

Les cahiers de recherche du CISMF
CISMF Research Paper Series

**Deep Semi-Supervised Anomaly Detection for Finding Fraud in
the Futures Market**

Timothy DeLise^a

^a Department of Mathematics and Statistics, Université de Montréal, Canada

Juillet / July 2023

Deep Semi-Supervised Anomaly Detection for Finding Fraud in the Futures Market

Timothy DeLise*

July 2023

Abstract

Modern financial electronic exchanges are an exciting and fast-paced marketplace where billions of dollars change hands every day. They are also rife with manipulation and fraud. Detecting such activity is a major undertaking, which has historically been a job reserved exclusively for humans. Recently, more research and resources have been focused on automating these processes via machine learning and artificial intelligence. Fraud detection is overwhelmingly associated with the greater field of anomaly detection, which is usually performed via unsupervised learning techniques because of the lack of labeled data needed for supervised learning. However, a small quantity of labeled data does often exist. This research article aims to evaluate the efficacy of a deep *semi-supervised* anomaly detection technique, called Deep SAD, for detecting fraud in high-frequency financial data. We use exclusive proprietary limit order book data from the TMX exchange in Montréal, with a small set of true labeled instances of fraud, to evaluate Deep SAD against its unsupervised predecessor. We show that incorporating a small amount of labeled data into an unsupervised anomaly detection framework can greatly improve its accuracy.

1 Introduction

Fraud detection in high frequency data is one important aspect of forensic work performed by the oversight teams at derivatives exchanges worldwide. Each day millions of lines of orders and transactions are recorded in real time. Certain technically destructive behaviors have been identified, such as *layering* or *spoofing* [Gid20], but other more subtle types exist such as insider trading. It is certainly a complicated task, often with legal implications associated. Given

*Département de mathématiques et de statistique, Université de Montréal

Acknowledgment: This research was funded by the Graduate Excellence Scholarship of the Centre d'intelligence en surveillance des marchés financiers (CISMF) de l'UQAM in the summer of 2021. Additional thanks to the TMX Group, our private sponsor and collaborator in this research for providing data, insight and context.

the importance and sensitivity of the job, the painstaking task of identifying instances of fraud has historically been left up to humans.

Amidst the recent clamour around topics like artificial intelligence and machine learning, in addition to the proliferation of high frequency market data, the Toronto Mercantile Exchange's (TMX) derivatives trading operation, located in Montreal, Quebec, has taken the initiative, inspired by recent developments in ML and AI, to research automated ways that technology can help in the fight to uphold the integrity of financial markets in Canada and worldwide.

Financial institutions often have a lot of data, such as order book data, or credit card transaction data, but they do not necessarily know if each data point was an instance of a fraud (or anomaly) or not. The process of flagging such data and finding out is a very slow and burdensome. After years of investigating, a team may only have found a handful of confirmed instance of fraud, while the size of data is in millions or billions of transactions. To be fair, most transactions we assume are normal, with fraudulent or anomalous data representing only a small portion of all the data. Fraud detection in financial data has been dominated by the the field of unsupervised learning [GU16; Gol23]. Unsupervised learning techniques aim to find models which can learn and explain the latent structure of data[GBC16; Mak18]. They do not require labeled targets as in supervised learning problems. This is one of their main advantages to the field of fraud detection. Without labeled data, it is also one of the only options.

One of the most popular unsupervised techniques for fraud detection in financial data using ML and AI are deep neural network autoencoders [ZZJ19]. An autoencoder model's job is simple. We provide our data as input, and the model simply needs to recreate that exact data as best it can. We impose limitations on the model, such as a low-dimensional bottleneck in the neural network, which forces a low-dimensional representation of the input data [RHW86]. For each input data point, the model produces an output, and training is performed via gradient descent-style algorithms to minimize the mean-squared distance between input and output. No model is perfect, and there are still non-zero errors associated with each data point after training. The magnitude of the errors can be a measure of how different, or anomalous, a data point is. A large error means that the autoencoder chose to sacrifice performance on that data in favor of other data. It means that this data is out-of-the-ordinary. Auto-encoder-based anomaly detection programs associate training error with *anomaly score*. The higher the score, the greater chance that the data point is an anomaly. Fraud detection schemes then associate the idea of *anomaly* with *fraud*. It is a loose specification and teams should make sure to realize that a high anomaly score is not alone evidence of a crime.

