THE CAPCO INSTITUTE
JOURNAL
OF FINANCIAL TRANSFORMATION

GOVERNANCE OF TECHNOLOGY

Data and AI governance

SARAH GADD

BALANCING
INNOVATION & CONTROL

#59 JUNE 2024

a wipro company

THE CAPCO INSTITUTE

JOURNAL OF FINANCIAL TRANSFORMATION

RECIPIENT OF THE APEX AWARD FOR PUBLICATION EXCELLENCE

CONTENTS

GOVERNANCE OF TECHNOLOGY

GOVERNANCE OF SUSTAINABILITY

GOVERNANCE OF CORPORATES

DEAR READER,

In my new role as CEO of Capco, I am very pleased to welcome you to the latest edition of the Capco Journal, titled **Balancing Innovation and Control**.

The financial services and energy sectors are poised for another transformative year. At Capco, we recognize that this is a new era where innovation, expertise, adaptability, and speed of execution will be valued as never before.

Success will be determined based on exceptional strategic thinking, and the ability to leverage innovative new technology, including GenAI, while balancing a laser focus on risk and resilience. Leaders across the financial services and energy industries recognize the transformative benefits of strong governance while needing to find the optimal balance between innovation and control.

This edition of the Capco Journal thus examines the critical role of balancing innovation and control in technology, with a particular focus on data, AI, and sustainability, with wider corporate governance considerations. As always, our authors include leading academics, senior financial services executives, and Capco's own subject matter experts.

I hope that you will find the articles in this edition truly thought provoking, and that our contributors' insights prove valuable, as you consider your institution's future approach to managing innovation in a controlled environment.

My thanks and appreciation to our contributors and our readers.

Annie Rowland, **Capco CEO**

DATA AND AI GOVERNANCE

SARAH GADD | Chief Data Officer, Bank Julius Baer[1]

ABSTRACT

Data governance has come a long way from its inception in the 1980s, transitioning from a necessary overhead to a vital business capability enabling intelligence at scale. This article discusses the data governance journey to data governance 3.0, the role data products can play in risk-managed business self-service with a future view, and the lessons we can learn that will help move AI governance from infancy to value enabler at scale.

1. INTRODUCTION

Data governance – involving Excel spreadsheets and checklists to capture the business concepts represented by the data – has been around since the 1980s. It was viewed then as a "necessary overhead" and had no link back to the actual data. In essence, as Hinkle (2020) notes, it was "a process for cataloging large quantities of transactional data."

A Chief Data Officer's role in that foundational period was to simply collate concepts and create inventories of these concepts. Updates were done infrequently, sometimes annually, through manual reviews, while data ownership was seen as a "technology problem" with little in the way of business accountability for the data being created.

This status quo remained in place until the early 2000s, right up to when the "digital transformation" and the "big data frenzy" came into being. This quickly led to what became known as "data governance 2.0" – essentially, to a new paradigm where "data as an asset" principles were created to enable modern, data-driven businesses.

Distilled, this new era can be explained by the phrase coined by Clive Humby in 2006: "Data is the new oil", which like oil, is "valuable, but if unrefined it cannot really be used" [Watts (2021), Talagala (2022)]. Data governance 2.0 embraced collaboration, broke down organizational silos, and spread accountability across more data governance specific roles alongside business ownership.

In 2018, the Wall Street Journal ran the headline "Global reckoning on data governance" [Loftus (2018)]. That was the time when data breaches at a number of global organizations resulted in decreased revenues due to reputational damage, making headlines around the world. On May 25th, 2018 the E.U.'s General Data Protection Regulation (GDPR) came into effect [E.U. (2018)], leaving many companies struggling to meet compliance standards.

That same year also saw artificial intelligence (AI) governance become a hot regulatory topic, with the European Commission working on developing the "Assessment list for trustworthy AI" (ALTAI), released in June 2020 [E.C. (2020)]. At the end of 2019, the Hong Kong Monetary Authority (HKMA) published a report titled "Reshaping banking with artificial intelligence" [HKMA (2019)], as part of a series of studies on the opportunities and challenges of applying AI technology in the banking industry. The Bank of England and the Financial Conduct Authority launched the "Artificial intelligence public-private forum" (AIPPF) on October 12th, 2020. On April 21st, 2021, the AI Act was officially proposed, with an agreement being concluded on December 9th, 2023 [European Parliament (2023)], while the Monetary Authority of Singapore published a toolkit for assessment of AI by financial institutions in June 2023 [MAS (2023)].

[1] Contributor: Bea Schroettner, Certified Data Ethicist, Bank Julius Baer. Edited by: Natalie Martini.

The result? Nations across the world are either updating existing regulations on data privacy and copyright, looking to create new AI specific regulations, or are searching for ways to embrace guiding principles such as the G7 AI Code of Conduct (currently in development) [OECD (2023)].

We have now entered the era of "data governance 3.0". What does this look like?

At its core, this is about utilizing data science and improved technologies to treat data governance as a true enabler for organizations. Large language models (LLMs), AI, and active metadata,[2] breathe life into all of the artifacts that were captured over the last two decades. Data governance 3.0 is a living part of the organization, improving efficiency through integration and automation. Compliance, data quality, and effective data management are built in by design, not add-ons at the end of a process.

But what is "AI governance 1.0"?

In essence, this is about building the foundations that will enable safe, ethical, scalable use of AI, in a world of fast-evolving regulation and technology.

Exponentially increasing unstructured data volumes, computing power, and citizen analytics and data science capabilities, offer organizations the treasure of more and more data intelligence. But this all comes at a cost. As we saw in 2018, when data governance faced a global reckoning, the risks associated with providing AI tools without the culture or the knowledge is elevated. The hard lessons that were learnt from the data governance journey need to be implemented if we are to evolve AI governance. Focus needs to be on education, culture, and strategic alignment as key facets of successful AI governance.

In short, it is not just about governing the model underlying the AI solution. AI governance is everyone's role. Governance must operate in the delicate balance between regulation and risk mitigation on one side and enablement and innovation on the other.

If this balance is achieved, well-designed governance can generate tangible value while evolving with a future that remains unknown.

Peter Drucker, one of the 20th century's leading management theorists, put it well: "The greatest danger in times of turbulence is not the turbulence; it is to act with yesterday's logic" [McConnell (2020)].

2. DATA GOVERNANCE 3.0

The International Data Management Organization noted: "Data governance is defined as the exercise of authority and control (planning, monitoring, and enforcement) over the management of data assets. [...] Data governance focuses on how decisions are made about data and how people and processes are expected to behave in relation to data" [DAMA International (2017)].

