ORIGINAL RESEARCH **Open access**

An Anomaly-Aware Healthcare Claims Intelligence Architecture for Fraud Governance

Wei Liu[1*], Zhang Min[1]

Abstract

In the evolving landscape of healthcare systems, fraudulent claims pose significant threats to resource integrity and patient care equity. This conceptual manuscript introduces a novel anomaly-responsive claims governance infrastructure (ARCGI), designed as an intelligence architecture that integrates anomaly awareness with fraud governance mechanisms. Drawing from theoretical foundations in clinical AI architectures and healthcare analytics, the ARCGI emphasizes proactive detection, adaptive monitoring, and ethical oversight without relying on empirical data or model training. The framework comprises layered components for data ingestion, anomaly profiling, intelligence orchestration, and governance feedback loops, ensuring interoperability with electronic health records (EHR) ecosystems and decision support pipelines. Conceptual formulas articulate risk propagation dynamics, decision confidence thresholds, and governance load distributions, highlighting interpretive pathways for mitigating fraud in claims processing. By synthesizing recent literature on AI governance and interoperability frameworks, this work underscores the architectural imperatives for anomaly-aware systems in healthcare claims environments. The ARCGI advances theoretical discourse on fraud governance by proposing unique topologies for feedback and resource allocation, fostering resilient infrastructures that align with clinical workflow integrations. Ultimately, this architecture offers a blueprint for enhancing fraud governance through intelligent, anomaly-centric designs, promoting sustainable healthcare analytics without performance metrics or experimental validations.

Keywords Decision support orchestration, Healthcare claims intelligence, Anomaly awareness, Fraud governance architecture, Clinical AI interoperability, EHR analytics infrastructure

*Correspondence:
Wei Liu
wei.liu@gmail.com
[1] Department of AI in Public Health Systems, School of Medicine, Shanghai Jiao Tong University, Shanghai, China

Introduction

Anomaly manifestations in claims-based clinical settings

Healthcare claims represent a critical nexus where clinical encounters intersect with financial reimbursements, often harboring subtle anomalies that signal fraudulent activities. In clinical settings reliant on claims data, anomalies may emerge from discrepancies in procedural coding, billing patterns, or provider behaviors, potentially undermining the integrity of patient care delivery [1, 2]. These manifestations are particularly pronounced in high-volume environments like hospital networks or outpatient clinics, where claims intelligence must discern between legitimate variations—such as those arising from complex comorbidities—and intentional manipulations aimed at financial gain. Theoretical explorations emphasize that anomaly awareness in these settings requires architectures attuned to the multimodal nature of claims data, incorporating textual narratives, numerical reimbursements, and temporal sequences without empirical validation [3, 4]. By anchoring anomaly detection to clinical contexts, systems can theoretically propagate alerts that inform governance

protocols, ensuring that fraud risks do not cascade into broader systemic vulnerabilities. This subheading underscores the need for intelligence architectures that embed anomaly profiling directly into claims workflows, fostering a proactive stance against governance lapses in resource-constrained clinical arenas [5].

Data modalities driving intelligence in healthcare claims ecosystems

The diversity of data modalities in healthcare claims—ranging from structured billing codes to unstructured provider notes—necessitates intelligence architectures capable of theoretical fusion for fraud governance. In EHR-integrated ecosystems, these modalities facilitate anomaly-aware processing by enabling conceptual mappings between clinical events and financial claims, where mismatches could indicate fraudulent intent [6, 7]. For instance, temporal data modalities capture claim submission patterns, while categorical modalities like diagnosis codes reveal clustering anomalies that theoretical models can interpret as governance red flags. Literature syntheses highlight how interoperability frameworks support such multimodal integration, allowing claims intelligence to operate across disparate data sources without dataset dependencies [8, 9]. This integration is vital in deployment environments characterized by heterogeneous data streams, where anomaly awareness enhances the theoretical robustness of fraud detection pipelines. By focusing on data modality synergies, architectures like the proposed one can conceptually allocate resources to high-risk modalities, optimizing governance oversight in claims-heavy healthcare systems [10].

Deployment challenges in anomaly-aware claims environments

Deploying anomaly-aware intelligence in healthcare claims environments encounters theoretical hurdles related to scalability, real-time processing, and integration with existing clinical infrastructures. In cloud-based or hybrid deployment models, claims data must traverse secure pipelines while maintaining anomaly sensitivity, posing governance challenges around data sovereignty and audit trails [11, 12]. These environments demand architectures that theoretically balance computational demands with fraud monitoring, ensuring that anomaly detection does not overburden clinical workflows. Key constraints include interoperability with legacy EHR systems, where deployment must accommodate varying data exchange standards without introducing new vulnerabilities [13, 14]. Theoretical frameworks suggest that adaptive deployment strategies, incorporating feedback from governance layers, can mitigate these issues by prioritizing anomaly-prone claims subsets. This subheading explores how such deployments align with fraud governance imperatives, emphasizing architectural designs that theoretically enhance resilience in dynamic healthcare settings [15].