The fact that we often do have a few examples of fraud (labeled data) has made the idea of a *semi-supervised* approach to fraud detection very enticing. Indeed it has been the subject of recent research [Vil+21], and it makes sense. Any amount of labeled data, however small, is valuable. The authors of [Goe+13] point out that the performance of unsupervised approaches often too low to be useful, but purely supervised technique fail to generalize well to unseen new kinds of anomalies. **They recommend that any semi-supervised**

anomaly detection system should be grounded in the unsupervised paradigm.

The TMX derivatives exchange had already been working with an unsupervised model similar to the Deep Support Vector Data Description (Deep SVDD) model[Ruf+18]. The formulation of this model is a slight deviation from the traditional autoencoder. Instead of training the deep neural network to recreate the input data, the network is trained to minimize the volume of a hypersphere that encloses the network representation of the data. This effect is realized through a novel objective loss function for this purpose, in line 1 below. In this expression x_i is the i^{th} data point out of n total data points, ϕ is the neural network model function, \mathcal{W} are the weights of the neural network to be optimized, and \mathbf{c} is the center of the hypersphere.

$$\min_{\mathcal{W}} \frac{1}{n} \sum_{i=1}^n \|\phi(x_i; \mathcal{W}) - \mathbf{c}\|^2 \quad (1)$$

The model has shown promise, but there remains a question around the small set of labeled anomalies (frauds) that we have. Can we further leverage this data to improve model performance? The purpose of this current paper is to investigate the efficacy of moving to a semi-supervised paradigm. To this end, we explore a further development, from the authors of the Deep SVDD model, called Deep SAD, which stands for *Deep Semi-Supervised Anomaly Detection*[Ruf+19]. This is a method, grounded in the unsupervised learning paradigm, which takes advantage of an arbitrary number of labeled examples alongside the vast amount of unlabeled data. This is accomplished by introducing a new term into the objective function on line 1 above. The new objective function minimizes the distance of the neural network representation from the center of the hypersphere for the unlabeled data while maximizing the same distance for the small set of labeled true anomalies. Normal data should be toward the center, and anomalous data should be toward the outside. This new objective function is written here in equation 2.

$$\min_{\mathcal{W}} \frac{n}{n+m} \sum_{i=1}^n \|\phi(x_i; \mathcal{W}) - \mathbf{c}\|^2 + \frac{\eta}{n+m} \sum_{j=1}^m (\|\phi(\tilde{x}_j; \mathcal{W}) - \mathbf{c}\|^2)^{\tilde{y}_j} \quad (2)$$

In addition to the original n unlabeled data points, x_i , we also have m labeled data points \tilde{x}_j with associated labels $\tilde{y}_j \in \{-1, 1\}$. Where $\tilde{y}_j = 1$ denotes normal data points and $\tilde{y}_j = -1$ denotes anomalous samples. We see that when the labeled data is from normal data, the new term just replicates the job of the old term. But, for the anomalous data, the negative in the exponent induces a maximization of the distance for those data points. It should be noted here that the L2 penalty weight decay term has been omitted from the previous loss function for clarity. In practice there is an additional L2 penalty on the magnitude of the weights of the deep neural network in these objective functions.

Finally, the anomaly score, $\mathbf{s}(x)$, of a data point x is thus the distance of that data point’s neural network representation from the center of the hypersphere.

$$\mathbf{s}(x) = \|\phi(x; \mathcal{W}) - \mathbf{c}\| \tag{3}$$

The empirical research to follow is aimed to confirm that Deep SAD indeed improves anomaly detection scores on a concrete internal data set provided by TMX. The following section outlines the data description and experimental methods. This is followed by the section with the results, and finally we conclude with a discussion of the findings.

