USER CONSIDERATIONS WHEN APPLYING MACHINE LEARNING TECHNOLOGY TO ACCOUNTING TASKS

by

Liezl Smith

Thesis presented in partial fulfilment of the requirements of the degree of Master of Commerce (Computer Auditing) in the Faculty of Economic and Management Sciences at Stellenbosch University

Supervisor: Prof W Boshoff

Co-supervisor: Dr C Lamprecht

December 2018

Declaration

By submitting this dissertation electronically, I declare that the entirety of the work contained therein is my own original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third-party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Liezl Smith

Date: December 2018

ABSTRACT

Machine learning is a strategic technology that can have an important effect on business, as it is able to perform tasks efficiently that were previously only performed by humans. When implementing this technology in the relevant business processes and utilising it effectively, users have to understand both it as well as other aspects have to be considered. It was found that one area that is well suited to the adoption of machine learning, is accounting. In addition, prior research has shown a need for accounting users to be educated in machine learning as part of their professional training. Therefore, the aim of this study was to enhance users' understanding of machine learning technology specifically in the performance of accounting processes.

A grounded theory methodology was employed to identifying the accounting tasks machine learning could perform, to describe how machine learning functions and to identify the risks, benefits and limitations associated with the technology. Finally, steps and considerations when implementing machine learning technology in the accounting process were provided.

The findings of this research are that the user has a key role to play when using machine learning technology in the accounting processes and thus has to understand the technology, the risks and limitations, as well as the benefits of the technology. The risks discussed relate not only to machine learning technology but also to all the components that enable the functioning of the technology to ensure alignment with the accounting process goals.

Based on these findings, this research presents the user considerations and steps to take when implementing machine learning in selected accounting processes. These can be used to identify areas that may require attention when a business is adopting machine learning. One important consideration is the implementation of adequate data governance. This is because most of the risks identified for machine learning technology are data risks. Further research could therefore be directed at developing a data governance framework for machine learning technologies.

OPSOMMING

Masjienleer is 'n strategiese tegnologie wat 'n belangrike uitwerking kan hê op besigheid, aangesien dit take doeltreffend kan uitvoer wat voorheen net deur mense uitgevoer is. Wanneer hierdie tegnologie in die toepaslike besigheids prosesse geïmplementeer en doeltreffend benut word, moet gebruikers dit verstaan en verskeie ander aspekte oorweeg. Daar is bevind dat Rekeningkunde een area is wat goed geskik is vir die aanneming van masjienleer. Daarbenewens, het vorige navorsing bevind dat rekeningkundige gebruikers opgelei moet word in masjienleer as deel van hul professionele opleiding. Die doel van hierdie studie was dus om gebruikers se begrip van masjienleertegnologie te verbeter, spesifiek in die uitvoering van rekeningkundige prosesse.

'n Gefundeerde teorie navorsingsmetodologie is gebruik om die rekeningkundige take wat masjienleer kan uitvoer te identifiseer, te beskryf hoe masjienleer funksioneer en om die risiko's, voordele en beperkings wat met die tegnologie verband hou, te identifiseer. Ten slotte is stappe en oorwegings tydens die implementering van masjienleertegnologie in die rekeningkundige proses verskaf.

Die bevindinge van hierdie navorsing is dat die gebruiker 'n sleutelrol speel wanneer masjienleertegnologie in die rekeningkundige prosesse gebruik word en dus moet die gebruiker die tegnologie, die risiko's en beperkings, sowel as die voordele van die tegnologie verstaan. Die risiko's wat bespreek word, hou nie net verband met masjienleertegnologie nie, maar ook met al die komponente wat die funksionering van die tegnologie moontlik maak om belyning met die doelwitte van die rekeningkundige proses te verseker.

Op grond van hierdie bevindinge, bied hierdie navorsing die gebruikersoorwegings en die stappe om te neem wanneer masjienleer in geselekteerde rekeningkundige prosesse geïmplementeer word. Hierdie oorwegings en stappe kan gebruik word om areas te identifiseer wat aandag benodig wanneer 'n besigheid masjienleer implementeer. Een belangrike oorweging is die implementering van voldoende databeheer, aangesien die meeste van die risiko's wat vir masjienleertegnologie geïdentifiseer is, data-risiko's is. Verdere navorsing kan dus gerig word op die ontwikkeling van 'n data-beheerraamwerk vir masjienleertegnologieë.

ACKNOWLEDGEMENTS

I am forever grateful to all those persons who assisted and supported me in the completion of this work. I firstly want to thank my heavenly Father, this work was enabled by you alone. All glory be yours. In addition, I would like to express a special word of thanks to the following people:

- My husband, William, for every hour of selfless time you gave up to stand by and support me. This thesis would not exist without you.
- My grandmother, who has since passed away, for supporting me to carry on.
- My parents, who love and support me and brought me thus far.
- My supervisor, Stiaan Lamprecht. I have such respect for your wisdom and knowledge; thank you for your guidance and willingness to share your expertise with me.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1: Introduction

1.1 Background

In order to maintain a competitive advantage, Gartner Inc. advises that companies should examine the business impacts of strategic technologies, which may indicate a need to adjust business models and operations. Failure to examine these technologies may mean the loss of a competitive advantage (Cearly, Walker & Burke, 2016). One such strategic technology is machine learning, a branch of artificial intelligence. Machine learning has enabled many tasks that would previously have been performed by humans to be efficiently and accurately completed by a computer.

Cearly *et al.* (2016) state that machine learning technologies have been developed to assist systems to appear to understand, learn, predict and adapt. They have the potential to operate with little or no human guidance, which surpasses that of traditional rule-based algorithms. Accordingly, by implementing machine learning technologies, business may benefit from an increase in productivity and accuracy as well as substantial cost savings.

The data science needed to create machine learning systems is complex and most businesses will therefore choose to acquire packaged machine learning applications rather than developing their own (Cearly *et al.*, 2016). However, despite the complexity, management must understand the unique characteristics of machine learning technology to ensure that it is correctly implemented and aligned with the desired business outcomes (Gillion, 2017:3).

The financial division of a business is one of the areas that could benefit from the application of machine learning, specifically the area of accounting. In considering the training given to professional accountants, PwC (2015:16) recognises the need for undergraduate accounting programmes to include advanced topics on machine learning as part of the curriculum. Of the various artificial intelligence skills that PwC (2015:16) recommends, Sutton, Holt and Arnold (2016:68) argue that machine learning is a key stream of artificial intelligence for application in accounting.

This study addresses the need to assist users to understand machine learning technologies, specifically in the area of accounting, as explained in the next section.

1.2 Research focus

In this section, the research problem, research aim, research questions and research objectives will be addressed.

Machine learning techniques can assist accounting users in their decision-making. Although there are many types of machine learning techniques, they can be divided into two categories, namely, predictive techniques and explanatory techniques. The type of technique used will depend on the decision being made. For example, predictive techniques are able to predict outcomes that are based on patterns learnt by the machine learning model from the data although these patterns are often not explained to, or seen by, the user. Predictive machine learning techniques are more complex than explanatory machine learning techniques, which identify factors that are causally related to an outcome (Sainani, 2014:841; Sutton *et al.*, 2016:69).

A lack of ability on the part of the accounting decision maker to discriminate between explanatory and predictive machine learning technologies when deciding which alternative to use and rely on indicates a need for machine learning research in accounting (Sutton *et al.*, 2016:69). Assisting accounting decision makers to understand machine learning technologies and the issues that require consideration when implementing them would help to promote the application of these technologies in accounting and stimulate the need for research on these technologies.

Bräuning, Hüllermeier, Keller, & Glaum (2016:296) emphasised the importance of understandability and model simplicity in machine learning. Understanding the technology and the associated benefits and risks will assist users to select one that is appropriate to their needs. It will also make them aware of the issues to consider and the steps to take when implementing and using these technologies. Accordingly, the aim of this research is to enhance users' understanding of machine learning technology specifically in carrying out accounting processes.

To achieve this aim, the research focused on three accounting processes, identifying the tasks involved in each process. These three accounting processes were selected as they cover some of the main processes in the traditional record-to-report process which is illustrated in Figure 1. These tasks were cast as problems in the accounting processes that machine learning technology could solve. Subsequently, machine learning technologies were identified that could be applied to the identified tasks.

In order to achieve the research aim, three research questions were formulated for this study. Research objectives were then in turn set for each research question. The research questions and the corresponding research objectives, as well as the sections in which the respective findings are discussed, are indicated in Table 1.

Table 1: Research questions and corresponding research objectives

Research questions	Research objectives	Findings
1. Which machine learning technologies can be applied to existing accounting processes?	To outline the components of the accounting processes.	Section 2.2; Section 2.3; Section 2.4
	To identify the tasks that machine learning technology can perform.	Section 3.8
	To identify the machine learning technologies that can be applied to the accounting process tasks.	Section 3.8
2. How does the machine learning technology function and what are the risks and benefits associated with the technology?	To explain how the machine learning technology functions.	Section 3.3; Section 3.4; Section 3.5
	To identify risks, benefits and limitations associated with the machine learning techniques.	Chapter 4
3. What are the considerations and steps to take when implementing and using machine learning technologies?	To explain the role of the user when using the machine learning technology to address identified risks.	Section 5.2
	To identify the steps to take when implementing machine learning technology to ensure alignment with accounting process goals.	Chapter 5

1.3 Research design and methodology

As the research questions focus on gaining theoretical insights in order to develop a method for understanding the technology in an accounting context, the research design will be exploratory. An exploratory research design allows for the development of a grounded picture of the phenomena as well as the development of tentative theories or hypotheses (USC Libraries, n.d.).

A grounded theory methodology was employed to address the research problem. This is a methodology that is used to develop a theory, in this case theory about emergent technologies where there is not, as yet, an established theory (Sutton, Reinking & Arnold, 2011:46). Bryant (2002:35) argues in this regard that the grounded theory method is particularly suited to information systems technology research and highlights the mandate of research to develop both an understanding of discovered facts and adequate models for specified purposes. Furthermore, this method aligns with the aim of this research, namely, to enhance accounting users' understanding of machine learning technology.

The research was both qualitative in nature (Creswell, 2009:4) and non-empirical, as the existing literature was synthesised (Torraco, 2005:357) to achieve the objective of improving the understanding of this technology with respect to accounting tasks. An integrative literature review was therefore performed with the aim of enhancing user understanding of machine learning, specifically with regard to its application in accounting.

The literature review relied on Scopus, EBSCOhost, IEEE and AAA digital library databases (Sutton *et al.*, 2016:64), as well as available publications in the form of books, accredited journals and academic work on machine learning, including publications specifically in the field of finance.

As a starting point, the three accounting processes selected were investigated to determine which components presented problems that could be solved by machine learning. These problems are presented as tasks.

Having provided an overview of the selected accounting processes and tasks, the research identified various machine learning techniques that could be applied to perform the tasks described. This was done by identifying the learning problem, being the task that the machine learning needs to be able to perform (Someren & Urbancic, 2006:363), for each accounting task and then identifying the different types of machine learning technique to address each learning problem and explaining their capabilities.

The research then considered how machine learning techniques were already being used in accounting practice to address the identified tasks. Where no machine learning had yet been applied in practice, suitable machine learning techniques were identified in the research. When evaluating the available machine learning technologies, it was evaluated in the context of the different technologies that encompass artificial intelligence. A link was now presented between accounting processes and machine learning techniques.

Having selected the machine learning techniques to apply to the tasks identified in the accounting processes, the functioning of the applicable machine learning technology was explained. This was followed by a discussion of the risks, benefits and limitations associated with these technology that was grounded in the principals of King IV. Finally, in response to these identified risks, benefits and limitations, the issues to be considered when implementing the technology and the steps to take when using the technology were identified.

1.4 Research motivation

Bailey and Pearson (1983:532) identify a number of factors that influence user satisfaction when using new information technology products. These factors include understanding the system, the perceived usefulness of the system, and the congruence between what the user wanted and what the product provided.

More specifically, to assist users in deciding which machine learning techniques to apply, an understanding of the strengths and weaknesses of these techniques in the context of business is useful (Bose & Mahapatra, 2001:211). Therefore, users hoping to exploit the advantages of machine learning in their accounting processes would benefit from an understanding of the technology, its benefits and its applications, as well as the issues to consider and the way that the user can responds to the identified risks.

This research will address the need for research into strategic technologies, namely machine learning, and the need to assists users in understanding the technology. Assisting users to understand the technology will hopefully encourage them to consider implementing machine learning when a problem arises that machine learning could address.

1.5 Research scope

The focus of this research will be on the following specific processes in accounting:

1. Translation of manual and electronic documents into accounting software information
2. Reconciliation of financial information
3. Preparation of management accounts

1.6 Limitation of research

The research does not intend to address all the areas of accounting in which machine learning intervention is possible; it will only address the three accounting processes that have been identified. Only those tasks for which a suitable machine learning technology could be found at the time of this research were addressed. Furthermore, the research will only consider machine learning technologies appropriate for addressing the identified problem types in the accounting process and therefore does not intend to present an exhaustive list of machine learning technologies.

Considerable research has been performed in the areas of machine learning applied to auditing and the detection of fraud using such technologies, therefore these areas were not considered for this research.

In order to achieve the objective of understanding the machine learning technologies, the design and functioning of each technology will be explained. In doing so, the research is limited to explaining the design and functionality of the technology for the purposes of understanding it and identifying the associated risks, benefits and limitations. Therefore, this is not explained at the technical level required to develop the technology.

The risks and benefits identified for the machine learning technology will be those that are unique to the technology and not those that pertain to the environment in which machine learning operates such as database or accounting software risks or risks pertaining to supporting technologies such as cloud platforms. Hence, these risks and benefits are not addressed in this research.

1.7 Organisation of the research

The research consists of six chapters.

Chapter 2 identifies the various accounting processes that will be addressed and outlines the tasks involved in these processes. This assists in outlining the areas that can be addressed by machine learning technology.

Chapter 3 provides an overview of the machine learning techniques available to perform the tasks identified in the accounting processes. The chapter also provides an understanding of the way machine learning technology is structured and how machine learning techniques work. This chapter further discusses which techniques are currently applied to accounting process tasks and which tasks could be applied to the identified accounting processes tasks as based on the existing research.

Chapter 4 explains the risks, benefits and limitations of machine learning technologies.

Chapter 5 formulates the guidelines for implementing machine learning technologies in the accounting context.

Chapter 6 concludes with a summary of the main findings of the study for consideration by users when implementing machine learning technology in an accounting context.

Chapter 2: Accounting processes and accounting tasks

2.1 Introduction

The aim of this study was to enhance users' understanding of machine learning technologies specifically in the performance of accounting processes. This chapter provides a description of the three accounting processes that were selected, and highlights the tasks involved in each of these processes. The purpose of the chapter, then, is to investigate the tasks performed in each accounting process and the technology that is currently being employed in the process and, finally, to describe the tasks. The tasks which could be performed by machine learning are presented in chapter 3 section 3.8.

To gain an understanding of the accounting processes, numerous services offered by financial process software providers were considered. There is a paucity of published research on applications that employ machine learning techniques. While the reason for this is unclear, it is speculated that it may be due to a lack of reporting on such applications, as a result of the unwillingness to reveal these applications for competitive reasons (Amani & Fadlalla, 2017:39).

The three accounting processes identified for this research were broken down into a number of tasks, each representing a certain sphere of the administrative activity. After identifying the different tasks, specific consideration was given to identifying which technology was used in each accounting process. Krutova and Yanchev (2014:13) indicate that the technology applied in an accounting process entails a sequence of measures that assist in the introduction of resources at the input phase, the processing of the resources and the production of different levels of information sets at the output.

Once the technology was identified, it was important to identify the data and the structure of the data available for the task, as the data needed to be appropriate for machine learning techniques (Fedyk, 2016). The choice of machine learning technique was based partly on the differences in data characteristics or data type (Someren & Urbancic, 2006:371).

Accordingly, the product of the current chapter is the layout of the tasks performed in the three accounting processes selected, together with an understanding of the supporting technologies identified as useful for enabling machine learning techniques to be applied and the structure of the data available in each accounting process. These will be used to identify the relevant machine learning techniques in chapter 3.

The research is set out in accordance with the traditional record-to-report process, as illustrated in Figure 1. In line with this, firstly, the manual and electronic documents are translated into accounting information, secondly, financial information is reconciled, and finally, management accounts are prepared.

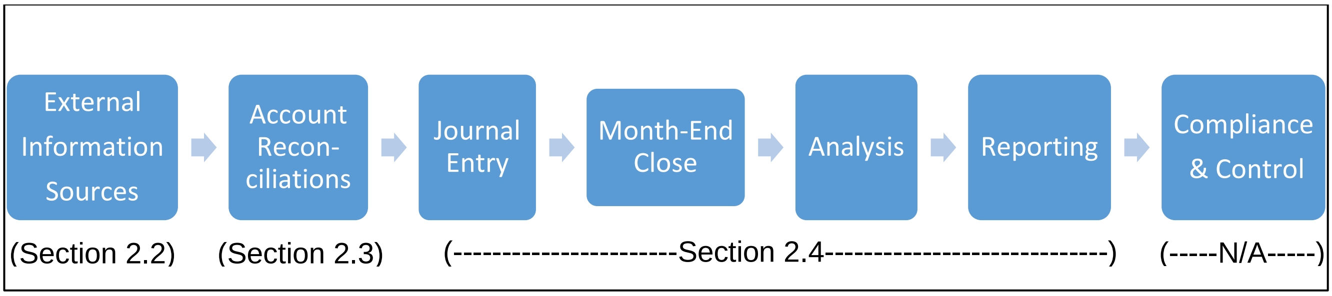

Figure 1: The traditional record-to-report process (Deming, n.d.:7) **(Adapted)**

Figure 1 above shows that the record-to-report process commences with the external information sources, which are covered in section 2.2, followed by account reconciliations, which are addressed in section 2.3, then journal entries and month-end closure, and finally analysis and reporting are performed, which are addressed in section 2.4. This study does not address the performance of compliance and control procedures.

2.2 Translation of manual and electronic documents into accounting information

The first process to be addressed is the translation of manual and electronic documents into accounting information. The purpose of converting documents into a digital form is to enable businesses to analyse the documented data efficiently and more affordably. In pursuance of this purpose, the process of detecting, extracting and processing data from documents needs to be efficient and accurate (Ming, Liu & Tian, 2003:489; Rhodes & Wheat, 2015:1). The efficiency and accuracy of the process may be increased by incorporating machine learning techniques.

2.2.1 Breaking down the translation of documents into tasks

This section describes each of the tasks in the process of translating documents into accounting information. Various documents need to be translated into electronic information for accounting purposes. An example of this process is illustrated in Figure 2 (Ming *et al.*, 2003:490) for a supplier invoice.

Task 1: Input

↓

Task 2: Pre-processing and standardisation

↓

Task 3: Document features extraction

↓

Task 4: Known type? → No → Update form feature library

↓ Yes

Task 5: Form data field extraction and data recognition ←

↓

Task 6: Validation of document data

↓

Task 7: Export: Processing and storage of data

Figure 2: Flow chart of tasks using a bill recognition system as an example (adapted from Ming *et al.*, 2003:490)

In Figure 2 the tasks in the documentation translation process are set out in the order in which the software or user performs them. The tasks are discussed in more detail as follows:

Task 1. Input: Documents such as invoices are uploaded into the translation software in the form of a paper document which is either scanned or loaded electronically, for example in the form of an email attachment (Kohlmaier, Hess & Klehr, 2006:1). The document format will affect the structure of the available data as described in section 2.2.3.

Task 2. Pre-processing and standardisation: Documents are pre-processed prior to data extraction. Various algorithms, which are described in section 2.2.2, are applied to correct the images and then standardising techniques are applied to these images, which include standardising the page size, page orientation, font, text size, colour, element positioning, page numbering, margins, column size, watermarks, line numbering and page breaks (Rhodes & Wheat, 2015:0029).

Task 3. Document features extraction: The components of the documents are translated by identifying the document template and the document schema. The *schema* identifies the different elements of information which are to be extracted from the document while the *template* sets out how the information is physically arranged in the document. Based on the identified template, the system will then have a set of rules for locating the information for each schema in the document (Sorio, 2013:25).

The data is extracted from the documents and converted to an electronic format using Optical Character Recognition (OCR) (Kohlmaier *et al.*, 2006:1; Rhodes & Wheat, 2015:0031). Accordingly, the document features may trigger algorithms which label identifiable data elements. For example, the presence of the words “Invoice number” may trigger the string of numbers following that header to be labelled as invoice number (Kohlmaier *et al.*, 2006:1). The technology used to perform the data and features extraction is described in section 2.2.2.

Additionally, the document translation software may extract other types of data from the document, for example the metadata, which identifies the document creator, time created, document file size, creation date and any modification dates, type of text in the file, version, font, name, and other data. These elements can be used as a database navigation key (Rhodes & Wheat, 2015:5).

Task 4. Form type feature library which includes document type recognition and classification: A library, knowledge repository or set of template-specific extraction rules may be kept of known document types, for which particular content, structure, form or other attributes have been established. Using the library, the document translation software can automatically determine where to find relevant data on a given document based on the document data type recognised (Ming *et al.*, 2003:493; Sorio, 2013:27; Rhodes & Wheat, 2015:0036).

The document type can be determined by the document translation software based on the extracted identifiable data elements such as the schema and the template by considering known document types in the form type feature library as described.

If, however, the document is an unknown type, the extracted results are displayed and the results need to be confirmed manually item by item. In this way, the user can train the system to identify and extract the key data. These features are subsequently stored in the form-type feature library. These stored rules are then able to interpret future documents from the same source (Ming *et al.*, 2003:490; Kohlmaier *et al*., 2006:2; Sorio, 2013:23).

Task 5. Form data field extraction and data recognition: Once the document type is established, the document translation software determines where to locate important data on the given document (Ming *et al.*, 2003:493; Rhodes & Wheat, 2015:0036). Thus the software is able to present structured data from the extraction, because it is able to determine which data is important and how that data should be recognised. For example, the identifying elements in an invoice, such as an invoice number, enable the software to identify the document as an invoice and process the data as invoice data (Kohlmaier *et al.*, 2006:2).

Task 6. Validation of document data: The software tests the validity, accuracy and completeness of the data by means of validation tests. These are, for example, able to ascertain whether there are inconsistencies between paired documents such as purchase orders and supplier invoices, or whether the invoice is a duplicate invoice. Error detection may then prompt the user to process the invoice manually (Kohlmaier *et al.*, 2006:1,2).

Task 7. Processing and storage of data: After the data extraction process, the user is shown the structured dataset generated. This interface may be interactive to allow entries in the generated dataset to be edited manually if necessary (Rhodes & Wheat, 2015:9). The user then confirms the data, and the interpreted data is converted into a file which is saved in a database in a specific document format such as Extensible Mark-up Language (XML), although there are various other documents formats that could be used (Rhodes & Wheat, 2015:0068). XML is a standard data exchange format with many variations, which makes the data structure complex (Lee, Tsatsoulis & Perry, 2009:1).

2.2.2 Technology used in document translation

Certain technologies can be greatly improved by incorporating machine learning techniques. The technologies identified in the document translation process are set out in accordance with the various tasks identified in section 2.2.1.

Task 1. Input: The technology employed to upload electronic documents include scanners, image capturing devices such as digital cameras and a variety of software, producing electronic documents in diverse formats.

Task 2. Pre-processing and standardisation: Images may be corrected by means of a trembling process to reduce noise in the image, and image angle testing may be conducted to correct for any slant in handwritten words or in entire images. This technology improves the data extraction task (Ming *et al.*, 2003:490-492).

Task 3. Document features extraction: Once documents are standardised, recognisable elements such as text, numbers and special characters can be automatically detected using optical character recognition (OCR) (Rhodes & Wheat, 2015:2).

OCR is a technique used to convert scanned documents into computer readable text (Larsson & Segerås, 2016:5). The three basic principles applied to OCR for recognising objects include integrity, purposefulness and adaptability (Emmanuel & Nithyanandam, 2014:439; ABBYY Technologies, n.d.). The last-mentioned, adaptability, may be assisted by the program being able to learn by itself, as enabled by machine learning.

Task 4. Form-type feature library which includes document type recognition and classification: The document type may be recognised based on its *form* using image classification, which sorts documents by appearance or pattern. This may incorporate machine learning techniques. Text classification can be used to classify the document type based on content, and both statistical and semantic text analysis are employed to classify text content (ABBYY, 2017).

