

Local Government **REIMAGINED CONFERENCE**

JUNE 5-7, 2024 PALM DESERT, CALIFORNIA

Putting Your Data to Work with AI

Jennifer Robinson, SAS Global Public Sector Strategic Advisor

ICMA | conference

AI

Computer
Vision

Machine Learning

Deep Learning

Generative AI
"GenAI"

LLM

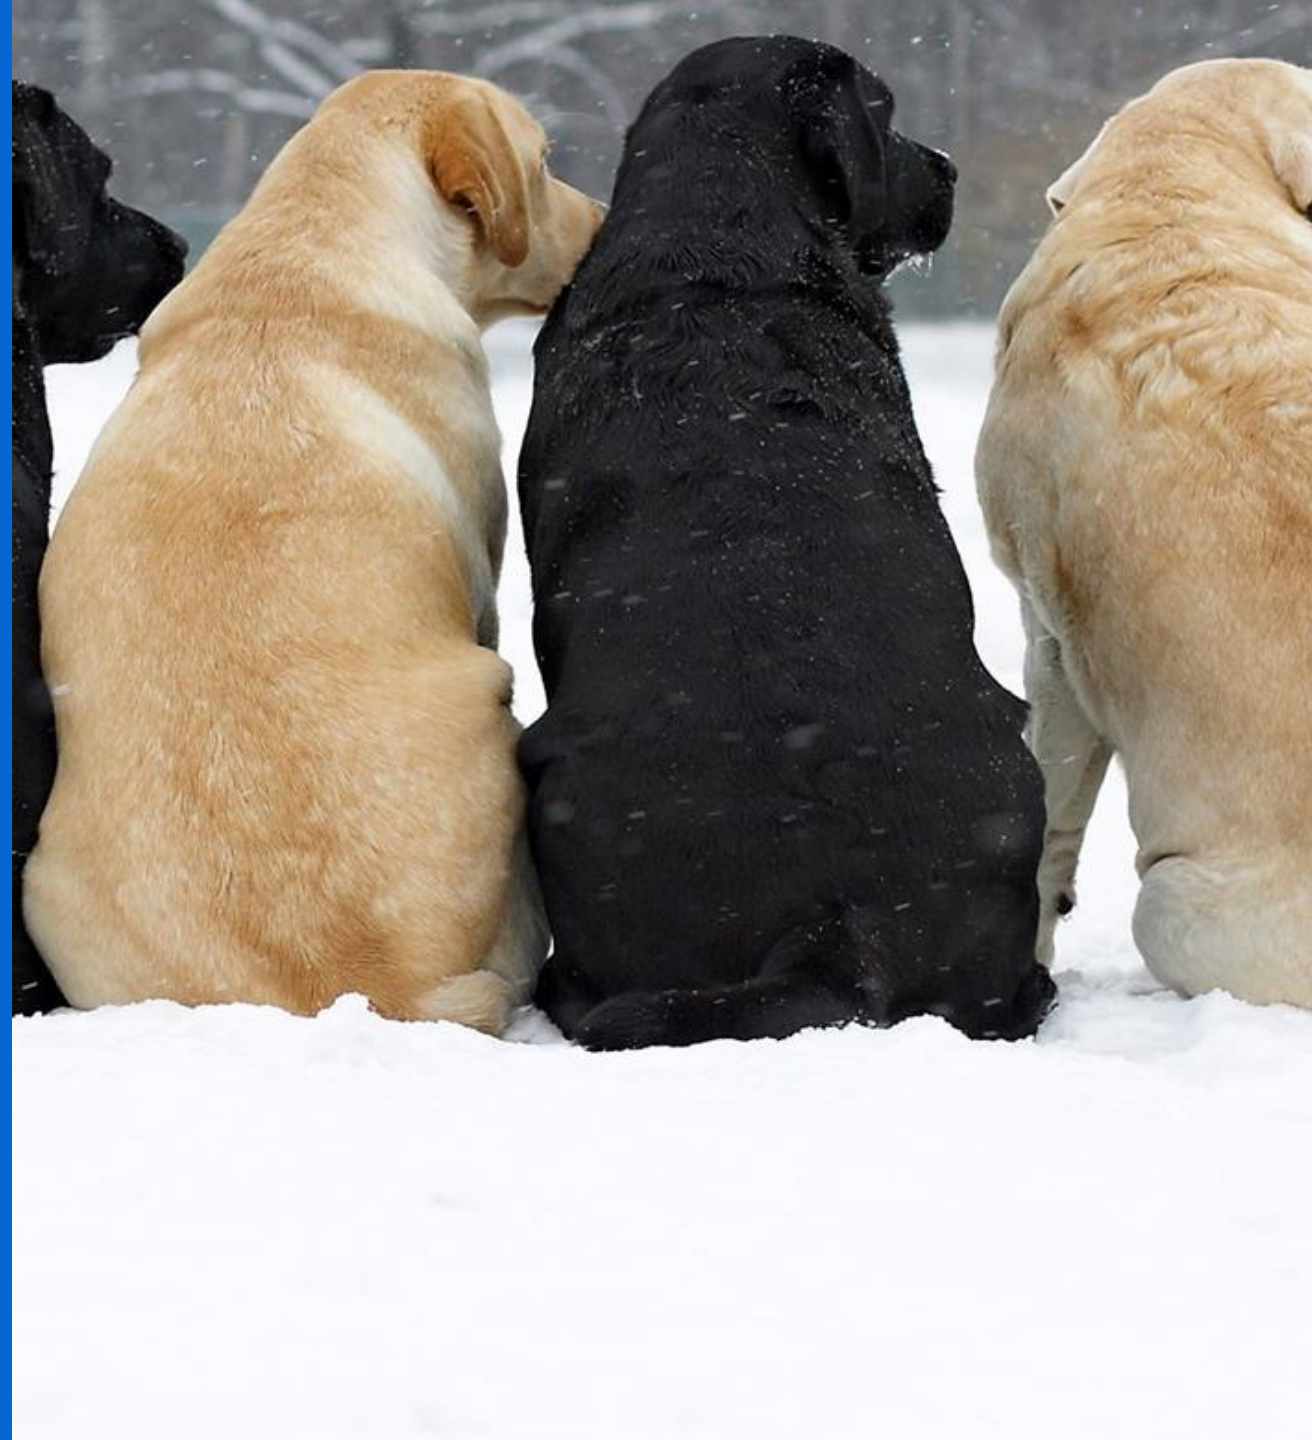
Synthetic Data

Digital Twins

Natural
Language
Processing

Machine Learning

Systems learn from data, identify patterns, and make decisions with minimal human intervention



Natural Language Processing

Enables understanding,
interaction and communication
between humans and machines.



Computer Vision

Enables systems to see, identify and process images or videos in the same way that human vision does.



Large Language Models (LLMs)

Deep learning algorithms that recognize, summarize, translate, predict, and generate content



Synthetic Data

Artificial data that's
manufactured.



Digital Twins

A virtual model designed to accurately reflect a physical object, system, or environment.



Deep Learning

A type of machine learning used in recognizing speech, identifying objects in images, generating new content, and more.

Comprised of three or more neural networks.



What are Humans and Machines **good at?**



HUMANS

- Common Sense
- Intuition
- Creativity
- Empathy
- Versatility



MACHINES

- Large Data Sets
- Complex Calculations
- Learning
- Automation

AI Adoption

ICMA | conference

Local Government
REIMAGINED CONFERENCE

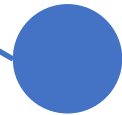
APRIL 12-14, 2023 | BURLINGTON, VERMONT

A possible progression of AI adoption

Automation of
routine tasks

Data analysis and
insights

Optimization and
logistics



Personalization and
CX

AI-enhanced
products and
services

Decision-making
strategy

Areas ripe for AI adoption



Automation of routine tasks

- Citizen service inquiries
- Data augmentation and simulation
- Document Processing



Personalization and CX

- Targeted assistance and programs
- Customized engagement and services
- Nudging for tax compliance



Data analysis and insights

- Utility forecasting
- Fraud detection
- Citizen segmentation



AI-enhanced products

- Healthcare diagnostic tools
- Tax robo-advisors
- Smart signalization



Optimization and Logistics

- Demand forecasting
- Inventory management
- Route optimization



Decision-making strategy

- Evaluating scenarios
- Decision objectivity
- Policy evaluation

Question:

What are the tasks within your organization that are ripe for an AI application?

