

Detection of Accounting Anomalies in the Latent Space using Adversarial Autoencoder Neural Networks

Marco Schreyer
University of St. Gallen
St. Gallen, Switzerland
marco.schreyer@unisg.ch

Timur Sattarov
Deutsche Bundesbank
Frankfurt am Main, Germany
timur.sattarov@bundesbank.de

Christian Schulze
University of St. Gallen
St. Gallen, Switzerland
christian.schulze@unisg.ch

Bernd Reimer
PricewaterhouseCoopers GmbH WPG
Stuttgart, Germany
reimer.bernd@pwc.com

Damian Borth
University of St. Gallen
St. Gallen, Switzerland
damian.borth@unisg.ch

ABSTRACT

The detection of fraud in accounting data is a long-standing challenge in financial statement audits. Nowadays, the majority of applied techniques refer to handcrafted rules derived from known fraud scenarios. While fairly successful, these rules exhibit the drawback that they often fail to generalize beyond known fraud scenarios and fraudsters gradually find ways to circumvent them. In contrast, more advanced approaches inspired by the recent success of deep learning often lack seamless interpretability of the detected results. To overcome this challenge, we propose the application of adversarial autoencoder networks. We demonstrate that such artificial neural networks are capable of learning a semantic meaningful representation of real-world journal entries. The learned representation provides a holistic view on a given set of journal entries and significantly improves the interpretability of detected accounting anomalies. We show that such a representation combined with the networks reconstruction error can be utilized as an unsupervised and highly adaptive anomaly assessment. Experiments on two datasets and initial feedback received by forensic accountants underpinned the effectiveness of the approach.

CCS CONCEPTS

• **Computing methodologies** → **Anomaly detection**; *Dimensionality reduction and manifold learning*; *Adversarial learning*.

KEYWORDS

anomaly detection, neural networks, accounting irregularities, fraud detection, forensic data analysis, computer assisted audit

ACM Reference Format:

Marco Schreyer, Timur Sattarov, Christian Schulze, Bernd Reimer, and Damian Borth. 2019. Detection of Accounting Anomalies in the Latent Space using Adversarial Autoencoder Neural Networks. In . ACM, New York, NY, USA, 11 pages.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD-ADF '19, August 05, 2019, Anchorage, Alaska

© 2019 Association for Computing Machinery.




Figure 1: Hierarchical view of an Accounting Information System (AIS) that records distinct layer of abstractions, namely (1) the business process, (2) the accounting and (3) technical journal entry information in designated tables.

1 INTRODUCTION

The Association of Certified Fraud Examiners estimates in its "Global Study on Occupational Fraud and Abuse 2018" [1] that organizations lose 5% of their annual revenues due to fraud. The term "fraud" refers to "the abuse of one's occupation for personal enrichment through the deliberate misuse of an organization's resources or assets" [48]. A similar study, conducted by the auditors of PwC, revealed that approx. 30% of the respondents experienced losses between \$100,000 USD and \$5 million USD due to fraud [39]. The study also showed that financial statement fraud caused by far the highest median loss of the surveyed fraud schemes¹.

At the same time, organizations accelerate the digitization of business processes [32] affecting in particular Accounting Information Systems (AIS) or more generally Enterprise Resource Planning (ERP) systems. Steadily, these systems collect vast quantities of business process and accounting data at a granular level. This holds in particular for the journal entries of an organization recorded in its general ledger and sub-ledger accounts. SAP, one of the most prominent enterprise software providers, estimates that approx.

¹The ACFE study encompasses an analysis of 2,690 cases of occupational fraud investigated between January 2016 and October 2017 that occurred in 125 countries. The PwC study encompasses over 7,228 respondents that experienced economic crime in the last 24 months.

77% of the world's transaction revenue touches one of their ERP systems [42]. Figure 1 depicts a hierarchical view of an AIS recording process of journal entry information in designated database tables.

In order to conduct fraud, perpetrators need to deviate from the regular system usage or posting pattern. Such deviations are recorded by a very limited number of "anomalous" journal entries and their respective attribute values. To detect potentially fraudulent activities international audit standards require the direct assessment of such journal entries [2], [19]. Nowadays, auditors and forensic accountants apply a wide range of data analysis techniques to examine journal entries during an audit. These techniques often encompass rule-based analyses referred to as "red-flag" tests (e.g. postings late at night, multiple vendor bank account changes) as well as statistical analyses (e.g. Benford's Law, time series evaluation). Nevertheless, the detection of traces of fraud in up to several hundred million journal entries remains a labor-intensive task requiring significant time effort and resources.

Driven by the recent technological advances in artificial intelligence [28] deep neural network-based techniques (e.g. deep autoencoder neural networks) have emerged into the field of forensic accounting and financial statement audits [43]. Such, approaches often lack a seamless interpretability of the detected "anomalous" journal entries selected for a detailed audit. This is a major drawback since auditors are required to test a representative sample of journal entries in order to reduce the "sampling risk" [20] of an audit. Ultimately, the testing of an individual entry while ignoring another one must be defensible in court [17].

To overcome this challenge we propose the application of Adversarial Autoencoder Neural Networks (AAEs) [31]. We demonstrate that such adversarial architectures are capable of learning a semantic meaningful representation of journal entries. The learned representation allows for an improved interpretability of the entries underlying generative processes as well as detected "anomalous" journal entries. In summary, we present the following contributions:

- We illustrate that AAEs can be used to learn a representation of journal entries observable in real-world ERP systems that partitions the entries into semantic meaningful groups;
- We demonstrate how such a learned representation can be used by a human auditor or forensic accountant to sample journal entries for an audit in an interpretable manner;
- We show that the learned representation combined with the magnitude of an entry's reconstruction error can be interpreted as a highly adaptive anomaly assessment of journal entries.

We envision this deep learning-based methodology as an important supplement to the auditors and forensic accountants toolbox [36]. The remainder of this work is structured as follows: In Section 2 we provide an overview of the related work. Section 3 follows with a description of the adversarial autoencoder network architecture and presents the proposed methodology to detect accounting anomalies. The experimental setup and results are outlined in Section 4 and Section 5. In Section 6 the paper concludes with a summary of the current work and future directions of research.

