

Dragana Radojičić
University of Belgrade -
Faculty of Economics and Business

Dragan Azdejković
University of Belgrade -
Faculty of Economics and Business

Jelena Stanojević
University of Belgrade -
Faculty of Economics and Business

UDK 004.415:519.876:005.8

APPLICATION OF MACHINE LEARNING METHODS TO BENFORD'S LAW ANALYSIS: POTENTIALS, AND LIMITATIONS

Abstract

Benford's Law has long served as a valuable statistical tool for detecting anomalies in numerical datasets, particularly in fraud detection, auditing, and forensic accounting. Traditionally, its application relies on hypothesis testing methods such as the Chi-square test, Kolmogorov–Smirnov test, or the Mean Absolute Deviation. Recent years have seen growing interest in enhancing Benford's Law-based analyses through the application of machine learning (ML) techniques, particularly in contexts where traditional statistical methods face limitations. This study surveys the current state of the art at the intersection of Benford's Law and machine learning, examining both supervised and unsupervised approaches for detecting deviations from expected digit distributions. We also review methodological advances across diverse domains such as financial auditing, digital forensics, and fraud detection, and outline promising directions for future research and applications. Special emphasis is devoted to future research, addressing its possible directions and limitations, along with data-related and methodological challenges. We conclude by outlining promising avenues for further study and practical implementation.

Key words: *Benford's Law, anomalies, forensic accounting, auditing, machine learning (ML) methods, anomaly detection algorithms*

JEL Classification: *M40, C15*

Introduction

Reports of economic offenses, including tax fraud, money laundering, corruption, and the illegal diversion of public resources, frequently appear in society. Such white-collar offenses require advanced techniques for effective identification and legal action. In this context, Varian (1972) proposed the application of Benford's Law as a diagnostic tool for identifying economic and financial datasets that warrant more detailed forensic scrutiny.

Benford's Law was originally documented by the astronomer and mathematician Simon Newcomb in 1881 and was subsequently rediscovered by the physicist Frank Benford in 1938. The law describes a specific logarithmic pattern in the distribution of leading digits across many naturally occurring numerical datasets. Contrary to the assumption of uniformity, the law indicates that lower digits appear as the first significant digit far more often than higher ones. Benford's Law has been applied across various disciplines, such as forensic accounting and fraud examination, to assist in the identification of anomalies and instances of potential data manipulation.

The application of Benford's Law in the domains of finance and accounting is widely recognized and utilized. Nigrini (1992) introduced Benford's Law for identifying fraud in tax and accounting records, and later supported this application with additional scientific evidence. In particular, Nigrini (1994), asserted that "individuals, either through psychological habits or other constraints peculiar to the situation, will invent fraudulent numbers that will not adhere to the expected digital frequencies", while Nigrini and Mittermaier (1997) advocated for its adoption as a standard analytical tool in auditing practice.

Over time, a distinct line of inquiry within accounting research has emerged, applying Benford's Law to identify fraudulent activity or the artificial manipulation of numerical data in accounting and financial documentation (Be-neish, 1999; Durtschi, Hillison, & Pacini, 2004; Asllani & Naco, 2014). Several empirical studies illustrate this approach. For example, Quick and Wolz (2003) analysed income and balance sheet data from various German companies between 1994 and 1998, finding that the first- and second-digit distributions generally conformed to Benford's Law, both in annual analyses and when aggregated over the entire study period. Tam Cho and Gaines (2007) examined financial transactions in the context of election campaign financing, identifying subsets of data warranting closer scrutiny. Rauch, Götsche, Brähler, and Engel (2011) investigated macroeconomic data reported to Eurostat by EU member states and found that Greece exhibited the greatest deviation from Benford's Law among eurozone countries. Alali and Romero (2013) analyzed more than a decade of financial statements from U.S. public companies and found significant deviations in the distribution of current assets like equipment, property, and accounts receivable.

In some studies, the methodological strengths of Benford's Law have been emphasised. For instance, Günnel and Tödter (2009) characterized it as a simple, objective, and effective approach for identifying anomalies within large datasets. Nonetheless, there remains considerable debate regarding the

best practices for its implementation. Günnel and Tödter (2009) argued that analysis should primarily focus on the first digit, whereas Ramos (2006) contended that the first three digits should be examined. According to Ramos, such an approach offers a more comprehensive “electrocardiogram” of a dataset, enabling the identification of potential fraudulent activities at a granular level.

