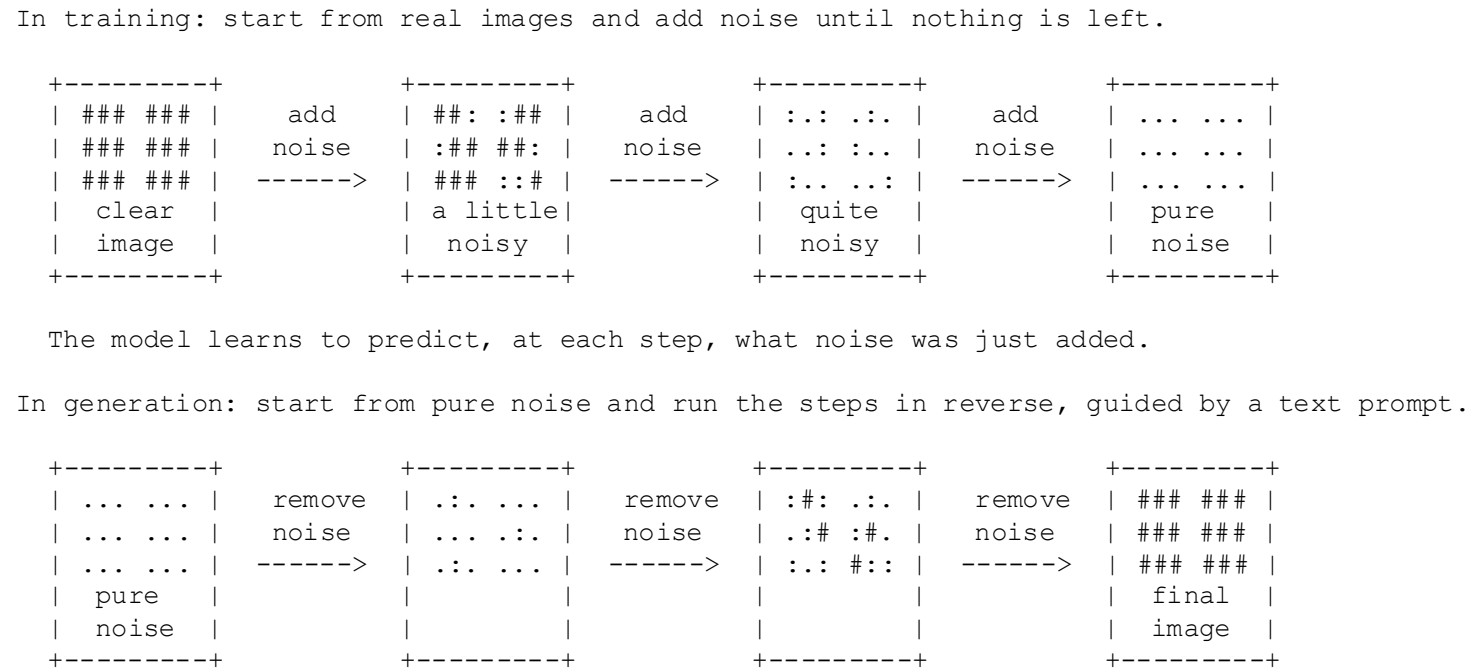
Affordances of AI: Perplexity



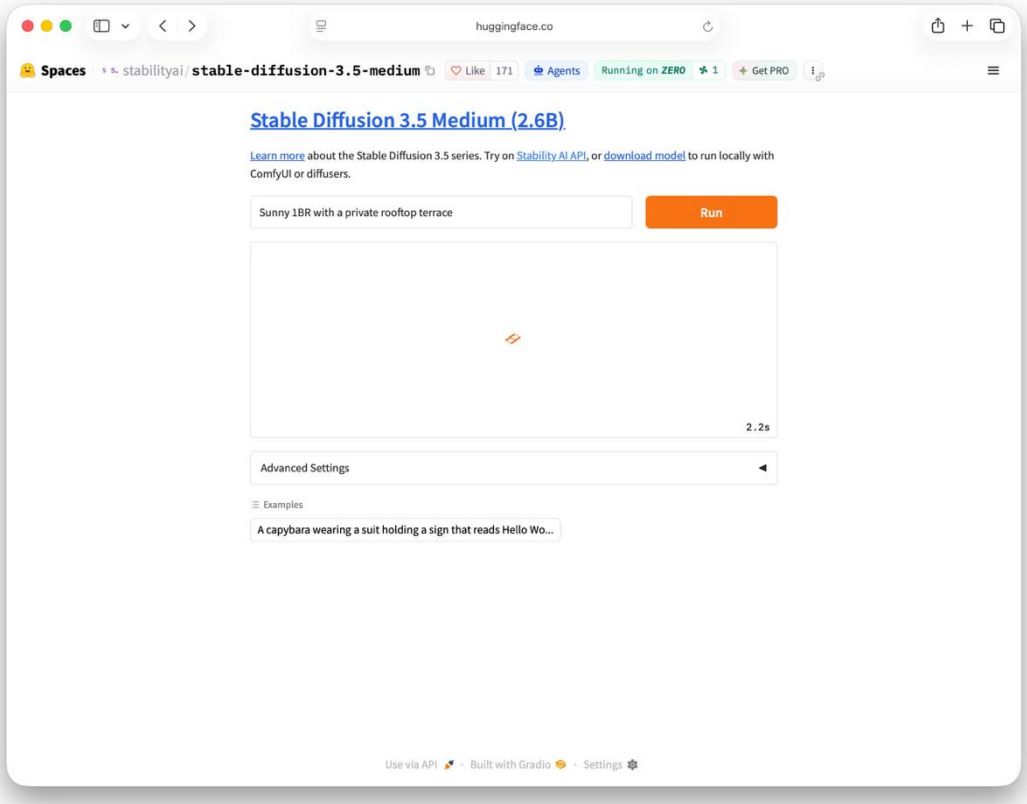
Perplexity in GSAPP EOYS work – 2020 to 2026. Results show text is more or less equally predictable pre- and post- consumer AI.

