

Next Data Paradigm: Using AI to Manage All Human Data

Roman Lukyanenko

McIntire School of Commerce

University of Virginia

romanl@virginia.edu

Abstract

In the age of smart IT, data management - the very foundation of information technology - remains laborious, inefficient, largely inaccessible, falling far short of its potential. The means of taking a major leap forward in data management is here. The rapid evolution of artificial intelligence presents a paradigm-shifting opportunity in digital storage and data management. This paper suggests how Agentic AI systems can revolutionize the ways organizations and people store, organize, and retrieve data. We propose AI to manage all data storage and retrieval needs of humans. By leveraging advanced machine learning, and autonomous decision-making capabilities, AI-driven data management promises to transform data management from an inefficient time-consuming process to an intelligent personalized service accessible to everyone.

Keywords

Data management, artificial intelligence, Agentic AI, Generative AI, AI ethics, large language models, large reasoning models, large action models, data modeling, data acquisition, data governance, data infrastructuring, data analytics, inclusive data management, green data management

1. Digital Neolith, Late Data Stone Age

From driverless cars to oh so realistic AI-generated deepfakes, to hyper addicting videos on YouTube, TikTok and Netflix, to helpful tools like MS Word and ChatGPT, the progress of information technology is dazzling. The pace of change is mesmerizing: “a year in machine learning [a foundation of modern AI] is a century in any other field” [151].

Marveling at the digital world, it is hard to believe, ours is probably the Data Stone Age. You read it right: Stone Age. How can this not be the pinnacle of human development, you may object? And yet, with all the advances in AI and other tech, we are likely only at the dawn of the IT revolution. The reason is, data management, the very foundation of information technology remains laborious, inefficient, largely inaccessible, and falling far short of its potential.

Data management is everywhere. Organizations, employees, and people going about their daily lives constantly manage digital data. And nearly every time, this is a struggle. Saving a picture on a smartphone, registering for a concert, making a social media post, backing up work files to the cloud, is not hard to do. Remembering the details in an email received a week ago, figuring out the most recent version of a shared document, finding all the relevant

customer data, securing private records from mischievous actors, is not so easy. As if by some cruel irony we create digital data with ease, only to struggle managing it afterwards.

While computing devices, cloud, AI, platforms, and apps have made remarkable progress, most data management tasks rely heavily on human intervention, and are grossly inefficient. Data management is commonly synonymous with effort and often, with frustration. In some projects, “data discovery, data ingestion, data cleansing, and data-pipe engineering” can take “months” [64:8]. It has been estimated that 80% of the time in analytics projects is spent on data management [60].

Poor data management causes massive financial, reputational losses. Even leading information technology companies struggle with data management, as evidenced by constant high-profile failures. Consider, for example, the IBM Watson - MD Anderson Cancer Center debacle that instead of curing cancer resulted in multimillion dollar losses [144, 161]. Due to frequent debacles, valuable opportunities routinely go untapped.

In the US alone, the estimated cost of poor data quality exceeds \$3 trillion [53, 141]. Organizations are said to be spending up to 30% of their revenue handling data quality issues [53]. These shocking figures are a vivid testimony of how important data has become, and how non-trivial it is to manage it.

The good news is we are in a Digital Neolith, the Late Data Stone Age. This means change is near. Artificial intelligence, leveraging such advances as Generative AI [8, 162] and Agentic AI [14, 54] is a natural response to the onslaught of data. It is a recipe to evolve into a different paradigm of data management, to usher in a new age of data.

This paper suggests a key feature of the Next Data Age: data management powered by AI. The idea is simple: for all data needs (creation, collection, storage, retrieval of data), a personal AI executive is there to handle it. An AI Data Executive (AIDE) dynamically adjusts to the preferences, skills and capabilities of the user, and handles all the files, documents, and information the user provides. It abstracts the handling and storage process, which is highly optimized to ensure intelligent retrieval and environmental efficiency. When the user wants to see or exchange the data, AIDE retrieves and provides it. This simple vision, however, harbors a lot of complexity, involving a host of technical, social, security and ethical challenges that need addressing. The paper outlines the general principles of AIDE and suggests open questions that need to be answered to make AI data management a staple of the Next Data Age.

2. Modern Data Management

2.1. Pains and Ills of Modern Data Management

Data management is laborious, inefficient, and falling far short of its potential. The problem becomes worse with time, as data volumes grow and reliance on digital data increases. Here are some observations on the state of data management today (summarized in Figure 1).

We begin with an unusual data management challenge, which unfortunately does not see the attention it deserves. We will start with the issue of inclusion.

Inclusion in data management is ensuring that anyone interested in collecting, storing and using digital data is able to do so with ease and benefit. Modern data management is not inclusive. Period. Consider just a few examples of data management inaccessibility. Most data management (e.g., setting up data collection interfaces, creating data models, scripting database schema, optimizing queries) relies heavily on human visual modality. This places a significant barrier for people with visual impairments. Furthermore, data management, especially on a scale, requires significant technical skills, such as knowledge of specialized software (e.g., Hadoop, NoSQL, data lakes), and languages (e.g., data modeling, python,

SQL). Some groups have been consistently marginalized and continue to be ill-represented in data management (e.g., women). Specific activities of data management [49, 53, 100], such as modeling, data collection and acquisition, data organization and curation have been called out for being poorly accessible to diverse users [12, 42, 56, 69, 72, 82–84, 102, 108, 109, 112, 118, 122, 145, 150, 152, 187].

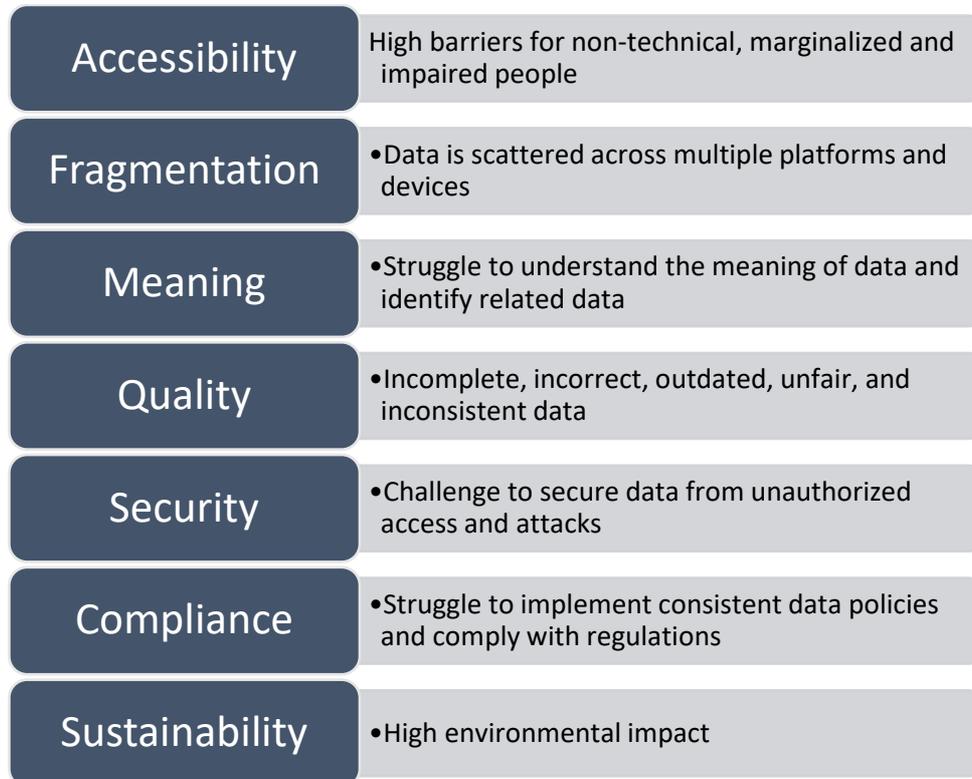


Figure 1. Data management challenges in the Late Data Stone Age

Inclusive data management is data management that accommodates different levels of skill and access across all data management activities. This includes all aspects of data management, such as modeling, acquisition, governance, infrastructure and consumption support (so called, MAGIC of data management). The present focus in accessibility is on issues of interface design (data acquisition aspect of MAGIC). In contrast, issues related to modeling, governance, infrastructure and consumption support are generally ignored. By lowering barriers to access, inclusive data management should provide more equitable and effective use of digital information, so everyone can benefit from the wealth of opportunities of the digital world.

Data fragmentation brings a host of challenges for people and organizations. Files and information records are scattered across multiple platforms and devices - personal computers, cloud drives, mobile phones, and organizational storage systems. Furthermore, data is not only standalone documents, pictures or video files. It is also information stored inside digital systems, such as platforms, apps, software packages scattered across the digital universe. Hence, customer data in an organization might be spread across a customer relationship management system, an email marketing platform, sales spreadsheets, and customer service logs.

This fragmentation increases risks and breeds inefficiency. Financial reports might rely on data from accounting software, supply chain management platforms, and sales forecasts - each stored in different formats and updated on separate schedules. A single discrepancy between these sources could skew the entire financial outlook but reconciling them requires

manual effort and cross-department coordination. The result is that an organization may miss critical context, such as a customer's recent complaint or a pending contract renewal, leading to poor service and lost business opportunities (or worse).

On a personal level, individuals face similar issues when trying to manage their private data. A person planning a trip might have flight details in their email, hotel reservations saved as PDFs in a cloud drive, and notes on sightseeing spots in a mobile app. If the person wants to create a single itinerary or share the plan with a friend, it would be an ordeal to manually gather this data from different platforms and formats. Worse, if any data is outdated or saved in a format that's no longer compatible with their current device or software, they may have to manually convert it, if this is at all possible.

At work or at home, people need to remember where they store specific files and manually update and transfer documents, pictures, videos between systems. Synchronization between platforms is often inconsistent, forcing people to track different versions of files and reconcile conflicts manually. Data retrieval depends on precise naming or folder structures, and search functions are often limited by incomplete metadata or poor indexing, making it difficult to find relevant information quickly.

Lack of contextual awareness of data by data managing services is a major cause of effort and failure. Data systems struggle to understand the deep meaning of data and identify semantic links. Current data management tools can store and retrieve files but struggle to recognize relationships between different data sets or formats. For instance, a presentation, a spreadsheet with financial data, and an email discussing the same project might be stored in completely different locations without any logical connection established between them. Users must manually consolidate and cross-reference these files, which increases the risk of overlooking critical information or making decisions based on incomplete data. Furthermore, data categorization and tagging remain largely dependent on user input, which is inconsistent and prone to error [6, 74, 111, 154].

Poor data quality - resulting from incomplete, outdated, or inconsistent data - can undermine decision making and actions taken with data. Issues such as duplicate records, incorrect customer information, and misclassified data increase risks and degrade operational efficiency [36, 93, 140]. Even leading information technology companies struggle with data quality issues. Thus, after multimillion dollar investments and much hype, MD Anderson Cancer Center scraped collaboration with an AI system Watson provided by IBM [144]. This is despite IBM's ambitious goal to use AI to cure cancer. The demise of the MD Anderson-Watson project was significantly influenced by issues of data quality. The project struggled with inconsistencies and errors in the data, which hindered the effectiveness of the Watson system in providing accurate and safe recommendations for cancer treatment [161].

