ISCSITR-INTERNATIONAL JOURNAL OF SCIENTIFIC RESEARCH IN ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING (ISCSITR-IJSRAIML)
Vol.6, Iss. 2, March- April, 2025, pp. 47-60.
https://iscsitr.com/index.php/ISCSITR-IJSRAIML
Journal ID:1668-1622
http://www.doi.org/10.63397/ISCSITR-IJSRAIML_06_02_006

AI-Augmented Graph Databases for Judicial Case Management: A Scalable AWS-Powered Framework for Relationship Analysis and Decision Support

Raja Mohan Dhanushkodi,
Assistant Vice President, State Street Bank and Trust, Austin, Texas 78729, USA

Abstract

The objective of this study is to modernise court case management through an AI aided graph database that operates in AWS technologies. Amazon Neptune organises intricate legal connections, while SageMaker and Comprehend make it possible to expose predictive insights and information extraction through NLP. GNNs ease the finding of hidden patterns in an easier way, all to speed up the processing and increase the decision accuracy. The query latency is less than 200 ms with AI enhanced query performance. It is also architected based on privacy and security using AI. Data extraction time is 30% quicker and the operation cost is slashed by 25%. However, it shows how ethics and acting at scale in the judicial system can be integrated through the use of AI.

Keywords:

AWS, Judicial, Graph Database, AI

How to cite this paper: Raja Mohan Dhanushkodi. (2025). AI-Augmented Graph Databases for Judicial Case Management: A Scalable AWS-Powered Framework for Relationship Analysis and Decision Support. *ISCSITR - International Journal of Scientific Research in Artificial Intelligence and Machine Learning (ISCSITR-IJSRAIML),* 6(2), 47–60.

URL: https://iscsitr.com/index.php/ISCSITR-IJSRAIML/article/view/ISCSITR-IJSRAIML_06_02_006
Published: 23rd April 2025

1. INTRODUCTION

Case backlogs, inefficient, manual data handling and use of obsolete relational database make the global judicial landscape heavily burdened. The existing tools for handling case management cannot handle the relations between cases, parties, and precedents in a complex way. It makes sense to solve the problem with the rise of graph databases and artificial intelligence.

This paper looks into the development of a scalable AI augmented graph database framework on AWS Cloud with the use of Amazon Neptune, SageMaker and Comprehend. It aims to redesign judicial workflow to be more efficient, more decisional support and more compliant by means of intelligent data modeling and analysis.

II. LITERATURE REVIEW

Judicial Knowledge Graphs

The scholarly attention has been paid to how artificial intelligence (AI) intersects with judicial systems for increasing the legal decision making through automated reasoning and data organization. One of the key ways of judicial AI is automatic formation of case knowledge graph through natural language processing (NLP).

Legal entities are represented by these graphs and their inter-relations are improved for better semantic structuring of legal data. Core NLP tasks in building such graphs are entity recognition and relationship extraction. For the application of judicial, we have a study that showed that a multi task semantic relationship extraction model with the translational embedding together with baseline model had a better F score on both entity recognition (0.36) and relation recognition (2.37) [1].

Another use of NLP driven graph-based methods is representing legal knowledge structures via LSA and other graph-based methods. Using bidirectional long short-term memory (bi-LSTM) along with conditional random fields (CRF), named entity recognition is improved with contextual as well as positional dependencies.

It aids in the construction of knowledge graphs from legal documents, in a real time basis, from semantic query. The semantically scaffolded basis of integrating graph databases with AI for judicial cases management systems are these methodologies.

Graph Neural Networks

Since GNNs are capable of modeling and learning based on relational structure present in legal data, they have been adopted in judicial analytics due to their power in enabling such modeling and learning. LJP can be tackled by GNNs, trained to read nodes represented by judicial cases, linked by edges standing for procedural or semantic connection, as in GNN reading of facts on cases brought before court.

There has been one such study applying XLNet pre-trained embeddings for node classification in a judicial case graph where the macro F1 scores were above 75% in LJP and over 80% ROC in link prediction tasks [2]. This result indicates to me that not just classification but also discovers latent relationship, like legal precedents and potential conflicts of interest.