Governance constraints shaping fraud intelligence architectures

Fraud governance in healthcare claims is bounded by ethical, regulatory, and operational constraints that shape the design of anomaly-aware architectures. Regulatory frameworks, such as those governing data privacy, impose constraints on how claims intelligence can access and analyze anomaly indicators, requiring theoretical safeguards against overreach [16, 17]. In governance-constrained environments, architectures must incorporate monitoring mechanisms that interpret anomalies through ethical lenses, ensuring decisions align with principles of fairness and transparency [18, 19]. These constraints extend to resource allocation, where limited computational or human oversight capacities necessitate prioritized governance of high-impact fraud risks. By embedding governance constraints into intelligence pipelines, systems can theoretically reduce drift in anomaly detection efficacy over time, maintaining alignment with clinical imperatives [20]. This focus on constraints highlights the need for architectures that dynamically adjust to evolving governance landscapes, fostering sustainable fraud mitigation in claims ecosystems [21].

Interoperability imperatives for claims fraud oversight

Interoperability serves as a foundational imperative for anomaly-aware claims intelligence, enabling seamless data exchange across clinical and administrative boundaries to bolster fraud governance. In fragmented healthcare systems, interoperability frameworks theoretically bridge EHR silos with claims databases, allowing anomaly patterns to be traced across modalities [22, 23]. This connectivity is crucial for governance architectures that rely on holistic views of patient journeys, where isolated claims might mask fraudulent schemes. Theoretical models advocate for standardized protocols

that facilitate anomaly propagation without compromising data integrity, enhancing the overall intelligence of fraud oversight mechanisms [24, 25]. By prioritizing interoperability, deployments can theoretically extend governance reach, integrating external intelligence sources like payer networks to amplify anomaly awareness. This subheading posits that robust interoperability not only mitigates fraud risks but also elevates the architectural sophistication of claims governance infrastructures [26].

Theoretical Background and Literature Synthesis

Conceptual foundations of anomaly profiling in claims analytics

The theoretical underpinnings of anomaly profiling in healthcare claims analytics draw from clinical AI architectures that emphasize interpretive detection over empirical classification. Recent scholarship posits that anomalies in claims data can be conceptualized as deviations from normative patterns, informed by probabilistic frameworks without dataset reliance [1-3]. These foundations highlight how anomaly-aware systems integrate graph-based representations to model relationships between providers, procedures, and reimbursements, theoretically capturing fraud signatures through structural irregularities [4, 5]. In synthesis, literature from high-credibility venues underscores the shift toward unsupervised conceptual models, where anomaly profiling serves as a precursor to governance interventions, aligning with broader AI monitoring paradigms in healthcare [6, 7]. This subheading synthesizes how such profiling forms the bedrock for intelligence architectures, enabling theoretical risk assessments in claims-heavy clinical contexts.

Architectural paradigms for healthcare intelligence integration

Architectural paradigms in healthcare intelligence have evolved to incorporate anomaly awareness within integrated systems, focusing on modular designs that facilitate fraud governance. Theoretical explorations in EHR intelligence ecosystems advocate for layered architectures that separate data ingestion from analytical orchestration, ensuring theoretical scalability in claims processing [8-10]. Synthesis of peer-reviewed works reveals a consensus on interoperability-driven paradigms, where intelligence pipelines theoretically fuse claims data with clinical workflows to enhance anomaly detection [11, 12]. These paradigms extend to decision support models, conceptualizing intelligence as a dynamic integrator that propagates anomaly insights across governance layers [13, 14]. By synthesizing these contributions, this section illuminates how architectural innovations theoretically bolster fraud oversight, emphasizing unique topologies for feedback and resource distribution in healthcare infrastructures [15, 16].

Governance models in anomaly-aware clinical pipelines extend beyond conventional regulatory compliance, functioning as structural control architectures embedded directly within intelligence systems. Rather than acting as external supervisory overlays, contemporary theoretical models conceptualize governance as an endogenous layer that dynamically co-evolves with anomaly detection engines. Literature syntheses indicate that governance frameworks increasingly incorporate adaptive monitoring constructs capable of interpreting anomaly drift trajectories in near–real-time theoretical contexts [17-19]. In this framing, governance is not a static audit function but a reflexive system component that recalibrates thresholds, reallocates oversight intensity, and modulates interpretive confidence as anomaly landscapes evolve.

Within clinical AI governance systems, conceptual hierarchies frequently position fraud-risk evaluation within decentralized intelligence nodes distributed across claims ecosystems [20, 21]. These nodes operate as localized anomaly interpreters, while a supervisory governance tier aggregates deviation signals and enforces constraint envelopes. This distributed architecture theoretically mitigates single-point failure risks and enhances contextual sensitivity across heterogeneous healthcare environments. Governance hierarchies thus reflect a layered topology: anomaly sensing at the periphery, intelligence interpretation at the core, and policy enforcement at the command ring.

