Supplementary Material

RCInvestigator: Towards Better Investigation of Anomaly Root Causes in Cloud Computing Systems

1. CASE STUDY

A. Expert Interview Settings

Procedure. This case study consists of three stages:

(1) **[10min]** introduce RCInvestigator
(2) **[50min]** investigate anomalies' root causes freely
(3) **[15min]** conduct a structured interview

Participants. This case study is conducted with four domain experts (EA, EB, EC, and ED) via one-on-one semi-structured interviews. All of the experts were employed by a leading cloud computing provider. EA (female) and EB (male) were data scientists with years of experience in conducting cloud platform incident analyses, while EC (male) and ED (male) were senior researchers specialized in AIOps, facilitating IT operations in cloud computing systems with AI-based approaches.

B. Structured Interview Form

Q1. Investigation framework: What is your opinion on the investigation framework, including its advantages and disadvantages? Which stage do you prefer and why? which stage can be improved and how?

Q2. Visualization and interactions: What is your opinion on the visualization and interaction designs? Please provide specific examples of designs that you find intuitive and helpful, and explain why they are effective. Additionally, please identify any designs that you consider complex, unnecessary, or missing, and offer suggestions for improvement.

Q3. Comparing with existing workflow: How do you evaluate the time performance of RCInvestigator compared to your original workflow? Please specify whether RCInvestigator helps to speed up your workflow, performs similarly, or becomes slower. Additionally, which process do you believe RCInvestigator affects the most in terms of speeding up or slowing down the workflow, and why?

C. Records of the Expert Interview

The original interview was conducted in Chinese, and we have translated it.

Q1 What is your opinion on the investigation framework, including its advantages and disadvantages? Which stage do you prefer and why? which stage can be improved and how?

EA: *I find the current framework to be effective. I can organize my understanding of the data into the knowledge graph during the building stage. I think this is useful. It allows me to systematically organize my experiences and either preserve them or share them with other team members. Additionally, during the analysis, I can identify suspicious key performance indexes without the need to write extra data scripts for data collection. This enhances my focus.*

EB: *This framework aligns with logic and my usage habits: identifying incidents, analyzing their causes, and summarizing the results. The framework frees me from the redundant and repetitive data scripting work. In terms of the stages, the building stage left a deep impression on me. I tried to add some facts to the graph with my understanding of the database and found that the knowledge graph could be easily expanded.*

EC: *I like this analysis framework. It's got a great feature: it lets you reuse the analysis experience. We often run into a problem where knowledge isn't shared well among us, so different analysts might know different things about the database. During the analysis, I just imported a ready-made knowledge graph and started analyzing it. With that, I can get possible causes via simple clicks. Plus, I don't have to keep switching between the analysis and data query panels. This makes the investigation more fluent than before.*

ED: *Overall, I think this analysis framework tackles some tedious issues I face at work. Rather than talking about my favorite stage of the analysis, I'd say I appreciate one feature, the reusable knowledge graph. It covers almost all the reasoning rules I use in my daily analysis. This way, I can write less repetitive data scripts, and the reasoning stays coherent. As for the disadvantages, there's nothing obviously wrong with it, but I find it a bit tricky to find patterns that align on different cards by time. If there were reference lines, it would be a lot better.*

Q2 What is your opinion on the visualization and interaction designs? Please provide specific examples of designs that you find intuitive and helpful, and explain why they are effective. Additionally, please identify any designs that you consider complex, unnecessary, or missing, and offer suggestions for improvement.

EA: *Generally, I find the design to be neat and intuitive. The KPIs are all laid out on different cards. These cards are sorted into different sections so it's easy to see how different clues fit together. The direct conclusions given by the machine agent help me value clues quickly.*

Plus, I can annotate my own thoughts and hypotheses. It's handy for me. Now, for what could be better, I'd like it if the system let me remove certain unexpected clues.

EB: *My first impression is that the overall design is clean and easy to understand. Through simple line charts, I can grasp what happened and thus focus all my efforts on the investigation rather than deciphering the charts. These charts are also well-organized. The only feature I think could be added is something like a minimap, so I can get a better understanding of the whole reasoning.*

EC: *I appreciate the overall design. It's neat and intuitive, showing only the necessary information and clues. The line charts are clear and well-structured. The layout is helpful. When I hit a mental block, I'll explore different directions to try and find some inspiration. On that note, I think the system could benefit from an analytical thumbnail to help me understand how deep into the analysis I currently am.*

ED: *The design issue I mentioned earlier concerns the reference lines. I believe everything else in the design is good. I find the detective board design to be interesting as it allows me to add my own descriptions, and there are also descriptions that are generated. The line plots are logically arranged and well thought out. Both of these design aspects are beneficial for review and understanding.*

Q3 How do you evaluate the time performance of RCInvestigator compared to your original workflow? Please specify whether RCInvestigator helps to speed up your workflow, performs similarly, or becomes slower. Additionally, which process do you believe RCInvestigator affects the most in terms of speeding up or slowing down the workflow, and why?

