Challenges and Innovations in Healthcare Fraud and Waste Detection Systems: A Systematic Review and Proposed Framework

Gelson André Schneider[1], **André Luis Korzenowski**[2]

[1]Federal Institute of Education of Tocantins – (IFTO)
Av. Bernardo Sayão, Lote 29B, Setor Santa Maria, Chácara Raio de Sol
77.760-000 – Colinas do Tocantins – TO – Brazil

[2]University of Vale do Rio dos Sinos - (UNISINOS) - PPG Production & System Eng
Av. Unisinos, 950. Bairro Cristo Rei 93022-750 – São Leopoldo – RS – Brazil

gelson.schneider@ifto.edu.br, akorzenowski@unisinos.br

Abstract. *Fraud and abuse in the healthcare sector cause considerable financial losses to public and private entities.This article conducted a systematic literature review to identify the challenges in fraud and waste detection systems. Ten articles were identified in the review that covered the IEEE, PubMed, Scopus, and Web of Science bases. Based on these studies, we developed a framework in which we categorized the works by anomaly point, contextual, and collective. We also identified the unsupervised learning techniques used and identified the following challenges dimensionality reduction, conceptual deviations caused by changes in fraudsters' behavior to bypass audit systems, and the need for real-time detection support.*

1. Introduction

Health support systems, such as hospitals, clinics, tests, and health professionals, can be privately contracted, insured, or fully funded by the state. As a result, the health sector receives large amounts of public and private investment, which makes it prone to fraud. The main global health frauds are listed in a 2017 OECD report [Gautam 2017]. In addition, healthcare systems are bureaucratic and complex, with many rules, regulations, and actors involved, making it difficult to uncover fraud in the different hierarchies of the system [Chen and Gangopadhyay 2013].

Fraud and abuse in the healthcare sector occur in a variety of ways, such as producing false invoices for services that were never provided to the client, submitting duplicate claims for the same individual service, misrepresenting the services provided to patients, submitting reimbursement claims for more than the services provided (Upcoding) and even submitting an invoice for services not provided.[Bhaskar et al. 2021, Abdallah et al. 2016].

In addition, types of fraud are becoming increasingly sophisticated, and fraudulent and non-fraudulent behaviors are rapidly becoming obsolete, with fraudsters developing smarter and more cautious approaches to avoiding audits [Shin et al. 2012]. Fraud detection can be seen as an outlier or anomaly detection problem [Ekin et al. 2018], anomaly detection refers to the problem of finding patterns in the data that do not conform to the expected behavior, these anomalies being categorized as point anomaly when an individual snapshot can be considered anomalous about the rest of the data, contextual anomaly is if a data instance is anomalous in a specific context, but not in another, it is called a contextual anomaly, the notion of context is induced by the structure of the dataset and must be specified as part of the problem formulation and collective anomaly is if a set of related data instances is anomalous about the rest of the data set [Chandola et al. 2009]. Thus, anomaly detection emerges as a key approach for developing effective fraud detection systems.

For a long time, manual fraud auditing techniques, such as discovery sampling, have been used to detect fraud [Tennyson and Salsas-Forn 2002]. However, to increase the effectiveness and speed of detection, over the years, computerized and automated fraud detection systems have been developed and improved, integrating a wide range of data mining methods involving statistical, mathematical, artificial intelligence, and machine learning techniques [Abdallah et al. 2016]. Two methods are commonly used, supervised methods, which aim to classify whether a given situation is a fraud or not according to labeled data. However, most of the data are not labeled and still have a large class imbalance, making unsupervised methods a potential solution for anomaly detection.

Health abuse fraud detection systems process large volumes of data to identify and alert about fraud and waste in healthcare. However, these systems are prone to failure due to the inherent complexity of the problem, which leads to low accuracy and frequent false alarms [Kemp et al. 2022].

- Large volumes of data often require the application of dimensionality reduction techniques to perform data mining tasks effectively. For example, distance-based anomaly detection methods, which are independent of the geometry of the space or the heterogeneity of the data, do not work well with high dimensionality [Houle 2013], as well as feature engineering techniques which is an important factor because it significantly affects the performance of the algorithms, and most of the literature on health fraud does not mention which data pre-processing techniques were used [Kumaraswamy et al. 2022];
- Concept drift, where the meaning of the data changes over time;
- Class imbalance is considered one of the most critical problems faced by fraud detection systems, where fraudulent or wasteful claims are expected to appear at a much lower rate than acceptable claims, and
- Support for online or real-time fraud detection that must be able to cope with limited resources (time and memory) to ensure that the detection process works efficiently [Abdallah et al. 2016].

