# Keeping Pace with AI Developments as the Legal and Technological Landscape Quickly Evolves

***Prepared for The Association of Corporate Counsel***

**Duane Pozza & Kat Scott (Wiley)**
**Jeff Dietz (Capital One)**

***September 14, 2023***

# Agenda

- The Basics of AI
- The Evolving Legal and Regulatory Landscape
- Key AI Risks and Principles
- Risk Management Approaches to AI Governance
- Practical Tips and Best Practices

# The Basics of AI

# AI and Its Beneficial Uses

## What is AI?

- Involves capacity to learn
- Generally based on large data sets
- Generative AI / large language models
- Deep learning & deep neural networks

"AI is the ability of a computer system to solve problems and to perform tasks that would otherwise require human intelligence."
*Interim Report of National Security Commission on AI (2019)*

## Potential Use Cases

- Predictive analytics
- Pattern recognition in large data sets
- Cybersecurity and fraud defense
- Voice interaction
- Detecting deepfakes
- Automated diagnostics and maintenance
- Object recognition and imagery analysis
- ***Buzzwords:*** AI, ML, algorithms, natural language processing, generative AI

# Generative AI

## What Is Generative AI?

- "Generative AI is a broad label that's used to describe any type of artificial intelligence (AI) that can be used to create new text, images, video, audio, code or synthetic data." *Techopedia*
- Examples
  - ChatGPT
  - Codex
  - Dall-E
  - Llama-2

## Awareness Trends

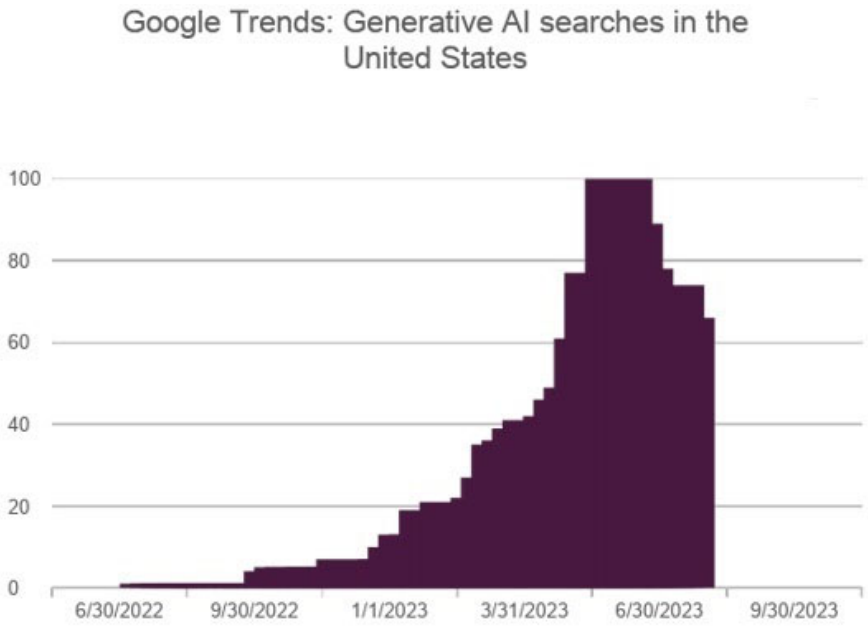

# The Evolving Legal and Regulatory Landscape

# Key Government Actors and Workstreams

## Federal

- White House
- Federal Trade Commission (FTC)
- NTIA
- Financial Regulators
- National Institute of Standards and Technology (NIST)
- Congress

## States

- Legislation targeting specific AI use cases
- State omnibus privacy laws & California's CPRA rulemaking

## International

- European Commission: Proposed AI Act
- Organization for Economic Cooperation and Development (OECD) AI Principles
- Other International Ventures