Implied in this definition is the alignment with a more traditional governance model, which lacks the dimension of what governance should be actively promoting: the desired outcome. In other words, to ensure that discoverable, curated, high-quality data is securely available to users – as and when they need it. Put differently, an "enabler" that brings together high-quality data and consumers of data to deliver trustworthy data-driven insights.

With the rise of big data alongside advances in computing power, the interest in generating insights from data has skyrocketed in the last decade. With the increased importance of data science and data-lead decision making, a range of data topics were pushed into focus, data quality being the most prominent [Brous et al. (2020)]. The fact that data scientists spent, and arguably still spend, a significant amount of their time cleaning and organizing (poorly governed) data [McKinsey (2020)] before any value generation, further highlighted the need to change the data governance approach. At the same time, highly publicized data breaches and failures reiterated in parallel the need for the gatekeeping aspect of the data governance role to become more prominent [Famularo (2019)].

Data governance 3.0 strives to achieve an effective way of balancing risk control with user-enabling innovation and insight generation. The ability to extract high-quality insights from data is maturing from being a competitive advantage to a necessary hygiene factor. George Fuechsel, an IBM programmer and instructor, is generally credited with coining the term "garbage in, garbage out" (GIGO) in the early 1960s [Awati (2023)], and 60 years later it still remains one of main hurdles for enabling data value generation, for both business intelligence as well as generative AI (GenAI).

[2] Metadata is a set of data that describes and gives information about other data, e.g., whether a piece of data is a personal identifier.

How can we realize the data governance 3.0 benefits? I believe that we need to stop thinking about data as just an "asset" and start thinking of data as a product ingredient, and as with all product ingredients, apply consumer safety standards. The core meaning of data hygiene has not changed, it is still the absolute need to understand the quality of a piece of data and what that piece of data can be used for. What has changed is the ability to use machine learning and LLMs to vastly improve data quality detection alongside robust data classification. One of the high barriers to data insights has been the ability to access the data itself, with estimates of between 50 to 70% of time being spent just getting access to the data you need to answer a question. Data access automation and attribute-based access control can now be realized by converting internal policies into sets of machine-readable rules, which, when overlayed with the attributes of the data consumer, their patterns of data usage, and the attributes of the data, can streamline data access greatly, thus reducing the time to answer the question (i.e., time to insights).

Data governance 2.0 moved from "concepts and cataloging" to physical data, while data governance 3.0 activates the physical data level by using data science approaches to understand data securely at scale. The governance roles that ensure the ownership and accountability for data need to remain in place but demand empowering through technological advancements, not manual exercises. Data governance 3.0 should embrace the use of technology from the moment data is created, to when that data is deleted (the data lifecycle). You need to augment governance through embedding AI/machine learning algorithms in the data lifecycle, so they can do what they are good at: dealing with vast amounts of data to classify, qualify, and enhance. By doing this, you provide the data governance "human in the loop" with fast insights they can use to make informed data governance decisions.

One practical example that applies in many companies is "entity resolution", i.e., when data is coming in at scale with different identifiers (including names) that actually represent the same thing – like client names, third-party names, and inventory items. There are proven machine learning techniques that will mine the data as it flows in and create clusters with a probability score on the data the machine believes represents the same entity. These clusters are presented to the human to validate and the machine continues to learn. In data governance 2.0, this was an extremely time consuming and often impossible task.

Some aspirational approaches to be considered across the data management lifecycle are (Figure 1):

- **Data creation:** pattern recognition to automatically classify whether a data attribute is a birthdate, social security number, third-party name, client ID number, or other personal identifying information that is deemed critical and needs a higher level of protection.
- **Data storage and archiving:** auto-classification of data (and records) that need to be stored in an archive for regulatory/business purposes and matching them to the length of retention that applies.
- **Data lineage:** tracking of data lineage to identify data that is not from authoritative sources versus the data that is, and where the data is being manipulated so it no longer represents the "truth".
- **Data usage:** use of a combination of machine-readable controls and attributes of the person trying to access the data (e.g., role, point of time location, normal access patterns, etc.) to provide the data as readable or obfuscated (with patterns intact for data scientists) in the environment needed by the user, whether for development, business intelligence, analytics, or data science.
- **Data quality:** applying AI to help facilitate the improvement of data quality, e.g., through data standardization, data validation, or data governance compliance checks and other features [Drenik (2023)]. Virtually all major data quality management tools already contain this functionality, which provides the standards and guardrails within which domains should operate. Setting up these tools is a critical enablement opportunity for a data governance function.
- **Data deletion:** machine-readable retention rules cross-referenced with legal hold information to enable the compliant on-time deletion of data and records, either from legal archives or from operational systems through API calls.

While data governance and data management efforts have traditionally been focused mainly on structured data, expanding this effort to unstructured data (e.g., documents, emails, and contracts) is of growing relevance. The amount of unstructured data is rapidly increasing – some estimate as much as 90% of a company's data to be unstructured [Violino (2023)]. By its definition, unstructured data does not follow a clear schema or data model and it may contain personal or sensitive data that is harder to spot.

Figure 1: Data lifecycle

The data lifecycle management framework will ensure governance and control of the data at the different data lifecycle stages

