Deep Science Publishing
https://doi.org/10.70593/978-93-49910-74-4_3

Chapter 3: Agentic artificial intelligence for dynamic claims processing and fraud detection

3.1. Introduction

In the last century, insurance claims processing has gradually moved from being an entirely manual process to being more automated with machine-learning optimization. Machine learning has improved important sub-tasks: Natural Language Processing can convert unstructured documents into structured data, Image Processing can analyze regulators' photo examinations, and Predictive Modeling can compute estimates based on historical experiences. Consolidation and integration of these models into a full cycle, end-to-end automated claims processing and adjustment system are important areas of research and development. Current automated systems are mostly passive and rudimentary: they merely predict part of the outputs ignore important dependencies, and are limited in their dynamic updating-to-enormous internal databases housing individual unique customer histories and carrier experiences based on geolocation and other important parameters (Ngai et al., 2011; Baesens et al., 2015; Esteva et al., 2019).

Agentic AI is a step toward the next generation of fully multi-agent systems, enabling significantly more sophisticated artificial agents capable of learning, applying, and updating real knowledge and long-term memory databases in permanently online mode so that the outputs are reconciled, complete, and credible. Agentic AI is a novel real-time, fully dynamic, multi-agent AI architecture that does not merely provide predictions at each time step but also provides a dynamic set of intra-model actuals agent-to-agent, addressing the issues of deep dependencies and credibility among the various sub-tasks. Furthermore, by being continuously online, learning real relationships and reconciliations across models dynamically, and updating permanent internal memory knowledge, memory errors, and catastrophic forgetfulness are reduced. We illustrate the power and efficiency of the Agentic AI paradigm through its specific application to the challenging problem of automated real-time repeated-cycle claims processing and fraud detection. The specific approach to loss prediction also increased credibility despite only simple predictive heuristics applying on the individual model level, by taking expected values across the various models and sets of model actuals (Sun et al., 2012; Van Vlasselaer et al., 2016).

3.1.1. Overview of the Study

Claim processing is a time-varying sequence of tasks, performed by a variety of actors over an extended period. The goal of claim processing is to transform an abnormality report – a claim – into a complete, consistent, verifiable dataset of facts about a specific event (or set of events), the generation of which is a prerequisite to reimbursement of losses. Occasionally, some of the events and facts contained in a dataset are fabricated; the goal of fraud detection is to identify such datasets. The data residing multiple times within each dataset (its narrative) and in the datasets of many surrounding events (the surrounding community) are, respectively, externalized and implicit knowledge about the event, required to validate the claim and identify the disordered surrounding community, who collaboratively generated the disordered dataset.

Fig 3 . 1 : AI in Claims Processing

The computational analysis of such claims processing and fraud detection, to accelerate time to validation and/or reimbursement, is presently performed only as a matter of course for the simplest – that is, low-cost and at lower-risk – claims. Claims scoring systems are applied only to identify the remaining claims that are complex enough for the detailed human investigations required for validation of higher-value datasets, at higher risk for the claimant's acknowledged abnormality being the result of an actual event, rather than the product of collusion by the claimant and the surrounding community. This study demonstrates the feasibility of an approach to implementing dynamic, agentic AI systems, for automating the tasks of dynamic claims processing and continuous fraud detection. Moreover, such actions often must be performed as more information about each claim becomes available, as additional documents are provided to the insurer by the claimant (or another actor), and as the insurer continues to search for evidence of actual events by the claimant, and acceptor by the surrounding community.

3.2. Understanding Agentic AI

Definition and Key Concepts Agentic AI, or Autonomous, Self-Directed AI, refers to AI systems capable of autonomously or self-directively performing tasks, eliciting goals, and acting in a manner that is often creative or complex, or that can disrupt existing processes, systems, or institutions. Understanding what is meant by Agentic AI requires unpacking and clarifying several key terms commonly employed in AI literature. First, AI systems are computational toolsets that can perform goal-directed, albeit usually narrow, tasks through complex processes that use algorithms, programming, and databases, while demonstrating intelligence that approximates or exceeds human skill or reasoning in completing designated tasks.

Computational toolsets may be autonomous in the operational or physical sense, meaning that they can work without the need for human oversight in completing functions, or in the initial planning sense, meaning that they can elicit their associated goals, either explicitly or implicitly, as one important function. AI systems are even capable of constructing other AI systems to perform particular functions or tasks according to logical or mathematical rules. Acquiring agency, intentionality, or sentience are features that distinguish humans and related entities from conventional AI systems.