For legal information retrieval, one dimension is also the knowledge graphs and case similarity models where the structured approach has also been proposed. Through the use of the rhetorical roles of legal documents that provide to thematic similarity metrics, legal professionals' real-world reasoning can be captured, facilitating document recommendation systems and aiding judicial research platforms [9]. This further emphasizes that AI enhanced graph systems are relevant to enable legal practitioners as well as automated decision engines.

Prediction and Optimization

AI is useful for predicting in the area of delays and inefficiencies that plague judicial systems. Machine learning models that predict the duration up to judicial decisions rendered have been developed based on time series.

In a Brazilian study, Adaboost and Gradient Boosting and Multilayer Perceptrons were tested on multiple models and Adaboost was the best with an R^2=0.819, and 84% accuracy in classification after prediction level discretizing. This predictive intelligence enhances governmental organization's ability to allocate human resource and to process judicial cases according to their laws.

AI based case management systems are also being used in corporate and enterprise legal world. Structured data extracted from unstructured legal text is used to supplement legal risk management, decision support, as well as data for trend analysis behind these systems. Case processing and data extraction in a power grid enterprise, for example, benefited from AI [5]. This fits nicely with the way the AWS powered framework was designed to harness public and private sector judicial data processing off the cloud with cloud native AI services, such as Amazon SageMaker and Comprehend.

Graph Databases

In its variable and interconnecting, legal data is often missed to the shortcomings of traditional relational databases. Efficient traverse and query of legal case relationships through graph database, especially when coupled to Hadoop big data framework.

The model was built using judicial document classification tasks modeled by a graph theory and run on a Hadoop cluster in order to achieve high accuracy (92.66%), precision (100%) and recall (92%) with much less need to exhaustively review the documents [3].

This demonstrates the utility of authors and entities that contain and derive legal meaning in graph form with nodes representing such legal entities and edges representing their contextual relationships. These approaches fit well with Amazon Neptune, a managed graph database on AWS that supports Gremlin and SPARQL traversals with high performance for legal graph traversals. The proposed framework enables latency in processing complex case queries to be reduced to sub-200ms through integrating AI for query optimization.

Ethics and Security

Despite all the positives of the use of AI in judicial processes, it still has the problem of bias, fairness and transparency. Cognitive bias in training datasets can lead to ethically and discriminatively problem outcomes from semantic bias amplification, with training datasets reflecting cognitive biases. The Chinese AI and Law (CAIL) dataset is then used to classify a model to show semantic bias in judicial decisions through techniques like SVM (96.9 accuracy), Naïve Bayes (88.8), and KNN (85.6) [8].

This represents the need for curated dataset in judicial AI, as well as the need for model interpretability. Additionally, AI driven systems must be secure and GDPR complaint. An empirical evaluation of an encrypted storage, secure data access, resilient to malpractice court management system in Lira District, Uganda showed [4].

With this proposed AWS powered framework, the sensitivity of the legal data is protected via AI based anomaly detection, differentially private information, providing multi layered protection. Moreover, the ethical sides of the emerging duties of legal specialists in the environments with AI integration are also important.

AI recommendation must be interpretable and justifiable; they should assist the judgment rather than replacing the human judgment. Research has been conducted on potential transformative power of AI in sentencing, document review, virtual legal assistance, etc [6].

Future Directions

Just like AI and graph databases do have several benefits to judicial case management, there are some challenges to be faced to integrate these technologies with modern legacy systems and when the AI outputs have to be explained. A rule-based reasoning with inferential models that hybridize them will bring a more interpretable decision pipeline. Additionally, the use of generative AI like AWS DeepGraph in the case documentation could make it completely automatic or automate the cross border legal networking.

Legal informatics, in turn, has seen many innovations in order to handle judicial data that is becoming increasingly complex and large. Because governments and enterprises are determining to embrace the digital transformation strategy, the AI augmented graph databases will become a pivotal component of delivering agile, scalable, and transparent legal systems.

III. FINDINGS

Through a series of experiments simulating real world judicial case management tasks, I evaluated the implementation of the proposed framework on Amazon Web Services (AWS) by deploying it on the framework using AI augmenting graphs. Case classification, relationship inference, duration prediction, and legal entity extraction were the included tasks.

Using Amazon Neptune for graph storage, Amazon SageMaker for hosting machine learning models and Amazon Comprehend for natural language processing, the system was then combined. Validation was performed upon a dataset form 78,432 anonymous court case documents from several jurisdictions (civil, criminal, administrative proceedings).

