Data & AI in Food Safety workshop: Framing & Foundations

June 5, 2026

Richard Todd
Chief Data Officer, City of Boulder

1) Welcome and introduction

Purpose and scope of this session

Grounding us all for today's workshop

I hope to:

- Establish a common understanding of AI and core concepts
- Explore trends in adoption and what is driving this latest wave
- Draw out implications for our work

Agenda

1) Welcome and introduction	3 minutes
2) What is it?	10 minutes
3) Why is it everywhere?	5 minutes
4) What does it mean for our work?	10 minutes
5) Questions!	10 minutes
6) Close and next steps	2 minutes

Level set: familiarity test

Level set: familiarity test

ChatGPT

Level set: familiarity test

Large language model (“LLM”)

Level set: familiarity test

$$\mathrm{w}^{(t+1)} = \mathrm{w}^{(t)} - \alpha \nabla f(\mathrm{w}^{(t)})$$

My approach

Setting solid foundations for the day!

Technology – including data and AI – is too important to be left to the technologists.

- We've come to empower you all to wrestle with this
- Conceptual, rather than technical

Broad before we go deep.

- We're going to hear about some fantastic, real-life food safety use cases later
- We will wrestle with specific scenarios in the hands-on session
- In this session we will focus more on the underlying technology and use examples from other settings

Proposed norms for today

To maximize our time together, I ask that we all adhere to the following norms.

No questions are stupid! We have come here to learn and to share.

- Please ask any urgent clarificatory questions in real-time
- Keep general questions to the designated space in each section, or the agenda item

As we wrestle with the ideas today, bring a generative mindset: positive and open.

- Constraints will come later!
- A range of critical safeguards will be addressed in a specific session

Questions?

2) Introduction to AI

What is AI?

“Artificial intelligence” is a broad and ambiguous term.

Two general approaches to defining it:

- Technological: as a collection of related technologies (regardless of how used)
- Functional: as the completion of a set of tasks that were traditionally human intelligence (regardless of how performed)

Both have merits, but both have drawbacks:

- Too focused on the arcane details of technology
- Liable to be misused (or oversold)

We will focus on three groups of tasks and the technologies that underpin them.