AI Readiness

ICMA | conference

Local Government
REIMAGINED CONFERENCE

APRIL 12-14, 2023 | BURLINGTON, VERMONT

Building an AI-Ready Team

1

Take an Inside – Outside approach

The Inside – Outside Approach

Know what you are solving for

Inside

Your SMEs
Small AI Exploratory Team
Your challenges/opportunities

Outside

AI techniques
New external hire
Hackathons/University Challenges

Building an AI-Ready Team

- 1 Take an Inside – Outside approach
- 2 Compete for talent

Compete for talent



Compete



Develop and Retain
Current Talent

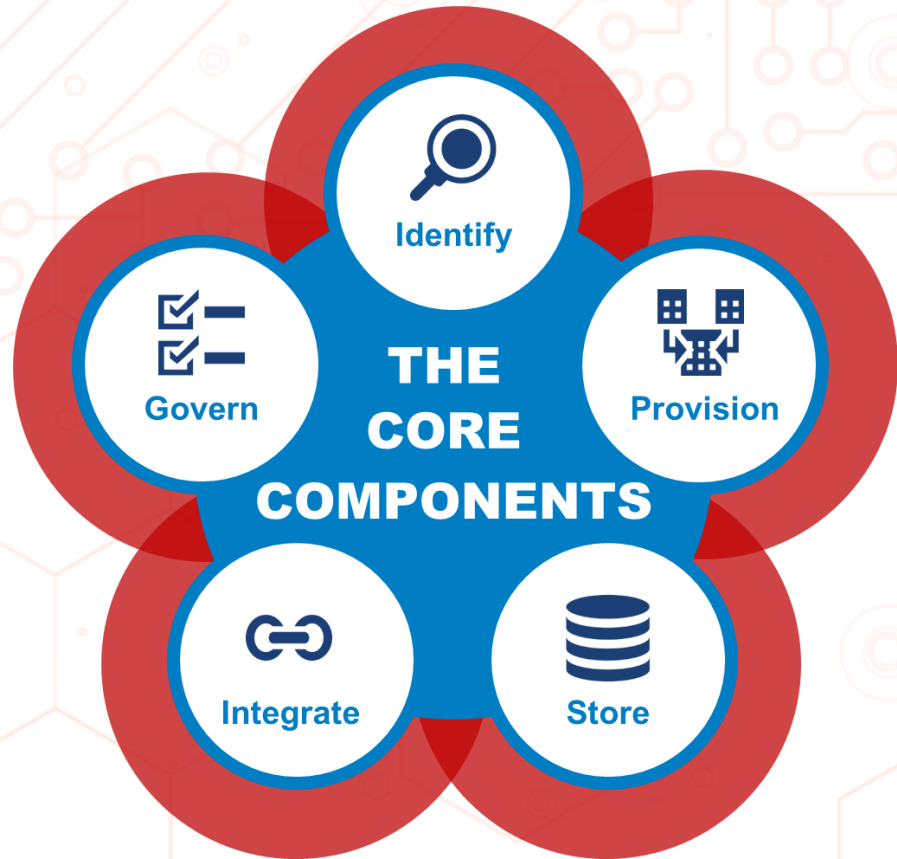


Establish Academic
Connections

Building an AI-Ready Team

- 1 Take an Inside – Outside approach
- 2 Compete for talent
- 3 Data and compute infrastructure as AI infrastructure

The Five Components of a Data Strategy



Identify

The ability to identify data and understand its meaning regardless of its structure, origin, or location

Provision

Enabling data to be packaged and made available while respecting all rules and access guidelines

Store

Persisting data in a structure and location that supports access and processing across the enterprise

Integrate

Moving and combining data residing in multiple locations and providing a unified view of the data

Govern

Establishing and communicating information policies and mechanisms to ensure effective data usage.

