

The Dentons logo consists of the word "DENTONS" in a bold, white, sans-serif font, enclosed within a white arrow-shaped graphic pointing to the right. The background of the slide is a vibrant purple with a large, abstract, painterly graphic on the right side featuring shades of orange, teal, and blue.

Decoding Responsible AI Governance

ACC San Diego

Peter Stockburger

Dentons

peter.stockburger@dentons.com

619.595.8018

San Diego Office Managing Partner


Lead US Artificial Intelligence

Co-Lead Global Autonomous Vehicles

Grow | Protect | Operate | Finance

Discussion

- ❖ Setting the table - the promise and peril of AI
- ❖ How is AI being regulated globally?
- ❖ What is the concept of “responsible” AI governance?
- ❖ Responsible AI governance in action – contracts, vendor management, internal guardrails
- ❖ Questions throughout



A close-up photograph of peacock feathers, showing the intricate, overlapping scales of the feathers. The colors are vibrant, ranging from deep blues and greens to bright yellows and oranges, creating a shimmering, iridescent effect. The feathers are arranged in a dense, overlapping pattern, with each feather having a distinct, fan-like shape. The lighting highlights the fine details of the feather structure, giving it a textured, almost crystalline appearance.

The promise and peril of AI

What is AI?

A machine-based system that can make predictions, recommendations, or decisions influencing real or virtual environments.

What is machine learning?

A set of techniques that can be used to train AI algorithms to improve performance at a task based on data (e.g., it looks for relationships and patterns).

What is deep learning?

A subset of machine learning that relies on layered neural networks to engage in higher levels of processing.

What is generative AI?

A class of AI models that can emulate the structure and characteristics of input data in order to generate synthetic data, including audio, video, pictures, and other digital content.

An aerial photograph of terraced rice fields, showing a series of concentric, wavy green terraces that follow the contours of a hillside. A small, irregularly shaped pond with brownish water is situated in the upper right quadrant of the image. The overall scene is lush and green, with some brownish paths or roads visible between the terraces.

**Why is there so much attention
on AI now?**

AI is not new.

The rate of acceleration is.

AI could add between 17 and 26 trillion annually to the economy.

There is great promise.

There is also great peril.

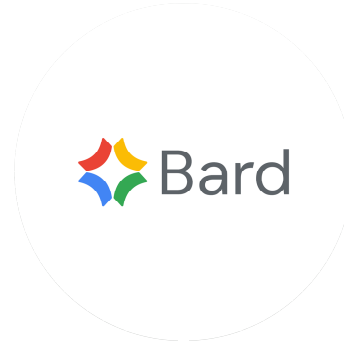
Generative AI Tools

Commercial Availability Is Booming

Chatbots



OpenAI
Chat GPT



Google
Bard



Meta
Llama



Anthropic
Claude



Microsoft
Bing

Enterprise



IBM
WatsonX



Microsoft
Copilot



Adobe
Sensei



ChatGPT
Enterprise



GitHub
Copilot X

The Promise and Peril

Multiple significant applications



Education



Employment



Healthcare



Insurance



Construction



Public Sector




The Promise and Peril

Multi-Faceted and Growing Risk

- **Accuracy**. Many models today generate inaccurate or invalid results.
- **Privacy**. Significant challenges exist relating to the types of data being processed. FTC investigations are ongoing. Global focus on data privacy risk.
- **IP Protections**. Multiple lawsuits are pending globally challenging LLMs on infringement of protected data. Using third-party tools without appropriate guardrails can also introduce significant risk.
- **Discrimination**. Any legal regime where discrimination is prohibited (e.g., employment, finance, housing) can present a legal risk.
- **Vendor Management**. You may only be as protected as your vendors are protected.
- **Human Safety**. AI used in medical care, housing, or other applications that have a physical component can introduce significant physical risk.

