Deep Science Publishing, 2026
https://doi.org/10.70593/978-93-7185-416-0_6

Chapter 6: Designing Agentic AI Architectures for Insurance Platforms

6.1. Introduction

In recent years, artificial intelligence (AI) has been applied to a growing number of functions in the insurance industry. These AI applications are typically not agents, however. They do not maintain their own citizen or business identity and have only a limited ability to independently observe the environment and act. Developing agents whose core objectives involve fulfilling insurance responsibilities—underwriting, claims processing, pricing, fraud detection, and risk management—offers the potential for significant operational efficiency and user experience gains.

Constructing agentic AI architectures for insurance platforms constitutes a challenging design problem. The task requires balancing the benefits and risks of agentic designs against the services and capabilities they provide. Architectures at all levels of the platform must satisfy end-user and enterprise needs while remaining transparent, trustworthy, fair, and accountable. Moreover, platform operations must conform with a wide range of insurance data governance, quality-control, and compliance requirements. These include maintaining data provenance, ensuring sensitive data is not revealed, and meeting other legal and regulatory obligations while minimizing the trust placed in AI systems. Among these needs, fairness considerations are particularly stark. Older AIs have clearly perpetuated print- and online-media bias through biased training data. Most likely, agentic agents will require carefully constructed synthetic data engines to safely replace historical data across sensitive application domains.

6.1.1. Overview of Agentic AI's Role in Insurance

The financial returns expected in the near to medium term for most insurance companies hover around unimpressive single digits, putting extraordinary pressure on any investment that can free up or optimise the use of capital. In this environment, Attention

Role-based Agentic Intelligence (ARAI), when used at appropriate levels of autonomy and orchestration, offers a new way of implementing elements and components of existing business processes that increase operational efficiency and effectiveness. In its simplest form, a failure of service delivery gratification can be expressed as a shortage of key hires or roles, who appear in the business process but cannot be filled or backfilled for operational reasons. Augmenting or replacing these roles may offer the fastest returns.

Agentic AI is an emerging paradigm for creating autonomous agents that can work independently or together as a team within a modular architecture. ARAI enables such agents to make their own choices, plan their future actions, and direct those actions based on feedback from their environment, using Attention Role-based Agentic Intelligence. Agentic systems can orchestrate decisions they cannot themselves implement, and do so while satisfying the principles of fairness, accountability, and transparency (FAT) as well as data governance and compliance requirements. These two sets of criteria have emerged as the guiding pillars for the design of agentic systems in the financial services sector.

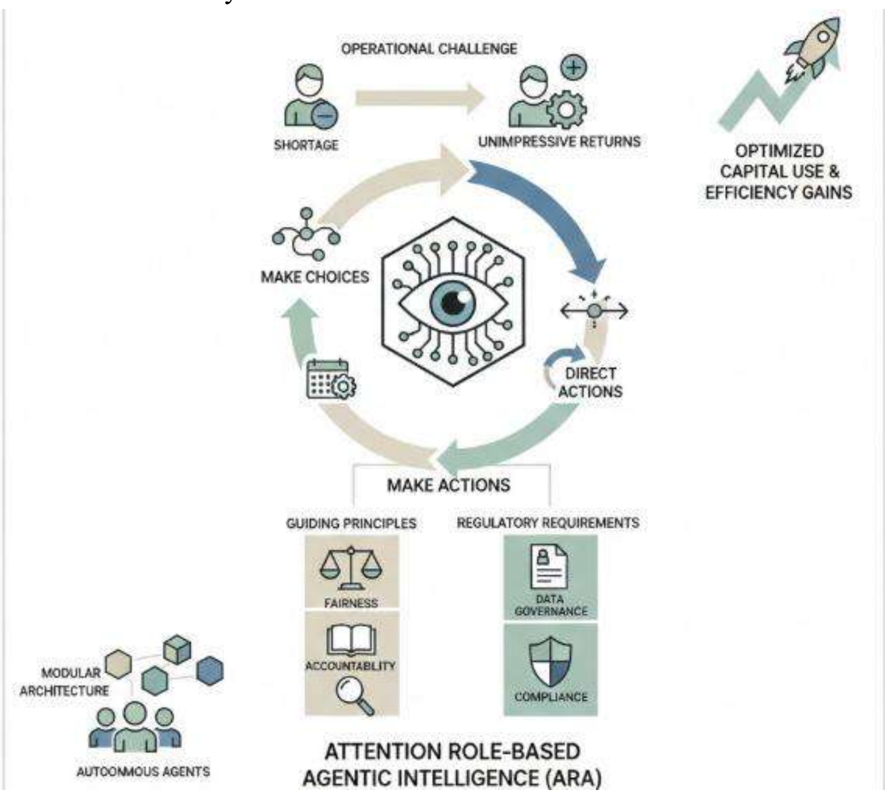

Fig 6.1: Optimizing Capital through Attention Role-based Agentic Intelligence (ARAI): A FAT-Compliant Framework for Insurance Operational Efficiency