Data quality is not only a technical issue of ensuring data is accurate, complete or timely – the common data quality considerations [16, 36, 99, 110, 111, 124, 164, 176, 178, 184]. Data is a social artifact [1, 2, 7, 132] and quality also involves fairness - whether data has been collected ethically, legally, and in accordance with relevant cultural norms. Present understanding of data fairness is scattered across specific areas of focus, such as informed consent issues in scientific research and consumer surveys [33, 50, 92], or legal evidence admissibility [67, 86, 94, 123]. Fairness is not well-incorporated in general data management frameworks, and hence, can be easily overlooked.

Data security further complicates data management. Securing data remains exceedingly difficult as data volumes and variety grow. Attacks and data breaches are on the rise, resulting in significant damages. The global cost of cybercrime is projected at \$10.5 trillion [120], growing at a rate of 15 percent annually [25]. For example, a data breach experienced by Equifax in 2017 was caused by inadequate data security measures that allowed hackers to access sensitive information of approximately 147 million people. Cyberattacks using stolen credentials increased 71% year-over-year [25].

harder the challenge to organize, curate, and retrieve data afterwards. Yahtzee! We are in a race against complexity.

Humans are naturally inclined to document, share, and store information as a means of communication, problem-solving, and self-expression [52, 115, 116]. The proliferation of smartphones, social media platforms, and cloud-based applications has made data creation almost frictionless, accelerating the growth of data volumes at an unprecedented rate. The more accessible and integrated these platforms become, the more individuals and organizations contribute to the data ecosystem, reinforcing a feedback loop of continuous data generation.

The more digital data becomes embedded in different facets of human existence, the more varied and specialized the types of data are produced. For instance, business data now includes not only structured information such as financial records but also unstructured data like customer feedback, social media interactions, and multimedia content [3, 37, 38, 41, 46, 85]. Similarly, personal data encompasses not just text-based information but also health metrics from wearable devices, location data from GPS tracking, and images from social platforms. The heterogeneity of this data creates significant challenges for integration, as different formats and sources require distinct handling and processing methods. Without intelligent systems capable of handling such diversity, the task of organizing and retrieving data becomes increasingly effortful and inefficient.

The challenge is further compounded by the fact that data creation often outpaces data management infrastructure. While storage capacity and data processing capabilities have improved, they have not kept pace with the exponential increase in data volume and variety. As a result, organizations and individuals frequently encounter data silos, inconsistent file structures, and incompatible formats. The rise of cloud computing and remote work has exacerbated this issue, as data is now distributed across multiple platforms and geographic locations. This decentralization creates gaps in data visibility and control, making it harder to establish a single source of truth. The more complex the data environment becomes, the greater the cognitive and operational burden on users to navigate and reconcile disparate data sources.

The cognitive load involved in data organization and retrieval increases with the diversity of data sources and formats. Human working memory is limited in its ability to process and categorize information [9, 119]. When faced with excessive or poorly organized data, people experience increased mental strain, leading to slower decision-making and higher error rates [80, 91, 131]. In organizational settings, this manifests as reduced productivity, miscommunication, and stress, especially for older employees [160, 166, 167] or employees with disabilities [165]. Even at the individual level, the effort required to consolidate personal data across devices and platforms - such as combining health data from a smartwatch with dietary records from a mobile app - creates barriers to effective data use and decision-making.

To sum up, as long as it remains easy and rewarding to create data, the volume and diversity of data will continue to grow, reinforcing the complexity of the data ecosystem. Effective data management strategies must therefore account for both the human tendency to generate data and the human limitations in handling this complexity. Addressing modern data management challenges requires more than just expanding storage or improving processing speeds; it necessitates the development of intelligent data management systems capable of contextual understanding and adaptive organization.

3. AI, the Ultimate Data Manager

3.1. Foundations

This paper advances a critical proposition for the Next Data Age: AI-driven data management as a fundamental data paradigm. At the core of this vision is the concept of an AI Data Executive (AIDE) - a personal, adaptive system designed to autonomously manage all aspects of data interaction, including creation, collection, storage, and retrieval. By dynamically adjusting to the user's preferences, skills, and capabilities, AIDE assumes full responsibility for organizing and handling files, documents, and information, thereby abstracting the complexities of data management into a seamless, optimized process. This intelligent system not only enhances efficiency in data retrieval but also prioritizes environmental sustainability through resource-aware storage mechanisms.

Despite the allure, the implementation of AIDE entails significant technical, social, security, and ethical challenges that must be rigorously addressed. Before articulating the architectural principles of AIDE, we make explicit its underlying core values. This is inspired by agile development and Agile Manifesto [17]. There may be multiple ways to implement AIDE, and we do not suggest for ours to be definitive or the best.¹ By laying out the foundational values of the system, however, we inspire future developments which can find better ways to realize these values into specific technological implementations. We propose Artificial intelligence to manage all data storage and retrieval needs of humans. This is needed to handle the excessive and rising complexity of data management. The resulting AI Data Executive (AIDE) is a specific system designed to achieve these goals. AIDE or any such system must adhere to the following foundational values:

Value 1. Human first, human is the end, not means.

The value “Human first, human as the end, not means” reflects a philosophical and ethical stance that prioritizes human well-being, dignity, and autonomy in any system, technology, or societal structure. It suggests that:

1. Human is always the priority. Human needs, rights, and values should be the primary consideration in any AI implementation. Rather than prioritizing efficiency, profit, or automation for its own sake, AI systems such as AIDE should be designed to serve and empower people.

2. “Human as the end, not means. Human users should never be treated merely as instruments or resources to achieve external goals (e.g., corporate profit, technological progress, or political power). Instead, all systems and advancements should ultimately serve human flourishing, rather than reducing people to tools for some other objective.

This value originates in Kantian ethics [180], which argues that human beings should always be treated as ends in themselves, never merely as means to another’s end.

Value 2. Safety, reliability, sustainability, transparency over raw performance.

The value “Safety, reliability, transparency over raw performance” insists that in the design and deployment of AI, certain foundational values should take precedence over speed, efficiency, or computational power.

¹ In general, there are multiple ways to implement any idea into a technological solution, the problem known as equifinality (a given end state or outcome can be reached through many different paths) [21, 65, 70, 106, 158]. The goal is to find the one that best fits the contextual requirements while maximizing adherence to the values of AIDE.

1. **Safety First.** Systems must be designed to prevent harm, both physical and digital. AI and automated systems should not pose risks to individuals, society, or the environment. Safety includes robustness against failures, adversarial attacks, and unintended consequences.
2. **Reliability.** A system should function consistently and predictably under various conditions. It should not behave erratically, make critical errors, or require excessive human intervention to correct unintended outcomes.
3. **Sustainability.** An implementation should prioritize the impact on the broader environment, including nonhuman animals and natural resources of the planet.
4. **Transparency.** The inner workings, decision-making processes, and limitations of a system should be understandable to users and stakeholders. Black-box AI models or opaque decision-making mechanisms can lead to mistrust, bias, and unintended harm. Transparency ensures accountability and fosters responsible use. Naturally, we recognize the impossibility of full transparency, especially when dealing with highly performant AI models (e.g., deep learning neural networks) [4, 4, 161].

This value is consistent with the foundations of design as outlined in Larsen et al. [90]’s design science validity. Any robust design according to design science validity should not only demonstrate superior performance (criterion validity) but the designers should also confidently know how the system works (causal validity), and where it is appropriate to use it (context validity).

Value 3. Data owners should be in full control over their data.

The value “Data owners should be in full control over their data” promotes data sovereignty, which asserts that individuals or entities that own data have the right to manage, access, and determine how their data is used, shared, and distributed. This value advocates for personal autonomy over digital data, ensuring that owners retain full authority over decisions regarding their data, rather than allowing third parties or external organizations to dictate its use without consent. This includes the ability to grant or revoke access, modify data, or delete it, all while maintaining transparency about how data is being handled.

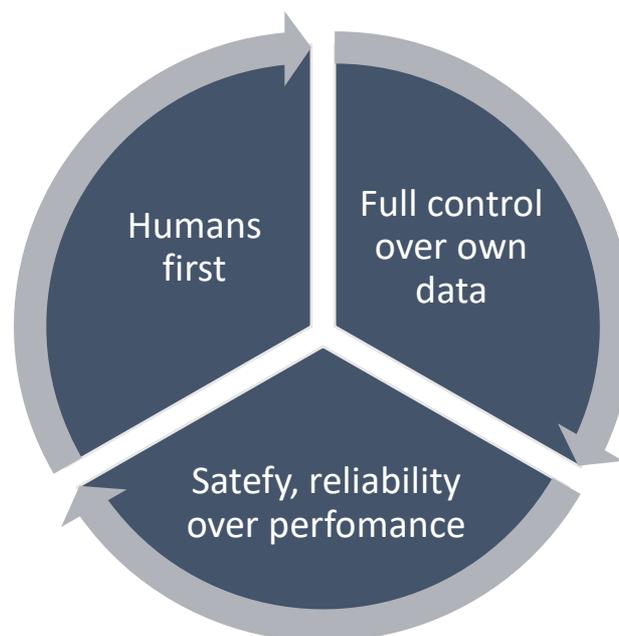


Figure 3. Foundational values of AIDE

The three foundational values of AIDE (Figure 3) promote a human-centric approach to technology and data management. Together, these principles advocate for ethical, responsible, and human-centered technological progress. We now show how these values are implemented.

3.2. Architectural Principles

The general architecture of an AI Data Executive (AIDE) is provided in Figure 4. It depicts a basic set up which can be specialized for particular needs, industries or in compliance with specific regulatory requirements. The aim of any architecture is to realize the three foundational values behind AIDE as best as possible. For general advice on how to implement abstract ideas into specific technological implementations (or to establish instantiation validity), see [10, 43, 44, 90, 104–107].

