QING HUANG
MARCO SCHREYER
NILSON ROMERO MICHILES JUNIOR
MIKLOS A. VASAERHELYI

A GRAPH SAYS MORE THAN A THOUSAND JOURNAL ENTRIES

Harnessing Graph Autoencoder Networks in Auditing

The integration of artificial intelligence into auditing opens up new pathways. By interpreting journal entries through the innovative lens of graphs, auditors are able to uncover complex accounting patterns and anomalies. This article illustrates how graph representation learning-based auditing can be applied in practice through two practical case studies.

1. INTRODUCTION

The foundations of modern accounting practices, including the recording of journal entries, are likely to have been established during the medieval period [1]. Luca Pacioli, a Franciscan friar, is credited with formalising double-entry bookkeeping in his influential work, *Summa de Arithmetica, Geometria, Proportioni et Proportionalita*, published in 1494 [2]. Pacioli [3] emphasised that accounts have "both likes and opposites", as it is the case with debit and credit, and that merchants should be mindful of their "dual aspects" [4]. Therefore, it is no surprise that accounting is often called the "language of business".

Accounting journal entries are the cornerstone of financial statements and financial reporting. Today, organisations record vast quantities of these entries digitally in Accounting Information Systems (AIS) [5], continuing to apply the same principles that were first established centuries ago [6]. As A. C. Littleton noted in 1928, a

> "journal entry is an important bookkeeping mechanism which serves as a means of converting a non-technical statement of a transaction into a technically-formed record" [7].

Each journal entry contains essential information about a business transaction, such as the date, the accounts to be credited and debited, the amounts involved, and a brief description.

The International Accounting Standard (IAS) 1 requires that financial statements "must 'present fairly' the financial position, financial performance and cash flows of an entity" [8]. Incorrect entries distort a company's reported financial position and performance. They may lead to misstated assets, liabilities, equity, revenues or expenses, affecting all three financial statements. As a result, the audit of the accurate and complete recording of entries is essential to the auditing process. Consequently, the International Standard on Auditing (ISA) 240 requires auditors to

> "test the appropriateness of [...] entries recorded in the general ledger [...] in the preparation of the financial statements" [9].

This audit procedure, usually referred to as journal entry testing (JET), examines entries for the end of the period and throughout the audit period and inquiries about any unusual transaction [10].

The ISA 240 highlights characteristics that indicate fraudulent journal entries.

> "[These] characteristics may include entries (a) made to unrelated, unusual, or seldom-used accounts, (b) made by individuals who typically do not post journal entries, (c) recorded at the end of the period or as post-closing entries that have little or no explanation or description, (d) made either before or during the preparation of the financial statements that do not have account numbers, or (e) containing round numbers or consistent ending numbers." [11]

In parallel, audits have become increasingly complex. Data-driven audit tasks, such as JET, call for the use of advanced

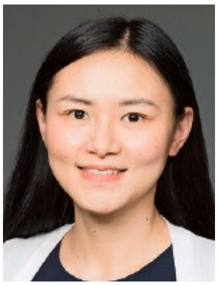

QING HUANG, PH. D., ASSISTANT PROFESSOR OF MANAGEMENT INFORMATION SYSTEMS, LEWIS COLLEGE OF BUSINESS, MARSHALL UNIVERSITY

MARCO SCHREYER, PH. D., AUDIT SPECIALIST, SWISS FEDERAL AUDIT OFFICE (SFAO), POSTDOCTORAL DAAD IFI FELLOW, INTERNATIONAL COMPUTER SCIENCE INSTITUTE (ICSI), BERKELEY

Figure 1: **JOURNAL ENTRY GRAPH REPRESENTATION**

Accounting Journal Entry (Tabular)

Entry-ID	Line-Item	Account-ID	Account	D/C	Amount
100011	0001	10001	Accounts Receivables	D	1,012.00
100011	0002	40101	Sales	C	1,000.00
100011	0003	20220	Value Added Tax	C	12.00

Accounting Journal Entry (Graph)

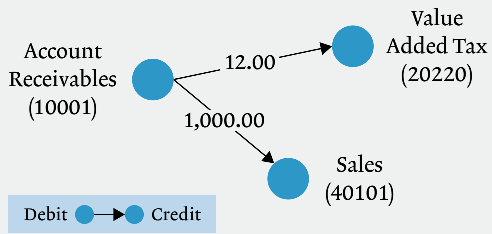

technological tools. The adoption of deep learning [12], a subfield of artificial intelligence, has the potential to enhance such audit procedures [13]. The following section introduces the idea of learning graph representations of journal entries.