1. Proof of Benford’s Law and basic properties of the distribution

For nearly nine decades, mathematicians and statisticians have advanced diverse explanations for the empirical regularity now known as Benford’s Law. Raimi (1976) reviewed a range of less formal hypotheses, including Goudsmit and Furry’s (1944) proposition that the phenomenon arises from “the way we write numbers” and Furlan’s (1948) suggestion that it reflects a fundamental “harmonic” property of nature. A rigorous mathematical proof was not presented until Hill (1995), who demonstrated the law’s applicability to domains such as stock market data, census figures, and certain accounting records. Hill (1995) emphasized that, similar to the normal distribution, Benford’s distribution is an empirically observable regularity. His proof rests on the principle that datasets adhering to Benford’s Law are formed by combining numbers from multiple different distributions. When distributions are randomly selected and samples are randomly drawn from each, the aggregated significant-digit frequencies converge toward Benford’s distribution, even if the individual source distributions deviate from it (Hill, 1995, 1998). This convergence arises from the combination of heterogeneous numerical sources, producing what Hesman (1999) described as a “distribution of distributions,” representing a universal manifestation of numerical randomness.

Further, Boyle (1994) contributed to the theoretical foundation of Benford’s Law by showing that datasets composed of random variables from varied sources, and subsequently multiplied, divided, or raised to integer powers, tend to follow Benford’s distribution. This insight clarifies why certain accounting datasets often conform closely to Benford’s Law, as accounting figures are frequently the product of multiplicative processes involving independent variables. For example, accounts receivable may be calculated by multiplying units sold (from one distribution) by price per unit (from another distribution), while the cost of goods sold often results from a combination of several values, each derived from different distributions.

1.1. Fundamental statistical characteristics of Benford's Law

Benford's Law (BL) describes the empirical phenomenon whereby, in many naturally occurring datasets in many naturally occurring datasets, whether derived from mathematical tables, real-world measurements, or a combination of both, the leading significant digits are not uniformly distributed, contrary to intuitive expectations. Instead, smaller digits occur with disproportionately greater frequency. Instead, smaller digits occur with disproportionately greater frequency. Specifically, BL asserts that the significant digits in many datasets follow a defined logarithmic distribution. In its most widely recognized formulation, which focuses on the first significant digit in base-10 numbers, this phenomenon is also referred to as the First-Digit Phenomenon (Benford, 1938):

$$P(D_1 = d_1) = \log_{10} \left(1 + \frac{1}{d_1} \right), \quad (1)$$

Where D_1 denotes the first significant decimal digit. For example, according to equation (1), $P(D_1 = 1)$ is approximately 0.301, while $P(D_1 = 9)$ is approximately 0.0457, indicating that the probability of the digit 1 appearing is approximately seven times greater than that of the digit 9.

In its more general form, Benford's Law describes the joint distribution of all decimal digits: for any positive integer m ,

$$P((D_1, D_2, \dots, D_m) = (d_1, d_2, \dots, d_m)) = \log_{10} \left(1 + \left(\sum_{j=1}^m 10^{m-j} d_j \right)^{-1} \right)$$

This holds for all m -tuples $(d_1, d_1, \dots, d_m)$, where $d_1 \in \{1, 2, \dots, 9\}$ and for $j \geq 2$, $d_j \in \{0, 1, \dots, 9\}$. Here, D_2, D_3, D_4 , etc., represent the second, third, fourth, and subsequent significant decimal digits.

Nevertheless, deviations from BL should not be automatically interpreted as evidence of manipulation or poor data quality. Such deviations may result from structural changes or data characteristics unrelated to fraud. However, significant departures from the expected distribution should be treated as red flags warranting further investigation.

Digit tests commonly include first-digit and first-two-digit (FTD) analyses, collectively referred to as first-order tests. The first-digit test is the most widely discussed but offers limited diagnostic precision. A poor fit to BL in the first-digit test may suggest anomalies, but it often results in impractically large audit samples. For this reason, the FTD test is preferred, as it offers

more targeted detection of potential biases while yielding smaller, more manageable samples.

