


Deloitte.

Artificial Intelligence (AI)

August 13, 2023




On the agenda today

1 | **What's happening in the world of AI**


2 | **Generative AI**

3 | **Trustworthy AI**

4 | **Art of the Possible**



What's happening in the World of AI



AI is inherently about augmenting humans with machines to reach greater heights

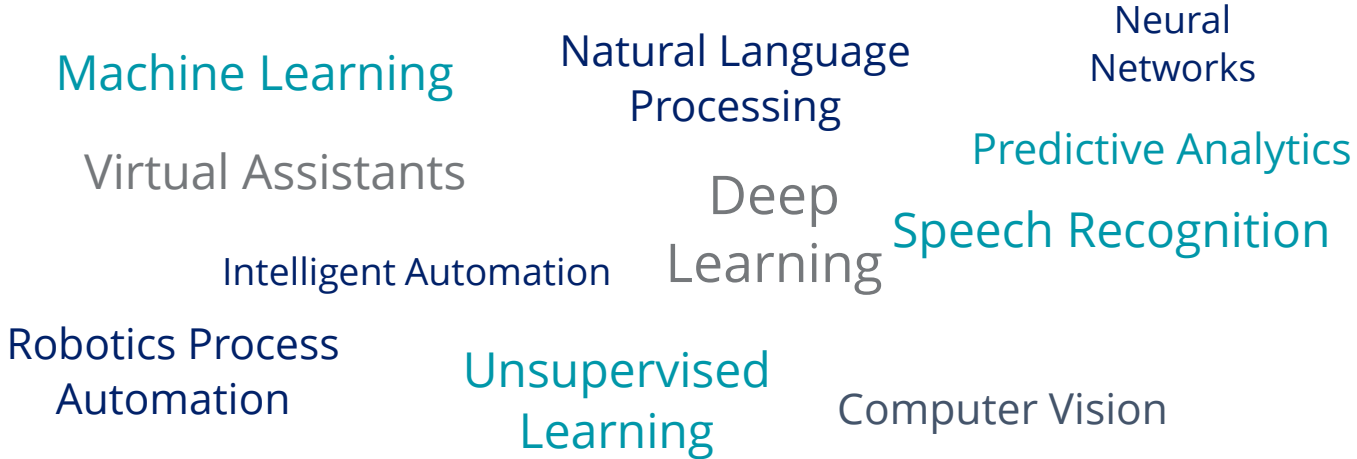
AI strives to mimic the way humans perceive information, devise insights based on experience, and make decisions accordingly



OXFORD DICTIONARY DEFINITION:

ARTIFICIAL INTELLIGENCE (AI) is the theory and development of computer systems able to perform tasks normally requiring human intelligence

AI encompasses many technologies that work together to build innovative solutions that transform society and business...



State-of-the-art and emerging capabilities in AI

What is happening in AI today

Generative AI	Natural Language Processing (NLP)	Computer Vision	Near real-time inferencing
<ul style="list-style-type: none"> AI systems can now compose text, audio, and images to a sufficiently high standard that humans have a hard time telling the difference between synthetic and non-synthetic outputs for some constrained applications of the technology (E.g., Deepfakes) Researchers are investing in technologies for detecting generative models; the Deepfake Detection Challenge data indicates how well computers can distinguish between different outputs 	<ul style="list-style-type: none"> Rapid progress in NLP has yielded AI systems with significantly improved language capabilities that have started to have a meaningful economic impact on the world Progress in NLP has been faster than that for the benchmarks that can test them; this can be seen in the rapid emergence of systems that obtain human level performance on SuperGLUE, an NLP evaluation suite developed in response to earlier NLP progress overshooting the capabilities being assessed by GLUE 	<ul style="list-style-type: none"> Computer vision has seen immense progress in the past decade, primarily due to the use of machine learning techniques (specifically deep learning) Increased investments by organizations in computational resources to train computer vision systems 	<ul style="list-style-type: none"> Computational infrastructure performance advances with graphics processing units (GPUs) and tensor processing units (TPUs) are providing near real time feedback on model accuracy and performance Across many industries/use-cases, AI is being inserted into operational activities to significantly improve performance

Source: Stanford University's Artificial Intelligence Index Report

Top AI Tech Trends



Move **from rules- and heuristic-based systems to pattern- and behavior-based** analyses and detections

- Unsupervised learning
- Ensemble models



Predictive AI – move AI models and apply intelligence in real time for faster detection of changes in operating environments and remedial actions




Increased detection of abnormal behavior and suspicious activity through **anomaly detection**

- Connecting the dots between various data sources; social media and digital profiles, personal identifiable information (PHI)/protected health information (PHI), Litigation and Regulatory Sanctions

All product names mentioned in this document are the trademarks or registered trademarks of their respective owners and are mentioned for identification purposes only. Deloitte & Touche LLP is not responsible for the functionality or technology related to the Vendor or other systems or technologies as defined in this work product.

