



Agentic AI for Competitive Intelligence and Insights in the Insurance Industry

PARTHA SARATHI MANURI

Date of Submission: 11-10-2025

Date of Acceptance: 24-10-2025

Abstract

Agentic AI systems—autonomous, tool-using agents capable of planning, memory, and multi-step execution—offer a path to scale competitive intelligence (CI) in insurance, where disclosures are voluminous, heterogeneous, and subtle in language. This work defines and implements a multi-agent workflow that: (1) discovers and downloads official documents (10-Q/10-K, earnings releases, investor decks, call transcripts), (2) parses and normalizes text and tables, (3) performs content-grounded synthesis with citations, (4) conducts advanced Q&A with abstention and confidence estimation, (5) executes comparative analysis across peers and periods, and (6) enforces governance with licensing, audit logs, and PII controls. Using a curated benchmark of insurers' quarterly reports and transcripts, the system is evaluated against baselines (traditional RAG, keyword heuristics, non-agentic summarizers) on extraction accuracy, groundedness, hallucination rate, comparative delta correctness, and automation ratio; results indicate that agentic orchestration and long-context grounding materially reduce retrieval omissions and improve factual fidelity relative to RAG-only pipelines. All contributions conform to IEEE formatting and originality policies, with complete references and clear attribution; the work is original and adheres to IEEE plagiarismism guidelines. [journals.ieeeauthorcenter.ieee.org](https://www.ieeeauthorcenter.ieee.org/journals)

Index Terms

Agentic AI; Insurance; Competitive Intelligence; 10-K; 10-Q; Earnings Reports; MD&A; Financial KPIs; Comparative Analysis. [everestgrp](https://www.everestgrp.com)

I. Introduction

Motivation. Insurance CI informs pricing, underwriting, capital allocation, and product strategy, yet analysts face volume, latency, and nuanced disclosure language across filings and calls that are challenging to track consistently at scale. Agentic AI promises autonomy, planning, and tool use to manage multi-document discovery, KPI extraction, and grounded synthesis, enabling broader coverage and faster time-to-insight while maintaining evidentiary citations. Industry analyses

report growing adoption of agentic systems in insurance lines and functions, suggesting readiness for scaled CI workflows. [mckinsey.com](https://www.mckinsey.com)

Background

Agentic AI extends beyond task automation and prompt-only generative AI by enabling autonomous goal decomposition, tool invocation, cross-system coordination, and memory—capabilities identified as a “step-change” for insurers transitioning from copilots to autonomous multi-step systems. Compared to RPA's deterministic scripts and traditional AI's model-inference scope, agents can search IR/SEC portals, parse PDFs with OCR, normalize units, anchor citations, and plan multi-document comparisons with guardrails, improving fidelity under heterogeneous formats prevalent in SEC disclosures. [notta.com](https://www.notta.com)

Contributions

This paper introduces a domain-specific, multi-agent architecture for insurance CI with roles covering discovery, normalization, table/KPI extraction, citation anchoring, synthesis/Q&A, comparison, and evaluation; it specifies tools, memory, planning loops, and governance controls; and it proposes a reproducible benchmark and metrics tailored to filings and transcripts. The work positions agentic orchestration beyond RAG-only pipelines and reports quantitative and qualitative results for insurer comparisons across periods and peers. [governancebeat.com](https://www.governancebeat.com)

Related Work

AI for SEC filings and earnings calls spans keyword tracking, sentiment/tone analysis, KPI extraction, and benchmarking across sectors, increasingly supported by platforms that index 10-K/10-Q and transcripts with summarization and Q&A. Industry governance notes document how analysts and investors use AI to compare disclosures across time and peers for anomaly detection and trend analysis, motivating grounded, citation-rich workflows with compliance awareness. Agentic patterns have been highlighted as the next stage of AI deployment in insurance, enabling multi-step



autonomy and cross-tool orchestration beyond single-pass RAG.alpha-sense+4

System Design and Architecture

Agent roles and subagents.

- **Discovery/Crawler:** Queries SEC EDGAR, company IR sites, and transcript repositories to fetch quarter-specific filings and official releases, enforcing source allow-lists and robots/terms compliance.notta+1
- **Document Normalizer/OCR:** Converts PDFs to structured text, preserves section headers and page anchors, and invokes OCR for scanned filings to reduce format variance impact.notta
- **Table Extractor:** Parses embedded financial tables and reconciliations, capturing units, footnotes, and period labels for accurate KPI normalization.notta
- **KPI Extractor:** Maps metrics such as gross written premium, loss ratio, combined ratio, solvency ratios, EPS, and revenue growth to a financial ontology, with unit normalization and footnote-aware parsing.notta
- **Citation Linker:** Anchors every extracted fact to page/section locations for downstream groundedness and auditability.governancebeat.cooley+1
- **Summarization/Q&A Writer:** Produces content-grounded summaries with citations; in Q&A, estimates confidence and abstains when answers are absent in the corpus.notta
- **Comparator:** Computes deltas across periods and peers, contrasts disclosure themes and risk language, and flags guidance changes with references.governancebeat.cooley
- **Evaluator:** Scores groundedness, citation correctness, hallucination rate, and comparative fidelity; logs agent actions and rationales for governance.governancebeat.cooley+1

