

Automatic electronic invoice classification using machine learning models

Chiara Bardelli, Alessandro Rondinelli, Ruggero Vecchio, Silvia Figini

Abstract Electronic invoicing has become mandatory for Italian companies since January 2019. Invoices are structured in a predefined xml template where the information reported can be easily extracted and analyzed. The main aim of this paper is to exploit the information structured in electronic invoices to build an intelligent system which can facilitate accountants work. More precisely, this contribution shows how it is possible to automate part of the accounting process: all sent or received invoices of a company are classified into specific codes which represent the economic nature of the the financial transactions. In order to classify data contained in the invoices a machine learning multiclass classification problem is proposed using as input variables the information of the invoices to predict two different target variables, account codes and the VAT codes, which composes a general ledger entry. Different approaches are compared in terms of prediction accuracy. The best performance is achieved considering the hierarchical structure of the account codes.

Key words: multiclass classification; text mining; accounting control system

Chiara Bardelli

PhD Program in Computational Mathematics and Decision Sciences, University of Pavia, Italy;
e-mail: chiara.bardelli01@universitadipavia.it

Alessandro Rondinelli, Ruggero Vecchio

Datev.it S.p.a.;

e-mail: alessandro.rondinelli@datev.it; ruggero.vecchio@datev.it

Silvia Figini

Department of Political and Social Sciences, University of Pavia, Italy;

e-mail: silvia.figini@unipv.it

1 Introduction

In the digitization era Artificial Intelligence systems can bring many opportunities in industrial and information areas to improve the efficiency and to deliver more value to business. On the basis of the results obtained in the research conducted by [6], the accountancy profession is one of the job which will be highly automated and digitized in the near future. Even if thanks to computer systems, the workload of accountants has been reduced [21], repetitive and monotonous routine tasks are still relevant part of the accountants work. For this purpose, Artificial Intelligence can be used to replace repetitive human activities with highly automated systems giving the possibility for the accountants to focus on more stimulating and motivating activities like decision-making, problem solving, advising and business strategy development([7], [10]).

One of the tasks of the accountants work which can be easily automated is the creation of the general journal which includes each financial transaction made by a business company. Among all the records of the financial transactions, we find those related to supplier and customer invoices. When an invoice is received or issued by a company, the bookkeeper or accountant creates the related accounting journal entries which describe the economic nature of the transactions using a specific list of codes (i.e. the Chart of Accounts and the codes related to VAT rules). This task can be treated using machine learning models: input variables include information extracted from the invoices and the characteristics of the companies, while the codes used for the creation of the accounting journal entry can be considered as the target variables. Our methodological approach considers three different steps in the development of an automated system which could propose possible codes for the journal entry: collection of invoices, processing of the textual content of the invoices and automatic classification of invoices to build the associated journal entry based on the training of a machine learning model.

Starting from January 2019 all Italian companies are required to use electronic invoices. They are composed of a structured and fixed xml template provided by "Agenzia delle Entrate", the Italian governmental agency which operates to ensure an high level of tax compliance. This allows to solve the issue of extracting information from documents which, thanks to electronic invoicing, follow a regular structure shared by all the Italian companies. As a consequence, the acquisition of the data, contained in the invoice, is handled more easily representing a strong encouragement to focus the research on the automation of the classification of transactions for the creation of the general journal.

This paper introduces a machine learning model to predict the elements of a bookkeeping entry starting from the content of the electronic invoice. Each line of the invoice is represented by a numeric vector that combines textual description, other information related to the line and the characteristics of the companies involved in the invoice. We test performances of our classifiers on two anonymized real-world accounting datasets which belong to two different Chartered Accounting firms. Data are collected and provided by Datev.it Spa, a company which develops software for professional accounting.

Our methodological approach considers: (i) the reconstruction of the training set, (ii) the definition of a complex structure of the account codes which are organized in a hierarchical taxonomy based on 5 different levels for a total of more than 200 different labels leading to an imbalanced classification problem and (iii) the strong heterogeneity in rules and methodologies used by different Chartered Accounting firms in the creation of the journal entries.

The paper is organized as follows: Section 2 reports the literature concerning the application of machine learning techniques to accounting systems; Section 3 defines the problem and the data available; Section 4 introduces the methods applied to reconstruct the training set and to solve our predictive problem; Section 5 shows the application and the most relevant results are displayed. Section 6 contains conclusions and further ideas of research.