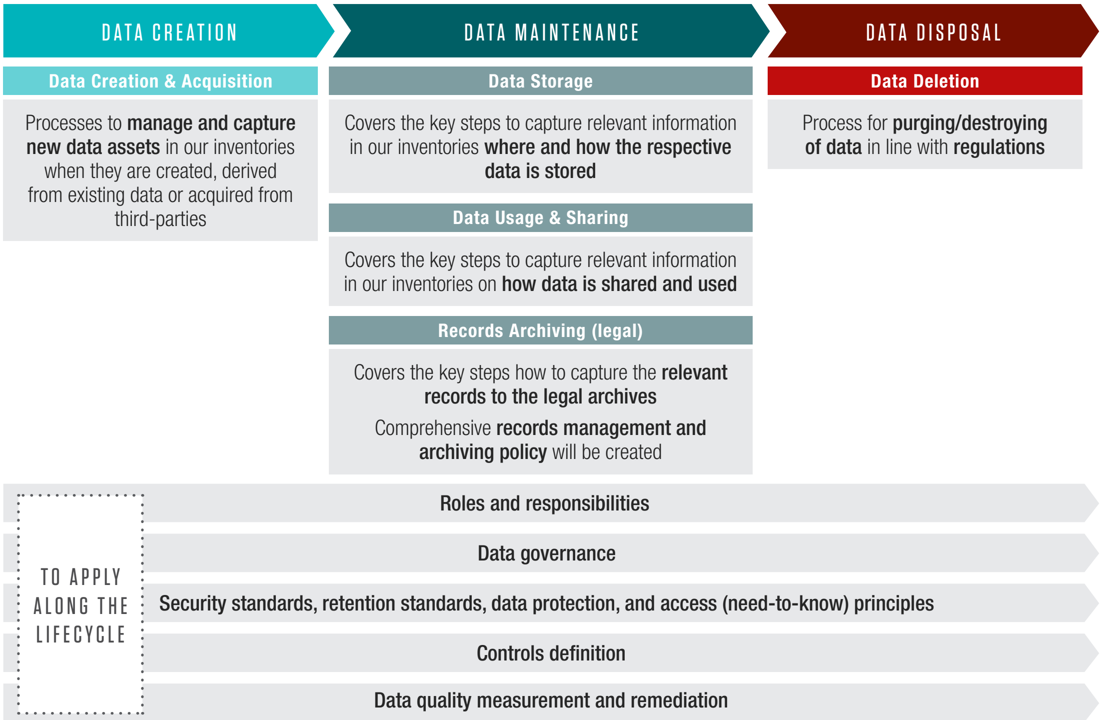

Source: Bank Julius Baer

Figure 2: Data ethics canvas

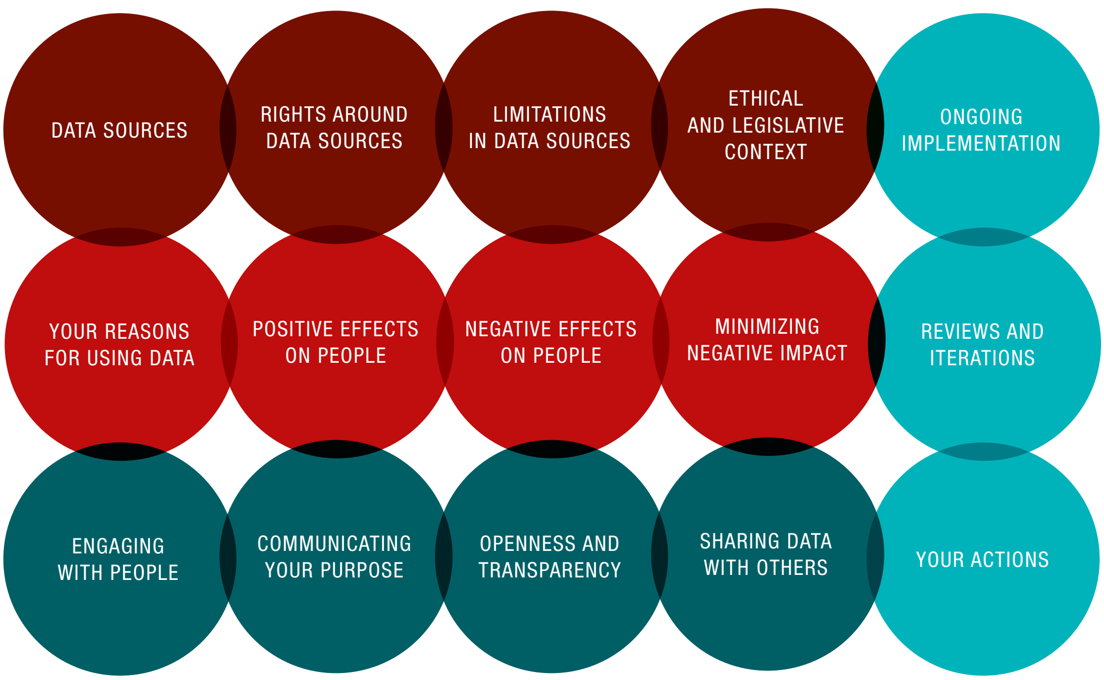

Source: Derived from Open Data Institute (ODI)
License: CC-BY-SA 4.0

Figure 3: Enabling AI at scale

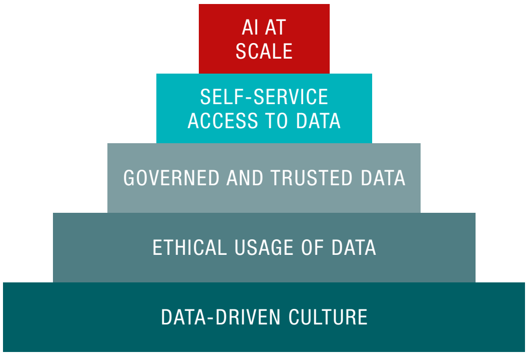

Source: Bank Julius Baer

GenAI and the power of LLMs have commoditized the capability of extracting value and insights from semi-structured and unstructured data, though trust in outcomes remains a challenge. "Garbage in, garbage out" also applies in the unstructured world, as do the normal dimensions of good data governance, including overall data hygiene [Abdullahi (2023), Rosencrance (2024)]. For example, unstructured data could be of questionable quality (e.g., multiple or duplicate versions of the same document). LLMs have enabled us to better identify and classify the data ingredients within the unstructured world. For instance, you can use models to identify and classify clauses in contracts, sensitive or personal identifying data, start dates, parties, and terms. These features make it far easier to mine data securely for intelligence. One could argue that wrangling unstructured data and applying governance is the larger value proposition and more likely to be a differentiator than the structured world of data.

With the incredible volume of data that enterprises are managing daily, the only way to curate the data ingredients is by using the data science tools and techniques that are available today and constantly evolving with the technology to bring further future value.

While AI ethics is a global topic, it is not a new concept for data. Hasselbalch and Tranberg (2016) was one of the early books to describe not only the privacy implications of the commercial exploitation of big data, but also the broader social and ethical implications. The Open Data Institute published a "data ethics canvas" in 2021, covering many of the aspects that are now in the news with AI ethics (Figure 2) [ODI (2021)].

Culture and ethics are vital aspects for successful data and AI governance and should be treated as critical governance dimensions (Figure 3). Everyone needs basic data literacy and awareness to understand the questions that should be asked when consuming or working with data.

So far, we have covered data governance 3.0 and adopting AI capabilities to build data governance by design across the data management lifecycle, as well as the importance of culture and ethics. Next, we will look at what happens if you do not have the resources and funding to invest in curating data ingredients, or perhaps even if you do. Enter "data products", which build on data governance 2.0 and 3.0, and in the immediate future will leverage the power of LLMs and GenAI to enable the business to self-serve automated data product creation.

3. DATA PRODUCTS

We have lived through the era of master data management, data warehouses, data lakes, and data lakehouses, and one challenge that consistently arises is "how do I keep all this data in sync"? Data is typically created in places that are fit for that type of data, whether that is software as a service, operational data stores, data integration layers, or even mainframes. The approach of a "one stop place for all data" has not worked, with many enterprises trying and failing [Woods (2016)].