What an AI model does: Prediction in Diffusion Models

How diffusion models learn and generate images

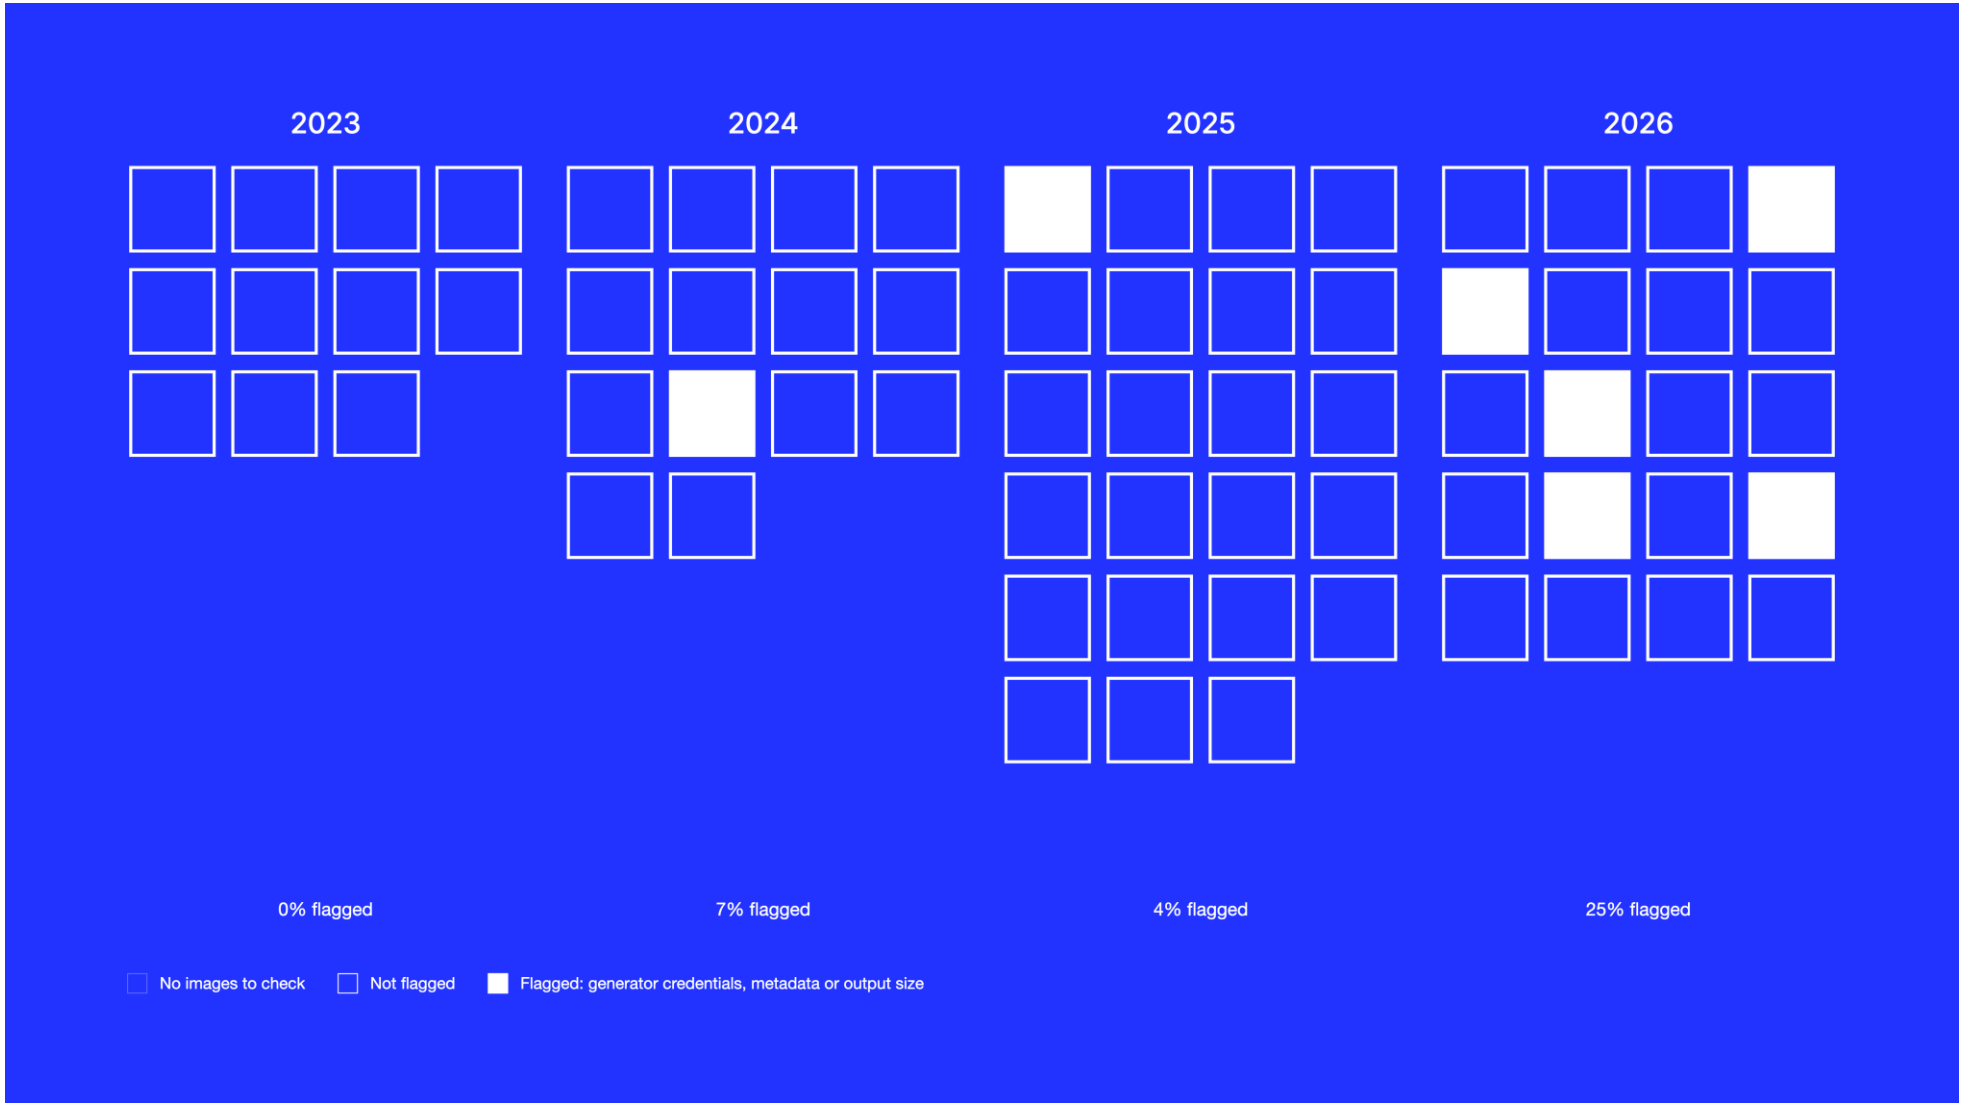


