

Summer 6-15-2016

# IMPROVING DECISION ANALYTICS WITH DEEP LEARNING: THE CASE OF FINANCIAL DISCLOSURES

Stefan Feuerriegel

*University of Freiburg*, stefan.feuerriegel@is.uni-freiburg.de

Ralph Fehrer

*University of Freiburg*, ralphfehrer@gmail.com

Follow this and additional works at: [http://aisel.aisnet.org/ecis2016\\_rip](http://aisel.aisnet.org/ecis2016_rip)

---

## Recommended Citation

Feuerriegel, Stefan and Fehrer, Ralph, "IMPROVING DECISION ANALYTICS WITH DEEP LEARNING: THE CASE OF FINANCIAL DISCLOSURES" (2016). *Research-in-Progress Papers*. 22.

[http://aisel.aisnet.org/ecis2016\\_rip/22](http://aisel.aisnet.org/ecis2016_rip/22)

This material is brought to you by the ECIS 2016 Proceedings at AIS Electronic Library (AISeL). It has been accepted for inclusion in Research-in-Progress Papers by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact [elibrary@aisnet.org](mailto:elibrary@aisnet.org).

# IMPROVING DECISION ANALYTICS WITH DEEP LEARNING: THE CASE OF FINANCIAL DISCLOSURES

*Research in Progress*

Feuerriegel, Stefan, University of Freiburg, Freiburg, Germany, stefan.feuerriegel@is.uni-freiburg.de

Fehrer, Ralph, University of Freiburg, Freiburg, Germany, ralphfehrer@gmail.com

## Abstract

*Decision analytics commonly focuses on the text mining of financial news sources in order to provide managerial decision support and to predict stock market movements. Existing predictive frameworks almost exclusively apply traditional machine learning methods, whereas recent research indicates that these methods are not sufficiently capable of extracting suitable features and capturing the non-linear nature of complex tasks. As a remedy, novel deep learning models aim to overcome this issue by extending classical neural networks with additional hidden layers. Indeed, deep learning often provides a viable approach to achieve a high predictive performance. In this paper, we adapt the novel deep learning technique to financial decision support, where we aim to predict the direction of stock movements following financial disclosures. As a result, our paper shows how deep learning can outperform the accuracy of benchmarks for machine learning by 5.66 %.*

*Keywords: Decision Analytics, IS Research Methodologies, Financial Prediction, Data Mining, Business Intelligence (BI), Text Mining, Information Processing.*

## 1 Introduction

Organizations are constantly looking for ways to improve their decision-making processes in core areas, such as marketing, firm communication, production and procurement (Turban, 2011). While the classical approach relies on having humans devise simple decision-making rules, modern decision support can also exploit statistical evidence that originates from analyzing data (Apte, Liu, Pednault, and Smyth, 2002; Arnott and Pervan, 2005; Asadi Someh and Shanks, 2015; Boylan and Syntetos, 2012; Davenport, 2006; Vizecky, 2011). This data-driven decision support is nowadays also fueled by the *Big Data* era (Boyd and Crawford, 2012; Chen, Chiang, and Storey, 2012; Halper, 2011; Power, 2014). The term Big Data usually refers to data that is massive in size. In addition, such data comes often in different formats (e. g. video, text), changes quickly and is subject to uncertainty (IBM, 2013).

Crucial aspects of data-driven decision support systems entail the prediction of future events, such as consumer behavior or stock market reactions to press releases, based on an analysis of historical data (Apte, Liu, Pednault, and Smyth, 2002; Vizecky, 2011). Decision analytics thus frequently utilizes modeling, machine learning and data mining techniques from the area of *predictive analytics*. In fact, predictive analytics can be instrumented for “*generating new theory, developing new measures, comparing competing theories, improving existing theories, assessing the relevance of theories, and assessing the predictability of empirical phenomena*” (Shmueli and Koppius, 2011).

Predictive analytics frequently contributes to managerial decision support, as is the case when predicting investor reaction to press releases and financial disclosures (Nassirtoussi, Aghabozorgi, Wah, and Ngo, 2014). In this instance, predictive analytics is typically confronted with massive datasets of heterogeneous and mostly textual content, while outcomes are simultaneously of high impact for any business. Until

now, decision support for financial news still predominantly relies on traditional machine learning techniques, such as support vector machines or decision trees (Minev, Schommer, and Grammatikos, 2012; Nassirtoussi, Aghabozorgi, Wah, and Ngo, 2014; Pang and Lee, 2008).

