International Journal of Graphic Design (IJGD)
Volume 3 No. 2 October 2025
E-ISSN: 2987-9434; P-ISSN : 2988-0343, Page 395-409
DOI: https://doi.org/10.51903/ijgd.v3i2.3710
Available online at: https://journal.stekom.ac.id/index.php/ijgd

Evidence-Grounded Accounting Disclosure Review Cards: A Visual Communication Framework for LLM-Style Explanations over SEC Financial Statements and Notes

Kai Zhang [1], Yuanzheng Chen[*2], Annie Qian[3]
[1]Financial Engineering, Baruch College, NY, USA
[2]Accounting, UIUC, IL, USA
[3]Human Centered Design & Engineering, Washington University, WA, USA
Email Address: kai.zhang.baruchmfe@gmail.com

Abstract. *This paper presents evidence-grounded accounting disclosure review cards, a visual communication framework for LLM-style explanations over SEC financial statements and notes. The revised study positions the card as an accounting evidence interface rather than as a test of real analyst performance or a live LLM pipeline. The objective is to organize structured statement facts, note-level disclosure evidence, ratio signals, citation trails, and review actions so that accounting claims can be inspected and challenged within a compact card. The evaluation combines the 2024 EDGAR quarterly master indexes, the 2024 SEC Financial Statement Data Sets, and the 2024 Q3-Q4 Financial Statement and Notes Data Sets. The master-index universe contains 1,241,395 filings, including 25,993 10-K/10-Q-family filings. The statement-level analysis uses 23,677 10-K/10-Q-family rows from 6,077 companies across 70 SIC groups, and 20,753 rows (87.65%) contain enough statement evidence to populate the core card. The statement-note paired analysis covers 11,880 filings from 5,811 companies, of which 10,391 (87.47%) are card-ready after note evidence is added. A stratified accounting review benchmark of 400 filings generates 2,400 evidence-checking tasks. The results support the feasibility of the proposed accounting evidence architecture and clarify that the reported outcomes describe data coverage, evidence readiness, and disclosure mapping rather than observed human speed, workload, trust, or live-model reliability.*

Keywords *: Accounting Disclosure, Accounting Review, Citation Trail, Evidence Grounding, Financial Statement Notes.*

INTRODUCTION

Accounting disclosure review requires reviewers to connect financial statement line items with the notes, policies, estimates, and filing metadata that explain those line items. A claim about liquidity, revenue, leverage, leases, taxes, fair value, or contingencies is useful only when the reviewer can see the underlying tagged fact and the related disclosure context. Public SEC filings are well suited to this design problem because they contain structured XBRL facts, presentation paths, filing identifiers, and note-level text blocks.

LLMs can turn filing data into fluent summaries, but a paragraph explanation can hide the relationship between a claim, the accounting metric, and the supporting disclosure. The design problem is therefore not only whether an LLM can write a plausible memo. It is also whether the interface makes accounting evidence visible at the moment of review. In accounting settings, that means keeping the statement line, the ratio or exception, the note topic, and the citation trail close enough that a reviewer can challenge the explanation rather than simply read it.

Received: August 2025; Revised: August 2025; Accepted: September 2025; Published: October 2025
*Corresponding author, kai.zhang.baruchmfe@gmail.com

This paper develops an accounting disclosure review card. The card contains a review tier, statement evidence badges, ratio chips, disclosure badges, citation trails, and a controlled review action. The revised evaluation does not report simulated analyst speed, trust, or workload. Instead, it asks whether the proposed card schema can be populated from real SEC statement and note data at a meaningful scale and whether the resulting evidence structure supports accounting review tasks.

The contribution is threefold. First, the paper defines a compact visual grammar for accounting evidence cards, with every card element bound to a statement fact, a disclosure category, or filing metadata. Second, it extends the data basis from a small selected filing set to a 2024 SEC data pipeline that includes EDGAR master-index coverage, statement facts, and note-level disclosure evidence. Third, it reports evidence-readiness, disclosure coverage, review-tier, and task-coverage results that are directly aligned with the revised accounting framing.