What is Generative AI



Overview of Generative AI

Generative AI can provide businesses with new opportunities to improve their products and services, automate repetitive tasks, and create new and innovative customer experiences

WHAT is Generative AI | artificial intelligence that creates **original content across various modalities** (e.g., text, images, audio, code, voice, video) that would have previously taken human skill and expertise to create

HOW does it work | Generative AI is powered by **foundation models** such as OpenAI's GPT-3 and NVIDIA's Megatron, which are trained on **vast amounts of data and computation** to perform a broad range of downstream tasks

WHY now | innovations in **machine learning** and the **cloud tech stack**, coupled with the **viral popularity** of publicly released applications, such as **ChatGPT** and **DALL-E2**, have propelled Generative AI into the zeitgeist

WHO is involved | **Big Tech** is building—and enabling access to—foundation models; **start-ups** are developing user applications on these underlying models; and **companies** are beginning to adopt

BUSINESS IMPACT | the **marginal cost of producing initial versions of knowledge-intensive content**—such as IT code, marketing copy, and creative design—**can fall to ~zero**

Recent News

- 2023-04-13: [AWS](#) are enabling creation and scaling of generative AI applications like chatbots, text, and image generation through Bedrock, a platform to customize pre-trained models with enterprise data and deploy them using AWS tools
- 2023-03-23: [OpenAI](#) enhanced ChatGPT with plug-ins support, allowing third-party services to interact with the bot and provide additional contextual information
- 2023-03-21: [NVIDIA](#) established NeMo and BioNeMo AI Foundations Cloud Services to develop, refine, and run domain-specific models for AI applications. NeMo service assists developers in customizing large language models, whereby BioNeMo service assists researchers in predicting and generating molecules, proteins, and DNA


All product names mentioned in this document are the trademarks or registered trademarks of their respective owners and are mentioned for identification purposes only. Deloitte & Touche LLP is not responsible for the functionality or technology related to the Vendor or other systems or technologies as defined in this work product.

Generative AI capabilities

Generative AI can produce a wide range of outputs depending on the specific application and type of data that is needed. Here are some common output types that are applicable to business

Text

Prompt: Explain the business impact of Generative AI in 50 words



Video

Prompt: Create a video of a teddy bear painting a portrait



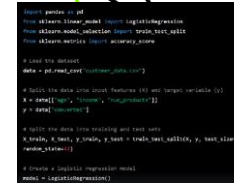
Image

Prompt: A bowl of soup that is a portal to another dimension as digital art



Code

Prompt: In python, code a program that predicts the likelihood of customer conversion



Audio

Generative AI-powered customer service agents

- Phone / Voice
- Multi-Lingual
- Multiple Tasks
- Empathy / Humor
- Ensures Resolution
- Systems Integration