Not all data is equal, and not all data has business value. Many enterprises focus on "critical" or "material" data, which at its core sounds good, but quite often the importance of the data is driven by the need for that data at any given time, which changes based on circumstances at that time.

"Data as a product" (DaaP) first appeared in 2019 as part of the "data mesh" concept defined by Zhamak Dehghani [Fowler (2019)]. Simply, a data product is a broad definition that includes any product or feature that utilizes data to facilitate a goal. Essentially, in addition to (or instead of) using the more manual data governance 2.0 approach, you could apply all the data governance approaches discussed in data governance 3.0 to a grouping of data ingredients rather than each individual ingredient. For example, if I wanted to create a data product that represented sales in the U.S., as the data owner for sales data, I could point to the individual sources for that data (whatever those might be) and put the quality measurement, security, data classification, on the product level instead of the individual attributes, and manage the data collection as a data product.

Figure 4: Example illustration of a data product framework

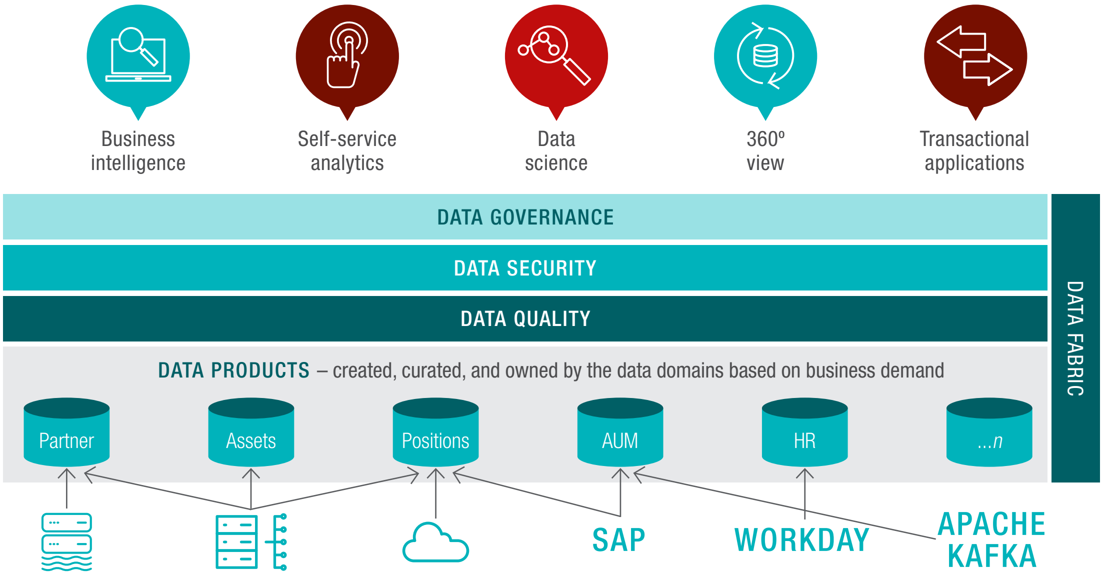

Authoritative strategic sources for the data used to create data products

Source: Bank Julius Baer

McKinsey & Co. define a data product as delivering "a high-quality, ready-to-use set of data that people across an organization can easily access and apply to different business challenges" [McKinsey (2022)]. Data products offer a number of benefits. First and foremost, they are built, curated, and maintained by subject matter experts (usually the data owners). As such, they represent an "official" version of the data. This reduces the effort for data analysts and data scientists, who would otherwise have to find the right data fields in raw data (data ingredients) and may integrate or manipulate them in a way that provides an answer that is not representative of the question.

Data products also save time through reuse – everyone can securely "shop for the product", not having to build it for themselves. Data products increase consistency across business intelligence and reporting, as people start with a "common view" and can then combine data products into a new data product to answer another business question. Like other products, data products can be advertised in a product store or a data catalogue and have an owner who can monitor usage and curates them throughout their lifecycle. Well-managed data products can reduce the time to implement use-cases significantly, by up to 90% [Desai et al. (2022)], and strengthen the concept of user self-service.

You can create data products off any kind of data store or combination thereof. This means you do not need to move all your enterprise data to one giant store somewhere, but rather you source the data ingredients for the data product from their "home store", for example, HR data from systems like Workday, finance data from SAP, etc. (Figure 4). Getting the data from the authoritative store for that data means that you do not need to constantly try to keep copied data in sync with the authoritative store, which, if not done properly, can result in data breaks and stale data being consumed by end-users.

Data governance still plays a major role in the creation of data products by, for example, providing clear standards around "metadata documentation, data classification, and data quality monitoring" [McKinsey (2023a)], as well as ways to enforce governance across the product lifecycle [Deighton (2023)] – where data governance 3.0 can play a role. In addition, the data owner needs to demonstrate data health for their product (just like health standards when you buy food products), otherwise how can users trust the data product they have been given?

The benefits derived from data products not only accrue to the consumers of the product, while driving business value, but also to data governance (Figure 5). They also help reduce risk [McKinsey (2022)]. By controlling how consumers can access