What an AI model does: Prediction in Diffusion Models



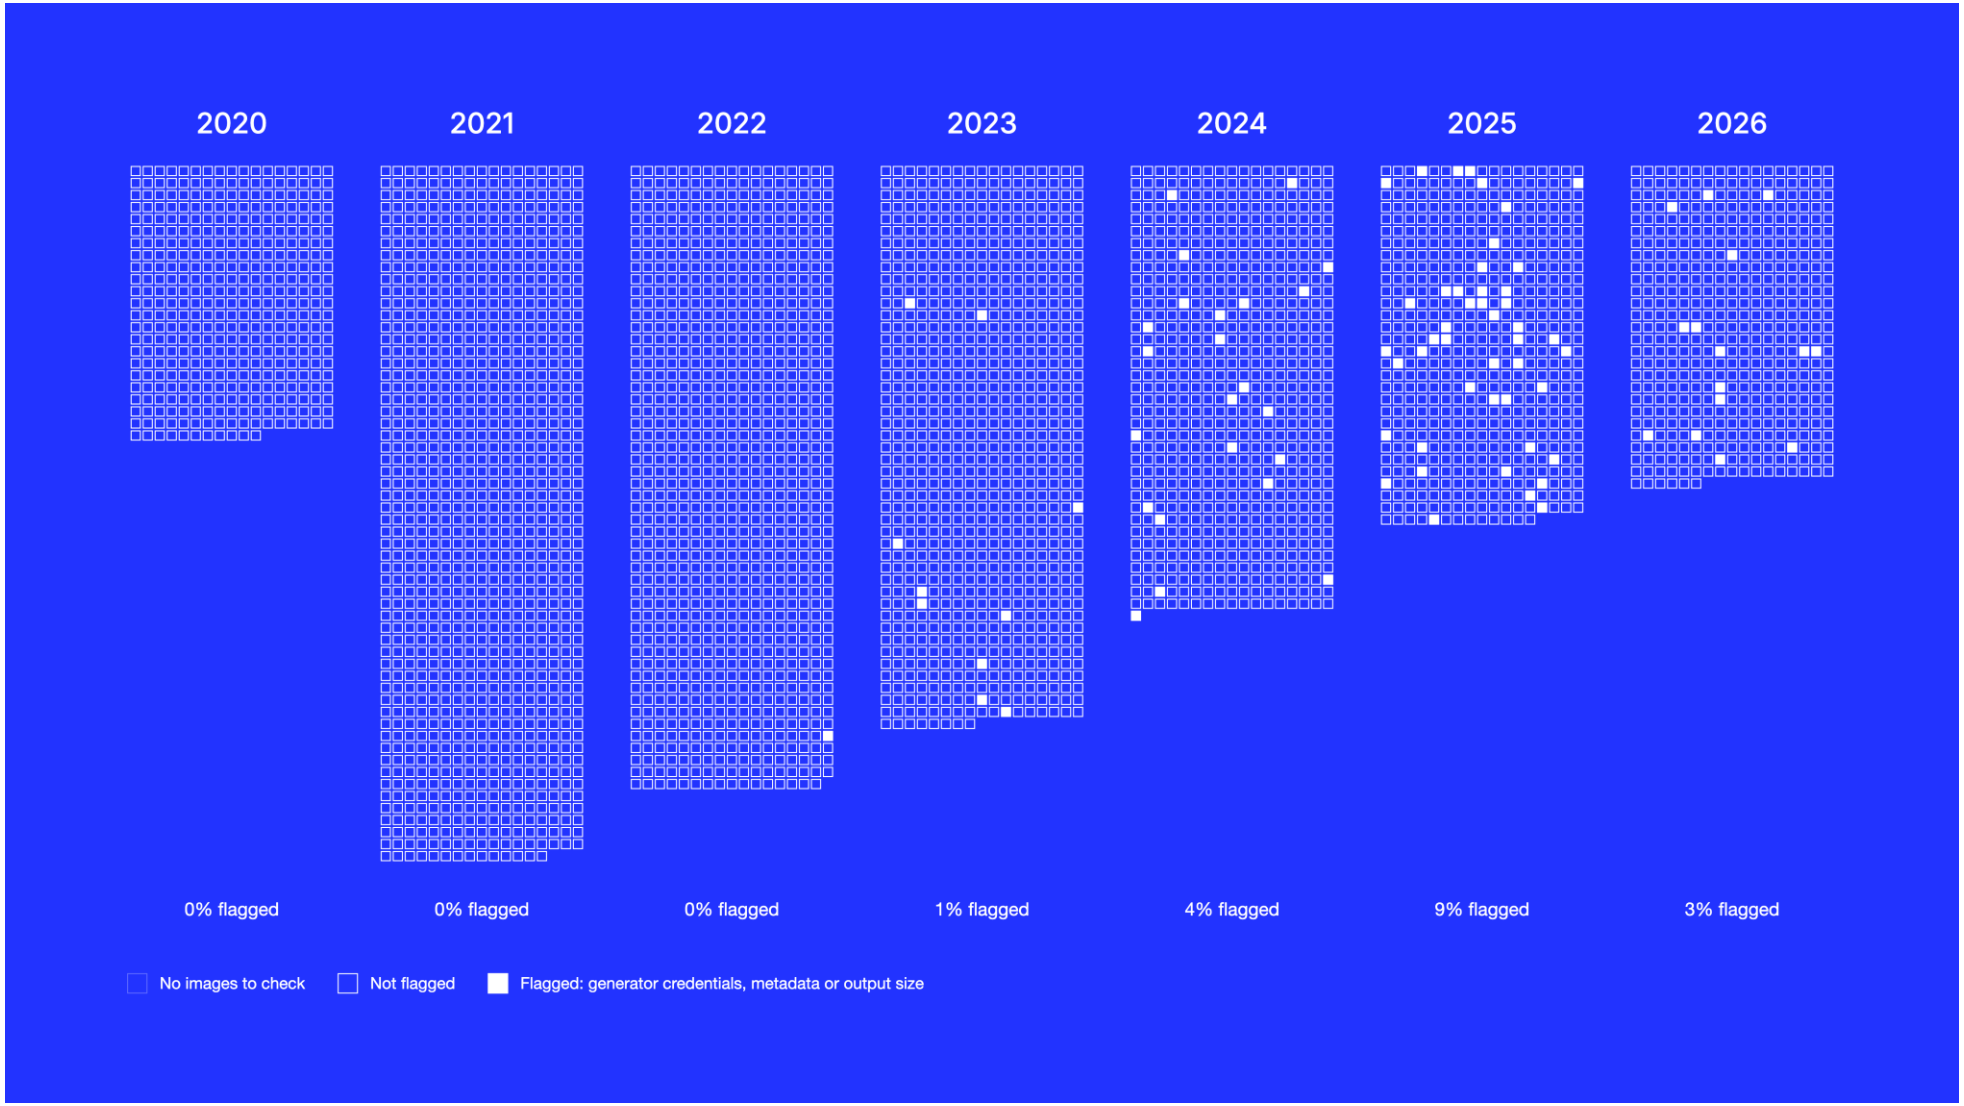
"Sunny 1BR with a private rooftop terrace" on stable diffusion 3.5

