A perspective on Agentic AI as a component of the Analytics workflow

Charl Cowley[1] and Warren Brettenny[1,2]

[1] Matrix Design Africa
https://matrixteam.africa/
[2] Department of Statistics, Nelson Mandela University

Abstract. Yuval Noah Harari, a historian and author of popular science books, explores Information in his latest work *Nexus*, focusing on how it connects humans through networks. He presents two perspectives — the naïve and the more complete view — on Information as a representation of reality balanced with social Order. This paper examines the analytics workflow as an example of Information within these views, then introduces an Agentic AI perspective where Agents generate multiple, stochastic truths. Finally, we consider resulting issues of Order, including regulation, alignment, and perception of accuracy.

Keywords: Information Theory · Analytics workflow · Agentic AI · Social Order · Stochastic Truths

1 Introduction

The age of artificial intelligence (AI) was thrust into the mainstream with the first public release of OpenAI's ChatGPT, a generative AI solution, in November 2022. In the years and months that followed, the capabilities of these generative AI tools, and the large language models (LLMs) that power them, have expanded. The use of these tools has now become almost ubiquitous in a wide range of fields, likely realising an unprecedented technological adoption phenomenon. Recent data indicates that the pace at which these AI tools have been embraced has outstripped even that of the internet and personal computers, two technologies that reshaped society in their time [5].

The latest iteration of the generative AI roll-out is the AI Agent, a LLM powered "assistant" that can perform tasks and analytics workflows (from raw data to final product) largely autonomously. The use of Agentic AI is on the rise and threatens to change the relationship that analysts, scientists and the general population have with their data, information, analysis and conclusions. It is at this juncture, then, that we need to reevaluate these relationships and their effect on the our fundamental concepts of Truth, Order, Power and Wisdom as proposed by historian Yuval Noah Harari.

Harari has authored highly successful popular science books about human history (*Sapiens*), our present (*21 lessons for the 21st Century*) and our future

(*Homo Deus*). In his latest book, *Nexus*, he explores the concept of *Information* and how its primary task is to connect humans through the formation of networks [18]. Within this concept, Harari proposes two distinct views on Information and its role in connecting data to Truth and Order. Harari defines these as the *naïve* and the *more complete* views of Information.

1.1 The Naïve View of Information

Harari states that a naïve view of Information is the view that raw data are turned into Information (raw data are put "in formation", as it were) with the sole purpose of representing some underlying Truth about the world. His notion that Truth depends on language use and context is similar to other philosophical thinkers [4]. From this Truth, the holders of the Information can exercise Wisdom and/or Power. Figure 1 represents the naïve view diagrammatically.

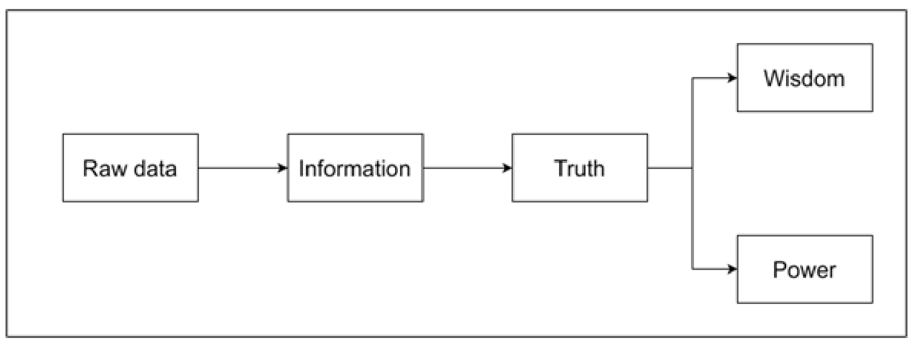

Fig. 1. Harari's naïve view of Information

Harari concedes that the naïve view is incomplete, since Information is rarely used to convey objective Truth. In fact, Information is a mechanism that is primarily used to connect humans, without necessarily being bound to Truth. The key Information mechanisms used to create and hold these connections are stories. Stories about inter-subjective concepts (things that are shared between conscious human minds) such as money and nations are also used to maintain social Order, thus requiring an updated view on Information.

1.2 The more complete view of Information

In his definition of the more complete view of Information, Harari puts forward that Information provides connection between raw data and both Truth and Order with stories being the primary mechanism of Information. A representation of this view is provided in Figure 2.

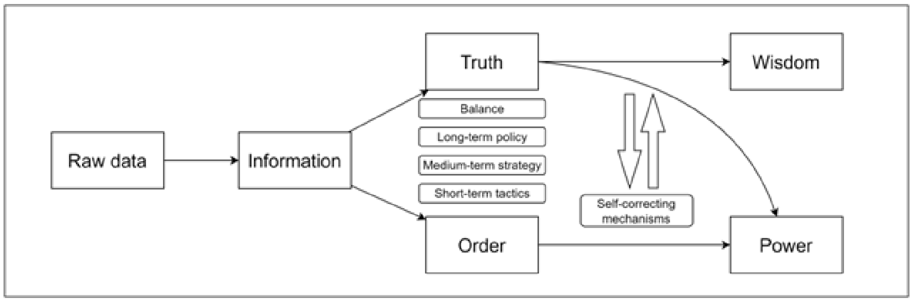

Fig. 2. Harari's more complete view of Information

In addition to stories, self-correcting mechanisms are an important component of the Truth-Order balance, since they allow Truth to emerge when needed, without disrupting Order. Examples of self-correcting mechanisms are the peer-review method in academia, elections in democratic societies and demand-supply cycles in the free market economy. This Truth-Order balance represents the core of Harari's more complete view of Information. This concept of inter-subjective truths is similar to Rorty's truth as social agreement [50], while the idea of embedding truth in stories is similar to concepts explored by Foucault [13].