2. JOURNAL ENTRIES AS GRAPH STRUCTURES

Risk-based auditing is focused on identifying areas of higher risk to ensure that resources are directed towards effectively identifying material misstatements [14]. Traditionally, auditors have approached journal entries transaction-by-transaction, examining individual entries for anomalies. Over time, this approach evolved into audit sampling as transaction volumes increased. While these methods are valuable, they often fail to capture the "bigger picture" of interconnected transactions. A more comprehensive approach [15] involves viewing each journal entry as part of an integrated financial network or financial graph [16].

Graph theory has been successfully applied in various domains, such as information systems, knowledge graphs, ecosystems, sociology, and biological networks [17]. By reinterpreting journal entries through the innovative lens of graphs [18], auditors gain a structured yet visual means of analysing financial accounting data [19]. In such a graph-based framework:

→ Nodes represent general or sub-ledger accounts on the organisation's balance sheet or profit and loss statement, such as accounts receivable or sales.

→ Edges represent the flow of financial resources within the organisation's accounts, such as debiting in the account receivable and crediting sales.

Figure 1 (left) illustrates a typical sales-related journal entry involving a debit of CHF 1,012 to account 10001 (Accounts Receivables), credit of CHF 1,000 to account 40101 (Sales), and credit of CHF 12 to account 20220 (Value Added Tax). This financial flow is captured in the graph, as also illustrated in *figure 1* (right), in which account 10001 (Accounts Receivables) serves as the source node, while accounts 40101 (Sales) and 20220 (Value Added Tax) act as target nodes. The graph-based interpretation allows for the visualisation of transaction flows, facilitating the detection of patterns and irregularities within the bookkeeping activities.

The analysis of journal entry graphs allows to examine the flow of funds across different accounts, cost centres, or vendors within an organisation. This approach helps uncover hidden relationships and potential anomalies that would be difficult to detect by analysing journal entries individually.

3. GRAPH AUTOENCODER NETWORKS

One method that allows auditors to model these relationships effectively is Graph Autoencoder Networks (GAENs), which was initially proposed by Kipf & Welling [20]. Such neural networks can be used to learn a model of the connections between accounts and the features associated with each account. Once learned, the model can be used to investigate complex patterns in accounting data and detect unusual activities [21].

GAENs learn from two matrices: the adjacency matrix, which maps connections between accounts; and the feature matrix, which adds details about the accounts and transactions, such as the type of account or the amount posted by a journal entry. Together, these matrices provide a comprehensive view of journal entries' structures and attributes, setting the foundation for more advanced learning techniques. *Figure 2* illustrates how a journal entry's graph structure is broken down into a feature and adjacency matrix.

NILSON ROMERO MICHILES JUNIOR, M. SC., INTERNAL AUDIT ANALYST, BRASIL FEDERAL DATA PROCESSING SERVICES, UNIVERSITY OF BRASILIA

MIKLOS A. VASAERHELYI, PH. D., KPMG DISTINGUISHED PROFESSOR OF ACCOUNTING INFORMATION SYSTEMS, RUTGERS BUSINESS SCHOOL, DIRECTOR RUTGERS ACCOUNTING RESEARCH CENTER

Figure 2: **GRAPH DECOMPOSITION INTO ADJACENCY AND FEATURE MATRICES**

Accounting Journal Entry (Graph)

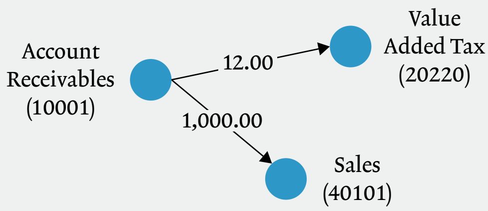

Account Features

	Item	D/C	Amount	...
10001	1	D	1,012.00	...
20220	3	C	12.00	...
40101	2	C	1,000.00	...
...	...	...	...	...