Image generation



Traces of AI tools in GSAPP CDP images – 2023 to 2026

Image generation



Traces of AI tools in GSAPP EOYS images – 2020 to 2026

Types of Models: Access and Parameters

Access

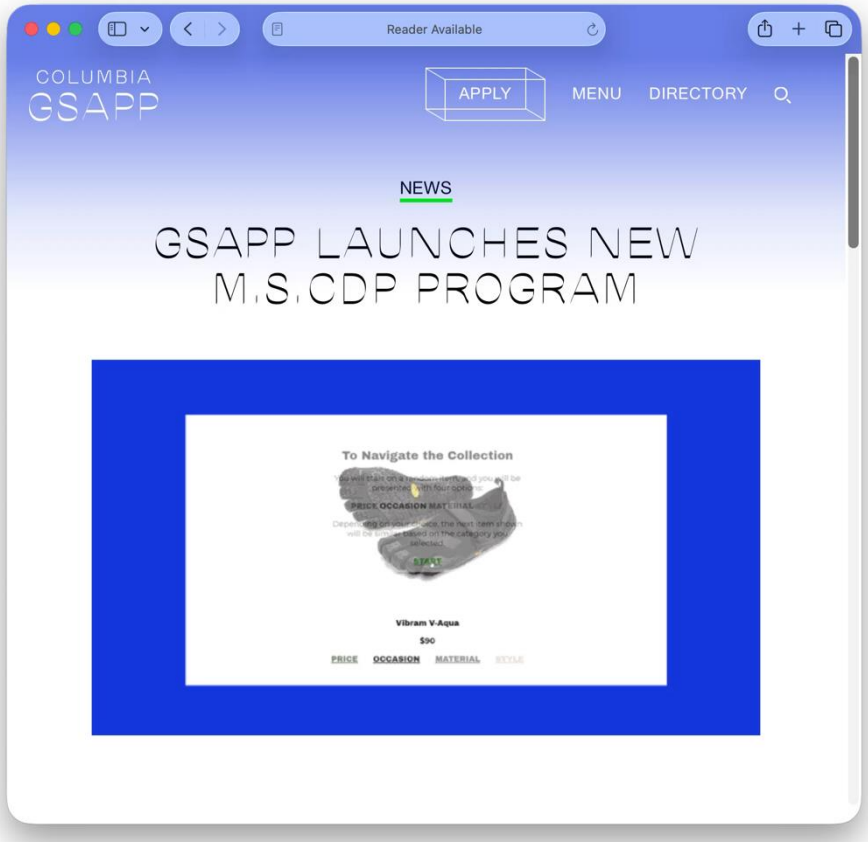
- **Frontier models** (Anthropic's Claude, OpenAI's ChatGPT, Google Gemini) — available only as a paid service, and expensive to train.
- **Open-weight models** (DeepSeek, Alibaba's Qwen, Moonshot's Kimi) — weights downloadable and runnable on your own hardware. Often much more efficient.
- **Open-source models** (Allen AI's OLMo, Nomic AI's Embed) — not common, but growing.