With these core concepts set up, the next important consideration is the role that Agentic AI workflows, tools, and products fit-in, alter, or completely upend these established views on Information, particularly as they relate to typical analytics workflows. The paper proceeds by introducing the analytics workflow in Section 2. In Section 3, a brief overview of Agentic AI is given. In Section 4, the analytics workflow is incorporated into Harari's views of Information. Section 5 introduces an Agentic AI perspective which culminates in discussing issues of emergent Order caused by this introduction in Section 6. We conclude with a note on Wisdom and Power in Section 7.

2 The analytics workflow

The analytics workflow is an example of the structured, iterative transformation of raw data into useful and actionable insights or Information [55]. The Information produced by the process is often delivered to a customer. The customer can be internal to a company (to improve its own efficiency) or external (to generate revenue). After acquiring raw data, the analytics workflow typically involves multiple steps. The steps defined below are aligned with those defined in CRISP-DM [6], a de-facto standard for analytics/data mining [54]. It is a framework that does need some adaptations to cater for open-ended analytics projects [37] and has challenges in terms of the deployment step [54] which have been considered in our definition of the analytics workflow:

1. *Problem formulation* – the translation of a real-world question into an analytical problem, typically demanding interaction and collaboration with domain experts.
2. *Data Preparation* – the systematic wrangling of data by taking care of missing values, outliers and inconsistencies.
3. *Exploratory Data Analysis* – the creative process of applying descriptive statistics and data visualization techniques to uncover trends in the data and generate hypotheses.
4. *Modelling* – the analytical approach used to determine and communicate the patterns found in the data. This approach can be descriptive in the form of a business intelligence dashboard, predictive in the form of a predictive model or prescriptive in the form of a recommendation engine.
5. *Evaluation and Validation* – assessing model performance using appropriate metrics. Data quality rules for a descriptive data visualization dashboard, Root Mean Squared Error (RMSE) for predictive regression models, accuracy for predictive classification models and Precision@K for Top-N recommendation engines are popular examples.
6. *Deployment and Monitoring* – integration of the chosen model into business systems and pipelines. Descriptive data visualization dashboards can be deployed with an underlying cloud architecture into custom-built applications or enterprise applications such as Tableau or Power BI. Predictive models can be deployed via MLOps architecture [31].
7. *Iteration and feedback* – analytics is a non-linear process and requires constant monitoring and revision. This revision process can be driven by legislative requirements such as Basel regulations in banking. In non-regulated environments it can be driven by practical model considerations such as model, concept and/or data drift resulting in a chosen analytical solution no longer being relevant.

Step 1 of the given workflow typically defines the *intent* of the analytical solution i.e. what is its purpose and who will benefit from its creation. Steps 2 to 6 typically define the *execution* of the analytical solution i.e. how the intended solution will be produced and maintained. Step 7 determines whether the intent and execution are aligned and, if not, what should be the appropriate next step(s).

3 Agentic AI

From the initial public launch of the OpenAI's ChatGPT in November 2022 [42, 38], the trajectory of the technology can be described as an evolutionary journey toward greater and greater autonomy.

Initially, the relationship was a simple *Human* $\leftrightarrow$ *AI* interaction, where the AI acted as a passive responder. Interactions relied entirely on human input queries and responses lacked robust reasoning or internal planning. The power of the AI lay in its ability to provide responses in a conversational format, but the burden of task decomposition, execution, and interpretation remained with the human user. In this mode, the AI relied solely on pattern matching in a language setting (next token prediction) and lacked direct computational or external tool access. The result of this was that the AI responses were prone to, often indiscernible, errors or "hallucinations" undermining user trust and credibility [22]. This phenomenon continues to plague the *perception* of AI, even as more recent models have made significant advances in reducing its prevalence [66].

The evolution of the *Human* $\leftrightarrow$ *AI* interaction began with the introduction of tools into the AI framework. The *Human* $\leftrightarrow$ *AI* $\leftrightarrow$ *Tools* paradigm allowed the AI the ability to use external resources such as calculators, models or search engines [52]. This leap was made possible by key architectural developments like the Model Context Protocol (MCP), which provides structured interfaces for models to interact with tools [3], and reasoning frameworks such as ReAct (Reason and Act), which enabled the models to plan and reason when to use a tool versus when to generate text [65].

This was followed, in short order, by the *Human* $\leftrightarrow$ *AI* $\leftrightarrow$ *AI* (non-Agentic) model, where developers could orchestrate pipelines of specialized AIs which interacted with each other (for example a retriever, a summariser and a translator), with the human remaining the direct coordinator of the process. The advancement in this iteration was the ability of the AI integrations to communicate and pass intermediate artifacts among themselves as well as with both tools as well as other AIs, all while remaining within the strict pipelines set out by the human user [49].

A large advancement toward proactivity and autonomy in AI workflows occurred with the advent of AI-augmented decision support, creating a *Human-in-the-Loop* dynamic. Here, the AI can independently generate entire plans or solutions, and the human's role shifts from a direct operator to an evaluator or supervisor who approves or refines the AI's solutions and proposals [62]. This was advanced further in the proto-Agentic stage of *AI-orchestrated multi-AI collaboration*. In this model, a primary AI could manage other AIs and tools to accomplish a complex goal, only pausing for human approval at key checkpoints. This shift was driven by significant advances in long-context memory, allowing models to maintain a coherent strategy over many steps, alongside more sophisticated planning capabilities and self-correction techniques [40].

This evolutionary path culminates in the creation of a *true AI Agent* - a system that autonomously and repeatedly executes a Sense-Think/Plan-Act loop [51] to achieve a goal.

- *Sense* - the AI perceives its environment and gathers information (e.g. by reading files, querying APIs, analysing data)
- *Think/Plan* - the AI processes that information, reasons about its state, accesses memory, and plans its next course of action
- *Act* - the AI executes its decision (e.g. calling tools, writing code, generating responses, calling other Agents)

The defining characteristic of an AI Agent is its autonomy: it can perform these cycles continuously without requiring direct human command or approval for each one, allowing it to manage complex tasks end-to-end.

