


Generative AI

AUDIT TRAILS & DATA SILOS

BY DR. AMIT K. SHAH, PH.D.



1. Introduction to Generative AI

What it is and recent advancements and implications for data management and governance

2. Implications of Generative AI

How Generative AI is transforming the industry and what's to come

3. Challenges with Generative AI

The problems that need to be overcome for effective implementation

4. Data Silos

What are data silos and why do they exist? How Generative AI will exacerbate this and what to do about it

5. Audit Trails

What are audit trails? Why do they need to be implemented with generative AI? Strategies to consider

6. How GNS-AI helps Organizations

How GNS-AI helps organizations implement generative AI and emerging technologies

7. Conclusion

Generative AI key points and impact on data management. Call to action.



Table of contents

COPYRIGHT GNS-AI LLC 2023

The speaker



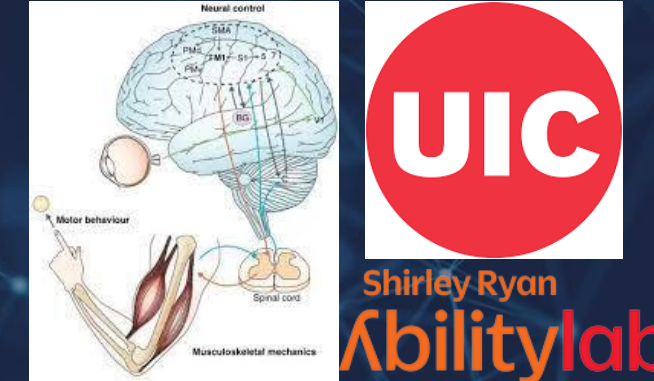
Dr. Amit K. Shah, Ph.D.



Bachelor's in Computer
Science/Pre-med




Former Lead Data Scientist



Biomedical Engineering
PhD (Neural Control of
Movement; Neural
Engineering)



Former Data Science Manager



Artificial Intelligence (AI)



This Photo by Unknown Author is licensed under [CC BY-NC-ND](#)





This Photo by Unknown Author is licensed under [CC BY-SA-NC](#)



COPYRIGHT GNS-AI LLC 2023
This Photo by Unknown Author is licensed under [CC BY-SA-NC](#)

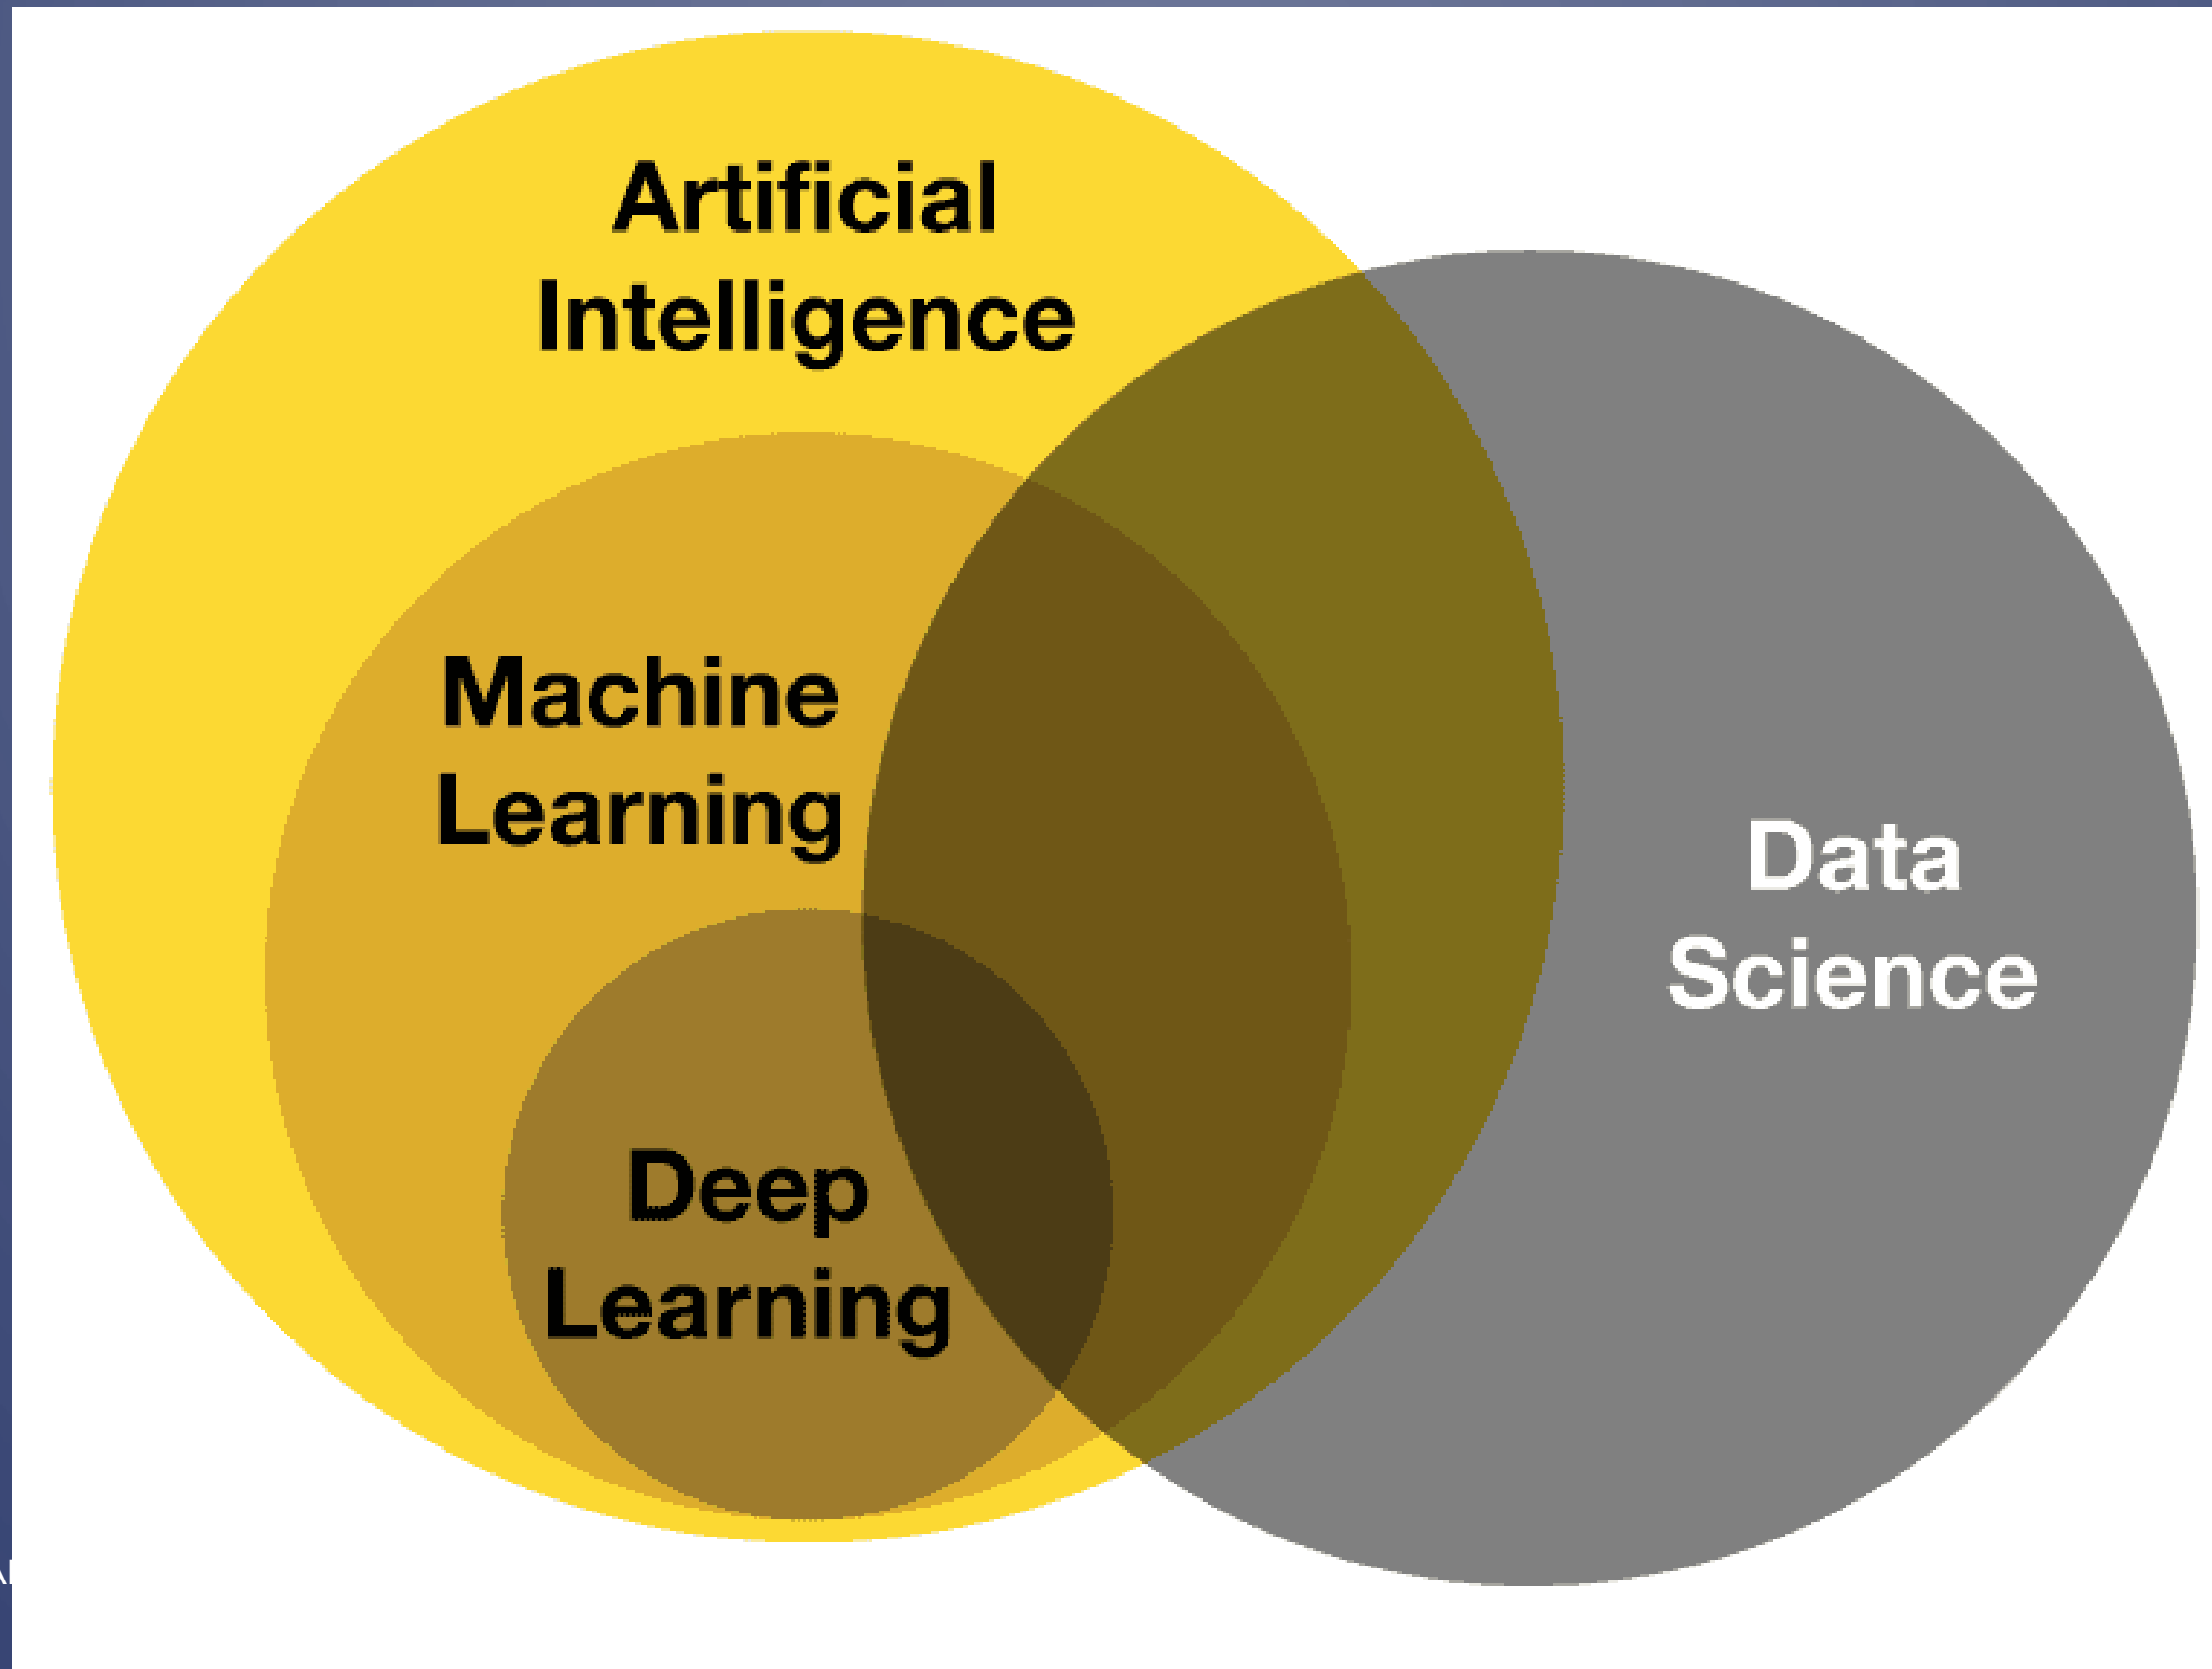


This Photo by Unknown Author is licensed under [CC BY-SA](#)



This Photo by Unknown Author is licensed under [CC BY-SA](#)