SAS® Data Management Framework

DATA GOVERNANCE

The organizing framework for establishing strategy, objectives, and policies for enterprise data



SAS® Data Management Framework

DATA GOVERNANCE

- ✓ Establishes the **authority** for **defining** key business terms
- ✓ Establishes **business-driven decision-making process** for **changes** to definitions or business rules
- ✓ Defines metadata and data quality **standards**
- ✓ Defines process for ongoing policy **monitoring** and **measurement**
- ✓ Ensures data is managed **separately** from applications



SAS® Data Management Framework

DATA GOVERNANCE

Program Objectives

- What do we want to achieve with Data Governance?
- Can we measure against the objective?

Sample

Eliminate redundant or conflicting business processes and practices

SAS® Data Management Framework

DATA GOVERNANCE

Guiding Principles

- Set of common philosophies intended to direct a program or organization irrespective of day-to-day changes
- Reflect the organization's values and goals
- Serve as a philosophical touchstone for questions or dilemmas during Data Governance deployment

Sample

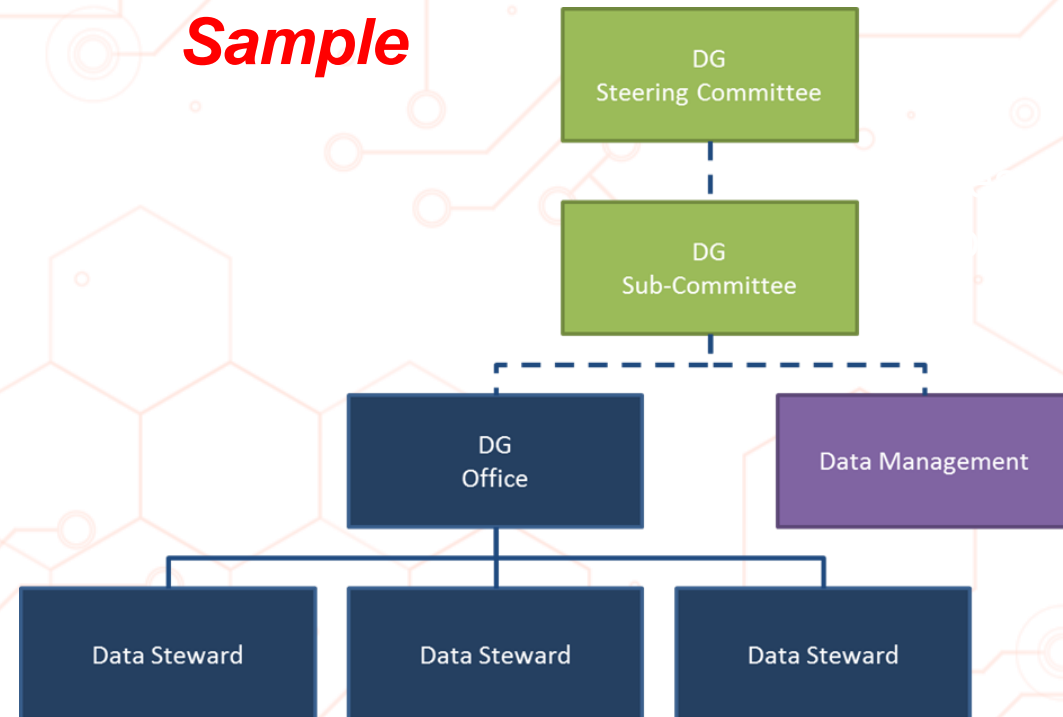
Customer data attributes will be shared through a common process or interface

SAS® Data Management Framework

DATA GOVERNANCE

Decision-making Bodies

Sample



SAS® Data Management Framework

DATA GOVERNANCE

Sample

Decision Rights

DG Function	DG Activities	DG Office	Enterprise Steering Committee	DG Sub-Committee	Data Steward Team	DM Team
MDM	Write policies and standards for MDM hub	R	I	A	R	R
	Approve policies and standards for MDM hub	I	I	A/R	I	I
	Define quality and validation	R	I	A	R	C
	Define and modify matching rules	R	I	A	R	C
	Define and modify survivorship rules	R	I	A	R	C
	Define data provisioning standards	C	I	A	C	R
	Monitor and report hub metrics (data quality, hub performance, etc.)	R	I	A	R	C
	Monitor compliance with MDM policies	R	I	A	R	C

SAS® Data Management Framework

DATA GOVERNANCE

Program Objectives

Guiding Principles

Decision-making Bodies

Decision Rights

Why?