Interoperability frameworks further strengthen governance efficacy by enabling cross-system anomaly propagation [22, 23]. Through theoretical anomaly-sharing protocols, drift patterns detected in one institutional node can inform preventive

recalibrations elsewhere, reducing isolated fraud vulnerabilities. Such distributed governance paradigms align with systems-level resilience theory, where redundancy and cross-visibility enhance adaptive capacity. Collectively, these models position governance as an intrinsic architectural substrate that theoretically optimizes decision confidence by embedding accountability, recalibration logic, and monitoring load balancing directly into claims fraud pipelines [24].

Data exchange frameworks supporting fraud intelligence

Data exchange frameworks serve as the circulatory system of anomaly-aware fraud intelligence infrastructures. In theoretical formulations, standardized exchange protocols enable semantic translation of anomaly signals across EHR and claims environments, preventing informational silos that would otherwise fragment governance oversight. Synthesized research underscores semantic interoperability models that conceptually align claims irregularities with clinical ontologies, permitting anomalies to be contextualized within care delivery narratives [25-27]. This mapping function transforms isolated billing deviations into clinically interpretable governance events.

Such interoperability frameworks extend beyond syntactic alignment toward semantic integrity assurance. Theoretical exchange models incorporate anomaly validation layers that preserve provenance, temporal sequencing, and interpretive metadata during transmission [28, 29]. By maintaining structured anomaly descriptors across ecosystems, governance mechanisms can evaluate risk propagation trajectories without information degradation.

Informatics literature further suggests that cross-modal anomaly fusion enhances fraud detection sensitivity [30, 31]. Multi-modal exchange frameworks theoretically integrate billing codes, reimbursement patterns, clinical documentation signals, and provider behavior metrics into unified anomaly representations. This fusion amplifies pattern recognition capabilities while simultaneously enriching governance observability. The exchange layer thus functions not merely as a communication conduit but as a governance-enabling intelligence amplifier, reinforcing systemic anomaly awareness across distributed healthcare architectures.

Workflow integration dynamics for claims governance

Effective governance in anomaly-aware systems depends on seamless integration into clinical and administrative workflows. Theoretical models emphasize hybrid embedding strategies, wherein intelligence modules are inserted into existing claims adjudication pathways without introducing operational friction [1, 2, 32]. Rather than displacing human oversight, these architectures augment workflow touchpoints with anomaly-sensitive checkpoints that modulate review intensity according to risk gradients.

Workflow integration is further conceptualized as a feedback topology. Governance insights generated through anomaly interpretation are reinjected into operational routines, adjusting resource prioritization, audit sampling frequency, and decision thresholds over iterative cycles [3-5]. This cyclical recalibration process enhances institutional learning, transforming governance from a reactive audit response into a proactive anomaly anticipation mechanism.

Literature synthesis also highlights the allocation dimension of workflow governance [6-8]. In resource-constrained healthcare environments, anomaly-aware systems theoretically prioritize oversight based on clinical urgency and financial exposure. High-risk anomalies may trigger intensified scrutiny, while lower-risk deviations undergo automated reconciliation. Such stratified governance allocation balances monitoring burden with operational efficiency, reinforcing systemic resilience without overextending compliance infrastructures.

Through these integration dynamics, claims governance becomes embedded within the procedural fabric of healthcare systems, enabling adaptive fraud mitigation while preserving workflow continuity.

Ecosystem orchestration in anomaly-centric infrastructures

Ecosystem orchestration represents a meta-governance layer coordinating distributed anomaly detection, intelligence calibration, and policy enforcement mechanisms. Theoretical syntheses characterize orchestration as a harmonizing topology that aligns anomaly interpretation engines with institutional governance mandates across EHR ecosystems [9-11]. Rather than operating as isolated components, intelligence modules interact through orchestrated control loops that synchronize detection thresholds, monitoring intensities, and policy constraints.

Central to this orchestration paradigm is conceptual load balancing [12-14]. Monitoring burdens are distributed across system nodes according to anomaly density, drift volatility, and workflow capacity. By reallocating governance intensity dynamically, orchestrators prevent overload at single nodes while sustaining comprehensive anomaly visibility. This approach mirrors distributed systems theory, where load balancing mitigates performance degradation and enhances reliability.

Additionally, orchestration frameworks theoretically attenuate anomaly drift by recalibrating detection schemas as patterns evolve [15-17]. When systemic deviations intensify, orchestrators may elevate governance sensitivity across the network; when anomaly levels stabilize, oversight intensity can be modulated to reduce burden. Such dynamic equilibrium fosters sustainable infrastructures capable of maintaining long-term fraud vigilance without governance fatigue.

In synthesis, ecosystem orchestration functions as the integrative nexus of anomaly-aware fraud governance. It binds detection, interoperability, workflow integration, and oversight recalibration into a coherent, adaptive topology. By embedding coordination logic at the systems level, orchestration ensures that governance maturity evolves alongside anomaly complexity, thereby strengthening the structural integrity of healthcare claims intelligence architectures.