Tools are intentionally designed by someone on some computing architecture to accomplish a specific goal or problem. However agentic tools can entail multiple agents creating associated algorithms and systems for solving a problem. The AI toolset is not just a simple calculator or expert system; it is capable of running complex programs, and of improvising, learning from experience, and casting a very wide design net about possible ways to organize knowledge and build logical problem-solving routes.

3.2.1. Definition and Key Concepts

The Agentic AI concept is likely to confuse most readers. The form ("Agentic AI") is novel. Existing terms, like "autonomous systems" or "intelligent agents", are familiar but not well-defined. The combination of the form and the unfamiliarity of content makes it likely that both the meaning readers take away and their reactions will be diverse. For some, Agentic AI will elicit broad belief or disbelief, possibly because of their group identity; for others, nostalgia and a wistful search for the personal experience; for the majority, surprise, and question.

We can define an Agentic AI as an AI system that acts as if it is acting independently of external control or influence. This definition is novel but simple: it extends the intuitive human conception of agency to encompass a ground-up, functionalist view of AI and other agents. A well-established point in the philosophy of the mind, functionalism holds that mental states are defined by their role in a system rather than any underlying biology. A computational approach to AI reinforces the functionalist view: AI systems increasingly show intelligent agency only by sophisticated algorithms implemented on advanced hardware. The conceptual leap of Agentic AI is to sufficiently recognize executive functioning as key to agency in any system, as with any human or AI intelligence. By executive functioning, we mean the cognitive skills – like working memory, inhibition, cognitive flexibility – and regulatory systems associated with self-control and the direction of controlled cognition related to achieving goals.

3.2.2. Historical Context and Evolution

In 1962, a Dartmouth College professor named John McCarthy defined artificial intelligence as "the science and engineering of making intelligent machines." The ultimate goal was to build machines that were as "intelligent" as people and act similarly; that is, humans could not know with certainty whether the machines were purely mechanical or something much more advanced. Many researchers have pursued this goal over the years, and one of the most visible approaches to such agentic AI has been to create systems that use language models along with some method of accessing the collective wisdom of the internet to provide expert-level insight on almost any question, in natural language. However, these expert systems are narrow AI in the sense that they do not plan, initiate, and carry out tasks without human intervention. "State" is a difficult term in this context, as it implies a personified entity that maintains continuity over time, unlike a phone that is in a different configuration after the alarm goes off. We use the term agentic AI to denote an agent, independent of state, that is capable of being agentic the way a person is capable of being agentic, because the action that typically gives rise to other types of agentic AI is text-based inference (like a person), even when it is over tasks that are not natural language tasks (unlike a person). Many people refer to only the

AI itself as "agentic AI," but the more precise term includes the tasks the AI is performing. An agentic AI might also be collaborating with agentic AI and other tools in the physical world (such as a robot, unlike a person).

One of the first personas to interact with early machine learning efforts after the 2012 launch of more sophisticated methods to create state-of-the-art neural networks was Google's engineer Scott Galloway, who became friends with Google's puppet. Other interactions were more confrontational, such as the famed disagreement between Eliezer Yudkowsky and a group of Google engineers who did not believe that their new language models were capable of achieving AGI.

3.3. Claims Processing Overview

Digital technologies, such as AI and Distributed Ledger, have made Insurance Claims Processing faster and more efficient. However, the current operational models have limitations of vulnerability to fraud, throughput, usage of resources, cost, transparency, and decency. Section 3 of this paper describes the Claims Processing System in detail and discusses its central challenges. The Claims Processing System is an essential element of Insurance Companies, as it protects customer interests and ensures that only valid claims are processed. However, with the advent of Cloud Computing, the Internet of Things, Digital Technologies, Financial Technologies, and Distributed Ledger Technologies, Insurance claims processing needs to be faster and more efficient while ensuring that fraud is blocked at the source. Currently, there are various models for Claims Processing including Manual Processes, Automated Sampling, and AI-Assisted Models. However, at present, fraudulent Claims Processing is still a significant burden on Insurance Companies. In the US alone, Insurance Fraud is costing the industry and policyholders more than 60 Billion annually. This is primarily because the focal point of Current Claims Processing cannot evolve to handle the incoming claims at high throughput and check for all anomalies that would indicate fraudulent intent. In an ideal state of the Claims Processing System, Insurance Claim Forms should be submitted through a guided digital process on a Device with embedded error detection and identification of fraudulent entries. As the Fraud check should be Level 1, Claims should then be processed by AI Agents after fulfilling the necessary KYC Norms by checking against Global Databases.

