



HAL
open science

Causal Discovery Under Local Privacy

Ruta Binkyte, Carlos Pinzón, Szilvia Lestyán, Kangsoo Jung, Héber Hwang
Arcolezi, Catuscia Palamidessi

► **To cite this version:**

Ruta Binkyte, Carlos Pinzón, Szilvia Lestyán, Kangsoo Jung, Héber Hwang Arcolezi, et al.. Causal Discovery Under Local Privacy. Proceedings of Machine Learning Research, 2024, 236, pp.325-383. hal-04617032

HAL Id: hal-04617032

<https://hal.science/hal-04617032>

Submitted on 19 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Causal Discovery Under Local Privacy

Ruta Binkyte

Carlos Pinzón

Szilvia Lestyán

Kangsoo Jung

Héber H. Arcolezzi

Catuscia Palamidessi

RUTA.BINKYTE@INRIA.COM

CARLOS.PINZON@INRIA.FR

SZILVIA.LESTYAN@INRIA.FR

GANGSOO.ZEONG@INRIA.FR

HEBER.HWANG-ARCOLEZI@INRIA.FR

CATUSCIA.PALAMIDESSI@INRIA.FR

Editors: Francesco Locatello and Vanessa Didelez

Abstract

Differential privacy is a widely adopted framework designed to safeguard the sensitive information of data providers within a data set. It is based on the application of controlled noise at the interface between the server that stores and processes the data, and the data consumers. Local differential privacy is a variant that allows data providers to apply the privatization mechanism themselves on their data individually. Therefore it provides protection also in contexts in which the server, or even the data collector, cannot be trusted. The introduction of noise, however, inevitably affects the utility of the data, particularly by distorting the correlations between individual data components. This distortion can prove detrimental to tasks such as causal discovery. In this paper, we consider various well-known locally differentially private mechanisms and compare the trade-off between the privacy they provide, and the accuracy of the causal structure produced by algorithms for causal learning when applied to data obfuscated by these mechanisms. Our analysis yields valuable insights for selecting appropriate local differentially private protocols for causal discovery tasks. We foresee that our findings will aid researchers and practitioners in conducting locally private causal discovery.

Keywords: local differential privacy, d -privacy, causal discovery.

1. Introduction

The notion of causality is gaining popularity in machine learning (ML) because of its benefits for accuracy (Richens et al., 2020), robustness (Tople et al., 2020; Schölkopf et al., 2021), explainability (Madumal et al., 2020) and fairness (Loftus et al., 2018). Many applications of causality in ML rely on knowing the “causal structure” in the data (Binkyte-Sadauskienė et al., 2022; Kyono and Van der Schaar, 2021). Causal discovery algorithms are computational methods that aim to infer causal relationships from observational data (Nogueira et al., 2021; Spirtes and Glymour, 1991). Those algorithms mostly rely on correlations between the various components (*variables*) of the data. These correlations can be affected by the application of data-privatization mechanisms aiming at protecting the privacy of the data providers. However, protecting data privacy is a legal obligation in Europe and many other countries worldwide. Answering to this necessity, numerous privatization methods have been developed to maximize the trade-off between a good level of data privacy and utility.

Differential privacy (DP) (Dwork et al., 2006) is one of the most popular data-privatization approaches. Depending on the trust model, DP can be further classified into *central* and *local*. Central DP, which is the original notion of DP, assumes the existence of a trusted server where the

data is aggregated. Data consumers (analysts) cannot access the data set directly but only query it via the server, which is supposed to obfuscate the answer by controlled noise, before reporting it to the analysts. The DP property establishes a bound on the ratio of the probability of getting the same reported answer from two adjacent databases, namely, two databases that differ for just one record. The bound is expressed in terms of a parameter ϵ , which represents the level of privacy. DP is used nowadays in a variety of applications from programming languages (Reed and Pierce, 2010) to social networks (Narayanan and Shmatikov, 2009) and geolocation (Machanavajjhala et al., 2008).

One limitation of the central DP model is that the server or the data collector cannot always be trusted: they may collude with an attacker, or just be unable to protect the data from security breaches. For this reason, local DP (LDP) has been proposed as an alternative model (Kasiviswanathan et al., 2008; Duchi et al., 2013). In LDP, the individual data are obfuscated directly at the end of the data provider, before even being collected. The main advantage of LDP is that users are more willing to share their data when they don't need to rely on the trustworthiness of the data collector and the server. This model has become popular, especially thanks to the fact that has been adopted and promoted by High Tech leading companies such as Google (Erlingsson et al., 2014), Microsoft (Ding et al., 2017) and Apple (Differential Privacy Team, 2017).

A variant of DP called *d-privacy* (also known as *metric privacy*), was introduced in (Chatzikokolakis et al., 2013). *d-privacy* is suitable for domains provided with a notion of distance. Like in central and local DP, *d-privacy* imposes a bound on the probability that the same result is obtained from two different objects (the arguments of the mechanism). However, in contrast to DP, this bound does not depend only on the parameter ϵ , but also on the distance between the objects. This means that the noise can be calibrated depending on how large the range in which we want to achieve indistinguishability is. In contrast, LDP requires indistinguishability between any pair of elements in the domain. *d-privacy*, therefore, is particularly useful in those applications in which hiding an element within a group of neighbors is a sufficient measure of privacy protection. *d-privacy* has been applied especially in the local model, and in particular, in the context of location privacy, where it takes the name of *geo-indistinguishability* (Andrés et al., 2013).

In general, the addition of noise tends to reduce the utility of the information that can be extracted from the data. Many privatization approaches and denoising techniques have been optimized for the summary statistics of the individual variables in the data, such as average values. However, notions of utility also depend on the correlation between the various components of the data, especially in the case of causal discovery. Some approaches to cope with this problem have been proposed in the global DP setting when the full unobfuscated data set is available. For example, the collected data may be synthesized using generative algorithms such as GAN (Jordon et al., 2019) or Bayesian Networks (Zhang et al., 2017). However, little or no instances of relation-preserving local DP mechanisms are known for causal discovery. Under the local setting, the data are already obfuscated before they get to the central server, and, therefore, the methods used in global DP are not applicable.

In this work, we experimentally assess the impact of state-of-the-art LDP and *d-privacy* mechanisms on the structural accuracy of causal discovery from the data. More precisely, as the LDP representative, we consider the *k*-Ary Randomized Response (*k*-RR, Section 3.2.1) (Kairouz et al., 2016). As the local *d-privacy* representative, we considered the Geometric mechanism (Section 3.2.2). We conduct extensive experiments on both real and synthetic data sets, and we evaluate their impact on 9 causal discovery algorithms, including constraint-based, score-based, and causal asymmetry-based methods. In summary, the two main contributions of the paper are the following:

- The paper systematically compares the performance of different locally differentially private mechanisms, specifically Geometric and k -RR, in the context of causal discovery tasks. With our findings, we highlight the advantages of using Geometric privatization methods over k -RR, shedding light on the impact of noise levels on algorithm performance.
- We introduce a unified privacy measure from an attacking perspective, allowing for the comparison of two distinct privacy notions: LDP and local d -privacy. This measure facilitates the assessment of privacy-utility trade-offs in real-world tasks such as causal discovery.

These contributions collectively enhance our understanding of locally differentially private mechanisms in the context of causal discovery and offer valuable insights into their application in real-world scenarios. Indeed, we hope this work can aid practitioners in collecting multidimensional user data in a privacy-preserving manner by providing insights into which locally private mechanism and causal discovery algorithms are best suited to their needs.

2. Related Work

Causal discovery with DP is an emerging research area that aims to combine the benefits of both identification of causal relationships among variables and privacy-preserving data analysis. The goal is to discover causal relationships between variables while preserving the privacy of sensitive data. One explored approach in the literature for differentially private causal discovery was to incorporate DP mechanisms directly into existing causal discovery algorithms (Kusner et al., 2016; Xu et al., 2017; Wang et al., 2020; Ma et al., 2022). These algorithms introduce controlled noise during the causal learning process to ensure privacy protection.

However, these existing differentially private causal discovery algorithms assume the centralized DP model, which requires collecting users’ original data. The approach adopted in this paper is to leverage the concept of local DP (Kasiviswanathan et al., 2008; Duchi et al., 2013) for causal discovery (respectively local d -privacy (Chatzikokolakis et al., 2013)). In recent years, there have been several works on the local DP setting (e.g., see (Wang et al., 2017; Erlingsson et al., 2014; Differential Privacy Team, 2017; Ding et al., 2017; Kairouz et al., 2016; Duchi et al., 2013; Acharya et al., 2019; Arcolezi et al., 2022; Kikuchi, 2022; Cormode et al., 2018) and references within), and applying them to causal discovery involves sanitizing the data at the individual level. Parallel to our work, Agarwal and Singh (2021) study a class of corruptions, such as measurement error, missing values, discretization, and differential privacy in the US Census. However, their goal is to learn a causal parameter (average treatment effect) from corrupted data and they conduct experiments only in an aggregated setting. Similarly, Ohnishi and Awan (2023) offer causal inferential methodologies to analyze locally differentially private data. Mooij et al. (2016) experiment with causal discovery with small amount of noise added to the data. However, the noise is not produced by the privatization mechanism. These goals differ from our work, they investigate the effect of noise in causal effect estimation of a treatment (or intervention) when randomized experiments are impossible to conduct, thus statistical theory is needed. Our work solely focuses on causal discovery, that is the inference of causal *relations*, causal *directions* among a set of variables (i.e., “how the change in X influences Y ?” versus “is X the cause of Y ?”). To the authors’ knowledge, this is the first work that thoroughly explores and analyzes the impact of locally differentially private mechanisms on causal discovery.

3. Preliminaries

3.1. Privacy Notions

3.1.1. LOCAL DIFFERENTIAL PRIVACY

One privacy model considered in this paper is LDP (Kasiviswanathan et al., 2008; Duchi et al., 2013), which is formally defined as follows.

Definition 1 (ϵ -Local Differential Privacy) *Let $\epsilon > 0$ be a parameter representing the level of privacy loss. A randomized mechanism \mathcal{M} satisfies ϵ -local-differential-privacy (ϵ -LDP) if, for any pair of input values $v_1, v_2 \in \text{Domain}(\mathcal{M})$, and any possible output x of \mathcal{M} , the following holds (where $\mathbb{P}[e]$ represents the probability of the event e):*

$$\mathbb{P}[\mathcal{M}(v_1) = x] \leq e^\epsilon \cdot \mathbb{P}[\mathcal{M}(v_2) = x].$$

In essence, LDP guarantees that it is unlikely for the data aggregator to infer the true value from the reported data. The privacy loss ϵ controls the privacy-utility trade-off. Note that lower values of ϵ result in tighter privacy protection. Similar to global DP, LDP also has several fundamental properties, such as robustness to post-processing and composition (Dwork et al., 2014).

3.1.2. LOCAL d -PRIVACY

d -Privacy assumes that the domain of the mechanism \mathcal{M} is provided with a notion of distance d .

Definition 2 *A mechanism \mathcal{M} satisfies d -privacy, with privacy parameter ϵ , iff for all values, $v_1, v_2 \in \text{Domain}(\mathcal{M})$ and all possible outputs x , the following inequality holds:*

$$\mathbb{P}[\mathcal{M}(v_1) = x] \leq e^{\epsilon d(v_1, v_2)} \cdot \mathbb{P}[\mathcal{M}(v_2) = x].$$

In essence, in the local model d -Privacy guarantees, like in LDP, that it is unlikely for the data aggregator or an attacker to infer the true value v from the reported data. But in this case, it is because it is made indistinguishable from all the other values in the neighborhood. In other words, nearby secrets should look almost identical to any observer.

3.2. Privacy Mechanisms

In this section, we describe the various discrete multidimensional mechanisms used in this paper. Visually, Figure 6 in Appendix A depicts the 4 mechanisms applied to a single point in a 4D space with shape $(2, 5, 5, 5)$, denoting the number of categories or bins per dimension.

