Deep Science Publishing
https://doi.org/10.70593/978-93-49910-38-6_2

DeepScience
Open Access Books

Chapter 2: Agentic artificial intelligence in enterprise advisory: From automation to autonomy

2.1 Introduction

AI has gone through distinct phases as it moves from automation to agentic systems. Application focus addresses individual automation use-cases with limited capabilities to extend impact beyond, often with insufficient company-wide buy-ins. Integrating AI into end-to-end processes with better solutions for data accounting and governance deals with more complex cross-departmental workloads. AI's capabilities increase significantly when integrated into organization-wide or ecosystems-transcending multi-agent systems addressing professional domains like finance. However, authority allocation becomes a critical concern at these scales, reversing established human hierarchies with AI recommendations lying at the core of partial autarchies. This paper discusses how AI systems figuratively migrate away from corporate premises and discusses the possible effects. The first goal is to help P with technical-scientific knowledge development for generative AI. Within the enterprise advisory domain, the human-involved information works within the framework of cyber-physical-social eco-systems. The second goal is to help consequential info enterprises with requesting, inquiring, and managing generative AI services to develop artifacts, text documents, models, and software. In P's practical experience, given a conceivable multi-modal capabilities implementation, the outputs of the advisor manager are structured, prescriptive data-transformation queries following the firm's simulation & digital-twin enterprise framework. Against the synthetic pretext of an absurdly simple practical solution, the proposed system is possible regarding a multi-modal LLM/GNN. This results in visible technology transfer as the float-out of multi-modal transformation production systems. The proposed implementation adheres to the principle of unrestricted availability of AI-powered tools, following dislikes towards monopolies of automated kitchen bots. All advisory firms offer basic tasks/proficiency generative data doodles as off-the-shelf tools free of further cost. With

proper P training, integration of the final output with the production system states is partly automated. Local implementation is also possible using GPT variants and stable-diffusion-like systems with open-source release API wrappers. The planned automation proposal is reversible with a ‘back2human’ feature. Finally, within a market political-social research pilot study, generative-aided returns of non-loss-making production decisions are less than a 1% increase from pre-AI levels.

Fig 2.1: Agentic AI in Enterprise Advisory

2.1.1. Background and Significance

The rise of AI as a transformative technology is being leveraged by fast-evolving advisory firms and corporations in disruption-proofing the industry. It will have sweeping repercussions on enterprise services and the practice of enterprise advisory. It proposes an AI maturity model that moves from Automation to Autonomy, providing specific use cases in enterprise services. By surfacing gaps in aggregate AI maturity, the proposed approach aligns enterprise services to stay ahead of the competition.

Research around AI has recently focused on its transformative impact. In response, several technology firms have pivoted from enhancing existing software applications to unleashing citizen-scale co-pilots. Consultancies, along with corporations adopting these

technologies, are grappling with their significance, severity and scope of workplace changes, as well as aspirations and expectations regarding these impacts. Discussions have already surfaced around ‘how to survive the disruption’. In response, new AI advisory practices are being rapidly built by consulting firms, both to safeguard their own firm against upheaval, as well as advising their clients on AI assessment management.

It is contended that AI’s emergence as a transformational technology will have significant repercussions on (1) enterprise services including audit, assurance, tax and consulting; and (2) the systems, methods and governance of enterprise advisory, fundamentally changing how working adults in advisory perform the work of thought leadership, analytical insight and collaborative teamwork. Those enterprise services assurance and advisory firms assessed as least agile in deploying human-centered generative AI to new opportunity spaces risk falling behind. This text attempts to unpack intelligence on AI, its current impacts and aspirations, what it means to become nimble in deploying newly agentic AI in enterprise advisory.

2.2. The Evolution of AI in Business

The evolution of AI refers to the ongoing efforts to develop and improve artificial intelligence technologies (Konsisme, 2023; KPMG, 2024; Aladebumoye, 2025). The first wave started in the early 1950s through symbolic AI, a logical and rule-based approach. Early successes included game-playing computer programs and expert systems. However, these narrow methods proved difficult to scale and were limited to narrow functionalities. The first wave of AI faded in the late 1970s, leading to the so-called "AI winter" and 30 years of decline in the industry. In the early 2000s, the second wave, statistical learning methods, came into play. AI powered up Big Data analytics through probabilistic statistical methods. These approaches demonstrated successes such as image classification. However, due to a lack of representation of knowledge across different domains, the second wave is difficult to adapt into general-use AI. At present, the third wave of AI emerged, also known as deep learning. It has opened new doors to real-world industrial applications and research.