The system can be trained to process flexible or irregular document layouts by incorporating machine learning techniques together with natural language processing (NLP) (ABBYY, 2017). NLP converts human language into a format that computers are able to recognise and use (Collobert & Weston, 2008:160).

Task 5. Form data field extraction and data recognition: Template-specific extraction rules are algorithms which enable the system to extract information from documents. These rules predict which field a specific data item refers to, for example a particular item may be the invoice number field and another the date field (Sorio, 2013:17).

Case-based reasoning, based on previous cases, can be used to decide which techniques to use to extract the data fields (Hamza, Belaïd & Belaïd, 2007:327; Larsson & Segerås, 2016:8). Watson (1999:307) describes case-based reasoning as a method which attempts to solve cases by using solutions observed in similar previous cases.

Task 6. Validation of document data: As with any data input, the accuracy of the data will need to be verified. Arithmetic validation rules can be applied to ensure the accuracy of the data (OCREX, 2017). In addition, machine learning can be used to determine whether the data from the document is correct, as demonstrated by Larsson and Segerås (2016:37).

Task 7. Processing and storage of data: Bose and Mahapatra (2001:212) indicate that data warehousing technology enables the organisation and storage of large amounts of financial information in a form that can be analysed using machine learning techniques .

The classification of the problems or tasks that can be addressed by machine learning, as well as the description of the specific techniques available for use, are described in chapter 3 section 3.8.

2.2.3 Data available in document translation

Printed documents may have weak structures caused by the placement of the text on the page. This has to be addressed using template specific extraction rules. Other documents may be in electronic and structured formats such as XML, from which the information in the document can be easily determined (Sorio, 2013:13).

Documents, and therefore the data used in training machine learning technology, are easily accessible and the input data either already exists in digital form or it can be easily digitised (SMACC, 2017:6). However, variations in the components of data input indicate the need for solutions to be adaptable. This consideration is discussed in chapter 5 section 5.4.1.

2.3 Reconciliation of financial information

This section will address the second accounting process, which is the reconciliation of financial information. The purpose of the account reconciliation process is to verify the integrity of a business's account balances. An example of the process is illustrated in Figure 3.

Task 1: Data preparation

↓

Task 2: Comparison

↓

Task 3: Matching

↓

Task 4: Exceptions identified

↓

Task 5: Exceptions investigated

↓

Task 6: Corrective action

↓

Task 7: Documentation

Figure 3: Tasks in the reconciliation accounting process (adapted from BlackLine, 2014)

These tasks will be performed for all the reconciliations a business has to perform as part of its overall financial reporting process and will depend on the industry in which the business operates and the nature of the transactions. Not all reconciliations are financial, for example businesses may require industry-specific reconciliations, but these will still support the reporting process (Trintech, 2017). Table 2 (Trintech, 2017) presents a summary of different types of reconciliation.

Table 2: Types of reconciliation

Category	Type of reconciliation
Bank reconciliations	Bank accounts Credit cards Gift cards Currencies
Other reconciliations	Inventory Supplier statements Goods receipt invoice receipt Suspense accounts System to system Intercompany Payroll Inventory
Industry specific reconciliations	Airliners (bag drop) Insurance (policy) Hospital (patient count)

Source: Adapted from Trintech (2017)

2.3.1 Breaking down the account reconciliation process into tasks

BlackLine (2014), a company which is endorsed by SAP for its accounting automation software, describes the account reconciliation process. The tasks performed in the account reconciliation process are as follows:

Task 1. Data preparation: Data is collected and processed into an appropriately comparable format, which depends on the technology employed to perform the comparison.

Task 2. Comparison: The transactional data contained in the account balance and the information produced by an independent system are compared. For example, for the bank general ledger account, a comparison is made between the information contained in the general ledger bank account and the information as per the bank statements, which are produced by an independent system.

Task 3. Matching: Transactional data records contained in the one set of information are matched to the corresponding record in the second set of information. These records are linked and marked as reconciled, verifying the integrity of the data.

Task 4. Exceptions identified: Discrepancies are identified where differences are detected between the data being compared in each set of information, or where data available in one system has no corresponding data in the other system.

Task 5. Exception investigation: The discrepancies identified are investigated by scrutinising the origin of the data records and inspecting the supporting sources for the existing data in order understand the cause of the discrepancies.

Task 6. Corrective action: Once the cause of discrepancies is determined, corrective action is taken. This may involve making journal entries to correct balances or transactional errors.

Task 7. Documentation: The investigation process is documented together with the corrective action and any supporting documentation, after which all of this information is stored for audit purposes.

2.3.2 Technology used in the account reconciliation process

The technologies required in the reconciliation process are set out in accordance with the different tasks identified in section 2.3.1.

Task 1. Data preparation: The datasets that are compared may be imported from a variety of sources, including accounting software packages, enterprise resources planning (ERP) software (described in section 2.4.2), external platforms such as banking or supplier platforms or data converted from scanned manually generated documents (BlackLine, 2014). The technology used in the preparation of data is discussed in section 2.2.2 above.

Task 2 and 3. Comparison and matching: Technologies employed to perform data matching have included record linkage approaches. These record linkage approaches include rules-based approaches, which rely on heuristics. However, the drawback of this method is that heuristics developed for one application are not likely to work for another. For automated reconciliation, deep learning and statistical methods are recommended (Chew &

Robinson, 2012:326). The problems that can be addressed with machine learning are described in chapter 3 section 3.8.

Task 4 to 7. Exception investigation, correction and documentation: The software will produce a report displaying both the unmatched and the matched items (Chew, 2014:2). The user will have to investigate, correct and document any unmatched items.

2.3.3 Data available in the account reconciliation process

An account reconciliation will consist of two transaction datasets, each of which may be from a different source and in a different format (Chew, 2014:1). Before the software can perform a reconciliation, the transactional data must be converted into the correct format.

2.3.3.1 Sources of data

Financial datasets may be sourced from the entity's own ERP system, accounting software package, or other internal system, while other datasets may comprise third-party documents, such as bank statements and supplier statements obtained from external platforms such as those of banks and suppliers (BlackLine, 2014).

2.3.3.2 Data format

The data, namely, the transactions being reconciled, consist of a set of transaction features. These transaction features may comprise only one data field or may be separated into a number of fields. The features represent the information that makes up the transaction and enables the software to distinguish one transaction from another. For example, the feature may comprise the transaction description, transaction date, account number or transaction value, represented as separate fields or one single field, depending on the source from which the data is obtained (Chew, 2014:3).

The data in each dataset may appear in a structured format such as categorical or continuous data or it may be in an unstructured format, with free-form text. Free-form transaction descriptions mean that the descriptions are at the user's discretion and not limited by a category such as an account number or invoice number, thus making this unstructured format more difficult to reconcile as result of the variation (Chew & Robinson, 2012:324).

2.4 Preparation of management accounts

The final process to be addressed is the preparation of management accounts. Management accounts are a summary of a business's accounting data which is prepared for the firm's management. The purpose of preparing management accounts is to provide information that can assist management with decision-making (Whittington, 2007:198; Gorbunova & Bochkarev, 2011:25). An example of the process is illustrated in Figure 4.

Task 1: Set up of the chart of accounts

↓

Task 2: Validation of data

↓

Task 3: Account allocation

↓

Task 4: Report templates selection

↓

Task 5: Posting of key journals

↓

Task 6: Period selection

↓

Task 7: Report generation

↓

Task 8: Report descriptions

Figure 4 Tasks in the management reporting process (adapted from DBASS Chartered Accountants, 2014)

Management accounts have three components, namely, financial reporting, performance measurement and cost planning (Cokins, 2013:27). Figure 5 illustrates the different components and gives examples of the different types of report that make up each component, in line with Trigo, Belfo and Estébanez (2014:120).

The reports contained in management accounts may be either descriptive and retrospective, providing an overall view of the business's historical performance, or they may be predictive and prospective, making use of historical data to predict future events (Amani & Fadlalla, 2017:35; Appelbaum, Kogan, Vasarhelyi & Yan, 2017:36). The differences inherent in the reports affect the tasks involved in their preparation.

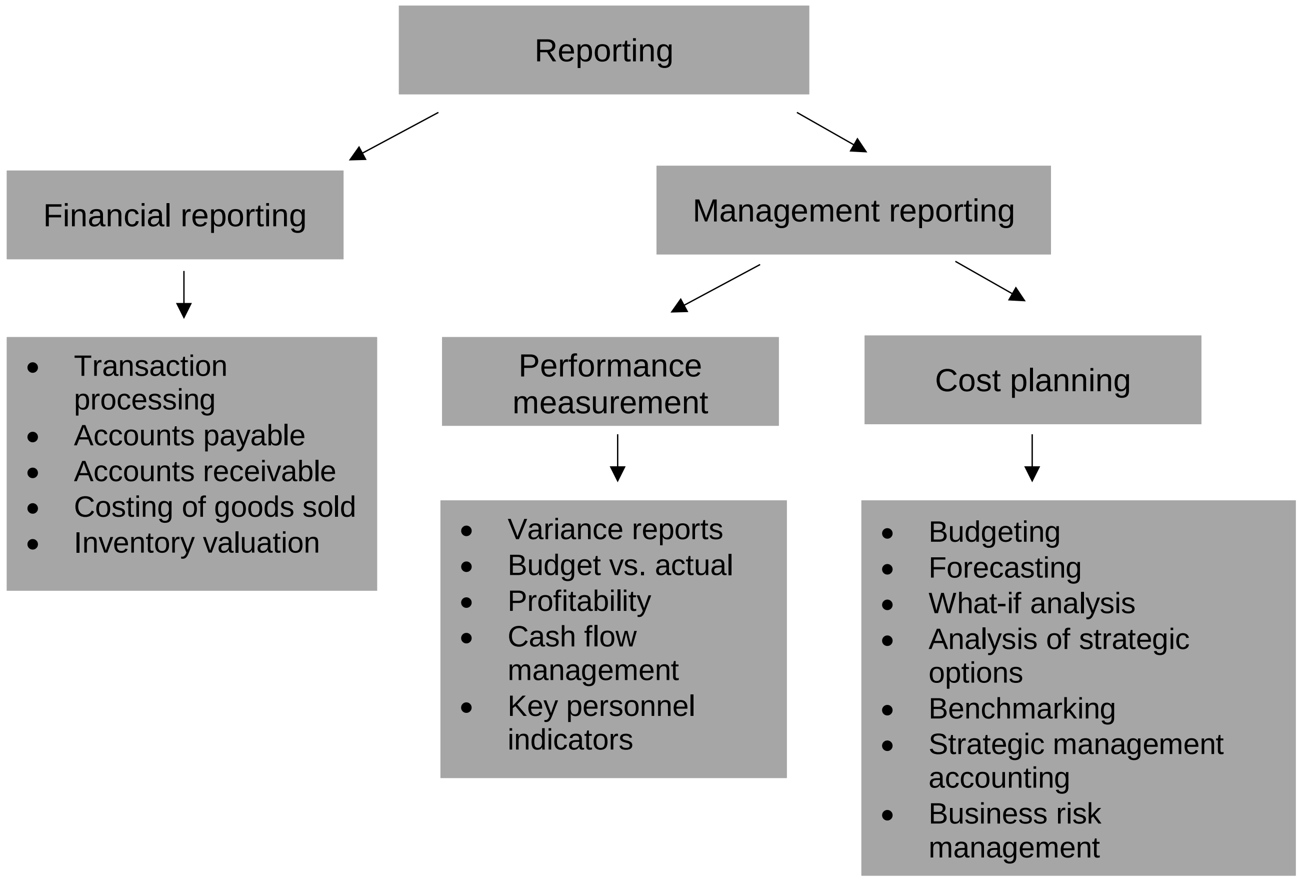

Figure 5: The components of management accounts (adapted from Cokins, 2013:27)

Business needs may vary, and small businesses in particular are unlikely to produce all the reports mentioned here. Those that are produced would be dependent on management's decision-making needs (CPA Australia Ltd, 2011:4).

2.4.1 Breaking down the management accounts reporting process into tasks

The first set of reports included in the management accounts is the financial reports, which are generated using internal data (Appelbaum *et al*., 2017:35). In order to produce an accurate set of financial reports, the process followed in the Sage Line 50 software was considered (DBASS Chartered Accountants, 2014). The tasks included in the preparation of financial reports are as follows:

Task 1. Set up of the chart of accounts: The chart of accounts (COA) is an organised listing of the individual accounts that are used to record transactions and make up the reporting lines of the ledger system. The COA therefore provides the structure for classification of financial information. Classification structures will vary from company to company. Management accounts could, for example, break down financial information into underlying segments such as departments, products, employees, geographical locations, projects and customers (Cooper & Pattanayak, 2011:3).

Task 2. Validation of data: The data used to prepare the management accounts must be accurate, valid and complete. Therefore, prior to preparing the accounts, all transactions should be posted and reconciliations performed as described in sections 2.2 and 2.3 above.

Task 3. Account allocation: The transactions used in preparing the reports may be stored in a database and will need to be allocated to the specific chart of accounts by the user. The account allocation will be based on the nature of the transaction.

Task 4. Report templates selection: Reporting templates may be used in the preparation of management accounts. These are templates used for creating financial reports with predefined fields and formats. The templates can be predefined or blank and customisable (Roy, 2005:3).

Task 5. Posting of key journals: Period end journals may need to be posted, including for example wage journals, stock journals, prepayments and accrual journals.

Task 6. Period selection: The period for which the management accounts are prepared is selected.

Task 7. Report generation: The reports that make up the management accounts are produced by summarising the transactions for each reporting line in accordance with the chart of accounts mapped onto reporting templates.

Task 8. Report descriptions: Notes and descriptions may be added to the quantitative information contained in the management accounts to assist users' understanding of the information.

Next, performance measurement is carried out using insights and inferences, as well as analysis of the processes or events that have taken place, to evaluate corporate performance. This process uses mainly internal data, although some external data such as industry benchmark information is also used to process cost analysis reports.

Cost planning reports are subsequently produced using both the financial reports and the analysis performed. Among other things, these cost planning reports forecast the future business performance, determine a budget to achieve the desired forecast and evaluate strategic options (Appelbaum *et al*., 2017:35).

2.4.2 Technology used in the management accounts reporting process

In this section, the technologies required for preparing the management accounts are discussed in line with the various tasks identified in section 2.4.1. In this process, technology is used to prepare the data, carry out the matching, produce the report of any exceptions identified and process any corrective actions.

User guidelines are discussed in chapter 5. As part of these guidelines, section 5.5 discusses the investment in technology required to enable the use of machine learning techniques. This section identifies those technologies that may enable the use of machine learning techniques in reporting.

Management accounts can be prepared either using information from the business's accounting information system or extracting it from the accounting module of a larger enterprise resource planning application.

An **accounting information system (AIS)** is an application which works together with other information technology systems to record accounting transactions. An AIS collects, stores and processes financial and accounting data which then used to prepare reports for use by management; data may also include nonfinancial transactions that may impact the processing of financial transactions (Belfo & Trigo, 2013:537).

An **enterprise resource planning (ERP)** system is an information system that integrates resources, business process activities and information. It consists of various modules, of which accounting is one. The processes performed in the accounting module include the creation of a chart of accounts, posting to journals and the generation of financial statements (Adhitama, Sarno & Sarwosri, 2016:20).

An ERP system enables management to access the operational data required for decision-making and business control (Appelbaum *et al.*, 2017:31).

Extensible Business reporting language (XBRL) assists in the integration of applications. XBRL, a XML-based language, is a global standard for communicating business information. Accordingly, its use will ensure interoperability, thus allowing accounting information systems be more integrated with other systems. This is important as the automation of management accounting reporting will be dependent on the level of integration of the different applications producing information for the reports (Belfo & Trigo, 2013:542).

Natural language generation (NLG) used in producing **report descriptions** is a technology that converts structured data into written or spoken language. By incorporating an inference engine, the NLG system can perform tasks like summarising large amounts of data, explaining why datasets change, and making recommendations (Yseop, 2017:6).

2.4.3 Data available in the management reporting process

During the analysis of the document translation and reconciliation processes, the recommended format identified for data was **XBRL**. The Companies and Intellectual Property Commission (CIPC) acknowledges that XBRL can be used to integrate back-end processes in companies when automating the preparation of financial statements (Companies and Intellectual Property Commission, 2017b).

The CIPC is the central government agency in South Africa responsible for the registration of all companies. The CIPC indicates that the use of XBRL enables the automatic verification of compliance by means of a validation engine, with the aim of improving the efficiency and accuracy of the reporting process (Companies and Intellectual Property Commission, 2017b).

Big data may enable the preparation of more complex reports as it ensures sufficient data for making decisions. Big data can consist of data gathered inside the business – because this data is usually stored in a database it is generally structured – and external data, which is gathered from outside the business. External data is generally unstructured and therefore often requires analytical tools to extract the information required for decision-making (Appelbaum *et al*., 2017:35).

2.5 Conclusion

This chapter described the three accounting processes, as well as the tasks performed in each process, the supporting technologies that would enable the inclusion of machine learning techniques in the process and the data available in the respective accounting processes.

An understanding of the tasks performed is the key to identifying the areas that can be addressed by machine learning technology. Each of the tasks that can be addressed using machine learning techniques will present a learning problem. Their machine learning solutions will be identified in chapter 3 and will be organised in the same sequence as the identified tasks.

Chapter 3: Overview of machine learning

3.1 Introduction

Having identified the different tasks in the accounting processes in chapter 2, this chapter describes the tasks that are considered problem areas and which are to be addressed by machine learning technology. The chapter also considers areas where machine learning techniques may be employed to enhance the capabilities of existing technologies in the respective processes (identified in chapter 2).

The objectives of this chapter of the study are to describe the components of the machine learning technology and the different machine learning techniques, as well as to identify the learning problems that machine learning techniques can address and that can be applied to each of the identified learning problems. This is in line with one of the research objectives of this study: to identify machine learning techniques that can be applied to the tasks in the accounting process.

The first section, section 3.2, provides a context for machine learning, which is considered important in view of the aim identified for this study, namely, to enhance users' understanding of machine learning technology specifically in the performance of accounting processes.

The next section provides a description of the types of machine learning and the components of a machine learning architecture. Various machine learning techniques are then described and, in conclusion, the final section indicates the learning problems that can be applied to these techniques.

The machine learning techniques identified in this chapter are then discussed in more detail in chapter 4, where the risks and benefits associated with the specific machine learning techniques are described, as well as the risks, benefits and limitations of machine learning technology. These identified risks, benefits and limitations are then used in chapter 5 to formulate guidelines for implementing machine learning technology in an accounting context.

3.2 Context and framework of machine learning

Machine learning is one of the technologies that can be used in the accounting process to assist in the automation of tasks. Tools used in automation can be divided into different classes, namely, rules-based automation (robotic), knowledge-based cognitive intelligence and artificial intelligence. Combining artificial intelligence tools such as machine learning with robotic tools can automate the processing of unstructured inputs from beginning to end (Everest Group, 2014:11).

Machine learning is a subset of artificial intelligence, where patterns in data are learnt and applied in a changing environment. The technology does not require all possible situations to be known during development. Machine learning can be used in two ways: to detect the patterns that explain a process, known as explanatory machine learning technology, and to make predictions; this is known as predictive machine learning technology (Ayodele, 2010a:2; Sainani, 2014:841).

Machine learning technology is able to predict solutions or detect patterns despite uncertainty. This differs from knowledge-based systems which can only solve problems using stored knowledge and facts, as well as heuristics and other elements such as models and known patterns provided by human experts (Valavanis, Kokkinaki & Tzafestas, 1994:114).

Figure 6 (Sutton *et al.*, 2016:62) illustrates the different branches of artificial intelligence and therefore provides the context of machine learning in relation to other artificial intelligence technologies.

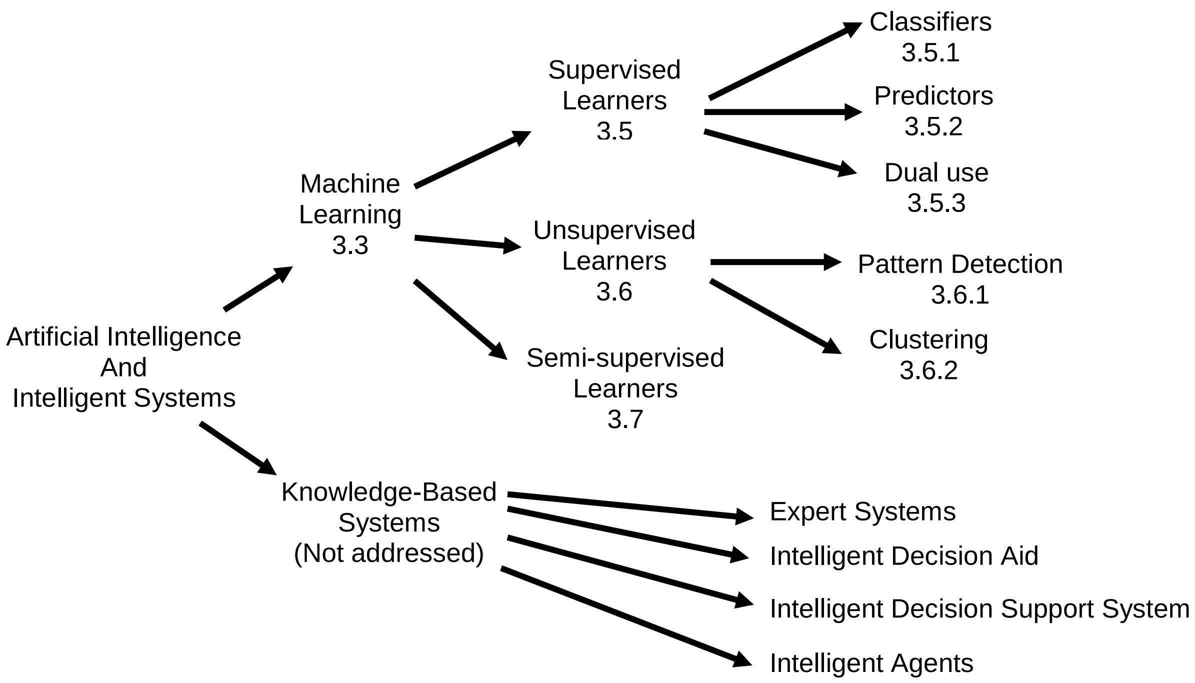

Figure 6: The artificial intelligence tree: the many branches of artificial intelligence application (adapted from Sutton *et al.*, 2016:62)

It is important to be able to distinguish between the different artificial intelligence technologies, as different technologies may be able to address different user needs in business processes (Everest Group, 2014). However, knowledge-based systems are beyond the scope of this study.

Having provided a context for machine learning technology, the different types of machine learning technique are discussed in the next section. This is intended to contribute to understanding the framework for machine learning technology and was considered important, as non-expert users of machine learning techniques do not usually understand the terms that distinguish the different types of machine learning (Someren & Urbancic, 2006:366).

3.3 Types of machine learning algorithm

Having provided the context for machine learning within the field of artificial intelligence, as well as a description of the different components of machine learning technology, this section describes the different types of machine learning algorithm. The different machine learning types are distinguished by considering the objective of the algorithm, how the machine learning algorithm learns, as well as the structure and volume of the data used for learning (Ayodele, 2010:19; Castle, 2018:1).

Supervised learning algorithms require training. The algorithm is trained by using a labelled dataset which consists of examples of input data as well as the labels which indicate predicted targets or output data. Labels assist the algorithm in determining which features are important. The algorithm then generalises the training set by mapping the inputs to the correct responses, which enables it to produce output for new inputs (Ayodele, 2010:19; Castle, 2018:1; Larsson & Segerås, 2016:11; Marsland, 2009:6; SMACC, 2017:9).

Unsupervised learning algorithms do not require training. The input data is unlabelled, meaning the predicted values are not provided, which may be because they are unknown. The algorithm needs to determine the links between the inputs provided to identify patterns or commonalities that can be used to categorise new data or solve problems (Ayodele, 2010:19; Larsson & Segerås, 2016:11; Marsland, 2009:6; SMACC, 2017:9).

Semi-supervised learning algorithms are trained using a combination of labelled and unlabelled data to generate an appropriate function. The labelled portion indicates patterns which may exist, while the unlabelled data, usually the larger portion of the data, is used to establish perceived or unknown patterns for the data (Ayodele, 2010:19; Castle, 2018:1).

