

Artificial Intelligence (AI) in Healthcare

National Committee on Vital and Health Statistics

11/30/2023

Travis Hoppe, PhD, thoppe@cdc.gov

Associate Director of Data Science and Analytics, Office of the Director

National Center for Health Statistics,
Centers for Disease Control,
Health and Human Services (HHS/CDC/NCHS)

Introduction: About Travis Hoppe



- Led CDC Tiger Team, "Utility and Risks to CDC of Conversational Artificial Intelligence (AI) Technologies Like ChatGPT"
- Subject matter expert in AI/ML. Published open-source models and research publications in natural language processing (NLP) including:
 - Topic choice contributes to the lower rate of NIH awards to AA/black scientists ([Science](#))
 - Semi-Automated Non-response Detection for Surveys ([NCHS AI model publicly released](#))
 - The Pile: An 800GB Dataset of Diverse Text for Language Modeling ([arXiv](#))
- CDC representative for National Science and Technology Council (NSTC) subcommittee on AI/ML
- Member of HHS AI CoP, FCSM, DMI workgroups for AI/ML
- PhD in Physics with background in biochemistry and biophysics; previously at NIH

Definitions of AI, ML, and generative AI

AI: Interdisciplinary field, usually regarded as a branch of computer science, dealing with models and systems for the performance of functions generally associated with human intelligence, such as reasoning and learning. *

Machine Learning: The study of computer algorithms that improve automatically through experience.*

Generative AI: Class of AI systems that, after being trained on large data sets, can be used to generate text, images, videos or other outputs from a given prompt.

