

Learning from Litigation: Graphs for Retrieval and Reasoning in eDiscovery

Sounak Lahiri^{1*}, Sumit Pai^{2*}, Tim Weninger^{3*}, Sanmitra Bhattacharya^{4†},

^{*}Deloitte & Touche Assurance & Enterprise Risk Services India Private Limited,
Bangalore, India

[†]Deloitte & Touche LLP, New York City, USA

^{*}University of Notre Dame, Notre Dame, USA

¹solahiri@deloitte.com, ²sumpai@deloitte.com, ³tweninger@nd.edu, ⁴sanmbhattacharya@deloitte.com

Abstract

Electronic Discovery (eDiscovery) requires identifying relevant documents from vast collections for legal production requests. While artificial intelligence (AI) and natural language processing (NLP) have improved document review efficiency, current methods still struggle with legal entities, citations, and complex legal artifacts. To address these challenges, we introduce DISCOGraph (DISCOG), an *emerging* system that integrates knowledge graphs for enhanced document ranking and classification, augmented by LLM-driven reasoning. DISCOG outperforms strong baselines in F1-score, precision, and recall across both balanced and imbalanced datasets. In real-world deployments, it has reduced litigation-related document review costs by approximately 98%, demonstrating significant business impact.

1 Introduction

During legal proceedings, such as investigations, regulatory reviews, and litigation, parties engage in a legal process called *discovery*, formally requesting relevant documents from opposing parties. Traditionally, this involves manually sifting through vast document repositories, a slow and costly process prone to human error. Electronic discovery (eDiscovery) encompasses the collection, review, and organization of digital documents, such as emails, contracts, and articles, to identify those relevant to discovery requests. Technology-assisted review (TAR) typically involves iterative workflows in which skilled professionals annotate documents for relevance guiding supervised learning models in prioritizing documents for review. Early TAR workflows relied on Boolean text queries but have since evolved to incorporate ranked retrieval, relevance feedback, and active learning techniques (Sansone and Sperli, 2022). Recently, *predictive coding*, which trains binary text classifiers to determine whether a document is relevant to a production re-

quest, has gained widespread use (Brown, 2015). Large Language Models (LLMs) have also been explored for document relevance classification in eDiscovery (Pai et al., 2023). However, these text-only models struggle to effectively capture entities, citations, and other complex legal information frequently found in legal production requests, limiting their adoption. To address these challenges, we introduce DISCOGraph (**DISCOG**), a novel *emerging* approach that constructs a knowledge graph from the complex structural information within document corpus and leverages it to enhance document classification and ranking.

DISCOG frames the eDiscovery problem as a link prediction task within a knowledge graph, augmented by a Large Language Model (LLM) for reasoning. The graph consists of documents (*e.g.*, email subjects and bodies from the EDRM corpus), topic statements, senders, and receivers. Keywords and keyphrases extracted from documents serve as additional nodes, with semantically similar keywords linked to enhance structural richness. Document relevance is determined through link prediction between document and topic nodes, where a document is classified as relevant if a link exists between them. To model the knowledge graph, DISCOG employs representation learning techniques, including Knowledge Graph Embedding (KGE) methods such as TransE (Bordes et al., 2013) and ComplEx (Trouillon et al., 2016), as well as Graph Neural Networks (GNNs) like GraphSAGE (Hamilton et al., 2018). The trained model ranks documents by prediction probability, selecting the top K documents (determined by a pre-defined recall threshold, typically 80% (Halskov and Takeda, 2013)) for further reasoning via LLMs. Fig. 1 provides an overview of the DISCOG framework.

Due to the confidentiality of legal discovery processes, direct experimentation on real-world litigation is not feasible. Instead, we evaluate DISCOG