2 Literature review

In accounting systems machine learning techniques have been recently applied for different purposes trying to improve the efficiency and the precision of monotonous and repetitive tasks.

When invoices are not required to fill a fixed template, one of the most laborious step of the bookkeeping process in terms of working hours is the extraction of structured information from the printed documents. Classical deterministic OCR (Optical Character Recognition) techniques are usually applied to this problem. Early approaches exploit the structure of documents to extract fields based on their position or detected forms ([20],[5]). However, in order to generalize extractions patterns with a flexible approach which can be applied to any template, some attempts have been made recently to combine deterministic OCR methods to machine learning models in presence of large quantities of historical data. In [9] and [22] deep neural networks have been trained on a large dataset to extrapolate specific fields from the images of the invoices. In [14] the density of the black pixels and the density of image edges are used as input variables to train a classifier which predict the class of the invoice. In both works the need of a large historical dataset is highlighted as a possible drawback for these techniques.

Another aspect of interest which has been tackled with machine learning models is the anomaly detection problem in the accounting journal. With the increasing complexity of business processes and the growing amount of structured accounting data, identifying erroneous or fraudulent business transactions (and corresponding journal entries) represents a critical challenge for accountants and auditors. In [16] and [23] a novel unsupervised approach based on neural networks is proposed to detect anomalous journal entries and potentially fraudulent activities in large scaled accounting data. In [17] a similar architecture is applied also to a small subset of data to investigate if accurate results can be obtained also in case of more limited information. These methods, in contrast with the classical approach based on the experience of chartered accountants like static red-flag tests, are able to detect novel

schemes of fraud based on historical data. Therefore, the opportunity to develop an automated and high precision detection of accounting anomalies can save work time.

As last area of application of machine learning models to accounting system there are at least two works ([1],[2]) which try to develop a classifier able to predict the details of an accounting journal entry, in particular they focus on the prediction of the account codes. The performances of the data driven algorithms are compared to a deterministic approach based on a rule induction system built on the experience of accountants. In both works the difference in accuracy between the machine learning classifier and the rule induction classifier can not be considered statistically significant, concluding that there is potential for machine learning in this area but still does not outperform the existing deterministic implementation. A possible extension could be to combine the results of rule induction system to the predictions of the machine learning classifier. In this work we extend the research of ([1],[2]) to the prediction of the entire journal entry, considering as target variable not only the account codes but also the VAT codes.

3 Dataset description

The accounting general journal is composed of a list of journal entries which collects financial transactions of a business firm and classifies them into specific codes. A consistent part of the journal entries derives from the recording of the customer and supplier invoices of a company. Thanks to the electronic invoicing introduced recently in Italy, the extraction of the relevant fields from the xml template can be easily handled.

Given all the information about a single line of an invoice and the characteristics of the companies which are involved in the invoice, the aim is to construct the journal entry related to a specific line of the invoice, thus predicting the account codes and the VAT codes. Multiple lines of an invoice can be associated to the same journal entry with the same codes. For the sake of simplicity, in this study, we assume that lines inside an invoice are independent one from each other (ignoring the fact that lines which belong to the same invoice are all influenced by the context of the invoice itself).

In our classification task, we construct the prediction rule given the training sample $\left\{ (y_i^{(1)}, y_i^{(2)}, \mathbf{x}_i) \right\}_{i=1}^n$ where:

- $y_i^{(1)}$, $i = 1, \dots, n$, are categorical observations which represent the account codes associated to the i -th line of the invoice. The account codes belong to the Chart of Accounts which has a particular structure: codes are organized in a hierarchical structure and only the accounts which are the leaves of the tree are used as tag in the journal entry. In our problem we consider only the accounts in the leaves of the hierarchical tree;

- $y_i^{(2)}$, $i = 1, \dots, n$ represent the VAT codes to predict associated to the i -th line of the invoice. This target variable is composed of two different sub-codes: one related to the tax rate applied to the line of the invoice, the other one related to tax rule. In our problem, this two codes are considered as a unique variable to predict;
- \mathbf{x}_i is the vector of predictors related to the content of the invoice and the characteristics of subjects involved.