First: sorting out the terms

AI is broader than ChatGPT, Claude.ai

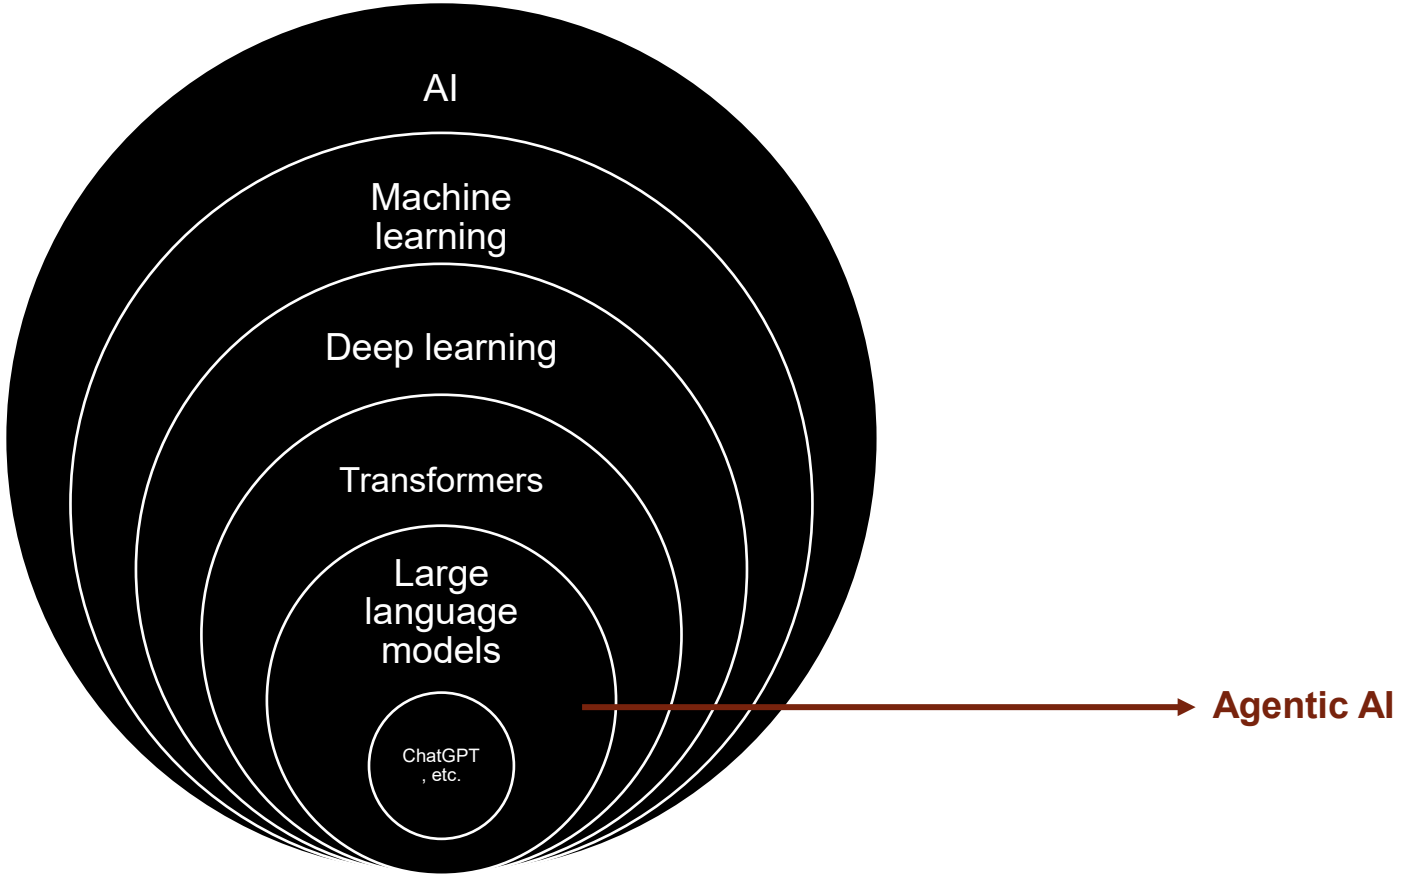

Second: a brief history of AI in three leaps

We will explore three major milestones; the use-cases that they unlocked and implications.