To understand how these challenges are currently addressed in the literature, we carried out an systematic review and presented its methodology in the following section.

2. Methodology

Considering the objectives of this research, a systematic literature review methodology was adopted and this section presents how the SBR was conducted. The systematic review was carried out using the PRISMA method [Moher et al. 2010].

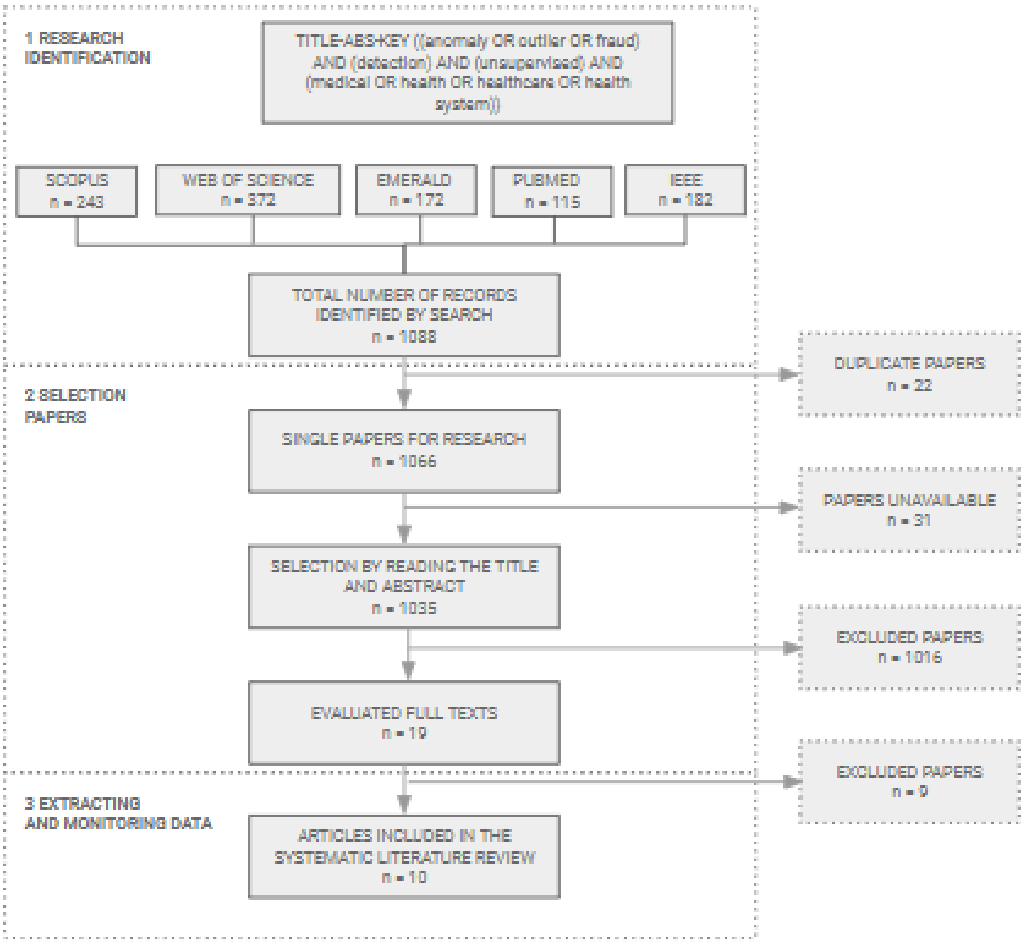

Figura 1. Prisma method flowchart

The systematic review was conducted on the CAPES journal platform using the Emerald Insight, IEEE, Scopus, and Web of science databases, as well as the PubMed database. The search String was used: ((anomaly OR outlier OR fraud) AND (detection) AND (unsupervised) AND (medical OR health OR helthcare OR health system)) for the abstract, title, and keywords fields.