Several goodness-of-fit measures are used to evaluate conformity with BL. The Z-statistic evaluates each of the 90 possible FTD combinations separately, taking into account the deviation magnitude, expected proportion, and sample size n :

$$Z = \frac{|OP - EP| - \frac{1}{2n}}{\sqrt{\frac{EP(1 - EP)}{n}}} \quad (3)$$

Where OP is the observed proportion, EP the expected proportion, and n the total sample size. The continuity correction term $1/2n$ is applied when smaller than $|OP - EP|$ to ensure a positive numerator. While effective for individual combinations, the Z-statistic is sensitive to large n , producing statistically significant results for trivial deviations.

The Chi-square (χ^2) test evaluates the overall fit of observed versus expected frequencies:

$$\chi^2 = \sum_{i=1}^K \frac{(OC - EC)^2}{EC} \quad (4)$$

where OC is the observed count, EC the expected count, and $K = 90$ for FTD analysis. The degrees of freedom equal $K - 1$.

The Mean Absolute Deviation (MAD) metric addresses large- N sensitivity by excluding N from its computation:

$$MAD = \frac{\sum_{i=1}^K |AP - EP|}{K} \quad (5)$$

Nigrini (2020) proposes interpretation thresholds for MAD values, categorizing them into close conformity, acceptable conformity, marginal conformity, and nonconformity.

In conclusion, while multiple statistical measures can assess conformity with Benford's Law, the FTD test—supplemented by Z-statistic, χ^2 , and MAD analyses—offers a balanced approach for detecting anomalies in large datasets. Although conformity tests yield valuable insights, a deeper analysis and more holistic interpretation are necessary.

2. Integration of Benford's Law into machine learning algorithms for anomaly and fraud detection

In the era of big data, the increasing complexity of financial information demands sophisticated tools for efficient anomaly detection. Benford's Law serves as a useful initial diagnostic method in auditing, and further investigation should be conducted using sophisticated auditing strategies. Recent advances in machine learning algorithms and the methodology of Benford's Law can intersect in meaningful ways, particularly in data analysis and fraud detection. When integrated with machine learning, Benford's Law exhibits potential to provide a powerful statistical feature for detecting deviations. By embedding these metrics into predictive models, these measures can potentially more effectively uncover irregularities, such as falsified financial statements, abnormal transaction patterns, irregular election data, or manipulated scientific results, etc. This intersection has attracted growing attention not only from the academic community but also from practitioners and industry, given its importance in financial fraud detection. Scientists have shown increasing interest in integrating Benford's Law with machine learning. Recent investigations have focused on integrating these approaches to enhance anomaly and fraud detection across various domains. Frameworks combining statistical analysis of digit distributions with advanced ML methods have been studied, leading to more interpretable detection systems. Detecting financial statement fraud is crucial for maintaining the integrity of financial systems and ensuring their proper functioning. However, traditional methods such as audits, internal controls, and regulatory oversight often prove to be insufficient.

2.1. Current advances in the application of Benford's Law within machine learning

Despite the existence of regulatory frameworks, financial fraud remains widespread due to the increasingly sophisticated nature of fraudulent schemes and the limitations of current detection methods (Beneish, 1999). As a result, there is a growing need for advanced, data-driven approaches to more effectively detect and prevent financial statement fraud. Traditional methods for detecting financial statement fraud primarily rely on audits and forensic accounting techniques. While these approaches are important, they are inherently limited by the scope of human judgment and expertise. Rule-based systems, dependent on preset thresholds and the detection of statistical anomalies, such as excessive revenue growth or irregular changes in financial ratios, are constrained by notable limitations. Fraud schemes often

evolve specifically to bypass these static detection mechanisms, as perpetrators adapt their methods to avoid triggering known alerts (Perols, 2011). Additionally, such systems tend to produce a high volume of false positives, overwhelming auditors and regulators and making it more difficult to isolate actual instances of fraud (Ngai et al., 2011; Hasan et al., 2024). Another significant challenge is the growing complexity of financial transactions and global supply chains, which have outpaced the capabilities of traditional detection methods.

Despite these advancements, current machine learning approaches still face notable limitations. One of the primary issues is the imbalance in fraud datasets, as fraudulent instances represent only a small fraction of all financial records. This imbalance can compromise model performance, often leading to a high rate of false negatives where fraudulent activities go undetected (Hossain et al., 2024). In addition, many machine learning models, particularly those based on deep learning, are often perceived as “black boxes,” offering little transparency or insight into how decisions are made. This lack of interpretability creates difficulties for auditors and regulators, who require clear, justifiable outcomes to support their assessments (Molnar, 2019).

