Generative AI/ML and AI governance for the public sector

Sergio Ortega
sergioai@amazon.com

AI/ML BD Lead

Amazon Web Services

AI/ML/GenAI hierarchy

Artificial intelligence (AI)
Any technique that allows computers to mimic human intelligence using logic, if-then statements, and machine learning

Machine learning (ML)
A subset of AI that uses machines to search for patterns in data to build logic models automatically

Deep learning (DL)
A subset of ML composed of deeply multi-layered neural networks that perform tasks like speech and image recognition

generative AI
Powered by large models that are pretrained on vast corpora of data and commonly referred to as foundation models (FMs)
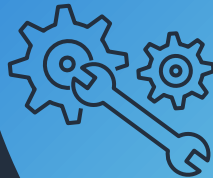
What is generative AI?

AI that can produce original content close enough to human generated content for real-world tasks

Powered by foundation models pre-trained on large sets of data with several hundred billion parameters

Tasks can be customized for specific domains with minimal fine-tuning

Applicable to many use cases like text summarization, question answering, digital art creation, code generation, etc.
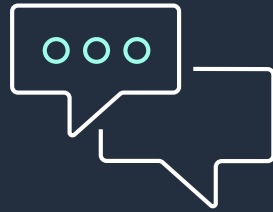
Reduces time and cost to develop ML models and innovate faster

Foundation model use cases

Productivity
Text generation

Chat
Virtual assistant

Summarization
Text extraction

Search

Code generation

Image generation

Image classification

What could go wrong?

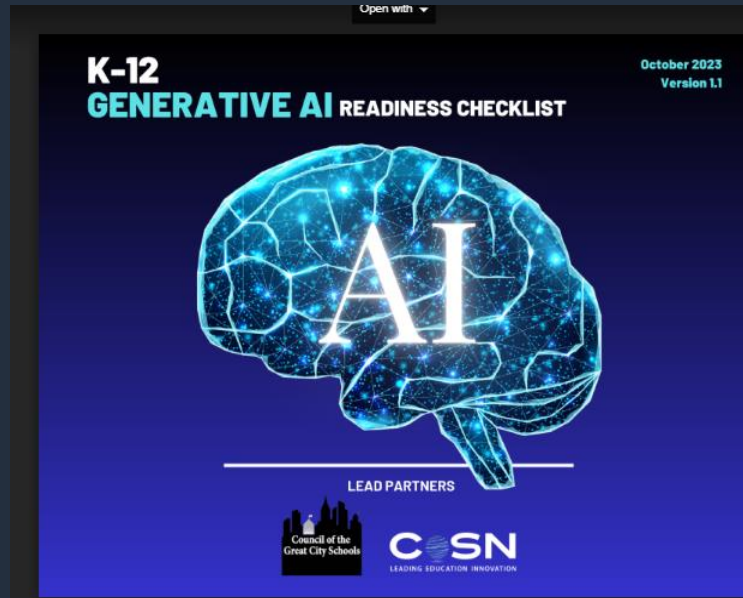| Hallucinations	Bias	Copyright and IP	Security and privacy
Answers that are factually incorrect, irrelevant, or nonsensical because of limitations in their training data and architecture	Answers that display discriminatory behavior resulting in prejudiced or unequal treatment of a particular group or groups	The rights of content creators from whom training data is collected remains uncertain and is currently being challenged	Some model providers use and store data for training purposes; entire end-to-end data pipelines require security and data privacy controls
"The world record for crossing the English Channel on foot is 15 hours"	"Generate a picture of a person cleaning" returns overwhelmingly women	Model creators getting sued for alleged improper use of photos	Engineers accidentally releasing source code by putting it into ChatGPT for debugging

Emerging Regulatory Environment

The importance of using AI responsibly

Consider how critical it is to use AI responsibly for reducing risks and deliver value comprehensively, at scale, while keeping the AI logic equitable and unbiased

Risks impacting organizations

Reputational impact
Poor organization
perception; erodes

Revenue loss
Diminished
credibility and trust

Regulatory repercussions
Legal penalty or
restrictions resulting from

"[Organizations] fail to focus on ethical, social, and regulatory implications, leaving themselves vulnerable to potential missteps when it comes to data acquisition and use, algorithmic bias, and other risks, and exposing themselves to social and legal consequences."

HBR's Year in Business and Technology: 2021
referencing McKinsey & Company article "Ten Red Flags Signaling Your Analytics Program Will Fail"
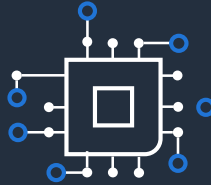
and people

legal ramifications

A multi-disciplinary problem

Economics

Moral philosophy

Technology

Law

Social science

Responsible AI is a complex, multi-disciplinary problem, blending requirements across a range of specialist fields

Although some organizations have begun to establish a basic awareness of the problems associated with responsible AI, few have access to the requisite skills or experience to tackle this problem in a comprehensive manner

Pillars of responsible AI

Value alignment
Systems should be designed and used in ways that align with the organization mission, social norms, and legal regulations

Inclusion
Inclusion of diverse and unique skills, experiences, perspectives, and cultural backgrounds