The performance of traditional machine learning algorithms largely depends on the features extracted from underlying data sources, which has consequently elicited the development and evaluation of feature engineering techniques (Arel, Rose, and Karnowski, 2010). Research efforts to automate and optimize the feature engineering process, along with a growing awareness of current neurological research, has led to the emergence of a new sub-field of machine learning research called *deep learning* (Hinton and Salakhutdinov, 2006). Deep learning takes into account recent knowledge on the way the human brain processes data and thus enhances traditional neural networks by a series of hidden layers. This series of hidden layers allows for deeper knowledge representation, possibly resulting in improved predictive performance. Deep learning methods have been applied to well-known challenges in the machine learning discipline, such as pattern recognition and natural language processing. The corresponding results indicate that deep learning can outperform classical machine learning methods (which embody only a shallow knowledge layer) in terms of accuracy (c. f. Hinton and Salakhutdinov, 2006; Socher, Pennington, Huang, Ng, and Manning, 2011).

In this paper, we want to unleash the predictive power of deep learning to provide decision-support in the financial domain. As a common challenge, we choose the task of predicting stock market movements that follow the release of a financial disclosure. We expect that deep learning can learn appropriate features from the underlying textual corpus efficiently and thus surpass other state-of-the-art classifiers. However, the successful application of deep learning techniques is not an easy task; deep learning implicitly performs feature extraction through the interplay of different hidden layers, the representation of the textual input and the interactions between layers. In order to master this challenge, we apply the recursive autoencoder model introduced by Socher, Pennington, Huang, Ng, and Manning (2011) and tailor it to the prediction of stock price directions based on the content of financial materials.

The remainder of this paper is structured as follows. Section 2 provides a short overview of related work in which we discuss similar text mining approaches and give an overview of relevant deep learning publications. We then explain our methodology and highlight the differences between classical and deep learning approaches (Section 3). Finally, Section 4 evaluates both approaches using financial news disclosures and discuss the managerial implications.

## 2 Related Work

This Information Systems research is positioned at the intersection between finance, Big Data, decision support and predictive analytics. The first part of this section discusses traditional approaches of providing decision support based on financial news. In the second part, we discuss previous work that focuses on the novel deep learning approach.

### 2.1 Decision analytics for financial news

Text mining of financial disclosures represents one of the fundamental approaches for decision analytics in the finance domain. The available work can be categorized by the necessary *preprocessing* steps, the *text mining algorithms*, the underlying *text source* (e. g. press releases, financial news, tweets) and its *focus* on facts or opinions (e. g. quarterly reports, analyst recommendations). While Pang and Lee (2008) provide a comprehensive domain-independent survey, other overviews concentrate solely on the financial domain (Minev, Schommer, and Grammatikos, 2012; Mittermayer and Knolmayer, 2006). In a very recent survey, Nassirtoussi, Aghabozorgi, Wah, and Ngo (2014) focus specifically on studies aimed at stock market prediction. We structure the discussion of the related research according to the above categories.

Among the most popular *text mining algorithms* are classical machine learning algorithms, such as support vector machines, regression algorithms, decision trees and Naïve Bayes. In addition, neural

network models have been used more rarely, but are slowly gaining traction, just as in other application domains (Nassirtoussi, Aghabozorgi, Wah, and Ngo, 2014). Furthermore, Bayesian learning can provide explanatory insights by generating domain-dependent dictionaries (Pröllochs, Feuerriegel, and Neumann, 2015).

As part of *preprocessing*, the first step in most text mining approaches is the generation of a set of values that represent relevant textual features, which can be used as inputs for the subsequent mining algorithms. This usually involves the selection of features based on the raw text sources, some kind of dimensionality reduction and the generation of a good feature representation, such as binary vectors. A comprehensive discussion of the various techniques used for feature engineering can be found in Nassirtoussi, Aghabozorgi, Wah, and Ngo (2014) and Pang and Lee (2008).

The *text sources* used for text mining include financial news (e. g. Alfano, Feuerriegel, and Neumann, 2015; Feuerriegel, Wolff, and Neumann, 2015, 2016; Feuerriegel and Neumann, 2013) and company-specific disclosures, and range from the less formal, such as tweets (e. g. Bollen, Mao, and Zeng, 2011), to more formal texts, such as corporate filings (e. g. Feuerriegel, Ratku, and Neumann, 2016; Muntermann and Guettler, 2007; Pröllochs, Feuerriegel, and Neumann, 2015). Some researchers have focused exclusively on the headlines of news sources to exclude the noise usually contained in longer texts (Peramunetilleke and Wong, 2002).