EA: *RCInvestigator can help me boost my analysis efficiency and speed up my work. In the old way, whenever I thought of some possible elements, I first had to write a query script, then visualize it, and only after that could I proceed to the next discussion and inference. If I encountered any code issues while writing the script, I'd need to look back at the database, which made it hard to focus. In RCInvestigator, I can concentrate more on thinking and inferring, which can help increase my speed. So, I think the reasoning process changes the most.*

EB: *This new workflow may speed up my work in two ways: first, by reusing previous experience, it reduces the redundant and laborious data acquisition tasks; second, it opens up my thinking by showing me different angles and possible factors, which can help enhance my reasoning efficiency.*

EC: *During my investigation, I've found that RCInvestigator recommends clues directly related to the root cause right from the first step, which means it has the potential to enhance the efficiency of daily analysis. Its ability to help narrow down the scope of clues that need manual checking during the reasoning phase is significant. Compared to the previous analysis that relied solely on experience, RCInvestigator seems to leverage the correlations in data to help uncover some KPIs that might not have been considered important before.*

ED: *When analyzing root causes, being able to get multiple clues with just a few mouse clicks is obviously faster than writing code for each key performance index. And, having these key performance indexes organized categorizingly and accompanied by descriptions further reduces the workload of organizing and understanding KPI information. So overall, I think RCInvestigator can help speed up the efficiency of analysis.*

2. DESIGN ITERATIVE PROCESS

In the design process of RCInvestigator, we traversed three phases. Initially, we conducted extensive research on existing commercial tools, open-source tools, and visual analysis systems for AIOps solutions to understand the research context and the pain points of the root cause analysis issue. In the second phase, we engaged in collaborative discussions with four experts to identify the pain points in the current workflow and jointly distilled the user requirements. In the third phase, we initially designed a preliminary version of the visual analysis system based on the user requirements and discussed potential workflows. Subsequently, we refined the design through three iterations to arrive at the final version. In this section, we introduce the design and how we improve it based on experts' comments. Since the experts' comments encompass a wide range of aspects, we list the experts' comments according to different design considerations below to more clearly demonstrate how we have revised the design.

Based on three design challenges and the four-stage workflow, we detailed user requirements into five design tasks.

T1. (Building) Build persistent investigation knowledge interactively . The system must provide an interactive way for analysts to build, edit, and persist their investigative knowledge about system components and their relationships. This structured knowledge serves as the foundation for the machine's automated clue collection.

T2. (Monitoring) Obtain an overview of cloud computing anomalies. The system must provide an interactive way for analysts to build, edit, and persist their investigative knowledge about system components and their relationships. This structured knowledge serves as the foundation for the machine's automated clue collection.

T3. (Reasoning) Extend analysis scope based on recommendations. This task focuses on the human agent's role in reasoning. The system must present machine-recommended clues in an interpretable manner, allowing the analyst to validate them as evidence. Based on this evidence, the analyst must be able to easily form new hypotheses and direct the next step of the investigation.

T4. (Reasoning) Explore time-oriented data based on investigation knowledge. This task defines the machine agent's role. Guided by the analyst's current hypothesis and the knowledge model, the system must automatically search the vast dataset to collect and recommend the most relevant causal clues. It should rank these clues based on their correlation to the anomaly to intelligently prune the search space.

T5. (Concluding) Share intuitive investigation results. The system must facilitate the creation of an intuitive and expressive summary of the entire investigation. This involves capturing the final root cause, key evidence, and reasoning steps in a structured format for sharing and future reference.

A. The Initial Design

Fig. S1 displays the illustration of the initial design of RCInvestigator.

- T1: The left panel was designed to display information about the database schema.
- T2: The left panel also supported this task by providing an overview of anomalies, including incident logs and KPI lines.
- T3: The middle panel served as the main reasoning workspace. It used a card-based design where each card represented an attribute (time series). Here, users could explore possible causes by analyzing temporal data, create new clues through interactions (such as merging or extending cards), and visualize cause relations using a stream graph.
- T4: The right panel displayed machine-extracted potential clues. These clues were listed from top to bottom according to their anomaly score, allowing the user to select them for investigation in the main reasoning panel.
- T5: The analysis board itself can be recorded as an image or a json object.