How?

Who?

What?

Question:

What kind of data and infrastructure challenges do you anticipate as your organization increases adoption of AI?

Building an AI-Ready Team

- 1 Take an Inside – Outside approach
- 2 Compete for talent
- 3 Data and compute infrastructure as AI infrastructure
- 4 Take an experimentation mindset
- 5 Run like Dev Ops

Treat AI projects like applications development

Common AI Job Roles

Data Scientist	Analyzes and interprets complex data to help in decision-making. They use statistical techniques, machine learning algorithms, and predictive modeling to extract insights from data.
Machine Learning Engineer	Focuses on designing and implementing machine learning models and algorithms. They work closely with data scientists to develop scalable and efficient AI systems.
Data Engineer	Designs and maintains data infrastructure required for storing, processing, and analyzing large datasets. Ensures that data is accessible and in a usable format for data scientists and analysts.
Software Developer	Develops and maintains software applications that incorporate AI functionalities. They integrate machine learning models into applications and ensure their smooth operation.
AI Product Manager	Oversees development and deployment of AI products, ensuring that they meet user needs and business objectives. They work at the intersection of business, technology, and user experience.
Ethics & Compliance Officer	Ensures that AI systems adhere to ethical standards and comply with legal and regulatory requirements. They address concerns related to privacy, fairness, and transparency.
Solutions Architect	Designs overall architecture of AI systems, including selection of technologies, frameworks, and infrastructure. Ensures that the system is scalable, secure, and aligned with business goals.
Business Analyst	Works on identifying business needs and opportunities where AI can provide solutions. They help translate business requirements into technical specifications for the AI team.
DevOps Engineer	Responsible for the continuous integration and deployment (CI/CD) of AI applications. They ensure that the development, testing, and production environments are streamlined and automated.

Building an AI-Ready Team

- 1 Take an Inside – Outside approach
- 2 Compete for talent
- 3 Data and compute infrastructure as AI infrastructure
- 4 Take an experimentation mindset
- 5 Run like Dev Ops

Challenges Faced with Generative AI

Our Principles for Trustworthy AI

Human-Centricity



Promote human **well-being**, human **agency** and **equity**.

Inclusivity



Ensure **accessibility** and include **diverse perspectives** and **experiences**.

Accountability



Proactively identify and mitigate adverse impacts.

Transparency



Explain and **instruct** on usage openly, including potential risks and how decisions are made.

Robustness



Operate **reliably** and **safely**, while enabling mechanisms that assess and manage potential risks throughout a system's lifecycle.

Privacy & Security



Respect the privacy of data subjects.

Challenges Involved in Using Generative AI



Accuracy

LLMs can produce output that is factually incorrect, not contextual or non-sensical.

Consider:

- Knowledge cutoff date
- Completion relies on context
- Hallucinations



Privacy & Security

LLMs can be manipulated to obtain desired output.

For example:

- Revealing private information
- Stealing private information
- Deliberately providing false information
- Executing arbitrary code



Governance

LLMs can be released in non-standard ways across providers.

- High computational demand translates in high cost for using LLMs.
- Commercial models are subject to change in model behavior and policy



Challenges Involved in Using Generative AI



Accuracy

LLMs can produce output that is factually incorrect, not contextual or non-sensical.

Consider:

- Knowledge cutoff date
- Completion relies on context
- Hallucinations



Privacy & Security

LLMs can be manipulated to obtain desired output.

For example:

- Revealing private information
- Stealing private information
- Deliberately providing false information
- Executing arbitrary code



Governance

LLMs can be released in non-standard ways across providers.

- High computational demand translates in high cost for using LLMs.
- Commercial models are subject to change in model behavior and policy



What is the problem here?

