“How can process mining be used to identify Robotic Process Automation opportunities?”

UNIVERSITY OF TWENTE.

Wytze Jan Haan

s1561731

w.j.haan@student.utwente.nl

Faculty of Behavioural, Management and Social Sciences

February 2021

Confidentiality note: Due to the potentially sensitive nature of information, the financial institution at which the research was conducted is kept anonymous. Divisions, departments or specific processes are also not named for this reason.

Abstract

This thesis aims to create a framework via which organizations can use process mining to find and prioritize processes suitable for improvement through Robotic Process Automation (RPA). Synergies between process mining and RPA are explored, finding that the two technologies supplement each other but also work at different levels of process abstraction. Process mining can enhance the implementation of Robotic Process Automation by increasing process understanding, checking process quality, evaluating the impact of implementation, and by being used as a tool to discover new RPA opportunities.

Based on literature a framework is developed for discovering new RPA opportunities through process mining. The developed framework proposes several indicators to measure the potential value of automating each process-step. Using these values, and an estimated cost of implementation, the process-steps can be prioritized to see which should be focused on first. This is an improvement over the currently common manual process selection based on ease of implementation, and can be used to support the business case for RPA. The framework is then applied in a use case at a financial institution, proving to be effective at discovering new opportunities. Limitations exist in data quality, ability to verify the results, and applying the framework in different cases with input from more experts. These areas warrant further research to confirm the added value of the framework.

Table of Contents

1. Introduction

Innovation is essential for companies and organizations to be effective at carrying out their mission. This is no less true in the financial sector; traditional financial businesses must adapt their internal processes and services to keep up.

This report details a research conducted at a financial institution, with a research focus on innovation through technology in order to enhance internal business processes. With the recent drive for process digitization, new information technologies can offer added value. The financial institution has been exploring some of these technologies, and is looking to explore the technologies of process mining and Robotic Process Automation.

Process mining is a technology which creates a complete visual model of a process based on digital event logs available in information systems. This model reflects how the process is truly being carried out, as opposed to theoretically, and allows for analyses regarding performance and conformance. This data driven approach has several advantages over traditional methods of analysis, which make use of interviewing employees carrying out the process and inspecting a small sample of process data (Davenport, 2019). It allows for gaining insight into a process based on complete and quantified data, enhancing the quality of analysis regarding performance and conformance. This is also how it varies from other data-driven analysis approaches such as data mining and statistical models, which only focus on in- and output. Process mining explores the entire process from input to output.

Robotic Process Automation (RPA) is a form of automation whereby a ‘virtual worker’, or robot, is created to carry out tasks normally done by a human. This means the robot simulates the way a human interacts with the system; with mouse clicks and keystrokes, copy and pasting information, etc. Unlike regular digital automation, RPA works on the UI level of systems. Complicated back-end systems, dependent on legacy software, are expensive to adapt for regular automation. As RPA works on the UI level, these changes are not required, making it a much cheaper and safer solution suitable for small automation projects. The bot is programmed by a human to follow certain pathways based on a set of decision rules.

Currently, the institution finds new applications for RPA manually, by talking with process owners and pitching the potential benefits of implementing a bot. Implementations are generally successful and add value, but it can still be difficult to build a strong business case or determine the highest value RPA opportunities. The current method has worked to gain some support for RPA. A more systematic and data-driven approach may help accelerate institution-wide adoption and in proving its value.

There is large speculated potential synergy between process mining and robotic process automation. The financial institution at which this research is conducted is interested in exploring the synergy and added value of these technologies. Specifically, the research is aimed at answering the question: *“How can process mining be used to discover Robotic Process Automation opportunities?”.* The research is carried out by first studying literature and consulting internal experts, and then applying the gained knowledge in a use case.

The results of literature research are elaborated on in chapters 2 and 3. In chapter 4, the gained knowledge is used to develop a framework for identifying RPA opportunities, with an approach for implementation developed in chapter 5. This is then applied to a use case in chapter 6 and 7. In chapter 8 the results of the framework, approach and use case are discussed.

2. Synergies between RPA and process mining

2.1 Added value of process mining

To better understand when process mining truly adds value in process analysis, it is required to gain insight into the current most prevalent process analysis method and its limitations. This was done via literature research and by interviewing an external consultant familiar with process mining. Through the literature review became apparent that the most prevalent method is to hold interviews with employees executing the process and to hold workshops, such as brown paper sessions, to map the process (Davenport, 2019). The external consultant added that there may also be inspection of random data samples, common practice when conducting audits, or measurement of KPI's on output.

Although relatively easy to conduct, there are significant limitations to this method. Interviews and workshops yield largely qualitative process data, and the qualitative process data are rough estimates based on the general perceived process. The reliability of this data is questionable. When there is a large amount of people involved, there will be a large range of answers, and piecing together the process flow may prove to be challenging. One employee may conduct a process different than another employee. Inspection of process data can only be done for small samples, and will not yield a complete picture. KPI's do yield qualitative, reliable and accurate data on some performance aspects; however they generally give little information on the process flow itself.

A process analysis in which these limitations significantly reduce the quality or do not yield the desired output, are instances where process mining has the most added value. In general this will mean large, multi-actor processes; this is where traditionally getting a complete picture is most difficult. (van der Aalst, 2012). These processes also contain more variations that may be missed when only inspecting samples, and answers on flow vary. Added value in process mining is also present when insight into the process flow is desired; measurements solely focusing in- and output are easier to obtain via other methods.

It is important to note that process mining still requires the analyst to talk to those executing the process; process mining only shows what is happening when, not how or why. There is also a large dependency on data; although event logs are generally available, their quality is far from guaranteed. The first time a process is mined, data preparation is the most time consuming part of the project (Leshob, 2018). Afterwards, it is solely the analysis.

There are three forms of process mining: process discovery, performance mining, and conformance checking. Discovery entails constructing a process diagram based on solely the event logs, not having a theoretical process diagram beforehand. Performance mining uses a prior existing model in combination with event logs to gain insight into the performance of a process. Conformance checking uses a prior existing model and event logs to check whether the process is executed according to regulations (Garcia, 2019).

There are still challenges remaining in process mining, some of which will be addressed in the report. Much like any other type of data analysis, process mining is dependent on the quality of input data. It can be challenging to find high quality and complete data for mining. Another challenge is concept drift, which entails the process changing during the time period that is being mined (R'bigui, 2017). Both of these challenges are expanded on and addressed in chapter 3.2. A final challenge is combining process mining with other types of analysis and software, which is the main purpose of the research.

2.2 Challenges in RPA

In order to better understand the synergy between Robotic Process Automation and process mining, it is first needed to gain insight into the current challenges in the field of RPA. This was done via literature research. The challenges can be summarized in the following four points: Process quality, impact evaluation, process understanding, and process discovery.

Process Quality

RPA projects automate tasks that are part of a larger end-to-end process, which can enhance efficiency and performance. However, it does not make any changes to the way the process is carried out. This means that when automating a 'bad' process, it only creates an automated 'bad' process that enhances its inefficiencies. This will generally not solve problems within the process and only produce limited benefits; an inefficient process should not be automated, it should be re-engineered. As such, to maximize the value of RPA projects, it is desirable to know that the processes being automated are already efficient processes in general.

Process Understanding

According to the H2 2017 Global Intelligent Automation Report, 38% of RPA projects fail due to the process that is attempted to automate being more complex than first thought. Essentially, the difficulty of programming an RPA bot is largely dependent on the complexity of the task. If the task is mostly executed in the same way, according to a few simple rules, it is easier to automate. If the task is complex, with many exceptions and unclear rules, it is more difficult to automate. More difficult automation tasks take longer, making them more expensive, and are more likely to fail (Sobczak, 2019). As a result, one should understand in a high level of detail how the process is carried out exactly before starting to automate it, as this determines the cost of automation and risk of failure.

Impact Evaluation

To better understand the impact of an automation project, and thus evaluate automation's value for future applications, it is desirable to have quantified data on process performance. Several measurements may be used as indicators of the process performance, such as a measure of a process output or input. Ideally, after successful RPA implementation, one would see not only a change in output, but also a change in process flow; the activity taking less time and having to be repeated less often due to errors. Gaining these insights into the process flow is still difficult to achieve, especially as the initial process flow might also not have been completely clear to begin with (Suri, 2017).

Process Discovery

RPA has already demonstrated that it is a solution capable of achieving great benefits (Anagnoste, 2018). However, like any solution, it should only be applied to fitting problems. What characteristics make a process suitable for RPA is a topic that has already been explored, and will be further discussed later in the research. What organizations struggle with, however, is actually *finding* these suitable processes. There is no large overview of all processes detailing if RPA could be of value there. This step in implementation is currently the least supported step by RPA vendors (Enriquez et al, 2020).

2.3 Addressing RPA challenges with process mining

As one might suspect by the topic of this research, process mining is theorized to be able to address these problems. Some even call RPA and process mining "a match made in heaven" (Geyer-Klingeberg, 2018), stating that it can solve all these problems. Although this may technically be true, an investigation into how well process mining would address this problem yields some interesting results. This also includes evaluating the added value of process mining over current traditional methods for addressing these problems.

Process Quality

The solution to this challenge is apparent; analyze the end-to-end process of the tasks that are to be automated before starting the automation project. This is one of the main purposes of process mining, and it contains definite advantages over traditional methods in the right setting. When process mining truly adds value is something that is explored later, but a short answer is that generally processes in which RPA can be of value are also processes which can be analyzed well by using process mining.

Process Understanding

The challenge of understanding exactly how a process works with regard to paths, complexity, and variation, is one that seems to be perfectly suited for process mining. Process mining gives a visual representation with all activities and paths between them, as well as measuring variations. There is, however, a complication: for RPA this needs to be known at a very high level of detail.

First and foremost, process mining at this level of detail no longer allows the organization to capitalize on the value process mining offers. Second, one would need a completely different dataset to address this challenge, whereas the other challenges could all be addressed with the same dataset. This means one would have to use process mining with the sole purpose of understanding a single task, which generally is not worth the effort.

There are other technologies more suited to addressing this challenge, such as task mining. Task mining collects UI and system level information as an employee carries out their tasks, creating a process diagram based on this data. This could then also be used as an input for programming an RPA bot. It is similar to process mining in some ways but considered a different technology in this research.

Impact Evaluation

Evaluating the impact of the RPA implementation requires an understanding and measurement of the performance before and after implementation, preferably using the same method. The analysis is similar to one that can be done for the process quality, but does not require the same depth and detail. For impact evaluation, it also holds true that there are a large variety of methods that may be suitable; a benefit of process mining is that a large amount of quantified data is made available, such as process frequencies and throughput times, which are needed to calculate the potential benefit of RPA.

Process mining does also have the added benefit here of requiring much less time the second time around than the first time; the data structure remains the same before and after RPA, thus cutting down immensely on data preparation time. Other methods tend to require the same amount of effort as the first analysis, making them less efficient in comparison.

One could argue that the value of an impact evaluation is limited; the benefits of the RPA project are there, whether they are known or not. Process mining before and after the RPA project just to measure impact may be considered too time consuming; however, if process mining was done before the project, it takes little time to do so again afterwards and can definitely add value for future RPA projects.

Process Discovery

The most prevalent current method for finding RPA use cases is by gathering employee suggestions for tasks they think could be automated. Employees have to give information on some aspects that indicate it is a suitable candidate for RPA, such as complexity or repetitiveness. Another employee then has to look through all the suggestions to pick those with the highest potential value. (Enriquez, 2020; Hatfield, 2020). This is an arduous process and relies heavily on the input of employees, who may not always have the time knowledge needed.

Process mining gives an overview of tasks within the process, as well as data on their performance and interaction. According to the case company's RPA team, if this data can be used to determine if a task is suitable for Robotic Process Automation, then process mining would add great value to solving this challenge compared to current solutions. Not only this, it would also increase the value of process mining as it both find problems and also identify where specific solutions could be implemented. If, and how, this can be done is the main area of this research.

3. Process mining and RPA opportunities

3.1 Goal

The broad synergies between RPA and process mining have been examined; one of these synergies deems further research to better asses its potential. This is the specific use of process mining to discover RPA opportunities.

3.2 Assumptions, Definitions, and Scope

To better approach the research question, some decisions regarding scope are written below. The goal is to keep the research concrete and practical in nature; although there are many theoretical possibilities, the focus is kept on what is currently possible and implementable.

The research involves a use case at a financial institution that is kept anonymous. The case institution has a size between 1,000 and 10,000 employees, and the focus of process mining and RPA are on internal or business-to-business processes. Examples of such processes are onboarding new employees, invoicing, and administrative tasks. Teams and managers focus on process improvement, with a small, dedicated RPA team working on implementing RPA bots throughout the organization. Process mining has been used for a single digit number of use cases, and the institution is looking to expand its use if it keeps proving to be of value.

RPA

RPA is a wide field with a range of technical possibilities. An important distinction is the use of programmable bots and the use of intelligent bots. Programmable bots require a human to program their behavior according to rules, instructions, and parameters. Intelligent bots make use of artificial intelligence and machine learning to learn how a process is carried out by watching an employee (Hawkins, 2018). Due to the large amount of complications that artificial intelligence and machine learning add to a project, the scope is kept to programmable bots. A final remark regarding RPA is that the scope is automating a process performed by one employee, thus automating *single-actor processes*.

Process Mining

Like RPA, process mining includes an array of technical possibilities. Again, the choice is made to not make use of machine learning or artificial intelligence based tools. As the research is conducted at a financial institution already making use of a process mining tool, "Fluxicon Disco", the capabilities of this tool will be used as the general capabilities of process mining. This does present a limitation to the validity of the research, as it could be that certain functionalities of other process mining software allow for much better discovery of RPA opportunities. Fluxicon Disco is not considered as one of the product leaders in 2020

(Modi, 2020), which suggests other process mining software might have more capabilities that would enhance the strength of process mining for the purposes of this research.

Data

As has been touched upon earlier, process mining requires event logs as data input. The quality of the process mining results is completely dependent on the event logs; low quality data will yield low quality results. Data that is rich in information will further enhance results by providing extra attributes to analyze. It is assumed that complete and accurate event logs are available for the process, containing a case ID, start time, end time, and activity. This is realistic as most information systems store this information, although it may take time some to extract this data and prepare it into a readable process mining format. It is not assumed that other information related to automation is present in the logs, for example no fields specifying what percentage of the process is already automated.

Automation data is unlikely to already be present. It could be solved by artificially enriching data with automation information, but this defeats the purpose of using process mining to discover automation opportunities. This is because it would require manual research into each of the process steps and obtaining relevant automation information, while the aim of the research is to do this via process mining and avoid manual research.

Process Granularity

Many of the proposed synergies between RPA and process mining are based on the fact that both technologies have a focus on digital processes; process mining discovers and analyses them, then RPA can automate them based on those results. There is, however, one complicating factor that throws a spanner in the works. It has already been touched upon lightly before, having to do with the level of detail at which a process is approached.

Granularity is the concept of the level of detail at which data is stored; high, or coarse-grained, granularity means it is stored at very high level of detail, and low, or fine-grained, granularity at low level of detail (Keet, 2013). Important to note here is that the higher the granularity, the more difficult it becomes to recognize broader concepts and understand underlying relationships for these concepts. The same holds true for processes.