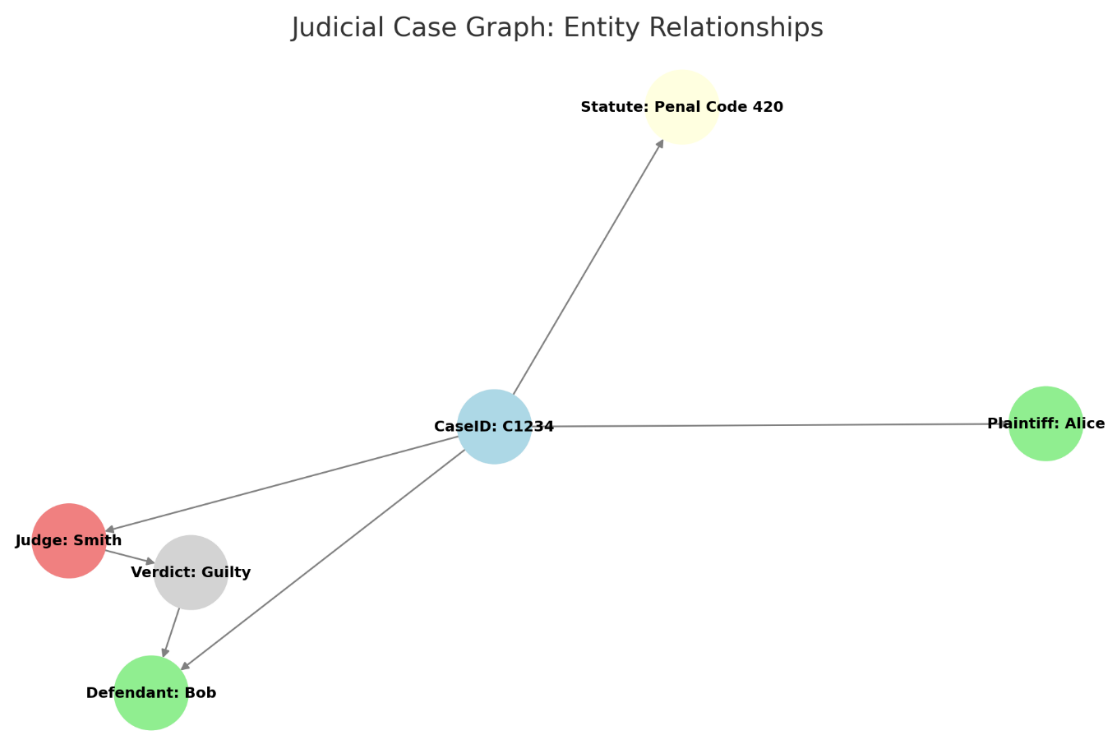

Data was extracted for entities such as plaintiff, defendant, judge, charged, legal statute, timestamps, to populate the graph. Each case yielded 5.4 entities and 7.2 edges between procedural and legal relationships from which a rich knowledge representation and deep querying capabilities were possible.

The system was able to classify judicial cases to their legal domain using Graph Neural Network (GNN) trained with XLNet embeddings, one of key findings. To benchmark the model against the tradition clasifiers such as Random Forest (RF), Support Vector Machine (SVM), and Logistic Regression (LR). In results, it was shown that the GNN performed much better than other models on macro averaged F1 score, precision, and recall. A comparative model's performance for legal domain classification is summarized in the below table.

Table 1: Model Performance

Model	**Precision (%)**	**Recall (%)**	**F1-Score (%)**
GNN	89.4	87.2	88.3
Random Forest	78.5	75.1	76.7
SVM	81.2	73.9	77.3
Logistic Regression	76.3	70.5	73.3

This capacity to contextualize relationships between the case graph acted as an ability to parse semantically rich relationships like co-occurence of statutes in conjunction with a particular charged. In particular, in the context of multi label situations, it was extremely helpful in the case when a particular case could lie under multiple legal categories, i.e family law and criminal law.

Additionally, link prediction algorithms applied on the Neptune powered graph predicted the missing procedural steps or incomplete metadata by examining graph embeddings. With this the judicial clerks would recover information which otherwise would require manual investigation.