# Key Federal AI Developments

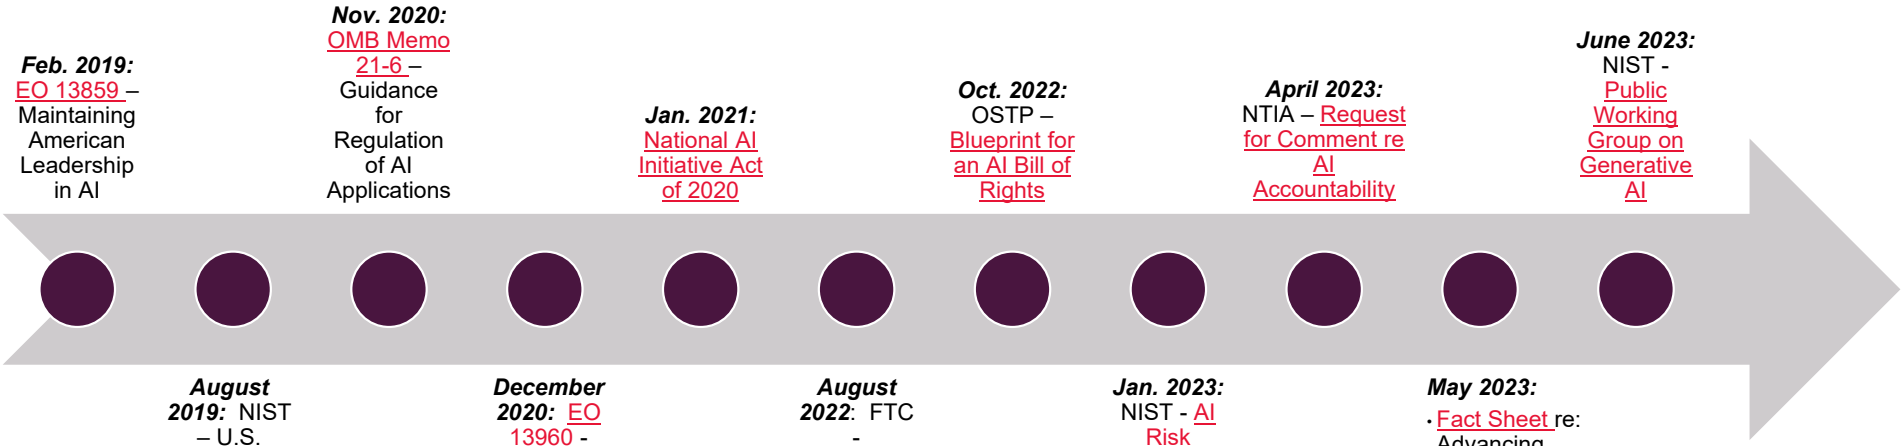

# Examples of AI-Specific Laws

## State

- Omnibus Privacy Laws & Regs
  - Opt-out rights for certain automated processing/profiling
  - Data protection assessment requirements for certain automated processing/profiling
  - Evolving expectations in CA
- Targeted AI Laws
  - The New York City AI Ordinance deals with use of AI in hiring
  - The California Bot Law creates chatbot disclosure requirements
  - CA Deepfake Political Ad Law covers the use of deepfakes in political ads
  - The Illinois AI Video Interview Act deals with use of AI analysis of video interviews
  - Laws that have established councils/commissions to study AI, e.g., Alabama and Illinois

## Federal

- National AI Initiative Act of 2020
- Key Examples of Proposed Legislative Approaches
  - Federal privacy legislation (ADPPA)
  - SAFE Innovation Framework for AI (Sen. Schumer)
  - Legislation focused on federal government and contractor use of AI, e.g.,
    - Oversee Emerging Technology Act (S. 1577)
    - Assuring Safe, Secure, and Ethical Systems for AI (Assess AI) Act (S. 1356)
- Multiple hearings focused on AI, e.g.,
  - May 2023 - Oversight of A.I.: Rules for Artificial Intelligence
  - June 2023 - Artificial Intelligence and Human Rights