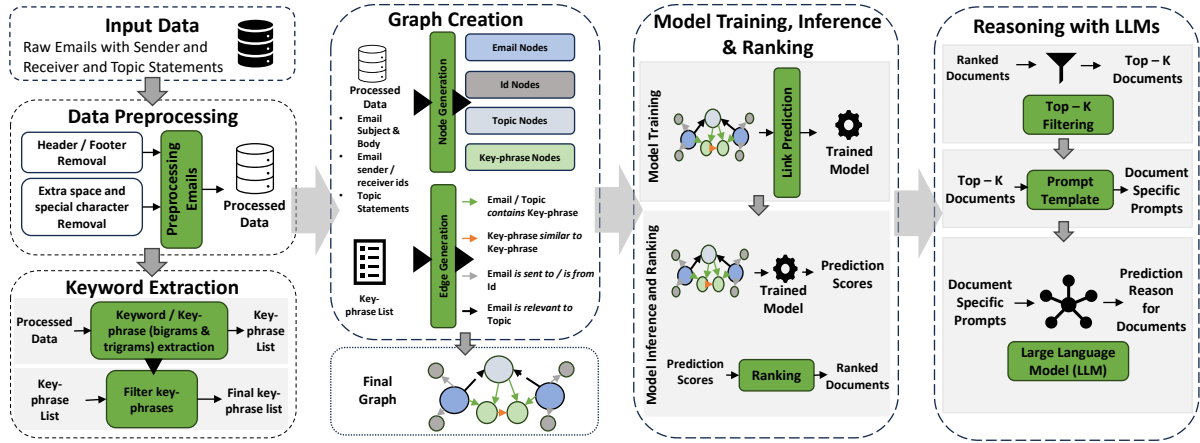


Figure 1: DISCOG: A heterogeneous graph-based approach for predictive coding and ranking in eDiscovery

using the publicly available Electronic Discovery Reference Model (EDRM) Enron Emails Dataset, previously used in the Text Retrieval Conference (TREC) Legal Track (2009–2011)¹ (Hedin et al., 2009; Grossman et al., 2011). This dataset, which includes production requests and human-labeled relevance judgments, remains a benchmark for NLP research on LLM applications (Li et al., 2024; Chen et al., 2024; Huang et al., 2024) and graph-based methods (Shakiba, 2023; Nouranizadeh et al., 2024). By demonstrating DISCOG’s effectiveness on this established benchmark, we showcase its potential for real-world eDiscovery tasks.

2 Related Work

Prior research on the EDRM Enron dataset has primarily employed traditional information retrieval (IR) techniques (Grossman et al., 2011; Robertson and Zaragoza, 2009), where queries are executed against a document index to generate ranked lists of relevant documents.

Transformer-based architectures (Vaswani et al., 2017) have transformed NLP by enabling cross-domain knowledge transfer with limited training data (Raffel et al., 2020). Models such as Contextualized Late Interaction over BERT (ColBERT) (Khattab and Zaharia, 2020) and its improved variant, ColBERT v2 (Santhanam et al., 2022), leverage contextualized embeddings and late interaction mechanisms to enhance document ranking. For adaptation to the legal domain, (Yang et al., 2021) pre-trained BERT on legal data and fine-tuning based on human review for active learning. Recently, large language models (LLMs) have

been used in several use-cases for identify relevancy based on semantic relations and generating responses along with appropriate reasoning. (Pai et al., 2023) experimented with out of the box and fine-tuned LLMs for classification of documents relevant to a topic. (Bron et al., 2024) additionally proposed active learning methods to rank the classifications obtained from LLMs. However, despite their strength in text processing, these models often overlook relational dependencies crucial in legal contexts.

To address this limitation, legal data can be structured as graphs, where documents, topics, and entities form nodes, and relationships define edges. Graph-based methods have been widely applied in areas such as social networks and biomedical research, offering structured representations of interconnected data (Cimiano and Paulheim, 2017). Graph representation learning captures latent semantic relationships by embedding nodes and edges into low-dimensional spaces, optimizing them for classification and link prediction tasks. (Tang et al., 2024b) proposed a text-attributed case graph (TACG) with downstream applications using graph attention trained with contrastive learning methods. (Tang et al., 2024a) builds on top of (Tang et al., 2024b) with updated attention layer to deal with both nodes and edges and graph augmentation technique for better learning. (Tang et al., 2024c) creates a Global Case Graph and employs inductive graph learning for various use-cases. (T.y.s.s et al., 2025) and (Louis et al., 2023) provides similar approaches for Statutory Articles.

Two main approaches dominate graph representation learning: (1) Knowledge Graph Embedding (KGE) models and (2) Graph Neural Networks

¹<https://trec-legal.umiacs.umd.edu/>

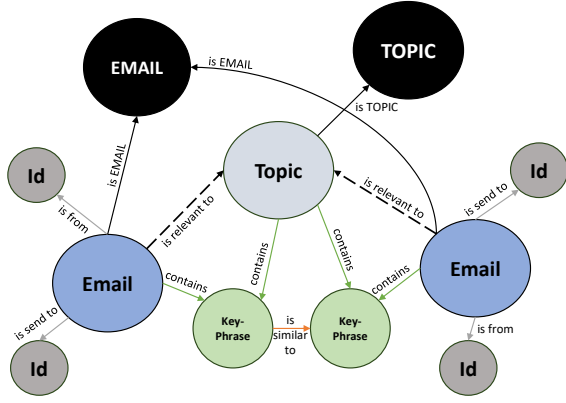


Figure 2: Graph Structure - Schematic Diagram

(GNNs). KGE models, including TransE (Bordes et al., 2013), ComplEx (Trouillon et al., 2016), RotatE (Sun et al., 2019), and DistMult (Yang et al., 2015), generate embeddings through lookup tables and optimize them using scoring functions. GNNs, in contrast, aggregate node features from their neighborhoods over multiple hops (n -hops), enabling more expressive representations (Zhou et al., 2021).