An implementation of the proposed methodology is available at <https://github.com/GitiHubi/deepAD>.

2 RELATED WORK

The literature survey presented hereafter focuses on (1) the detection of fraudulent activities in Enterprise Resource Planning (ERP) data and (2) the detection of financial fraud using deep Autoencoder Neural Networks (AENs) [18] as well as Generative Adversarial Networks (GANs) [29].

2.1 Fraud Detection in Accounting Data

The task of detecting fraud and accounting anomalies has been studied both by practitioners [48] and academia [3]. Several references describe different fraud schemes and ways to detect unusual and "creative" accounting practices [45].

The forensic analysis of journal entries emerged with the advent of Enterprise Resource Planning (ERP) systems and the increased volume of data recorded by such systems. Bay et al. in [7] use Naive Bayes methods to identify suspicious general ledger accounts, by evaluating attributes derived from journal entries measuring any unusual general ledger account activity. Their approach is enhanced by McGlohon et al. applying link analysis to identify (sub-) groups of high-risk general ledger accounts [33]. Kahn et al. in [25] and [26] create transaction profiles of SAP ERP users. The profiles are derived from journal entry based user activity pattern recorded in two SAP R/3 ERP system in order to detect suspicious user behavior and segregation of duties violations. Similarly, Islam et al. in [21] use SAP R/3 system audit logs to detect known fraud scenarios and collusion fraud via a "red-flag" based matching of fraud scenarios. Debreceeny and Gray in [12] analyze dollar amounts of journal entries obtained from 29 US organizations. In their work, they search for violations of Benford's Law [8], anomalous digit combinations as well as an unusual temporal pattern such as end-of-year postings. More recently, Poh-Sun et al. in [44] demonstrate the generalization of the approach by applying it to journal entries obtained from 12 non-US organizations. Jans et al. in [22] use latent class clustering to conduct a uni- and multivariate clustering of SAP ERP purchase order transactions. Transactions significantly deviating from the cluster centroids are flagged as anomalous and are proposed for a detailed review by auditors. The approach is enhanced in [23] by a means of process mining to detect deviating process flows in an organization procure to pay process. Argyrou et al. in [4] evaluate self-organizing maps to identify "suspicious" journal entries of a shipping company. In their work, they calculated the Euclidean distance of a journal entry and the code-vector of a self-organizing maps best matching unit. In subsequent work, they estimate optimal sampling thresholds of journal entry attributes derived from extreme value theory [5].

Concluding from the reviewed literature, the majority of references either draw from accounting and forensic knowledge about historical fraud schemes or non deep-learning based techniques to detect financial fraud. However, driven by the recent success of deep learning techniques, which are potentially misused by fraudsters, we see a high demand for auditors to likewise enhance their examination methodologies.




Figure 2: The adversarial autoencoder architecture [31], applied to learn the journal entries characteristics and to partition the entries into semantic meaningful groups. The adversarial autoencoder network architecture imposes an arbitrary prior distribution $p(z)$ on the discrete latent code vector z , e.g., a mixture of Gaussians. With progressing training the encoder learns an aggregated posterior distribution $q_\theta(z|x)$ that matches the imposed prior to fool the discriminator network d_ϕ .

2.2 Anomaly Detection using Deep Learning

Nowadays, deep learning inspired methods are increasingly used for novelty and anomaly detection in financial data [10, 37].

Renström and Holmsten in [40] evaluate AENs to detect fraud in credit card transactions. Similarly, Kazemi and Zarrabi [24] and Sweers et al. [46] train and evaluate a variety of variational AEN architectures. Pumsirirat and Yan in [38] compare the anomaly detection performance of AENs based on three datasets of credit card transactions. Wedge et al. [47] use AENs to learn behavioral features from historical credit card transactions. Paula et al. in [35] use AENs in export controls to detect traces of money laundry and fraud by analyzing volumes of exported goods. Similarly, Schreyer et al. in [43] utilized the reconstruction error of deep AENs to detect anomalous journal entries in two datasets of real-world accounting data.

More recently, GANs are utilized in the context of fraud detection. Fiore et al. in [13] train such networks to generate mimicked anomalies, which were used to augment training data to improve credit card fraud detection classifiers. Choi et al. in [11] train ensembles of generative models to successfully detect anomalies in credit card transactions. Zheng et al. in [51] train LSTM-AENs in an adversarial training set up to detect fraudulent credit card transactions. In another study, Zheng et al. in [50] propose generative denoising GANs to detect telecommunication fraud in the transactions of two financial institutions.

To the best of our knowledge, this work presents the first deep-learning inspired methodology trained in an adversarial training setup to detect anomalous journal entries in real-world accounting data.

3 METHODOLOGY

To detect anomalous journal entries one first has to define "normality" with respect to accounting data. We assume that the majority of

journal entries recorded within an organizations' ERP system relate to regular day-to-day business activities. In order to conduct fraud, perpetrators need to deviate from the "normal". Such a deviating behavior will be recorded by a very limited number of journal entries and their respective attribute values. We refer to journal entries exhibiting such deviating attribute values as *accounting anomalies*.

3.1 Accounting Anomaly Classes

When conducting a detailed examination of real-world journal entries, recorded in large-scaled ERP systems, two characteristics can be observed: First, journal entry attributes exhibit a high variety of distinct attribute values, e.g., due to the high number of vendors or distinct posting amounts, and second, journal entries exhibit strong dependencies between certain attribute values e.g. a document type that is usually posted in combination with a certain general ledger account. Derived from this observation and similarly to Breunig et al. in [9] we distinguish two classes of anomalous journal entries, namely *global* and *local anomalies*:

Global accounting anomalies are journal entries that exhibit unusual or rare individual attribute values. Such anomalies usually relate to skewed attributes, e.g., rarely used ledgers, or unusual posting times. Traditionally, "red-flag" tests performed by auditors during an annual audit, are designed to capture this type of anomaly. However, such tests often result in a high volume of false-positive alerts due to rare but regular events such as reverse postings, provisions and year-end adjustments usually associated with a low fraud risk [43]. Furthermore, when consulting with auditors and forensic accountants, "global" anomalies often refer to "error" rather than "fraud".

