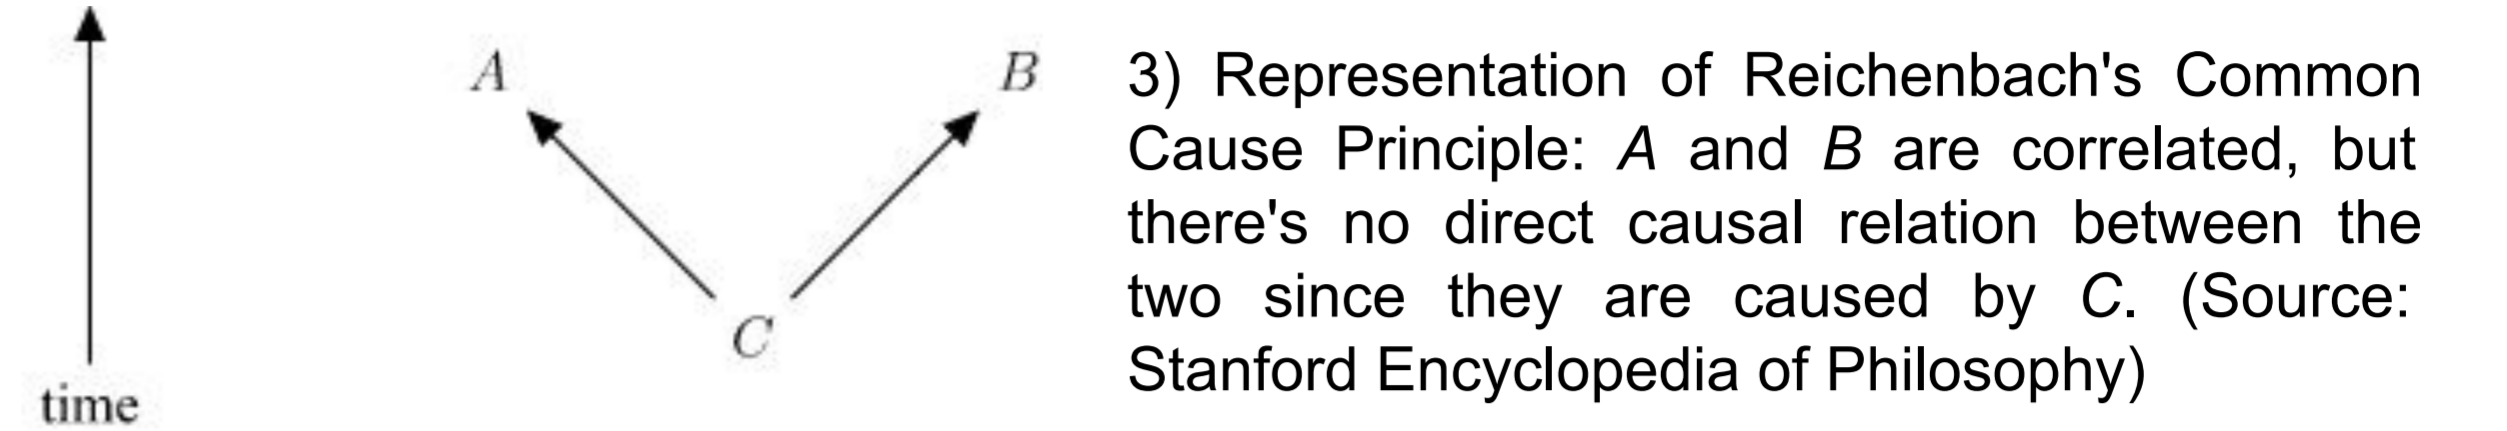


Motivation

The Principle of Causality is fundamental to virtually all scientific studies. However, often we are unable to experiment with a system to witness the causes and effects it contains, so we must infer causality from collected data.

In this work, we describe a method for inferring causal relations through the computation of quantities called Mutual Informations (MIs). The graph can be pruned by further computing another quantities called Conditional Mutual Informations (CMIs). This shall serve as a first step for future analysis that involve the determination of root causes in real world systems, such as assembly lines. In fact, in this work, we shall apply the described methods to construct a network of causal relations of an assembly line from publicly available data [1].