Hype Cycle for Artificial Intelligence, 2023



Plateau will be reached:

- less than 2 years
- 2 to 5 years
- 5 to 10 years
- ▲ more than 10 years
- ⊗ obsolete before plateau
- As of July 2023

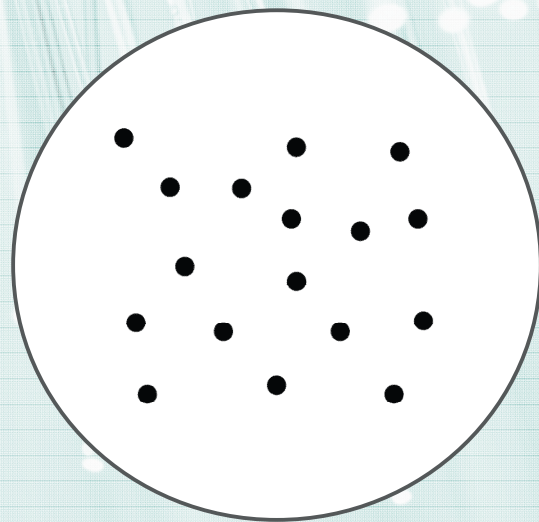
[gartner.com](https://www.gartner.com)

Source: Gartner
© 2023 Gartner, Inc. and/or its affiliates. All rights reserved. 2079794

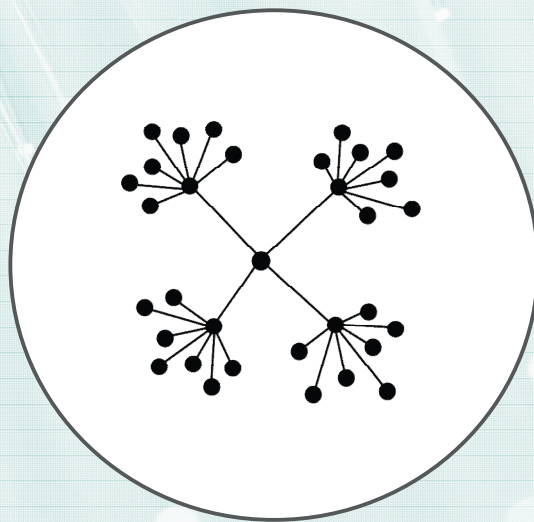
Gartner

AI will network, forming ecosystems

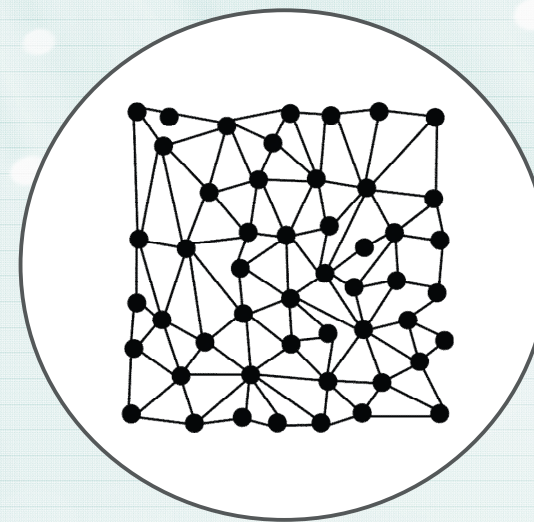
SILOED



NETWORK



ECOSYSTEM



AI will become more intelligent than we are

Artificial Narrow Intelligence (ANI)

ANI describes AIs that are good at a particular task at a level equal or better than a human being.

EXAMPLE

Virtual assistants, such as Siri or Alexa.



Artificial General Intelligence (AGI)

AGI is an AI that can perform any task that a human being can. This is what most of us think of when we think of AI.

EXAMPLE

David, the child-like android from the 2001 movie Artificial Intelligence.




Artificial Super Intelligence (ASI)

This is an intelligence that surpasses anything that humans can do.


EXAMPLE

Marvel's J.A.R.V.I.S. (Just A Rather Very Intelligent System)




AI is on the trajectory to become Autonomic


AUTOMATIC




AUTOMATED



AUTONOMOUS



AUTONOMIC



**How do you govern systems that are
on a path to governing themselves?**

Download the report at:

www.dentons.com



A satellite view of the Earth showing the Americas, with the text "How is AI being regulated today?" overlaid in white. The image shows the continent of North America on the right and South America on the left, with the Atlantic Ocean to the east and the Pacific Ocean to the west. The text is centered horizontally across the middle of the image.