Local accounting anomalies are journal entries that exhibit an unusual or rare combination of attribute values while their individual attribute values occur quite frequently, e.g., unusual combinations of general ledger accounts or user accounts used by

several accounting departments. This type of anomaly is significantly more difficult to detect since perpetrators intend to disguise their activities by imitating a regular activity pattern. As a result, such anomalies usually pose a high fraud risk since they correspond to processes and activities that may not be conducted in compliance with organizational standards.

We aim to learn a model that detects both classes of anomalous journal entries in an unsupervised manner. Thereby, the learned model should partition the population of journal entries into semantic meaningful classes that allows for an increased interpretability of the detection results. To achieve this two-fold objective we utilize Adversarial Autoencoders (AAEs), a deep neural network architecture introduced by Makhzani et al. [31]. We provide preliminaries of Autoencoder Neural Networks (AENs) and Generative Adversarial Networks (GANs) that constitute the AAEs in the following. A more detailed presentation can be found in [15].

3.2 Autoencoder Neural Networks

Formally, let X denote a set of N journal entries $x^1, x^2, \dots, x^n$, where each journal entry x^i consists of K attributes $x_1^i, x_2^i, \dots, x_j^i, \dots, x_k^i$. Thereby, x_j^i denotes the j -th attribute of the i -th journal entry. The individual attributes x_j describe the journal entries accounting specific details, e.g., the entries fiscal year, posting type, posting date, amount, general-ledger. Hinton and Salakhutdinov in [18] introduced AENs, a special type of feed-forward multi-layer network that can be trained to reconstruct its input. Formally, AENs are comprised of two nonlinear functions referred to as encoder q_θ and decoder p_θ network [41]. The encoder function $q_\theta(\cdot)$ maps the input $x \in \mathcal{R}^k$ to a code vector $z \in \mathcal{R}^m$ referred to as *latent space representation*, where usually $k > m$. This latent representation is then mapped back by the decoder function $p_\theta(\cdot)$ to a reconstruction $\hat{x} \in \mathcal{R}^k$ of the original input space. In an attempt to achieve $x \approx \hat{x}$ the AEN is trained to minimize the dissimilarity of a given journal entry x^i and its reconstruction $\hat{x}^i = p_\theta(q_\theta(x^i))$ as faithfully as possible. Thereby, the training objective is to learn a set of optimal model parameters θ^* by minimizing the AENs reconstruction loss, formally denoted as:

$$\arg \min_{\theta} \|x^i - p_\theta(q_\theta(x^i))\|. \quad (1)$$

3.3 Generative Adversarial Neural Networks

Goodfellow et al. introduced GANs in [16], a framework for training deep generative models using a mini-max game. The objective is to learn a generator distribution $q(x)$ that matches the real data distribution $p_d(x)$ of journal entries. Instead of trying to explicitly assign probability to every x^i in the data distribution, the GAN aims to learn a set of parameters θ of a generator network q_θ that generates samples from the generator distribution $q(x)$ by transforming a noise variable $z \sim p_n(z)$ into a sample $q_\theta(z)$. Thereby, the generator is trained by playing against an adversarial discriminator network d_ϕ that aims to learn a set of parameters ϕ to distinguish between samples from the true data distribution p_d and the generator's distribution $q(z)$. Both networks establish a min-max adversarial game. A solution to this game can be, expressed as:

$$\min_{q_\theta} \max_{d_\phi} \mathbb{E}_{x \sim p_d(x^i)} [\log d_\phi(x^i)] + \mathbb{E}_{z \sim p_n(z^i)} [\log(1 - d_\phi(q_\theta(z^i)))]. \quad (2)$$

3.4 Adversarial Autoencoders

The AAE architecture, as illustrated in Fig. 2, extends the concept of AEN by imposing an arbitrary prior on the AENs latent space using a GAN training setup [31]. This is achieved by training the AAE jointly in two phases (1) a reconstruction phase as well as (2) an adversarial regularization phase.

In the reconstruction phase, the AAEs encoder network $q_\theta(z|x)$ is trained to learn an aggregated posterior distribution $q(z)$ of the journal entries X over the latent code vector Z . Thereby, the learned posterior distribution corresponds to a compressed representation of the journal entry characteristics. Similarly to AENs, the decoder network $p_\theta(\hat{x}|z)$ of the AAE utilizes the learned latent code vector representations Z to reconstruct the journal entries $\hat{X}$ as faithfully as possible to minimize the AAEs reconstruction error.

In the regularization phase, an adversarial training setup is applied where the encoder network $q_\theta(z|x)$ of the AAE functions as the generator network. In addition, a discriminator network $d_\phi(z)$ is attached on top of the learned latent code vector Z . Similarly to GANs, the discriminator network of the AAE is trained to distinguish samples of an imposed prior distribution $p(z)$ onto Z from the learned aggregated posterior distribution $q(z)$. In contrast, the encoder network is trained to learn a posterior distribution $p(z) \approx q(z)$ that fools the discriminator network into thinking that the samples drawn from $q(z)$ originate from the imposed prior distribution $p(z)$.

3.5 Accounting Anomaly Detection

In order, to detect interpretable accounting anomalies in real-world ERP datasets we propose a novel anomaly score utilizing the introduced AAE architecture. The score builds on the regularisation applied throughout the AAE training process, namely the reconstruction error loss, denoted by Eq. (1), and the adversarial loss, denoted by Eq. (2), described in the following.

The reconstruction loss promotes the AAE to learn a set of non-overlapping latent journal entry representation. However, this may result in a highly "fractured" latent space in which deviating representations are learned for similar journal entries. The additionally applied adversarial loss prevents the fracturing problem. It forces the learned representations z to reside within the high probability density regions of the imposed prior distribution $p(z)$. To partition the latent space into semantic regions we impose a multi-modal prior, e.g., a mixture of Gaussians. Thereby, the interaction of both regularising losses forces the AAE to learn groups of semantic similar journal entries located in close spatial proximity to the modes of the imposed prior.

