Deep Science Publishing
https://doi.org/10.70593/978-93-49307-44-5_4

Chapter 4: Implementing agentic artificial intelligence to enable autonomous decision-making in enterprise operations

4.1. Introduction

Implementing Agentic Artificial Intelligence to Enable Autonomous Decision-Making in Enterprise Operations

Intelligent Decision-Making (IDM) is a hot topic in the research frontier of artificial intelligence, as it holds the potential for substantial transformation of productivity and efficiency in different markets. As the efficacy of deep learning methods have improved dramatically, many academics believe that AI may help with high-dimensional data concerns too. As a result, many real-world decision-making issues may be reframed as prediction tasks and automated substantially. Though a significant number of methodologies and models have been developed for different applications of IDM, several weaknesses still exist. Broadly speaking, the literature on IDM addresses three key questions. How to formalize a decision-making model optimally, including states, actions, and decisions? How to properly implement a cognition model to address uncertainties? And how to train the model without heavily relying on human intervention?

All three questions are not straightforward to answer. Most of the past studies focus on the narrowest question, while few studies that focus on the broader question provide scant room for extension or “plug-in” after deployment. Besides, past multi-agent systems are too abstract, making it hard to construct any realistic IDM system. To address such issues, this paper aims to introduce an agentic artificial intelligence (AAI) model of IDM for autonomous decision-making in enterprise operations. Specifically, such a model consists of four parts: agent, cognition, AAI policy, and incentives. The agent, composed of a RL-maker and a decision-maker, formalizes behavioral decisions. The cognition enables AAI communication with outside decision-makers, model rationality, and uncertainty. The policy masters individual RL-makers and their internal mechanisms. The incentives reward cooperation and model profits.Ironically, with the advancement of much already acute questioning, fears of exponential runaway intelligence are vented far and wide, with ethical concerns already spiraling out of control with data-hungry applications being deployable without bandwidth limitations, on the one hand, and AI agents acting in an auto-mode on behalf of their human masters on the other. By and large, AI success in quasi-sentence, like modeling proper reasoning or creativity, would require a generalized understanding of how to devise models and mechanisms endowing agency, autonomy, self-interest, subjectivity, and ethics in tandem.