Account Adjacencies

	10001	20220	40101	...
10001	**1**	**1**	**1**	...
20220	0	0	**1**	...
40101	...	...	...	...
...	0	0	0	...

The core idea behind GAENs is that they consist of two main components: an encoder network and a decoder network. The encoder learns a simplified version of the financial graph by capturing the essential relationships between accounts. It compresses this information and passes it on to the decoder. The decoder then attempts to reconstruct the original graph from this compressed representation. Successful reconstruction indicates that the model has captured the journal entries' underlying structures. At the same time, difficulty in reconstructing certain parts may suggest the presence of unusual or irregular journal entries. For example, a high-reconstruction error can indicate a journal entry anomaly that warrants further investigation.

Figure 3 illustrates the GAEN architecture, consisting of a graph encoder and a graph decoder neural network that learn from account features and adjacencies. The account features are used to capture the attributes of each account, such as transaction amounts and document types. In contrast, the account adjacencies represent how accounts are connected through debits and credits within the journal entries. The graph encoder (left) compresses these inputs into simplified graph representations. The graph decoder (right) then attempts to reconstruct the original financial graph from this compressed data. As the learning process progresses, the GAEN becomes more proficient at grouping similar journal entry graph structures into distinct clusters.

Figure 3: **GRAPH AUTOENCODER NEURAL NETWORK (GAEN)** [22]

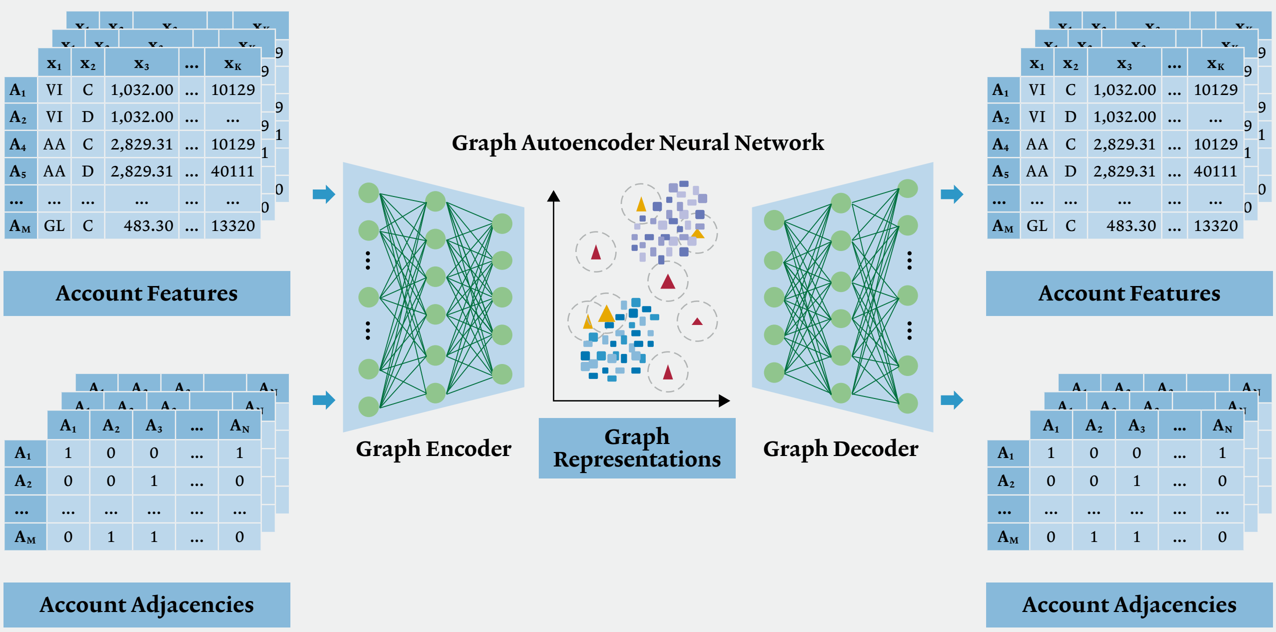

The learned representations enable underlying patterns within the data to be uncovered, providing insights into common accounting patterns ("conformance checking") and deviations from these norms ("anomaly detection").