Essentially, a process consists of a sequence of activities. Each activity, when zoomed in on, is also a process that again consists of a sequence of activities. These activities, again, are a process. Granularity describes at which level of ‘zooming’ a process is being considered. Understanding at which level of granularity a process is taking place is key to determining the added value of process mining and RPA. There is, however, no existing framework or approach to define process granularities. Business Process Management makes use of process decomposition, which decomposes processes into smaller atomic components, but does not provide an unambiguous method to do so (Caeteno, 2010).

The concept of process granularity in workflow systems, highly related to the topic of process mining and RPA, has also been explored and developed into a framework (Vanderfeesten et al, 2008). However, this is at a much higher level of granularity than this research and thus cannot be applied effectively.

An attempt at granularity definitions is made with the purpose of differentiating processes for this research, thus using characteristics which are relevant for process mining and Robotic Process Automation. One concept which has already been touched upon is that of single-actor and multi-actor processes; a single-actor process can be carried out by one person, a multi-actor process requires multiple people. A multi-actor process consists of activities, where each activity is a single-actor process.

RPA generally automates single-actor processes, as it can only work on a single computer. It is possible to automate a multi-actor (two-three person) process with one RPA bot if those process steps all flow into each other and concern the same data or objects. However, this could also be seen as automating a sequence of single-actor activities. If a long sequence of activities can be automated, other more traditional automation methods become more suitable (Penttinen, 2018). As discussed in chapter 2.1, process mining adds most value when mining end-to-end, multi-actor processes.

Steps in a single-actor process may take place in one piece of software, or multiple. One of the strengths of RPA is that it works on the GUI level and can thus easily switch between software. For the creation of event logs, however, it is desirable to have a start time and end time of each activity (single-actor process). Thus, each single-actor process should start and end in the same software to increase the quality of event logs. This is not absolutely necessary, as with some data preparation steps, start and end times could be connected from different software. Steps in a single-actor process can thus happen in different software.

This leads to the following definitions as defined by the author for the purposes of readability and clarity in the research:

Process – A multi-actor process.

Activity – A single-actor process, starting and ending in the same software.

Step – Actions taken to complete an activity.

So a process, such as “register new user”, would consist of activities like “fill out user form”, which would consist of steps like “enter name”. A process view has low level of granularity, whereas the steps view has the highest level of granularity.

This approach allows for a clear formulation on the granularity scope of the research: improving a multi-actor end-tot-end process by automating single-actor activities. This means the process is process mined, and then individual activities are evaluated on their RPA potential.

As a consequence of this low granularity, process mining will not give much insight into the activity itself. To gain more insight, it would be necessary to talk to those executing the process. This is a general limitation of process mining.

Concept Drift

A final topic that needs to be addressed is that of concept drift. Concept drift is the idea that a process changes as it is being analyzed; either gradually or abruptly. It is one of the challenges that still needs to be solved in process mining (Bose et al, 2014). Abrupt changes will either already be known, or are fairly easy to recognize within process mining software, and are thus not considered a threat to validity in this research. Gradual concept drift is difficult to deal with, but an attempt will be made to take it into account.

3.3 Methodology

With the scope, assumptions and definitions set, it is possible to formulate an approach to answer the research question: *How can process mining be used to discover RPA opportunities?*

In chapter 4, based on literature review, the concept of RPA opportunities will be further defined and decomposed into variables. This will be a step towards making the quality of an RPA opportunity measurable. Then, these variables will be concretized to be made measurable through process mining if possible, which will be called a framework. This should allow for the measurement of all activities in a process on their RPA potential.

A framework cannot simply be applied in every situation, and some preparatory work is required to ensure it is applicable and adds value. As such, in chapter 5, a step-by-step approach is written with the goal of eliminating threats to the validity or effectiveness of the framework's output. The steps will also aim to reduce the amount of work required to implement the framework.

The framework and approach are then validated through consulting with an RPA expert from the host institution by discussing completeness, practicality and added value. This allows for incorporating practical insights that might have been missed in literature research.

In chapter 6 the framework will then be tested by applying it to a dataset from the financial institution at which this research is conducted. In chapter 7 these case results will be verified through consultations with the institution's RPA team and limitations identified. Finally, the results of the research will be discussed in chapter 8.

4. Developing the framework

4.1 Defining RPA opportunities

According to the institution's RPA team, determining whether an activity is an opportunity for Robotic Process Automation gives rise to two major questions: "Can we automate this activity, and how difficult is it?" as well as "What is the gained value of automating this activity?" These reflect the general RPA approach of looking at the costs and benefits for each activity (Bellam, 2018).

The driver of cost for RPA projects is the *technical suitability* of the activity; the easier it is to automate an activity the lower the costs. This is because it takes less time and expertise to automate and maintain, as well as the project having a lower chance of failing.

The drivers for benefits of RPA projects are a bit more diverse and depend on the business targets. Although saving costs is an important aspect, the majority of factors focus on *added value* of the automation project such as lower processing time and fewer human errors (Radke, 2020).

4.1.1 Technical Suitability

With the scope of RPA limited to programmable bots, the technical suitability of an activity is a topic that has already been explored extensively. It is also important to note that a bot is not just programmed and then works ad infinitum; it must be maintained. Software updates may break the bot or cause it to behave incorrectly, which must be fixed by the programmer. The more difficult the bot was to program, the more difficult it is to maintain. The following variables summarize aspects of an activity that determine its technical suitability:

Rule based – Because an activity is not carried out exactly the same way each time, the bot needs to decide which steps to execute and in which order. These decisions are based on parameters and rules which the programmer must define, such as a decision tree. If these decisions are difficult to define and program, this increases the complexity of the project and decreases its technical suitability.

Low variations – Each variation in which an activity can be executed has to be programmed manually, taking up more of the programmer's time. High variations also means the bot is more difficult to maintain, as a small change in the UI means having to update each activity variation. The programmer will generally start with the most frequent paths, the "happy flow", and expand from there. Activities with few variations are therefore more suitable for RPA than activities with more variations.

Structured readable input – Each activity has a form of data as input, which may be a picture, slip of paper, email, excel file or more. The bot will need to read this data and process it in order to execute the activity steps. If the bot cannot read the input, it will pass the activity to a human worker. If the bot incorrectly reads the input, it will carry out the process incorrectly and cause errors. Structured, digital input is easily correctly read by bots, making activities with this input more technically suited than other data input types. For example, an excel file with a set format is easy to read because it is digital and always structured in the same way. The bot knows in which cells to look for which data. A handwritten note is very difficult, as handwriting is difficult to read for a computer and the data can be in a different place each time. More advanced software and machine learning techniques do allow for more complex input, but this increases the dependencies of the bot and the time taken to program it.

Mature – If the way an activity is executed changes, either due to software being discontinued, new software acquired or the overarching process changed, the bots will also need to be reprogrammed accordingly. If these changes are significant enough, the bot will need to be programmed completely from scratch, and initial investment will be lost. This concerns the maturity of *activities*. Thus, activities that are expected to change in the near future, or are still changing, are less well suited for automation.

With regards to the maturity of the *process* it also holds true that they are less suited to automation. This is because the input and output of an activity are dependent on other activities within the same process. Thus, if the process changes, or some activities within the process change, this will also impact other activities in their execution.

Translation into indicators

The aspects *rule based*, *low variations*, *structured readable input*, and *activity maturity* are all very activity specific. They depend largely on how the activity is carried out at a high level of granularity, thus requiring information at a high level of granularity. As discussed in section 3.2, process mining functions on a low level of granularity and will thus not yield the required information to measure or assess these aspects. This means that process mining is not suited to assess the technical suitability, or costs, of an RPA project for a specific activity.

The aspect of *process maturity* is one that exists at a lower level of granularity and should be further explored. Process maturity is closely related to concept drift, as it entails that the process is changing as it is measured. A mature process is one that is no longer changing, or only changing very slowly; it is thus also less affected by concept drift. Similarly, if concept drift is a low threat, the process must be mature. If concept drift is a high threat, and the process thus immature, the validity of the entire process mined model is called into question. This plays a much larger role in the project than simply for assessing individual activities, and will be addressed in the general approach instead of under the aspect of technical suitability.

4.1.2 Added Value

As mentioned earlier, there are a more diverse set of aspects to consider here. The importance of each aspect is determined by the business needs of the organization; some activities may be automated to save costs, whereas others may be automated to increase compliance. Below are the most common variables that are considered when determining the worth of an automation project, and each is translated into a measurable indicator.

Human error prone – Making mistakes in an activity has two negative downsides; first of all it decreases the service levels of the activity by not producing the desired results. Second, it causes rework as the activity, plus all activities completed before discovering the error, need to be repeated. Eliminating human error can thus significantly improve performance and add value to a process where human errors are made. RPA is not capable of 'correcting' other mistakes such as systematic ones.

Process mining analysis does not give a direct measurement of human errors, or errors in general, as this data is not expected to be present in the dataset. This means the variable will need to be measured via an indicator.

If an error is made when executing an activity, it essentially means that the activity must be repeated. As such, error is a cause of repetitions and measuring the repetitions can be an indicator for measuring the errors made. However, there can also be different causes for repetition, such as the price of a quotation being adjusted later in the process.

With modern software being designed to disallow most errors, human errors are the most common reason for mistakes being made in a process. This is expected to happen at a rate of about 0.005 for routine simple steps when uninterrupted (Magrabi et al, 2010). The institution's RPA team estimates the amount of steps per activity between 20 and 30; thus an expected error rate per activity is between 0.105 and 0.140.

As a result, any activity which is repeated in less than 10% of cases is not considered prone to human errors. There might also be other reasons for activities to be repeated, which means that any activity which is repeated in more than 10% of cases *may* be prone to human errors. After further research, it proved too difficult to compensate for other reasons of repetitions and thus not possible to isolate human error repetitions. Of note is that a human error should cause exactly one repetition to fix the error, any other repetitions should have other causes.

Human Error Indicator = number of times activity is executed / number of cases activity is carried out for

High frequency – The fixed cost of programming an RPA bot is high, the variable cost of running the bot, or any amount of copies, is very low. This means that the value of a bot is increased when used to execute an activity which is executed in often. Automating an activity that is only carried out a few times a year simply will never be worth it.

Frequency simply has to do with the amount of times an activity is carried out over the course of the dataset time range, which is something that is always and directly measured by process mining. It should be noted that this is about the absolute frequency, and not relative frequency of an activity. It does not matter if an activity is carried out in 0.1% of process instances or 99%, as long as they are both carried out the same amount of times over the same period. The only difference is that if the activity only comprises of 0.1% of process instances, there are most likely other activities that are more valuable to automate.

Frequency Indicator = Number of times activity is executed / Dataset time range in years

Time sensitive – A bot can work much faster than a human and is only limited by the speed of the software UI, which is about four times faster than a human (Gielen, 2019). Besides this, a bot can work 24/7 without breaks. As a result, it carries out an activity much faster than a human, thus increasing the performance of activities and adding value.

The key to reducing the throughput time for an activity lies in two aspects; the waiting time and the execution time. As a bot works about four times faster than a human, it reduces the execution time by 75%. Besides this, the cost of scaling bots is very low, which means that there should be enough bots to carry out the required amount of activities. As a result, it can be assumed that waiting times are reduced to a negligible amount.

Time Reduction Indicator = 0.75 average execution time + average waiting time*

Human productivity – All work carried out by a bot, no longer needs to be carried out by a human. This means employees can focus on more meaningful tasks that make better use of their human capital. This increases the efficiency of employees while also increasing their satisfaction levels.

The amount of work saved, and thus can be spent on other activities, can best be expressed in terms of Full Time Employee (FTE), where 1 FTE equals the amount of time a full time employee works during a year. At the financial institution, this is 36 hours per week for 46 weeks, equaling 1656 hours. Essentially 95% of time of an activity is freed up, with an estimated 5% amount of time is required to handle exceptions the RPA bot is not programmed for. Most process mining technologies can directly show the total amount of time an activity takes for the given dataset, or else it can be calculated by multiplying the average time for the activity by the amount of times the activity is carried out.

*FTE's Saved Indicator = (0.95*Total activity execution time) / (1656 * Dataset time range in years)*

Cost reduction – Although costs are also reduced by increasing compliance and decreasing waste, the main factor is decreased employee costs. It is estimated that bots generally cost about the same as one-third full time employee (Lakshmi et al, 2019; Anagnoste, 2017), although this of course depends on the complexity of the bot and thus difficulty of maintenance. This ignores the initial cost of having to implement the RPA bot. As the focus is on added value, cost reduction only takes into account the saved employee hours and not cost of maintenance or programming a bot.

Much like the productivity variable, this is completely dependent on the amount of FTE's that can be saved by implementing a bot. It therefore uses the same indicator. It should be noted however, that the two are mutually exclusive in terms of benefits; it is not possible to both save costs on employees and increase human productivity costs. It is possible to split the benefits, for example if 2 FTE's can be saved by RPA, 1 saved FTE can be used for cost reduction and 1 saved FTE for increasing human productivity.

If it is chosen to use an amount of FTE's saved for cost reduction, the reduced costs can easily be calculated using the cost of an FTE.

*Reduced costs = FTE's saved*cost of FTE*

Irregular labor – Scaling up or down an activity performed manually is cost intensive; new employees need to be hired or current employees need to be pulled away from other tasks.

The inexperienced employees are less efficient and more prone to error. With bots, however, this is not an issue. If a single bot has already been programmed for an activity, it is possible to simply create 10 identical copies of this bot and run them. When quick changes in required capacity are predictable, they are less costly than when unpredictable.

The irregularity of labor can best be approached by measuring the fluctuations in the number of times an activity is performed over a certain time period. Two approaches can be taken in this regard, the first placing the emphasis on sudden changes in labor demand, and the second placing emphasis on gradual changes in labor demand.

The first method, focusing on sudden changes, can be calculated by using the amount of times an activity was executed in a period and the amount of times an activity was executed in the period before this. A reasonable length for a period would be 1 month, as this should average out random weekly fluctuations that happen in every process. This can be calculated for each month except the first, yielding a maximum and average.

Sudden fluctuation indicator = (number of times activity is executed period x) / (number of times activity is executed period x-1)

For this it is important that the average over the entire dataset should be close to 0, else the activity is growing or shrinking in magnitude. This could indicate that the process has not matured, threatening the suitability of RPA in this process and the validity of process mining due to concept drift.

The second method, focusing on gradual changes, can be calculated in a similar way. Here, however, the minimum of times an activity is executed during the entire dataset is used. This yields a factor of largest growth during the entire dataset period.

Gradual fluctuation indicator = (number of times activity is executed period x) / (minimum number of times activity is executed in any period)

In this case only the maximum is of importance, showing a factor of multiplication by which the activity grows and shrinks over the dataset period. Of course many more interesting statistical analyses can be conducted over both the set of gradual and sudden fluctuation indicators, but the added value over time taken to do these analyses is expected to be limited.

4.2 Measurement through process mining

With each of the variables having been made measurable, it is time to investigate the best method of doing so using process mining software. The software used by the institution is Fluxicon's 'Disco', and will thus also be used in this approach. However, the functionalities offered by 'Disco' are also offered by many other vendors and should not limit external validity.

