## FOR IMMEDIATE RELEASE

**Chicago, IL – March 13, 2025 – Booth Students Launch AI Tool to Transform Corporate Compliance, Fraud Detection, and Operational Efficiency**

Bloomberg, March 13, 2025

A team of students from the University of Chicago Booth School of Business today unveiled a new application of AI poised to revolutionize how corporations manage compliance, detect fraud, and streamline their operations. By harnessing Microsoft's cutting-edge GraphRAG technology and integrating it over corporate data, this tool delivers unprecedented data analysis capabilities to organizations both large and small. Its capabilities extend significantly beyond conventional compliance monitoring and operational analytics by offering deeper, more accurate insights.

The tool takes conventional data analysis a step further by not only examining the data but also interpreting the results in the context of the organization. It enables users to uncover hidden insights and relationships from otherwise siloed or disconnected pieces of information, understanding complex data traditional AI solutions often miss. This leads to more informed decision-making and faster, smarter business operations, saving hundreds of man-hours of work and freeing employees up to accomplish other tasks.

**Figure 1: GraphRAG in action – An example knowledge graph constructed from a document corpus, with nodes representing entities and colored clusters indicating communities of related information.** GraphRAG's indexing process partitions the data clusters using the Leiden algorithm and generates summaries for each, enabling insight into the dataset's structure even before querying.

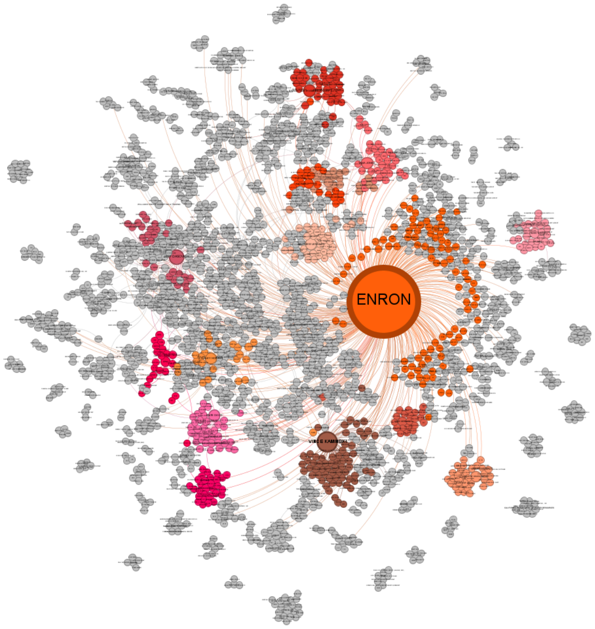

## CAPABILITIES:

**Track Compliance.** Automatically monitor and verify adherence to regulatory requirements across communications and documentation.

**Identify Fraud.** Detect suspicious activity and potential fraud by recognizing patterns and anomalies within corporate data.

**Uncover Inefficiencies.** Pinpoint operational bottlenecks and optimize workflows with targeted insights into communication and task management.

**Enhance Decision-Making.** Generate real-time, data-driven insights that reveal how people, information, and influence flow through your organization, supporting strategic initiatives and improving overall business performance.

## ADVANTAGES:

**The tool's workflow is simple and automated**. Compliance departments, HR, operations teams, internal audit groups, and others can simply upload their data in the form of JSON, CSV, or other structured file types to the tool and receive tailored, actionable insights after indexing is complete.

**The tool can be run locally by existing employees**. Typically, organizations without the ability to analyze these types of data internally will outsource to consulting firms who specialize in taking on these types of projects. This requires legal contracts, manpower, and large budgets. Requiring around $60 in API costs (GPT 4o Mini)  per 100,000 emails indexed, plus negligible costs for custom extraction tuning and querying, this tool offers cost-effective, rapid insights compared to outside intervention, which enables continuous monitoring, scalability, and adaptability.

**The tool is configurable.** A one-size-fits-all solution typically fits no one. We have built in custom extraction tuning to ensure that any type or format of communication can be indexed properly to draw the most pertinent insights into your organization.

## REAL WORLD USE CASES:

**Financial Fraud Detection:** In banking and finance, fraudulent schemes often involve complex networks of transactions and communications. Our tool can ingest emails, transaction logs, and customer data to build a comprehensive knowledge graph of relationships (accounts, transfers, personnel). By enabling multi-hop reasoning, the system can **uncover hidden links** – e.g. connecting an employee's email about an "urgent deal" to a series of unusual transactions across accounts, revealing a possible money laundering ring. Traditional fraud detection systems might flag isolated anomalies, but GraphRAG surfaces the *context* and network behind those anomalies. This leads to earlier detection of complex fraud patterns that involve collusion or cross-silo activity. Moreover, the structured output provides an **audit trail** for investigators:

they can trace how a suspicious entity is connected through the graph (for example, mapping out all intermediaries between a corrupt official and shell companies). Financial institutions face billions in fraud losses annually; our tool offers a proactive approach to identify threats that would otherwise go unnoticed until too late.