→ *Conformance Checking:* The GAEN might learn that one cluster corresponds to journal entries in which accounts receivable are debited and sales are credited. Evaluating whether these journal entries conform to the general sales accounting pattern helps determine whether they are recorded in compliance with established accounting policies, improving the statements' reliability.

→ *Anomaly Detection:* The GAEN might also detect an instance in which the credit is split between sales revenue and miscellaneous income. Such a deviation from the established sales cluster may indicate an anomaly. This anomaly could signal a misclassification or an attempt to adjust revenue, warranting further review by an auditor to ensure the accuracy and integrity of the entry.

In brief, examining graph representations provides insights into the interconnected nature of accounts, helping to ensure adherence to accounting practices and identify potential risks.

4. GRAPH REPRESENTATION AUDITING PROCESS

Graph Autoencoder Networks offer a novel extension to the conventional graph representation of journal entries [23], leveraging machine learning to enhance the analysis of accounting data. By learning how accounts interact with one another and how organisational resources flow, GAENs provide an advanced model that can be used to improve the accuracy of financial audits by offering insights into unusual transactions and identifying potential risk areas. The process of applying GAENs to auditing can involve the following five steps.

4.1 Initialisation. In the first step, the initial GAEN model is set up to define how journal entries will be transformed into graphs. The initialisation involves determining which accounts will serve as nodes and how journal entries, representing transactions, will be modelled as edges between these nodes. The setup ensures that the graph structure captures the inherent relationships in the accounting data.

4.2 Data selection. In the second step, the appropriate dataset of journal entries is selected for training the GAEN model. This dataset should include transactions across multiple accounts, providing a rich source of information. The quality of the data chosen will determine whether the model will be able to effectively learn common and uncommon accounting behaviours, improving the model's ability to identify irregularities.

4.3 Graph construction. In the third step, journal entries are converted into a graph format. Each journal entry is represented as an adjacency matrix that defines the relationships between accounts, while the feature matrix adds additional details such as transaction amounts or account types. The matrices enable the model to learn from the graph structure and attributes.

4.4 Model learning. In the fourth step, the GAEN learns to represent journal entries as graphs by processing the adjacency and feature matrices. This learning process allows the model to capture the relationships between accounts, gradually building an understanding of how the entries establish connections between the different accounts and their typical financial flows.