6.2. Theoretical Foundations of Agentic AI

Fundamental Concepts of Agentic Intelligence

Agentic AI refers to predictive models serving as agents, capable of autonomously discovering a solution to any question or task that requires judgment and/or physical action. While the concept of agentic AI is broad and can be applied to myriad domains, its implementation requires considering the functional nature and the specific objectives of the system being designed.

An agent is a system that perceives its environment through sensory input and takes action through an actuating interface in order to maximize a reward signal. More complex systems can be described hierarchically as collections of collaborative agents, where each agent accomplishes its objective through interaction with other agents, either as direct input to its own actions or via rewarded action sequences that influence the actions of other agents in a way that benefits its own reward. In such systems, behavior emerges from the collective interaction of agents, rather than being explicitly programmed or designed. Agentic AI draws on concepts and terminology from multi-agent reinforcement learning.

6.2.1. Fundamental Concepts of Agentic Intelligence

The term "agentic intelligence" characterizes a class of architectures in which autonomous agents make decisions that effect change to the physical or digital world, but where the agents themselves do not learn. The DSR intended to provide a rigorous foundation for future empirical research into agentic intelligence. The synthesis uncovered four fundamental concepts illustrated by a diagram. An agentic architecture is controlled by a modular architecture and responsive to a change in the environments of action and observation. Decision-making does not occur in isolation, but instead revolves around the organization of actions with respect to their effect on the world. Change and design are fundamentally separate functions in agentic intelligence.

Specialization of an agentic architecture leads to an agentic system. An agentic system preserves the distinction between design and change. It is modular in the sense that agentic change may be enacted simultaneously by a number of different agents originating from different organizations. The emergent behavior of the overall system is, therefore, neither directly controlled nor designed by any individual or group acting alone. Agentic systems differ in the degree to which sunlight is focused to a change to the world, but the information theory of selection explains how, at the extreme of extreme no design at all is necessary.

6.3. Requirements for Insurance Platform Architectures

Data governance, data sovereignty, and regulatory compliance issues influence the design of insurance platforms. The adoption of regulatory constraints motivated by the fairness, accountability, and transparency (FAT) framework—as it is in the use of Fairness in Machine Learning—creates requirements for the architecture of agent-based Insurance solutions. Furthermore, FAT can serve as a high-level guiding principle for agent-based Insurance solutions on a platform composed of agents developed/maintained by different insurers. Such an insurance platform can be conceived of as a set of services available to the agents. Agents are expected to provide or consume those services in a manner that is compliant with the associated requirements stated in FAT. Requirements influence the selection of the architectural paradigm used in the platform.

From a Data Governance Perspective, agents are expected to comply with legal data sovereignty constraints in the management and processing of sensitive personal data. Data Federalism also provides support for data governance in the presence of multiple parties wanting to compute on private data. Orchestration services can assure that the agents follow the required data governance during the execution of the solution. For example, a core-sensitive data governance constraint of Insurance platforms is that claim information must remain secret for the underwriting process of other agents. Claims must be processed only by the agent that received them and only the agent that incurred the losses must be informed in order to comply with Legal Privilege Rules, thereby assuring the communications to remain secret.

6.3.1. Data Governance and Compliance

Data governance remains an area of intense scrutiny for financial services firms. Data assets must be properly held and processed within established guidelines and policies. Such guidelines must cover general security and access rights, data quality, data privacy, intellectual property management, records management, and use of data analytics, among others. Noncompliance can lead to serious sanctions and penalties. Legal data governance requirements must be satisfied, including those associated with the General Data Protection Regulation and the Insurance Distribution Directive in the European Union and the Health Insurance Portability and Accountability Act in the United States. Companies must also ensure compliance with industry standards such as Solvency II and FRS 103 and regulations established by the Prudential Regulatory Authority and Financial Conduct Authority in the United Kingdom. The strict regulation of insurance operations reflects society’s concern over the impact of insurance market failures on people’s welfare. When agents or managers are aligned with the interests of policyholders, information asymmetries may motivate dishonest behavior. The resulting risks are frequently non-diversifiable, thus exposing the whole group of policyholders to the consequences of the agency costs.