LITERATURE REVIEW

The study connects visualization design, explainable AI, and accounting evidence. Visualization theory supports stable grouping, direct labeling, and overview-to-detail movement in information displays (Cleveland & McGill, 1984; Heer et al., 2010; Munzner, 2014; Shneiderman, 1996; Ware, 2020). Those principles are especially relevant to financial statements because reviewers often move from a headline issue to a tagged line item and then to a note disclosure. The card uses small graphic units rather than large dashboard views because accounting review is often a verification task: a reader must decide whether the evidence shown actually supports the claim.

Explainable-AI research (Chen & Chan, 2023) argues that explanations should answer user questions, communicate evidence, and support appropriate reliance (Abdul et al., 2018; Kizilcec, 2016; Liao et al., 2020; Ribeiro et al., 2016). In the present setting, the relevant user questions are accounting questions: Which statement fact changed? Which note topic explains it? Is the ratio or exception grounded in tagged evidence? What filing source would a reviewer check next? The card treats these questions as interface requirements rather than as properties of generated prose alone.

Accounting standards and audit-evidence concepts further motivate the evidence-first design. FASB's conceptual framework emphasizes useful financial information and the qualitative characteristics that support decision use. PCAOB AS 1105 explains audit evidence and the need for sufficient appropriate evidence when performing audit procedures. The proposed card is not an audit procedure and does not provide an audit opinion, but it borrows the practical

idea that a financial reporting claim should be traceable to appropriate evidence. This makes accounting disclosure review a natural domain for the UI/UX contribution (Kuhn et al., 2024).

The SEC data model makes this extension practical. The Financial Statement Data Sets provide flattened SUB, NUM, TAG, and PRE files for XBRL-tagged financial statement facts, while the Financial Statement and Notes Data Sets provide text and detailed numeric information from statements and notes. EDGAR master indexes provide filing-level metadata, including CIK, company name, form type, filing date, and filing path. Together, these sources allow a card to bind a visual element to a statement fact, a disclosure topic, and a filing identifier.

METHODS

A. Data sources and selection

Table 1 summarizes the data sources used in the revised study. The EDGAR master indexes define the 2024 filing universe and make the selection process transparent. The Financial Statement Data Sets provide statement-level facts for all four 2024 quarters. The note-level extension uses the Q3-Q4 Financial Statement and Notes Data Sets to map statement claims to accounting disclosure evidence. The scope therefore moves beyond a small selected 10-K set while still keeping the note analysis tied to available note files.

Table 1. Source data and role in the revised accounting evidence evaluation

Data source	Coverage	Records used	Use in study
EDGAR master index, 2024 Q1-Q4	All EDGAR filing metadata in the four quarterly master indexes	1,241,395 filings; 25,993 10-K/10-Q-family filings	Defines the filing universe and supports transparent form and quarter selection.
SEC Financial Statement Data Sets, 2024 Q1-Q4	SUB, NUM, PRE, and TAG files for XBRL-tagged financial statement facts	23,677 10-K/10-Q-family rows from 6,077 companies	Populates statement-line badges, ratios, accounting signals, and presentation paths.
SEC Financial Statement and Notes Data Sets, 2024 Q3-Q4	Text and detailed numeric information from financial statements and notes	11,880 statement-note paired filings from 5,811 companies	Populates note-level disclosure badges and citation trails for accounting review tasks.

The master-index screening in Table 2 shows that the 2024 filing universe contains more than one million EDGAR records and 25,993 10-K/10-Q-family filings. These counts are used for transparency rather than for performance claims: they show the universe from which the accounting form families were identified.

Table 2. EDGAR master-index coverage by SEC quarter

Quarter	Master-index filings	Unique CIKs	10-K/10-Q-family filings
2024q1	370,303	79,731	6,329
2024q2	316,023	72,339	7,244
2024q3	270,929	62,127	6,242
2024q4	284,140	62,327	6,178

The statement-level filing set is summarized in Table 3. Across the four 2024 Financial Statement Data Sets, 20,753 rows (87.65%) contain enough tagged statement facts and presentation paths to populate the core accounting card.