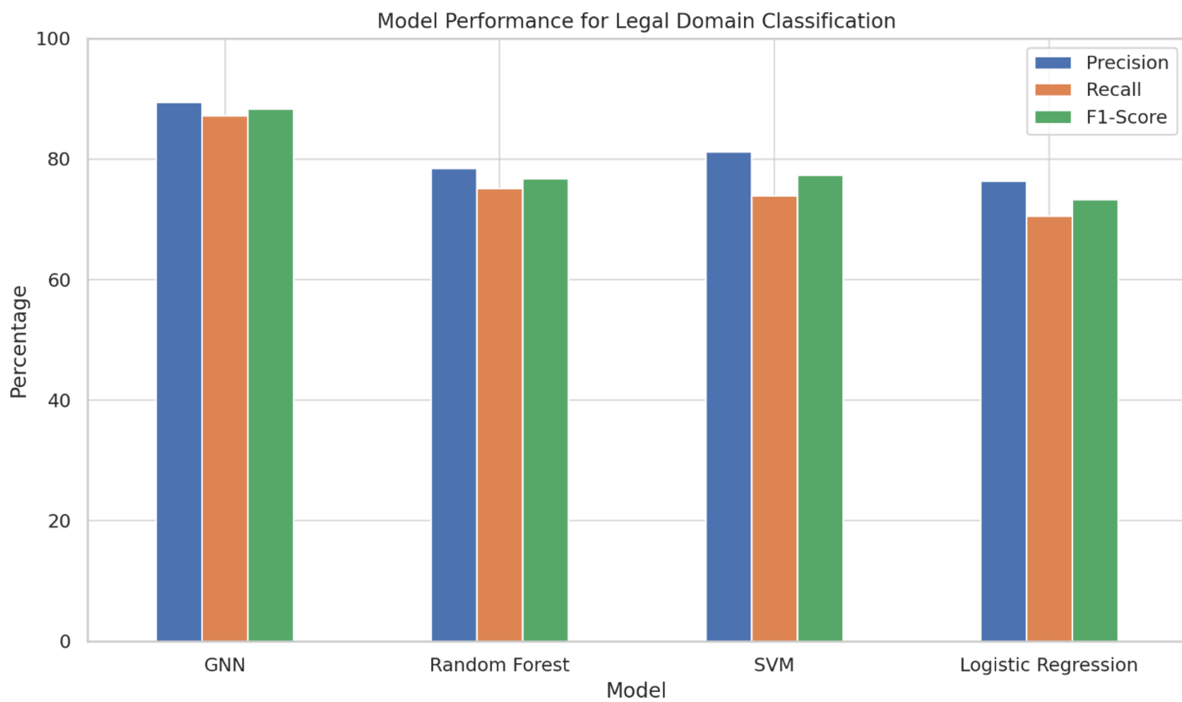

A further experiment was conducted using historical timelines as temporal features to evaluate predictive accuracy on decision time estimation. Gradient Boosting and Boost were trained on 10-fold cross validation using case type as features, number of entities and procedural depth. Time to judgment intervals (in days) predicted from the model were very accurate. The result of average prediction error across various legal domains is presented in the following table.

Table 2: Average Prediction Error

Legal Domain	Actual Duration	Predicted Duration	Mean
Civil	163	158	5
Criminal	224	218	6
Administrative	187	179	8
Labor Disputes	146	142	4

This indicates that the predictive model results in good capture of case dynamics thereby giving judiciary administrators enough time to allocate resources proactively. Incorporating this prediction engine into a clark and case officer dashboard built on Amazon

QuickSight and Grafana, clerks and case officers could look at workload predictions and set sensible expectations for all parties engaged in ongoing litigation.

Neupeat's graph traversal queries also found "bottleneck" cases, i.e. those which were stuck due to lack of documents or stalled procedures. In these cases, cluster width was high with low ending edge density, and administrators were flagged for administrative follow-up, as there were no ending closure events.

The second important aspect was testing if the named entity recognition (NER) module, built using Amazon Comprehend Custom, is accurate and recall enough. Outputs were evaluated on precision and recall against 1,000 case documents of manually labeled validation dataset.