In contrast with the first-wave symbolic AI and the second-wave statistical learning methods, deep learning is much closer to mimicking human cognitive processes and functions. Deep learning networks consist of interconnected neurons. Different processes in a deep learning network resemble various cognitive capabilities of learning from good and bad examples. This approach excels in tasks such as image processing, speech recognition, language understanding, and language generation where the strong drive for automation is observed in the real world.

2.2.1. Historical Context

Agentic AI has existed for some time. Fundamental elements of computer programming, such as loops and since if-else-then statements, allow for the automating of simple processes as well as more complex ones. Actions can be sequenced, restricted, and transitioned based on input variables. Models of perception and cognition exist and can be implemented into computer programming as well. From there, the actions of machines could become more complex, mimicking autonomic insight and inference. Machines can be programmed to not only act, but to reason, analyze, choose, plan, learn, and reflect. Such 'mind-like' systems could be constructed to improve, evolve, and refine their own performance and models of themselves—and out-perform their designers as a result.

This is also a future of frustrations and fears. Dilemmas surrounding existential risk, the future of work, and the distribution of power and agency multiply. Rigid controls could either be imposed, stifling innovation, or inadequately understood systems could escape control and wreak havoc. Mistrust of human actors could also spread. Embezzlement, insider trading, or even nefarious collusion could shred the underpinnings of corporate reputation (or of civilization) as data and decision-capacity become decoupled from human agents.

Still, the future with which to contend bears strong stylistic resemblances to familiar pasts. Many of the practices and politics of traditional enterprise advisory already provision sharply agentic AI within the governance of top-tier corporations and major firms. Data transformation, use case assessment, and integration of predictive analytics are embodied in business analyst as well as enterprise data architect roles throughout companies and CSUs. Note that such forms of internal enterprise advisory are outside the consultancy ecosystem; the "Big Four" consulting firms have all acquired M&A service lines steeped in data engineering situated within functional advisory.

2.2.2. Current Trends

In the business environment the term 'Artificial Intelligence' is in the process of being redefined (Paulson and Partners, 2025; Wolters Kluwer, 2025). It describes a collection of technologies that can discover patterns in electronic data and attempt to simulate scenarios in response, allowing tasks to be undertaken that were previously solely in the realm of human professionals. The tasks of professionals in the business and public service sectors are increasingly becoming amenable to automation, and this is already having a profound effect on their role in corporations. So, with the increasing availability and understanding of better predictive and agentic AI tools, how might the roles of professionals in business advance in the near future? In this paper the rationales for, and

current manifestations of, agentic AI in management, finance, consultancy, accounting and auditing, law and the public services are examined. The likely business structures and agentic AI tools of this future environment are also outlined. Finally, it is proposed that by implementing agentic AI while proactively managing decision rights redistribution, it is possible for firms to 'avoid the dystopia' and bring about a future in which agentic AI enhances rather than diminishes society. The terms 'artificial intelligence' and 'AI' encompass a range of tools, devices and predictions that can be summarized as 'capital-intensive systems for predicting the future'. Primarily in the form of self-learning simulation engines backed by estimates of counterfactuals and man-made 'fitness functions', present-day AI can take actions that are hard for that future to argue with, simultaneously enabling the use of better predictive models and leading to emergent agentic outcomes.

2.3. Understanding Agentic AI

A coherent definition of agency and exploration of human and AI agency rescales how the problems of AI-hyper-automation are characterized, understood, and solved. Understanding human agency and the crucial role in it of agency mechanisms will shape the expectations for AI. AI will grow and reach levels of agency that make agency mechanisms of interest and concern. Specifically, considering and mapping how agentic AI mechanisms resemble or differ from human ones will enable AI systems to be designed to take on responsibilities and act for the desired outcome and to monitor agentic AI decision-making and behavior for how an unexpected outcome/value misalignment could arise. With this view in mind, several key distinctions in agency are explored in more detail, including agency constraints, agency levels, and agency mechanisms. A taxonomy is presented of considerations for mapping agentic AI and agency mechanisms of interest in respect to desired outcomes and value alignment.