The next step is to define the list of eligible papers, which filters will be adopted with inclusion and exclusion criteria, which have the aim of eliminating papers that do not meet the research objectives. The first filter to be adopted is the following inclusion criteria (IC): IC1: complete works and IC2: works published in English, applying this first filter we obtained the following number of documents per database: Emerald Insight 172, IEEE 182, Scopus 243, Wef of Science 372 and PubMed 115 for a total of 1088 documents. The second filter was made up of the following exclusion criteria (EC): EC1: duplicate papers; EC2: papers not available to download in full; EC3: reading of the title and abstract and evaluation of the subject, according to the objective and research questions defined in this systematic review and finally reading of the full text and its adherence to the systematic review, we obtained the following numbers Emerald Insight 0, IEEE 3, Scopus 3, Wef of Science 2 and PubMed 2 in a total of 10 documents. Figure 1 shows a detailed flowchart of the PRISMA method.

Despite the number of papers returned by the search string, only ten papers were eligible for the review. A large number of papers were eliminated because they dealt with the use of unsupervised methods to detect anomalies in the production process, equipment, and machines, as many abstracts contained the term “equipment health”or “machine health”. The selected documents are thoroughly analyzed, and their insights are discussed in the following section.

3. Results

In this section we present the compilation of our systematic review, in Table 1 we present the studies, identifying the types of anomaly (point, contextual or collective), the technique applied and the challenge addressed.

Tabela 1. Type of Anomaly X Challenge

Paper	Type of Anomaly	Technique	Challenge
[Bauder et al. 2018]	Point	**LOF**, IF, URF, AE, KNN	-
[Bhaskar et al. 2021]	Point	**IF**, K means, LOF	-
[Nagata et al. 2021]	Point	**OCSVM**, IF, LOF, RC	-
[Putina et al. 2020]	Point	**RHF**, IF, LOF, HBOS, OCSVM, KNN	-
[Dik et al. 2018]	Point	**RDUFL**, FCM, NC	-
[Davidow and Matteson 2022]	Point	**FAMDAD**, IF, SPAD	Dimensionality reduction
[Massi et al. 2020]	Contextual	**K-means**	Concept drift
[Settipalli and Gangadharan 2021]	Collective	**PSPGA**	Concept drift
[Kemp et al. 2022]	Collective	**Association rules**	Concept drift
[Dos Santos et al. 2018]	Collective	**DDC**, OCSVM, IF, LOF, GM, RC	-

We present in Figure 2 our proposal for a framework based on the systematic review insights to be used as a guide in the development of solutions for health abuse fraud detection systems. The framework is developed based on the type of anomaly and which challenge can be addressed; in cases where the challenges are not mentioned, the main techniques applied are shown. Below, we explain each type of detection in more detail.

3.1. Point anomalies detection

An anomaly is considered a point when an individual snapshot can be considered anomalous compared to the rest of the data [Chandola et al. 2009]. According to the author, this anomaly is the simplest to detect and has been the focus of most research into anomaly detection, as can be seen in our current systematic review, in which six of the ten documents analyzed fall into this category, as can be seen in the Table 1.

However, as useful as these models may be if they are the only ones used, they will not reduce the overall scale of healthcare fraud. If these models are the only ones used in the fight against healthcare fraud, it is unlikely that providers involved in fraud who take a more cautious approach will be identified [Musal 2010]. On the other hand, the chance of finding false positives using this approach is very low. The field of medicine is inherently heterogeneous in terms of both health conditions and their presentation and treatment. Increasing volumes of data and rapidly changing patterns bring challenges that require innovative solutions [Kemp 2023].

With the large volume of data, databases often contain high dimensions, i.e. an instance with a large number of attributes, but anomaly detection methods based on