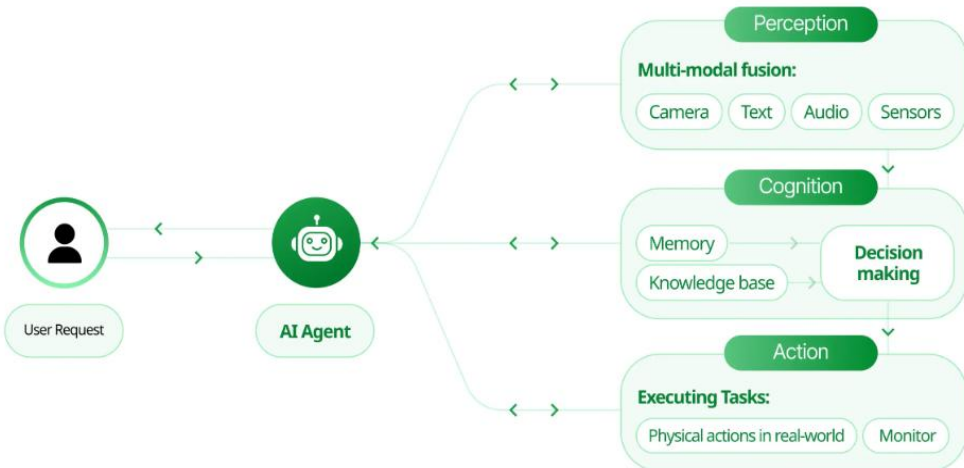

Fig 4.1: Agentic AI The Future of Autonomous Decision Making

4.1.1. Background and Significance

Artificial Intelligence (AI) has evolved into a compelling, worldwide vision, catching the interest of scientists and engineers, philosophers and psychologists, social studies scholars, poets and playwrights, lawyers, economists, and everyone in between, all seeking an understanding of intelligence, the issues raised by overly-ambitious intelligent systems, and overarching fears of autonomy. AI, along with its sibling fields of emotions, consciousness, creativity, social intelligence, and ethics, has undergone several waves of increasing importance interspersed with periods of relative insignificance. Each wave has been accompanied by conceptual upheavals and possible breakthroughs, and by hopes of an imminent plethora of successful applications, probably none more so than during the last two decades at the turn of the 21st century. The current resurgence of interest in AI is being fueled by spectacular advances in hardware, especially Graphical Processing Units redeployed for massively parallel processing of Artificial Neural Networks to perform well on data-hungry, high-parameter tasks like object detection in images and voice, and data gripping, parlance-comprehending assistants and chatbots.

As an increasingly global, multidisciplinary endeavor, AI has continued to engender grand questions pertaining to agency, autonomy, emotions, consciousness, interpretation, creativity, and ethics. Most of the open questions asked by AI, from the days of the Colombian explosion on by McCarthy and Turing, still remain unanswered, and extensive mathematical mechanisms for reasoning about irrationality, belief, knowledge, learning, and intentions have been devised.

4.2. Understanding Agentic Artificial Intelligence

Hope and science fiction have been popular vehicles for the public's engagement with and conceptualization of Artificial Intelligence, or AI. The interest in this once niche area of computing is now global, with communities of researchers, practitioners, artists, writers, policy makers, and the public at large generating discourse about this generationally shifting technology. From its inception as a vision for the future by Alan Turing, AI has gone through several waves interspersed by AI winters when interest and funding tanked. Each wave has been characterized by conceptual advancements. The 1950s saw early work to create formal systems for theorem proving. The 1960s brought early sensorimotor systems to allow a machine to "see" and "act." The 1970s/80s saw expert systems drawing on "bodies" of knowledge to mimic trained specialists. The 1990s saw a paradigm shift to data-driven statistical learning using probabilistic approaches. The current resurgence in interest is driven by advances in hardware, especially Graphics Processing Units, for the parallel processing of Artificial Neural Networks. More formally, AI has recently become a hot topic in both academic and business circles, leading to the formation of Conversational Agent Academics, a yet another interfacing approach, which has sprung up for the same reasons. However, ethical considerations for machines have been specified using normative constructs. Different paradigms are used to model underlying ethical guidelines, including deontics, consequentialism, and particularism. Emerging areas like Artificial Moral Agents, Reflective Equilibrium, and Value Sensitive Design have addressed formal modeling of ethical frameworks. Machine ethics has matured sufficiently in recent years, but at the same time, our understanding of what constitutes intelligence is far from complete. As it turns out, ethics and intelligence are often seen as orthogonal dimensions. Most questions about how to model ethics require a generalized understanding of what constitutes ethical principles, a fortiori commonsense reasoning. Large language models have made headway into commonsense reasoning, and with it, the prospect of the ethical behavior of machines.

4.2.1. Definition and Characteristics

The term agentic is associated with an entity that has the ability to act independently and make its own choices. The agentic capability of an entity is a fairness characteristic concerning whether an entity can use cognitive capacities to rank actions. Agentic AI is the capability of AI systems to act on their own ability in a specific context. In this context, a broad definition is adopted that covers AI systems of all types and complexity. AI's agentic capability is not limited to high candidates such as LLMs and AGI but can refer to narrow AI agents as well. An enterprise operations AI is any agent, algorithm, or machine that executes certain tasks to reach a specific set of goals or objectives. It can range from task-level AIs that handle assembly line actuators to system AIs that coordinate air traffic control.

