ESCUELA TÉCNICA SUPERIOR DE INGENIERÍA (ICAI)

Máster en Big Data: Tecnología y Analítica Avanzada

TRABAJO FIN DE MASTER

Causal Modeling in Manufacturing: Analysis and Applications

Autor: Daniel Bicand Fernández

Director: Cristian Alberch Gracia

Madrid

2023

Declaro, bajo mi responsabilidad, que el Proyecto presentado con el título

Causal Modeling in Manufacturing: Analysis and Applications

en la ETS de Ingeniería - ICAI de la Universidad Pontificia Comillas en el

curso académico 2022/23 es de mi autoría, original e inédito y

no ha sido presentado con anterioridad a otros efectos.

El Proyecto no es plagio de otro, ni total ni parcialmente y la información que ha sido

tomada de otros documentos está debidamente referenciada.

Fdo.: Daniel Bicand Fernández Fecha: 06/ 06/ 2023

Autorizada la entrega del proyecto

EL DIRECTOR DEL PROYECTO

Fdo.: Cristian Alberch Gracia Fecha: 12/ 06/ 2023

Firmado por ALBERCH GRACIA, CRISTIAN (FIRMA) el día 12/06/2023 con un certificado emitido por AC DNIE 005

ESCUELA TÉCNICA SUPERIOR DE INGENIERÍA (ICAI)

Máster en Big Data: Tecnología y Analítica Avanzada

TRABAJO FIN DE MASTER

Causal Modeling in Manufacturing: Analysis and Applications

Autor: Daniel Bicand Fernández

Director: Cristian Alberch Gracia

Madrid

2023

Memory Index

Figure Index

CHAPTER 1. INTRODUCTION

Manufacturing processes play a critical role in the production of goods and services, and their efficiency and effectiveness can have a significant impact on business success. Traditionally, manufacturing projects that had the aim of improving production processes using data have relied on statistical methods to identify correlations and patterns in data. However, the limitations of statistical inference, particularly when it comes to identifying causal relationships, have become increasingly apparent in recent years.

Causal inference can provide a more rigorous framework for understanding the underlying mechanisms that drive manufacturing processes. By accounting for confounding variables and identifying the causal effect of interventions, causal inference can help manufacturing organizations make more informed decisions and improve their operations.

This thesis will begin by providing an overview of causal inference and an explanation on the different current approaches and limitations there are in this field. After explaining the foundations, the thesis will explain the relationships and synergies of applying causal inference in manufacturing. A high-level methodology for applying causal techniques in manufacturing will be explained. The thesis will then present a case study that illustrates the application of causal inference methods to a real-world manufacturing project and will evaluate the effectiveness of these methods in improving process efficiency and product quality.

CHAPTER 2. CAUSAL INFERENCE

2.1 CORRELATION VS CAUSATION:

Correlation and causation are two concepts that are often used interchangeably, but they refer to different types of relationships between variables. Correlation refers to a statistical relationship between two variables, where the occurrence of one variable is related to the occurrence of another variable. Causation, on the other hand, refers to a relationship where one variable directly influences or causes a change in another variable. Whilst causation implies correlation between two variables, the inverse, need not hold true. It is possible that variables are correlated but there is no causal relationship between these.

Machine learning models are designed to identify patterns and relationships in data, and they rely heavily on association to make predictions and generate insights. These models use statistical algorithms to identify correlations between variables, and predict outcomes based on these correlations. However, while machine learning models can be effective at identifying associations between variables, they do not necessarily capture causation. This is because statistical relationships between variables do not always imply causation. For example, a machine learning model might identify a correlation between the number of ice cream sales and the number of shark attacks on humans, but this does not mean that eating ice cream causes sharks attacks to increase. Rather, both variables are likely to be influenced by a third variable, such as temperature or seasonality.

There are three different ways to exemplify the difference between correlation and causation between variables:

1. When variables have coincidentally similar behavior

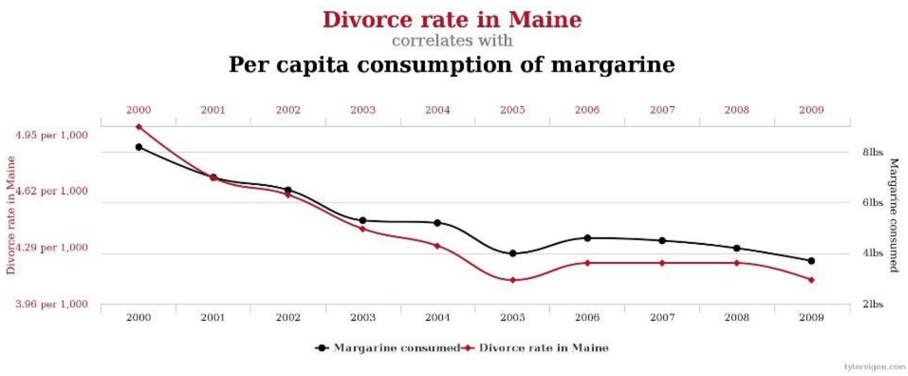

Figure 1. Correlation between divorce rate in Maine and per capita consumption of margarine [3]

In a study conducted by Tyler Vigen *[3]*, he realized that there was a 99% correlation between the divorce rate in Maine and the per capita consumption of margarine. This type of correlation is known as spurious correlation because it has appeared by random coincidence.

2. When the causal direction is unknown:

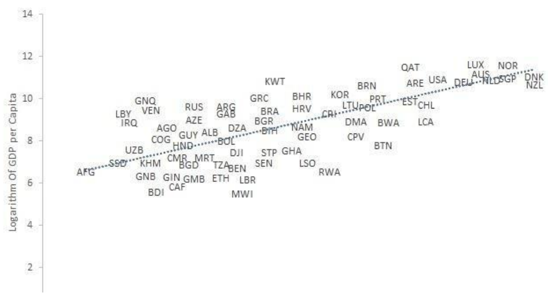

Figure 2. Correlation between corruption and GDP per capita of a country [4]

There are several studies that show the correlation between high levels of corruption in a country as measured by the Global Corruption Index (GCI) and having a low GDP *[4]*. Some of the conclusions of these studies imply that by fighting against corruption, then the GDP of a country will increase, but this is unclear.

In this case, correlation is not able to show the direction of what variable is causing the other. It could be the case that in poor countries, people must rely on corruption to survive, therefore the solution would not be to act against corruption but trying to increase each of the countries GDP's.

Causal analysis, as opposed to correlation analysis alone, can show the direction and strength of the relations and give insights on how to solve the underlying root cause of the problem.

3. When there is a common cause (confounder):

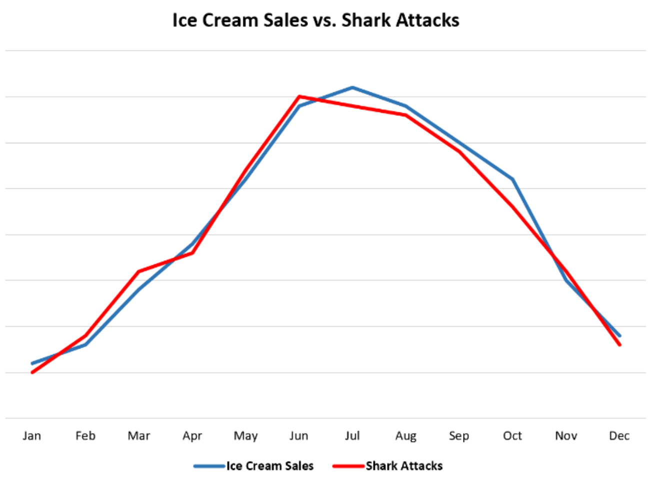

Figure 3. Correlation between shark attacks and ice cream sales

Variables that are caused by the same variable usually are correlated, but that does not give any insight if the common causal variable is not observed. Looking at the case of shark attacks and ice cream sales, the only reason why the correlation between both variables exists is due to the variable "temperature" which influences both variables in a common way.