Table 3. Financial Statement Data Set coverage by quarter and form family

Quarter	Form family	Filings	Companies	SIC groups	Statement-card ready	Mean accounting signals	Ready rate
2024q1	10-K	4,037	3,995	69	3,439	1.93	85.19%
2024q1	10-Q	1,125	1,037	62	1,039	2.21	92.36%
2024q2	10-K	1,137	1,085	60	1,083	2.72	95.25%
2024q2	10-Q	5,498	5,383	70	4,796	2.07	87.23%
2024q3	10-K	490	470	57	450	2.44	91.84%
2024q3	10-Q	5,484	5,309	70	4,782	2.09	87.20%
2024q4	10-K	434	423	50	395	2.16	91.01%
2024q4	10-Q	5,472	5,288	70	4,769	2.09	87.15%

B. Card schema and accounting review tasks

Figure 1 shows the proposed accounting disclosure review card. Table 4 lists the components and their data bindings. The card is designed as a controlled evidence surface for LLM-style explanations: the explanation may be written in prose, but the card requires each major accounting claim to appear with statement evidence, disclosure evidence, and a citation trail.

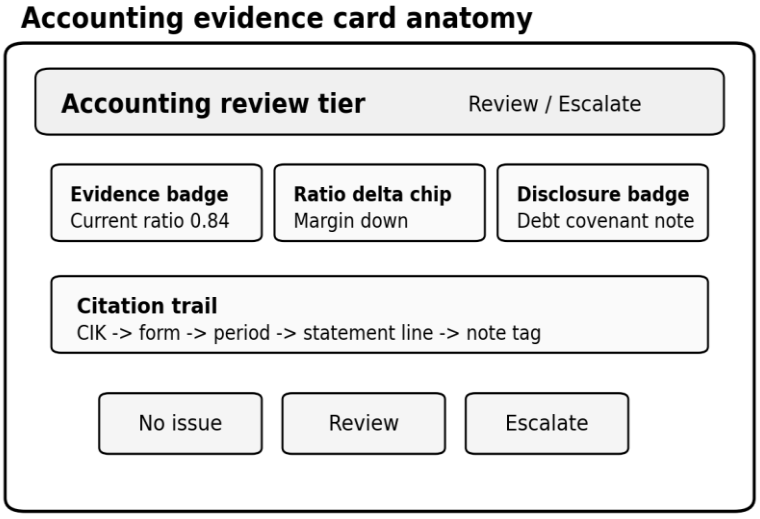

Figure 1. Anatomy of the proposed accounting disclosure review card

Figure 2 summarizes the data-to-card pipeline. The revised benchmark creates evidence tasks from the same fields that populate the card. Table 5 defines the six accounting review task types used in the benchmark. These tasks are evidence-availability and grounding tasks, not observed human task-performance trials.

Table 4. Accounting disclosure review card components

Component	Encoding	Data binding	Purpose
Accounting review tier	Low, Moderate, High, or Critical with redundant text label	Liquidity, profitability, leverage, negative equity, custom-tag, note-complexity, and disclosure-breadth signals	Provides a first-pass accounting review priority without claiming a market or audit opinion.
Statement evidence badge	Compact line-item badge with metric, value, and statement source	NUM facts joined to SUB and PRE statement metadata	Keeps ratio or balance-sheet claims close to the tagged statement facts that support them.
Ratio delta chip	Small change indicator for liquidity, leverage, profitability, or revenue context	Derived ratios from current assets, liabilities, equity, revenue, and net income	Makes accounting changes visible without forcing readers to search through prose.
Disclosure badge	Note-category badge for revenue, debt, leases, fair value, taxes, segments, contingencies, or policies	TXT tags from the Financial Statement and Notes Data Sets	Connects line-item review to accounting-policy and footnote evidence.
Citation trail	CIK -> form -> period -> statement line -> note tag	EDGAR metadata plus FSDS/Notes accession and tag information	Supports verification and auditability of the explanation surface.
Review action	No issue, Review, or Escalate action area	Accounting review tier and evidence sufficiency flags	Turns the explanation into a controlled review action rather than an open-ended recommendation.