3.2.1. k -ARY RANDOMIZED RESPONSE (k -RR)

Randomized Response (RR) was proposed in (Warner, 1965) with the aim of providing “plausible deniability” to individuals responding to embarrassing (binary) questions in a survey. Kairouz et al. (2016) generalized RR to domains of arbitrary size k (with $k \geq 2$), and proposed the so-called k -RR mechanism, which is one classical technique for achieving LDP on categorical/discrete data. Given a data domain V , and the privacy parameter ϵ , let $k = |V|$ and $p := \frac{e^\epsilon}{k-1+e^\epsilon} \in (0, 1)$. For each $v \in V$, let $\eta_{\neq v} \in V$ be a uniform random variable (i.e., exogenous noise with uniform distribution) over $V \setminus \{v\}$. We let k -RR : $V \rightarrow V$ be the random variable given by:

$$k\text{-RR}(v; \epsilon) := \begin{cases} v, & \text{with probability } p \\ \eta_{\neq v}, & \text{with probability } 1 - p. \end{cases}$$

This mechanism satisfies ϵ -LDP (Kairouz et al., 2016), because $\frac{p}{q} = e^\epsilon$, where $q := (1-p)/(k-1)$. When collecting data in practice, one is often interested in multiple attributes of a population, i.e., *multidimensional data*. We assume there are d attributes with domains A_1, A_2, \dots, A_d , where each A_i is a discrete set of finite size $k_i = |A_i|$. Each data provider u_j for $j \in \{1, 2, \dots, n\}$ contributes to the data set with a tuple (record) $\mathbf{v}^{(j)} = (v_1^{(j)}, v_2^{(j)}, \dots, v_d^{(j)})$, where $v_i^{(j)}$ represents the value of the attribute A_i . We now describe the two main known methods for applying k -RR on multidimensional data (Arcolezi et al., 2022; Kikuchi, 2022; Domingo-Ferrer and Soria-Comas, 2022).

k -RR Component-wise (k -RR C-wise). This is a naive approach that applies k -RR independently on each attribute. More precisely, k -RR C-wise splits the privacy budget ϵ among the d attributes uniformly or proportionally to their size, and reports each attribute in A_i using k_i -RR parameterized with ϵ_i -LDP, for $\sum_{i=1}^d \epsilon_i = \epsilon$. In this paper, we set $\epsilon_i = \epsilon \cdot \frac{k_i}{k_1 + k_2 + \dots + k_d}$.

k -RR Combined (k -RR Comb). This mechanism considers the Cartesian product $A_1 \times A_2 \times \dots \times A_d$ as a single attribute and sanitizes it using k -RR parameterized with ϵ -LDP, where $k = k_1 \cdot k_2 \cdot \dots \cdot k_d$.

3.2.2. BOUNDED GEOMETRIC MECHANISM

The geometric mechanism is the discrete analogous of the Laplace mechanism. The output Y is related to the input X by the formula:

$$\mathbb{P}[Y = y | X = x] = p_{\max} \exp(-\epsilon |y - x|) \quad (1)$$

for some parameters ϵ that represents the level of privacy. p_{\max} is a normalization factor, i.e., it is chosen so that $\sum_y \mathbb{P}[Y = y | X = x] = 1$. This formula is valid in 1D, in which $|\cdot|$ denotes the absolute value, as well as in multidimensional Euclidean space, in which x and y are discrete vectors and $|\cdot|$ denotes the Euclidean norm, or any other p -norm chosen in advance (see Figure 7 for a comparison). From the definition of the geometric mechanism, it is immediate that it satisfies local d -privacy with privacy parameter ϵ , where the metric d is the chosen p -norm based distance.

In this paper, we are interested in bounding the geometric mechanism so that the output domain equals the input domain, as in k -RR. There are three natural ways to do it, namely (1) clipping, (2) replacing samples that are out of the box with uniform noise, and (3) resampling whenever a sample is out of the box. Let us review them in more detail.

The method (1), clipping, consists of replacing all the output values that lie outside the box with the closest values that lie inside the box, i.e., with the maximum or minimum values of the domain in the 1D case. In this case, the two extremes of the box may increase their probabilities excessively, and the property that the output y with maximum probability is always $y = x$ might be lost, especially when the input x is close to the border. In method (2), whenever the output y is outside the box, it is replaced with a uniform sample from the box. In terms of the probability distribution of the mechanism, this method crops it from the background (two tails in the 1D case), and rescales the cropped distribution by adding a constant. This addition results in combinations of exponential terms with additive constants, which adds complexity to the formulas unnecessarily and

distorts the exponential shape and its decay properties. Instead, in method (3), which corresponds to sampling as many times as necessary until the output is inside the box, the cropped distribution is simply multiplied by a constant. This preserves the main shape of the distribution while also keeping the formulas relatively simple. For this reason, we prefer method (3) over the other two.

Notice that bounding is not symmetric, except for the input in the center of the box. This means, that we should have different values of p_{\max} or ϵ for different values of x so that the bounded summation is 1 on all x . As it will be justified in Section 4, we opt for fixing p_{\max} , so the formula that characterizes the bounded geometric mechanism becomes:

$$\mathbb{P}[Y = y|X = x] = p_{\max} \exp(-\epsilon_x |y - x|)$$

where both x and y are constrained to a fixed bounded discrete set, and ϵ_x are chosen so that $\sum_y \mathbb{P}[Y = y|X = x] = 1$. These values always exist (assuming $p_{\max} \geq 1/k$), and we provide an algorithm for finding them.

The computation of ϵ_x for every x is not possible symbolically through a formula. It is required that $\sum_y \mathbb{P}[Y = y|X = x] = 1$, or equivalently, $\sum_y \exp(-\epsilon_x |y - x|) = \frac{1}{p_{\max}}$, where both x and y are constrained to a fixed bounded discrete set. In the 1D case, the domain is a set of k contiguous integers and for the smallest value of x , only one tail of the geometric distribution intersects the domain, which allows us to write $\frac{1}{p_{\max}} = \sum_y \exp(-\epsilon_x |y - x|) = \sum_{\delta=0}^{k-1} \exp(-\epsilon_x \delta) = \frac{1 - \exp(-k\epsilon_x)}{1 - \exp(-\epsilon_x)}$. However, there is no analytical solution for ϵ_x from this formula. Moreover, for the remaining values of x , the expression becomes more complex, as an additional term is added for the second tail, and even more for the multidimensional case.

Nevertheless, the computation of each ϵ_x can be carried out numerically by exploiting the fact that $\sum_y \exp(-\epsilon_x |y - x|)$ is decreasing on ϵ_x . At one extreme, if $\epsilon_x \rightarrow 0$, the sum approaches k , and at the other, if $\epsilon_x \rightarrow \infty$, the sum approaches 1. This implies, first, that there is a unique point ϵ_x for which this function crosses the threshold $\frac{1}{p_{\max}}$, and more importantly, that we can use a binary search to compute ϵ_x . In the multivariate domain, the summations still satisfy the monotonicity property. Therefore, this method can be used to implement the multidimensional geometric distribution. Similar to k -RR, we compare two versions of the Geometric mechanisms, i.e., Geo Comb and Geo C-Wise.

3.3. Causality Notions

3.3.1. CAUSAL GRAPH

A directed acyclic graph (DAG) $\mathcal{G} = (\mathbf{V}, \mathcal{E})$ is composed of a set of variables/vertices \mathbf{V} and a set of (directed) edges \mathcal{E} between them such that no cycle is formed. Let \mathbb{P} be the probability distribution over the same set of variables \mathbf{V} . \mathcal{G} and \mathbb{P} satisfy the Markov condition if every variable is conditionally independent of its non-descendants given its parents. Assuming the Markov condition, the joint distribution of variables $V_1, V_2, \dots \in \mathbf{V}$ can be factorized as:

$$\mathbb{P}[V_1, V_2, \dots, V_d] = \prod_i \mathbb{P}[V_i | Pa(V_i)]. \quad (2)$$

where $Pa(V_i)$ denotes the set of parents of V_i . A partially directed acyclic graph (PDAG) is a special type of DAG that contains directed and undirected edges.

3.3.2. CAUSAL DISCOVERY ALGORITHMS

Causal discovery is concerned with the identification of causal relations from the data. More precisely, it aims to learn the fully directed DAG or partly directed PDAG that best describes the given data set. Several causal discovery algorithms exist for a wide range of different assumptions, for a survey see (Glymour et al., 2019).

4. Tuning the Level of Privacy

The parameter ϵ in LDP does not have the same meaning as the ϵ in d -privacy, i.e., they represent different levels of privacy. In order to compare the mechanisms of these two families, we need to tune the respective ϵ 's so as to represent the same level of privacy. To avoid confusion for the readers that know the standard notion of DP, and are not so familiar with LDP, it is important to remind that the standard notion for privacy in the local framework is not the same as in the central one: In central DP, the challenge for an attacker is to distinguish between two adjacent data sets, i.e., data sets that differ for presence or absence of one record. In other words, the attacker wants to infer whether or not a certain record is in the data set or not. In LDP, on the contrary, the aim of the attacker is to infer the true value of the individual data provider.

To measure the level of privacy, therefore, we consider the probability that an attacker has to infer the true value from the reported value. Naturally, the attacker will put her bet on the value that has the maximum posterior probability, given the obfuscated value (Arcolezi et al., 2023; Chatzikokolakis et al., 2023). We note that this measure of privacy is directly related to the notion of *advantage of an attacker* in security, and to the notion used to assess the vulnerability of the training set in ML.

In both k -RR and d -privacy, the value that has the highest probability to be reported is the true value itself, hence the level of privacy provided by these mechanisms (assuming a uniform prior) is the probability to report the true value. Specifically, the level of privacy provided by k -RR with parameter ϵ is:

$$Priv_{k\text{-RR}}(\epsilon) := \frac{e^\epsilon}{k-1+e^\epsilon}.$$

whereas, for a Geometric with parameter ϵ' , the level of privacy is:

$$Priv_{Geo}(\epsilon') := \mathbb{P}_{\max} \cdot e^{\epsilon'-0} = \mathbb{P}_{\max}.$$

where p_{\max} is the normalization factor used in the definition of the geometric mechanism (Equation (1)). Tuning the parameters of k -RR and L to provide the same level of privacy means adjusting the above ϵ and ϵ' so that $Priv_{k\text{-RR}}(\epsilon)$ and $Priv_{Geo}(\epsilon')$ give the same result.

5. Experimental Results

In this section, we empirically assess how locally private mechanisms impact causal discovery. We evaluate the performance of 9 causal discovery algorithms in multidimensional, two-dimensional, real and synthetic data sets obfuscated using the various mechanisms described in Section 3.2. We start by applying the causal discovery algorithms to discretized non-privatized data. Then we select the algorithms that performed best at a particular data set and apply them to the privatized versions of this data set. We measure the effect of each privatization method on the algorithms by comparing the Structural Hamming Distance (SHD) score and the F1 score or Accuracy on the non-privatized and privatized data. We use the Benchpress causal discovery benchmarking framework (Rios et al., 2021) to generate synthetic data and run causal discovery algorithms for multidimensional experiments. As

(L)DP mechanisms are randomized, we report average results over 5 runs. Due to space constraints, we have included all of our additional experiments in Appendix D.

5.1. Data Sets

We use real benchmark and synthetic data sets for the experiments. The details can be found in Table 1 and in Appendix B.

Name	Type	Nodes	Bins	Size	Origin
Sachs	real	11	10	902	(Sachs et al., 2005)
Human Stature	real	3	10	898	(Han et al., 2015)
Synth10	synthetic	10	10	5000	random DAG, IID, Linear, Gaussian
Synth5	synthetic	5	5	50000	random DAG, IID, Linear Gaussian
CEP	real	2	2-100	94-16382	(Mooij et al., 2016)

Table 1: Data sets used for causal discovery. For CEP the number of bins was determined by $\min(u, 100, u * 0.1)$, where u denotes the number of distinct values.

5.2. Causal Discovery Algorithms

We apply constraint-based and score-based causal discovery algorithms for multidimensional data. We select several well-known algorithms that can run on discretized data. For pairwise data sets, we apply algorithms that are capable of identifying the causal direction for two variables. We test the performance of the discrete and continuous data-specific versions of the algorithms, as well as various parameter values. The details can be found in Table 2 and in the Appendix C.

We have used two libraries for implementation: we used the Benchpress (Rios et al., 2021) package for the PC, FCI, FGES, Iterative MCMC and MMHC causal discovery algorithms and metrics. For the RECI, IGCI, CDS and ANM methods we used the Causal Discovery Toolbox (Kalainathan et al., 2020).

Algorithm	CI Test/Score	Parameter
PC (Spirtes and Glymour, 1991)	Gaussian, Chi-square	Alpha (0.001,0.05, 0.1)
FCI (Entner and Hoyer, 2010)	Fisher-Z, Chi-square	Alpha (0.01,0.05,0.1)
FGES (Ramsey et al., 2017)	BIC	Penalty discount (0.75,0.8,1,1.5)
Iterative MCMC (Kuipers et al. (2022))	BGe	Alpha (0.001,0.01,0.1)
MMHC (Tsamardinos et al., 2006)	BDe	Alpha (0.01,0.05, 0.1)
RECI (Blöbaum et al., 2018)	MSE	
IGCI (Danusis et al., 2012)	sp1	
CDS (Fonollosa, 2019)	std. dev.	Forced Decision
ANM (Hoyer et al., 2008)	HSIC	

Table 2: The structure learning algorithms.

5.3. Discretization

In order to apply the discrete mechanisms of interest to our data set, it was necessary to discretize the original continuous data. Discretization is a critical step in the process, as it plays a pivotal role in the subsequent data analysis. There are several approaches to discretizing data, each with varying effects on the quality of the results. Some of these methods yield higher average precision, up to the

highest possible (Pinzón et al., 2020), but rely on knowledge of properties about the underlying data distribution, such as quantiles or an estimation of the density function. However, in situations where the underlying data is sensitive and private, revealing such properties can risk privacy breaches, so it is safer to assume that they are unknown. To address this challenge, we opted for the simplest method of discretization, namely uniform bins within a fixed range. In practice, this fixed range corresponds to estimations of the minimum and maximum values of the population.