Figure 5: Data, AI, and GenAI governance – how they all fit together

WELL-GOVERNED DATA IS THE FOUNDATION OF GENERATIVE AI

“Garbage in, garbage out” still applies with GenAI, only on a massive scale

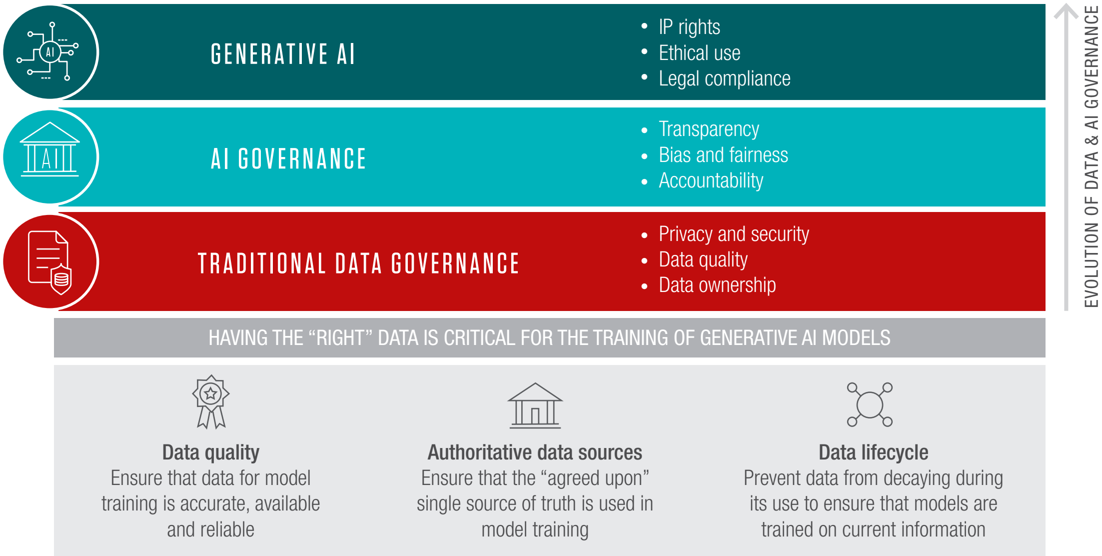

Source: Bank Julius Baer

your products, you can put the right safety standards on the products, which reduces the risk management complexity of doing this across all the data ingredients. This is analogous to shopping in a supermarket: when you buy a tin of soup, you trust the manufacturer, you trust the container, you do not need to review every single ingredient (though they are listed so you have the option), instead you trust that the tin of soup is going to be exactly what you thought it would be.

With the power of GenAI and LLMs to mine metadata and generate code, the ability for the business to create data products as code using natural language is emerging, further commoditizing data product generation with the business owners of the data being able to self-serve. Applying LLMs to a) allow domain experts who have business, but not necessarily coding skills, to specify the data products they wish to build and b) extract the relevant metadata to build the required data pipelines, overlaying security and governance, further builds upon the concepts of data governance by design.

4. AI GOVERNANCE: MACHINE LEARNING TO GenAI

AI governance can be defined as “a system of rules, practices, processes, and technological tools that are employed to ensure an organization’s use of AI technologies aligns with the organization’s strategies, objectives, and values; fulfills legal requirements; and meets principles of ethical AI followed by the organization” [Birkstedt et al. (2023)]. Part of AI governance is the layer resting on top of data governance, as AI solutions, at their core, consist of input data, the models or algorithms trained for specific tasks, and their output [IBM (n.d.)]. Model input and output are data and as such, benefit from a strong and effective data governance across both structured and unstructured data.

However, AI requires additional facets of governance. AI models often automate decisions and/or processes – due to the increasing complexity of AI solutions, understanding and explaining how a decision was arrived at can be challenging. Model output can sometimes display unwanted bias or could be discriminatory against certain groups. With GenAI, intellectual property violations have been widely reported in the media [Appel et al. (2023)].

Given the growing importance of AI solutions, governing bodies around the globe, as well as technical experts and corporations, are trying to define guardrails and legislation to mitigate the risks associated with AI, balanced with innovation and the benefits the solutions bring. This is the exact same goal mentioned previously regarding data governance: it strives to achieve an effective way of balancing risk control with user-enabling innovation and insight generation. There is general consensus around the broad areas that require

Figure 6: Ethics guidelines

REGULATORS CARE ABOUT ETHICS

E.U.'s "Ethics Guidelines for Trustworthy AI 7 Key Areas"

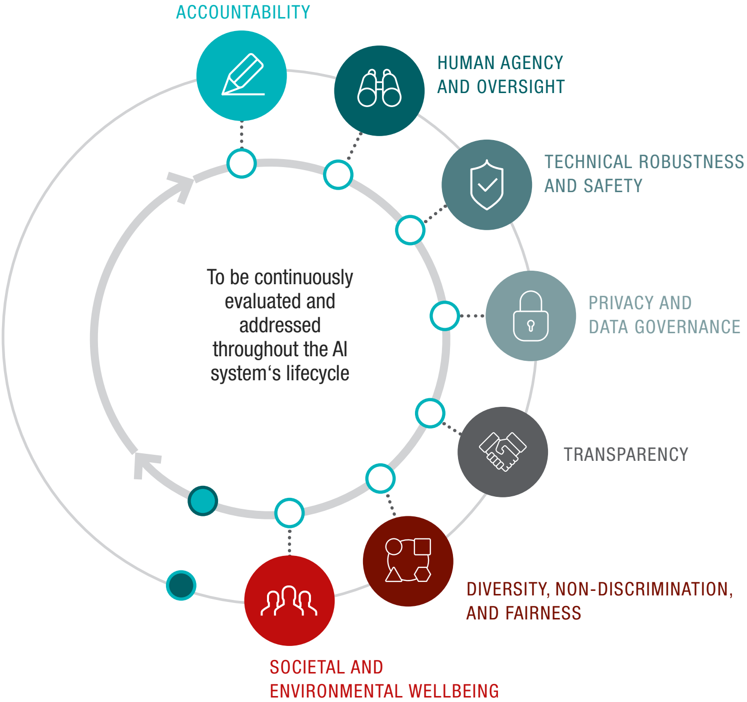

Trustworthy AI has components, which should be met throughout the system's entire lifecycle:

(1) It should be lawful, complying with all applicable laws and regulations;

(2) It should be ethical, ensuring adherence to ethical principles and values; and

(3) It should be robust, both from a technical and social perspective since, even with good intentions, AI systems can cause unintentional harm.

Source: Adapted from the E.U.'s Guidelines for Trustworthy AI

Figure 7: Keeping the end-to-end view in mind when building an AI governance framework

ENABLING SECURE ARTIFICIAL INTELLIGENCE

Build upon existing structure and capabilities, close gaps

DRIVERS OF CHANGE

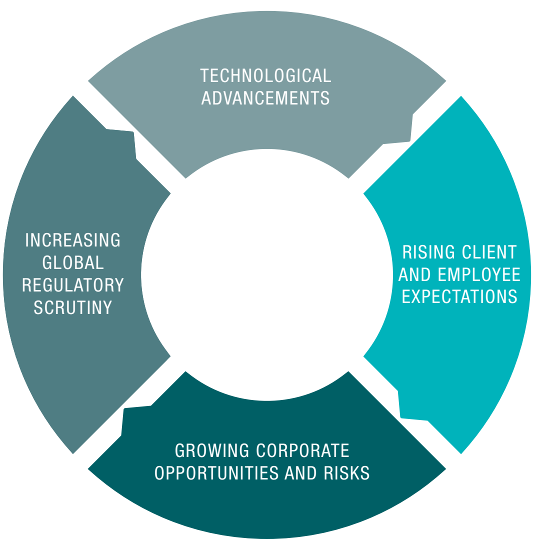

AREAS FOR ATTENTION

Strategy
- Strategic alignment
- Value capture with measurement
- Talent acquisition and retention

Governance
- Data governance
- AI governance
- AI ethics

Culture
- Awareness
- Literacy
- Training and skills development

IT
- Infrastructure requirements
- Machine learning operations framework

Risk controls
- Legal and infosec frameworks and policies
- Procurement third-party risk management
- Model risk management

Source: Bank Julius Baer

attention, with a spotlight on ethical considerations around topics such as fairness, bias, and explainability – topics that are also covered under data ethics.