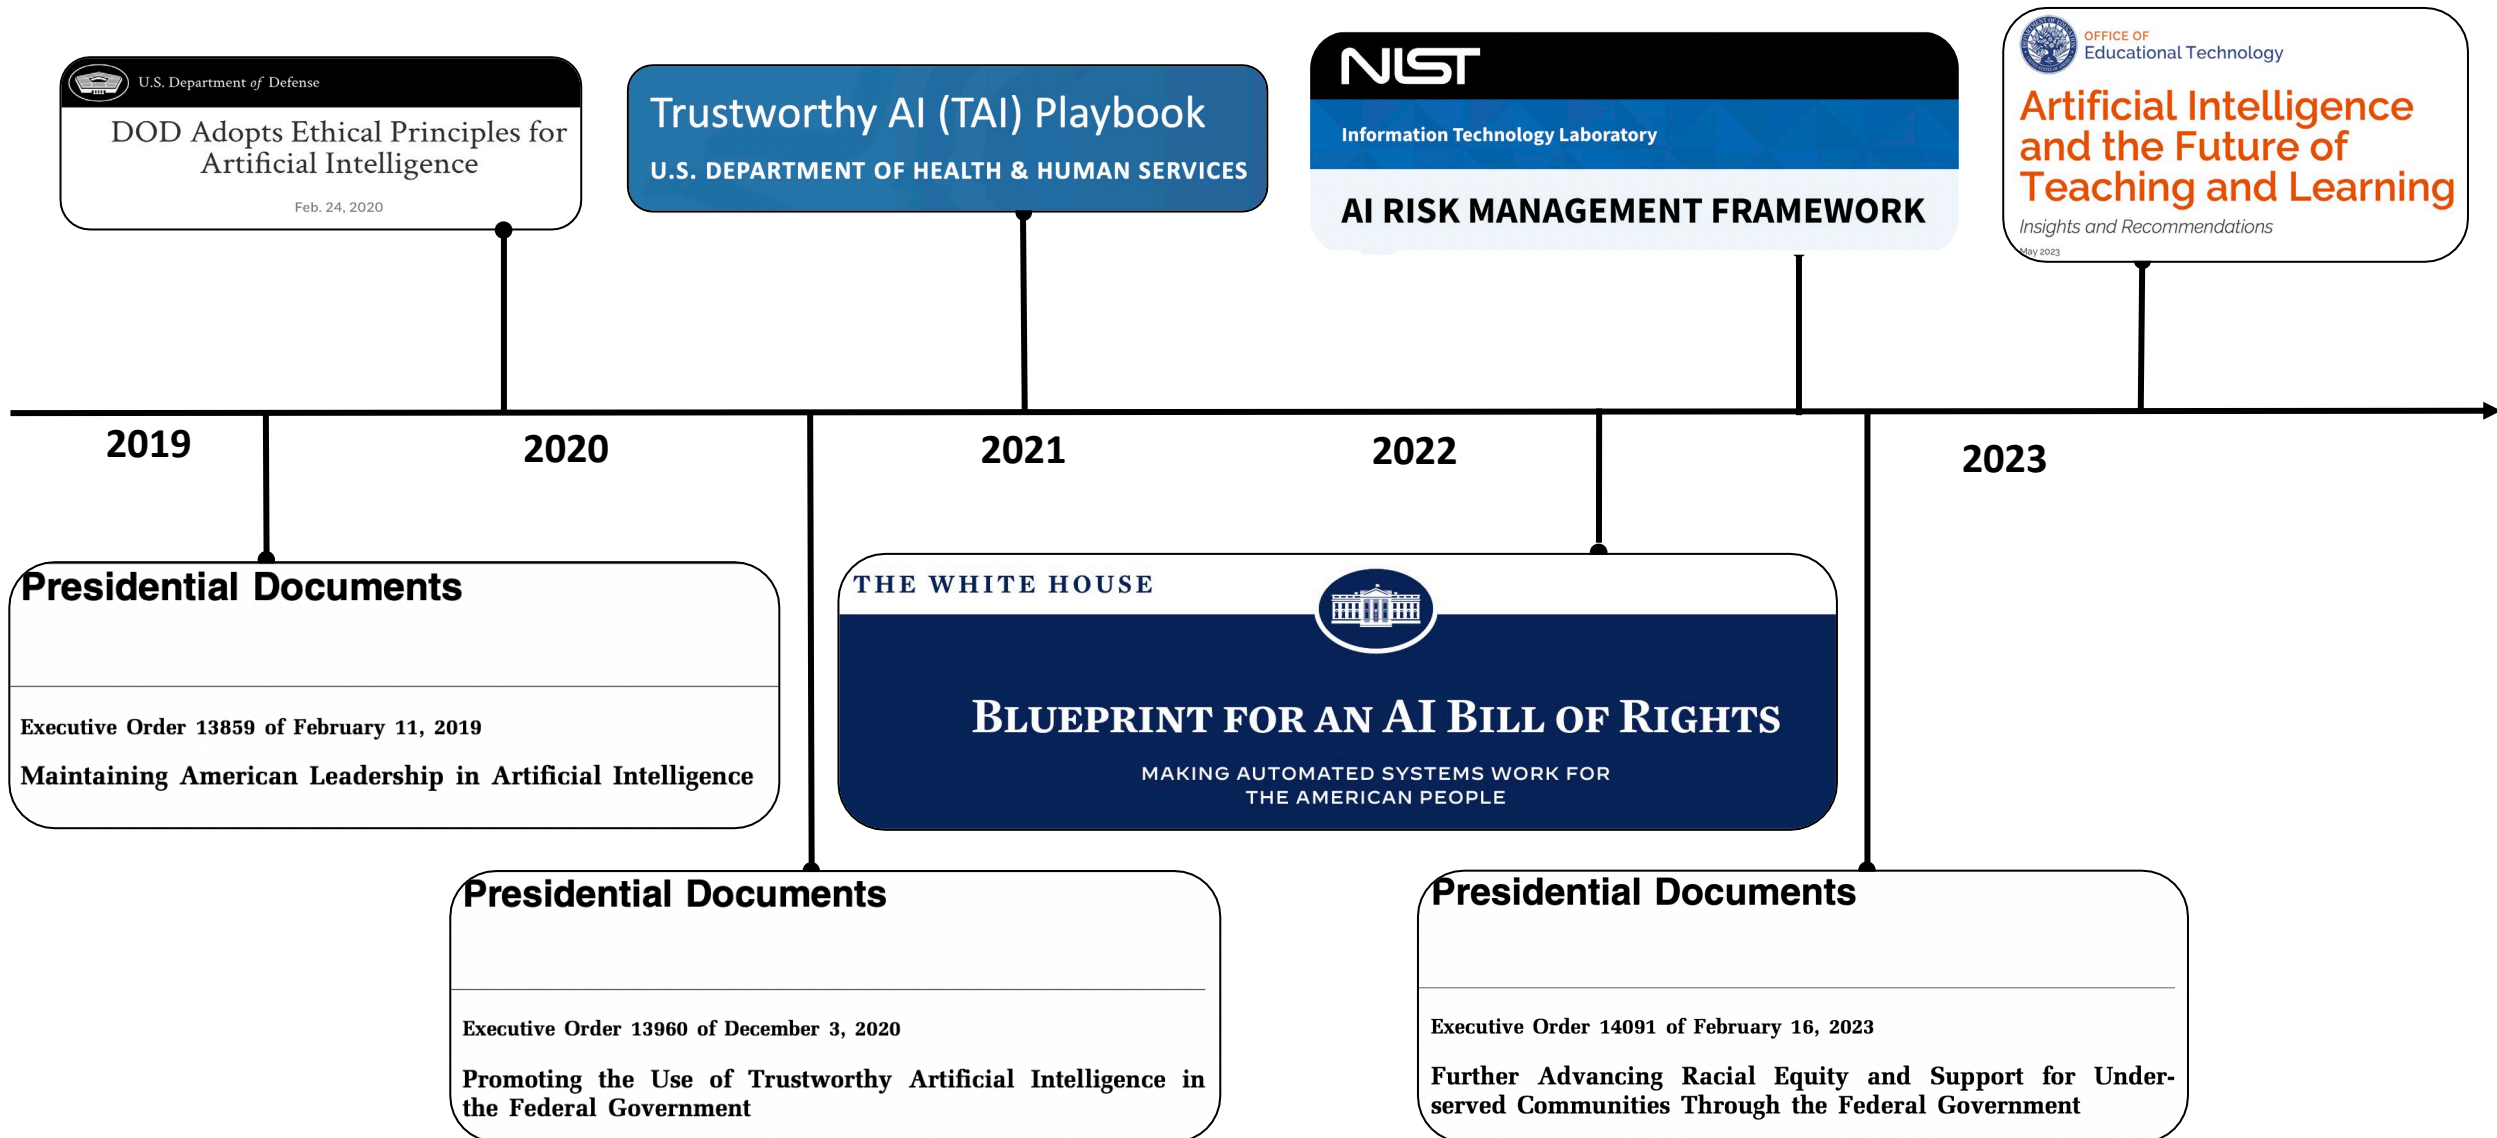


The Treachery of Images, René Magritte

* [National Institute of Standards and Technology Language of Trustworthy AI: In-Depth Glossary of Terms](#)

* <https://www.whitehouse.gov/pcast/briefing-room/2023/05/13/pcast-working-group-on-generative-ai-invites-public-input>

Trustworthy AI initiatives in the federal government



Recent AI initiatives in the federal government

Strengthening and Democratizing
the U.S. Artificial Intelligence
Innovation Ecosystem

*An Implementation Plan for a
National Artificial Intelligence Research Resource*

Jan.

2023

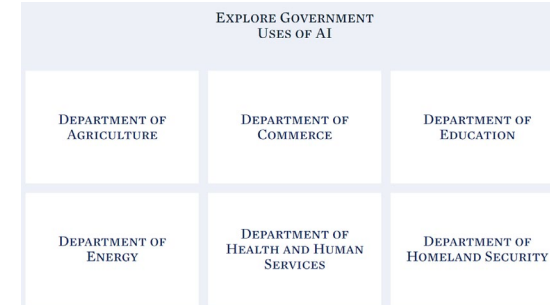
May

Administration Takes New
Steps to Advance Responsible
Artificial Intelligence
Research, Development,
and Deployment

PCAST Working Group on
Generative AI Invites
Public Input



Aug.



Oct.

Sept.

Administration Secures
Voluntary Commitments
from Eight Additional
Artificial Intelligence
Companies to Manage the
Risks Posed by AI

Safe, Secure, and
Trustworthy
Development and Use
of Artificial Intelligence
[EO 14110](#)

Nov-Dec

Advancing Governance,
Innovation, and Risk
Management for
Agency Use of
Artificial Intelligence
[Draft OMB Memo](#)



NCHS AI activities, per EO 13960 (Trustworthy AI)

Developed in house

- [SANDS](#): Semi-automated non-response detection, live model release
- [MedCoder](#): Coding cause of death information to ICD-10
- Detecting Stimulant and Opioid Misuse and Illicit Use in EHRs

Commercial-off-the-shelf or open-source implementation

- Whisper: Speech to text (CCQDER) improvement, 20x improvement
- Private AI: PII detection model used for FOIA requests