Having described the different machine learning technique types, the machine learning architecture is described in the next section. The architecture provides an overview of the components of the machine learning technology, including the different areas required to create a machine learning model that supports the use of the machine learning techniques.

3.4 Machine learning architecture

This section describes the different components of the machine learning architecture. This architecture will be adapted to the needs of the machine learning model as determined by the type of machine learning techniques used for a given task. The different machine learning techniques are subsequently described in the next section.

Gartner recommends that the following five functional areas be included in the machine learning architecture: data acquisition, data processing, data modelling, execution and deployment (Sapp, 2017:19). These areas are illustrated in Figure 7 which diagrammatically demonstrates the machine learning architecture.

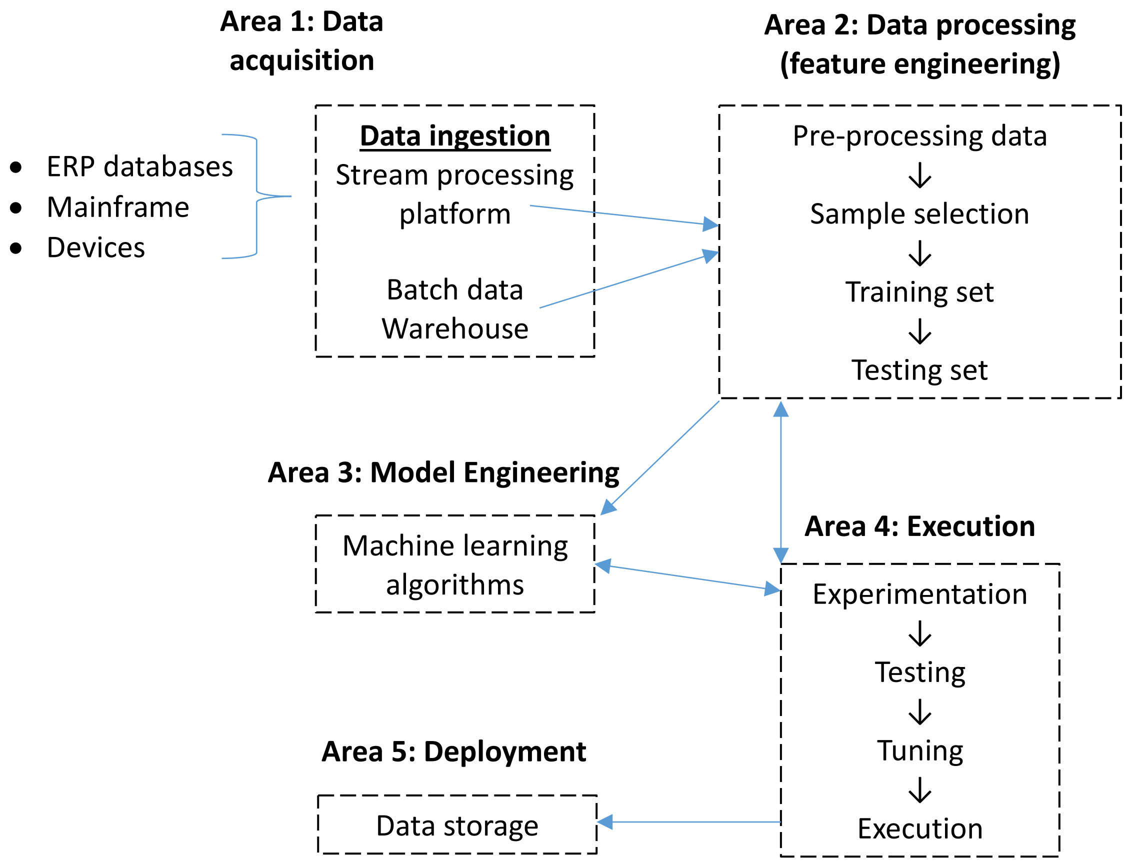

Figure 7: Machine learning architecture (adapted from Sapp, 2017:20)

The functional areas illustrated in Figure 7 of the machine learning architecture are described by Sapp (2017:21) as follows:

Area 1: Data acquisition

Encompasses the collection and preparation of data for processing from a variety of sources and ensures that the data is reliable and adaptable for processing.

Area 2: Data processing including feature analysis

This area normalises and transforms the data into a structure suitable for machine learning. In addition, training sets and testing sets are selected in this area. A training set is a dataset which is used by the algorithm to identify relationships (Dataiku, 2017:5), while a testing set is a dataset used to assess whether the algorithm functions as desired. Feature analysis is performed to assess which features of the data are required for training the machine learning algorithm.

Area 3: Data modelling

This area includes the selection of the machine algorithms and the adaptation of the algorithm to the identified learning problems.

Area 4: Execution

In this area the prepared data and the machine learning algorithm are brought together to train the machine learning algorithm, test the model and then make any necessary changes to the algorithm to ensure that the machine learning model operates in a way that addresses the learning problem.

Area 5: Deployment

In this area the outputs of the machine learning model are made available for use in the applicable business applications or are stored as data to be used in reporting for example.

Having described the different machine learning technique types and the machine learning technology architecture, the different machine learning techniques available to address the learning problems identified in the form of accounting tasks are described in the next section.

3.5 Description of the supervised learning techniques

Descriptions of each of the machine learning techniques are provided in sections 3.5 to 3.7. The techniques are organised by learning problem type, as shown in Figure 6 of section 3.2. In describing the different machine learning techniques, the term “features” in data science refers to the independent variables or predictor variables (Dataiku, 2017:5).

The techniques described are applied to the learning problems in section 3.8. The benefits and limitations of the different machine learning techniques are identified in chapter 4 section 4.7. This section describes the various supervised learning techniques, starting with classification algorithms.

3.5.1 Classification algorithms

Classification algorithms are supervised learners, and therefore their development consists of a two-step process consisting of training and testing. During training, the algorithm maps class labels to data features. These features can predict the class labels of new data by learning from a training dataset that will consist of a set of data records and associated class labels for each record. Once trained, the accuracy of the classifier is assessed by testing its ability to predict classes using the test dataset (Chadha & Singh, 2012:51).

The different supervised classification algorithms can be separated into three different types of technique, namely, logic-based techniques, perceptron-based techniques and statistical techniques (Kotsiantis, 2007:251). These are presented in Figure 8.

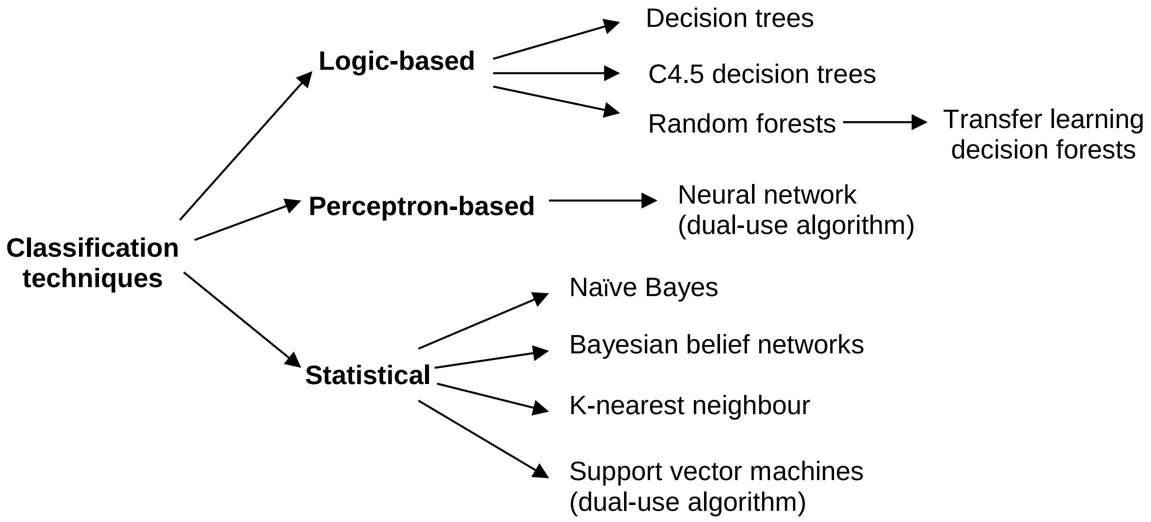

Figure 8: Types of classification machine learning technique (adapted from Kotsiantis, 2007:251)

Logic-based techniques use acquired knowledge from examples to classify data (Lopez De Mantaras & Armengol, 1998:99), whereas according to Kotsiantis (2007:254), perceptron-based techniques are based on the ability of a perceptron to determine the weights to assign to each identified feature and calculate the appropriate class using the features and assigned weights.

Also of importance here is Kotsiantis' (2007:257) description of statistical techniques as those techniques that make use of a probability model to determine the probability that an instance belongs to a particular class. The first of the techniques to be described will be the logic-based technique of the decision tree.

3.5.1.1 Logic-based: decision trees

The decision tree consists of nodes, each containing a question which relates to a particular feature. The algorithm starts at the root node, determines which features are present for that root node question and the, depending on the answer, moves on to the next node. The information required to train these decision trees takes the form of instances, which consist of a set of features.

The instance moves along the branches to assess different features at each node, ending at a leaf node. A leaf node is a group of features that are labelled as a particular class; the instance is then classified using the label assigned to that particular leaf (Kotsiantis, 2007:251; Marsland, 2009:133; Thomassey & Fiordaliso, 2006:410).

The decision tree may be considered as a set of if–then statements or rules. These rules are determined by the algorithm based on the training set (Samoil, 2015:35). Table 3 provides an example of a training set to which a decision tree can be applied.

Table 3: An example of a decision tree training set

Features				Class
Ft1	**Ft2**	**Ft3**	**Ft4**	**Label**
a1	a2	a3	a4	Yes
a1	a2	a3	b4	Yes
a1	b2	a3	a4	Yes
a1	b2	b3	b4	No
a1	c2	a3	a4	Yes
a1	c2	a3	b4	No
b1	b2	b3	b4	No
c1	b2	b3	b4	No

(Adapted from Kotsiantis, 2007:251)

Figure 9 presents an example of a decision tree for the training set example in Table 3 above. In the decision tree structure, the features are considered one at a time, followed by the assignment of a class or the consideration of another feature.

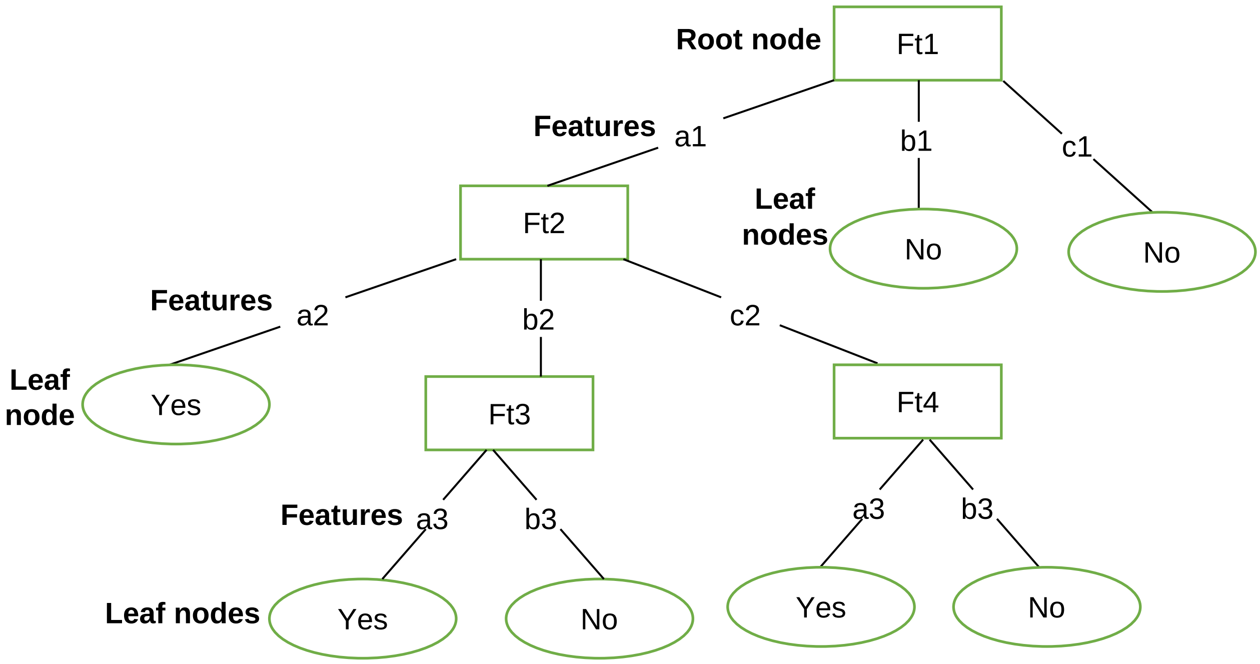

Figure 9: An example of a decision tree (adapted from Kotsiantis, 2007:251)

As can be seen from Figure 9 above, the classification consists of a number of decisions which occur at each node, ending at the leaf node. The leaf node does not require a decision but rather assigns the instance to the particular class label (Narasimha Murty & Susheela Devi, 2011:127). A more advanced decision tree is the C4.5 algorithm, which is described below.

3.5.1.2 Logic-based: C4.5 decision trees

In producing the ID3 decision tree algorithm, the basic decision tree algorithm is adapted to ensure that the correct features are selected at each stage of the tree. The ID3 algorithm ranks features in such a way that the more informative features are closer to the root. This is measured using "entropy", a term used to describe how informative a feature is (Thomassey & Fiordaliso, 2006:410).

The ID3 algorithm is further adapted to the C4.5 decision tree algorithm by pruning it to reduce the number of nodes without losing the ability to classify the instance. There are two types of pruning, prepruning and post pruning. Prepruning tries to determine when to stop building branches by assessing at which point enough features have been considered to reasonably classify the case without requiring further branches.

Post pruning takes place after an entire tree has been built, thus branches are removed at the end. C4.5 uses post pruning, because it takes an ID3 tree, converts it into a set of if–

then rules and subsequently prunes certain conditions if the accuracy of the rules is increased without them (Marsland, 2009:143; Thomassey & Fiordaliso, 2006:410).

3.5.1.3 Logic-based: random forests

This model consists of a number of decision trees, each composed of a subsample of features, and is usually weaker than a full decision tree. The average, or the weighted average, of the trees is determined and used to perform the classification, effectively combining the power of the individual trees which often produces a higher quality result (Bucheli & Thompson, 2014:4; Dataiku, 2017:7).

3.5.1.4 Logic-based: transfer learning decision forests

This model uses random forests, as described above, where the knowledge produced can subsequently be applied or transferred to a given target task. This generates a classifier that can be used to exploit the knowledge from other tasks to improve the ability of the classifier to perform a target task (Goussies, Ubalde, Fernandez & Mejail, 2014:4312).

3.5.1.5 Perceptron-based: neural networks

A neural network can provide prediction and classification solutions and is discussed in section 3.5.3.2.

3.5.1.6 Statistical: Naïve Bayes

The Naïve Bayes algorithm is a probabilistic model which determines the probability of different classes or outcomes, based on previously encountered examples. These examples are identified in the training data. For an instance or event that has been classified, the algorithm calculates the probability of each identified feature present in the instance and these probabilities are multiplied with each other for all possible classes or outcomes. The class or outcome with the highest probability is chosen, being the most likely outcome (Larsson & Segerås, 2016:12).

The Naïve Bayes algorithm is derived from the Bayesian theorem and assumes that the features in the instance are independent, which implies that the value of the features do not influence one another. When considering multiple features, the Naïve Bayes algorithm is a more simplified algorithm than the Bayesian theorem (Marsland, 2009:171).

The Bayesian theorem is used to calculate the probability of an event based on previous knowledge of the probability of an event. The algorithm determines the probability of the occurrence of, for example, event A, when having observed event B, taking into account the likelihood of event B being present when event A is observed, as well as the probabilities of event A and event B as follows:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

The variables in the above algorithm are as follows:

- P(A|B) is the probability of occurrence of event A, given that event B (evidence) occurred
- P(A) is the prior probability of event A
- P(B|A) the conditional probability or likelihood of the occurrence of event B, given event A
- P(B) is the prior probability of the occurrence of event B (Chadha & Singh, 2012:52; Samoil, 2015:16).

This algorithm can be used in a normal classification problem where there are multiple classes, say C_1, C_2, …, C_k. Naïve Bayes calculates the conditional probability, that is, the probability of a feature conditional on the observed features of an object with a set of multiple features (observed features) such as x_1, x_2, …, x_n belonging to a particular class C_i.

The algorithm calculates the probability of the class, given the observed features, by multiplying the probability of the features by the probability of the class, using the Bayes theorem as follows (Narasimha Murty & Susheela Devi, 2011:93):

$$P(C_i|x_1, x_2, \dots, x_n) = \frac{P(x_1, x_2, \dots, x_n|C_i) * P(C_i)}{P(x_1, x_2, \dots, x_n)}$$

To illustrate how this would be used to classify something based on its features, the example of classifying vegetables adapted from Larsson and Segerås (2016:12) is provided. Table 4 contains the training data followed by an example of the observed features and how they are used to classify the vegetables using Naïve Bayes.

Table 4: Naïve Bayes training data

Vegetable class	Long	Not Long	Purple	Not Purple	Total
Tomato	5	24	3	26	29
Brinjal	14	4	16	2	18
Total	19	28	19	28	47

Source: **Adapted from Larsson & Segerås (2016:12)**

To classify an unknown vegetable having the observed features or evidence Long and Not Purple, the probability of each respective vegetable given the features will need to be calculated, as shown in formula (1) and formula (2).

$$P(Tomato|Long\ and\ Not\ Purple) = \frac{P(Long|Tomato)*P(Not\ Purple|Tomato)*P(Tomato)}{P(Long)*P(Not\ Purple)} \quad (1)$$

$$P(Brinjal|Long\ and\ Not\ Purple) = \frac{P(Long|Brinjal)*P(Not\ Purple|Brinjal)*P(Brinjal)}{P(Long)*P(Not\ Purple)} \quad (2)$$

The probabilities calculated in formula (1) and formula (2) will be compared, and the highest probability will be selected as the chosen vegetable.

$$P(Tomato|Long\ and\ Not\ Purple) = \frac{{}^{5}/_{29} * {}^{26}/_{29} * {}^{29}/_{47}}{{}^{19}/_{47} * {}^{28}/_{47}} = \frac{0.0954}{0.2408} = 0.3961 \quad (1)$$

$$P(Brinjal|Long\ and\ Not\ Purple) = \frac{{}^{14}/_{18} * {}^{2}/_{18} * {}^{18}/_{47}}{{}^{19}/_{47} * {}^{28}/_{47}} = \frac{0.0331}{0.2408} = 0.1374 \quad (2)$$

Therefore, a long and not purple vegetable will most likely be a tomato based on the probability of the prior observed evidence using Bayes theorem and assuming that the features are independent.

3.5.1.7 Statistical: Bayesian belief networks (BBN)

Like the Naïve Bayes algorithm, a Bayesian belief network is modelled on the Bayesian theorem. However, in contrast to the Naïve Bayes algorithm which assumes that features are independent, a Bayesian belief network takes into consideration the probabilistic dependencies among features (Heckerman, 2008:33; Narasimha Murty & Susheela Devi, 2011:97).

Bayesian belief networks plot these probabilistic relationships using a graphical model. The graph presents a network of nodes, one for each feature, which are connected by lines going in a specific direction. The one feature needs to be present for the possibility of the other, therefore they are causally linked and the graph indicates this parent–child relationship. This graph is known as a directed acyclic graph (Witten, Frank, Hall & Pal, 2016:340) and is illustrated in Figure 10.

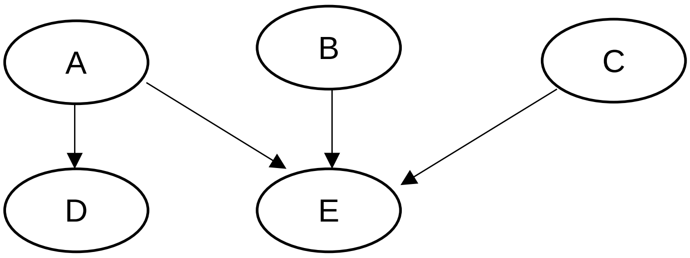

Figure 10: A Bayesian belief network showing causal relationships between events (adapted from Heckerman, 2008:45)

The model takes these dependencies into consideration by determining the joint probability of causal features. This is done by using the product rule, which states that the probability of both event A and event B is the probability of event A, given B multiplied by the probability of B (Witten *et al.*, 2016:337).

$$\therefore P(A,B) = P(A|B) * P(B)$$

By using this product rule and taking into account the causal relationship between features, the joint probability distribution for a set of features can be determined. The Bayesian belief network is able to identify conditional independencies; for example, D is conditionally independent of E in Figure 10. Therefore, the conditional probability of event D, given events A, B and E, will be as follows:

$$P(D|A,B,E) = P(D|A,B).$$

The conditional probability for each node X will be calculated as the joint probability of the parent events, that is, P(X|Parent events (X)), therefore excluding conditional independencies (Heckerman, 2008:45; Witten *et al.*, 2016:343).

Having established a Bayesian belief network, the probabilities of causal events can be determined or inferred using the probabilities of the observed events, taking into account only the joint probabilities and excluding the conditional independencies (Heckerman, 2008:46).

3.5.1.8 Statistical: *K*-nearest neighbour (kNN)

The nearest neighbour algorithm classifies instances or patterns according to the nearest known neighbour class by finding similarities in the instance being classified to patterns or features in the training set (Narasimha Murty & Susheela Devi, 2011:48). The *k*-nearest neighbour algorithm is derived from this nearest neighbour algorithm.

For the *k*-nearest neighbour algorithm, instead of only one nearest neighbour being detected by the algorithm, more than one to the amount of *k* nearest neighbours may be detected. The nearest neighbours detected may then indicate that the instance being classified is near to more than one class. The class for which the majority of nearest neighbours are identified will then be the class to which the instance is classified.

Figure 11 illustrates a linear example of a *k*-nearest neighbour classification. Where *k* is 3, three of the nearest neighbours will be identified. Instance P can therefore be correctly classified as belonging to the circles class, as the majority of the nearest neighbours are circles. If the basic nearest neighbour algorithm was used, then P would have been incorrectly classified in the class of the stars, even though point number 5 is an outlier (Narasimha Murty & Susheela Devi, 2011:51).

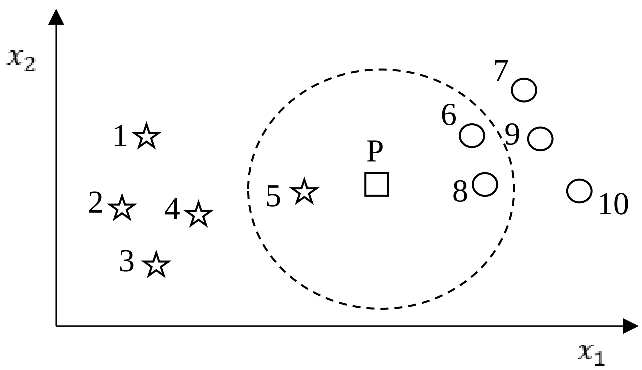

Figure 11: Accurate classification of P using the *k*-nearest neighbour algorithm (adapted from Narasimha Murty & Susheela Devi, 2011)

In the design of the *k*-nearest neighbour algorithm, an important factor is therefore the value of *k*. This will need to be an odd value in order to avoid an even split between two classes. The *k*-nearest neighbour algorithm may also be modified to assign a weight to each of the nearest neighbours based on their proximity to the instance being classified (Narasimha Murty & Susheela Devi, 2011:51).

3.5.1.9 Statistical: support vector machines

A support vector machine can provide prediction and classification solutions and is discussed in section 3.5.3.1.

3.5.2 Prediction algorithms

3.5.2.1 Conditional random fields

Where normal classifiers predict only one class at a time, conditional random fields use a graphical model to plot many interdependent variables, thus determining the conditional distribution for multiple predictions (Sutton & Mccallum, 2007:98; Witten *et al.*, 2016:407). According to Sutton and Mccallum (2011:269), conditional random fields are able to predict outputs by combining discriminative classification with graphical modelling.

As opposed to a generative model, which determines a joint probability distribution, conditional random fields use a discriminative model which calculates only the conditional distribution. A conditional distribution $p(y|x)$ where the probability of outputs y is calculated given specific inputs x, does not consider the modelling of the probability of the input $p(x)$ that would consider all the features dependent on x. Discriminative models are therefore simpler models than generative models (Sutton & Mccallum, 2011:269).