Hybrid and unsupervised learning models may provide new ways to detect fraud patterns that traditional models might miss, offering valuable insights for academic researchers and professionals in auditing, compliance, and corporate governance. Song et al. (2014) indicated that statistical and machine learning techniques, such as logistic regression, artificial neural networks, and support vector machines, utilize historical data to identify anomalies, thereby enhancing the accuracy and reliability of fraud detection systems. In an effort to improve fraud detection accuracy, Le, L., & Mantelaers, E. (2024) employ learning methods, in particular K-means, to establish a framework for auditors to apply once anomalies are detected using Benford's Law. Graph representation learning has gained attention for its ability to capture complex relationships in financial data. Li's study (Li, 2023) on the integration of multiple relational graphs for fraud detection highlights the effectiveness of these approaches in capturing the complex interconnections within financial data, thereby improving the accuracy and reliability of fraud detection systems.

3. Conceptual frameworks for integrating Benford's Law within machine learning

Financial statement fraud has long represented a significant challenge for businesses, regulators, and investors due to its detrimental effects on market

confidence and economic stability. Over the years, a variety of methods have been developed to detect fraudulent financial reporting, ranging from traditional statistical approaches to modern machine learning techniques. This section reviews these methodologies, traces their evolution, and identifies gaps in the current research, particularly in the application of machine learning and artificial intelligence.

As a gap in traditional approaches, we underline the following fact. While traditional methods provide useful insights, they are limited by their reliance on fixed rules and thresholds, which can be bypassed by increasingly sophisticated fraud tactics. Additionally, many of these approaches generate high false-positive rates, overwhelming auditors with cases that may not involve actual fraud. A major drawback is their inability to learn and adapt over time—an essential capability in an environment where fraudulent strategies are constantly evolving (Ngai et al., 2011). As a result, detecting financial statement fraud now requires the adoption of more advanced, data-driven approaches, such as machine learning and artificial intelligence (AI).

3.1. Towards the incorporation of Benford's Law into machine learning frameworks

Researchers often draw on the Fraud Triangle Theory, introduced by Donald Cressey in 1953, to understand the core factors that contribute to financial fraud. The theory suggests that three key elements must be present for fraud to occur: (1) pressure, often arising from financial hardship or unrealistic performance targets; (2) opportunity, usually enabled by weak internal controls; and (3) rationalization, where the individual justifies their fraudulent actions as necessary or harmless (Dorminey et al., 2012). This framework remains widely used in fraud detection research, as it provides a theoretical basis for understanding the motivations and conditions that drive fraudulent behaviour.

In addition to the Fraud Triangle, researchers have identified specific financial statement indicators that may signal potential fraud. These include sudden shifts in financial ratios, discrepancies between reported earnings and cash flows, and unusual patterns in accruals (Beneish, 1999). Machine learning models can be trained to detect such red flags by analyzing large volumes of financial data, offering greater accuracy and sensitivity in fraud detection compared to traditional methods.

While machine learning holds great promise for improving fraud detection, addressing the challenges of class imbalance, data availability, model interpretability, and overfitting is essential for developing robust and trustworthy systems. Ongoing research plays a critical role in refining these models, ensu-

ring they remain effective in real-world applications while maintaining transparency and trustworthiness (Pan, 2024; Alarfaj et al., 2022). This section explores key machine learning techniques applied in fraud detection, including logistic regression, decision trees, random forests, support vector machines (SVM), and deep learning models.

Logistic regression is a foundational machine learning technique widely used in financial fraud detection due to its simplicity and ease of interpretation. This statistical model estimates the probability of a binary outcome, fraud or no fraud, based on input variables such as financial ratios and corporate governance indicators. Studies by Dechow et al. (2011) and Perols (2011) have demonstrated the effectiveness of logistic regression as a baseline method for detecting earnings manipulation and financial misstatements. However, despite its widespread use, logistic regression assumes a linear relationship between predictors and the outcome, which may not adequately capture the complexity of sophisticated fraud schemes.