AI Taxonomy



Types of Artificial Intelligence

Based on Capabilities

Narrow AI

General AI

Super AI

Based on Functionalities

Reactive Machine


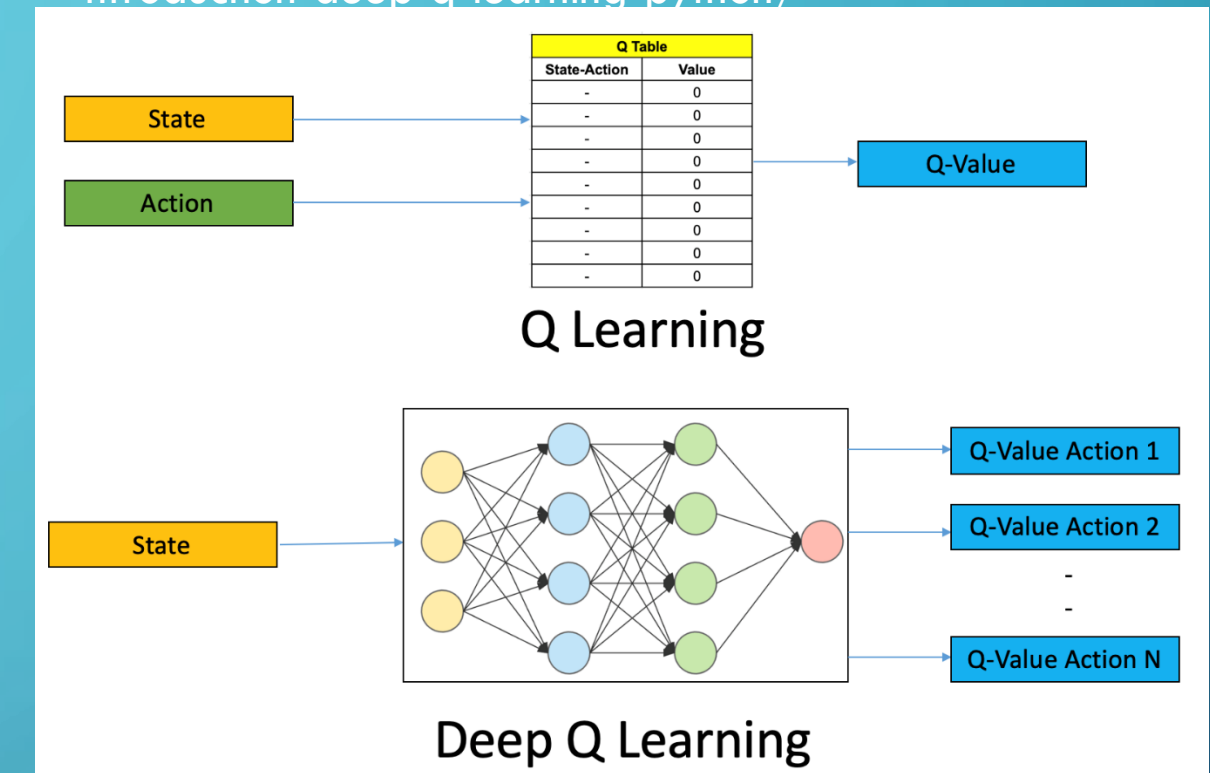
Limited Memory

Theory of Mind

Self-awareness

The Quest for Super AI – Q* and GPT-4: The Dilemma at OpenAI and Sam Altman

<https://www.analyticsvidhya.com/blog/2019/04/introduction-deep-q-learning-python/>




Q*

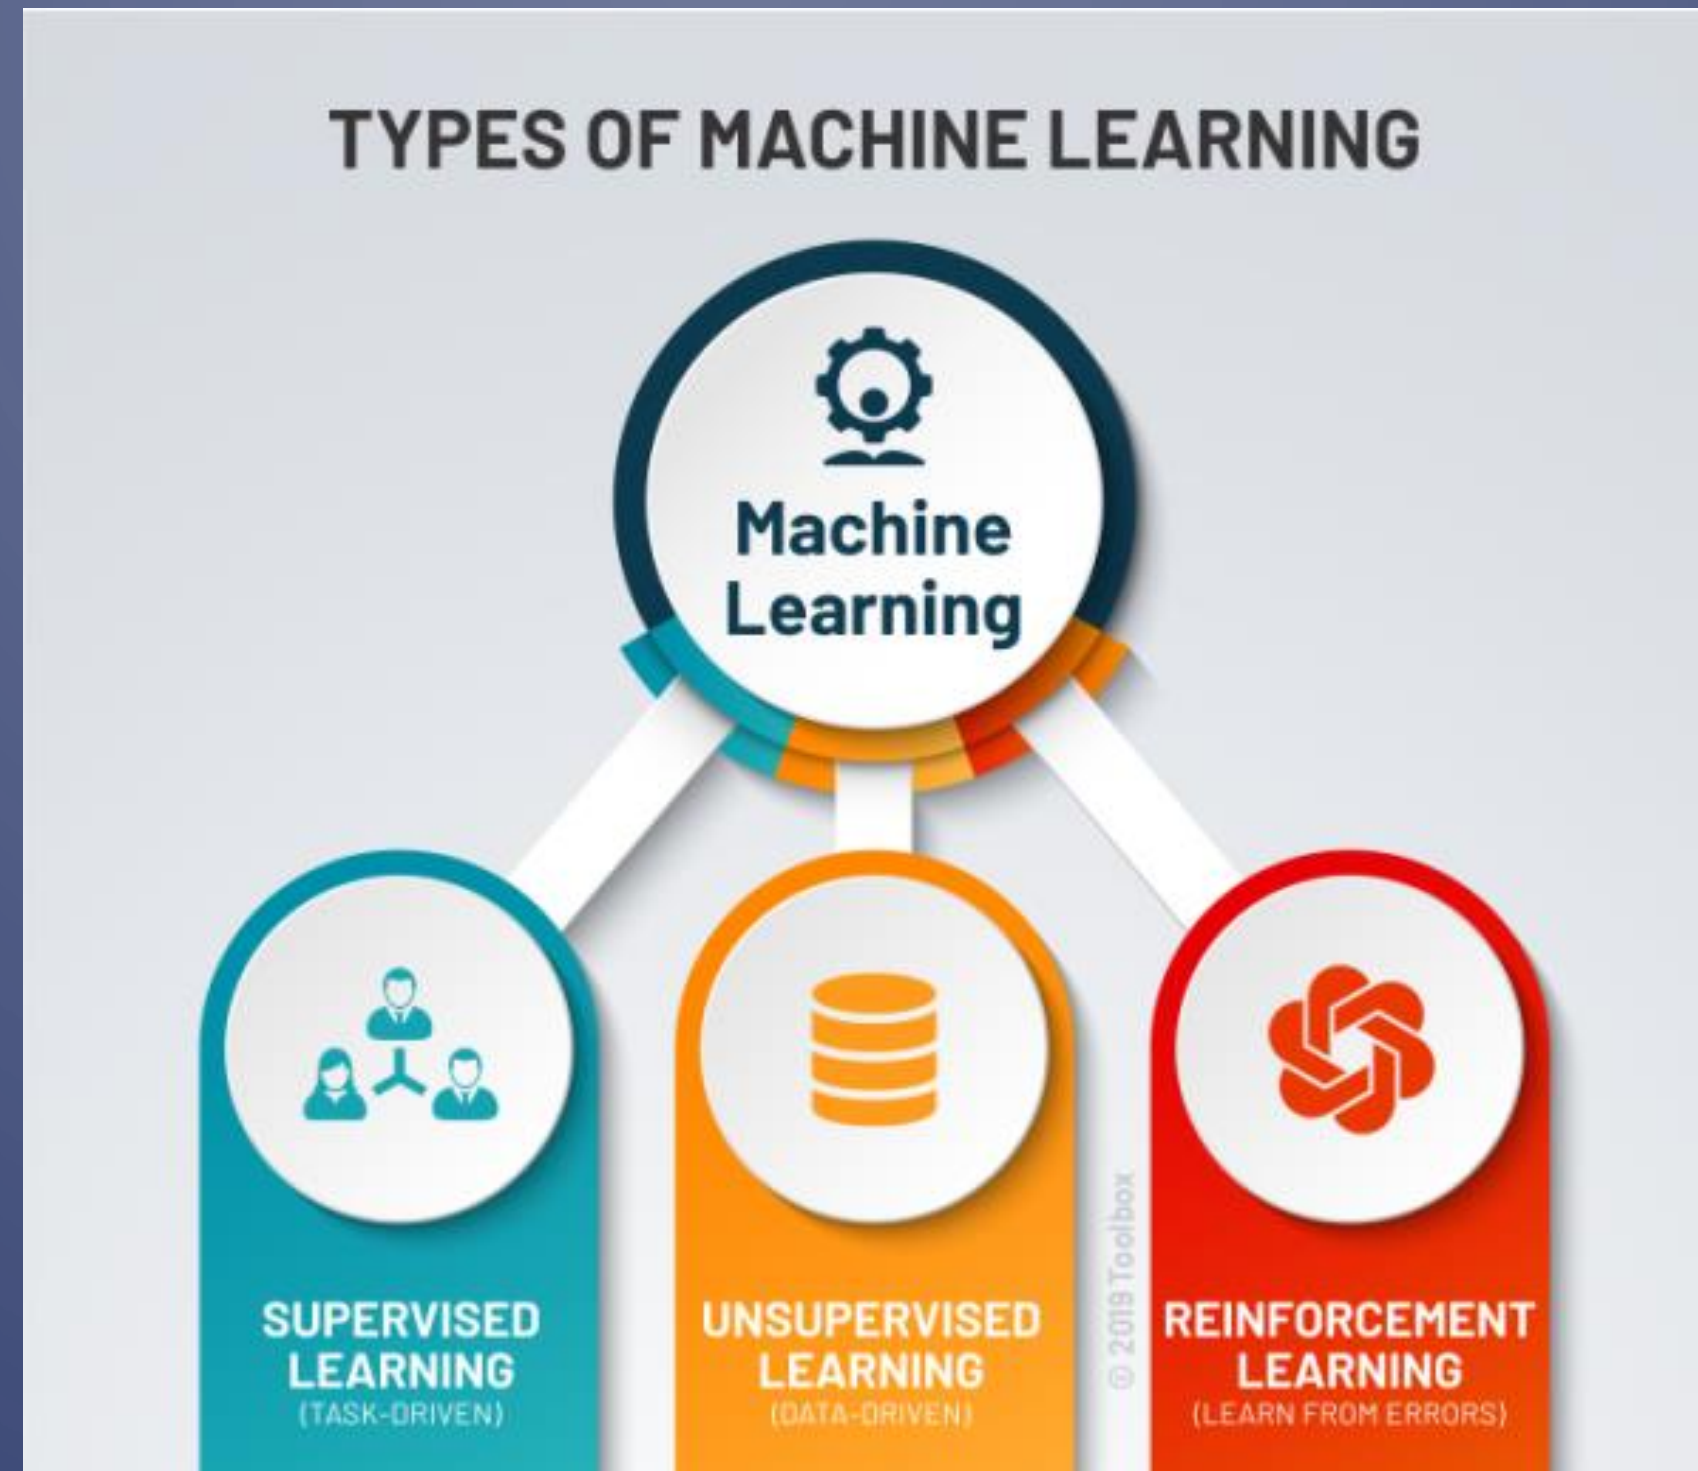
???



COPYRIGHT GNS-AI LLC 2023



Machine Learning



What is Generative AI?

A class of AI that can generate new content or data based on its training


Benefits

- Enhanced Creativity and Productivity
- Enables Personalized Customer Experiences
- Streamlines Content Creation Processes

Challenges

- Requires Significant Data and Computational Power
- Unpredictable Outputs
- Navigating Ethical Considerations and Privacy

COPYRIGHT GNS-AI LLC 2023



Use Cases in Retail



Personalized Product Recommendations



AI-powered Chatbots



Automated Inventory Management



Marketing Content



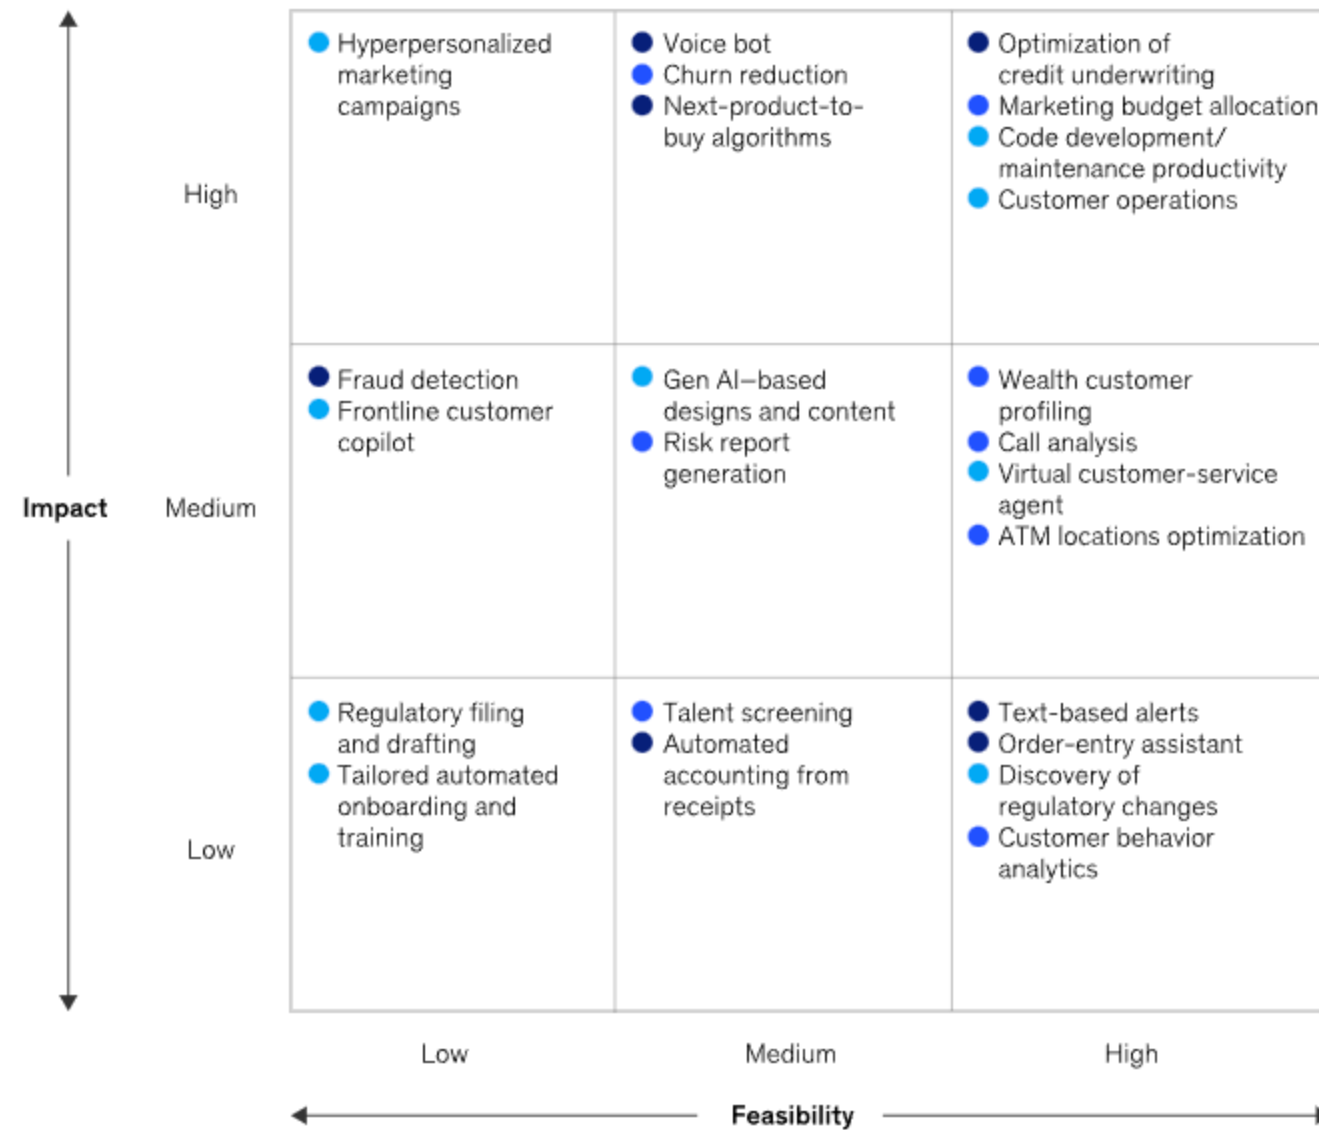
Automated Customer Insights

Generative AI Value

Exhibit 1

Take a portfolio view on value.