We mocked up a usage scenario on the design and received some comments from experts. Later, we will follow the narration of design components.

Layout: There was a divergence of opinion on the layout design. Some experts believed that the current card-style/chart group design offered sufficient flexibility, while others suggested trying a more structured visualization, similar to a dashboard, with a categorized layout. Some experts thought a structured layout could help users better organize clues.

Visualizations: The experts reflected that the flow chart seemed unnecessary because the links between cards already revealed the relations between different clues. Compared with the colorful flow chart, the links between cards were easier to read. Besides, the experts preferred concise charts, such as line charts.

Interactions: The experts commented that time series operations had unique interaction requirements; standard operations such as addition, subtraction, and aggregation were seldom used in routine analysis. They often defined frequently used series directly as Key Performance Indicators (KPIs). Moreover, aggregating time series of differing units, like the count of available nodes with the cumulative anomaly counts, was illogical. The interaction design should have emphasized practicality.

Knowledge: Experts, after reviewing the time series calculations, indicated that simply displaying the database schema was insufficient. In their routine analysis, they often utilized scripts to generate Key Performance Indicators (KPIs) that encapsulated valuable insights and expertise. They recommended an interface capable of displaying, editing, and enabling the iterative creation of these KPIs.

Recommendations: Although the recommended clues showed a strong correlation with the existing ones, their overwhelming display made it difficult to distinguish them in practical applications. Experts suggested categorizing these clues according to specific criteria like hierarchy and type to enhance clarity and usability.

Fig. S1. The initial design of RCInvestigator. The system is divided into three panels. The left panel displays the information of the database and the monitoring data. The middle panel is the main reasoning panel. The right panel displays detected potential clues.

B. 2nd Iteration

In this iteration, we incorporated modifications based on expert feedback. Fig. S2 shows the user interface. The modifications are as follows:

Layout: In the first design iterations, we encountered conflicting feedback on the investigation layout. Some experts initially advocated for a flexible, card-based design to support free-form exploration, while others preferred a more structured, table-based layout for clarity and to align with existing tools. To resolve this disagreement methodically, we mocked up both competing designs and conducted a scenario-based discussion, walking the experts through a typical investigation using each mockup. Therefore, for this round, we designed a new layout. The time series (KPIs) are now presented in a table format, with each row dedicated to a specific KPI and each column representing an entity.

Visualizations: We prioritized the use of line plots and Gantt charts for visualizations. While we retained the stacked area chart for its ability to illustrate entity relationships through color, we removed explicit links to simplify the user interface.

Interactions: The redesigned interface enables users to modify an existing KPI to create a new processed time series, ensuring that time series with incompatible units are not inadvertently aggregated. For example, the KPI 2 can be aggregated.

Knowledge: We enhanced the left panel by incorporating an illustration of the knowledge graph.

Recommendations: After consultation with experts, we identified four types of interactions and categorized recommended clues. These are now presented to users in a structured layout, improving accessibility and usability.

We presented the mock-up system to the experts, engaged in a discussion, and received valuable feedback for further enhancement.

Layout: During this comparative, scenario-driven review, the limitations of the structured layout became clear. Experts noted that the table-based design, which relied heavily on color encoding, was inadequate for revealing the complex, non-linear relations between different entities and time series. They emphasized that visualizing these explicit relations (clues) was critical for the reasoning process. Consequently, based on this direct comparison, all experts collectively agreed that the card-based design, while requiring careful layout management, offered a superior approach for visualizing and interacting with the investigative path. This led us to refine the card-based approach into the final hierarchical layout presented in the paper.

Recommendations and visualizations: Experts acknowledged that the current recommended clues organization has improved, aiding in the identification of some semantic aspects. However, experts conveyed that the 'recommended clues' displayed on the panel present an excess of details, which might lead to information overload. Experts suggested simplifying the presentation by omitting unnecessary details to enhance comparability. They also stressed the importance of prominently highlighting and distinguishing the investigation board from other recommendations, underscoring its role as the central feature.

Knowledge: Experts found the knowledge overview to be coherent, but they recommended a significantly larger panel for editing and creation. They also noted that the database display is somewhat obscured by the knowledge graph design, rendering it redundant.

Fig. S2. This diagram presents the second iteration of the design, divided into six sub-panels, each displaying incident logs, database information, a knowledge graph overview, KPI lines, an investigation board, and recommendations. The investigation board features a table-based layout, while the recommendation board organizes suggested clues into four distinct categories.