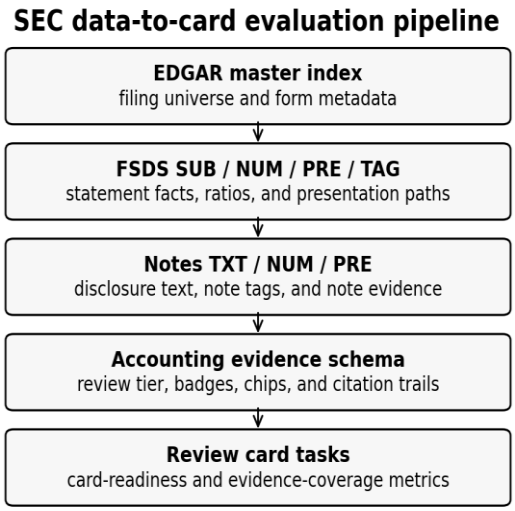

Figure 2. SEC filing-to-accounting-card evaluation pipeline

C. Measurement

The revised analysis reports four families of measures. First, statement-card readiness measures whether a filing contains enough statement facts and presentation information to populate the core card. Second, note-card readiness measures whether a statement-level record can be paired with note text and at least two disclosure categories. Third, disclosure coverage measures whether note tags support accounting areas such as revenue recognition, debt and

liquidity, leases, fair value, income taxes, segments, contingencies, and accounting policies or estimates. Fourth, the accounting review tier combines statement signals and disclosure complexity into Low, Moderate, High, or Critical review priority. The tier is a triage aid rather than an investment, audit, or enforcement label. It uses liquidity, profitability, negative equity, leverage, custom-tag intensity, disclosure breadth, and note complexity to indicate how much accounting review attention the card may require. This keeps the evaluation focused on data grounding rather than observed analyst behavior.

Table 5. Accounting review task types and evidence rules

Task type	Objective	Evidence rule	Eligible filings	Coverage rate
statement_line_check	Verify that a ratio or balance-sheet claim is supported by tagged statement facts	>=4 statement facts plus presentation path	10,396	87.51%
liquidity_going_concern_check	Inspect current ratio together with liquidity- or going-concern disclosure evidence	Statement facts plus debt/liquidity or going-concern disclosure tag	9,151	77.03%
revenue_policy_check	Check whether revenue or margin claims are linked to revenue-recognition disclosure evidence	Statement facts plus revenue-recognition disclosure tag	8,258	69.51%
debt_lease_check	Connect leverage, debt, and lease facts to the corresponding disclosure area	Statement facts plus debt or lease disclosure tag	9,722	81.84%
fair_value_tax_segment_check	Confirm whether complex disclosure areas are visible as evidence badges	At least one fair-value, tax, or segment disclosure area	11,876	99.97%
accounting_review_action	Map evidence signals to No issue, Review, or Escalate	Statement and note evidence sufficient to produce a card-ready review action	10,391	87.47%

RESULTS

A. Statement-level accounting evidence coverage

The four 2024 Financial Statement Data Sets contain 23,677 10-K/10-Q-family filing rows from 6,077 companies and 70 SIC groups. As shown in Table 3, statement-card readiness is consistent across quarters and form families. Overall, 20,753 statement-level filings (87.65%) contain enough tagged statement evidence and presentation information to populate the core card. Figure 3 visualizes the quarter and form-family coverage.

B. Statement-to-note evidence pairing

The note-level analysis links statement rows to Q3-Q4 note records. Table 6 shows 11,880 paired statement-note filings from 5,811 companies. Of those records, 10,391 (87.47%) are card-ready after note evidence is added. The annual form family has higher average text-fact counts,

while the 10-Q family contributes most of the paired records. Figure 4 shows that the card-ready rate remains close to or above 87% in the largest groups.