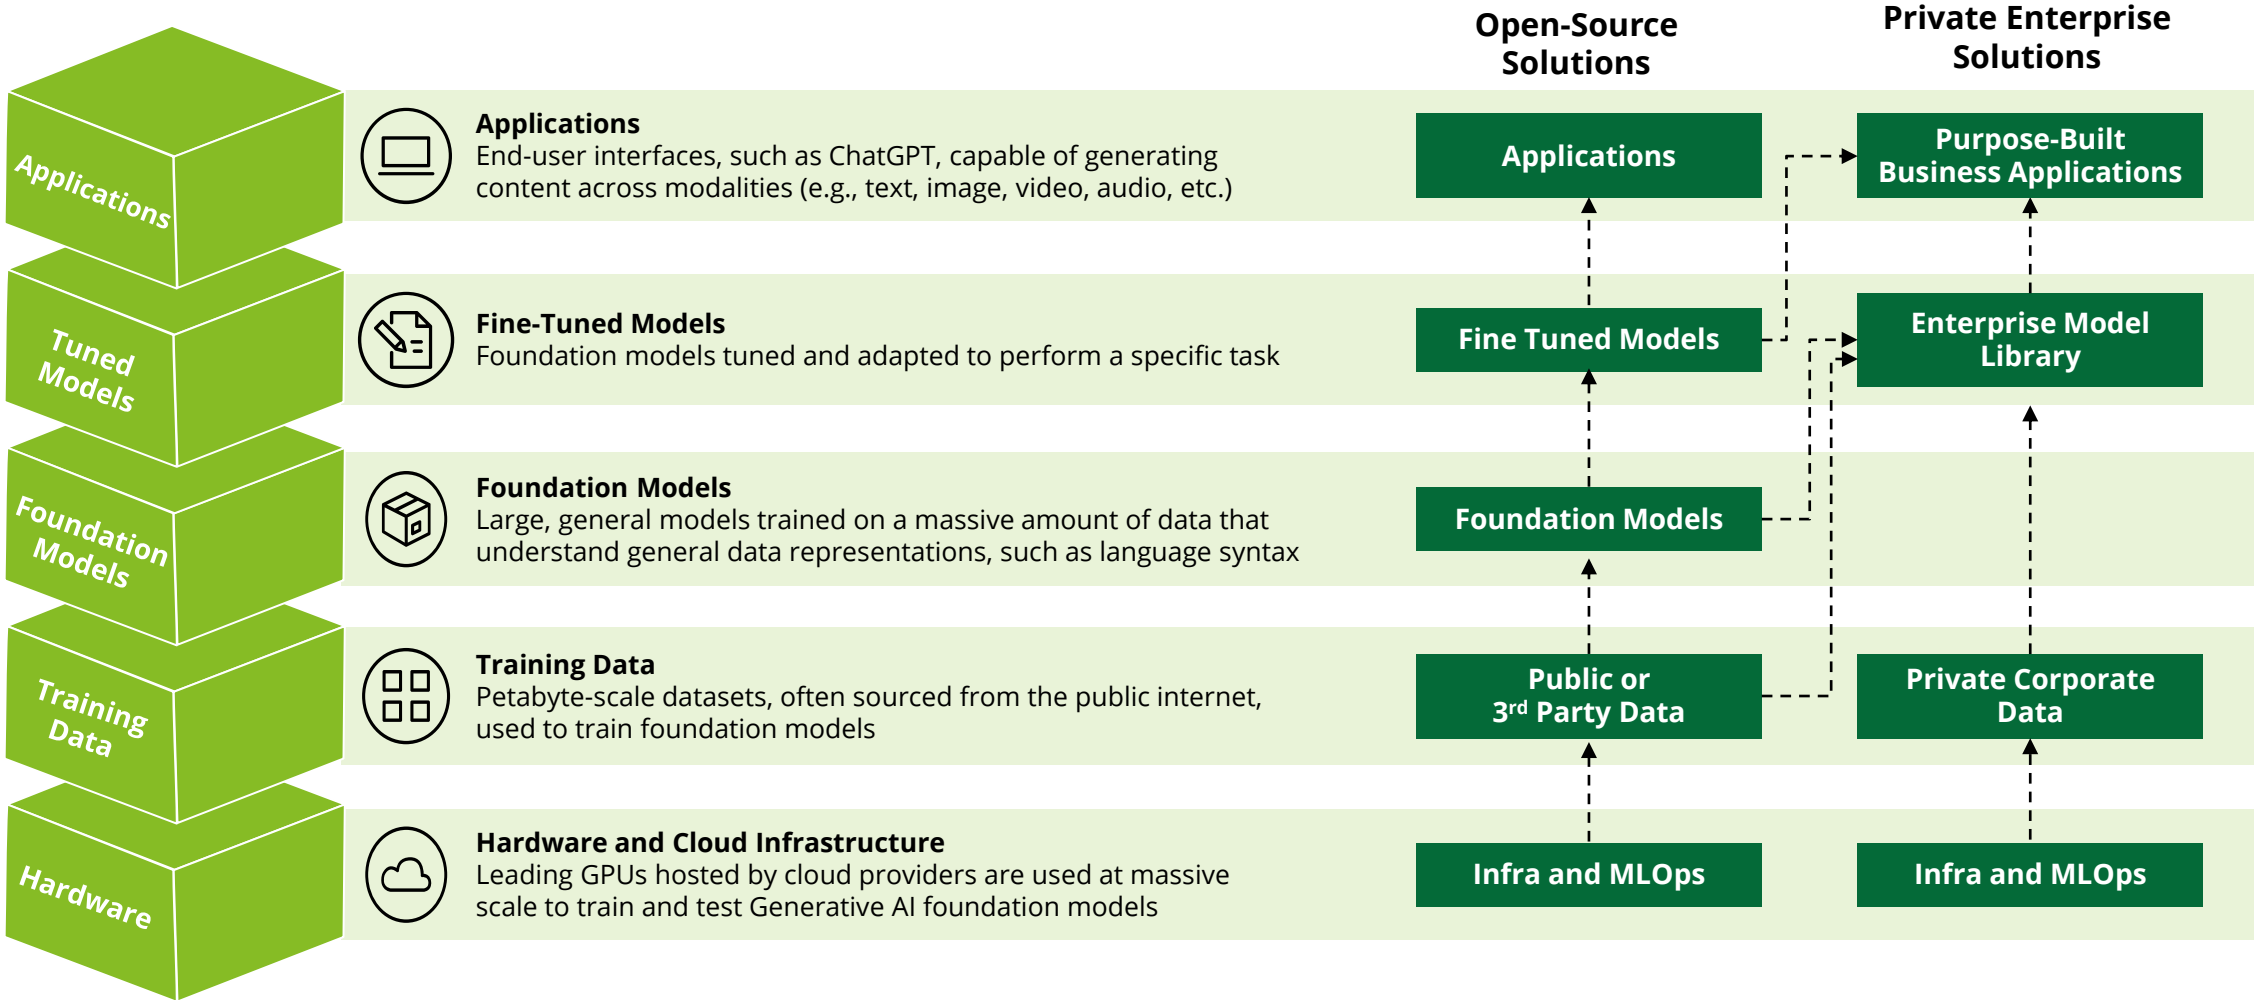
3D

Prompt: A beautiful dress made from garbage bags, on a mannequin. Studio lighting, high quality, high resolution




Generative AI components for an enterprise

Generative AI solutions rely on the latest advancements in hardware and AI combined with dedicated 'human-in-the-loop' training programs to ensure the models create output consistent with the targeted tasks.



How do we measure and mitigate AI Risk

...across the three lines of defense?




Generative AI comes with risks and limitations

There are several limitations to consider when using Generative AI

Bias in; bias out. If the training data is biased (e.g., over/under-representation of a population cohort, sexism, racism), then outputs generated could exhibit biases as well. Bias reductions in the training data and/or human supervision during model training is needed

Bias



Foundation models generally offer a pay-as-you-go billing mechanism, and the cost per use of sophisticated models is materially significant. Fine tuning the biggest model and running large documents through several times could easily run up a bill of tens of US \$1000s

Cost



Is the AI being used in a manner consistent with the purpose of the overall exercise? Is a human being brought into the loop to decide whether the AI's suggestion needs adjustment before actual use? Submitting an AI-generated essay for a high-school assignment may not be ethical

Ethical Use



Models might output facts that are factually false. Sources and citations are unavailable for many models. Users should be conscious that outputs could be inaccurate and should perform due diligence to validate generated content.

Hallucination



SaaS-AI companies may save some or all of prompt payloads for future training. Therefore, confidential data will be used to train future versions of the base model – how will this affect your organization's competitiveness in the market?

IP Protection



It is critical to proactively minimize risk from malicious behavior on the network to maintain operations and customer trust. For example, a customer service bot revealing confidential information to a hacker either by prompt or unintentionally

Malicious behavior



Foundation Models (e.g., GPT-3) are comprised of billions of parameters (model size) and trained on petabytes of data. In theory, the larger the model, the better the output. Foundation Models take time to produce outputs, which may limit real-time use cases

Model Performance




SaaS-AI companies require to submit text as a payload to users' API call. The data could be crossing borders. Is this in accordance with data privacy laws and with a company's policies? Cloud service providers (e.g., AWS, GCP,) offer market-leading controls to manage data privacy of Foundation Models