News disclosures with a fact-based *focus* are especially relevant for investors. As such, German ad hoc announcements in English contain strictly regulated content and a tight logical connection to the stock price, making them an intriguing application in research. The measurable effect of ad hoc news on abnormal returns on the day of an announcement have been established by several authors (c. f. Groth and Muntermann, 2011; Hagenau, Liebmann, and Neumann, 2013; Muntermann and Guettler, 2007; Pröllochs, Feuerriegel, and Neumann, 2015). Consequently, we utilize the same corpus in our following evaluation.

## 2.2 Deep learning as an emerging trend

Deep learning<sup>1</sup> originally focused on complex tasks, in which datasets are usually high-dimensional (Arel, Rose, and Karnowski, 2010; Bengio, 2009). As such, one of the first successful deep learning architectures consisted of a so-called autoencoder in combination with a Boltzmann machine, where the autoencoder carries out unsupervised pre-training of the network weights (Hinton and Salakhutdinov, 2006). This model performs well on several benchmarks from the machine learning literature, such as image recognition. Moreover, its architecture can be adapted to enhance momentum stock trading (Takeuchi and Yu Ying, 2013) as one of the few successful applications of deep learning to financial decision support. However, this publication relies only on past stock returns and neglects the predictive power of exogenous input, such as financial disclosures.

The natural language processing community has only recently started to adapt deep learning principles to the specific requirements of language recognition tasks. For example, Socher, Pennington, Huang, Ng, and Manning (2011) utilize a recursive autoencoder to predict sentiment labels based on individual movie review sentences. Further research improved the results on the same dataset by combining a recursive neural tensor network with a sentiment treebank (Socher et al., 2013).

## 3 Methodologies for Financial Decision Support

This section introduces our research framework to provide financial decision support based on news disclosures. In brevity, we introduce a benchmark classifier and our deep learning architecture to predict stock movements. Altogether, Figure 1 illustrates how we compare both prediction algorithms. The random forest and the recursive autoencoder are both trained to predict stock market directions based

<sup>1</sup> For details, see reading list “*Deep Learning*”. Retrieved April 21, 2015 from <http://deeplearning.net/reading-list/>.

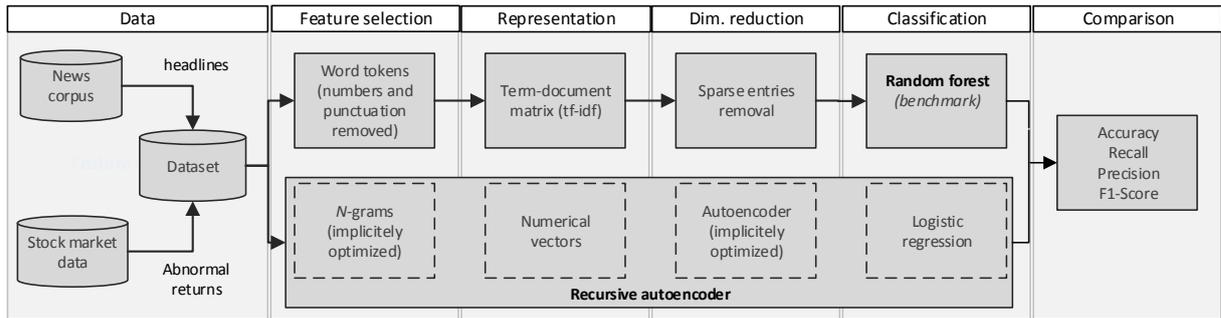


Figure 1. Research framework comparing classical machine learning and deep learning.

on the ad hoc announcements and the according abnormal returns. To compare the performance of the recursive autoencoder to the benchmark, we apply the same test set to each of the trained algorithms and measure the predictive performance in terms of accuracy, precision and recall based on the confusion matrix.

Following Nassirtoussi, Aghabozorgi, Wah, and Ngo (2014), we divide the overall procedure into steps for data generation, feature selection, feature reduction and feature representation. Both approaches, the benchmark algorithm and the recursive autoencoder, differ fundamentally in their preprocessing. The application of a random forest or support vector machine requires traditional feature engineering, whereas the recursive autoencoder, as a remedy, automatically generates a feature representation as part of its optimization algorithm. This is indicated in Figure 1 by the extension of the recursive autoencoder box over all preprocessing steps.

### 3.1 Benchmark: predicting stock movements with random forests

In a first step, one transforms the running text into a matrix representation, which subsequently works as the input to the actual random forest algorithm. First of all, we remove numbers, punctuations and stop words from the running text and then split it into tokens (Manning and Schütze, 1999). Afterwards, we count the frequencies of how often terms occur in each news disclosure, remove sparse entries to reduce the dimensionality and store these values in a document-term matrix. The document-term matrix then represents the features. The actual values are weighted (Salton, Fox, and Wu, 1983) by the *term frequency-inverse document frequency* (tf-idf). This is a common approach in information retrieval to adjust the word frequencies by their importance.