Finally, we can use Reichenbach's Common Cause Principle [3] in order to further prune the network. This principle states that if two events are correlated, then either they are directly causally linked or there's a common cause between the two (see Figure 3).



This principle can be included in a new quantity, called conditional mutual information, that is, the mutual information between X and Y conditioned to Z. It is defined as

$$CMI(X, Y | Z) = - \iiint_{X, Y, Z} f_{x, y, z}(x, y, z) \log \left(\frac{f_{x, y | z}(x, y | z)}{f_{x | z}(x | z) \cdot f_{y | z}(y | z)} \right) dx dy dz.$$

This quantity works similarly to MI and allows to exclude edges between X and Y, if conditioned by Z, when CMI is null. We could also compute CMIs conditioned to more than one node, but it might not be worth it due to the time and memory required.

Introduction

The question at hand is if statistical analysis is sufficient for establishing a causal relation or if human judgment is always a necessity. There are attempts, such as Pearl's [2], to axiomatize causal inference. In this work, we shall use some causal arguments in order to construct a faithful causal network with the objective of reducing, as much as possible, the need for a human element in the process.

The first thing we need is to study the dependence between each pair of variables. One of the simplest ways of doing so is to measure correlations. This method is however, insufficient. We want a metric that can encompass different types of distributions. A more robust approach is to compute pairwise mutual information instead. This quantity, computed for two continuous random variables, X and Y, is defined as


$$MI(X, Y) = - \iint_{X, Y} f_{x, y}(x, y) \log \left(\frac{f_{x, y}(x, y)}{f_x(x) \cdot f_y(y)} \right) dx dy,$$

where f_x , f_y are the probability density functions for X and Y, respectively, and $f_{x, y}$ is the joint probability density function of X and Y. This is always a non-negative quantity, with equating 0 in the case of statistical independence for all values of X and Y. In fact, we shall assume that $MI(X, Y)=0$ implies there is no causal relations between X and Y.

Mutual information is tightly related with the concept of entropy, which for a continuous random variable X is defined as

$$H(X) = - \int_X f_x(x) \log(f_x(x)) dx.$$

We can write $H(X, Y) = H(X) + H(Y) - MI(X, Y)$, evidenced by Figure 1). However, the continuous version of entropy has the caveat that it can be negative (see Figure 2).



1) Venn diagram indicating the mutual information - $I(X;Y)$ - corresponds to the entropy of the information that is both in X and Y. (Source: Wikipedia)

2) The orange distribution has entropy 0, the green has a positive entropy, and the blue negative, since we can state more precisely in this last one where a random sample will lie.