Conditional random fields can determine the probability of possible label sequences which are interdependent given an observation sequence, as recommended by Lafferty, Mccallum and Pereira (2001:282) for segmenting and labelling sequence data, thus taking context into account when predicting the outputs. This context provides information that contributes to predicting the outputs (Witten *et al.*, 2016:406).

For this reason conditional random fields are useful for modelling multifaceted outputs consisting of interrelated parts, for example for identifying the parts of a sentence (Sutton & Mccallum, 2007:106) or image captioning where the components of the sentence or picture are related and are indicative of what is being classified (Arnab, Zheng, Jayasumana, Romera-Paredes, Larsson, Kirillov, *et al.*, 2018:38).

3.5.3 Dual-use algorithms: classification and prediction

3.5.3.1 Support vector machines

A support vector machine is a binary classifier that aims to separate data into two classes, based on the case features. A set of features that describe one case is known as a vector. In order for a support vector machine to perform at its best, the optimal hyper plane needs to be identified which separates the two classes of vectors.

The optimal hyper plane would be one which has a maximum margin from each of the classes, where this margin is the distance from the hyper plane to the closest vectors in each class. This maximum margin is determined by considering the vectors closest to the hyper plane in each class. The closest vectors are known as the support vectors.

The support vector machine firstly maps the vectors into a multidimensional (*N*-dimensional) space and then determines the hyperplane which separates the vectors into two classes. New instances can then be classified into these two classes (Ayodele, 2010b:25).

Although N-dimensions are used to map vectors, for illustration purposes Ayodele (2010b:26) recommends using a two-dimensional example as shown in the Figure 12. In the example there are two features for each vector, one represented by the X_1 axis and one represented by the X_2 axis. The vectors are plotted accordingly and each class is represented by a shape, resulting in two classes, circles and squares. The maximum margin is used to determine the hyperplane, which is at the optimal distance from the support vectors.

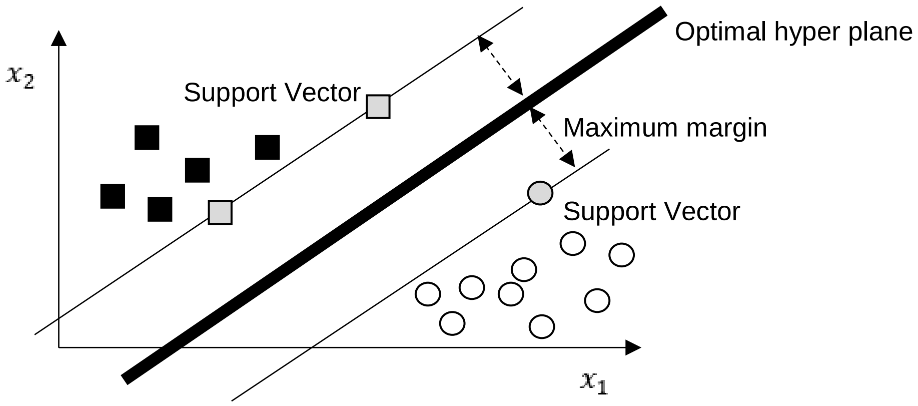

Figure 12: A two-dimensional example of a support vector machine (adapted from Ayodele, 2010b:26; Kotsiantis, 2007:261)

Despite the linear hyperplane shown in the two-dimensional figure, support vector machines usually create non-linear class boundaries, which increases the possibilities for which the support vector machines can be used (Witten *et al.*, 2016:252).

3.5.3.2 Artificial neural networks

An artificial neural network is a combination of mathematically generated neurons, which operate in a similar manner to the human brain. These neurons are each assigned a weight based on what the artificial neural network learns, collectively forming part of a mathematical function (SMACC, 2017:9).

Narasimha Murty and Susheela Devi (2011:169) describe the functioning of a neuron as follows: The neuron receives input from a combination of other neurons, for which each input carries a specific weight. The total inputs received are summed by adding all the input values adjusted for their respective weights. If the cumulative input exceeds a threshold, the neuron will produce an output. Figure 13 illustrates an artificial neuron, where the output value will be calculated as follows: $a = x_1w_1 + x_2w_2 + x_3w_3$.

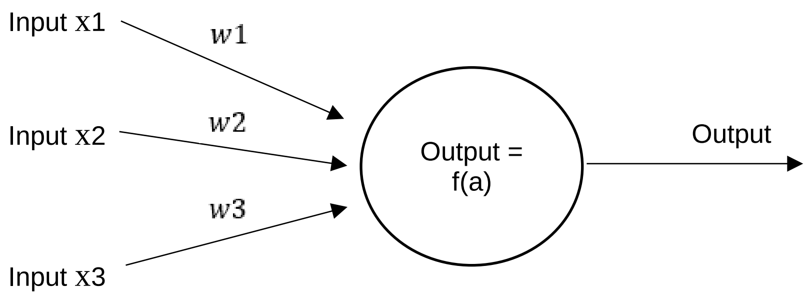

Figure 13: A single neuron (adapted from Narasimha Murty & Susheela Devi, 2011:169)

According to Narasimha Murty and Susheela Devi (2011:169), the training of a neural network takes place as follows:

- Random weights are assigned to every link or neuron in the network.
- Inputs are given to the artificial neural network for which an output is produced.
- If the output is correct, then no changes are made to neuron weights.
- If the output is incorrect, the error is used to adjust the weights in the network and the process repeated until the correct output is produced for those inputs.

The structure of the neural network, and therefore the arrangement of the neurons, will depend on the learning task and the data available (SMACC, 2017:9). A network consisting of an input and an output layer will be able to classify linearly separable classes – this is known as a feed-forward network.

For more complex classification, a multilayer network will be required. In addition to an input and an output layer, these networks have a hidden layer which enables non-linear classification. The structure of this network is illustrated in Figure 14 (Narasimha Murty & Susheela Devi, 2011:174).

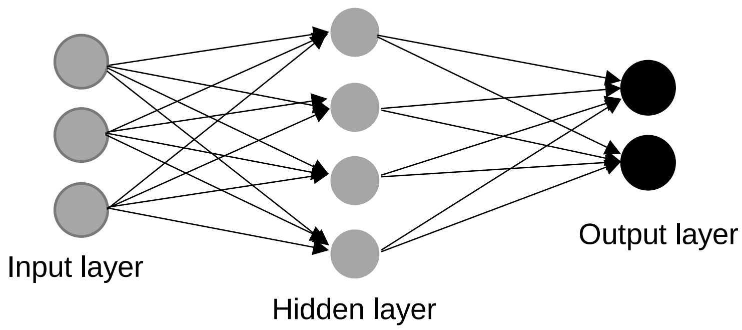

Figure 14: A multilayer neural network with a hidden layer (adapted from Narasimha Murty & Susheela Devi, 2011:174)

The artificial neural network can be used to extract information from unfamiliar documents, using its knowledge to determine the probability of the identified information being similar to known information. OCR is used to extract the information and the characteristics of the information are used to determine what the information is. SMACC (2017:10) provides an example of how the artificial neural network extracts information from an invoice, as shown in Figure 15.

Invoice example

The invoice company
Fictitious Street
Stellenbosch
19582

Further invoices details XXX

The number *19582* appears on the invoice, what does this number mean?

A word above the number: *Stellenbosch 19582*
So the word provides context: based on probability it is a postal code

Line above: *Fictitious Street*
This may be an address, which increases the probability of 19582 being a postal code

The address is also located below a company name, which further increases the probability of it being a postal code.

Figure 15: An artificial network classifying elements of an invoice (adapted from SMACC, 2017:10)

Neural networks can be trained using supervised learning, since both the inputs and outputs are supplied to the neural network. However, neural networks can also be designed using unsupervised learning, which would require self-organisation, as is seen in self-organising maps (Hadzic, Dillon & Tan, 2007:225; Kohonen, 1990:1464).

3.5.3.3 Convolutional neural networks

Convolutional neural networks are explained by Albawi, Mohammed and Al-Zawi (2017:1) as a particular type of deep neural network consisting of multiple layers. These layers are structured in such a way that the network is able to handle complex data such as images, and to classify these images based on the combined features identified by the different layers of the network. The data input passes through layers that include convolutional, pooling and fully connected layers, resulting in an output.

When classifying an image, a convolutional neural network considers each pixel as an input. Unlike in other neural networks, the neurons in the convolutional layer do not connect to all the inputs (pixels), they only obtain input from regions in the picture. This is done by breaking the image down into smaller pieces, consisting of sets of pixels, and then systematically connecting the neurons to these smaller pieces. Therefore the same set of neurons is re-used to detect portions of the picture piece by piece (Albawi *et al.*, 2017:2; Witten *et al.*, 2016:438).

The convolutional layer passes each piece of the image through a set of filters where each filter looks for different aspects in the image. Different filters detect different features of an image such as edges and shapes or higher-level features such as faces. Each filter needs to be trained to effectively perform its given task, thus these filters are prepared by means of supervised learning.

The elements detected in the convolutional layer are then reduced in the pooling layer by means of a threshold. As each of the filters are neurons, each neuron will produce a weighted output. These outputs will also be analysed systematically by dividing outputs into sub-regions. The features in each sub-region with the maximum weighted output will be selected by the pooling layer – this is called maximum pooling. In this way the pooling layer summarises the outputs of the convolutional layer and makes the features detected more robust (Albawi *et al.*, 2017:5; Krizhevsky, Sutskever & Hinton, 2012:4) .

There may be multiple layers of convolution and pooling layers, which each filter and measure different aspects such as edges, curves, faces and hands and so on. All successful detections are then combined by the fully connected layer, which uses logic to identify the image being classified. It is important to note that features in the image cannot be spatially dependent. So, for example, in a face detection application, it should not matter where the faces are located in the images, the focus is simply on detecting the faces (Albawi *et al.*, 2017:1).

Figure 16 illustrates the functioning of a convolutional neural network, where the convolutional layer interprets sections of the image at a time, followed by a reduction of the detections in the pooling layer and then finally a classification of the findings by the fully connected layer (Krizhevsky *et al.*, 2012:5; Lawrence, Giles, Tsoi & Back, 1997:103)

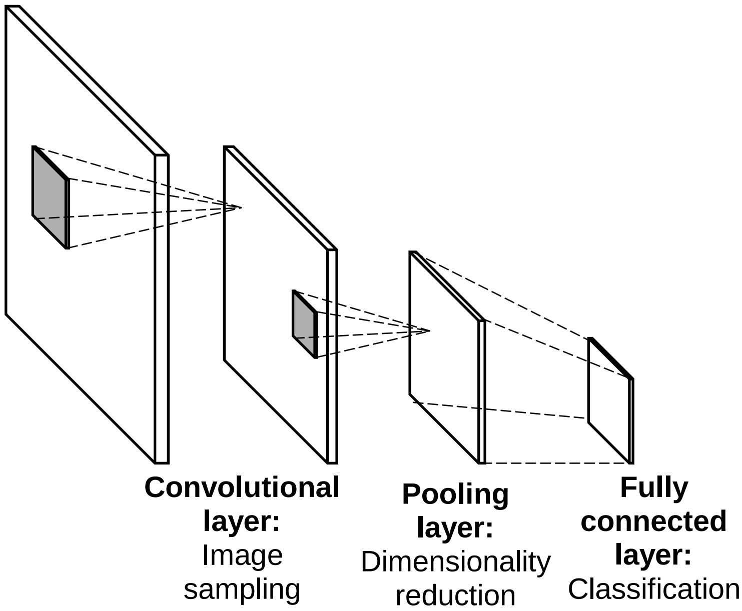

Figure 16: A high-level diagram of the convolutional neural network used for image classification (adapted from Krizhevsky *et al.*, 2012:5; Lawrence *et al.*, 1997:103)

Oquab, Bottou, Laptev and Sivic (2014:1717) indicate that convolutional neural networks used to classify images can also be used in combination with other technologies to determine document types, when sorting documents by their appearance or patterns.

3.6 Description of unsupervised learning techniques

Having described the various supervised learning techniques in section 3.5, this section describes the relevant unsupervised learning techniques.

3.6.1 Pattern detection

3.6.1.1 Association rules

Association rules determine the associative relationships between data, where the occurrence of one feature may indicate the possible occurrence of another feature (Narasimha Murty & Susheela Devi, 2011:55). Instead of predicting a particular class, association rules are able to predict combinations of features and which features are commonly associated with each other, irrespective of class (Witten *et al.*, 2016:79).

Association rules need to be measured in order to determine whether they can be relied upon, and to do this the coverage of the association rule is considered. This is known as the support and thus the accuracy of the association rule is determined, which is called the confidence (Witten *et al.*, 2016:79).

Berka and Rauch (2010:11) provide an example of how the support and confidence can be determined. In the example, a is the number of classes where both X and Y are present, b are the classes where only X is present, similarly c where only Y is present and d where neither classes have X or Y.

The support is determined as follows:

$$P(X \wedge Y) = \frac{a}{a+b+c+d}$$

The confidence is determined as

$$P(X|Y) = \frac{a}{a+b}$$

Usually the support and confidence for the association rule will need to exceed a specified minimum threshold for the rule to be considered for pattern detection. Since some associations may be indicative of others, often only the strongest rules are selected in order to reduce the number of rules (Witten *et al.*, 2016:79).

3.6.2 Clustering

Clustering is an unsupervised machine learning method which divides instances into groups or clusters. The following different groups may be identified:

- **Exclusive** – each instance belongs to only one cluster.
- **Non-exclusive** – one instance may belong to more than one cluster.
- **Probabilistic or fuzzy** – there is a certain probability or degree of membership of each cluster to which an instance belongs.
- **Hierarchical** – instances are divided into high-level broader clusters, each of which are refined into smaller subclusters up to individual instances level (Thomassey & Fiordaliso, 2006:411; Witten *et al.*, 2016:88).

Clustering is often the first stage of a hybrid approach, which consists of more than one machine learning technique. The next stage may be a decision tree or rule set which is derived based on the features of the determined clusters, and this rule set then allocates new instances to the appropriate clusters (Thomassey & Fiordaliso, 2006:413; Witten *et al.*, 2016:88).

3.6.2.1 Self-organising maps

Self-organising maps are a form of neural network that uses unsupervised learning. The objective of the self-organising map is to produce its own representation or self-organisation of the given data, since outputs are not provided (Hadzic *et al.*, 2007:225; Kohonen, 1990:1464).

According to Ayodele (2010b:37), the self-organising map aims to learn the structure of the data by identifying clusters of data and linking similar clusters to each other. This results in feature mapping, where neurons representing similar features are located close to each other in a network.

In order to form the network, the neurons of self-organising maps are typically organised into a two-dimensional grid, with connections between the neurons in the grid (Ayodele, 2010b:37; Marsland, 2009:208). This forms the output map. The inputs form another layer, as is the case with neural networks. Each node in the input layer is fully connected to the output map neurons (Ehsani, Quiel & Malekian, 2010:411). This is illustrated in Figure 17.

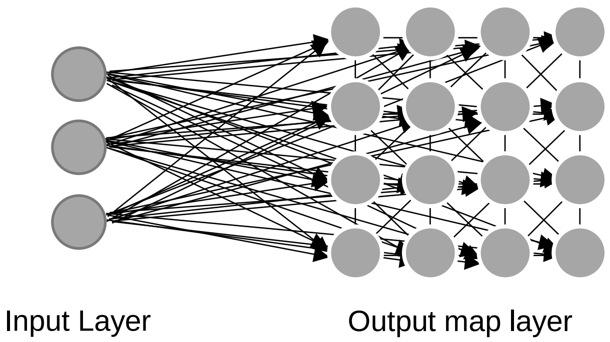

Figure 17: The self-organising map network (adapted from Hadzic *et al.*, 2007:228)

The objective the self-organising map is to plot input patterns onto a self-organising map. Warwick (2012:97) describes the self-organising map model as follows:

1. Self-organising maps start with an **untrained map**, which consists of any number of neutrons arranged in a grid. This grid usually represents two dimensions, even though the input patterns are usually more highly dimensional than 2D, so a self-organising map enables dimensionality reduction (Kohonen, 1998:1; Marsland, 2009:208).

2. Between the neurons in the network are **lateral connections,** which means that the output from each neuron forms further inputs to each of the other neurons in turn. This is in contrast to artificial neural networks where connections only occur between the different layers of inputs and outputs and not within the same layer (Marsland, 2009:208).

3. Each of the signals from the inputs and from the surrounding neurons will have a **weight** attached to them, which may be randomised to begin with between 0 and 1. The location of each neuron in the map is important as neurons located close to each other need to respond to similar input patterns (Marsland, 2009:210).

4. The self-organising map is organised using **competitive learning,** where the neurons compete to best represent the input data (Ayodele, 2010b:39; Hadzic *et al.*, 2007:228; Marsland, 2009:210).

5. When a particular **input pattern** (feature combination) is presented to the algorithm, one of the neurons will present an output which is higher in weight than the other neurons (different features trigger different weights resulting in an answer for each neuron). The neuron with the highest response is considered the winning neuron.

6. The weights of the **winning neuron** are adjusted so that its weight for that input pattern is higher than before. The weights of the neighbouring neurons are also adjusted upwards but to a lesser degree for that particular input pattern. The neurons therefore adapt to different input patterns, thus ensuring they will be able to recognise and classify these input patterns (Kohonen, 2013:52).

7. After various adaptations, neighbouring neurons are organised onto the two-dimensional map, representing similar features as is required for **feature mapping**. This effectively creates clusters of similar neurons (Kohonen, 1998:1; Marsland, 2009:208).

8. Each time the self-organising map is presented with an input pattern a specific region of the neuron grid will respond. And since neurons are grouped according to their similarities for identifying a specific input pattern, by means of feature mapping, the

response of a particular region of the neurons will indicate the type of input pattern being presented, thus enabling classification or identification of the input.

3.6.2.2 K-means clustering

K-means clustering divides the data into *k* number of categories. In order to perform k-means clustering, the number of clusters, that is *k*, needs to be specified and a random initial central data point (centroid) needs to be selected for each cluster. The data is then grouped based on the distance of each data point from the initial centre.

Once initial assignment of data to clusters and distances has been done, the mean distance from the central data point for each cluster can be calculated and the centre moved to the mean distance point.

The algorithm runs again until the cluster centres no longer need to move (Ayodele, 2010b:27; Marsland, 2009:196). Figure 18 represents an example of a K-means flow diagram (Ayodele, 2010b:28).

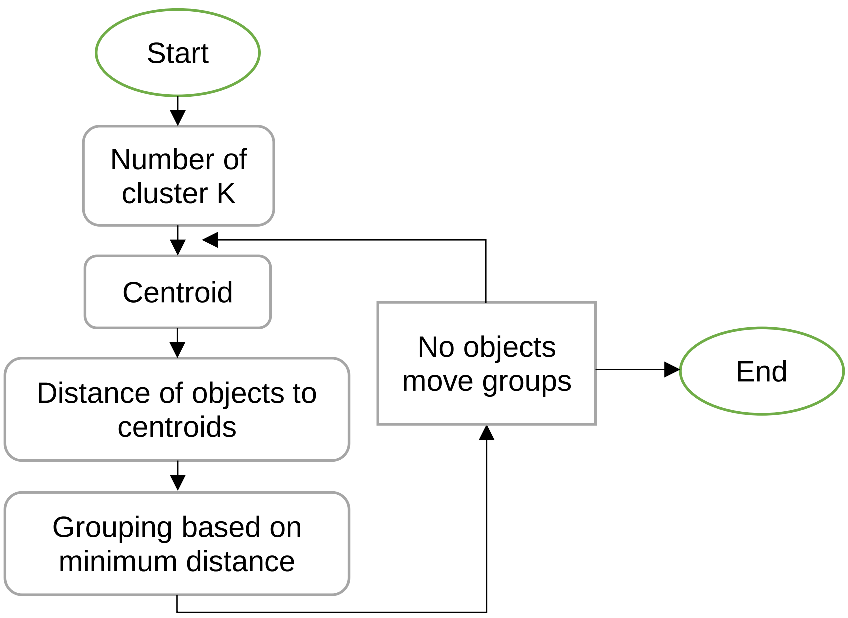

Figure 18: An example of K-means follow diagram (adapted from Ayodele, 2010b:28)

The steps in a K-means clustering algorithm are as follows:

1. Choose a value for K.
2. Determine the centre coordinate or centroid (Takaki, Petersen & Ericson, 2018) for each group.
3. Calculate the distance of each data object from each centroid.
4. Place data points in clusters based on closest centroid, look at minimum distance.
5. Calculate the mean distance of cluster from the centroid.
6. Move the centroids for each group to the mean distance point.
7. Repeat grouping of clusters, calculation of mean distance, and movement of centroids until centroids no longer need to move.

3.7 Description of the semi-supervised learning techniques

Having described the various unsupervised learning techniques in section 3.6, this section describes the relevant semi-supervised learning techniques.

3.7.1 Semi-supervised clustering

As Jain, Jin and Chitta (2014:1) state, clustering algorithms are unsupervised machine learning algorithms that work to find a partition in the dataset. Semi-supervised clustering assists the algorithm to find a better quality partition by providing the algorithm with any prior knowledge about the data. The clustering algorithm is then guided by the prior knowledge to find the partition in the data.

Prior knowledge refers to labelled data, as is required for supervised learning, whereas unsupervised learning algorithms function using unlabelled data (Zheng, Zhou, Deng & Yang, 2017:7447). The labelled data may indicate constraints such as which data must or cannot be clustered together, or prior knowledge may mean increasing or decreasing known distances between data points (Bezerra, Mattoso & Xexéo, 2006:88).

Having described the different machine learning techniques, certain tasks which are considered problem areas to be addressed by machine learning are described in the next section, together with the applicable machine learning technique that can be applied to each of the identified learning problems.

3.8 Tasks that can be addressed by machine learning technology

In order to achieve the objectives of this section the learning problem for each identified task was determined using the findings of prior research. This enabled the identification of the machine learning technique that could be applied. The process of selecting machine learning techniques to apply to identified problems has been proven to be significantly difficult, especially when matching the problems identified with techniques intended to solve them (Someren & Urbancic, 2006:365).

The process of matching each problem to the technique to solve it involves firstly understanding the task and then defining the learning problem. Identifying the learning problem enables a developer to identify the information and algorithm required to address the problem (Saitta & Neri, 1998:137; Someren & Urbancic, 2006:366). This process of understanding the task and defining the learning problem is crucial as there are often many solutions available for addressing a learning problem (Someren & Urbancic, 2006:369).

The different types of learning solutions that are available are defined as follows (Amani & Fadlalla, 2017:34):

- **Classification** is suitable for mapping data into two or more categories, each with its own distinct attributes (Larsson & Segerås, 2016:11).
- **Clustering** is suitable for separating data into classes or groups that are similar in some meaningful way (Larsson & Segerås, 2016:11).
- **Prediction** is suitable for producing a forward-looking numerical prediction (forecasting) or non-numerical prediction (classification).
- **Outlier detection** is suitable for finding the items or events that significantly deviate from the expected pattern or other data considered normal in the dataset.

Having defined the different types of learning solution above, Table 5 sets out the different problem areas and then designates the learning solution type for each area. The problems are organised by task areas, as described in chapter 2. The key findings of existing machine learning techniques identified in research to address the specific learning problem areas are then provided.