Decision trees are a widely adopted machine learning technique in fraud detection, known for their ability to classify data by recursively splitting it into subsets based on the most informative features, forming a hierarchical, tree-like structure. Their high level of interpretability makes them especially valuable to auditors and regulators who need clear explanations for model decisions. Kirkos et al. (2007) applied decision trees to detect financial statement fraud among Greek firms, achieving promising results. However, decision trees can be prone to overfitting, particularly when dealing with complex datasets containing many variables, which can limit their ability to generalize to new data (Ngai et al., 2011).

Random forests address the overfitting challenges commonly associated with decision trees by constructing multiple trees from different subsets of the data and averaging their predictions. This ensemble approach enhances model robustness and has proven effective in a variety of applications, including fraud detection. In a comparative study of machine learning algorithms, Perols (2011) found that random forests outperformed other models in detecting financial statement fraud, notably achieving lower false-positive rates. Additionally, random forests handle high-dimensional data well, making them suitable for large datasets with many features. However, like decision trees, they can be difficult to interpret, particularly when a large number of trees are used (Breiman, 2001).

Support Vector Machines (SVMs) are a widely used machine learning algorithm in the field of fraud detection. They work by identifying the optimal hyperplane that separates fraudulent from non-fraudulent cases within a multidimensional feature space. One of the key strengths of SVMs is their ability

to handle high-dimensional data, making them well-suited for analyzing complex financial datasets. Cecchini et al. (2010) demonstrated the effectiveness of SVMs in detecting financial misstatements, particularly when combined with feature selection methods. However, SVMs can be computationally intensive and generally offer less interpretability than simpler models such as decision trees or logistic regression (Phua et al., 2010).

In recent years, neural networks, particularly deep learning models, have attracted growing interest for their potential in detecting financial statement fraud. These models consist of multiple layers of interconnected neurons, enabling them to capture complex, non-linear relationships between input variables and outcomes. Their ability to detect subtle patterns in large datasets makes them especially effective at identifying sophisticated fraud schemes (Zhou & Kapoor, 2011). However, deep learning models are often criticized for their lack of transparency, frequently being referred to as “black boxes,” which presents a major limitation for auditors and regulators who require a clear understanding of the model’s decision-making process (Molnar, 2019). The use of deep learning in fraud detection is still relatively new, with limited research fully exploring its capabilities. Some studies have applied convolution neural networks (CNNs) and recurrent neural networks (RNNs) to time-series financial data, achieving promising results. These models are capable of recognizing temporal patterns, allowing for the detection of fraudulent behavior that evolves over time. Nevertheless, challenges such as limited interpretability and the need for large, high-quality labeled datasets remain significant barriers to the widespread adoption of deep learning in this domain.

There is growing interest in developing hybrid models that combine multiple machine learning techniques to improve accuracy and reduce false positives. For example, Chen, Li, and Wu (2020) proposed a hybrid approach that merges logistic regression with decision trees and random forests. This model leverages the interpretability of logistic regression along with the strong predictive performance of ensemble methods, resulting in a more robust solution for detecting financial fraud. Additionally, hybrid frameworks that incorporate both supervised and unsupervised learning have been explored to detect not only known fraud patterns but also previously unidentified anomalies (Ngai et al., 2011). Toosi, F. G. (2025) investigates the Degree of Benford’s Law Existence (DBLE) in two language model methodologies and finds that model weights, regardless of optimizer, conform to Benford’s Law. Monkam, G. F., Yan, J., & Bastian, N. D. (2025) proposed a forensic-based framework to detect poisoning in machine learning models, showcasing enhanced detection performance and providing a robust, interpretable method for safeguarding models in real-world settings.

3.2. Advances in the state-of-the-art integration of Benford's Law with machine learning

In general, the object of analysis is the dataset which incorporates a combination of key financial and non-financial variables relevant to predicting financial statement fraud. Financial indicators included revenue growth rate, profit margin, current ratio, and total debt-to-equity ratio—all commonly linked to potential signs of fraudulent reporting (Batista et al., 2004). In addition, non-financial features such as audit fees are included, as they provide valuable insight into corporate governance practices that may influence the credibility of financial disclosures (Padmaja et al., 2007). The selection of these variables was informed by prior research that has consistently identified them as reliable indicators of accounting manipulation (Setiawan, 2023).