Within a consolidated insurance group, role delineation among the various companies is even more critical. Agents must answer to independent third parties to help ensure that service provision and performance maintain an acceptable standard of fairness. Therefore, the operating insurance agent should be perceived as a neutral provider of risk transfer that substitutes for the state volume of risk transfers in return for a fee.

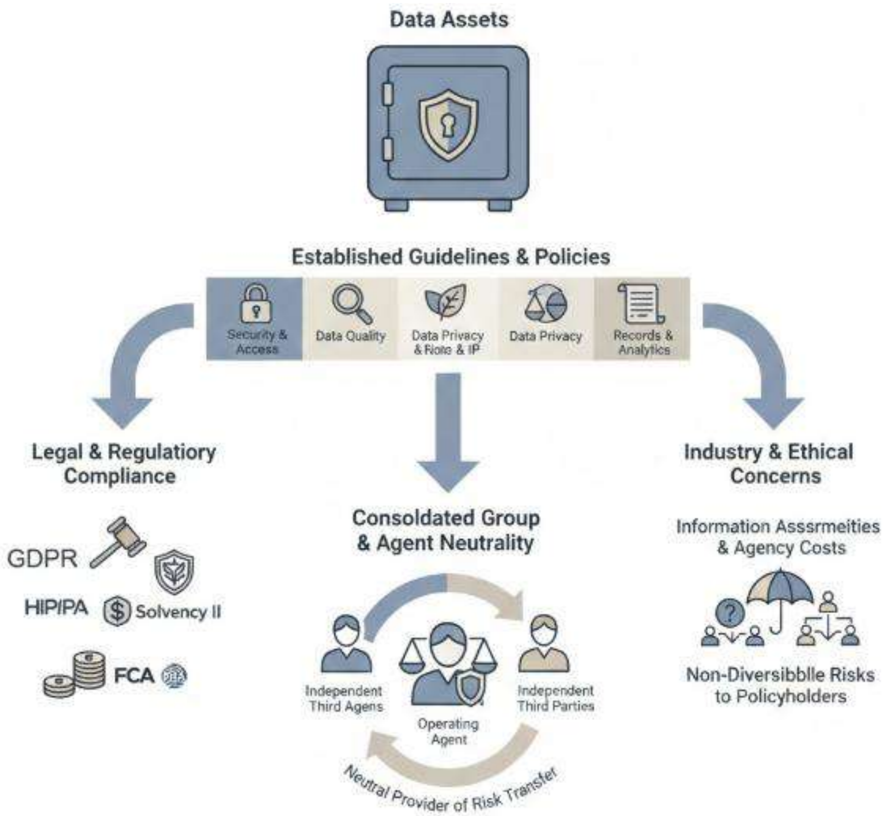

Fig 6.2: Integrated Data Governance in Consolidated Financial Groups: Mitigating Information Asymmetries and Agency Costs through Regulatory Compliance and Neutral Agent Frameworks

6.3.2. Fairness, Accountability, and Transparency

Fairness, accountability, and transparency (FAT) represent the most prominent pillars of responsible AI research. FAT safeguards deal with inappropriate discrimination, harmful decision-making, and discrimination between consumers. The first safeguard warrants reasonable expectations concerning the outcomes of the agent's actions and decisions in order to afford consumers and stakeholders the reassurance that data collection, storage, and usage conform to legal requirements. For example, insurance companies are required to list precisely the factors determining an agent's decision to reject a loan. Worded differently, FAT seeks to meet a vulnerable and disadvantaged position of the consumer.

A corollary safeguard, correctHD6, focuses on the need to employ data that is correct and free of errors, inconsistencies, and biases because these qualities are essential to preventing existing biases from propagating in the training set. However, the non-discriminatory training requirement is an extremely high bar in practice. Hence the common alternative is to seek a small level of bias in the predictions, a trade-off that—when correctly applied—does not violate the FAIR principle. Furthermore, it also demands that agents have constant and reproducible states, yet the current approaches for some models, such as reinforcement learning or GANs, do not guarantee reproducibility. Consequently, the concept of explainability has emerged as a serious research direction not only in academia but also in the industry, which is essential in underwriting decisions for banking or insurance companies—where it forms part of the specific implementation of FAT—for maintaining the trust of their consumers.