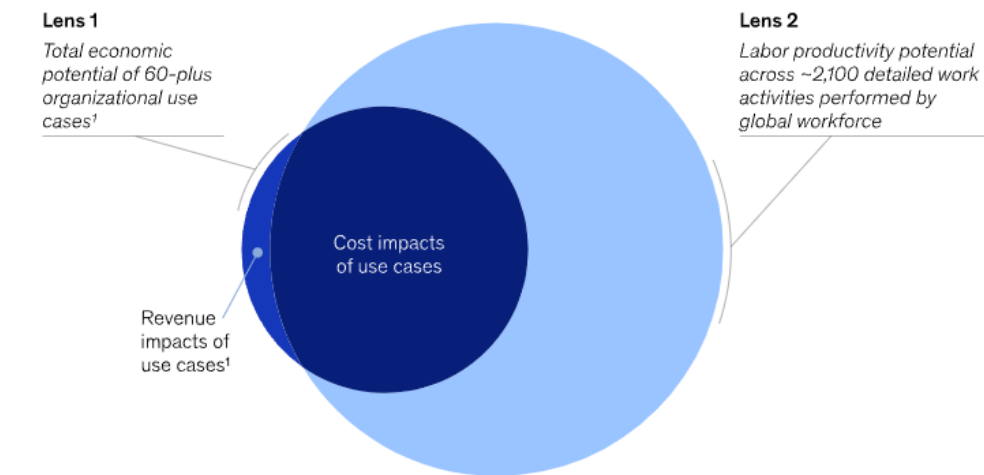
Illustrative banking use cases portfolio ● Generative AI ● Business intelligence and analytics ● Classical AI/ML



McKinsey & Company

Exhibit 1

The potential impact of generative AI can be evaluated through two lenses.



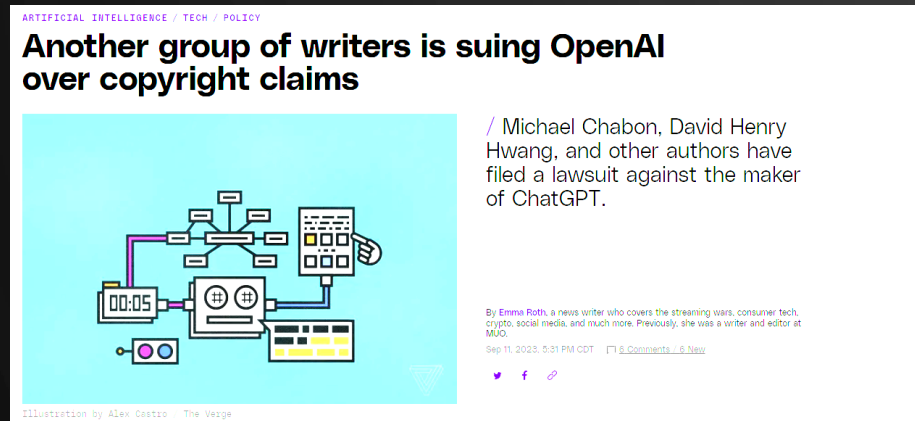
*For quantitative analysis, revenue impacts were recast as productivity increases on the corresponding spend in order to maintain comparability with cost impacts and not to assume additional growth in any particular market.

McKinsey & Company

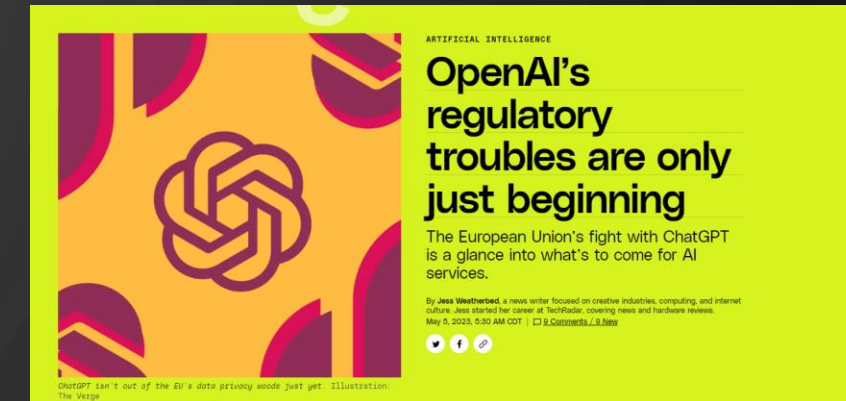
<https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#business-value>

“Our latest research estimates that generative AI could add the equivalent of \$2.6 trillion to \$4.4 trillion in annual economic benefits across 63 use cases.” McKinsey – June 14, 2023 – [The economic potential of generative AI – the next productivity frontier](#)

Ethical, Legal and Regulatory Issues



Intellectual Property



Regulatory Issues



Legal and Compliance Risks

Lack of Stability in Model Outputs

Why can't we trust ChatGPT?



New research reveals that the AI models behind ChatGPT have exhibited extreme instability, which may affect future use and monetization

17:18, 23/07/23

TAGS: [artificial intelligence](#) [OpenAI](#) [ChatGPT](#)

A study published by Berkeley and Stanford universities last week revealed instability in the outputs of GPT-4, the latest generative artificial intelligence model used by OpenAI's ChatGPT. The study highlighted significant changes in GPT-4's performance over a brief three-month period, particularly in relatively simple tasks. Notably, the model showed a dramatic decline in accuracy in identifying prime numbers, plummeting from 97.6% in March to a mere 2.4% in June. Surprisingly, the GPT-3.5 model, upon which the free public version runs, actually exhibited an improvement in this aspect.

OpenAI acknowledged the research and said it was aware of the reported regressions. Logan Kilpatrick, the company's head of developer relations, tweeted that their team was actively investigating the matter.

<https://www.calcalistech.com/ctechnews/article/ryray95q3>

Implications of Unstable Outputs for Business & Technology

- **Reduced Reliability:** Challenges in depending on AI for consistent results

- **Decision-making Risks:** Potential errors in automated decision processes

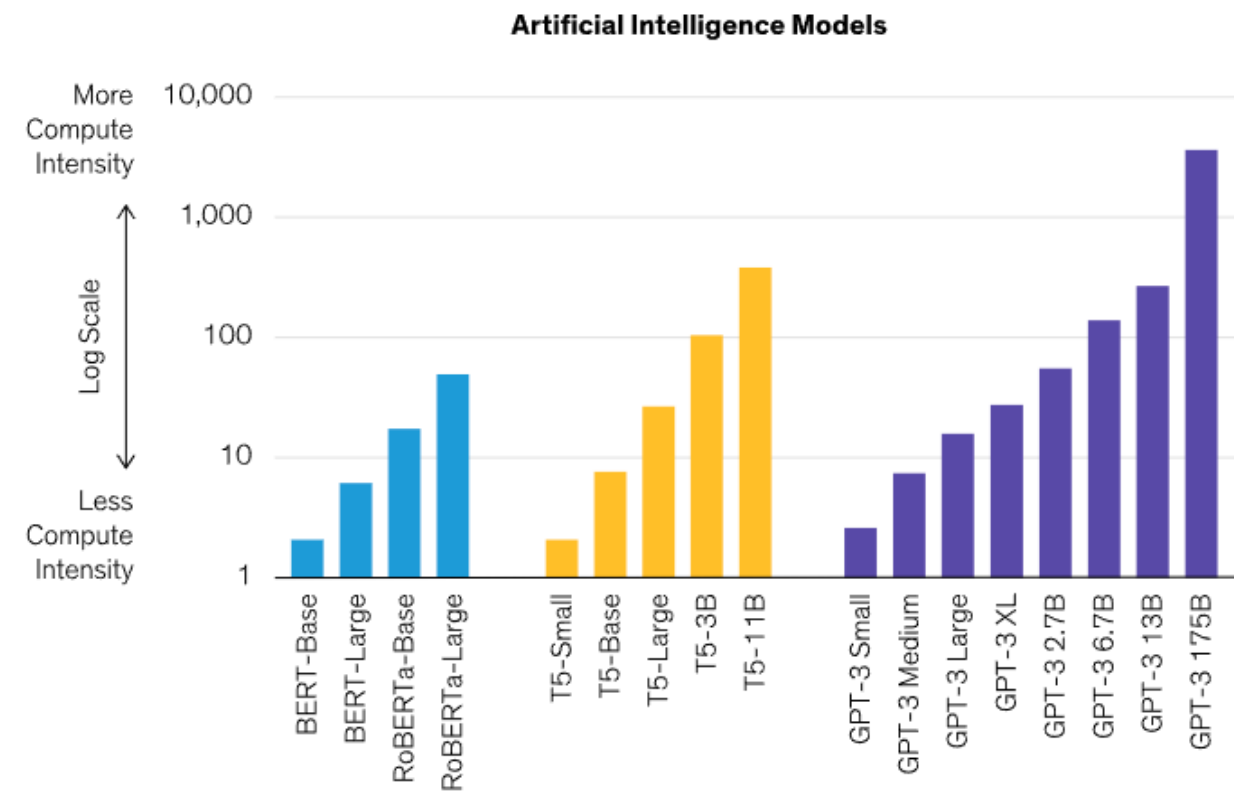
- **User Trust Issues:** Eroding confidence in AI systems among users

COPYRIGHT GNS-AI LLC 2023

Challenges in Generative AI: Computational Power and Data Requirements Are Drastically Increasing

AI Models' Computational Complexity Requires Plenty of Power

Training time for AI models in petaFLOP/s-days*




Historical analysis does not guarantee future results.

*FLOPs (floating-point operations per second) is a measure of compute performance used in deep-learning models that require floating-point operations. PetaFLOP/s-days represents the number of days required to train a particular model, assuming that machine training the model performed a fixed amount of computation (1,015 neural net operations per second for one entire day).

As of June 30, 2023

Source: Bank of America, NVIDIA and AllianceBernstein (AB)





<https://epochai.org/blog/trends-in-training-dataset-sizes>

Data Silos – The Necessity of Data Integration



Data silos



https://www.researchgate.net/figure/The-graph-of-compartmental-periodic-unpredictable-function-f-t-The-length-of-step-h_fig1_369220320

Data silos occur when data is isolated within different departments or systems, hindering organization-wide access and analysis

Impact on Generative AI Products

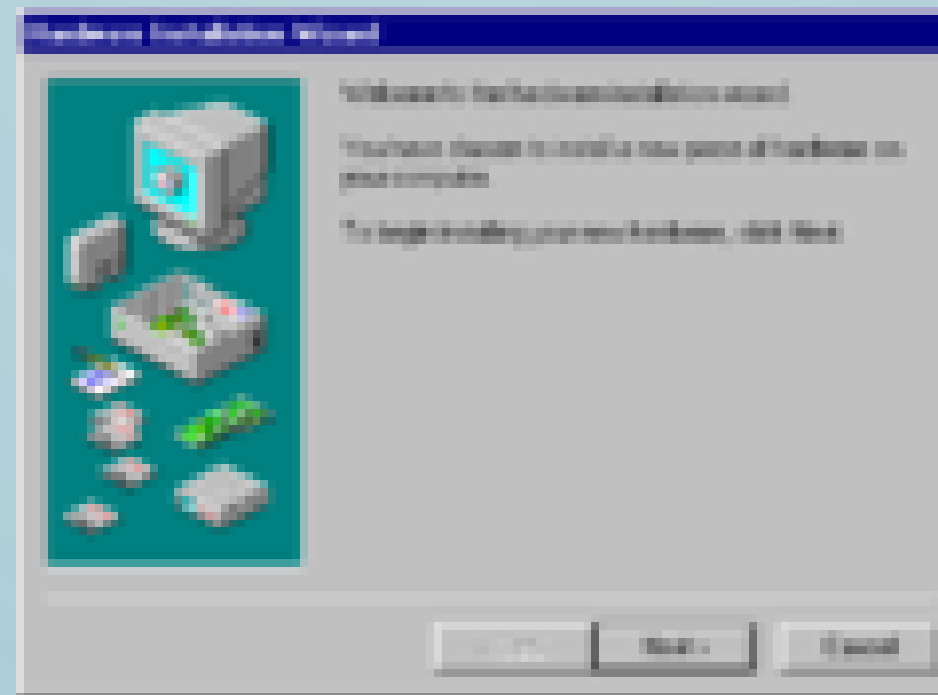
- Limited Data Access: Restricts AI's learning potential.
- Biased Outputs: Increases risk of skewed AI decisions.
- Innovation Stifled: Prevents comprehensive insights for AI development.