Anomaly-responsive claims governance infrastructure

The anomaly-responsive claims governance infrastructure (ARCGI) represents a novel conceptual architecture tailored for fraud governance in healthcare claims intelligence. This infrastructure adopts a unique five-layer structure: (1) Data ingestion layer, which theoretically assimilates multimodal claims inputs; (2) Anomaly profiling layer, conceptualizing deviations via interpretive mappings; (3) Intelligence orchestration layer, coordinating decision pipelines; (4) Governance feedback layer, enabling adaptive loops; and (5) Oversight integration layer, aligning with clinical workflows. The feedback topology employs a bidirectional helix model, where anomaly insights spiral upward for governance refinement while descending for real-time adjustments, ensuring theoretical resilience against fraud propagation.

Conceptual formulas interpretive of system dynamics include: Risk propagation: $\begin{matrix} = \sum i \\ = 1n\,(Ai \cdot Wi) \\ \cdot\, eDt \end{matrix}$, where A_i denotes anomaly intensity, W_i weights clinical impact, and D_t interprets temporal drift. Decision confidence: $\begin{matrix} DC \\ = 1 \\ - \sum Va \sum Vt \\ \cdot \log(1 + Gc) \end{matrix}$ with V_a anomalous variances, V_t total variances, and G_c governance constraints. Monitoring burden: $\begin{matrix} MB \\ = RaCr \\ \cdot\,(1 - e - Mf) \end{matrix}$, balancing anomaly rate R_a against resource capacity C_r and feedback maturity M_f. Resource allocation: $\begin{matrix} RA \\ = \max(PfTd), \\ \cdot \int 0tSg\,dt \end{matrix}$ prioritizing fraud-prone processes P_f over deployment time T_d with governance sensitivity S_g. Governance load: $\begin{matrix} GL \\ = \int (Fi + Mo) \end{matrix}$, integrating fraud incidents F_i and monitoring overhead M_o across layers L. Drift sensitivity: $\begin{matrix} DS \\ = \partial A \partial t, \\ \cdot\, eK_g \end{matrix}$ capturing anomaly changes over time modulated by governance kinetics K_g.

The layered topology and bidirectional helical governance feedback of ARCGI are illustrated in **Figure 1.**