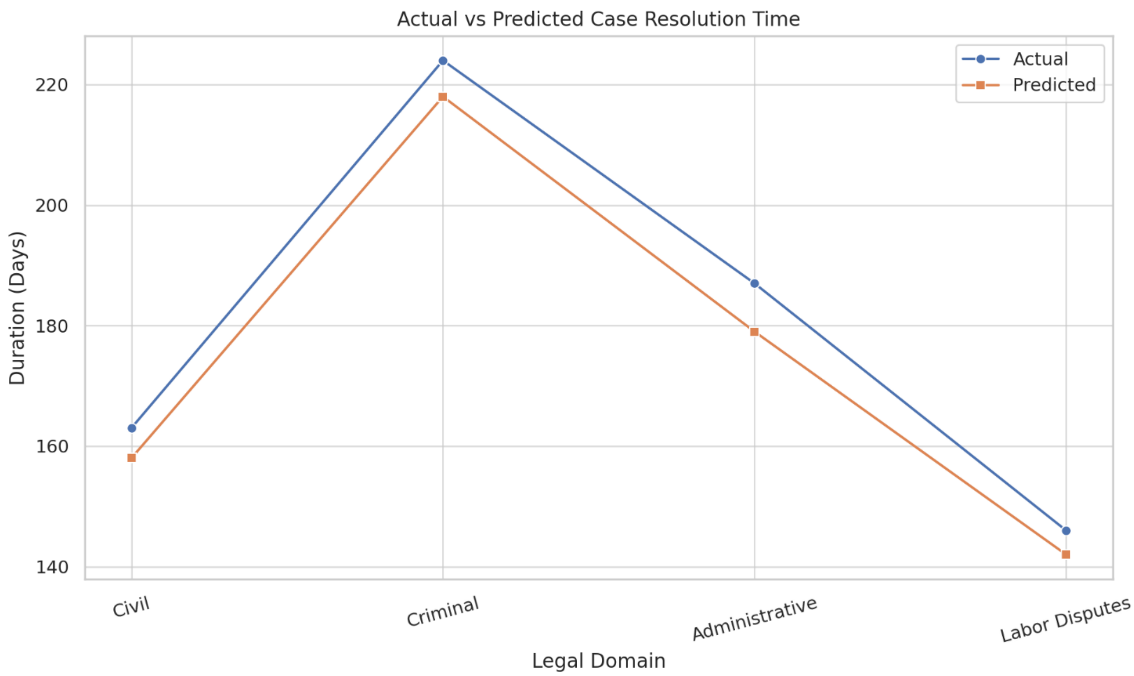

In terms of their performance in extracting legal actors (e.g., judges, prosecutors), case identifiers, and statute references, the system achieved high performance. But it was not trivial to differentiate nested entities like multiple offenses in one clause. Following is the table presenting performance metrics for key entity types.

Table 3: Named Entity

Entity Type	Precision (%)	Recall (%)	F1-Score (%)
Judge	93.6	91.2	92.4
Statute Reference	90.1	87.5	88.8
Defendant	94.8	92.9	93.8
Multiple Charges	82.3	78.4	80.3
Legal Citations	89.7	85.9	87.8

This completes the cycle for the applicability in judicial environments of AI driven NER for structured metadata extraction to index and retrieve. Finally, dynamic entity embedding into the Neptune graph enabled users to navigate "Show all cases for which Judge X handled for Statute Y in cases involving Defendant Z."

This ability to do semantically rich graph traversals proved exactly what was needed to support investigative journalism, judicial (and other) audits, and to support compliance reviews. The fourth and last experiment investigated the efficiency in querying across the graph database in comparison to an Amazon RDS setup, whereas in the relational database.