Privacy




Models are good at understanding text but they struggle when the data are in irregular formats, or when the position of the text on the page (e.g., info graphic, PPT slide) is relevant to the context and understanding. Other emphasis generators such as bolded text, font color, etc. don't play a role yet

Text Formatting

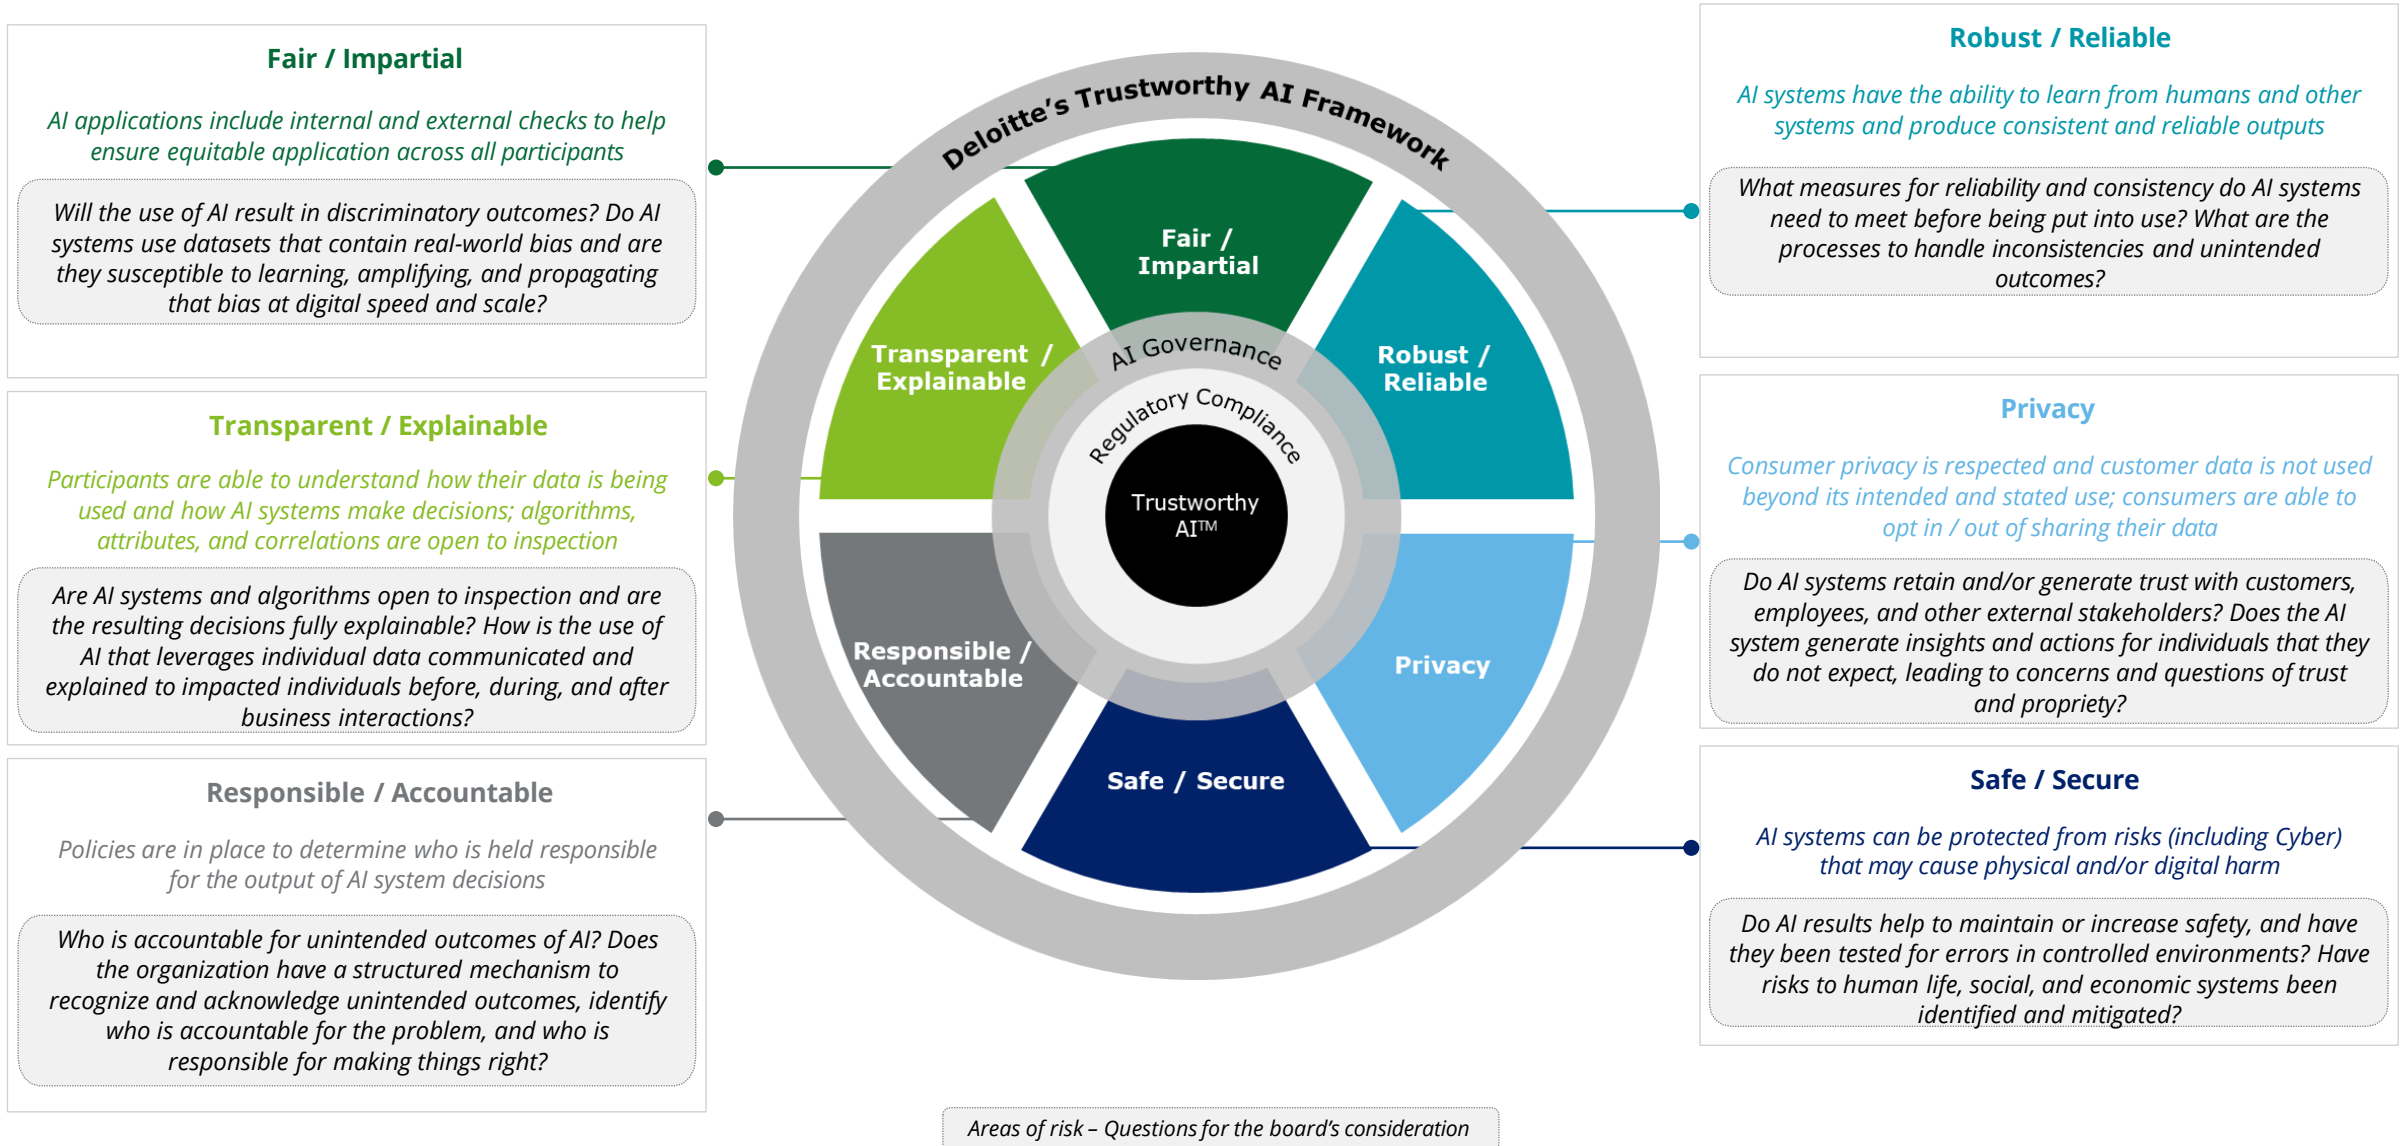


Most models have a 2k token size limit. Some larger ones can process 4k tokens in a single call. 2k tokens are approximately 2-2.5 pages. This limit makes it difficult to process larger documents

Token Size Limits



Deloitte's Trustworthy AI™ framework is a first step in diagnosing the trustworthiness of AI solutions



Elements of an AI risk governance and operating framework

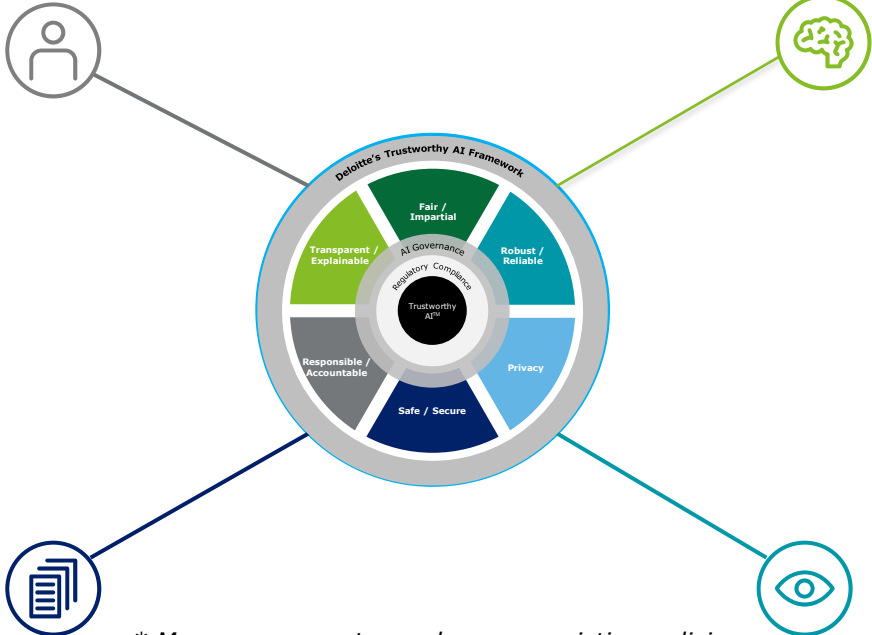
Implementing an AI/ ML governance framework can leverage existing risk management processes and stakeholders