The performance of the model is usually assessed using a range of evaluation metrics, including accuracy, precision, recall, F1-score, and the AUC-ROC curve (Joudaki et al., 2014). Accuracy reflects the overall proportion of correct predictions, while precision and recall measure the model's effectiveness in correctly identifying fraudulent cases. The F1-score, which balances precision and recall, is particularly useful in scenarios with class imbalance—a common challenge in fraud detection (Polvimoltham & Sinapiromsaran, 2021). The AUC-ROC curve evaluates the model's ability to distinguish between fraudulent and non-fraudulent instances, with a higher area under the curve indicating better discriminatory performance (Phua et al., 2004).

Given the inherent class imbalance in fraud detection datasets, techniques such as SMOTE (Synthetic Minority Over-sampling Technique) are usually applied to increase the number of fraudulent instances by generating synthetic samples. Additionally, undersampling of the majority class is considered a complementary approach to mitigate imbalance, though this carries the risk of discarding informative data (Leevy et al., 2018). To ensure reliable model evaluation, the dataset is split into training and testing sets, usually using an 80/20 ratio, and K-fold cross-validation is implemented to enhance the model's ability to generalize to unseen data (Pozzolo et al., 2014).

An important step in analysis is comparison with existing methods. Conventional approaches, such as rule-based systems and statistical techniques like Benford's Law and Z-score analysis, are employed as benchmark models for comparison. While these methods are well-established in auditing and fraud detection practices, they often fall short in terms of scalability and adaptability to evolving fraud patterns.

An important area for future research is to extend the forensic framework combined with Benford's Law to larger-scale and real-world datasets, including those from autonomous systems. Additionally, the integration of advanced explainable AI techniques will be a key focus to improve interpretability.

Conclusion

The aforementioned studies reveal several key insights. First, machine learning models, especially ensemble methods, outperform traditional rule-based approaches in accuracy, recall, and AUC-ROC. Second, non-financial indicators such as corporate governance factors often play a more crucial role in detecting fraud than conventional financial metrics. Third, addressing class imbalance through techniques like SMOTE and undersampling is vital for improving recall and ensuring fraud cases are not missed. Lastly, while complex models like random forests excel in predictive accuracy, they pose interpretability challenges; however, tools like feature importance analysis and SHAP help provide transparency into their decision-making processes. The results of this type of study offer several valuable practical implications for auditors, regulators, financial institutions and forensic accountants. Machine learning models can be integrated into existing audit systems to identify high-risk financial statements, allowing auditors to prioritize and concentrate their investigative efforts more efficiently. By relying on model-generated risk scores, regulatory agencies can utilize machine learning models to detect potentially fraudulent financial statements earlier in the reporting cycle. Financial institutions can integrate fraud detection models into due diligence processes to assess risk, support informed decisions, and reduce exposure to financial loss. Forensic accountants can use machine learning to faster identify suspicious patterns in financial statements.

References

- Alali, F., & Romero, S. (2013). Benford's law: Analyzing a decade of financial data. *Journal of Emerging Technologies in Accounting*, 10(1), 1–39. <https://doi.org/10.2308/jeta-50450>
- Alarfaj, F. K., Malik, I., Khan, H. U., Almusallam, N., Ramzan, M., & Ahmed, M. (2022). Credit card fraud detection using state-of-the-art machine learning and deep learning algorithms. *IEEE Access*, 10, 39700–39715. <https://doi.org/10.1109/ACCESS.2022.3167668>
- Asllani, A., & Naco, M. (2014). Using Benford's law for fraud detection in accounting practices. *International Journal of Economics and Business Research*, 7(2), 165–176. <https://doi.org/10.1504/IJEBR.2014.060248>