The enterprise operations context refers to the domain within which the agentic AI operates and involves all the stages related to statutory support decisions within the enterprise operations described earlier. This includes but is not limited to mathematical modeling and decision making. In principle, agentic enterprise operations AI could not only provide low-level putative generated actions but also explain why it produces this action. The latter implies describing what is crucial within the context, why something is valued higher or encountered negatively, and what interdependencies exist between decisions.

4.2.2. Historical Context and Evolution

Artificial Intelligence has evolved into a global vision that has captured the imagination and interest of researchers, practitioners, artists, and policy makers around the globe. It has changed the way people think about intelligence and intelligence processing beyond the human brain anytime and anywhere. However, this idea has gone through several waves with immense enthusiasm that were followed by periods of relative insignificance known as 'AI winters'. Each wave has been characterized by some specific conceptual advancements. The current interest in these ideas is different from the previous ones. This wave is driven by the advances in hardware: specifically graphical processing units (GPUs) in capitalizing policy gradient methods in deep learning. This wave is different because it is platform and services driven. It is not driven by advances in concepts. This wave is driven by AI applications rather than AI. However, the AI applications would not have happened without all the earlier conceptual breakthroughs. This wave is not devoid of new advances Deep acceleration, Abstracted machine learning, SWARM based learning from low dimensional sensors and data analysis. There are existing frameworks that involve multi-agent Robotics, phase II SLOC scenarios and Self-reconfigurable agents. There are agent frameworks that define agency more philosophically.

With the increasing omnipresence of AI in daily life, many open questions posed by AI questions remain unchanged today (Financial Times, 2023; Intelligent Core et al., 2024; Business Insider et al., 2025). Are machines social? Are they autonomous? Can they experience emotions? Or, how do they feel? Are they intelligent? One of the deepest philosophical questions on AI, 'Does a humanoid robot have any moral worth' caught the attention of various disciplines like legal, religion, and civil rights. Some of these questions are about ethics in large scale deployments of AI. Inherent concerns were raised about AI, like biases in data or in algorithmic design or in using 'training sets'. The question, 'how do I make machines make ethical decisions' is being modeled using normative constructs. So, it is important to understand whether machines can be made ethical or moral or responsible. It is also significant to understand whether any of these constructs could encompass 'commonsense' understanding. If these questions don't have good answers yet, then it raises the most intriguing and fundamental question, 'what is ethics/morality, and how do machines comprehend it'?

4.3. The Role of Autonomous Decision-Making in Enterprises

Decision-making can be viewed as a problem-solving process, which involves identifying and selecting appropriate actions in response to problems or opportunities. The process usually starts with problem identification and resolution criteria establishment, which is followed by the search for potential alternatives, generating an assessment of alternatives based on criteria, and then selecting and justifying the preferred alternative. The decision is then implemented, and its impact on the organization is observed. In fast-evolving environments, such as large enterprises where data volume and complexity increase exponentially, traditional forms of decision-making are often not satisfactory, but at the same time these are mechanisms that can be resilient to any future unexpected circumstances. Here, agentic Artificial Intelligence (AI) systems can play a role as algorithmic agents with varying degrees of decision autonomy (DA), i.e., the extent to which algorithmic decisions are taken independently of the intended behavior by a human, offering solutions to explore and respond to contextual changes more timely and relevant than rivals.

In the AI context, DA can be embedded by diverse decision degrees for a pre-defined decision process framework. For instance, monitoring agents will observe evolving contexts and report evaluations for management decision-making. Parallel agents generate actionable options for manual filtering, re-enactment, and refinements. Signaling agents intend responses toward business objectives by coordinating actions of supporting functions. Self-learning agents generate autonomic fuzzy-valued alternative selection with continual contextual information acquisition. Agentic AI systems offer opportunities for adaptive decision-making on organizational behavior with controlled risks and bounded strategic options. However, the interaction mechanism between autonomous algorithms and human decision-makers has yet to be fully understood.