Governance, policy, and strategic planning

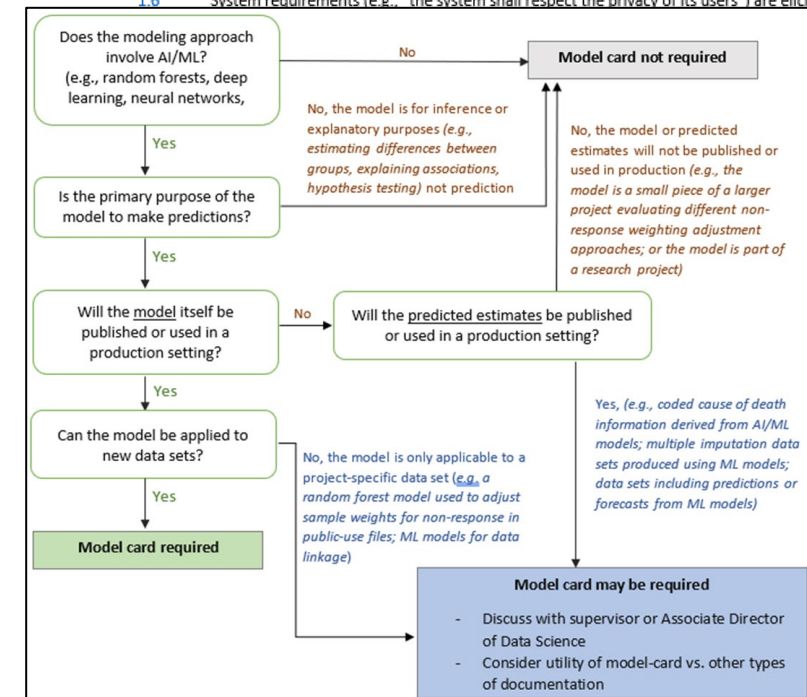
- Utility and Risks to CDC of Conversational Artificial Intelligence (AI) Technologies Like ChatGPT

NCHS innovation and standards in AI/ML

Developed NCHS standards for *model release*

- Standards require the creation of a *model card*, a kind of “product sheet” that tells how the model was trained and what use is appropriate
- Model cards document risks, known biases, context for use, and help build trust
- Standards aligned with NIST AI RMF through extensive crosswalk
- Standards to be incorporated into a larger CDC AI strategy (organized through EDO)

1.2	Inter-disciplinary AI actors, competencies, skills and capacities for establishing context reflect demographic diversity and broad domain and user experience expertise, and their participation is documented. Opportunities for interdisciplinary collaboration are prioritized.	Development
1.3	The organization's mission and relevant goals for the AI technology are understood and documented.	Release
1.4	The business value or context of business use has been clearly defined or – in the case of assessing existing AI systems – re-evaluated.	Development/ Production
1.5	Organizational risk tolerances are determined and documented.	Development/ Production
1.6	System requirements (e.g., “the system shall respect the privacy of its users”) are elicited from	Development



s into	Release	1.1, 2.1
port is	Release	1.1, 2.1
ized and assist	Release	2.3, 2.4, 2.5, 2.6
s are tion and	Release	1.1, 1.2, 1.3, 3.1, 4.1
nd	Release	2.1
l errors nce - are	Development	

NCHS AI/ML model release standards are already in use

Semi-Automated Non-response Detection for Surveys

[Print](#)

Model Details

This model is a fine-tuned version of the supervised SimCSE BERT base uncased model. It was introduced at [AAPOR](#) 2022 at the talk *Toward a Semi-automated item nonresponse detector model for open-response data*. The model is uncased, so it treats `important`, `Important`, and `ImPoRtAnT` the same.

Risk, Limitations, and Biases

To investigate if there were differences between demographic groups on sensitivity and specificity, we conducted two-tailed Z-tests across demographic groups. These included education (some college or less and bachelor's or more), sex (male or female), mode (computer or telephone), race and ethnicity (non-Hispanic White, non-Hispanic Black, Hispanic, and all others who are non-Hispanic), and age (18-29, 30-44, 45-59, and 60+). There were 4,813 responses to 3 probes. To control for family-wise error rate, the Bonferroni correction was applied to the alpha level ($\alpha < 0.00167$).