As a result, the AAE learns with progressing training, a model that disentangles the underlying generative processes of journal entries in the latent space z . Each group of learned representations corresponds to a distinct generative process of journal entries, e.g., depreciation postings or vendor payment postings. To detect potential accounting anomalies, we investigate the individual entries of each group in terms of potential "violations" of one of the two




Figure 3: Exemplary distribution of the 'account key' (technically: 'KTOSL') attribute values (left), the log-normalized 'local' and 'foreign' currency amount (technically: 'DMBTR' and 'WRBTR') attribute values (center), as well as the 'posting key' (technically: 'BSCHL') and 'general ledger account' (technically: 'HKONT') (right) observable in Dataset B.

applied regularising losses. We hypothesize, that anomalous journal entries can be captured by either (1) their latent divergence from the modes of the imposed prior or (2) an increased reconstruction error. Thereby, the type of violation also reveals the anomaly class of the investigated entry, as described in the following:

Mode Divergence (MD): Journal entries that exhibit anomalous attribute values (global anomalies) result in an increased divergence from the imposed multi-modal prior, e.g., as in this work the divergence to the modes of an imposed mixture of multivariate isotropic Gaussians $\mathcal{N}(\mu, I)$, where $\mu \in \mathcal{R}^m$ defines the τ modes of the distinct Gaussians denoted by $\mu = \{\mu^1 \dots \mu^\tau\}$. Throughout the AAE training, the entries will be "pushed" towards the high probability density regions of the prior by the regularization. In order to be able to discriminate between the imposed prior and the learned aggregated posterior the AAE aims to keep the majority of the entries within the high-density regions (modes) of the prior. In contrast, representations that correspond to rare or anomalous journal entries will tend to differ from the imposed modes and be placed in the priors low-density regions. We use this characteristic and obtain an entry's x^i mode divergence D as the Euclidean distance of the entry's learned representation z^i to its closest mode μ^τ . Formally, we derive the mode divergence as denoted by $D_{\theta^*}^\tau(z^i; \mu) = \min_{\tau} \|z^i - \mu^\tau\|^2$ under optimal model parameters θ^* . Finally, we calculate the normalized mode divergence MD as expressed by:

$$MD_{\theta^*}^\tau(x^i) = \frac{D_{\theta^*}^\tau(z^i; \mu) - D_{\theta^*, \min}^\tau}{D_{\theta^*, \max}^\tau - D_{\theta^*, \min}^\tau}, \quad (3)$$

where $D_{\min}$ and $D_{\max}$ denotes the min- and max-values of the obtained mode divergences given by D_{θ^*} and closest mode τ .

Reconstruction Error (RE): Journal entries that exhibit anomalous attribute value co-occurrences (local anomalies) tend to result in an increased reconstruction error [43]. This is caused by the compression capability of the AAE architecture. Anomalous and therefore unique attribute co-occurrences exhibit an increased probability of getting lost in the encoders "lossy" compression. As a result, their low dimensional representation will overlap with regular entries in the latent space and are not reconstructed correctly by the decoder. Formally, we obtain the reconstruction error E of each entry x^i and its reconstruction $\hat{x}^i$ as the squared-difference

denoted by $E_{\theta^*}^\tau(x^i; \hat{x}^i) = \frac{1}{k} \sum_{j=1}^k (x_j^i - \hat{x}_j^i)^2$ under optimal model parameters θ^* . Finally, we calculate the normalized reconstruction error RE as expressed by:

$$RE_{\theta^*}^\tau(x^i; \hat{x}^i) = \frac{E_{\theta^*}^\tau(x^i; \hat{x}^i) - E_{\theta^*, \min}^\tau}{E_{\theta^*, \max}^\tau - E_{\theta^*, \min}^\tau}, \quad (4)$$

where $E_{\min}$ and $E_{\max}$ denotes the min- and max-values of the obtained reconstruction errors given by E_{θ^*} and closest mode τ .

Anomaly Score (AS): Quantifying both characteristics for a given journal entry, we can reasonably conclude (1) if the entry is anomalous and (2) if it was created by a "regular" business activity. To detect global and local accounting anomalies in real-world audit scenarios we propose to score each journal entry x^i by its normalized reconstruction error RE regularized and normalized mode divergence MD given by:

$$AS^\tau(x^i; \hat{x}^i) = \alpha \times RE_{\theta^*}^\tau(x^i; \hat{x}^i) + (1 - \alpha) \times MD_{\theta^*}^\tau(x^i), \quad (5)$$

for each individual journal entry x^i and optimal model parameters θ^* and closest mode τ . We introduce α as a factor to balance both characteristics.

4 EXPERIMENTAL SETUP

In this section, we describe the experimental setup and model training. We evaluate the anomaly detection performance of the proposed scoring based on two datasets of journal entries.

4.1 Datasets and Data Preparation

In general, SAP ERP systems record journal entries and their corresponding attributes predominantly in two database tables: (1) the table "Accounting Document Headers" (technically: "BKPF") contains the meta-information of a journal entry, such as document id, type, date, time, or currency, while (2) the table "Accounting Document Segments" (technically: "BSEG") contains the entry details, such as posting key, general ledger account, debit-credit information, or posting amount. In the context of this work, we extract a subset of the most discriminative journal entry attributes of the "BKPF" and "BSEG" table.

In our experiments we use two datasets of journal entries: a real-world and a synthetic dataset referred to as *dataset A* and




Figure 4: Exemplary AAE latent space distribution of dataset B with progressing network training: imposed prior distribution $p(z)$ consisting of a mixture of $\tau = 10$ Gaussians (left), learned aggregated posterior distribution $g_\theta(z|x)$ after 100 training epochs (center), learned aggregated posterior distribution $g_\theta(z|x)$ after 2,000 training epochs (right).

dataset B in the following. Dataset A is an extract of an SAP ERP instance and encompasses the entire population of journal entries of a single fiscal year². Dataset B is an excerpt of the synthetic dataset presented in [30]³. The majority of attributes recorded in ERP systems correspond to categorical (discrete) variables, e.g. posting date, account, posting type, currency. We pre-process the categorical journal entry attributes to obtain a binary ("one-hot" encoded) representation of each journal entry.