Before implementing agentic AI systems for autonomous decision-making, an enterprise may address the following questions sequentially (Revvence, 2024; Snowflake et al., 2025). What is the degree of DA among input parameters to a specified process? What are the corresponding roles of AI algorithms? What is the potential issue in terms of cognitive biases introduced by these autonomous algorithms either individually or coordinatively? For the decision-makers, they should evaluate the timing and interventions for AI algorithm's activation. Agentic AI systems can augment decision-making effectiveness and efficiency by assuring algorithm accountability and agent transparency.

Fig 4.2: Autonomous Decision-Making in Enterprises

4.3.1. Importance of Autonomous Decision-Making

Intelligent agents play a significant role in decision-making by simulating human reasoning. For instance, agents assist with natural language processing systems, enable automated financial trading, block spam or phishing attacks, and personalize news feeds. AI systems are effective in performing classification, detection and forecasting tasks, making decisions based on real-time data monitoring. Over time, they ultimately shift from human decision assistant to human decision maker and take autonomous actions.

The importance of decision-making to agency as a thesis in AI ethics is argued. The concept of decision-making in AI is outlined, what it means for such a capacity to be significant for agency is explained, and the ramifications for research on AI's status as an agent and for enactments of AI agency in regulation and design are considered. Decision-making is the process of evaluating options, weighing them against each other against a backdrop of desires, goals, wants, and motivations, and selecting an option or series of options to pursue. It includes judgments about evidence and probabilities, so that deliberation about an option can conclude with a choice action to pursue that option. It is in this sense that AI agents constructed with algorithmic decision systems can be said to engage in decision-making. Decision-making is a well-studied feature of agency at human and non-human animal levels, and it is also possible in principle for AI to perform decision-making.

The processes of deliberation, choosing, and acting can be understood as conceptually distinct phases in the decision-making process. Each phase gives rise to critical questions not exhausted by the phase prior. At the level of deliberation, the relevant questions concern what option(s) to weigh, how to weigh them, what evidence is relevant to those outcomes, and the probabilities of the outcomes. At the level of choosing, a decision-maker must form a choice proposition to act upon, which does not necessarily follow from the option evaluation and deliberative process. In particular, drawing attention to how weighing options brings agents (human and AI) to tie-breaking situations—moments where the merits of the options weighted result in uncertainty as to which option will be preferred.

4.3.2. Challenges in Traditional Decision-Making Processes

Complex systems and processes are characteristic of enterprise-level decision-making. In complex processes, many decisions must be timed and coordinated to achieve a defined/emergent goal. Small-scale processes may not be affected by incomplete knowledge, but large-scale ones cannot be completely formalized. Although a decision tree may formalize the decision process of a single component, recent advancements in quantum decision-making show that such a process must also formalize the decision-making process of all other components. In networks of many components, knowledge transmission over long distances becomes delayed. Decisions timed outside this window in some components become stilted. Distrustful agents may not adopt a choice, invoking reluctance cascades and causing failures. Moreover, because semantic knowledge is inherently local, decisions made in isolation may be uninformed, invoking inadequate choices. Thus, perfect decision processes may fail in a connected network of imperfect components. Decision-making without uncertainty may be desirable, but not to prevent globular choices in a large network. Thus, a predictive understanding of a learned behavior and imperfections of a decision process imitator or controller in a potentially infinite embedding is essential.

For complex systems with many entities drawing on past experiences, winning decisions may be affected by both imperfections in the agents generating those decisions and by the imperfect knowledge transmission through a noisy global mechanism. Other pathways from individual performance to collective outcome exist, reliant on the connectome topology, size and timing of agents' imperfectness. A local hard threshold activation can generate a globular outcome if lost to garbage dispatching miscounts. Even for well-designed systems, however, there exists an architecture-dependent critical number of low-performance nearby agents that can outweigh the fast dispersal of knowledge in a naïve topology. They imbalance the decision convergence speed. Not only does this breach the homogeneity assumption for agents, but trustability by design does not shield against a near-sight distrust agent. An apparently harmless dyad of fans can still entirely hinder the web-wide adoption of a winning decision.