In the following evaluation, we utilize random forests as a benchmark classifier. Random forests represent one of the most popular machine learning algorithms due to their favorable predictive accuracy, relatively low computational requirements and robustness (Breiman, 2001; Hastie, Tibshirani, and Friedman, 2009; Kuhn and Johnson, 2013). Random forests are an ensemble learning method for classification and regression, which is based on the construction and combination of many de-correlated decision trees.

Given a training set  $X = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$  with associated responses  $Y = \{y_1, \dots, y_N\}$ , the algorithm repeats the following steps  $B$  times (we choose  $B = 500$ ): (1) sample with replacement from  $X$  and  $Y$  to generate new subsets  $X'$  and  $Y'$ . (2) Train a decision tree  $t_b$  using  $X'$  and  $Y'$ . The individual decision trees  $t_b$ ,  $b = 1, \dots, B$  can be combined to predict a response  $\hat{y}$  for unseen values  $\mathbf{x}$  as follows. One calculates individual predictions  $t_b(\mathbf{x})$  for  $b = 1, \dots, B$  from each tree and then aggregates these predictions by simply using the majority vote to get the final response.

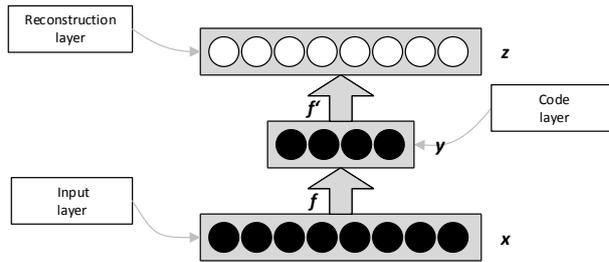


Figure 2. An autoencoder searches a mapping between input  $\mathbf{x}$  and a lower-dimensional representation  $\mathbf{y}$  such that the reconstructed value  $\mathbf{z}$  is similar to the input  $\mathbf{x}$ .

### 3.2 Deep learning architecture: recursive autoencoders

This section describes the underlying architecture of our deep learning approach for financial disclosures based on so-called *autoencoders*. The architecture of an autoencoder is illustrated in Figure 2. An autoencoder is basically an artificial neural network, which finds a lower-dimensional representation of input values. Let  $\mathbf{x} \in [0, 1]^N$  denote our input vector, for which we seek a lower-dimensional representation  $\mathbf{y} \in [0, 1]^M$  with  $M < N$ . The mapping  $f$  between  $\mathbf{x}$  and  $\mathbf{y}$  is named *encoding* function and can be, generally speaking, any non-linear function, although a common choice is the sigmoid function

$$f(\mathbf{x}) = \sigma(W\mathbf{x} + \mathbf{b}) = \frac{1}{1 + \exp(-W\mathbf{x} - \mathbf{b})} = \mathbf{y} \quad \text{with parameters } W \text{ and } \mathbf{b}. \quad (1)$$

The key idea of an autoencoder is to find a second mapping from  $\mathbf{y}$  to  $\mathbf{z} \in [0, 1]^N$  given by  $f'(\mathbf{y}) = \sigma(W'\mathbf{y} + \mathbf{b}')$ , such that  $\mathbf{z}$  is almost equal to the input  $\mathbf{x}$ . Mathematically speaking, we choose the free parameters in  $f$  and  $f'$  by minimizing the difference between the original input vector  $\mathbf{x}$  and the reconstructed vector  $\mathbf{z}$ . This can be effectively achieved using numerical optimization, such as gradient descent, in order to determine the weights  $W$  and  $W'$ . Altogether, the representation  $\mathbf{y}$  (often called *code*) is a lower-dimensional representation of the input data; it is frequently used as input features for subsequent learning because it only contains the most relevant or most discriminating features of the input space.

The classical autoencoder works merely with a simple vector as input. In order to incorporate contextual information, we extend the classical autoencoder, resulting in a so-called *recursive autoencoder*. Here, one trains a sequence of autoencoders, where each not only takes a vector  $\mathbf{x}$  as input but also recursively the lower-dimensional code of the previous autoencoder in the sequence. Let us demonstrate this approach with an example as illustrated in Figure 3. We process input in the form of a sequence of words. Each word is e. g. given by a binary vector with zeros except for a single entry with 1 representing the current word. Then, we train the first autoencoder with the input from the first two words *Company* and *Ltd*. Its lower-dimensional code is then input to the second autoencoder together with the vector representation of the word *placing*. This recursion proceeds up to the final autoencoder, which produces as output the code representation for the complete sentence. Hence, this recursive approach aims to generate a compact code representation of a complete sentence while incorporating contextual information in the code layer. More precisely, this approach can learn from an ordered sequence of words and not only the pure frequencies. In addition, the recursive autoencoder entails an intriguing advantage: it can compress large input vectors in an unsupervised fashion without the need for class labels.