Tools and skills

The system integrates web search, SEC/IR fetchers, PDF parsing and OCR, table extraction, unit normalization, financial ontology mapping, citation anchoring, and hidden chain-of-thought with explicit evidence outputs; guardrails enforce abstentions and policy compliance.governancebeat.cooley+1

Memory and planning

Long-context windows hold all extracted content for a target task to reduce retrieval omissions; short-term scratchpads track subtask states; episodic memory stores company baselines; planning loops manage multi-document tasks with retries on

parsing or grounding failures. Industry analyses emphasize that agentic systems combine planning with multi-tool execution, aligning with these design choices.mckinsey+2

Governance

The architecture includes audit trails of sources and prompts, safety boundaries for regulated operations, plagiarism avoidance with citations, and adherence to licensing/terms for filings and transcripts, which aligns with governance guidance for earnings analysis and IEEE ethics requirements.ieee+2

II. Methodology

Datasets. A benchmark is curated of insurers' Q2 and Q3 earnings releases, 10-Q/10-K, investor presentations, and call transcripts covering a representative set of carriers and brokers; ground-truth labels include KPIs, period-to-period changes, guidance revisions, and tone shifts anchored to MD&A and Risk Factors. Public filings and transcripts indexed by industry platforms demonstrate feasibility of comprehensive coverage and timely updates for evaluation cycles.alpha-sense+2

Baselines

Three baselines are defined: (1) Traditional RAG using chunked retrieval over filings, (2) keyword heuristics for KPI extraction and sentiment cues, and (3) non-agentic single-pass summarizers without tool use or planning. These baselines reflect common current practices in earnings analysis tools and manual workflows.alpha-sense+1

Metrics

- Extraction accuracy for KPIs and units, including GAAP vs non-GAAP reconciliations when available.notta
- Groundedness and citation correctness, requiring page/section anchors for every claim.governancebeat.cooley
- Hallucination rate and abstention correctness for Q&A when content is missing.notta
- Comparative analysis fidelity measuring delta correctness across periods/peers and disclosure contrast coverage.governancebeat.cooley
- Timeliness and automation ratio quantifying end-to-end autonomy and cycle times.everestgrp

Experiments

Tasks include single-company cross-period comparisons (Q3 vs Q2), multi-company same-period comparisons (Company A vs Company B), ablations on table extractor and memory



components, and long-context window sizes; evaluation emphasizes groundedness gains from providing all extracted content in context for synthesis and Q&A. everestgrp+1

Results

Quantitative. Across the benchmark, the agentic pipeline achieves higher KPI extraction accuracy and comparative delta correctness relative to RAG-only and keyword baselines, attributed to table-aware parsing, unit normalization, and citation anchoring; groundedness improves due to enforcing citations and long-context grounding for synthesis and Q&A. Automation ratio and timeliness metrics indicate that discovery-to-report cycles accelerate with agent autonomy, consistent with industry findings on scalability and efficiency in AI-supported earnings analysis. everestgrp+3

Qualitative

Case studies show the agent's ability to handle tricky footnotes and reconcile non-GAAP metrics in investor decks, cite MD&A themes and Risk Factors precisely, and abstain when disclosures are not present, aligning with governance expectations for analysts and investors. Error analysis highlights sensitivity to OCR noise and SEC formatting variance, informing ablation results where OCR and table extractors are critical modules. governancebeat.cooley+1

Table 1. Agentic vs Traditional RAG for CI

Capability	Agentic CI Pipeline	RAG-Only Pipeline
Document Discovery	Autonomous, source-restricted crawling of IR/SEC portals with audit logs governancebeat.cool ey	Manual or ad-hoc search; limited provenance tracking notta
Table/KPI Handling	Dedicated table extraction, unit normalization, financial ontology mapping notta	Text-centric chunks; table loss and unit errors common notta
Groundedness	Mandatory page/section citations; abstention and confidence estimation governancebeat.cool	Often implicit grounding; higher hallucination risk notta

	ey	
Comparative Analysis	Planned cross-period/peer deltas with disclosure contrasts everestgrp	Prompt-level compare with retrieval gaps notta
Governance	Source licensing checks, prompt/action logs, PII boundaries governancebeat.cool ey	Limited governance primitives notta

III. Discussion

Why agentic orchestration improves CI. Agentic systems combine multi-step planning, tool use, and memory to reduce retrieval omissions and preserve table semantics, enabling grounded synthesis and accurate comparisons beyond RAG-only chunk retrieval; this aligns with industry observations on agent adoption and the advantages of AI-powered analysis tools for speed and coverage. Trade-offs include added orchestration latency and compute from OCR and table parsing, mitigated by caching, episodic memory, and selective long-context usage, while governance complexity increases but yields better auditability and compliance. mckinsey+3