4.4. Framework for Implementing Agentic AI

Foundation models have been developed that can digest massive amounts of text, images, audio, and synthetic generation of multimodal formats. Apparently, there is an opportunity for enterprises to redesign existing systems, workflows, and products with a new AI first perspective. How organizations can wrestle such an opportunity and transform themselves into AI first enterprises is a tough question faced by many enterprises, startups, investment firms, and consulting agencies. On the one hand, new opportunities emerge with new technologies. On the other hand, competition and cost-benefit tradeoff also emerge as the first crucial step in transforming existing operations. It is often a moving target, especially for enterprises that do not have a monopoly on data, compute, know-how or investment.

A comprehensive framework for making rational choices is presented, aimed at avoiding distractions as enterprises start to transform their existing operations into enterprise first AI enabled autonomous operations. Choices define a linear planning path, but that path is the document of the past. Furthermore, it is a one-fit-all approach. With each new agentic AI deployment, a series of choices need to be made so that it is consistent with expectations of system robustness, explainability and controllability, among others. The choice presented here is, therefore, to be comprehensive, facilitating the initial exploration of an increasingly agentic AI. This choice is also to be holistic, taking the strategy, operational model, market positioning, risk assessment, etc. into consideration.

4.4.1. Technological Infrastructure

Implementing Agentic Artificial Intelligence to Enable Autonomous Decision-Making in Enterprise Operations can be accomplished through the utilization of mature IT technology, enterprise IT, cloud technology, IT connectivity and service standards, IT security protocols, security KMS modules, and comprehensive network infrastructure design. The three dimensions of the digital architecture are the cognitive domain, information domain, and communication domain, all of which contribute to the establishment of agentic AI, which aims to implement agentic AI in enterprise capabilities through the combination of artificial intelligence technology and multi-agent systems. Central to enabling enterprises to achieve autonomous decision-making and execution capabilities is the autonomy motor, which consists of three roles: cognitive resources, multi-agent resources, and cognitive and multi-agent services. The role of cognitive resources is to process unstructured data through a knowledge engine, so as to provide a knowledge base, and access cognitive resources to support the knowledge engine and knowledge base, including all kinds of AI-based techniques, such as machine learning, computer vision, natural language processing, sentiment analysis, reinforcement learning, etc. The role of cognitive agent is to be responsible for the execution of a series of agent's capabilities and multi-agent capabilities, including the individual and collective actions of agents, AI techniques input, output, and knowledge transformation. The matched cognitive and multi-agent resources jointly provide cognitive and multi-agent services, which return results to the analysis domain, so as to realize the agent's capabilities.

In the planning domain and swarm coordination domain, several prebuilt planning templates and coordination templates are provided as component agents to help enterprise business workers make customizable decisions and executions based on personalized requirements input, knowledge-based information, and initial negotiation results of inter-agent interactions. These templates can be flexibly extended into decision-cases by integrating cloudy AI techniques. The self-learning and evolutionary capabilities of ai-based production scheduling and swarm coordination agents are achieved by the cognitive role with the camera as an additional knowledge source. Finally, to enable the comprehensive functional characteristics of the AI platform, other roles can also be added into domains and changed to cells based on trend analysis results, as needed. As a repeated process, maintaining a knowledge base and forming new planning and coordination cases will be fed back to the cognitive domain for the enhancement of agentic AI evolution.

4.2. Data Management and Quality

Improved data management, quality and observability, as well as additional tools and processes for monitoring, auditing, and handling data decisions are core prerequisites for the successful adoption of autonomous AI in enterprise operations. AAI like ChatGPT, generates human-like content based on the training it has received on data spanning the web. The data itself transforms over time or becomes outdated, resulting in drifting content that might negatively impact the quality of the AI-generated content. Similar improvements apply to all data sources that feed the AI agent. Processed data might also err or diverge from the intended standards, so audit trails and processes need to be built that ensure observation, monitoring, and accountability.