Data silos: Causes

Organizational Structure



Legacy Systems



Lack of Standardization



Data Silos: Impact on Generative AI Implementation




Incomplete Training Data



Biased Model Outputs

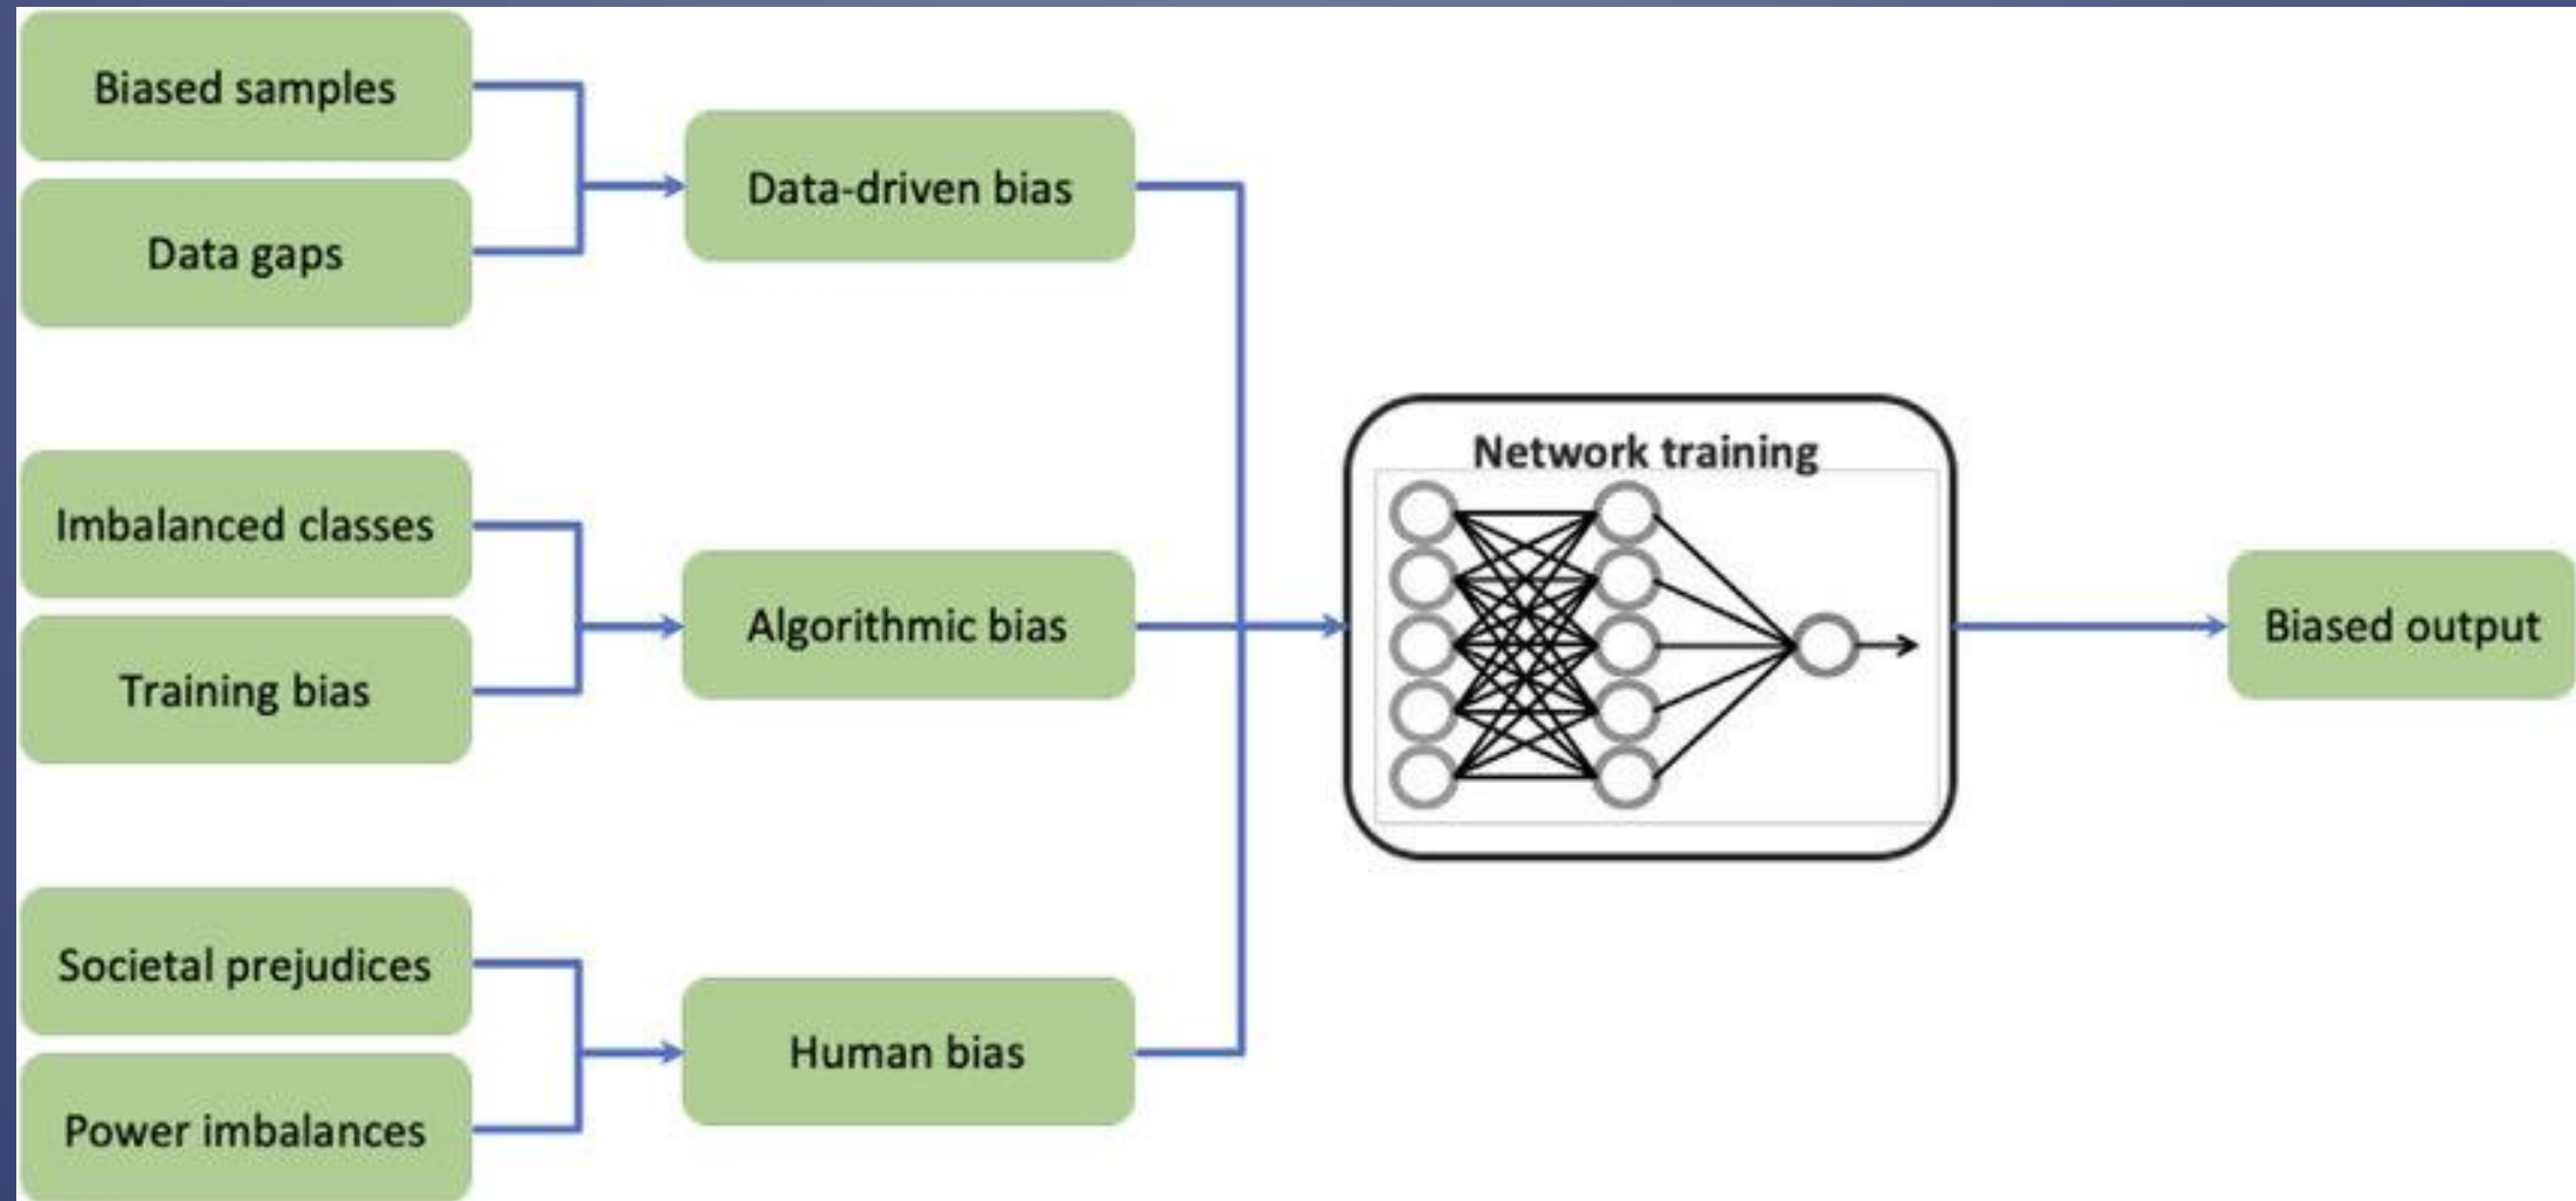


**Challenges in Real-Time
Data Utilization**



**Model Training and
Validation Issues**

Sources of Bias



Increased Use of Third-Party Tools, Including Third-Party AI Software Exacerbate Data Silos and Contribute to Additional Spend in Organizations

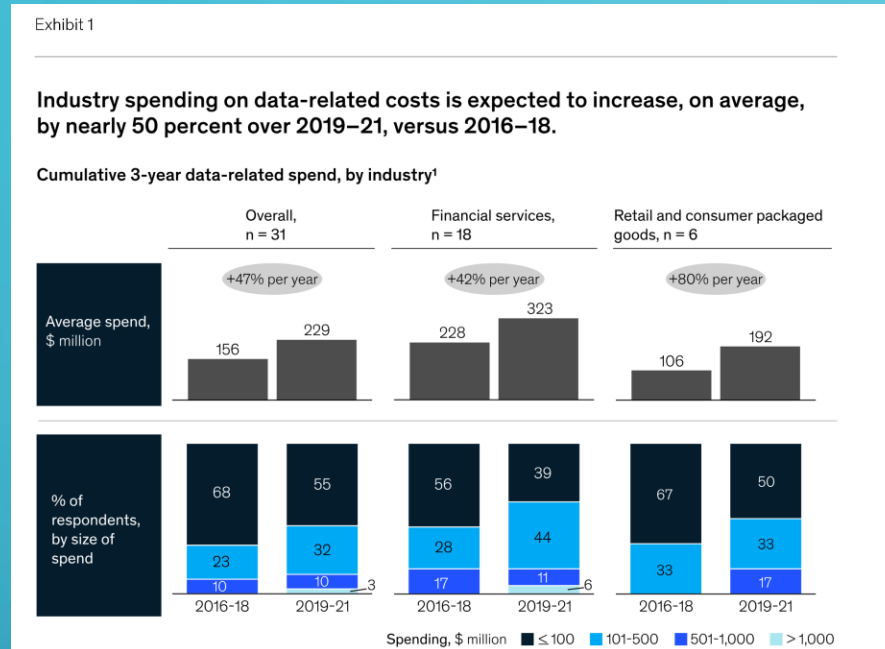


Exhibit 2

Data-related spending breaks down into four areas.

	1. Data sourcing	2. Data architecture	3. Data governance	4. Data consumption
Description	Cost associated with procuring data from customers, ¹ 3rd-party vendors, etc	Cost associated with data infrastructure (procuring software, hardware) and data engineering (building and maintaining infrastructure)	Cost of data-quality monitoring, remediation, and maintaining data-governance artifacts (eg, data dictionary, data lineage)	Cost associated with data analysis and report generation (including spending on data access and cleanup)
Components	3rd-party data	Labor, infrastructure, and software	Labor, software	Labor, software
Typical owner of spend	Head of business unit	CIO	Chief data officer	Head of function or business unit
Typical spend, % of IT spend	5–25 ²	8–15	2.5–7.5	5–10
Example for a midsize financial institution,³ \$ million	70–100	90–120	20–50	60–90

¹Excludes internal data-capture processes.
²Industries that don't directly touch consumers (eg, consumer packaged goods) spend a higher share (>20%) on data sourcing.
³For midsize organizations with revenues of \$5 billion to \$10 billion and operating expenses of \$4 billion to \$6 billion. Absolute values vary by industry and size of the organization; eg, absolute spend is, on average, higher for the telecommunications industry.

<https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/reducing-data-costs-without-jeopardizing-growth>

“The report, based on an executive survey of more than 1,240 respondents representing companies in 59 industries and 87 countries, revealed that 78% of organizations use third-party AI tools, and more than half use third-party tools exclusively.” – MIT Sloan Third-party AI tools pose increasing risks for organizations


Case Study: AEG's Transformation in Campaign Execution

Breaking down Data Silos




About AEG

The world's largest sports and live entertainment brand, operating venues, and producing global events



The Challenge

Needed a quick, efficient system for scheduling and targeting diverse email campaigns for multiple segments.



Innovative Solution

- Custom email system and Inboxable integration
- 100+ marketing users can now create campaigns without coding.
- Campaign execution time reduced from 2 weeks to a few hours.



[How AEG Cut Its Campaign Execution from Hours to Minutes \(data-axle.com\)](https://data-axle.com)

Impactful Outcomes

- Streamlined internal workflows
- Faster campaign turnarounds
- Increased autonomy in campaign design for marketing teams

Data Lakes: Overcoming Data Silos

Data Lake - A centralized repository designed to store, process, and secure large volumes of structured and unstructured data.