Among GNNs, GraphSAGE (Hamilton et al., 2018) constructs node embeddings by aggregating sampled neighbor information, while Graph Attention Networks (GAT) (Veličković et al., 2018) enhance this by assigning attention scores to different neighbors. Relation Graph Convolutional Networks (RGCNs) (Schlichtkrull et al., 2017) further extend GCNs by incorporating different edge types, making them well-suited for heterogeneous legal data. These graph-based approaches provide a structured way to model complex dependencies in legal discovery, addressing limitations of purely text-based methods.

3 Methodology

This study tackles predictive coding in eDiscovery by constructing a heterogeneous knowledge graph from documents, emails, topic statements, and metadata (e.g., email IDs). Semantic relationships are derived from keywords and keyphrases, and link prediction techniques classify document relevance by predicting links between document and topic nodes. We employ both Knowledge Graph Embedding (KGE) methods (e.g., TransE, ComplEx) and Graph Neural Networks (GNNs) (e.g., GraphSAGE). While KGE methods learn low-dimensional node and edge embeddings, GNNs

aggregate features from neighboring nodes to enhance link prediction accuracy. The trained models rank documents by relevance, with an LLM providing reasoning for predictions—addressing both classification and interpretability in legal document review.

3.1 Dataset

The EDRM Enron Emails dataset, used in the TREC Legal Track (2009–2011), contains 455,449 emails and 230,143 attachments (Grossman et al., 2011). As a case study, we focus on production requests from the 2009 and 2011 tracks, covering seven topics in 2009 and three in 2011. Each topic includes a seed document set for training and *qrels*, which provide human-assessed relevance judgments for evaluation. The full topic details and data distribution are provided in Appendix A.1.

3.2 Baselines for Predictive Coding

We evaluate DISCOG against two widely used predictive coding baselines in eDiscovery:

BM25L: A standard IR model that ranks documents based on query relevance (Lv and Zhai, 2011). In our setup, the topic statement serves as the query, and BM25L computes a relevance score for each document based on term frequency, document length, and other factors.

ColBERT v2: A Transformer-based retrieval model optimized for passage ranking. We use a pretrained ColBERT v2 model with frozen weights and a downstream classifier to refine relevance predictions, leveraging ColBERT’s contextualized embeddings for improved predictive coding.

3.3 DISCOG Graph (DISCOG)

DISCOG employs a graph-based predictive coding approach in three stages: (1) it constructs a heterogeneous knowledge graph from extracted keywords, documents, topics, senders, and receivers; (2) it applies predictive coding using KGE methods (TransE, ComplEx) and GNN models (GraphSAGE, GAT, RGCN) to learn relationships and improve classification accuracy; (3) the trained model ranks documents by predicted relevance to topic statements, capturing complex relational dependencies to enhance predictive performance.

3.3.1 Graph Construction

To harness relational structures in text data, we construct a heterogeneous knowledge

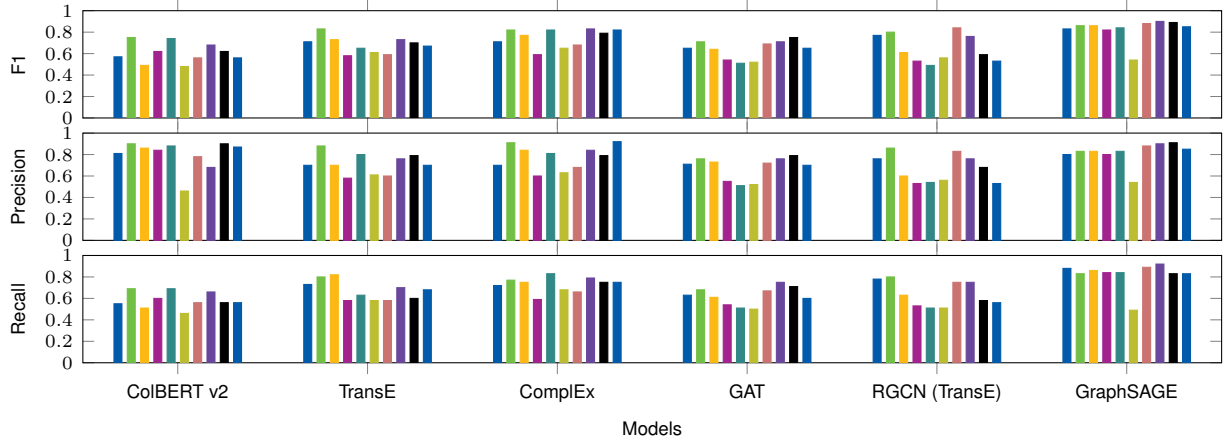


Figure 3: Predictive coding performance of baselines and graph-based models. Grouped bars represent Topics 201–403 in numeric order left to right; their specific identification is relevant this illustration.