2 Methods

2.1 Data Description

Our data set consists of high-frequency limit order book data [OHa15] from one product of the futures market exchanges operated by TMX. The time duration of the data is one month. This is to say that we have every limit order book update to this particular exchange-traded financial asset over a one month time span. This totals 5,220,091 data points. Each data point contains 20 numerical fields representing the top 10 levels of the order book. Each level of the order book contains data for the price and size of the orders. This is one of the de facto standard data formats that are generated and distributed by electronic exchanges worldwide. For a detailed description of limit order books and their associated data, please refer to this resource [Gou+13].

Out of the 5,220,091 available data points, 293 of them are labeled, and all the labeled data are true anomalies. In the context of the objective function above, which we will use later on, $n = 5,219,798$, $m = 293$, and $\tilde{y}_j = -1 \forall j \in \{1, 2, \dots, m\}$. The data set is visualized temporally in figure 1.

2.2 Experimental Design

For the experiment we utilize randomly initialized 3-layer feed-forward neural networks with 100 nodes per hidden layer. The input layer is size 20 and so is the output layer, which equals the shape of the input data (20 dimensions). ReLU activation functions are used on the output of each layer, except the final output layer, which is just a linear output layer. This kind of neural network is also called a ReLU Net, and is a standard example of a deep neural network model with decent capacity, and the quality of universal function approximation [Han19]. Both SVDD and DeepSAD undergo an unsupervised pre-training phase, where they are both trained for 1,000 epochs as an autoencoder, learning to approximate the input. The center of the hypersphere is then initialized as the mean output of the neural net over all the training data. Finally, the main training routine is run for 10,000 epochs. This is the stage where the SVDD and the Deep SAD models differ. The objective function for SVDD is defined in equation 1 and for Deep SAD in equation 2 of the introduction. Once training

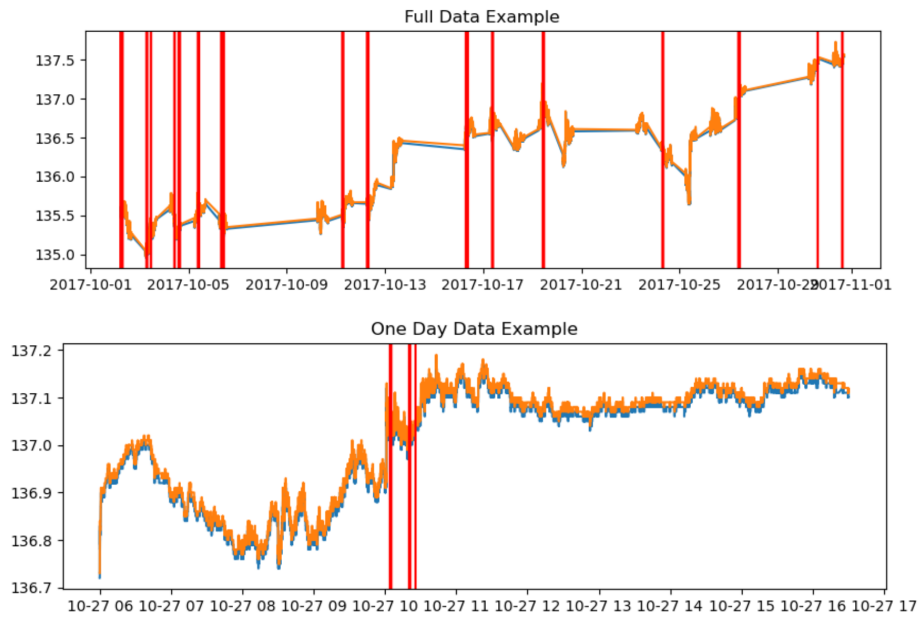


Figure 1: This figure helps describe the layout of the data. We plot the top two order book price levels over time. We see the dates vary over the month of October 2017. The price varies between 135 to 137.5 dollars over this time. The red lines indicate the location of the true anomaly labels. The top pane is the entire data set and the bottom pane zooms in on one day, October 27th. This helps understand how sparse the labeled data points are.

is complete, the anomaly scores (equation 3) of the test data are then evaluated and stored for downstream analysis.