C. 3rd Iteration

Fig. S3 displays the refined system. In this iteration, we refined the design with a focus on four main aspects:

Layout: We reintroduced a card-based layout. We adopted an algorithm to position related cards close, akin to a force-directed graph approach.

Recommendations: We employed small multiples to streamline recommendations, setting them apart from the primary plots on the investigation board. Besides, time series within the same category are aligned on a common timeline for straightforward comparison.

Visualizations: We simplified the visualizations by eliminating complex plots and color schemes. In this system, we only used straightforward line plots and Gantt charts.

Knowledge: A dedicated panel has been added for users to create and edit the knowledge graph.

We implemented the system according to the design and utilized it to analyze real-world scenarios. We presented these scenarios to experts and engaged in discussions on system improvement. The feedback is primarily focused on two areas:

Layout: In scenarios where investigations involve multiple aspects across different entities, the presence of long-linked cards was noted to increase mental effort, especially when sharing results, causing a loss of context. After discussions with experts, we collectively identified that recommendations, when displayed separately from the investigation, lead to more frequent context-switching and panel-switching efforts. Consequently, experts have expressed a desire for a more semantic layout that provides context, aids in locating the investigation, and integrates recommendations seamlessly.

Semantic enhancement: Experts found the simplified visualization to be much easier to understand but noted that it might still impose a reading burden on users. One expert, EA, stated, “Natural language is more easy to understand than any other representation when there are multiple plots.” Specifically, the experts suggested enhancing the visualization with additional text to summarize time series data or guide users throughout their investigation.

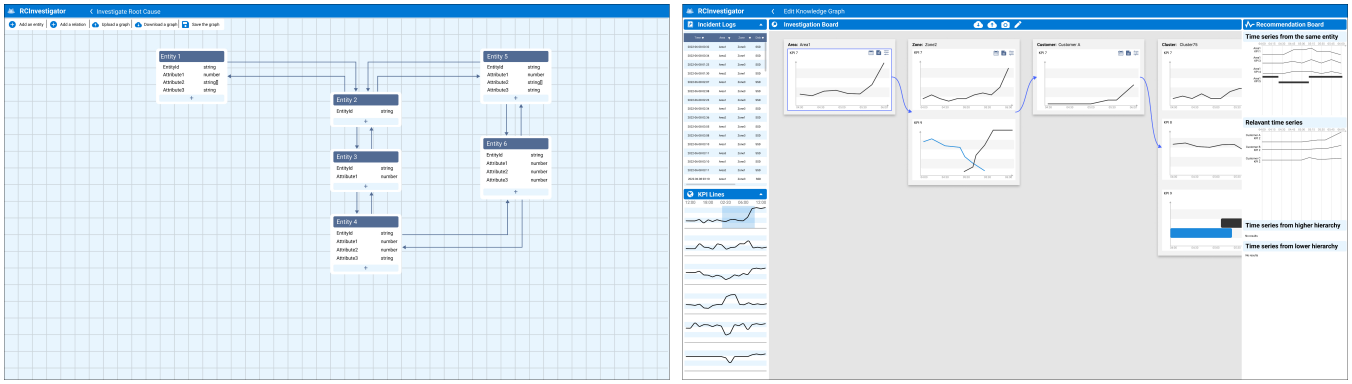

Fig. S3. The system is divided into two interfaces: one for constructing and editing the knowledge graph, and the other for analyzing root causes. The root cause analysis interface is further divided into three panels. The leftmost panel displays two manifestations of anomalies (incident logs and KPI lines), the center panel serves as the analysis area, and the rightmost panel showcases the currently recommended suspicious clues.

D. Final Design

During our discussions with experts, we engaged in brainstorming sessions to enhance the semantic aspects of not just the layout, but also the visualization. We discovered that the 'investigation board' itself was already metaphorically tailored. As depicted in Fig. S4D, the cards resemble clues, the links represent logic, and the annotations meet the semantic requirements effectively.

Consequently, we optimized the design, resulting in the final system. The system comprises three main boards: the building board (Fig. S4A), the monitoring board (Fig. S4B), and the investigation board (Fig. S4C). We strategically separated monitoring KPIs and incident logs to help users maintain focus on the investigation during the reasoning process.

Furthermore, we introduced a semantic layout algorithm designed to position cards within a balanced and free-form structure. To further enhance the semantic clarity, we encoded orientations for different categories of time series and their respective entities. For a detailed introduction, please refer to the main text (Section IV-C Expand Relevant Clues) of the paper.