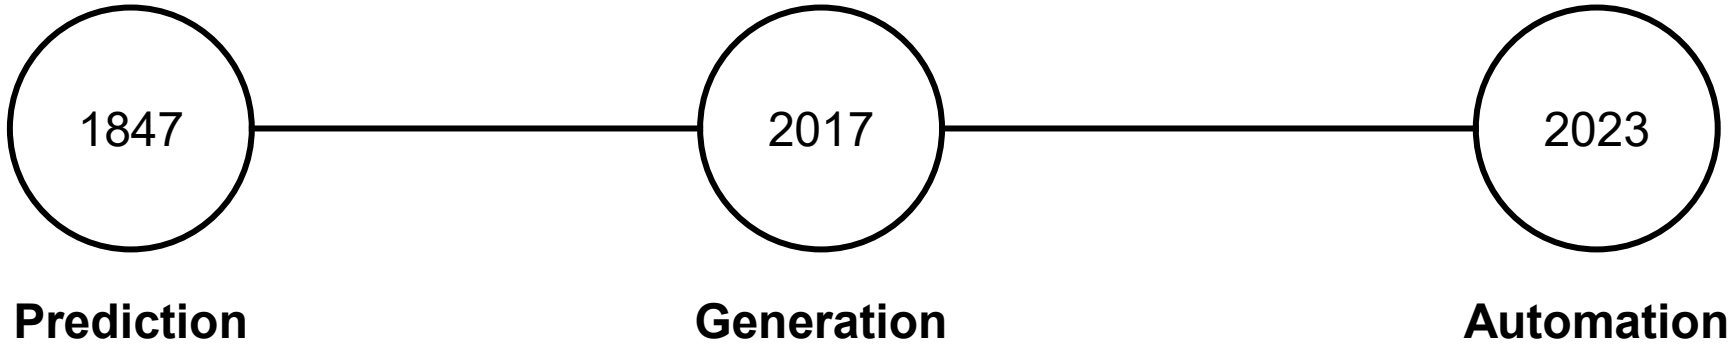

Prediction: the 1847 discovery

The foundation of computational optimization, which underpins prediction.

Augustin Louis Cauchy develops gradient descent, a method of mechanically identifying minima.

- Pick a starting point. Feel the slope under your feet. Take a small step in the steepest downhill direction. Then feel again. Then step again. Repeat!

Versions of this mathematical optimization technique underpin Machine Learning, and all modern AI.

Image credit: Wikipedia

Prediction: the use-case

Predicting missing or future values using historic patterns.

Prediction can be hugely valuable, allowing us to better plan for the future or organize our work.

- Consider how early signals from public-reported data might predict risk

Deployment of automated sepsis prediction system in Baltimore/DC area hospitals flagged issue 2-3 hours before clinicians, associated with 18% reduction in mortality.

Mexican SEDESOL social safety net deployed machine learning to identify likely eligibility for programs where clients had not yet made applications.

The UK Food Standards Agency built the "Food Hygiene Rating Scheme – AI" tool to help local authorities prioritize which food establishments to inspect, predicting non-compliance risk from business history, ownership, prior inspection outcomes, and local socio-economic factors.

Adams, R., Henry, K. E., Sridharan, A., Soleimani, H., Zhan, A., Rawat, N., ... & Saria, S. (2022). Prospective, multi-site study of patient outcomes after implementation of the TREWS machine learning-based early warning system for sepsis. Nature medicine, 28(7), 1455-1460.

Sankaran, K. et al. (2017). Applying Machine Learning Methods to Enhance the Distribution of Social Services in Mexico. arXiv

UK Government, "FSA: Developing an AI-based Proof of Concept that prioritises businesses for food hygiene inspections," GOV.UK case study. https://www.gov.uk/ai-assurance-techniques/fsa-developing-an-ai-based-proof-of-concept-that-prioritises-businesses-for-food-hygiene-inspections-while-ensuring-the-ethical-and-responsible-use-of-ai

Prediction: implications

What should we take from this?

The math is not new; the availability of data and computing power is.

AI tasks are framed as minimizing error from a target value.

Optimization toward what? The importance of training data.

Generation: the 2017 discovery

Massively parallelizable architecture that unlocked huge models for language.

Researchers at Google developed transformer architecture which models language relationships in parallel.

Before this, models read text one word at a time, in order, storing what came before; slow and memory-intensive.

Architecture that could process in parallel meant that models can be trained on the entire internet and robustly modeling complex relationships across human language was possible, leading to Large Language Models ("LLMs").