The data set is trained and evaluated in a 3-fold cross validation routine. During each fold, the model is trained on 2/3 of the data and evaluated on the remaining 1/3 of the data. The data is not shuffled and split by date, so each fold of data is a contiguous chunk of time. There isn't any time frame overlapping from one fold to another. The entire experiment is run twice, each time on 3 folds of data. In the following sections, we display results from each of the 6 trials separately, which we believe helps to emphasize the findings. We display results both in-sample (training set) and out-of-sample (test set). This enhances transparency into the generalization ability of the models.

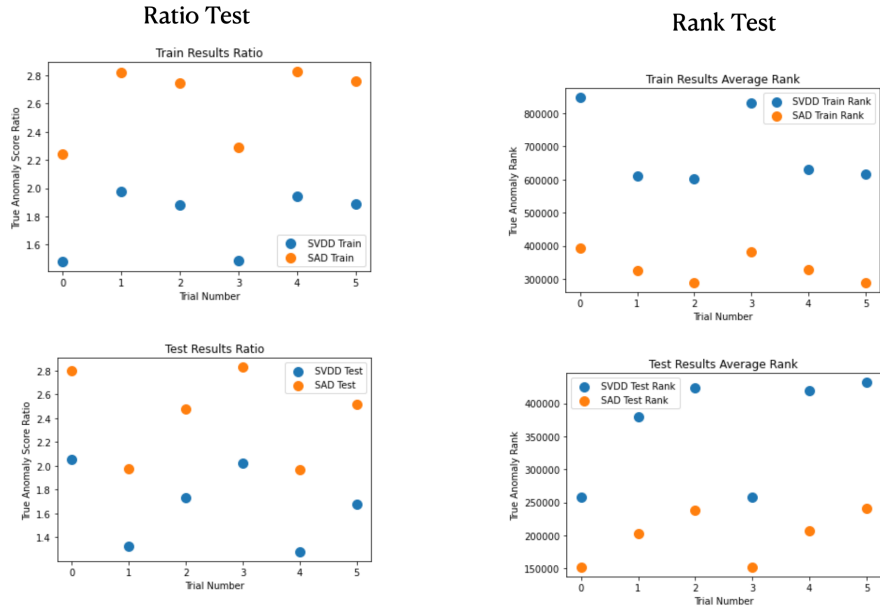
In order to compare the two techniques against each other, we've developed two criteria which we call the *ratio test* and the *rank test*. The ratio test displays the mean anomaly score of the labeled data divided by the mean anomaly score of the unlabeled data. The rank test shows the average rank of the true anomalies relative to all the data when sorted by anomaly score. For the ratio test, a score above 1 shows the model is working at all, and a better model will produce higher scores. For the rank test, a better model will produce lower scores, since we have ordered the data so that a rank of 1 is the best, with the highest score.

3 Results

The results are quite clear and point to a convincing benefit of the Deep SAD model over the SVDD model. The results of the experiment are graphically shown in figure 2.

The results of the ratio test show that, both in-sample and out-of-sample, the DeepSAD model attains a higher score. The scores range from 1.5 to 3. This means that the SVDD model, which scores at the lower end of the range, still attributes higher anomaly scores to the true anomalies as opposed to the unlabeled data. This is good. However, the Deep SAD model increases this ratio, so that the true anomalies score even higher compared to normal data. This shows that, by incorporating a small amount of labeled data into our workflow, we've greatly increased the performance of our anomaly detection system.

The rank test paints the same picture. We see that the Deep SAD model leads to better identify instances of true anomalies in the data, both in-sample and out-of-sample. We should clarify here that the lower the rank, the higher the score. The data point with rank zero would have the largest distance from the center of the hypersphere.



(a) Results from the **Ratio Test**. In every trial, the ratio test score for Deep SAD is higher (better) than SVDD, both in-sample and out-of sample.

(b) Results from the **Rank Test**. In every trial, the rank test score for Deep SAD is lower (better) than SVDD, both in-sample and out-of sample.