The data set considered in this study is the combination of two data sources:

- Customers and suppliers xml invoices of different companies
- Accounting journal entry related to the recording process of xml invoices (account codes and VAT codes)

The match of these two different sources is possible only at document level: the general journal records can be directly associated only to their original invoice. On the other hand, it is not possible to recreate directly the link between a single line of an invoice and the related journal entry. This problem of data reconstruction can be addressed exploiting the information about the amounts which are included both in the xml invoices (detailed amounts) and in the accounting journal entry (aggregated amounts). Starting from this point, it is possible to translate the problem at hand into a combinatorial optimization task as the knapsack representation with equality constraints as explained in Section 4.1.

In our machine learning algorithm we consider features related to the content of the invoice and characteristics of the companies. In particular, the information considered are:

- textual description of the line of the invoice;
- codes associated to the line (like tax rate);
- information about activities performed by companies:
 - ATECO code (classification of economic Italian activities) provided by ISTAT (the Italian Statistics Agency);
 - ISA categories based on the level of fiscal reliability;
 - type of supplier: person, Italian firm, European company, extra-European company;
 - type of accounting used by the company: ordinary or simplified;
 - tax regime.

Machine learning classifiers are trained and tested on two anonymized real-world datasets of two different accounting firms which include invoices from January 2019 to March 2020. The first one contains about 32172 electronic invoices which include more than 320000 lines to classify. The second one is composed by 34932 invoices and more than 200000 lines. The number of distinct codes which the algorithm has to predict is shown in Figure 1. The most complex problem, in terms of accuracy, seems to be the one related to the prediction of the account codes for the received invoices since, in this scenario, the number of different categories to predict is very high as reported in Figure 1a.



Fig. 1: Number of distinct codes used for sent and received invoices in the two datasets analyzed

4 Methodological proposal

This section reports methods applied to derive the training set, the pre-processing approaches adopted and the predictive classification algorithm used to predict the account codes and of VAT codes of the journal entry.

4.1 Knapsack problem

The reconstruction data problem can be represented as a multi knapsack problem with equality constraints [13] considering a single invoice at a time.

Let $i = 1, \dots, N$ be the lines of the invoice and $j = 1, \dots, M$ be the index that identifies a single entry of the accounting general journal. The combinatorial optimization problem can be formulated as follows:

$$\max \sum_{j=1}^M \sum_{i=1}^N p_i z_{ij} \quad (1)$$

$$\text{s.t.} \quad \sum_{i=1}^N c_i z_{ij} = b_j, \quad \forall j \quad (2)$$

$$\sum_{j=1}^M z_{ij} = 1, \quad \forall i \quad (3)$$

where c_i is the detailed amount related to the i -th line of the invoice and b_j is the aggregated total of the j -th journal entry. In our setting the vector of weights p_i is

equal to 1, for each i . The value of $z_{i,j}$ is 1 if the i -th line of the invoice is associated to the j -th journal entry and 0 otherwise.

The problem has been resolved through an heuristic which stops when the first feasible solution is found.

First, the heuristic sorts the detailed amounts of invoice lines in a vector in decreasing order and the aggregated amounts in an increasing order vector. Then, the algorithm starts to match first values in the first positions of the two vectors. Notice that lines of invoice with negative amounts have been excluded from the analysis (the algorithm can not converge in presence of negative values).

The algorithm has successfully matched the 61% of the lines of our two initial datasets without any ambiguity. The 11% of the lines has been matched using the first solution proposed by the heuristic. These two sources of data have been used for the training phase of the classifier and the evaluation of its performances. The other 28% of the lines has been excluded from the analysis since the heuristic has not converged because of the high number of lines in the invoice or because of the presence of negative amounts in the invoice.

4.2 Data pre-processing

Most of the information extracted from the invoice, to construct the feature space of our problem, are categorical variables which are included in the model through the one-hot encoding representation. In addition, textual descriptions of the lines of invoices contain helpful information for the creation of the journal entry. In order to process textual data, descriptions have been previously cleaned through standard pre-processing steps [12]: punctuation, numbers, stop words and words of 2 characters have been removed from text. Finally, textual information has been tokenized to create array of words.