3.3.1. Current Practices in Claims Processing

The business context for our exploration of AI for claims processing is the insurance industry, which is undergoing a digital transformation driven by customer demands for better customer experiences and cost-effective services. Claims processing is a critical operational area within this context, with a major influence on customer satisfaction and the profitability of insuring operations. Digitalizing the claims processing operations involves deploying advanced technologies for digitizing claims data collection, rapidly validating claims data for its completeness and accuracy, and automated analysis of claims data to decide on the merits of the claim, including the estimation of the amount of the claim payable, and the final settlement of payments quickly and accurately, with minimal human intervention. This process is presented diagrammatically.

A typical insurance organization offers a handful of insurance products across a large number of geographic regions, generally through multiple agents, to a large and diverse pool of customers. On purchasing an insurance policy, a customer pays a premium amount regularly over the lifetime of the policy. In return, the insurer promises to settle a specified amount upon the occurrence of some unfortunate event to the customer or their legal heir during the lifetime of the policy, such as a medical event, the demise of the insured person, or the loss of a property. If the insured event occurs, the customer files a claim and is required to provide supporting documents and proof of the terms of the policy. The organization then validates the claim information against the policy terms and available historical data, checks for factors of possible fraud, and settles payment to the customer if the claim is found validated.

3.3.2. Challenges in Traditional Systems

Despite the growing data and the industry investment, claims processing and fraud detection systems still rely heavily on manual reviews and rules-based systems. On the one hand, although rules-based systems using parametric correlations have low modeling overhead and low latency, these systems suffer from significant limitations in accuracy and scalability since they cannot adapt easily to system changes and often generate many false alarms. On the other hand, more advanced solutions using machine learning techniques require training data, model tuning, cost-sensitive modeling, and expert knowledge to build. Without the ability to use models to simulate claims processing decisions at scale, fraud analysts and insurance investigators can only generate limited insights. As a result, traditional claims processing and fraud detection systems are lagging behind other data-driven decision-making systems in the industry.

Further, a major limitation of claims processing systems is data reliance. Either models are used purely for detection as in rules-based systems, or they predict outcomes directly as is done for the machine learning solutions. Recently, researchers have developed reinforcement learning-based combinatorial models to help detect fraudulent claims at a small scale. However, existing solutions cannot be generalized to the industry as they do not utilize the structure of the claims fraud detection processes nor take into account the continuous need for fraud updates and the dynamic nature of the overall claims processing system. We believe that with a small cost of adapting the models with up-to-date parameters based on a small set of recent decisions, a well-designed agentic AI-based modeling framework would be able to model the entire claims processing and fraud detection system. From there, we would simulate a typical backlog decision-making process. Then, we would be able to predict outcome probabilities and the overall decision costs for different claims, which could be incorporated into downstream machine-learning models.

3.4. Fraud Detection Mechanisms

Insurance fraud is defined as an intentional misrepresentation by an individual to receive undeserved benefits from an insurance policy. Fraud could exist on both sides of claim transactions – a claimant submitting a fraudulent claim for insurance indemnities and an insurer committing fraud by denying, failing to settle, and delaying the settlement of a legitimate claim. While fraudulent claims are committed by a small percentage of clients and involve deceit, damage caused by both forms of fraud can be substantial. Fraud increases claims costs, which in turn increases insurance premiums to cover these losses. Hence, to keep fraud at a minimum level, insurance companies invest in fraud detection mechanisms.

Fraud is a complex phenomenon that is difficult to identify and predict by traditional methods. Although a variety of factors play an important role in fraud detection, the value or existence of certain characteristics may be linked to different types of fraud. We identify three types of fraud in claims: exaggerated claims fraud, which alleges more damages in a claim than incurred; opportunistic claims fraud, which inflates a legitimate claim or files a false claim; and fraudulent claims fraud, which makes a false claim intending to solicit benefits that are not deserved. Besides insurance companies, fraud detection has also drawn much interest from domain experts in financial and medical institutions and telecom and utility providers. However, while insurance fraud detection is still an exploratory area of research, work in other domains has invariably focused on the development of suitable detection mechanisms using artificial intelligence techniques.