4.5 Representation analysis. In the fifth step, the learned representations of journal entries are examined. The model

Figure 4: **REPRESENTATION SPACE ANALYSIS**

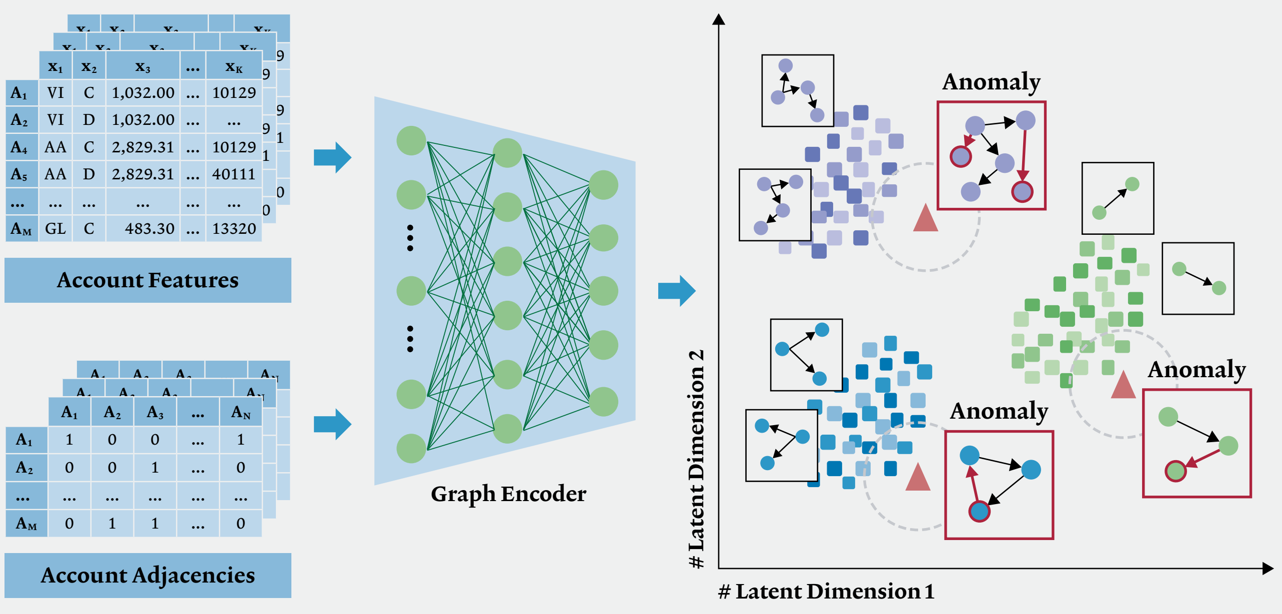

organises these representations into clusters, in which similar entries correspond to identical accounting practices. Journal entry representations that do not fit into these clusters may indicate unusual activity, highlighting unusual accounting practices.

Figure 4 illustrates the representation analysis of a GAEN applied to learn from journal entry data. As the model learns the journal entry representations, distinct clusters emerge, each representing a typical accounting pattern or structure.

As described earlier, these cluster visualisations enable conformance checking, verifying whether journal entries fall into the expected, well-structured clusters based on regular accounting behaviours. Additionally, the model identifies anomalies by detecting deviations from these clusters. Journal entries not conforming to the established clusters are flagged as potential anomalies, signalling transactions that may require further investigation. For example, a regular cluster might represent vendor payments in which accounts payable is debited and a bank account is credited. However, if an entry debits accounts payable but credits an unusual account such as intercompany transactions, this could suggest an unusual activity or misclassification, prompting auditors to investigate further.

In short, by applying GAENs, this provides an effective way to analyse the structure and content of financial accounting data. By learning graph representations of journal entries, these GAENs allow for the identification of typical and atypical accounting behaviours. The ability to detect patterns and flag irregularities makes GAENs a valuable addition to the auditing toolkit in order to uncover hidden risks that might otherwise go unnoticed in audits.

5. EMPIRICAL CASE STUDY: JOURNAL ENTRY TESTING

An empirical case study [24] was conducted to evaluate the application of GAENs for testing and analysing journal entries in the context of audits. The study focuses on two datasets: a case study dataset (referred to as Dataset A) sourced from the EY Academic Resource Center (EYARC) and a real-world dataset from a partner company with $681 million in gross revenue (referred to as Dataset B).

- → Dataset A consisted of approx. 40,000 journal entry line items, corresponding to 18,934 journal entries across 74 general ledger accounts.
- → Dataset B consisted of approx. 1.3 million journal entry line items, corresponding to 128,045 journal entries across 441 general ledger accounts.

The primary objective of the case study was to explore how GAENs are able to model the graph pattern established by journal entries, identify unusual patterns, and assist auditors in detecting potential anomalies that might not be immediately visible using traditional methods.

5.1 Synthetic Dataset A results. The model's anomaly detection capabilities were tested on Dataset A, a controlled set of journal entries with artificially introduced anomalies. These anomalies were divided into two categories: ten global anomalies (unusual individual transactions) and ten local anomalies (unusual patterns between accounts within transactions). The results showed that the GAEN effectively identified the global anomalies in which individual transactions differed from typical patterns. The model's success in detecting these anomalies was mainly due to its ability to analyse the graph structure of journal entries, highlighting transactions with unusual relationships between accounts that may represent potential risks.

5.2 Real-world Dataset B results. The model's capability for practical application was tested by having expert auditors review flagged transactions in Dataset B. The focus was on identifying naturally occurring transactions that deviated from the norm. The model flagged several transactions as anomalies. Upon review, auditors confirmed that many of these transactions were unusual or rare. They found issues such as duplicate records, self-postings (where the same account was debited and credited), and unexplained reversals. These anomalies, which might have gone unnoticed using traditional auditing methods, were effectively identified by the GAEN model due to its ability to analyse distinct graph structures in journal entries.

Across both Dataset A and Dataset B, the GAEN model demonstrated its ability to detect anomalies by analysing the underlying graph structures of journal entries.

5.3 Valuable lessons for the practical application of GAENs in auditing. The capability of GAEN to capture complex relationships within accounting data is of great value to practical applications.