In order to find the best method of measuring the indicators, Disco's functionalities were first explored. This was done by first attending a two-day Disco training course, and then exploring independently afterwards. An overview of relevant functionalities can be found in appendix B.

In addition to this, practicality was kept in mind for the measurement of the variables. Although 'Disco' yields quantified and accurate data, the operations performed on those measurements to translate them into the RPA suitability indicators are based on rough estimates. Besides this, life is stochastic and it is uncertain that the measurements hold true for the future executions of the process. Therefore, in some cases, accuracy may be forsaken for the case of practicality.

Human Error Prone – Dividing the absolute frequency by the case frequency would yield a mean number of repetitions. Means are notorious for being skewed by high outliers and may not be the most accurate measurement.

Mean repetitions = absolute frequency / case frequency

Much more reliable would be to see what percentage of cases involve repetition of the activity. This can be achieved via an endpoint filter, where the activity is defined as both the start and endpoint. Disco now only displays cases where the activity was repeated. Taking the number of cases with repetition and dividing it by the total amount of cases in the activity will yield an error rate, which is a more accurate measurement than simply taking mean repetitions.

Error rate = cases with repetition / total cases

In case of human errors, it would make sense that the activity only has to be repeated once; it is assumed that a human error is fixed on the first attempt to do so. If possible to filter for only cases repeated exactly once, using these for the error rate would be even more accurate.

High frequency – As mentioned before, this is a simple statistic which measures how often an activity is executed. The way to measure this within Disco is simply by selecting *absolute frequency* to be shown in the model view and noting the values for each activity. The values should then be divided by the dataset time range in years to give a yearly representation. *Absolute frequency* is chosen over *case frequency*, as repetitions are not included in *case frequency* but do factor into how often an activity is performed.

Time sensitive – There are two indicators that can be used for the time reduction for an activity; the mean or the median. These can both be selected as measurements for the model view. There is, however, a complication: there are generally multiple pathways leading into an activity, each with a different mean or median queuing time. Each pathway also has a different amount of cases going through it.

Taking the weighted average of all incoming queueing times would provide the most accurate overall queuing time for an activity. However, especially in complicated processes, there may be a high number of incoming pathways. This means visually inspecting a so-called 'spaghetti diagram', which can be an arduous and time-consuming effort. It is proposed that using the median of the two or three most frequent pathways and taking their weighted average will still give a strong indication of overall queuing time while saving time and keeping the approach practical. This is investigated in appendix C, which yields that using three most frequent pathways creates satisfactory results close to the true weighted average.

*Time Reduction = (0.75*median activity time) + weighted average of median 3 most frequent queuing times*

FTE's saved – FTE's saved replaces both *human productivity* and *cost reduction*, as this is the measurement used for both these indicators. To calculate FTE's saved, the activity can be displayed in *total duration* in the model view. The dataset time range should be known beforehand for data preparation. If not, it can also be found in the global statistics view, which shows the start and end date for the dataset.

*FTE's Saved Indicator = (0.95*Total activity execution time) / (Dataset time range in years)*

Irregular Labor – This is one of the more challenging indicators to measure through 'Disco'. The global statistic view does offer the option of displaying *events over time*, which gives an graphical overview of the number of events per time unit distributed over the course of the entire dataset time range. This only indicates the amount of activities performed, which does not indicate the labor or incoming cases.

There is also the option of displaying *active cases over time*, which displays a graph of how many cases are active per time unit. This is shown over the course of the entire dataset time range. This gives a much stronger indication of the amount of labor that has to be performed and how this changes over time. Limitations are that this can also be caused by lower productivity where cases take longer to finish. It also does not show active cases per activity, but instead for the entire process. Still, if there are significant fluctuations in this graph this could indicate irregular labor. As exact measurements can not be depended on due to the above reasons, only a visual inspection of the *active cases over time* will be used as an indicator for irregular labor.

5. Framework approach

With a clear method on how RPA suitability can be assessed using process mining, it is possible to determine an exact approach for how to do so. The approach is written with the aim of ensuring that the framework is applied correctly and its results add value while keeping the framework practical. The framework and approach are then validated through informal discussion with the institution's RPA expert, after which limitations of the framework and approach are discussed.

5.1 Creating an approach

The first steps will aim to eliminate threats and limitations identified earlier in the research. These are chosen as the first steps because if they yield that the threats or limitations cannot be mitigated, the project should not continue. These threats are process maturity, concept drift, and automating a process that should not be automated. Process maturity and concept drift are addressed in step 1. Checking if automation is the right solution is part of step 2.

Next steps to include based on the framework are process mining the process, calculating the metrics for value of each step, and assessing the technical complexity of each step. Process mining the process is part of step 2, as this is where the general value of process mining can also be leveraged when checking if the process should be automated.

The next logical step would then be to calculate the metrics for each activity, which can be a time consuming process. In order to keep the framework practical, two extra steps are included to reduce wasteful work. These are discarding non-RPA activities (step 3) and optionally discarding low frequency activities (step 4). Carrying out these steps takes little time and ensures that metrics are not being calculated for activities where RPA cannot or will not be implemented anyway. Once this is done, the metrics for the left over activities are calculated in step 5.

As it is unlikely that all possible activities will be automated, it is efficient to look at the most valuable activities first. Therefore, the next step (step 6) is to prioritize the activities based on potential value. This will allow for effectively assessing the technical complexity of each activity in step 7 and thus increasing the practicality of the framework. A selection of activities can then be automated in step 8.

Step 1 – Assess process maturity

Concept drift threatens the validity of process mining, and will thus also threaten the validity of the RPA suitability indicators measured using process mining. Furthermore, a process not being mature undermines the value of RPA, as the RPA implementation may have to be reengineered in the near future. Dealing with concept drift is beyond the scope is beyond this research, but a proposed solution which may reduce risks is simply by talking to process experts. By consulting them on if the process has been changing over the course of the past few months, or if any changes are on the roadmap, an indication of maturity can be found. It is also recommended to use a data range of maximum one year, and no more than three years old.

Step 2 – Check on good process

If the process has been considered mature and concept drift a low threat, the next step is to load the dataset into the process mining tool. Then, the process must be analyzed to determine if it is a 'good process'. The ambiguous word 'good' is used here on purpose as the definition is highly dependent on the process and its domain. The goal is to determine that there are no significant problems in the process that will not be solved by automation. Finding wasteful steps is one example, as automation will not reduce waste. It also means checking that the process effectively yields the desired output, as automation will not change the output.

Performing this analysis is the more traditional purpose of process mining, and the depth of this analysis depends on the desires of the organization. If there are no glaring bottlenecks or other problems, the next step can be taken. If these problems do exist, they must be fixed before returning to step 1 with the newly generated data from the improved process.

Step 3 – Discard non-RPA activities

At this point, some insight into how the activities are executed will have been obtained. With this knowledge, it should be possible to strike out a few activities as being impossible to automate using RPA. This could be physical activities such as 'deliver goods', or activities requiring high human cognition such as 'assess extent of damages based on video evidence'. It is recommended to start with a list of all activities, and remove all non-RPA activities.

Step 4 – Discard infrequent activities

An optional step to reduce work in case many activities are left. Conventionally speaking, an activity which takes places less than 500-1,000 times per year is much less likely to be worth automating. Removing these activities from the list will reduce the amount of activities to inspect for the RPA opportunity metrics.

Step 5 – Calculate metrics

With a list of frequent activities that should be possible to automate with RPA, the next step is to evaluate each of the activities according to the indicators as described in the section "measurement through process mining" (p.12-15). *Irregular labor* can of course only be evaluated once for the entire process, and not individual activities.

Step 6 – Prioritize RPA opportunities

Based on the discovered metrics and value drivers of the process, it should be possible to list the activities in order of highest added value.

Step 7 – Assess technical complexity

Starting with the activity highest on the list, and working down from there, each activity needs to be further investigated on its technical suitability regarding RPA. This includes evaluating them based on the earlier defined indicators *rule based*, *low variations* and *structured readable input*. This will mean talking to an employee carrying out the activity and potentially consulting an RPA expert.

Step 8 – Implement RPA

With a clear overview of activities regarding added value and technical suitability, all information needed to choose which activities to automate using RPA is present. Of course, it may be that no activities are suitable, or all activities are. In case many activities are suited for RPA, it may be worth exploring more traditional forms of automation.

5.2 Validation of Framework

To validate the framework and approach of implementation, it was discussed with the institution's RPA expert. This was chosen because much of the framework was based on theory from literature, and the insights of putting RPA into practice were still lacking. Due to limited time and contacts, only a single expert opinion was taken into account.

Before the discussion, important aspects to asses the quality of the framework were identified. Practicality of the framework has already been touched upon before; application of the framework should not be an arduous process or too time consuming. Besides this, the framework must offer added value over the traditional approach in order to be worth applying. Multiple speculated sources of added value have been derived from literature; which ones hold true and offer value to the institution is something the expert will be able to offer insight into. Finally, the framework must also be complete and cover enough relevant aspects for RPA candidates.

As such, three criteria were established and discussed; completeness, practicality, and added value. With regards to completeness, the expert estimated that the framework covers around 90% of the topics relevant when considering an RPA candidate. The 10% left would be a multitude of small aspects, of which they doubted the added value of further exploring. This due to these aspects only being relevant in specific cases and hurting the practicality of the approach. The framework was thus considered complete.

Practicality of the eight step approach was assessed to be high, as multiple steps were included to reduce unnecessary work. The only questionable aspect of practicality comes in the form of process mining in general, which can be a time-consuming activity in and of itself. Whether this is worth it, is something that is further explored in the evaluation of the research.

Added value of the framework touched on the concepts *process quality*, *impact evaluation* and *process discovery* based on chapter 2.2. Process quality is something that was not considered before when considering activities for RPA and taking this into account would add value to the overall effectiveness of RPA solutions. Interestingly, the expert was very happy to hear about the concept of process maturity; this is a challenge they had ran into often but not known what it was called or how to learn more about this. This was already considered a valuable contribution to their RPA efforts. Process quality of the activity itself is something they did not consider either, as this is something that would be discovered as they code the RPA bot. If it is found that the activity was not carried out efficiently, this is something that can be fixed by directly programming the bot to take a different approach. So, with regards to process quality, the framework adds value when looking at the broader process, but not the activity itself.

Impact evaluation was something that their current approach already somewhat included. RPA software gives an overview of metrics of the RPA bot, with statistics on performance of the activity such as frequency and throughput time. Limitations are that there were only rough estimates of these statistics before RPA was implemented, and that the resulting statistics could not be placed in the process context. Saving an hour in an activity is a huge improvement for a process that takes several hours in total, but less significant if the process takes weeks. These limitations would be addressed via the framework and thus it adds some value.

Process discovery is the main value adding part of the framework. Having a scored overview of the candidates can be a huge help in finding the best suited ones and prioritizing which ones should be addressed first. The expert here also identified the current challenge of business cases, where the RPA team must convince the process owner that implementing RPA would be worth the time and resources. This is generally done by submitting a business case, outlining the costs, benefits, and risks, for the process owner to approve. Costs are well known ahead of time, but benefits are generally estimated and could be considered unreliable. Risks are also broad. With this framework, a reliable and data-based estimate of the benefits can be given. The framework also addresses some risks (such as process quality). The expert explained that this is a significant area for added value in the framework and could drive the implementation of this approach.

Based on these findings, it was decided that further work on developing the framework was not required. The framework was considered to be complete, practical and add value. Before putting it to the test, however, its limitations were also discussed.

5.3 Limitations of Framework

A general limitation of process mining is also of relevance to this framework. This is the fact that in order to be able to properly understand the generated process diagrams, one must talk with the process experts to know what each activity means and how it (should) relate to other activities. This means it does not eliminate the phase of talking to process experts, although it does change the topic of conversation. The focus can now be placed on how the activity is carried out, instead of trying to reconstruct the entire process and measure performance.

A second limitation is the framework’s inability to address the technical suitability of activities, which is directly linked to the cost of implementing RPA. The most significant variable regarding cost is simply the time it takes to program the bot; the more technically suited the process is for RPA, the less time it takes to program the bot. Maintenance is a second factor, but each bot takes the same relative time to maintain; about 10% of programming time according to the RPA team.

Technical suitability of an activity can also be evaluated based on a quick demonstration of the activity being carried out, taking maybe 10-15 minutes. Although perhaps not as accurate and fast as doing it via process mining (if this were possible), this still yields insight into technical complexity. Having to do this after applying the framework, where the most value-adding activities are investigated first, means that this limitation can be mitigated at low cost and effort.

6. Applying the Framework

With more insights into the validity and limitations of the framework, the next step is to put the framework into practice. The goal here is twofold; the first is to identify RPA opportunities within a process for the institution. The second is to evaluate the effectiveness of the framework by comparing the results to those obtained via more traditional methods. Applying the framework to a process will also yield more insight into the practicality of the approach and may bring more limitations or strength to light.

6.1 Data and Data Quality

A candidate process was identified within the institution, being an end-to-end multi-actor process carried out around 12,000 times per year. It consists of 16 activities, named “A” through “P” due to confidentiality. Event logs for this process were stored in three separate software systems, meaning some data preparation steps had to be undertaken. This was done with the support of one of the institution’s data scientists.

First and foremost, the different datasets had to be merged into one complete event log to be loaded into process mining tool ‘Disco’. This involved introducing matching primary and foreign keys to link ID’s of all activity logs. After this, unneeded data columns were deleted to decrease file size. Process experts were consulted to translate activity names into a logical and readable name. Finally, some formatting rules were applied to make sure all date-time fields used the same format. It was here that it was discovered that there were some poor aspects in the data quality. All three systems stored dates of events in format day-month-year, but only one also stored timestamps containing hours and seconds. Besides this, all three systems only stored the end date, and not a start date. This strongly impacts the results of process mining.

The impact of only having end dates is that process mining is not able to discover the time an activity takes, only the time between the completion of two activities. Thus, queueing time and activity time are displayed as a single time in a path between two activities. This has a large impact on measurement of the *FTE’s saved* indicator, which relies on the activity time. It has a small impact on the *Time Sensitive* indicator, which uses a slightly reduced activity time and the queueing time. As a solution for this limitation, an estimate of the activity time is used. These estimates are provided by those executing the process, and although less accurate than if measured by process mining, should still give a ‘close enough’ indication to continue the research.

The result of two systems only storing dates, not timestamps, is more significant. This essentially means that process mining discovers the times between activities in batches of 24 hours. Activities executed on the same day have 0 milliseconds between them, the next day is logged as 24 hours, then 48 hours, etc. This creates a huge error margin in the results, impacting the measurements of *FTE’s saved* and *Time Sensitive.* Using the activity time estimates mitigates most of this impact on *FTE’s saved*. The *Time Sensitive* indicator is still made much more unreliable though, effectively introducing a ±12 hours error margin. This affects eight of the 16 activities: *E, I, J, L, M, N, O, and P*. The significance of this error margin depends on the magnitude of the queueing times and will be evaluated per activity.

It was attempted to obtain more complete timestamps for the activities, but this simply had not been logged in the systems and could thus not be added. Ideally, a different process and thus dataset would be chosen to obtain higher quality data. Finding a dataset, obtaining permissions to mine it and then preparing the dataset is a lengthy process, easily taking two to three weeks. Due to time constraints for this research, the choice was made to continue with the low quality dataset. By taking into account the uncertainty, meaningful results should still be possible to obtain. It also did not impact the entire dataset.