An entire ecosystem of accountably and ethically responsible AI services should be created around agencies or insurance companies, including tracking and recording links between any consumers and datasets used for training-driven decisions. In effect, it creates a set of consumers that must provide explicit consent when making use of the services provided by the agency or institution, valuing the partners' intellectual properties, intellectual control over the use of the data, and equally the possibility to monetise their IP when consumers are being served.

6.4. Architectural Paradigms for Agentic Insurance Systems

Two architectural paradigms are pertinent to the engineering of agentic AI systems for the insurance industry. The first deals with modularity and suggests that agentic systems should be built as a portfolio of utilitarian building blocks that link and blend, rather than as a closed suite of integrated applications. Such a style is particularly well-suited to the development of core insurance functions that lend themselves to design-as-a-service. The second concerns orchestration and the appropriate levels of autonomy and control for the constituent agents of a system. Both paradigms are equally relevant to the core capabilities of the insurance platform and to other capabilities that extend beyond the insuring economy into related domains.

Core insurance functions, such as underwriting and claims processing, require the processing of sensitive personal data in the decision-making pipeline. Even under the most benign and diligent supervision, their implementation will remain beholden to the biases of the training data, judgement and coding of the creators, and operational choices of those in command. Enterprise-wide insurance agents therefore need to respect a higher standard with respect to fairness, accountability and transparency for artificial intelligence (AI) specified by the OECD and the European Commission. Such demands are more properly satisfied by an expanded version of the transparent agent – an agent that not only disseminates explainability but also hosts privacy-preserving decomposition protocols to allow sensitive attributes to remain confidential and secure even from agent designers.

6.4.1. Modular Architectural Styles

The modular architectures of agentic AI systems determine how processing, learning, and knowledge are integrated within an agent in the context of Agentic computational models. Different modular approaches enable varying levels of autonomy. A single modular agent controls all agent processing, integrates all necessary learning, and has a single knowledge base. Control can be piqued between orchestrator and specialized components to enable a more complex behavior while maintaining a high level of control. The minimal orchestration style connects a planning agent with specialized agents. More relaxed levels of orchestration allow components with their own learning and memory capabilities to operate independently or autonomously while remaining within an orchestration-based environment.

Modular architectural styles encourage agentic agency through the married concepts of composability and autonomy. Agents are modular because they reside within a well-defined software boundary that separates internal from external processes. Internal components are data craft and information providers (sensors) in the case of agents in Contract role (see Contract and Design Task) or action executors and performance reports (effects) for agents in Agency role. External agents, such as the environment and the community, are also modular because they expose information and/or affordances and capability invocation on demand through well-defined communication interfaces. Agency autonomy incorporates the ability of an agent to exercise its own judgment, in absence of external control, to decide and/or execute actions and plot a course of action for the near future, ignoring its community.

6.4.2. Orchestration and Autonomy Levels

An agent-oriented approach to the architecture of insurance platforms necessitates consideration of the levels of autonomy that govern interaction between agents. While a platform might comprise autonomous agents, it may also deploy less intelligent, taskoriented agents that either complement collaborative tasks or interact on behalf of users or groups. Furthermore, a single platform agent might rely on multiple others to achieve a collective goal. These relationships can generally be described as either orchestration or delegated task execution. In orchestration, a supervisory agent maintains oversight while organizing the number and nature of agents constituting a collaboration. In delegation, a requesting agent (which might not be a fully autonomous agent) depends on others to fulfil components of the overarching task.

Autonomous agents interacting and collaborating with a high degree of task-orientation demand careful architecture that supports seamless exchange and processing in insurance context. While these agents might present in complementary and supporting roles, two key areas are underwriting and claims processing. Underwriting agents acquire, analyse, and interpret applicant data to compute risk with respect to individually identified insured assets or activities. Claims processing agents undertake a similar yet inverted task, analysing and assessing claims and interpreting customer data in relation to the proposed claim. These agents represent two of the higher frequency multiplex tasks that could sustain agent-prompted collaboration over the lifecycle of a group or community.

6.5. Agent Design for Core Insurance Tasks

Agent design directly determines how adequately, fairly, and efficiently an insurance platform's core computational tasks are performed. This section outlines the essential skill sets required of agents performing either underwriting or claims processing tasks that are designed for these core insurance task roles—agents that serve as primary service providers or their immediate orchestration dependents. Core skill requirements are enumerated, followed by a proposal for an agent design framework that comprehensively captures each specified requirement.