**Corporate Compliance & Risk Monitoring:** Large enterprises generate massive communication trails (emails, chat transcripts, reports). Ensuring **regulatory compliance** and internal policy adherence (e.g. anti-corruption, insider trading rules, data privacy) is a major challenge. Our tool can function as an intelligent compliance assistant: it can answer questions like "*Who discussed project X outside approved channels?*" or "*Were any executives communicating with competitors around the bidding period?*". By mapping the *who, what, when* of communications, the tool spots **anomalies and conflicts of interest**. For instance, if an employee in accounting suddenly appears in frequent email contact with a vendor right before a contract award, our tool would flag that subgraph for review. Compared to simple keyword scanning (which often misses nuanced phrasing or indirect references), graph-based RAG understands context – it knows if a discussion about "offshore partnerships" involves people in the tax department and coincides with certain file exchanges, possibly indicating a compliance risk. Importantly, explainability is a key asset here: compliance officers can get a clear narrative of how a potential violation was detected, increasing trust in AI monitoring. This addresses a pain point in regulated industries: Black-box AI is often not acceptable, whereas a graph-backed AI provides **transparent, explainable insights**.

**Legal Discovery and Investigations:** In legal cases (litigation or corporate internal investigations), teams must sift through troves of documents and correspondence to find relevant evidence. Our tool can accelerate e-discovery by organizing information into an entity-relation graph. A lawyer could query, for example, "*Find communications linking Person A, Project B, and Issue C*" – the tool would retrieve and even summarize the chain of emails or memos connecting those dots, something that could take humans weeks. It can answer high-level "global" questions like *"What are the main themes present in this corpus of documents?"*, which naive search tools fail to adequately address. By capturing the **relationships** between documents in an entire document collection, this approach can help legal teams quickly identify clusters of interest (e.g. all documents related to a specific merger) and drill down efficiently. The result is faster case building and the ability to uncover subtle connections (perhaps a minor figure who serves as a link between major parties) that keyword search would miss.

## FREQUENTLY ASKED QUESTIONS

### Who is going to use this and why?

We envision any team within an organization that is concerned with compliance, fraud, or efficiency to get great value from this tool. Early adopting teams would be compliance departments, HR, and operations teams. Typically firms without the ability to analyze these items internally will outsource to consulting firms, which make up a $316.5B market[1]. Teams that do have the time and money to do this internally often have to have specialists on the payroll. This tool can eliminate, or alleviate, a large portion of that spend.

**Has this been done before? How does it compare to existing solutions?**
Standard approaches to compliance monitoring rely on comparing semantic similarity of email text to a list of unapproved terms. While typically done in real-time, this approach analyzes the email text in a vacuum, without context of who is speaking with whom, or their prior conversation history. Our solution harnesses the capability of knowledge graph-based RAG, along with automatically tuned extraction prompts, to analyze each email within the greater context of the organization using existing GraphRAG python libraries published by Microsoft. Table 1 in the appendix below compares our solution to some current tools and methods.

**How can users trust that the results are accurate and unbiased?**
Our solution builds on Microsoft GraphRAG's libraries, utilizing LLM's to generate custom entity extraction prompts based on the format and nature of the data being processed. In the case of email data, the user simply uploads either a JSON or CSV file of their target email corpus (The downloading of emails is left to the user, because each organization's productivity suite may be different). Our solution is meant to be a tool for internal audit teams to help flag and pinpoint potentially non-compliant activity, which can then be further investigated by a direct audit of an employee's full communication activity. As such, the results generated by our tool should still be reviewed by a human and investigated according to each company's compliance policies.

**How does it work? Does it use any third-party models, or is it proprietary?**
GraphRAG constructs a knowledge graph using an LLM, where nodes represent key entities and edges define their relationships. It then partitions this graph into a hierarchy of closely related communities. An LLM generates summaries for these communities in a bottom-up manner, incorporating lower-level insights into higher-level summaries. These hierarchical summaries offer a global understanding of the corpus. For query answering, GraphRAG employs a map-reduce approach: community summaries generate partial answers in parallel (map step), which are then combined to produce a final response (reduce step).

**Can I see an example of its output?**
Yes. We can provide sample analytics dashboards that highlight compliance warnings, flagged fraud cases, and recommendations for operational improvement. These outputs are customizable to an organization's specific requirements. See Figure 2 for an example output. In this case, the tool was fed email data from Enron (renamed VoltGlass to avoid look-ahead bias), and is able to recognize potential manipulation of financial metrics using only the data and no outside or historical knowledge of the scandal.

**What kind of content can the tool analyze?**
The tool works best with semi-structured data formats such as CSV or JSON. Our tool samples

[1] Allied Market Research, Sept 2023. Available at: https://www.alliedmarketresearch.com/management-consulting-services-market-A19875

the data and uses an LLM to dynamically adjust the entity extraction prompt used in the indexing of the data, so almost all formats are able to be analyzed. Currently our tool is limited to only text inputs, meaning that attachments or media within the email body cannot be processed; however, it is able to extract the names of attachments.

**Does this mean the end of employee privacy?**
Not at all. While the system can provide insights into corporate processes and relationships, it is designed to comply with privacy regulations. Each organization determines the scope of data analyzed based on its policies and legal requirements.