To allow for a detailed analysis and quantitative evaluation of the experiments we inject a small fraction of synthetic global and local anomalies into both datasets. Similar to real audit scenarios this results in a highly unbalanced class distribution of anomalous vs. regular day-to-day entries. The injected global anomalies consist of attribute values not evident in the original data while the local anomalies exhibit combinations of attribute value subsets not occurring in the original data. The true labels are available for both datasets. Each journal entry is labeled as either (1) synthetic *global anomaly*, (2) synthetic *local anomaly* or (3) non-synthetic *regular entry*. The following descriptive statistics summarize both datasets:

- **Dataset A:** contains a total of 307,457 journal entry line items comprised of six categorical and two numerical attributes. The encoding resulted in a total of 401 encoded dimensions for each entry $x^i \in \mathcal{R}^{401}$; and, in total 95 (0.03%) synthetic anomalous journal entries have been injected into the dataset. These entries encompass 55 (0.016%) global anomalies and 40 (0.015%) local anomalies.
- **Dataset B:** contains a total of 533,009 journal entry line items comprised of six categorical and two numerical attributes. The encoding resulted in a total of 618 encoded dimensions for each entry $x^i \in \mathcal{R}^{618}$. In total 100 (0.018%) synthetic, anomalous journal entries have been injected into the

dataset. These entries encompass 70 (0.013%) global anomalies and 30 (0.005%) local anomalies.

Figure 3 illustrates an exemplary distribution of the attributes primarily investigated during and audit, namely the 'account key' (technically: 'KTOSL') attribute values, the log-normalized 'local' and 'foreign currency amount' (technically: 'DMBTR' and 'WRBTR') attribute values, as well as the 'posting key' (technically: 'BSCHL') and 'general ledger account' (technically: 'HKONT') observable in Dataset B.

4.2 Adversarial Autoencoder Training

Our architectural setup follows the AAE architecture [31] as shown in Fig. 2, comprised of three distinct neural networks that are trained in parallel. The encoder network q_θ uses Leaky Rectified Linear Unit (LReLU) activation functions [49] except in the last "bottleneck" layer. Both the decoder network p_θ and the discriminator d_ϕ network use LReLU in all layers except the output layers where a Sigmoid activation function is used. Table 1 depicts the architectural details of the networks which are implemented using PyTorch [34].

Training stability is a main challenge in adversarial training [6] and we face a variety of collapsing and non-convergence scenarios. To determine a stable training setup we sweep the learning rates of the encoder and decoder networks through the interval $\eta \in [10^{-05}, 10^{-02}]$, and the learning rates of the discriminator network through the interval $\eta \in [10^{-07}, 10^{-03}]$. Ultimately, we use the following constant learning rates to learn a stable model of each dataset:

- **Dataset A:** $\eta = 10^{-4}$ for the encoder and the decoder network, $\eta = 10^{-5}$ for the discriminator network; and,
- **Dataset B:** $\eta = 10^{-3}$ for the encoder and the decoder network, $\eta = 10^{-5}$ for the discriminator network.

We train the AAE with mini-batch wise SGD for max. 10,000 training epochs and apply early stopping once the reconstruction loss converges. In accordance with [49] we set the scaling factor of the LReLU to $\alpha = 0.4$ and initialized the AAE parameters as described in [14]. A mini-batch size of 128 journal entries is used

²In compliance with strict data privacy regulations, all journal entry attributes of dataset A have been anonymized using an irreversible one-way hash function during the data extraction process. To ensure data completeness, the journal entry based general ledger balances were reconciled against the standard SAP trial balance reports e.g. the SAP 'RFBILA00' report.

³The original dataset is publicly available via the Kaggle predictive modeling and analytics competitions platform and can be obtained using the following link: <https://www.kaggle.com/ntnu-testimon/paysim1>.

Table 1: Neurons per layer ℓ of the distinct networks that comprise the AAE architecture [31]: encoder q_θ , decoder p_θ and discriminator d_ϕ neural network.

Net	Dataset	$\ell = 1$	2	3	4	5	6	7	8
$q_\theta(z x)$	A	256	128	64	32	16	8	4	2
$p_\theta(\hat{x} z)$	A	2	4	8	16	32	64	128	256
$d_\phi(z)$	A	128	64	32	16				
$q_\theta(z x)$	B	256	64	16	4	2	-	-	-
$p_\theta(\hat{x} z)$	B	2	4	16	64	256	-	-	-
$d_\phi(z)$	B	256	64	16	4	1	-	-	-

in both the reconstruction and the regularization phase. We use Adam optimization [27] and set $\beta_1 = 0.9$ and $\beta_2 = 0.999$ in the optimization of the network parameters. In the reconstruction phase, we use a combined loss function $\mathcal{L}_\theta$ to optimize the encoder q_θ and decoder net p_θ parameters. For each journal entry we calculate (1) the cross-entropy reconstruction error $\mathcal{L}_\theta^{CE}$ of the categorical attribute value encodings x_{cat}^i , e.g., the encoded general ledger account id, and (2) the mean-squared reconstruction error $\mathcal{L}_\theta^{MSE}$ of the numerical attribute value encodings x_{con}^i , e.g., the encoded posting amount, formally expressed by:

$$\mathcal{L}_\theta(x^i; \hat{x}^i) = \gamma \mathcal{L}_\theta^{CE}(x_{cat}^i; \hat{x}_{cat}^i) + (1 - \gamma) \mathcal{L}_\theta^{MSE}(x_{con}^i; \hat{x}_{con}^i) \quad (6)$$

were the parameter γ balances both losses. In this initial work, we set $\gamma = \frac{2}{5}$ in all our experiments to account for the higher amount of categorical attributes in both datasets. In the regularization phase, we calculate the adversarial loss, according to equation 2, when optimizing the parameters of the discriminator d_ϕ .