5.3.1 Improved decision-making. Flagged anomalies identified by GAENs can guide the performance of more targeted audit procedures. Audit efficiency and effectiveness can be enhanced by concentrating resources on areas with the highest risk of misstatements or fraudulent activities. GAENs can also assist in prioritising work by indicating which anomalies require immediate attention based on the severity or frequency of irregularities. These insights facilitate informed decisions regarding which transactions to investigate further, enhancing the quality of audit outcomes as a result.

5.3.2 Interpretation challenges. One challenge identified is the difficulty in interpreting the output from GAENs, especially for individuals without the required technical background. Understanding why certain transactions are flagged as anomalies can be complex. To bridge this gap, training programmes focused on understanding graph structures and based on visual aids will be necessary. Visual tools that intuitively present graph structures, such as node-link diagrams or cluster visualisations, can be used to ensure that advanced machine learning techniques can be effectively utilised in auditing practices.

As demonstrated in *figure 5*, a dynamic visualisation for monitoring atypical or unusual transactions, classified by

Figure 5: **GRAPH NEURAL NETWORK JOURNAL ENTRY ANOMALY DETECTION DASHBOARD**

$78.38 Mio. Total Journal Entry Amount	**$0.06 Mio.** Total Anomaly Amount	**0.0007** % Anomaly Count	**2.42 Mio.** Total Journal Entry Count	**468** Total Anomaly Count	**0.0002** % Anomaly Count