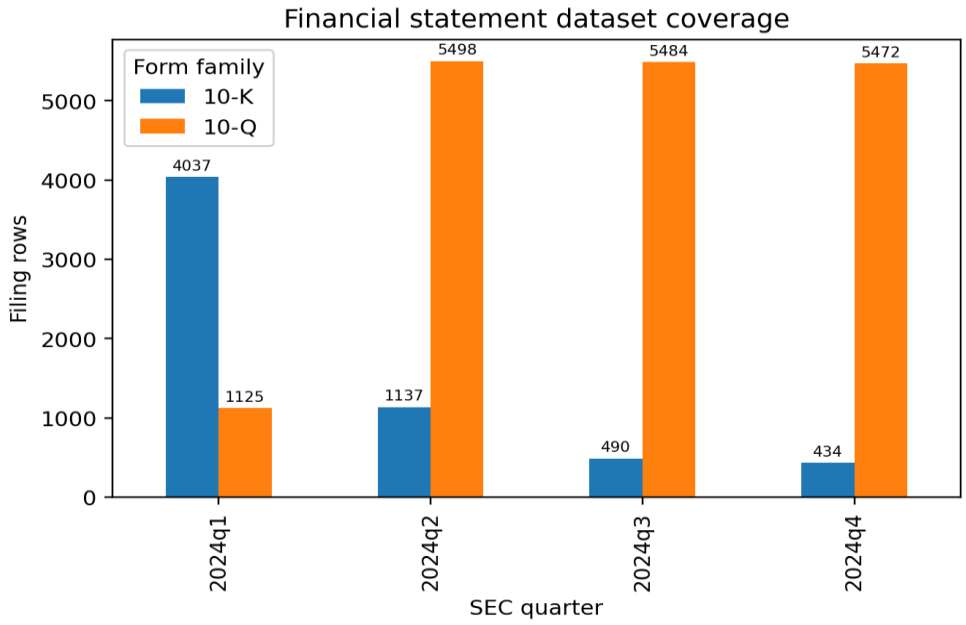

Figure 3. Financial statement dataset coverage by quarter and form family

Table 6. Statement-to-note paired coverage by quarter and form family

Quarter	Form family	Paired filings	Companies	Mean text facts	Mean disclosure areas	Card-ready filings	Ready rate
2024q3	10-K	490	470	136.71	6.94	450	91.84%
2024q3	10-Q	5,484	5,309	105.55	6.46	4,779	87.14%
2024q4	10-K	434	423	157.15	7.21	395	91.01%
2024q4	10-Q	5,472	5,288	107.68	6.47	4,767	87.12%

C. Accounting disclosure coverage

Table 7 reports coverage for the accounting disclosure categories used by the card. Income taxes, accounting policies and estimates, fair value, commitments and contingencies, and debt/liquidity appear in a large share of paired filings. Revenue recognition, leases, and segments also appear frequently enough to support distinct disclosure badges. Figure 5 shows the same coverage pattern visually.

D. Review tiers and task coverage

The accounting review tier is associated with greater evidence complexity. Table 8 shows that Critical-tier filings have higher card-ready rates, more statement signals, broader disclosure coverage, and higher median text-fact counts than lower-tier records. Figure 6 shows the

monotonic increase in median text facts from Low to Critical. This result supports the card design choice of keeping disclosure badges and citation trails visible for higher-priority reviews.

Table 7. Accounting disclosure area coverage in the statement-note paired dataset

Disclosure area	Filings with area	Coverage rate
Income taxes	11,847	99.72%
Accounting policies/estimates	11,087	93.32%
Fair value	10,970	92.34%
Commitments/contingencies	10,458	88.03%
Debt and liquidity	10,420	87.71%
Revenue recognition	8,944	75.29%
Leases	7,040	59.26%
Segments	6,591	55.48%

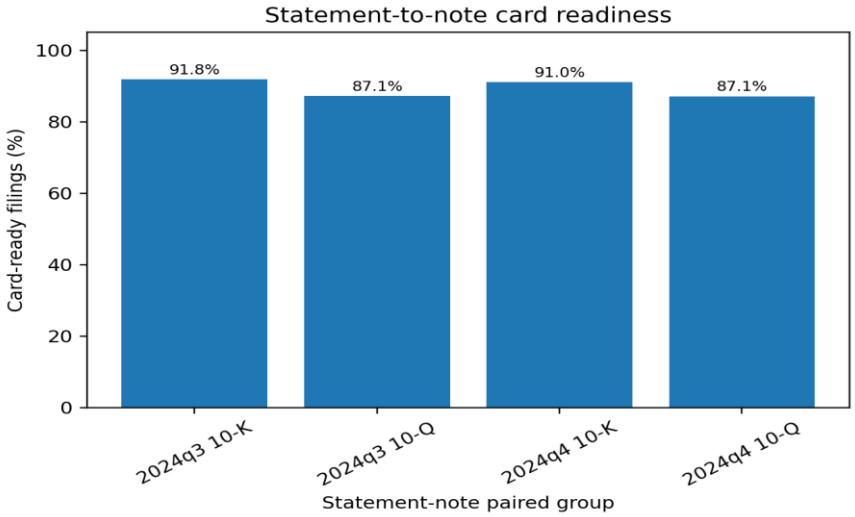

Figure 4. Statement-to-note card readiness by quarter and form family

A stratified accounting review benchmark was then drawn from card-ready filings. Table 9 summarizes the benchmark: 400 filings, 346 companies, 51 SIC groups, and 2,400 generated evidence-checking tasks. The benchmark is large enough for prototype evaluation of the card schema while keeping the unit of analysis clearly tied to evidence availability and disclosure mapping.