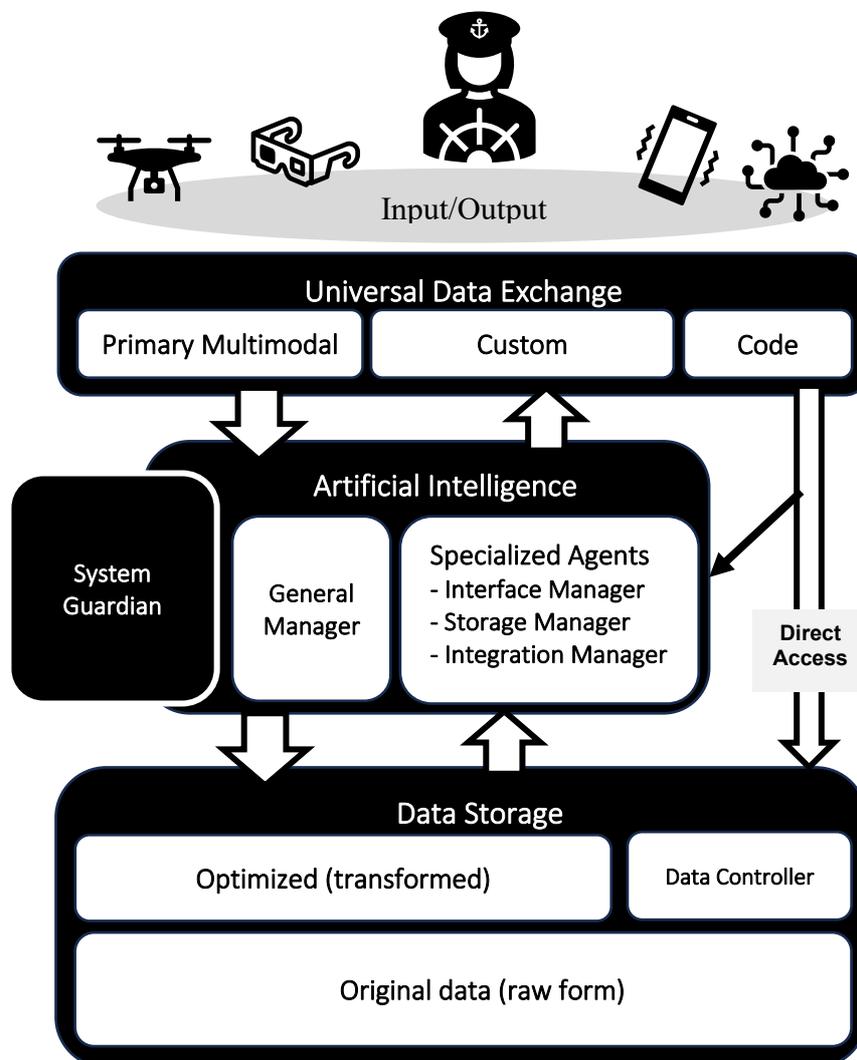


Figure 4. Architecture of an AI Data Manager

3.2.1. System Owner

A human figure at the top is the **System Owner**, a person, a team of people, or their designated proxies, who own the data. Recall that foundational value is to ensure the System Owner is always in full control of their data, is always fully informed of any decision made and that no decision can be made without informed consent of the System Owner. The rest of the architecture implements these ideals.

The AI Data Executive is a multi-layered, modular system. This ensures its ability to evolve and be flexible to the changing user needs, technological developments and regulatory requirements [51, 87, 109, 125, 153, 188]. Data flows with the support of AI from its origin – the data creator, be it the human or human-controlled process - to storage. Data also flows back to data consumer – the owner or the devices authorized by the System Owner to receive the defined information.

A System Owner is free to utilize a diverse range of devices, such as smartphones, laptops, drones, 3D printers, and tactile interfaces, to both collect and display information. All this is managed and stored centrally, facilitated by artificial intelligence.

While the System Owner has the flexibility to choose which device to use, the devices place physical limitations on the kind of information that can be collected and displayed. On the other hand, the properties of the devices make certain type of data collection and presentation possible. For example, a smartphone can gather location data, a drone can capture aerial imagery, and a 3D printer can produce physical models based on processed data.

The devices themselves may or may not be part of AIDE. Different options are possible. Hence, it is possible to develop some devices that would provide additional capabilities to take advantage of AIDE. For instance, a drone could be specifically designed to supplement AIDE, allowing it to send real-time landscape imagery directly into the system for analysis and transmission to other devices (e.g., 3D printer). This way the System Owner can orchestrate a multistage information exchange to automate some complex tasks (e.g., building a 3D image of the landscape for construction purposes). At the same time, general computational devices, such as laptops and smart phones, could simply install AIDE as a software package or an app, thereby permitting these devices to be integrated into the AIDE's ecosystem.

In a device that runs AIDE the collected data is processed and standardized through Universal Data Exchange. This way data remains accessible and interoperable across multiple devices. This enables AIDE to manage data consistently, even if the devices vary in capability and design.

3.2.2. Universal Data Exchange

All collected information is transmitted through a **Universal Data Exchange (UDE)**, which standardizes and translates the data into a unified storage format. The UDE acts as a central hub, facilitating data integration and enhancing the overall efficiency and adaptability of the system. The output capabilities of UDE transform data into actionable formats suited to the receiving device's functionality. For example, data collected from a smartphone on air quality can be visualized on a laptop dashboard or translated into haptic feedback on a tactile interface to alert visually impaired users about hazardous conditions.

The **Universal Data Exchange** has three components: Primary Multimodal, Custom and Code. In using various devices, System Owner can leverage the specific functionalities of these three components, provided the devices permit the integration with these components. Not all devices will be able to support each component. For example, some devices may not allow the user to directly write code or manipulate custom graphic user interface objects. In this context, the role of Primary Multimodal, which is the AI-powered component, is to automatically discover the physical capabilities of each device and adjust data collection and

presentation options accordingly. In contrast, Code and Custom components are not AI-based and may require System Owner to manually negotiate the capabilities of a particular device in order to enable these components.

Primary Multimodal is the main interface for collecting and presenting data. It is AI driven and seeks to collect and provide data in any possible mode. The aim is to support the known human sensory systems [159]:

- **Vision** (Visual System): Graphic use interface, the standard and most used type of interface today. Principles of user interaction with graphical interfaces are established [155, 168]. Advances in AI, including natural language processing, and computer vision, allow for greater adaptivity and customization of the graphic interface [156].
- **Hearing** (Auditory System): Voice and sound interface. More recent, but now becoming widely available due to progress in speech recognition, text-to-speech, natural language processing [156].
- **Taste** (Gustatory System): Gustatory interface is presently a futuristic idea of providing real experience of taste. Taste simulations are already being developed [126, 174].
- **Smell** (Olfactory System): Olfactory interface permits detecting odors and producing odors. Such interface is already being considered [126].
- **Touch** (Tactile System): Allows users to feel sensations like pressure, temperature, and pain through receptors in the skin. These interfaces are emerging, including 3D surface generation, Braille systems, and realistic video games [18].

AIDE's Primary Multimodal interface allows users to interact with the AI system through different modalities depending on the context, user preferences and abilities. The interface dynamically adjusts to the user's input style and the task at hand. For example, a user could start interacting with the system through a graphical interface, such as a dashboard displaying data visualizations. If the user needs deeper insights, they could switch to natural language by asking follow-up questions in plain language, like "Why did sales drop last quarter?" This transition between graphical and natural language inputs allows the interface to adapt to varying levels of technical expertise and preferred ways to handle data.

An optional component of AIDE is **Brain-Computer Interface** (BCI), a system that establishes a direct communication pathway between the brain and an external device, such as a computer or prosthetic limb. BCIs work by detecting and interpreting neural signals, typically electrical activity from the brain, and translating them into commands that control an external device or software. The goal of BCIs is to enable individuals to interact with machines or environments without relying on traditional muscle-based pathways. It is understandable that BCI is a controversial technology and is not for everyone. However, BCI plays a critical role for supporting people with some disabilities and those undergoing rehabilitation that can be facilitated with BCI. Hence AIDE supports BCI.

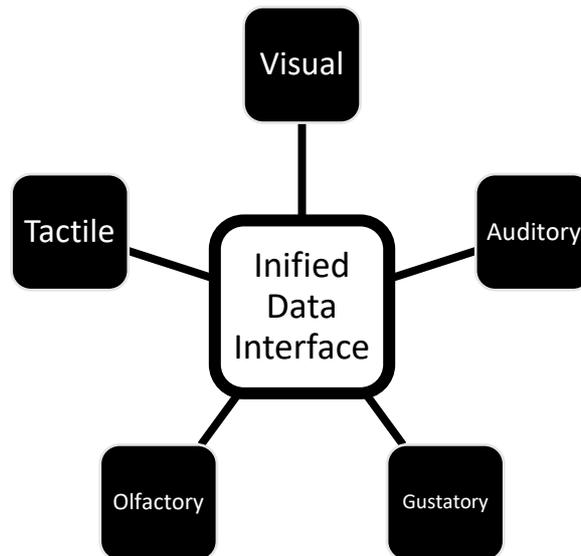


Figure 5. Integrated Sensory Interface of UDE

The optimal implementation of AIDE is to ensure that all sensory systems are integrated (see Figure 5). This way, the interface, driven by AI, can adapt based on environmental and situational factors. In a manufacturing setting, for instance, a technician might use haptic feedback to adjust machinery settings directly through a tactile interface, while also receiving real-time diagnostic reports through a graphical or natural language display. Only an integrated interface would permit people with different skills, abilities, and needs to manage data effectively.

An important component of AIDE is **Custom Interface**. Unlike Primary Multimodal, Custom Interface is not driven by AI's adaptivity, but is rather a blank slate, allowing the user to design their own interfaces by customizing and fixing the modalities and presentations upon their choice. This interface lacks the adaptivity of the Primary Multimodal interface but ensures the ability of the System Owner to have full control over all interface components. AI would still be responsible for processing the data generated by Custom Interface. The output of data is controlled by the user and can be presented either through Custom Interface or through Primary Multimodal, depending on the type of data and the rules set up by the System Owner.

Both Primary and Custom interfaces are able to execute programming code, such as Python or SQL. However, Primary Multimodal interface is driven by AI, and hence, code interpretation and execution are AI-mediated. To allow for direct control over data using code, **Code Interface** is provided. The interface seeks to support common programming languages and frameworks, including ability to download and install new languages and frameworks. The aim of the interface is to collect and retrieve data using code. Such systems as Apache Spark illustrate the potential to use different languages, including SQL, Java, Scala, Python and R for data manipulation. They directly inspire AIDE.

Code Interface allows "power users" familiar with programming to write code in their preferred language without interference from AI that might attempt to "optimize" or modify the execution path. For example, a user might write a data manipulation job in Python to process a large dataset and then switch to SQL to run some queries on the processed data. Direct code execution ensures that the execution engine of the Code Interface processes the code exactly as written. The Code Interface always becomes available for any device that permits entering text (or voice) or supports low/no code graphics.

The aim of Universal Data Exchange is to be responsive and intuitive, regardless of the complexity of the task or the user's preferred mode of interaction. The flexibility across modalities ensures that the system is not only user-friendly and inclusive but also capable of supporting sophisticated, context-specific workflows.

3.2.3. Artificial Intelligence and Data Flow

The AI layer of AIDE is made of several AI components. An AI functions as a tool to augment human data management rather than replace it. Any implementation of the system remains aligned with human values, goals and regulatory regimes. The aim is to build the AI layer based on the sound, reliable state-of-the-art engineering principles [29].

The Artificial Intelligence layer's job is collecting, storing, processing, organizing data, and deriving insights from data to support owner's decision making and actions. Effectively, it connects System Owner, via Universal Data Exchange with the Storage Layer. The AI layer interacts with the modalities of the Universal Data Exchange. These modalities in turn are able to pass to AI any external sources provided by the owner (e.g., video file).

Among the core AI components, a **General Manager** functions as a central orchestrator, overseeing and coordinating the activities of multiple specialized AI agents to achieve complex, high-level objectives. Acting as a meta-controller, the manager assigns tasks, monitors progress and dynamically adjusts the roles and priorities of specialized agents based on real-time feedback and changing conditions.

The functions of the General Manager leverage such AI technologies as hierarchical reinforcement learning (enables autonomous decomposition of long-horizon tasks into simpler subtasks) [134] and meta-learning (facilitating the learning process itself) [181].