Table 5: Accounting problem types and recommended machine learning techniques

Description of the learning problem	Solutions to the learning problem	Machine learning techniques	Source
Translation of manual and electronic documents into accounting information			
Task 3. Document features extraction			
OCR requires adaptability, a characteristic which can be enabled by means of machine learning.	Classification	Transfer learning decision forests	ABBYY Technologies, (n.d.); Goussies *et al.* 2014:4309)
Task 4. Document type recognition and classification(			
Image classification can be used to detect the document type, which can be enhanced by means of machine learning.	Classification	• Convolutional neural networks • New document class: k-nearest neighbour • Similar known documents: support vector machine	ABBYY (2017); Oquab *et al.* (2014:1717); Sorio (2013:23); Sorio, Bartoli, Davanzo & Medvet (2010:187); Witten *et al.* (2016:523)
Irregular document layout classification using NLP combined with machine learning to train the system to process flexible or irregular document layouts.	Classification	Convolutional neural networks	ABBYY (2017); Chen, Wang, Fan, Sun, & Satoshi (2015:436)
Text classification is used to classify text, using both statistical and semantic text analysis.	Clustering	• Parallelisation MapReduce k-nearest neighbour • Semi-supervised clustering	ABBYY (2017); Du (2017:195); Zhang, Tang, & Yoshida, (2015:152)
Task 6. Validation of document data			
Validation of document information can apply machine learning to determine	Classification	• Naïve Bayes • Support vector machine	Larsson & Segerås (2016:33)

Description of the learning problem	Solutions to the learning problem	Machine learning techniques	Source
whether the extracted data from the document is correctly classified or not.			
Removing of duplicate entries and linking documents may be achieved by using approximate string matching, making use of machine learning for string classification.	Classification	• Naïve Bayes • Decision trees • Support vector machine • Artificial neural network	Amtrup, Thompson, Kilby & Macciola, (2015:24); Larsson & Segerås (2016:18); De Leone & Minnetti (2015:2); Samoil (2015:16)
Reconciliation of financial information			
Task 3. Matching			
Matching records or record-linkage has been performed using a variety of machine learning techniques.	Classification	• Naïve Bayes • Decision trees • Support vector machine • Artificial neural network	Chew & Robinson (2012:328); Samoil (2015:16)
Preparation of management accounts			
Task 3. Account allocation			
Account allocation may be performed by incorporating machine learning, which learns to predict the account allocation based on probability and can recommend which accounts to post to.	Classification and clustering	**Classification:** Naïve Bayes **Clustering:** K-means clustering Random forests	Bengtsson & Jansson (2015:40); Brady Leider, Resnick, Natalia Alfonso & Bishai (2017:354); SMACC (2017:12); Takaki & Ericson (2018:1)
Task 7. Report generation			

Description of the learning problem	Solutions to the learning problem	Machine learning techniques	Source
Error detection in financial data and fraud detection can be performed by incorporating machine learning to identify irregularities in datasets.	Classification Outlier detection and clustering	**Classification**: • Bayesian belief network and a decision table • Naïve Bayes hybrid model **Outlier detection**: • Association rules **Clustering**: • K-means clustering • Self-organising maps	Ahmed, Mahmood & Islam (2016:283); Alpar & Winkelsträter (2014:2261); Hajek & Henriques (2017:146); Kokina & Davenport (2017:117)
Task 8. Report descriptions			
Report descriptions may incorporate machine learning techniques in natural language generation technologies to enable a reasoning process to be applied to the reported data and thus produce required explanations in natural language.	Prediction	Conditional random fields	Gardent & Perez-Beltrachini (2017:15); Lafferty *et al.* (2001:283); Yseop (2017:7)
Tasks pertaining to performance measurement			
Forecasting performance reports can be prepared using predictive analytics, which may employ machine learning algorithms. These predictive analytics can be used for forecasting the business's financial performance.	Prediction	• Support vector machine • Artificial neural network • C4.5 decision trees • Bayesian belief network	Appelbaum *et al.* (2017:36)

Source: Own observation

Table 5 above demonstrates which accounting processes tasks can be performed or assisted by machine learning techniques. These techniques were described in sections 3.5, 3.6 and 3.7 of this chapter. It is clear from this table that there may be more than one method available to address a specific task. It is important to note that in certain instances a combination of methods may be selected (Someren & Urbancic, 2006:380).

3.9 Conclusion

This chapter provided a context, framework and recommended architecture for machine learning technology, as well as a description of the functioning of various machine learning techniques that can be applied to the different accounting tasks.

The functioning of the different machine learning techniques was described as this contributes to accounting users' understanding of machine learning technologies and it supports users in understanding the limitations of such technologies. These limitations are discussed in chapter 4 and, it is hoped, will ultimately assist users in selecting machine learning technologies that are appropriate for their needs.

The findings in this chapter were presented as a table, indicating the different tasks to which the various machine learning techniques can be applied and the accompanying machine learning techniques that are available to perform these tasks. The benefits of these techniques and the limitations to their use will be discussed in chapter 4 section 4.7.

Chapter 4 also describes the risks, benefits and limitations of the use of machine learning technology in an accounting context. Subsequently, guidelines for implementing machine learning technology in an accounting context, taking into account these risks and benefits, are provided in chapter 5.

Chapter 4: Risks, benefits and limitations when implementing machine learning

4.1 Introduction

Chapter 3 placed the different machine learning techniques in context and discussed which of the techniques could be applied to the accounting tasks identified in chapter 2. In this chapter the risks, benefits and limitations of machine learning technology, as well as the benefits and limitations of the identified machine learning techniques, are discussed.

This chapter identifies the risks, benefits and limitations in relation to machine learning in line with King IV (Institute of Directors of Southern Africa (IODSA), 2016:30), which states that while risks may be negative they also inherently present certain opportunities, which could support the business in achieving its objectives. The principles of King IV were therefore considered in preparing this chapter.

This chapter is structured in terms of the stages of the software development life cycle and the data science life cycle, both of which were considered when evaluating the different risks, benefits and limitations of machine learning, as according to Sapp (2017:15), the data science life cycle often overlaps with machine learning. Furthermore, the stages of this combined life cycle in which the accounting user has to be involved were considered.

The findings of this chapter, being a summary of the risks, benefits and limitations of machine learning and the life cycle stages in which the accounting users need to be involved, are addressed in chapter 5, where guidelines for implementing the machine learning techniques in an accounting context are provided.

4.2 Machine learning technology risks pertaining to the accounting objectives

Gillion (2017:8) states that in all businesses the objective of accounting processes is to produce high quality accounting information for decision-making. As part of identifying the risks in the planning phase of the life cycle, the objectives of the technology were considered within the context of the accounting process, that is, the accounting objectives.

The framework selected for identifying the accounting objectives was the *Conceptual Framework for Financial Reporting*, as approved by the International Accounting Standards Board (2018:6). This framework describes the objectives of financial reporting by providing a description of the qualitative characteristics of financial information produced by the accounting processes (International Accounting Standards Board, 2018:14).

4.2.1 Qualitative characteristics for financial reporting

In addition to providing a description of the fundamental qualitative characteristics, the framework describes the factors that enhance these characteristics. These qualitative characteristics and factors are described in this section and, for the purposes of this study, are designated the accounting objectives. The applicable risks, benefits and limitations of each objective when using machine learning technology to produce the financial information are identified in section 4.2.2.

Objective 1: Relevance. Information needs to be relevant to the decisions users are making. Information influences decisions if it can be used to predict future outcomes or to confirm prior evaluations.

Objective 2: Materiality. Information is material if ignoring it or misstating it could affect decisions. Materiality will be determined by the nature or the magnitude of the information and materiality will be unique to every business.

Objective 3: Faithful representation. Information must represent the substance of the matter being presented and not just the form. To do this the information should be complete, neutral and free from error.

Objective 4: Comparability. Accounting information needs to be comparable and enable users to identify the similarities and differences in information. Consistency helps to achieve this goal.

Objective 5: Verifiability. The information needs to be able to be verified in some way, either by direct observation or by being able to recalculate the outputs using the known inputs and methods used.

Objective 6: Timeliness. Information needs to be available to users in time to be able to make the required decisions.

Objective 7: Understandability. Information needs to be presented and classified clearly and concisely.

Objective 8: Cost vs benefit. The framework also takes into account the **cost constraint** of useful financial information, as opposed to the benefits to the user.

In the next section, each of the identified accounting objectives for the applicable risks, benefits and limitations of machine learning are identified.

4.2.2 Machine learning risks and benefits per accounting objective

The risks and limitations that affect the respective objectives are presented in Table 7 and are organised according to the accounting objectives; if a risk pertains to more than one objective, both objectives are noted. The risks the applicable type of consideration has been indicated in the "Type of consideration" column. The identified risks and limitations are addressed using user considerations in chapter 5.

Table 6: Machine learning risks mapped to accounting objectives

Objective number	Risk	Source	Type of consideration
1 (Relevance)	The risk that irrelevant data is included in the dataset. This may be due to outliers, which are values which are far removed from other observations in the data. These increase the risk of misleading representations.	Brownlee (2013:1)	Data
1 & 2 (Relevance & Materiality)	The risk that outliers are relevant to decision-makers and not identified as relevant by the machine learning model. This could impact decision-making.	Brownlee (2013:1)	Model
2 & 3 (Materiality & Faithful Representation)	The limitation relating the fact that machine learning uses probability to identify patterns used to make predictions. This means there is always a margin of error in the predictions made by the machine learning model.	Ayodele (2010a:2); Sainani (2014:841); Vihinen (2012:3)	Model

Objective number	Risk	Source	Type of consideration
	The risk that the predictive accuracy of the machine learning model is not applicable to the task, such as is the case when using the information for compliance tasks.	Gillion 2017:7)	
2 & 3 (Materiality & Faithful Representation)	The risk that the machine learning algorithm is inaccurate owing to insufficient data for training.	Burrell (2016:5)	Training set
3 (Faithful Representation)	The risk that important features are missing from the training data, resulting in not all relevant features being considered when executing solutions. Similarly, missing features in the input data may inhibit model performance.	Amani & Fadlalla (2017:47); Barreno, Nelson, Joseph & Tygar (2010:126)	Feature selection
3 (Faithful Representation)	The risk that there are errors in the data used to train the machine learning algorithm, which may result in incorrect processing and outputs. Similarly, if the data integrity of inputs is not maintained it may have a negative impact on the model.	Appelbaum *et al.* (2017:40); Gillion (2017:9); Sculley, Holt, Golovin, Davydov, Phillips, Ebner, Chaudhary, Young, *et al.* (2015:2500)	Data
3 (Faithful Representation)	The limitation relating to the fact that machine learning algorithms adopt bias to generalise the data, as well as the risk of machine learning algorithms adopting societal bias. This increases the risk of errors or misleading results.	Dietterich & Kong (1995:2); Gillion (2017:7)	Algorithm; Feature selection

Objective number	Risk	Source	Type of consideration
4 (Comparability)	The limitation relating to the fact that machine learning models struggle to transfer solutions from one learning problem to another, thus limiting consistency of information.	The Royal Society (2017:30)	Model
4 (Comparability)	The limitation posed by changes in prediction behaviour owing to changes in feature weights when further features are taught to the machine learning algorithm. The risk that changes in algorithm behaviour cannot be monitored owing to complex or incorrect design.	Sculley *et al.* (2015:2495)	Model
5 & 7 (Verifiability & Understand-ability)	The limitation relating to the fact that users are unable to understand how information is generated by the machine learning technology owing to the complexity of the algorithms, thus making information difficult to verify. This also highlights the limitation of interpretability, as the knowledge that machine learning uses or discovers in order to perform its tasks may not always be available to users.	Ayodele (2010a:2); Sainani (2014:841); The Royal Society (2017:30)	Model
6 (Timeliness)	The risk of increased learning times for machine learning models as the size and complexity of the datasets increase.	Ghanem (2012:161)	Model
7 (Understandability)	The risk that the users do not understand how the machine learning algorithm functions and processes information owing to a lack of technical skills.	Burrell (2016:4)	User
8 (Cost vs benefit)	The risk of costs exceeding the financial benefits to the business, since machine learning requires advanced data	Gillion (2017:9); Sapp (2017:13)	Infrastructure

Objective number	Risk	Source	Type of consideration
	integration tools and infrastructure, which may present significant costs to the business to acquire.		

Source: Own observation

Having identified the relevant machine learning risks and limitations, the benefits of machine learning organised by accounting objective are presented in Table 8. Some benefits address more than one accounting objective, in which case both objectives are noted.

Table 7: Machine learning benefits per accounting objective

Objective number	Benefits	Source
1 & 7 (Relevance & Understand-ability)	Machine learning is able to provide the user with valuable information that they would otherwise not have had access to. The technology is also able to learn which information is relevant to users.	Sapp (2017:12)
2 & 3 (Materiality & Faithful Representation)	Machine learning models can be programmed to include error checks, thus reducing the possibility of omitting information or possible misstatements of information.	Sorio (2013:51)
2 & 3 (Materiality & Faithful Representation)	By automating accounting information processing, the risk of human error will be eliminated; this will increase the accuracy of information and reduce the risk that material information is misstated or incomplete.	Aberdeen Group (2017:2)
4 (Comparability)	Automation models such as machine learning can execute repeated tasks consistently, thus ensuring better comparability of the information produced by the tasks.	Gillion (2017:6); Ventana Research (2016:7)
6 (Timeliness)	Machine learning can process data more efficiently than previous data tools were able to, thus making useful information available faster and improving the business's ability to respond to information.	Aberdeen Group (2017:3); Sapp (2017:12)
8 (Cost vs benefit)	Machine learning will reduce the number of manual tasks required in a process which will save users time. It also increases the efficiency of processing and supports better decision-making. This will leading to cost savings.	Gorbunova & Bochkarev (2011:33)

Source: Own observation

The above findings in relation to the risks, limitations and benefits of machine learning will be further expanded upon for other areas of the machine learning technology life cycle. The risks and limitations identified were mapped to the machine learning life cycle components. These components of the machine learning architecture and the specific risks to consider for each component are described in the next section.

4.3 Technology governance of the machine learning life cycle

When preparing this chapter, principle 12 of King IV (Institute of Directors of Southern Africa (IODSA), 2016:41) was considered in particular. This requires businesses to govern technology in a way that supports the business in achieving its objectives. Accordingly, this principle supports one of the stated objectives of this study, namely, to identify the steps to take when implementing machine learning technology so as to ensure alignment with the goals of the accounting process.

The aim of information technology governance, as described by Alreemy, Chang, Walters and Wills (2016:907), is to ensure compatibility between the goals of the business and a satisfactory level of risk with the use of the emerging technologies. In order to achieve this aim, Alreemy *et al.* (2016:907) highlighted COBIT 5 as a framework which could be used to implement technology governance.

The COBIT 5 framework assists businesses in achieving governance objectives and IT enterprise management. COBIT 5 (ISACA, 2012:19) expands on the traditional software development life cycle stages that need to be managed. The stages that COBIT 5 describes are plan, design, build or acquire and implement, use or operate, evaluate or monitor, and lastly, update or dispose.

The technology life cycle stages together with the software development life cycle stages are illustrated in Table 6. Sapp (2017:17) states that a slightly adapted life cycle is required when developing machine learning to enable more of a focus on model evaluation and tuning, and therefore the traditional software development life cycle has been adapted.

The combination of the two different life cycles was done to ensure that all possible areas of the machine learning technology were considered when identifying the different risks, benefits and limitations. As one of the stated objectives of this study is to explain the role of the user when using the machine learning technology to address identified risks, Table 6 also indicates which tasks of the data science life cycle accounting users need to be involved in, as recommended by Sapp (2017:16).

Table 8: Technology life cycle user involvement

Technology life cycle	Data science life cycle stage	Task	User involved	Related risks, benefits & limitations
Plan	1. Problem understanding	Determine problem objective	×	Section 4.2
		Define success criteria	×	
		Assess constraints	×	Sections 4.4 and 4.5
	2. Data understanding	Assess available data	×	
		Obtain data (access)		Section 4.10
		Explore data	×	Section 4.4
Design	3. Data preparation	Filter data		
		Clean data		
		Training & testing set selection	×	
	4. Modelling	Select algorithm		Sections 4.4; 4.6; 4.7 and 4.9
Build/acquire		Build model		
	5. Evaluation of results	Select/train model		
		Validate/test/tune model		
Use and Evaluate		Explain model	×	Sections 4.7 and 4.8
	6. Deployment	Deploy model		
Monitor		Monitor and maintain	×	Sections 4.4; 4.5; 4.9 and 4.10
Update/dispose		Terminate	×	

Source: Adapted from Sapp (2017:16)

The planning and design phases require the identification of the accounting objectives and technology architecture (Suer, ITIL & Nolan, 2015:1). The risks relating to these objectives are described in section 4.2. The risks relating to the technology architecture are described in section 4.4.

When training the machine learning model during the build/acquire phase, one of the key requirements is the capabilities of the infrastructure (Sapp, 2017:26). Therefore the risks relating to the business infrastructure are described in section 4.5 and the specific risks when acquiring the machine learning technology are described in section 4.6. For the use and evaluate phase, the benefits and limitations of the different machine learning techniques are described in section 4.7 and the user-related risks are described in section 4.8.

Security risks for the monitor and maintain phase are described in section 4.10 and further monitoring risks are described in sections 4.4 and 4.5. Maintenance risks are described in section 4.9 and finally, the update/dispose phase is addressed in section 4.10 as part of security risks.

4.4 Machine learning architecture risks

This section describes the risks pertaining to the identified components of the machine learning architecture that the user must consider. The different machine learning architecture components were described in chapter 3 section 3.4. For ease of reference, Figure 7 in chapter 3 section 3.4 is repeated as Figure 19.

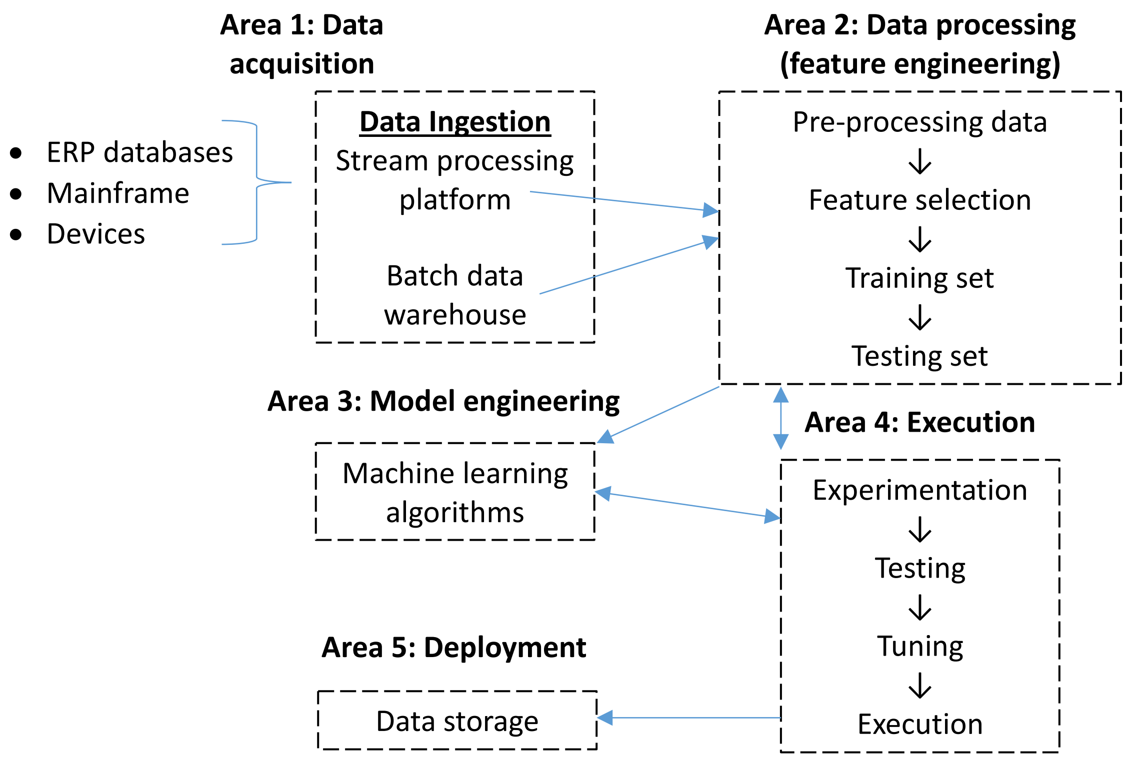

Figure 19: Machine learning architecture (copy of Figure 7) (adapted from Sapp, 2017:20)

The component specific risks for each area, as illustrated above, are provided in Table 9 and are grouped by component.

Table 9: Machine learning architecture risks

Identified components	Component specific risks	Source
Area 1: Data acquisition and Area 2: Pre-processing		
Data ingestion	The risk that the input data source is unstable. Obtaining input from other systems may be convenient but if not monitored there may be unexpected changes in the quality of the input data over time.	Sculley *et al.* (2015:2496)
	The risk that important features are missing from the input data.	Amani & Fadlalla (2017:47)
Area 2: Data processing (feature engineering)		
Pre-processing of data	The risk of errors in the training data used to train the model.	Breck, Polyzotis, Roy, Whang & Zinkevich (2018:1)
	The risk of errors in the input data remaining undetected, thus affecting the quality and the integrity of the data and the outputs.	Appelbaum *et al.*, (2017:40; Gillion (2017:9); Sculley *et al.* (2015:2500)
Feature engineering (feature analysis and selection)	The risk that input features that have little modelling benefit are included in the training data. These features may increase the sensitivity of the technology to changes in the inputs, even though they could be excluded with no disadvantages. This is known as overfitting.	Hawkins (2004:1); Sculley *et al.* (2015:2496)
	The risk that features selected lead to discriminatory predictions or outcomes.	Sapp (2017:13)
	The risk that important features are missing from the training data.	Amani & Fadlalla (2017:47)
Sample selection: training set	The risk that the training set is not large enough. This may result in some necessary features not being represented in the training set or a possible class imbalance where the training set represents a large number of one class and very few of the other class.	Japkowicz & Stephen (2002:435)

Identified components	**Component specific risks**	**Source**
	The risk of the unethical use of data for training or use that infringes on privacy when using confidential data to train machine learning technology.	Gillion (2017:9)
Sample selection: testing set	Risk that real-world data is unavailable for testing purposes.	Ahmed *et al.* (2016:278)
Area 3: Model engineering		
Algorithm and model	The risk that the algorithm has adopted inappropriate bias. This increases the risk of errors. Bias is however necessary in order to be able to generalise beyond the training data.	Dietterich & Kong (1995:11)
	The risk that the algorithm type applied is unable to discover the required pattern. The model is therefore not appropriate for the identified learning problem.	Someren & Urbancic (2006:371)
	The risk that changes in algorithm behaviour cannot be monitored owing to complex or incorrect design.	Sculley *et al.* (2015:2494)
	The risk that the machine learning models are not achieving the desired business objectives, as these change, and therefore that the model is not adaptable to changing business needs.	Amani & Fadlalla (2017:48); Gillion (2017:7); Sapp (201718)
Area 4: Execution		
Experimentation (training)	The risk that the company is unable to process training or experimentation owing to immense computer and storage requirements.	Sapp (2017:6)
Testing	The risk that errors in the algorithm go undetected owing to users not being involved in testing. In the case of accounting tasks, it would be accounting users not being involved in the testing of the accounting machine learning models.	Gillion (2017:10)
	The risk that incorrect assessment measures are used to determine the adequacy of the machine learning model.	Amani & Fadlalla (2017:47)
Tuning	The risk of overfitting the algorithm, where a machine learning model is too closely linked to the actual training data from used to train it.	Witten *et al.* (2016:286).

Identified components	Component specific risks	Source
Execution	The risk that processing power and hardware are not able to meet execution needs.	Gillion (2017:9); Sapp (2017:26)
	The risk of overreliance on models and the tasks performed by machine learning.	Gillion (2017:7)
Area 5: Deployment		
Data storage	The risk that the output may be used by the incorrect systems or unauthorised users, presenting a security risk.	Sculley *et al.* (2015:2495)
	Interoperability of the machine learning technology with the business applications providing the inputs or applications using the outputs.	Daecher & Schmid (2016:42)

Source: Own observation

While considering the risks relating to the machine learning architecture, risks pertaining to the company system that supports this architecture were also identified. These include data availability, interoperability, processing power and hardware requirement risks. Accordingly, the next section, which assesses business infrastructure risks, considers the machine learning business support system .

4.5 Business infrastructure risks when building machine learning models

This section describes the risks relating to the business infrastructure when building and integrating machine learning technology into the accounting processes. Two different ways were identified to integrate the technology into the accounting processes. Firstly, businesses could decide to develop machine learning solutions for specific accounting processes themselves (Gillion, 2017:9). Alternatively, the technology could be integrated into the enterprise accounting software. This is generally done by the software service provider and sometimes without the user realising it (Sapp, 2017:9). The specific risks pertaining to acquiring machine learning technology from a service provider are described in section 4.6.

According to The Royal Society (2017:48), an environment which allows for the effective use of data will be crucial to enabling machine learning, as the technology requires large amounts of data to create the machine learning methods and to train the machine learning systems and such large amounts of data require complex infrastructure. An example of the business infrastructure components required to support machine learning and data availability is given in Figure 20 (Sculley *et al.*, 2015:2497).