Table 8. Accounting review tier profile

Review tier	Filings	Card-ready rate	Mean statement signals	Mean disclosure areas	Median text facts
Low	4,606	82.33%	1.36	5.96	81
Moderate	3,469	87.86%	1.98	6.79	97
High	2,544	90.92%	2.88	6.68	104
Critical	1,261	98.18%	3.65	7.41	121

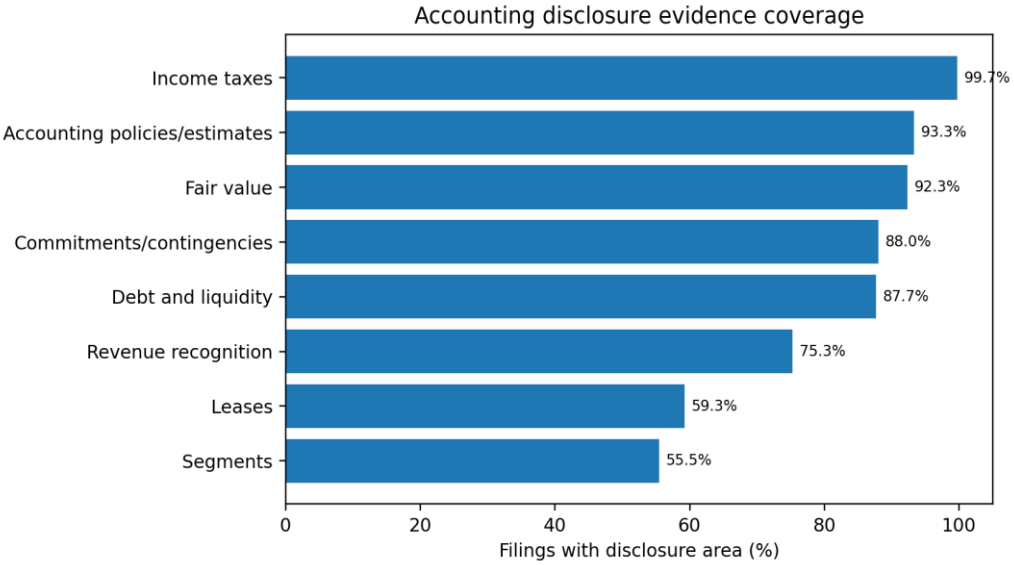

Figure 5. Accounting disclosure evidence coverage by category

DISCUSSION

The results support a focused and defensible claim about accounting evidence architecture. The card should be understood as an accounting evidence architecture: it specifies how statement facts, disclosure topics, and citation trails can be organized for review. The findings show that this architecture can be populated at scale from SEC statement and note data, but they do not show that real analysts work faster, make fewer errors, experience lower workload, or trust the system more.

Table 9. Stratified accounting review benchmark summary

Metric	Value
Benchmark filings	400
Benchmark companies	346
SIC groups	51
Form families	4
Task types	6
Generated accounting review tasks	2,400
Mean text facts per filing	110.36
Mean disclosure areas per filing	6.71
High or Critical review tier	50.50%

This reframing directly strengthens the accounting contribution. Instead of treating risk as a trading-desk triage score, the card treats review priority as a function of accounting signals and disclosure complexity. Liquidity, profitability, negative equity, leverage, custom tags, revenue recognition, leases, fair value, taxes, segments, contingencies, and accounting policies become visible evidence units. That change makes the design more relevant to accounting disclosure review and better aligned with a paper that includes accounting expertise.