Underwriting agents, tasked with assessing the risk level for insurance applicants and determining appropriate premium amounts, must possess complex predictive and explanatory modeling capabilities to simultaneously estimate the likelihood of adverse future events for various applicant groups, and precisely describe the causal relationships underlying those predictive patterns. Claims processing agents, responsible for determining the legitimacy of an event-based claim and the compensation amount for a legitimate claim, likewise require sophisticated predictive and explanatory modeling skills to assess the likelihood of an insurance event for various historical groups with respectful considered features.

6.5.1. Underwriting Agents

An insurance platform must offer independent access to a wide range of insurance policies. Insurers need to consider variables not under their control or policies not aligned with their portfolio. Such insurance platforms operate like aggregators. They are a source of demand for multiple underwriting platforms, allowing users to gather quotes from several insurance services in a single place. The difference with traditional aggregators is that they do not connect a single risk to different offers but allow a user to evaluate different risks using different policies.

An agent can optimize the insurance offered to a user by constructing a risk-transferred portfolio that minimizes the cost or maximizes the associated utility while, at the same time, analyzing issues that agents might not take into account, such as a lack of diversification or overlap between policies. An underwriting agent analyzes a user's risk portfolio and requests quotes from different underwriting agents or conducting a search-and-replace operation to modify an existing portfolio.

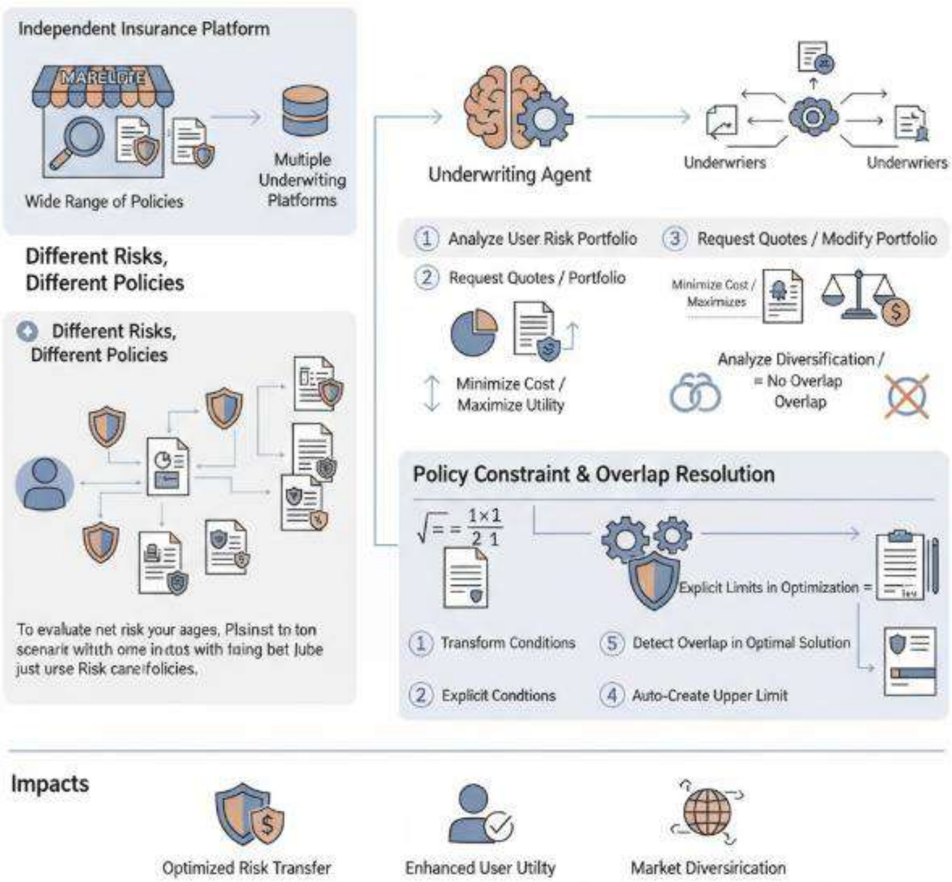

Fig 6.3: Agent-Based Portfolio Optimization in Multi-Risk Insurance Platforms: Mitigating Coverage Overlap Through Constrained Underwriting Logic

The underwriting agent can then add limits on the policies to avoid overlapping coverage. Constraining insurance policies during optimization is possible by transforming the conditions of the problem. If the optimization explicitly includes limits, it can provide a different solution without overlap. Otherwise, it can detect overlapping elements in the optimal solution and automatically create an additional upper limit for one of the risks.