With data preparation finished, the eight steps of the framework could be applied.

6.2 Application

Step 1 – Assess process maturity

The data range was chosen to be one year, stemming from data less than two years old. This already mitigated some risks regarding process maturity. During the data preparation stage, process experts were already consulted. During this phase, they were also asked if the process had been changing recently, or was to undergo significant changes in the near future. The answer here was no, and that the process had been carried out like this for more than a decade. Although this area could be explored further, it was considered enough for the means of this research.

Step 2 – Check on good process

As a first check on good process, the process mining results were compared to the theoretical design of the process. Around 75% of instances matched the theoretical process. Another ~15% of instances deviated from the theoretical process with negligible negative consequences, and did not hinder the quality of the executed process. The remaining 10% consisted of mostly single variations that were hard to analyze regarding impact. Overall, after consulting process owners and experts, the conclusion was that the process met its requirements and did not require significant alterations to further improve.

Step 3 – Discard non-RPA activities

By quick inspection of the activity names, and by already having some background information gained during data preparation, it is easy to disqualify some process steps as RPA candidates. Amongst others, the reasons were that they were physical activities, activities that legally require a person to execute them, or activities that required human-level intelligence. The discarded activities were: *F, G, J, L, M*, and *P*. This also discards four of the eights activities with low data-quality, decreasing the data quality impact on the results.

Step 4 – Discard infrequent activities

After discarding non-RPA activities, 10 activities remained. As this was not a large number of activities to analyze, and to make results as complete as possible, it was chosen not to discard any infrequent activities.

Step 5 – Calculate metrics

Calculation of the metrics was done via inspection of the process model generated through process mining in the tool ‘Disco’. Diagrams for each activity have been provided in Appendix D. A low detail level of the process is given below in figure 1.

A higher level detail diagram is used when finding the three most frequent pathways for queueing time calculations, this is not provided in the report as this ‘spaghetti’ diagram is difficult to read.

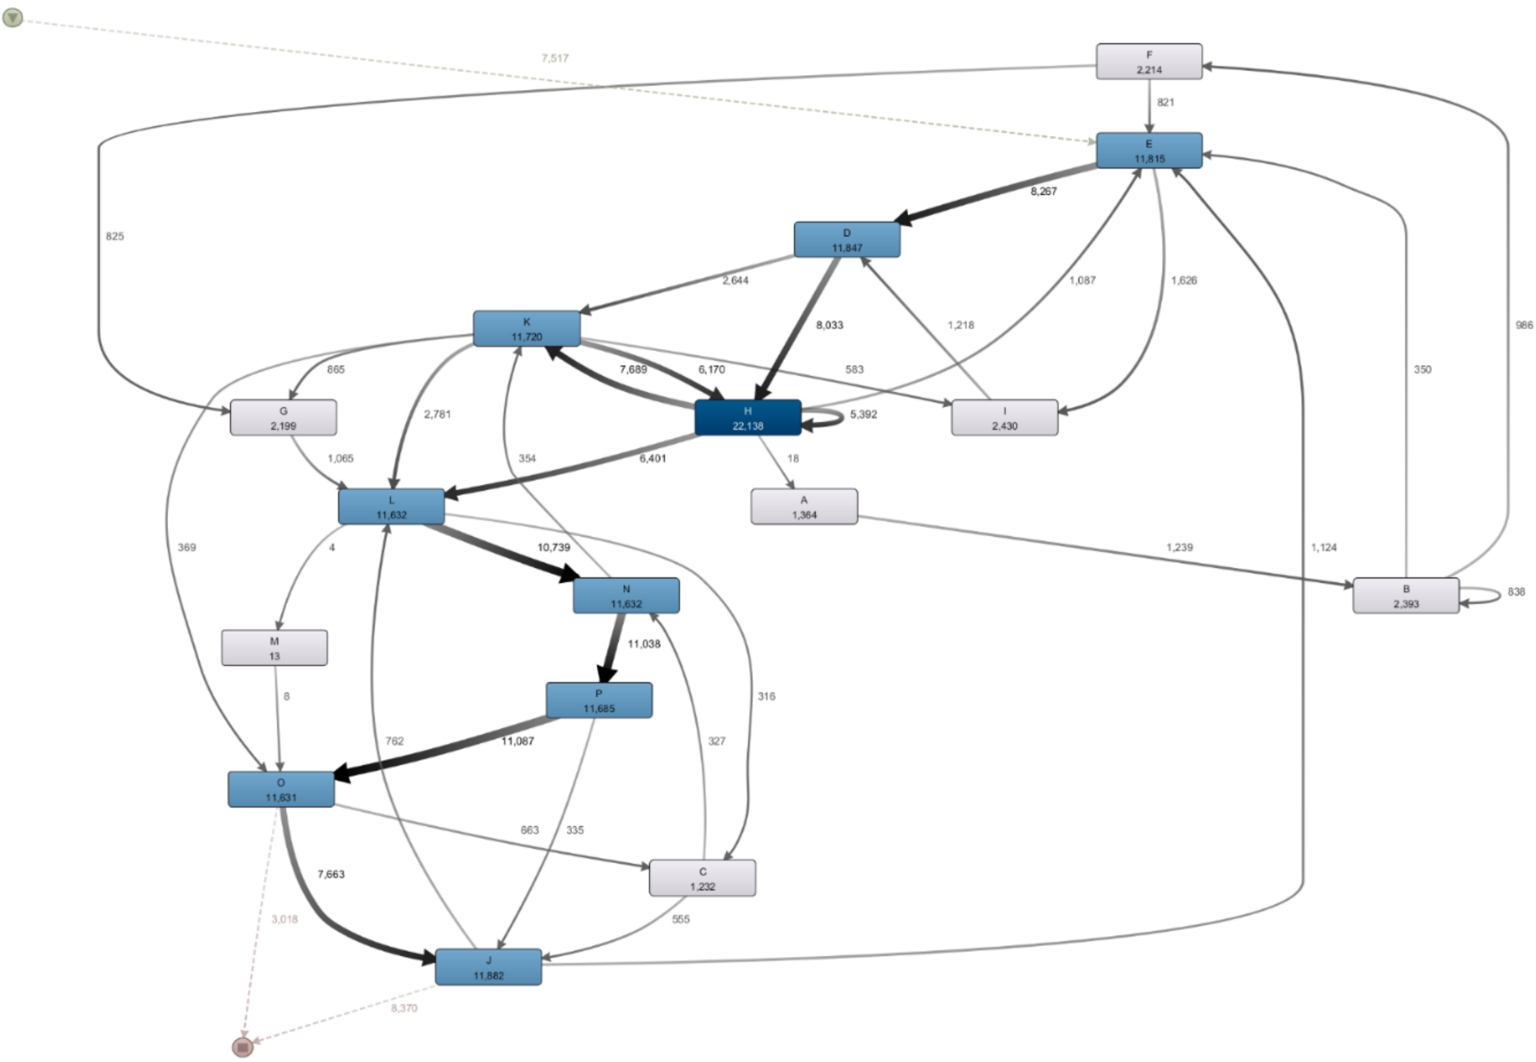

Figure 1: Overview of process, not showing all pathways

An oversight of calculated metrics are given per indicator below.

Human Error Prone

	A	B	C	D	E
Absolute frequency	1364	2393	1232	11847	11815
Case frequency	1314	1302	1212	11816	11784
Repeated cases	18	623	13	31	31
Mean repetitions	**1.04**	**1.84**	**1.02**	**1.00**	**1.00**
Error rate	**1.37%**	**47.85%**	**1.07%**	**0.26%**	**0.26%**

Table 1: Human Error Prone Metrics Activities A-E

	H	I	K	N	O
Absolute frequency	22168	2430	11720	11632	11631
Case frequency	9578	2430	11689	11592	11591
Repeated cases	8462	0	31	39	39
Mean repetitions	**2.31**	**1.00**	**1.00**	**1.00**	**1.00**
Error rate	**88.35%**	**0.00%**	**0.27%**	**0.34%**	**0.34%**

Table 2: Human Error Prone Metrics Activities H-O

Absolute and case frequencies were found by clicking on the activity itself. The repeated cases were found by applying an endpoint filter, where the activity was both the starting and finishing point and using the 'trim longest' setting. This displays all cases where the activity was performed at least once. A follower filter was then applied on top of this, where the activity had to be eventually followed by itself. This means that all cases where the activity was performed only once are filtered out, and only cases where the activity was repeated are shown. The case frequency of the activity in this filtered state thus displays the amount of cases where the activity was repeated at least once.

High Frequency

	A	B	C	D	E
Absolute Frequency	1364	2393	1232	11847	11815

Table 3: Frequency Metric Activities A-E

	H	I	K	N	O
Absolute Frequency	22168	2430	11720	11632	11631

Table 4: Frequency Metric Activities H-O

Time Sensitive

In the framework, median activity time is taken based on the process mining data. However, in this data, there are no starting times for activities and all activity times are thus displayed as *instant* (0 milliseconds). Given this data, process mining only displays the time between the ending of two activities as a pathway. This is the sum of both queueing time and activity time. Although this could still work for the time reduction indicator calculations, the FTE's saved indicator is reliant on having separate activity times. Thus, another method must be found for finding an activity time.

For practical reasons, it was chosen to simply ask employees carrying out the activities for an estimate in the time it takes on average. If a range was given, the middle value was chosen. For example, if the answer was 15-20 minutes, 17.5 minutes was chosen as the value. This is less accurate and reliable than finding the value through process mining, and taking more time as well, this does allow for continuing the evaluation of the framework.

In order to find then the queueing time, the median time of a pathway was taken and the estimated activity time subtracted from it. For example, in figure 2, the pathway in the process mining diagram displays a median duration of 4.6 days (110.4 hours) for activity *L*. If activity *L* is estimated to take 1 hour, the queueing time is calculated as:

$$110.4 - 1 = 109.4 \; hours$$

This is done for each of the three most frequent pathways, where a negative time is not possible and rounded up to 0. Finally, to express the time saved as a percentage of the total process, the value is divided by the median total process time (632.1 hours).

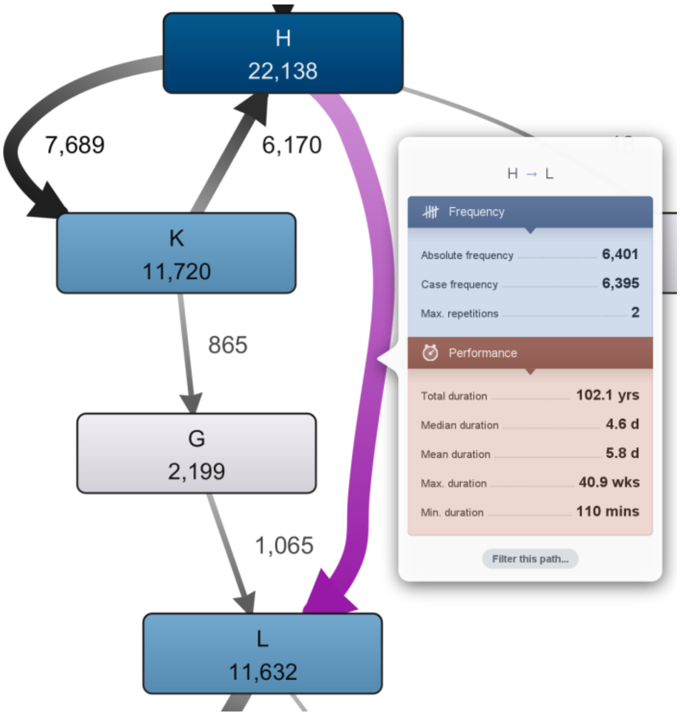

Figure 2: Queuing time activity L

This approach yields the following results:

Activity	A	B	C	D	E*
Median activity time (mins)	20	10	5	10	20
Queue time 1 (hrs)	0	0	98	14	0
Queue time 2 (hrs)	1	0	0	14	624
Queue time 3 (hrs)	144	504	3	278	134
Total queue frequency	1347	2200	1119	9715	9462
Time reduction (%)	**0.35%**	**4.00%**	**12.16%**	**3.23%**	**13.63%**

Table 5: Time Sensitive Metrics Activities A-E

Activity	H	I*	K	N*	O*
Median activity time (mins)	10	10	10	5	10
Queue time 1 (hrs)	46	0	0	24	0
Queue time 2 (hrs)	0	396	0	5	19
Queue time 3 (hrs)	13	62	12	3	7
Total queue frequency	19595	2305	10687	11225	11614
Time reduction (%)	**2.89%**	**5.08%**	**0.08%**	**3.67%**	**0.15%**

Table 6: Time Sensitive Metrics Activities H-O
* metrics with 24h increments in time measurements, thus 12 hour error margin

FTE's saved

Here the estimated activity time is used again. The framework assumes the 'total activity duration' can be found in 'Disco', but this is displayed as 0 milliseconds due to the limited data. Instead, this is calculated by taking the estimated activity time and multiplying it by the amount of times an activity is executed (absolute frequency). This yields the following results:

Activity	A	B	C	D	E
Median activity time (mins)	20	10	5	10	20
Frequency	1364	2393	1232	11847	11815
Hours saved (hours)	455	399	103	1975	3938
FTE's saved (FTE's)	**0.27**	**0.24**	**0.06**	**1.19**	**2.38**

Table 7: FTE's Saved Metrics Activities A-E

Activity	H	I	K	N	O
Median activity time (mins)	10	10	10	5	10
Frequency	22168	2430	11720	11632	11631
Hours saved (hours)	3695	405	1953	969	1939
FTE's saved (FTE's)	**2.23**	**0.24**	**1.18**	**0.59**	**1.17**

Table 8: FTE's Saved Metrics Activities H-O

Step 6 – Prioritize RPA opportunities

With regards to human error prone activities, there is a very strong binary divide between them. *B* and *H*, with 48% and 88% error rate respectively, score extremely high. This is much higher than an expected 10% human error rate, and implies that there may be another cause for repetitions occurring. All other activities score around 1% or lower, implying that human error occurs very infrequently. Activities *H* and then *B* may thus warrant further inspection as an RPA opportunity, whereas the other scores are negligible.

The other three metrics are much less binary in division and displayed next to each other in table 9. Microsoft Excel's "conditional formatting" has been applied in each column, where scores are highlighted from dark green to yellow depending on their relative score, with dark green implying a high score and yellow a low score. The higher a metric, the more value implementing RPA could have.

Activity	Frequency	Time reduction	FTE's saved
A	1364	0.35%	0.27
B	2393	4.00%	0.24
C	1232	12.16%	0.06
D	11847	3.23%	1.19
E	11815	13.63%	2.38
H	22168	2.31%	2.23
I	2430	5.08%	0.24
K	11720	0.08%	1.18
N	11632	3.67%	0.59
O	11631	0.15%	1.17

Table 9: Metrics per activity in conditional formatting

The conditional formatting allows for easy visual comparison between activities for prioritization. Other more involved approaches could also be taken, such as converting each metric into z-scores via standard deviations and then tallying the total z-scores per activity. Although a more formal approach, the trade-off between extra time taken and added value was not deemed worth it. The precise order is not so important, only that the higher potential value activities are somewhere in the top.