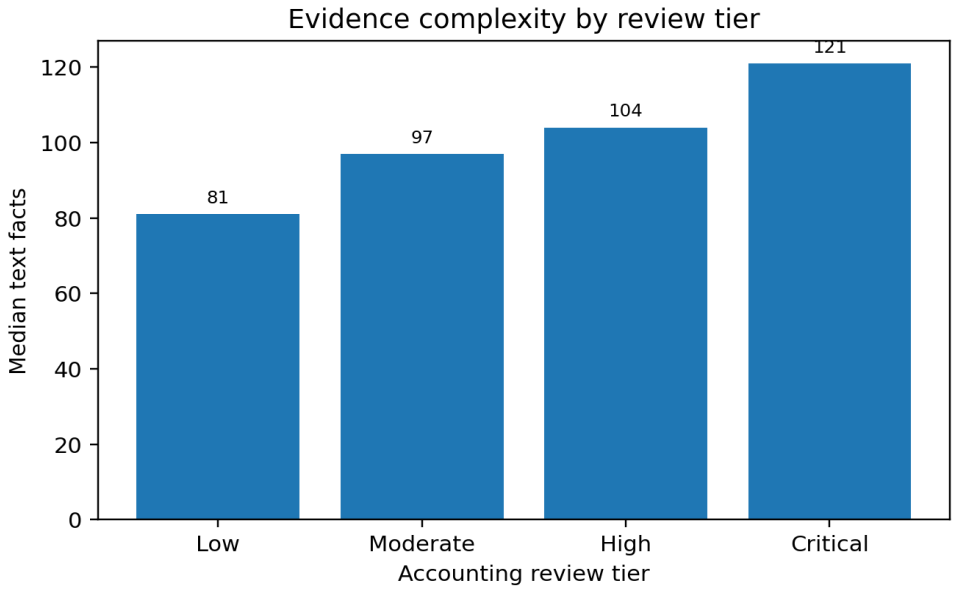

Figure 6. Evidence complexity by accounting review tier

The data also clarify why a visual card is useful. Financial reporting evidence is distributed across statement line items, presentation structures, and notes. A paragraph memo can mention these facts, but it often forces the reader to infer which number belongs to which disclosure. The proposed card reduces that ambiguity by placing the metric, disclosure category, and citation path in the same visual neighborhood. The design therefore supports challengeability: the reader can inspect the basis for the claim without depending on fluent prose alone.

The revised LLM framing is also more precise. The paper evaluates an LLM-style explanation surface, not a live retrieval-augmented LLM system. A live system would need additional tests for retrieval errors, hallucinated citations, prompt variability, and model drift. The present study focuses on the interface and data-binding layer that such a system would need before generated explanations could be considered evidence-grounded.

Limitations

The study has four limitations. First, it does not include live human participants; conclusions are about evidence coverage and card readiness rather than analyst performance. Second, it does not evaluate a live LLM pipeline; the contribution is the controlled explanation surface and accounting evidence schema. Third, the note-level analysis uses the 2024 Q3-Q4 Financial Statement and Notes Data Sets, so Q1-Q2 note coverage is not part of the note-pairing results. Fourth, the accounting review tier is a transparent triage rule for interface evaluation and should not be read as an audit conclusion, investment recommendation, or prediction of financial misstatement.

These limitations are also boundaries that make the revised claims clearer. The results show that the proposed card can be populated with real SEC statement and note evidence and that its task taxonomy covers major accounting disclosure areas. A later human-subject study or live LLM evaluation could build on the same card schema, but those claims are not required for the present design contribution.

CONCLUSION

This paper introduced evidence-grounded accounting disclosure review cards for LLM-style explanations over SEC financial statements and notes. The revised framework replaces trading-desk performance claims with an accounting evidence-mapping evaluation. Across 2024 SEC filing data, the card schema supports statement-line badges, ratio chips, disclosure badges, citation trails, and review actions at a scale suitable for prototype evidence evaluation. The strongest design principle is evidence-first visual hierarchy: every accounting claim needs a tagged statement fact, every statement fact benefits from nearby disclosure context, and every explanation needs a citation trail that can be checked. In this form, the card is not a persuasive memo or a black-box recommendation. It is a compact visual artifact for accounting disclosure review, designed to make LLM-style explanations easier to inspect, challenge, and govern.