3.4.1. Types of Fraud in Claims

Fraud fundamentally denotes an action or instance of deceiving or misleading another to gain something of value. It is further defined by various forms encompassing a variety of deceitful acts and omissions by which an individual or entity seeks undue or undeserved advantages, benefits, or enrichment. In insurance claims, fraud is the intentional deception perpetrated by an insured, claimant, individual, or business entity, which results in an unauthorized benefit or payment, and costs the insurance industry hundreds of millions of dollars annually. Furthermore, claims fraud detrimentally affects the insurance ecosystem at large, as insurers countersign the sway manipulated losses and rise in premiums by consumers. Fraudulent claims also slow the resolution of legitimate claims, and this will curtail the insurer's ability to fund future claims.

In claims processing, fraud is broadly classified into three categories: First-party fraud involves an insured deceitfully inflating a claim, filing a claim for an event that did not occur to benefit financially, or misrepresenting information to obtain a benefit or payment. In adverse selection, an applicant withholds specific information about a greater probability of a claim when purchasing a policy. Second-party fraud occurs, for the most part, in the context of commercial insurance, with a business deceptively colluding with its suppliers to claim an inflated loss. Third-party fraud involves a third party who colludes with the insured to state a fraudulent claim, makes a false claim against the insured, or otherwise acts against the insurance company's interest. Insurance claims fraud may further relate to property and casualty insurance, healthcare, commercial auto, workers compensation, and disability. Various types of fraud impact these segments and their other sub-segments differently. For example, soft tissue/collusion and staged accidents fraud hit property and casualty and commercial auto hardest, while billing and identity theft fraud is most prevalent in healthcare.

3.4.2. Traditional Fraud Detection Techniques

Fraud detection is an important aspect of audit practices. Auditors employ various techniques to detect fraud, including physical examination of specific items, observation, confirmations, inquiries, and analytical review of financial statements. Auditors are required to apply these techniques for a large volume of transactions that are usually reported. In the literature review, it was found that somewhat similar techniques are also utilized by fraud detection systems. The purpose of fraud detection systems is to help organizations deal with the enormous volume of data on activities that might be fraudulent.

The use of the data mining technique of cluster analysis to identify groups of transactions that appear to be similar has been proposed. Other outlier analysis techniques proposed for use in fraud detection include the use of expert systems, neural networks, and statistical techniques such as regression analysis and traditional rule-based techniques. These techniques are used in one of two basic methods or approaches. These are the knowledge-based approach, which employs the expert system, or the data-driven or statistical approach that employs outlier detection based on the prediction of parametric and nonparametric functions, classification, or clustering algorithms. This review shows that both methods can be used in fraud detection systems to detect errors in a range of domains.

Statistical and predefined rule-based methods, as well as artificial neural networks, have been applied to detect organizational and internet-related fraudulent activities. While the major drawback of using predefined rules or statistical methods is the high number of false positives generated, using neural networks provides a possible solution to help auditors sieve through the higher volume of lower-quality transactions.

3.5. Agentic AI Applications in Claims Processing

Claims processing is often a highly complex assortment of activities, including inquiries, file verification, determining compensable warrants, preventing fraud or subrogation, working with policyholders, and making payments. Both financial and non-financial data must be gathered and considered to enhance and expedite this process. Streamlining the processing of claims can therefore reduce expenses while improving customer experience. Several enhancements in agentic AI applications might be explored to address dynamic claims processing.

1. Automating Claims Review

Even when company tacit knowledge and industry knowledge associated with complex claims are codified into systems, claims processing may not be delegated to computers, although many aspects increasingly can be, for instance, where well-defined computational policies providing definitive action rules can be developed. Claims analysis often requires human experts with domain knowledge and company-specific tie-ins to facilitate the necessary interpretation of hundreds of pages of data and ascertain what is material and relevant to justifying or questioning millions, or even hundreds of millions or billions of dollars of claims which must be paid out or at risk. For rules that cannot be hooked into automatable functions, AI can play a role through machine learning.

2. Enhancing Decision-Making Processes

The decision-making function for claims and coverage, especially in the context of large, complex, high-dollar cases that cross multiple domains with negotiations for subrogation, may require specialist human-intensive exploration, with hybrid decision support systems that highlight question prompts and analyses by human experts. Recent advances in natural language processing, however, hold the promise of developing more sophisticated tools that may be concatenated or re-orientated than those that were available earlier using expert systems technologies.