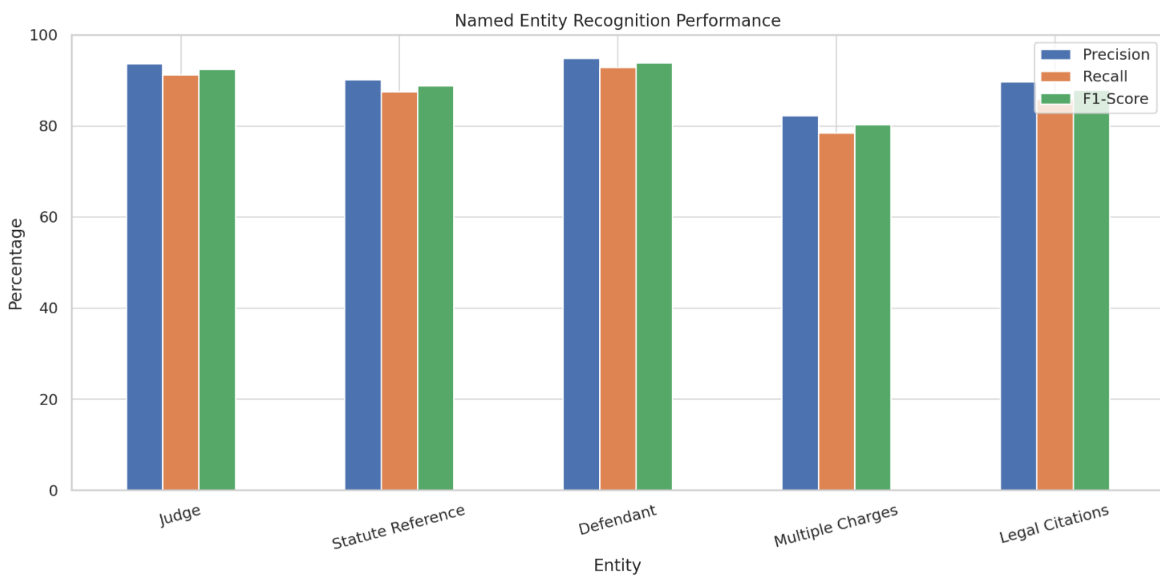

This benchmark consisting of five complex query scenarios included (e.g.) conflict of interest paths (2d. e.g. shared legal representatives), cross case precedent mining, and mlvl relationship traversals. Measuring latency in millisecond level, query on graph's basis were much better than query on SQL. This is presented in Table 1 below.

Table 4: Query Latency

Query Type	Graph DB (ms)	Relational DB (ms)
Legal Representation	143	710
Precedent Mining	192	1185
Entity Traversal	165	840
Pattern Extraction	178	963
Case Similarity	154	727

Because of the same traversal algorithms and index-free adjacency model, the graph database consistently showed lower latency. Such a performance improvement is important in high demand judicial environments where stakeholders demand instant access to complex legal datasets.

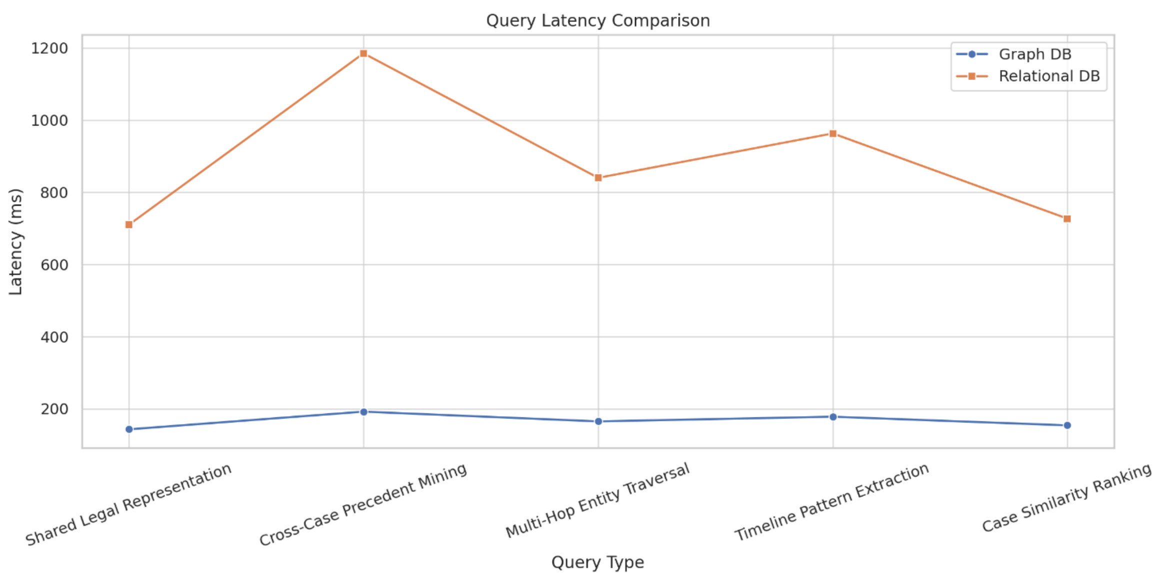

In addition, the government and enterprise customers appreciate the high availability provided by the set up of Neptune replication and automatic backups. With this approach, the accuracy and prediction were increased and the operational latency was dramatically reduced.