Roles & Responsibilities

- LoBs (Users of AI/ML)
- AI / ML Code or Model Owners
- Implementation Leads
- Development Leads
- Data Governance Managers
- Validation and Challenge Leads
- Peer Reviewers
- Governance Officers

Policies & Procedures

- Enterprise AI Governance Policy
- Data use, privacy and security policies
- Procedures (Development, etc.)
- AI system and process controls framework
- Documentation Templates
- Testing Playbook



** Many components can leverage existing policies, processes, and stakeholder responsibilities. For example, AI that is used for machine learning models can leverage existing Model Risk Management (MRM) cadences, if updated appropriately*

It is critical that AI / ML governance integrate with related existing frameworks (data governance, cyber risk, MRM, etc.)

Process & Technology

- AI /ML Inventory process
- Risk Identification and Profiling
- Development and use protocols
- Monitoring and Effective Challenge
- Reporting to Sr. Mgmt and stakeholders
- Human vs. automated controls
- Engines, platforms, and supporting technology

Cross Function & Compliance

- Privacy
- Data Governance
- Information Security
- Global Entity Management
- Legal
- Procurement (Vendor Risk Management)
- Business Continuity

Building a Trustworthy AI environment

Changes driven by the adoption of AI call for re-imagining governance processes, mechanisms and operational controls.



Key questions for organizations:

- Who is responsible for managing risks related to automated processes?
- How to adapt governance processes to account for the complexities of automation?
- How to respond quickly to automation-induced errors?