Specialized AI agents are designed to handle specific data management functions - such as data analysis, language processing, decision support, or automation - while the General Manager agent integrates their outputs, resolves conflicts, and ensures alignment with overarching goals and AIDE's values. This architecture enables adaptive, context-aware decision-making, where the general AI agent synthesizes insights from diverse agents, identifies gaps or redundancies, and orchestrates coordinated responses. Such a system is informed by research in multi-agent systems and autonomous coordination [40, 76, 183], where central intelligence manages distributed, task and domain-specific expertise.

Consider the tasks for one specialized agent. The *Specialized Storage Agent* is tasked with continuously updating the Optimized (transformed) storage. This agent functions as an intermediary between the raw data in the original storage and the AI processing layer, ensuring that the data remains clean, structured, and up to date. The agent can employ data cleansing techniques such as outlier detection, imputation of missing values, and normalization of data formats to maintain consistency. Moreover, it could use real-time data streaming and batch processing strategies to handle large-scale data updates.

The specialized AI agent would also need to handle data augmentation and fusion from multiple sources. Related data often originates from heterogeneous sources sensor networks, user interactions, external databases, and more. The agent would be responsible for aligning and merging these disparate data streams, resolving conflicts, and filling gaps to create a unified, enriched dataset. Research on data fusion techniques (Dong & Srivastava, 2015) suggests strategies for resolving inconsistencies and redundancies to improve data reliability and downstream AI performance.

The Storage Agent should leverage advanced machine learning techniques to dynamically classify and organize files, moving beyond traditional hierarchical folder structures. By employing contextual understanding, AI can analyze not only file content but also metadata, creation context, and usage patterns to create meaningful organizational structures. For example, an AI system could recognize that a financial report and a sales forecast are related

through shared data points and automatically group them together under a broader “Business Strategy” category. This allows the system to maintain a more fluid and adaptive file organization system, ensuring that files are easier to find and logically connected based on real-world relationships rather than rigid folder paths.

Predictive categorization enables AI storage systems to anticipate user needs by learning from individual and organizational file management behaviors. For instance, if a user frequently accesses market analysis reports at the start of each quarter, the AI system can proactively suggest or organize those files for easy access. Semantic indexing further enhances this capability by understanding the meaning and relationships between files rather than relying solely on keywords. If a user searches for “customer engagement,” the AI can retrieve not only files explicitly titled with those terms but also related documents like user feedback reports, customer service logs, and marketing plans, based on semantic similarity. This deeper understanding allows for more intuitive and effective search and retrieval experiences.

Another specialized agent is Security. The AI agent would require robust governance and monitoring to ensure the integrity and security of the optimized storage. The agent would implement anomaly detection models to identify and isolate corrupted data, unauthorized access attempts, or system malfunctions. For example, a cybersecurity AI system could use the agent to detect unusual network traffic patterns and update firewall rules accordingly. The agent’s ability to maintain data integrity would be guided by research on secure machine learning (Papernot et al., 2018).

An agentic AI storage system would also incorporate autonomous management features to streamline operations and enhance system reliability. Intelligent backup and redundancy mechanisms can automatically create and manage backup copies based on the importance and frequency of use of a file. For example, mission-critical project files could be backed up more frequently and stored in geographically distributed locations to minimize the risk of data loss. Dynamic security and access management would allow the AI to adjust file permissions based on contextual understanding of user roles and behaviors. If an employee from the finance department attempts to access sensitive HR files, the AI could flag the attempt or restrict access automatically. Furthermore, resource optimization features would enable the AI to predictively manage storage allocation, migrating infrequently used files to lower-cost storage while keeping frequently accessed files in high-performance storage.

3.2.4. Two Guardrails – System Guardian and Direct Access

To promote the autonomy and control of the System Owner, AIDE implements a semi-independent System Guardian. It also allows to bypass AI entirely via Direct Access.

A **System Guardian** is a specialized entity designed specifically to protect the System Owner by monitoring and evaluating AIDE’s behavior. It is an independent system component whose aim is not data management, but ensuring the system operates within acceptable moral, legal and ethical frameworks that prioritize System Owner’s well-being.

This Guardian is a neurosymbolic AI system [22]. It relies on deterministic, hard-coded rules, with some AI-driven learning and adaptability. The deterministic principles provide a stable foundation for enforcing key ethical boundaries, security constraints, and operational rules. This ensures transparent and predictable behavior aimed at promoting System Owner’s interests.

Meanwhile, the AI component introduces adaptive intelligence, enabling the Guardian to identify emerging risks, interpret complex scenarios, and recommend adjustments. The AI component of the Guardian should be kept to a minimum - just enough to permit flexible and adaptive behavior, but not enough to introduce unwanted uncertainty or opaqueness of decision logic that comes with the heavy reliance on AI.

The System Guardian's primary role is to mediate between the system's internal processes and the System Owner's values, priorities, and goals. Drawing from research in value alignment and AI ethics [31, 147], the Guardian monitors potential conflicts or risks and enforces constraints on the potentially unsafe operations of the AI layer of AIDE. For instance, if AI's action could unintentionally compromise user privacy (e.g., by exchanging private information with an external system without explicit owner's consent), the Guardian intervenes and stops the action, while altering the System Owner. Its deterministic core defines fundamental limits - such as ensuring data integrity or preventing harmful outcomes - while the AI-driven component permits flexible responses in nuanced contexts. This hybrid design ensures the Guardian remains both trustworthy and responsive, ultimately protecting the System Owner's autonomy and well-being.

AIDE implements an always-available **Direct Access** to data, an ability to bypass AI by providing a channel from the Universal Data Interface to the underlying storage. This includes access to data in its original raw format. This assures transparency, control, and reliability.

As with Custom and Code interfaces, Direct Access ensures that users are not solely dependent on AI for managing data, which can introduce biases, distortions, or filtering based on the model's training data or internal heuristics. By accessing raw data directly, users can validate AI-generated insights, detect inconsistencies, and apply independent analytical methods when AI outcomes appear unreliable or unclear. This approach enhances the user's autonomy and decision-making capacity by enabling a fallback mechanism that is not mediated by AI.

The System Guardian along with the AI core have visibility of the data moving through Direct Access but is incapable of interfering with it (the noninterference is assured through immutable rules). The Guardian's access to the bypass actions of the System Owner is required for the Guardian to learn the patterns of user's direct actions. These learned patterns can later be enacted by the Guardian's in a semi-automatic manner, saving the owner's future effort in manual intervention. The AI agents also have visibility of Direct Access, but this is done for the purposes of learning to improve their data management.

Direct Access is automatically implemented when using Code and Custom interface and is an option in Primary Multimodal interface.

3.2.5. Data Storage Layer

The Storage Layer is the place where data that needs to persist beyond a live session is stored. It consists of three main components: original data, optimized (transformed) data and control data.

The original, unprocessed data is always kept in the **Original Container**. Storing data in its original, unprocessed format is essential for ensuring data integrity, reproducibility, and flexibility in future analysis. When data is preserved in its raw state, it remains free from potential distortions introduced by preprocessing, transformation, or AI-driven interpretation, allowing for independent validation and reprocessing as needed. Retaining raw data enables traceability, allowing the System Owner or an entrusted party to audit processing steps and verify outcomes against the source data. Moreover, in fields such as scientific research and financial modeling, maintaining the original data is critical for reproducibility and ensuring that findings can be independently verified [68].

While original data may not be altered, it should be compressed as long as no signal is lost (so-called lossless compression). The compression should not use any proprietary algorithms, to ensure that data can always be uncompressed using open-source, freely available software. It is also advisable to store the compression algorithm, in Data Controller Container (described below).

The **Optimized Data Container** keeps the output of preprocessing, cleaning, and transformation processes, designed to structure and refine raw data into a format that is both meaningful and efficient for AI processing. This layer ensures that data is free from noise, inconsistencies, and missing values. For example, in natural language processing (NLP), tokenization, lemmatization, and stop-word removal are essential preprocessing steps that enhance the AI model's ability to understand and generate human-like text. Similarly, in machine learning models, data normalization and encoding of categorical variables ensure that numerical inputs are scaled appropriately, preventing models from being influenced disproportionately by large or small values [55, 88, 148]. As data cleaning and preparation can take up to 80% of the time, it is essential to optimize data for retrieval. Without this layer, AI systems would struggle to handle raw, unstructured data, leading to increased latency, poor model performance, and reduced predictive accuracy.

A **Data Controller Container** is an advanced meta-data repository. It serves as a centralized storage place that manages and organizes critical metadata, versioning, and information about the stored raw and processed data. It functions as a control hub, maintaining an auditable record of data transformations, access logs, and algorithmic processes applied to the data. This includes the storage of design patterns and essential code used in data management, such as compression algorithms, encryption methods, and indexing strategies. This ensures that data operations are not only traceable but also reversible when necessary. The container also facilitates version control by managing historical snapshots of the data and the code used to process it, allowing for rollback in case of processing errors. By consolidating this information, the Data Controller enhances data governance and compliance with regulatory requirements, such as GDPR and HIPAA, ensuring that data handling meets established standards of transparency and accountability.

4. Potential Challenges and Mitigation Strategies

While the potential of AI in storage management is significant, several challenges must be addressed to ensure its successful usage. One of the foremost concerns is *privacy*. AI-driven storage systems are aimed at handling vast amounts of sensitive data, raising concerns about unauthorized access and data breaches. To mitigate these risks, AIDE should implement robust encryption techniques and require explicit user consent before processing or storing sensitive information. Additionally, differential privacy methods and access control mechanisms can enhance data security.

Another challenge is *algorithmic bias*, which can lead to unfair or suboptimal decision-making in AI-driven storage management. Bias may stem from training data that lacks representativeness [117, 127, 128] or algorithms that reinforce existing prejudices [35, 59, 61, 117, 130]. To counteract this, AI systems should be carefully pre-trained using diverse datasets that reflect varied user needs. Improved bias mitigation techniques with the focus on data [23, 27, 71, 97, 149, 182] and algorithms [5, 23, 27, 62, 117, 170] is needed. Notably, AI bias sometimes originates in the underlying reality which is discriminatory and biased [161]. Hence organizational and psychology research on causes, identification and mitigation of organizational and personal biases [11, 32, 63, 79, 80, 172] can also support the technical developments of AIDE.

System reliability is another critical factor, as AI-based storage solutions must function seamlessly to prevent data loss or workflow disruptions. Without adequate safeguards, unexpected AI errors or failures could lead to significant productivity setbacks. To address this, implementing fallback mechanisms such as redundant system backups and human oversight options can ensure continuity. Regular system monitoring, error detection, and fail-safe protocols are essential to enhance the resilience of AI-driven storage. In general, as

Bostrom [30] and others argue [26, 58, 66, 78, 143, 171], AI needs to be based on solid engineering and moral principles.

An important facet of reliability is transparency of the decision logic used by AI. We already noted transparency challenges which have been dramatically exacerbated with the introduction of deep learning neural networks and large language models. There are many open questions related to transparency, such as reverse engineering neural networks [4, 129] and explainability of large language models [39, 114, 121, 186]. A promising opportunity is model invariant approaches; these have become especially important as modern AI models grow in complexity [4, 13, 103, 139].