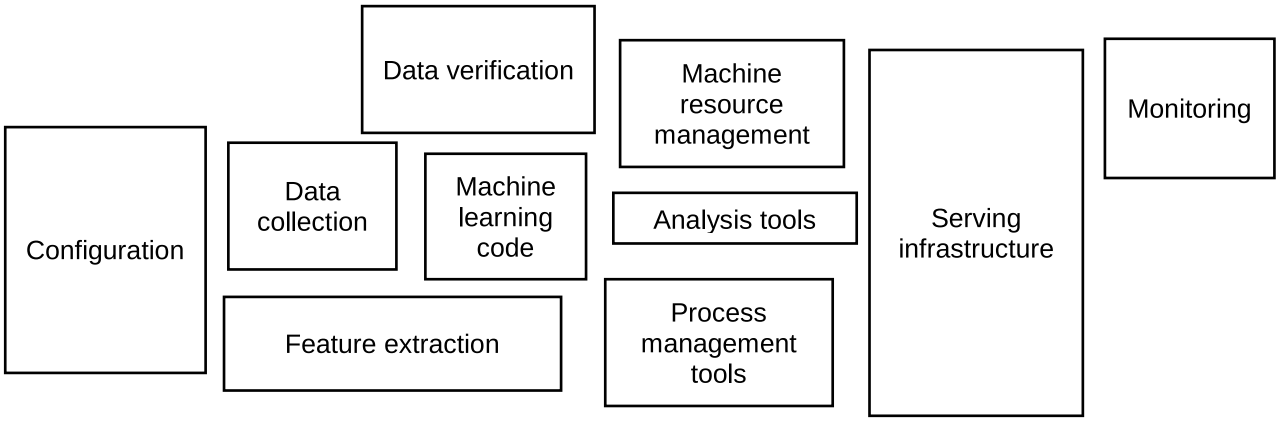

Figure 20: Machine learning support infrastructure (adapted from Sculley *et al.*, 2015:2497)

Figure 20 illustrates some of the components that business infrastructure required to support machine learning should have. The business will also need to consider the risks related to such an infrastructure platform. The risks attached to each element of the infrastructure platform are discussed below.

4.5.1 Configuration of machine learning architecture

Configuration determines how data is selected (data collection), the features used (feature extraction), algorithm settings (machine learning code) and verification methods (data verification). There is a risk that configuration options have been incorrectly selected, leading to errors (Sculley *et al.*, 2015:2499). Specific configuration errors relate to the risks identified for the different machine learning architecture components already described in section 4.4.

4.5.2 Interoperability of analysis tools

The risk that other data analytics platforms and the selected machine learning framework are not interoperable (Sapp, 2017:13).

4.5.3 Serving architecture

Machine learning algorithms require extensive computing power and data management capabilities, therefore the risks that may need to be considered regarding the serving architecture of the business are as follows:

- **Data management capabilities**
  The risk that the architecture is not able to supply the required large and varying amounts of data (Sapp, 2017:12). This in turn presents two more risks, namely, that the amount of data available is insufficient and that the data is not managed adequately to enable it to be used by the machine learning technology.

- **Computing power**
  The risk that the architecture does not have adequate processing power (Sapp, 2017:13).

- **Scalability**
  The risk that the infrastructure is not scalable to accommodate the changing business needs, including the needs of the machine learning model (Corless, De Villiers, Garibaldi & Norton, 2018:6).

- **Flexibility**
  The risk that the infrastructure is not flexible enough to accommodate the changing learning patterns and processing demands of the machine learning model (Sapp, 2017:18).

- **High costs of modernising core systems**
  The risk of disruption and high costs to the business resulting from the need to update or replace core systems to support machine learning technology, as these systems enable the underlying data and processes on which the machine learning technology relies (Buchholz, Jones & Krumkachev, 2016:49).

4.5.4 Monitoring

There is a risk that the infrastructure does not support the monitoring of the machine learning models (Sapp, 2017:18).

4.6 Acquiring machine learning from a service provider

Certain risks and benefits need to be considered when acquiring machine learning enabled software or obtaining machine learning tools or pre-trained models from a service provider.

Such risks may include the following:

- The software purchased may not meet business-specific needs, especially where there are regulatory requirements for accounting (Gillion, 2017:9).
- Software developers may not develop specialist accounting tasks if the market demand for such products is low (Gillion, 2017:9).
- Off-the-shelf models may still require significant computing power, depending on the model (Sapp, 2017:34).
- Professionals who purchase machine learning algorithms may not have the basic understanding of machine learning technology to use the technology effectively to achieve the business objectives (Sapp, 2017:13).
- Outsourced models may be maliciously trained, for example the backdoored neural network described by Gu, Dolan-Gavitt and Garg (2017:1), which performs badly on specific attacker-chosen inputs, thus posing a security risk.

Benefits may include:

- The skills barrier to incorporating machine learning into a business presented by purchased software is significantly lower (Sapp, 2017:13).
- Fewer integration challenges are experienced compared to the in-house development of machine learning technology (Sapp, 2017:35).
- Implementation costs are lower, especially when considering infrastructure risks as described.

4.7 The benefits and limitations of various machine learning techniques

Specific benefits and limitations were identified for each of the machine learning techniques described in chapter 3. The findings are presented in Table 10.

Table 10: Benefits and limitations of respective machine learning techniques

Machine learning technique	Benefits	Limitations	Research
Decision trees (section 3.5.1.1)	• Easy to interpret	• Bias towards a certain category • Overfitting	Marsland (2009:133); Thomassey & Fiordaliso (2006:410); Samoil

Machine learning technique	Benefits	Limitations	Research
		• Too simple a model for complex data	(2015:35, 42); Dataiku (2017:7)
C4.5 decision trees (section 3.5.1.2)	• Can process data with missing features • Can process continuous feature values by using thresholds to create value intervals	• Pruning may lead to loss of accuracy	Dataiku (2017:7) Samoil (2015:17)
Random forests (section 3.5.1.3)	• Greater performance than individual decision trees • Quick to train	• Difficult to interpret • Overfitting	Dataiku (2017:7)
Naïve Bayes (section 3.5.1.6)	• Performs well in multi class prediction • Where variables are independent it performs better than other classification models and less training data is required • The algorithm is scalable and can adapt additional features	• The algorithm assumes that features are independent, while this is unlikely • If a feature has a category which was not observed in training dataset, then a zero probability will be assigned to that category, thus resulting in the algorithm not being able to make a prediction, known as zero frequency.	Samoil (2015:10) Witten *et al.* (2016:99); Samoil (2015:16); Larsson & Segerås (2016:34)
Bayes belief network (section 3.5.1.7)	• Can handle missing data • Provides knowledge about causal relationships between variables • Provides a method for avoiding data overfitting	• Complex calculations are required to train the network, which are expensive and take time • Highly dependent on the prior knowledge used to assume prior probabilities	Heckerman (2008:33); Niedermayer (2008:128)

Machine learning technique	Benefits	Limitations	Research
k-nearest neighbour (kNN) (section 3.5.1.8)	• Handles database noise or outliers better than a nearest neighbour algorithm by considering more than one item to determine a class	• Poor interpretability • Complex which makes it slow	Kotsiantis (2007:263); Witten *et al.* (2016:87, 136, 141)
Conditional random fields (section 3.5.2.1)	• Can make predictions despite haphazard complex features in the input sequence • These discriminative models are more suited than other techniques where there are overlapping features	• Trade-off between large feature datasets, which are more accurate but require more memory to store and increased risk of overfitting	Witten *et al.* (2016); Sutton & Mccallum (2007:282); Sutton & Mccallum (2007:293)
Support vector machines (section 3.5.3.1)	• Reduced risk of overfitting • Data training is simple • In the case of large datasets, it is able to simplify the problem	• Poor interpretability	Kotsiantis (2007:263); Karamizadeh, Abdullah, Halimi, Shayan & Rajabi(2014:65); Witten *et al.* (2016:257)
Artificial neural networks (section 3.5.3.2)	• Very effective at handling complex tasks	• Poor interpretability • The process is extremely computing intensive and requires modern, powerful computers	Kotsiantis, (2007:263); Dataiku (2017:7); SMACC (2017:9)
Convolutional neural networks (section 3.5.3.3)	• Ideal for image recognition	• Takes a long time to train • Poor interpretability	Dataiku (2017:7)

Machine learning technique	Benefits	Limitations	Research
Association rules (section 3.6.1.1)	• Association rules can be assessed by the coverage and accuracy of the rule; assisting decision-making on which rules to use	• The number of rules discovered may be excessive • The algorithm takes a long time to produce outputs	Witten *et al.* (2016:79); Kaur (2014:2322)
Self-organising maps (section 3.6.2.1)	• Easy to understand	• The neural networks process is extremely computing intensive and requires modern, powerful computers • Need a value for every dimension of the map so it is difficult to obtain adequate data	Ayodele (2010b:46); SMACC (2017:9)
K-means clustering (section 3.6.2.2)	• Simple and effective method of clustering.	• Difficult to determine the correct number of clusters, therefore this algorithm will need to be run a number of times to get the correct number, taking lots of time	Ayodele (2010b:30); Witten *et al.* (2016:144)
Semi-supervised clustering (section 3.7.1)	• More accurate than unsupervised clustering	• The rate at which these algorithms function is highly dependent on the size of the available labelled data	Zheng *et al.* (2017)

Source: Own observation

The benefits and limitations presented in Table 10 were used to identify how these machine learning techniques address the accounting objectives identified in section 4.2.1. Table 11 presents the constraints and benefits of the different machine learning techniques and maps these to the different accounting objectives, based on the definition of these as provided in section 4.2.1. The table also indicates in which accounting tasks the different machine learning techniques can use, as identified in table in chapter 3 section 3.8.

Table 11: Benefits and limitations of machine learning techniques mapped to objectives

Machine learning techniques	Benefits (section 4.7)	Accounting objective (section 4.2)	Limitations (section 4.7)	Accounting objective (section 4.2)	Tasks in the accounting process (Chapter 3 section 3.8)
Transfer learning decision forests and random forests	• Greater performance than individual decision trees	Materiality & faithful representation	• Poor interpretability	Verifiability & Understandability	• Adaptability of OCR • Account allocation
	• Fast training	Timeliness	• Overfitting	Relevance	
Support vector machine	• Reduced risk of overfitting	Relevance	• Poor interpretability.	Verifiability & Understandability	• Image classification • Validation of document information • Forecasting performance • Removing of duplicate entries and linking documents • Matching records or record-linkage
	• Data training is simple	Understandability			
	• Simplifies the problem	Understandability			
Convolutional neural networks	• Ideal for image recognition	NA	• Poor interpretability	Verifiability & Understandability	• Irregular document layout classification using NLP
k-Nearest neighbour	• Adaptable to outliers	Faithful representation	• Poor interpretability	Verifiability & Understandability	• Image classification

Machine learning techniques	Benefits (section 4.7)	Accounting objective (section 4.2)	Limitations (section 4.7)	Accounting objective (section 4.2)	Tasks in the accounting process (Chapter 3 section 3.8)
			• Complex which makes it slow	Timeliness	• Text classification
Semi-supervised clustering	• More accurate than unsupervised clustering	Materiality & Faithful representation	• Training rate may be slow	Timeliness	• Text classification
Naïve Bayes	• Multi-class prediction	Comparability	• Requires independent variables	Faithful representation	• Validation of document information • Removing of duplicate entries and linking documents • Matching records or record-linkage • Account allocation • Error detection in financial data and fraud detection
	• Excellent classifier for independent variables	Faithful representation	• Training set sensitive	Materiality & faithful representation	
	• The algorithm is scalable and adaptable	Comparability			
Artificial neural network	• Handles complex tasks effectively	Relevance	• Poor interpretability	Verifiability & Understandability	• Removing of duplicate entries and linking documents • Matching records or record-linkage • Forecasting performance
			• Computing intensive	Cost saving	

Machine learning techniques	Benefits (section 4.7)	Accounting objective (section 4.2)	Limitations (section 4.7)	Accounting objective (section 4.2)	Tasks in the accounting process (Chapter 3 section 3.8)
Bayesian Belief network	• Can handle missing data	Relevance & Materiality	• Computing intensive.	Cost saving	• Error detection in financial data and fraud detection • Forecasting performance
	• Provides knowledge about causal relationships between variables	Relevance & Understandability			
	• Provides a method for avoiding data overfitting	Relevance	• Dependent on the prior knowledge	Materiality & Faithful representation	
Association rules	• Association rules can be assessed by the coverage and accuracy of the rule, assisting decision-making on which rules to use	Verifiability & Understandability	• Excessive output	Relevance	• Error detection in financial data and fraud detection
			• Requires lots of time	Timeliness	
K-means clustering	• Simple and effective method of clustering	Understandability	• Requires lots of time	Timeliness	• Account allocation • Error detection in financial data and fraud detection

Machine learning techniques	Benefits (section 4.7)	Accounting objective (section 4.2)	Limitations (section 4.7)	Accounting objective (section 4.2)	Tasks in the accounting process (Chapter 3 section 3.8)
Self-organising maps	• Easy to understand	Understandability	• Computing intensive	Cost saving	• Error detection in financial data and fraud detection
			• Requires adequate data	Materiality & Faithful representation	
Conditional random fields	• Manages haphazardly complex features	Relevance	• Trade-off between accuracy which requires memory and overfitting	Relevance & Faithful representation	• Report descriptions
	• Suited to overlapping features	Faithful representation			
C4.5 decision trees	• Can process data with missing features	Relevance & Materiality	• Pruning may lead to loss of accuracy	Faithful representation	• Forecasting performance
	• Can process continuous values	Relevance			

Source: Own observation

The table enables the user to consider the accounting objectives and assess the constraints of the technology. This forms part of the accounting user's role when implementing and using machine learning technology (Sapp, 2017:16).

4.8 User-related risks

This section describes the risks related to the users of machine learning technology. Certain user risks have already been identified when examining other stages of the machine learning life cycle. In addition to these risks and benefits, further areas will be addressed in this section.

The user-related risks that have already been identified are the following:

- The risks relating to users not having the necessary technical skills and interpretability of machine learning algorithms as linked to accounting objective 7 in section 4.2.2.
- The risk that errors in the algorithm go undetected owing to accounting users not being involved in testing (Gillion, 2017:7), as identified in section 4.4.
- The risk of societal and unethical or discriminative biases in the machine learning model owing to bias in the training data (Gillion, 2017:7), as identified in section 4.4.

Apart from these risks, there may be further risks related to the users of machine learning technology in the accounting process. These include the following:

- The risk of overreliance on machine learning models because the limitations of the machine learning models are not understood (Gillion, 2017:7; Vihinen, 2012:3).
- The risk that accounting users are unable to adapt to the new ways of thinking required for machine learning, leading to value loss for the business as users are unable to utilise the machine learning capabilities (Gillion, 2017:10).

4.9 Maintenance risks

- The risk that the desired functioning of the machine learning model is not maintained owing to changes in the model caused by input data (Sculley *et al.*, 2015:2500). This may be the result of a lack of monitoring of the changes in the machine learning model.

- As described in section 4.4, there is a risk that the machine learning models are not updated regularly to remain relevant to the changing business environment (Amani & Fadlalla, 2017:48).

- The risk of significant maintenance costs owing to the complexities that machine learning models present. Furthermore, maintenance is not only has to be performed on the machine learning code but also on the entire machine learning system, as data influences and changes the machine learning model (Sculley *et al.*, 2015:2494).

4.10 Security risks

Barreno *et al.* (2010:126) suggest that from a machine learning security perspective, the risk is that an attacker may attempt to use the adaptive aspect of a machine learning model to cause problems. This would generally be achieved by targeting the data used by the machine learning model. The following security risks were identified:

- **Erroneous data not detected leading to errors**
  If malicious false negative input data is processed by the machine learning model it could lead to the production of erroneous information, thus affecting the integrity of the information available to the accounting user. For example, the data features could be set up in such a way that the algorithm is unable to classify the data as erroneous and thus processes it as correct.

- **System compromise**
  Risk of unauthorised access entering the system via viruses from input false negative data.

- **Corruption of model during training**
  Risk of malicious data disrupting the machine learning process training with false positive data, causing the machine learning model to operate in a manner that differs from the objectives set by the accounting user; for example, causing the algorithm to classify correct data as incorrect or irrelevant.

- **Disruption of service attack**
  Risk of malicious data disrupting the operation of the machine learning model, in most cases with false positive data, also known as a denial of service attack. This may take the form of the machine learning model receiving an overwhelming amount of false data to the point that the algorithm is unable to process all the false inputs, resulting in downtime of the system.

- **Eavesdropping on training data**
  Risk of an attacker eavesdropping on all network traffic while the learner gathers training data, thus being able to determine which data the business has available for training purposes.

Apart from the security risks posed by malicious attackers, further security risks may include:

- **Data protection rights**
  Risks surrounding privacy of data used in training machine learning algorithms, where non-authorised users have access to sensitive information (Gillion, 2017:9), as well as risks related to infringing data protection requirements (European Union, 2018:3).

- **Data ownership risks**
  Risks pertaining to access to sensitive data or data not owned by the company that is legally protected, as a variety of data is necessary for training machine learning models. This may include data that the user or service provider may not necessarily own (The Royal Society, 2017:49).

- **Acquiring machine learning software**
  Security risks identified in the machine learning model supply chain are described as part of the risks that are present when acquiring machine learning technology from a service provider in section 4.6.

- **Application and software risks**
  Security risks relating to the application or system in which the machine learning model operates. These include policy enforcement risks, confidentiality risks, access control risks and data transmission risks. These risks are however not unique in a machine learning environment (Demchenko, De Laat & Membrey, 2014:110).

- **Secure disposal of assets**

  The risks that machine learning technology is not securely disposed of or terminated at the end of its useful life, which may lead to unauthorised access to data (ISACA, 2012:165) or unintended code behaviour, where a machine learning code is not properly removed from the system (Sculley *et al.*, 2015:2498).

Apart from the possible security risks, machine learning can also hold benefits for security. The ability to adapt to changing and complex situations has meant that machine learning has also become a fundamental tool for computer security (Barreno *et al.*, 2010:121).

4.11 Conclusion

This chapter identified the risks, benefits and limitations associated with machine learning and specific machine learning techniques as identified in chapter 3. The risks were mapped to each stage of the technology life cycle, thus indicating where user involvement was required. These risks include risks relating to achievement of the accounting objectives, machine learning architecture risks, business infrastructure risks, user-related risks and security risks.

On the other hand, the benefits identified included ways in which machine learning could assist the user in achieving the accounting objectives. The limitations of machine learning in regard to achieving the accounting objectives were also described. In addition, the benefits and the limitations of the different machine learning techniques identified in chapter 3 were discussed.

The findings of this chapter, namely, the risks, benefits and limitations of machine learning technology, especially in an accounting context, are used in chapter 5 to develop guidelines for implementing machine learning techniques in an accounting context. A summary of the identified risks per category and the relevant guidelines presented in chapter 5 are listed in Table 12.

Table 12: Identified risks and relevant user considerations

Risk category	Accounting objective	Risk	Consideration	Guideline section
Data ingestion		Unstable input data source	Monitor stability of input data	5.4.2
Data ingestion	Faithful representation	Missing features in input data	Monitor quality of input data and user involvement in preparation of data	5.4.2
Pre-processing data	Faithful representation	Errors in input data	Monitor quality of input data and user involvement in preparation of data	5.4.2
Pre-processing data	Faithful representation	Errors in the training data	Data validation techniques on training data	5.4.2
Pre-processing data/security	Faithful representation	Data integrity, privacy and data protection	Access controls	5.4.2
Model	Relevance & Materiality	Impact of outliers	Users to specify when outliers are relevant	5.4.1
Feature selection	Faithful representation	Redundant features in the data.	Enquire which features have been included and perform leave-one-feature-out training	5.4.3
Feature selection	Faithful representation	Discriminatory features or unethical data in training set	Enquire which features have been included and remove or filter discriminatory features	5.4.3
Feature selection	Faithful representation	Missing features in training data	Enquire about feature selection techniques and analysis tools	5.4.3
Training and testing set	Materiality & Faithful representation	Insufficient real-world data for training and testing	Synthetic data and data preparation tools	5.4.3
Algorithm	Faithful representation	Incorrect bias	Bias will need to be assessed by technical professionals	5.4.4
Algorithm/ maintenance	Comparability	Undetected changes in behaviour	Enquire about monitoring mechanisms	5.4.4
Model		Does not address learning problem	There needs to be a degree of repeatability or structural pattern in the learning problem	5.4.1

Risk category	**Accounting objective**	**Risk**	**Consideration**	**Guideline section**
Model	Verifiability & Understandability	Complex algorithm affecting interpretability	Model certification and input to output mappings	5.4.1
Model/ maintenance		Not adaptable to changing business needs	Regular recalibration of model	5.4.1
Model	Comparability	Solutions not transferable to new problems	Improve model interpretability	5.4.1
Model	Materiality & Faithful representation	Predictive accuracy of the model	Confidence levels or probability measurements	5.4.1
Model		Scalability of the model	Consider incorporating cloud technologies	5.4.1
Model/ maintenance	Cost vs benefit	Monitoring and maintenance costs of the model	Real-time monitoring and automated responses and obtaining independent assurance as well as appropriate system design	5.4.1
Model and experimentation	Timeliness	Long training and operating times	Enquire about operating times and volume of transactions that model can process	5.4.1
Tuning/testing		Overfitting	Separate dataset for testing and evaluation of the model as well as cross-validation techniques	5.4.5
Training		Undetected accounting errors in model	Accounting users should be involved in testing	5.4.5
Testing/ user risk		User not involved in testing	Accounting users should be involved in testing	5.4.5
Testing		Incorrect assessment measures	Testing environment should mimic real environment and guaranteed performance levels	5.4.5
Execution		Overreliance on models	Be aware of limits of the model.	5.4.1
Experimentation/ infrastructure		Insufficient storage	Consider cloud technology	5.5

Risk category	Accounting objective	Risk	Consideration	Guideline section
Experimentation/ infrastructure/ service provider		Insufficient processing power	Ensure adequate on-site power or use of cloud technology	5.5
Infrastructure	Cost vs benefit	High costs of adoption	Cost-benefit analysis	5.5
Data storage/ security		Unauthorised access to output	Settings, encryption, read only rights	5.7
Data storage/ infrastructure		Interoperability	Update data analytic architectures and define data and service standards	5.5
Infrastructure		Incorrect configuration options selected	Adopting good configuration principles	5.5
Infrastructure		Insufficient supply of data	Agreements with external parties, assessing external data quality and data analytics architectures and investment in enabling technologies	5.5
Infrastructure		Not scalable or flexible	Modernise infrastructure and consider cloud technologies	5.5
Infrastructure	Cost vs benefit	Outdated core systems leading to costs of updating	Cost-benefit analysis	5.5
Infrastructure		Lack of monitoring of machine learning model	Ask technical professionals about monitoring capabilities of the model and obtaining independent assurance	5.5
Using service provider		Does not meet accounting needs	Use reputable suppliers and select providers that interoperate with multiple frameworks	5.5.1
Using service provider		Unavailable in the market	Users perform cost-benefit analysis before in-house development	5.5.1
Using service provider/ security	Faithful representation	Security risks such as backdoors	Use reputable suppliers, integrity in transit guarantees and machine learning models come with digital signatures, independent assurance and the use of a service level agreement	5.5.1

Risk category	Accounting objective	Risk	Consideration	Guideline section
Using service provider/ user risks	Understandability	Insufficient technical skills to understand model	Train users in machine learning basics	5.6
User risks		Unethical or discriminative models	Enquire which features have been included and remove or filter discriminatory and unethical features	5.4.3
User risks		Overreliance on models	Be aware of the limits of the model	5.4.1
User risks		Users not adaptable leading to value loss	Accountants need to adapt their thinking and improve their communication skills	5.6
Security	Faithful representation	Unauthorised access-viruses	Virus detection system	5.7
Security	Faithful representation	Malicious training data	Incorporate identification capabilities into algorithm	5.7
Security		Denial of service attack	Machine learning adaptive security	5.7
Security	Faithful representation	Access to training data	Access controls	5.7
Security		Data privacy and protection	Access controls	5.4.2
Security		Data protection risks where not owned by company	Access controls, encryption, scanning of data for threats and read-only rights	5.4.2 & 5.7
Security/ terminate/ maintenance		Inadequate disposal of data or assets	Secure disposal of data and review codes to identify unnecessary codes	5.7

Source: Own observation

Chapter 5: Guidelines for implementing machine learning in an accounting context

5.1 Introduction

In chapter 4 the risks, benefits and limitations when implementing machine learning were identified. These findings are used to achieve a further objective of this study, namely, to identify the steps to take when implementing machine learning technology to ensure alignment with the goals of the accounting process.