The only parameter we can freely choose in this process is, therefore, the number of bins and it should be chosen taking into account that more bins imply more accurate information being revealed. Moreover, the number of dimensions of the data, which corresponds to the number of columns in the data set, also plays a role in the choice of the number of bins, as it increases exponentially the total number of bins. We chose between 5 and 10 bins for data sets with 3 or more dimensions, and for the CEP data set, which has two dimensions but contains several different data sets, we applied a dynamic number of bins (see Table 1). Some of these data sets were already discretized (e.g., had only 2 distinct values), and some had continuous data. We determined the number of bins by $\min(u, 100, u * 0.1)$, where u denotes the number of distinct values in a given data set.

5.4. Evaluation Metrics

For the data sets with more than two-dimensions we used structural hamming distance (SHD) to measure the difference between the ground truth adjacency matrix and the output of the causal discovery algorithm. It assigns a distance of 1 for every missing, redundant or reversed edge in the graph. Intuitively, SHD provides a number of edges that are need to be added, removed and re-directed to make the two graphs identical. We have also calculated the F1 score, that combines the precision and recall of a model, and is used to evaluate the recovery of the skeleton of the DAG. In case of the CEP data set, we have applied the same method as in (Mooij et al., 2016). Forced-decision: given a sample of a pair (X, Y) the methods *must* decide on a causal direction. Then, we evaluate the weighted¹ accuracy of the decisions. We also calculate the confidence intervals assuming a binomial distribution using the method by Clopper and Pearson (1934).

5.5. Results on Multidimensional Data

We report the results for the algorithms that performed the best on the discretized, but not privatized data. PC algorithm performed the best on most of the data sets. The iterative MCMC algorithm was performing better on the data sets with 10 or more nodes. Both data sets with 10 or more nodes show that causal discovery algorithms in general perform better under geometric privatization methods rather than k -RR. For the Sachs data set (Figure 1), PC and GES algorithms perform almost the same on Geo C-wise and Geo Comb. The performance on data privatized with the geometric mechanism is very close to the performance on the original data without the noise. For the Synth10 data set (Figure 3), the performance on data privatized with geometric mechanisms with $p_{\max} = 0.5$ outperforms the results on the original data. However, this result can also be accidental. For the Synth10 data set, we also observe a slightly better performance when Geo Comb is applied as compared to Geo C-wise. Performance is better with k -RR C-wise privatization than with k -RR Comb privatization on the Sachs and Synth10 data sets. For smaller multidimensional data sets (Figures 2 and 4) the variation of the performance is too large to draw reliable conclusions. This is probably due to the high influence of chance on recovering the data structure when the true graph is small. However, we still observe a slight advantage in applying geometric mechanisms to Synth5

1. Not all pairs can be considered as independent. Weights' list was acquired from the authors' website.

and Human Stature data sets. We can also observe slightly better SHD results with K -RR C-wise privatization than with K -RR Comb privatization on Synth5 and Human Stature data sets.

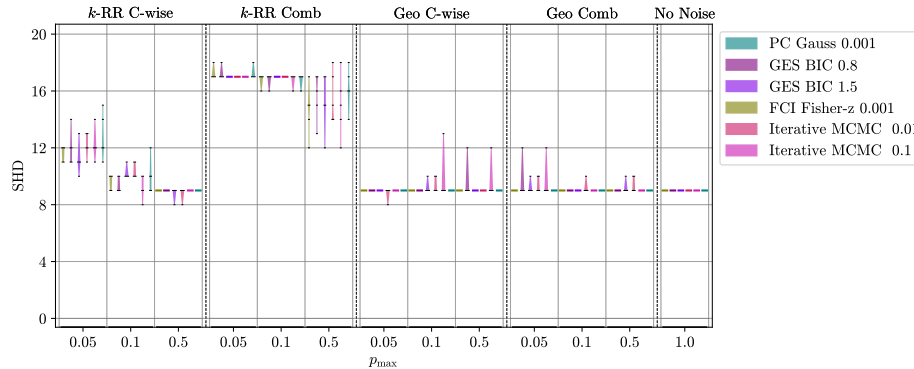


Figure 1: Sachs data, SHD. The results for PC algorithm with Gaussian CI test and alpha value 0.001; GES algorithm with BIC score and penalty discount values 0.8 and 1.5; FCI algorithm with Fisher-z CI test and alpha values 0.001; Iterative MCMC algorithm with BGe score and alpha values 0.01 and 0.1. The width of each bar varies for different values on the y-axis proportionally to the number of samples attaining that value.

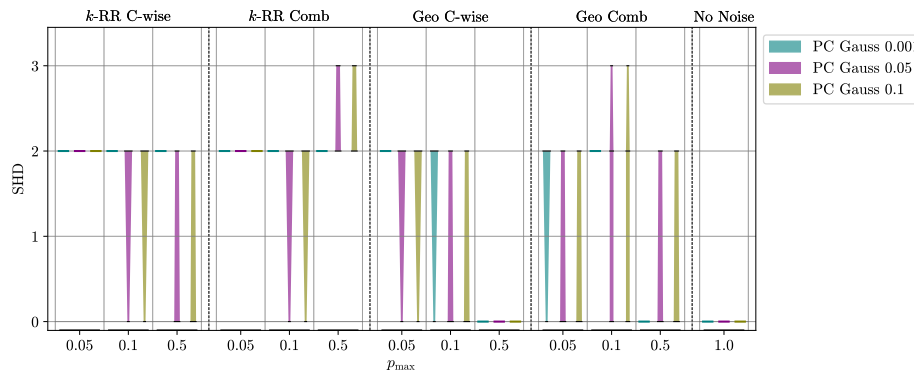


Figure 2: Human Stature data, SHD. Results for PC algorithm with Gaussian CI test and alpha values 0.001, 0.05 and 0.1. The width of each bar varies for different values on the y-axis proportionally to the number of samples attaining that value.

In our additional experiments in Appendix D.1, we observe similar results when measuring the F1 score for the causal discovery of an undirected graphs (Figures 8, 35, 90, 145).

5.6. Results on Two-dimensional Data

We report the results of all causal discovery algorithms applied for the CEP data set. In Figure 5, we show the results before (“No Noise”) and after privatization. It is evident that, similar to previous experiments, the Geometric mechanism consistently outperforms k -RR, with notable improvements, especially in the case of RECI, where the accuracy surpasses the baseline. We hypothesize that this phenomenon could be attributed to the potential data augmentation properties of noise addition, although further research is required to confirm this. The CDS algorithm performs similarly after

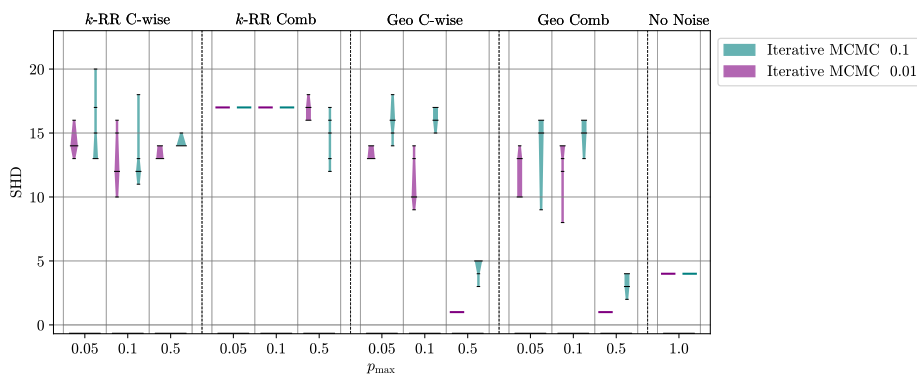


Figure 3: Synthetic data, 10 nodes, SHD. The results for Iterative MCMC algorithm with BGe score and alpha values 0.01 and 0.1. The width of each bar varies for different values on the y-axis proportionally to the number of samples attaining that value.

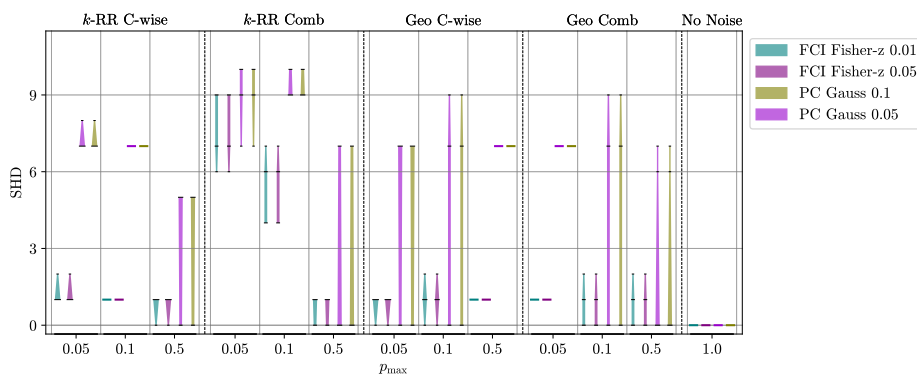


Figure 4: Synthetic data, 5 nodes, SHD. The results for FCI algorithm with Fisher-z CI test, alpha values 0.01 and 0.05; PC algorithm with Gaussian CI test, alpha values 0.1 and 0.05. The width of each bar varies for different values on the y-axis proportionally to the number of samples attaining that value.

privatization, except when applying the k -RR Comb mechanism. But k -RR Comb generally has the poorest performance (also with Sachs and HS data sets), and we think this is due to the available small sample size, and because the mechanism is affected by the curse of dimensionality. ANM exhibited unsatisfactory performance even before noise introduction, and its performance deteriorated further (sometimes falling below chance levels) after privatization.

6. Discussion

Our results consistently demonstrate that **geometric privatization methods (both component-wise and combined) exhibit higher accuracy in terms of SHD compared to k -RR methods (both component-wise and combined)**. In the case of geometric noise, the algorithms do not seem to perform much worse as the noise increases. This can be expected because this privatization method is not disruptive of the correlations in the data. It would be an interesting extension to also evaluate its effect on the model parameters. On the other hand, k -RR noise deteriorates the data structure and more noise results in worse performance of the causal discovery algorithms. We observe similar

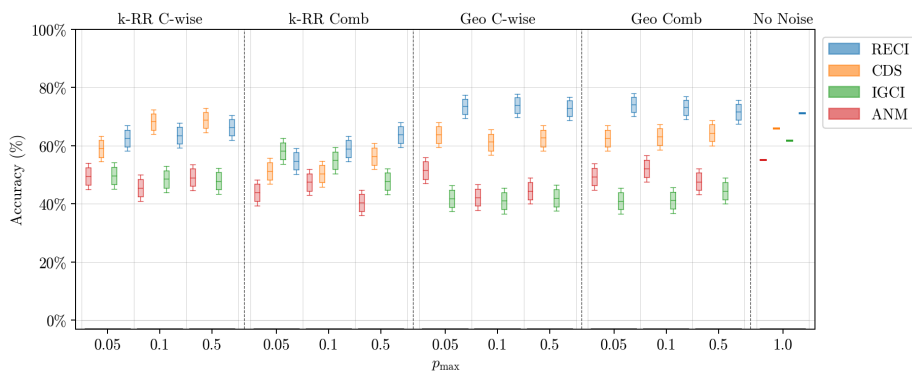


Figure 5: CEP data set with 2 nodes, weighted accuracy. Box whiskers are at 95%, body is at 80% confidence.

results when measuring the performance of causal discovery algorithms with the F1 score. The reason the geometric noise has less negative impact on causal discovery algorithms is that geometric noise in general tends to substitute a data point with one that is “similar” (in the sense of being not too distant, numerically). More precisely, the closer the point, the more likely it is chosen for replacement. In contrast, k -RR substitutes (with a certain probability) the original data point with any other point in the domain, chosen with uniform probability, regardless of the distance. Hence, the causal relation is preserved better by the geometric noise, especially when the relation is preserved by proximity, in the sense that if two data points are related, also their immediate neighbours are. We observe some dependence between the higher parameter alpha (PC) or penalty discount (GES) parameters and better F1 scores on the noisy data in the experiments on multidimensional data. Higher parameter values result in sparser graphs and help avoid spurious edges in the graphs. We also observe that algorithms that are less accurate on the original data are also less sensitive to data privatization. More precisely, when applied to privatized data, their performance drops less compared to the baseline on the original data (the detailed results can be found in Appendix D). However, the algorithms which are best on the original data still provide best overall results under geometric noise (despite being more sensitive to k -RR noise).

Although this paper focuses on empirical studies, we would like to extend the discussion with some theoretical considerations. For this aim, we considered two viewpoints: (1) how the LDP noise affects the independence tests, and (2) how the noise affects the causal discovery algorithms that are not based on independence tests (e.g., IGCI and RECI).