Enabling transformation across governance processes

First Line of Defense

- Enable increased FLOD testing by model owners using stress testing and continuous testing to automatically validate and monitor models
- Establish thresholds and identify key data points to better define model parameters and refine model development processes

Second Line of Defense


- Review in-depth model documentation and see high-level overviews of all models through development of governance-dedicated dashboards
- Establish an AI risk strategy including forward thinking risk taxonomies, setting AI risk appetite, identifying KRIs and formulating testing strategies
- Enable real-time notification of model issues and root cause analysis of model failures, reducing issue remediation timelines

Third Line of Defense

- Promote transparency and accountability with internal audit teams through sharing of model results and data points
- Establish audit trails by enabling approval option and feedback management for models going into deployment and generating reporting around AI/ML solutions

Art of the Possible

AI considerations for the Insurance sector



Insurance industry example AI benefits



Automate Processes and Work




Improve Efficiency and Productivity



Enhance Creativity



Optimize / Reduce Costs



Enable Experiences



Fraud claims detection



Predictive analytics around customer behavior



Enhanced underwriting processes and risk evaluation



Claims processing and fraud detection



In-App virtual assistants



Accelerated document processing

Insurance industry AI use cases

Examples of Leveraging AI by Insurance Companies

-
- 01 Underwriting and Risk Assessment:**
 - Analysis of vast amounts of structured and unstructured data to evaluate risks more accurately
 - Pattern recognitions by AI leads to informed decisions on policy pricing, coverage limits, and eligibility
 - 02 Claims Processing and Fraud Detection:**
 - Automate claims handling by extracting relevant information, validating policy coverage, and assessing damage or loss
 - Analysis of historical data, claim patterns, and anomalies to detect potential fraud
 - 03 Customer Service and Engagement:**
 - AI-powered chatbots provide customer support, answer policy-related queries, and assist with claims submissions
 - AI systems handle a large volume of inquiries, offer real-time assistance, and provide 24/7 support in a human-like manner
 - 04 Personalized Insurance Products:**
 - Analysis of customer data to develop tailored insurance products and pricing models
 - Enables insurers to offer personalized coverage options and pricing structures
 - 05 Risk Prediction and Prevention:**
 - Analysis of historical data, sensor readings, and other relevant information to identify patterns, trends, and potential risks
 - Usage of insights to offer proactive risk prevention measures and risk management advice to policyholders
 - 06 Actuarial Analysis and Pricing:**
 - Assists actuaries in analyzing complex datasets and building predictive models for pricing insurance products
 - Machine learning algorithms identifies risk factors, estimates claim probabilities, and optimizes pricing models
 - 07 Data Analytics and Fraud Analytics:**
 - Processes and analyzes vast amounts of structured and unstructured data to generate actionable insights
 - Identification of market trends, assessment of customer behavior, and enhancement in decision-making across functions



This presentation contains general information only and Deloitte/Deloitte Risk & Financial Advisory is not, by means of this presentation, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This presentation is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor.

Deloitte/Deloitte Risk & Financial Advisory shall not be responsible for any loss sustained by any person who relies on this presentation.

As used in this document, “Deloitte” means Deloitte & Touche LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Common AI terms and definitions

Term	Definition
Artificial Intelligence	The capability of a computer system to imitate human intelligence. Using math and logic, computer systems simulate the reasoning process that humans.
Computer Vision	Computer vision focuses on enabling computers to identify and understand objects and people in images and videos.
Deep Learning	A type of machine learning in which artificial neural networks process, learn, and make decisions based on unstructured data.
Generative AI	AI models that create original content including text, images, and video based on patterns in large quantities of training data.
Machine Learning (ML)	A subset of AI that uses algorithms to enable computer systems to learn without human instruction and develop their own intelligence.
Neural Networks	A machine learning process that teaches computers to process data using interconnected nodes or neurons in a layered structure that resembles the human brain.
Natural Language Processing (NLP)	A machine learning technology that gives computers the ability to interpret, manipulate, and comprehend human language such as emails, text messages, and social media.
Supervised Learning	This form of machine learning requires human input (labeled or reference data) during the training process to help the machine identify patterns
Unsupervised Learning	This form of machine learning requires little or no human involvement. The machine draw conclusions on its own from patterns that it finds in the training data.