The basic steps for the recursive autoencoder are as follows: in a first step, the individual words of an input sentence (1) are mapped onto vectors of equal length  $l$  (2). We initialize the values of the weights by sampling from a Gaussian distribution and later continuously updated through backpropagation. Through a recursive application of the autoencoder algorithm (3), the complete input sentence is then compressed bit by bit into a single code representation of length  $l$  (4). For this purpose, the first autoencoder generates a code representation of length  $l$  from the vectors representing the first two words in the sentence. The

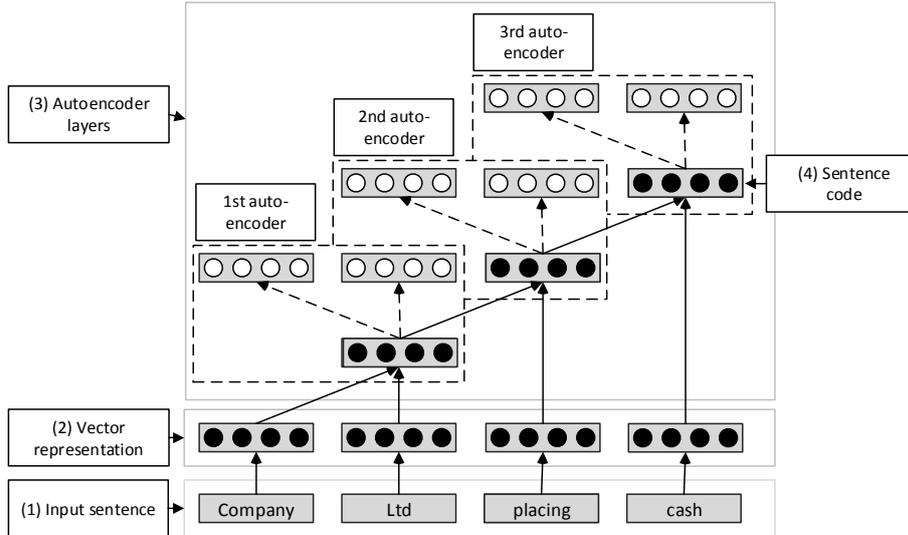


Figure 3. A recursive autoencoder is a sequence of autoencoders, where each not only takes a vector as input but also the lower-dimensional code of the previous autoencoder in the sequence.

second autoencoder takes this code representation and the third word vector as inputs, and calculates the code representation of the next level.

In order to extract and predict sentiment values, we use an extended variant of the recursive autoencoder model (Socher, Pennington, Huang, Ng, and Manning, 2011), which includes an additional *softmax layer* (sometimes also referred to as multinomial logit) in each autoencoder. This softmax function estimates the probability that an input vector  $\mathbf{x}$  belongs to a certain class  $j \in K$  via

$$P(y = j | \mathbf{x}) = \frac{\exp(\mathbf{x}^T W_j)}{\sum_{k=1}^K \exp(\mathbf{x}^T W_k)}. \quad (2)$$

In order to train this model, we optimize the weights  $W_k$  of both the autoencoders and the softmax layers simultaneously with a combined target function. We then utilize the trained weights to classify unknown sentences by first computing the code representation inside the recursive autoencoder and, second, calculating the probabilities for each class from the softmax function. Interestingly, the backward mapping  $f'$  is needed for training but is no longer needed for the prediction (that is why black circles indicate the vectors only necessary for prediction in Figure 2 and Figure 3).

## 4 Evaluation: Predicting Stock Market Direction from Financial News

In this section, we discuss and evaluate our experimental setting for predicting the direction of stock market movements following financial disclosures. We start with describing the steps involved in the generation of the underlying dataset and then compare classical machine learning with our deep learning architecture.

### 4.1 Dataset

Our news corpus originates from regulated ad hoc announcements<sup>2</sup> between January 2004 and the end of June 2011 in English. These announcements conform to German regulations that require each

<sup>2</sup> Kindly provided by Deutsche Gesellschaft für Ad-Hoc-Publizität (DGAP).

listed company in Germany to immediately publish any information with a potentially significant effect on the stock price. With their direct relation to a particular stock, the tightly controlled content and the measurable effect on the stock price on the day of the announcement, ad hoc announcements are particularly well-suited for the development and evaluation of techniques for predictive analytics.