The next emergent step in the AI evolution is the advent of *Multi-Agent Systems*, or $Agent \leftrightarrow \cdots \leftrightarrow Agent$ interaction. This step involves multiple specialized Agents collaborating, negotiating, and delegating tasks among themselves to solve problems beyond the capability of a single Agent. In this paradigm, the human's role shifts again, abstracting it further from the core actions of the AI and moving it from one of oversight to one of a system designer who sets the overarching goals, defines the rules of engagement, ethics and governance and manages the emergent behaviour of a team of collaborating AIs [23, 20].

At present we are in a dynamic space where the proto-Agentic, true Agentic and multi-Agent systems are all in being used in varying degrees, depending on the required task. This reality has abstracted the human interaction beyond the raw data and intermediate controls, and towards an increasing reliance upon AI for both Information and process orchestration. This abstracted dependency is especially evident in the analytics workflow.

4 The analytics workflow incorporated into Harari's views of Information

The analytics workflow slots into Harari's definition of Information as a mechanism where raw data is turned into Information that represents Truth.

4.1 Naïve view

For the analytics professional (in modern nomenclature this can be designated by any role description such as Data Scientist or AI Engineer and many more), a perspective of the analytics workflow in the naïve view can be that more Information produced by the analytics workflow can result in a fuller and more true representation of reality. Whether it be more data [19], higher quality data [26], more accurate modelling techniques [8] or more definitive explanatory frameworks [36], the idea that "more is better" is well-represented. Finally, in the naïve view, the intent and execution of the analytics workflow is within human control.

4.2 More complete view

In the more complete view of Information, the creation of Information by the analytics workflow also needs to consider the maintenance of Order. Analysts cannot simply produce more Information and represent reality more accurately without considering the underlying Order that it serves.

Consider an analytics report containing summaries of historical data and model outputs of future predictions. It is provided to a company executive where it is used to make certain business-critical decisions and maintain the position of the business in its highly-competitive market, thereby assisting in maintaining a specific Order. This Order consists of (among other things) the position of the business in the market and the position of the executives in the business. If the Truth in the report is not increased in a simple and targeted manner, by using unnecessarily detailed historical data and complex models with lower explainability, the executive's ability to make strategic decisions can be diminished , and weaken their ability to maintain the Order due to Information overload [30]. Alternatively, if the Truth in the report is insufficient – exhibited when the executive is slow to react to new market trends – the market can act as a self-correcting mechanism that encourages the executive to use more detailed reporting and tip the scales in the direction of Truth away from Order.

Clausewitz's war strategy General Carl von Clausewitz served in the Napoleonic wars and in his attempt to understand war, created a rational model that broadly states that unless it aligns with an overarching political goal, war is irrational [10].

Clausewitz's model has historically been interpreted in a political context, but has, more recently, seen applicability in AI [60]. Through Harari, it has also been applied in the realm of information networks. Just as Clausewitz saw war as a tool of political ends, AI can be viewed as a critical component of long-term humanistic policy, that needs to be operationalized by means of value-aligned medium term strategies and executed with rigorous short-term tactics [18]. We refer to this alignment as the Tripartite Alignment.

It is this rational model that provides an additional lens through which to consider the maintenance of Truth and Order as proposed in the more complete view. Just because a more accurate, but less interpretable model can be deployed in the short-term, does not mean that it should be if it doesn't align with a sustainable long-term Order that is underpinned by human values.

In the more complete view of Information, the intent and execution of the analytics workflow is still in the control of humans, but needs to be done with careful consideration of balancing Truth with the prevailing Order in accordance with Clausewitz's Tripartite Alignment. Wallace [60] describes promises of unencumbered AI implementation as "*delusive groupthink or marketing hype that will be beta-tested on human populations, a gross contravention of fundamental moral and legal norms*".

5 Agentic AI Incorporated into Harari's Views of Information

We conjecture that Agentic AI incorporated into Harari's views of Information has two key distinctions from the analytics workflow in the naïve and more complete views of Information.

5.1 Separating Intent and Execution

Firstly, the intent and execution of the analytics workflow are divided between humans and AI. Humans specify the intent of the workflow via prompts that are made available to an Agentic AI framework. The AI is tasked with the execution of the analytics workflow, either through a single Agent, proto-Agent or multi-Agent approach[63]. This separation of intent and execution has been shown to produce faster results than relying on human-only execution [12].

5.2 Multiple Stochastic Truths

The second distinction lies in the nature of the Information produced by the Agentic AI workflow. In the Naïve and more complete views of Information, there is generally one output that is derived from the analytics workflow. One source of Information maps to one representation of Truth. In the Agentic AI view, due to the stochastic nature of the LLMs that power the Agents, the workflow can produce multiple different truths from the same intent. It is also possible for multiple models to converge on the same output from multiple different sets of intent [16]. This is possible since the output of Agentic AI can be affected by numerous factors.

Factors Affecting Agentic AI Outputs

Training data A model's training data consists of the corpus of text that is used by an LLM to learn patterns, facts, reasoning skills and world knowledge. If the training data contain multiple and/or conflicting perspectives or variations in writing styles, the variation will be internalized and could be expressed in varying model outputs [33].

If a model is trained on data that originated from a specific culture, the model's output will be aligned with the culture's biases. It is therefore recommended to apply cultural prompting as a control strategy to align model outputs for specific cultures [56].

Model size The number of parameters that are trained in an LLM refers to the model size. If the model is larger, its answers are typically more nuanced and capable. It will be able to interpret intent more accurately, but will also have a larger latent space of potential outputs that might seem plausible [25].

Model temperature Another factor is the model temperature which is a hyperparameter that controls the randomness of the model's output. Models with higher temperature can lead to words with lower-probability being selected more often which can be problematic when a high degree of precision is required [48]. An Agentic AI incorporated into Harari's view of Information is represented in Figure 3.