As determined in chapter 4, principle 12 of King IV (Institute of Directors of Southern Africa (IODSA), 2016:41) requires businesses to govern technology in a way that supports the business in achieving its objectives. Therefore, the findings of chapter 4 together with the principles of King IV will be used to develop guidelines for implementing machine learning technology in the accounting processes.

To achieve the objective, this chapter is structured in the form of steps that users could take when implementing machine learning technology. These steps are aligned to the stages of the data science life cycle as described in section 4.3 of chapter 4. The first step involves assigning the responsibility for implementing machine learning technology and for the governance of the technology. The second step requires users to consider the impact of the machine learning technology on the accounting objectives. This is important as the accounting objectives are in fact the goals of the accounting process.

The third step involves users considering different aspects of the machine learning architecture, while the fourth step requires users to consider the various requirements of the business infrastructure. The fifth step involves determining user requirements and the final step involves the consideration of security requirements.

5.2 Step 1: Assigning responsibility for implementing machine learning technology

According to King IV (Institute of Directors of Southern Africa (IODSA), 2016:62), those responsible for the governance of the business should set the direction for technology governance. This may include drawing up a policy that describes the direction of the businesses approach to the technology and incorporates plans for managing the risks surrounding machine learning technology as part of the business's risk management procedures.

The responsibility for implementing the technology can then be delegated to management. When users develop their own machine learning models or acquire software packages that incorporate machine learning technology, the associated risks will need to be addressed.

When considering the role of the user in implementing machine learning technology, there are certain tasks in the data science life cycle that Sapp (2017:16) recommends accounting users be involved in. These tasks are listed in Table 6 in section 4.3. An extract of this table is provided as Table 13, where the tasks that users need to be involved in are mapped to the user considerations section.

Table 13: Sections applicable to accounting user tasks in data science life cycle

Task	Accounting user involvement	Related risks, benefits & limitations	Consideration type	Relevant section
Determine problem objective	✕	Section 4.3	Accounting objectives	5.3
Define success criteria	✕	Section 4.3	Testing	5.4.5
Assess constraints	✕	Sections 4.4 and 4.5	Machine learning technique limitations	5.3
Assess available data	✕	Sections 4.4 and 4.5	Data	5.4.2
Explore data	✕	Section 4.4	Data	5.4.2
Training & testing set selection	✕	Section 4.4	Feature selection	5.4.3
Explain model	✕	Sections 4.7 and 4.8	Model	5.4.1
Monitor and maintain	✕	Sections 4.4; 4.5; 4.9 and 4.10	Infrastructure	5.5
Terminate	✕		Security	5.7

Source: Own observation

Table 13 will assist users to understand their role in implementing machine learning technology and what they should consider when addressing the identified risks. The sections that follow present these considerations and the guidelines that address the identified risks.

Table 12 in chapter 4 section 4.11 provides a summary of the identified risks per category and the relevant guideline section.

5.3 Step 2: Consider the accounting objectives

Risks are assessed based on their significance to the accounting users in terms of the specific processes the machine learning technology will be used for. The most significant risks for the purposes of user considerations are those that affect the achievement of the accounting objectives. The risks identified in chapter 4 section 4.2.2 pertaining to the accounting objectives may have their origin in the machine learning architecture or the machine learning model or they may be user related.

Table 7 in chapter 4 section 4.2.2 indicates the type of consideration pertaining to each risk. When considering the accounting objectives, users need to take into account the specific risks and benefits described in chapter 4 section 4.2.2, as well as the benefits and limitations of the machine learning techniques, as described in chapter 4 section 4.7. Table 11 in chapter 4 section 4.7 maps the respective benefits and limitations of the different machine learning techniques to the accounting objectives for each accounting task.

In general, accounting tasks are suited to machine learning owing to the following accounting attributes, as described by SMACC (2017:6):

- Financial information has an organised data structure.
- Data inputs such as invoices and bank accounts are readily available.
- Data inputs are easy to transform into a digital form.
- Accounting has rules that must be followed for the verification of data.
- Accounting lends itself to the processing time that is sometimes required for processing that uses machine learning. Outputs do not have to be instantly available, rather they should be available in a timely manner.

When a user has established that the learning problem is suited to machine learning and the objectives a set by the accounting users, there will be additional steps in the form of considerations and controls that users will have to put in place in order to implement machine learning in the accounting process. These steps are expanded on in the following sections.

5.4 Step 3: Consider the machine learning model and architectural components

This section sets out the user considerations pertaining to the various components of the machine learning architecture. Most of the risks that affect the achievement of the accounting objectives are risks pertaining to components of the machine learning architecture. These risks include the machine learning model, data considerations, feature selection, the training set and algorithm considerations and testing, as set out in Table 7 of chapter 4 section 4.2.2.

This section is organised according to the different machine learning architecture components and their respective considerations.

5.4.1 Machine learning model considerations

The considerations listed below will need to be taken into account by accounting users when asking technical professionals about the machine learning model.

- **The limits of the model**

  Gillion (2017:7) recommends that users be aware of the limits of the model to ensure that these models are not overly relied on and that human involvement is retained in the accounting decision processes.

- **The predictive accuracy of the model**

  One control to assist users in understanding the predictive accuracy of a machine learning model and the implications for decision-making would be to provide explicit confidence levels or a measure of the probability of the model outputs (Gillion, 2017:7; Vihinen, 2012:3).

- **The impact of outliers on the model**

  Users should specify whether outliers are relevant in the task being performed, such as in error detection in task 7 of the preparation of management accounts as described in chapter 2 section 2.4.3.

- **Adaptability of the model to new problems**

  The Royal Society (2017:30) suggests that improving the interpretability of the model will increase the possibility of being able to transfer the model between learning problems.

- **The interpretability of the model**

  The Royal Society (2017:94) makes the following recommendations to address the problem of interpretability:

  - Model certification – this would indicate the competence of the machine learning model.
  - Input to output data mappings – this would indicate the influence of the different inputs on the outputs.

- **The adaptability of the model to changing business needs**

  Amani and Fadlalla (2017:48) recommend that models be recalibrated regularly to ensure that they remain valid and are able to perform the required tasks over time.

- **Model training and operating times**

  To address risks related to timeliness, users will need to ask technical professionals about the speed at which the model operates, the volume of transactions the model can process and the training time required to train the model.

- **The scalability of the model**

  Sapp (2017:2) recommends that to assist model scalability, cloud-based capabilities be incorporated when designing machine learning models. This is because the cloud platform has elastic characteristics that assist in scaling algorithms.

- **Appropriateness of the machine learning approach**

  The learning problem will need to have a certain amount of repeatability or a structured pattern in order to be suited to machine learning (Gillion, 2017:7; Someren & Urbancic, 2006:371).

- **Monitoring requirements**

  Users will want to ask technical professionals about capabilities for monitoring the machine learning models (Sapp, 2017:18). Sculley *et al.* (2015:2500) recommend real-time monitoring of the entire machine learning system and state that an automated response is important to sustain the reliability of the system.

  The following specific items will need to be monitored by users, based on the risks identified in chapter 4:
  - Unauthorised access to the machine learning model
  - Unauthorised or malicious data being processed by the machine learning technology
  - Quality of input data
  - Errors and exceptions
  - Outliers
  - Algorithm changes.

- **Assurance requirements**

  The business may consider periodic independent assurance on the effectiveness of the machine learning model, including where the technology is provided by a service provider (Institute of Directors of Southern Africa (IODSA), 2016:63).

- **Maintenance costs**

  Maintenance costs may be high as a result of a number of factors, including but not limited to configuration problems, changes in features or data changes affecting algorithm performance and data dependencies. Ensuring that a system is adequately designed and monitored can reduce unnecessary maintenance costs caused by these risk factors (Sculley *et al.,* 2015:2494).

5.4.2 Data considerations

Data considerations pertain to risks such as errors in training or input data, missing features in the input data, the impact of outliers and risks relating to unstable input data sources. Specific considerations and possible controls users can employ to address the identified data risks are described below.

To address the risk of errors in the training data, Breck *et al.* (2018:2) recommend using data validation techniques to monitor the data quality. Similarly, Sculley *et al.* (2015:2500) recommend that in order to maintain a well-functioning system, some input data should be tested to address the risk of errors, changes in the data or incomplete data.

In addition, it may be necessary to monitor the stability of the input data, especially in situations where data is produced by other machine learning algorithms. This may happen, for example, when integrating machine learning into the various accounting processes (Sculley *et al.*, 2015:2496).

Accounting users may need to be directly involved in managing the inputs or outputs of machine learning models, such as exception-handling or preparation of the data inputs (Gillion, 2017:10). Users may also need to enquire about the input data sources required to use the model. This data could come from a number of sources and may be structured or unstructured (Sapp, 2017:6).

Lastly, adequate controls will need to be put in place for the governance of data, including access controls to maintain the integrity of input data and to protect privacy rights (Sapp, 2017:23).

5.4.3 Feature selection and training and testing set considerations

Feature selection and training set risks include risks of redundant, discriminatory and unethical or missing features. In addition, there is the risk of insufficient data for training leading to an imbalance of features.

To identify redundant or discriminatory features, users may want to ask technical professionals which features have been selected to train algorithms to perform the accounting tasks and use their accounting knowledge to determine whether those features are relevant. To be able to perform this function users may require an understanding of machine learning techniques (Gillion, 2017:10).

Sapp (2017:23) also recommends that users should try and remove features affecting privacy or ethical rights from the dataset used to train machine learning models. This can be done by filtering data that may infringe on privacy rights or support unethical predictions.

Furthermore, there may be undetected unnecessary features in the data set which can be addressed by evaluations where leave-one-feature-out trainings are done (Sculley *et al.*, 2015:2496). Users may also need to ask technical professionals about the feature selection techniques, variable clustering, and analysis tools used to ensure that no important features are missing from the training data set (Amani & Fadlalla, 2017:47).

Data for training may also be insufficient, in which case Ahmed *et al.* (2016:285) recommend the use of synthetic data to train models. Furthermore, to ensure that data is available for machine learning models, users may want to consider obtaining self-service data preparation tools to support technical professionals in preparing and manipulating data (Sapp, 2017:23).

5.4.4 Algorithm considerations

Algorithm considerations will be based on the identified risks, which included incorrect algorithm bias and undetected changes in algorithm behaviour. Accounting users may not be directly involved in addressing these risks but they will still need to consider the risks and ensure that they are addressed by technical professionals.

Accounting users may want to ask about the way changes in algorithm behaviour are monitored, including the available detection mechanisms and whether these have been taken into consideration in the model design (Sculley *et al*., 2015:2494).

Lastly, in order to avoid errors, appropriate data tools may need to be employed to detect and diagnose overly strong, weak or inappropriate biases in the machine learning algorithm (Dietterich & Kong, 1995:11).

5.4.5 Testing considerations

The main risks regarding testing considerations are that the wrong assessment measures are used to test the machine learning model and that errors go undetected because accounting users are not involved in testing the model. There is also the risk of overfitting, where the machine learning model predictions are too closely linked to the specific training data used to train it.

With regard to testing the machine learning models, Gillion (2017:10) recommends that accountant users should be involved in training and testing the models and possibly in auditing the machine learning algorithms. In addition, Sapp (2017:27) recommends that the testing environment should be as close to the real environment as possible.

Users will need to ask technical professionals about how and whether the credibility of the model has been assessed to address overfitting. If overfitting is addressed during the development of the machine learning model, it will require the tuning of the model and cross validation methods, as well as the use of a large dataset separate from the training or testing data (Witten *et al.*, 2016:286).

The Royal Society (2017:112) recommends asking technical professionals about the guaranteed minimum level of performance of the model, which could be achieved by including the theoretically worst possible observable data during the training phase.

5.5 Step 4: Consider infrastructure needs

In order to support the machine learning architecture, the machine learning infrastructure which supports the architecture will need to be considered. The risks identified in this regard are insufficient storage, insufficient processing power, high adoption costs, lack of interoperability, and insufficient supply of data, scalability and flexibility and lack of monitoring capabilities. Based on the identified business infrastructure risks, users should consider the following infrastructure requirements:

- **Enabling technologies**

  In chapter 2, various technologies were identified that would enable the use of, or could be combined with, machine learning technologies in the accounting processes. Users should consider whether investing in these technologies is necessary to be able to use machine learning technology in the desired processes.

- **The storage requirements of the model**

  One option that enables that economical storage of data is cloud computing; however, it does expose the business to additional risks that are outside the scope of this study (Richins, Stapleton, Stratopoulos & Wong, 2017:74).

- **Processing power and hardware requirements**

  Businesses will need the correct infrastructure in terms of hardware and processing power to ensure that machine learning processing can be properly executed. This will infrastructure will depend on how advanced the machine learning technology is and may mean ensuring that the business has enough power on the premises or obtain the service from a cloud service provider (Gillion, 2017:9; Sapp, 2017:26).

- **The financial effects of adopting a machine learning model**
  The business will need to perform a cost-benefit analysis to assess the economic benefits of replacing existing processes with machine learning technology, and thereby assess the business case for adopting machine learning. This will depend on whether users develop the technology in house or purchase the technology as part of accounting software (Gillion, 2017:9).

- **Interoperability of data analytics architecture**
  Sapp (2017:1) recommends updating data analytics architectures to support data preparation for machine learning algorithms and to ensure adequate data supply and interoperability.

- **Standards that assist interoperability**
  Daecher and Schmid (2016:42) recommend defining data and service standards to help ensure interoperability when implementing new technology in a business.

- **Sufficient data supply**
  Accounting users will need to consider new ways of accessing data. The Royal Society (2017:49) highlights the need for users to enter into agreements with external parties to access the data required for machine learning models. Furthermore, users may need to monitor the quality of data obtained from external sources.

- **Configuration requirements**
  For configuration, Sculley *et al.* (2015:2499) recommend that there are certain principles of good configuration in machine learning systems that should be adhered to, including enabling transparency with regard to number of features used, data dependencies, detection of unused settings, and controls to ensure that omissions or errors are detected. Finally, a full code review of the system configuration should be performed.

- **Scalable and flexible infrastructure**
  When a business is considering the adoption of machine learning technology and in ensuring its infrastructure is scalable and flexible, Corless *et al.* (2018:8) maintain it should continually modernise its infrastructure to incorporate the use of cloud technologies (Sapp, 2017:19).

- **Integration requirements.**
  To integrate the technology into the existing system, Sapp (2017:13) highlights the need for businesses to have data integration tools and to ensure they have a thorough data integration strategy.

5.5.1 Service provider and purchased machine learning considerations

The risks of purchasing machine learning technology from a service provider include the risk that the technology required to meet business needs is not available, as well as the inherent security risks when using such technology.

In situations where certain specialised accounting products are not supplied by service providers in the market, it may be that the cost of producing the product exceeds the benefit of solving the business problem. In such cases, users will need to perform a cost-benefit analysis before developing in-house solutions (Gillion, 2017:9).

Users should also ensure that reputable software suppliers are chosen, especially when regulatory or legal requirements, as required with accounting information, are at issue (Gillion, 2017:9).

When selecting a machine learning platform provider, users need to ensure that they select one than interoperates with multiple frameworks as this will assist the business in incorporating additional machine learning business solutions as these become available (Sapp, 2017:2).

Even if a reputable service provider is chosen, there are still security risks posed by malicious attackers on outsourced machine learning technology. This emphasises the need to (Gu *et al.*, 2017:11)

- ensure that the channels used to obtain machine learning technology provide guarantees of integrity in transit

- o ensure that repositories require the use of digital signatures for the machine learning models.

Businesses will also need to consider assurance requirements such independent evaluation of service providers' internal control environments. Furthermore, business may want to have a service level agreement in place to set out the responsibilities of the service provider in terms of errors, upgrades, downtime, security and integration. This would support the achievement of the King IV recommendation, which requires the performance of and risks related to outsourced service providers to be managed (Institute of Directors of Southern Africa (IODSA), 2016:62).

5.6 Step 5: Consider user requirements

Users also need to be aware of their own requirements when implementing machine learning technology, especially when considering the risks related to interpretability, errors not being detected and overreliance on machine learning models. Gillion (2017:10) recommends the following considerations be taken into account with regard to user requirements:

- The skills required by accountants may need to be adapted to machine learning. Although accountants will not be able to train machine learning models, which requires a deep understanding and knowledge of machine learning techniques, they may need a basic understanding of machine learning to be able to perform their role when working with experts.

- Furthermore, Gillion (2017:10) recommends that accountants should change their way of thinking, improve their critical thinking skills and communication abilities and become more adaptable to change.

5.7 Step 6: Consider the security requirements

The considerations and recommendations below address the security requirements for machine learning and are focused mainly on data governance and protection, as most of the security risks identified for machine learning technology pertain to malicious data.

- Best practices for data protection may need to be included in a governance policy ("Considerations for senstive data within machine learning datasets", 2017).

- Barreno *et al.* (2010:134) recommend the following defences when considering the security of machine learning technology:
  - Identify the components of the system to which access needs to be controlled, such as the training data, feature selection and model code, and ensure there are adequate access controls to sensitive technology resources (Oracle Corporation, 2018:5).
  - Limit the feedback and therefore the output from the machine learning model to which unauthorised users have access.
  - Build resilience into the learning algorithm to identify contaminated training data or input data; in some cases increasing the complexity of the learning algorithm may defend against security attacks.
  - A virus detection system may be able to reduce the risk of a virus infection.

- Ensure adequate security considerations are in place when acquiring machine learning technology from a service provider, as described in section 5.5.1.

- Adequate controls surrounding the governance of data need to be in place, including access controls to maintain data integrity and to protect privacy rights on sensitive data (Sapp, 2017:23).

- For sensitive data where ownership of the data is at risk, there may be settings that enable safe data sharing and use (The Royal Society, 2017:51). These settings may include encrypting sensitive data fields, processes that scan for sensitive and risky data and providing certain users with read-only rights ("Considerations for sensitive data within machine learning datasets", 2017).

- For data disposal, controls must ensure that machine learning technology is adequately disposed of or data securely deleted at the end of its useful life (ISACA, 2012:165). In addition, users must ensure that machine learning code is periodically examined to determine any unnecessary code that can be removed (Sculley *et al.*, 2015:2498).

- The Oracle Corporation (2018:5) recommends identifying potential security risks through alerts from system and application logs, as these indicate user activities and changes in security configurations.

- Users should be allowed the minimal rights and permissions required to complete their required actions (Oracle Corporation, 2018:5).

- Finally, users may want to consider incorporating machine learning-based adaptive intelligence into their internal control environment as part of their risk assessment procedures and internal controls to provide an intelligent security framework (Oracle Corporation, 2018:10).

5.8 Conclusion

The guidelines presented in this chapter provided the various considerations to be made when implementing machine learning technology, based on the risks, benefits and limitations identified in chapter 4. These guidelines can assist users in determining whether to proceed with implementing machine learning technology, as well as in aligning the technology to the accounting process goals, and highlight the user's role when implementing this technology.

Chapter 6: Conclusion

Prior research has shown a developing need for users to obtain an understanding of machine learning. For professional accountants, PwC (2015:16) even recommends that undergraduate accounting programmes should include advanced topics on machine learning as part of the curriculum.

The aim of this study was to enhance users' understanding of machine learning technology specifically in the performance of the accounting processes. This was achieved by identifying the accounting tasks that machine learning could perform and describing how this technology functions, as well as the risks, benefits and limitations associated with machine learning, including those that have a specific impact on the achievement of the accounting objectives. Based on the risks identified, steps to take when implementing machine learning technology in the accounting process were developed.

This study focused on three accounting processes, namely, the translation of manual and electronic documents into accounting information, the reconciliation of financial information and the preparation of management accounts. As demonstrated in chapter 2, these processes consist of numerous tasks, many of which are enabled by existing technologies. Without the capabilities of these technologies, much of the functionality of machine learning could not be utilised.

Having identified the accounting tasks, certain of these tasks presented learning problems to which machine learning techniques could provide a solution. Chapter 3 discussed the learning problems that could be addressed by machine learning, as well as the different machine learning techniques available to address these problems. It was shown that there may often be more than one machine learning technique available to address a learning problem and, in certain cases, the most beneficial solution may even be a combination of various machine learning techniques.

Subsequently, each of the functions of the relevant machine learning techniques was discussed with the aim of providing accounting users with an understanding of them. The design and functionality of the technology was explained not only for the purpose of understanding it but also for identifying the associated risks, benefits and limitations. It is, however, important here to bear in mind that the technology was not explained at the technical level required to develop the technology.

The study then considered the risks, benefits and limitations of the machine learning technology. In trying to assist users in understanding the technology, chapter 4 considered the risks and benefits of each of the machine learning techniques and mapped those to the different accounting objectives as determined by the *Conceptual Framework for Financial Reporting,* as approved by the International Accounting Standards Board. These risks were summarised and linked to the relevant user considerations. The majority of the risks identified in chapter 4 were data risks, which included risks pertaining to data ingestion, pre-processing of data, impact on the machine learning model, feature selection, the training and testing set, infrastructure requirements, data security risks, and termination and maintenance risks.

Finally, the identified risks, benefits and limitations were used to develop guidelines for accounting users when implementing and using machine learning technology, including the areas where user involvement is required. The findings of chapters 4 and 5 have shown that users would need to give broader consideration to other components in addition to the actual machine learning model. One important broader consideration was the user requirements, which once again highlighted and confirmed the need for users to have an understanding of the technology, thus re-establishing the motivation for this study.

When evaluating the considerations and possible options that users have available to address the identified risks, it was interesting to note that in certain instances machine learning technology gives rise to risks and is also able address the risk. For example, in the case of security risks, as described in chapter 4, users could consider incorporating machine learning-based adaptive intelligence in their security framework (Oracle Corporation, 2018:10).

In summary, the user has a key role to play when implementing and using machine learning technology in the accounting processes and should be equipped with an understanding of the technology and the risks and limitations, as well as the benefits of using the technology. In doing so, consideration should be given not only to machine learning technology but also to addressing the risks pertaining to all the components that enable the functioning of the technology in order to ensure alignment with the accounting process goals.

Further research which may be of value is the adaptation of a data governance framework applicable to machine learning technology. It should be noted that this study has just emphasised the need for adequate data governance and provided data governance considerations. Further research may therefore be required to assess the impact on an existing business when implementing machine learning in its current accounting processes and applying the steps provided in this research. There may also be a need to assess the impact of machine learning on the accounting profession by considering specific case studies of companies that have adopted the technology.

List of references

ABBYY. 2017. [Online], Available: https://www.abbyy.com/en-apac/flexicapture/features/.

ABBYY Technologies. n.d. *What is OCR and OCR technology?* [Online], Available: https://www.abbyy.com/en-apac/finereader/what-is-ocr/ [2017, August 06].

Aberdeen Group. 2017. *The financial close: automation, efficiency, and the emergence of RPA*. [Online], Available: http://www.kofax.com/~/media/Files/Kofax/whitepaper/rp-the-financial-close-automation-efficiency-and-the-emergence-of-rpa-aberdeen-en.pdf [2017, August 20].

Adhitama, M.A., Sarno, R. & Sarwosri. 2016. Account charting and financial reporting at accounting module on Enterprise Resource Planning using tree traversal algorithm. In Surabaya, Indonesia: IEEE *Proceedings of 2016 International Conference on Information and Communication Technology and Systems, ICTS*. 20–25.

Ahmed, M., Mahmood, A.N. & Islam, M.R. 2016. A survey of anomaly detection techniques in financial domain. *Future Generation Computer Systems*. 55:278–288.

Albawi, S., Mohammed, T.A. & Al-Zawi, S. 2017. Understanding of a Convolutional Neural Network. In Anatalya: IEEE *International Conference on Engineering and Technology (ICET)*. 1–6.

Alpar, P. & Winkelsträter, S. 2014. Assessment of data quality in accounting data with association rules. *Expert Systems with Applications*. 41(5):2259–2268.

Alreemy, Z., Chang, V., Walters, R. & Wills, G. 2016. Critical success factors (CSFs) for information technology governance (ITG). *International Journal of Information Management*. 36(6):907–916.

Amani, F.A. & Fadlalla, A.M. 2017. Data mining applications in accounting : A review of the literature and organizing framework. *International Journal of Accounting Information Systems*. 24:32–58.

Amtrup, J.W., Thompson, S.M., Kilby, S. & Macciola, A. 2015. *Patent No. 9058580 B1*. United States.