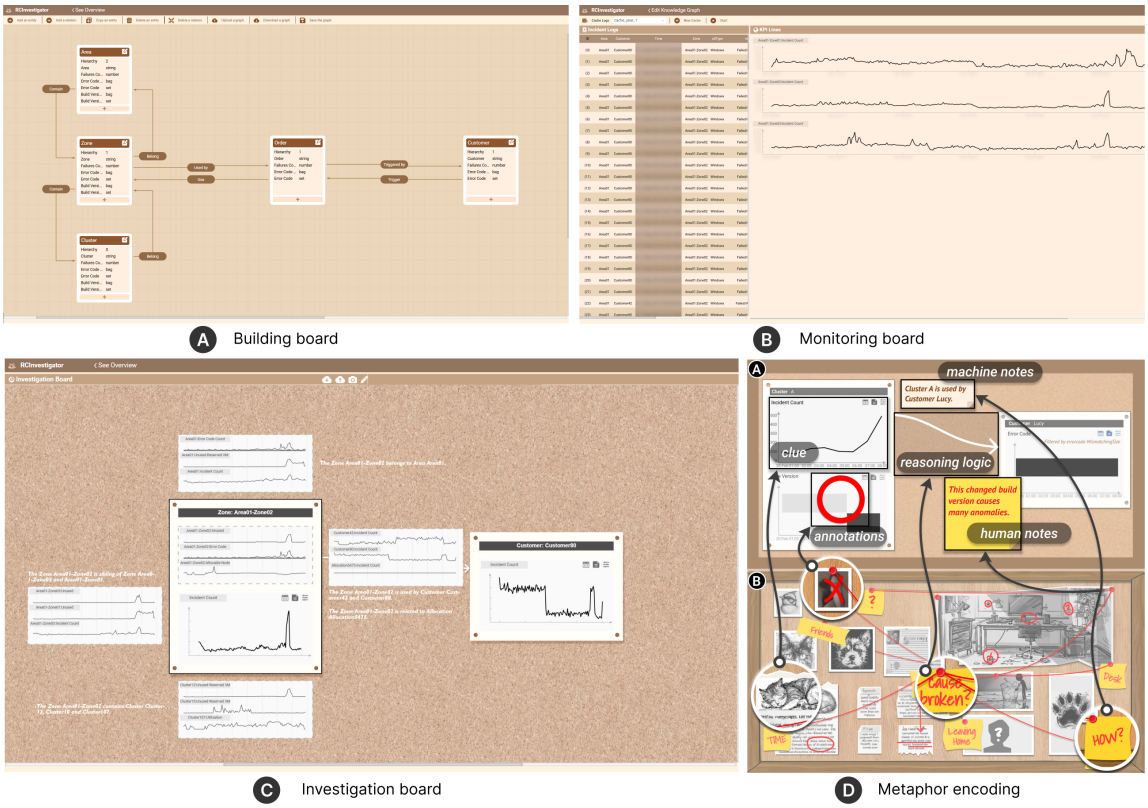

Fig. S4. This presents the RCInvestigator design overview. It has three panels: (A) The building board, for constructing and editing the knowledge graph; (B) The monitoring board, for displaying records of monitored anomalies and KPI lines; (C) The investigation board, the primary panel for assisting users in analyzing root causes. Diagram (D) illustrates the semantic metaphors that inspired the overall design. It maps elements inherent to the investigation board to the elements involved in the analysis process.

3. QUANTITATIVE EVALUATION

A. Model Efficacy Test

To provide a quantitative comparison and test our model's precision/recall, we conducted an experiment with a semi-synthetic dataset.

Data Generation. We generated datasets of 100 and 1,000 time series. Each dataset contained one known root cause and a 10% rate of relevant attributes (ground truth). The remaining 90% of series were irrelevant noise and event data (80% time series, 20% events). Here we put the visualization of 100-series dataset in Fig. S5.

Baseline. We compared our change-point relevance model against a standard pure correlation (Pearson) baseline, a common automated RCA approach.

Results. On the 100-series test, our model achieved Precision=0.90 / Recall=0.82, versus the baseline's PR=0.80 / RE=0.73. This advantage scales: on a 1,000-series test, our model achieved PR=0.94 / RE=0.93, while the baseline remained at PR=0.80 / RE=0.79.