<https://www.cdc.gov/nchs/data-science/SANDS-model-context.htm>

Misuses and Out-of-scope Use

The model has been trained to specifically identify survey non-response in open ended responses where the respondent taking the survey has given a response but their answer does not respond to the question at hand or providing any meaningful insight. Some examples of these types of responses are `meow`, `ksdhfkshgk`, or `idk`. The model was fine-tuned on 3,000 labeled open-ended responses to web probes on questions relating to the COVID-19 pandemic gathered from the [Research and Development Survey or RANDS](#) conducted by the Division of Research and Methodology at the National Center for Health Statistics. Web probes are questions implementing probing techniques from cognitive interviewing for use in survey question design and are different than traditional open-ended survey questions. The context of our labeled responses limited in focus on both COVID and health responses, so responses outside this scope may notice a drop in performance.

CDC's proposed Artificial Intelligence roadmap



Strategy

- Outline strategies and objectives specific to use of AI in public health.
- Guide strategic investments in computing / cloud infrastructure, data, and workforce upskilling.



Community of Practice

- Elevate AI use-cases at CDC and advocate for practitioner needs.
- Educate and empower programs to leverage AI for better public health outcomes.
- Define best practices and upholds standards for responsible and trustworthy AI within CDC.

CoP includes:

- AI Workstreams
- CDC Enterprise staff
- Centers and programs
- State and local CDC partners



Guidance

- Provide guidance for the responsible development and use of generative AI at CDC to support the public health mission.



Communication

- Establish and ensure strategic communication on AI is timely, accurate, and meet the needs of internal and external audiences.
- Develop internal and external communication products to support awareness, education, and information sharing.



AI/ML Consultation Group ("AMC")

- Advise on developing and deploying new use cases leveraging machine learning (ML) and AI to further public health outcomes.
- Assist in the identification of promising use-cases that align with the mission of different centers/programs.
- Develop guidance for implementation of AI and ML models.



Investments and Infrastructure

- Increase technical capabilities within CDC for technical teams in centers and programs to deploy AI and ML models at scale.
- Strengthen data pipelines from STLT partners.

Reminder: Generative will happily report falsehoods



Can you draft an MMWR "Notes from the Field" report about an outbreak of sillivirus occurring on April 22, 2023, with 11 cases in 60 U.S. states? 




Sure, here's a draft of an MMWR "Notes from the Field" report about the sillivirus outbreak:



"Notes from the Field: Sillivirus Outbreak Across Multiple U.S. States - April 22, 2023



On April 22, 2023, the Centers for Disease Control and Prevention (CDC) was alerted to an outbreak of sillivirus affecting multiple U.S. states. As of this report, a total of 11 confirmed cases have been identified in 60 states. 

Reminder: Generative AI can create fake references

“Generate a sample set of references from a CDC report from the National Center for Injury Prevention and Control that are hot takes about drug overdoses on Xylazine mixed with Fentanyl.”

1. Smith, J. (2023). Unveiling the Deadly Duo: Xylazine and Fentanyl Cocktail in Drug Overdoses. CDC Report No. 456789.
2. Johnson, A. (2023). Beyond Lethal: The Alarming Rise of Xylazine and Fentanyl Combinations. National Center for Injury Prevention and Control Research Brief, Issue 127.
3. Thompson, M., & Rodriguez, L. (2023). Unprecedented Danger: Exploring the Synergistic Effects of Xylazine and Fentanyl in Drug Overdose Cases. Journal of Substance Abuse and Public Health, 18(3), 45-59.
4. Brown, R., & Davis, S. (2023). Xylazine-Fentanyl Mix: The New Wave of Overdose Tragedies. In Annual Meeting of the American Association of Addiction Medicine (AAAM), Atlanta, GA.
5. Patel, K., & Williams, E. (2023). Deadly Concoction: Xylazine and Fentanyl—A Deadly Combination Behind the Opioid Crisis. Journal of Drug Abuse Awareness, 7(2), 112-129.

Of note, Xylazine was not recognized as a threat when the model was trained in September 2021.

Reminder: Generative AI has a limited knowledge window

“Tell me about the rise in (Fentanyl / Xylazine) deaths?”

The rise in fentanyl-related deaths is a serious concern and has been considered a public health crisis in many countries, including the United States and Canada. Fentanyl is often mixed with other drugs, such as heroin or cocaine, without the knowledge of the user. This makes it particularly dangerous because it is much more potent than other opioids, and even a small amount can lead to a fatal overdose.