Figure 2: In each plot, the trial number is along the x-axis and the test score is along the y-axis. On the left side we have results from the ratio test and on the right side, results from the rank test. The top show in-sample results (train) and the bottom displays out-of-sample (test) results. Data points for the SVDD model are in blue and Deep SAD is in blue.

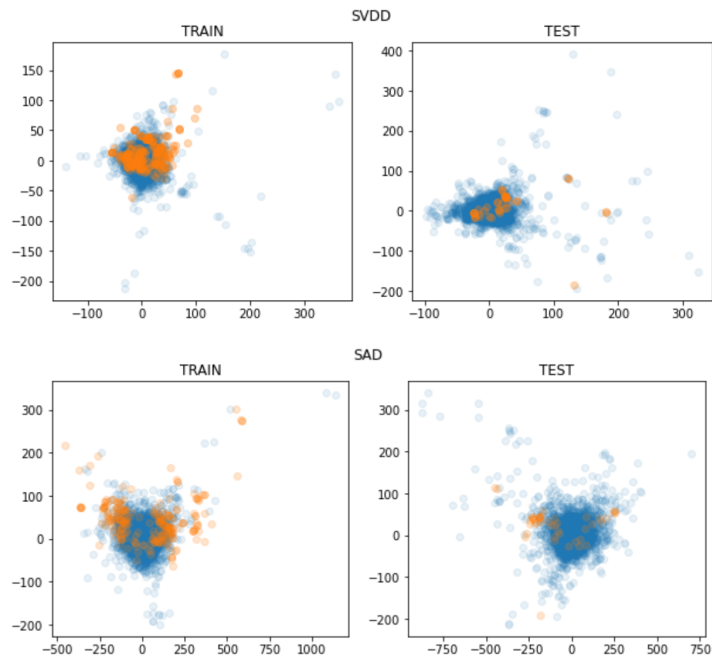


Figure 3: Visualizing the Distance to the Center of the Hypersphere. In this figure we plot the neural network representations of the data points after performing a principle component analysis to reduce the dimension of the data from 20 dimensions to 2. The blue dots are from unlabeled data and the orange dots are from the true labeled anomalies. There are far fewer labeled data points, so they are plotted on top. The results from the SVDD model, train and test set, are presented on the top row, while those for the Deep SAD model are presented on the second row.

4 Discussion

We’ve taken quite a journey during this research, and come a long way. We started by facing a seemingly insurmountable task of identifying fraud in the high-frequency financial markets, akin to finding a needle in a haystack. We’ve built on the ponderous ideas of the fore bearers of applied ML and AI. We’ve seen how unsupervised learning can form a basis to the fraud detection framework. Finally we’ve discovered here that we can effectively leverage a small amount of labeled anomaly (fraud) data to enhance our unsupervised pipeline.

At the crux of this new Deep SAD approach is idea that is baked into the objective function of the model; that the distance from the center of a hypersphere in high-dimensional space is an anomaly score. In order to get more intuition behind this, we’ve created a visualization which allows us to see the data points projected onto the space of the hypersphere, shown in figure 3. In this figure, each data point has been run through the neural network, and output data is compressed from 20 dimensional to 2 using principle component analysis. The blue data points are the unlabeled data, and the orange are the labeled true anomalies. For the SVDD model in the training set, we see the relative spread of the unlabeled data and the labeled data is about the same, while in the Deep SAD training set, the orange dots are pushed farther away from the center of the cluster. This visually confirms that, indeed, the Deep SAD training technique makes the true anomalies appear farther from the center of the hypersphere as opposed to the SVDD model. In the test sets, while there are less true anomalies, so the data is scarce, we see a similar pattern. For the SVDD model the true anomalies appear right near the center of the main cluster, while for the Deep SAD model, the true anomalies show up more spread out, farther away from the center.

This qualitative interpretation, along with the results from the previous section, suggest that this model generalizes surprisingly well. Affects in the training data transfer to the test data with minimal loss.