Ranking Analysis. We also added two plots (Fig. S6) to the revision. They show the ranking of recommended clues, colored by their type: Root Cause (Red), True Positive (Green), False Positive (Blue), and False Negative (Orange). This visualization clearly shows that our model consistently ranks the Root Cause (Red) and True Positives (Green) at the top, while the baseline buries them in noise.

B. Layout Algorithm Test

We tested the interaction latency and layout performance of our two-step DAGre layout. Specifically, we measured the "time-to-add" latency—the time from a user's expand click to the moment a new entity card and its edges are fully rendered on the Investigation Board.

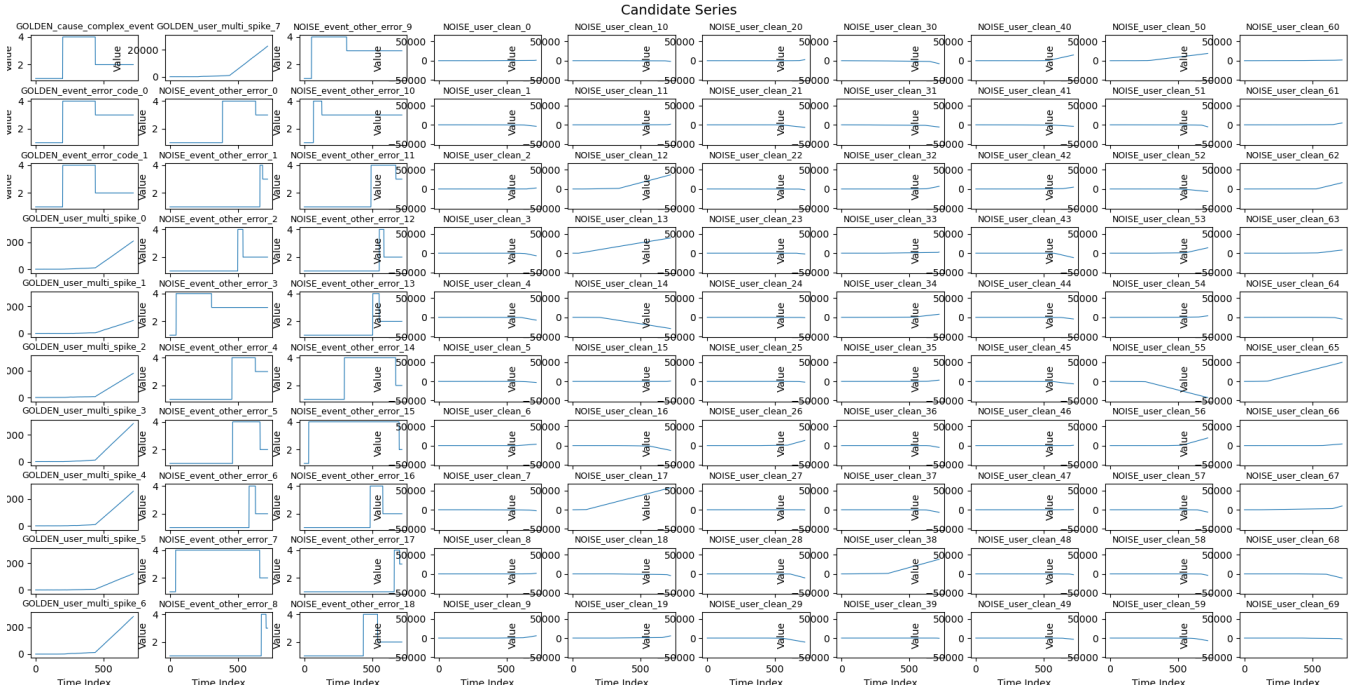

Fig. S5. The generated 100 attributes.