Semantic transparency, ensuring that AI agents fully understand the meaning of data they manage, remains a challenge despite impressive progress in AI. Present AI-powered search and organization tools are limited in their ability to understand the context and nuances of different types of data, leading to incomplete or irrelevant search results [89, 95, 175].

Hallucination, tendency of AI to make up nonsensical data remains a challenge [34, 77, 185]. Some argue, the issues of semantics can never be fully resolved, as human meaning is not a web of statistical probabilities [19, 48, 177]. As Chomsky and colleagues contend, “machine learning systems can learn both that the earth is flat and that the earth is round. They trade merely probabilities that change over time. For this reason, the predictions of machine learning systems will always be superficial and dubious” [48]. Not everyone shares this pessimistic view. Some argue AI does not need to think like humans to be useful [24]. More generally, research on AI semantics continues to advance progress [34, 77, 173, 185]. Such areas as *large reasoning models* hold special promise as they combine statistical probabilities with logical reasoning [136]. Large action models, the key technology of Agentic AI [28, 45],

Another key challenge is *data and system interoperability*. AI-driven storage must integrate seamlessly with various existing platforms, file formats, and enterprise software. Incompatibility issues can lead to inefficiencies, data silos, and operational disruptions. To mitigate this, AIDE should adopt standardized data exchange protocols and API-driven solutions to facilitate seamless integration across different digital ecosystems. Ensuring cross-platform compatibility and scalability is crucial for maximizing the benefits of AI storage management.

Another important issue is coordinating multiple AIDE systems, each belonging to their respective owners. Such a need would be prominent in organizational settings, where employees and departments may have their own AIDEs. This challenge can be supported by ongoing research on coordinating multi-agent systems [40, 76, 183]. Developing high-level coordinating systems specifically tasked with coordinating multiple AIDEs is another opportunity.

Under the assumption that AIDE would be widely used, *environmental sustainability* remains an important challenge. While the goal of AIDE is reducing environmental impact, which we expect to happen due to streamlining of data management activities, a rebound effect can be expected. Rebound or take-back effect occurs when efficiency gains lead to increased consumption, potentially offsetting the initial savings [169]. Consequently, conserving energy and reducing environmental impact should always be a moving target, ideally staving off any rebound effects. One of the long-term solutions that addresses green data management is *fundamentally green technology*. This is technology that even at massive scale does not significantly pollute. For example, humans use air for verbal communication. This does not appear to have any notable environmental impact.

Quantum computing, based on the most abundant substance in the universe, elementary particles, holds futuristic promise to provide fundamentally green energy [101]. However, for this to happen, the significant inefficiencies of current quantum solutions must be resolved [15, 96, 137, 146].

Device capability limitations represent a challenge for AIDE, as the accuracy, resolution, and scope of data collection and presentation are inherently tied to the physical and technical constraints of the devices. For example, drones with limited flight time due to battery capacity may not capture comprehensive data over large areas, and tactile interfaces may lack the precision needed for detailed feedback. Moreover, the processing power and sensor range of smartphones and laptops could restrict the complexity of real-time analysis. A major challenge and opportunity is providing interfaces for the presently underutilized human senses, such as olfactory, gustatory and tactile [18, 126, 174]. Much of reality is inaccessible to humans due to their sensory limitations (capable of communication in a narrow band of audiovisual spectra). Celebrated biologist, E. O. Wilson [179] calls humans “sensory idiots” as 99% of species have greatly superior sensory systems. The olfactory, gustatory and tactile interfaces may allow humans to rediscover the world they live in. Another challenge is integrating the sensory modalities into an integrated, universal interface.

Ethical considerations also play a crucial role in AI adoption for storage management. Issues such as data ownership, accountability, and transparency must be well implemented to foster trust among users. Establishing clear guidelines on AI responsibilities, ensuring users have control over their data, and maintaining open communication about how AI makes decisions can help mitigate ethical concerns. Regulatory compliance with data protection laws should also be a priority. Advancing research on general ethics – which is a juvenile field [138] - and ethics in AI [26, 58, 66, 78, 143, 171] is paramount to ensure AIDE truly holds humans first in all its decisions.

A further challenge is *cybersecurity threats*. AI systems that handle large volumes of sensitive data may become prime targets for cyberattacks, including ransomware, data poisoning, and adversarial AI threats. To counter these risks, organizations should invest in advanced threat detection mechanisms, implement multi-layered security architectures, and conduct regular security audits. It is paramount to stay far ahead of the cybersecurity threats and anticipate developments as much as possible. Criminals are already harvesting personal data with the hope that smarter AI and quantum computers may permit decrypting the data decades from now [157]. Futureproofing and staying ahead of the developments in cybersecurity is a paramount objective.

Scalability present another significant hurdle. As data storage needs grow, AI solutions must be able to scale without performance degradation. Ensuring that AI algorithms can handle increased workloads and diverse data structures efficiently is crucial. Implementing modular and flexible AI architecture that allows for incremental updates and performance optimization will help maintain scalability while keeping pace with evolving storage demands. Here again, future-proofing any storage solutions is critical, as data volumes may suddenly explode requiring immediate solutions. Of particular promise is work on scalable and distributed storage [57, 163].

Organizations and regulators must also navigate the *economic and organizational implications* of AIDE adoption. AI-driven storage systems have the potential to significantly reduce time spent on data management. This efficiency increase allows employees to focus on higher-value work, improving overall productivity. Nonetheless, issues of job displacement, trust, integration into existing organizational routines, implications for organizational strategy and general market dynamics, must be thoroughly investigated to address any negative consequences.

By addressing these challenges proactively and leveraging AI’s capabilities strategically, organizations and individuals can unlock the full potential of agentic AI in storage management.

5. Conclusion

Digital data generation has grown exponentially in recent years, creating unprecedented challenges in storage, organization, and retrieval. Traditional data management systems rely heavily on manual human intervention, leading to inefficiencies, accessibility challenges, and significant risks. Artificial intelligence driven by such breakthroughs as deep learning, large language models, Agentic AI, presents a solution that can address the fundamental limitations of data management by introducing intelligent, autonomous data management capabilities at scale and for everyone.

Agentic AI represents a paradigm shift in digital storage, moving beyond traditional file management paradigms. By introducing intelligent, autonomous systems capable of understanding context, predicting user needs, and dynamically managing digital assets, we can create more efficient, secure, and user-friendly data management for everyone.

6. References

- [1] Aleksi Aaltonen, Cristina Alaimo, Elena Parmiggiani, Marta Stelmaszak, Sirkka L Jarvenpaa, Jannis Kallinikos, and Eric Monteiro. 2023. What is Missing from Research on Data in Information Systems? Insights from the Inaugural Workshop on Data Research. *Communications of the Association for Information Systems* 53, 1 (2023), 17.
- [2] Aleksi Aaltonen and Marta Stelmaszak. 2024. Data Innovation Lens: A New Way to Approach Data Design as Value Creation. Available at SSRN 4574855 (2024).
- [3] Ahmed Abbasi, Yilu Zhou, Shasha Deng, and Pengzhu Zhang. 2018. Text Analytics to Support Sense-Making in Social Media: A Language-Action Perspective. *MIS Quarterly* 42, 2 (2018), 1–38.
- [4] Amina Adadi and Mohammed Berrada. 2018. Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE access* 6, (2018), 52138–52160.
- [5] Julius A Adebayo. 2016. FairML: ToolBox for diagnosing bias in predictive modeling. (2016).
- [6] Eugene Agichtein, Carlos Castillo, Debora Donato, Aristides Gionis, and Gilad Mishne. 2008. Finding high-quality content in social media. In *Proceedings of the international conference on Web search and web data mining*, 2008. ACM, Palo Alto, California, USA, 183–194.
- [7] Cristina Alaimo and Jannis Kallinikos. 2024. *Data rules: Reinventing the market economy*. MIT Press, Cambridge, MA.
- [8] Maryam Alavi, Dorothy E Leidner, and Reza Mousavi. 2024. A Knowledge Management Perspective of Generative Artificial Intelligence. *Journal of the Association for Information Systems* 25, 1 (2024), 1–12.
- [9] John Robert Anderson. 1985. *Cognitive psychology and its implications*. WH Freeman/Times Books/Henry Holt & Co.
- [10] Ofer Arazy, Nanda Kumar, and Bracha Shapira. 2010. A theory-driven design framework for social recommender systems. *Journal of the Association for Information Systems* 11, 9 (2010), 455–490.
- [11] A. Archambault, F. Gosselin, and P. G. Schyns. 2000. A natural bias for the basic level? *Twenty-Second Annual Conference of the Cognitive Science Society* (2000), 585-590 1075.
- [12] Lauryn Arora, Akansh Choudhary, Margi Bhatt, Jayakumar Kaliappan, and Kathiravan Srinivasan. 2024. A comprehensive review on NUI, multi-sensory interfaces and UX design for applications and devices for visually impaired users. *Frontiers in Public Health* 12, (2024), 1357160.
- [13] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, and Richard Benjamins. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* 58, (2020), 82–115.

- [14] Aaron Baird and Likoebe M Maruping. 2021. The Next Generation of Research on IS Use: A Theoretical Framework of Delegation to and from Agentic IS Artifacts. *MIS quarterly* 45, 1 (2021).
- [15] Karen Barad. 2007. *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. Duke University Press, Durham, North Carolina.
- [16] C. Batini and M. Scannapieca. 2006. *Data quality: concepts, methodologies and techniques*. Springer.
- [17] Kent Beck, Mike Beedle, Arie Van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler, James Grenning, Jim Highsmith, Andrew Hunt, and Ron Jeffries. 2001. Manifesto for agile software development. *Agile Alliance* (2001). Retrieved from [https:// agilemanifesto.org](https://agilemanifesto.org)
- [18] Mohamed Benali-Khoudja, Moustapha Hafez, Jean-Marc Alexandre, and Abderrahmane Kheddar. 2004. Tactile interfaces: a state-of-the-art survey. 2004. Citeseer, 23–26.
- [19] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 2021. 610–623.
- [20] Daniel J Bernstein and Tanja Lange. 2017. Post-quantum cryptography. *Nature* 549, 7671 (2017), 188–194.
- [21] Ludwig von Bertalanffy. 1968. General system theory: Foundations, development, applications. *Braziller, New York* (1968).
- [22] Bikram Pratim Bhuyan, Amar Ramdane-Cherif, Ravi Tomar, and TP Singh. 2024. Neuro-symbolic artificial intelligence: a survey. *Neural Computing and Applications* 36, 21 (2024), 12809–12844.
- [23] Tomas J Bird, Amanda E Bates, Jonathan S Lefcheck, Nicole A Hill, Russell J Thomson, Graham J Edgar, Rick D Stuart-Smith, Simon Wotherspoon, Martin Krkosek, and Jemina F Stuart-Smith. 2014. Statistical solutions for error and bias in global citizen science datasets. *Biological Conservation* 173, (2014), 144–154.
- [24] Johan J Bolhuis, Stephen Crain, Sandiway Fong, and Andrea Moro. 2024. Three reasons why AI doesn't model human language. *Nature* 627, 8004 (2024), 489–489.
- [25] Emily Bonnie. 2025. 110+ of the Latest Data Breach Statistics [Updated 2025]. *Secureframe*. Retrieved April 1, 2025 from <https://secureframe.com/blog/data-breach-statistics>
- [26] Jason Borenstein and Ayanna Howard. 2021. Emerging challenges in AI and the need for AI ethics education. *AI and Ethics* 1, 1 (2021), 61–65.
- [27] Vadim Borisov, Tobias Leemann, Kathrin Seßler, Johannes Haug, Martin Pawelczyk, and Gjergji Kasneci. 2022. Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems* (2022).
- [28] P. Bornet, J. Wirtz, T. H. Davenport, D. De Cremer, B. Evergreen, P. Fersht, R. Gohel, S. Khiyara, P. Sund, and N. Mullakara. 2025. *Agentic Artificial Intelligence: Harnessing AI Agents to Reinvent Business, Work and Life*. Irreplaceable Publishing, NA. Retrieved from <https://books.google.com/books?id=UfRTEQAAQBAJ>
- [29] Nick Bostrom. 1998. How long before superintelligence? *International Journal of Futures Studies* 2, 1 (1998), 1–9.
- [30] Nick Bostrom. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, Oxford, England.
- [31] Nick Bostrom and Eliezer Yudkowsky. 2018. The ethics of artificial intelligence. In *Artificial intelligence safety and security*. Chapman and Hall/CRC, 57–69.
- [32] Marilyn B. Brewer. 1979. In-group bias in the minimal intergroup situation: A cognitive-motivational analysis. *Psychological Bulletin* 86, 2 (1979), 307–324.
- [33] Kristin Briney. 2015. *Data Management for Researchers: Organize, maintain and share your data for research success*. Pelagic Publishing Ltd, Exeter, UK.
- [34] Joy Buchanan, Stephen Hill, and Olga Shapoval. 2024. ChatGPT Hallucinates Non-existent Citations: Evidence from Economics. *The American Economist* 69, 1 (2024), 80–87.
- [35] S Buranyi. 2017. Rise of the racist robots—how AI is learning all our worst impulses. *The Guardian*. August 8, (2017).