Appelbaum, D., Kogan, A., Vasarhelyi, M. & Yan, Z. 2017. Impact of business analytics and enterprise systems on managerial accounting. *International Journal of Accounting Information Systems*. 25(April):29–44.

Arnab, A., Zheng, S., Jayasumana, S., Romera-Paredes, B., Larsson, M., Kirillov, A., Savchynskyy, B., Rother, C., et al. 2018. Conditional Random Fields Meet Deep Neural Networks for Semantic Segmentation. [Online], Available: http://www2.maths.lth.se/vision/publdb/reports/pdf/arnab-etal-2017.pdf. [2018, March 23].

Ayodele, T.O. 2010a. Introduction to Machine Learning. In Intech *New Advances in Machine Learning*. 1–8.

Ayodele, T.O. 2010b. Types of Machine Learning Algorithms. In Y. Zhang (ed.). Intech *New Advances in Machine Learning*. 19–48.

Bailey, J.E. & Pearson, S.W. 1983. Development of a Tool for Measuring and Analyzing Computer User Satisfaction. *Institute of Management Sciences*. 29(5):530–545.

Barreno, M., Nelson, B., Joseph, A.D. & Tygar, J.D. 2010. The security of machine learning. *Machine Learning*. 81(2):121–148.

Belfo, F. & Trigo, A. 2013. Accounting Information Systems: Tradition and Future Directions. *Procedia Technology*. 9:536–546.

Bengtsson, H. & Jansson, J. 2015. Using Classification Algorithms for Smart Suggestions in Accounting Systems. Unpublished masters thesis. Gothenburg: Chalmers University of Technology. [Online], Available: http://publications.lib.chalmers.se/records/fulltext/219162/219162.pdf [2018, February 7].

Berka, P. & Rauch, J. 2010. Machine Learning and Association Rules. In *19th International Conference On Computational Statistics COMPSTAT 2010*. 1–29. [Online], Available: https://www.rocq.inria.fr/axis/COMPSTAT2010/TU_Berka-Rauch_paper.pdf. [2018, March 23].

Bezerra, E., Mattoso, M. & Xexéo, G. 2006. Semi-Supervised clustering of XML documents: Getting the most from structural information. In Atlanta *Proceedings of the 22nd International Conference on Data Engineering Workshops*. 88.

BlackLine. 2014. *What are account reconciliations*? [Online], Available: https://www.blackline.com/blog/account-reconciliation/ [2017, June 11].

Bose, I. & Mahapatra, R.K. 2001. Business data mining: a machine learning perspective. *Information & Management*. 39:211–225.

Brady, E.S., Leider, J.P., Resnick, B.A., Natalia Alfonso, Y. & Bishai, D. 2017. Machine-learning algorithms to code public health spending accounts. *Public Health Reports*. 132(3):350–356.

Bräuning, M., Hüllermeier, E., Keller, T. & Glaum, M. 2016. Lexicographic preferences for predictive modeling of human decision-making: A new machine learning method with an application in accounting. *European Journal of Operational Research*. 258:295–306.

Breck, E., Polyzotis, N., Roy, S., Whang, S.E. & Zinkevich, M. 2018. *Data Infrastructure for Machine Learning*. [Online], Available: http://www.sysml.cc/doc/9.pdf. [2018, May 20].

Brownlee, J. 2013. *How to Identify Outliers in your Data*. [Online], Available: https://machinelearningmastery.com/how-to-identify-outliers-in-your-data/ [2018, May 28].

Bryant, A. 2002. Re-grounding Grounded Theory. *Journal of Information Technology Theory and Application*. 4(1):25–42. [Online], Available: http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1186. [2017, August 7].

Bucheli, H. & Thompson, W. 2014. *Statistics and Machine Learning at Scale*. [Online], Available: https://media.bitpipe.com/io_12x/io_123306/item_1144704/Statistics and Machine Learning at Scale-discovery.pdf [2018, May 27].

Buchholz, S., Jones, B. & Krumkachev, P. 2016. *Reimagining core systems Emerging Technologies*. [Online], Available: http://d27n205l7rookf.cloudfront.net/wp-content/uploads/2016/02/DUP_TechTrends2016.pdf [2018, May 27].

Burrell, J. 2016. How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*. 3(1):1–12.

Castle, N. 2018. *What is Semi-Supervised Learning?* [Online], Available: https://www.datascience.com/blog/what-is-semi-supervised-learning [2018, February 15].

Cearly, D.W., Walker, M.J. & Burke, B. 2016. *Top 10 Strategic Technology Trends for 2017*. [Online], Available: https://www.gartner.com/doc/3471559?ref=unauthreader&srcId=1-6595640685 [2017, March 08].

Chadha, P. & Singh, G.N. 2012. Classification Rules and Genetic Algorithm in Data Mining. *Global Journal of Computer Science and Technology Software & Data Engineering*. 12(15):51–54. [Online], Available: https://globaljournals.org/GJCST_Volume12/6-Classification-Rules-and-Genetic-Algorithm.pdf [2018, January 23].

Chen, L., Wang, S., Fan, W., Sun, J. & Satoshi, N. 2015. Deep learning based language and orientation recognition in document analysis. In Tunis, Tunisia: IEEE *2015 13th International Conference on Document Analysis and Recognition*. 436–440.

Chew, P.A. 2014. *Patent No. 8,639,596 B2*. United States.

Chew, P.A. & Robinson, D.G. 2012. Automated account reconciliation using probabilistic and statistical techniques. *International Journal of Accounting & Information Management* 20(4):322–334.

Cokins, G. 2013. Top 7 trends in management accounting. *Strategic Finance*. 95(6):21–30.

Collobert, R. & Weston, J. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning* 160–167.

Companies and Intellectual Property Commission. 2017. *Why the CIPC Decided to Mandate XBRL*. [Online], Available: http://www.cipc.co.za/files/8214/9139/3411/Why_XBRL_2.pdf [2017, August 10].

Considerations for senstive data within machine learning datasets. 2017. [Online], Available: https://cloud.google.com/solutions/sensitive-data-and-ml-datasets [2018, June 16].

Cooper, J. & Pattanayak, S. 2011. *Chart of Accounts : A Critical Element of the Public Financial Management Framework*.

Corless, K., De Villiers, J., Garibaldi, C. & Norton, K. 2018. *Reengineering technology*. [Online], Available: https://www2.deloitte.com/content/dam/insights/us/articles/Tech-Trends-2018/4109_TechTrends-2018_FINAL.pdf [2018, March 27].

CPA Australia Ltd. 2011. *Dashboard Reporting*. Melbourne.

Creswell, J.W. 2009. *Qualitative, Quantitative, and Mixed Methods Approaches*. Third ed. Los Angeles: Sage Publications.

Daecher, A. & Schmid, R. 2016. *Internet of Things: From sensing to doing*. [Online], Available: https://www2.deloitte.com/content/dam/Deloitte/global/Documents/Technology/gx-tech-trends-2016-innovating-digital-era.pdf [2018, May 27].

Dataiku. 2017. *Machine learning Basics: An Illustrated Guide for Non-Technical Readers*. [Online], Available: https://pages.dataiku.com/machine-learning-basics-thank-you?submissionGuid=80d21f82-ac46-45d0-969a-cd9914d06af9 [2018, February 06].

DBASS Chartered Accountants. 2014. *How to Produce Valuable Management Accounts Using Sage Line 50*. [Online], Available: http://www.dbass.ie/advice-resources/advice/how-to-produce-valuable-management-accounts-using-sage-line-50 [2017, August 20].

De Leone, R. & Minnetti, V. 2015. *Electre Tri-Machine Learning Approach to the Record Linkage Problem*. [Online], Available: http://arxiv.org/abs/1505.06614.[2017, November 22]

Demchenko, Y., De Laat, C. & Membrey, P. 2014. Defining Architecture Components of the Big Data Ecosystem. In Minneapolis, MN: IEEE *2014 International Conference on Collaboration Technologies and Systems (CTS)*. 104–112. [Online], Available: https://ieeexplore.ieee.org/abstract/document/6867550/.[2018, May 21]

Deming, Z. n.d. [Online], Available: https://www.blackline.com/assets/docs/uploads/continuous-accounting-ebook.pdf.[2017, November 21].

Dietterich, T.G. & Kong, E.B. 1995. *Machine Learning Bias, Statistical Bias, and Statistical Variance of Decision Tree Algorithms*. *Citeseer*. Vol. 255. 0-13.

Du, J. 2017. Automatic Text Classification Algorithm based on Gauss Improved Convolutional Neural Network. *Journal of Computational Science*. 21:195–200.

Ehsani, A.H., Quiel, F. & Malekian, A. 2010. Effect of SRTM resolution on morphometric feature identification using neural network-self organizing map. *GeoInformatica*. 14(4):405–424.

Emmanuel, A.S. & Nithyanandam, S. 2014. An Optimal Text Recognition and Translation System for Smart phones Using Genetic Programming and Cloud. *International Journal of Engineering Science and Innovative Technology (IJESIT)*. 3(2):437–443.

European Union. 2018. *Seven Steps for business to get ready for the General Data Protection Regulation*. [Online], Available: https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-business-7-steps_en.pdf [2018, June 12].

Everest Group. 2014. *Service Delivery Automation (SDA) Market in 2014: Moving Business Process Services Beyond Labor Arbitrage*. (White paper EGR-2014-1-R-1264). [Online], Available: http://www.everestgrp.com/wp-content/uploads/2014/10/Service-Delivery-Automation-Market-in-2014-Everest-Group-Report.pdf. [2017, November 15]

Fedyk, A. 2016. *How to Tell If Machine Learning Can Solve Your Business Problem*. [Online], Available: https://hbr.org/2016/11/how-to-tell-if-machine-learning-can-solve-your-business-problem [2017, July 29].

Gardent, C. & Perez-Beltrachini, L. 2017. A Statistical, Grammar-Based Approach to Microplanning. *Computational Linguistics*. 43(1):1–30.

Ghanem, K. 2012. A simple process to speed up machine learning methods: Application to Hidden Markov Models. In S. Vaidyanathan & D. Nagamalai (eds.). Dubai *The First International Conference on Advanced Information Technologies & Applications*. 161–171.

Gillion, K. 2017. *Artificial intelligence and the future of accountancy*.

Gorbunova, A. & Bochkarev, S. 2011. Automation of the functions of management accounting and controlling for industrial enterprises. In U. Fissgus, B. Krause, A. Kostygov, A. Leonid Mylnikov, & D.-I. Vladimir Alikin (eds.). Koethen: Perm State Technical University *Solutions of applied problems in control, data processing and data analysis*. 23–34. [Online], Available: http://www.inf.hs-anhalt.de/~zischner/ic2/doc/Proceedings_of_the_Conference.pdf. [2017, June 12].

Goussies, N.A., Ubalde, S., Fernandez, F.G. & Mejail, M.E. 2014. Optical Character recognition using transfer learning decision forests. In Paris, France: IEEE *IEEE International Conference on Image Processing(ICIP)*. 4309–4313. [Online], Available: http://ieeexplore.ieee.org.ez.sun.ac.za/document/7025875/. [2017, August 27]

Gu, T., Dolan-Gavitt, B. & Garg, S. 2017. *BadNets: Identifying Vulnerabilities in the Machine Learning Model Supply Chain*. [Online], Available: http://arxiv.org/abs/1708.06733 [2018, June 16].

Hadzic, F., Dillon, T.S. & Tan, H. 2007. Outlier Detection Strategy using self-organising maps. In X. Zhu & I. Davidson (eds.). Hershey: IGI Global *Knowledge Discovery and Data Mining: Challenges and Realities*. 224–243.

Hajek, P. & Henriques, R. 2017. Mining corporate annual reports for intelligent detection of financial statement fraud: A comparative study of machine learning methods. *Knowledge-Based Systems*. 128(2017):139–152.

Hamza, H., Belaïd, Y. & Belaïd, A. 2007. A case-based reasoning approach for invoice structure extraction. In Parana, Brazil: IEEE *Ninth International Conference on Document Analysis and Recognition, ICDAR*. 327–331.

Hawkins, D.M. 2004. The Problem of Overfitting. *Journal of Chemical Information and Computer Sciences*. 44(1):1–12.

Heckerman, D. 2008. Innovations in Bayesian Networks. In D.E. Holmes & L.C. Jain (eds.). *Innovations in Bayesian Networks*. Berlin Heidelberg: Spring-Verlag. 33–82.

Institute of Directors of Southern Africa (IODSA). 2016. *King IV Report on corporate governance for south africa 2016*. [Online], Available: https://c.ymcdn.com/sites/www.iodsa.co.za/resource/resmgr/king_iv/King_IV_Report/IoDSA_King_IV_Report_-_WebVe.pdf [2018, May 13].

International Accounting Standards Board. 2018. *Conceptual Framework for Financial Reporting*. IFRS Foundation.

ISACA. 2012. *COBIT 5 Enabling Processes*. [Online], Available: http://www.isaca.org/COBIT/Pages/COBIT-5-Framework-product-page.aspx [2018, April 04].

Jain, A., Jin, R. & Chitta, R. 2014. Semi-supervised Clustering. In C. Hennig, M. Meila, F. Murtagh, & R. Rocci (eds.) *Handbook of Cluster Analysis*. 1–29.

Japkowicz, N. & Stephen, S. 2002. The class imbalance problem: A systematic study. *Intelligent Data Analysis*. 6(5):429–449.

Karamizadeh, S., Abdullah, S.M., Halimi, M., Shayan, J. & Rajabi, M.J. 2014. Advantage and drawback of support vector machine functionality. In *1st International Conference on Computer, Communications, and Control Technology, Proceedings*. 63–65.

Kaur, G. 2014. Association rule mining: A survey. *International Journal of Computer Science and Information Technologies*. 5(2):2320–2324. [Online], Available: http://sci2s.ugr.es/keel/pdf/specific/report/zhao03ars.pdf. [2018, June 4].

Kohlmaier, K., Hess, E. & Klehr, B. 2006. *Patent No. 2006/00899087 A1*. United States.

Kohonen, T. 1990. The self-organizing map. *Proceedings of the IEEE*. 78(9):1464–1480.

Kohonen, T. 1998. The self-organizing map. *Neurocomputing*. 21(1–3):1–6.

Kohonen, T. 2013. Essentials of the self-organizing map. *Neural Networks*. 37:52–65.

Kokina, J. & Davenport, T.H. 2017. The Emergence of Artificial Intelligence: How Automation is Changing Auditing. *Journal of Emerging Technologies in Accounting*. 14(1):115–122.

Kotsiantis, S.B. 2007. Supervised Machine Learning: A Review of Classification Techniques. *Informatica*. 31:249–268.

Krizhevsky, A., Sutskever, I. & Hinton, G.E. 2012. ImageNet Classification with Deep Convolutional Neural Networks. *Advances In Neural Information Processing Systems*. 1–9.

Krutova, A.S. & Yanchev, A. V. 2014. Accounting process technology: Structural and Logical Aspects. *Practical Science Edition “Independent Auditor”*. 2(8):12–19.

Lafferty, J., Mccallum, A. & Pereira, F. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. *Proceedings of the 18th International Conference on Machine Learning: (ICML)*. 282–289.

Larsson, A. & Segerås, T. 2016. Automated invoice handling with machine learning and OCR. KTH Royal Institute of Technology. [Online], Available: https://www.diva-portal.org/smash/get/diva2:934351/FULLTEXT01.pdf. [2017, June 7].

Lawrence, S., Giles, C.L., Tsoi, A.C. & Back, A.D. 1997. Face recognition: A convolutional neural-network approach. *IEEE Transactions on Neural Networks*. 8(1):98–113.

Lee, D., Tsatsoulis, C. & Perry, S.M. 2009. *Patent No. 7,596,545 B1*. United States.

Lopez De Mantaras, R. & Armengol, E. 1998. Machine learning from examples: Inductive and Lazy methods. *Data & Knowledge Engineering*. 25:99–123.

Marsland, S. 2009. *Machine Learning An Algortihmic Perspective*. Boca Raton: CRC Press.

Ming, D., Liu, J. & Tian, J. 2003. Research on Chinese financial invoice recognition technology. *Pattern Recognition Letters*. 24:489–497.

Narasimha Murty, M. & Susheela Devi, V. 2011. *Pattern Recognition: An Algorithmic Approach*. I. Mackie (ed.). London: Springer.

Niedermayer, D. 2008. An Introduction to Bayesian Networks and their Contemporary Applications. *Studies in Computational Intelligence*. 156:117–129.

OCREX. 2017. *Features | Receipt Scanner App | Accounting Invoice Software*. [Online], Available: https://autoentry.com/features/#accuracy-verified-extraction. [2017, August 09].

Oquab, M., Bottou, L., Laptev, I. & Sivic, J. 2014. Learning and transferring mid-level image representations using convolutional neural networks. In Columbus, OH, USA: IEEE *2014 IEEE Conference Computer Vision and Pattern Recognition*. 1717–1724.

Oracle Corporation. 2018. *Machine learning-based adaptive intelligence: The future of cybersecurity Executive summary*. [Online], Available: http://www.oracle.com/us/solutions/cloud/future-of-cyber-security-4302684.pdf [2018, June 12].

PWC. 2015. Data driven: What students need to succeed in a rapidly changing business world. [Online], Available: https://www.pwc.com/us/en/faculty-resource/assets/pwc-data-driven-paper-feb2015.pdf. [2017, March 5].

Rhodes, J.J. & Wheat, D.A. 2015. *Patent No. 20150117721 A1*. United States. [Online], Available: https://www.google.com/patents/US20150117721. [2017, August 11].

Richins, G., Stapleton, A., Stratopoulos, T.C. & Wong, C. 2017. Big Data Analytics: Opportunity or Threat for the Accounting Profession? *Journal of Information Systems*. 31(3):63–79.

Roy, E. 2005. *Patent No. 2005/0102212 A1*. United States. [Online], Available: https://patents.google.com/patent/US20050102212A1/en [2017, October 29].

Sainani, K.L. 2014. Explanatory Versus Predictive Modeling. *American Academy of Physical Medicine and Rehabilitation*. 6(9):841–844.

Saitta, L. & Neri, F. 1998. Learning in the “Real World”. *Machine Learning*. 30:133–163.

Samoil, L.A. 2015. Multiple Entity Reconciliation. Stockholm: KTH Royal Institute of Technology. [Online], Available: http://www.diva-portal.org/smash/get/diva2:928531/FULLTEXT01.pdf. [2018, January 18]

Sapp, C.E. 2017. *Preparing and Architecting for Machine Learning*. [Online], Available: https://www.gartner.com/binaries/content/assets/events/keywords/catalyst/catus8/preparing_an. [2018, April 22].

Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., et al. 2015. *Hidden technical debt in machine learning systems*. [Online], Available: http://papers.nips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf. [2018, May 20].

SMACC. 2017. *Are SMEs able to keep an eye on their corporate finances in the same way as large corporates do?* [Online], Available: https://www.smacc.io/wp-content/uploads/2017/04/finance-with-artificial-intelligence.pdf [2018, February 05].

Someren, M. Van & Urbancic, T. 2006. Applications of machine learning: matching problems to tasks and methods. *The Knowledge Engineering Review*. 20(4):363–402.

Sorio, E. 2013. *Machine Learning Techniques for Document Processing and Web Security*. Published doctors dissertation. Trieste: University of Trieste. [Online], Available: http://hdl.handle.net/10077/8533 [2017, August 29]

Sorio, E., Bartoli, A., Davanzo, G. & Medvet, E. 2010. Open world classification of printed invoices. In Manchester, United Kingdom: ACM *Proceedings of the 10th ACM symposium on Document engineering*. 187–190.

Suer, M., ITIL & Nolan, R. 2015. *Using COBIT to deliver information and data governance.* [Online], Available: http://www.isaca.org/COBIT/focus/Pages/Using-COBIT-5-to-Deliver-Information-and-Data-Governance.aspx [2018, May 13].

Sutton, C. & Mccallum, A. 2007. An Introduction to Conditional Random Fields for Relational Learning. In L. Getoor & B. Taskar (eds.). Cambridge, Massachusetts: The MIT Press *Introduction to Statistical Relational Learning*. 93–128.

Sutton, C. & Mccallum, A. 2011. An Introduction to Conditional Random Fields. *Foundations and Trends in Machine Learning*. 4(4):267–373.

Sutton, S.G., Reinking, J. & Arnold, V. 2011. On the Use of Grounded Theory as a Basis for Research on Strategic and Emerging Technologies in Accounting. *Journal of Emerging Technologies in Accounting*. 8(1):45–63.

Sutton, S.G., Holt, M. & Arnold, V. 2016. “The reports of my death are greatly exaggerated”: Artificial intelligence research in accounting. *International Journal of Accounting Information Systems*. 22:60–73.

Takaki, J. & Ericson, G. 2018. *Assign Data to Clusters*. [Online], Available: https://cloud.google.com/blog/big-data/2018/01/problem-solving-with-ml-automatic-document-classification [2018, February 07].

Takaki, J., Petersen, T. & Ericson, G. 2018. *K-Means Clustering*. [Online], Available: https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/k-means-clustering [2018, February 07].

The Royal Society. 2017. Machine learning: the power and promise of computers that learn by example. Report. [Online], Available: https://royalsociety.org/~/media/policy/projects/machine-learning/publications/machine-learning-report.pdf. [2018, April 6].

Thomassey, S. & Fiordaliso, A. 2006. A hybrid sales forecasting system based on clustering and decision trees. *Decision Support Systems*. 42(1):408–421.

Torraco, R.J. 2005. Writing Integrative Literature Reviews: Guidelines and Examples. *Human Resource Development Review*. 4(3):356–367.

Trigo, A., Belfo, F. & Estébanez, R.P. 2014. Accounting Information Systems: The Challenge of the Real-time Reporting. *Procedia Technology*. 16:118–127.

Trintech. 2017. *There’s More to Operational Reconciliations than Matching*. [Online], Available: https://www.trintech.com/resources/record-report-insights/theres-more-to-operational-recs-than-matching/. [2017, July 23].

USC Libraries. n.d. [Online], Available: http://libguides.usc.edu/writingguide/researchdesigns. [2018, Januaray 18].

Valavanis, K.P., Kokkinaki, A.I. & Tzafestas, S.G. 1994. Knowledge-based (expert) systems in engineering applications: A survey. *Journal of Intelligent and Robotic Systems*. 10(2):113–

145.

Ventana Research. 2016. *The Value of Continuous Accounting for Business*. [Online], Available: https://www.ventanaresearch.com/white_paper/office_of_finance/the_value_of_continuous_accounting_for_business [2017, July 30].

Vihinen, M. 2012. How to evaluate performance of prediction methods? Measures and their interpretation in variation effect analysis. *BMC Genomics*. 13(Supplement 4):1–10.

Warwick, K. 2012. *Artificial Intelligence: the basics*. New York: Routledge.

Watson, I. 1999. Case-based reasoning is a methodology not a technology. *Knowledge-Based Systems*. 12(5–6):303–308.

Whittington, G. 2007. *Profitability, Accounting Theory and Methodology: The Selected Essays of Geoffrey Whittington*. Abingdon: Routledge. [Online], Available: https://books.google.co.za/books?id=mH9nysHp0TgC&pg=PA198&dq=define+%22Management+accounts%22&hl=en&sa=X&ved=0ahUKEwjU0ZX4tubVAhVIL8AKHR_lB2oQ6AEILzAB#v=onepage&q&f=false [2017, August 20].

Winkler, W.E. 2014. Matching and record linkage. *Wires Computational Statistics*. 6(5):313–325. [Online], Available: http://onlinelibrary.wiley.com/doi/10.1002/wics.1317/full. [2017, November 22].

Witten, I.H., Frank, E., Hall, M.A. & Pal, C.J. 2016. *Data mining: Practical Machine Learning Tools and Techniques*. Fourth ed. Cambridge: Morgan Kaufman.

Yseop. 2017. *Integrating Natural Language Generation Software Into Finance Workflows*. [Online], Available: http://compose.yseop.com/downloads/integrating-natural-language-generation-finance-workflows/ [2018, September 10].

Zhang, W., Tang, X. & Yoshida, T. 2015. TESC: An approach to text classification using Semi-supervised Clustering. *Knowledge-Based Systems*. 75:152–160.

Zheng, J., Zhou, Y., Deng, T. & Yang, X. 2017. A self-trained semi supervised fuzzy clustering based on label propagation with variable weights. *Proceedings of the 29th Chinese Control and Decision Conference, CCDC 2017*. 7447–7452.