REFERENCES

Abdul, A., Vermeulen, J., Wang, D., Lim, B. Y., & Kankanhalli, M. (2018). Trends and trajectories for explainable, accountable and intelligible systems: An HCI research agenda. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3173574.3174156

Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., Avila, R., Babuschkin, I., Balaji, S., Balcom, V., Berner, C., Brockman, G., ... Zoph, B. (2023). GPT-4 technical report. arXiv. https://doi.org/10.48550/arXiv.2303.08774

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610-623. https://doi.org/10.1145/3442188.3445922

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., ... Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877-1901.

Cleveland, W. S., & McGill, R. (1984). Graphical perception: Theory, experimentation, and application to the development of graphical methods. Journal of the American Statistical Association, 79(387), 531-554. https://doi.org/10.1080/01621459.1984.10478080

Financial Accounting Standards Board. (2021). Concepts Statement No. 8: Conceptual framework for financial reporting, Chapter 1, The objective of general purpose financial reporting, as amended. FASB.

Financial Accounting Standards Board. (2024). Concepts Statement No. 8: Conceptual framework for financial reporting, Chapter 3, Qualitative characteristics of useful financial information, as amended. FASB.

Heer, J., Bostock, M., & Ogievetsky, V. (2010). A tour through the visualization zoo. Communications of the ACM, 53(6), 59-67. https://doi.org/10.1145/1743546.1743567

International Organization for Standardization. (2019). ISO 9241-210:2019: Ergonomics of human-system interaction-Part 210: Human-centred design for interactive systems. ISO.

Jason Kuhn, Yushan Chen, & Evelyn Chan. (2024). AI-Driven Mobile UI Pattern Recognition and Design Topic Mining on RICO: Semantic Clustering and Screenshot-Based Topic Classification. Journal of Advanced Computing Systems , 4(5), 67-83. https://doi.org/10.69987/JACS.2024.40506

Kizilcec, R. F. (2016). How much information? Effects of transparency on trust in an algorithmic interface. Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, 2390-2395. https://doi.org/10.1145/2858036.2858402

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors, 46(1), 50-80. https://doi.org/10.1518/hfes.46.1.50_30392

Liao, Q. V., Gruen, D., & Miller, S. (2020). Questioning the AI: Informing design practices for explainable AI user experiences. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, 1-15. https://doi.org/10.1145/3313831.3376590

Munzner, T. (2014). Visualization analysis and design. CRC Press.

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

Nielsen, J. (1994). Enhancing the explanatory power of usability heuristics. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 152-158. https://doi.org/10.1145/191666.191729

Public Company Accounting Oversight Board. (2024). AS 1105: Audit evidence. PCAOB.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International

Conference on Knowledge Discovery and Data Mining, 1135-1144. https://doi.org/10.1145/2939672.2939778

Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. Proceedings of the IEEE Symposium on Visual Languages, 336-343. https://doi.org/10.1109/VL.1996.545307

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12(2), 257-285. https://doi.org/10.1207/s15516709cog1202_4

Tufte, E. R. (2001). The visual display of quantitative information (2nd ed.). Graphics Press.

U.S. Securities and Exchange Commission. (2024a). Financial Statement Data Sets. https://www.sec.gov/data-research/sec-markets-data/financial-statement-data-sets

U.S. Securities and Exchange Commission. (2024b). Financial Statement Data Sets documentation. https://www.sec.gov/files/financial-statement-data-sets.pdf

U.S. Securities and Exchange Commission. (2024c). Financial Statement and Notes Data Sets. https://www.sec.gov/data-research/sec-markets-data/financial-statement-notes-data-sets

U.S. Securities and Exchange Commission. (2024d). EDGAR Application Programming Interfaces. https://www.sec.gov/search-filings/edgar-application-programming-interfaces

Ware, C. (2020). Information visualization: Perception for design (4th ed.). Morgan Kaufmann.

World Wide Web Consortium. (2023). Web Content Accessibility Guidelines (WCAG) 2.2. W3C Recommendation. https://www.w3.org/TR/WCAG22/

Yushan Chen, & Evelyn Chan. (2023). Multimodal UI Representation Learning: Ablation of Screenshot, Wireframe, and View-Hierarchy Proxies on an Uploaded 168-Screen Dataset. Journal of Advanced Computing Systems , 3(1), 1-15. https://doi.org/10.69987/JACS.2023.30101