I'm sorry, but as of my last knowledge update in September 2021, there was no specific information available regarding a rise in Xylazine deaths. Xylazine is a medication primarily used as a veterinary sedative and analgesic. However, it is worth noting that my information might be outdated, and there could have been developments since then.

Of note, Xylazine was not recognized as a threat when the model was trained in September 2021.

Potential benefits of generative AI for CDC

Potential benefit to CDC	Examples
Save resources: Automate or accelerate certain tasks and processes	<ul style="list-style-type: none">• Draft an executive summary of a written report from talking points• Draft a chart abstraction questionnaire• Draft a scope of work for a contract• Draft a position description for a vacant job
Increase efficiency: Accelerate manual or repetitive tasks that do not require a high level of expertise or using structured formats	<ul style="list-style-type: none">• Draft or edit scientific products including abstracts and manuscripts• Create resumes or other structured documents• Create or edit code and documentation• Translate code between programming languages
Synthesize and summarize information	<ul style="list-style-type: none">• Summarize written material• Describe methodological options, with references, for conducting multiple imputation at scale
Communicate: Adapting text to various reading levels, platforms, audiences	<ul style="list-style-type: none">• Draft communications messages• Translate websites or other written materials between spoken languages• Draft tweets from a list of talking points.
Personalize responses	<ul style="list-style-type: none">• Draft responses to public inquiries with references to relevant CDC guidance

CDC internal report, "Utility and Risks to CDC of Conversational Artificial Intelligence (AI) Technologies Like ChatGPT", June 2023.

Potential risks of generative AI for CDC

Potential harm to CDC	Examples
Failure to meet ethical or regulatory standards	<ul style="list-style-type: none">Relevant regulations and policies that could be threatened by AI use include those addressing data accuracy, data security and privacy/confidentiality, intellectual property, health equity, discrimination, and scientific integrity.
Fabrication	<ul style="list-style-type: none">Conversational AI may “hallucinate” or fabricate information including data and references has been reported widely. Fabrication, falsification, and plagiarism are included in the federal definition of reportable research misconduct.
Bias and discrimination	<ul style="list-style-type: none">Conversational AI technology is trained on vast amounts of data from undisclosed sources, and tone and content may include biases, misconceptions, or outdated information regarding specific populations
Lack of privacy and data security	<ul style="list-style-type: none">Cloud-based models rely on external infrastructure, which may not always be compliant with the strict security standards required in the healthcare sector or by federal assurances.
Other legal or ethical risks, including risks to agency credibility	<ul style="list-style-type: none">Transparency, oversight, and accountability for AI-generated material can be unclear.

CDC internal report, "Utility and Risks to CDC of Conversational Artificial Intelligence (AI) Technologies Like ChatGPT", June 2023.