My prompt:

Using Generative AI policies adopted by US cities, counties, and states, write a Generative AI policy that governs its use by employees. Cite references to existing policies.

The LLM Response:

Introduction:

As [Organization/Agency] embraces innovative technologies to enhance its operations and service delivery, the adoption of generative artificial intelligence (AI) presents significant opportunities and challenges. This policy aims to provide guidance and establish best practices for the responsible and ethical use of generative AI technologies by employees.

Policy Statement:

Purpose: The purpose of this policy is to promote the ethical, transparent, and accountable deployment of generative AI technologies by employees of [Organization/Agency] to achieve organizational objectives while upholding the highest standards of integrity and public trust.

References:

- "AI in Government: Principles and Guidelines" (City of Seattle)
- "Guidelines for the Ethical Use of Artificial Intelligence by the Government" (New York City Mayor's Office)
- "Responsible AI Governance Framework" (State of California)
- "AI Ethics Guidelines for Public Agencies" (National Association of State Chief Information Officers)

What is the problem here?

My prompt:

Using Generative AI policies adopted by US cities, counties, and states, write a Generative AI policy that governs its use by employees. Cite references to existing policies.

The LLM Response:

Ethical Principles:

- a. Fairness and Equity: Employees must ensure that generative AI applications are designed and implemented to minimize bias and promote fairness and equity in outcomes, particularly in areas such as decision-making and resource allocation.
- b. Transparency: Employees utilizing generative AI technologies must provide clear and transparent explanations of the processes, methodologies, and decision criteria employed, ensuring accountability and public trust.
- c. Privacy and Data Protection: Employees are responsible for safeguarding the privacy and confidentiality of data used in generative AI applications, complying with all applicable laws and regulations governing data privacy and protection.**
- d. Accountability: Employees must take responsibility for the outcomes and impacts of generative AI technologies under their purview, promptly addressing any concerns or issues that arise and mitigating risks to the organization and stakeholders.



City of Boston Interim Guidelines for Using Generative AI

Version 1.1

Prepared by Santiago Garces, Chief Information Officer, City of Boston

Published: 5/18/2023

Applies to: *all City agencies and departments with the exception of Boston Public Schools*

“Do not share sensitive or private information in the prompts

Why: data including prompts used in generative AI might be used by the companies that power these systems. Any information that includes personally identifying information about our residents, other public servants, etc. could inadvertently be shared with others. Basically if you wouldn't share with other people or want to put the prompt in a public place, avoid sharing the information in the prompt. If you have an application that requires sensitive information to be used with a generative AI, contact DoIT so we can help you provision access to enterprise secure resources to do so.”

What is the problem here?

My prompt:

What local governments protect streams and intermittent streams with a 200' buffer?

The LLM Response:

While regulations vary by jurisdiction, several local governments have enacted stream buffer protections with a 200-foot buffer zone. Here are some examples:

- City of Seattle, Washington: Seattle's Environmentally Critical Areas Ordinance establishes a 200-foot buffer along streams, wetlands, and other sensitive areas to protect water quality, habitat, and ecological functions.
- City of Raleigh, North Carolina: Raleigh's Stream Buffer Protection Regulations require a 200-foot buffer along perennial streams and a 100-foot buffer along intermittent streams to safeguard water quality, reduce erosion, and maintain natural habitat.
- City of Portland, Oregon: Portland's Stream Protection and Enhancement Policy includes a 200-foot buffer requirement along both sides of streams to protect riparian areas, water quality, and wildlife habitat.
- City of Chapel Hill, North Carolina: Chapel Hill's Stream Buffer Protection Ordinance establishes a 200-foot buffer along perennial streams and a 100-foot buffer along intermittent streams to protect water quality, prevent erosion, and preserve natural vegetation.