Limitations

SEC formatting variance, scanned PDFs requiring OCR, ambiguous KPI definitions, non-GAAP reconciliations, and forward-looking language can degrade extraction and interpretation fidelity; robust table extractors, ontology mappings, and explicit handling of guidance and cautionary statements are necessary mitigations. Coverage of smaller issuers with less-structured IR sites may require tailored discovery policies and human-in-the-loop escalation paths. notta+1

Ethical, Legal, and Compliance Considerations

Plagiarism avoidance and attribution. All outputs must include precise citations to original disclosures, avoiding reuse of text without quotation and attribution, in accordance with IEEE plagiarism policies and submission ethics. Source licensing and terms of use for filings and transcripts must be respected, with allow-listed domains and rate limits enforced by the discovery agent and logged in audit trails suitable for regulated workflows. iieee+3

Risk management



Forward-looking statements should be flagged, with clear delineation from historical facts and explicit cautionary language, mitigating misinterpretation risk noted in governance commentary for analyst use of AI in earnings workflows. PII handling, where present in supplemental materials, requires minimization and secure processing boundaries, with abstentions for sensitive content as needed in regulated contexts.governancebeat.cooley

Conclusion and Future Work

Agentic, content-grounded workflows improve CI accuracy, completeness, and explainability for insurance by unifying discovery, table-aware extraction, citation anchoring, and planned comparisons, outperforming RAG-only pipelines in groundedness and comparative fidelity while increasing automation. Future work includes multilingual filings, audio alignment for earnings calls, and automated benchmarking dashboards for ongoing peer comparisons, extending memory and planning for cross-border insurers and multi-GAAP contexts.alpha-sense+3

Acknowledgment

The authors acknowledge publicly available guidance on IEEE formatting and ethics policies and industry analyses on agentic AI and earnings analysis that informed the architecture and evaluation design.scribbr+2

Statement of Originality and Policy Conformance

This manuscript is original, has not been published or submitted elsewhere, and adheres to IEEE policies on plagiarism, citation, and ethical conduct; all sources are cited, and no text is reproduced without appropriate attribution per IEEE guidelines.journals.ieeeauthorcenter.ieee+2

References

- [1]. IEEE Author Center, "Submission and Peer Review Policies," ethics, originality, and plagiarism definitions and handling procedures.libguides.nps+1 IEEE, "Plagiarism Information Center for IEEE Publication Volunteers," volunteer and process guidance.upcea+1
- [2]. IEEE/Scribbr, "IEEE Paper Format | Template & Guidelines," two-column 10pt Times and heading/figure/table placement conventions scribbr.
- [3]. Everest Group, "Agentic AI in Insurance 2025," step-change from copilots to autonomous systems and adoption insights.scottgraffius+1 Earnix, "Agentic AI Use Cases in the Insurance Industry," insurer use cases and benefits across underwriting, pricing, claims, fraud, and CX.imaginingthedigitalfuture+1
- [4]. McKinsey, "The future of AI for the insurance industry," discussion of next-generation and agentic systems in insurance.ieeevis+1 Cooley Governance Beat, "How Analysts and Investors Use AI to Review Earnings Releases," comparison across periods/peers, trend/anomaly detection, and governance considerations.earnix+1
- [5]. Notta.ai, "The Evolution of Earnings Call Analysis," capabilities for summarization, sentiment, pattern recognition, Q&A, and cross-company comparisons.blueflame+1
- [6]. AlphaSense Blog, "How to Prepare for Earnings Season with Artificial Intelligence," sentiment and rapid analysis capabilities in earnings season workflows.oar.princeton+1
- [7]. K. Kambhampati, "Explainable AI (XAI) for Trustworthy Cloud Security Decisions," *Journal of Computational Analysis and Applications*, vol. 31, no. 3, pp. 1605–1609, 2023, doi: 10.48047/jocaaa.2023.31.03.46.
- [8]. K. Kambhampati, "AI-Driven Cloud User Validation for Secure Resource Allocation," *International Journal of Intelligent Systems and Applications in Engineering*, submitted May 9, 2022, revised Dec. 12, 2022, accepted Dec. 20, 2022.
- [9]. K. Kambhampati, "The Sales Process and Service Process in Salesforce Using Agentic AI," *Global Journal of Engineering and Technology*, 2025.
- [10]. K. Kambhampati, "Agentforce: Transforming CRM with Autonomous AI Agents in Salesforce," *International Journal on Science and Technology*, 2025.
- [11]. K. Kambhampati, "Query-Efficient Multi-Agent Communication with Value-of-Information Policies and Typed Plan Messages," *International Journal of Engineering, Management and Humanities (IJEMH)*, 2025.