By agency, there is reference to all aspects that relate to the capacity to act in the physical world and the cognitive circuitry and mechanisms that enable that capacity. An agent exists in an environment with an action space. All parts of an agency are predictive elements that can be described in terms of constraints, levels, and mechanisms. High-agency entities respond not only to complex stimuli in the environment, but they also reflect and actively re-examine how they perceive the environment, themselves, and the world. High agency enables open-ended response capability, including non-obvious solution-generation and exploration. The more comprehensively and deeply recognized the input and forms of internal states of a network, the more effectively the systems' fundamental outputs will align to expected. There are a range of levels at which agency may be realized and constraints on agency determining at what level agency is effective in a particular environment.

Many of these levels are encompassed or paralleled by concepts within complexity theory. A high-agency AI might progressively formalize and instrument longholders' value constraints, expertly explain critique to normholders, unearth tacit beliefs and contradictions, generate antithetical evidence, or materially change costs associated with inaction. There is an expansion in the interpretation of agency to include many aspects of AI intention, belief, and self-understanding. Embedding agentic AI largely corresponds to performance guarantees for the pre-conditions of functioning of the agentic AI systems as intended.

Fig 2.2: Understanding Agentic AI

2.3.1. Definition and Characteristics

Recent technological advances in AI have allowed for impressive breakthroughs in how machines can mimic human cooperation and be used to automate increasingly complex and cognitively demanding tasks. But this effort to enhance the capabilities and reach of AI goes hand in hand with a wish to leverage this technology's enormous potential for productivity and insight as an enterprise, even a source of competitive differentiation. Me´ dia-chrono envisages the emergence of "agentic" AI across every corner of the enterprise, individually specialized and tasked, cooperating with lower debate co-agents but also with human agents tasked with more supervisory or executive roles.

Agentic AI represents a new class of AI systems that operate with a great deal of autonomy, including autonomously making plans and operating upon them in real

business environments, learning either from interactions with the environment or instead through imitation of human or agent-agent interactions, and, as a result intermediately, changing business processes of their own accord. These AI systems often bring new capabilities in terms of scale, reach, and dimensionality of tasks or breadth of operations due to the differential framing of those tasks or new knowledge representation that enables them to settle previously mathematically intractable problems. Currently, there is great interest in multimodal models that jointly reason over language, vision, 3D, and sometimes occluded objects with complex interactions or novel robotic capabilities and contents from different modalities, for instance, reasoning about language, chemistry, and potential business opportunities.

For a very long time, “intelligent” systems in the accounting, finance, supply chain, tax, operational departments of multinational corporations and banks have been steadily replacing humans in basically data cleaning tasks. Hence, besides empirical evidence, basic principles or rules used, as well as limitations used by such AI systems, are remarkably simple given recent accomplishments in neurosciences, computation, algorithmic complexity, and mathematics. Even though AI systems are incredibly powerful and scalable, enabling new jobs and the coevolution of new infrastructures, processes, activities, and even niches for novel businesses, and are a long-desired aim of founding fathers like Alan Turing and John von Neumann who both postulated the inevitability of intelligent machines.

2.3.2. Comparison with Traditional AI

The way enterprises adopt and leverage Agentic AI is truly new and unique. It builds upon the existing Basement of traditional AI and then iterates and innovates from there. In this context, this new portion of Agentic AI-technologies will be compared with the older Basement of traditional AI, using the AI spectrum put forth in. The Agentic AI indicators discussed above will be mapped onto the AI spectrum and juxtaposed against the field's current Basement of fully on-prem and non-agentic AI technologies. This enables a thorough analysis of how Agentic AI fundamentally reformulates the classic and current understanding of AI.

At the bottom of the AI spectrum is the Basement which encompasses fully non-agentic AI, fundamentally limited to supervised inference or traditional automation. Such non-agentic AI solutions rely on simple and fully passive ask-do response and prediction-response loops, and during implementation their current behavior cannot be improved or changed. Consequently, these non-agentic AI technologies yield risk management concerns like operational control, transparency, accountability, or blind spots, as their decision-making mechanisms cannot be altered as a function of the environment. In the Enterprise advisory context, purely non-agentic AI typically refers to classic,

constraining AI implementations such as data analysis and visualization systems, color-sorting cameras, recommender engines, or text generation and prediction models.