Detected Journal Entry Graph Anomalies

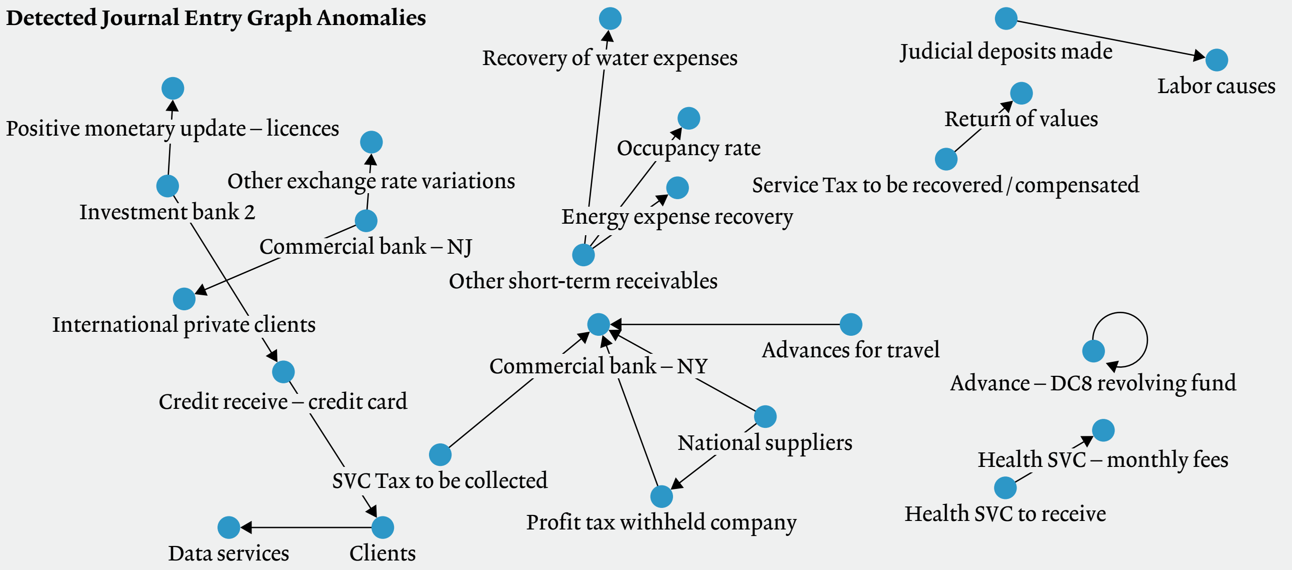

Journal Entry Overview							
Year	**Month**	**Account Number**	**Account Type**	**Account Class**	**Amount**	**Preparer**	**Anomaly Score**
2023	October	3506956	Revenue	Gross Service Revenue	38,542.87	192462	100.00 %
2023	September	4923830	Revenue	Gross Service Revenue	36,970.83	188711	97.18 %
2023	July	2800515	Revenue	Gross Service Revenue	33,031.02	218245	92.39 %
2023	December	2476410	Assets	Current Assets	6,830.642	313658	79.96 %
2023	October	4968625	Liabilities	Current Liabilities	22,230.47	164137	79.06 %
2023	December	3636944	Liabilities	Current Liabilities	18,917.7	256366	77.10 %
2023	August	1633920	Liabilities	Current Liabilities	16,340.52	255734	76.46 %
2023	September	2400193	Liabilities	Current Liabilities	15,035.01	178353	75.64 %
2023	June	2640457	Assets	Current Assets	658.9465	41379	75.44 %

GAEN based on a predefined anomaly score threshold, allows auditors to quickly prioritise high-risk accounts. The top of the dashboard in *figure 5* displays key metrics for the audited data (amounts are shown in Brazilian currency), including total amount, detected anomaly amount, percentage of anomalous amounts, total transaction count, total anomalous transactions, and percentage of anomalous transactions. The middle section of *figure 5* shows the detected journal entry anomaly graphs, where nodes represent ledger account and edges indicate connections between them. Below the anomaly graph is an overview table detailing specific journal entry anomalies. This promises to optimise the efficiency and effectiveness of audits in detecting fraud and ensuring accounting compliance.

In summary, the GAEN model recognised unusual relationships between accounts and transactions, providing insights that complement traditional auditing methods. These findings highlight the practical benefits of applying GAENs, while underscoring the need for strategies to address challenges in interpreting machine learning outputs.

6. CONCLUSION AND OUTLOOK

Graph representation learning in financial auditing has shown promising results, particularly in enhancing anomaly detection by analysing journal entries' graph structures. By transforming traditional journal entries into graph representations, it becomes possible to gain insights into the underlying relationships between accounts, leading to more efficient and effective audits. The GAEN model's ability to identify unusual patterns and relationships has the potential to improve the quality of audits by focusing on high-risk areas and enabling more targeted audit procedures.

Integrating these advanced machine learning techniques into auditing practices presents challenges, however – particularly regarding interpretability and the level of technical expertise required. Addressing these challenges through comprehensive training programmes and the development of intuitive visual tools is crucial to making these technologies accessible and practical for auditing professionals.

Looking forward, GAENs have the potential to become an integral part of AI-assisted or co-piloted auditing [25]. As GAEN models evolve, they will offer more sophisticated tools for analysing finanial data, ultimately leading to more accurate and efficient audits. Integrating GAENs with other data analytics tools and automating aspects of the auditing process could further enhance the effectiveness of audit procedures. The combination of traditional auditing expertise with advanced deep learning technologies has the potential to transform financial auditing into a more data-driven discipline [26] that is better equipped to address the complexities of modern financial statements and their underlying journal entries [27]. ■