graph consisting of four node types: *documents/emails*, *senders/recipients*, *topics*, and **keywords/keyphrases**. Keywords/keyphrases are a combination of unigrams, bigrams and trigrams and extracted from documents using the subject and body and from topic statements using KeyBERT (Grootendorst, 2020) and used as distinct nodes. To reduce noise, we retain only keywords appearing in at least five documents. These are connected to one another based on semantic similarity obtained by a cosine similarity score of 0.75 and above.

Most knowledge graph embedding methods are transductive, making inference on unseen nodes challenging (Costabello et al., 2023). To address this, we introduce two master nodes: **DOCUMENT** and **TOPIC**, linking all documents and topics to their respective master nodes. The master node **DOCUMENT** is connected to all nodes obtained from documents and ensures that no isolated nodes are present during inference. Similar connections are followed for topic nodes. These master nodes are only used during transductive embedding generation and are unnecessary for graph neural networks, which are inductive and handle unseen nodes inherently. A schematic diagram of the graph is shown in Figure 2

The graph incorporates links from the seed and qrels sets. For knowledge graph embedding, only positive links—indicating document relevance—are included to align with the open-world assumption (Costabello et al., 2023). In contrast, graph neural networks leverage both positive and negative links, improving their ability to distinguish relevant from non-relevant documents.

3.3.2 Predictive Coding

DISCOG formulates predictive coding as a link prediction task within a knowledge graph. Two modeling approaches are employed: Knowledge Graph Embeddings (KGE) and Graph Neural Networks (GNNs).

For **KGE-based prediction**, TransE and ComplEx learn low-dimensional node embeddings by minimizing triplet loss with multi-class negative log-likelihood (Costabello et al., 2019). Training considers only relevant links from the seed set, represented as triples $\langle Document_i, relevant_to, Topic_j \rangle$. During inference, confidence scores for predicted links are calibrated using ground-truth labels, with a classification threshold optimized for F1-score on the validation set.

For **GNN-based prediction**, node embeddings are initialized using Sentence Transformers and refined via GraphSAGE, GAT, and RGCN, integrated with TransE. Unlike KGE, GNN training incorporates both relevant and non-relevant links, assigning edge values of 1 (relevant) and 0 (non-relevant). A classification head predicts edge labels, and edge scores are thresholded to optimize macro average F1-score during inference.

Both approaches enable DISCOG to classify documents as relevant or non-relevant, leveraging graph structure to enhance predictive coding in eDiscovery.

3.3.3 Ranking and LLM Prediction

Documents are ranked based on edge scores, normalized via min-max scaling for KGE methods, while GNNs use classification probabilities directly. Performance is evaluated using Recall@ k , and re-

sults are benchmarked against BM25L and BERT with a classifier.

Finally, building on [Pai et al. \(2023\)](#), we apply LLMs to explain predictions. The top K ranked documents are selected, and GPT-3.5-turbo is queried Out-Of-Box (OOB) with a prompt designed to validate graph model predictions and generate reasoning. The prompt is upgraded from the work in [Pai et al. \(2023\)](#), to incorporate the prediction results from the GNN or KGE based method, along with the keywords identified from the document, and the overall LLM task is modified to validate the model’s prediction along with a reason to support its decision.

4 Results

We evaluate DISCOG using emails from the EDRM Enron Dataset, excluding attachments. This section details the graph construction, predictive coding and ranking outcomes, and an analysis of cost savings and business impact.

4.1 Heterogeneous Information Network

The final graph consists of 455,449 email nodes, ten topic nodes, and 34,134 keyword nodes extracted from emails and topic statements. Additionally, 103,926 distinct sender/receiver IDs are included. Edges are formed based on email-to-keyword associations, with similar keywords linked. The number of **Emails relevant to Topic** edges varies per topic, determined by the seed and qrels sets used for model training and evaluation.

4.2 Predictive Coding Results

We evaluate classification and ranking performance using qrels. Since BM25L is a ranking algorithm, it is excluded from classification evaluation. The classification results, summarized in Fig. 3, show that the GNN-based GraphSAGE model consistently outperforms others, including RGCN and GAT.

Most topics exhibit highly skewed distributions of relevant and non-relevant cases, leading to lower performance for baseline and KGE-based approaches. Despite this, GraphSAGE maintains strong performance across topics, with the exception of a single topic (#206), which has the fewest relevant seed cases.