2.4. Applications of Agentic AI in Enterprise Advisory

Agentic AI applications may be found throughout varied Enterprise Advisory work processes: marginalized consultation services for businesses solely focused on enterprise responsibility would necessarily and sufficiently have access to means for bookkeeping and bookkeeping review, which might be based on Agentic AI means adapted to their conditions. Thus, bookkeeping processes could be better managed by the relevant Agentic AI means, thereafter systematically generating reports on the enterprise's conduct options on the counts that are relevant for compliance with constraints on its operation, and updating its noncompliance with the broad margins of these constraints. And since Agentic AI means would certainly evolve in their common management, these AI means could even broker thorough Agentic AI entrants for the knowledge-works for compliance review, if relevant noncompeting Agentic AI means co-ordinarily operated on compliance review for free Agentic AI means credibly accessible to the broader business community. Therefore, to the extent that bookkeeping processes would have attained their discretion, business consideration advice and report preparation could be effectively and continuously fully executed by the relevant, off-the-shelf Agentic AI means.

On the other hand, high-priced, intensive, authoritative enterprise strategy advisory services would not only be variously motivated, ontologically quite capable, and necessarily in unique business situations, but also found in means which would themselves enjoy greatly more discretion in their consideration options. And having unique enterprise positioning, they would necessarily rely solely on enterprise secret means for posts, and therewith necessarily not have their documents fully remotable. On the other hand, no supervision would thus be conceived at all throughout by the AI means. And neither their conclusions nor the actual drafts of the report would be capturable in this way. As desired, generations of new drafts from the Agentic AI might then be executed as an independent professional and dismissive of them. In contrast, the many awareness meetings and the elaborative report drafts delivered throughout the work process would be unrecorded. In particular, the essential meetings deliberated upon the results and new directions of the project as well as the final meetings would be unremovable.

2.4.1. Risk Management

In the past five to ten years, organizations have rapidly started using available AI technologies to solve a broadening range of challenges around operations and customer experience. Improvements in the performance of AI technologies, together with cost reductions, have encouraged systems integrators and other large technology firms to package and offer AI products “off the shelf.” The availability of affordable data has driven many small and medium-sized enterprises to build AI capabilities. As a result, enterprise activity around AI is proliferating significantly with a growing number of AI projects initiated within organizations and with competitive pressures to move quickly. While AI technologies promise to add significant business value, the probabilistic nature of AI results in a class of risks and challenges that are far greater than traditional technologies. Organizations started recognizing these challenges as they reviewed their hiring practices, communicated with clients and customers, underwent technology and service transformations, and mitigated many other organizational risks. Typically, organizations have established internal audit functions and risk oversight committees accountable for organizations-wide identification and management of operational, compliance, reputational, strategic, and financial risks. With enterprise AI activity proliferating, organizations must also ensure that AI-related risks are identified, assessed, mitigated, tracked, and reviewed, and that the incidence of problematic AI is reduced. Five major risks associated with the adoption and growth of AI within organizations: (1) regulatory and compliance, (2) reputational, (3) user trust, (4) financial, and (5) societal. Five major challenges organizations face in governing and managing AI risk: (1) AI risk awareness, (2) AI risk valuation, (3) AI risk identification, (4) AI risk mitigation & control, and (5) AI risk culture, monitoring & assurance. AI technologies can pose a significant risk to organizations if not used responsibly and must be understood, reviewed, and overseen at the organizational level.