Causal analysis can show this type of relations when correlation stays blind to this type of associations.

2.2 *HOW CAUSAL RELATIONS ARE INFERRED:*

There are multiple approaches to infer causal relationships between variables. The easiest way to understand the underlying logic behind causal inference is by examining the formula for the Average Treatment Effect (ATE):

$$ATE = E[Y_{T=1} - Y_{T=0}]$$

Y = Outcome

T = Treatment

The ATE formula essentially calculates the difference in outcomes between two treatment variables. For example, let's consider an outcome where Y=1 represents a person getting cured and Y=0 represents a person not getting cured. Additionally, let´s assume T=1 represents giving a person a treatment and T=0 represents giving them a placebo. By comparing the outcomes of two groups separated by treatment variable, we can infer if the treatment is causing the desired outcome of curing people. The following graph illustrates this example.

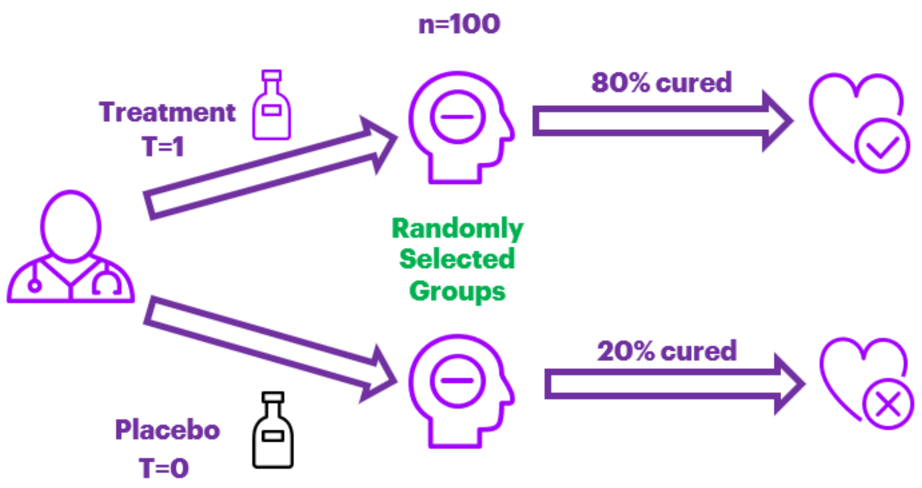

Figure 4. Treatment and Placebo diagram

By giving different treatments to two samples of the population of size 100, it is observed that 80% of the sample of people that received the actual treatment is cured vs the 20% of the people that received the placebo treatment. This implies that there is a difference of 60% between both groups and the ATE is of 0.6, which is significantly high so we could conclude that the treatment is causing people to get cured.

This technique to infer causal knowledge is known as **Randomized Control Trials** (RCT) and it is considered the golden rule for causal inference. It is important to note that both selected groups are completely random samples to test the treatment. If the group that received T=1 was formed just by girls and the group that received T=0 was formed by boys, it would be impossible to assess if the effect was caused by the treatment T or the difference in gender.

Randomized controlled trials (RCTs) have been used for a long time to study cause-and-effect relationships. However, there are cases in which conducting RCTs is not possible, such as in manufacturing. Testing different product configurations and analyzing multiple variables can be expensive and impractical as the search space can be enormous. This is why, inferring causal knowledge in **manufacturing requires a different and more sophisticated causal inference technique**.

In situations where randomized control trials cannot be executed, using observational data becomes a viable alternative for causal analysis. By leveraging existing data on production processes, product quality, and other relevant variables, using a machine learning approach can help manufacturers gain valuable insights into the factors that impact product outcomes.

While inferring causal relations just on observational data may have some limitations, such as potential biases or confounding factors, analyzing it carefully can provide valuable insights for optimizing manufacturing processes and improving product quality **without disrupting production**. In chapter 2.4 this thesis will delve deeper into the limits of building causal knowledge on observational data and the different approaches that try to solve this complex problem.

2.3 CAUSAL STRUCTURES:

In the previous chapters, we have explored the concept of causality and the traditional methods used for inferring causal relationships. Building upon that foundation, this chapter will focus on causal structures.

Causal structures represent the underlying mechanisms that govern how variables and events interact between each other. Understanding them is crucial for comprehending the complex dynamics of cause and effect. These structures are represented in the form of graphs, and they obey some properties.
The two most important properties are:

1) **Directedness**. Causal structures are directed, meaning that the relationships between variables or events have a specific directionality. There must be a direction to represent that one variable or event (the cause) influences or affects another variable or event (the effect). In association this is not the case because the relationships are symmetric or bidirectional.

2) **Acyclicity**: Causal structures shall-not have cycles or loops, meaning that a variable or event cannot cause itself, either directly or indirectly through a chain of causal relationships. This property helps establish a clear sequence of events over time and ensures that causality flows in a consistent and one-way direction.

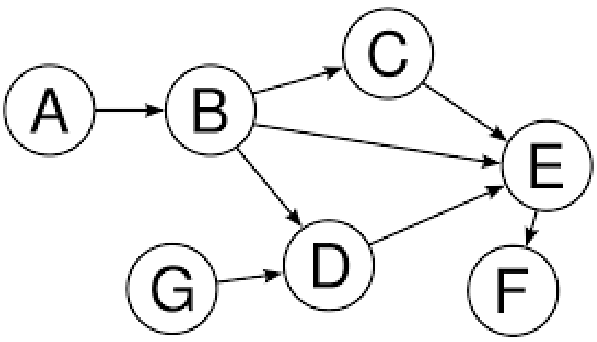

Figure 5. Example of a Directed Acyclic Graph (DAG)

These two properties define that causal structures are Directed Acyclic Graphs, or DAG's.

DAGs are graphical representations of causal structures that consist of nodes and arrows. The nodes represent variables or features, and arrows depict the direction of causality, indicating which node is the cause and which is the effect. Arrows also can have a weight that represents the strength of the relation. In DAG's, there are three basic structures: forks, chains, and colliders. **Forks** occur when a node has multiple outgoing edges, causing the graph to split into two or more branches. **Chains** are sequential connections between nodes, where each node has only one incoming and one outgoing edge, creating a linear structure. **Colliders**, also known as "v-structures," are nodes with multiple incoming edges that converge into a single outgoing edge, forming a "V" shape. Colliders are special because they have a unique property: they are blocked by their descendants, meaning that they do not directly affect each other but can be affected by a common ancestor. Understanding these basic structures is crucial in identifying and interpreting causal relationships within DAGs.

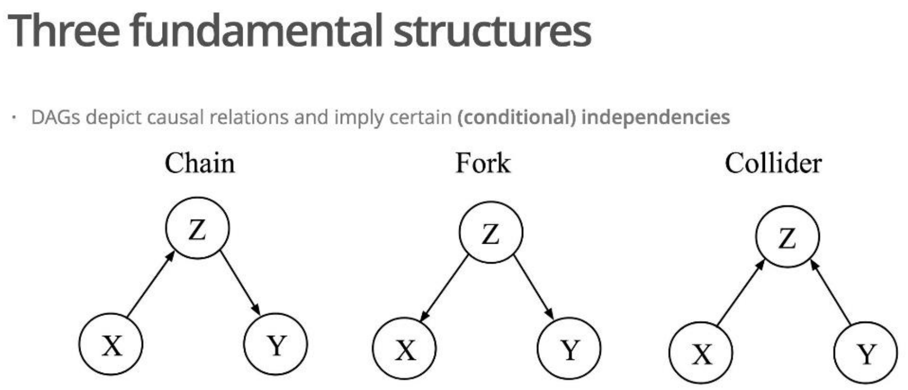

Figure 6. Chain, Fork and Collider structures in DAGs

2.4 CAUSAL INFERENCE LIMITATIONS:

Although causal inference can be a powerful technique, it has some underlying limitations especially when inferring causal relations without conducting trials, limited to observational data only. Some of these limitations are:

- **Selection bias**: Occurs when the process of selecting study participants is not random, leading to a biased sample that may not accurately represent the target population. This can result in misleading or incorrect causal conclusions. This thesis has already mention selection bias when explaining what RCT was. In figure 4, if there was a difference in gender between the placebo and the testing group, the selection of the groups could be causing a potential bias in the effect of the treatment.