6.1. Independence Test

In the main body of our paper, we used the Fisher Z-test because we observed that it performs better than the χ -square test (see Appendix D). This is in line with the results of Gaboardi and Rogers (2018), which show that locally adding Laplacian noise (the continuous version of the geometric mechanism) changes the χ -square statistic so that it is no longer a χ -square random variable. Indeed, our results shown in Figures 1, 4, and 8 show that the output of causal discovery algorithms that use the Fisher Z-test had little to no change in 3 out of 4 settings (the exception is the k -RR Comb method that gives consistently bad results that is due to the curse of dimensionality). We did not find any previous work on the effect of LDP on the Fisher Z-test, but this could be an interesting line of work whether the Fisher Z test is generally robust to LDP. In (Gaboardi and Rogers, 2018), the authors design new hypothesis testing algorithms to compensate for the noise and to get a more

precise estimation of the χ -square. We think that it would also be an promising research direction to explore how these modifications could be applied to perform causal discovery with χ -square tests on the locally privatized data. The paper by [Gopi et al. \(2020\)](#) proves the existence of an algorithm to test independence on the data resulting from the application of the k -RR mechanism, which coincides 2/3 of the time with the test on the distribution of the original data. This has not been used for causal discovery on data sanitized with k -RR yet. The paper of [Kap et al. \(2021\)](#) explores the effect of “natural” noise (e.g., measurement error) on the performance of causal-discovery on ANMs that use independence scores, like HSIC in our paper. Their results indicate that, although ANMs are based on noise analysis, certain types of noise can hinder the detection of the causal direction. A study of whether LDP noise can give rise to the same effect could be an interesting direction for future research.

6.2. Other Tests

RECI defines causality using polynomial regression. We have not found any results in the local privacy setting about how the LDP noise influences polynomial regression. This could be one interesting research direction. *IGCI* uses a 1-spacing entropy estimation ([Mooij et al., 2016](#)), in which, after adding noise, the values are sorted and the average distance of neighboring values is computed. We corroborated in Figure 5 that adding noise degrades the performance of the IGCI algorithm, as shown by [Mooij et al. \(2016\)](#). *CDS* uses the variance of the conditional probability after discretizing the input into 13 equally spaced bins based on the standard deviation of the distribution. After adding noise, if the bins remained unchanged, the noise should affect the conditional probabilities, however, since the standard deviation increases, the bins change as well, and compensation may explain why the output of the CDS were not strongly affected by the LDP noise.

7. Conclusion and Future work

In this work, we investigated the challenging problem of preserving causal structure when learning over locally differentially private data. To allow the comparison between two distinct privacy notions, namely LDP and local d -privacy, we introduced a unified privacy measure based on an attacking perspective. We performed extensive experiments on both synthetic and real-world data sets comparing the privacy-utility trade-off of 9 causal discovery algorithms when applied to locally private data. Overall, our results demonstrate that locally d -private mechanisms offer a more promising approach for tackling this problem by preserving the causal structure of multidimensional data at an equivalent level of privacy. Based on the findings of this paper, there are several areas that could be explored for future work. Some potential avenues for further research include investigating the same problem on continuous data and quantifying the effect of the sample size on the variability of the output metrics. Another possible extension is exploring the effect of privatization on the parameters of the causal model. Finally, we identify the need for designing a locally private mechanism specifically for causal discovery tasks.

Acknowledgments

This work was supported by the European Research Council (ERC) project HYPATIA under the European Union’s Horizon 2020 research and innovation programme. Grant agreement n. 835294. And this work was also supported by the ELSA – European Lighthouse on Secure and Safe AI funded by the European Union. Grant agreement n. 101070617. H.H. Arcolezi has been partially supported by MIAI @ Grenoble Alpes (ANR-19-P3IA-0003).

References

- Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Hadamard response: Estimating distributions privately, efficiently, and with little communication. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 1120–1129. PMLR, 16–18 Apr 2019.
- Anish Agarwal and Rahul Singh. Causal inference with corrupted data: Measurement error, missing values, discretization, and differential privacy. *arXiv preprint arXiv:2107.02780*, 2021.
- Miguel E. Andrés, Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security - CCS '13*. ACM Press, 2013. doi: 10.1145/2508859.2516735.
- Héber H. Arcolezi, Jean-François Couchot, Bechara Al Bouna, and Xiaokui Xiao. Improving the utility of locally differentially private protocols for longitudinal and multidimensional frequency estimates. *Digital Communications and Networks*, 2022. doi: 10.1016/j.dcan.2022.07.003.
- Héber H. Arcolezi, Sébastien Gambs, Jean-François Couchot, and Catuscia Palamidessi. On the risks of collecting multidimensional data under local differential privacy. *Proc. VLDB Endow.*, 16(5):1126–1139, jan 2023. ISSN 2150-8097. doi: 10.14778/3579075.3579086.
- Rūta Binkytė-Sadauskienė, Karima Makhoulouf, Carlos Pinzón, Sami Zhioua, and Catuscia Palamidessi. Causal discovery for fairness. *arXiv preprint arXiv:2206.06685*, 2022.
- Patrick Blöbaum, Dominik Janzing, Takashi Washio, Shohei Shimizu, and Bernhard Schölkopf. Cause-effect inference by comparing regression errors. In *International Conference on Artificial Intelligence and Statistics*, pages 900–909. PMLR, 2018.
- K. Chatzikokolakis, G. Cherubin, C. Palamidessi, and C. Troncoso. Bayes security: A not so average metric. In *2023 IEEE 36th Computer Security Foundations Symposium (CSF) (CSF)*, pages 159–177, Los Alamitos, CA, USA, jul 2023. IEEE Computer Society. doi: 10.1109/CSF57540.2023.00011.
- Konstantinos Chatzikokolakis, Miguel E Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. Broadening the scope of differential privacy using metrics. In *Privacy Enhancing Technologies: 13th International Symposium, PETS 2013, Bloomington, IN, USA, July 10-12, 2013. Proceedings 13*, pages 82–102. Springer, 2013.
- David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.
- Charles J Clopper and Egon S Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934.
- Graham Cormode, Somesh Jha, Tejas Kulkarni, Ninghui Li, Divesh Srivastava, and Tianhao Wang. Privacy at scale: Local differential privacy in practice. In *Proceedings of the 2018 International Conference on Management of Data*, pages 1655–1658, 2018.

- Povilas Danušis, Dominik Janzing, Joris Mooij, Jakob Zscheischler, Bastian Steudel, Kun Zhang, and Bernhard Schölkopf. Inferring deterministic causal relations. *arXiv preprint arXiv:1203.3475*, 2012.
- Apple Differential Privacy Team. Learning with privacy at scale. In *Apple Machine Learning Journal*, volume 1. Apple, 2017.
- Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, pages 3574–3583, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Josep Domingo-Ferrer and Jordi Soria-Comas. Multi-dimensional randomized response. *IEEE Transactions on Knowledge and Data Engineering*, 34(10):4933–4946, 2022. doi: 10.1109/TKDE.2020.3045759.
- John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438. IEEE, October 2013. doi: 10.1109/focs.2013.53.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pages 265–284. Springer Berlin Heidelberg, 2006. doi: 10.1007/11681878_14.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Doris Entner and Patrik O Hoyer. On causal discovery from time series data using fci. *Probabilistic graphical models*, pages 121–128, 2010.
- Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1054–1067, New York, NY, USA, 2014. ACM. doi: 10.1145/2660267.2660348.
- José AR Fonollosa. Conditional distribution variability measures for causality detection. *Cause Effect Pairs in Machine Learning*, pages 339–347, 2019.
- Marco Gaboardi and Ryan Rogers. Local private hypothesis testing: Chi-square tests. In *International Conference on Machine Learning*, pages 1626–1635. PMLR, 2018.
- Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.
- Sivakanth Gopi, Gautam Kamath, Janardhan Kulkarni, Aleksandar Nikolov, Zhiwei Steven Wu, and Huanyu Zhang. Locally private hypothesis selection. In *Conference on Learning Theory*, pages 1785–1816. PMLR, 2020.
- Hao Han, Yeming Ma, and Wei Zhu. Galton’s family heights data revisited. *arXiv preprint arXiv:1508.02942*, 2015.

- Patrik Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems*, 21, 2008.
- Ronald C Johnson, Gerald E McClearn, Sylvia Yuen, Craig T Nagoshi, Frank M Ahern, and Robert E Cole. Galton’s data a century later. *American Psychologist*, 40(8):875, 1985.
- James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. Pate-gan: Generating synthetic data with differential privacy guarantees. In *International conference on learning representations*, 2019.
- Peter Kairouz, Keith Bonawitz, and Daniel Ramage. Discrete distribution estimation under local privacy. In *Int. Conf. on Machine Learning*, pages 2436–2444. PMLR, 2016.
- Diviyam Kalainathan, Olivier Goudet, and Ritik Dutta. Causal discovery toolbox: Uncovering causal relationships in python. *The Journal of Machine Learning Research*, 21(1):1406–1410, 2020.
- Benjamin Kap, Marharyta Aleksandrova, and Thomas Engel. The effect of noise level on the accuracy of causal discovery methods with additive noise models. In *Benelux Conference on Artificial Intelligence*, pages 120–140. Springer, 2021.
- Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 531–540. IEEE, October 2008. doi: 10.1109/FOCS.2008.27.
- Hiroaki Kikuchi. Castell: Scalable joint probability estimation of multi-dimensional data randomized with local differential privacy. *arXiv preprint arXiv:2212.01627*, 2022.
- Jack Kuipers, Polina Suter, and Giusi Moffa. Efficient sampling and structure learning of bayesian networks. *Journal of Computational and Graphical Statistics*, 31(3):639–650, 2022.
- Matt J. Kusner, Yu Sun, Karthik Sridharan, and Kilian Q. Weinberger. Private causal inference. In Arthur Gretton and Christian C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 1308–1317, Cadiz, Spain, 09–11 May 2016. PMLR.
- Trent Kyono and Mihaela Van der Schaar. Exploiting causal structure for robust model selection in unsupervised domain adaptation. *IEEE Transactions on Artificial Intelligence*, 2(6):494–507, 2021.
- Joshua R Loftus, Chris Russell, Matt J Kusner, and Ricardo Silva. Causal reasoning for algorithmic fairness. *arXiv preprint arXiv:1805.05859*, 2018.
- Pingchuan Ma, Zhenlan Ji, Qi Pang, and Shuai Wang. Noleaks: Differentially private causal discovery under functional causal model. *IEEE Transactions on Information Forensics and Security*, 17: 2324–2338, 2022. doi: 10.1109/TIFS.2022.3184263.
- Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber. Privacy: Theory meets practice on the map. In *Proceedings of the IEEE 24th International Conference on Data Engineering (ICDE)*, pages 277–286, 04 2008. doi: 10.1109/ICDE.2008.4497436.

- Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. Explainable reinforcement learning through a causal lens. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 2493–2500, 2020.
- Christopher Meek. *Graphical Models: Selecting causal and statistical models*. PhD thesis, Carnegie Mellon University, 1997.
- Joris M Mooij, Jonas Peters, Dominik Janzing, Jakob Zscheischler, and Bernhard Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *The Journal of Machine Learning Research*, 17(1):1103–1204, 2016.
- Arvind Narayanan and Vitaly Shmatikov. De-anonymizing social networks. *Proceedings - IEEE Symposium on Security and Privacy*, 04 2009. doi: 10.1109/SP.2009.22.
- Ana Rita Nogueira, João Gama, and Carlos Abreu Ferreira. Causal discovery in machine learning: Theories and applications. *Journal of Dynamics & Games*, 8(3):203, 2021.
- Yuki Ohnishi and Jordan Awan. Locally private causal inference. *arXiv preprint arXiv:2301.01616*, 2023.
- Carlos Pinzón, Camilo Rocha, and Jorge Finke. An approach to optimal discretization of continuous real random variables with application to machine learning, 2020.
- Joseph Ramsey, Madelyn Glymour, Ruben Sanchez-Romero, and Clark Glymour. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International journal of data science and analytics*, 3:121–129, 2017.
- Jason Reed and Benjamin Pierce. Distance makes the types grow stronger a calculus for differential privacy. *Sigplan Notices - SIGPLAN*, 45:157–168, 09 2010. doi: 10.1145/1932681.1863568.
- Jonathan G Richens, Ciarán M Lee, and Saurabh Johri. Improving the accuracy of medical diagnosis with causal machine learning. *Nature communications*, 11(1):3923, 2020.
- Felix L. Rios, Giusi Moffa, and Jack Kuipers. Benchpress: a scalable and versatile workflow for benchmarking structure learning algorithms for graphical models, 2021.
- Karen Sachs, Omar Perez, Dana Pe’er, Douglas A Lauffenburger, and Garry P Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):523–529, 2005.
- Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021.
- Peter Spirtes and Clark Glymour. An algorithm for fast recovery of sparse causal graphs. *Social science computer review*, 9(1):62–72, 1991.
- Shruti Tople, Amit Sharma, and Aditya Nori. Alleviating privacy attacks via causal learning. In *International Conference on Machine Learning*, pages 9537–9547. PMLR, 2020.

Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65(1):31–78, 2006.

Lun Wang, Qi Pang, and Dawn Song. Towards practical differentially private causal graph discovery. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 5516–5526. Curran Associates, Inc., 2020.

Tianhao Wang, Jeremiah Blocki, Ninghui Li, and Somesh Jha. Locally differentially private protocols for frequency estimation. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 729–745, Vancouver, BC, August 2017. USENIX Association. ISBN 978-1-931971-40-9.

Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, March 1965. doi: 10.1080/01621459.1965.10480775.

Depeng Xu, Shuhan Yuan, and Xintao Wu. Differential privacy preserving causal graph discovery. In *2017 IEEE Symposium on Privacy-Aware Computing (PAC)*, pages 60–71, 2017. doi: 10.1109/PAC.2017.24.