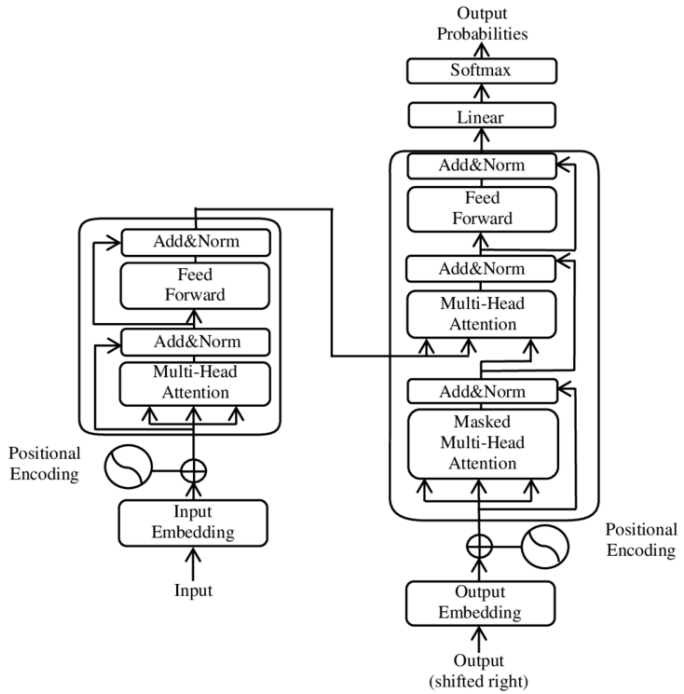

Image reproduced from Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.

Generation: the use case

Generating new text (or images…) from prompts.

Once you can model patterns across all language, you can produce new language that fits the patterns.

Generative AI can rapidly produce drafts, summaries, or creative content.

- Drafting inspection narratives, summarizing complaint streams, translating documents, explaining technical material to the public, generating training scenarios…

The UK civil service deployed Generative AI to 20,000 civil servants, and estimated productivity impacts to be 2 weeks per person per year.

UK Government News: https://www.gov.uk/government/news/landmark-government-trial-shows-ai-could-save-civil-servants-nearly-2-weeks-a-year

Generation: implications

What should we take from this?

This [extraordinary] technical leap is only nine years old.

The model was trained on one objective: predict the next token.

Fluency has decoupled confidence from correctness.

Automation: the ~2023 discovery

LLMs displaying step-by-step reasoning that drove the leap to agentic AI.

Building on the transformer, AI developers combined LLMs with tool use, memory, and step-by-step reasoning, creating systems that can take sequential actions.

Before this, models could only respond to a prompt with text; they couldn't act in the world or string multiple steps together toward a goal.

Letting models observe, decide, and act in loops means that AI can now execute multi-step workflows on its own; 'agentic' AI.

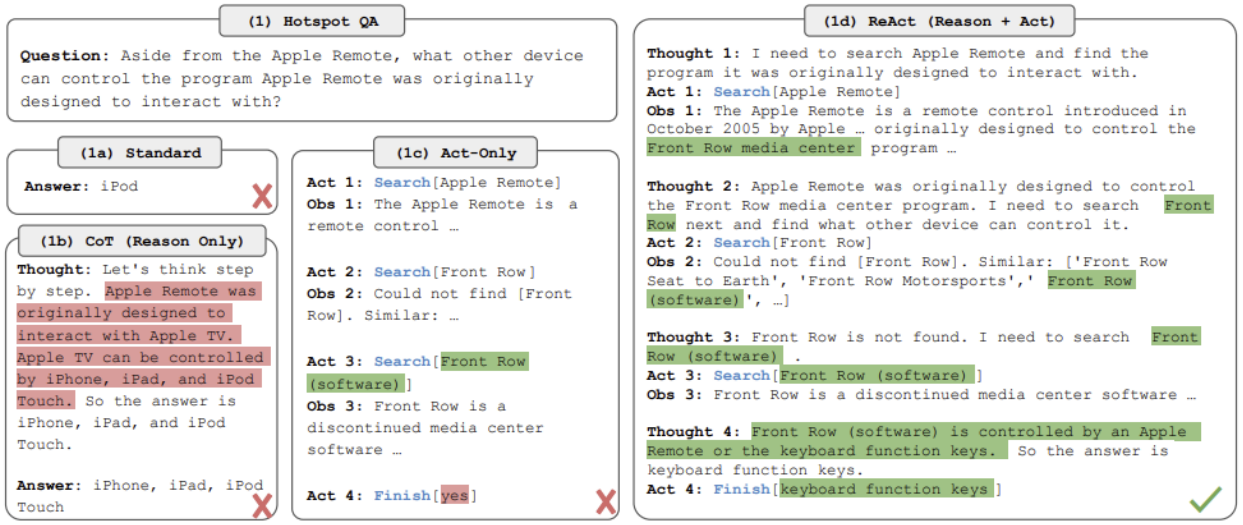

Image reproduced from Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.