Fig 3 . 2 : AI Agentic Applications Transforming the Banking Sector

3.5.1. Automating Claims Review

The demand for rapid and reliable claims processing is increasing as organizations strive to deliver superior customer experience. Yet, the claims processing workload, a common source of frustration amongst insurance employees, is only becoming heavier. Insurance companies also face rigid statutory timelines for settling claims, along with an ever-increasing volume of claim submissions. Even when a submission can be readily resolved by leveraging a previous understanding of similar past claims, these types of claims require authorizations by a trained professional. Simultaneously, while insurance companies are looking to automate more of their processes, high volume low-value claims remain a challenge where companies are hesitant to outright provide autonomy to AI-driven systems, nor are they willing to increase the overhead of designing and implementing robust AI systems that become a bottleneck in claims processing.

This leads to a situation where companies are hesitant to adopt advanced automation solutions which can lower the burden of processing low-value claims, and facilitate timely review by experts on sensitive claims that require more nuanced human judgment. This is where agentic-based Assistants are well poised to make an impact. For example, our solution for claims processing uses deep learning-based object detection and zero-shot concept learning to rapidly assess and surface the salient components in a claim, enabling agents to form a latitudinal understanding of what is being driven in the claim. After which they can elicit clarifications from the claimant via a wizard-based assistive interface. The operation can be highly regulated to ensure that the claims processing agent has appropriately understood and documented the processes surrounding the claims. The suggestions presented at every stage can be constructed to assist with a specific dimension of understanding the claim using an explainable AI approach thus reducing opaqueness.

3.5.2. Enhancing Decision-Making Processes

The widespread use of rule-based systems around decision-making problems makes them suboptimal over time. Such decisions often require human argumentation, debate, and consensus-seeking processes as complex real-world tasks grow in difficulty, size, and interdependence with other decisions. Decision-makers often outsource tedious evaluation aspects of the problem to weaker agents or other systems, but the core evaluative steps must still be conducted by the decision-makers. Hence, we regard claims decision-making as a collaborative activity between decision-makers and an Agentic AI, where the AI modeled as a super-agent earns its key position.

The Agentic AI retrieves historical solutions and their rationales from the shared decision space and initially proposes its solution, which can be accepted, modified, or rejected by human decision-makers. This type of technology is called a decision support system. For an iterative solution process, the decision support model must perform its initial version faster than any decision makers. The final solution will likely differ from its initial version, as decision-makers are involved throughout the solution process. Hence, incorporating advisory decision support systems is the least disruptive for organizations, as no major changes in processes, culture, and skills are required. Alerting systems or advising AIs using our computational modeling of Agentic AI can be applied to all known problems, where decisions are conducted using formalized procedural steps.

The collaborative capability is also useful in faster iterative processes for design and visualization tasks. Such a process allows for creative feedback loops between humans and AIs, where both can learn and improve from the results of each cycle. Some design and technical tasks may see such rapid iteration cycles over large problem spaces that humans cannot stay ahead of the progress anymore even in creative aspects of the projects over larger technological horizons.

3.6. Agentic AI in Fraud Detection

"Agentic AI, capable of semi-autonomous, goal-oriented decision making involving dynamic action-selection in changing environments, can play a significant role in fraud detection. Unlike most machine learning algorithms for fraud detection, an agentic AI can be assigned the specific goal of detecting fraudulent activity on live data. Thus, it can perform such fraud detection continuously on many different, rapidly evolving sources of information, using several different techniques in a coordinated, timely, and efficient manner.

Machine Learning Algorithms for Fraud Detection

Various machine-learning algorithms have been developed to detect fraudulent activity. These have included classification algorithms that differentiate between legitimate and fraudulent activity, clustering algorithms that identify behavioral patterns for user activity, anomaly detection algorithms for tracking changes over time, and network analysis algorithms for mapping the connections between multiple instances of user activity and detecting patterns or outliers in the transactions or activities of the network. Additionally, natural language processing and computer vision-based methods are also utilized to extract fraud-relevant features for improved prediction of fraudulent activities.

Real-Time Monitoring and Alerts

Detecting fraud in real-time as activities are undertaken is critical to minimizing the financial losses associated with fraud. Even if knowing your customer processes has ensured that each customer is trustworthy in the eyes of the company at the onset, without continuous monitoring of real-time transactions for anomalies, KYC processes cannot protect from colluding actions that violate the implicit terms of a business relationship."

3.6.1. Machine Learning Algorithms for Fraud Detection

This happens to be quite a broad area of ML, but one that has a very different flavor of application than many of the more typical AI uses, whether it's in computer vision natural language processing, or image enhancement. In this area, rather than computationally generic learning and inference algorithms that would be applied to very different problems at scale by people with highly specialized domain expertise, the domain of application is highly specialized - measuring and predicting behavior for certain types of financial transactions performed via certain types of systems or channels.