The E.U., to name just one example, proposed in their Guidelines for Trustworthy AI, that all AI solutions be lawful, ethical, and robust technologically and socially (Figure 6). Regarding ethics, it specifies four ethics principles: respect for human autonomy, prevention of harm, fairness, and explicability. It also suggests seven requirements to realize these principles: human agency and oversight; technical robustness and safety; privacy and data governance; transparency; diversity, non-discrimination, and fairness; societal and environmental wellbeing; and accountability [AI HLEG (2019)].

For organizations, establishing AI governance to meet the requirements of regulators and legislators globally is critically important for ensuring that implemented solutions will withstand the test of time, and will not fail to evolve alongside regulatory requirements. A recent study by Ernst & Young found that while organizations and regulatory bodies broadly agree on the areas of focus for trustworthy AI, the importance of the individual principles is weighted differently [EY (2023)]. In addition, the regulatory landscape is still in flux, so the full scope of final legislative requirements cannot fully be judged yet.

As learned on the data governance journey, to govern AI efficiently within organizations requires a cross functional approach [Schneider et al. (2023)]. In addition to basic governance steps, such as defining principles for good model development and specifying AI principles that align with corporate values, experts from different domains need to collaborate make AI governance useful across a model's lifecycle (Figure 7).

Complementing the data governance experts, who curate and can help identify high-quality data sources, experts from the legal, information security, and IT domains are needed to ensure the proper operation of models. Data scientists and model risk managers need to monitor and validate model performance throughout the model lifecycle. Beyond the operational, governance bodies need to be established or upskilled to check for ethical considerations and risks associated with models [Blackman (2022)]. In addition, policies and controls need to be updated, or newly created, to address AI-related risks and to provide guidance to those working with the models. It is essential that this does not create additional overhead for data users, who face pressure from business management to provide information or solutions fast.

Applying what we learnt from data governance 1.0, we cannot start by only looking at the models and risks, we must include the consumer and business perspectives, and leverage technology as an enabler. Based on industry experience, only 10-20% of AI ideas and early proof of concepts actually make it all the way to production. Taking the right AI governance steps early in the process can increase the chances for success.

A number of points that could be considered at the ideation stage of a potential use-case include:

- **Strategy:** discuss if the use-case is really a direction you want to go as a firm. What would be the unintended outcomes if successful (impact on people, clients, and risk) and decide whether to park or move forward.
- **Governance:** could the use-case cause ethical issues or lead to a negative impact on society. Do you have the data to support the use-case, or only "some of it", which means outcomes will not be representative.
- **Culture:** do the people working on the use-case have the right skills and know the questions they should be asking from both ethical and risk perspectives.
- **IT:** can the data the team will be using legally be used in the environment they are looking to experiment in. Are they building in an environment that could never be productionalized.
- **Risk and controls:** does the team have the rights to use the input or output data, or could intellectual property be violated, or could data from third-parties be used in a way that would violate contracts.

A lightweight AI governance tollgate process at the idea stage can help avoid wasting time and resources on use-cases that will never go to production. Experimentation and exploration are nevertheless important, but should be understood by the business to be exactly that to set expectations.

Similar to the concept of security by design in software engineering, and the adoption of technology to enable data governance 3.0, AI can help AI governance, especially when it comes to managing the risks associated with the models themselves.

AI does not only support the identification of models across the organization. Embedding machine readable controls at the different stages of the model development lifecycle into the models themselves, unlocks the capability to validate models against different regulations, as well as continuously testing the models once in production to monitor their performance

Figure 8: Areas where data literacy and awareness are particularly critical

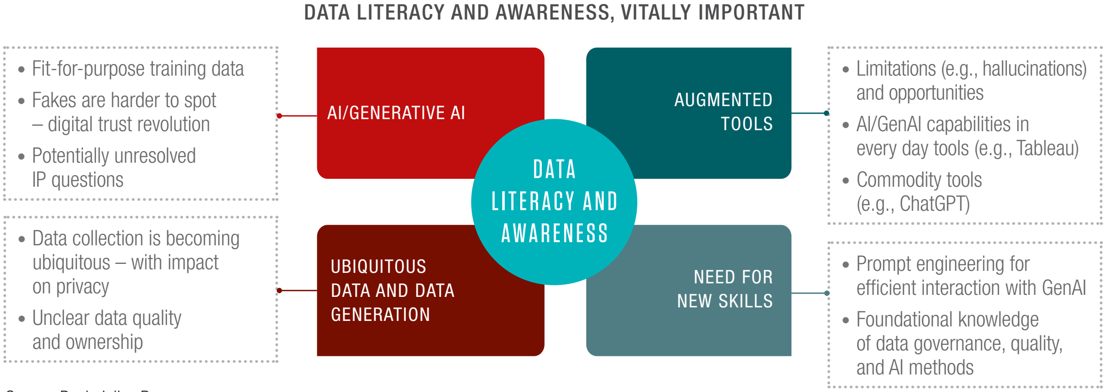

Source: Bank Julius Baer

and to ensure they do not "drift" into a state that is no longer within the expected risk tolerances [Aristi Baquero et al. (2020)]. The controls themselves become part of the model, embedded in place during model versioning, testing, and release cycles, and can be enhanced with further controls, as regulation and societal needs change.

Products are emerging that provide sets of machine-readable controls that represent different regulations or standards, and these controls can be downloaded and embedded into the models themselves as part of the model development lifecycle. Many of the vendors who are providing pre-built models/AI capabilities are being pressured to provide evidence of compliance with regulatory demands. With the regulatory space evolving so rapidly, creating the controls in a way that they can easily be embedded into models will benefit organizations that are "training their own" models, and vendors who can build in the controls as part of the offered solutions. It will be extremely difficult for enterprises to scale their use of AI without embedding controls as part of the model design.

The launch of ChatGPT in November 2022 brought GenAI and its capabilities to the forefront of public attention. Hailed as a tipping point for AI just two weeks after its launch [Mollick (2022)], McKinsey & Co. estimate the value potential of GenAI to be between U.S.$2.6 trillion and U.S.$4.4 trillion annually [McKinsey (2023b)]. GenAI solutions are usually built on foundation models, which are "pre-trained on large, unlabeled datasets and capable of a wide array of applications [… and] can then be fine-tuned for specific tasks" [IBM (n.d.)].