- **Confounding variables:** Refers to the presence of uncontrolled variables that are associated with both the treatment and the target variable, leading to spurious associations. Confounders can distort causal inference and may need to be accounted for to establish true causal relationships. The shark attacks and ice cream sales correlation example perfectly illustrate this limitation, since the relation between both variables just makes sense when including the confounding variable (temperature).

- **Data quality and measurement errors:** As in every Machine Learning model, inaccuracies in data collection, measurement, or reporting can introduce errors and bias into the analysis, leading to incorrect causal inferences.

- **Generalizability**: Findings from one study or population may not be directly applicable to other populations or settings, limiting the external validity and generalizability of causal inferences.

2.5 *GENERAL APPROACHES TO INFERRING CAUSAL RELATIONS WITH OBSERVATIONAL DATA:*

There are several approaches to discovering causality when relying on observational data. These approaches aim to find the best possible configuration of a Directed Acyclic Graph (DAG) given a set of data. However, the number of possible DAGs is very large, making it difficult to find the optimal configuration. In fact, finding the optimal DAG is considered an NP-hard problem. Therefore, the only feasible option is to attempt to find a good enough solution heuristically or by using traditional Machine Learning methods.

The four groups of approaches to solutions are:

- **Constraint-based**
- **Score-based**
- **Functional**
- **Gradient-based**

Constraint-based methods:

The approach involves analyzing the statistical independence between triplets of variables in the system. By analyzing these independence relationships, the methods can begin to infer the causal relationships between variables. Essentially, if variables A and B are independent, but variable A and C are not independent, this suggests that there may be a causal relationship between A and C. Variables are analyzed in triplets because as explained in 2.3, there are only 3 possible structural combinations when analyzing the relations between 3 nodes (Chains, Forks, and Colliders).

Constraint-based methods are a useful tool for inferring causal relationships between variables, particularly when other information or prior knowledge about the system is not available.

The most used algorithm that follows this approach is the **PC algorithm** (Sprites & Glymour, 1991).

The algorithm works as follows:

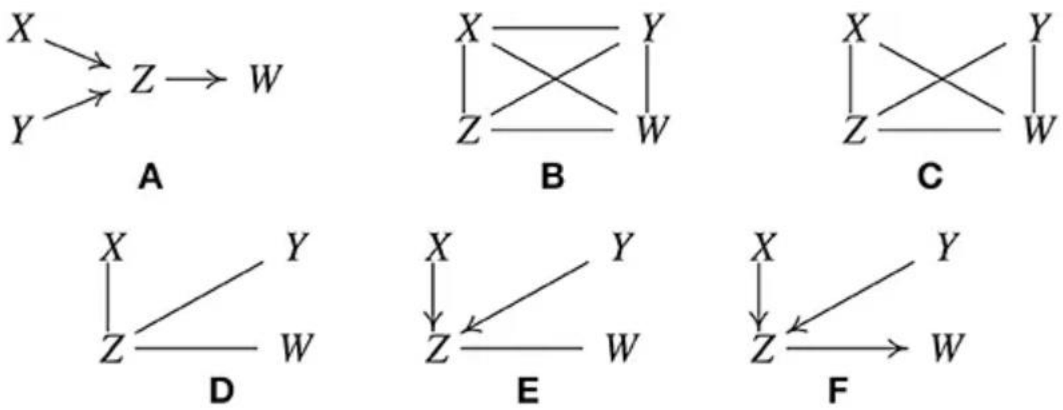

Figure 7. PC algorithm iterations

It begins with a fully connected undirected graph (B). Then on the next step, the edges between variables that are unconditionally independent are removed (C). After that, the ones that are also conditionally independent are removed (D). Once the graph is pruned, then the directions are inferred using as base the grounding blocks of DAGs, colliders, chains, and forks (E and F).

Score-based methods:

Score-based methods for causal discovery aim to identify the best causal structure of a given system by iteratively generating candidate graphs, evaluating how well each one explains the data, and selecting the best one.

These methods start with an initial graph structure, such as a fully connected or empty graph, and then modify the graph by adding, removing, or reversing edges. After each modification, the resulting graph is evaluated based on how well it explains the observed data, using some criterion or scoring function. Common scoring functions are the Bayesian Information Criterion (BIC), the Akaike Information Criterion (AIC), or the Maximum Likelihood Estimation (MLE).

One well-known example of a score-based method for causal discovery is the Greedy Equivalence Search GES algorithm.

The GES algorithm starts with an empty graph and iteratively adds, removes, or reverses edges based on a set of conditional independence tests until it arrives to a high score where the structure had converged on the best causal structure for the system. The score-based approach is a powerful tool for causal discovery, but it is resource intensive, meaning that it is computationally expensive, and it requires a careful tuning of the scoring function to converge to an optimal solution.

Functional methods:

Functional methods for causal discovery aim to identify the underlying causal relationships between variables by analyzing their functional dependencies. This involves estimating the parameters of a statistical model that describes the relationship between variables, and then using this model to infer the causal structure of the system. Functional approaches also use a score function to estimate the goodness of the solution, but the mechanics of functional methods are different from score-based methods. Rather than searching through candidate graphs to find the best fit to the data, functional methods make use of the distributional imbalances present in the data to detect the causal connections between variables.

The Linear Non-Gaussian Acyclic Model (LiNGAM) algorithm is a classic functional method that uses the non-Gaussianity of the data to identify the causal structure. LiNGAM exploits the non-Gaussian nature of the data to determine the probable causal relationships between variables. More specifically, when two variables have a non-Gaussian correlation, it implies a potential direct causal connection between them. On the other hand, if two variables have a Gaussian correlation, it indicates a possible indirect or confounding relationship between them.

Gradient-based methods:

Gradient-based methods for causal discovery are the latest and most advanced methods developed to date. In fact, for the use case project that will be presented in chapter 4, the algorithm used to infer the causal structure followed the gradient optimization method. Gradient descent approaches are particularly useful for solving optimization problems, where the objective function can be expressed as a gradient of a scalar function. The gradient is a vector that points in the direction of steepest increase of the function. By iteratively updating the parameters in the direction of the negative gradient, the algorithm can converge to the optimal solution. This approach has been applied in a wide range of applications in Machine Learning and Artificial Intelligence, including deep learning, image recognition, and natural language processing.

The algorithm that was used in the case study is called **NoTEARS** and it was developed by Xun Zheng, Bryon Aragam, Pradeep Ravikumar and Eric P. Xing. Quantum Black by McKinsey added an implementation of this algorithm to the python library we used called CausalNex.

2.6 NOTEARS ALGORITHM:

The NoTEARS algorithm is considered to be the current state-of-art algorithm for building causal knowledge on observational data. There has been several updates and improvements in the algorithm like DyNoTEARS, which include causal relationships on the temporal domain.

The NoTEARS algorithm was revolutionary because it was the first algorithm to frame structure learning as a purely continuous optimization problem. This was a perfect approach to search optimal solutions in the huge DAG solution space.