- Batista, G. E., Prati, R. C., & Monard, M. C. (2004). A study of the behavior of several methods for balancing machine learning training data. *ACM SIGKDD Explorations Newsletter*, 6(1), 20-29. <https://doi.org/10.1145/1007730.1007735>
- Beneish, M. D. (1999). The detection of earnings manipulation. *Financial Analysts Journal*, 55(5), 24–36. <https://doi.org/10.2469/faj.v55.n5.2296>
- Benford, F. (1938). The law of anomalous numbers. *Proceedings of the American Philosophical Society*, 78(4), 551–572. <https://www.jstor.org/stable/984802>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5-32. <https://doi.org/10.1023/A:1010933404324>
- Boyle, J. (1994). An application of Fourier series to the most significant digit problem. *The American Mathematical Monthly*, 101(9), 879–886. <https://doi.org/10.2307/2975123>
- Cecchini, M., Aytug, H., Koehler, G. J., & Pathak, P. (2010). Detecting management fraud in public companies. *Management Science*, 56(7), 1146-1160. <https://doi.org/10.1287/mnsc.1100.1157>
- Chen, Y., Li, H., & Wu, D. (2020). A hybrid fraud detection model based on logistic regression and machine learning techniques. *Expert Systems with Applications*, 139, 112837. <https://doi.org/10.1016/j.eswa.2019.112837>
- Dal Pozzolo, A., Caelen, O., Le Borgne, Y. A., Waterschoot, S., & Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert systems with applications*, 41(10), 4915-4928.
- Dechow, P., Ge, W., & Schrand, C. (2010). Understanding earnings quality: A review of the proxies, their determinants, and their consequences. *Journal of Accounting and Economics*, 50(2-3), 344-401. <https://doi.org/10.1016/j.jacceco.2010.09.001>
- Durtschi, C., Hillison, W., & Pacini, C. (2004). The effective use of Benford's law to assist in detecting fraud in accounting data. *Journal of Forensic Accounting*, 5(1), 17–34.
- Dorminey, J., Fleming, A. S., Kranacher, M. J., & Riley Jr, R. A. (2012). The evolution of fraud theory. *Issues in Accounting Education*, 27(2), 555-579. <https://doi.org/10.2308/iace-50131>
- Furlan, S. (1948). *La legge matematica dei numeri*. Trieste, Italy: Edizioni Lint.
- Goudsmit, S. A., & Furry, W. H. (1944). Significant figures of numbers in statistical tables. *Nature*, 154, 800–801. <https://doi.org/10.1038/154800b0>
- Günnel, S., & Tödter, K. H. (2009). Does Benford's law hold in economic research and forecasting? *Empirical Economics*, 36(3), 627–650. <https://doi.org/10.1007/s00181-008-0212-0>
- Hasan, L., Hossain, M. Z., Johora, F. T., & Hasan, M. H. (2024). Cybersecurity in accounting: Protecting financial data in the digital age. *European Journal of Applied Science, Engineering and Technology*, 2(6), 64-80. [https://doi.org/10.59324/ejaset.2024.2\(6\).06](https://doi.org/10.59324/ejaset.2024.2(6).06)

- Hesman, T. (1999, August 16). There's something about Benford's law. *The Washington Post*, p. A03.
- Hill, T. P. (1995). A statistical derivation of the significant-digit law. *Statistical Science*, 10(4), 354–363. <https://doi.org/10.1214/ss/1177009869>
- Hill, T. P. (1998). The first digit phenomenon. *American Scientist*, 86(4), 358–363. <https://doi.org/10.1511/1998.4.358>
- Hossain, M. Z., Johora, F. T., Raja, M. R., & Hasan, L. (2024). Transformative impact of artificial intelligence and blockchain on the accounting profession. *European Journal of Theoretical and Applied Sciences*, 2(6), 144-159.
- Joudaki, H., Rashidian, A., Minaei-Bidgoli, B., Mahmoodi, M., Geraili, B., Nasiri, M., & Arab, M. (2014). Using data mining to detect healthcare fraud and abuse: A review of literature. *Global Journal of Health Science*, 7(1), 194. <https://doi.org/10.5539/gjhs.v7n1p194>
- Kirkos, E., Spathis, C., & Manolopoulos, Y. (2007). Data mining techniques for the detection of fraudulent financial statements. *Expert Systems with Applications*, 32(4), 995-1003. <https://doi.org/10.1016/j.eswa.2006.01.018>
- Le, L., & Mantelaers, E. (2024). Benford's Law and Beyond: A framework for auditors. *Maandblad voor Accountancy en Bedrijfseconomie*, 98(7), 427-438.
- Leevy, J. L., Khoshgoftaar, T. M., Bauder, R. A., & Seliya, N. (2018). A survey on addressing high-class imbalance in big data. *Journal of Big Data*, 5(1), 1-30. <https://doi.org/10.1186/s40537-018-0151-6>
- Molnar, C. (2019). Interpretable machine learning: A guide for making black box models explainable. Retrieved from <https://christophm.github.io/interpretable-ml-book/>
- Monkam, G. F., Yan, J., & Bastian, N. D. (2025). A Forensic Analysis Framework for Machine Learning Model Poisoning Detection. *Security and Privacy*, 8(5), e70079.
- Newcomb, S. (1881). Note on the frequency of use of the different digits in natural numbers. *American Journal of Mathematics*, 4(1), 39–40. <https://doi.org/10.2307/2369148>
- Ngai, E. W., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559-569. <https://doi.org/10.1016/j.dss.2010.08.006>
- Nigrini, M. J. (1992). The detection of income tax evasion through an analysis of digital frequencies. *Journal of the American Taxation Association*, 14(1), 52–71.
- Nigrini, M. J. (1994). Using digital frequencies to detect fraud. *The White Paper*, 8(2), 3–6.
- Nigrini, M. J., & Mittermaier, L. J. (1997). The use of Benford's law as an aid in analytical procedures. *Auditing: A Journal of Practice & Theory*, 16(2), 52–67. <https://doi.org/10.2308/aud.1997.16.2.52>