Parameters

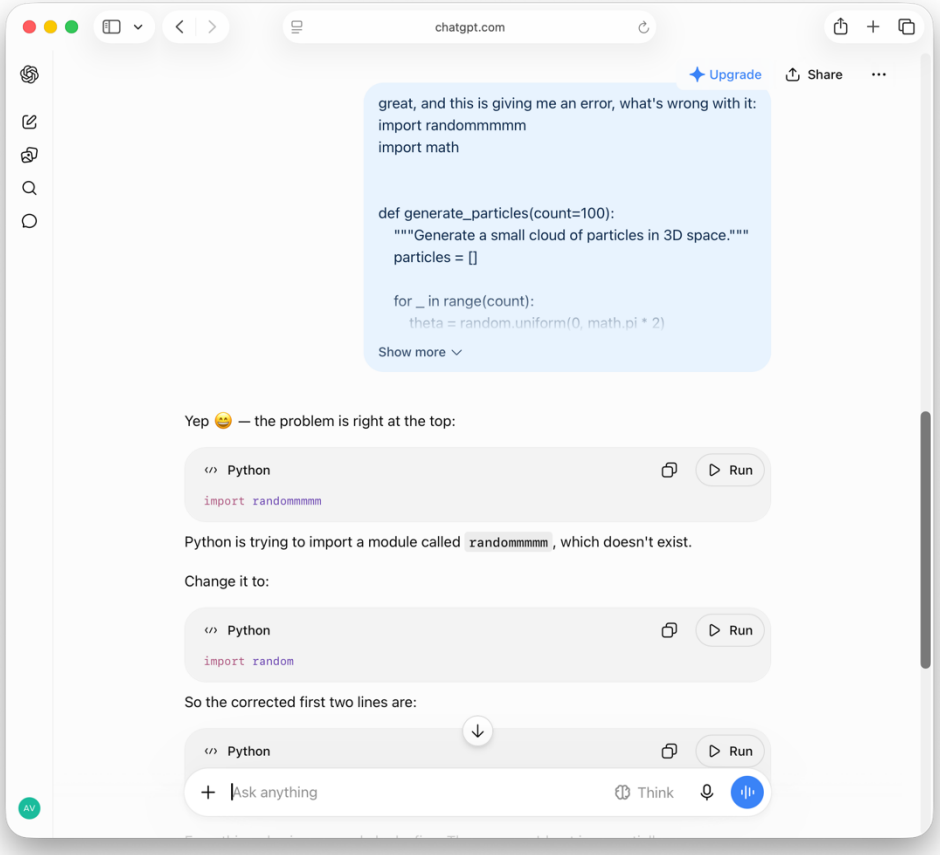
- Models often improve with size, measured in the billions of parameters.
- For example, Alibaba's Qwen3.5 comes in 8 variations from 0.8 billion to 397 billion parameters.
- The smaller models can easily run on most modern computers. The larger need a larger scale of compute than a personal computer.
- While it is just speculation, the largest "frontier" models out now (such as Claude 5.1 Fable, GPT-6 Astra) are estimated to be 1 to 10T parameters.