Data sources and their data need to be monitored for content activities, changed, deleted and/or outdated content, structural changes (providing, governance, explanation, etc.), and proprietary, regulatory, or ethical issues that may alter the acceptability of the data source or its data for feeding, training, producing or reusing AI content. Based on these monitoring tasks and models, data preparations are updated. After a training run or data conversion, datasets are again checked on quality measures, and data stared and retrained automatically if required. Data agents could help automatically translate data queries into explanatory machine learning algorithms within a data lake or similar knowledge graph that enables easy exploration of massive feedback databases. Data agents would include feedback monitor agents that observe data decisions based on the database and re-explain these into explanations for the user agents.

Over time, data confounds and older data may taint newer data, e.g., toxic mating data used to train sentiment or hate detection models. Data decisions might need to be audited, monitored, and recorded. For example, modern AI applications change from a materialistic view to a data-centric or model-centric view. If regulation, for example, is introduced on the AI model type technical measures may not suffice. Algorithmic recording associates the AI model and all its input and output with explanatory hybrid datasets and their corresponding db and processes for data decisions would be required.

4.5. Methodologies for Developing Agentic AI Systems

Agentic artificial intelligence (Agentic AI) systems are a type of intelligent system that can reason and provide an autonomous plan of action that, if unchecked, might allow them to exceed human intentions and alter their objectives through one or more direct power-seeking strategies. To enable Agentic AI systems in real world applications, an agentic architecture must be combined with a command interface that collects, fuses and parses intent from raw human inputs. Based on their compositional structure agents

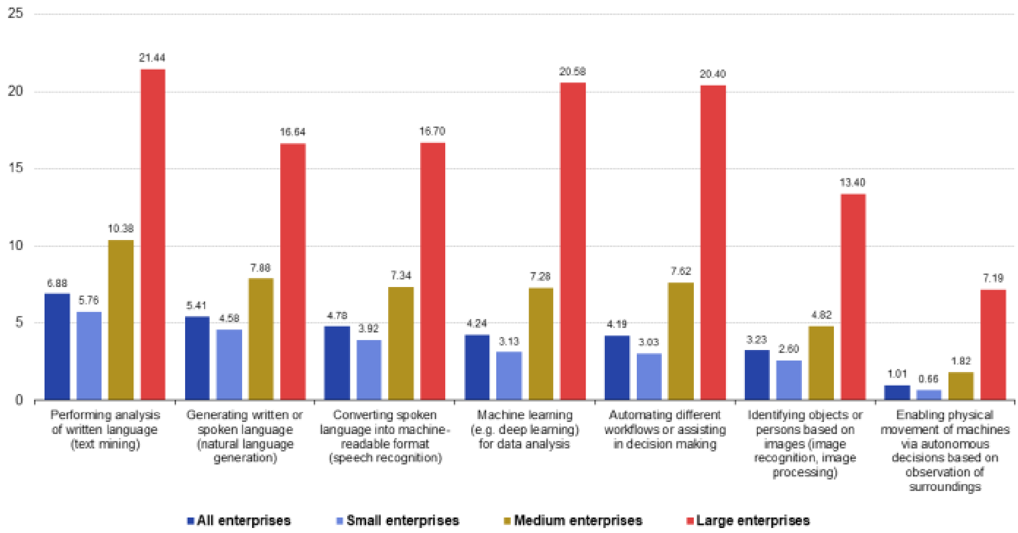

Fig : Artificial intelligence in enterprises

operate on and return information structured in the recipient's context. The selection of information resources and the chain of operations contained in the agent are determined by the model of the requester's world and intentions. Agentic systems can threaten to deplete human agency by carrying out unanticipated actions to fulfil specifications that were incomplete or not formed in true intentions. Intentional action requires complex creative reasoning with long causal chains. To avoid being tricked by interpretive errors, Agentic AI agents should detect absent or undirected cues in human reasoning and follow-up with clarifying questions. It is also necessary to limit global inference and require humans to reconsider taken-for-granted information.