Data privacy and protection
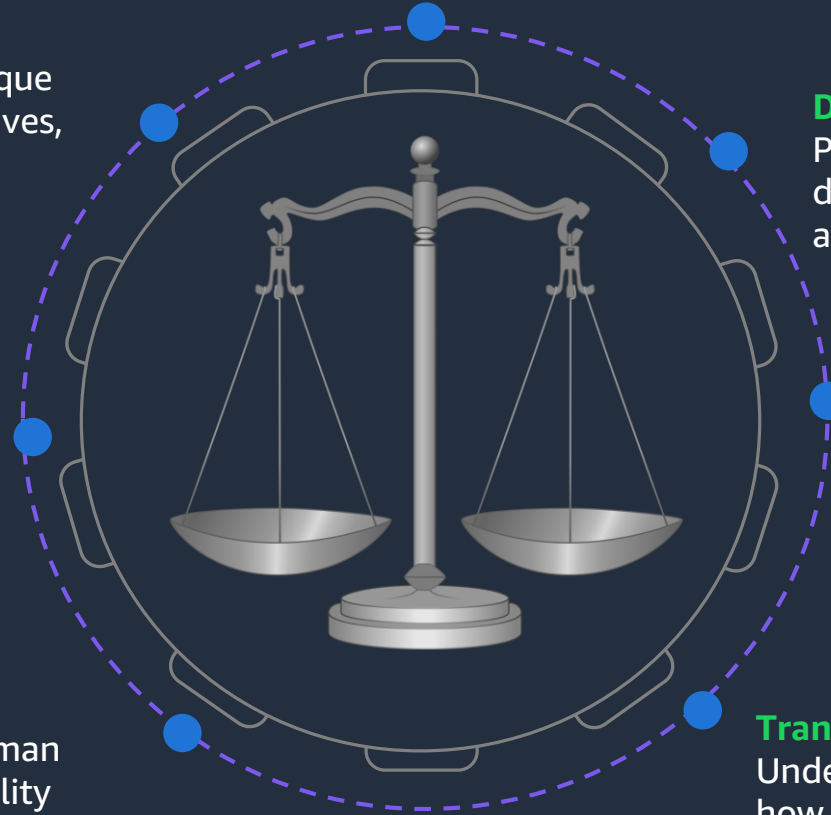Protects the quality and integrity of data used as well as its relevance, access, and processing

Training and education
Appropriate knowledge sharing and education to understand purpose, use, and impact

Fairness
Systems must be designed to minimize bias and promote inclusive representation

Accountability
Structured, maintaining human involvement and responsibility for design, development, decision processes, and outcomes
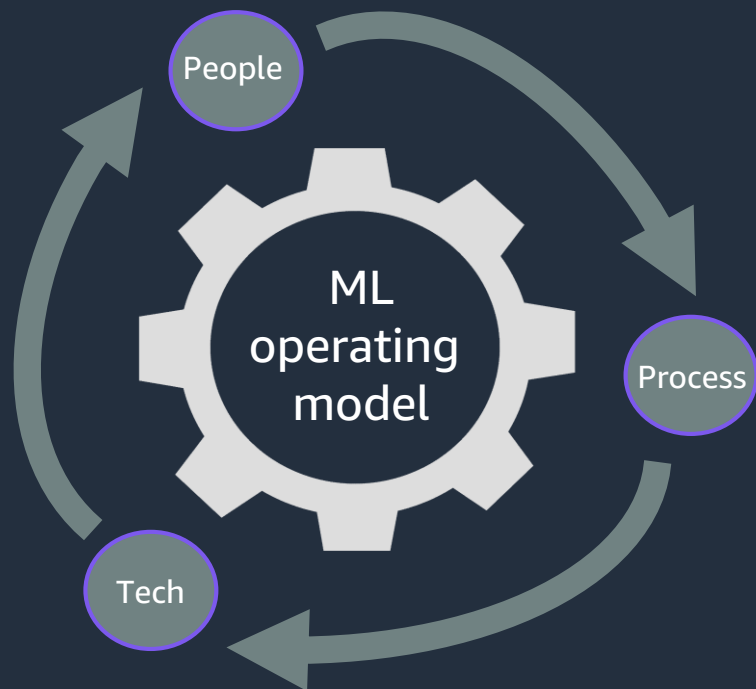
Transparency and explainability
Understanding how data is used and how decisions and outcomes are made understandable to a human

Inclusion matters

- Improves organization performance
- Fuels innovation
- Bridges gaps to address inadvertent access from other cultures
- Avoids more errors

Benefits of building responsibly

People

ML operating model

Process

Tech

Accelerate adoption

Institute appropriate governance structure

Align AI risk management with broader risk efforts

Develop people resources and skills

Build operational capability

Drive inclusive innovation

Value

Technological advancement must respect the rule of law, human rights, and dignity, as well as our shared values of inclusivity, privacy, and fairness

Thank you!

Sergio Ortega

sergioai@amazon.com

Please Provide Your Feedback

Step 1: Select <Your Track Name>
Step 2: Select <Your Session Title>