Based on these results, the order of activity inspection was as follows:

Rank	Activity
1	E
2	H
3	C
4	N
5	D
6	K
7	O
8	I
9	B
10	A

Table 10: Prioritization of activities to inspect as RPA opportunity

Step 7 – Assess technical complexity

Ideally, the approach here would be to, for each activity, sit down with someone from the RPA team and an employee that carries out the activity. This would allow for closer inspection of how the activity is carried out and the accompanying complexities for implementing RPA. This takes time from many employees at the institution, as well as some lead time in which the meetings would need to be planned. Multiple employees indicated to simply not have the time in the near future for these meetings. As a result, another approach needed to be taken.

The main purpose behind these meetings would be to gain insight into the cost of implementing RPA for each activity. Some basic understanding of each activity was already gained during the data preparation phase, and RPA thought at least (partially) possible. A meeting was held with an employee from the institution's RPA team to gain more understanding of the costs.

Essentially, at the institution a simple RPA project takes around 1 week, and more complex projects 2 to 3 weeks. Maintenance is estimated at around 10% of implementation time on a yearly basis. Translating this into hours based on a 36-hour work week, a simple project takes around 36 hours to implement plus 4 hours to maintain for one year. A total of 40 hours. Assuming a 'worst' case scenario of 4 weeks, this translates into 160 hours to implement RPA for an activity. Other costs, such as licensing, do not increase as these are fixed costs already paid.

One can simply compare this to the hours saved calculated as part of the FTE's saved indicator to gain at least partial insight into whether RPA could be worth it based on employees being freed up to work on other tasks or simply to save costs.

An oversight of the estimated hours saved per activity is given in table 11.

Activity	Hours saved (hrs)
A	455
B	399
C	103
D	1975
E	3938
H	3695
I	405
K	1953
N	969
O	1939

Table 11: Estimated hours saved per activity

With the 160 hours 'cost' of a complex RPA project, it becomes clear that automation for each possible activity would be net beneficial, as they save significantly more hours than it would cost to automate. It is recommended to start with the highest hours saved activities and work from there. This assumes that each activity can be automated using RPA, but it might quickly come to light that there are more complicated factors playing a role that would make it extremely difficult or impossible. This has to be assessed per activity by having someone from the RPA team meet with an employee carrying out the process.

Step 8 – Implement RPA

This would be the next step to take but is considered outside of the scope for this research.

7. Results of applying the framework

7.1 Results and verification

After applying the framework, it became clear that there are ten activities suited for RPA and may yield benefits for the institution. These activities and their potential benefits are displayed in table 12 below.

Activity	Time reduction	FTE's saved (FTE)
A	0.35%	0.27
B	4.00%	0.24
C	12.16%	0.06
D	3.23%	1.19
E	13.63%	2.38
H	2.31%	2.23
I	5.08%	0.24
K	0.08%	1.18
N	3.67%	0.59
O	0.15%	1.17

Table 12: Potential RPA benefits per activity

This provides a strong guide for where to start trying to implement RPA within the process.

Thus, out of 16 initial activities, 10 were identified as RPA candidates. Ideally, it would be known for each found candidate whether they are also RPA candidates in reality, which would allow evaluation of the framework's accuracy. However, this is currently unknown and other methods to verify and evaluate the results must be found. A control group would provide significant value here, where the process is analyzed in the traditional manner and the results could be compared. This is a time-consuming process and could not be done within the timeframe of this research.

Instead, a more crude but effective approach was taken, where the end-to-end process was discussed with an RPA expert and a process expert to get an overview of the activities and identify RPA candidates among them. This took about 45 minutes and yielded a list of 6 activities. This is similar to the current way of finding RPA opportunities used at the financial institution, except the focus was on an end-to-end process. Normally, smaller process chunks would be discussed as a general overview of the process is not present.

The 6 identified opportunities were A, B, D, E, H, and I. This overlaps with the activities discovered through the framework, meaning no activities were missed, and the framework even discovers more potential opportunities. These candidates were not 'found' via the traditional method as the activities seemed too complex and the reward unknown. Low technical complexity was the driving factor for identifying tasks, as the added value was very difficult to estimate without analytical tools such as process mining.

7.2 Further limitations of framework

During the application phase, some more limitations of the framework came to the surface. Specifically, these have to do with the strength of the indicator measurements.

The first weakness found was for the human error indicator, which was measured via the relative frequency of repeated activities. The difficulty here is that there are simply too many other reasons an activity may be repeated. There may be changes later in the process which require a quick update on a form, there may sub-stages of an activity that can be completed independently but registered as the same activity, or an employee starts an activity, stops to do something else, and finishes it at a later stage. This makes it nigh impossible to determine which amount of the repetitions are due to human error. This is something that could better be evaluated via inspection of the activity when also assessing technical complexity.

A second weakness was found in the time-saved indicator, where it is calculated that RPA removes all the waiting time for an activity to occur. The underlying assumption is that the sole reason an activity is not being carried out, is that the employee is preoccupied with another task. However, there may be other reasons as well; the employee could be waiting for additional required information, a decision by management, or there could be a minimal legal amount of time required to pass before continuing with the process. How much queueing time in the model is truly waiting time depends on the process itself and may be accounted for in calculations based on process information.

8. Discussion

As can be concluded from the results, the approach using the process mining framework adds value to the discovery of RPA opportunities over the traditional method. After comparing the output of the framework and the traditional approach, and discussing these with an RPA expert, several of the hypothesized advantages were confirmed. First and foremost, process mining gives a complete overview of the end-to-end process, which is generally lacking with the traditional approach. Just having this overview is already immensely useful according to the RPA expert.

Besides this overview, a reliable measurement of the added value of automating an activity is given through the framework. This has two advantages; it allows for prioritization of candidates, and it provides a strong basis for a business case for RPA. It can be difficult to convince stakeholders of the use of RPA; being able to show that a certain amount of hours can be saved or the process sped up by a certain amount of time makes this much easier.

Third, by providing strong insights and measurements into the process before RPA, process mining not only allows to ensure a 'good' process is being automated, but it also enables the process owner to accurately compare the before and after situation.

Although there are some limitations to the framework, its reliability and accuracy still seem satisfactory. The framework's practicality allows it to be used effectively even by inexperienced process miners, as the method's wholistic approach accounts for pitfalls such as concept drift and low process quality. This was verified by having a new intern at the financial institution read this report after completing a short process mining course. The intern could answer questions on the approach to a satisfactory degree and felt confident in being able to apply it themself with a little practice.

This does not mean no further validation is required. The framework has only been applied in a single case and takes input from a single RPA expert. The results of implementing RPA for the found activities are not available yet, and thus it is difficult to measure the added value and effectiveness of the framework. The indicators of the framework also include some estimates, such as the speed of an RPA bot compared to a human worker, and what proportion of an activity will be an exception to be handled by a human worker. Gathering more data on the effectiveness of RPA solutions and applying the framework to more cases are two areas for further research.

There are two drawbacks to using process mining for discovering RPA opportunities; it requires process mining to be implemented within the organization, and it is more time-consuming than the traditional approach due to the time needed for obtaining and preparing the data. This is on top of the general reliance on the availability of high quality data, as discussed in chapters 3.2 and 6.1.

8.1 Process mining solely for RPA opportunities

Purchasing and implementing process mining technology within an organization is a costly change. Doing this only for finding RPA opportunities is not expected to be worth it, unless RPA plays an extremely large part in the organization.

If process mining is already implemented within an organization, using it to discover RPA opportunities can be worth it. It can take several days to properly acquire, prepare, and analyze the data. Thus, one should be careful in selecting the process to mine and increase the chances that multiple candidate tasks are found. If it is already known that there is already a high level of automation, many activities that can not be automated, or any other data that implies low value of RPA, a traditional approach may be more cost effective.

8.2 General process mining and RPA opportunities

There are many more benefits to process mining beyond its synergies with Robotic Process Automation. If these benefits are of enough value to an organization, process mining should be implemented. Here, its synergies with RPA can provide extra value on top of the regular process mining benefits.

Acquiring data and preparing it is required for any process mining task, and only needs to be done once. Multiple analyses can then be made with the same dataset, each focusing on different aspects. This is extremely beneficial to do, as it increases the efficiency of process mining. This is also the case when looking for RPA opportunities. Process mining just for finding candidate tasks may not always be worth it, but if a process is mined for any other purpose, the framework designed in this report can also be applied with relative ease. Essentially, organizations should aim to capitalize on as many benefits of process mining within the same dataset as possible. If a dataset is readily available for analysis for process mining, it will generally be worth to do so for RPA opportunities.

References

Anagnoste, S. (2017). Robotic Automation Process - The next major revolution in terms of back office operations improvement. *PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON BUSINESS EXCELLENCE*, *11*(1), 676–686. https://doi.org/10.1515/picbe-2017-0072

Anagnoste, S. (2018). Robotic Automation Process - The operating system for the digital enterprise. *PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON BUSINESS EXCELLENCE*, *12*(1), 54–69. https://doi.org/10.2478/picbe-2018-0007

Asatiani, A., & Penttinen, E. (2016). Turning robotic process automation into commercial success - Case OpusCapita. *Journal of Information Technology Teaching Cases*, *6*(2), 67–74. https://doi.org/10.1057/jittc.2016.5

Bellam, S. (2018). Robotics vs Machine Learning vs Artificial Intelligence: Identifying the Right Tools for the Right Problems. *Credit & Financial Management Review*, *24*(2), 1–10. http://ezproxy2.utwente.nl/login?url=https://search.ebscohost.com/login.aspx?direct=true&db=bsh&AN=131055941&site=ehost-live

Bose, R. P. J. C., Van Der Aalst, W. M. P., Zliobaite, I., & Pechenizkiy, M. (2014). Dealing with concept drifts in process mining. *IEEE Transactions on Neural Networks and Learning Systems*, *25*(1), 154–171. https://doi.org/10.1109/TNNLS.2013.2278313

Caetano, A. (2010). Business Process Decomposition. *Enterprise Modelling and Information Systems Architectures*, *X*(X), 1–11.

Cooper, L. A., Holderness, D. K., Sorensen, T. L., & Wood, D. A. (2019). Robotic process automation in public accounting. *Accounting Horizons*, *33*(4), 15–35. https://doi.org/10.2308/acch-52466

Davenport, T. H., & Spanyi, A. (2019). What Process Mining Is, and Why Companies Should Do It. *Harvard Business Review Digital Articles*, 2–6.

Enriquez, J. G., Jimenez-Ramirez, A., Dominguez-Mayo, F. J., & Garcia-Garcia, J. A. (2020). Robotic Process Automation: A Scientific and Industrial Systematic Mapping Study. *IEEE Access*, *8*, 39113–39129. https://doi.org/10.1109/ACCESS.2020.2974934

Gao, J., van Zelst, S. J., Lu, X., & van der Aalst, W. M. P. (2019). Automated robotic process automation: A self-learning approach. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics): Vol. 11877 LNCS*. https://doi.org/10.1007/978-3-030-33246-4_6

Garcia, C. D. S., Meincheim, A., Faria Junior, E. R., Dallagassa, M. R., Sato, D. M. V., Carvalho, D. R., Santos, E. A. P., & Scalabrin, E. E. (2019). Process mining techniques and applications – A systematic mapping study. *Expert Systems with Applications*, *133*, 260–295. https://doi.org/10.1016/j.eswa.2019.05.003

Gartner. (2020, October). *Top Strategic Technology Trends for 2021.* Gartner.

Geyer-Klingeberg, J., Nakladal, J., Baldauf, F., & Veit, F. (2018). Process mining and Robotic process automation: A perfect match. *CEUR Workshop Proceedings*, *2196*, 124–131.

Gielen, R. (2019). *RPA: faster, cheaper and more efficient*. Cegeka. https://www.cegeka.com/en/be/blog/rpa-faster-cheaper-efficient

Hatfield, J. (2020, June 1). *How to Speed Your RPA Implementation 20x with UiPath Automation Hub.* UiPath. https://www.uipath.com/blog/accelerate-rpa-implementation-uipath-automation-hub

Hawkins, I. (2018). *A guide to Robotic Process Automation (RPA)*. Process Excellence Network. https://www.processexcellencenetwork.com/rpa-artificial-intelligence/articles/a-guide-to-robotic-process-automation-rpa#:~:text=Robotic Process Automation is a,applications and uses keyboard shortcuts.

Keet C.M. (2013) Granularity. In: Dubitzky W., Wolkenhauer O., Cho KH., Yokota H. (eds) Encyclopedia of Systems Biology. Springer, New York, NY. https://doi.org/10.1007/978-1-4419-9863-7_65

Leno, V., Dumas, M., Maggi, F. M., & La Rosa, M. (2018). Multi-Perspective process model discovery for robotic process automation. *CEUR Workshop Proceedings*, *2114*, 37–45.

Leshob, A., Bourgouin, A., & Renard, L. (2018). Towards a Process Analysis Approach to Adopt Robotic Process Automation. *Proceedings - 2018 IEEE 15th International Conference on e-Business Engineering, ICEBE 2018*, 46–53. https://doi.org/10.1109/ICEBE.2018.00018

Leopold, H., van der Aa, H., & Reijers, H. A. (2018). Identifying candidate tasks for robotic process automation in textual process descriptions. In *Lecture Notes in Business Information Processing* (Vol. 318). https://doi.org/10.1007/978-3-319-91704-7_5

Magrabi, F., Li, S. Y. W., Day, R. O., & Coiera, E. (2010). Errors and electronic prescribing: A controlled laboratory study to examine task complexity and interruption effects. *Journal of the American Medical Informatics Association*, *17*(5), 575–583. https://doi.org/10.1136/jamia.2009.001719

Modi, A., Harpreet, K., Utkarsh, S. (2020). Process Mining – Technology Vendor Landscape with Products PEAK Matrix® Assessment 2020. *Everest Group PEAK Matrix Assessment.* https://www2.everestgrp.com/reportaction/EGR-2020-38-R-3576/Marketing

Naga Lakshmi, M. V. N., Vijayakumar, T., & Sai Sricharan, Y. V. N. (2019). Robotic process automation, an enabler for shared services transformation. *International Journal of Innovative Technology and Exploring Engineering*, *8*(6), 1882–1890.

Penttinen, E., Kasslin, H., & Asatiani, A. (2018). How to choose between robotic process automation and back-end system automation? *26th European Conference on Information Systems: Beyond Digitization - Facets of Socio-Technical Change, ECIS 2018*.