The algorithm works as follows:

The first step is to build the data matrix and the DAG matrix. In this case, X is the data matrix and is composed by *n* rows and *d* columns. *n* is the number of observations of the data and *d* is the number of features (nodes of the causal structure). The *W* matrix is the weight matrix, and it encodes the causal weights between variables. This is a non-symmetrical matrix, so the causal direction is established as rows point to columns. The size of the matrix is *d* x *d*, and the direction is encoded as row features cause column features. The causal weight encodes how much a feature is causing another feature to occur.

$$X \in \mathbb{R}^{n \times d}$$
$$\mathrm{W} \in \mathbb{R}^{d \times d}$$

The next step is to build the cost function that it is going to be optimized. To understand the cost function, it is essential to understand the basic assumption that is made before optimizing:

$$X \approx XW$$

In a noiseless case, this equation tends to be equal when *W* is optimal. It seems to be counterintuitive, but it is the grounding block of this algorithm and the reason why it works. Once this assumption is made then the cost function can be easily understood because it penalizes the difference between *X* and *XW*.

The causal weight loss function looks like this:

$$F(W) = \frac{1}{2n} \parallel X - XW \parallel_F^2 + \lambda \parallel W \parallel_1$$

The first term of *F* penalizes the suboptimality of *W*. As we know from the assumption, in the optimal case, the first term will tend to be zero. The second term is a regularization term of norm one that penalizes having non-important weights. This is a widely used technique that helps the equation to converge to a solution faster and reduces non crucial weights directly to zero instead of having causal relations with and insignificant causal weight.

The non-DAGness behavior of the weight matrix penalization function (h) that looks like this:

$$h(W) = Tre^{W \odot W} - d$$

This function computes the DAG-ness of W. The higher the value of *h* the more cycles the matrix contains. The way to do this is by computing the Trace of the exponential matrix of the Hadamard product of the *W* matrix. The computation and convergence of this function is slow and complex. The higher the dimension of *W* the slower the computation, so there is a limit to how many features can be included in the causal structure. This thesis will explain the consequences of this limit in chapter 3.

Finally, the general loss function is built as a mixture of function F and function h, so both the DAG-ness and the weight optimization are considered.

$$L^{\rho}(W,\alpha) = F(W) + \frac{\rho}{2}|h(W)|^2 + \alpha h(W)$$

Once the L function is built, then the gradient descent method is applied to optimize for a solution of W.

There are some limitations to this algorithm, and it does not always work perfectly. The main limitation occurs because the algorithm converges to what is known as a Markov Equivalence class. This is because the following expressions are mathematically equivalent:

Mathematically Equivalent

$$X_1 = w_{12}X_2 + \varepsilon_1$$

$$X_2 = w_{21}X_1 + \varepsilon_2$$

Figure 8. Mathematically equivalent expressions

The weights of the *W* function can converge to a DAG graph that has some rotated causal directions. In the following figure, one can observe the Markov equivalence classes. Graphs that belong to the same Markov class are enclosed within the same box.

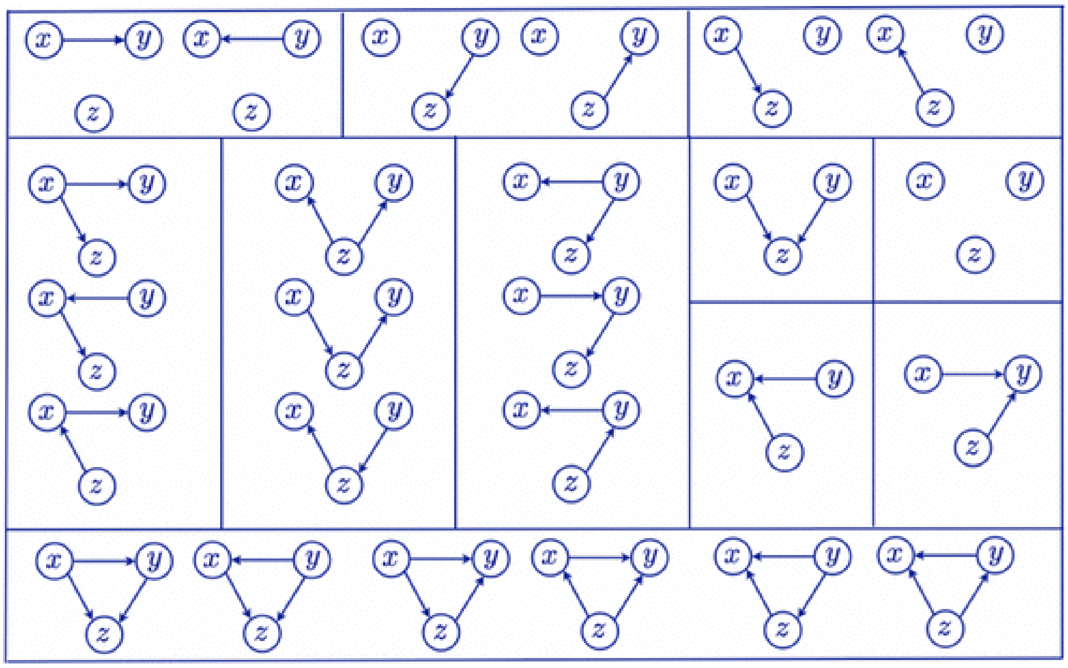

Figure 9. Markov Equivalent Classes

As it can be observed, graphs that belong to the same class are equivalent if the direction of the causes are not considered.

In most of the cases, the algorithm converges to an optimal solution, but there are exceptions where some of the causal directions are wrongly encoded. Usually, domain understanding will help finding these mistakes, but in some cases contacting with a subject matter expert is needed. We will talk about the importance of SME in chapter 3.

CHAPTER 3. CAUSALITY IN MANUFACTURING

In this chapter, the thesis is going to focus on explaining the implementation of causal inference in manufacturing. Firstly, it will introduce the general applications and then describe a general methodology explaining the key principles that must be considered.

3.1 GENERAL APPLICATIONS:

Causal analysis can have a wide range of applications in the manufacturing sector. Although the focus of this thesis is to explain a general methodology to apply causal analysis, it is important to explain the power of causal analysis by using two general use cases where causality can bring great insight and potentially help as a key tool in solving or addressing problems.

The first and main one is **fault detection**:

Every manufacturing line can arrive to a fault situation where the process is stopped due to a misbehavior of an element in the line. Fault detection is the process of identifying when a process is not behaving as expected and determining which are the variables that are causing the problem. It can be hard to figure out how a complex production system works, especially when it's too complicated to make a model of it from scratch and the underlying physics of the system are hard to parameterize and compute. Thanks to Machine Learning techniques, using data together with ML tools can help to show the system's dynamics better. As we have already explained in chapter 2, most of the traditional ML methods are not enough to find the cause of a problem because the models make conclusions using correlation and association, not causation. Causal analysis can help identify how the system works and how problems spread through the system. So, when a system is failing and the reasons that cause this fault are hard to see for an expert, applying causal discovery can be the solution.

Causal inference is also a helpful technique when applying **root cause analysis** to any manufacturing plant.

Industrial processes are becoming more complex and interconnected, with many different parts and controls. If something goes wrong in one part of the process, it can affect the entire production line and cause problems like a reduction of the product quality and incremental costs. It's important to quickly find the causes and fix them to keep the process running smoothly. Usually, experts that oversee the manufacturing line have built an intuition on how each production step is interconnected and are able to quickly identify the causes and act accordingly. But when this is not the case, causal structures can show in the underlying levers that regulate a manufacturing system. It can also provide deep insights and motivate changes and experiments in the process to improve the overall productivity and effectiveness. Understanding the system and how each part interacts with another also reduces the search space when trying to optimize for a variable inside the system.