2.4.2. Strategic Decision Making

The quality of products and services, as well as their acceptance on the market, is often determined by strategic decisions made early in the management process. The degree of independence and transparency of these decisions is of utmost importance to protect the sensitive business knowledge proposed by a corporation. However, the increasing global competition requires a marked improvement of these competencies, shifting recent methods and approaches towards automated systems. Active research is currently being carried out in this domain. The purpose of this study is to investigate potential applications, requirements and implications of Autonomous Collaborating Agents in strategic enterprise decision making problems. The proposed approach consists of two components: a knowledge-based system to process high-level knowledge about the

problem domain and the intelligent decision maker's preferences and concerns, and an application-specific agent architecture to fully automate formal model generation and agent orchestration based on the knowledge processed earlier. At present, and in the near future, many organizations are confronted with a growing number of decisions requiring their attention. Even if some low-level decisions are automatically processed, there is still essential decision knowledge left to maintain and automate more intelligent business solutions. Furthermore, decision outcomes are equally sensitive, and the enterprise landscape equally complicated. Attention must therefore not only apply to data privacy, but also to the independence of this extremely sensitive business knowledge. The investigation reveals various intriguing possibilities for further research and has implications for the agent community at large. The architecture is currently under implementation. Based on a commonly accepted formal language for the definition of decision problems and corresponding languages for algorithms, agents will be implemented and connected. Furthermore, platforms for the execution of decision problems as well as for knowledge processing will be developed. Finally, a number of application cases in relevant domains will be analyzed and illustrated with a simplified proof-of-concept implementation.

2.4.3. Operational Efficiency

Enterprise knowledge is growing massively and is already a huge asset for companies across the globe. However, it is fragmented across large and disparate repositories, which makes it difficult to leverage. Creating one user-friendly knowledge base for employees starts with a description of the current situation, followed by a design proposal for a solution and a choice of paths through complex evaluations, chunks of data, and investments. Teams of people with complementary skills work together to collect the knowledge and convert, refine, and curate it, before deploying it in an Enterprise Search system with which employees can query and browse.

First issues arise when the digital transition is initiated, such as consolidating the repositories estimated to keep over fifty different, mostly text-based but also spreadsheets and diagrams archives intact, migrating them into local indexes, quite some textual data will be lost in conversion. Displaying summary metadata that will eventually grow to thousands of entries in chapters in a book-like portal is a challenge. Meanwhile, budgets are limited, requiring a make-or-buy decision, and companies usually have to choose from another vendor in the development phase.

Also, there are legislative issues governing the coping of competitive knowledge. When conducting due diligence on a prospective vendor, it is impossible to check anything without knowing the vendor's strategy. Therefore, an encapsulation solution that restricts refinement or direct replication of the indexed knowledge while allowing querying and

browsing is required. Moreover, large language models are notoriously bad at reframing questions in a factual manner and producing accurate answers.

2.5. Benefits of Implementing Agentic AI

The excitement about Agentic AI in enterprise advisory emerges from its promise of changing how knowledge-intensive work is performed. In particular, these systems represent a conceptual leap from automation to autonomy. The advent of advanced generative models is giving rise to a new economy of enterprise applications. Companies already see the potential of automating mundane tasks such as email composition and calendar scheduling thanks to large language models. Beyond these narrowly-skilled assistants, more powerful models are being unleashed, marking the dawn of agentic AI. These semi-autonomous knowledge agents are capable of reasoning and researching independently, operating and coordinating over existing digital environments—much like human agents. It holds the promise of changing entirely how enterprise knowledge is created or structured: from unique, semi-structured human-curated databases to self-operational digital knowledge systems populated with up-to-date, rich knowledge, interoperating over heterogeneous sources of information and knowledge on their own.

Even though the underlying models, techniques, and infrastructure for agentic AI are rapidly evolving and becoming democratized, so is its low-hanging fruit: agentic assistants capable of detailed reasoning and retrieval over enterprise data. Nevertheless, to leverage the full potential, organizations need orientation on how to pursue enterprise agents at scale, design for reliability and alignment, and address adoption readiness in their workforce. Decompiling enterprise agent applications into organizational-level modules, and developing frameworks, identifiers, and facets to characterize each application's capabilities and development priorities are crucial. Agentic systems are not a silver bullet. The impact of agentic applications on the advisory workforce depends on various context-specific factors, including design decisions, task environment, training datasets, and cultures. If left unchecked, this frontier technology poses serious reputational risks, including biases, privacy loopholes, indecorous behavior, and misinformation.