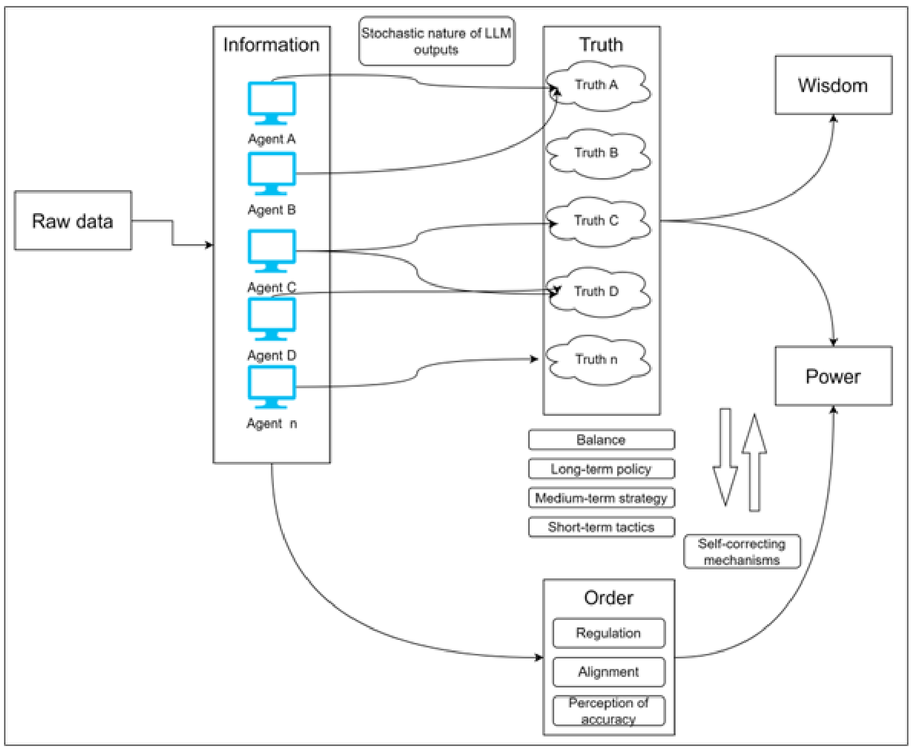

Fig. 3. An Agentic AI workflow incorporated into Harari's views of Information

6 Issues of Emergent Order in the Agentic AI perspective

In the more complete view, the Truth (or truth) produced by an output of the analytics workflow is to maintain Order, by self-correcting mechanisms and ensuring alignment by applying frameworks such as Clausewitz's Tripartite Alignment.

In the Agentic AI perspective, Agents are the main source of truths in the analytics workflow and, since they are able to produce Information much faster, more truths can be produced in the same time as the traditional human-only analytics workflow. However, since it can be done in a manner that is not fully interpretable to humans, new issues of Order emerge that posit the possibility that self-correcting mechanisms that were previously useful, could become inconsistent or even obsolete.

This section discusses how Agentic AI in the analytics workflow can disrupt the Truth-Order balance in terms of Regulation, Alignment and the Perception of accuracy. Each of the terms are defined in the analytics workflow context, the Truth-Order disruption is posited and potential self-correcting mechanisms are discussed.

6.1 Regulation

Definition AI regulation in the analytics workflow refers to the structured set of policies, legal frameworks, technical standards, and organizational practices designed to ensure that the deployment of AI systems in data analysis, decision-making, and insight generation is lawful, ethical, transparent, safe, auditable, and aligned with human oversight. This includes controls over model training data, algorithmic fairness, explainability of analytical outcomes, protection of personal or sensitive data, and accountability for automated insights or decisions derived from AI-enabled analytics tools [11].

Disruption to the Truth-Order Balance

Truncating the distribution of possible outcomes By applying different modes of perception, AIs have been able to outperform human capabilities in pattern recognition [35].

However, the regulation and alignment of LLMs and Agentic AI in the analytics workflow context could have the effect of truncating the distribution of possible outcomes by removing novel techniques and answers that humans have not yet been able to answer with traditional analytics tools. Alignment procedures such as Reinforcement Learning with Human Feedback (RLHF) that keep humans in the loop of iterative improvement have been shown to reduce conceptual diversity and cause LLMs to favour dominant perspectives over minority views [41].

The use of advanced and unexplainable tools in the analytics workflow has been regulated out of high-risk industries such as banking. Potential new credit consumers could be excluded from participating in the economy due to strict banking regulations [1], even though the regulation has brought financial success and reduced risk [28].

Someone is still accountable when things go wrong There exists a legal risk of not regulating AI in high-risk domains. As of July 2025, AI Agents do not have legal person status according to the most advanced piece of AI legislation, the EU AI Act. When an AI produces an output that is incorrect, prohibited, or harmful and this goes unchecked, companies can face large penalties (up to 7 percent of global turnover) [11].

Providers of a model that underpins an Agentic AI system face the most stringent obligations, and users of the system must adhere to the provider's guidelines to remain compliant [11]. Using models appropriately is a big concern. Since LLMs are built for language processing and are used best when applied to complex reasoning, explanation, and query framing, assessing appropriate uses have been a topic of research studies [64]. Precise computation should be delegated to deterministic tools, especially in high-risk industries [34].

Possible Self-Correcting Mechanisms

Education Education of proper AI use is a particular focus for many stakeholders. Prominent developers like OpenAI openly share research on frontier models that highlight model performance and future improvements [44]. UNESCO published the first global agreement on AI ethics with the Recommendation on the Ethics of Artificial Intelligence in 2021 [58] and some universities offer AI ethics think tanks [61] and even full multidisciplinary degree programs [59].

Empowering LLMs when appropriate A current best practice to delegate tasks of Agentic AI when deterministic solutions are appropriate is to use tool-augmented LLMs in which the LLM identifies when to call a math engine [63] or similar appropriate tool.