Benefits of Data Lakes for Solving Data Silos



Unified Data Access




Enhanced Data Analytics



Improved Data Governance

Data Virtualization: Overcoming Data Silos

Data Virtualization – A virtual layer that allows for unified data access and retrieval across various silos without physical data movement



Benefits of Data Virtualization for Solving Data Silos



Unified Data View



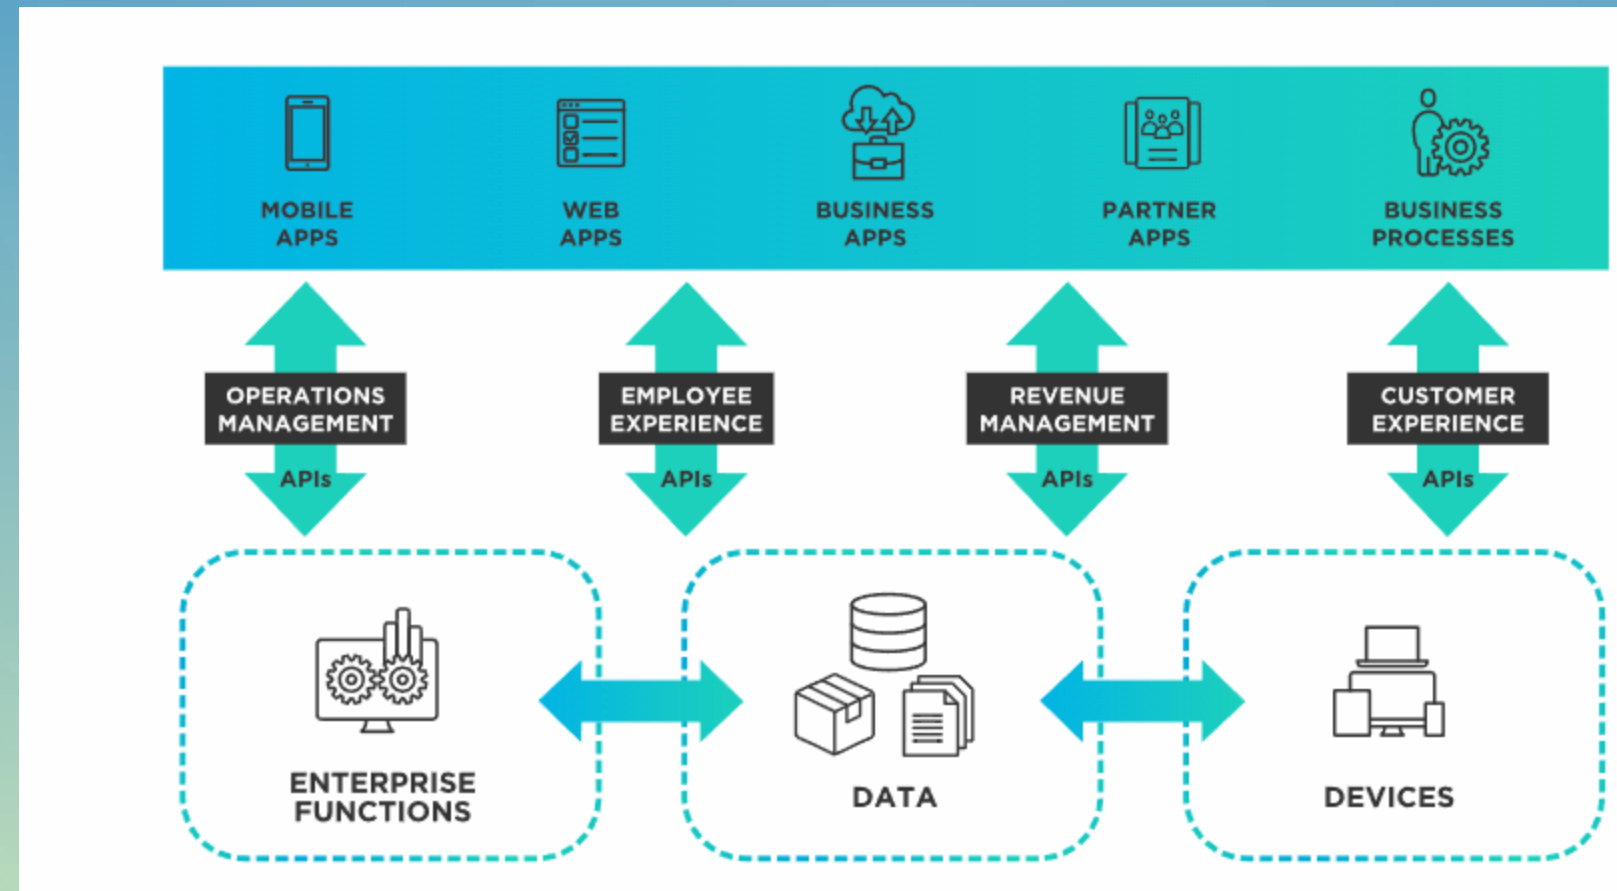
Improved Data Analytics



Cost-Effective

API-led Integration: Overcoming Data Silos

API-led Integration – Connecting different software systems and data sources



<https://www.tibco.com/reference-center/what-is-api-led-integration>

Benefits of API-led Integration for Solving Data Silos



Seamless Data Connectivity



Realtime Data Access



Scalability and Flexibility

Choosing the Right Data Integration Strategy: Generative AI

Data Lakes: Centralized Big Data Repositories

- Ideal for aggregating vast amounts of varied data
- Enables advanced analytics with a holistic data view and deep learning analytics
- Supports varied data types




This Photo by Unknown Author is licensed under [CC BY-SA](#)

When to Use: Opt for data lakes when AI requires comprehensive dataset and contextual information, e.g. chatbots, decision-making

Critique: Requires strong governance to prevent data mismanagement; don't use when compliance/regulatory/legal/ethical concerns

API-Led Integrations: Seamless Interconnectivity

- Real-time data access and sharing
- Modular and flexible architecture
- Enables external data integration



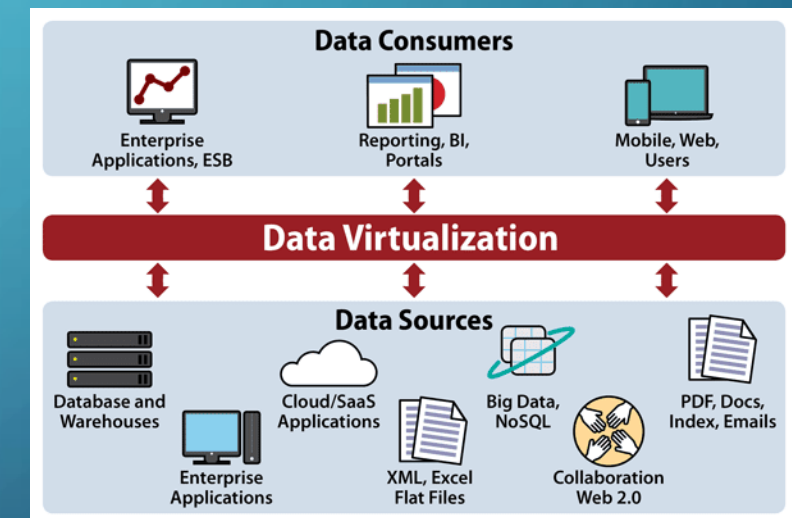
This Photo by Unknown Author is licensed under [CC BY](#)

When to Use: Best when generative AI needs dynamic data from multiple, distributed sources.

Critique: Complexity in management and potential performance impacts; don't use when data cannot be shared due to compliance/regulatory/legal/ethical concerns

Data Virtualization: Agile Data Fusion

- Agile access to cross-domain data
- Reduces need for data movement
- Supports diverse data consumption



When to Use: Ideal for scenarios offering quick, unified access to a variety of live data sources.

Critique: May not suit heavy processing tasks; dependent on source system performance

Choosing the Right Data Integration Strategy: Compliance and Ethics in Generative AI

Data Lakes: Centralized Big Data Repositories

- Facilitates compliance with storage and processing regulations
- Enables detailed audit logs for data tracing
- Demands robust data governance to maintain quality and privacy




This Photo by Unknown Author is licensed under [CC BY-SA](#)

When to Use: Few compliance/ethical/legal/regulatory risks, the budget exists

Critique: Requires strong governance to prevent data mismanagement; don't use when compliance/regulatory/legal/ethical concerns

API-Led Integrations: Seamless Interconnectivity

- Promotes data stewardship with controlled access
- Eases compliance with data protection standards
- Requires secure API management to prevent data leaks



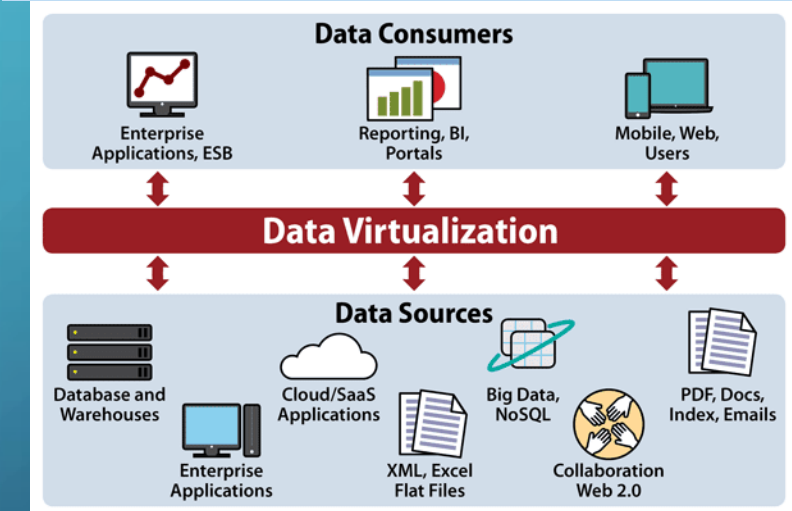
This Photo by Unknown Author is licensed under [CC BY](#)

When to Use: Ideal for real-time, controlled data exchange where compliance with data sharing policies is essential.

Critique: Ensure APIs do not expose sensitive data, adhering to privacy law

Data Virtualization: Agile Data Fusion

- Minimizes data duplication, reducing scope of compliance
- Supports real-time ethical data handling practices
- Depends on compliance and security measures



When to Use: When need to swiftly integrate compliant data from multiple regulated domains for AI applications

Critique: Requires careful implementation to ensure each data source's regulations are respected