The view from CDP

The view from CDP

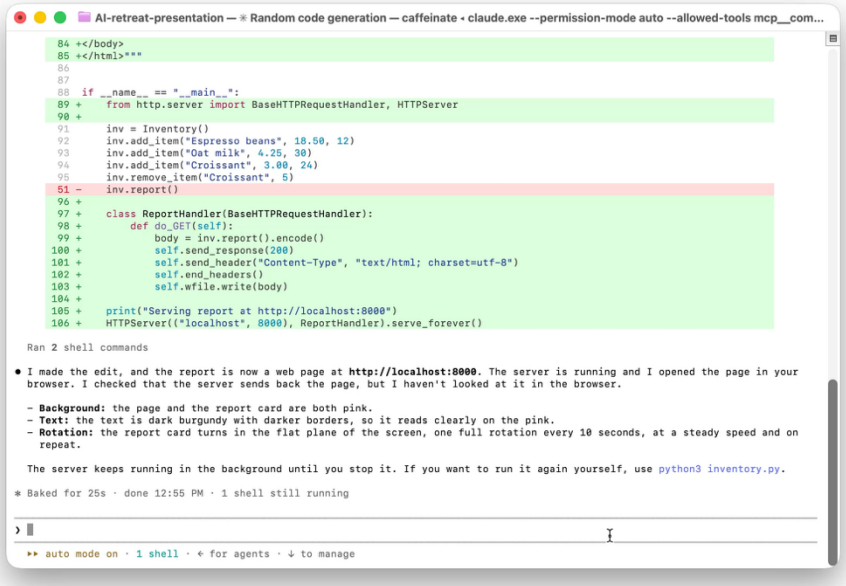


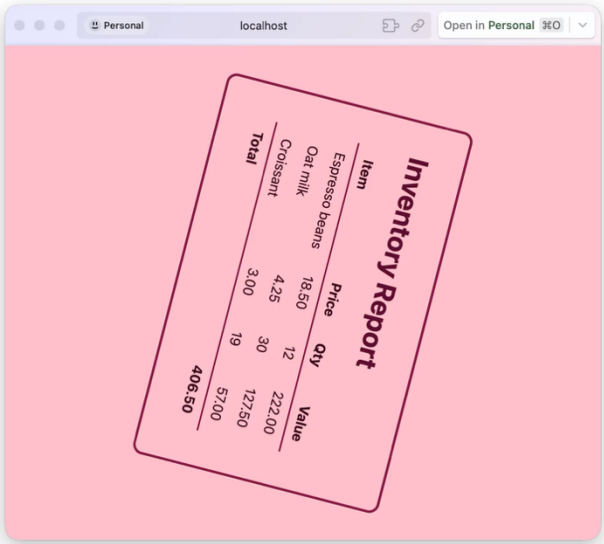
CDP launched July 2022



Chat GPT launched November 2022

The view from CDP

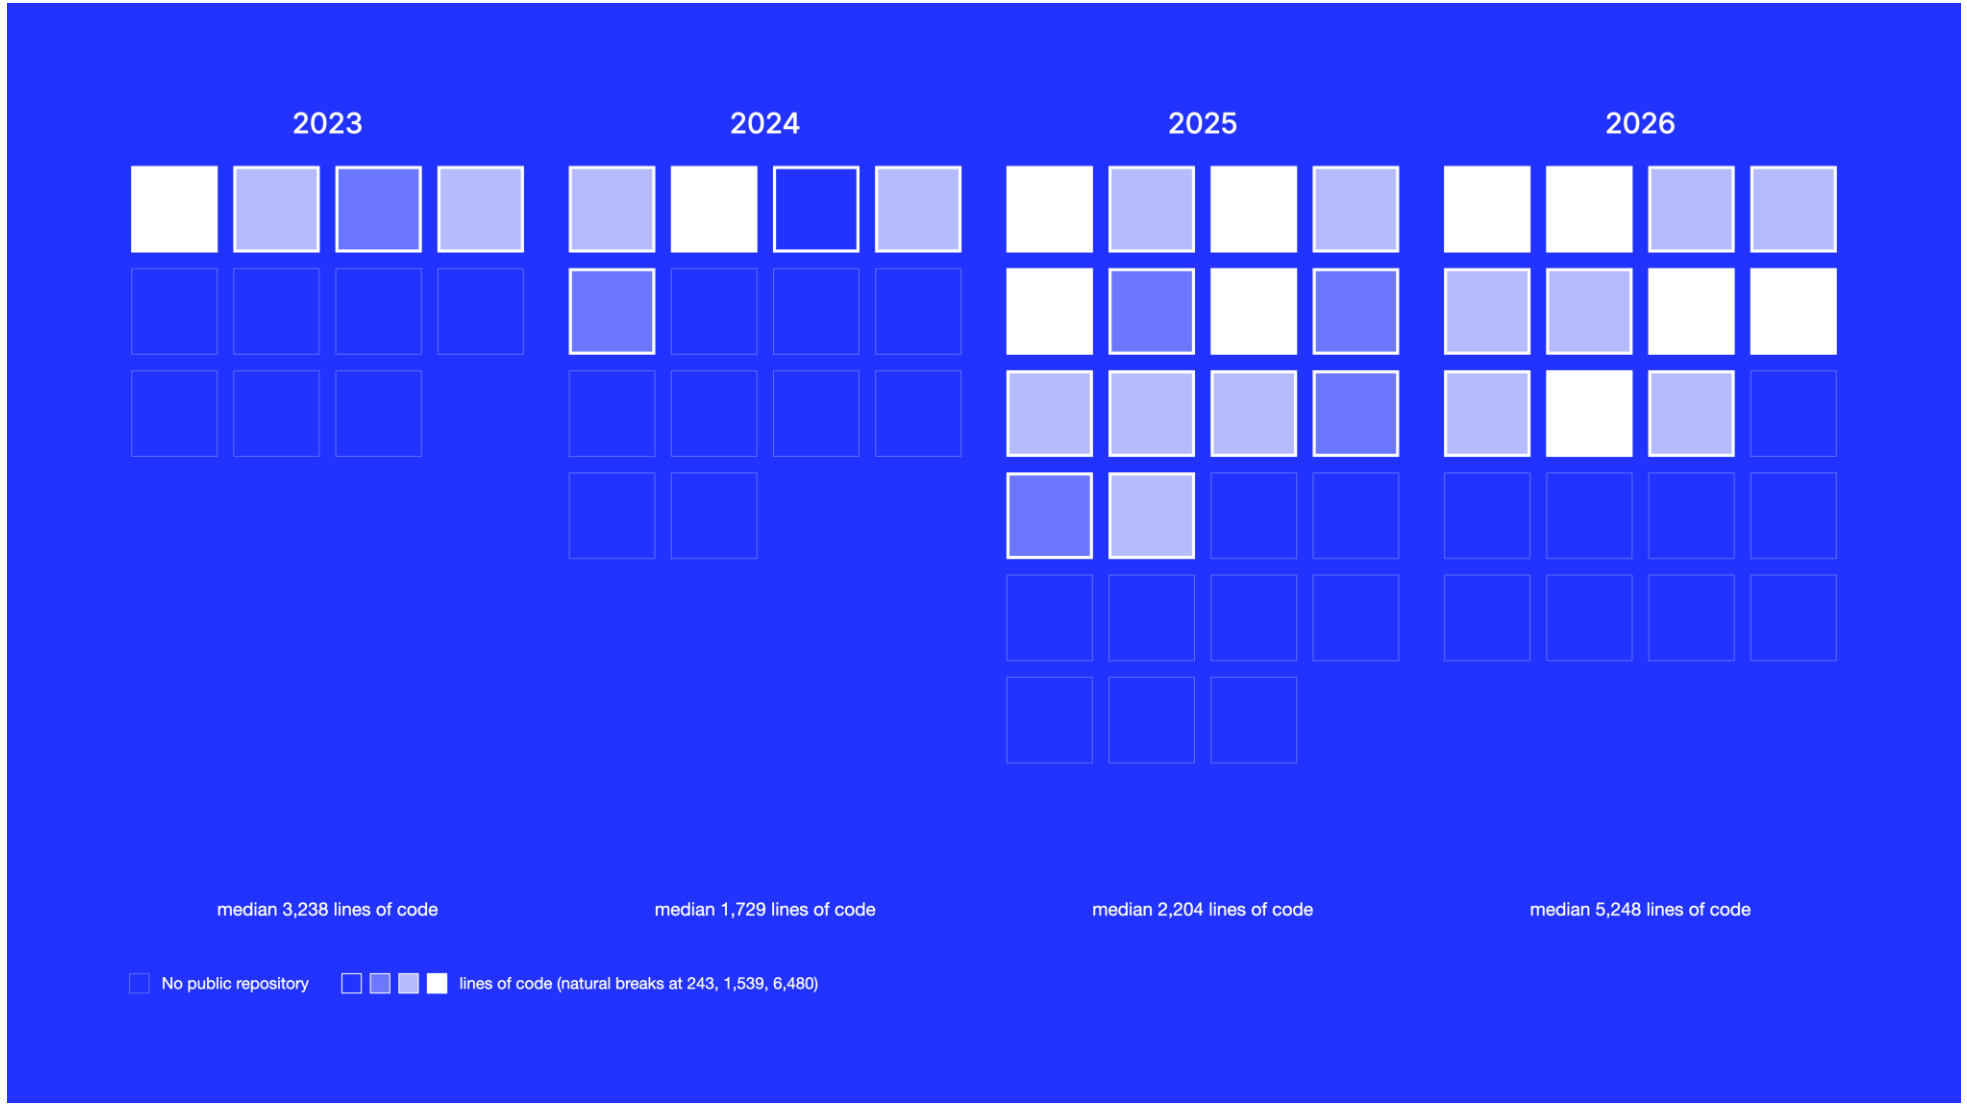




Coding "agents" like Claude code run an AI model in a sandbox