4.3 Ranking Results

Following the TREC 2009 and 2011 Legal Track evaluation scheme, we assess ranking performance




Figure 4: Recall@ k plots for topics at different values of k ranging from 0 to total count of emails in the dataset.

using F1-score, precision, and recall at various cut-off values of k , where k represents the number of reviewed documents. BM25L generates a natural ranking, while graph-based methods first classify documents before ranking them by confidence scores. Metrics are evaluated at thresholds: 2000 through 200000, as shown in Figs. 4 and 5.

GraphSAGE consistently outperforms other methods in ranking, achieving high recall across most topics. While knowledge graph models like RGCN and TransE show higher recall at lower k values, overall, GraphSAGE demonstrates superior performance across both balanced and imbalanced datasets.

4.4 LLM Reasoning Results

We use GPT-3.5 Turbo out-of-the-box (OOB) without instruction tuning to generate reasoning for predictions. Keywords extracted from documents guide the LLM, improving interpretability. Sample outputs are shown in Table 1.

Table 1: Reasoning results sample - Two examples have been shown for topic statements Online Trading and Derivative Trading (The predictions are taken from the Graph Model)

Topic Statement	Document	Pred.	LLM Answer & Reason
All documents or communications that describe, discuss, refer to, report on, or relate to the design, development, operation, or marketing of enrononline, or any other online service offered, provided, or used by the Company (or any of its subsidiaries, predecessors, or successors-in-interest), for the purchase, sale, trading, or exchange of financial or other instruments or products, including but not limited to, derivative instruments, commodities, futures, and swaps.	subject: re: Aquila EOL trade assignments The EOL trade assignment letters are prepared off the same form we do our regular assignments. there is no special language in the assignment form about online issues.	Relevant	Yes, the AI model is correct. The email discusses the use of a form for trade assignments, which is related to the trading of financial instruments and products, including derivatives and futures, on EnronOnline or other online services offered by Enron. Therefore, it is relevant to the topic identified by the keywords.
	subject: online trading per mark's conversation, you were going to get back to us on the scoop on France and utility counterparties. any update?	Non-Relevant	Yes, the AI model is correct in its prediction that the email is not relevant to the given topic. The email is about a conversation regarding France and utility counterparties, and there is no mention of any online service offered, provided, or used by Enron for the purchase, sale, trading, or exchange of financial or other instruments or products. Therefore, it does not fall under the scope of the given topic.
All documents or communications that describe, discuss, refer to, report on, or relate to whether the purchase, sale, trading, or exchange of over-the-counter derivatives, or any other actual or contemplated financial instruments or products, is, was, would be, or will be legal or illegal, or permitted or prohibited, under any existing or proposed rule(s), regulation(s), law(s), standard(s), or other proscription(s), whether domestic or foreign.	subject: Credit Susie financial products ("csfb") susan: csfb had a name change earlier this year. can you please see that the blue file is changed? also, please let the other paralegals know what to do in an instance such as this. thanks. sara	Non-Relevant	Yes, the AI model is correct in identifying the email as not relevant to the given topic. The email is a simple request to change a file related to a name change in Credit Suisse Financial Products and does not discuss or describe any purchase, sale, trading, or exchange of financial instruments or products, nor does it mention any legal or regulatory issues related to them. Therefore, it is not relevant to the given topic.
	subject: re: ECT Nymex hedge exemption applications ECT has a policy that prohibits employees from trading in any commodities that ECT trades. please make sure any future trading you may do complies with this policy. thank you.	Relevant	No, the AI model is not correct. The email is not relevant to the given topic. The email only discusses the company policy on trading commodities and does not provide any information or discussion on the legality or permissibility of financial instruments or products. The keywords identified in the email are not directly related to the topic.

Due to the augmentation of the Graph model's prediction result and the keywords identified from the documents, the LLM is better able to justify the reasoning behind the graph model's prediction. In most cases, the LLM agrees with the prediction and formulates a reason based on the observed keywords and its similarity with the topic statement. However, in instances where graph model misclassifies documents, the LLM can correct errors, as seen in the fourth example. This provides a second level of check to correct misclassifications by the graph model. By combining the Graph model with LLM-based reasoning, DISCOG enhances analysis accuracy, with the LLM acting as a validation and correction mechanism.