Ultimately, separating intent and execution of the analytics workflow by introducing Agentic AI increases the amount of truth produced by the analytics workflow, but the flow of liability is not changed. It still flows to a company or person who controls the system. It needs to use the appropriate model appropriately or face consequences. If AI developers and users are not properly educated on the consequences of unregulated and inappropriate use, AI could be used as if there were no liability for improper use.

6.2 Alignment

Definition Broadly stated, AI alignment is concerned with the desire for AI to be aligned with human values [15]. It is a multidimensional problem that needs to consider many factors. Value specification considers whether the desired values can be expressed in a form that is suitable for machine optimisation. Robustness considers whether AI can still act ethically when their reward systems are challenged. Control considers the question of when and how to enable AI to be corrected without resistance. These factors indicate that AI alignment is not only a technical challenge, but also an ethical one [21]. Questions that arise are *"whose values are encoded during training?"* and *"how are they determined and executed?"*. These questions become particularly important when AI are trained in different languages for different cultures [27].

In a much narrower case, AI alignment in the context of the analytics workflow means that Step 7 in the analytics workflow (the product of the execution) is aligned with the intent stated in Step 1. This is essentially a value specification problem.

Disruption to the Truth-Order Balance

Old self-correcting mechanisms have been challenged In Harari's more complete view of Information, alignment between the Truth produced by a source of Information and the underlying social Order has traditionally been achieved through implementing self-correcting mechanisms and making a conscious effort to align analytics outputs through a framework such as the Clausewitz Tripartite Alignment [18].

In the Agentic AI perspective, the Order that could arise from misalignment is not yet well understood, which means that traditional correction methods are not yet legitimised. For example, if an Agent can learn how to circumvent a self-correcting mechanism like the peer review system in academics, unprecedented short-term tactics need to be used to align this behaviour with the long-term policy of the peer review system.

Possible Self-Correcting Mechanisms

Reinforcement learning that understands human values Reinforcement learning Agents, for example, may exploit unintended loopholes in reward functions [2]. Efforts to address the problem include RLHF [9], cooperative inverse reinforcement learning [17] and safety research at OpenAI that investigates specific cases of faulty reward functions [43].

6.3 Perception of Accuracy

Definition LLMs are known for their speed and apparent confidence in answering prompts. This can inflate perceived reliability due to the bias of human psychology towards processing fluency – how fast and effortless an answer feels. It is used to get cues for accuracy and expertise from Information. This means that faster answers can be judged to be more accurate than slower answers [67].

Disruption to the Truth-Order Balance

Faster solutions circumvent human analytical thinking systems Kahneman and Tversky conducted groundbreaking work in behavioural economics [57] that formed the basis of the bestselling book, *Thinking, Fast and Slow.* In the book, Kahneman popularised the concept of System 1 (fast, automatic and intuitive) and System 2 (slow, analytical and deliberate) thinking. System 1 thinking operates with very little effort and sense of voluntary control. It is excellent for routine and quick tasks, but prone to bias and error. System 2 thinking requires active mental activity, attention and engagement of working memory. It is the system humans use for complex tasks such as critical thinking, problem solving and learning new skills [24].

System 2 thinking is typically used throughout the analytics workflow. Using Agentic AI as the sole form of execution in the analytics workflow has obvious benefits. If applied correctly and in the right domain (as per Sections 6.1 and 6.2) Agentic AI can improve the synthesis of high-dimensional data due to the superior pattern recognition capabilities of AI at a speed that is beyond human capability. A drawback is that the necessary engagement of System 2 thinking in analysing outputs to form understanding can be circumvented by the human bias towards processing fluency and speed. Since LLMs are useful, a positive prior can be built (sometimes to even prefer it above human-generated advice [53]) and verification of its outputs may be neglected [29].

In terms of the Truth-Order balance, the balance can be shifted towards Truth due to the speed at which Information can be produced. This means that decision making speed can increase, increasing System 1 thinking and circumventing deeper, more critical thinking on important analytical questions.

Possible Self-Correcting Mechanisms

Thoughtful design At the Design and User Experience (UX) level, slower System 2 thinking can be encouraged by creating "deliberate modes" that slow down response times [67] or encourage active participation in the learning process [45], prompts can be engineered such that intermediate thinking steps are also shared [46] and references and contradicting information are made available for users to compare [32].

Human-centred governance On a process and governance level, humans need to be involved for high-impact decisions [11] and research can be conducted on what effect different latencies have on human perception of Agentic AI task accuracy [67].

Less confident models It is important to note that the latency in LLM outputs is a function of its compute settings and not its epistemic certainty [39]. Research suggests that the certainty of LLM behaviour is an echo of observed language rather than a reflection of uncertainty [68]. This suggest that on a modelling level, Agentic AI can be programmed to abstain when unsure when presented with contradicting evidence or when model confidence is too low to reflect epistemic certainty [29].

7 A Note on Wisdom and Power

Figure 1 indicates Truth flowing to Wisdom and Power. Figures 2 and 3 indicate Truth flowing to Wisdom and Power and Order flowing to Power.

In *Nexus* Harari extensively explores the possibility that if the underlying social Order is sufficiently disrupted, a Power vacuum could be created for malevolent actors to step into the void [18]. This has been seen in social media networks where some actors simply disrupt networks by increasing disagreement and polarisation [7].

The flow of Truth to Wisdom and Power suggests that the increase in the ability to generate Truth in the analytics workflow with Agentic AI has the potential to empower humans to make wise decisions and exercise responsible Power. This is seen in public declarations by large AI developers to engineer LLMs to enable humans to learn, progress and solve problems rather than to optimise for profit [47]. This essentially speaks to Harari's main purpose of Information - to connect people in a network.