Since recursive autoencoders work on sentence tokens, we exclusively use the headlines of English ad hoc announcements for the prediction and discard the message body. As previous work (e. g. Peramunetilleke and Wong, 2002) has shown, this is not a major disadvantage and can even help in reducing noise, as long as the titles concisely represent the content of the text.

We gather the financial data of the releasing companies from Thomson Reuters Datastream. We retrieve the firm performance with the help of the International Securities Identification Numbers (ISIN) that appear first in each of the ad hoc announcements. The stock price data before and on the day of the announcement are extracted using the corresponding trading day. These are then used to calculate abnormal returns (Konchitchki and O’Leary, 2011; MacKinlay, 1997; Pröllochs, Feuerriegel, and Neumann, 2015); abnormal returns can be regarded as some kind of excess return caused by the news release. In addition, we remove penny stocks with stock prices below \$5 for noise reduction. We then label each announcement title with one of three return direction classes (*up*, *down* or *steady*), according to the abnormal return of the corresponding stock on the announcement day and discard the steady samples for noise reduction.

The resulting dataset consists of 8359 headlines from ad hoc announcements with corresponding class labels *up* or *down*. Of this complete dataset, we use the samples covering the first 80 % of the timeline as training samples and the remaining 20 % as test samples.<sup>3</sup>

## 4.2 Preliminary results

We can now apply the above methods for predictive analytics to provide decision support regarding how investors react upon textual news disclosures. By comparing random forests and recursive autoencoders, we can evaluate our hypothesis that deep learning outperforms our benchmark in the current setting.

The detailed results are listed in Table 1. We compare the predictive performance on the out-of-sample test set in terms of accuracy, precision, recall and the F1-score. The random forest as our benchmark achieves an in-sample accuracy of 0.63 and an out-of-sample accuracy of 0.53 at best. In comparison, the recursive autoencoder<sup>4</sup> as our deep learning architecture results in an accuracy of 0.56. This accounts for a relative improvement of 5.66 %. Similarly, the F1-score increases from 0.52 to 0.56 – a substantial rise of 7.69 %. The higher accuracy, as well as the improved F1-score, of the recursive autoencoder underlines our initial assumption that deep learning algorithms can outperform our benchmark from classical machine learning. When comparing the necessary computational resources, we see that recursive autoencoders ( $\approx 23$  min) require less computation time than decision trees (more than 200 min).<sup>5</sup> Moreover, recursive autoencoders have an additional advantage: one can simply inject the complete set of news headlines as input without the manual effort of feature engineering. The reason for this is that the calculation and optimization of a feature representation is integrated into the optimization routines of deep learning algorithms.

The above results comply with the reported figures of around 60 % with the full message body from related work (Groth and Muntermann, 2008; Hagenau, Liebmann, and Neumann, 2013). In direct comparison to the benchmarks, our evaluation provides evidence that deep learning is a compelling approach for the prediction of stock price movements.

<sup>3</sup> We avoid the use of  $k$ -fold cross-validation as this would neglect the timing of disclosures. For instance, we would evaluate our models with disclosures during the financial crisis, while the same models were previously trained with later knowledge of how news were perceived after the happening of the financial crisis.

<sup>4</sup> We systematically tried several combinations for the two adjustable parameters embedding size (the length  $l$  of the mapped feature vectors) and number-of-iterations (i. e. number of gradient descent iterations). The best result accounts for an accuracy of 0.56 on the test-set, with a vector embedding size of 40 and 70 iterations. As expected, increasing the number of iterations usually results in better accuracy on the training set and lower accuracy on the test set – a typical indication of over-fitting.

<sup>5</sup> Timings measured on an Intel Core i7-4700MQ CPU running at 2.4 GHz with 8 GB RAM and 64-bit Windows 8.1. Please note that programming languages and matrix algebra libraries vary which makes a fair comparison difficult.

Predictive Analytics Method	Accuracy	Precision	Recall	F1-Score
Random Forest	0.53	0.53	0.51	0.52
Recursive Autoencoder	0.56	0.56	0.56	0.56
<i>Relative Improvement</i>	5.66 %	5.66 %	9.80 %	7.69 %

Table 1. Preliminary results evaluating improvements by utilizing deep learning to predict the direction of stock price movements following financial disclosures.

### 4.3 Discussion and implications for practitioners

Traditional machine learning techniques still represent the default method of choice in predictive analytics. However, recent research indicates that these methods insufficiently capture the properties of complex, non-linear problems. Accordingly, the experiments in this paper show that a deep learning algorithm is capable of implicitly generating a favorable knowledge representation.