- [36] Andrew Burton-Jones and Olga Volkoff. 2017. How can we develop contextualized theories of effective use? A demonstration in the context of community-care electronic health records. *Information Systems Research* 28, 3 (2017), 468–489.
- [37] Elizabeth Byrne, Judith A Vessey, and Lauren Pfeifer. 2018. Cyberbullying and social media: Information and interventions for school nurses working with victims, students, and families. *The Journal of School Nursing* 34, 1 (2018), 38–50.
- [38] Jaclyn Cabral. 2011. Is generation Y addicted to social media? *The Elon Journal of Undergraduate Research in Communications* 2, 1 (2011), 5–14.
- [39] Erik Cambria, Lorenzo Malandri, Fabio Mercorio, Navid Nobani, and Andrea Seveso. 2024. Xai meets llms: A survey of the relation between explainable ai and large language models. *arXiv preprint arXiv:2407.15248* (2024).
- [40] Yongcan Cao, Wenwu Yu, Wei Ren, and Guanrong Chen. 2012. An overview of recent progress in the study of distributed multi-agent coordination. *IEEE Transactions on Industrial Informatics* 9, 1 (2012), 427–438.
- [41] Arturo Castellanos, Alfred Castillo, Roman Lukyanenko, and Monica Chiarini Tremblay. 2017. Understanding Benefits and Limitations of Unstructured Data Collection for Repurposing Organizational Data. In *Information Systems: Research, Development, Applications, Education: 10th SIGSAND/PLAIS EuroSymposium 2017, Gdansk, Poland, September 22, 2017, Proceedings*, Stanisław Wrycza and Jacek Maślankowski (eds.). Springer International Publishing, Cham, 13–24. https://doi.org/10.1007/978-3-319-66996-0_2
- [42] Arturo Castellanos, Monica Tremblay, Roman Lukyanenko, and Binny M Samuel. 2020. Basic Classes in Conceptual Modeling: Theory and Practical Guidelines. *Journal of the Association for Information Systems* 21, 4 (2020), 1001–1044.
- [43] Leona Chandra Kruse and Stefan Seidel. 2017. Tensions in Design Principle Formulation and Reuse. In *DESRIST 2017, 2017*. .
- [44] Leona Chandra Kruse, Stefan Seidel, and Sandeep Puro. 2016. Making Use of Design Principles. In *LNCS 9661, 2016*. Springer, Berlin / Heidelberg, 37–51.
- [45] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, and Yidong Wang. 2023. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology* (2023).
- [46] Hsinchun Chen, Roger HL Chiang, and Veda C. Storey. 2012. Business Intelligence and Analytics: From Big Data to Big Impact. *MIS Quarterly* 36, 4 (2012), 1165–1188.
- [47] Wallace Chipidza, Yan Li, Atefeh Mashatan, Ozgur Turetken, and Lorne Olfman. 2023. Quantum Computing and IS-Harnessing the Opportunities of Emerging Technologies. *Communications of the Association for Information Systems* 52, 1 (2023), 7.
- [48] Noam Chomsky, Ian Roberts, and Jeffrey Watumull. 2023. Opinion | Noam Chomsky: The False Promise of ChatGPT. *The New York Times*. Retrieved March 13, 2024 from <https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html>
- [49] Cecil Eng Huang Chua, Marta Indulska, Roman Lukyanenko, Wolfgang Maass, and Veda C. Storey. 2022. Data Management. *MIS Quarterly Online* (2022), 1–10.
- [50] Russell M Church. 2002. The effective use of secondary data. *Learning and motivation* 33, 1 (2002), 32–45.
- [51] Carlo Combi, Barbara Oliboni, and Francesca Zerbatto. 2019. A modular approach to the specification and management of time duration constraints in BPMN. *Information Systems* 84, (2019), 111–144.
- [52] Caren Cooper. 2016. *Citizen science: How ordinary people are changing the face of discovery*. Abrams, New York NY.
- [53] DAMA, D. Henderson, and S. Earley. 2024. *DAMA-DMBOK: Data Management Body of Knowledge* (2nd ed.). Technics Publications, Sedona, AZ.
- [54] Thomas Davenport, Abhijit Guha, Dhruv Grewal, and Timna Bressgott. 2020. How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science* 48, (2020), 24–42.

- [55] Pablo Duboue. 2020. *The Art of Feature Engineering: Essentials for Machine Learning*. Cambridge University Press, Cambridge, UK.
- [56] Owen Eriksson, Paul Johannesson, and Maria Bergholtz. 2019. The case for classes and instances—a response to representing instances: the case for reengineering conceptual modelling grammars. *European Journal of Information Systems* 28, 6 (2019), 681–693.
- [57] Ixchel M Faniel and Ann Zimmerman. 2011. Beyond the data deluge: A research agenda for large-scale data sharing and reuse. *International Journal of Digital Curation* 6, 1 (2011), 58–69.
- [58] Luciano Floridi and Josh Cowls. 2021. A unified framework of five principles for AI in society. In *Ethics, Governance, and Policies in Artificial Intelligence*. Springer, 5–17.
- [59] Batya Friedman and Helen Nissenbaum. 1996. Bias in computer systems. *ACM Transactions on Information Systems (TOIS)* 14, 3 (1996), 330–347.
- [60] Armand Ruiz Gabernet and Jay Limburn. 2017. Breaking the 80/20 rule: How data catalogs transform data scientists’ productivity. *IBM Cloud Blog*. Retrieved January 1, 2023 from <https://www.ibm.com/cloud/blog/ibm-data-catalog-data-scientists-productivity>
- [61] Megan Garcia. 2016. Racist in the machine: The disturbing implications of algorithmic bias. *World Policy Journal* 33, 4 (2016), 111–117.
- [62] Joey F George, Kevin Duffy, and Manju Ahuja. 2000. Countering the anchoring and adjustment bias with decision support systems. *Decision Support Systems* 29, 2 (2000), 195–206.
- [63] Thomas Gilovich, Dale Griffin, and Daniel Kahneman. 2002. *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press, Cambridge, UK.
- [64] Davide Grande, Jorge Machado, Bryan Petzold, and Marcus Roth. 2020. Reducing data costs without sacrificing growth | McKinsey. *McKinsey & Company*, 1–13.
- [65] Shirley Gregor and Dirk Hovorka. 2011. Causality: The elephant in the room in information systems epistemology. In *19th European Conference on Information Systems*, 2011. Helsinki, Finland, 1–10.
- [66] Thilo Hagendorff. 2020. The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines* 30, 1 (2020), 99–120.
- [67] Rockne Harmon. 2005. Admissibility standards for scientific evidence. In *Microbial Forensics*. Elsevier, 381–392.
- [68] Alan Hevner, Jeffrey Parsons, Alfred Benedikt Brendel, Roman Lukyanenko, Verena Tiefenbeck, Monica Chiarini Tremblay, and Jan vom Brocke. 2024. Transparency in Design Science Research. *DSS* 182, 1 (2024), 1–19.
- [69] Emily L Horton, Ramkesh Renganathan, Bryan N Toth, Alexa J Cohen, Andrea V Bajcsy, Amelia Bateman, Mathew C Jennings, Anish Khattar, Ryan S Kuo, and Felix A Lee. 2017. A review of principles in design and usability testing of tactile technology for individuals with visual impairments. *Assistive technology* 29, 1 (2017), 28–36.
- [70] D. S. Hovorka and M. Germonprez. 2009. Tinkering, tailoring and bricolage: Implications for theories of design. *Proceedings of the 15th Americas Conference on Information Systems* (2009).
- [71] Ellen M. Hufnagel and Christopher Conca. 1994. User Response Data: The Potential for Errors and Biases. *Information Systems Research* 5, 1 (March 1994), 48–73.
- [72] Merete Hvalshagen, Roman Lukyanenko, and Binny M Samuel. 2023. Empowering Users with Narratives: Examining The Efficacy Of Narratives For Understanding Data-Oriented Conceptual Models. *Information Systems Research* 34, 3 (2023), 890–909.
- [73] IEA. 2023. Data centres & networks. *IEA*. Retrieved September 13, 2024 from <https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks>
- [74] Panagiotis G Ipeirotis, Foster Provost, and Jing Wang. 2010. Quality management on amazon mechanical turk. In *Proceedings of the ACM SIGKDD workshop on human computation*, 2010. ACM, 64–67.