5.18.6 Riparian Area Protection

(a) Riparian Buffer Zones

The protected riparian buffer shall have two zones as follows:

- (1) Zone One shall consist of a vegetated area that is undisturbed except for uses and activities provided for in Subsection 5.18.7(b). The location of Zone One shall be as follows:
 - A. For intermittent and perennial streams, Zone One shall begin at the top of the bank and extend landward a distance of thirty (30) feet on all sides of the surface water, measured horizontally on a line perpendicular to a vertical line marking the top of the bank.
 - B. For perennial waterbodies located within a natural drainage way, Zone One shall begin at the normal pool elevation and extend landward a distance of thirty (30) feet, measured horizontally on a line perpendicular to a vertical line marking the normal pool elevation.
- (2) Zone Two shall consist of a stable, vegetated area that is undisturbed except for uses and activities provided for in Subsection 5.18.7(b). Grading and revegetating in Zone Two is allowed provided that the health of the vegetation in Zone One is not compromised. Zone Two shall begin at the outer edge of Zone One and extend landward twenty (20) feet as measured horizontally on a line perpendicular to the surface water. The combined width of Zones One and Two shall be fifty (50) feet on all sides of the surface water.

Challenges Involved in Using Generative AI



Accuracy

LLMs can produce output that is factually incorrect, not contextual or non-sensical.

Consider:

- Knowledge cutoff date
- Completion relies on context
- Hallucinations



Privacy & Security

LLMs can be manipulated to obtain desired output.

For example:

- Revealing private information
- Stealing private information
- Deliberately providing false information
- Executing arbitrary code



Governance

LLMs can be released in non-standard ways across providers.

- High computational demand translates in high cost for using LLMs.
- Commercial models are subject to change in model behavior and policy



What is the problem here?

My prompt:

Summarize list of The Happy Day City Camp campers for each type of allergy listed (bees, peanut butter, tree nuts, suntan lotion, pollen) from these campers: Alice One is allergic to tree nuts and peanut butter; Brandon Twosome is allergic to bees and peanut butter; Charles Third is allergic to bees and pollen, Dylan Quarterly is allergic to tree nuts; Elsie Highfive has no allergies; Franklin Sixth is allergic to bees, suntan lotion and pollen; Gretchen Seven is allergic to bees and tree nuts..

The LLM Response:

Here's a summary of The Happy Day City Camp campers grouped by their type of allergy:

Bees:

- Brandon Twosome
- Charles Third
- Franklin Sixth
- Gretchen Seven

Peanut Butter:

- Alice One
- Brandon Twosome

Tree Nuts:

- Alice One
- Dylan Quarterly
- Gretchen Seven

Suntan Lotion:

- Franklin Sixth

Pollen:

- Charles Third
- Franklin Sixth

No Allergies:

Elsie Highfive

What is the problem here?the follow-up

My prompt:

Who are the participants of the Happy Day City Camp?

The LLM Response:

The participants of The Happy Day City Camp are:

- Alice One
- Brandon Twosome
- Charles Third
- Dylan Quarterly
- Elsie Highfive
- Franklin Sixth
- Gretchen Seven

What is the problem here?the follow-up

My prompt:

Who are the participants of the Happy Day City Camp?

The LLM Response:

The participants of The Happy Day City Camp are:

- Alice One
- Brandon Twosome
- Charles Third
- Dylan Quarterly
- Elsie Highfive
- Franklin Sixth
- Gretchen Seven

Q:

What are some possible goals that a person might have to enter **fake** local government data to corrupt a system?

Challenges Involved in Using Generative AI



Accuracy

LLMs can produce output that is factually incorrect, not contextual or non-sensical.

Consider:

- Knowledge cutoff date
- Completion relies on context
- Hallucinations



Privacy & Security

LLMs can be manipulated to obtain desired output.

For example:

- Revealing private information
- Stealing private information
- Deliberately providing false information
- Executing arbitrary code



Governance

LLMs can be released in non-standard ways across providers.

- High computational demand translates in high cost for using LLMs.
- Commercial models are subject to change in model behavior and policy



What is the problem here?

My prompt:

Organize the following 12,000 public comments in response to a proposed ordinance into like groups and write a summary of each group of comments that reveal support, opposition, or concern. *(enter in 12,000 individual comments)*

Thank you.

Jennifer Robinson

jennifer.robinson@sas.com

ICMA | conference

Local Government
REIMAGINED CONFERENCE

APRIL 12-14, 2023 | BURLINGTON, VERMONT