Jun Zhang, Graham Cormode, Cecilia M. Procopiuc, Divesh Srivastava, and Xiaokui Xiao. Privbayes: Private data release via bayesian networks. *ACM Trans. Database Syst.*, 42(4), oct 2017. ISSN 0362-5915. doi: 10.1145/3134428.

Appendix A. Privacy Mechanisms

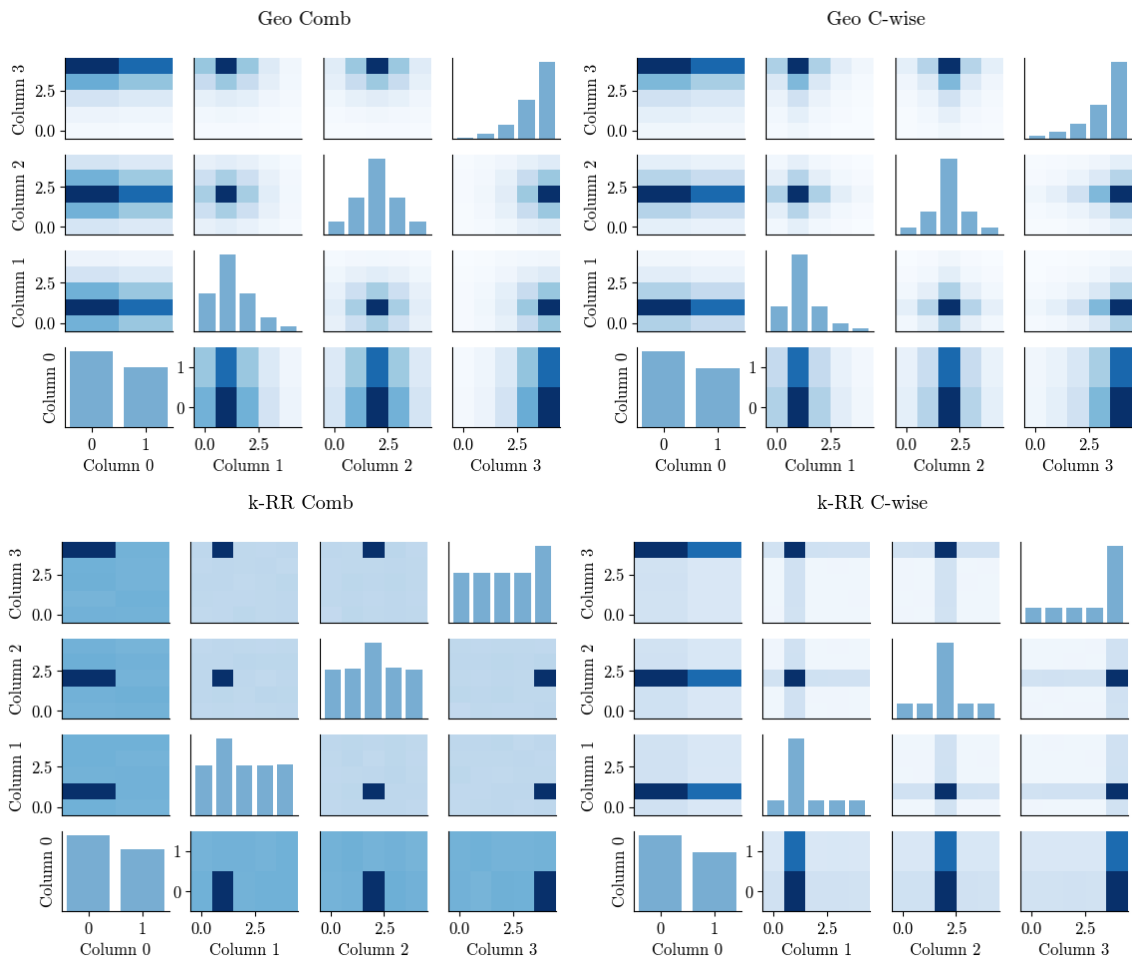


Figure 6: Illustration of 4 multidimensional mechanisms discussed in this paper: 4D bounded Geometric, 4x1D bounded Geometric, 4D k -RR and 4x1D k -RR.

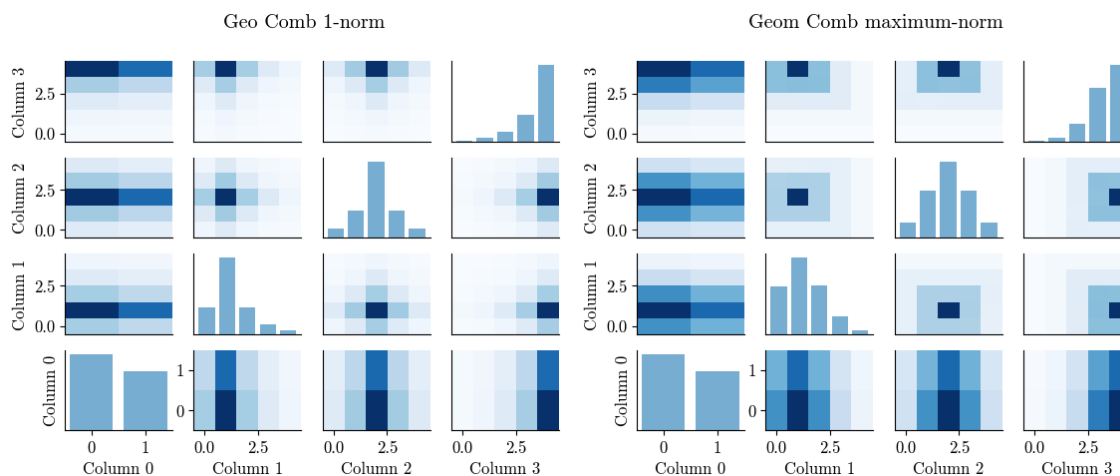


Figure 7: Comparison between Manhattan ($p = 1$) and Chebyshev ($p = \infty$) distances for bounded geometric mechanisms. Refer to Figure 6 for euclidean ($p = 2$).

Appendix B. The Data Sets

The Sachs data set measures the expression levels of various proteins and phospholipids within human cells. It was originally generated by [Sachs et al. \(2005\)](#). The data set consists of 11 variables and 902 samples. Sachs is a popular benchmarking data set in causal discovery, because of availability of the ground truth causal structure.

Human Stature data is a classic historical data set collected by the statistician Francis Galton and first used for regression analysis ([Johnson et al., 1985](#)). Later it has been re-used as one of the benchmark data sets for causal discovery. The data set consists of four variables: father height, mother height, gender and child height, and has 898 samples. We remove the binary gender variable for our experiments. We do it, because when applied to binary data, geometric noise becomes equivalent to k -RR method.

Synth10 and Synth5 are synthetic data sets with 10 and 5 nodes respectively. The background structure DAG is generated randomly using the benchpress framework ([Rios et al., 2021](#)). We specify the number of nodes and the maximum number of parents for each node. The data is generated to using a generation process compatible with the underlying structure DAG.

Appendix C. The Algorithms

The Peter and Clark (**PC**) ([Spirtes and Glymour, 1991](#)) algorithm is a constraint-based method with two primary stages. The initial stage, known as “adjacency search”, involves identifying the undirected skeleton of the Directed Acyclic Graph (DAG). The second stage focuses on estimating a completed partially directed acyclic graph (CPDAG). PC can be applied to linear, Gaussian data (the Fisher Z test), discrete multinomial data (the Chi Square test) and mixed multinomial/Gaussian data (the Conditional Gaussian test). PC uses an alpha parameter which is a cutoff, which signifies the threshold at which test results are considered indicative of dependence in a statistical test of independence, typically defaults to 0.05. When using a higher alpha value, PC leads to a sparser

graph. In other words, a higher alpha makes the test more stringent, and it requires stronger evidence to conclude that variables are dependent, resulting in fewer edges in the graphical model.

The **FCI** (Fast Causal Inference) (Entner and Hoyer, 2010) algorithm is a constraint-based method designed to work with sample data, and it can also consider optional background knowledge. In the large sample limit, FCI provides an equivalence class of Conditional Bayesian Networks (CBNs) that encompass the set of conditional independence relations believed to be valid in the population, even when there are hidden confounding variables. However, FCI has limitations and is most suitable for data sets with several thousand variables. When applied to realistic sample sizes, it can be inaccurate in determining both adjacencies and orientations. FCI consists of two phases: the adjacency phase and the orientation phase. During the adjacency phase, the algorithm begins with a complete undirected graph and then conducts a series of conditional independence tests. These tests lead to the removal of edges between pairs of variables that are determined to be independent, given some subset of the observed variables. The conditioning sets that result in the removal of an edge are stored. By the end of the adjacency phase, the undirected graph correctly represents the set of adjacencies among variables, but all edges remain unoriented. FCI then proceeds to the orientation phase, where it uses the stored conditioning sets to orient as many edges as possible, adding directionality to the graph.

FGES (Ramsey et al., 2017) is an enhanced and parallelized variant of the Greedy Equivalence Search (GES) algorithm, initially developed by Meek (1997) and later studied by Chickering (2002). GES is a Bayesian algorithm that uses a heuristic approach to explore the space of Conditional Bayesian Networks (CBNs) and identify the model with the highest Bayesian score. Specifically, GES commences its search with an empty graph and proceeds with a forward stepping search, where it adds edges between nodes to maximize the Bayesian score. This process continues until no further single edge addition improves the score. Subsequently, it performs a backward stepping search, eliminating edges until no single edge removal can enhance the score. These algorithms are capable of handling both continuous data, utilizing the Structural Equation Modeling Bayesian Information Criterion (SEM BIC) score, and discrete data, making use of the Bayesian Dirichlet equivalent uniform (BDeu) score. FGES takes the penalty discount parameter. Higher penalty discount yield sparser graphs.

Iterative MCMC (Kuipers et al., 2022) is a hybrid optimization technique based on Markov chain Monte Carlo (MCMC) methods. The algorithm's initial step involves generating a skeleton, obtained through the Greedy Equivalence Search (GES) algorithm. Subsequently, it conducts a score-based search within the space defined by this initial skeleton, exploring various Directed Acyclic Graphs (DAGs).

The Max-min hill-climbing (**MMHC**) (Tsamardinos et al., 2006) method is a hybrid approach that follows a two-step process. Firstly, it estimates the skeleton of a Directed Acyclic Graph (DAG) using an algorithm known as Max-Min Parents and Children. Then, it applies a greedy hill-climbing search to determine the orientation of edges within the graph based on Bayesian scoring. MMHC is particularly suitable for domains with a high number of dimensions.

RECI (Regression Error based Causal Inference) (Blöbaum et al., 2018) addresses non-deterministic and nonlinear relations and allows dependency between cause and noise. The algorithm's key idea is to fit regression models in both possible directions and compare the MSE. No independence tests are used, but the assumptions on the model depend on the regressor used for the model. In our experiments we used a polynomial regressor with degree 3 after rescaling to $[0, 1]$.

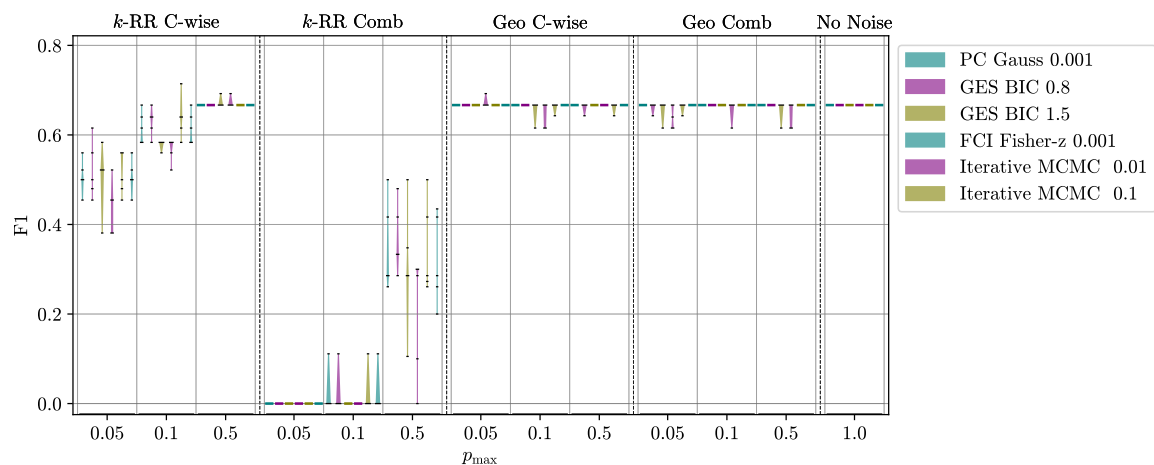


Figure 8: Sachs data, F1.

IGCI (Information Geometric Causal Inference) (Danušis et al., 2012) is a pairwise causal discovery model that able to determine the causal relationship in a deterministic setting $Y = f(X)$ (where f is invertible), under the ‘independence assumption’ $Cov[\log f', p_X] = 0$. In our experiments we have used a Gaussian reference measure² and the sp1 or “1-spacing” method for entropy estimation used in Mooij et al. (2016).