Radke, A. M., Dang, M. T., & Tan, W. K. A. (2020). Using robotic process automation (RPA) to enhance item master data maintenance process. *Logforum*, *16*(1), 129–140. https://doi.org/10.17270/J.LOG.2020.380

R'Bigui, H., & Cho, C. (2017). The state-of-the-art of business process mining challenges. *International Journal of Business Process Integration and Management*, *8*(4), 285–303. https://doi.org/10.1504/IJBPIM.2017.10009731

Santos, F., Pereira, R., & Vasconcelos, J. B. (2019). Toward robotic process automation implementation: an end-to-end perspective. *Business Process Management Journal*, *26*(2), 405–420. https://doi.org/10.1108/BPMJ-12-2018-0380

Suri, V. K., Elia, M., & van Hillegersberg, J. (2017). Software bots -The next frontier for shared services and functional excellence. In *Lecture Notes in Business Information Processing* (Vol. 306). https://doi.org/10.1007/978-3-319-70305-3_5

Sobczak, A. (2019). DEVELOPING A ROBOTIC PROCESS AUTOMATION MANAGEMENT MODEL. *BUDOWA MODELU ZARZĄDZANIA ROBOTYZACJĄ PROCESÓW BIZNESOWYCH.*, *2*(52), 85–100. http://10.0.60.251/ie.2019.2.06

Van Der Aalst, W. (2012). Process mining: Overview and opportunities. *ACM Transactions on Management Information Systems*, *3*(2). https://doi.org/10.1145/2229156.2229157

Vanderfeesten, I., Reijers, H. A., & van der Aalst, W. M. P. (2008). Evaluating workflow process designs using cohesion and coupling metrics. *Computers in Industry*, *59*(5), 420–437. https://doi.org/10.1016/j.compind.2007.12.007

Wanner, J., Hofmann, A., Fischer, M., Imgrund, F., Janiesch, C., & Geyer-Klingeberg, J. (2020). Process selection in RPA projects - Towards a quantifiable method of decision making. *40th International Conference on Information Systems, ICIS 2019*.

Zaharia-Radulescu, A.-M., Pricop, C. L., Shuleski, D., & Ioan, A. C. (2017). RPA AND THE FUTURE OF WORKFORCE. In Popa, I and Dobrin, C and Ciocoiu, CN (Ed.), *PROCEEDINGS OF THE 11TH INTERNATIONAL MANAGEMENT CONFERENCE: THE ROLE OF MANAGEMENT IN THE ECONOMIC PARADIGM OF THE XXIST CENTURY (IMC 2017)* (pp. 384–392).

Appendix

Appendix A – Systematic Literature Research

A.1 Research question and definitions

Research question: What are the value drivers for a financial institution to implement RPA?

Value drivers – Factors that increase the added value of an activity, process, product or service.

Financial institution – An organization that offers financial services, such as banks, insurance companies and brokerage firms.

Robotic Process Automation – A tool based on written scripts to execute specific keystrokes on the UI of applications, to execute activities and processes automatically that are normally executed by human employees.

A.2 Inclusion and exclusion criteria

Inclusion criteria:

Nr	Criteria	Reason for inclusion
1	“Robotic Process Automation” in abstract	Value drivers must be related to RPA
2		
3		
4		

Exclusion criteria:

Nr	Criteria	Reason for exclusion
1	Pre 2010 articles	RPA was not developed before this time
2	Product based business	Interested in service based value
3	RPA2.0	Research scope is RPA, which is ready to be implemented. RPA2.0 is not.
4	“RPA” or “Robotic Process Automation” not in title	Looking for articles very specifically about RPA.
5	RPA in education	Focus on financial services
6	RPA in healthcare	Focus on financial services
7	Impact assessment	Not directly relevant to the value drivers of RPA

A.3 Databases

Databases to use will be SCOPUS, Web of Science and Business Source Elite. Google Scholar may be used as a fall back in case of insufficient results or to gain access to articles found through the other databases.

A.4 Search terms and strategy

Search Term	Synonyms
“RPA”	“Robotic Process Automation”, “Software robotics”

“Value driver”	“Add* value”, “important”, “competitive advantage”, “grow*”, “efficien*”, “risk”
“Financial institution”	“Bank”, “fintech”, “investment firm”, “broker”, “insur*”

At first, a search will be started focusing on just RPA. RPA is in nature tailor made for businesses, and articles exploring the technology will most likely already focus on the value it can add. If there are too many results and I need to specify further, I will add “value driver” or synonyms to the search string. Only if there are still too many results will “financial institution” or synonyms also be added.

A.5 Search log

Platform	Search strings	Hits	Date search	Date range	Scope
SCOPUS	RPA	11,092	12-Aug-20	any	Title, keywords, abstract
SCOPUS	"RPA" and not "DNA"	8,821	12-Aug-20	any	Title, keywords, abstract
SCOPUS	Robotic Process Automation	6,095	12-Aug-20	any	Title, keywords, abstract
SCOPUS	"Robotic Process Automation"	159	12-Aug-20	any	Title, keywords, abstract
SCOPUS	"Robotic Process Automation" AND value	31	12-Aug-20	any	Title, keywords, abstract
SCOPUS	"Robotic Process Automation" AND efficien*	34	12-Aug-20	any	Title, keywords, abstract
SCOPUS	"Robotic Process Automation" AND "competitive advantage"	1	12-Aug-20	any	Title, keywords, abstract
Web of Science	"Robotic Process Automation"	71	12-Aug-20	any	Title, keywords, abstract
Web of Science	"Robotic Process Automation" AND value	10	12-Aug-20	any	Title, keywords, abstract
Web of Science	"Robotic Process Automation" AND efficien*	12	12-Aug-20	any	Title, keywords, abstract
Web of Science	"Robotic Process Automation" AND "competitive advantage"	1	12-Aug-20	any	Title, keywords, abstract
Business Source Elite	"Robotic Process Automation"	353	12-Aug-20	Any	
Business Source Elite	"Robotic Process Automation"	59	12-Aug-20	Any	Scholarly journals
Business Source Elite	"Robotic Process Automation" AND value	15	12-Aug-20	Any	Scholarly journals
Business Source Elite	"Robotic Process Automation" AND efficien*	19	12-Aug-20	any	Scholarly journals
Business Source Elite	"Robotic Process Automation" AND "competitive advantage"	0	12-Aug-20	any	Scholarly journals

Figure 3: Search log

The results from “Robotic Process Automation” searches from SCOPUS, Web of Science and Business Source Elite (scholarly journals only) were compiled in Mendeley. 52 duplicates were removed, ending with 187 articles. An inspection of 25 abstracts yielded that many articles not containing “Robotic Process Automation” or “RPA” in the title were too general to be of use for this study, so this was an applied exclusion criterium. This left 96 articles.

This was then further narrowed down to articles containing “value”, “efficient*”, “benefit” or “competition?*” in the abstract. This left 42 articles. The abstracts of these 42 articles were read, yielding 16 articles related to value drivers of RPA. 6 of these articles were behind a paywall and could not be accessed, even after several attempts to circumvent the paywall. This left 10 articles to read, see figure 3.

	Authors	Title	Year
1	Cooper, L.A.; Holderness, D.K.; Sorensen, T.L.; Wood, ...	Robotic process automation in public accounting	2019
2	Zaharia-Radulescu, Adrian-Mihai; Pricop, Calalin Liviu; S...	RPA AND THE FUTURE OF WORKFORCE	2017
3	Wanner, J.; Hofmann, A.; Fischer, M.; Imgrund, F.; Ja...	Process selection in RPA projects - Towards a quantifiable method of decision making	2020
4	Leopold, H.; van der Aa, H.; Reijers, H.A.	Identifying candidate tasks for robotic process automation in textual process descriptions	2018
5	Leno, V.; Dumas, M.; Maggi, F.M.; La Rosa, M.	Multi-Perspective process model discovery for robotic process automation	2018
6	Naga Lakshmi, M.V.N.; Vijayakumar, T.; Sai Srichar...	Robotic process automation, an enabler for shared services transformation	2019
7	Geyer-Klingeberg, J.; Nakladal, J.; Baldauf, F.; Veit, F.	Process mining and Robotic process automation: A perfect match	2018
8	Santos, F.; Pereira, R.; Vasconcelos, J.B.	Toward robotic process automation implementation: an end-to-end perspective	2019
9	Asatiani, A.; Penttinen, E.	Turning robotic process automation into commercial success - Case OpusCapita	2016
10	Gao, J.; van Zelst, S.J.; Lu, X.; van der Aalst, W.M.P.	Automated robotic process automation: A self-learning approach	2019

Figure 4: Final Articles

A.6 Concept matrix

	IT Infrastructure	Cost-saving	Quality (errors, time)	Performance increasing	Scaling challenges
1					
2					
3					
4					
5					
6					
7					
8					
9					
10					

Appendix B – Overview of Disco Functionalities

Model view

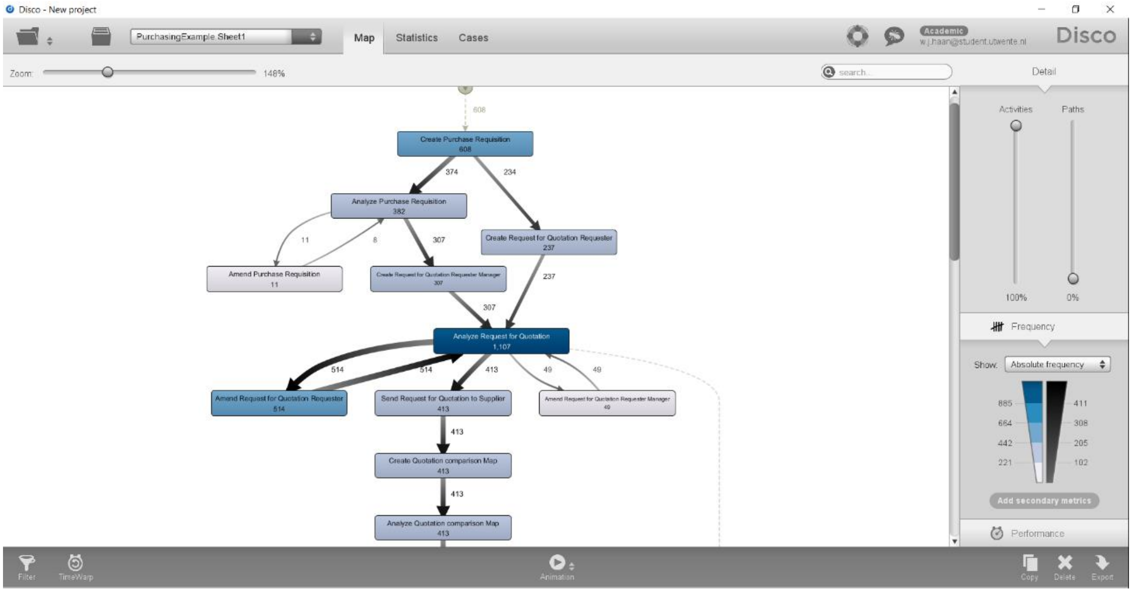

Figure 5: model view

The model view gives a visual overview of the process. Pathways and activities are highlighted based on their values. The values that can be shown are shown in the table below.

Frequency	
	Absolute frequency
	Case frequency
	Max repetitions
Performance	
	Total duration
	Median duration
	Mean duration
	Maximum duration
	Minimum duration

An overview of these measurements is also given for each activity when clicked on, as can be seen in the figure below. *Start frequency* or *end frequency* may be an added measurement for this action, which shows how often this activity is the starting point and ending point respectively for the process.

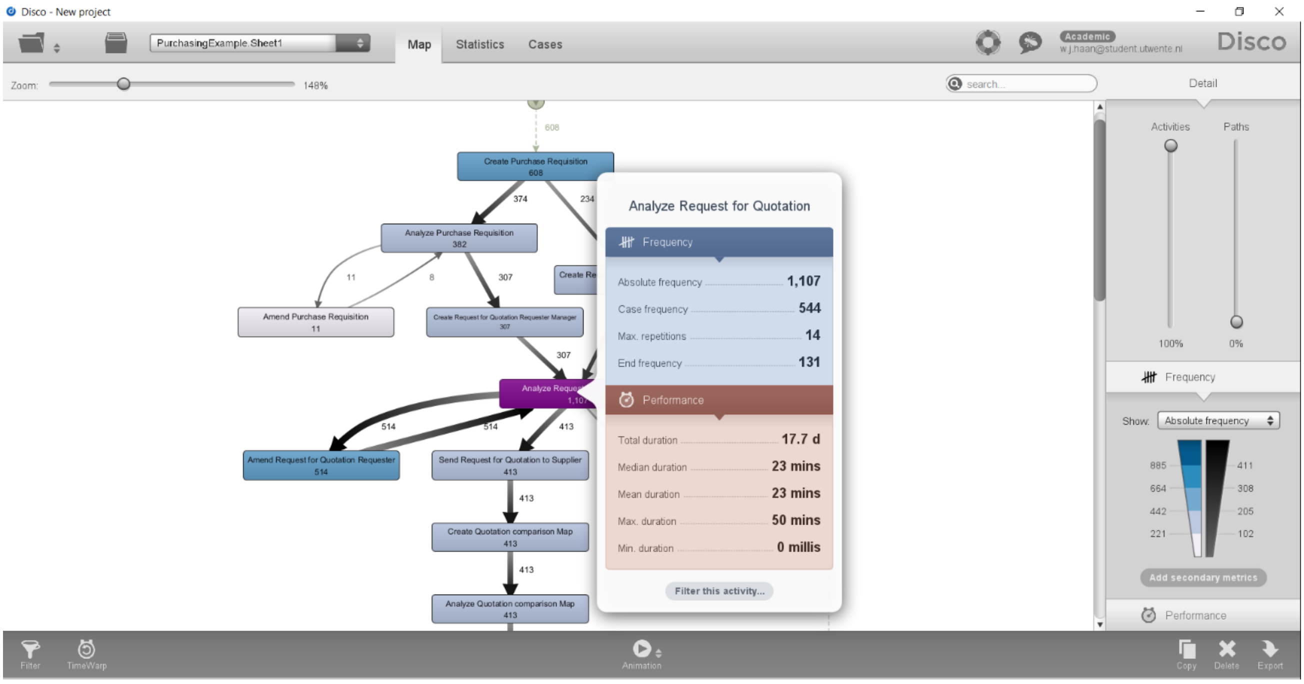

Figure 6: Activity overview

Statistics View

The statistics global view offers a wealth of information regarding the process, which can be beneficial to better understand how much RPA might improve the process. For example, if it is calculated that implementing RPA for a certain activity reduces its throughput time from 2 days to 1 day, it is beneficial to know if the total process takes 3 days or 50 days. This way process improvement is also made measurable.