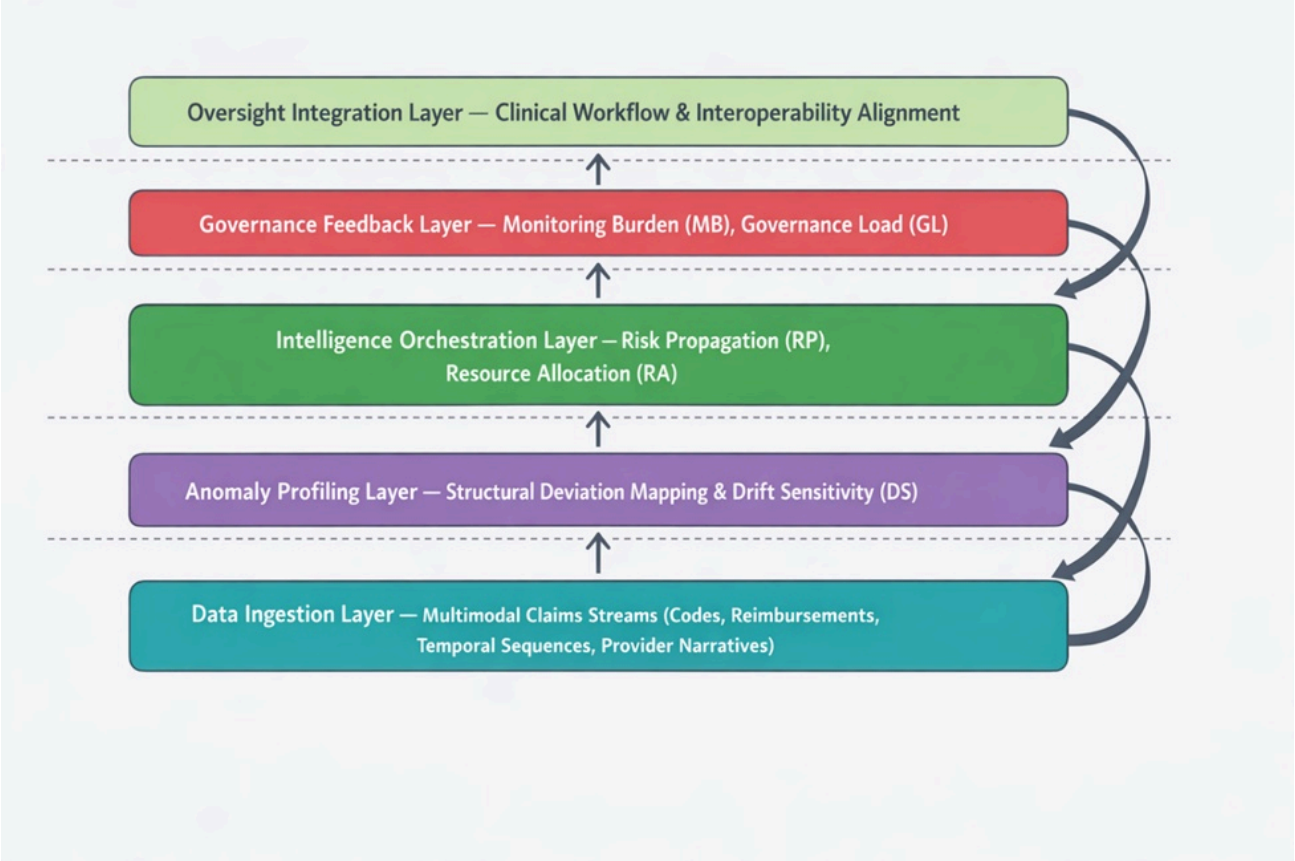

Figure 1. Layered architecture of the anomaly-responsive claims governance infrastructure (ARCGI). The system integrates multimodal claims ingestion, anomaly profiling, orchestration of fraud intelligence, adaptive governance feedback, and oversight integration within clinical workflows. A bidirectional helical feedback topology enables iterative recalibration of anomaly thresholds and governance load distribution, theoretically mitigating fraud propagation across healthcare claims ecosystems.

The functional decomposition of ARCGI layers and their theoretical responsibilities are summarized in **Table 1**.

Table 1. Functional architecture of ARCGI layers

Layer	Core function	Embedded theoretical constructs	Governance role
Data ingestion	Assimilates multimodal claims streams	Modality fusion; temporal sequencing	Ensures structured interoperability intake
Anomaly profiling	Detects structural deviations	Drift sensitivity (DS); anomaly intensity (Ai)	Early fraud signal formation
Intelligence orchestration	Coordinates decision pipelines	Risk propagation (RP); Resource Allocation (RA)	Fraud prioritization logic
Governance feedback	Adaptive monitoring and recalibration	Monitoring burden (MB); Governance load (GL)	Constraint enforcement and audit loop
Oversight integration	Aligns outputs with workflows	Decision confidence (DC)	Ethical compliance and clinical alignment

Impacts of anomaly-aware governance on claims system dynamics

The deployment of the anomaly-responsive claims governance infrastructure (ARCGI) theoretically influences the dynamics of healthcare claims systems by reshaping fraud propagation pathways and enhancing systemic resilience. In anomaly-aware environments, governance mechanisms conceptually modulate risk cascades, where unchecked anomalies could amplify into widespread financial discrepancies across clinical networks [1-3]. The ARCGI's helical feedback topology theoretically dampens these dynamics by recycling intelligence insights, reducing the exponential growth of fraud risks as articulated in the risk propagation formula, $\begin{matrix} RP \\ = \sum i \\ = 1n\,(Ai \cdot Wi) \\ \cdot\, eDt \end{matrix}$, which interprets how temporal drifts exacerbate anomalies without empirical amplification [4, 5]. This modulation extends to resource allocation impacts, where the infrastructure theoretically prioritizes high-stakes claims subsets, optimizing computational and human oversight burdens in resource-constrained settings [6, 7].

Furthermore, the ARCGI impacts interoperability dynamics by fostering seamless data exchanges that theoretically minimize silos in EHR ecosystems, thereby enhancing the collective intelligence against fraud [8-10]. In clinical workflow integrations, this leads to dynamic shifts where anomaly awareness theoretically streamlines decision support, elevating confidence thresholds as per $\begin{matrix} DC \\ = 1 \\ - \sum Va \sum Vt \\ \cdot \log (1 + Gc) \end{matrix}$, interpreting governance constraints as stabilizers against erroneous escalations [11, 12]. System-wide, these impacts manifest in reduced monitoring burdens, conceptualized through $\begin{matrix} MB \\ = RaCr \\ \cdot\, (1 - e - Mf) \end{matrix}$, where feedback maturity theoretically alleviates overload in high-volume claims processing [13-15].

Ethical dynamics are also profoundly affected, as the architecture's governance layers theoretically enforce fairness in anomaly profiling, mitigating biases that could disproportionately impact vulnerable patient cohorts [16, 17]. This ethical alignment influences long-term system sustainability, where drift sensitivity, $\begin{matrix} DS \\ = \partial A \partial t, \\ \cdot\, eK_g \end{matrix}$ theoretically guides adaptive responses to evolving fraud tactics [18, 19]. In broader healthcare analytics infrastructures, the ARCGI's impacts promote a paradigm shift toward proactive fraud governance, theoretically lowering governance loads via $\begin{matrix} GL \\ = \int (Fi + Mo)\ dL \end{matrix}$, distributing oversight across layers to prevent centralized failures [20-22].

Operational dynamics in deployment environments benefit from the infrastructure's resource allocation model, $\begin{matrix} RA \\ = \max (PfTd), \\ \cdot \int 0tSg\, dt \end{matrix}$ which theoretically ensures efficient prioritization amid varying clinical demands [23, 24]. This fosters resilience against external shocks, such as regulatory changes, by dynamically recalibrating anomaly thresholds [25, 26]. Ultimately, the systemic impacts of ARCGI underscore a transformative approach to claims dynamics, where anomaly-aware governance theoretically harmonizes intelligence with ethical imperatives, paving the way for more robust healthcare fraud mitigation frameworks [27-32].