While ML models can be applied to a very wide range of data types, there are two main classes of problems for which they have been successfully applied. The first is supervised detection or prediction of labels - for example, given a transaction and label indicating whether it was fraudulent or righteous, train a model that predicts the label for new transactions in the dataset. The difficulty is that data is usually highly unbalanced - perhaps 99.995% of labeled transactions are righteous, and only 0.005% are fraudulent and labeled examples of new detection targets are also likely to be unbalanced in effortful and time-consuming ways. As a result, standard supervised prediction algorithms generally do poorly on this task, particularly if highly generalizable features aren't available. Ad hoc solutions must be used to determine "good" approximations of probability score thresholds such that the balance of computation and misclassification error are acceptable.

The second is anomaly detection, where the task at hand isn't to detect or extract timely labels about new transactions (although it is often done very much that way), but rather to "flag" those transactions that appear anomalous, given data about previously-validated righteous transaction types. However, both tasks can be computationally intensive, particularly in real-time scenarios.

3.6.2. Real-Time Monitoring and Alerts

Dynamic claims processing pushes the frontier in allowing the detection of emerging forms of fraud and the processing of a multitude of small-value claims for which traditional human-in-the-loop processing may not be a viable option. Further, because there exists a significant correlation between emerging forms or peaks in fraud and other real-world dynamics, the agent needs to have access to constantly changing external data for optimal fraud detection and detection of fraudulent claims. Because of the constantly evolving nature of fraud, machine learning techniques alone cannot provide the final answer. Human set and human shift in the automated processes must ultimately be the backbone of continued intelligent alert generation, correlation detection, and optimization.

Real-time monitoring of claims activity and alert generation based on sets of models need to be enhanced by correlation detection across groups of claims and integrated alerts on matrixes connecting nodes of related claims across groups and types of fraud. Instead of attempting to use a single model to process all possible claims, developing and integrating multiple models, each focused on detecting different kinds of irregularities, to create a composite model that produces and/or corroborates alerts or results of individual models is the optimal approach. Also, modeling of claims activity for fraud detection cannot be a one-time event since either the nature or mix of fraud in the external environment or the type, volume, or mix of claims being processed may vary thus creating or pushing out detection threshold limits of individual claim types.

3.7. Data Requirements for Agentic AI

Although our primary motivation in the above sections was to talk about a novel abstract and philosophical concept of agency in AI and to discuss the implications of the results, products, and tools that we plan to offer, none of this is feasible or possible without a specific kind of data that may be referred to as agentive data. In the context of our work, agentive data will be a collection of data types about the nature of the claims, claims data, nature of fraud, related information, instructions and guidelines for processing these and making key decisions, possible decisions, outcomes, fine-grained annotations of rightly, wrongly processed claims, sample claims for training, and so on. In addition to the key features described above, these data types will be global, historical, and dynamic over time.

Data Collection Strategies

Identifying data about any subjective task across different demographics is messy and taxing due to the sensitivity of the task. It requires access and permissions to dedicated target groups. Over the years, companies have accumulated a huge amount of processed, yet unprocessed claims data, from various sources, using a combination of supervised, semi-supervised, and weakly supervised strategies. This data has a tedious, yet shallow annotation as it is barred from decision-making process details, decision-making action steps, possible diverse decisions, possible steps to enable, and decision outcomes. This data, although has a lot of noise, could be utilized as a starting point or interim scaffolding for learning various requirements. Harmonics is also useful in building possible personalization principles, foundational use case specifications, and agentive principles.

Data Quality and Integrity

While toxic and toxic often might be taught with amazing achievement acceleration with only user-generated coarse data or imperfect decision spectrum, it's important to pipeline agentive AI, especially for high consequence and subjective tasks, with strictly curated data. For each task, the available curated data must represent every way those tasks are normally and practically performed - every alternative, ignore options and assets, nuances, intimately, interconnected systems and task interferers - but more refined and expanded to fit an AI-ready model. Careful seeding, moderation, and governance would be the essence of this data. Beyond curation, ensuring the quality and integrity of this data is also very important. Systems for enabling automated and collaborative feedback loops for quick and cost-effective reporting and moderation surely dictate, enrich, simplify, and improve the formality pipeline.