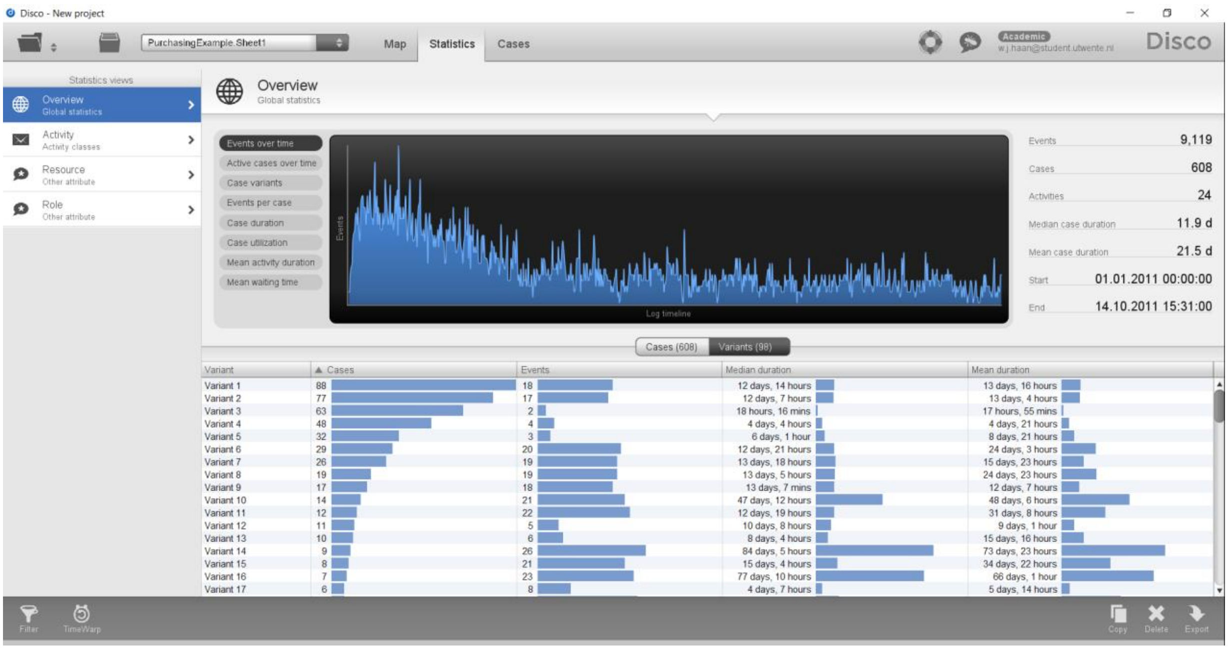

Figure 7: Statistics global view

There is also a statistics activity overview, which yields information on the activities in relation to each other. This can be of benefit when there are many activities to inspect, as it allows for easy discovery of highest scoring activities for example for "aggregate duration" which could translate into *total FTE's*.

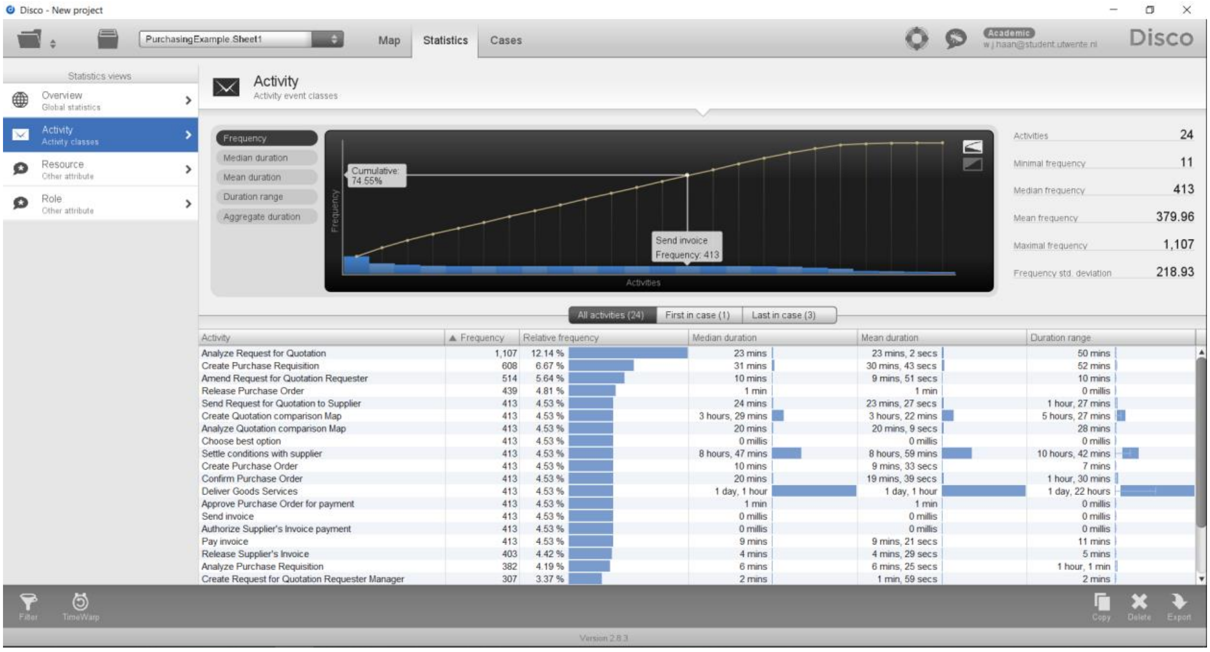

Figure 8: Statistics activity view

Cases view

There is also the cases view, which displays variants and cases. A variant is a group of cases that are all executed in the same order. This information is not directly relevant for RPA opportunities, however can be hugely important for more general automation types. If there are few variants, are a few which make up most of the cases, the process may be suitable for more traditional types of automation.

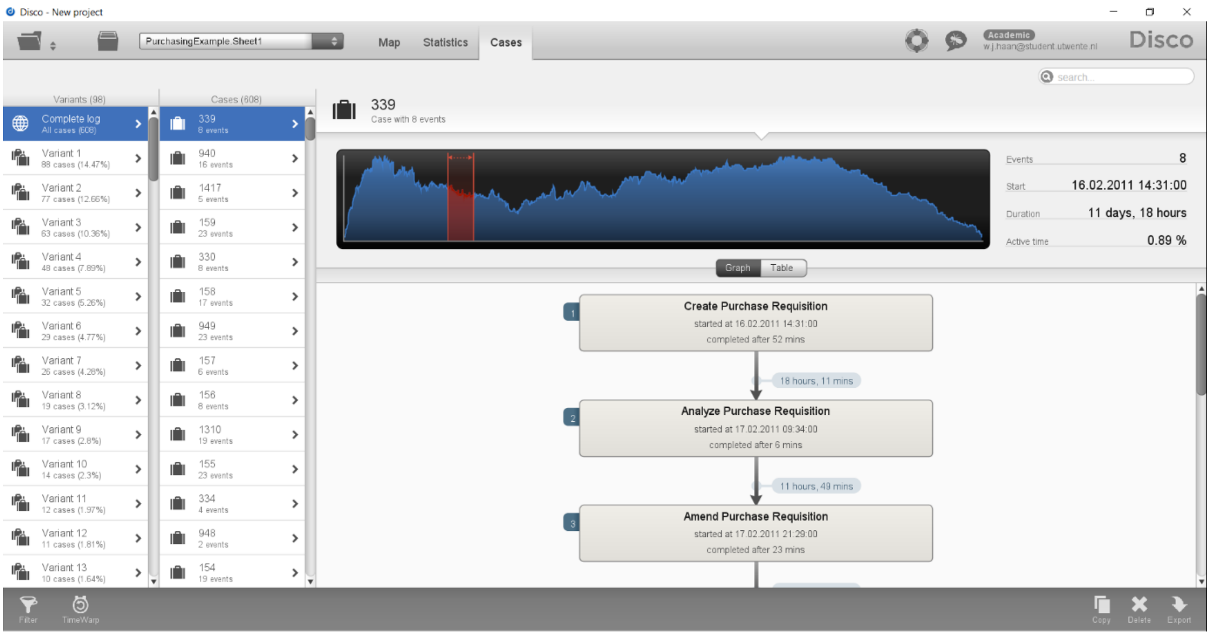

Figure 9: Cases view

Filters

Finally there is the functionality of filters. Filters are ways of reducing the data you are exploring according to certain rules. Essential in process analysis, as it allows for the inspection of incidents of interest to the analyst; for example, a product arriving before even being ordered. Although very powerful and a core aspect of process mining, these will not be elaborated on unless used in the analysis.

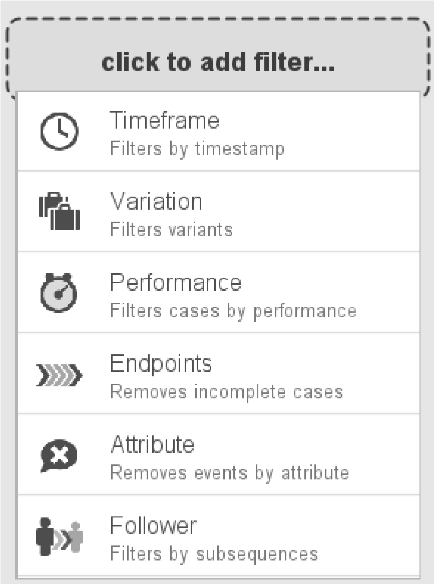

Figure 10: Filters

Appendix C – Weighted Average Queuing Times Investigation

It is investigated whether taking the weighted average of the two- or three- most frequent pathways, as opposed to all pathways into an activity, yields a 'close enough' measure to justify the practicality of the approach. Results of the faster approach must lie within 25% of the true value to be satisfactory.

A publicly available dataset on a refund process is used as data. Each activity investigated must contain at least five incoming pathways. The seven most frequent activities meeting this requirement are investigated. The summarizing results can be found in the table below, the raw data in the section 'Raw Data' further below. Based on the summarizing results, using the 3 most frequent pathways yields results close enough to the complete average.

Activity	Nr pathways	Avg 2 most frequent	Avg 3 most frequent	Avg all	Percent Dif 2	Percent Dif 3
1	9	44.34	39.62	36.07	22.9%	9.8%
2	5	13.52	21.57	25.36	-46.7%	-14.9%
3	7	1.47	3.01	2.92	-49.7%	2.9%
4	6	27.64	25.78	29.67	-6.8%	-13.1%
5	8	161.40	129.00	128.66	25.4%	0.3%
6	9	64.76	64.90	65.41	-1.0%	-0.8%
7	7	11.99	14.11	18.58	-35.5%	-24.1%

Tabel 1: Summarizing Results of 7 activities in "Refund Process"

Raw Data

Request Signed L1 (1)		Request rejected L1 (2)		Order Sent back (3)	
Frequency	Median (hr)	Frequency	Median (hr)	Frequency	Median (hr)
366	69.40	85	24.10	268	1.49
214	1.48	70	0.67	35	1.30
173	23.80	27	67.80	28	19.70
169	0.15	7	127.20	7	0.02
24	192.00	1	2.50	6	2.50
7	0.92			2	0.02
2	21.50				
1	0.38				
1	1.78				

Pick-up prepared (4)		Product received (5)		Missing documents requested (6)	
Frequency	Median (hr)	Frequency	Median (hr)	Frequency	Median (hr)
328	21.60	67	141.60	1	988.80
29	96.00	1	1488.00	1	20.00
26	0.12	334	122.40	5	18.80
5	47.60	10	163.20	1	5.30
3	506.40	1	196.80	1	96.00
1	0.03	2	0.68	9	20.80
		9	91.20	11	70.00
		2	151.20	40	22.30
				366	69.40

Shipment via logistics partner (7)

Frequency	Median (hr)
8	168.00
81	50.60
263	0.10
1	434.40
1	288.00
73	24.10
1	3.60