In order to transform textual data information into a suitable feature space we compare two different procedures:

- Bag of Words (BoW) approach [11], a simple way to encode the array of words into a binary vector. The main drawback is that the length of the feature space grows linearly with the number of distinct words, leading to infeasible dimensions of the feature space. Different methods can be adopted for the dimensionality reduction [18], in our case we include in the vocabulary words with a frequency higher than 0.1% in the collection of documents.
- Word2Vec algorithm [15], a language modeling technique which maps similar sentences into similar numeric vectors of fixed size.

This two different procedures have been applied on the data at hand before running the classification algorithms.

4.3 Classification algorithm

The aim of this work is to understand if standard rules and logic of the accounting process can be learned from a machine learning classifier which is trained on the real data of an accounting firm. Two different models have been trained separately to predict:

- the account codes related to the economic nature of the transaction;
- the VAT codes related to the tax rates coupled with tax regulations applied to the invoices.

The input variables chosen for the two models are slightly different. The textual data are part of feature space in both cases, but the set of variables which describe the characteristics of the companies changes depending on which target variable we aim to predict.

For each data set two different classification algorithms have been compared: Random Forests (RF) [3] and AdaBoost [8]. The multi-class version of both classifiers has been applied to predict separately account codes and VAT codes.

5 Empirical Analysis

The different nature of sent and received invoices led us to split data in 2 different sub-datasets and analyze them separately to obtain 2 sub-classification problems. To evaluate the performances of all the possible approaches described in Section 4, the datasets of sent and received invoices have been split into training set (80 %) and test set (20 %).

Tables 1 and 2 report accuracy for the different combinations of algorithms applied respectively to sent invoices and received invoices for the prediction of the account codes. As we expected, the algorithms trained on received invoices show, in general, an accuracy rate lower than sent invoices. This is mainly due to the fact that received invoices are more diversified in terms of content and number of possible account codes associated. As far as the pre-processing algorithm for textual data is concerned, Word2Vec obtains better results compared to Bag of Words both when it is used with Random Forests and AdaBoost. A better performance of the Word2Vec can be observed both in results of both the two Chartered Accounting firms. This can be explained by the limited vocabulary used in Bag of Words model: a selection of words has been made following parameters of Section 4.2 in order to reduce the dimensions of the feature space and keep computational costs under control. Fixing the method for preprocessing textual data, performances of Random Forests and AdaBoost are equivalent when applied to sent invoices; on the other hand, in case of received invoices, the first algorithm shows better results both in case of Word2Vec and Bag of Words. From these results it seems that Random Forests is able to better handle classification problems with an higher number of distinct categories to predict.

Table 1: Performances of the classifiers which predict account codes evaluated on test set (sent invoices)

	Dataset 1			Dataset 2		
	Precision	Recall	F1-score	Precision	Recall	F1-score
RF + BoW	96.9 %	97.2 %	97.0 %	98.3 %	98.3 %	98.0 %
RF + Word2Vec	97.1 %	97.3 %	97.1%	98.3 %	98.3 %	97.9 %
AdaBoost + Bow	96.9 %	97.2 %	97.0%	98.3 %	98.3 %	98.0 %
AdaBoost + Word2Vec	97.1 %	97.4 %	97.2 %	98.4 %	98.4 %	98.1 %

Table 2: Performances of the classifiers which predict account codes evaluated on test set (received invoices)

	Dataset 1			Dataset 2		
	Precision	Recall	F1-score	Precision	Recall	F1-score
RF + BoW	83.2 %	84.4 %	83.1%	87.2 %	88.3 %	86.8%
RF + Word2Vec	87.3 %	88.4 %	87.4 %	90.2 %	91.2 %	90.2%
AdaBoost + BoW	41.4 %	30.0 %	33.8%	75.5 %	76.3 %	75.8%
AdaBoost + Word2Vec	78.5 %	78.3 %	78.1 %	81.2 %	80.4 %	80.6 %

Tables 3 and 4 display accuracy rates related to the prediction of the VAT codes. It is remarkable to note that the number of distinct VAT codes is smaller than the number of distinct account codes, especially in case of received invoices as shown in Figure 1. As a consequence, the precision of all classifiers for the VAT codes is on average higher with the respect to the results of prediction of the account codes. In particular, as regards VAT codes, Word2Vec results the best approach to transform textual data in a numeric vector in case of received invoices. In case of sent invoices, there is no differences between the two textual approaches, probably due to the fact that prediction of the VAT codes is not highly influenced by the textual information contained in the invoice but other input variables of the invoice are sufficient to obtain high prediction accuracy. As regards the machine learning model, AdaBoost seems to perform slightly better for sent invoices but it is equivalent to Random Forests in case of received invoices.