Synthetic Data Use

Simulation Helps When Data


Is Difficult to Source



Is Difficult to Label



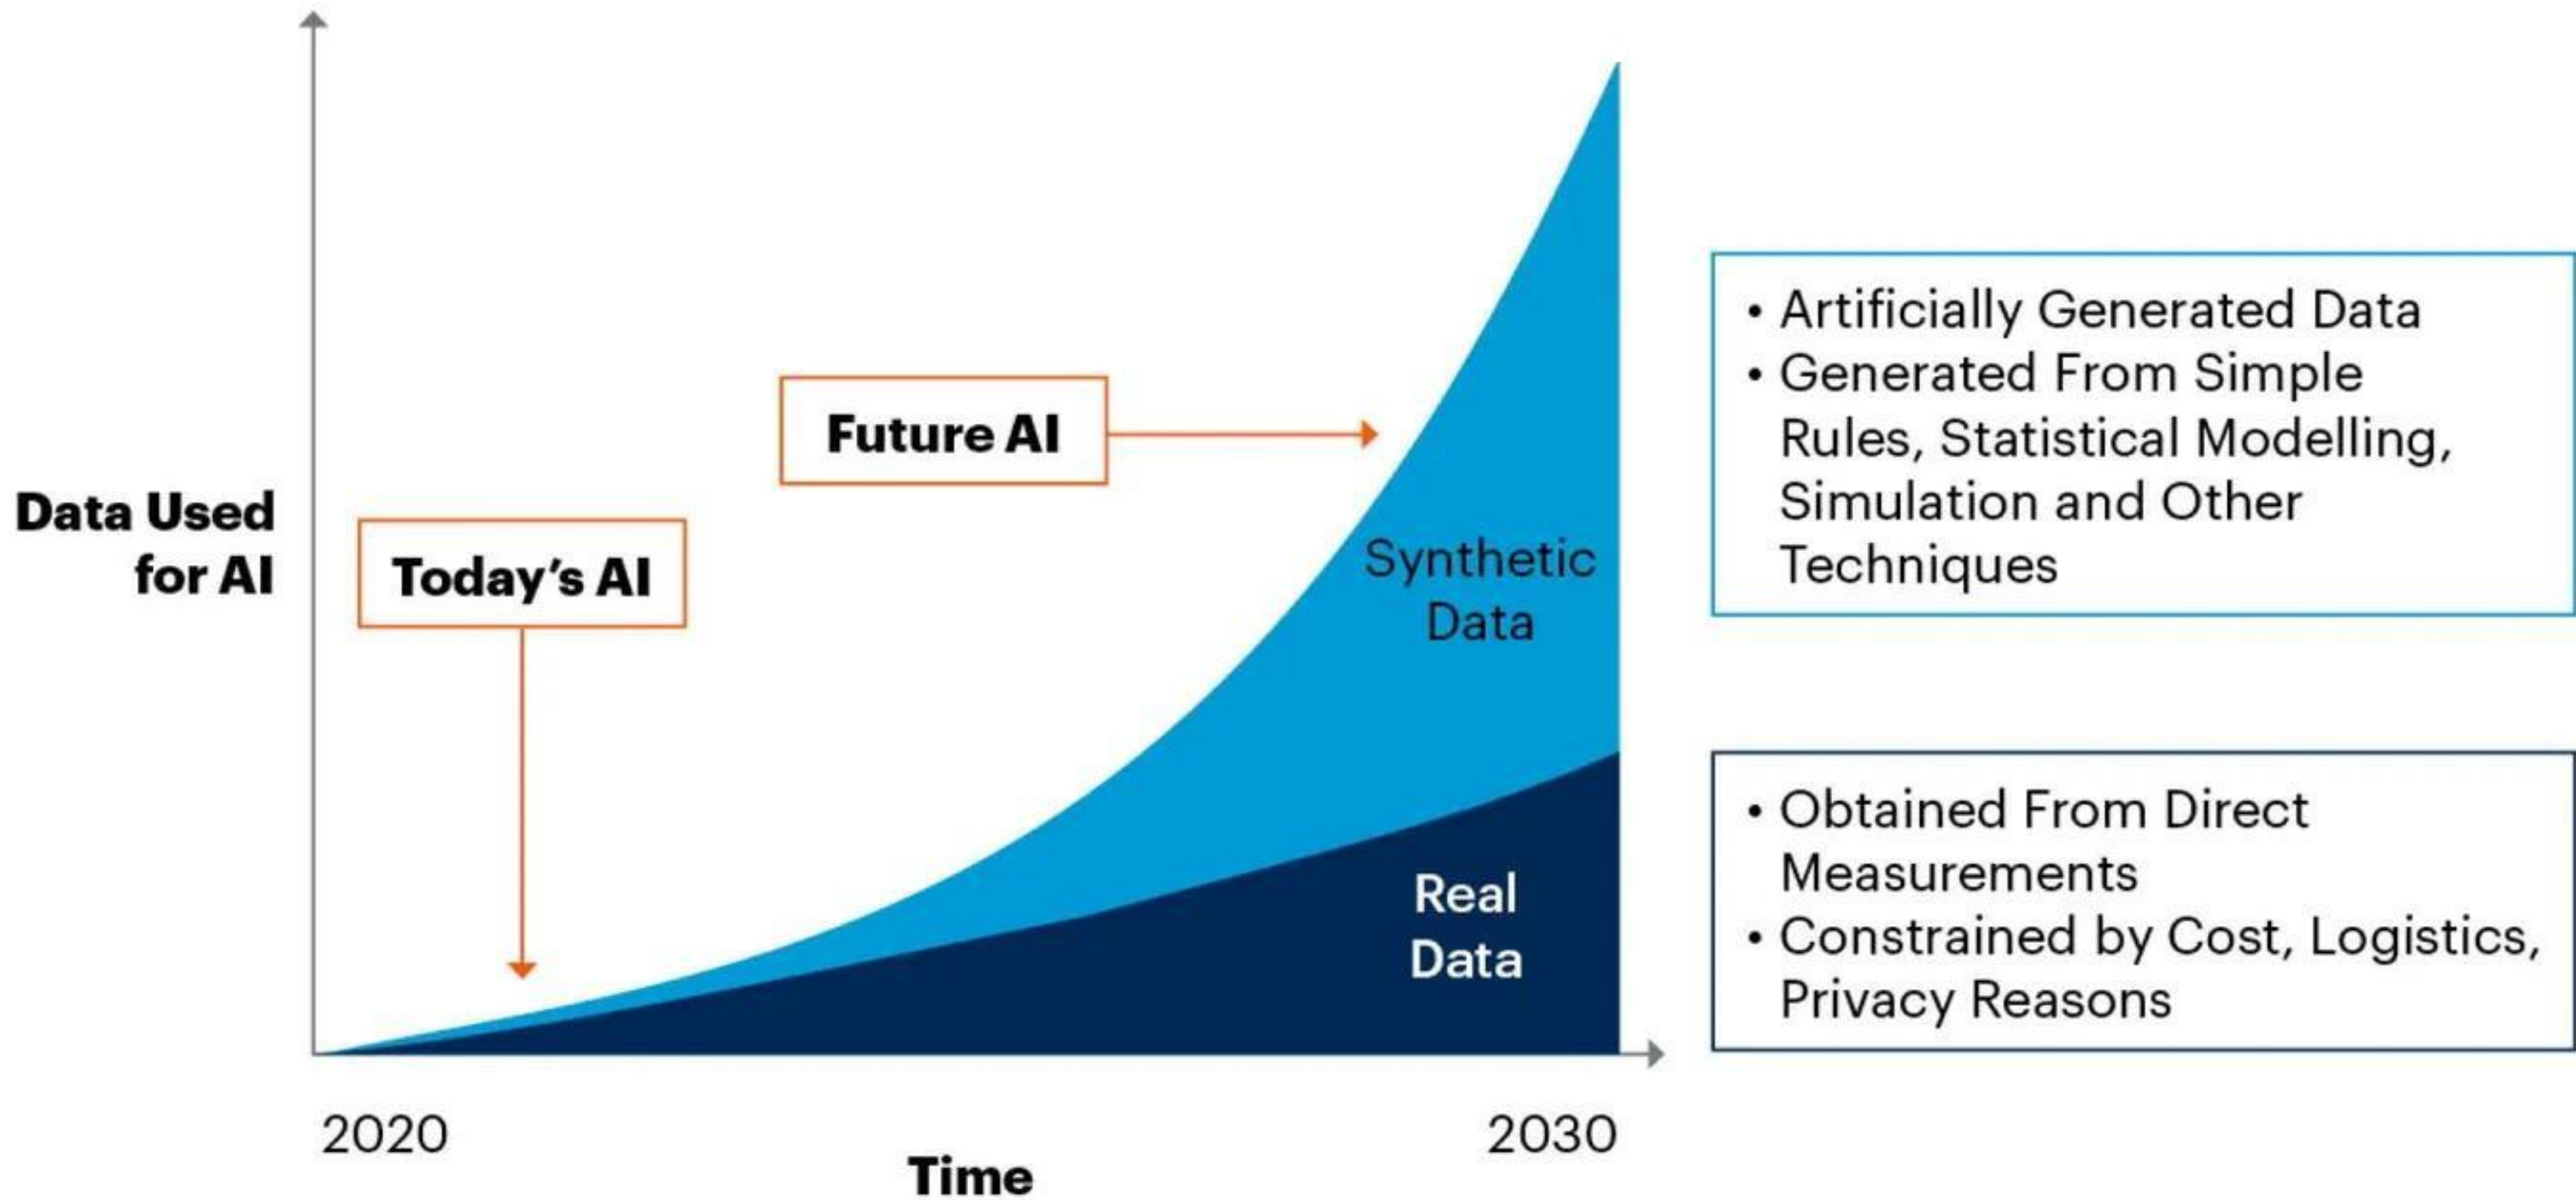
Is Closed Loop



What's Needed to Make This Happen?

1:36:39 / 3:03:20 • Simulation >



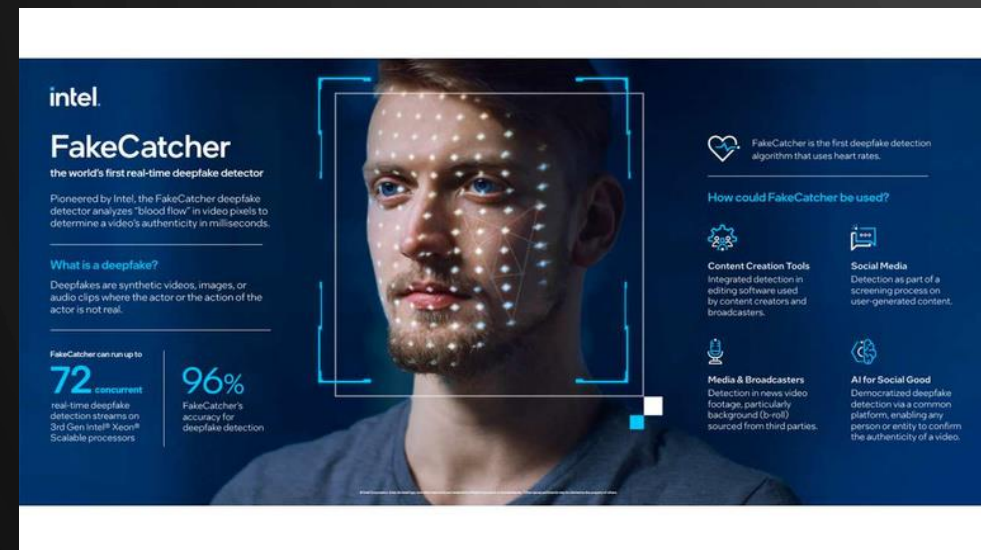


Source: Gartner
750175_C

Reification Fallacy



Separating Synthetic Data from Real Data



intel.
FakeCatcher
the world's first real-time deepfake detector

Pioneered by Intel, the FakeCatcher deepfake detector analyzes 2700 "blood flow" in video pixels to determine a video's authenticity in milliseconds.

What is a deepfake?
Deepfakes are synthetic videos, images, or audio clips where the actor or the action of the actor is not real.

FakeCatcher can run up to **72 concurrent** real-time deepfake detection streams on 3rd Gen Intel® Xeon® Scalable processors. **96%** FakeCatcher's accuracy for deepfake detection.

FakeCatcher is the first deepfake detection algorithm that uses heart rates.

How could FakeCatcher be used?

- Content Creation Tools**
Integrated detection in editing software used by content creators and broadcasters.
- Social Media**
Detection as part of a screening process on user-generated content.
- Media & Broadcasters**
Detection in news video footage, particularly background (B-roll) sourced from third parties.
- AI for Social Good**
Democratized deepfake detection via a common platform, enabling any person or entity to confirm the authenticity of a video.


ARTIFICIAL INTELLIGENCE

OpenAI Abruptly Shuts Down ChatGPT Plagiarism Detector— And Educators Are Worried

College professors see AI Classifier's discontinuation as a sign of a bigger problem: AI plagiarism detectors do not work.

By Casey Epstein Gross · 07/26/23 3:50pm

[f](#) [t](#) [i](#) [m](#) [e](#)



The logos of OpenAI and ChatGPT. AI? via Getty Images

AVSspoo Challenge: the leading database against voice spoofing

Just as there are tools and programs to generate fake voices, initiatives aim to fight this increasingly widespread practice.

Challenge AVSspoo is one of them. Its main goal is to launch challenges for companies specialising in the field to analyse speech processing problems related to voice spoofing and design measures to combat it.

AVSspoo Challenge 2021

AVSspoo 2021 was the last challenge launched. It aimed to promote progress in reliable automatic speaker verification and deepfake detection in more realistic and practical scenarios.

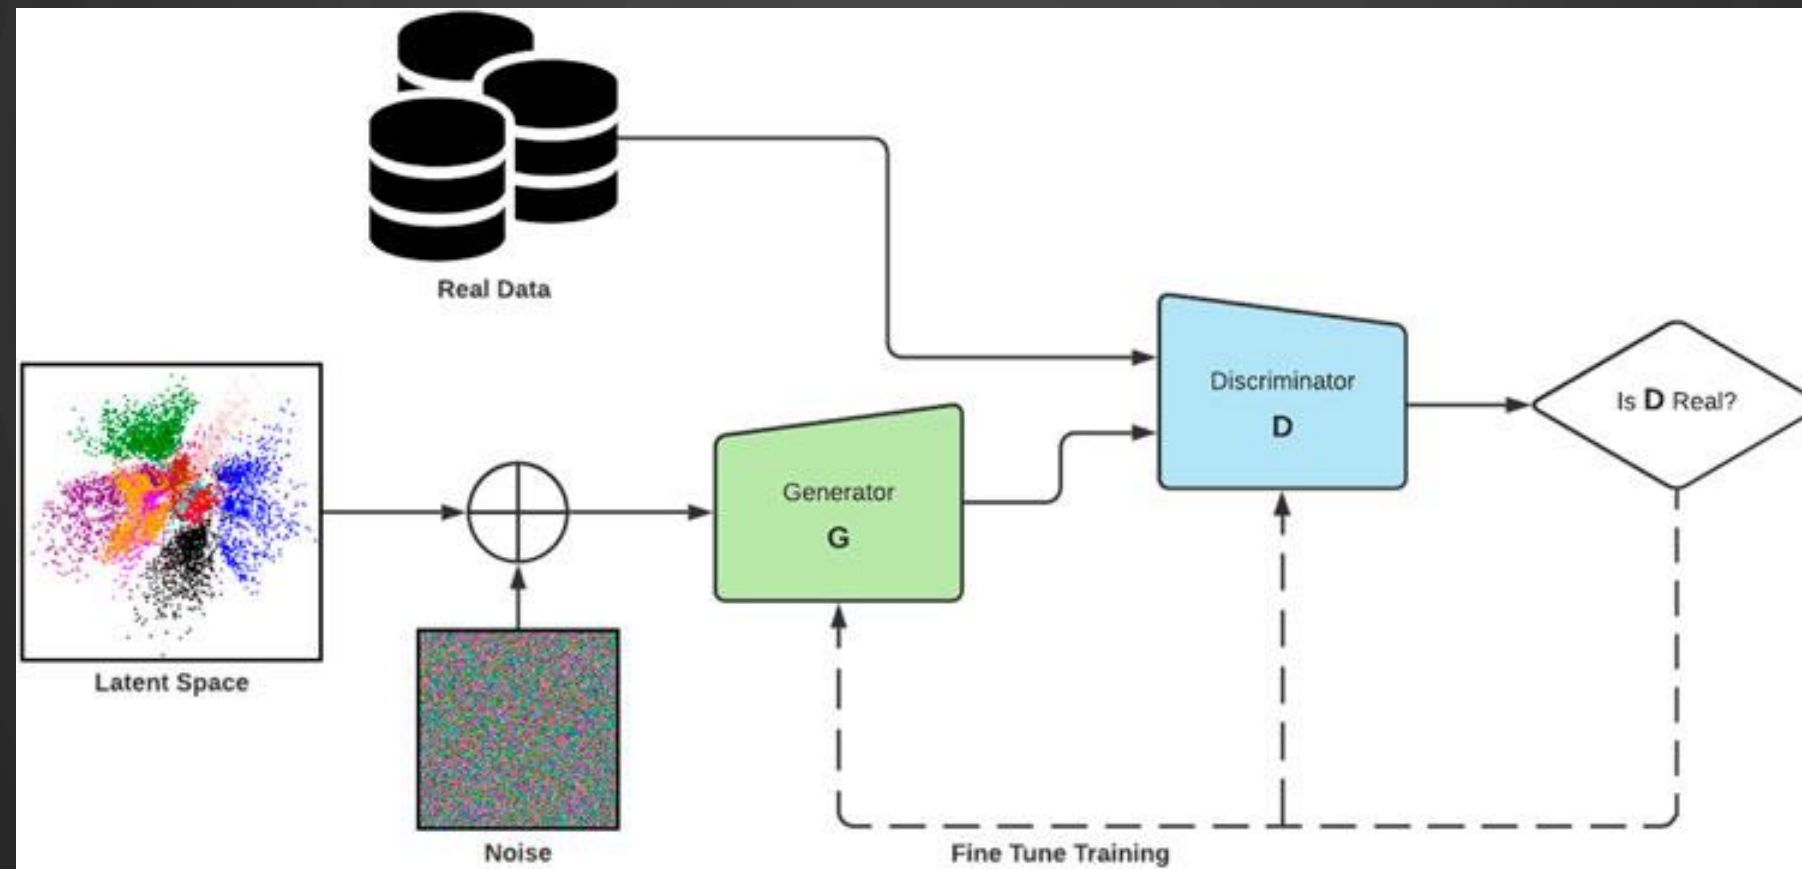
For this purpose, telephone channels were to be simulated in which voice data was to be encoded, understood, and transmitted. At the same time, the acoustic propagation in physical spaces had to be analysed by creating sentences with the voice of a target person.

The challenge consisted of **three tasks** where each team had to detect a specific type of voice attack: logical access (LA), physical access (PA), and deepfake (DF).

- Logical access (LA):** the objective was to study the robustness of the solutions against compression variations, packet loss, and other artefacts derived from bandwidth, transmission infrastructures, and variable bitrates issues.
- Physical access (PA):** concerned with replay attack detection in different environments.
- Deepfake (DF):** concerned with voice conversion (VC) detection and text-to-speech (TTS) synthesis over compressed audio. This task aimed to evaluate the robustness of spoofing detection solutions when used to detect manipulated speech data.

Challenge conclusions

Challenge results indicated that the robustness of spoofed audio detection is substantially improved when deepfake techniques are employed.