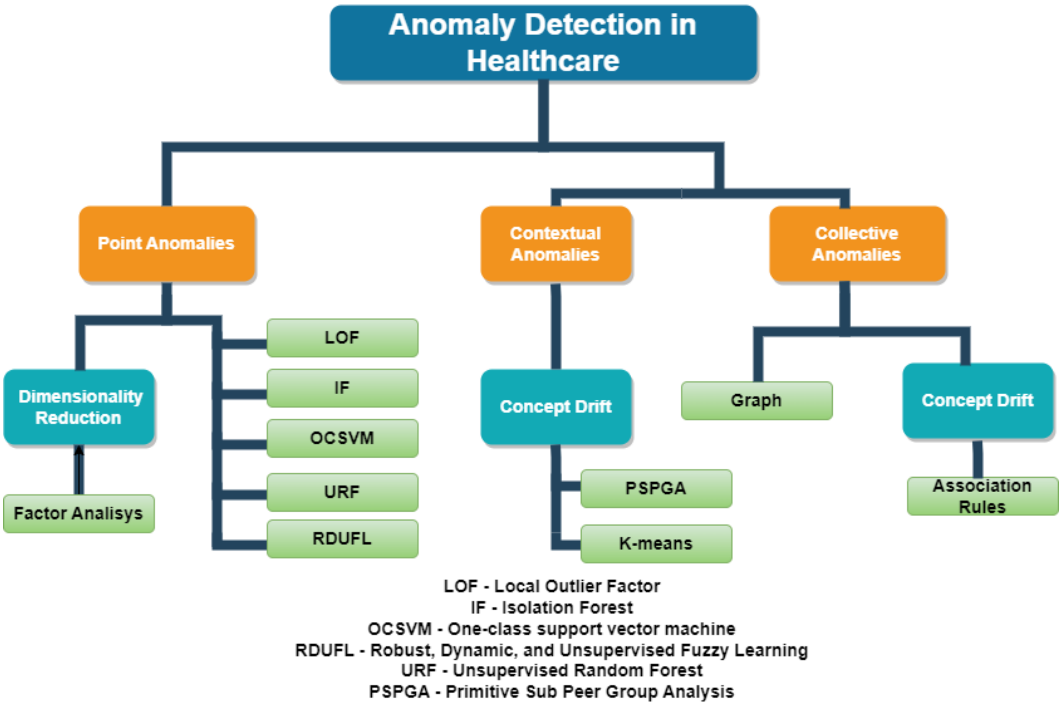

Figura 2. Framework proposal

Distance/Proximity/Nearest-Neighbor as K-NN, Kth-NN, Local Outlier Factor (LOF) and K-means which are commonly used in point anomaly detection are agnostic about the geometry of the space or the heterogeneity of the data, so they don't work very well on these data sets [Houle 2013].

Dimensionality reduction is a challenge addressed in this work, in our systematic review we found only one document that brings this discussion to light, [Davidow and Matteson 2022] the document proposes a model for detecting point anomalies but uses a method based on Factor Analysis to reduce dimensionality in mixed data, i.e. continuous, discrete and categorical attributes. The method creates a continuous, reduced-dimensional embedding that can be combined with detection algorithms for a low-dimensional continuous domain such as the Isolation Forest (IF), Local Outlier Factor, and DBSCAN.

3.2. Contextual anomaly detection

To mitigate the shortcomings of anomaly detection through point detection, we can resort to contextual anomaly detection. If a data instance is anomalous in a specific context, but not in another, it is called a contextual anomaly, the notion of context is induced by the structure of the dataset and must be specified as part of the problem formulation [Chandola et al. 2009]. Each data instance must have contextual attributes that are used to determine the context (or neighborhood) of that instance, for example in time series data, time is a contextual attribute that determines the position of an instance in the entire sequence. And behavioral attributes that define the non-contextual characteristics of an instance.

In this context, we classified two documents in our systematic review as contextual anomaly detection approaches. In [Massi et al. 2020] the researchers carry out data mining to detect fraud/anomalies relating to diagnosis related group (DGR) Upcoding. For this, an unsupervised K-means algorithm was used to identify and group locally consistent and locally similar hospitals according to the characteristics and behavior in the treatment of a specific disease, for the authors the K-means algorithm is more robust to variable cluster densities and high data dimensionality, and the simplest to be included in a broader algorithm that automatically selects the best subsets of resources and the best number of clusters. Finally, the proposed model was found to be promising in detecting anomalous DRG coding behavior and is easily transferable to all diseases and contexts of interest.

Already using the medicare dataset, the study carried out by [Settipalli and Gangadharan 2021] demonstrates the unsupervised primitive sub-peer group analysis (PSPGA) technique based on primitive group analysis (PGA) [Bolton et al. 2001], PGA performs a local pattern analysis, the pattern of similarity in the behavior of the objects can only be drawn when it is compared with the behavior of the objects in its peer group, rather than with that of the entire population. The authors' idea in developing PSPGA is to group providers with a similar profile using PGA and then form subgroups by analyzing the similarity in the claims of each primitive peer group. In this way it is possible to distinguish a conceptual deviation from an anomalous deviation, this distinction significantly reduces false positives, and the proposal showed better results than the Kmeans, NaiveBayes, SVM, Decision tree, and KNN techniques when compared through sensitivity analyses. In future work, the authors suggest applying PGA to an online fraud detection system.