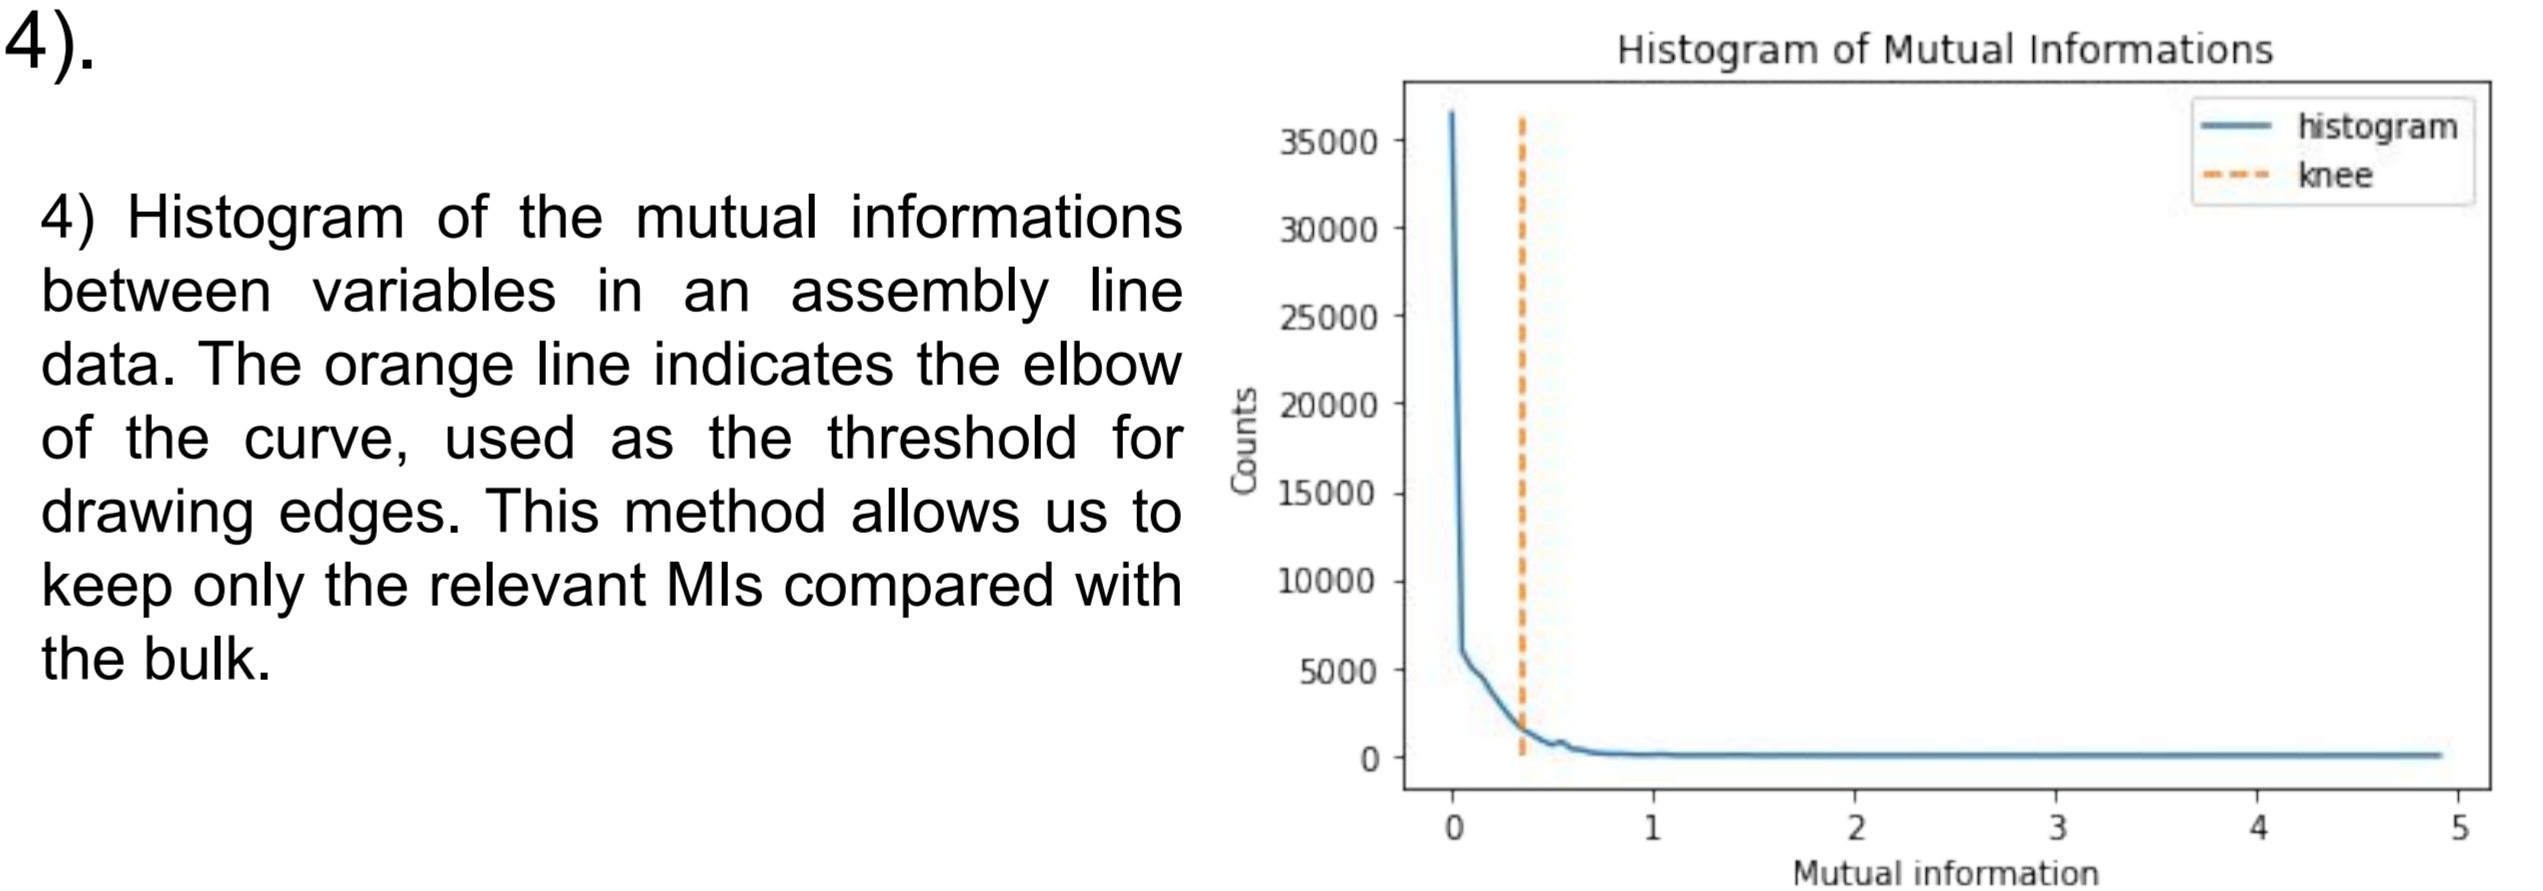
In order to construct a causal network, we shall take each variable as a node of the network and draw edges whenever MI exceeds a given threshold. We can use time in order to decide the directions of the edges and at the same time exclude edges between simultaneous variables.