2.5.1. Enhanced Decision-Making

Recent advances in Artificial Intelligence (AI) lead to an increase of automated decision-making. Driven by AI-based techniques such as machine learning, AI is becoming increasingly capable of completing complex, cognitive tasks previously thought necessary only for humans. Hence, decisions can be classified into unstructured, semi-structured, and structured decisions. Traditionally, automated decision-making was

applied to structured problems, whereas decision support systems enhanced decision-making for unstructured problems. However, advances in AI now increasingly enable the automation of even complex cognitive tasks, such as driving a car. Hence, AI has the potential to address even semi-structured and unstructured decisions, such as in the context of loan approval and recruitment choices.

Fig : Agentic AI The Future of Business Process Automation

However, while many AI applications succeed in automating simple cognitive tasks, the skill of knowledge workers is a major problem of economic and social dimension: Knowledge workers should train, challenge and evolve AI. Knowledge workers create the labels for AI that is the basis for its training. They should be able to challenge the AI's recommendations regarding performance and ethical concerns, such as bias, explainability and fairness. Further, they should be able to provide additional and complementary information, such as real-time developments, that is not part of AI's input data. Negative impacts occur at low automation levels: Instead of positively influencing decision-making via appropriate automation, automation can have a range of negative impacts, reducing positive features of human decision-making such as engagement and due diligence.

Many state-of-the-art decision support systems include already automated, AI-based recommendations. Hence, they are subject to negative impacts mentioned before, such as automation bias and deskilling. Existing inquiries either focus exclusively on decision support systems with automated decision-making or disregard these systems, which represents a significant gap in the literature. This research aims to provide a systematic inquiry into decision support systems, focusing on the way in which AI can be integrated into decision support systems to preserve the benefits of automation while counteracting

its side-effects. In doing so, a new class of decision support systems, intelligent decision assistance systems, will be conceptualized that reduce automation-induced disadvantages while preserving decision support levels.

2.5.2. Cost Reduction

Major advisory firms rely on extensive human involvement to manage client projects. Writing laborious statements of work, creating project plans, and identifying qualified consultants are all expensive manual processes that influence client outcomes. At large firms with thousands of employees, automation opportunities abound. Newer boutique firms, which generally have fewer than 200 total employees, modes of organizational activity are less repetitively structured, are more bespoke, and thus less amenable to automation via procedural AI. However, as these firms evolve, the emergence of highly experienced employees makes generative AI increasingly useful in reducing costs for routine administrative and staffing-related tasks.

With a superior understanding of the domain and capable of much more advanced reasoning, widely available foundation models provide transformative new avenues for automating enterprise tasks addressing more cognitively demanding question types. Some elaboration is warranted regarding the operational requirements of this innovative approach to automating enterprise tasks. First, these AI agents represent a major change in technological capability compared to traditional automation approaches. As a result, human participation and control of specifications, instructions, and oversight will be very different from previous generations of automation. Such profound changes often raise apprehension, skepticism, or opposition. Nevertheless, the capabilities of agentic AI are expected to radically change the nature of enterprise consulting and thus provide competitive first mover advantages and opportunities irreplicable via human analysts.

2.5.3. Increased Agility

Enterprises are becoming increasingly pressured to support and develop Agile concepts, capabilities, and methodologies. Agility is not a software tool or methodology; it is a mindset, a state of mind, a way of thinking, operating, and behaving. Businesses implementing or adopting Agile concepts, methodologies, or ways of working need to be aware that agility is brittle and requires agile tools, leaders, people, and cultures. If these prerequisites of agility are not met, the benefits will be short-lived and fraught with the risk of agony, timidity, and opposed outcomes.

The Agile Manifesto emerged in February 2001 and has become as significant to the modern practice of software development as the Agile Development Methodologies and

the concepts surrounding iteration, incremental development, and Agile development. Agility in enterprise software development is a term that has become widespread but does not yet have a single clear, widely adopted definition. Indeed, agility is often considered a buzzword or a keyword and is easily confused with the keyword agility, which refers to the velocity, nimbleness, and quickness of a business.

An enterprise is a complex socio-technical system composed of interacting agents who design, develop, install, maintain, upgrade, or decommission computerized systems, along with the underlying system parts plus other components such as people, processes, and policies. This wider enterprise context also includes the system environment(s) or domain(s) outside the enterprise but does have a specific domain. The creation, growth, and evolution of the computer-based enterprise system(s) in varied risk situations can be informal. All this does require enterprise agility within the boundaries of soundness and safety, but the literature does indicate that enterprise agility is a slippery concept that can mean different things in different contexts.