Appendix D – Process Diagrams

Appendix D.1 - High level model

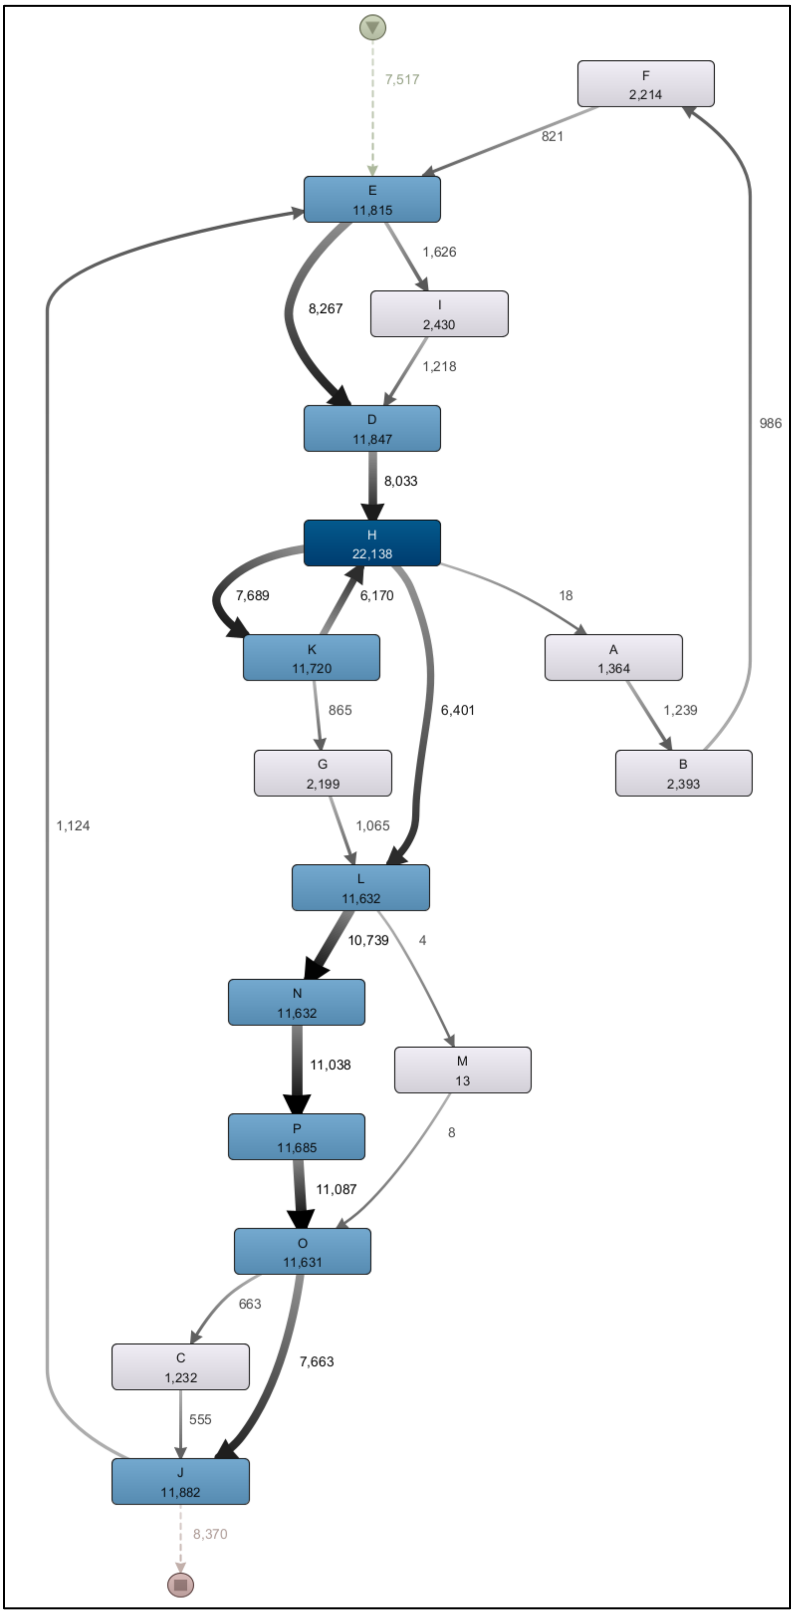

Figure 11: High level model of complete process

Appendix D.2 - Detailed Model

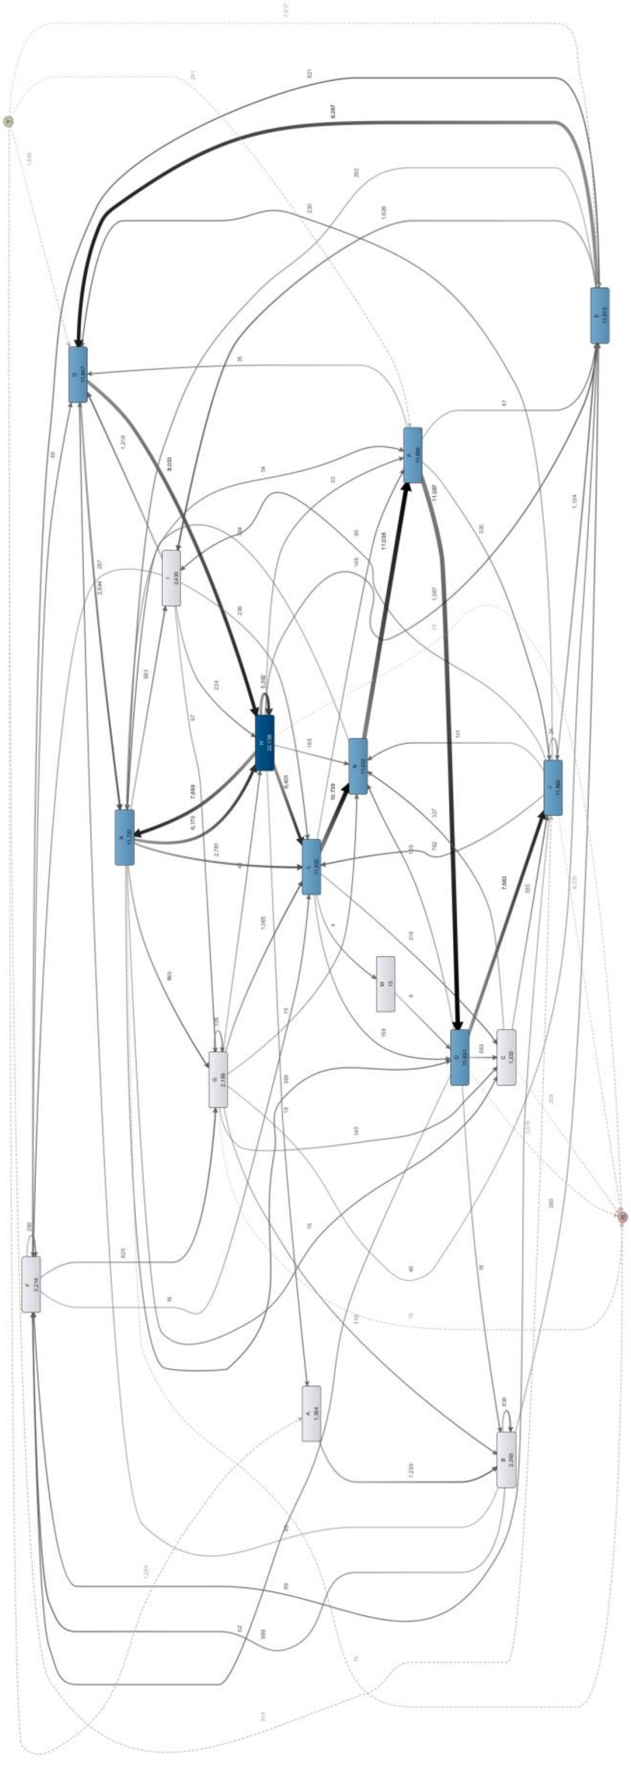

Appendix D.3 - Activity A

Figure 12: Metrics Activity A

Appendix D.4 – Activity B

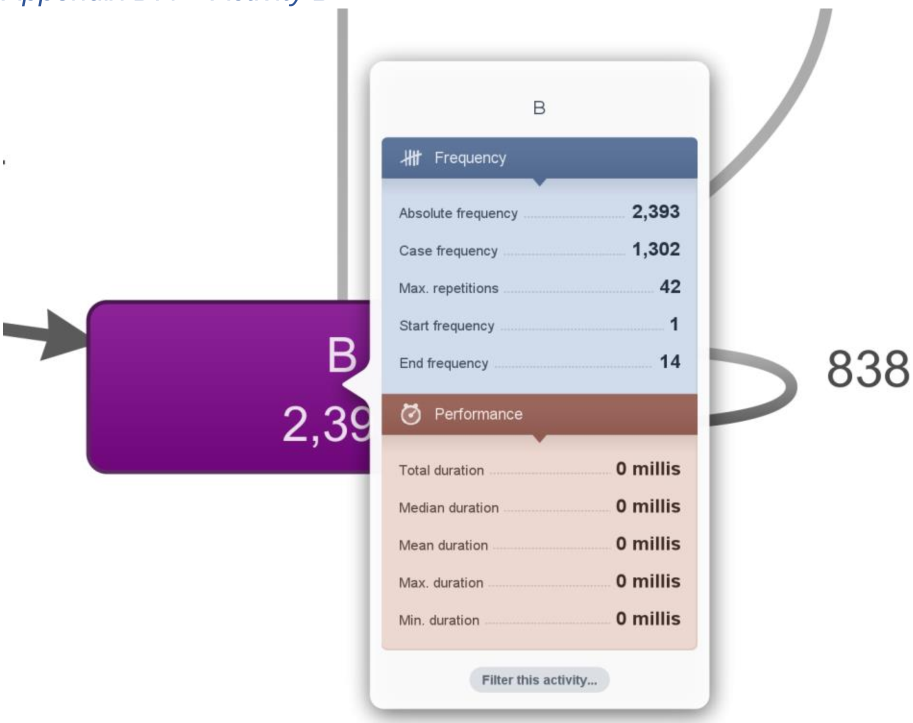

Figure 13: Metrics Activity B

Appendix D.5 – Activity C

Figure 14: Metrics Activity C

Appendix D.6 – Activity D

Figure 15: Metrics Activity D

Appendix D.7 – Activity E

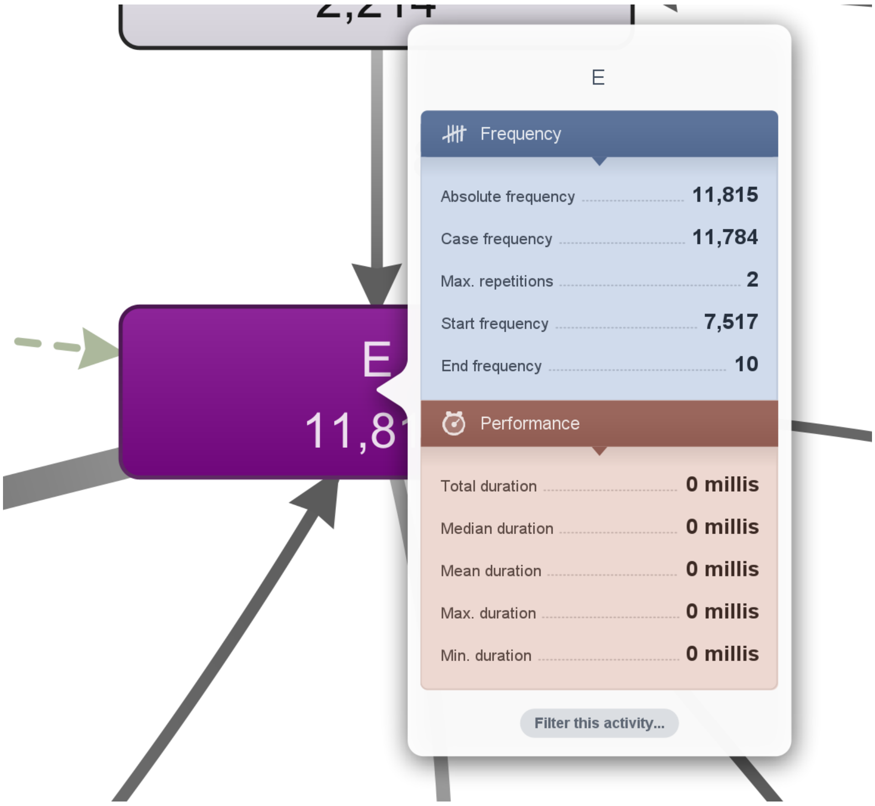

Figure 16: Metrics Activity E

Appendix D.8 – Activity H

Figure 17: Metrics Activity H

Appendix D.9 – Activity I

Figure 18: Metrics Activity I

Appendix D.10 – Activity K

Figure 19: Metrics Activity K

Appendix D.11 – Activity N

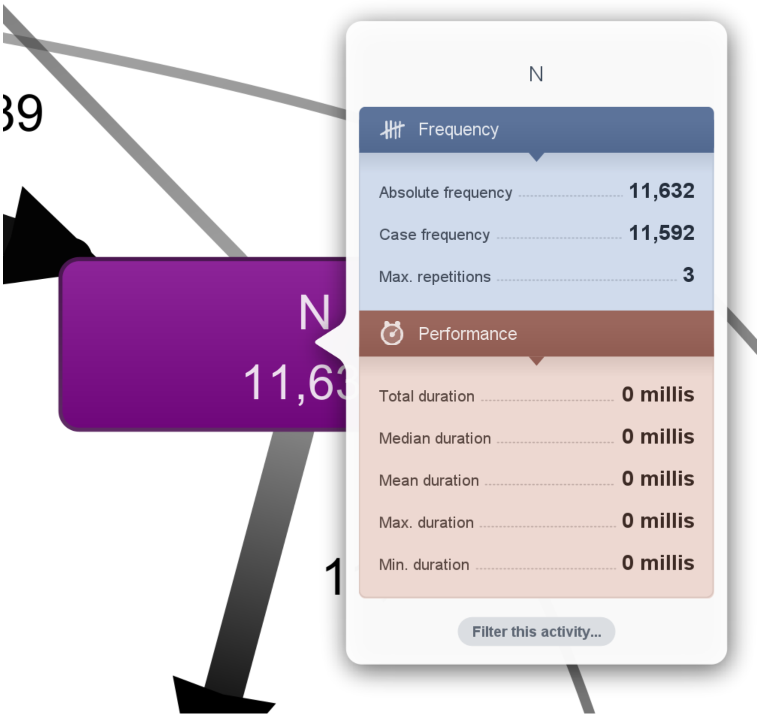

Figure 20: Metrics Activity N

Appendix D.12 – Activity O

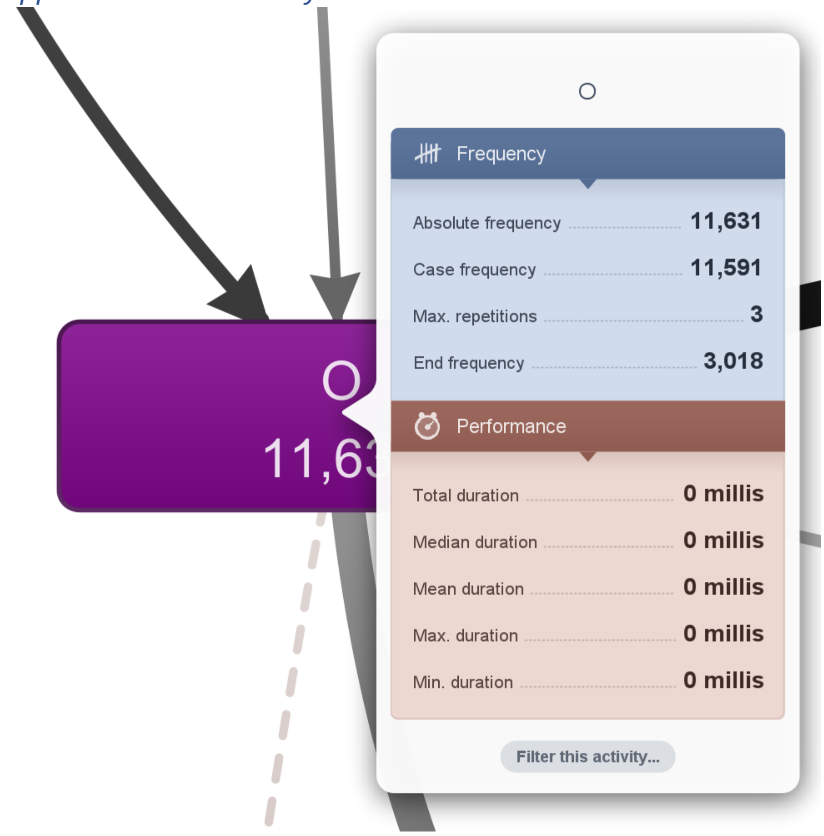

Figure 21: Metrics Activity O

Appendix E – Reflection on Professional Functioning

The project, focused on writing a graduation thesis at an organization, was combined with a regular internship in order to benefit more from the experience of a professional organization. This all happened during the height of the corona-pandemic, meaning both the internship and project were done completely from home. This offered new and interesting challenges to overcome, accompanied by the regular challenges of starting an internship at a new organization.

Starting at a new organization virtually was something I knew in advance would most likely happen due to the governmental restrictions at the time. As a result, I could already prepare a little and ask for some advice from friends. I made an effort to schedule a short call with every colleague to get to know each other a little and build a small network. The team was extremely friendly and welcoming, making it very easy for me to integrate and start asking questions. It took a little time to understand the team dynamics and find the pathways used when getting in touch for different matters, as there were multiple possible approaches such as texting, chat messages, emails and calls. I could observe some from colleagues, and experience taught the rest to get appropriate responses.

The team's purpose included a fair amount of network building, which proved to be a challenge for my to do virtually, as you don't 'run into' people and meetings are very result-oriented. This is where my teammates helped tremendously, taking me with them into meetings and introducing me to others that might be able to help in my projects. This also helped get a taste of the organization's culture and dynamics. Building and maintaining the built network is something that I need to improve on more in the future, and was also part of the feedback I received from my boss.

One of my stronger feedback points was a pro-active approach and readiness to take on unfamiliar tasks. The position focused on innovation, meaning that many obstacles were new to everyone and a solution was not readily available. A pro-active approach was thus valuable to pick up and address potential obstacles to continue progressing the project. Thanks to the strong support from teammates to help point me in the right direction, strong results were booked and I am proud of the work I was able to contribute.

Combining the internship with writing the thesis was another test of my functioning, as it meant multiple objectives to juggle and prioritize. It required me to use my time effectively. Luckily I was also given the required time and freedom by my boss to allow me to work on my thesis in internship 'downtime'. There were a couple of very long weeks required to get everything finished, but I think the two strengthened each other as many insights gained during the internship contributed to the thesis, besides of course carrying out the thesis research at the same organization.

Just like many other students and employees, transitioning to working completely from home took a little time to adjust to and do effectively. A major challenge here was living in a single student room; meaning almost all my life was lived in a single space. I was very used to just relaxing in my room and studying at the university, so in the beginning I would sometimes 'forget' I was working for a short period.

I did some research and rearranged things; creating an area for work, an area for relaxation and an area for sleeping. This allowed me to separate work from personal life, and remained focused while working. I quickly became conditioned to working at my desk, much like studying in the library, not losing focus or ‘automatically’ doing relaxing activities. This really helped my productivity and is something I still strictly maintain.

It did become monotonous after a while, so I spent about a month working from my parent’s house. This helped break up the routine and refreshed me to stay productive and in a healthy mindset.

Overall, I am very happy with my performance both at the internship and with the thesis. I took the time to pursue and implement feedback I received. There are still areas for improvement; my thesis could have been written faster and with a more strict planning. I also could have approached more people to critically asses what I had written. I also still have feedback from my boss to work with, especially with regards to networking and using different emotions in meetings to get desired effects.