- [75] Anne Ixmeier, Philipp Staudt, Marie-Claude Boudreau, Gilbert Fridgen, Marc Körner, Jyoti Mishra, and Richard Watson. 2024. Green Is Everywhere And Nowhere Should Information Systems Scholarship Engage in Sustainability? In *ECIS 2024*, 2024. .
- [76] Nick R Jennings. 1993. Commitments and conventions: The foundation of coordination in multi-agent systems. *The knowledge engineering review* 8, 3 (1993), 223–250.
- [77] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys* 55, 12 (2023), 1–38.
- [78] Anna Jobin, Marcello Lenca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (2019), 389–399.
- [79] Lee Jussim, Thomas E. Nelson, Melvin Manis, and Sonia Soffin. 1995. Prejudice, stereotypes, and labeling effects: Sources of bias in person perception. *Journal of personality and social psychology* 68, 2 (1995), 228–246.
- [80] Daniel Kahneman. 2011. *Thinking, fast and slow*. Macmillan Publishing Co., London, UK.
- [81] Arpan Kumar Kar, Wu He, Fay Cobb Payton, Varun Grover, Adil S Al-Busaidi, and Yogesh K Dwivedi. 2024. How could quantum computing shape information systems research—An editorial perspective and future research directions. *International Journal of Information Management* (2024), 102776.
- [82] Akif Khan and Shah Khusro. 2021. An insight into smartphone-based assistive solutions for visually impaired and blind people: issues, challenges and opportunities. *Universal Access in the Information Society* 20, 2 (2021), 265–298.
- [83] Hyun K Kim, Sung H Han, Jaehyun Park, and Joohwan Park. 2016. The interaction experiences of visually impaired people with assistive technology: A case study of smartphones. *International Journal of Industrial Ergonomics* 55, (2016), 22–33.
- [84] Hyung Nam Kim, Tonya L Smith-Jackson, and Brian M Kleiner. 2014. Accessible haptic user interface design approach for users with visual impairments. *Universal access in the information society* 13, (2014), 415–437.
- [85] Brent Kitchens, Steven L Johnson, and Peter Gray. 2020. Understanding Echo Chambers and Filter Bubbles: The Impact of Social Media on Diversification and Partisan Shifts in News Consumption. *MIS Quarterly* 44, 4 (2020).
- [86] John E Kurtz and Emily M Pintarelli. 2024. The Daubert standards for admissibility of evidence based on the Personality Assessment Inventory. *Psychological Injury and Law* 17, 2 (2024), 105–116.
- [87] Aaron M Kushner and Zhibin Guan. 2011. Modular design in natural and biomimetic soft materials. *Angewandte Chemie International Edition* 50, 39 (2011), 9026–9057.
- [88] Kai R Larsen and Daniel S. Becker. 2020. *Automated Machine Learning for Business: An Introduction to Accurate, Easy, and Fast Analytics*. Oxford, UK, New York NY.
- [89] Kai R Larsen, Dirk Hovorka, Alan Dennis, and Jevin D West. 2019. Understanding the Elephant: The Discourse Approach to Boundary Identification and Corpus Construction for Theory Review Articles. *Journal of the Association for Information Systems* 20, 7 (2019), 15.
- [90] Kai R Larsen, Roman Lukyanenko, Roland Mueller, Veda C. Storey, Jeffrey Parsons, Debra Vander Meer, and Dirk Hovorka. 2025. Validity in Design Science. *MIS Quarterly* (Forthcoming 2025), 1–40.
- [91] Nilli Lavie, Aleksandra Hirst, Jan W De Fockert, and Essi Viding. 2004. Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General* 133, 3 (2004), 339.
- [92] Margaret Law. 2006. Reduce, reuse, recycle: Issues in the secondary use of research data. *IASSIST Quarterly* 29, 1 (2006), 5–5.
- [93] Yang W. Lee. 2003. Crafting Rules: Context-Reflective Data Quality Problem Solving. *Journal of Management Information Systems* 20, 3 (2003), 93–119.
- [94] Olivier Leroux. 2004. Legal admissibility of electronic evidence. *International Review of Law, Computers & Technology* 18, 2 (2004), 193–220.

- [95] Jingjing Li, Kai R. Larsen, and Ahmed Abbasi. 2020. TheoryOn: Designing a Construct-Based Search Engine to Reduce Information Overload for Behavioral Science Research. *MIS quarterly* 44, 4 (2020), 1733–1772.
- [96] Daniel A Lidar and Todd A Brun. 2013. *Quantum error correction*. Cambridge university press, Cambridge.
- [97] S.L. Lohr. 2009. *Sampling: Design and Analysis*. Cengage Learning. Retrieved from <https://books.google.com/books?id=aSXXKbyNIMQC>
- [98] Peter Loos, Wolfgang Nebel, Jorge Marx Gómez, Helen Hasan, Richard T Watson, Jan vom Brocke, Stefan Seidel, and Jan Recker. 2011. Green IT: a matter of business and information systems engineering? *Business & Information Systems Engineering* 3, 4 (2011), 245–252.
- [99] Roman Lukyanenko. 2016. Information Quality Research Challenge: Information Quality in the Age of Ubiquitous Digital Intermediation. *J. Data and Information Quality* 7, 1–2 (2016), 3:1–3:3. <https://doi.org/10.1145/2856038>
- [100] Roman Lukyanenko. 2024. The MAGIC of Data Management: Understanding the Value and Activities of Data Management. *arXiv preprint arXiv:2408.07607* (2024).
- [101] Roman Lukyanenko. 2025. Quantum Data Management to Reduce Environmental Impact of IT. In *SIG GREEN 2024, 2025*. Bangkok, Thailand, 1–5.
- [102] Roman Lukyanenko, Dominik Bork, Veda C. Storey, Jeffrey Parsons, and Oscar Pastor. 2023. Inclusive Conceptual Modeling: Diversity, Equity, Involvement, and Belonging in Conceptual Modeling. In *ER Forum 2023, 2023*. Springer, Lisbon, Portugal, 1–4.
- [103] Roman Lukyanenko, Arturo Castellanos, Veda C Storey, Alfred Castillo, Monica Chiarini Tremblay, and Jeffrey Parsons. 2020. Superimposition: Augmenting Machine Learning Outputs with Conceptual Models for Explainable AI. In *1st International Workshop on Conceptual Modeling Meets Artificial Intelligence and Data-Driven Decision Making, 2020*. Springer, Vienna, Austria, 1–12.
- [104] Roman Lukyanenko, Joerg Evermann, and Jeffrey Parsons. 2014. Instantiation Validity in IS Design Research. In *DESRIST 2014, LNCS 8463*. Springer, 321–328.
- [105] Roman Lukyanenko, Joerg Evermann, and Jeffrey Parsons. 2015. Guidelines for Establishing Instantiation Validity in IT Artifacts: A Survey of IS Research. In *DESRIST 2015, LNCS 9073*. Springer, Berlin / Heidelberg.
- [106] Roman Lukyanenko and Jeffrey Parsons. 2020. Design Theory Indeterminacy: What is it, how can it be reduced, and why did the polar bear drown? *Journal of the Association for Information Systems* 21, 5 (2020), 1–30.
- [107] Roman Lukyanenko, Jeffrey Parsons, and Binny M Samuel. 2018. Artifact Sampling: Using Multiple Information Technology Artifacts to Increase Research Rigor. In *Proceedings of the 51st Hawaii International Conference on System Sciences (HICSS 2018), 2018*. Big Island, Hawaii, 1–12.
- [108] Roman Lukyanenko, Jeffrey Parsons, and Binny M Samuel. 2019. Representing Instances: The Case for Reengineering Conceptual Modeling Grammars. *European Journal of Information Systems* 28, 1 (2019), 68–90.
- [109] Roman Lukyanenko, Jeffrey Parsons, Veda C. Storey, Binny M Samuel, Oscar Pastor, and Jabbari Araz. 2024. Universal Conceptual Modeling: Principles, Benefits, and an Agenda for Conceptual Modeling Research. *Software & Systems Modeling* (2024), 1–35.
- [110] Roman Lukyanenko, Jeffrey Parsons, and Yolanda Wiersma. 2014. The Impact of Conceptual Modeling on Dataset Completeness: A Field Experiment. In *Proceedings of the International Conference on Information Systems (ICIS), 2014*. 1–18.
- [111] Roman Lukyanenko, Jeffrey Parsons, Yolanda Wiersma, and Mahed Maddah. 2019. Expecting the Unexpected: Effects of Data Collection Design Choices on the Quality of Crowdsourced User-generated Content. *MISQ* 43, 2 (2019), 634–647.
- [112] Roman Lukyanenko, Binny M Samuel, Jeffrey Parsons, Veda C. Storey, and Oscar Pastor. 2023. Datish: A Universal Conceptual Modeling Language to Model Anything by Anyone. In *ER Forum 2023, 2023*. Springer, Lisbon, Portugal, 1–14.

- [113] Roman Lukyanenko and Ron Weber. 2022. A Realist Ontology of Digital Objects and Digitalized Systems. In *“Digital First” Era — A Joint AIS SIGSAND/SIGPrag Workshop, 2022*. Virtual Workshop, 1–5.
- [114] Haoyan Luo and Lucia Specia. 2024. From understanding to utilization: A survey on explainability for large language models. *arXiv preprint arXiv:2401.12874* (2024).
- [115] Viktor Mayer-Schönberger and Kenneth Cukier. 2013. *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt.
- [116] Andrew McAfee and Erik Brynjolfsson. 2017. *Machine, platform, crowd: Harnessing our digital future*. WW Norton & Company, New York, NY.
- [117] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [118] Oussama Metatla, Alison Oldfield, Taimur Ahmed, Antonis Vafeas, and Sunny Miglani. 2019. Voice user interfaces in schools: Co-designing for inclusion with visually-impaired and sighted pupils. 2019. 1–15.
- [119] George Miller. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The psychological review* 63, (1956), 81–97.
- [120] Steve Morgan. 2018. Cybercrime To Cost The World \$10.5 Trillion Annually By 2025. *Cybercrime Magazine*. Retrieved from <https://cybersecurityventures.com/hackerpocalypse-cybercrime-report-2016/>
- [121] Fuseini Mumuni and Alhassan Mumuni. 2025. Explainable artificial intelligence (XAI): from inherent explainability to large language models. *arXiv preprint arXiv:2501.09967* (2025).
- [122] Emma Murphy, Ravi Kuber, Graham McAllister, Philip Strain, and Wai Yu. 2008. An empirical investigation into the difficulties experienced by visually impaired Internet users. *Universal Access in the Information Society* 7, (2008), 79–91.
- [123] Heather Murphy. 2018. She Helped Crack the Golden State Killer Case. Here’s What She’s Going to Do Next. *The New York Times*. Retrieved December 20, 2018 from <https://www.nytimes.com/2018/08/29/science/barbara-rae-venter-gsk.html>
- [124] R. Ryan Nelson, Peter A. Todd, and Barbara H. Wixom. 2005. Antecedents of information and system quality: an empirical examination within the context of data warehousing. *Journal of Management Information Systems* 21, 4 (2005), 199–235.
- [125] Donald A. Norman. 2002. *The design of everyday things*. Basic Books, New York, NY.
- [126] Marianna Obrist, Carlos Velasco, Chi Thanh Vi, Nimesha Ranasinghe, Ali Israr, Adrian D Cheok, Charles Spence, and Ponnampalam Gopalakrishnakone. 2016. Touch, taste, & smell user interfaces: The future of multisensory HCI. 2016. 3285–3292.
- [127] Shawn Ogunseye and Jeffrey Parsons. 2016. Can Expertise Impair the Quality of Crowdsourced Data? In *SIGOPEN Developmental Workshop at ICIS 2016*, 2016. .
- [128] Shawn Ogunseye, Jeffrey Parsons, and Roman Lukyanenko. 2017. Do Crowds Go Stale? Exploring the Effects of Crowd Reuse on Data Diversity. In *WITS 2017*, 2017. Seoul, South Korea.
- [129] Seong Joon Oh, Bernt Schiele, and Mario Fritz. 2019. Towards reverse-engineering black-box neural networks. *Explainable AI: interpreting, explaining and visualizing deep learning* (2019), 121–144.
- [130] Cathy O’Neil. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books, New York NY.
- [131] Fred Paas, Alexander Renkl, and John Sweller. 2003. Cognitive load theory and instructional design: Recent developments. *Educational psychologist* 38, 1 (2003), 1–4.
- [132] Min-Seok Pang, Atreyi Kankanhalli, Margunn Aanestad, Sudha Ram, and Likoebe M Maruping. 2024. DIGITAL TECHNOLOGIES AND THE ADVANCEMENT OF SOCIAL JUSTICE: A FRAMEWORK AND AGENDA. *MIS Quarterly* 48, 4 (2024).
- [133] Jiyong Park, Kunsoo Han, and Byungtae Lee. 2023. Green cloud? An empirical analysis of cloud computing and energy efficiency. *Management Science* 69, 3 (2023), 1639–1664.