environment where they can edit files and run commands on your computer.

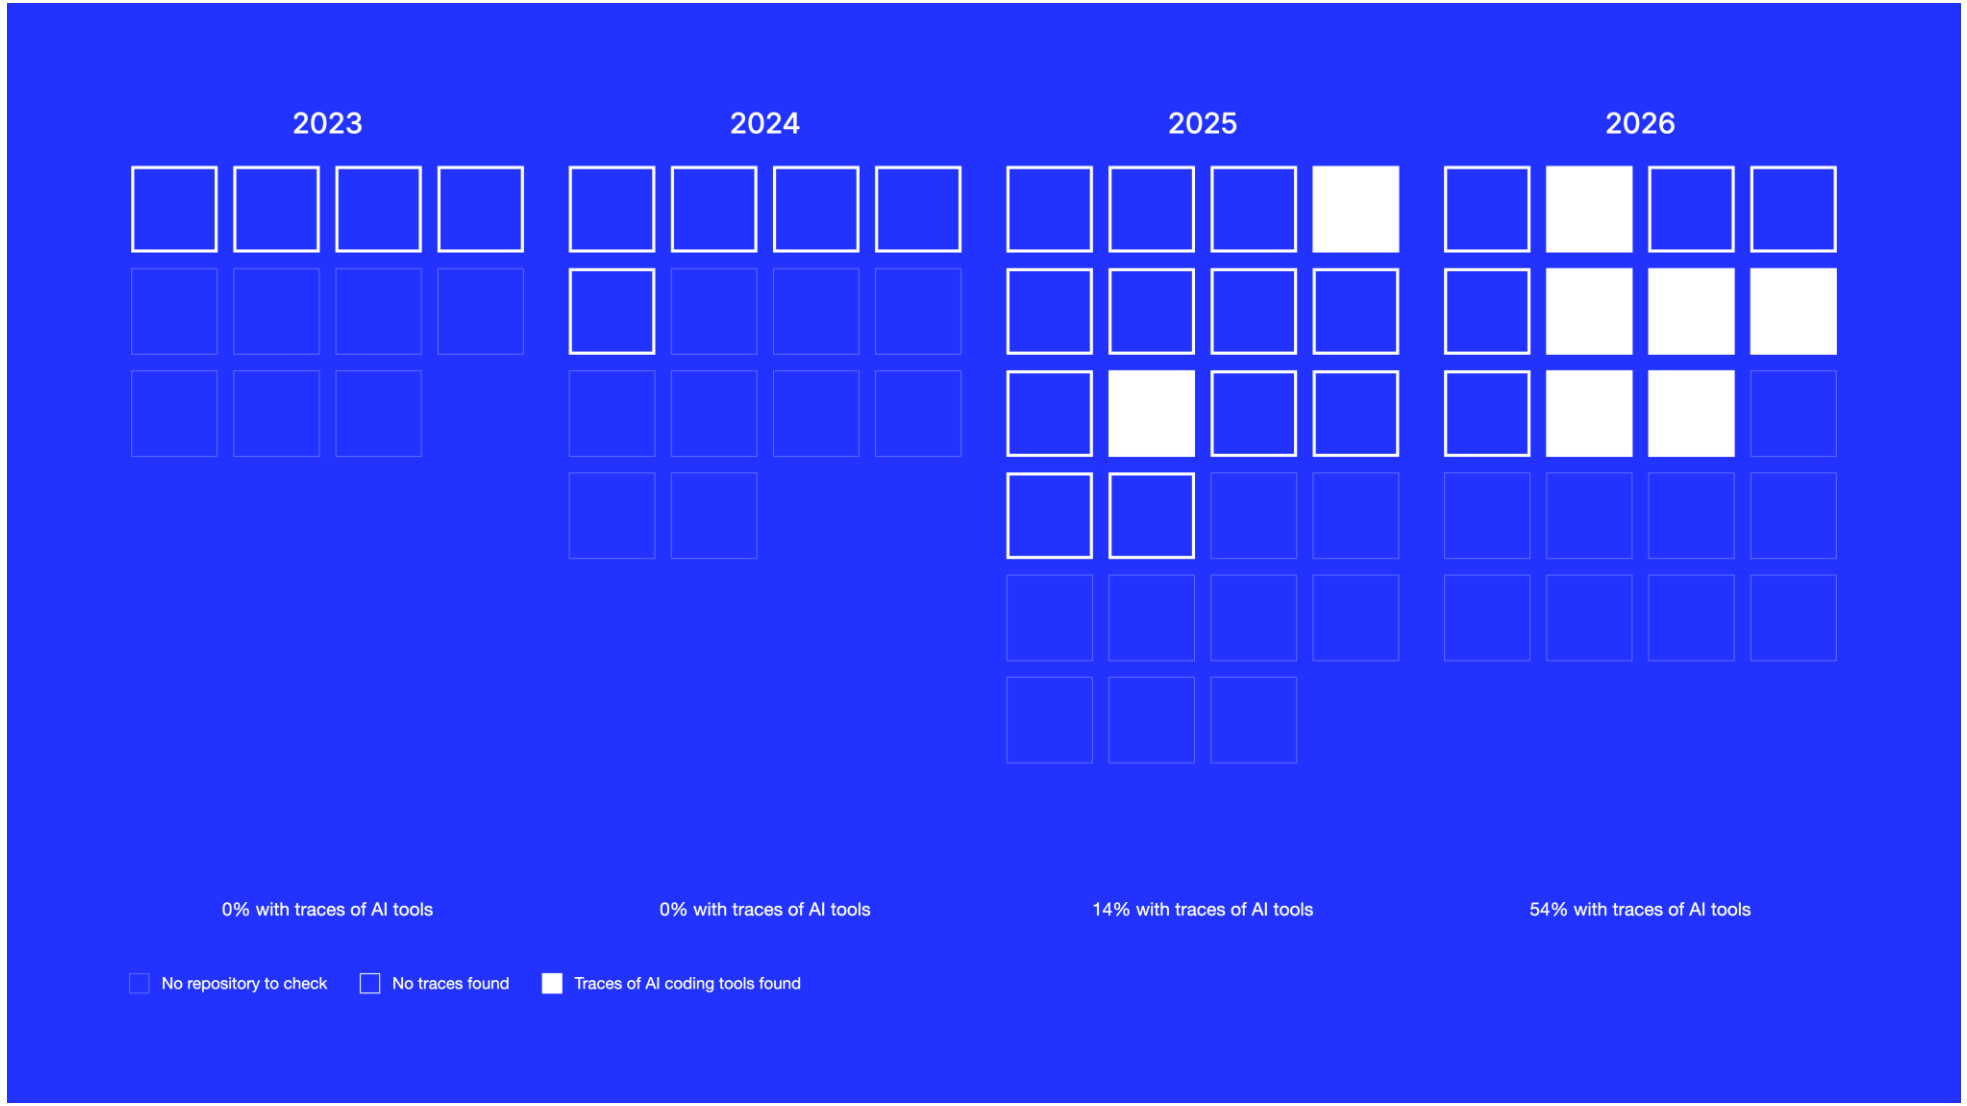
*You may have heard of "**Vibe coding**": a software development practice where a person builds apps or websites by giving natural language prompts to an AI model rather than writing code manually.*