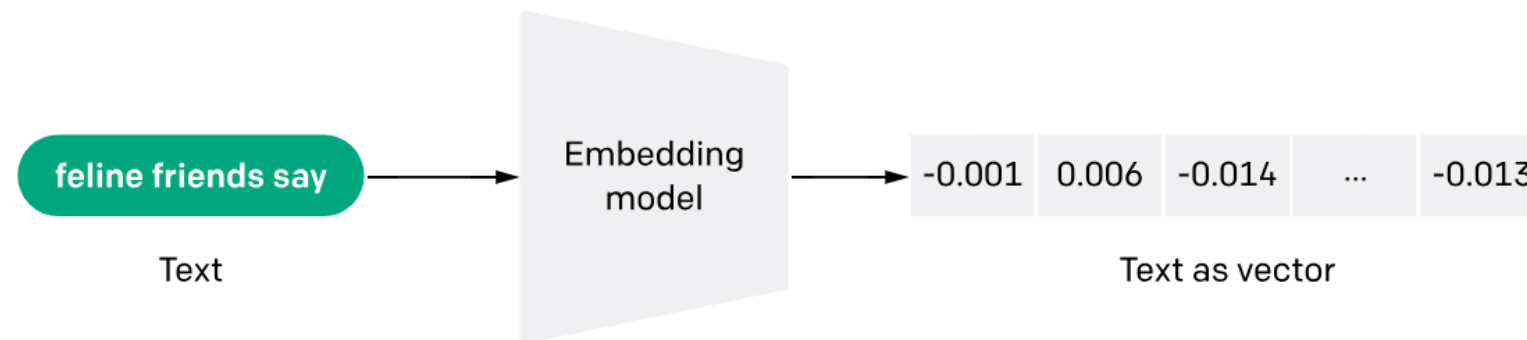
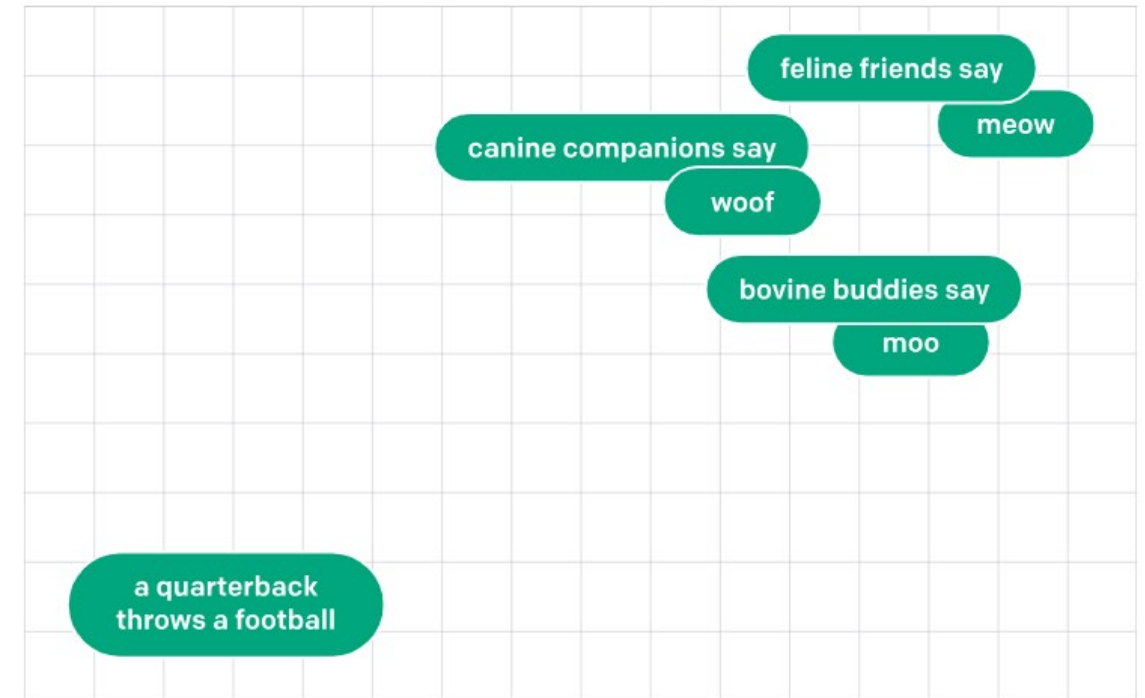
CDC should anticipate malicious use by external actors

Potential action	Examples
Generating misinformation	<ul style="list-style-type: none">• External actors might generate large volumes of text-based communications that include unintentional errors or propagate misinformation.• Misinformation may use various strategies to appear credible, such as using scientific language or citing references.
Impersonating CDC	<ul style="list-style-type: none">• External actors might use generative AI to imitate CDC or other public health authorities (“deepfakes”) or hallucinated written material falsely attributed to CDC.• CDC may need to expend resources explaining or debunking information falsely attributed to CDC, including requests to verify statements, publications, or guidance.
Spamming CDC	<ul style="list-style-type: none">• In platforms for public communication where CDC is required or obligated to respond, external actors might spam CDC with large volumes of credible inquiries.• This includes public comments (including on government documents or at open government meetings), public inquiries to CDC (including CDC-Info and NIP-Info), social media comments and inquiries, and media or other requests for CDC response.• Recently, volume of public comments at a CDC ACIP in late 2022 increased from an anticipated 600 comments to 127,000 comments (unknown proportion AI).

Generative AI powering search & literature review

“[...] tremendous opportunity for expanding scope of literature review.”, comments from NCVHS member.

RAG (retrieval augmented generation)




Thank you.

Travis Hoppe, thoppe@cdc.gov

Associate Director of Data Science and Analytics, NCHS/CDC/HHS

For more information, contact CDC
1-800-CDC-INFO (232-4636)
TTY: 1-888-232-6348 www.cdc.gov

A decorative footer consisting of several overlapping diagonal stripes in various shades of blue, ranging from dark navy to a lighter sky blue, creating a modern, abstract geometric design.