In general, we observe that results for the two different accounting firms (Dataset 1 and Dataset 2) are consistent with each other concluding that, at least in this two cases, it is possible to choose a unique type of classifier and train it on historical data of a specific accounting firm. This can be considered an advantage from the business value point of view, since we have a tool adaptable to both the accounting firms, which learn new predictions from own historical accounting database.

To improve performances of models for the prediction of the account codes for received invoices (which is the most problematic target variable considering the accuracy rate of Table 2) we investigated also the hierarchical structure of the Chart of Accounts. Random Forests algorithm has been trained for each level of the hi-

Table 3: Performances of the classifiers which predict VAT codes evaluated on test set (sent invoices)

	Dataset 1			Dataset 2		
	Precision	Recall	F1-score	Precision	Recall	F1-score
RF + BoW	98.5 %	98.8 %	98.6 %	99.8 %	99.8 %	99.8 %
RF + Word2Vec	98.0 %	98.5 %	98.2%	99.6 %	99.6 %	99.6 %
AdaBoost + Bow	98.6 %	98.8 %	98.7%	99.7 %	99.7 %	99.7 %
AdaBoost + Word2Vec	98.4 %	98.7 %	98.5 %	99.6 %	99.7 %	99.6 %

Table 4: Performances of the classifiers which predict VAT codes evaluated on test set (received invoices)

	Dataset 1			Dataset 2		
	Precision	Recall	F1-score	Precision	Recall	F1-score
RF + BoW	93.2 %	93.4 %	93.0%	90.4 %	90.7 %	90.4%
RF + Word2Vec	94.0 %	94.3 %	94.0 %	92.2 %	92.4 %	92.2 %
AdaBoost + Bow	92.6 %	92.7 %	92.2%	89.0 %	89.5 %	89.0%
AdaBoost + Word2Vec	94.0 %	94.4 %	94.1 %	92.0 %	92.3 %	92.0%

erarchy in order to show the improvement of the performances and to motivate a deeper study on hierarchical classification algorithms. Figure 2 shows the accuracy rates for the different levels of the target variable. For both datasets of the two accounting firms the accuracy rates (in terms of Recall) improve at lower levels. This is due to the fact that account codes associated to lower levels of the hierarchy are more generic and easier to predict. This aspect motivates the development of an algorithm which guides the classification output through the hierarchical structure of the account codes.

6 Conclusions

The part of the bookkeeping process which produces a journal entry is often a time-consuming process. Classifying the transactions reported in invoices into specific codes can be translated in a machine learning classification task, whose predictions can facilitate the work of accountants.

This work introduces a methodological proposal to handle accounting data. First of all, the problem of the reconstruction of the training set is tackled, using an heuristic to solve the knapsack problem. Then, different classification approaches on two real data sets are tested; higher accuracy is achieved using the classification algorithm Random Forests combined with the pre-processing technique Word2Vec for the prediction of the account codes. Concerning VAT codes, the classifier AdaBoost




Fig. 2: Accuracy computed at different levels of the account codes (target variable) for received invoices datasets. Model performances improve at lower levels since the target variable is more generic and the number of distinct categories is lower.

combined with Word2Vec performs slightly better. Results show that account codes are the most problematic target variable, due to the high number of distinct labels which can be employed in the classification.

A possible workaround to solve this problem is to consider a hierarchical classification framework which reflects the structure of the Chart of Accounts: the target variable is structured in a predefined hierarchy which takes the form of a rooted tree. This is also motivated by the results obtained for the classification at different levels of our target variable. The application of hierarchical classification algorithm ([19],[4]) can be considered a possible method to improve the results. Since the two target variables predicted in our problem, VAT codes and account codes, are dependent one from each other, another possible approach to improve the accuracy of the account codes can be the application of a multi label classification algorithm. In this case, we would not have two separately classifiers for the two target variable, but a unique classifier which exploit the information extracted from text and from other categorical variables to predict VAT codes and account codes together.