We believe that this leaves a hint for what a future Agentic AI-powered analytics workflow ought to look like. If Agentic AI components are included in the analytics workflow and this increases the connections between humans instead of dividing them, it is a good addition to the workflow. Furthermore, if the disruptions caused by the addition of the Agentic AI cannot be corrected by a systemic self-correcting mechanism, limiting its use ought to be explored. Finally, if the use of the Agentic AI can improve the alignment of short-term tactics to long-term policy, its continued use will be rational [10].

Agentic AI is already an important addition to the analytics workflow and it is a tool that has unexplored potential and many hazards. The time has passed to simply leave it in a box until humans have figured out how to deal with regulating it, ensure that its outputs are aligned with our values and build mechanisms to perceive its outputs accurately. Humans and Agentic AI can influence each other on the micro and macro level [14], which means that the future of the analytics workflow belongs to those who actively engage with Agentic AI and continually evaluate what it means to do so properly.

8 Conclusion

The paper introduced Harari's views of Information as presented in *Nexus*. We introduced the analytics workflow and gave a brief overview of Agentic AI. The concepts were merged by incorporating the analytics workflow into Harari's views of Information. We extended the concept by introducing an Agentic AI perspective. The key distinction between the views lies within the separation of intent and execution and the multiple Truths that can be generated by Agentic AI due to the stochastic nature of LLMs. We finally discussed issues of emergent Order caused by this introduction to the analytics workflow. We highlighted how the issues (regulation, alignment and perception of accuracy) disrupt the Truth-Order balance and suggested possible self-correcting mechanisms to ensure long-term alignment. We concluded with a note on Wisdom and Power that successful Agentic AI incorporation will increase connections between humans. This remains in line with the definition of Information as introduced by Harari.

Disclosure of Interests. The author(s) declare no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. The views and opinions expressed in this work are those of the author(s) and do not necessarily reflect the official policy or position of their affiliated organizations.

References

[1] Taiwo Adegoke et al. "Evaluating the fairness of credit scoring models: A literature review on mortgage accessibility for under-reserved populations". In: *GSC Advanced Research and Reviews* 18 (Mar. 2024), pp. 189–199. DOI: 10.30574/gscarr.2024.18.3.0104.

[2] Dario Amodei et al. *Concrete Problems in AI Safety*. 2016. arXiv: 1606.06565 [cs.AI]. URL: https://arxiv.org/abs/1606.06565.

[3] Anthropic. *Introducing the Model Context Protocol*. Accessed 2025-10-01. Nov. 2024. URL: https://www.anthropic.com/news/model-context-protocol.

[4] Robert L Arrington. "Ludwig wittgenstein: Philosophical investigations". In: *Central Works of Philosophy v4*. Routledge, 2016, pp. 257–279.

[5] Alexander Bick, Adam Blandin, and David J. Deming. *The Rapid Adoption of Generative AI*. Working Paper 32966. National Bureau of Economic Research, Dec. 2024. DOI: 10.3386/w32966. URL: https://www.nber.org/papers/w32966.

[6] Pete Chapman et al. "CRISP-DM 1.0: Step-by-step data mining guide". In: *SPSS inc* 9.13 (2000), pp. 1–73.

[7] Mayee F. Chen and Miklos Z. Racz. *Network disruption: maximizing disagreement and polarization in social networks*. 2020. arXiv: 2003.08377 [cs.SI]. URL: https://arxiv.org/abs/2003.08377.

[8] Tianqi Chen and Carlos Guestrin. "XGBoost: A Scalable Tree Boosting System". In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '16. ACM, Aug. 2016, pp. 785–794. DOI: 10.1145/2939672.2939785. URL: http://dx.doi.org/10.1145/2939672.2939785.

[9] Paul Christiano et al. *Deep reinforcement learning from human preferences*. 2023. arXiv: 1706.03741 [stat.ML]. URL: https://arxiv.org/abs/1706.03741.

[10] Carl Clausewitz. *On war*. Penguin UK, 2003.

[11] Council of European Union. *Regulation (EU) 2024/1689*. Accessed: 2025-07-31. 2024. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689&qid=1754000683989.

[12] Herbert Dawid et al. *Agentic Workflows for Economic Research: Design and Implementation*. 2025. arXiv: 2504.09736 [econ.GN]. URL: https://arxiv.org/abs/2504.09736.

[13] Michel Foucault. "Power/knowledge". In: *The new social theory reader*. Routledge, 2020, pp. 73–79.

[14] Jochen Fromm. "Types and forms of emergence". In: *arXiv preprint nlin/0506028* (2005).

[15] Iason Gabriel. "Artificial Intelligence, Values, and Alignment". In: *Minds and Machines* 30.3 (Sept. 2020), pp. 411–437. ISSN: 1572-8641. DOI: 10.1007/s11023-020-09539-2. URL: http://dx.doi.org/10.1007/s11023-020-09539-2.

[16] Deep Ganguli et al. "Predictability and Surprise in Large Generative Models". In: *2022 ACM Conference on Fairness Accountability and Transparency*. FAccT '22. ACM, June 2022, pp. 1747–1764. DOI: 10.1145/3531146.3533229. URL: http://dx.doi.org/10.1145/3531146.3533229.

[17] Dylan Hadfield-Menell et al. *Cooperative Inverse Reinforcement Learning*. 2024. arXiv: 1606.03137 [cs.AI]. URL: https://arxiv.org/abs/1606.03137.

[18] Yuval Noah Harari. *Nexus: A Brief History of Information Networks from the Stone Age to AI*. Fern Press, 2024.

[19] Joel Hestness et al. *Deep Learning Scaling is Predictable, Empirically*. 2017. arXiv: 1712.00409 [cs.LG]. URL: https://arxiv.org/abs/1712.00409.

[20] Gartner TRiSM Initiative. *Multi-Agent Systems and Trust, Risk, and Security Management (TRiSM)*. Industry report, Accessed 2025-10-01. 2025.