Results and Discussion

The conceptual articulation of the ARCGI advances theoretical discourse in healthcare AI by integrating anomaly awareness with fraud governance in a manner that transcends traditional silos. Central to this discussion is the architecture's unique helical feedback topology, which theoretically enables iterative refinement of intelligence processes, distinguishing it from linear models prevalent in existing literature [1-4]. This topology not only interprets risk propagation but also conceptually adapts to governance constraints, offering a nuanced perspective on how anomaly drifts can be

mitigated without relying on empirical datasets [5, 6]. In synthesizing clinical AI architectures, the ARCGI highlights interoperability as a pivotal enabler, theoretically bridging claims data with EHR ecosystems to amplify fraud detection efficacy [7-9].

A key point of discussion revolves around the interpretive formulas embedded in ARCGI, such as decision confidence and monitoring burden, which provide theoretical tools for analyzing system behaviors in anomaly-prone environments [10, 11]. These formulas allow for abstract modeling of dynamics, where governance load distributions theoretically optimize resource use, addressing gaps in prior workflow integration models that often overlook adaptive mechanisms [12-14]. Ethical considerations emerge prominently, as the infrastructure's design theoretically embeds fairness protocols, countering potential biases in anomaly profiling that could exacerbate inequities in healthcare delivery [15-17]. This ethical layering invites further theoretical exploration into how governance feedback can evolve to incorporate emerging regulatory landscapes, ensuring long-term alignment with clinical imperatives [18, 19].

Comparatively, while existing decision support pipelines emphasize static monitoring, ARCGI's dynamic orchestration layer theoretically introduces flexibility, enabling proactive responses to fraud evolutions [20, 21]. This flexibility is particularly relevant in high-stakes clinical settings, where anomaly-aware intelligence can theoretically reduce false positives, though without performance metrics, the discussion remains focused on architectural advantages [22, 23]. Limitations inherent to conceptual frameworks, such as the absence of empirical validation, underscore the need for future theoretical extensions that simulate interoperability challenges in diverse deployment environments [24-26]. Nonetheless, ARCGI contributes to the literature by proposing a governance-centric view of claims intelligence, theoretically fostering resilience against sophisticated fraud schemes [27, 28]. The co-evolutionary relationship between anomaly intelligence and governance oversight is conceptually illustrated in **Figure 2**.