2.6. Conclusion

This paper has examined and assessed the socio-technical aspects of the advent of AI in enterprise advisory, focusing on the shift in agency from humans and hyper automation to AGI and autonomy. There are many definitions of AI, but this paper has focused on one special class of AI, in particular, on how it is becoming increasingly present in advisory roles. As associated control and regulation challenges grow in importance, the paper has coined the term "Agentic AI." This paper has used historical analysis and examination of two provocations to explain and assess the characteristics and utility of Agentic AI in an enterprise advisory context.

Paradigmatic examples of Agentic AI advisory systems have been presented. Powered by generative AI and natural language processing advances, Agentic AI represents a major shift in enterprise advisory systems towards autonomy, or, merely that autonomous intelligence replaces human intelligence as the basis for decision-making in enterprise advice. The "recommendations first" paradigm of Agentic AI-based advisory seeks to provide an autonomous search strategy to minimize the resolution time of queries. The engagement with Agentic AI advisory brings about new considerations and implications for the wisdom, robustness, and trustworthiness of the advice provided.

AI-based systems are, however, not solely responsible for mistrust and misinformation. The repeated reports of AI agents race conditions and "AI misapplications gone wrong" illustrate the risk of overly participative interaction and reduced explainability of intelligent systems. In conclusion, while predominantly within acceptable boundaries, there are a multitude of considerations that warrant due diligence.

In summary, this paper has defined and characterized Agentic AI in enterprise advisory; and explored implications for organizations, owners, and relevant actors and stakeholders. Granted the complexity of Agentic AI, it is important to keep expectations realistic and recognize ambiguity and multiple characteristics of phenomena like IT-enabled intelligence, also including limitations and risks.

2.6.1. Future Trends

It is often claimed that AI will revolutionise the economy and entire industries in some fundamental way, and that this transformation is just around the corner. However, these bold predictions come with carefully drawn qualifications that describe the limits of current technology or embed it within a complex sociotechnical landscape. AI systems are capable of learning from experience and making predictions and decisions. They do not act with genuine agency but perform probabilistic inference under conditions of uncertainty. As analytical tools or decision support systems, they are assisting and augmenting human intelligence in many domains. The success of these systems in the workplace depends not only on compatibility with managerial strategies or the skill of the data science team, but also on trade-offs in time, cost, and risk, on the data ecosystem, organisational infrastructure and culture, and the design of the human–computer interface. These factors co-evolve with AI capabilities, at different speeds, leading to hybrid intelligence or intelligence augmentation, rather than a wholesale replacement of human roles. This tends to be optimistic about the speed and scale of transformation and its global equity. AI fills labour shortages and improves productivity growth; the best and brightest are flocking to this field. Development has slowed in the last few years because of a lack of compute power and data. Recent innovations will not spark a new gold rush; in-house tech capabilities will grow, closing some of the gap between large and small firms and rich and poor countries. The apparent stability over recent decades will return, as firms adapt their operations to the mitigation of bias and explainability. Companies will jostle for positions in newly emerging market structures, opening new avenues for research on matters like the causes of quasi-monopoly. Cyber security will remain an issue, with deployment limited in sensitive situations where lives hang in the balance.

References

Aladebumoye, T. (2025). *Integrating AI-Driven Tax Technology into Business Strategy*. International Journal of Research and Review, 12(1), 224–231. IJRR Journal

Konsisme (2023). *Embracing the Power of AI: Transforming Corporate Taxation*. Retrieved from https://www.konsise.com/embracing-the-power-of-ai-transforming-corporate-taxation/ konsise.com

KPMG. (2024). *AI in Tax*. Retrieved from https://kpmg.com/xx/en/what-we-do/services/ai/ai-services/ai-in-tax.html KPMG

Paulson and Partners. (2025). *AI in Tax Advisory: A Revolution in the Making*. Retrieved from https://paulsonandpartners.com/ai-in-tax-advisory-a-revolution-in-the-making/ Paulson and Partners –

Wolters Kluwer. (2025). *Artificial Intelligence (AI) in Tax and Accounting*. Retrieved from https://www.wolterskluwer.com/en/know/artificial-intelligence-tax-accounting