How is AI being regulated today?

The Global Challenge

- There is a [growing consensus globally](#) about the need for coordinated AI regulation. What does that look like? Three current approaches to global AI regulation are forming.
- [Risk Based Approach](#). This is the model being adopted in the draft EU AI Act, and is being followed by Canada and Brazil.
- [Light Touch / Market Driven](#). This is the approach being followed by the UK and several states in the ASEAN region.
- [Central Management](#). This approach is being driven by China, with various regulations in place that apply to AI.
- [US](#). A moving target...



Federal & State Approach

Federal Approach

- No federal AI law...yet.
- Senate and House are holding multiple AI hearings and learning sessions.
- Multiple bills pending in the Senate and House.
- White House has been leaning in with an AI bill of rights, instruction to federal agencies, and a new Executive Order issued on 10.30.
- NIST AI Risk Management Framework is driving the conversation. New guidance from Executive Order.
- FTC, EEOC, DOJ, CFPB are enforcing AI misuse using existing legal structures.

State & Local Approach

- Since August, governors in CA, VA, WI, OK, PA, and NJ have announced executive orders around exploring AI in the state.
- The governments of Seattle, NYC, San Jose, and Santa Cruz County have issued independent policies and guidelines on the use of AI by local public employees, in NYC private entities.
- States like WA and CA are adopting “responsibility” frameworks. Others are forming task forces / working groups to study the use of AI.
- New bills to address specific industry that represents a “high risk” to the population (e.g., healthcare, finance, education, etc.)

A photograph of a slot canyon with smooth, undulating rock walls and a bright light source at the end of the passage. The walls are illuminated with warm, golden light, creating a dramatic and atmospheric scene. The text "What is 'responsible' AI?" is overlaid in white, bold font across the center of the image.

What is "responsible" AI?

Decoding Responsible AI

- States, regulators, and industry want to see organizations take a “**responsible**” or “**trustworthy**” approach to AI.
- **OECD, NIST, EU, and White House** have adopted “responsibility” frameworks.
- The **Dentons Responsible AI Framework** benchmarks responsibility against five key principles.
- The responsibility principles **are not a checklist** for all AI use, but rather a lens through which you can view AI risk.
- Think of a responsible AI governance framework as “brakes” to a race car.



Valid, Reliable, and Robust



Safe, Secure, and Resilient



Transparent, Explainable, Interpretable



Privacy Enhanced, Fairness



Accountable



Decoding Responsible AI Governance

Identifying Uses, Building Structure



Identifying Uses, Building Structure

- **Identifying Uses**. Know the who, what, and where before starting. Who is using AI within the organization? Why? What are the future use cases? **Map risk and opportunity**.
- **Committee Development**. Successful organizations identify an AI risk committee or group to drive the AI risk governance process. Ensure cross-functionality, and clear decision-making. Example – ticketing system.
- **Policy Development**. Build out a “responsible” AI policy, along with sub-policies for departments depending on risk. Develop an acceptable use policy in the immediate for generative AI risk. What are the standards for HR when deploying for job applicants? Programmers for coding? Marketing for creative problem solving? IT for operational software?
- **Keep The Human-in-the-Loop**. Keeping the human-in-the-loop is important to mitigate immediate risk across the AI spectrum. New York example.




Decoding Responsible AI Governance
Department Guidelines - Sample Vendor
Diligence Questions

Responsible AI In Action

Sample Vendor Diligence Questions

These questions could be sent directly as part of an AI risk assessment or tacked on to existing security or privacy questionnaires. Will this questionnaire be made a mandatory part of the procurement processes, in aligned with the responsible AI policy?