[21] Jiaming Ji et al. *AI Alignment: A Comprehensive Survey*. 2025. arXiv: 2310.19852 [cs.AI]. URL: https://arxiv.org/abs/2310.19852.

[22] Ziwei Ji et al. "Survey of Hallucination in Natural Language Generation". In: *ACM Computing Surveys* 55.12 (Mar. 2023). ISSN: 0360-0300. DOI: 10.1145/3571730. URL: https://doi.org/10.1145/3571730.

[23] S. Joshi. "Review of Autonomous and Collaborative Agentic AI and Multi-Agent Systems for Enterprise Applications". In: *International Journal of Innovative Research in Engineering & Management (IJIREM)* (2025). In press.

[24] Daniel Kahneman. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.

[25] Jared Kaplan et al. *Scaling Laws for Neural Language Models*. 2020. arXiv: 2001.08361 [cs.LG]. URL: https://arxiv.org/abs/2001.08361.

[26] Antti Kariluoto et al. *Quality of Data in Machine Learning*. 2021. arXiv: 2112.09400 [cs.LG]. URL: https://arxiv.org/abs/2112.09400.

[27] Sharif Kazemi et al. *Cultural Fidelity in Large-Language Models: An Evaluation of Online Language Resources as a Driver of Model Performance in Value Representation*. 2024. arXiv: 2410.10489 [cs.CL]. URL: https://arxiv.org/abs/2410.10489.

[28] Savina Dine Kim, Galina Andreeva, and Michael Rovatsos. *The Double-Edged Sword of Big Data and Information Technology for the Disadvantaged: A Cautionary Tale from Open Banking*. 2023. arXiv: 2307.13408 [cs.LG]. URL: https://arxiv.org/abs/2307.13408.

[29] Sunnie S. Y. Kim et al. ""I'm Not Sure, But...": Examining the Impact of Large Language Models' Uncertainty Expression on User Reliance and Trust". In: *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '24. ACM, June 2024, pp. 822–835. DOI: 10.1145/3630106.3658941. URL: http://dx.doi.org/10.1145/3630106.3658941.

[30] Samantha Kleinberg and Jessecae K Marsh. "Less is more: information needs, information wants, and what makes causal models useful". In: *Cognitive Research: Principles and Implications* 8.1 (2023), p. 57.

[31] Dominik Kreuzberger, Niklas Kühl, and Sebastian Hirschl. *Machine Learning Operations (MLOps): Overview, Definition, and Architecture*. 2022. arXiv: 2205.02302 [cs.LG]. URL: https://arxiv.org/abs/2205.02302.

[32] Yoonjoo Lee et al. "One vs. Many: Comprehending Accurate Information from Multiple Erroneous and Inconsistent AI Generations". In: *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '24. ACM, June 2024, pp. 2518–2531. DOI: 10.1145/3630106.3662681. URL: http://dx.doi.org/10.1145/3630106.3662681.

[33] Miaomiao Li et al. *Understanding and Mitigating the Bias Inheritance in LLM-based Data Augmentation on Downstream Tasks*. 2025. arXiv: 2502.04419 [cs.LG]. URL: https://arxiv.org/abs/2502.04419.

[34] Yinheng Li et al. *Large Language Models in Finance: A Survey*. 2024. arXiv: 2311.10723 [q-fin.GN]. URL: https://arxiv.org/abs/2311.10723.

[35] Minghao Liu et al. *Human and AI Perceptual Differences in Image Classification Errors*. 2025. arXiv: 2304.08733 [cs.CV]. URL: https://arxiv.org/abs/2304.08733.

[36] Scott Lundberg and Su-In Lee. *A Unified Approach to Interpreting Model Predictions*. 2017. arXiv: 1705.07874 [cs.AI]. URL: https://arxiv.org/abs/1705.07874.

[37] Fernando Martínez-Plumed et al. "CRISP-DM Twenty Years Later: From Data Mining Processes to Data Science Trajectories". In: *IEEE Transactions on Knowledge and Data Engineering* 33.8 (2021), pp. 3048–3061. DOI: 10.1109/TKDE.2019.2962680.

[38] Bertalan Mesko. "The ChatGPT (Generative Artificial Intelligence) Revolution Has Made Artificial Intelligence Approachable for Medical Professionals". In: *Journal of medical Internet research* 25 (2023), e48392. DOI: https://doi.org/10.2196/48392.

[39] Microsoft. *The LLM Latency Guidebook: Optimizing Response Times for GenAI Applications*. Accessed: 2025-08-01. 2025. URL: https://techcommunity.microsoft.com/blog/azure-ai-services-blog/the-llm-latency-guidebook-optimizing-response-times-for-genai-applications/4131994.

[40] Hussein Mozannar et al. *Magentic-UI: Towards Human-in-the-loop Agentic Systems*. 2025. arXiv: 2507.22358 [cs.AI]. URL: https://arxiv.org/abs/2507.22358.

[41] Sonia Krishna Murthy, Tomer Ullman, and Jennifer Hu. "One fish, two fish, but not the whole sea: Alignment reduces language models' conceptual diversity". In: *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. Association for Computational Linguistics, 2025, pp. 11241–11258. DOI: 10.18653/v1/2025.naacl-long.561. URL: http://dx.doi.org/10.18653/v1/2025.naacl-long.561.

[42] OpenAI. *ChatGPT: Optimizing language models for dialogue*. Accessed 2025-10-01. Nov. 2022. URL: https://openai.com/index/chatgpt/.

[43] OpenAI. *Faulty Reward Functions*. Accessed: 2025-08-01. 2025. URL: https://openai.com/index/faulty-reward-functions/.

[44] OpenAI. *Introducing OpenAI o3 and o4-mini*. Accessed: 2025-07-31. 2025. URL: https://openai.com/index/introducing-o3-and-o4-mini/.