5 Deployment and Business Impact

DISCOG seamlessly integrates with existing eDiscovery solutions, significantly reducing the manual review workload. The heterogeneous graph can be constructed from similar databases on any system and stored on prem or in dedicated databases. The models used for prediction are light-weight and can be run on any infrastructure, with or without GPUs, while the LLM can be used from cloud services or open-source depending on the use-case and cost availability. Experiments on the ENRON dataset show that DISCOG achieves 80% recall while requiring review of less than 10% of the dataset. The approach scales efficiently to large eDiscovery datasets with minimal modifications, reducing false positives while maintaining low false




Figure 5: Ranking performance with prediction metrics as a function of k .

negative rates.

According to market reviews in 2023, the document review process constitutes approximately 66% of the total expenditure in the eDiscovery business², with the cost per document review ranging between \$0.50 to \$1.00, varying depending on the experience level of the reviewer and even higher for onsite reviews³. Leveraging DISCOG, deployable on-premise or on a low-cost cloud instance, significantly reduces costs by reducing the number of documents requiring manual review bringing down the overall cost to 10%-20% of the traditional process. For a database with millions of documents,

²<https://complexdiscovery.com/a-2022-look-at-ediscovery-processing-task-spend-and-cost-data-points/>

³<https://edrm.net/2023/12/shaping-ediscovery-strategies-winter-2024-pricing-report/>

DISCOG eliminates majority of the documents, thereby reducing the database size from millions to approximately 10,000 - 20,000 documents, because of its high recall rate. This in turn reduces the overall review cost of the entire corpus to 1%-2% of its original cost, achieving approximately 98% cost reduct. Further explanation and calculations of the cost saving is added in Appendix A.4.

6 Conclusions

We introduce **DISCOG**, a graph-based approach for predictive coding in eDiscovery, outperforming existing solutions in both classification and ranking tasks. Our analysis demonstrates its high accuracy, recall, and substantial cost savings compared to industry-standard methods.

Future work will focus on benchmarking the system's interpretability against manual reasoning and improving further scalability with open source LLMs for on-prem deployments. This hybrid approach aims to provide clear, interpretable justifications for the graph model's predictions, further improving the review process and fostering greater trust in automated document review systems.

References

- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Michiel P. Bron, Berend Greijn, Bruno Messina Coimbra, Rens van de Schoot, and Ayoub Bagheri. 2024. Combining large language model classifications and active learning for improved technology-assisted review. *CEUR Workshop Proceedings*, 3770:77–95. Publisher Copyright: © 2024 Copyright for this paper by its authors.; 8th International Workshop and Tutorial on Interactive Adaptive Learning, IAL@ECML-PKDD 2024 ; Conference date: 09-09-2024.
- Shannon Brown. 2015. Peeking inside the black box: A preliminary survey of technology assisted review (tar) and predictive coding algorithms for ediscovery. *Suffolk J. Trial & App. Advoc.*, 21:221.
- Ruizhe Chen, Tianxiang Hu, Yang Feng, and Zuozhu Liu. 2024. [Learnable privacy neurons localization in language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 256–264, Bangkok, Thailand. Association for Computational Linguistics.
- Philipp Cimiano and Heiko Paulheim. 2017. [Knowledge graph refinement: A survey of approaches and evaluation methods](#). *Semant. Web*, 8(3):489–508.
- Luca Costabello, Alberto Bernardi, Adrianna Janik, Sumit Pai, Chan Le Van, Rory McGrath, Nicholas McCarthy, and Pedro Tabacof. 2019. [AmpliGraph: a Library for Representation Learning on Knowledge Graphs](#).
- Luca Costabello, Adrianna Janik, Eda Bayram, and Sumit Pai. 2023. [Knowledge graph embeddings for nlp: From theory to practice](#).
- Maarten Grootendorst. 2020. [Keybert: Minimal keyword extraction with bert](#).
- Maura R. Grossman, Gordon V. Cormack, Bruce Hedin, and Douglas W. Oard. 2011. [Overview of the trec 2011 legal track](#). In *Text Retrieval Conference*.
- Jakob Halskov and Hideki Takeda. 2013. [When to stop reviewing documents in ediscovery cases: The lit i view quality monitor and endpoint detector](#). pages 227–232.
- William L. Hamilton, Rex Ying, and Jure Leskovec. 2018. [Inductive representation learning on large graphs](#). *Preprint*, arXiv:1706.02216.
- Bruce Hedin, Stephen Tomlinson, Jason R Baron, and Douglas W Oard. 2009. Overview of the trec 2009 legal track. In *TREC*.
- Baixiang Huang, Canyu Chen, and Kai Shu. 2024. [Can large language models identify authorship?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 445–460, Miami, Florida, USA. Association for Computational Linguistics.
- Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 39–48.
- Qinbin Li, Junyuan Hong, Chulin Xie, Jeffrey Tan, Rachel Xin, Junyi Hou, Xavier Yin, Zhun Wang, Dan Hendrycks, Zhangyang Wang, Bo Li, Bingsheng He, and Dawn Song. 2024. [Llm-pbe: Assessing data privacy in large language models](#). *Proc. VLDB Endow.*, 17(11):3201–3214.
- Antoine Louis, Gijs van Dijk, and Gerasimos Spanakis. 2023. [Finding the law: Enhancing statutory article retrieval via graph neural networks](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2761–2776, Dubrovnik, Croatia. Association for Computational Linguistics.
- Yuanhua Lv and ChengXiang Zhai. 2011. [When documents are very long, bm25 fails!](#) In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’11*, page 1103–1104, New York, NY, USA. Association for Computing Machinery.
- Amirhossein Nouranizadeh, Fatemeh Tabatabaei Far, and Mohammad Rahmati. 2024. [Contrastive representation learning for dynamic link prediction in temporal networks](#). *Preprint*, arXiv:2408.12753.
- Sumit Pai, Sounak Lahiri, Ujjwal Kumar, Krishanu Bakshi, Elijah Soba, Michael Suesserman, Nirmala Pudota, Jon Foster, Edward Bowen, and Sanmitra Bhattacharya. 2023. [Exploration of open large language models for eDiscovery](#). In *Proceedings of the Natural Legal Language Processing Workshop 2023*, pages 166–177, Singapore. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: Bm25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Carlo Sansone and Giancarlo Sperl . 2022. Legal information retrieval systems: State-of-the-art and open issues. *Information Systems*, 106:101967.
- Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia.