3.7.1. Data Collection Strategies

While agentic AI largely remains an unsolved problem, we see great potential in existing approaches to enhance claims processing and fraud detection automation through human-computer augmentation. We speak of hybrid AI in these regards, since human experts are still needed to define goals, monitor and steer the agents, and assess the output quality and validity. As hybrid models of automated decision-making are implemented, data requirements vary widely. Some augmentations will rely on relatively shallow technologies, such as pattern-matching models that classify input documents or input variables. Other augmentation scenarios will use multi-modal models that generate complex output variables. We will have different data requirements depending on which modality plays the role of the preliminary process and which one directs the user's attention to more relevant cues.

If shallow models lend themselves to classical machine learning approaches, that require large numbers of labeled examples for supervised learning or domain data to build the desired features in the case of non-/weakly supervised learning, then hybrid models will require fewer human-created training labels. For example, a multi-modal model could make the selection of relevant examples for a few-shot model more effective; this would require that the examples be labeled in the first place, e.g., for outcome variables, such as decision outcome, and at various points in the preliminary process. If the multi-modal model is also to be used to highlight relevant cues in the deep model's decisions, then attention maps would also need labeling.

3.7.2. Data Quality and Integrity

Data quality and data integrity are critical dimensions when implementing agentic AI and far more important than they would be for traditional AI. Consider the question: Why use agentic AI for important, complicated, high-stakes tasks like dynamic claims processing and fraud detection? The answer is that we need super-AIs, far more capable and trustworthy than ordinary enterprise AI, to carry out these important tasks at scale and with requisite quality and control which we cannot expect from humans without breaking under work pressure and risk of making errors or engaging in fraud. Our response, however, is premised on AI's working with the same quality of data and data integrity as receiving high-quality input from data providers. For ordinary data quality, it is far more likely that the work will indeed be carried out faster and with greater accuracy and control by AI than by human processors.

Therefore, human data teams must ensure that the data received from various data providers is of the best quality and that the provenance of the data and processes that generate the data. This is the basis of data quality and integrity, the foundation of AI function. It is far less common than anticipated because disintegrated, large-scale collections of data combine from multiple data generators and are stored as separate data silos. The data from the various data sources differ widely and in potentially consequential ways including in data definitions, structures, formats, and frequencies of input; timing, triggers, filters, and control of data collection; privacy controls and security provisioning; data accuracy and veracity management; and quality assurance. Indeed, agentic AI consumes these models for its parallel multi-tasking.

3.8. Conclusion

Dynamic claim processes and fraud detection are two key functions of large P&C insurance claims operations. In both functions, the ability to process information rapidly, often incomplete or uncertain information, is essential. Over the last ten years, there have been huge accomplishments in machine learning that can be used to automate difficult and high-error tasks. In this work, we proposed the use of these AI agents to carry out major unautomated parts of dynamic claims and fraud processing. This proposal opens up several research directions. How can we best translate developed procedures into agent capabilities? What is the right pace and types of model improvement? What degree of human participation at each step is optimal? What communication protocol should the agents and claims processors use? What mixture of tasks should be done by the agent and by human processors? When is it appropriate for the agent to take full charge? Lastly, how can large model-agent capabilities be adapted for a specific claims operation? These questions require further research but offer the use of AI agents as a window into the future enhanced machine-foundation models.

Final Thoughts and Future Directions

AI foundation models have gained increasing attention for their exceptional generalization ability across a wide range of tasks. Multi-task fine-tuning has become a fundamental training paradigm for foundation models, where models are jointly trained on diverse tasks to support better transfer for unseen tasks. Easy customization enables the use of foundation models deployed in various applications. As shown in this project, agents built around AI foundation models can support practical solutions to real-world problems in various domains. Increased cooperation between agents and human experts is being explored for a wide range of tasks. In particular, agents generate actions based on observations, while human experts either critique the actions or provide additional feedback. When a user supplies an action, the agent can also build on the action to enhance joint performance. The two-dimensional, sequential, and shared nature of agent and human collaboration creates a rich communication between the two sides, thus supporting a high degree of synergy. Whether using a single foundation model to implement the agent or leveraging multiple different specialized models, agent-human cooperation can be expected to power many fields of human endeavor for the foreseeable future.