The creation of the general ledger entries starting from an xml invoice seems to be a promising field in which machine learning models can be adopted to automate this repetitive and monotonous part of the bookkeeping process. These preliminary

results support a deeper study in more advanced techniques which can solve some of the existent problems and improve the accuracy of the classifiers.

References

1. Bengtsson, H., Jansson, J.: Using classification algorithms for smart suggestions in accounting systems. Master's thesis (2015)
2. Bergdorf, J.: Machine learning and rule induction in invoice processing: Comparing machine learning methods in their ability to assign account codes in the bookkeeping process (2018)
3. Breiman, L.: Random forests. *Machine learning* **45**(1), 5–32 (2001)
4. Cesa-Bianchi, N., Gentile, C., Zaniboni, L.: Incremental algorithms for hierarchical classification. *Journal of Machine Learning Research* **7**(Jan), 31–54 (2006)
5. Cesarini, F., Gori, M., Marinai, S., Soda, G.: Informys: A flexible invoice-like form-reader system. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **20**(7), 730–745 (1998)
6. Frey, C.B., Osborne, M.A.: The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change* **114**, 254–280 (2017)
7. Gulin, D., Hladika, M., Valenta, L.: Digitalization and the challenges for the accounting profession. Available at SSRN 3492237 (2019)
8. Hastie, T., Rosset, S., Zhu, J., Zou, H.: Multi-class adaboost. *Statistics and its Interface* **2**(3), 349–360 (2009)
9. Holt, X., Chisholm, A.: Extracting structured data from invoices. In: *Proceedings of the Australasian Language Technology Association Workshop 2018*, pp. 53–59 (2018)
10. ICAEW: Artificial intelligence and the future of accountancy. Tech. rep., Information Technology Faculty Chartered Accountants (2018)
11. Joachims, T.: *Learning to classify text using support vector machines*, vol. 668. Springer Science & Business Media (2002)
12. Khan, A., Baharudin, B., Lee, L.H., Khan, K.: A review of machine learning algorithms for text-documents classification. *Journal of advances in information technology* **1**(1), 4–20 (2010)
13. Kozanidis, G., Melachrinoudis, E., Solomon, M.M.: The linear multiple choice knapsack problem with equity constraints. *International Journal of Operational Research* **1**(1-2), 52–73 (2005)
14. Palm, R.B., Winther, O., Laws, F.: Cloudscan-a configuration-free invoice analysis system using recurrent neural networks. In: *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, vol. 1, pp. 406–413. IEEE (2017)
15. Rong, X.: word2vec parameter learning explained. arXiv preprint arXiv:1411.2738 (2014)
16. Schreyer, M., Sattarov, T., Borth, D., Dengel, A., Reimer, B.: Detection of anomalies in large scale accounting data using deep autoencoder networks. arXiv preprint arXiv:1709.05254 (2017)
17. Schultz, M., Tropmann-Frick, M.: Autoencoder neural networks versus external auditors: Detecting unusual journal entries in financial statement audits. In: *Proceedings of the 53rd Hawaii International Conference on System Sciences* (2020)
18. Sebastiani, F.: Machine learning in automated text categorization. *ACM computing surveys (CSUR)* **34**(1), 1–47 (2002)
19. Silla, C.N., Freitas, A.A.: A survey of hierarchical classification across different application domains. *Data Mining and Knowledge Discovery* **22**(1-2), 31–72 (2011)
20. Tang, Y.Y., Suen, C.Y., De Yan, C., Cheriet, M.: Financial document processing based on staff line and description language. *IEEE transactions on systems, man, and cybernetics* **25**(5), 738–754 (1995)
21. Tekbas, I., Nonwoven, K.: The profession of the digital age: Accounting engineering. *IFAC Proceedings Volumes, Project: The Theory of Accounting, Engineering* (2018)

Automatic electronic invoice classification using machine learning models

13

22. Wang, Y., Gui, G., Zhao, N., Yin, Y., Huang, H., Li, Y., Wang, J., Yang, J., Zhang, H.: Deep learning for optical character recognition and its application to vat invoice recognition. In: International Conference in Communications, Signal Processing, and Systems, pp. 87–95. Springer (2018)
23. Zupan, M., Letinic, S., Budimir, V.: Accounting journal reconstruction with variational autoencoders and long short-term memory architecture