CDS (Conditional Distribution Similarity Statistic) (Fonollosa, 2019) first normalizes the conditional distribution $P(Y|X = x)$ (for all x) to have zero mean and unit variance, then quantizes it. In our experiments as conditional distribution variability measure we used standard deviation of the preprocessed conditional distributions. The lower the standard deviation, the more likely the pair to be $X \rightarrow Y$.

ANM (Hoyer et al., 2008) assumes that $Y = f(X) + E$, where f is nonlinear. The causal inference bases itself on the independence between X and E . We used a Gaussian process regression for the prediction and normalized HSIC for the evaluation of the causal direction.

Appendix D. Additional Experiments

We perform experiments using real and synthetic data. Data sets are distinguished into two main groups. The first category is pairwise data, which have two variables A and B where A causes B or B causes A . The task is to determine the causal direction between the two variables. The second category is the data that has more than two variables. The task here is to determine the causal structure (the skeleton) and the causal direction between the pairs within this structure.

D.1. F1 Score results Sachs data set

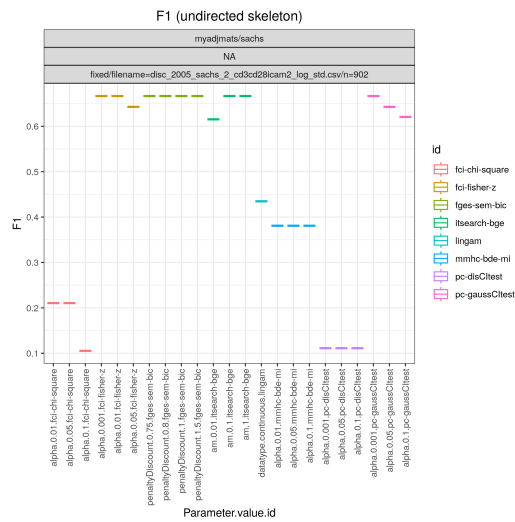


Figure 9: F1 Scores on the Sachs data set. Discretized, no noise.

2. Our experiments with the uniform reference measure produced almost identical results, thus we exclude it from this paper.

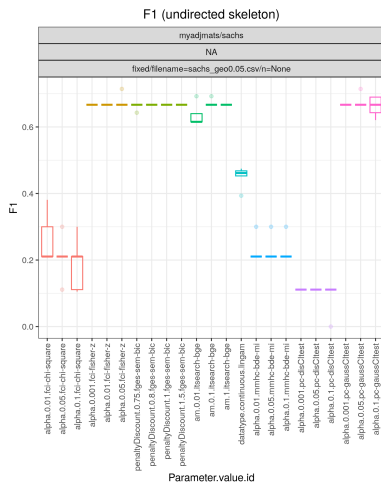


Figure 10: Sachs data, Geo C-wise mechanism, max probability 0.05.

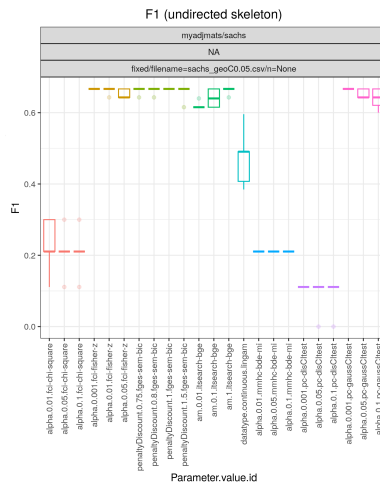


Figure 11: Sachs data, Geo Comb mechanism, max probability 0.05.

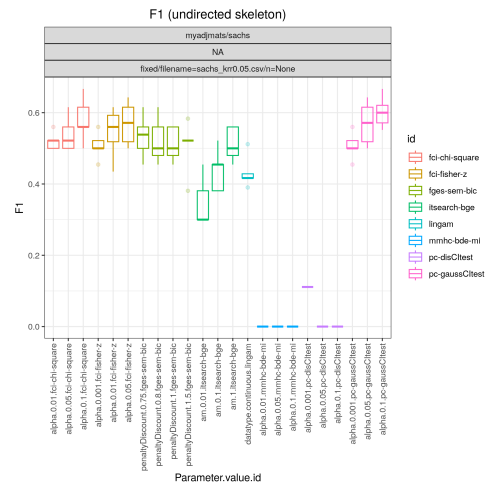


Figure 12: Sachs data, k -RR C-wise mechanism, max probability 0.05.

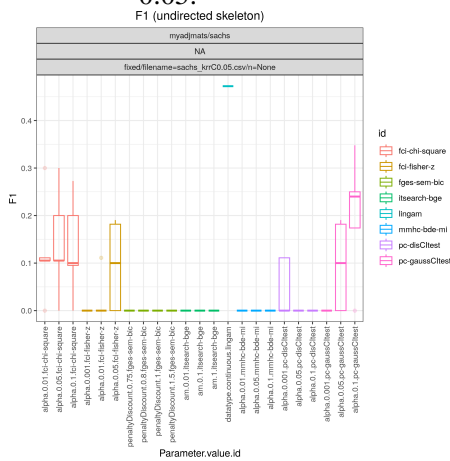


Figure 13: Sachs data, k -RR Comb mechanism, max probability 0.05.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

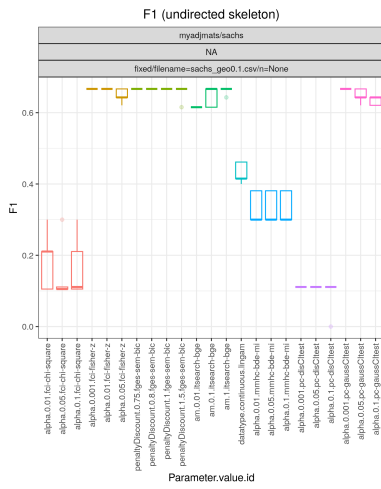


Figure 14: Sachs data, Geo C-wise mechanism, max probability 0.1.

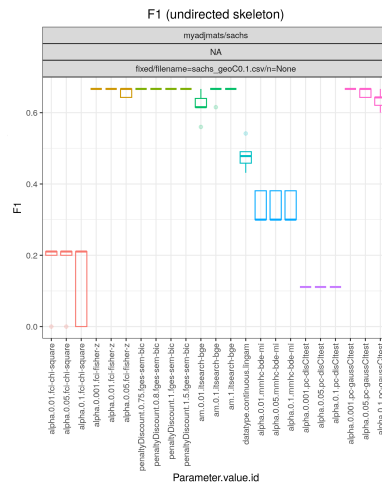


Figure 15: Sachs data, Geo Comb mechanism, max probability 0.1.

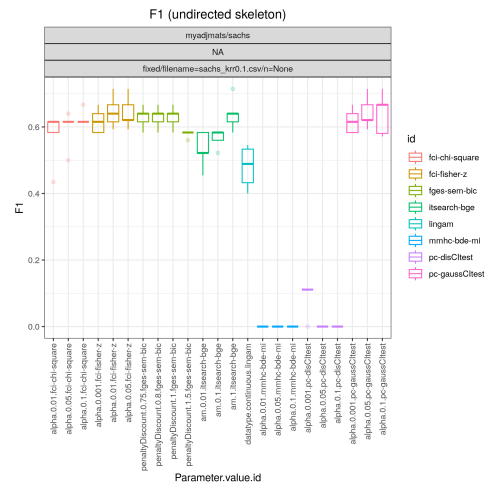


Figure 16: Sachs data, k -RR C-wise mechanism, max probability 0.1.

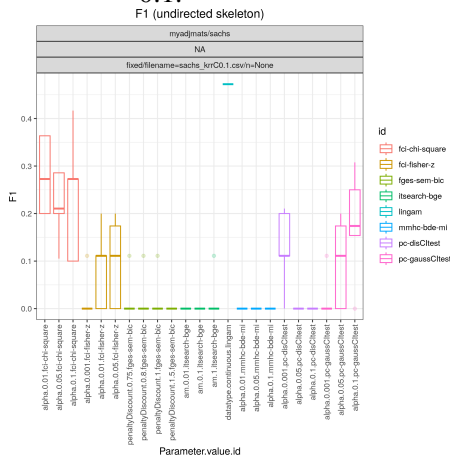


Figure 17: Sachs data, k -RR Comb mechanism, max probability 0.1.

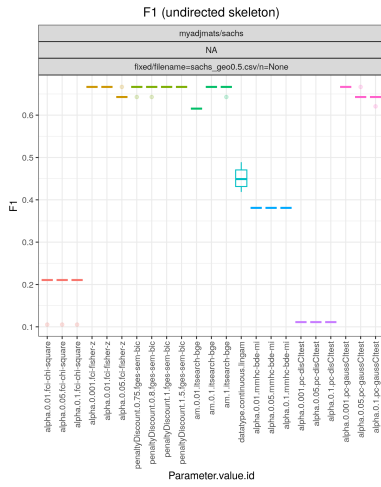


Figure 18: Sachs data, Geo C-wise mechanism, max probability 0.5.

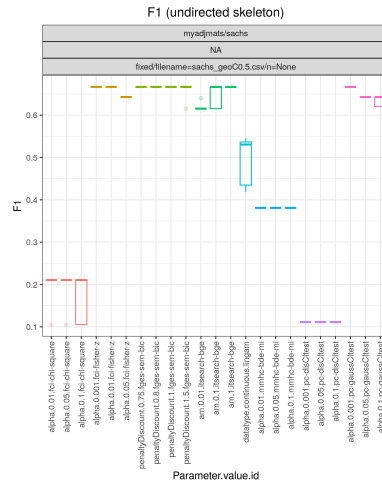


Figure 19: Sachs data, Geo Comb mechanism, max probability 0.5.

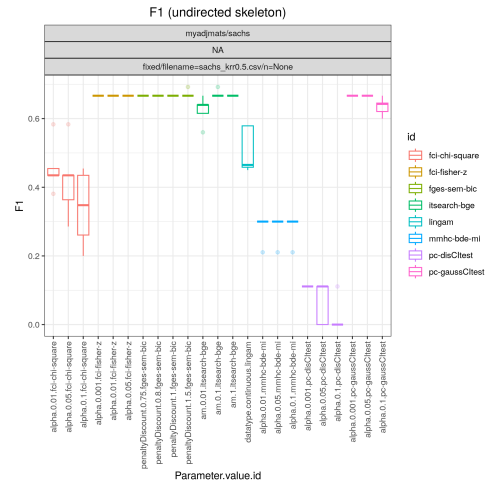


Figure 20: Sachs data, *k*-RR C-wise mechanism, max probability 0.5.

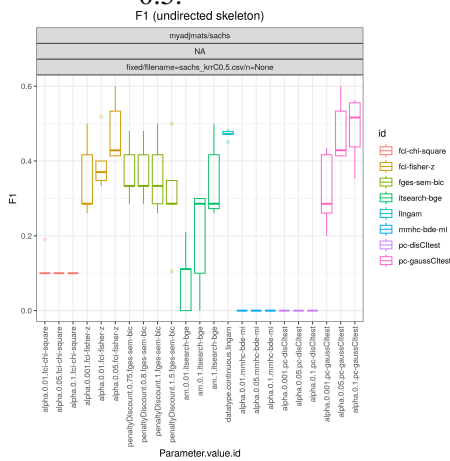


Figure 21: Sachs data, *k*-RR Comb mechanism, max probability 0.5.

D.2. SHD Score results Sachs data set

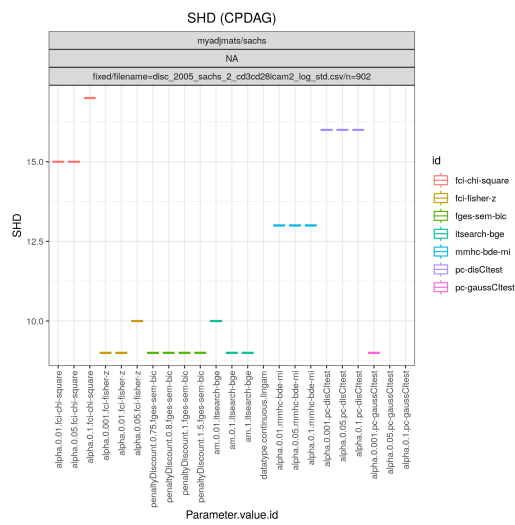


Figure 22: SHD Scores on the Sachs data set. Discretized, no noise.

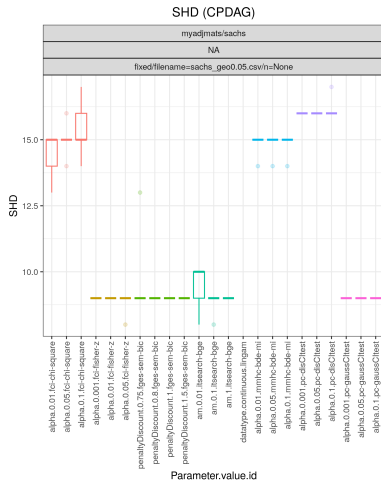


Figure 23: Sachs data, Geo C-wise mechanism, max probability 0.05.