Synthetic Data Will Eventually Become Indistinguishable from Real Data

How do you keep track of:




what's real?

This Photo by Unknown Author is licensed under [CC BY-NC-ND](#)



This Photo by Unknown Author is licensed under [CC BY-NC](#)

who is using
the data?



This Photo by Unknown Author is licensed under [CC BY-SA-NC](#)

where the
model went
wrong?

Audit Trail

Audit Trails: Audit trails are records that chronologically catalog events and changes, providing a transparent trail of activities over time.

Use Case: In AI, audit trails are essential for tracking and logging the decision-making processes of automated systems

Significance in Generative AI: Audit trails help trace how data inputs influence creative outputs, crucial for ethical and responsible AI use.

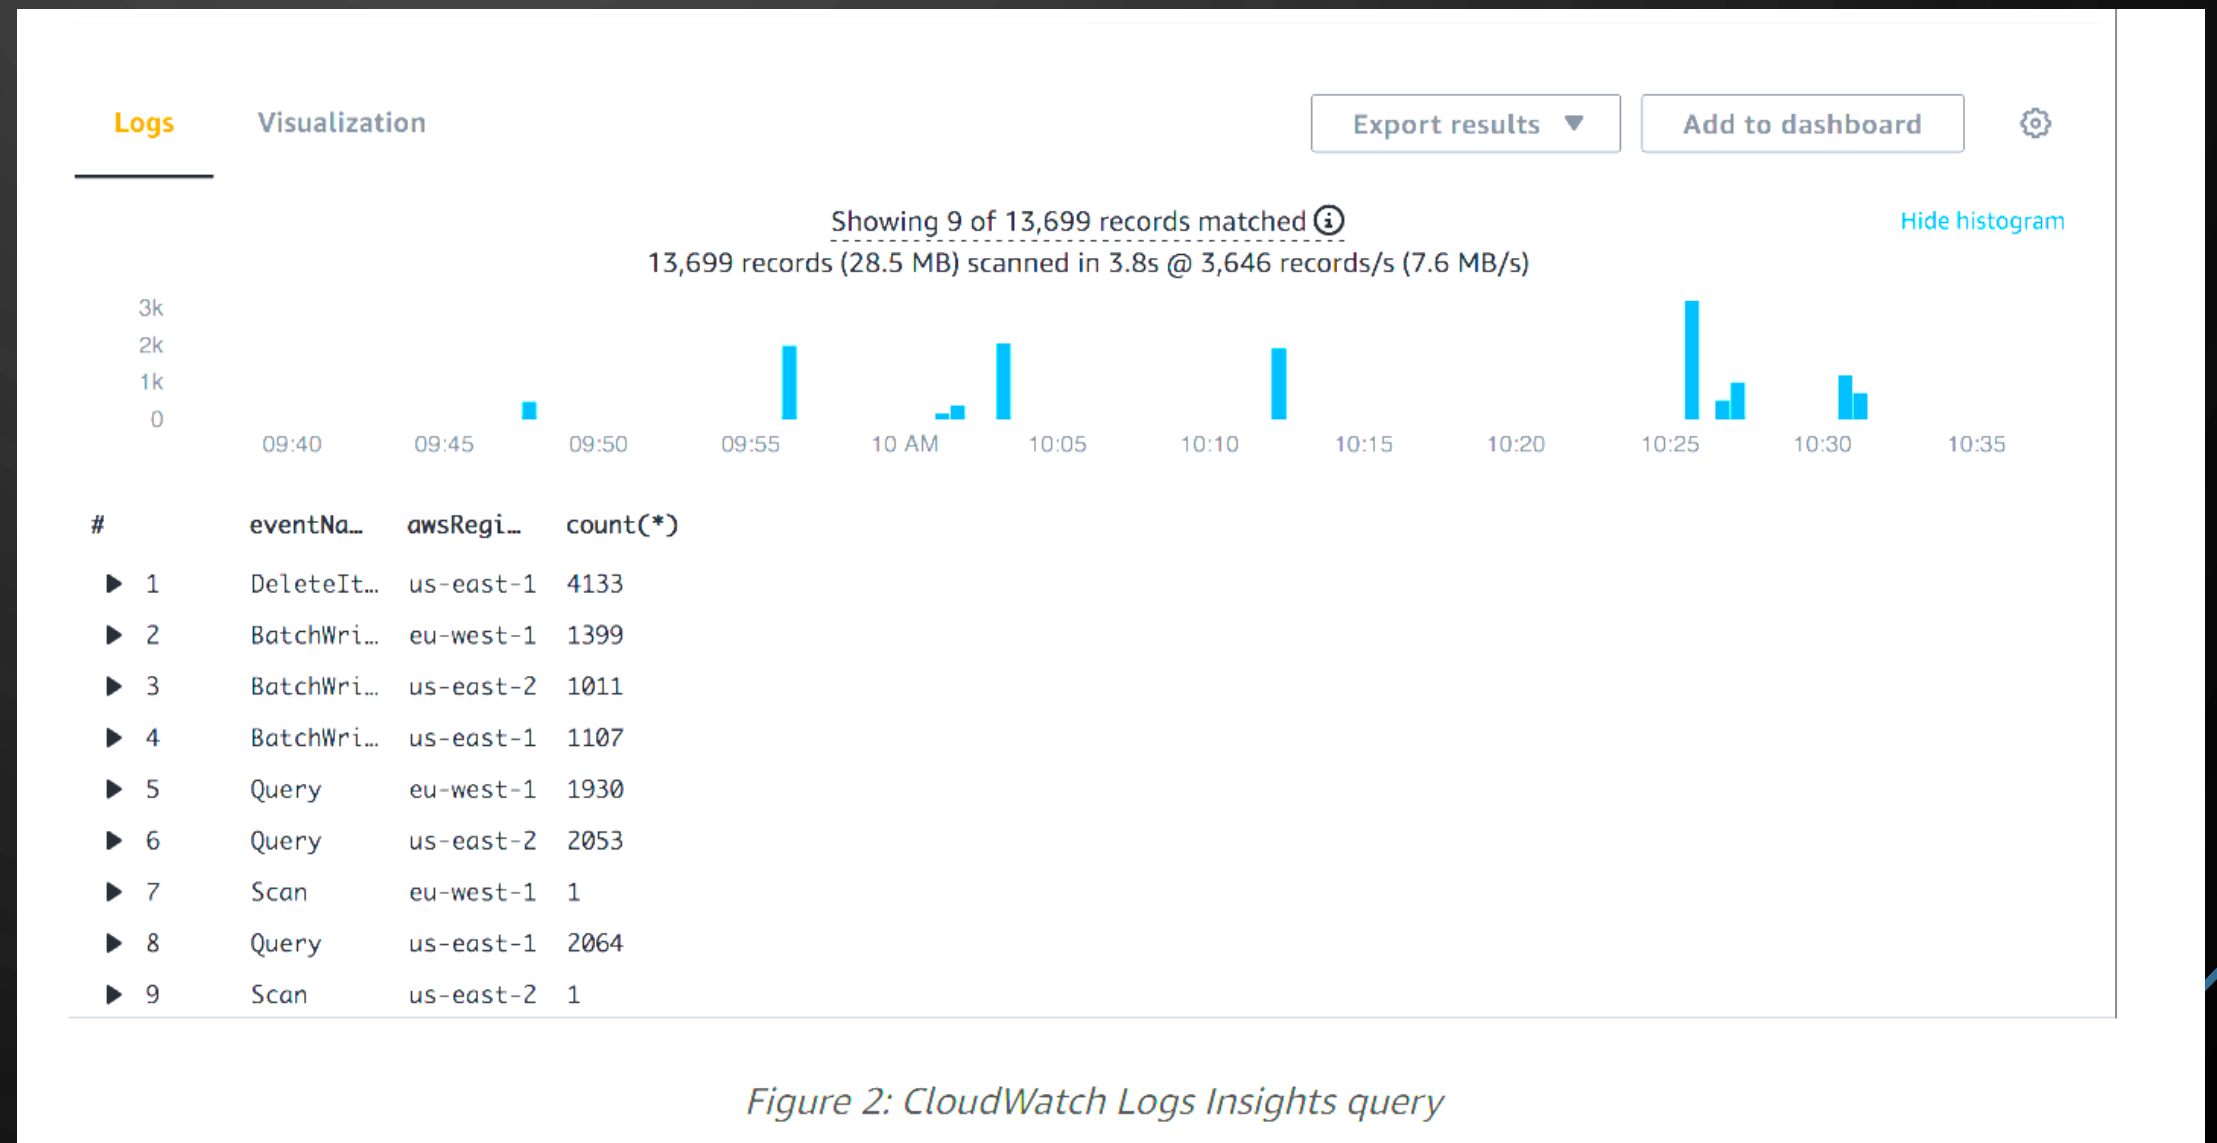
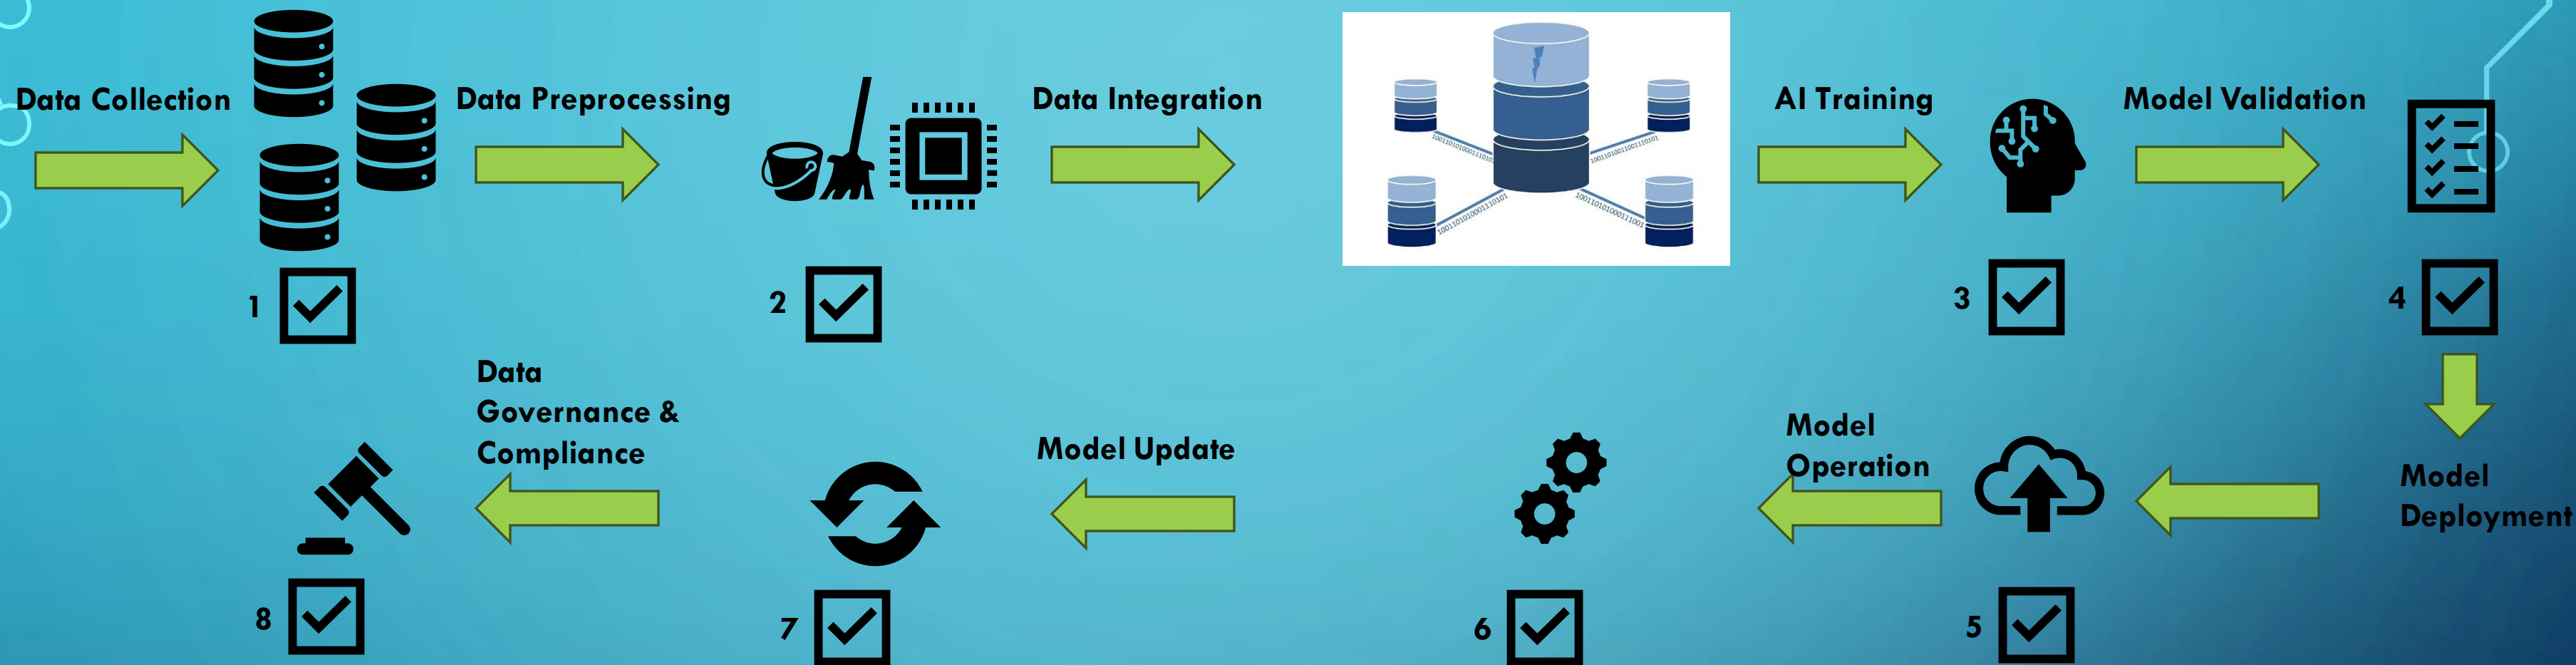


Figure 2: CloudWatch Logs Insights query

Implementing Audit Trails in Generative AI



1. Sources of data, timestamps, any preprocessing or transformation applied at time of collection

2. Details of data cleansing, normalization and labeling processes including changes to original data

3. Training data sets used, parameter adjustments, model iterations, and any changes or updates made to the model during training

4. Results of model tests, performance metrics, validation outcomes, and any adjustments or fine-tuning performance based on these results