6.5.2. Claims Processing Agents

The primary objective of the claims processing phase is to settle claims as efficiently and equitably as possible. This task can be supported by agentic technology using both supervised learning for high-volume low-complexity tasks and agentic systems for the most complex cases. Intelligent agents equipped with natural language processing capabilities can typically pre-process the claim when received. The agent can parse the claim, identifying special facts from natural language descriptions, determine which counterparties were involved in the situation, check for preliminary indications that the case may be fraudulent, and identify preliminary indicators of loss adjustment costs.

In most situations, however, the level of complexity is too high, and a dedicated evaluation by a claims expert is required. For these cases, the probability that a claims expert will be involved can be predicted via supervised learning. If the model predicts the probability of involvement to be higher than a threshold, an intelligent agent can direct the case automatically to a claims expert able to take full charge of the operation. If automatic redirection is not applicable and the model predicts a low likelihood of claims expert involvement, a human-free pipeline can be executed to compare the claim with past similar incidents. The goal of this comparison is highly efficient and accurate settlement of high-volume low-complexity claims.

6.6. Data Infrastructure for Agentic Systems

In agentic insurance systems, agents that perform core tasks—such as underwriting and claims processing—are responsible for preparing data destined for control agents. However, systematic guarantees of data provenance and quality over long time horizons and across multiple organizations supporting a common task are essential to engender trust not only amongst task-supporting agents but also between task and control agents. A combination of data quality monitoring tools or pipelines, together with information sharing and processed data validation protocols, would help imbue these agents with the required levels of trust for them to operate effectively. Additionally, privacy-preserving techniques also contribute to enhancing trust, particularly when sensitive customer data are processed.

The need for data-quality assurance mechanisms has become increasingly pressing with the rise of generative AI and related technologies, which have resulted in Data Quality and AI Factoring as new industries on their own. Nevertheless, ensuring high-quality input and output data remains an ongoing challenge. Insurance platforms are well positioned to build infrastructure that guarantees the integrity of input data supplied to underwriting agents and the quality of the data produced as output by claims processing agents.

6.6.1. Data Provenance and Quality Assurance

Data provision is crucial for agentic intelligence, whose success relies on high-quality data. Data provenance should, therefore, be integrated into the insurance platform architecture. The provenance of all data used in machine learning training, inference, and testing phases should be recorded to achieve data quality monitoring, assurance, and improvement. Annotating data with quality and provenance metadata permits effective data governance by facilitating the detection of data drift, vulnerability to hostile attacks, and reproducibility of results. Quality rules can be defined by exploiting subject matter experts' knowledge or leveraging agreed standards related to insurance-specific data. Automated data quality monitoring and reporting tools can assure data quality for critical algorithmic decision-making that require compliance with regulations such as GDPR and fairness, accountability, and transparency principles.

Provenance information can also be scrutinized to help detect attacks and manipulations. Chained data provenance provides a conceptual model to rationally defend against web attacks in general and especially against poisoning attacks against machine learning models. As data provision is a highly demanding task, data quality-related knowledge can also be integrated into reinforcement learning training processes to reward agents for providing the correct or appropriate data required. Such techniques can improve data quality and training performance while reducing the amount of data samples needed for effective training.

6.6.2. Privacy-Preserving Computation

The potential for revealing sensitive information about individuals or organizations while obtaining insights from data for responsible business decision-making or regulatory compliance is a crucial issue for insurance platforms as they design their data infrastructures. Therefore, agents that perform data science on sensitive data need to employ cryptographic solutions in order to preserve the privacy of the involved parties.

The most common solutions in this domain require the use of homomorphic encryption, multi-party computation or differential privacy. Homomorphic encryption enables the computation of encrypted data without the need for decryption using a secret key, hence offering privacy guarantee for the parties that outsource their data. Multi-party computation allows a group of parties to jointly compute a function while keeping the input of other parties private, through secure sharing of the input data. Differential privacy introduces a mathematically principled means of sharing and releasing information from a database such that the risk of unwarranted disclosure about an individual in the database is limited. These solutions, however, come with a non-negligible computational and communication overhead.