Automation: the use case

Automating procedural workflows.

AI can be used to implement adaptive automation workflows, automating repetitive tasks, reducing human effort and (potentially) error.

A UK pilot estimated that about 143 million complex, repetitive citizen-facing transactions could be automated across central government, potentially saving over 1,200 person-years annually.

In June 2025, the U.S. FDA launched "Elsa," a generative AI tool deployed agency-wide to identify high-priority inspection targets by analyzing adverse event reports, compliance data, Form 483 observations, and historical inspection outcomes. In December 2025, the agency expanded Elsa to "agentic AI" for safety reviews, inspections, and compliance.

Straub, V. J., Hashem, Y., Bright, J., Bhagwanani, S., Morgan, D., Francis, J., ... & Margetts, H. (2024). AI for bureaucratic productivity: Measuring the potential of AI to help automate 143 million UK government transactions. arXiv preprint arXiv:2403.14712.

U.S. FDA, "FDA Launches Agency-Wide AI Tool to Optimize Performance for the American People," June 2, 2025. https://www.fda.gov/news-events/press-announcements/fda-launches-agency-wide-ai-tool-optimize-performance-american-people

Automation: implications

What should we take from this?

The breadth of applicability was a surprise; general agentic systems turned out to perform well on range of tasks.

Stakes change from ‘wrong words’ to ‘wrong actions’.

Human oversight needs to be carefully designed-in.

Second: a brief history of AI in three leaps

Three major milestones; technology underpinning three types of use-cases.

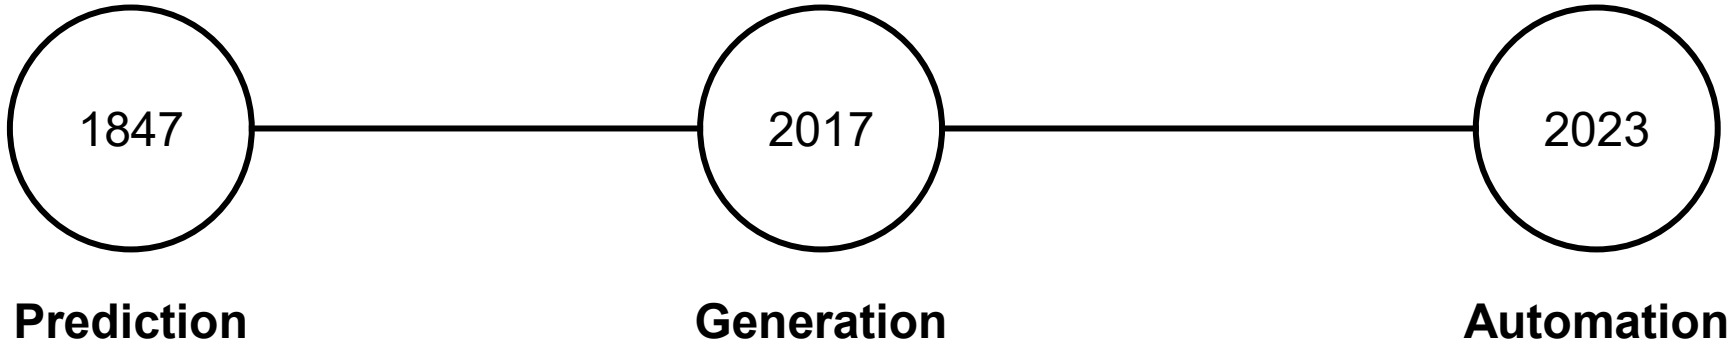

Questions?

3) Why is it everywhere?

Why is it everywhere?

Our old friends, supply and demand.

Connections to what we've discussed already

- Widespread (and growing…) availability of data and computing power leading to more opportunities for prediction, larger models
- Surprising (?) generalizability of models and applicability of tasks. Credible estimates that 23%[1] of tasks can be feasibly (and profitably) automated within ten years.