To partition the learned journal entry representations, we sample from a prior distribution $p(z)$ comprised of a mixture of τ multi-variate isotropic Gaussians $\mathcal{N}(\mu, I)$, where $\mu \in \mathbb{R}^2$. Thereby, τ is a hyperparameter we evaluate when sampling of $\tau \in \{5, 10, 15\}$ Gaussians. Figure 4 shows an exemplary prior consisting of $\tau = 10$ Gaussians as well as the learned aggregated posterior distributions after 100 and 2,000 training epochs.

5 EXPERIMENTAL RESULTS

In this section, we first assess the semantic partitioning of the journal entries by the imposed prior distributions. Afterward, we examine the anomalies detected of each semantic partition.

Semantic partitioning: We qualitatively review the latent space partitioning of the journal entries and assess the accounting specific semantics that is learned by each mode. Figure 5 shows the partitioning result of dataset A, where $\tau = 5$ Gaussians (see appxs. for results of varying τ and dataset B). It can be observed that the AAE learned a rather clean separation of the regular journal entries. The review of the journal entries accounting specific semantics captured by each mode and dataset revealed:

- **Dataset A:** The entries of each partition exhibit a high semantic similarity while each partition corresponds to a general accounting process, such as (1) automated payment run entries postings, (2) outgoing customer invoices, and (3) material movements.

- **Dataset B:** Similarly, the entries of each partition exhibit a high semantic similarity and correspond to the following general accounting processes (1) foreign and domestic invoice postings, (2) purchase of goods, (3) manual payments.

The experimental results when imposing $\tau \in \{5, 10, 15\}$ Gaussians on the latent space of each dataset are presented in the appendix of this work. The results show that the AAE is capable of learning a semantic partition of a given set of journal entries according to that disentangles the entries underlying generative processes. The learned partition provides the auditor a holistic view on a given set of accounting data subject to audit. Furthermore, it allows to effectively obtain a representative and interpretable sample of the data and thereby reduces the audits sampling risk.

Table 2: Mean anomaly score AS obtained per journal entry class and $\alpha = 0.8$ by imposing a prior distribution consisting of $\tau = 5$ (10, and 15) mixture of Gaussians and training the AAE for 5,000 (10,000, and 15,000) epochs (variances originate from the distinct parameter initialization seeds).

Class	Data	AS, $\tau = 5$		AS, $\tau = 10$		AS, $\tau = 15$	
global	A	0.295	± 0.233	0.448	± 0.207	0.532	± 0.244
local	A	0.248	± 0.276	0.275	± 0.143	0.446	± 0.202
regular	A	0.045	± 0.076	0.053	± 0.085	0.110	± 0.034
global	B	0.508	± 0.249	0.442	± 0.245	0.437	± 0.241
local	B	0.357	± 0.260	0.164	± 0.148	0.273	± 0.228
regular	B	0.046	± 0.061	0.070	± 0.041	0.028	± 0.029

Anomaly detection: In addition, we analyze the anomaly detection capability of the proposed anomaly score. Table 2 depicts the mean anomaly score AS obtained for each journal entry class by imposing a prior distribution consisting of $\tau \in \{5, 10, 15\}$ mixture of Gaussians and training the AAE for 5,000 epochs. The quantitative results show the distinct journal entry classes (global, local, and regular entries) can be distinguished according to their anomaly score in both datasets. Figure 5 exemplary shows the anomaly scores obtained for dataset A (see appxs. for results of dataset B) of each journal entry and corresponding partition τ as well as the distribution of the obtained individual anomaly scores. Figure 6 illustrates the change in anomaly scoring when varying the α parameter of the AS. It can be observed that increasing α (and therefore the weight on the reconstruction error of the score) improves the ability to detect local accounting anomalies in the dataset.

We also qualitatively evaluate the characteristics of the anomalies detected in each partition. Therefore, we review journal entries that correspond to a high anomaly score but have not been synthetically injected as anomalies into the evaluated datasets. Thereby, we interpret the detected anomalies of each mode τ in the context of the modes regular entries:

- **Global anomalies** exhibit a low semantic similarity to the regular entries of a mode. The detected entries correspond to rarely observable attribute values and accounting "exceptions", e.g., unusual purchase order amounts or high depreciation, year-end as well as impairment postings.
- **Local anomalies** exhibit high semantic similarity to the regular entries of a mode. The detected entries correspond




Figure 5: Learned AAE latent space representations of the journal entries contained in dataset A after training the AAE for 5,000 epochs and imposing a mixture of $\tau = 5$ Gaussians (left), the anomaly scores AS obtained for $\alpha = 0.8$ of each journal entry x_i and corresponding mode μ_τ (center), the anomaly score distribution (bold line defines the median, upper and lower bound define the 0.05 and 0.95 quantile of the distribution) obtained of each journal entry class with progressing network training (right).




Figure 6: Anomaly scores AS obtained by the application of distinct *RE* vs. *MD* balance factors α after training the AAE for 5,000 epochs on dataset A (see appxs. for results of varying τ and dataset B) and imposing a mixture of $\tau = 5$ Gaussians. It can be observed that decreasing α results in an improved detection of global anomalies (left). In contrast, increasing α results in an improved detection of local anomalies (right).

to rarely observable attribute value combinations, e.g., system users that switched departments, postings exhibiting unusual general ledger account combinations.

In summary, these results lead us to conclude that the proposed anomaly score can be utilized as a highly adaptive anomaly assessment of financial accounting data. It furthermore provides the ability to interpret the detected anomalies of a particular mode in the context of the modes regular journal entry semantics. Initial feedback received by auditors on the detected anomalies underpinned not only their relevance from an accounting perspective.

6 SUMMARY

In this work, we showed that Adversarial Autoencoder (AAE) neural networks can be trained to learn a semantic meaningful representation of journal entries recorded in real-world ERP systems. We

also provided initial evidence that such representations provide a holistic view of the entries and disentangle the underlying generative processes. We believe that the presented approach enables a human auditor or forensic accountant with the ability to sample journal entries for a detailed audit in an interpretable manner and therefore reduce the "sampling risk". In addition, we proposed a novel anomaly score that combines an entry's learned representation and reconstruction error. We demonstrated that the scoring can be interpreted as a highly adaptive and unsupervised anomaly assessment to detect global and accounting anomalies.