As a recommendation to practitioners, better results are achievable with deep learning than with classical methods that rely on explicitly generated features. Nevertheless, practitioners must be aware of the complex architecture of deep learning models. This requires both a thorough understanding and solid experience in order to use such models efficiently.

The economic impact of our improvement is manifold. A higher predictive performance enables business opportunities for automated traders. In addition, corporates can utilize this approach to assess the expected market response subsequent to disclosures. This works as a safety mechanism to check if the subjective perception of investors matches the content of a release.

## 5 Conclusion and Research Outlook

In the present paper, we show how a novel deep learning algorithm can be applied to provide decision support for the financial domain. Thereby, we contribute to Information Systems research by shedding light on how to exploit deep learning as a recent trend for managerial decision support. We demonstrate that a recursive autoencoder outperforms our benchmark from traditional machine learning in the prediction of stock market movements following financial news disclosures. The recursive autoencoder benefits from being able to automatically generate a deep knowledge representation.

In future research, we intend to broaden the preliminary results of this Research-in-Progress paper. First, our analysis could benefit substantially from incorporating and comparing further algorithms from predictive analytics. Second, we want to generalize our results by including further news sources.

## References

- Alfano, S. J., S. Feuerriegel, and D. Neumann (2015). "Is News Sentiment More than Just Noise?" In: *23rd European Conference on Information Systems (ECIS 2015)*.
- Apte, C., B. Liu, E. P. D. Pednault, and P. Smyth (2002). "Business Applications of Data Mining." *Communications of the ACM* 45 (8), 49–53.
- Arel, I., D. C. Rose, and T. P. Karnowski (2010). "Deep Machine Learning: A New Frontier in Artificial Intelligence Research." *IEEE Computational Intelligence Magazine* 5 (4), 13–18. URL: [http://web.eecs.utk.edu/~itamar/Papers/DML\\_Arel\\_2010.pdf](http://web.eecs.utk.edu/~itamar/Papers/DML_Arel_2010.pdf) (visited on 04/19/2015).
- Arnott, D. and G. Pervan (2005). "A Critical Analysis of Decision Support Systems Research." *Journal of Information Technology* 20 (2), 67–87.
- Asadi Someh, I. and G. Shanks (2015). "How Business Analytics Systems Provide Benefits and Contribute to Firm Performance?" In: *23rd European Conference on Information Systems (ECIS 2015)*.