Market forces

- Expectation -> investment -> market share
- Free at the point of use

1 - Svanberg, Maja, Wensu Li, Martin Fleming, Brian Goehring, and Neil Thompson, Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?, 2024. Working Paper.

But… is it everywhere?

Adoption rates are faster than previous major technologies, but not yet saturated.

Indicative industry analysis suggested a thirty-four percentage point gap between AI adoption rates in technology companies and in regulated industries, such as government, healthcare and finance.

- Technology investment gaps mean that this is likely to be even lower in public sector organizations.

Shadow AI; a leading cyber security firm reported that 68% of analyzed employees use free-tier AI tools via personal accounts, with 57% inputting sensitive data.

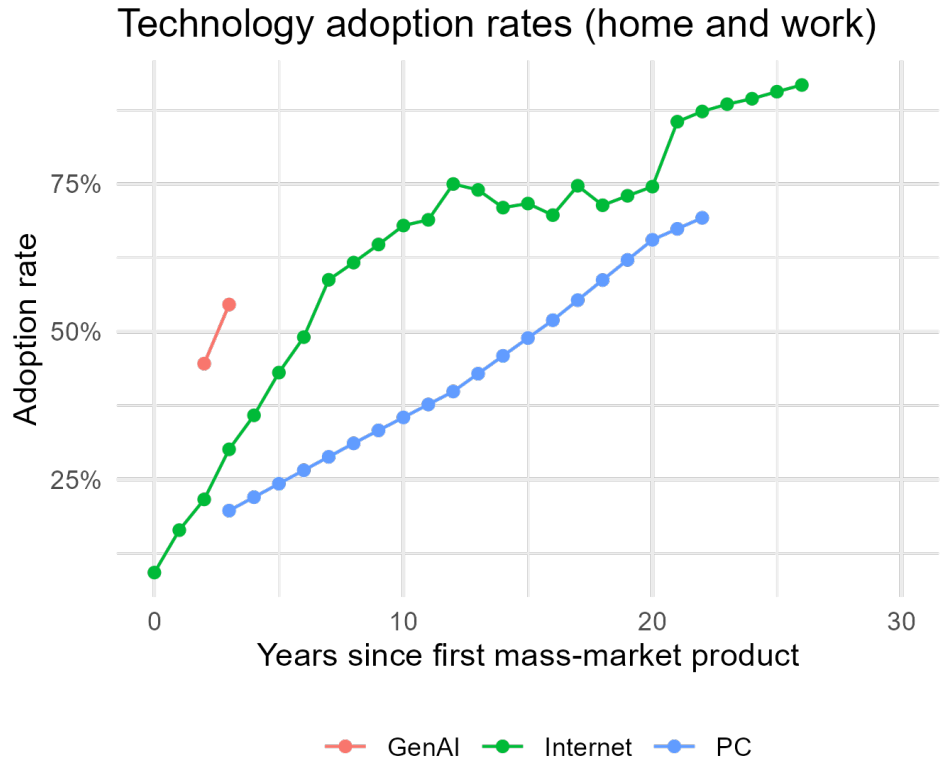

Sources: Real-Time Population Survey, ITU, Federal Reserve Bank of St. Louis

Deskpro's State of AI in Support Operations discussed at https://www.businesswire.com/news/home/20251112702468/en/New-Study-Reveals-Tech-Industry-Leads-AI-Adoption-in-Support-Operations-at-92-Regulated-Industries-Only-at-58-Due-to-Compliance-Constraints
Menlo Security: How AI is shaping the modern workplace. Available at https://www.menlosecurity.com/resources/how-ai-is-shaping-the-modern-workspace-report

Questions?

4) What does it mean for our work?

Major themes

Dimension	Implication for our work
Breadth of applicability.	Huge opportunity (bigger hype)
Speed of adoption.	Changing expectations among stakeholders
Built on data, expanded scope	Growing importance of data governance
Questions over accountability; new regulation.	AI governance practices
Fear of human replacement.	Change management

1 - Svanberg, Maja, Wensu Li, Martin Fleming, Brian Goehring, and Neil Thompson, Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?, 2024. Working Paper.