1. Is AI being used by the vendor as a component of the services being delivered? If so, how?
2. What is the underlying technology architecture supporting the vendor's use of AI? Is it powered by another company's AI, built by the vendor, or a hybrid?
3. Are the models the vendor is using open source or proprietary?
4. What types of data will be processed by the vendor through your AI models? For example, will the vendor be using our sensitive data to train their AI models?
5. Does the vendor have processes for identifying and mitigating risk within the AI systems they use to deliver the services?
6. How was the AI trained? On what data sets? How does the vendor continue to train and fine tune the AI model used?
7. How does the vendor improve its AI over time? How is the AI updated?
8. How does the vendor ensure the AI that is being used to provide services does not result in unlawful or discriminatory outcomes?
9. How does the vendor ensure the AI does not produce incorrect or misleading results?
10. How can the vendor ensure transparency around how the AI generates its outputs? Are there model cards available?
11. How does the vendor ensure security and privacy by design within the AI is contemplated?
12. Does the vendor have any internal policies, governance procedures, or standards regulating its use of AI? If so, please provide.



Decoding Responsible AI Governance
Department Guidelines - Sample Vendor
Agreement Language

Responsible AI In Action

Sample Vendor Contract Language

1. Vendor has established and will maintain standard operating policies and procedures to ensure Vendor's creation, deployment, or use of artificial intelligence (AI) to provide the Services is done so in a commercially responsible and trustworthy manner ("**AI Governance Procedures**"). Vendor's AI Governance Procedures shall comply with all Applicable Laws, and shall align with industry best practice, including but not limited to the AI Risk Management Framework promulgated by the United States Department of Commerce's National Institute for Standards and Technology, as amended from time to time.
2. Vendor's AI Governance Procedures shall include, at a minimum, a repository of written policies and procedures, a formal review and approval process, a revision management process, a change control process, a risk assessment process, testing processes, acceptable tolerances for applicable AI models, decommissioning processes, ongoing monitoring and periodic review of models, and other appropriate components.
3. Vendor shall provide Company with appropriate documentation upon request, including where applicable information providing: (1) basic information about how the AI was developed, model date, version, type, and architecture details; (2) basic information about how the AI is trained, parameters used in the AI, and how the AI will be processing Company data; and (3) any other appropriate information to provide Company with sufficient information to analyze the risk associated with such AI.
4. Vendor shall not use AI to generate any content, text, audio, visual, or other material for Company without written prior approval.
5. Vendor will not: (1) use or disclose any information or data collected regarding Company's use of Vendor's Services to create additional data, datasets and/or use such data to teach, enhance, improve, or add to any algorithm, model, or similar automated tool owned or operated by Vendor, or make any such data available to a third party or be used for a commercial benefit to Vendor without prior written approval of Company.
6. Vendor shall ensure any AI used to provide the Services is designed and deployed with capabilities enabling the automatic recording of events ("**AI Logs**") while the AI is operating. Vendor will provide Company copies of any AI Logs relevant to the Services upon reasonable written request.

The background features a close-up of peacock feathers with iridescent colors ranging from gold and green to deep blue and purple. A large, semi-transparent purple shape with a rounded right edge is overlaid on the left side of the image, serving as a background for the text.

Decoding Responsible AI Governance

Key Takeaways



Key Takeaways

- **Harness The Promise, Mitigate The Peril.** Don't be afraid to leverage AI to drive market share, revenue, efficiencies. Watch out for the peril. Recognize the peril, build in guardrails.
- **Map, Manage, Cultivate.** The first step in the journey is often mapping existing use cases, future use cases, and where AI may have the most benefit to the organization. From that knowledge base, developing guardrails and management of risk is paramount. Cultivating a culture of responsible AI governance is also critical.
- **Mitigate Immediate Risk, Build The “Brakes” For The Racecar.** The immediate risks around rogue uses of AI, or tools that don't comport with legal risk should be addressed. But go deeper, and build the brakes for the race car of the future.
- **Stay Informed.** Commercial tools, terms, and legal standards are constantly shifting. Implementing a responsible framework now will allow you to be nimble in the future.

Download the report at:

www.dentons.com



DENTONS

Thank you

Across over 80 countries, Dentons helps you grow, protect, operate and finance your organization by providing uniquely global and deeply local legal solutions. Polycentric, purpose-driven and committed to inclusion, diversity, equity and sustainability, we focus on what matters most to you. www.dentons.com.

© 2023 Dentons. Dentons is a global legal practice providing client services worldwide through its member firms and affiliates. This publication is not designed to provide legal advice and you should not take, or refrain from taking, action based on its content. Please see dentons.com for Legal Notices.