We plan to conduct a more detailed investigation of the journal entries' latent space disentanglement. Given the tremendous amount of journal entries annually recorded by organizations, an automated semantic disentanglement improves the transparency of entries to be audited and can save auditors considerable time.

ACKNOWLEDGMENTS

We thank the members of the statistics department at Deutsche Bundesbank and PwC Europe's Forensic Services for their valuable review and remarks. Opinions expressed in this work are solely those of the authors, and do not necessarily reflect the view of the Deutsche Bundesbank or PricewaterhouseCoopers (PwC) International Ltd. and its network firms.

REFERENCES

- [1] ACFE. 2018. *Report to the Nations on Occupational Fraud and Abuse, The 2018 Global Fraud Study*. Association of Certified Fraud Examiners (ACFE). <https://s3-us-west-2.amazonaws.com/acfe-public/2018-report-to-the-nations.pdf>
- [2] AICPA. 2002. *Consideration of Fraud in a Financial Statement Audit*. American Institute of Certified Public Accountants (AICPA). 1719–1770 pages. <https://www.aicpa.org/Research/Standards/AuditAttest/DownloadableDocuments/AU-00316.pdf>
- [3] Farzaneh A. Amani and Adam M. Faddalla. 2017. Data mining applications in accounting: A review of the literature and organizing framework. *International Journal of Accounting Information Systems* 24 (2017), 32–58.
- [4] Argyris Argyrou. 2012. Auditing Journal Entries Using Self-Organizing Map. In *Proceedings of the Eighteenth Americas Conference on Information Systems (AMCIS)*. Seattle, Washington, 1–10.
- [5] Argyris Argyrou. 2013. Auditing Journal Entries Using Extreme Value Theory. *Proceedings of the 21st European Conference on Information Systems* 1, 2013 (2013).
- [6] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein gan. *arXiv preprint arXiv:1701.07875* (2017).
- [7] Stephen Bay, Krishna Kumaraswamy, Markus G Anderle, Rohit Kumar, David M Steier, Almaden Blvd, and San Jose. 2006. Large Scale Detection of Irregularities in Accounting Data. In *Data Mining, 2006. ICDM'06. Sixth International Conference on*. IEEE, 75–86.
- [8] Frank Benford. 1938. The Law of Anomalous Numbers. *Proceedings of the American Philosophical Society* 78, 4 (1938), 551–572.
- [9] Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. 2000. LOF: Identifying Density-Based Local Outliers. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*. 1–12.
- [10] Raghavendra Chalaphaty and Sanjay Chawla. 2019. Deep learning for anomaly detection: A survey. *arXiv preprint arXiv:1901.03407* (2019).
- [11] Hyunsun Choi and Eric Jang. 2018. Generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392* (2018).
- [12] R. S. Debreceeny and G. L. Gray. 2010. Data mining journal entries for fraud detection: An exploratory study. *International Journal of Accounting Information Systems* 11, 3 (2010), 157–181.
- [13] Ugo Fiore, Alfredo De Santis, Francesca Perla, Paolo Zanetti, and Francesco Palmieri. 2017. Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences* (2017).
- [14] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. 249–256.
- [15] I. Goodfellow, Y. Bengio, and A. Courville. 2016. *Deep Learning*. MIT Press.
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [17] Thomas W Hall, James E Hunton, and Bethane Jo Pierce. 2002. Sampling practices of auditors in public accounting, industry, and government. *Accounting Horizons* 16, 2 (2002), 125–136.
- [18] Geoffrey E Hinton and Ruslan R Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *science* 313, 5786 (2006), 504–507.
- [19] IFAC. 2009. *International Standards on Auditing 240, The Auditor's Responsibilities Relating to Fraud in an Audit of Financial Statements*. International Federation of Accountants (IFAC). 155–197 pages.
- [20] IFAC. 2009. *International Standards on Auditing 530, Audit Sampling*. International Federation of Accountants (IFAC). 441–457 pages.
- [21] Asadul Khandoker Islam, Malcolm Corney, George Mohay, Andrew Clark, Shane Bracher, Tobias Raub, and Ulrich Flegel. 2010. Fraud detection in ERP systems using Scenario matching. *IFIP Advances in Information and Communication Technology* 330 (2010), 112–123.
- [22] Mieke Jans, Nadine Lybaert, and Koen Vanhoof. 2010. Internal fraud risk reduction: Results of a data mining case study. *International Journal of Accounting Information Systems* 11, 1 (2010), 17–41.
- [23] Mieke Jans, Jan Martijn Van Der Werf, Nadine Lybaert, and Koen Vanhoof. 2011. A business process mining application for internal transaction fraud mitigation. *Expert Systems with Applications* 38, 10 (2011), 13351–13359.
- [24] Zahra Kazemi and Houman Zarrabi. 2017. Using deep networks for fraud detection in the credit card transactions. In *2017 IEEE 4th International Conference on Knowledge-Based Engineering and Innovation (KBEI)*. IEEE, 0630–0633.
- [25] Rq Khan and Mw Corney. 2009. A role mining inspired approach to representing user behaviour in ERP systems. In *Proceedings of The 10th Asia Pacific Industrial Engineering and Management Systems Conference*. 2541–2552.
- [26] Roheena Khan, Malcolm Corney, Andrew Clark, and George Mohay. 2010. Transaction Mining for Fraud Detection in ERP Systems. *Industrial Engineering and Management Systems* 9, 2 (2010), 141–156.
- [27] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [28] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436.
- [29] Chongxuan Li, Kun Xu, Jun Zhu, and Bo Zhang. 2017. Triple Generative Adversarial Nets. *JAMA Internal Medicine* 177, 3 (mar 2017), 326. [arXiv:1703.02291](https://arxiv.org/abs/1703.02291)
- [30] E. A. Lopez-Rojas, A. Elmir, and S. Axelsson. 2016. PaySim: A financial mobile money simulator for fraud detection. In *The 28th European Modeling and Simulation Symposium-EMSS, Larnaca, Cyprus*.
- [31] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. 2015. Adversarial Autoencoders. *arXiv* (2015), 1–10. [arXiv:1511.05644](https://arxiv.org/abs/1511.05644)
- [32] Shahar Markovitch and Paul Willmott. 2014. Accelerating the digitization of business processes. *McKinsey & Company* (2014), 1–5.
- [33] Mary McGlohon, Stephen Bay, Markus G Mg Anderle, David M Steier, and Christos Faloutsos. 2009. SNARE: A Link Analytic System for Graph Labeling and Risk Detection. *KDD-09: 15th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2009), 1265–1273.
- [34] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. (2017).
- [35] Ebberth L. Paula, Marcelo Ladeira, Rommel N. Carvalho, and Thiago Marzagão. 2017. Deep learning anomaly detection as support fraud investigation in Brazilian exports and anti-money laundering. In *Proceedings - 2016 15th IEEE International Conference on Machine Learning and Applications, ICMLA 2016*. 954–960.
- [36] Isabel Pedrosa and Carlos J Costa. 2014. New trends on CAATs: what are the Chartered Accountants' new challenges? *ISDOC '14 Proceedings of the International Conference on Information Systems and Design of Communication, May 16–17, 2014, Lisbon, Portugal* (2014), 138–142.
- [37] Marco AF Pimentel, David A Clifton, Lei Clifton, and Lionel Tarassenko. 2014. A review of novelty detection. *Signal Processing* 99 (2014), 215–249.
- [38] Apapan Pumsirirat and Liu Yan. 2018. Credit card fraud detection using deep learning based on auto-encoder and restricted boltzmann machine. *International Journal of advanced computer science and applications* 9, 1 (2018), 18–25.
- [39] PwC. 2018. *Pulling Fraud Out of the Shadows, The Global Economic Crime Survey 2018*. PricewaterhouseCoopers LLP. <https://www.pwc.com/gx/en/forensics/global-economic-crime-and-fraud-survey-2018.pdf>
- [40] Martin Renström and Timothy Holmsten. 2018. Fraud Detection on Unlabeled Data with Unsupervised Machine Learning.
- [41] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. 1985. *Learning internal representations by error propagation*. Technical Report. California Univ San Diego La Jolla Inst for Cognitive Science.
- [42] SAP. 2019. *SAP Global Corporate Affairs, Corporate Factsheet 2019*. <https://www.sap.com/documents/2017/04/4666ecdd-b67c-0010-82c7-eda71af511fa.html>
- [43] Marco Schreyer, Timur Sattarov, Damian Borth, Andreas Dengel, and Bernd Reimer. 2017. Detection of anomalies in large scale accounting data using deep autoencoder networks. *arXiv preprint arXiv:1709.05254* (2017).
- [44] Suwardy Seow, Poh-Sun; Sun, Gary; Themin. 2016. Data Mining Journal Entries for Fraud Detection : a Replication of Debreceeny and Gray ' S (2010). *Journal of Forensic Investigative Accounting* 3, 8 (2016), 501–514.
- [45] Tommie. Singleton and Aaron J. Singleton. 2006. *Fraud auditing and forensic accounting*. John Wiley & Sons.
- [46] Tom Sweers, Tom Heskes, and Jesse Krijthe. 2018. Autoencoding Credit Card Fraud. (2018).
- [47] Roy Wedge, James Max Kanter, Santiago Moral Rubio, Sergio Iglesias Perez, and Kalyan Veeramachaneni. 2017. Solving the " false positives" problem in fraud prediction. *arXiv preprint arXiv:1710.07709* (2017).
- [48] Joseph T. Wells. 2017. *Corporate Fraud Handbook: Prevention and Detection*. John Wiley & Sons.
- [49] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. 2015. Empirical Evaluation of Rectified Activations in Convolution Network. *ICML Deep Learning Workshop* (2015), 1–5. [arXiv:1505.00853](https://arxiv.org/abs/1505.00853)
- [50] Panpan Zheng, Shuhan Yuan, Xintao Wu, Jun Li, and Aidong Lu. 2018. One-class adversarial nets for fraud detection. *arXiv preprint arXiv:1803.01798* (2018).
- [51] Yu-Jun Zheng, Xiao-Han Zhou, Wei-Guo Sheng, Yu Xue, and Sheng-Yong Chen. 2018. Generative adversarial network based telecom fraud detection at the receiving bank. *Neural Networks* 102 (2018), 78–86.