**Where does the data go, and how private is it?**
The tool can be run locally or within a secure cloud environment, depending on an organization's needs. The current library by Microsoft supports OpenAI models, which comply with OpenAI's enterprise data privacy policies here: https://openai.com/enterprise-privacy/. There is potential to run the code with a locally-run LLM given sufficient local hardware capabilities, as well.

**What are the general limitations of our approach using GraphRAG?**
With large corpuses of text, finding the "needle in the haystack" becomes increasingly difficult. As a result, our tool is best employed periodically and should be incorporated into weekly or monthly audit cycles to ensure the most accurate and granular responses.

**What about cost constraints and computational overheads?**
The current methodology employed by our tool is expensive relative to naive RAG-based methods. Entity extraction and community summarization, although tailored to the input data for efficiency, is still token intensive, requiring around 15 times as many input tokens to process compared to the token length of the input data. As a result, the indexing process is relatively slower and may be limited by OpenAI's limitations on daily API calls and input/output tokens. Additionally, each query produced by the user requires a similar multiple of token requirements compared to traditional RAG.

## APPENDIX

**Table 1: Overview of GraphRAG Versus Traditional RAG in Fraud Detection & Compliance Monitoring Approaches.** GraphRAG's blend of structured knowledge and AI reasoning offers superior context integration, the ability to answer broad and complex questions, and improved result accuracy & explainability. This leads to more actionable insights and fewer false alarms, which is a decisive edge for businesses seeking to leverage AI for risk management.

| **Feature/ Criteria** | **Traditional Solutions (Keyword Search, Rule-Based, Baseline RAG)** | **GraphRAG (Our Approach)** |
|---|---|---|
| Data Processing | Isolated text chunks or keywords, little context integration. | Holistic knowledge graph integrates context and relationships across corpus. |
| Query Capability | Limited to exact matches or top-K similar snippets. | Global (dataset-wide) insight queries and local deep dives via graph traversal. |
| Detection of Complex Patterns | Difficult – requires manual correlation of multiple sources. | Automatic multi-hop reasoning connects entities, revealing hidden patterns |
| Accuracy of Answers | Prone to missing context or hallucinating if using LLM without grounding. | High accuracy. LLM is grounded by graph facts, yielding comprehensive, factual answers |
| Explainability | Low – "why was this flagged?" often unclear. | High – graph provides a transparent chain of reasoning for each result |
| Adaptability | Rigid – new rules/keywords must be added manually, slow with new data. | Flexible – LLM extraction generalizes to new data, graph updates can be automated for real-time monitoring. |
| Proven Performance | Often siloed results, misses big-picture queries. | Demonstrated ~70–80% win rate in answer quality vs. baseline RAG on complex queries |

**Figure 2: Example Output (Dashboard Format, Incidences sorted by Risk Level)**

Query: "Review discussions of Value at Risk (VAR). Are there any potential instances of employees trying to manipulate it?"

```
### Dashboard of Top Findings on Value at Risk (VAR) Discussions

#### Summary of Findings

The discussions surrounding Value at Risk (VAR) within the Voltglass community reveal
several instances where employees have expressed concerns about the accuracy and
integrity of VAR calculations. There are indications of potential manipulation or at
least attempts to influence VAR metrics, particularly in relation to trading positions
and market conditions.

#### Key Incidents

1. **Concerns Over VAR Calculations**

      Employees have raised issues regarding the accuracy of VAR calculations,
particularly in relation to the influence of market volatility and trading positions.
For instance, John Arnold highlighted that his VAR was significantly affected by
spread levels and volatility, suggesting that minor changes could lead to substantial
discrepancies in reported VAR [Data: Sources (44192, 51020, 26201); Relationships
(34902, 34903)].

   **Employees Involved:** John Arnold, Frank Hayden

   **% Risk:** 25%

2. **Pressure to Maintain VAR Levels**

    John Lavorato requested to maintain current VAR levels until a specified date,
indicating a desire to manage perceptions of risk rather than allowing natural
fluctuations to reflect in the VAR [Data: Sources (41807, 51020); Relationships
(11310, 116197)]. This raises questions about whether there was an intention to
manipulate the reported risk levels.

   **Employees Involved:** John Lavorato, Rick Buy, Ted Murphy

   **% Risk:** 20%

3. **Increased VAR Due to Specific Positions**

    Kirstee Hewitt noted that a copper option position may have increased the VAR by
more than $500,000, which suggests that specific trading decisions could be
influencing the overall risk profile in a way that might not accurately reflect the
underlying market conditions [Data: Sources (11137, 34819); Relationships (45766,
11137)].

   **Employees Involved:** Kirstee Hewitt, Vince Kaminski

   **% Risk:** 15%

4. **Manipulation of Factor Loadings**

  Ted Murphy discussed the importance of factor loadings in VAR calculations and
expressed concerns that the process was not being adequately controlled, which could
lead to misrepresentation of risk [Data: Sources (27043, 43840); Relationships (34902,
34903)].

   **Employees Involved:** Ted Murphy, Rick Buy, John Lavorato

   **% Risk:** 10%
```