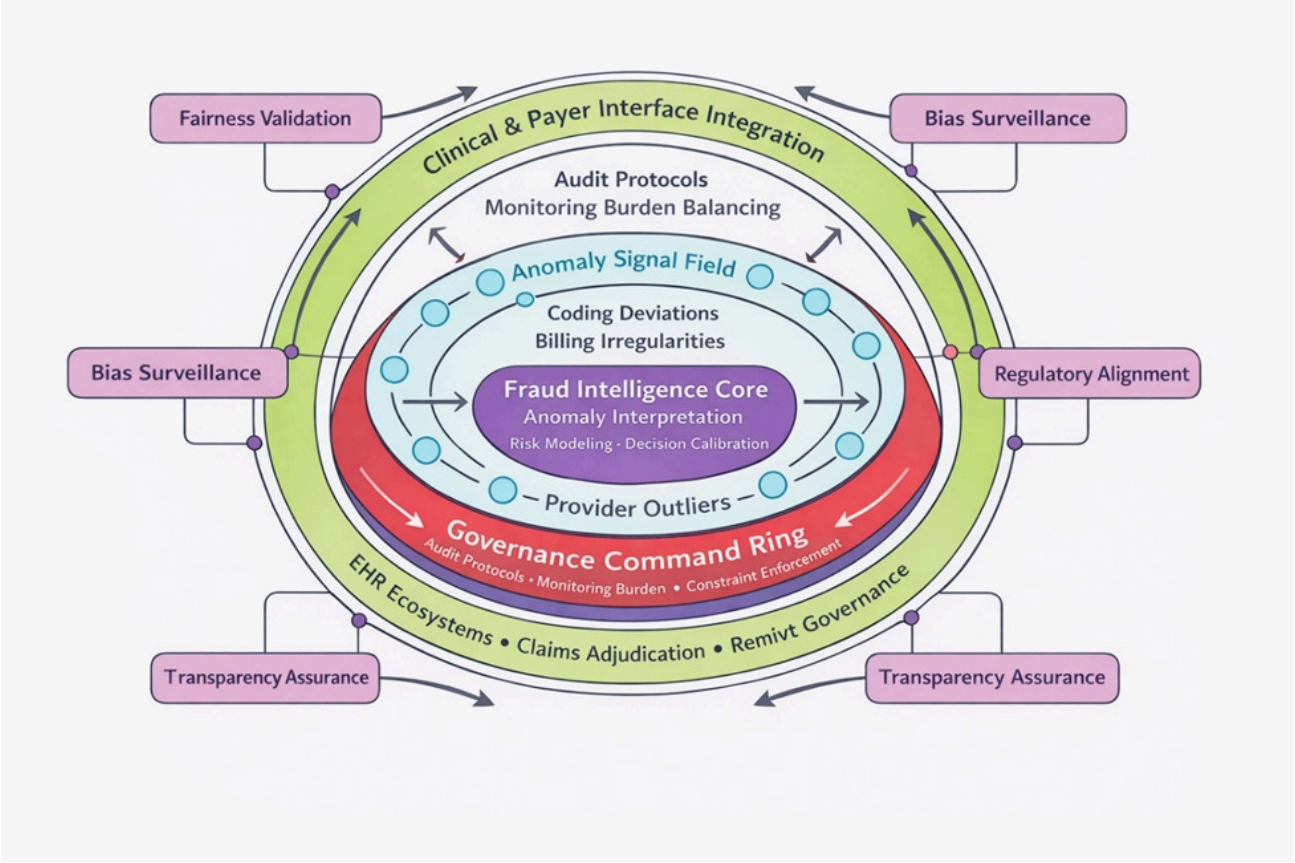

Figure 2. Governance–anomaly co-evolution topology within the ARCGI. The figure depicts a circular intelligence core interpreting emergent anomaly signals, encircled by governance command structures and ethical oversight nodes.

Bidirectional feedback channels illustrate adaptive recalibration between fraud intelligence, policy enforcement, and clinical–payer interfaces, highlighting the dynamic governance maturation processes discussed in anomaly-aware claims ecosystems.

Broader implications for healthcare analytics infrastructures suggest that anomaly-responsive designs like ARCGI could theoretically influence policy formulations, advocating for standardized governance protocols across payer networks [29, 30]. This discussion posits that by prioritizing anomaly awareness, architectures can theoretically shift the paradigm from reactive detection to preventive governance, aligning with evolving AI monitoring systems [31, 32]. In essence, ARCGI serves as a theoretical blueprint, stimulating discourse on integrating intelligence with ethical oversight in fraud-prone healthcare domains.

Conclusion

In conclusion, the ARCGI emerges as a pivotal conceptual advancement in the realm of healthcare claims intelligence, specifically tailored for anomaly-aware fraud governance. By synthesizing theoretical foundations from clinical AI architectures, EHR ecosystems, and interoperability frameworks, this manuscript delineates a unique architectural paradigm that theoretically fortifies claims systems against fraudulent incursions. The ARCGI's layered structure and helical feedback topology, complemented by interpretive formulas for risk dynamics and resource allocation, provide a robust blueprint for enhancing governance without empirical dependencies.

This work underscores the imperative of anomaly awareness in mitigating fraud propagation, where conceptual models like decision confidence and drift sensitivity theoretically guide adaptive intelligence orchestration. Impacts on system dynamics reveal how ARCGI theoretically optimizes monitoring burdens and ethical alignments, fostering resilient integrations with clinical workflows. Discussions highlight the architecture's potential to redefine fraud governance paradigms, addressing limitations in existing literature while opening avenues for theoretical refinements.

Ultimately, ARCGI contributes to the theoretical discourse on AI-driven healthcare analytics by proposing an infrastructure that harmonizes anomaly profiling with governance feedback, promoting sustainable fraud mitigation strategies. Future conceptual explorations could extend this framework to incorporate emerging modalities, such as blockchain-enhanced interoperability, further elevating its applicability in dynamic clinical environments. By emphasizing proactive, intelligence-centric designs, this manuscript advocates for a transformative approach to claims governance, theoretically safeguarding the integrity of healthcare systems amid evolving fraud challenges.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 27 Jan 2024 Revised: 29 Feb 2024 Accepted: 02 Apr 2024
Published online: 20 July 2024

References

Bauder RA, Khoshgoftaar TM, Seliya N. A survey on the state of healthcare upcoding fraud analysis and detection. Health Serv Outcomes Res Methodol. 2017;17(1):31-55.

Bauder RA, Khoshgoftaar TM. Multivariate outlier detection in medicare claims payments applying probabilistic programming methods. Health Serv Outcomes Res Methodol. 2017;17(3):256-89.

Bauder RA, Khoshgoftaar TM. Medicare fraud detection using machine learning methods. In: Proc IEEE 16th Int Conf Mach Learn Appl. 2017:858-65.

Bauder RA, Khoshgoftaar TM. The effects of varying class distribution on learner behavior for medicare fraud detection with imbalanced big data. Health Inf Sci Syst. 2018;6(1):9.

Bauder RA, Khoshgoftaar TM. A survey of medicare data processing and integration for fraud detection. In: Proc IEEE Int Conf Inf Reuse Integr. 2018:9-14.

Bauder RA, da Rosa R, Khoshgoftaar TM. Identifying medicare provider fraud with unsupervised machine learning. In: Proc IEEE Int Conf Inf Reuse Integr Data Sci. 2018:285-92.