From a governance perspective, foundational models introduce another layer of complexity that can be addressed across three broad categories:

- Foundation models and commercial applications are usually trained outside an organization's control, so there is no understanding as to whether the data used for training is representative and legally allowed to be used [Bommasani et al. (2023), Heaven (2023)], for example with regards to the use of plagiarized content or content that is created on copyrighted materials. Companies offering foundation models only publish select information and justify the lack of transparency with the protection of trade secrets, as well as the risk of bad actors gaming or hijacking the models [OpenAI (2023)].
- LLMs are statistical models at their core and come with certain limitations that are impactful, but the limits are not well understood. This is especially apparent in GenAI with hallucinated answers, which are statistically probable, but completely untrue.
- With the promise and popularity of GenAI, many vendors are adding components to their offerings (e.g., copilots), which increases the pressure on third-party risk management and technology providers to stay on top of these "features" being released into existing tools (some providers, such as Microsoft, offer to shield users of their models from possible lawsuits).

Regulators and organizations alike are keen to capitalize on the benefits that GenAI solutions offer but are also trying to understand and guardrail its specific risks. While the E.U. has included foundation models in its risk-based approach to AI regulation [European Council (2023)], the risk-profile is still evolving, which will add further turbulence to the AI regulatory space.

Outside of the risks mentioned above, the most important topic to cover when considering a corporate governance approach to GenAI, is awareness building with end-users coupled with data literacy (Figure 8). They need to understand the limitations (e.g., hallucinations, plagiarism, etc.) and the risks (potential loss of personal data) of foundation-model-based solutions, alongside their accountability (e.g., checking the correctness of content). Upskilling on how to most effectively interact with the solutions (e.g., prompt engineering) can help drive user-developed solutions for the areas they are the experts in, while understanding the risks involved.

5. CONCLUSION: APPLYING DATA GOVERNANCE LESSONS TO AI GOVERNANCE

The evolution of data governance taught us key lessons that we can now leverage as enterprise AI governance matures. From concepts through to actionable metadata linked to physical data, we learned that technological advancements far outpace our ability to govern through traditional methods.

Similar to "security by design", which we see embedded in software engineering around the globe, governance needs to be built-in as part of the design and become a natural part of the ecosystem. In the case of data, the physical data assets themselves need to contain the metadata that enables the identification of risk, privacy, security, quality, and usability aspects of that data to enhance business and shareholder value. In the case of AI governance, the AI governance controls need be incorporated as part of the code of the model, generating the artifacts and evidence needed for model validation, trustworthiness, and ongoing monitoring.

The capabilities of LLMs will enable faster evolution of regulation and expectations on reporting. Today, regulatory and governing bodies work in the world of analogue rules and principles that are open to interpretation when being implemented by the organizations in scope. Since 2018, we have seen a number of regulatory bodies explore a more digital machine-readable approach to rules and regulation [PwC (2021), Ledger and McGill (2023)], which I expect will be further enabled through the strengths that LLMs have to turn unstructured non-digital content into machine-readable content.

All the data modeling work that regulated companies have undertaken to meet data regulations would reap even greater benefits, given that digital regulatory reporting (DRR) requires common data models to be effective and to ensure all parties are "speaking" the same language. This is just one example where the evolution from data governance 1.0 to 3.0 shows us that we need to create the building blocks of the future today.

Organizations that have not linked data governance to physical data, or have not captured the needed metadata, or not modeled the data, are going to have a much harder time trying to meet future demands while generating business and shareholder value. Real-time financial and regulatory reporting may have sounded like an unachievable goal ten years ago, and it may be another ten-plus years before it becomes a reality, but it is certainly something that companies need to be creating the foundation for.

With the fast evolution of AI regulation taking shape around the world [IAPP (2023)], it is obvious that the "global reckoning on AI governance" is coming in the not-too-distant future, where we are seeing an initial divergence on a global scale (similar to the divergence that took place on data protection when the GDPR was introduced in Europe). Certain countries/geo-political alliances will take on more risk, regulate less, trying to leverage the capabilities of AI to upskill populations and improve economic conditions. On the other extreme, we have the heavily regulated E.U., which will struggle to innovate under the burden of expansive regulation [Greenacre (2023), Jorge Ricart and Alvarez-Aragones (2023)].

Converting regulation into machine-readable control frameworks, which can be modified, enhanced, and added to, enables the controls to mature and shape alongside the regulation. It is key to embed these controls as part of the AI development lifecycle, so as the controls change, they can easily be applied to both existing and new AI models. For example, you may have an E.U. AI Act set of controls as code, which can be called from different points in the AI development lifecycle and post-production for ongoing monitoring. This is not a new concept. In data lifecycle management there are several machine-readable controls that are applied at different stages, from data privacy classification to when data can be erased – pieces of callable code, ranging from standard scripts to running a machine learning algorithm at a certain time in the data lifecycle – a data governance 3.0 lesson we can leverage to take AI governance from infancy to a value-generating, scalable asset.

I will close with another Peter Drucker quote: "The relevant question is not simply what shall we do tomorrow, but rather what shall we do today in order to get ready for tomorrow" [Power (2018)].

REFERENCES

Abdullahi, A., 2023, "Five ways to improve the governance of unstructured data," TechRepublic, August 28, http://tinyurl.com/yc5vy46j

AI HLEG, 2019, "Ethics guidelines for trustworthy AI," High-Level Expert Group on Artificial Intelligence, April 8, http://tinyurl.com/nyenub54

Appel, G., J. Neelbauer, and D. A. Schweidel, 2023, "Generative AI has an intellectual property problem," Harvard Business Review, April 7, http://tinyurl.com/234z7jn2

Aristi Baquero, J., R. Burkhardt, A. Govindarajan, and T. Wallace, 2020, "Derisking AI by design: how to build risk management into AI development," McKinsey & Co., August 13, http://tinyurl.com/5488ahyc

Awati, R., 2023, "Garbage in, garbage out (GIGO)," TechTarget, June, http://tinyurl.com/ms6pyay7

Birkstedt, T., M. Minkkinen, T. Anushree, and M. Mäntymäki, 2023, "AI governance: themes, knowledge gaps and future agendas," Internet Research 33:7, 133-167

Blackman, R., 2022, "Why you need an AI ethics committee," Harvard Business Review, July-August, http://tinyurl.com/yanv59cu

Bommasani, R., K. Klyman, S. Longpre, S. Kapoor, N. Maslej, B. Xiong, D. Zhang, and P. Liang, 2023, "The Foundation Model Transparency Index," Standford University

Brous, P., M. Janssen, and R. Krans, 2020, Data governance as success factor for data science. Responsible design, implementation and use of information and communication technology, Springer

DAMA International, 2017, DAMA – DMBOK. Data managment body of knowledge, Technics Publications