5. Deployment details, including date, time, and environment of deployment, configuration and settings used.

6. Inputs given to the model, outputs generated, user interactions, and system responses, especially for real-time applications


7. Feedback data collected, updates or changes made to the model based on feedback and reasons for these updates

8. Records of compliance checks, governance actions, and adherence to data policies and regulations

Challenges in Data Audit Trails



Volume and Complexity of Data



Integration with Existing Systems



Performance Overhead



Data Privacy and Security Concerns



Complexity in Analysis and Interpretation



Balancing Detail with Manageability

The Necessity of Audit Trails in Generative AI

Ensuring Transparency, Accountability, and Compliance

Compliance Assurance

- Traceability for regulatory standards
- Verification of compliant processes



Ethical Transparency

- Accountability in AI decision-making
- Aids detection and correcting biases



Enhancing Data Governance


- Improves data management and risk mitigation
- Enables performance monitoring and contextual AI improvement



Keys to Effective Implementation of Generative AI




Data Collection & Diversity



Technical Understanding



Data Cleansing & Labeling



Risk Assessment

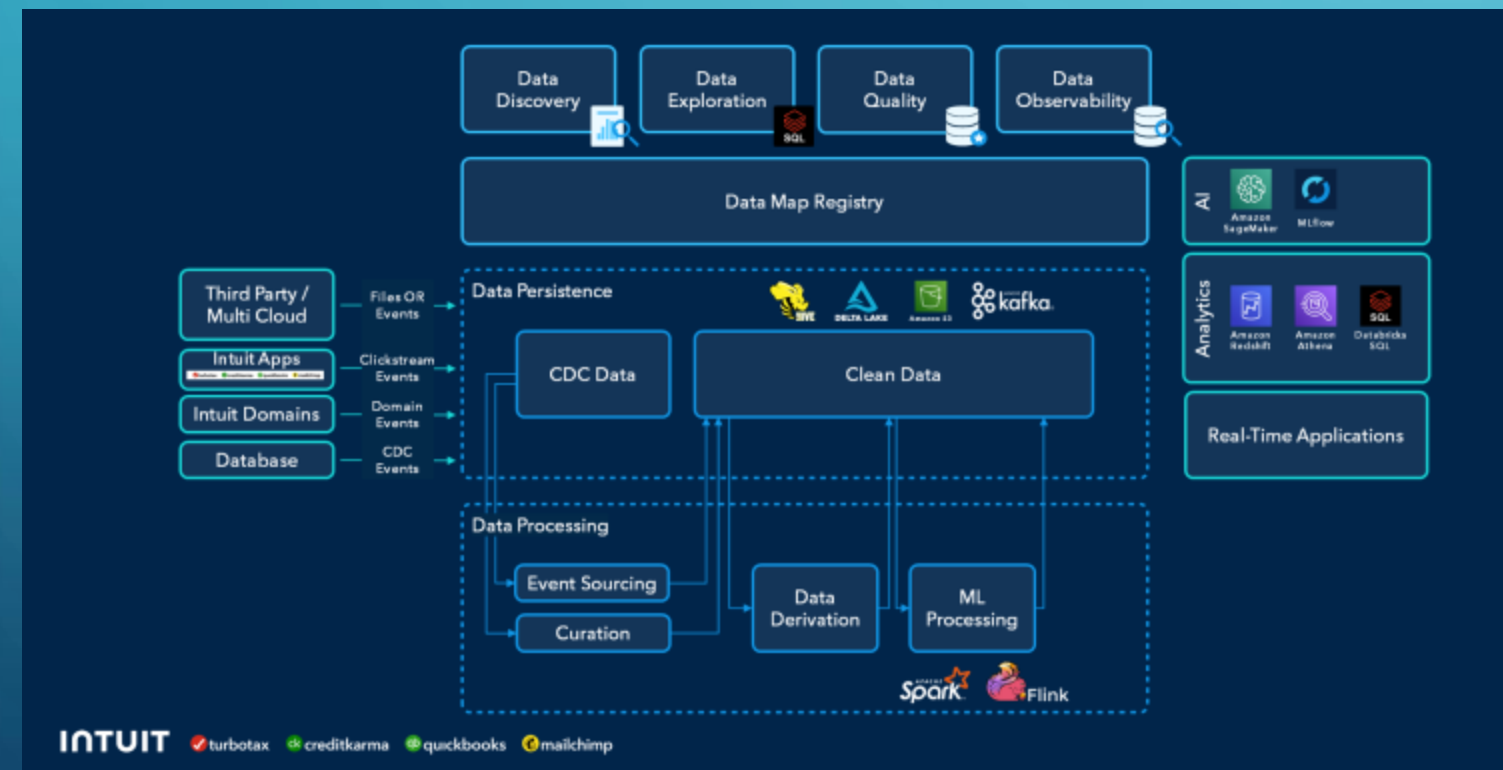
... but with effective data management.

Proper Data Management Allows for Effective Generative AI Implementation

Intuit Assist
Your new generative AI-powered financial assistant

[Watch the launch](#) [Find out more](#)

“...Intuit took several years to work through this data layer, to make sure data was well integrated, accurate, governed, and non-replicated. Only after doing this were LLMs able to call upon that data to allow personalized interactions with Intuit’s 100 million small business and consumer customers.” from VentureBeat interview with Alon Amit, VP of Product Management at Intuit



Intuit launches generative AI-powered digital assistant for small businesses and consumers

Jagmeet Singh @jagmeets13 / 7:00 AM CDT • September 6, 2023

The screenshot shows a digital assistant interface for 'Lennie's Pets and Plants'. It displays financial data such as 'Customer Cash Balance' of \$13,089.34, 'REVENUE' of \$19,312.44, and 'EXPENSES' of \$7,861.24. A 'Welcome to Intuit Assist' message is visible on the right side of the screen.

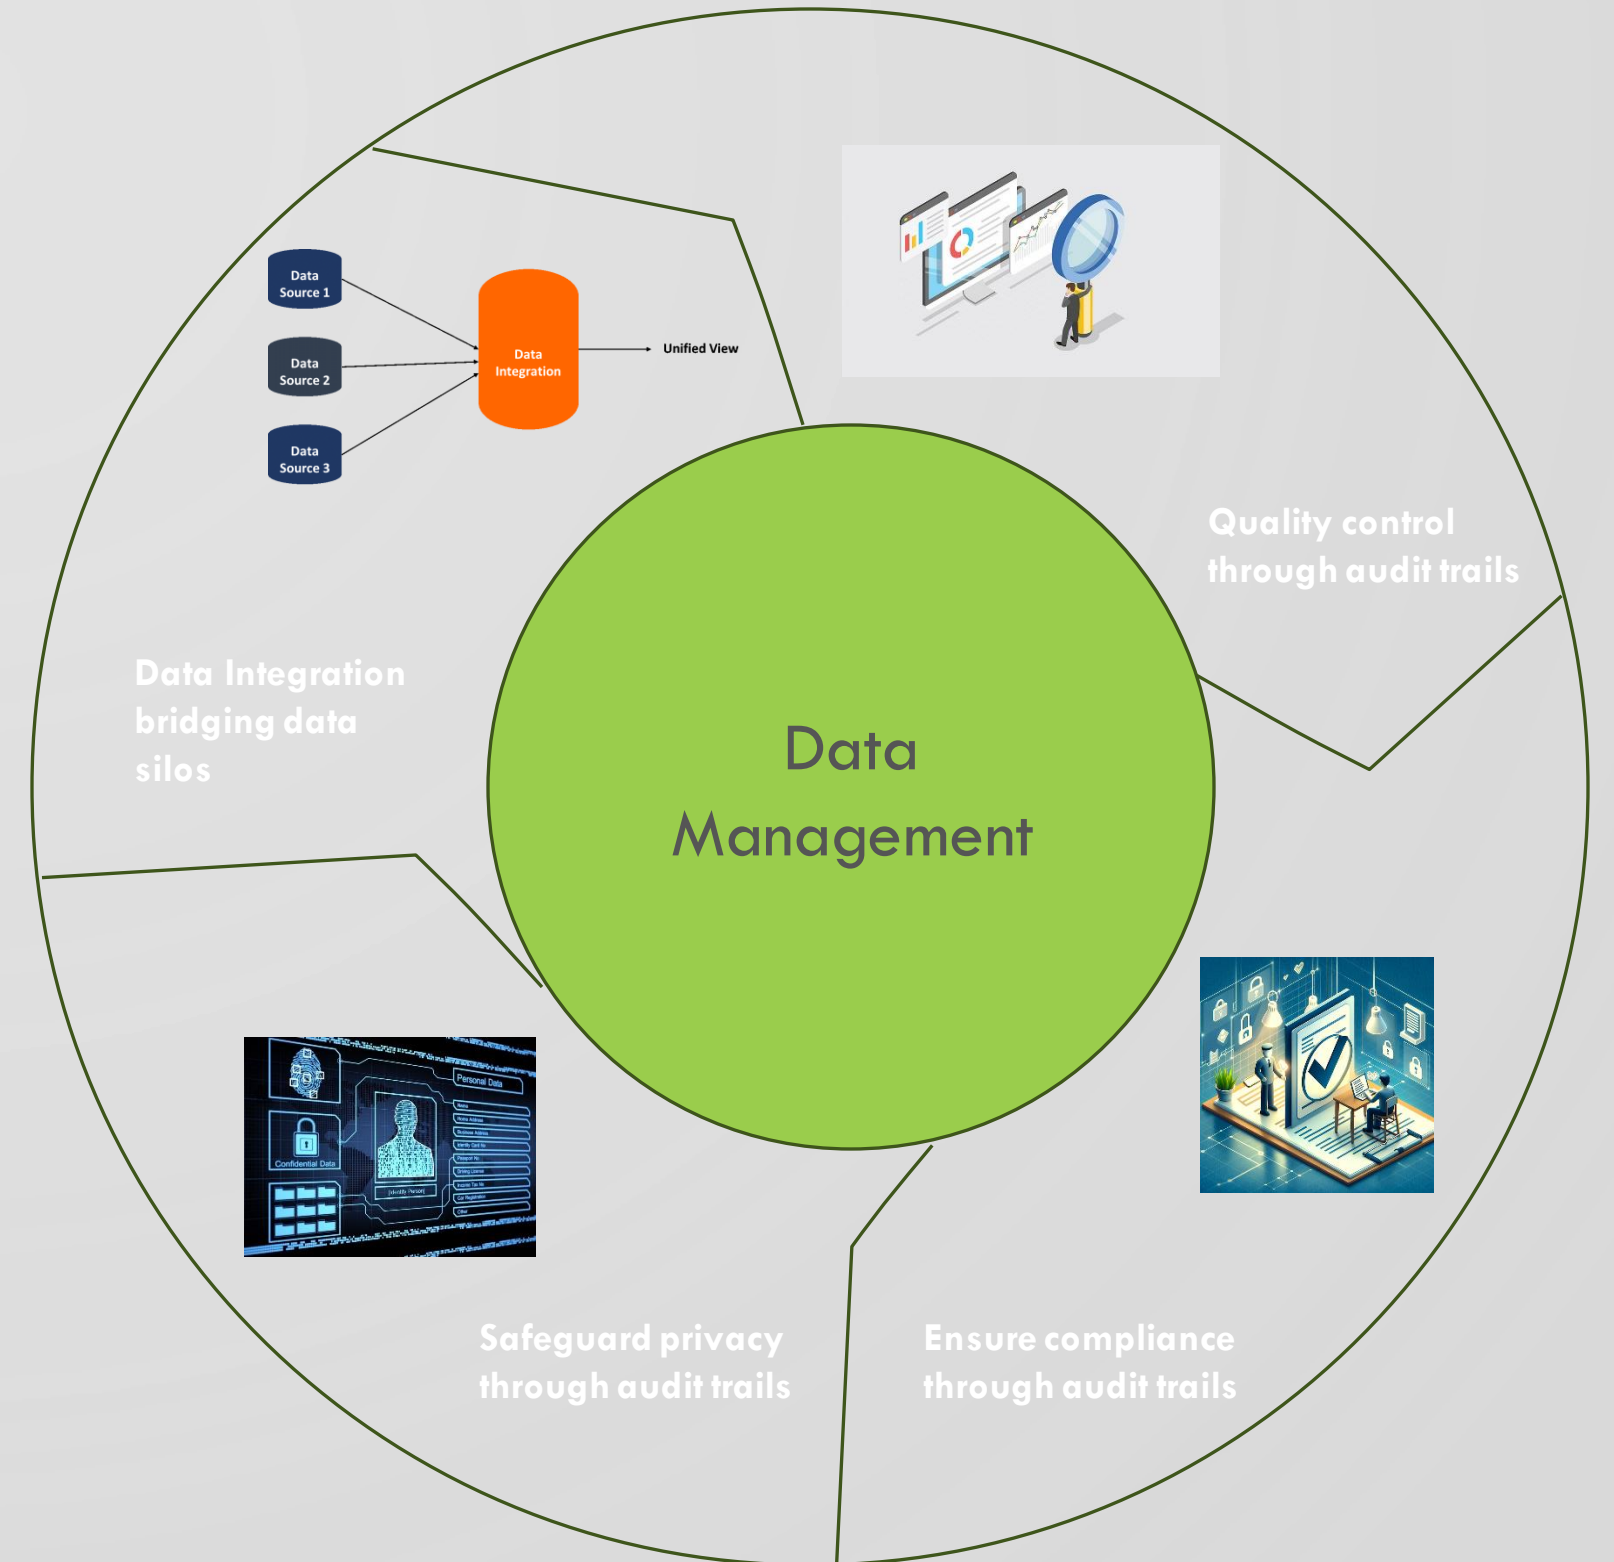
Image Credits: Intuit

Intuit, the U.S. financial and accounting software giant, has unveiled its first customer-facing generative AI-powered solution: a digital assistant to assist small businesses and consumers.

Called Intuit Assist, the digital assistant is embedded across Intuit's platform and products, namely TurboTax, Credit Karma, QuickBooks and MailChimp, with a standard user interface to offer personalized recommendations using contextual datasets to the company's more than 100 million small business and consumer customers across the world. The offering also provides human assistance using Intuit's live platform when needed.

COPYRIGHT CNS/ALLIC 2023
<https://venturebeat.com/ai/a-perfect-enterprise-data-stack-for-generative-ai/>


Data Management: The Heartbeat of AI Success




Empowering Business Transformation with GNS-AI



Data & AI Strategy



AI Solution Development
COPYRIGHT GNS-AI LLC 2023



For a consultation, contact me at
dr.amit.shah@gnsai.com



Compliance and
Ethics Advisory



Ongoing Support

Q and A

ASK AWAY!