It is well known that one of the possible problems in the health system is that systematic changes occur very frequently. In this sense, contextual anomaly detection can be applied to the challenge of detecting concept changes, since this anomaly may or may not be fraud. For example, in a context where concept changes are treated as non-fraud and sudden changes as fraud, it is important to understand the distinction between them to greatly reduce false alarms in the detection process [Settipalli and Gangadharan 2021].

Although contextual analyses naturally have a higher false positive rate, they can identify frauds that are not well known [Musal 2010]. Thus, by detecting contextual anomalies, it is possible to detect changes in concept and sudden changes, resulting in the discovery of fraud that attempts to pass for a service within the standards. However, contextual models are likely to show more false positives than methods based on finding extreme values (point), which increases costs for the health system [Massi et al. 2020]. Despite advances, it is very difficult to quantify the number of fraudulent cases that go undetected. In this sense, the two aforementioned studies seek to present solutions for identifying the challenge of concept drift.

3.3. Collective anomaly detection

If a set of related data instances is anomalous about the rest of the data set, it is called a collective anomaly. The individual data instances in a collective anomaly may not be anomalies on their own, but their occurrence together with the set is anomalous [Chandola et al. 2009]. In this context, we have classified two documents in our systematic review as being Collective anomaly detection approaches:

i) The task of sequential pattern mining or association rule mining, which aims to find associations between data instances, is thus a technique for detecting collective anomalies. At work [Kemp et al. 2022] sequential pattern mining is used to identify and group courses of cancer treatments using radiotherapy that may be comparable, the method seeks to find patterns within treatments, calculate the costs of possible additional claims or unusual Upcoding, and rank providers based on possible recoverable costs. According to the authors, the method was able to identify anomalous claims as well as the patterns in which they were anomalous. This work addressed the challenge of concept drift.

ii) To analyze drug prescriptions for automatic detection of atypical prescriptions, i.e. normal and abnormal doses of each drug regarding dosage and frequency based on prescription history data., [Dos Santos et al. 2018] developed an unsupervised method called Density-Distance-Centrality (DDC). The method is based on graph centrality and is calculated using the PageRank algorithm so that higher scores are more frequent or normal prescriptions and lower scores are anomalies, i.e. possible errors in the prescribing dose and/or frequency. Overdoses or underdoses are probably prescriptions whose centrality score is below an average centrality index for each drug. According to the authors, the method was able to quickly create a distribution pattern to detect anomalous prescriptions.

We must understand that a point anomaly detection or collective anomaly detection problem can be transformed into a contextual anomaly detection problem by incorporating context information. The techniques used to detect collective anomalies are very different from those used to detect point and contextual anomalies [Chandola et al. 2009]. We can see this in the example of the contextual anomaly detection carried out by [Massi et al. 2020] using the k-means clustering technique, a technique that is commonly used to detect point anomalies.

3.4. Framework proposal

The proposed framework will be validated by means of an artifact, in which a positivist paradigm will be used, with an exploratory approach, of an applied nature using the methodology Design Science Research, [Dresch et al. 2020] which offers a method that establishes and operationalizes research, with the necessary rigor, when the desired goal is an artifact or a recommendation to solve problems, evaluate what has been designed and communicate the results. Where artifacts refer to entities (tangible or intangible) that provide applicability in solving problems or improving understanding of a phenomenon.

4. Conclusions

We present a systematic literature review to identify the challenges facing fraud and waste detection systems. In addition, we sought to understand how these identified challenges are currently being addressed. Thus, we identified the following challenges: concept drift and dimensionality reduction and we related them to the anomalies, point, contextual, and collective.

Dimensionality reduction was addressed in a study on point anomalies, in which the technique can significantly improve the algorithms, but we didn't find any feature engineering and/or data pre-processing techniques, an important step for the replicability of studies, especially with real data. This is related to the type of study, as the studies normally use small, pre-processed data sets to test the algorithms.

Concept drift was found in three documents, two of which were related to the detection of contextual anomalies and one to a collective anomaly. Advances in the techniques show that contextual and collective methods tend to reduce false positives, a problem caused by the dynamic nature of health data, i.e. concept drift. On the other hand, the detection of anomalies in real time was not mentioned in any of the studies, which is still an open gap in fraud and waste detection systems. In addition, for real applications in fraud and waste detection systems, it is necessary to improve the presentation of the outputs of the anomaly detection algorithms so that they can be more interpretable by the auditors, as already mentioned by [Massi et al. 2020]. Recent advances point to innovative techniques, but tests on real databases are needed.

Referências

Abdallah, A., Maarof, M. A., and Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68:90–113.