3.2 BUILDING THE CAUSAL STRUCTURE:

There are several things that one must keep in mind when building any causal structure given a set of data. To build a robust and resilient causal structure, a two-step approach that tries to minimize the bias and errors in causal structures defining any manufacturing process is going to be explained. The first step focuses on the reduction of the error introduced by confounding effects. The second step focuses on the reduction of the run-time and complexity of the algorithm that oversees building the causal structure. In chapter 4, a real use-case of this methodology is provided.

3.2.1THE DANGERS OF CONFOUNDERS:

Confounders are variables that affect both the cause and the effect variables in a study. Since these affect simultaneously both variables, if these are not considered, the causal analysis can be greatly biased. Therefore, it is important to account for confounders in causal inference to ensure that the causal relations extracted from the data are representative of the reality and it is not just a relationship caused because there is a hidden confounder that hasn't been taken into account.

Together with the shark and ice cream example, there is another example that is widely used to explain the confounders effect. This example states the following: sleeping wearing shoes is causing people to wake up with a headache.

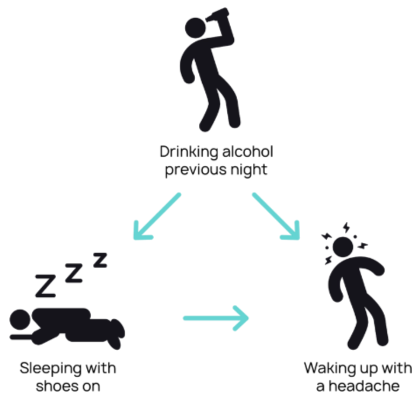

Figure 10. Confounder variable example

A statistical analysis was made to find if sleeping wearing shoes was causing people to wake up with headaches. When observing just the cause and effect, it seems to be clear that sleeping with shoes causes headache, but the reason both variables are related is because there is a common confounder, in this case drinking alcohol the previous night. If this confounder is not taken into account, then the inferences obtained conducting causal inference are incorrect. This example clearly shows the importance of including all the possible confounders in a study when conducting causal inference.

The way to fight against the negative effect of confounding variables in a manufacturing system is to enclose the system by adding all the variables and features that affect the system in any possible way. When including all the variables in a system, then, all the system confounders are observed, and the causal structure includes them in the graph removing the confounding effect.

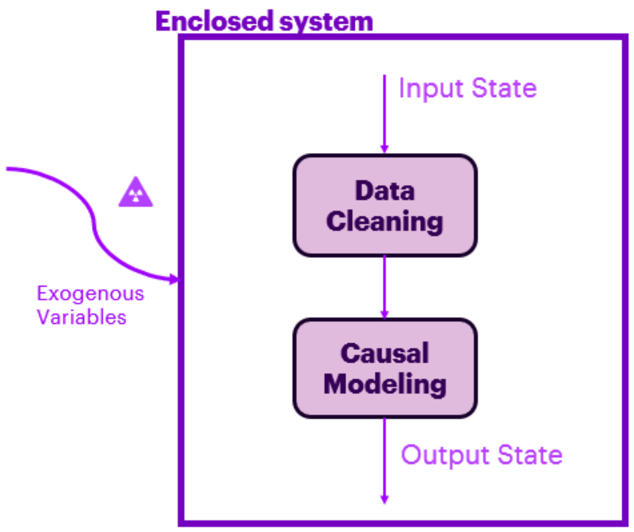

Figure 11. Enclosed system diagram

Figure 11 explains how by enclosing the system, the causal structure built on observational data is more resilient to biases caused by not including confounding variables. There is also a possibility that an exogenous variable, which is not considered as a variable of the system, may be influencing the system in an unobserved manner. For example, imagine you are analyzing a factory that oversees manufacturing chairs. You measure the quality of the chairs at the end of the process and track how all the machines are working. If the quality of the chairs depends solely on how the machines behave, then your analysis is fair. But if how the wood was chopped an important factor in making a good chair, and you don't consider it in your analysis, then you'll have a biased result due to the confounding effect. How the wood was initially chopped is an example of an exogenous variable. By paying attention to these exogenous confounding variables, you can improve your analysis and reduce errors and biases in your results.

Therefore, before building a causal structure based on observational data, it is important to **enclose the system subject to study**. This means including all the system variables and the possible exogenous variables that can have a confounding effect in the target variable. This is the first step of the general methodology for conducting causal inference on manufacturing processes.

3.2.2 THE COMPLEXITY PROBLEM:

As this thesis has explained in chapter 2, algorithms that build causal structures just relying on observational data are complex algorithms that take some time to converge to a solution. The complexity of these algorithms depends on the quantity of nodes (variables) that are forming the graph. This is because most of the algorithms work by analyzing the relationships between variables one to one, so the number of iterations increases exponentially with the number of variables.

Some manufacturing process need large complex machines to work and the amount of sensors that parametrize the system can be really large. In those cases, pre-processing techniques are required to reduce the number of variables. Pre-processing and variable reduction is a typical step that must be done in almost every data science project, but there are differences in what to consider when doing pre-processing for a prediction algorithm and doing it for a causal inference algorithm.

In Chapter Two, the distinction between association and causation in the relationship between variables was explained. In the context of predictive modeling, association relations are sufficient, which means that grouping associated variables during pre-processing does not affect the accuracy of the model, even if the variables are not causally related. One example of this approach is seen in the application of principal component analysis (PCA) to a dataset, where the principal components are formed by combinations of the original variables, and their combinations depend on the original associations between them. However, it is important to note that this approach is not suitable for analyzing causality. In causality it is important to keep the raw variables to understand what is really happening. If different variables are combined, the causation between those variables will be diluted and the whole purpose of the analysis is lost.

So, to obtain un-biased results, it is important to keep in mind that creating artificial variables by combining them can reduce information on causal relationships. Normally, the best way to deal with this situation is with the help of a Subject Matter Expert (SME), we will explain in the next section (3.3) of this chapter how an SME can help to empower causal analysis.

After studying convergence times and the complexity of the NoTEARS algorithm (which is the one that was used for the case study that will be explained in chapter 4), the recommended maximum number variables to include into the analysis is in the order of less than 100. If more than 100 variables are included, the convergence time breaks and increases exponentially.

3.3 THE IMPORTANCE OF SMEs:

Once the dataset is prepared, then the causal structure can be built. After building the causal structure, the next step is to look carefully at the relationships and check for possible errors. As we explained in chapter 2, the algorithms that extract causal knowledge solely on observational data approximate an NP algorithmic problem.

An NP algorithmic problem is a type of problem in computer science that is difficult to solve quickly, but easy to verify a potential solution once it is found. It stands for "nondeterministic polynomial time" and refers to problems that can be solved in a reasonable amount of time by using heuristic searches or random exploration. This implies that finding the exact optimal solution may take an impractical amount of time as the problem size increases so we must settle for approximations to optimal solutions. NP algorithmic problems are commonly found in areas like optimization, cryptography, and scheduling. Markov equivalence classes are a materialization of the limitations of this algorithm. SMEs are the helpers that find better conversions to solutions by providing industry insights and guidance.

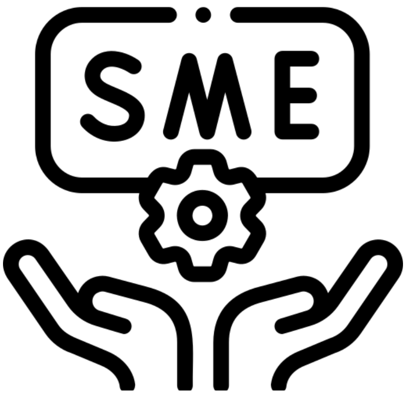

Figure 12. SME icon

In the context of a work or project, a SME (Subject Matter Expert) is a person who has extensive knowledge and expertise in a particular field or subject. They are considered the go-to person for information and guidance in their area of expertise. SMEs are typically hired or consulted to provide specialized knowledge, advice, and support to a project or team. They may be brought in to solve complex problems, develop strategies, or provide technical guidance. Their expertise is crucial in ensuring the success of a project and achieving the desired outcomes. They can be helpful in different ways when building a causal inference pipeline.