2022. [Colbertv2: Effective and efficient retrieval via lightweight late interaction](#). *Preprint*, arXiv:2112.01488.
- Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2017. [Modeling relational data with graph convolutional networks](#). *Preprint*, arXiv:1703.06103.
- Ali Shakiba. 2023. [Correlation clustering algorithm for dynamic complete signed graphs: An index-based approach](#). *Preprint*, arXiv:2301.00384.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. [Rotate: Knowledge graph embedding by relational rotation in complex space](#). *Preprint*, arXiv:1902.10197.
- Yanran Tang, Ruihong Qiu, Yilun Liu, Xue Li, and Zi Huang. 2024a. [Casegcn++: Graph contrastive learning for legal case retrieval with graph augmentation](#). *CoRR*, abs/2405.11791.
- Yanran Tang, Ruihong Qiu, Yilun Liu, Xue Li, and Zi Huang. 2024b. [Casegcn: Graph neural networks for legal case retrieval with text-attributed graphs](#). In *ECIR*.
- Yanran Tang, Ruihong Qiu, Hongzhi Yin, Xue Li, and Zi Huang. 2024c. [Caselink: Inductive graph learning for legal case retrieval](#). In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '24*, page 2199–2209, New York, NY, USA. Association for Computing Machinery.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). *Preprint*, arXiv:1606.06357.
- Santosh T.y.s.s, Hassan Sarwat, and Matthias Grabmair. 2025. [QABISAR: Query-article bipartite interactions for statutory article retrieval](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 1496–1502, Abu Dhabi, UAE. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#). *Preprint*, arXiv:1710.10903.
- Bishan Yang, Wen tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. [Embedding entities and relations for learning and inference in knowledge bases](#). *Preprint*, arXiv:1412.6575.
- Eugene Yang, Sean MacAvaney, David D. Lewis, and Ophir Frieder. 2021. [Goldilocks: Just-right tuning of BERT for technology-assisted review](#). *CoRR*, abs/2105.01044.
- Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2021. [Graph neural networks: A review of methods and applications](#). *Preprint*, arXiv:1812.08434.

A Appendix

A.1 Dataset

We primarily concentrate on production requests from the TREC Legal Tracks of 2009 and 2011, which include seven distinct topics for 2009, **Pre-pay Transactions (201)**, **FAS 140 (202)**, **Financial Forecasts (203)**, **Disposal of Documents (204)**, **Energy Loads (205)**, **Company’s Financial Condition (206)**, and **Football Activities (207)** and three distinct topics for 2011, **Online Trading (401)**, **Derivative Trading (402)**, and **Environmental Impact (403)**. The distribution of the seed and qrels sets for each topic is shown in Fig. 6.

A.2 Hyper-parameter Tuning

Hyper-parameter tuning was performed for all models, with a focus on optimizing epochs, learning rate, and batch size. For the KGE methods, the number of epochs ranged from 300 to 600 to achieve reasonable validation loss results. For GNN models, the number of epochs varied from 50 to 150 for GraphSAGE, and from 1000 to 2000 for GAT and RGCNs. Lower learning rates were applied for imbalanced data distributions, with fewer epochs for balanced datasets. The learning rate was adjusted within the range of 0.001 to 0.0001, while batch sizes varied from 128 to 1024 for GNN methods and around 100,000 for KGE methods. The hidden layer vector dimensions for GNNs were also tuned, with values ranging from 32 to 256.