# Examples of Generally Applicable Laws with Impacts on AI

- April 2021 – FTC Blog Post providing guidance and signaling enforcement regarding discrimination in AI
  - Fair Credit Reporting Act (FCRA)
  - Equal Credit Opportunity Act (ECOA)
  - Section 5 of the FTC Act
- April 2023 – Joint Statement (FTC, DOJ, CFPB, EEOC) outlining a commitment to enforce their respective laws and regulations to promote responsible innovation in automated systems

# Looking Ahead on AI Policy

- AI legal frameworks are developing in real time
  - Lawsuits and regulatory actions look backwards – so what companies do now will be under scrutiny in the future
- Companies also must be ready for a patchwork of laws
- There are many open questions:
  - Will AI regulation happen?
    - If so: *Who* will regulate AI?
    - Fragmentation among federal regulators *and* potentially among states is on the horizon
  - Will regulation be prescriptive?
    - White House Office of Science and Technology Policy is seeking input on a National AI Strategy, and one question it has is: "Can inspiration be drawn from analogous or instructive models of risk management in other sectors, such as laws and policies that promote oversight through registration, incentives, certification, or licensing?"
    - Should foundation models be treated differently?
    - Or will new regulations be outcome-oriented?
  - Will regulation be technology-neutral?
    - Principles-based regulations – like UDAAP laws – do not need to focus on a particular technology

# Key AI Risks and Principles

# Trustworthy AI Principles

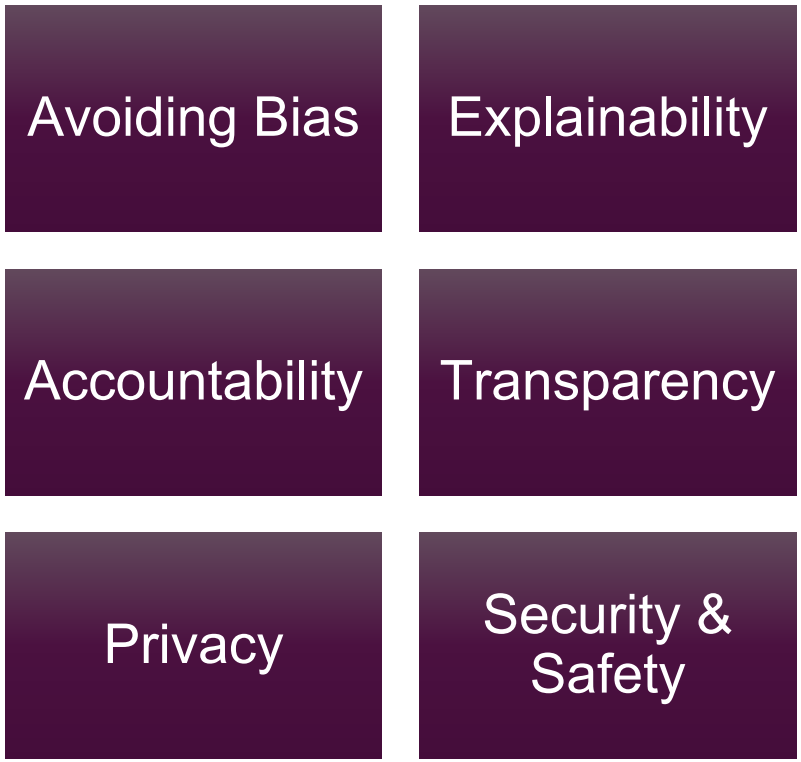


- There is significant consensus forming around key principles for "trustworthy" and "responsible" AI
- Examples:
  - NIST
  - OECD

# Avoiding Bias

## Key Issues

- Mitigating harmful AI bias
- Existing laws
  - Anti-discrimination laws and regulations apply (e.g., ECOA)
  - Sector-specific laws: credit, housing, employment
- Biased data sets vs. algorithmic bias
- Third party testing proposals
- Data can be used not only to avoid bias but to promote fairness