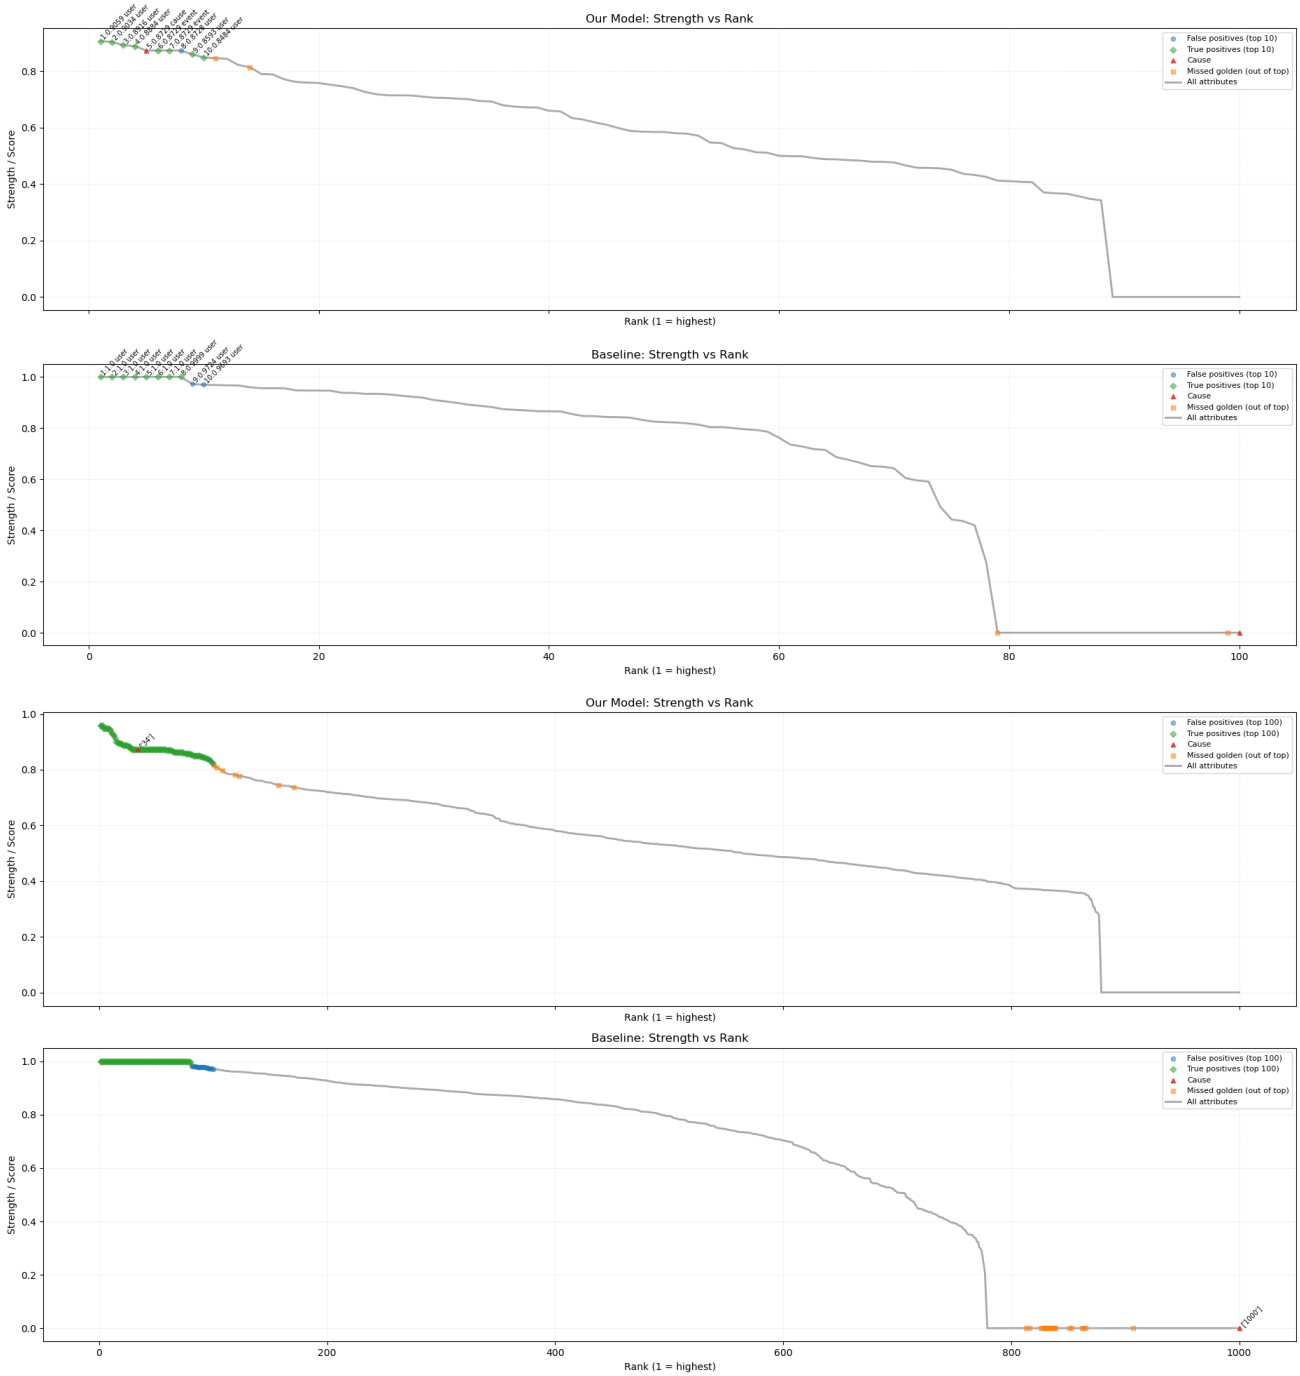
Our Model: Strength vs Rank
Strength / Score
Rank (1 = highest)
False positives (top 10)
True positives (top 10)
Cause
Missed golden (out of top)
All attributes
Baseline: Strength vs Rank
Strength / Score
Rank (1 = highest)
False positives (top 10)
True positives (top 10)
Cause
Missed golden (out of top)
All attributes
Our Model: Strength vs Rank
Strength / Score
Rank (1 = highest)
False positives (top 100)
True positives (top 100)
Cause
Missed golden (out of top)
All attributes
Baseline: Strength vs Rank
Strength / Score
Rank (1 = highest)
False positives (top 100)
True positives (top 100)
Cause
Missed golden (out of top)
All attributes