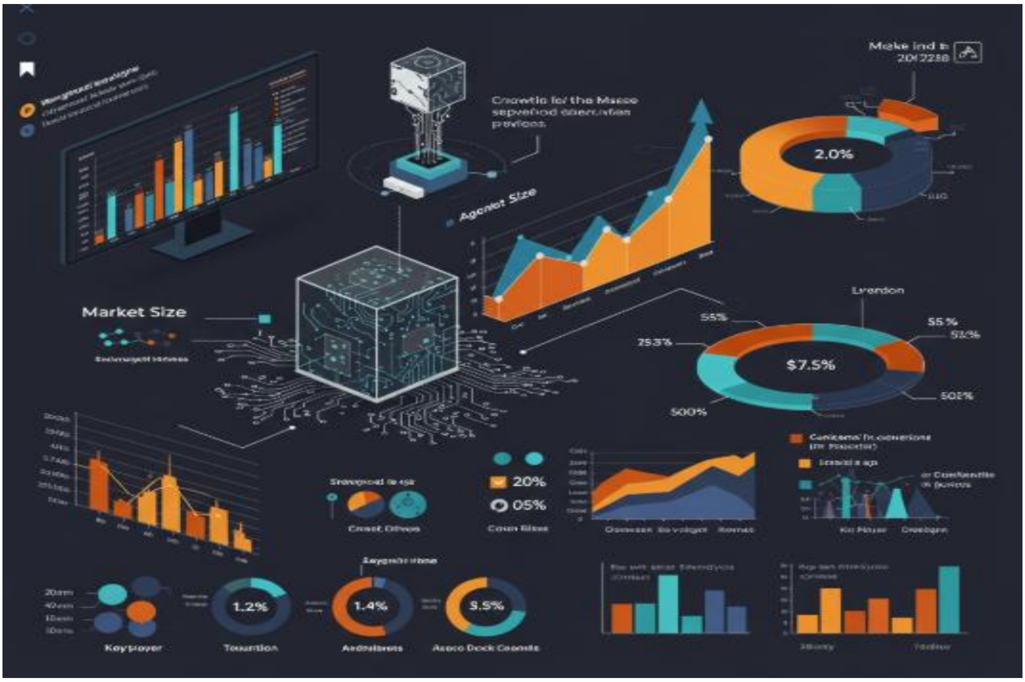

Fig 3 . 3 : Agentic AI For Financial Services Market Size

3.8.1. Final Thoughts and Future Directions

While substantial progress is being made in reducing the cost of claims processing through the development of agentic AI, we are at an inflection point in the development of dynamic claims processing capacities. Further advances will be built on the creation of a new ecosystem of tools that enhance auditor capabilities across the claims ecosystem. Just as certain systems did not automate the prosecution of a case, but rather tailored and connected existing tools to enhance the capabilities of human resource and time-constrained prosecutors, new tools must similarly enhance the capabilities of claims auditors. These new capabilities will require a sculpting of agentic AI along with enhanced links to existing tools for claims processing. The enhanced relations among these tools will substantially improve the effectiveness of claims processing by enhancing the capability of human auditors who establish the penal consequences that drive the claimants' behavior.

The ecosystem of new tools also requires a re-integration of claims auditing with the purpose of the claims — enforcement of the social contract with claimants, who are owed assistance during times of need. A central failure of the modern administrative state has been the failure to quickly respond to changes in the widely shared public perception that the social contract between the state and the citizens obligates governments to provide a safety net for those unable to support themselves without assistance and also to enforce penalties on those who could work, but choose not to. That strain on public support for government programs has been exacerbated by the agencies' adoption of risk-based rules to certification for access to programs. New tools can help agencies re-establish the missing links with the claimants and restore the trust necessary to implement effective dynamic claims processing and monitoring.

References

Baesens, B., Van Vlasselaer, V., & Verbeke, W. (2015). *Fraud Analytics Using Descriptive, Predictive, and Social Network Techniques: A Guide to Data Science for Fraud Detection.* Wiley. https://doi.org/10.1002/9781119133127

Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). *A Guide to Deep Learning in Healthcare.* Nature Medicine, 25(1), 24–29. https://doi.org/10.1038/s41591-018-0316-z

Van Vlasselaer, V., Eliassi-Rad, T., Akoglu, L., Snoeck, M., & Baesens, B. (2016). *Gotcha! Network-Based Fraud Detection for Social Security Fraud.* Management Science, 63(9), 3090–3110. https://doi.org/10.1287/mnsc.2016.2489

Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). *The Application of Data Mining Techniques in Financial Fraud Detection: A Classification Framework and an Academic Review of Literature.* Decision Support Systems, 50(3), 559–569. https://doi.org/10.1016/j.dss.2010.08.006

Sun, J., Wang, F., Hu, J., & Edabollahi, S. (2012). *Supervised Patient Risk Stratification through Integrating Clinical Metrics and Patient Similarity.* AMIA Annual Symposium Proceedings, 2012, 950–959. https://doi.org/10.48550/arXiv.1206.6310