Firstly, they can help with the selection of variables and the data pre-processing. As it was explained in last section, cleaning the dataset, and reducing the amount of variables is a complex task that needs precision and knowledge of how the system subject to study works. The SME can help with the selection of important variables and cleaning the independent variables that do not affect the system. It is important the data scientist and the SME work collaboratively, as the data scientist can clean and prune a dataset without losing information, the SME can help and guide the decisions with his knowledge of the system.

It is important to note that the data cleaning and pre-processing phase is common to almost all the use cases, but when the number of variables surpass the hundredths then it becomes a risky phase because eliminating important information can affect the whole causal structure.

Secondly, the SME can also guide the causal structure algorithm to converge to optimal solutions. This is done by providing constrains to the algorithm, therefore decreasing the number of possible solutions and reducing the number of operations (algorithm converges faster). The SME can add constrains to the causal structures in three different ways:

- **Adding a causal relation**: The SME can force a relation to appear in the causal structure. Maybe, it is the case that it is known for sure that the behavior of a variable directly causes the behavior of another variable.
  (E.g., Variable X1 is causing variable X2)

Figure 13. Adding a causal relation between X1 and X2

- **Removing a causal relation**: The SME can also remove causal relations in the causal structure. It can be known that a variable is certainly not causing other variables to behave differently.
  (E.g., Variable X1 is not causing variable X2)

Figure 14. Removing the causal relation between X1 and X2

- **Changing causal directions**: The last action an SME can force into a causal structure is to change the direction of causation between two variables. This action is helpful when correcting the problems with

Markov equivalent classes. Sometimes, if the NoTEARS algorithm is used to build the causal structure, it can happen that some of the directions are wrongly encoded. It is not frequent because there are not too many combinations to a possible solution that ensures the DAG property and has equivalent Markov classes. The SME can identify some relations that do not make sense and rearrange the direction. (E.g., Variable X1 is not causing variable X2, it is the other way around)

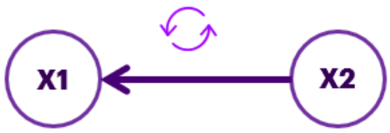

Figure 15. Changing the direction of the causal relation

3.4 OVERALL METHODOLOGY:

In this section (3.4) the thesis is going to explain the overall methodology and the different steps that must be followed to ensure the development of a correct and unbiased causal structure when studying a manufacturing system.

- **Step 1 (Data Collection):** Collect all the available data. It is important to include to the causal structure all the possible confounders. Be aware of exogenous variables that may be affecting the system. Use the SMEs to grant that every step of the system is covered with the data.

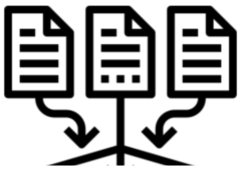

Figure 16. Data collection icon

- **Step 2 (Data Cleaning):** Once all the data is collected, if the number of variables is higher than 100 then the algorithm is going to take a long time to converge. Also, it is recommended to use around 3000 rows to train the structure. Reduce all the redundant variables applying correlation analysis. Also, iterate with the SME and delete variables that are not providing any information. The creation of features is also a good idea when the features represent a physical item of the system. Once the dataset is cleaned and reduced to a reasonable amount of variables the causal structure can be trained.

Figure 17. Data cleaning icon

- **Step 3 (Structure training):** Train the causal structure with any of the available algorithms. NoTEARS is currently one of the most advanced ones so it will give the better results. Once the structure is trained, save it, and plot it.

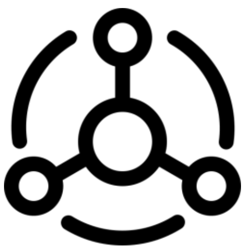

Figure 18. Structure training icon

- **Step 4:** Show the causal structure to the SME. It might happen that some of the causal associations are not there. The SME can add constrains to the causal structure by deleting, forcing, or changing the

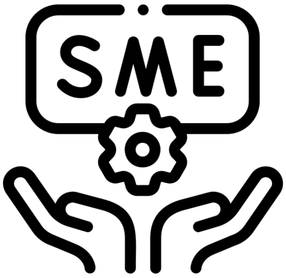

Figure 19. SME icon

direction of the causal relations. Once the constrains are added, the structure can be retrained. Iterate between **Step 3** and **Step 4** until a reasonable causal structure is achieved.

- **Step 5:** Visualization is the last part of the step. To extract the maximum amount of insights and give the client a clear understanding of their system there are different tools to visualize the structure. CausalNex, the library that includes the NoTEARS algorithm, has a powerful and very customizable tool to create structure visualizations.

Figure 20. Data visualization icon

CHAPTER 4. USE CASE – TISSUE MACHINE

Chapter 3 presented a general methodology for developing causal inference analysis in a manufacturing project. In this chapter, we will explore how this methodology was implemented in a real-life project. Firstly, the context of the project will be established, giving an overview of its purpose and objectives. Then, we will detail the steps taken to apply the methodology, providing a comprehensive description of the process. Finally, the results obtained from the application of the methodology will be presented, providing an evaluation of the project's success. Overall, this chapter serves to provide an in-depth analysis of the practical application of the methodology presented in chapter 3 and its effectiveness in achieving project objectives.

4.1 GENERAL CONTEXT:

A tissue paper manufacturer wanted to develop several projects to improve their production of roll paper. The manufacturer had a problem because there were times during the production phase that the roll paper broke, delaying production and increasing scrap. The initial project consisted of developing a predictive model to make predictions on when the roll paper was going to break 5 minutes before the break occurred. In addition to that, a **causal prescriptive model** was going to be developed to help understand the client which variables were causing the breaks to occur and help the workers achieve a better understanding of the whole system.

The roll paper machines are around 80 meters long, and the production is separated in many different phases. The first step is the creation of paper paste. This is done by processing wood pulp and refining it with different additives. Once the pulp is refined, it is sent to the tissue machine. The tissue machine oversees converting the paper paste to actual roll papers. This is the part that was analyzed during the project. The tissue machine is formed with several parts.

- **The headbox**, initial step where the pulp is formed into a paper sheet.
- **The press section**, it finishes the job of the headbox and completely press the pulp into a sheet. It squeezes out water and compress the fibers together.
- **The dryer section** is where the paper is dried out and the rest of the water is taken out of the paper. The paper is dried up in big rotating cylinders inside the Yankee. In the Yankee dryer, the steam-heated surface of the cylinder dries the sheet as it rotates. The Yankee dryer can be several meters in diameter and is designed to produce a high-quality, uniform finish on the paper sheet.
- The next step is the **creping process**. In the creping process, the dried paper sheet is scraped off the Yankee dryer with a sharp blade or creping blade, which causes the fibers in the paper sheet to loosen and stretch. This process creates the unique soft, fluffy texture and stretchy properties that are characteristic of tissue products.
- Finally, the paper is rolled in big reels preparing for the final step which consist of cutting and forming the rolling paper.