Fig. S6

As shown in Table S1, the test confirms our system is highly interactive for its intended analytical scope. For simple (6-10 nodes) and medium (44-50 nodes) graphs, which represent the scale of our case studies, the average latency is excellent at 159ms and 556ms, respectively. While the latency for an extreme 99-node graph (avg. 1.7-2.0s) approaches our 2-second interactivity budget, we argue that a single investigation board growing to 100 distinct entity cards is an unrealistic edge case for this domain. The performance for all typical scenarios remains robust.

Our two-step hierarchical layout is inherently designed to prevent overlap between entity cards. The rendered examples are shown in Fig. S7, Fig. S8, and Fig. S9.

	simple0	simple1	simple2	medium0	medium1	medium2	complex0	complex1	complex2
nodes	6	9	10	50	46	44	83	85	99
time(ms)	97.9	177.8	200	273.1	530.6	865.1	1502.1	1800.7	1995.5
avg time (ms)	158.5666667			556.2666667			1766.1		

Table S1

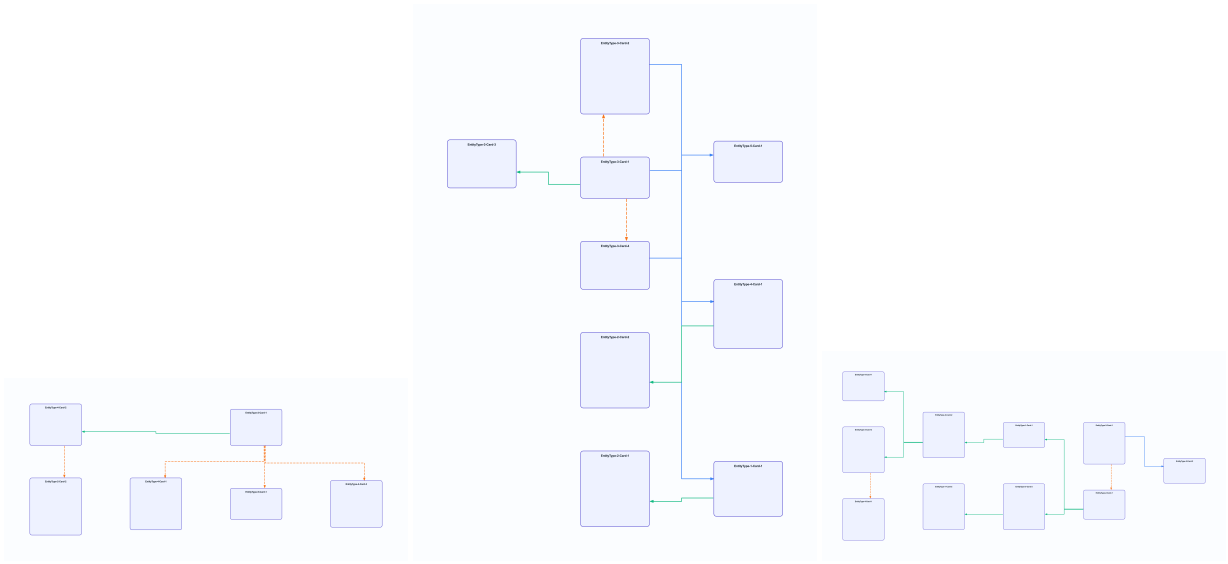

Fig. S7

Fig. S8

Fig. S9