A.3 Ablation Study

To assess the impact of different graph attributes, we conduct an ablation study that systematically evaluates the necessity of various node types within the graph. Given space constraints, this study focuses on three representative topics (401, 402, and 403, which are described as Online Trading, Derivative Trading, and Environmental Trading respectively). These topics capture a range of relevant and non-relevant distributions. However, the DISCOG




Figure 6: Distribution of relevant and non-relevant emails across the seed dataset and qrels for various topics.

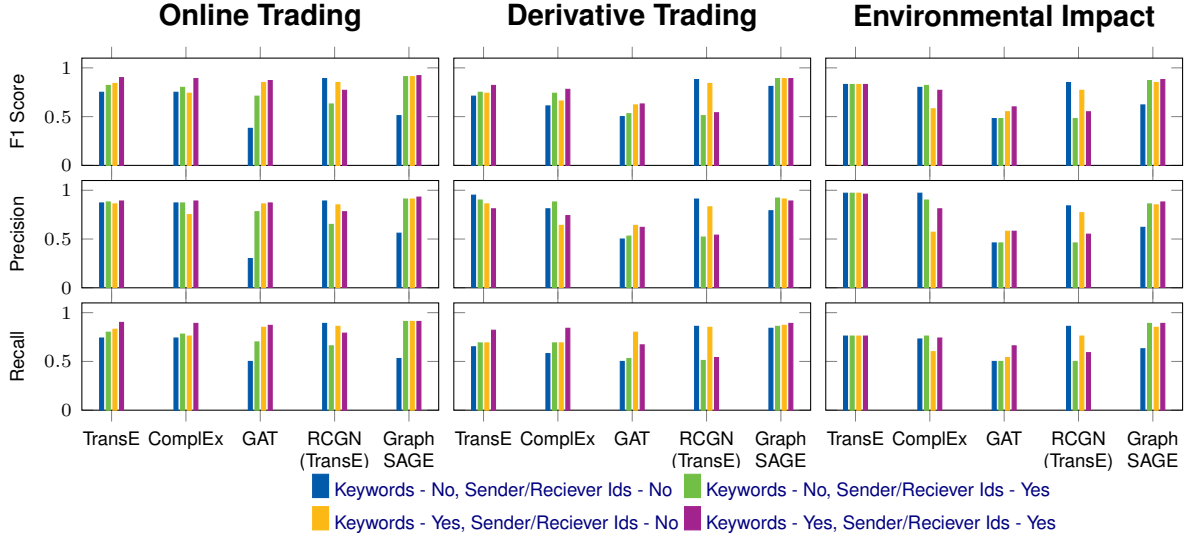


Figure 7: Ablation study of attributes added to the graph in the form of nodes.

methodology can be extended to other topics discussed in this paper.

In this ablation study, the base graph structure consists of two core node types: Emails and Topics, which remain constant across all experiments. To explore the effect of additional features, we incrementally add different combinations of nodes—specifically, keyword nodes and sender/receiver nodes. The influence of these additions is assessed by analyzing their impact on the predictive coding results, measured against the *qrels* set. Importantly, the model architecture and training hyperparameters are held constant across all experiments to ensure that observed differences are solely due to the variations in the graph structure.

The results, as shown in Fig. 7, reveal that incorporating keyword nodes and sender/receiver nodes, along with establishing similarity links between them, leads to a marked improvement in the overall performance metrics of the models. These improvements are consistent across several models, with the exception of the RGCN model, which shows little

to no performance gain from the additional graph attributes. This suggests that the effectiveness of graph augmentation may depend on the underlying model architecture, with some models being more sensitive to additional structural information than others.

A.4 Business Impact Calculations

Assuming review cost per document is \$0.5 - \$1.0, depending on the type of review, the cost of reviewing a million documents for any eDiscovery use case ranges between \$500,000 to \$1,000,000. With the use of DISCOG, the number of documents tagged for review is reduced to 10% -20% of the original corpus, which is approximately 10,000 - 20,000 documents from the entire corpus (assuming a cutoff of 20,000 documents requiring manual review). This cutoff is determined based on the Recall@k metrics, where the DISCOG method with GraphSAGE algorithm achieves over 80% recall. By analyzing the cost of \$0.50 to \$1.00 per document for 20,000 documents instead of the entire

corpus, the cost for the entire corpus is reduced to \$10,000 - \$20,000, which translates to a per document cost of \$0.01 to \$0.02 on average for the entire corpus. Consequently, our method requires only 1% to 2% of the manual review cost.