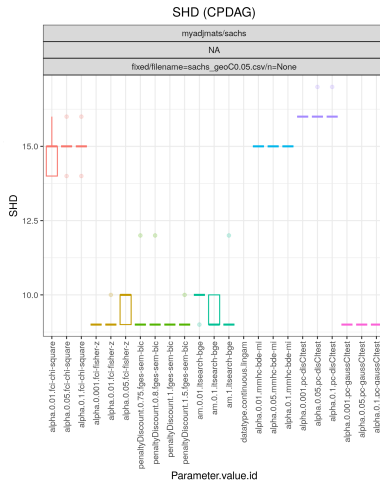


Figure 24: Sachs data, Geo Comb mechanism, max probability 0.05.

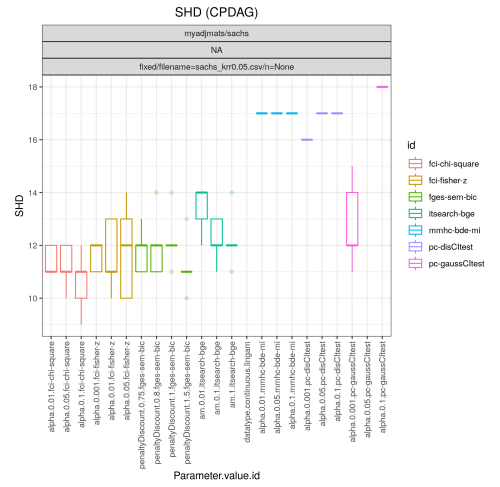


Figure 25: Sachs data, k -RR C-wise mechanism, max probability 0.05.

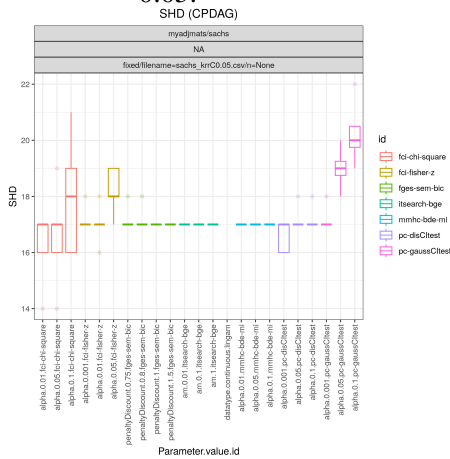


Figure 26: Sachs data, k -RR Comb mechanism, max probability 0.05.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

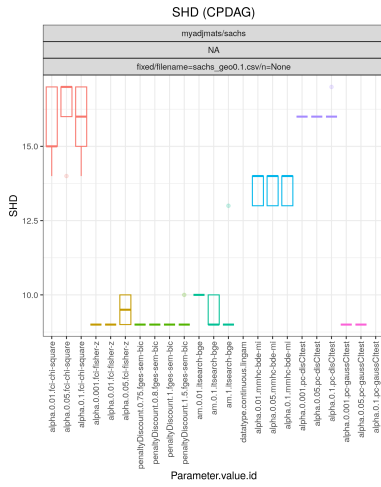


Figure 27: Sachs data, Geo C-wise mechanism, max probability 0.1.

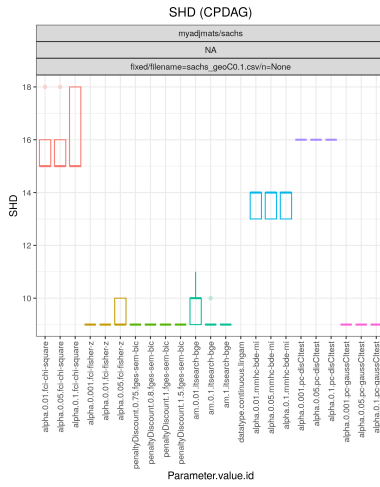


Figure 28: Sachs data, Geo Comb mechanism, max probability 0.1.

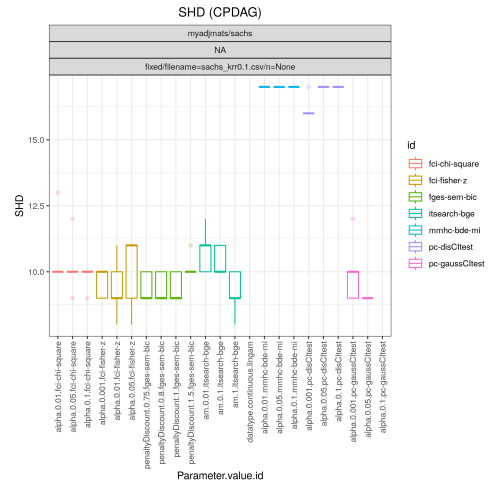


Figure 29: Sachs data, k -RR C-wise mechanism, max probability 0.1.

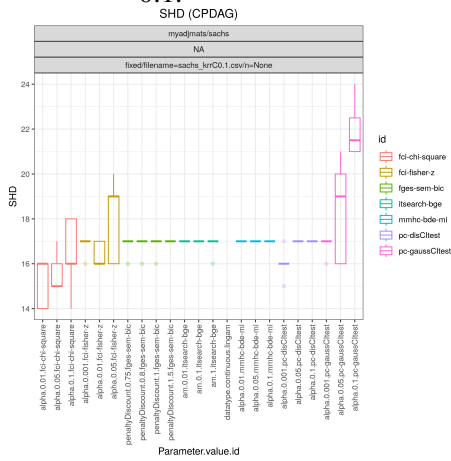


Figure 30: Sachs data, k -RR Comb mechanism, max probability 0.1.

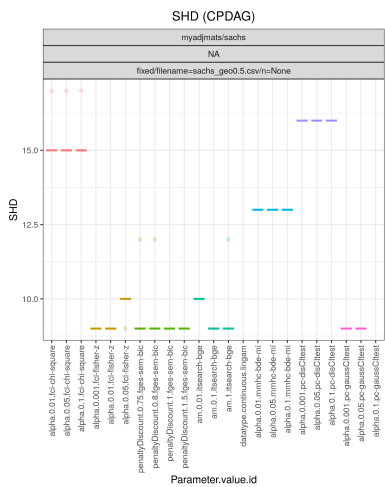


Figure 31: Sachs data, Geo C-wise mechanism, max probability 0.5.

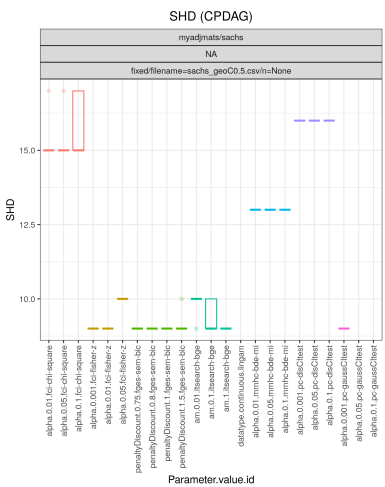


Figure 32: Sachs data, Geo Comb mechanism, max probability 0.5.

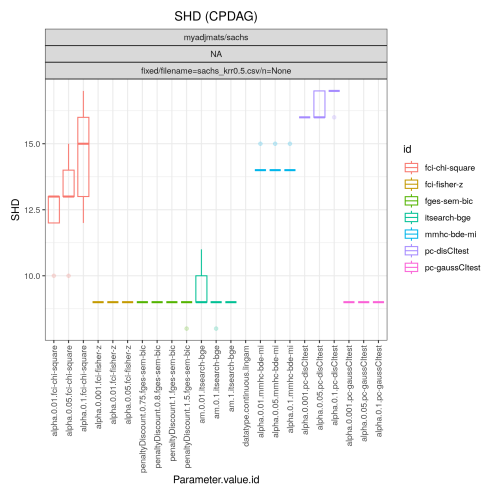


Figure 33: Sachs data, k -RR C-wise mechanism, max probability 0.5.

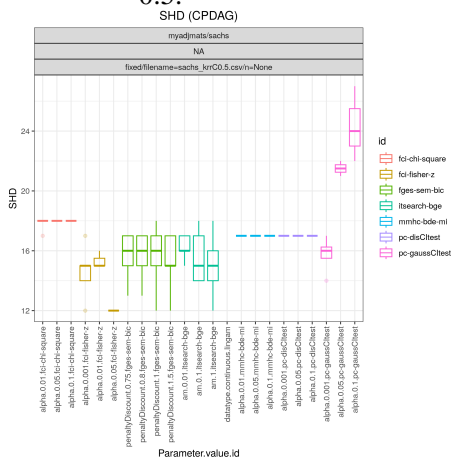


Figure 34: Sachs data, k -RR Comb mechanism, max probability 0.5.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

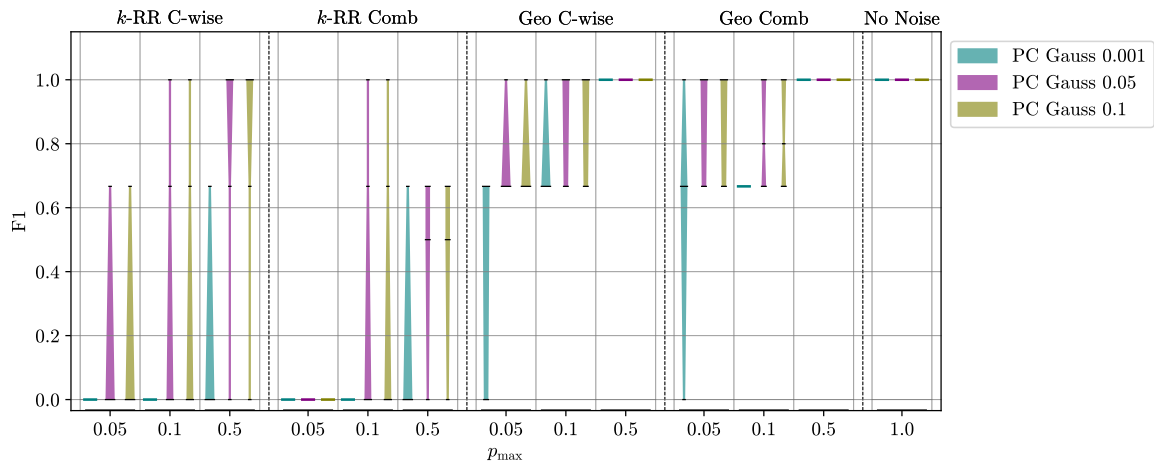


Figure 35: Human Stature data, F1.

D.3. F1 Score results Human Stature data set

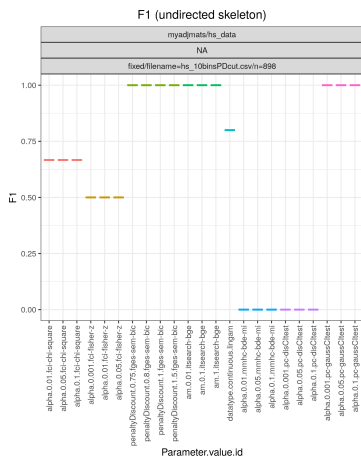


Figure 36: F1 Scores on the Human Stature data set. Discretized, no noise.

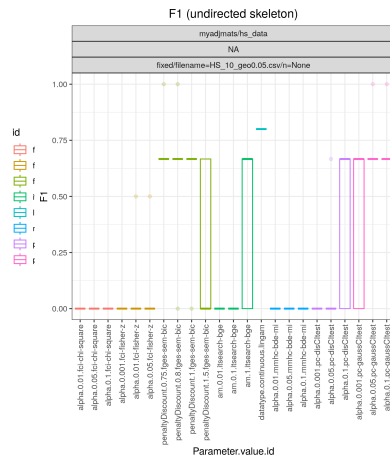


Figure 37: Human Stature data, Geo C-wise mechanism, max probability 0.05.

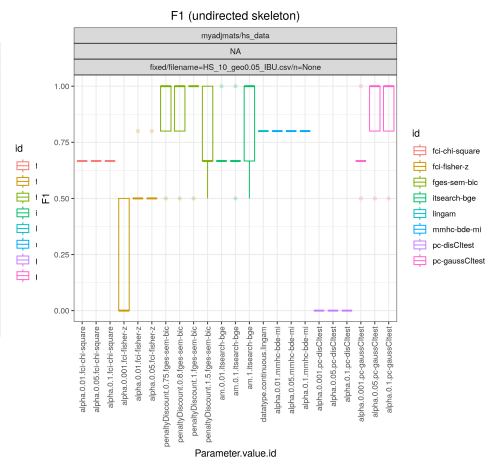


Figure 38: Human Stature data, Geo C-wise IBU mechanism, max probability 0.05.

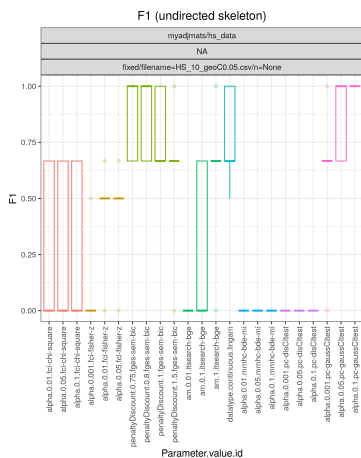


Figure 39: Human Stature data, Geo Comb mechanism, max probability 0.05.