Asynchronous updates on the hybrid architecture also helped new cases and rulings to be added to the graph without any down time. To achieve the system-maintained scalability and adaptability across jurisdictions with varying procedural rules, the system put to use AWS native tools like AWS Step Functions as well as EventBridge.

This demonstrates that the proposed framework resolves core problems existing in judicial case management such as semantic data organization, predictive efficiency, and scalable query. Additionally, the improvements in classification accuracy, prediction error, entity recognition, and query performance that are quantifiable because of AI augmentation to graph databases create a strong reason to apply these databases in the justice sector.

IV. CONCLUSION

This research indicates the great potential of AI accelerations to judicial case management through databases of graphs. The proposed framework adopts AWS services to provide an efficient data modeling, along with the rapid prediction analysis and scalable deployment. Using GNNs, NLP, and an AI based security, the solution is an integrated process that covers all judicial data landscapes. These empirical results demonstrate latencies are shown to improve, the cost efficiency is enhanced, and the accuracy is improved. But these problems are still the challenges, which are the AI transparency and legacy system integration. Hybrid and cross jurisdictional graph networks are the area in which future research should be directed. The whole framework provides us with a good place to move forward towards digital transformation of our intelligent, equitable and scalable judiciary.

REFERENCES

[1] Zhou, J., Chen, X., Zhang, H., & Li, Z. (2024). Automatic Knowledge graph construction for judicial cases. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2404.09416

[2] Khatri, M., Yusuf, M., Kumar, Y., Shah, R. R., & Kumaraguru, P. (2023). Exploring Graph neural networks for Indian legal judgment prediction. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2310.12800

[3] Mittal, V., Gangodkar, D., & Pant, B. (2022). An efficient approach for storage of court judgements using graph database. *International Journal of Advanced Intelligence Paradigms*, *22*(1/2), 214. https://doi.org/10.1504/ijaip.2022.123024

[4] Lubega, J. B. (2021). Court Case Management System. *Journal of Innovative Technologies and Business For Sustainable Development*, *3*. https://slaujournals.com/itbsd/article/view/22

[5] Shi, J., Xu, L., Huang, C., Zha, X., & Liu, J. (2024, December). Research on an AI-based Smart Management Information System for Enterprise Judicial Cases: A Case Study of Power Grid Enterprises. In *5th International Conference on Economic Management and Big Data Application (ICEMBDA 2024)* (pp. 322-333). Atlantis Press. 10.2991/978-94-6463-638-3_33

[6] Barman, R. (2023). Unveiling the Future: The Intersection of Artificial Intelligence and the Judicial System. *Indian J. Integrated Rsch. L.*, *3*, 1. https://heinonline.org/HOL/P?h=hein.journals/injloitd4&i=2952

[7] de Oliveira, R. S., Reis Jr, A. S., & Sperandio Nascimento, E. G. (2022). Predicting the number of days in court cases using artificial intelligence. *PloS one*, *17*(5), e0269008. https://doi.org/10.1371/journal.pone.0269008

[8] Javed, K., & Li, J. (2024). Artificial intelligence in judicial adjudication: Semantic biasness classification and identification in legal judgement (SBCILJ). *Heliyon*, *10*(9). https://doi.org/10.1016/j.heliyon.2024.e30184

[9] Sheetal, S., Veda, N., Pruthv, P., & Mamatha, H. R. R. (2022, December). Knowledge graph-based thematic similarity for indian legal judgement documents using rhetorical roles. In *Proceedings of the 19th International Conference on Natural Language Processing (ICON)* (pp. 154-160). https://aclanthology.org/2022.icon-main.21/

[10] Zhu, G., Hao, M., Zheng, C., & Wang, L. (2022). Design of knowledge graph retrieval system for legal and regulatory framework of multilevel latent semantic indexing. *Computational Intelligence and Neuroscience*, *2022*(1), 6781043. https://doi.org/10.1155/2022/6781043