- [134] Shubham Pateria, Budhitama Subagdja, Ah-hwee Tan, and Chai Quek. 2021. Hierarchical reinforcement learning: A comprehensive survey. *ACM Computing Surveys (CSUR)* 54, 5 (2021), 1–35.
- [135] G. Pitron. 2023. *The Dark Cloud: how the digital world is costing the earth*. Scribe Publications Pty Limited, London, UK. Retrieved from <https://books.google.com/books?id=0-6qEAAAQBAJ>
- [136] Aske Plaat, Annie Wong, Suzan Verberne, Joost Broekens, Niki van Stein, and Thomas Back. 2024. Reasoning with large language models, a survey. *arXiv preprint arXiv:2407.11511* (2024).
- [137] John Preskill. 2018. Quantum computing in the NISQ era and beyond. *Quantum* 2, (2018), 79.
- [138] James Rachels and Stuart Rachels. 2012. *The elements of moral philosophy 7e*. McGraw Hill, New York NY.
- [139] Arun Rai. 2020. Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science* 48, 1 (2020), 137–141.
- [140] T. C. Redman. 1996. *Data quality for the information age*. Artech House, Norwood, MA.
- [141] Thomas C Redman. 2016. Bad data costs the US \$3 trillion per year. *Harvard Business Review* 22, (2016), 11–18.
- [142] Roman Rietsche, Christian Dremel, Samuel Bosch, Léa Steinacker, Miriam Meckel, and Jan-Marco Leimeister. 2022. Quantum computing. *Electronic Markets* 32, 4 (2022), 2525–2536.
- [143] Lionel P Robert Jr, Gaurav Bansal, Nigel Melville, and Thomas Stafford. 2020. Introduction to the special issue on AI fairness, trust, and ethics. *AIS Transactions on Human-Computer Interaction* 12, 4 (2020), 172–178.
- [144] Jesus Rodriguez. 2017. Some AI Lessons from Watson’s Failure at MD Anderson. *Medium*. Retrieved July 28, 2024 from <https://jrodthoughts.medium.com/some-ai-lessons-from-watson-s-failure-at-md-anderson-9b895cf70840>
- [145] Ben Roelens and Dominik Bork. 2020. An evaluation of the intuitiveness of the PGA modeling language notation. In *Enterprise, Business-Process and Information Systems Modeling*. Springer, 395–410.
- [146] Bruce Rosenblum and Fred Kuttner. 2011. *Quantum enigma: Physics encounters consciousness*. Oxford University Press.
- [147] Stuart Russell. 2019. *Human compatible: AI and the problem of control*. Penguin UK, London England.
- [148] Stuart J Russell and Peter Norvig. 2016. *Artificial intelligence: a modern approach*. Pearson Education Limited, Malaysia.
- [149] Carl F Salk, Tobias Sturn, Linda See, Steffen Fritz, and Christoph Perger. 2015. Assessing quality of volunteer crowdsourcing contributions: lessons from the Cropland Capture game. *International Journal of Digital Earth* 1 (2015), 1–17.
- [150] Binny M Samuel, Vijay Khatri, and V. Ramesh. 2018. Exploring the Effects of Extensional Versus Intentional Representations on Domain Understanding. *MIS Quarterly* 42, 4 (2018), 1187–1209.
- [151] Katharine Sanderson. 2023. AI science search engines explode in number. *Nature* 616, (2023), 639–640.
- [152] Aylin Sarioglu, Haydar Metin, and Dominik Bork. 2023. How Inclusive is Conceptual Modeling? A Systematic Review of Literature and Tools for Disability-aware Conceptual Modeling. In *42nd International Conference on Conceptual Modeling (ER 2023)*, 2023. 1–15.
- [153] Cristina Sernadas, José Fiadeiro, and Amílcar Sernadas. 1990. Modular construction of logic knowledge bases: An algebraic approach. *Information Systems* 15, 1 (1990), 37–59.
- [154] Victor S Sheng, Foster Provost, and Panagiotis G Ipeirotis. 2008. Get another label? improving data quality and data mining using multiple, noisy labelers. In *ACM SIG KDD*, 2008. 614–622.
- [155] Ben Shneiderman. 2003. *Designing the user interface*. Pearson Education, Boston, MA.
- [156] Ben Shneiderman. 2022. *Human-centered AI*. Oxford University Press, Oxford, UK.
- [157] Harbaksh Singh. 2024. Managing the Quantum Cybersecurity Threat: Harvest Now, Decrypt Later. In *Quantum Computing*. CRC Press, New York NY, 142–158.

- [158] Lars Skyttner. 1996. General systems theory: origin and hallmarks. *Kybernetes* 26, 6 (1996), 16–22.
- [159] Christopher Upham Murray Smith. 2008. *Biology of sensory systems*. John Wiley & Sons, Hoboken, NJ.
- [160] Ida Someh, Michael Davern, Christoph F Breidbach, and Graeme Shanks. 2019. Ethical issues in big data analytics: A stakeholder perspective. *Communications of the Association for Information Systems* 44, 1 (2019), 34.
- [161] Veda C Storey, Roman Lukyanenko, Jeffrey Parsons, and Wolfgang Maass. 2022. Explainable AI: Opening the Black Box or Pandora’s Box? *Communications of the ACM* 66, 4 (2022), 1–6.
- [162] Veda C Storey, W. T. Yue, L. J. Zhao, and Roman Lukyanenko. 2025. Generative Artificial Intelligence: Evolving Technology, Growing Societal Impact, and Opportunities for Research. *Information Systems Frontiers* (2025).
- [163] Piethen Strengtholt. 2020. *Data Management at scale*. O’Reilly Media, Inc, New York.
- [164] Diane M. Strong, Yang W. Lee, and Richard Y. Wang. 1997. Data quality in context. *Communications of the ACM* 40, 5 (1997), 103–110.
- [165] Terrence Sullivan and Stephen Adler. 1999. Work, stress, and disability. *International journal of law and psychiatry* 22, 5–6 (1999), 417–424.
- [166] Stefan Tams. 2022. Helping older workers realize their full organizational potential: a moderated mediation model of age and IT-enabled task performance. *MIS Q* 46, 1 (2022), 1–30.
- [167] Stefan Tams, M Ahuja, Jason Thatcher, and Varun Grover. 2020. Worker Stress in the Age of Mobile Technology: The Combined Effects of Perceived Interruption Overload and Worker Control. *Journal of Strategic Information Systems* (2020), 1–40.
- [168] D. Tegarden, Binny M Samuel, Roman Lukyanenko, A. Dennis, and B. Wixom. 2025. *Systems Analysis and Design: An Object-Oriented Approach with UML* (7th ed.). Wiley, Hoboken, NJ.
- [169] Joan Thiesen, Torben S Christensen, Thomas G Kristensen, Rikke D Andersen, Brit Brunoe, Trine K Gregersen, Mikkel Thrane, and Bo P Weidema. 2008. Rebound effects of price differences. *The International Journal of Life Cycle Assessment* 13, (2008), 104–114.
- [170] Monica Chiarini Tremblay, Kaushik Dutta, and Debra Vandermeer. 2010. Using data mining techniques to discover bias patterns in missing data. *Journal of Data and Information Quality (JDIQ)* 2, 1 (2010), 2.
- [171] Andreas Tsamados, Nikita Aggarwal, Josh Cows, Jessica Morley, Huw Roberts, Mariarosaria Taddeo, and Luciano Floridi. 2021. The ethics of algorithms: key problems and solutions. *AI & SOCIETY* (February 2021). <https://doi.org/10.1007/s00146-021-01154-8>
- [172] Amos Tversky, Daniel Kahneman, and P Moser. 1990. Judgment under uncertainty: Heuristics and biases. *Rationality in action: Contemporary approaches* (1990), 171–188.
- [173] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30, (2017).
- [174] Chi Thanh Vi, Damien Ablart, Daniel Arthur, and Marianna Obrist. 2017. Gustatory interface: The challenges of ‘how’ to stimulate the sense of taste. 2017. 29–33.
- [175] Gerit Wagner, Roman Lukyanenko, and Guy Pare. 2022. Artificial intelligence and the conduct of literature reviews. *Journal of Information Technology* 37, 2 (2022), 209–226.
- [176] Yair Wand and Richard Y. Wang. 1996. Anchoring data quality dimensions in ontological foundations. *Communications of the ACM* 39, 11 (1996), 86–95.
- [177] Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2023. Against predictive optimization: On the legitimacy of decision-making algorithms that optimize predictive accuracy. *ACM Journal on Responsible Computing* (2023).
- [178] Richard Y. Wang and Diane M. Strong. 1996. Beyond accuracy: what data quality means to data consumers. *Journal of Management Information Systems* 12, 4 (1996), 5–33.
- [179] Edward O Wilson. 2014. *The meaning of human existence*. WW Norton & Company, New York NY.

- [180] Allen W Wood. 2008. *Kantian ethics*. Cambridge University Press Cambridge.
- [181] Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. 2018. Bayesian model-agnostic meta-learning. *Advances in neural information processing systems* 31, (2018).
- [182] Tilmann Zäschke, Stefania Leone, Tobias Gmünder, and Moira C Norrie. 2013. Optimising conceptual data models through profiling in object databases. 2013. Springer, 284–297.
- [183] Chongjie Zhang and Victor Lesser. 2013. Coordinating multi-agent reinforcement learning with limited communication. 2013. 1101–1108.
- [184] Ruoqing Zhang, Marta Indulska, and Shazia Sadiq. 2019. Discovering data quality problems. *Business & Information Systems Engineering* 61, 5 (2019), 575–593.
- [185] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, and Yulong Chen. 2023. Siren’s song in the AI ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219* (2023).
- [186] Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology* 15, 2 (2024), 1–38.
- [187] Tiejun Zhu and Yujin Yang. 2024. Research on immersive interaction design based on visual and tactile feature analysis of visually impaired children. *Heliyon* 10, 1 (2024).
- [188] Srđan Živković and Dimitris Karagiannis. 2015. Towards metamodelling-in-the-large: Interface-based composition for modular metamodel development. In *Enterprise, Business-Process and Information Systems Modeling*. Springer, 413–428.