APPENDIX

Experimental Results - Dataset A




Figure 7: Learned AAE latent space representations of the journal entries contained in dataset A after training the AAE for 10,000 epochs and imposing a mixture of $\tau = 10$ Gaussians (left), the anomaly scores AS obtained for $\alpha = 0.8$ of each journal entry x_i and corresponding mode μ_τ (center), the anomaly score distribution (bold line defines the median, upper and lower bound define the 0.05 and 0.95 quantile of the distribution) obtained of each journal entry class with progressing training (right).




Figure 8: Learned AAE latent space representations of the journal entries contained in dataset A after training the AAE for 15,000 epochs and imposing a mixture of $\tau = 15$ Gaussians (left), the anomaly scores AS obtained for $\alpha = 0.8$ of each journal entry x_i and corresponding mode μ_τ (center), the anomaly score distribution (bold line defines the median, upper and lower bound define the 0.05 and 0.95 quantile of the distribution) obtained of each journal entry class with progressing training (right). It can be observed, based on the anomaly scores obtained with progressing training of mode $\tau = 12$, that mode stability wasn't reached entirely after training the AAE for 15,000 epochs.

Experimental Results - Dataset B




Figure 9: Learned AAE latent space representations of the journal entries contained in dataset B after training the AAE for 5,000 epochs and imposing a mixture of $\tau = 5, 10$ and 20 Gaussians (top to bottom) (left), the anomaly scores AS obtained for $\alpha = 0.8$ of each journal entry x_i and corresponding mode μ_τ (center), the anomaly score distribution (bold line defines the median, upper and lower bound define the 0.05 and 0.95 quantile of the distribution) obtained of each journal entry class with progressing training (right).