Algorithm

The method to compute either entropies, MIs, or CMIs is very similar. We define probability distribution functions (pdf's) histogramming the variable data, substitute them in the formula and sum over all the bins multiplied by their distance. This method is dependent on the number of bins, so we used Freedman–Diaconis rule [4] to compute the optimum bin distance for integration. The code is available in [5].


This code was developed in *Python*, for which there are already packages that compute these quantities. However, our code was twice as fast as *SciPy's differential_entropy* function when tested for uniform, normal, or exponential distributions and provided more accurate results for the first two. Furthermore, it allows us work with sparse pdf's decreasing the consumption of memory.

Edges are draw between pairs of non-simultaneous variables whose MI exceeds a threshold. The threshold is chosen by the Kneedle method [6] as the elbow of the distribution of MIs (see Figure 4).



4) Histogram of the mutual informations between variables in an assembly line data. The orange line indicates the elbow of the curve, used as the threshold for drawing edges. This method allows us to keep only the relevant MIs compared with the bulk.

Figure 5, below, represents the causal relations between variables in a real assembly line using this method. Each node represents a station of the assembly line and can contain more than one variable, thus the graph contains multiple edges. Some stations don't take numerical measurements, thus no connections appear.



5) Continuous causal relations in an assembly line with 4 lines and 52 stations.

References

- <https://www.kaggle.com/c/bosch-production-line-performance/overview>
- Pearl, J. (2000). *Causality: Models, Reasoning and Inference*. Cambridge University Press
- Reichenbach, H. (1956). *The direction of time*. University of California Press, Berkeley.
- Freedman, D.; Diaconis, P. (1981). "On the histogram as a density estimator: L2 theory". *Probability Theory and Related Fields*. doi:10.1007/BF01025868
- <https://github.com/F-Barroso/Information>
- Satopää, V., et. al. (2011). Finding a "Kneedle" in a Haystack: Detecting Knee Points in System Behavior. 31st International Conference on Distributed Computing Systems Workshops, doi:10.1109/ICDCSW.2011.20.