Bauder, R., da Rosa, R., and Khoshgoftaar, T. (2018). Identifying medicare provider fraud with unsupervised machine learning. In *2018 IEEE international conference on information Reuse and integration (IRI)*, pages 285–292. IEEE.

Bhaskar, A., Pande, S., Malik, R., Khamparia, A., et al. (2021). An intelligent unsupervised technique for fraud detection in health care systems. *Intelligent Decision Technologies*, 15(1):127–139.

Bolton, R. J., Hand, D. J., et al. (2001). Unsupervised profiling methods for fraud detection. *Credit scoring and credit control VII*, pages 235–255.

Chandola, V., Banerjee, A., and Kumar, V. (2009). Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3):1–58.

Chen, S. and Gangopadhyay, A. (2013). A novel approach to uncover health care frauds through spectral analysis. In *2013 IEEE International Conference on Healthcare Informatics*, pages 499–504. IEEE.

Davidow, M. and Matteson, D. S. (2022). Factor analysis of mixed data for anomaly detection. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 15(4):480–493.

Dik, A., Jebari, K., and Ettouhami, A. (2018). An improved robust fuzzy algorithm for unsupervised learning. *Journal of Intelligent Systems*, 29(1):1028–1042.

Dos Santos, H. D., Ulbrich, A. H. D., Woloszyn, V., and Vieira, R. (2018). Ddc-outlier: preventing medication errors using unsupervised learning. *IEEE journal of biomedical and health informatics*, 23(2):874–881.

Dresch, A., Lacerda, D. P., and Junior, J. A. V. A. (2020). *Design science research: método de pesquisa para avanço da ciência e tecnologia*. Bookman Editora.

Ekin, T., Ieva, F., Ruggeri, F., and Soyer, R. (2018). Statistical medical fraud assessment: exposition to an emerging field. *International Statistical Review*, 86(3):379–402.

Gautam, A. (2017). *Tackling wasteful spending on health*. OECD.

Houle, M. E. (2013). Dimensionality, discriminability, density and distance distributions. In *2013 IEEE 13th International Conference on Data Mining Workshops*, pages 468–473. IEEE.

Kemp, J. (2023). *Unsupervised learning for anomaly detection in Australian medical payment data*. PhD thesis, UNSW Sydney.

Kemp, J., Barker, C., Good, N., and Bain, M. (2022). Sequential pattern detection for identifying courses of treatment and anomalous claim behaviour in medical insurance. In *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 3039–3046. IEEE.

Kumaraswamy, N., Markey, M. K., Ekin, T., Barner, J. C., and Rascati, K. (2022). Healthcare fraud data mining methods: A look back and look ahead. *Perspectives in health information management*, 19(1).

Massi, M. C., Ieva, F., and Lettieri, E. (2020). Data mining application to healthcare fraud detection: a two-step unsupervised clustering method for outlier detection with administrative databases. *BMC medical informatics and decision making*, 20:1–11.

Moher, D., Liberati, A., Tetzlaff, J., Altman, D. G., et al. (2010). Preferred reporting items for systematic reviews and meta-analyses: the prisma statement. *Int J Surg*, 8(5):336–341.

Musal, R. M. (2010). Two models to investigate medicare fraud within unsupervised databases. *Expert Systems with Applications*, 37(12):8628–8633.

Nagata, K., Tsuji, T., Suetsugu, K., Muraoka, K., Watanabe, H., Kanaya, A., Egashira, N., and Ieiri, I. (2021). Detection of overdose and underdose prescriptions—an unsupervised machine learning approach. *PloS one*, 16(11):e0260315.

Putina, A., Sozio, M., Rossi, D., and Navarro, J. M. (2020). Random histogram forest for unsupervised anomaly detection. In *2020 IEEE International Conference on Data Mining (ICDM)*, pages 1226–1231. IEEE.

Settipalli, L. and Gangadharan, G. (2021). Healthcare fraud detection using primitive sub peer group analysis. *Concurrency and Computation: Practice and Experience*, 33(23):e6275.

Shin, H., Park, H., Lee, J., and Jhee, W. C. (2012). A scoring model to detect abusive billing patterns in health insurance claims. *Expert Systems with Applications*, 39(8):7441–7450.

Tennyson, S. and Salsas-Forn, P. (2002). Claims auditing in automobile insurance: fraud detection and deterrence objectives. *Journal of Risk and Insurance*, 69(3):289–308.