Deighton, A., 2023, "Three reasons why your organization needs a data product strategy," Forbes, February 10, http://tinyurl.com/yx6f9awu

Desai, V., T. Fountaine, and K. Rowshankish, 2022, "A better way to put your data to work," Harvard Business Review, July-August, http://tinyurl.com/33b7h8ap

Drenik, G., 2023, "Data quality for good AI outcomes," Forbes, August 15, http://tinyurl.com/v77nxrnx

E.C., 2020, "Assessment list for trustworthy artificial intelligence (ALTAI) for self-assessment," European Commission, July 17, http://tinyurl.com/3szc5aen

European Council, 2023, "Artificial Intelligence Act: Council and Parliament strike a deal on the first rules for AI in the world," December 9, http://tinyurl.com/2s36kakp

European Parliament, 2023, "EU AI Act: first regulation on artificial intelligence," December 19, http://tinyurl.com/bdkkhb4r

E.U., 2018, The General Data Protection Regulation applies in all Member States from 25 May 2018," European Union, May 24, http://tinyurl.com/muybvc3p

EY, 2023, "The artificial intelligence (AI) global regulatory landscape," Ernst & Young, http://tinyurl.com/y377fxxf

Famularo, A., 2019, "The evolution of data governance," Forbes, March 11, http://tinyurl.com/5xxcbxuy

Fowler, M., 2019, "How to move beyond a monolithic data lake to a distributed data mesh," martinfowler.com, May 20, http://tinyurl.com/sft7e8sh

Greenacre, M., 2023, "EU/US divergence in data protection holds lessons for global regulation of artificial intelligence, experts say," Science Business, September 28, http://tinyurl.com/mrxtxyep

Hasselbalch, G., and P. Tranberg, 2016, Data ethics – the new competitive advantage, Publishare

Heaven, W. D., 2023, "The inside story of how ChatGPT was built from the people who made it," Technology Review, March 3, http://tinyurl.com/2vh7ma9r

Hinkle, O., 2020, "The evolution of data governance," Dataversity, May 18. http://tinyurl.com/5eahdd8c

HKMA, 2019, "Reshaping banking with artificial intelligence," Hong Kong Monetary Authority and PwC, http://tinyurl.com/4fpn2kxb

IAPP, 2023, "Global AI legislation tracker," International Association of Privacy Professionals, http://tinyurl.com/mwum42j7

IBM, n.d. "What is an AI model?" http://tinyurl.com/4us7xcf8

Jorge Ricart, R., and P. Alvarez-Aragones, 2023, "The geopolitics of generative AI: international implications and the role of the European Union," Elcano Royal Institute, November 27, http://tinyurl.com/2s42wb8c

Ledger, M.-A., and A. McGill, 2023, "Digital regulatory reporting. The evolution of global initiatives," Bank of England, http://tinyurl.com/mr238drk

Loftus, T., 2018, "The morning download: Facebook at center of global reckoning on data governance," Wall Street Journal, March 19, http://tinyurl.com/2mz2szwy

MAS, 2023, "Veritas initiative," Monetary Authority of Singapore, October 26, http://tinyurl.com/2t5b2hsp

McConnell, J., 2020, "Leadership everywhere means reversed leadership," Global Peter Drucker Forum, September 30, http://tinyurl.com/yckkv9mw

McKinsey, 2020, "Designing data governance that delivers value," McKinsey & Co., June 26, http://tinyurl.com/4rnj9fw2

McKinsey, 2022, "How to unlock the full value of data? Manage it like a product," McKinsey & Co., June 14, http://tinyurl.com/5n6wxjtx

McKinsey, 2023a, "Demystifying data mesh," McKinsey & Co., June 28, http://tinyurl.com/bdfxrmja

McKinsey, 2023b, "The economic potential of generative AI: the next productivity frontier," McKinsey & Co., June 14, http://tinyurl.com/3ypwe2y5

Mollick, E., 2022, "ChatGPT is a tipping point for AI," Harvard Business Review, December 14, http://tinyurl.com/mskfh6pt

ODI, 2021, "Data ethics canvas," June 28, Open Data Institute, http://tinyurl.com/2s3ww7a7

OECD, 2023, "G7 Hiroshima Process on generative artificial intelligence (AI)," September 7, http://tinyurl.com/yck6jjcz

OpenAI., 2023, "GPT-4 technical report," http://tinyurl.com/4e7ez627

Power, D. J., 2018, "What are Drucker's views on planning and decision making?" DSSResources, April 23, http://tinyurl.com/2hvay9yw

PwC, 2021, "Risk and regulatory outlook 2021 – key developments in Southeast Asia – digitalising regulatory reporting," http://tinyurl.com/k27ccpsk

Rosencrance, L., 2024, "Five strategies for managing unstructured data," Techopedia, January 16, http://tinyurl.com/3kwx9hmv

Schneider, J., R. Abraham, C. Meske, and J. vom Brocke, 2023, "Artificial intelligence governance for businesses," Information Systems Management 40:3, 229-249

Talagala, N., 2022, "Data as the new oil is not enough: four principles for avoiding data fires," Forbes, May 2, http://tinyurl.com/4zjzza7d

Violino, B., 2023, "Eight tips for unleashing the power of unstructured data," CIO Magazine, November 28, http://tinyurl.com/bde7hb42

Watts, M., 2021, "Why data is the new oil," Futurescot, November 17, http://tinyurl.com/39k5tfs4

Woods, D., 2016, "Why data lakes are evil," Forbes, August 26, http://tinyurl.com/4az9h2hp

ABOUT CAPCO

Capco, a Wipro company, is a global technology and management consultancy focused in the financial services industry. Capco operates at the intersection of business and technology by combining innovative thinking with unrivalled industry knowledge to fast-track digital initiatives for banking and payments, capital markets, wealth and asset management, insurance, and the energy sector. Capco's cutting-edge ingenuity is brought to life through its award-winning Be Yourself At Work culture and diverse talent.

To learn more, visit www.capco.com or follow us on Facebook, YouTube, LinkedIn and Instagram.

WORLDWIDE OFFICES

APAC
Bengaluru – Electronic City
Bengaluru – Sarjapur Road
Bangkok
Chennai
Gurugram
Hong Kong
Hyderabad
Kuala Lumpur
Mumbai
Pune
Singapore

MIDDLE EAST
Dubai

EUROPE
Berlin
Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
Glasgow
London
Milan
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA
Charlotte
Chicago
Dallas
Hartford
Houston
New York
Orlando
Toronto

SOUTH AMERICA
São Paulo

WWW.CAPCO.COM

CAPCO
a wipro company