The view from CDP



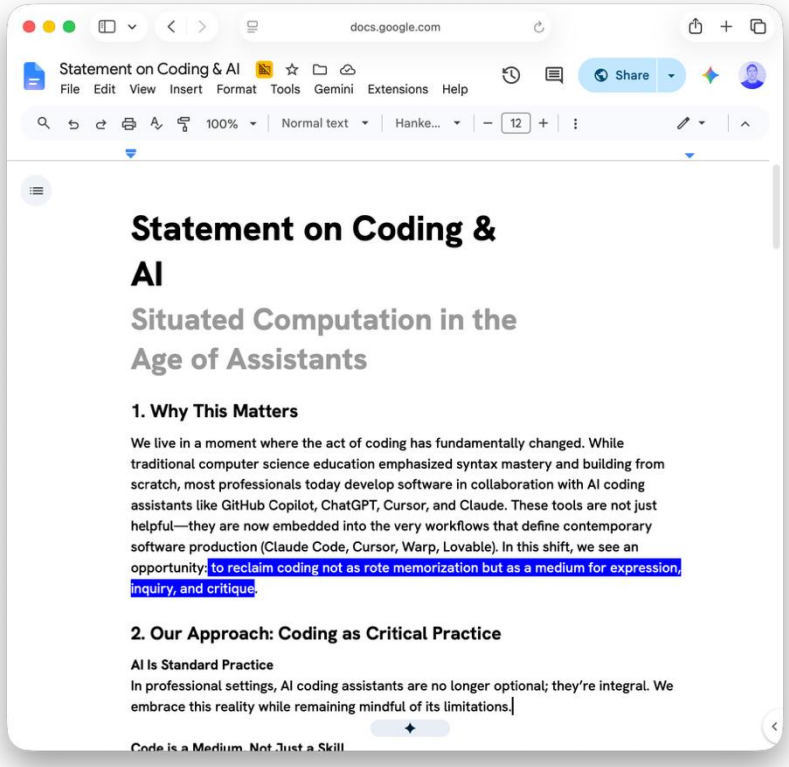
Lines of code in GSAPP CDP work – 2023 to 2026

The view from CDP



Traces of AI tools in GSAPP CDP code – 2023 to 2026

The view from CDP



CDP Statement on Coding + AI. 2025.

The view from CDP



Simulations, Models, Twins. http://simmodeltwin.net

The view from CDP

Takeaways from CDP Five-Year Review

Summarized responses to AI questions, from interviews with 16 key alumni

Q: "Have AI tools changed your practice or work significantly?"

yes, significantly	■■■■■■■■■■	15 / 16	94%
no, avoids them	■	1 / 16	6%

Q: "Fundamentals first, AI extends ambition, emphasis on the ideas."

Does that resonate? How would you change the program?

agree	■■■■■■■■	12 / 16	sounds good!
agree, with worry	■	2 / 16	sounds good, but are we deskilling? is vibe coding a hireable skill?
other	■	2 / 16	one unsure; one would refocus around AI

Questions?

References, visualizations, slide pdf and further reading at ***retreat.adamvosburgh.com***
password: ***gsapp-ai***