Remember our quick history tour?

As opportunities grow, the demands upon us increase, rather than decrease.

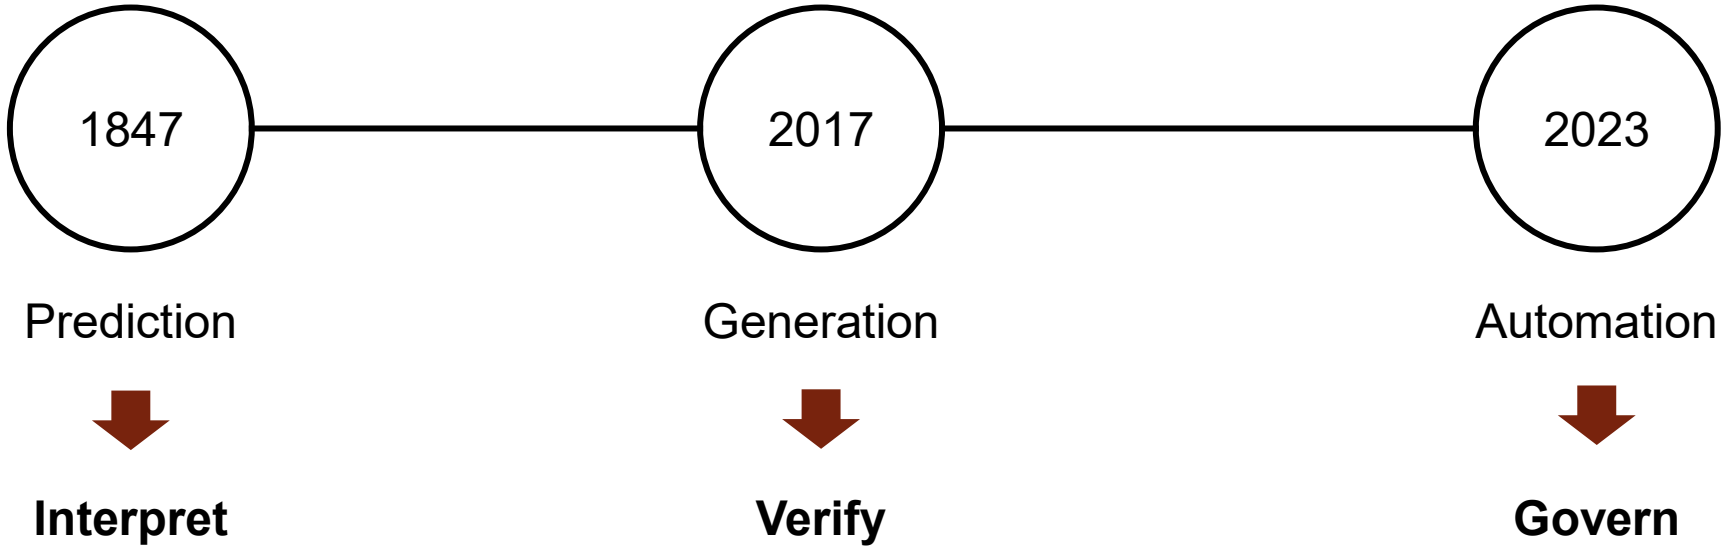

Questions?

5) Questions!

6) Close and next steps

Purpose and scope of this session

Did we get here?

I hoped to:

- Establish a common understanding of AI and core concepts
- Explore trends in adoption and what is driving this latest wave
- Draw out implications for our work

Looking forward

Two questions to carry through the rest of the day.

How can these tools further the food safety mission?

What would it take to trust them enough to act on?

Remember: too important to be left to the technologists!

Coming up

Grounding us all for today's workshop

All to come:

- Data Culture & Data Quality with Professor Kowalcyk
- AI governance with the entire panel
- Rich case studies from a range of experts
- Hands-on storyboarding and feedback

Thank you.

Richard Todd
Chief Data Officer, City of Boulder