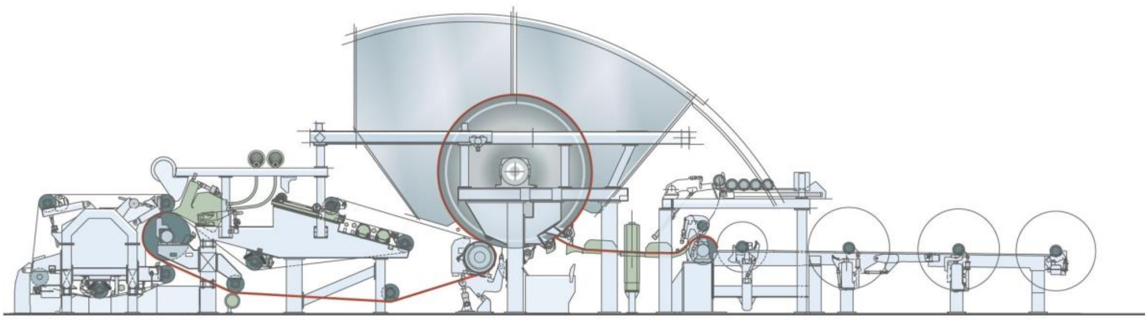

Figure 21. Tissue paper machine diagram

All these steps inside the tissue machine were sensorized and the process was parametrized from start to finish. Records of sensor data were taken each 5 seconds. From time to time, the process failed and the paper that was getting rolled in the big reels broke before arriving to an optimal length. When this occurred, the process had to be stopped and the paper rolled up to that moment had to be recycled and reprocess again, so breaks caused high costs to the client.

Machine operators where in charge of labeling when the breaks occurred. With the timestamps provided by the machine's operators and all the data retrieved from sensors, the idea was to create a predictive and prescriptive model to help the client optimize the manufacturing process. To support the prescriptive model, a causal structure map was provided to explain which variables were causing the process to break. Some of the production conditions could be modified by machine operators so the causal map would provide insights on how to tweak them to reduce the number of breaks.

4.2 DATA PREPARATION:

Since the tissue machine is a large complex machine with streaming data from instrumentation, the available data to construct the causal model consisted in more than 1100 entries recorded from different sensors. In some way, this are good news because it is ensured that most of the variables affecting the system are tracked. This helps to solve the problem of confounder bias in causal structures. But there was a problem to solve because building a causal map with that many data entrances is no possible due to the complexity of the algorithms that infer causal structures on observational data. As the thesis explained on chapter 3, the first step when building the causal structure is to get all the available data, and the second step is pruning and rearranging the data so it can be adapted for the algorithms. In this case, step 1 was easily completed since the manufacturing plant was greatly monitored. Step 2 was harder to complete since the number of variables had to be reduced from around 1000 to a list of around 100.

The approach that was followed was first to reduce redundant data by merging into one variable correlated measures of different sensors that measured the same phase of production. Also, by carrying out uni-variate analysis, data that has no variance was removed because it provided little insight. Then, with the help of the SME, other variables were discarded since they were not important and did not provide additional information.

After carefully merging variables and working on feature engineering, a set of 80 variables with one binary target variable indicating if a break occurred was provided to build the causal structure.

One of the variables, indicated the type of tissue paper that was being built during the time the sensors were collecting data. It was a categorical variable. In total, 12 different types of paper were being built, each one with significantly different characteristics and hence, different causal structures. The decision to split the data in 12 different datasets was made, one for each grade (type of tissue paper). In total, 12 different causal structures had to be built.

4.3 STRUCTURE TRAINING:

Having the 12 different datasets prepared, the next step was building the structures and iterating with the SMEs to obtain a reasonable optimal solution. With the first iteration, we realized that the causal weights pointing to the break variable were not high enough. We expected to obtain high causal weights pointing to the break variable, but that was not the case. We realized that the break variable in the dataset was not optimally tagged.

Break variable was set to 0 when the process was running correctly, and it was set to 1 five minutes before each break occurred. NoTEARS algorithm works better with continuous data, so having a binary target variable was limiting the algorithm to perform well. The problem was solved by making the break variable to increase in small steps from 0 (five minutes before break) to 1 (when the break occurred). The results obtained with this modification were significantly better.

There were several iterations over the causal structures adding and deleting causal relations. The results obtained consisted of a bundle of various DAGs, the biggest one contained the break variable. This means that there were other causal structures independent of break causing. Those structures were not considered for the project. Some of the causation weight between variables were small, so to clean the causal graph from spurious relations we decided to apply a threshold on causal weight.

Once the SMEs were in agreement with the obtained results, a custom visualization was created where the relations were shown in a clear manner.

4.4 VISUALIZATION:

It was decided that there were two graphs that could show the insights extracted from the causal maps. The first one was directly to show the resulting causal graph with the different connections and causal relations. The second one was conducting a Pareto analysis and showing the top variables that were causing the break to occur in each grade.

4.4.1 CAUSAL MAPS:

The causal map visualization consisted of making an easy to understand and clean representation of the DAGs. To do so, a color code was used to separate the variables in the different stages of production. Also, the thickness of the arrows showing the causal direction was dependent on the causal weight. The thicker the arrow, the bigger the causal weight. Finally, if a variable was causing the break (directly or indirectly), the edges and arrows coming out from that node were painted in blue. This representation was done for each of the 12 different causal graphs we obtained. Here is an example of one of them:

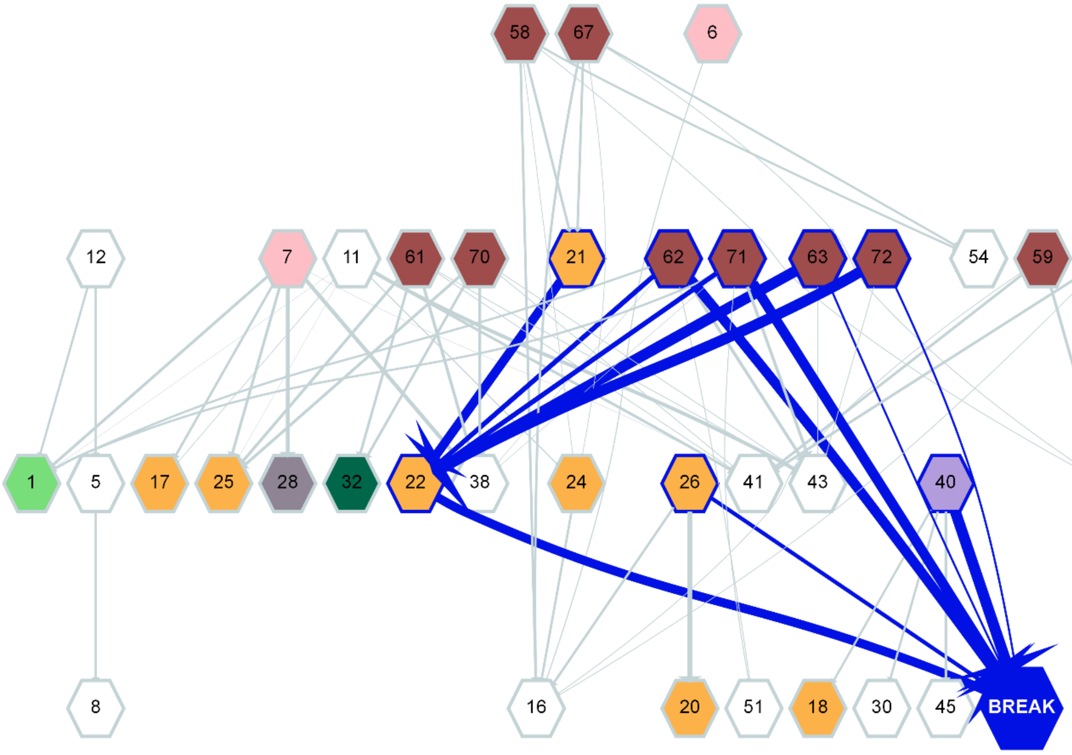

Figure 22. Causal graph showing relations of the tissue making process.