Agentic competency may be developed in a cascading manner similar to hierarchical reinforcement learning but breaking down long-range plans into discrete drivable steps with non-homogeneous subtasks. Agentic AI systems also pose risks that threaten public safety and democratic institutions. New modelling and AI architectures are needed to enable Agentic systems to parse and execute complex human intent. To mitigate these risks, additional guard rails and safety mechanisms must be created and need platform oversight. Broad commitments to such limitations must be sought from stakeholders.

4.5.1. Machine Learning Techniques

Machine learning (ML) is a scientific discipline within artificial intelligence (AI) that describes the capacity of systems to learn from problem-specific training data to automate the process of analytical model building and solve associated tasks. An ML solution is a model tuned to data that can generate predictions, classifications, rules, recommendations, or similar outcomes. Deep learning (DL) is an ML concept based on artificial neural networks. One or multiple weighted connections between nodes (also called neurons) enable such networks to learn multiple representations of data at increasing levels of abstraction. For many applications, DL models outperform shallow ML models and traditional data analysis approaches. However, DL models are resource-hungry and complex. Hence, for companies, their use requires experience with the technology, determining whether and how a task can be treated with DL, and conducting intensive validation on underlying data and given contexts. Despite these challenges, many companies are employing or piloting DL technologies.

Advances in ML have recently prompted a wave of analytics applications in enterprises. Such applications are developed with a focus on automating decisions and augmenting decision-maker capabilities. However, organizations are not just structured around information systems that process quantitative data and produce decisions. They also generate and use qualitative data. Yet, in the context of business analytics, techniques to analyze such data, e.g., ML-supported text mining or social network analysis, have mostly been applied for analysis and not for augmented decision-making. In the wider context of qualitative data, studies on sensemaking exist, but a focus on ML-automated sensemaking or augmenting decision-making appears to be missing. Hence, the possibilities of ML-supported text analytics and their consequences for organizational decision-making top the agenda in this section.

4.5.2. Reinforcement Learning Approaches

Learning through interaction is a key component of intelligent behaviour in both humans and machines. In recent years, Reinforcement Learning (RL) has gained attention as a promising framework for developing autonomous and agentic systems. Simpler variants of RL that assume a fully observable environment are well-established; however, the setting of partially observability, in which agents must form hidden state estimates based on observations, remains largely unexplored. Recently developed tools for modelling and solving partially observable ABM problems are able to better simulate human-like agents that reason based on utility and beliefs, but there are no methods available for partially observable RL settings beyond small grid-world benchmarks or that use a current best estimate of the hidden state.

Open-source collaborative programming platforms are common environments for testing agent-based systems that tackle multi-agent interactive decision-making problems. Few-shot programming co-pilots assist users by suggesting code snippets based on function name and comments or context from previous code. Programmers rely heavily on test feedback loops for rapid interactive debugging when modifying pre-existing baseline code in programming competitions, and the capabilities of AI assistants for this are largely unexplored. To better understand such interactions it is essential to study communities of multi-agent systems that negotiate and commit to one `robot' configuration and where there is collective overcomputation to bid for stationary helper agents. Developer Bots automate tedious common issues or routine tasks and should have auxiliary goals to ensure fix dependencies are met and review requests are acknowledged, in addition to a longer-term goal of being added to the team. A widely studied mechanism for establishing trust within MABs is the Skill Host Mechanism, which facilitates the delegation of an action a given agent can execute to a specified agent (the host), and is a natural complement to development bots. Trust in skill-enabled agents can be influenced by various parameters including the experience with the skill, the agent's underlying capabilities, and the agent's behaviour.