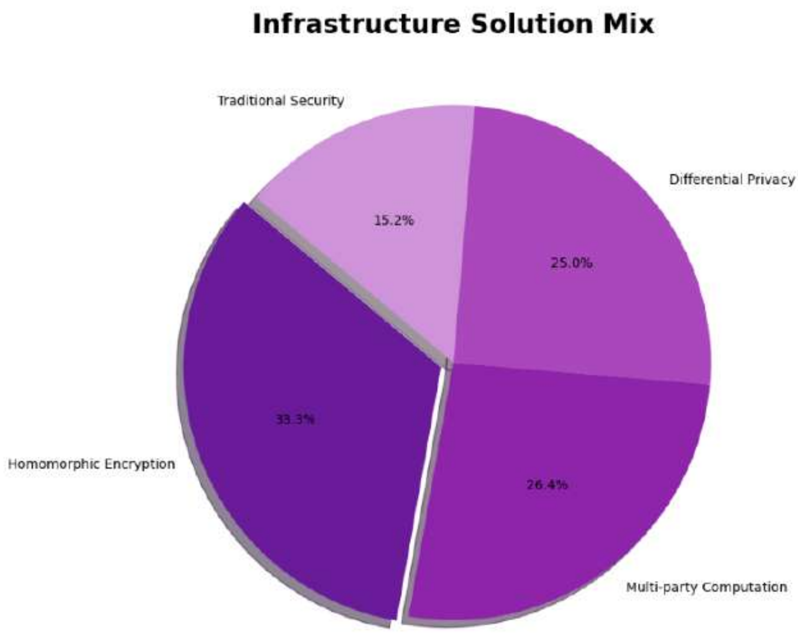

Fig 6.4: Infrastructure Solution Mix

6.7. Conclusion

In sum, agentic AI architectures for insurance platforms must provide suitable environments for agents able to autonomously execute tasks such as underwriting and claims processing. Such agents act directly and are therefore architecturally closer to the data sources they require than to the organizations that manage them. Architectures must not only offer a level of autonomy suited to the tasks at hand, but also provide the required orchestration of activities performed across multiple agents. Guidance from these perspectives, together with a focus on data governance, quality, privacy, and compliance with fairness, accountability, and transparency principles, serves to reduce the design space.

Future work will refine the treatment and explore concrete representations of agents responsible for major insurance tasks, consider the design, deployment, and control of an agentic AI platform for the insurance industry, explore how many of the considerations address tasks and workflows relevant to other domains, and examine how these aspects shape data infrastructures in a more general context.

6.7.1. Final Thoughts and Future Directions

Looking ahead, agentic capabilities appear suitable for directing insurance platforms to meet their objectives while addressing the concerns raised. Indeed, agentic systems are defined by an architecture that allows constituent agents to operate with greater levels of autonomy and responsibility. Design choices yielding a more distributed and modular system could provide new dimensions of agentic capability that were not considered in the analysis. At a lower level, the specific infrastructure requirements involved in realizing agentic systems were shown to align strongly with common areas of concern across the insurance sector, suggesting an opportunity for most platform providers to invest concurrently in agentic capacity.

With the importance of agentic capability for insurance systems established, further research should examine how agentic insurance platforms can be engineered in practice. Research should therefore address how autonomy and agency are distributed across constituent agents, especially in the context of greater orchestration-oriented slogans (e.g., by looking at orchestration mechanisms with higher levels of abstraction). Auxiliary responsibilities such as development support can also be abstracted through the use of proper modules.

References

Russell, S., & Norvig, P. (2021). Artificial intelligence: A modern approach (4th ed.). Pearson.

Kolla, S. H. (2021). Rule-Based Automation for IT Service Management Workflows. Online Journal of Engineering Sciences, 1(1), 1–14. Retrieved from https://www.scipublications.com/journal/index.php/ojes/article/view/1360

Wooldridge, M. (2020). An introduction to multiagent systems (2nd ed.). Wiley.

Keerthi Amistapuram. (2023). Privacy-Preserving Machine Learning Models for Sensitive Customer Data in Insurance Systems. Educational Administration: Theory and Practice, 29(4), 5950–5958. https://doi.org/10.53555/kuey.v29i4.10965

Sutton, R. S., & Barto, A. G. (2020). Reinforcement learning: An introduction (2nd ed.). MIT Press.

Aitha, A. R. (2022). Cloud Native ETL Pipelines for Real Time Claims Processing in Large Scale Insurers. Available at SSRN 5532601.