Notes: 1) Parker, L. M. Medieval Traders as International Change Agents: A Comparison with Twentieth Century International Accounting Firms. Accounting Historians Journal 16, 2 (1989), pp. 107–118. **2)** Sangster, A. The Printing of Pacioli's Summa in 1494: How Many Copies Were Printed? Accounting Historians Journal 34, 1 (2007), pp. 125–145. **3)** Smith, M. Luca Pacioli: The Father of Accounting. Available at SSRN 2320658 (2018). **4)** Littleton, AC. The Antecedents of Double-Entry Bookkeeping. Contemporary Studies in the Evolution of Accounting Thought (1968), pp. 21–29. **5)** Appelbaum, D., Kogan, A., Vasarhelyi, M. A. Big Data and Analytics in the Modern Audit Engagement: Research Needs. Auditing: A Journal of Practice & Theory 36, 4 (2017), pp. 1–27. **6)** Vasarhelyi, M. A., Halper, F. B. The Continuous Audit of Online Systems. Auditing – A Journal of Practice & Theory 10, 1 (1991), pp. 110–125. 7) Littleton, AC. The Evolution of the Journal Entry. Accounting Review (1928), pp. 383–396. **8)** International Accounting Standards Board (IASB). International Accounting Standard (ISA) 1 – Presentation of Financial Statements. International Accounting Standards Board (IASB), 2007. **9)** International Federation of Accountants (IFAC). International Standards on Auditing (ISA) 240: The Auditor's Responsibilities Relating to Fraud in an Audit of Financial Statements. International Federation of Accountants (IFAC), 2009. **10)** American Institute of Certified Public Accountants (AICPA). Statement on Auditing Standards No. 99: Consideration of Fraud in a Financial Statement Audit. American Institute of Certified Public Accountants (AICPA), 2002. **11)** Schreyer, M., Baumgartner, M., Ruud, T. F., Borth, D. (2022). Artificial Intelligence in Internal Audit as a Contribution to Effective Governance – Deep-Learning enabled Detection of Anomalies in Financial Accounting Data. Expert Focus 2022/January, pp. 39–44. **12)** LeCun, Y., Bengio, Y., Hinton, G. Deep Learning. Nature 521, 7553 (2015), pp. 436–444. **13)** See note 11 and Schreyer, M., Gierbl, A., Ruud, T. F., Borth, D. (2022). Artificial Intelligence Enabled Audit Sampling – Learning to Draw Representative and Interpretable Audit Samples from Large-Scale Journal Entry Data. Expert Focus 2022 /April, pp. 106–112. **14)** Marten, K.-U., Föhr, T. L., McIntosh, S. I. (2022). KI-basierte Datenanalysen und risikoorientierter Prüfungsansatz. WPg: Kompetenz schafft Vertrauen, 75(16), 898–908. **15)** Arya, A., Fellingham, J. C., Glover, J. C., Schroeder, D. A., Strang, G. (2000). Inferring transactions from financial statements. Contemporary Accounting Research, 17(3), pp. 366–385. Arya, A., Fellingham, J. C., Mittendorf, B., Schroeder, D. A. (2004). Reconciling financial information at varied levels of aggregation. Contemporary Accounting Research, 21(2), pp. 303–324. **16)** Boersma, M., Wolsink, J., Sourabh, S., Hoogduin, L. A., Kandhai, D. (2023). Measure cross-sectoral structural similarities from financial networks. Scientific Reports, 13(1), 7124. **17)** Xia, F., Sun, K., Yu, S., Aziz, A., Wan, L., Pan, S., Liu, H. (2021). Graph learning: A survey. IEEE Transactions on Artificial Intelligence, 2(2), pp. 109–127. **18)** Mochty, L., Wiese, M. (2011). Die Netzwerkstruktur der Buchhaltung als Grundlage des risikoorientierten Prüfungsansatzes (No. 188). IBES Diskussionsbeitrag. **19)** Guo, K. H., Yu, X., Wilkin, C. (2022). A picture is worth a thousand journal entries: accounting graph topology for auditing and fraud detection. Journal of Information Systems, 36(2), pp. 53–81. **20)** Kipf, N. T., Welling, M. (2016). Variational graph auto-encoders (2016). In Neural Information Processing Systems Workshop on Bayesian Deep Learning. **21)** Huang, Q., Schreyer, M., Michiles, N., Vasarhelyi, M. (2024). Connecting the Dots: Graph Neural Networks for Auditing Accounting Journal Entries. Available at SSRN 4847792. **22)** See note 21. **23)** Guo, K. H., Yu, X., Wilkin, C. (2022). A picture is worth a thousand journal entries: accounting graph topology for auditing and fraud detection. Journal of Information Systems, 36(2), pp. 53–81. **24)** See note 21 and Huang, Q. (2024). Three Essays on Anomaly Detection in Accounting: Channel Stuffing, Journal Entry Testing, and Graph Learning (Doctoral dissertation, Rutgers The State University of New Jersey, Graduate School-Newark). **25)** Gu, H., Schreyer, M., Moffitt, K., Vasarhelyi, M. (2024). Artificial Intelligence Co-Piloted Auditing. International Journal of Accounting Information Systems, 54, 100698. **26)** Sun, T. (2019). Applying deep learning to audit procedures: An illustrative framework. Accounting Horizons, 33(3), 89–109. **27)** Gierbl, A. S., Schreyer, M., Borth, D., Leibfried, P. (2021). Deep Learning in der Wirtschaftsprüfung: Disruptive Schlüsseltechnologie im Einsatz. Zeitschrift für Internationale Rechnungslegung (IRZ), 7/8, 349–355. C. H. Beck Vahlen Verlag.