As *Figure 22* shows, there are nodes with different colors, each one representing the different stages of production. The white variables are general variables or variables that didn't belong to a specific stage. Also, it can be observed how the arrows change in size depending on the strength of the causal relation. The most important ones were pointing directly to the variable break (target variable).

Thanks to this visualization, SMEs and manufacturers could gain deep insights on how the system really worked and the variables that were affecting the misbehavior of the manufacturing line. This visualization also serves as a starting point for looking into modifying some of the flexible parameters in the production line to increase the productivity of the system and reduce the break occurrence.

4.4.2 PARETO ANALYSIS:

The pareto analysis visualization served as a summary of graph visualization. This visualization compiles the 20 variables that were causing the biggest impact in break occurrence for each type of tissue paper and for all the tissue papers in common.

It showed a bar plot of the causal weights towards the target variable, a line plot with the causal weight accumulation. Also, a color code was added to differ between variables that were common across all the different types of papers (dark blue) and the specific variables for each type (light blue).

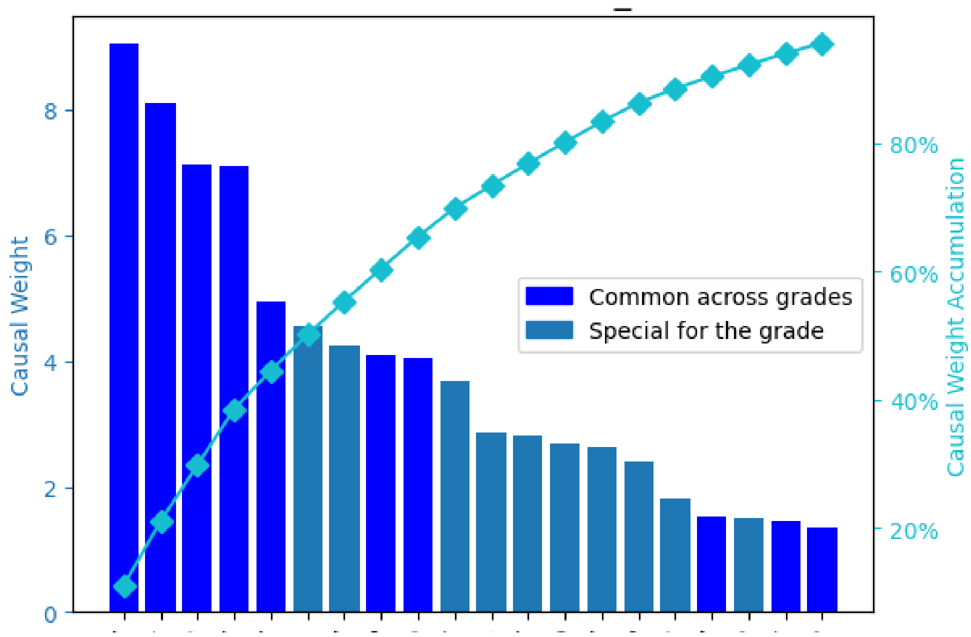

Figure 23. Pareto analysis visualization

CHAPTER 5. CONCLUSIONS AND FUTURE WORK

Applying causal inference to manufacturing projects is not an easy task. As it has been explained throughout this thesis, there are a lot of things that one must keep in mind when applying these techniques.
Although it is a complex task, as it has been shown in Chapter 4, it is doable, and it provides very powerful insights to SMEs and manufacturers.

The feedback we received from the client when we showcase the results were positive. The causal structures helped the SMEs understand more deeply how the production line behaves. Also, the visual representations where very useful to drive change and make manufacturers try new approaches. Causal maps can be considered as some kind of explainable AI, where the different levers that influence in production quality are clearly shown.

Usually, SMEs and manufacturers are quite reticent about black box machine learning models. It is hard for them to trust predictions when they don't understand how the model is constructed, but in the case of causal inference, this blockage to change don't appear because the solution is clear and understandable.

As future works, more causal inference algorithms could be tested, especially the DoWHY library from Microsoft.

CHAPTER 6. BIBLIOGRAPHY

[1] Correlation vs. Causation | Difference, Designs & Examples https://www.scribbr.com/methodology/correlation-vs-causation/

[2] Correlation vs Causation: Understand the Difference for Your Product https://amplitude.com/blog/causation-correlation

[3] Spurious correlations: Margarine linked to divorce? https://www.bbc.com/news/magazine-27537142

[4] What are the costs of corruption? https://blogs.worldbank.org/governance/what-are-costs-corruption

[5] Strengthening causal inference from randomized controlled trials of complex interventions. - https://gh.bmj.com/content/7/6/e008597

[6] Causal Structure. https://www.sciencedirect.com/topics/mathematics/causal-structure

[7] Causal inference & directed acyclic diagrams (DAGs). - https://bookdown.org/jbrophy115/bookdown-clinepi/causal.html

[8] Unlock the Secrets of Causal Inference with a Master Class in Directed Acyclic Graphs.

https://towardsdatascience.com/unlock-the-secrets-of-causal-inference-with-a-master-class-in-directed-acyclic-graphs-f2d3b40738e

[9] The Limits of Causal Inference. https://visionsinmethodology.org/wp-content/uploads/sites/4/2016/05/Wise_VIM_05_09_2016.pdf

[10] Causal Discovery Learning causation from data using Python. - https://towardsdatascience.com/causal-discovery-6858f9af6dcb

[11] Assessing bias: the importance of considering confounding. - https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3503514/

[12] Machinery Used In Tissue Paper Making. https://www.pulpandpaper-technology.com/articles/machinery-used-in-tissue-paper-making

ANNEX I: ALIGNMENT OF THE PROJECT WITH THE SDGS

This project, with the aim of increasing the efficiency and productivity of manufacturing processes can contribute significantly to the development of Sustainable Development Goals (SDGs), particularly Goals 8, 9, and 12.

- **Goal 8**: Decent Work and Economic Growth

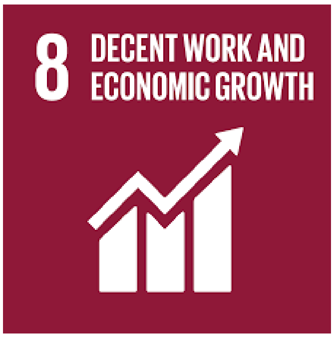

The eighth goal focuses on promoting inclusive and sustainable economic growth, full and productive employment, and decent work for all. By improving manufacturing processes' efficiency and productivity, the project can help generate more employment opportunities and enhance the quality of work. Increased efficiency often leads to cost savings, enabling businesses to expand and invest in workforce development. This, in turn, creates more job opportunities, reduces unemployment rates, and enhances overall economic growth.

- **Goal 9:** Industry, Innovation, and Infrastructure

Goal 9 emphasizes the development of resilient infrastructure, fostering sustainable industrialization, and encouraging innovation. A project targeting the improvement of manufacturing processes aligns perfectly with this goal. By implementing innovative technologies, streamlining operations, and optimizing resource utilization, the project can contribute to sustainable industrial growth. Efficient manufacturing processes reduce waste, minimize resource consumption, and enhance productivity, ultimately promoting the development of sustainable infrastructure and the adoption of cleaner technologies.

- **Goal 12:** Responsible Consumption and Production

Goal 12 aims to ensure sustainable consumption and production patterns. Manufacturing processes often have significant environmental impacts due to energy consumption, waste generation, and resource depletion. By increasing efficiency, the project can reduce the environmental footprint of manufacturing activities. It can optimize material usage, minimize waste generation, and implement recycling and reuse practices. By promoting sustainable manufacturing practices, the project contributes to responsible production, supports sustainable supply chains, and encourages the efficient use of resources.