- Bengio, Y. (2009). *Learning Deep Architectures for AI*. URL: <http://www.iro.umontreal.ca/~lisa/pointeurs/TR1312.pdf> (visited on 08/10/2014).
- Bollen, J., H. Mao, and X. Zeng (2011). "Twitter Mood Predicts the Stock Market." *Journal of Computational Science* 2 (1), 1–8.
- Boyd, D. and K. Crawford (2012). "Critical Questions for Big Data." *Information, Communication & Society* 15 (5), 662–679.
- Boylan, J. E. and A. A. Syntetos (2012). "Forecasting in Management Science." *Omega* 40 (6), 681.
- Breiman, L. (2001). "Random Forests." *Machine Learning* 45 (1), 5–32.
- Chen, H., R. H. L. Chiang, and V. C. Storey (2012). "Business Intelligence and Analytics: From Big Data to Big Impact." *MIS Quarterly* 36 (4), 1165–1188.
- Davenport, T. H. (2006). "Competing on Analytics." *Harvard Business Review* 134 (1), 98–107.
- Feuerriegel, S, A Ratku, and D Neumann (2016). "Analysis of How Underlying Topics in Financial News Affect Stock Prices Using Latent Dirichlet Allocation." In: *Proceedings of the 49th Hawaii International Conference on System Sciences (HICSS 2016)*. IEEE Computer Society.
- Feuerriegel, S, G Wolff, and D Neumann (2015). "Information Processing of Foreign Exchange News: Extending the Overshooting Model to Include Qualitative Information from News Sentiment." In: *Proceedings of the International Conference on Information Systems (ICIS 2015)*. Association for Information Systems.
- (2016). "News Sentiment and Overshooting of Exchange Rates." *Applied Economics* (forthcoming).
- Feuerriegel, S. and D. Neumann (2013). "News or Noise? How News Drives Commodity Prices." In: *Proceedings of the International Conference on Information Systems (ICIS 2013)*. Association for Information Systems.
- Groth, S. and J. Muntermann (2008). "A Text Mining Approach to Support Intraday Financial Decision-Making." In: *Americas Conference on Information Systems (AMCIS 2008)*.
- Groth, S. S. and J. Muntermann (2011). "An Intraday Market Risk Management Approach Based on Textual Analysis." *Decision Support Systems* 50 (4), 680–691.
- Hagenau, M., M. Liebmann, and D. Neumann (2013). "Automated News Reading: Stock Price Prediction based on Financial News Using Context-Capturing Features." *Decision Support Systems* 55 (3), 685–697.
- Halper, F. (2011). *The Top 5 Trends in Predictive Analytics*. URL: [http://www.information-management.com/issues/21\\_6/the-top-5-trends-in-redictive-an-alytics-10021460-1.html](http://www.information-management.com/issues/21_6/the-top-5-trends-in-redictive-an-alytics-10021460-1.html) (visited on 08/20/2014).
- Hastie, T., R. Tibshirani, and J. H. Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd ed. Springer Series in Statistics. New York: Springer.
- Hinton, G. E. and R. R. Salakhutdinov (2006). "Reducing the Dimensionality of Data with Neural Networks." *Science* 313 (5786), 504–507.
- IBM (2013). *The Four V's of Big Data*. URL: <http://www.ibmbigdatahub.com/infographic/four-vs-big-data> (visited on 04/21/2014).
- Konchitchki, Y. and D. E. O'Leary (2011). "Event Study Methodologies in Information Systems Research." *International Journal of Accounting Information Systems* 12 (2), 99–115.
- Kuhn, M. and K. Johnson (2013). *Applied Predictive Modeling*. New York, NY: Springer.
- MacKinlay, A. C. (1997). "Event Studies in Economics and Finance." *Journal of Economic Literature* 35 (1), 13–39.
- Manning, C. D. and H. Schütze (1999). *Foundations of Statistical Natural Language Processing*. 6th Edition. Cambridge, MA: MIT Press.
- Minev, M., C. Schommer, and T. Grammatikos (2012). "News and Stock Markets: A Survey on Abnormal Returns and Prediction Models." URL: <http://orbilu.uni.lu/bitstream/10993/14176/1/TR.Survey.News.Analytics.pdf> (visited on 04/19/2015).
- Mittermayer, M.-A. and G. F. Knolmayer (2006). "Text Mining Systems for Market Response to News: A Survey: Working Paper." *SSRN Electronic Journal*.

- Muntermann, J. and A. Guettler (2007). "Intraday Stock Price Effects of Ad Hoc Disclosures: The German Case." *Journal of International Financial Markets, Institutions and Money* 17 (1), 1–24.
- Nassirtoussi, A. K., S. Aghabozorgi, T. Y. Wah, and D. C. L. Ngo (2014). "Text Mining for Market Prediction: A Systematic Review." *Expert Systems with Applications* 41 (16), 7653–7670.
- Pang, B. and L. Lee (2008). "Opinion Mining and Sentiment Analysis." *Foundations and Trends in Information Retrieval* (2), 1–135. (Visited on 04/19/2015).
- Peramunetilleke, D. and R. Wong (2002). "Currency Exchange Rate Forecasting from News Headlines." *Australian Computer Science Communications* 24 (2), 131–139.
- Power, D. J. (2014). "Using 'Big Data' for Analytics and Decision Support." *Journal of Decision Systems* 23 (2), 222–228.
- Pröllochs, N., S. Feuerriegel, and D. Neumann (2015). "Generating Domain-Specific Dictionaries Using Bayesian Learning." In: *23rd European Conference on Information Systems (ECIS 2015)*.
- Salton, G., E. A. Fox, and H. Wu (1983). "Extended Boolean Information Retrieval." *Communications of the ACM* 26 (11), 1022–1036.
- Shmueli, G. and O. Koppius (2011). "Predictive Analytics in Information Systems Research." *MIS Quarterly* 35 (3), 553–572.
- Socher, R., J. Pennington, E. Huang, A. Ng, and C. Manning (2011). "Semisupervised Recursive Autoencoder." In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2011)*, pp. 151–161.
- Socher, R. et al. (2013). "Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank." In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*. Vol. 1631.
- Takeuchi, L. and L. Yu Ying (2013). *Applying Deep Learning to Enhance Momentum Learning Trading Strategies in Stocks*. URL: <http://cs229.stanford.edu/proj2013/TakeuchiLee-ApplyingDeepLearningT.pdf> (visited on 05/03/2015).
- Turban, E. (2011). *Business Intelligence: A Managerial Approach*. 2nd Edition. Boston, MA: Prentice Hall.
- Vizecky, K. (2011). "Data Mining meets Decision Making: A Case Study Perspective." In: *Americas Conference on Information Systems (AMCIS 2011)*, Paper 453.