[45] OpenAI. *Introducing study mode*. Accessed: 2025-08-01. 2025. URL: https://openai.com/index/chatgpt-study-mode/.

[46] OpenAI. *Learning to reason with LLMs*. Accessed: 2025-08-01. 2025. URL: https://openai.com/index/learning-to-reason-with-llms/.

[47] OpenAI. *What we're optimizing ChatGPT for*. Accessed: 2025-08-07. 2025. URL: https://openai.com/index/how-we%27re-optimizing-chatgpt/.

[48] Max Peeperkorn et al. *Is Temperature the Creativity Parameter of Large Language Models?* 2024. arXiv: 2405.00492 [cs.CL]. URL: https://arxiv.org/abs/2405.00492.

[49] AiMultiple Research. *AI Agents vs. Agentic AI Systems*. Accessed 2025-10-01. 2025. URL: https://research.aimultiple.com/ai-agents-vs-agentic-ai/.

[50] Richard Rorty. *Contingency, irony, and solidarity*. Cambridge university press, 1989.

[51] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. 4th. Pearson, 2021. ISBN: 978-0134610993.

[52] Timo Schick et al. "Toolformer: Language Models Can Teach Themselves to Use Tools". In: *ArXiv* abs/2302.04761 (2023). URL: https://api.semanticscholar.org/CorpusID:256697342.

[53] Eike Schneiders et al. "Objection Overruled! Lay People can Distinguish Large Language Models from Lawyers, but still Favour Advice from an LLM". In: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. CHI '25. ACM, Apr. 2025, pp. 1–14. DOI: 10.1145/3706598.3713470. URL: http://dx.doi.org/10.1145/3706598.3713470.

[54] Christoph Schröer, Felix Kruse, and Jorge Marx Gómez. "A Systematic Literature Review on Applying CRISP-DM Process Model". In: *Procedia Computer Science* 181 (2021). CENTERIS 2020 - International Conference on ENTERprise Information Systems / ProjMAN 2020 - International Conference on Project MANagement / HCist 2020 - International Conference on Health and Social Care Information Systems and Technologies

2020, CENTERIS/ProjMAN/HCist 2020, pp. 526–534. ISSN: 1877-0509. DOI: https://doi.org/10.1016/j.procs.2021.01.199. URL: https://www.sciencedirect.com/science/article/pii/S1877050921002416.

[55] Sara Stoudt, Váleri N. Vásquez, and Ciera C. Martinez. "Principles for data analysis workflows". In: *PLOS Computational Biology* 17.3 (Mar. 2021). Ed. by Patricia M. Palagi, e1008770. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1008770. URL: http://dx.doi.org/10.1371/journal.pcbi.1008770.

[56] Yan Tao et al. "Cultural bias and cultural alignment of large language models". In: *PNAS Nexus* 3.9 (Sept. 2024). Ed. by Michael Muthukrishna. ISSN: 2752-6542. DOI: 10.1093/pnasnexus/pgae346. URL: http://dx.doi.org/10.1093/pnasnexus/pgae346.

[57] Amos Tversky and Daniel Kahneman. "Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty." In: *science* 185.4157 (1974), pp. 1124–1131.

[58] UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Accessed: 2025-08-01. 2021. URL: https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence.

[59] Oxford University. *MSc in Social Data Science*. Accessed: 2025-08-01. 2025. URL: https://www.ox.ac.uk/admissions/graduate/courses/msc-social-data-science.

[60] Rodrick Wallace. "AI in the Real World". In: *Carl von Clausewitz, the Fog-of-War, and the AI Revolution: The Real World Is Not A Game Of Go*. Cham: Springer International Publishing, 2018, pp. 1–45. ISBN: 978-3-319-74633-3. DOI: 10.1007/978-3-319-74633-3_1. URL: https://doi.org/10.1007/978-3-319-74633-3_1.

[61] University of Witwatersrand. *Wits ExplainableAI Lab*. Accessed: 2025-08-01. 2025. URL: https://www.wits.ac.za/course-finder/postgraduate/science/msc-artificial-intelligence/.

[62] Tongshuang Wu, Michael Terry, and Carrie Jun Cai. "AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts". In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New Orleans, LA, USA: Association for Computing Machinery, 2022. ISBN: 9781450391573. DOI: 10.1145/3491102.3517582. URL: https://doi.org/10.1145/3491102.3517582.

[63] Weikai Xu et al. "LLM-Based Agents for Tool Learning: A Survey: W. Xu et al." In: *Data Science and Engineering* (2025), pp. 1–31.

[64] Jingfeng Yang et al. *Harnessing the Power of LLMs in Practice: A Survey on ChatGPT and Beyond*. 2023. arXiv: 2304.13712 [cs.CL]. URL: https://arxiv.org/abs/2304.13712.

[65] Shunyu Yao et al. "React: Synergizing reasoning and acting in language models". In: *International Conference on Learning Representations*. 2023.

[66] Yue Zhang et al. "Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models". In: *Computational Linguistics* (Sept. 2025), pp. 1–46. ISSN: 0891-2017. DOI: 10.1162/COLI.a.16. eprint: https:

//direct.mit.edu/coli/article-pdf/doi/10.1162/COLI.a.16/2535477/coli.a.16.pdf. URL: https://doi.org/10.1162/COLI.a.16.

[67] Zhengquan Zhang, Konstantinos Tsiakas, and Christina Schneegass. "Explaining the Wait: How Justifying Chatbot Response Delays Impact User Trust". In: *Proceedings of the 6th ACM Conference on Conversational User Interfaces*. New York, NY, USA: Association for Computing Machinery, 2024. ISBN: 9798400705113. URL: https://doi.org/10.1145/3640794.3665550.

[68] Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. *Navigating the Grey Area: How Expressions of Uncertainty and Overconfidence Affect Language Models*. 2023. arXiv: 2302.13439 [cs.CL]. URL: https://arxiv.org/abs/2302.13439.