## Resources & Best Practices

- ***NIST research***
  - Towards a Standard for Identifying and Managing Bias in Artificial Intelligence
  - Managing AI/ML Bias in Context
- ***FTC guidance***
  - Don't discriminate based on protected classes.
  - Focus on inputs as well as outcomes.
  - Ensure AI models are validated and revalidated to ensure that they work as intended, and do not unlawfully discriminate.
  - Ask questions before using algorithm.
    - How representative is your data set?
    - Does your data model account for biases?
    - How accurate are your predictions based on big data?
    - Does reliance on big data raise ethical or fairness concerns?

# Explainability

## Key Issues

- When do you need to explain how AI reached a conclusion? And can you?
- Different explanations may be appropriate for different audiences
- Explanability requirements in practice:
  - Credit decisions and adverse action notices

## Resources & Best Practices

- NIST's Four Principles of Explainable Artificial Intelligence:
  - Explanation
  - Meaningful
  - Explanation Accuracy
  - Knowledge Limit

# Accountability and Transparency

## Accountability

- Who is accountable for making sure nothing goes wrong with AI?
- Multiple participants in AI lifecycle: software developers, product developers, downstream operators
- Human oversight / “human-in-the-loop” is a key concept when dealing with algorithms
- Third-party vendor oversight is critical
- Even if legal responsibilities are in flux, negative outcomes will result in consequences

## Transparency

- When do you need to disclose that AI is being used?
- Particularly critical for audio or visual content – e.g., provenance or watermarking systems
- State law examples
  - The California Bot Law creates chatbot disclosure requirements.
  - The Illinois AI Video Interview Act deals with use of AI analysis of video interviews.

# Privacy and Security & Safety

## Privacy

- Are there privacy risks unique to AI?
- Federal Trade Commission (FTC)
  - Enforcement: Focus on deception – what is said in privacy policy?
  - Rules: Commercial Surveillance and Data Security ANPR
- State privacy laws
  - Omnibus privacy laws in VA, CO, and CT create "opt-out" right for automated profiling in furtherance of legally significant decisions
  - Forthcoming rulemaking in CA to establish access and opt-out rights
- What's next? Congress?

## Security & Safety

- Are there security and safety risks unique to AI?
- AI faces a range of cybersecurity threats – which pose both operational risk and legal risk in case of an incident
- Physical safety issues
- Reputational risks and IP-related issues
- Certain controls – e.g., internal access controls to personal data – serve critical functions for both cybersecurity and privacy

# AI Risks/Principles in Practice

- Potential tension between principles (e.g., explainability and security)
- Tension with traditional privacy principles (e.g., purpose specification and avoiding secondary use)
- Standards and measurements continue to be developed
- Will there be clear rules and guidance on issues like bias?

# Risk Management Approaches to AI Governance

# NIST AI Risk Management Framework (AI RMF)

- The AI RMF is a risk management resource for organizations designing, developing, deploying, or using AI systems
  - Provides voluntary guidance and risk management practices
- Frames AI related risks
  - 7 “trustworthy AI characteristics”
- Outlines the AI RMF “Core”
  - 4 “Functions,” along with “Categories” and “Subcategories” that help organizations address AI system risks as a practical matter

# AI RMF: Trustworthy AI Characteristics

- Valid and reliable
- Safe
- Secure and resilient
- Accountable and transparent
- Explainable and interpretable
- Privacy enhanced
- Fair – with harmful bias managed

# AI RMF: 4 Functions

- Govern: recommendations concerning high level processes and organizational schemes for fostering a culture of risk management throughout an organization
- Map: recommended methods for contextualizing and identifying AI system risks
- Measure: recommendations for assessing, analyzing, and tracking identified AI risks
- Manage: recommendations for allocating resources and prioritizing AI system risks