Nigrini, M. J. (2020). *Forensic analytics: Methods and techniques for forensic accounting investigations*. Hoboken, NJ: John Wiley & Sons.

Padmaja, T. M., Dhulipalla, N., Bapi, R. S., & Krishna, P. R. (2007, December). Un-balanced data classification using extreme outlier elimination and sampling techniques for fraud detection. In *15th International Conference on Advanced Computing and Communications (ADCOM 2007)* (pp. 511-516). IEEE. <https://doi.org/10.1109/ADCOM.2007.22>

Pan, E. (2024). Machine learning in financial transaction fraud detection and prevention. *Transactions on Economics, Business and Management Research*, 5, 243-249.

Phua, C., Alahakoon, D., & Lee, V. (2004). Minority report in fraud detection: Classification of skewed data. *ACM SIGKDD Explorations Newsletter*, 6(1), 50-59. <https://doi.org/10.1145/1007730.1007738>

Perols, J. (2011). Financial statement fraud detection: An analysis of statistical and machine learning algorithms. *Auditing: A Journal of Practice & Theory*, 30(2), 19-50. <https://doi.org/10.2308/ajpt-50009>

Polvimoltham, P., & Sinapiromsaran, K. (2021). Mass ratio variance majority under-sampling and minority oversampling technique for class imbalance. In *Fuzzy Systems and Data Mining VII* (pp. 152-161). IOS Press. <https://doi.org/10.3233/FAIA210198>

Quick, R., & Wolz, M. (2003). Benford's law in German balance sheets. *Schmalenbach Business Review*, 55(4), 285-303. <https://doi.org/10.1007/BF03396629>

Raimi, R. A. (1976). The first digit problem. *American Mathematical Monthly*, 83(7), 521-538. <https://doi.org/10.1080/00029890.1976.11994042>

Ramos, R. A. (2006). Applying Benford's law to detect fraudulent practices in accounting data. *Revista de Contabilidad y Auditoría*, 23, 45-59.

Rauch, B., Götsche, M., Brähler, G., & Engel, S. (2011). Fact and fiction in EU-governmental economic data. *German Economic Review*, 12(3), 243-255. <https://doi.org/10.1111/j.1468-0475.2010.00517.x>

Setiawan, R., Tjahjono, B., Firmansyah, G., & Akbar, H. (2023). Fraud detection in credit card transactions using HDBSCAN, UMAP and SMOTE methods. *International Journal of Science, Technology & Management*, 4(5), 1333-1339.

Song, X. P., Hu, Z. H., Du, J. G., & Sheng, Z. H. (2014). Application of machine learning methods to risk assessment of financial statement fraud: Evidence from China. *Journal of Forecasting*, 33(8), 611-626. <https://doi.org/10.1002/for.2298>

Tam Cho, W. K., & Gaines, B. J. (2007). Breaking the (Benford) law: Statistical fraud detection in campaign finance. *The American Statistician*, 61(3), 218-223. <https://doi.org/10.1198/000313007X223496>

Toosi, F. G. (2025). Benford's Law in Basic RNN and Long Short-Term Memory and Their Associations. *Applied AI Letters*, 6(3), e70002.

Varian, H. R. (1972). Benford's law. *The American Statistician*, 26(3), 65. <https://doi.org/10.1080/00031305.1972.10478954>