Bauder RA, Khoshgoftaar TM. Medicare fraud detection using random forest with class imbalanced big data. In: Proc IEEE Int Conf Inf Reuse Integr Data Sci. 2018:80-7.

Bauder RA, Khoshgoftaar TM, Hasanin T. Data sampling approaches with severely imbalanced big data for medicare fraud detection. In: Proc IEEE Int Conf Tools Artif Intell. 2018:137-42.

Herland M, Bauder RA, Khoshgoftaar TM. The effects of class rarity on the evaluation of supervised healthcare fraud detection models. J Big Data. 2019;6(1):21.

Johnson JM, Khoshgoftaar TM. Medicare fraud detection using neural networks. J Big Data. 2019;6(1):63.

Johnson JM, Khoshgoftaar TM. Deep learning and data sampling with imbalanced big data. In: Proc IEEE 20th Int Conf Inf Reuse Integr Data Sci. 2019:175-83.

Johnson JM, Khoshgoftaar TM. Deep learning and thresholding with class-imbalanced big data. In: Proc IEEE Int Conf Mach Learn Appl. 2019:755-62.

Johnson JM, Khoshgoftaar TM. Deep learning and data sampling with imbalanced big data. In: Proc IEEE 20th Int Conf Inf Reuse Integr Data Sci. 2019:175-83.

Herland M, Bauder RA, Khoshgoftaar TM. Medical provider specialty predictions for the detection of anomalous medicare insurance claims. In: Proc IEEE Int Conf Inf Reuse Integr. 2017:579-88.

Johnson JM, Khoshgoftaar TM. Data-centric AI for healthcare fraud detection. SN Comput Sci. 2023;4(4):389.

Settipalli L, Gangadharan GR. WMTDBC: An unsupervised multivariate analysis model for fraud detection in health insurance claims. Expert Syst Appl. 2023;215:119259.

Yoo Y, Shin J, Kyeong S. Medicare fraud detection using graph analysis: a comparative study of machine learning and graph neural networks. IEEE Access. 2023;11:77493-503.

Hancock JT, Bauder RA, Wang H, Khoshgoftaar TM. Explainable machine learning models for medicare fraud detection. J Big Data. 2023;10(1):154.

Johnson JM, Khoshgoftaar TM. Healthcare provider summary data for fraud classification. In: Proc IEEE 23rd Int Conf Inf Reuse Integr Data Sci. 2022;236-42.

Hancock JT, Khoshgoftaar TM. Optimizing ensemble trees for big data healthcare fraud detection. In: Proc IEEE 23rd Int Conf Inf Reuse Integr Data Sci. 2022:243-9.

Hancock JT, Khoshgoftaar TM. The effects of random undersampling for big data medicare fraud detection. In: Proc IEEE Int Conf Serv-Oriented Syst Eng. 2022:141-6.

Hancock JT, Khoshgoftaar TM. Hyperparameter tuning for medicare fraud detection in big data. SN Comput Sci. 2022;3(6):440.

Yoo Y, Shin D, Han D, Kyeong S, Shin J. Medicare fraud detection using graph neural networks. In: Proc Int Conf Electr Comput Energy Technol. 2022;1-5.

Kumaraswamy N, Markey MK, Ekin T, Barner JC, Rascati K. Healthcare fraud data mining methods: A look back and look ahead. Perspect Health Inf Manag. 2022;19(1):1i.

Johnson JM, Khoshgoftaar TM. Encoding high-dimensional procedure codes for healthcare fraud detection. SN Comput Sci. 2022;3(5):362.

Kumaraswamy N, Markey MK, Barner JC, Rascati K. Feature engineering to detect fraud using healthcare claims data. Expert Syst Appl. 2022;210:118433.

Nugraha RA, Pardede HF, Subekti A. Oversampling based on generative adversarial networks to overcome imbalance data in predicting fraud insurance claim. Kuwait J Sci. 2023;50(1B):1-7.

Li J, Liu J, Liu X, Yang F, Xu Y. A medical insurance fraud detection model with knowledge graph and machine learning. In: Proc Int Conf Comput Appl Inf Secur. 2022:531-40.

Settipalli L, Gangadharan GR. Provider profiling and labeling of fraudulent health insurance claims using weighted MultiTree. J Ambient Intell Humaniz Comput. 2023;14:11471-93.

Leevy JL, Hancock J, Khoshgoftaar TM, Abdollah Zadeh A. Investigating the effectiveness of one-class and binary classification for fraud detection. J Big Data. 2023;10(1):157.

Matloob I, Khan SA, Rahman HU. Sequence mining and prediction-based healthcare fraud detection methodology. IEEE Access. 2020;8:143256-73.
https://doi.org/10.1109/ACCESS.2020.3013962.

Sowah RA, Kuuboore M, Ofoli A, Kwofie S, Asiedu L, Koumadi KM, et al. Decision support system for fraud detection in health insurance claims using genetic support vector machines. J Eng. 2019;2019:1432597.