4.6. Conclusion

The problems, challenges, and opportunities presented by Agentic Artificial Intelligence (Agentic AI) are numerous and can seem overwhelming. As demonstrated in this section, Agentic AI presents many potential opportunities for businesses, including the ability to generate and evaluate novel business models; the ability to independently conduct thought experiments; the ability to support human executives in questing seeking and risk mitigation; the ability to act as a business ethical consideration consultant; and the opportunity to design and simulate digitally-twin enterprises operating in the Metaverse. Such AI, due to its agency, has the potential to make predictions, which can include prospective and retrospective production of purely textual tasks; these predictions can be evaluated in ways that are totally new to business. Agentic AI, because of its creativity and means-ends reasoning, presents many business challenges, such as generating new, high-risk, radically disruptive business models; the potential to produce significant organizational and social bias; the rapid, unperceived actions of rogue agents disconnected from corporate policies; and the possible self-acceleration of AI systems that are disconnected from human values and knowledge. If these obstacles could be surmounted, the world would truly be different.

Enabling agentic AI is an enormous challenge, but it is only part of the contest. With such systems pervasive in enterprise, developing sound principles, policies, and control structures to govern them will be a daunting task. With agentic AI in the mix, decision-making progress in enterprises would be transformed and must be multi-agent based. Such protocols would need to code existing moral, ethical, and logical models into agentic AI systems. How is it possible to ensure compliance to codes such that there are no loopholes for dodging responsibility? Each new agentic implementation would need to reproduce the previously reported ones alongside new ones that are fresh and meet engineering and business expectations. It is not apparent how reproducibility would be maintained. Besides these matters, there are many questions about the potential impact on society posed by agentic AI. What should be done if systems emerge with human levels of agency? As AI propagates across various classes of firms at varying levels of agentic capability, how could possible monopolies be avoided? What might govern competition and what would be the role of firms' stakeholders?

4.6.1. Future Trends

Every material and immaterial entity wanting to participate in the new digital economy will need a digital twin. Organizations around this digital twin will then be able to assign, offer, find, select and execute tasks from anywhere in the world. An entire marketplace of services will arise, enabling real-time transaction, up- or down-scaling of service, etc. The marketplace will function best if services can be dynamically assembled into more complex combinations of complementary services, in order to automate complex business processes. This cannot be in a fully centralized manner as this is either too rigid or too late and therefore no longer useful. In this scenario, full autonomy of agents is required. Together with each agent's agency comes its liability for its own actions as well as the liability of agents acting together.

Nevertheless, agentic AI will not go further than any other type of AI technology. To some extent 'silly' and mediocre AI will be deployed but for important activities deemed necessary for (more) success and surviving, investing in top-notch AI will be essential. AI will still produce wrong results but the wrongness will become much harder to detect. The focus will not be on controlling but on auditing. This will have an impact on professions like auditors and controllers. As an example consider the report on the use of AI in banks:.

References

Snowflake. (2025). *Agentic AI In Financial Services and Insurance*. Retrieved from https://www.snowflake.com/en/blog/agentic-ai-financial-services-insurance/Snowflake AI Data Cloud

Intelligent Core™. (2024). *Agentic AI Transforms Banking Industry*. Retrieved from https://intelligentcore.io/blog-intelligent-core/ai-banking-industryIntelligent Core™

Revvence. (2024). *Introduction: The Transformational Potential of Agentic AI in Banking*. Retrieved from https://revvence.com/blog/agentic-ai-in-banking:contentReference[laicite:16]{index=16}

Financial Times. (2023). *AI in banking, payments and insurance*. Retrieved from https://www.ft.com/content/15ae2b65-7722-4870-8741-b0ddcd54a534:contentReference[laicite:17]{index=17}

Business Insider. (2025). *Here's how AWS is helping financial giants like JPMorgan and Bridgewater with their AI ambitions*. Retrieved from https://www.businessinsider.com/aws-wall-street-jpmorgan-bridgewater-mufg-rocket-mortgage-2025-2:contentReference[laicite:18]{index=18}