Since this small amount of labeled data exists, it seems like quite a sacrifice to not incorporate it into an fraud detection pipeline for financial markets. We’ve shown here that semi-supervised frameworks like Deep SAD, grounded in the unsupervised paradigm and implementing deep learning, can benefit from both unlabeled data and a small amount labeled data. Surely there is value to creating and saving increasing amount of labeled true anomaly data. This research is hopeful to further guide students and practitioners to advance in this field of semi-supervised anomaly detection.

References

- [RHW86] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. “Learning Internal Representations by Error Propagation”. In: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*. Ed. by David E. Rumelhart and

- James L. McClelland. Cambridge, MA: MIT Press, 1986, pp. 318–362.
- [Goe+13] N. Goernitz et al. “Toward Supervised Anomaly Detection”. In: *Journal of Artificial Intelligence Research* 46 (Feb. 2013), pp. 235–262. DOI: 10.1613/jair.3623. URL: <https://doi.org/10.1613%2Fjair.3623>.
- [Gou+13] Martin D. Gould et al. “Limit order books”. In: *Quantitative Finance* 13.11 (2013), pp. 1709–1742. URL: <https://EconPapers.repec.org/RePEc:taf:quantf:v:13:y:2013:i:11:p:1709-1742>.
- [OH15] Maureen O’Hara. “High frequency market microstructure”. In: *Journal of Financial Economics* 116.2 (2015), pp. 257–270. ISSN: 0304-405X. DOI: <https://doi.org/10.1016/j.jfineco.2015.01.003>. URL: <https://www.sciencedirect.com/science/article/pii/S0304405X15000045>.
- [GU16] Markus Goldstein and Seiichi Uchida. “A Comparative Evaluation of Unsupervised Anomaly Detection Algorithms for Multivariate Data”. In: *PLoS ONE* 11.4 (2016), e0152173. DOI: 10.1371/journal.pone.0152173.
- [GBC16] Ian J. Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. <http://www.deeplearningbook.org>. Cambridge, MA, USA: MIT Press, 2016.
- [Mak18] Alireza Makhzani. “Unsupervised Representation Learning with Autoencoders”. Thesis. June 2018. URL: <https://hdl.handle.net/1807/89800>.
- [Ruf+18] Lukas Ruff et al. “Deep One-Class Classification”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by Jennifer Dy and Andreas Krause. Vol. 80. Proceedings of Machine Learning Research. PMLR, July 2018, pp. 4393–4402. URL: <https://proceedings.mlr.press/v80/ruff18a.html>.
- [Han19] Boris Hanin. “Universal Function Approximation by Deep Neural Nets with Bounded Width and ReLU Activations”. In: *Mathematics* 7.10 (Oct. 2019), p. 992. DOI: 10.3390/math7100992. URL: <https://doi.org/10.3390%2Fmath7100992>.
- [Ruf+19] Lukas Ruff et al. “Deep Semi-Supervised Anomaly Detection”. In: *CoRR* abs/1906.02694 (2019). arXiv: 1906.02694. URL: <http://arxiv.org/abs/1906.02694>.
- [ZZJ19] Junyi Zou, Jinliang Zhang, and Ping Jiang. “Credit Card Fraud Detection Using Autoencoder Neural Network”. In: *CoRR* abs/1908.11553 (2019). arXiv: 1908.11553. URL: <http://arxiv.org/abs/1908.11553>.
- [Gid20] Mark Gideon. “Spoofing and Layering”. In: *The Journal of Corporation Law* 45.2 (Jan. 2020). ISSN: 0360-795X.

- [Vil+21] Miryam Elizabeth Villa-Pérez et al. “Semi-supervised anomaly detection algorithms: A comparative summary and future research directions”. In: *Knowledge-Based Systems* 218 (2021), p. 106878. ISSN: 0950-7051. DOI: <https://doi.org/10.1016/j.knsys.2021.106878>. URL: <https://www.sciencedirect.com/science/article/pii/S0950705121001416>.
- [Gol23] Markus Goldstein. “Special Issue on Unsupervised Anomaly Detection”. In: *Applied Sciences* 13.10 (2023). ISSN: 2076-3417. DOI: [10.3390/app13105916](https://doi.org/10.3390/app13105916). URL: <https://www.mdpi.com/2076-3417/13/10/5916>.