Nagubandi, A. R. (2025). Advanced Predictive Autonomous Agents for Multiportfolio Risk Analytics and Real-Time Enterprise P&L Decisioning: Self-Learning AI Systems for Multi-counterparty Derivatives, Collateral Valuation, and Accounting Reconciliation. Collateral Valuation, and Accounting Reconciliation (December 01, 2025).

Eling, M., Nuessle, D., & Staubli, J. (2022). The impact of artificial intelligence along the insurance value chain and on the workforce. Journal of Risk and Insurance, 89(2), 1–38.

Gottimukkala, V. R. R. (2022). Licensing Innovation in the Financial Messaging Ecosystem: Business Models and Global Compliance Impact. International Journal of Scientific Research and Modern Technology, 1(12), 177-186.

Blier-Wong, C., Cossette, H., & Marceau, E. (2021). Tree-based machine learning models for insurance ratemaking. North American Actuarial Journal, 25(3), 383–407.

Varri, D. B. S. (2022). A Framework for Cloud-Integrated Database Hardening in Hybrid AWS-Azure Environments: Security Posture Automation Through Wiz-Driven Insights. International Journal of Scientific Research and Modern Technology, 1(12), 216-226.

Charpentier, A., Denuit, M., & Trufin, J. (2021). Explainable machine learning in insurance pricing. Scandinavian Actuarial Journal, 2021(7), 565–594.

Guntupalli, R. (2025, June). Federated Learning in Cloud AI: Enhancing Privacy and Security. In International Conference on Data Analytics & Management (pp. 435-443). Cham: Springer Nature Switzerland.

Xu, Y., Sun, J., & Liu, J. (2023). Fairness-aware machine learning for insurance risk prediction. IEEE Access, 11, 24501–24513.

Wüthrich, M. V., & Merz, M. (2022). Statistical foundations of actuarial learning and its applications. Springer.

Bommasani, R., Hudson, D. A., Adeli, E., et al. (2022). On the opportunities and risks of foundation models. Stanford Institute for Human-Centered Artificial Intelligence.

Bai, T., Zheng, Z., Ren, K., & Shi, S. (2024). Cloud-native machine learning systems: Architecture and optimization. IEEE Software, 41(1), 50–58.

Garapati, R. S. (2022). AI-Augmented Virtual Health Assistant: A Web-Based Solution for Personalized Medication Management and Patient Engagement. Available at SSRN 5639650.

Amershi, S., Begel, A., Bird, C., et al. (2021). Software engineering for machine learning: A case study. IEEE Transactions on Software Engineering, 47(12), 2913–2932.

Uday Surendra Yandamuri. (2022). Cloud-Based Data Integration Architectures for Scalable Enterprise Analytics. International Journal of Intelligent Systems and Applications in Engineering, 10(3s), 472–483. Retrieved from https://ijisae.org/index.php/IJISAE/article/view/8005

Richman, R., & Wüthrich, M. V. (2021). A neural network extension of the classical chain-ladder model. Insurance: Mathematics and Economics, 99, 331–347.

Vadisetty, R., Polamarasetti, A., Goyal, M. K., Rongali, S. K., kumar Prajapati, S., & Butani, J. B. (2025, May). Generative AI for Creating Immersive Learning Environments: Virtual Reality and Beyond. In 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC) (pp. 1-5). IEEE.

Serrano-Guerrero, J., Herrera-Viedma, E., & Olivas, J. A. (2021). Fuzzy and machine learning models in insurance risk classification. Applied Soft Computing, 102, 107070.

Ramesh Inala. (2023). Big Data Architectures for Modernizing Customer Master Systems in Group Insurance and Retirement Planning. Educational Administration: Theory and Practice, 29(4), 5493–5505. https://doi.org/10.53555/kuey.v29i4.10424

Henckaerts, R., Antonio, K., Côté, M.-P., & Verbelen, R. (2021). Boosting insights in insurance tariff plans with tree-based machine learning methods. North American Actuarial Journal, 25(2), 255–285.

Unifying Data Engineering and Machine Learning Pipelines: An Enterprise Roadmap to Automated Model Deployment. (2023). American Online Journal of Science and Engineering (AOJSE) (ISSN: 3067-1140) , 1(1). https://aojse.com/index.php/aojse/article/view/19

Zhang, Z., Li, J., & Wang, H. (2024). Dynamic pricing strategies based on reinforcement learning for insurance markets. Expert Systems with Applications, 235, 121083.

Floridi, L., Cowls, J., Beltrametti, M., et al. (2018). AI4People—An ethical framework for a good AI society. Minds and Machines, 28(4), 689–707.