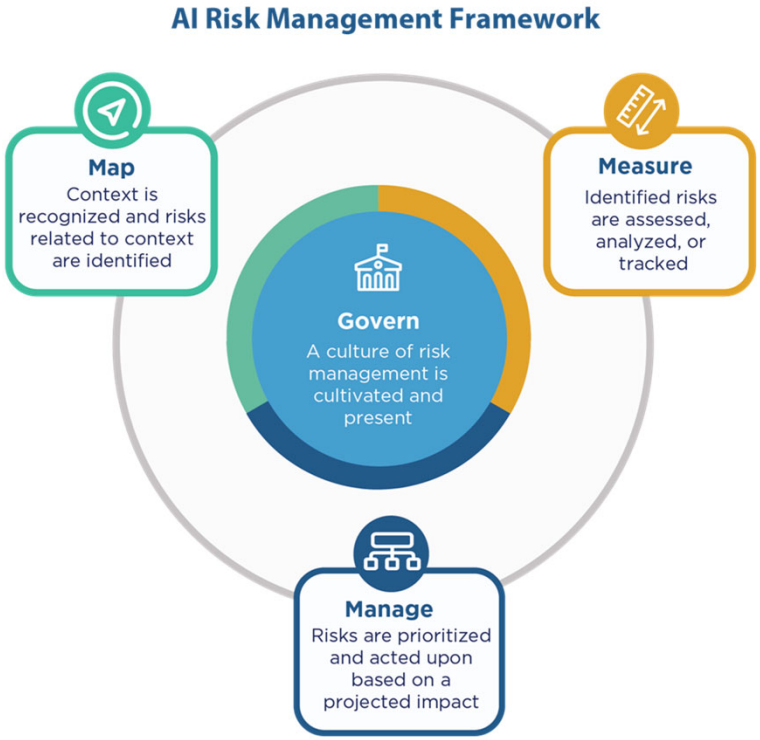

# Generative AI: What's All The Chatter About?

# Unique Considerations for Generative AI

## Background / Frame of Reference: ChatGPT

- **Based on OpenAI's GPT**: Generative Pretrained Transformer
- **GPT-3**: Jun '20, A large language model ("LLM") trained on vast quantities of text (100s of Billions of words), self-supervised learning to predict "what might come next?" from a sequence.
- **ChatGPT**: Nov '22, Chat bot on GPT-3.5, fine-tuned with reinforcement learning from human feedback.
- **GPT-4**: Feb '23, Next-generation model with dramatic quality improvements.

## Compare "Standard" AI / ML to Generative AI

- **Training**: Supervised v. Self-Supervised
- **Training Data Size**: Big v. Enormous
- **Training Data Set**: Known & Labelled v. ?
- **Use Case**: Specific v. Multiple

# Unique Considerations for Generative AI

## Legal Issues

- **Data Protection** (Who Can See / Own / Use Your Data?)
- **Intellectual Property** (Litigation; Ownership)
- **Privacy** (Training Data)
- **Accuracy** (Hallucinations)
- **Fairness** (Explainable?)
- **Cyber Risks** (e.g., deepfakes and social engineering)
- **Reputation Risks** (“creepy” factor? notification?)
- **Emerging ESG** concern around social disruption (ethics)

# Practical Tips and Best Practices

# Developing Best Practices

- Adopt risk-based approach to identify and proactively mitigate risks
  - Ensure review processes throughout lifecycle of AI
  - Establish clear responsibility and accountability within organization
  - Ensure both technical and non-technical personnel are engaged
  - Establish incident response protocols and training
- Consider corporate or industry-wide principles and policies
  - Establish generative AI policies
- Pay attention to “informal” guidance from regulators like the FTC

# Questions? Contact Us.

**Jeff Dietz**
**jeffrey.dietz@capitalone.com**

**Duane Pozza**
**dpozza@wiley.law**

**Kat Scott**
**kscott@wiley.law**

Capital One

wiley