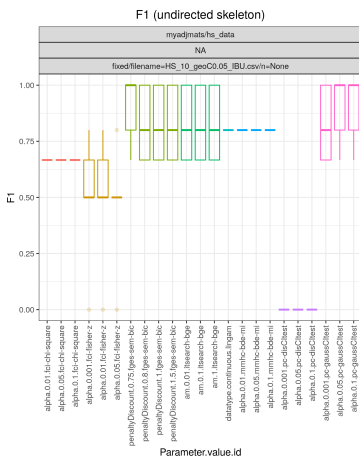


Figure 40: Human Stature data, Geo Comb IBU mechanism, max probability 0.05.

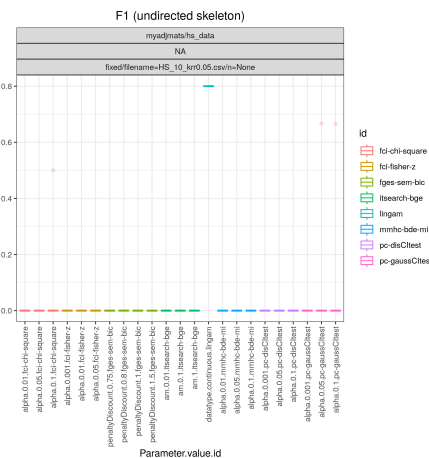


Figure 41: Human Stature data, k -RR C-wise mechanism, max probability 0.05.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

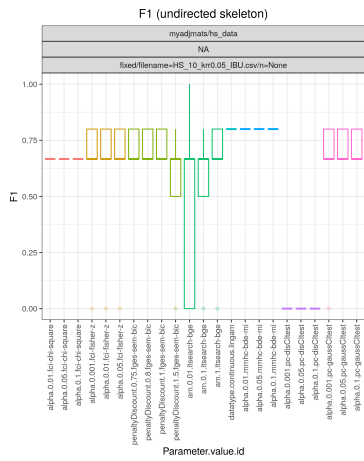


Figure 42: Human Stature data, k -RR C-wise IBU mechanism, max probability 0.05.

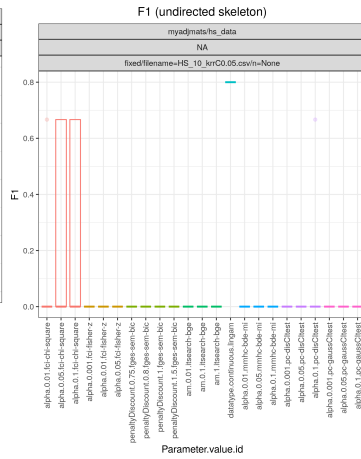


Figure 43: Human Stature data, k -RR Comb mechanism, max probability 0.05.

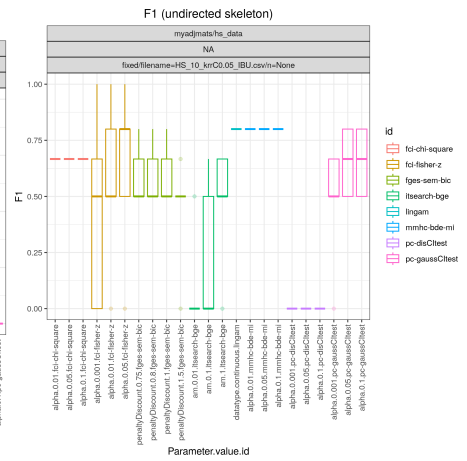


Figure 44: Human Stature data, k -RR Comb IBU mechanism, max probability 0.05.

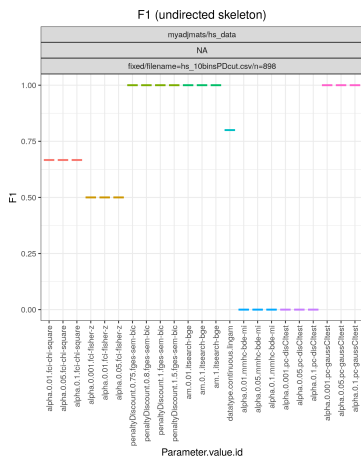


Figure 45: F1 Scores on the Human Stature data set. Discretized, no noise.

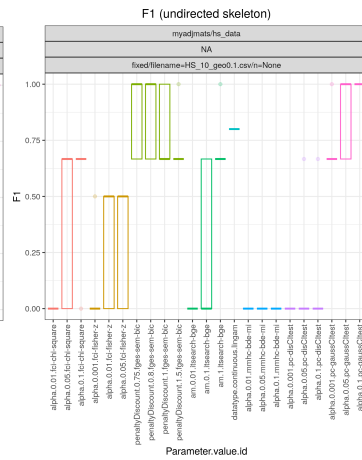


Figure 46: Human Stature data, Geo C-wise mechanism, max probability 0.1.

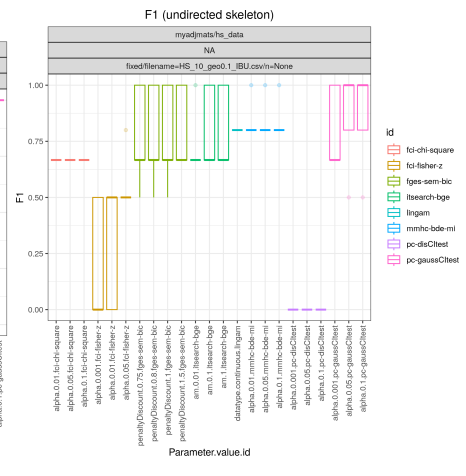


Figure 47: Human Stature data, Geo C-wise IBU mechanism, max probability 0.1.

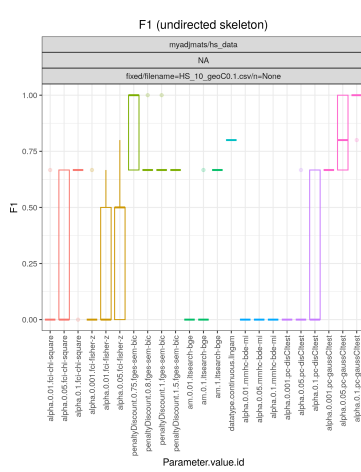


Figure 48: Human Stature data, Geo Comb mechanism, max probability 0.1.

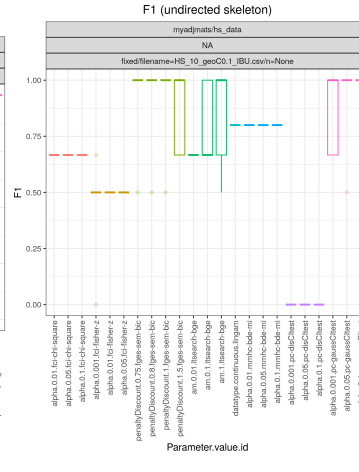


Figure 49: Human Stature data, Geo Comb IBU mechanism, max probability 0.1.

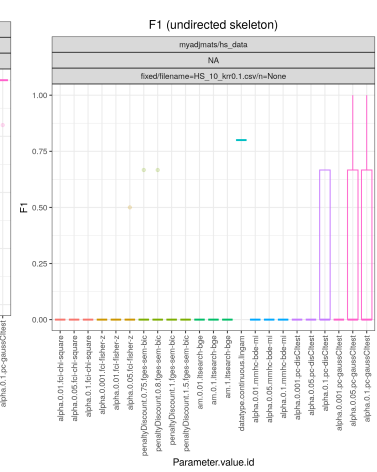


Figure 50: Human Stature data, k -RR C-wise mechanism, max probability 0.1.

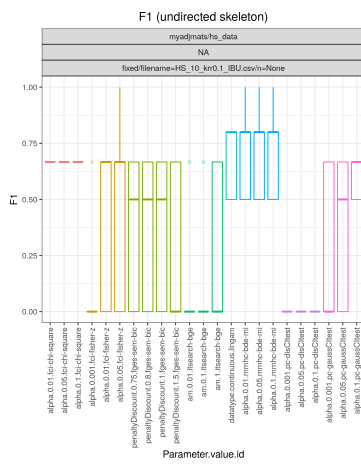


Figure 51: Human Stature data, k -RR C-wise IBU mechanism, max probability 0.1.

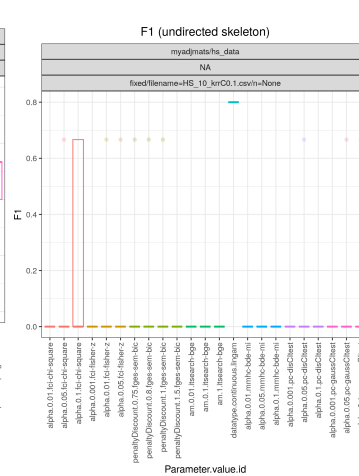


Figure 52: Human Stature data, k -RR Comb mechanism, max probability 0.1.

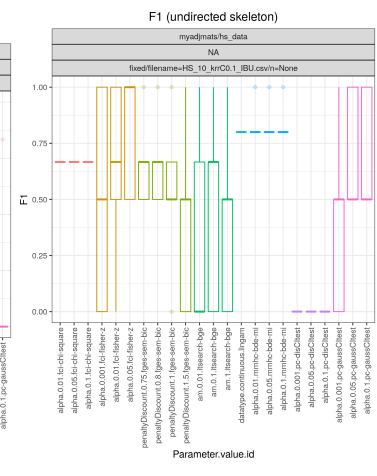


Figure 53: Human Stature data, k -RR Comb IBU mechanism, max probability 0.1.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

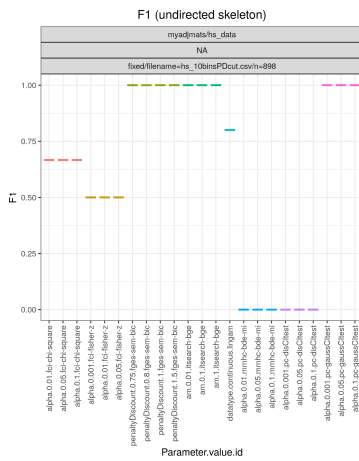


Figure 54: F1 Scores on the Human Stature data set. Discretized, no noise.

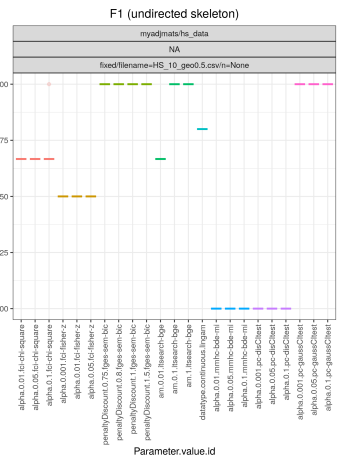


Figure 55: Human Stature data, Geo C-wise mechanism, max probability 0.5.

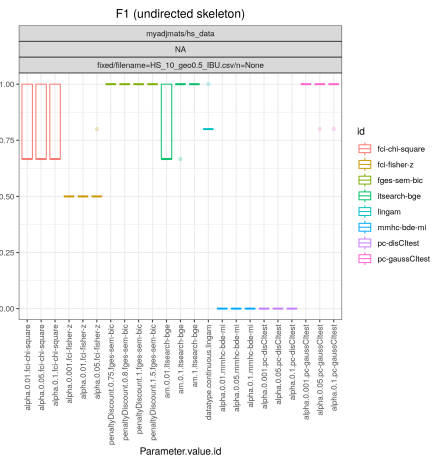


Figure 56: Human Stature data, Geo C-wise IBU mechanism, max probability 0.5.

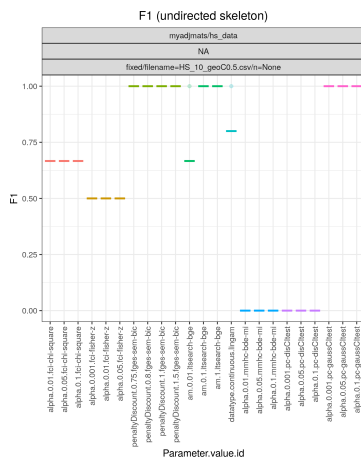


Figure 57: Human Stature data, Geo Comb mechanism, max probability 0.5.

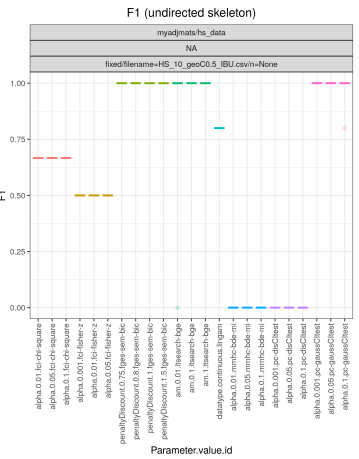


Figure 58: Human Stature data, Geo Comb IBU mechanism, max probability 0.5.

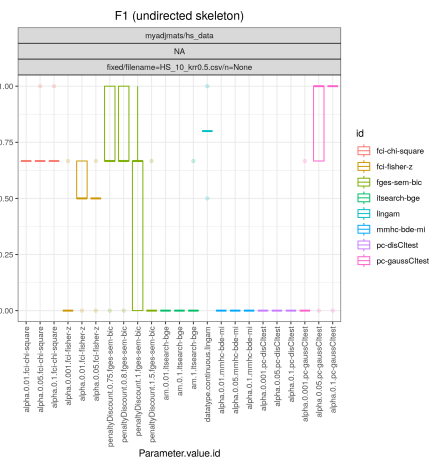


Figure 59: Human Stature data, k-RR C-wise mechanism, max probability 0.5.

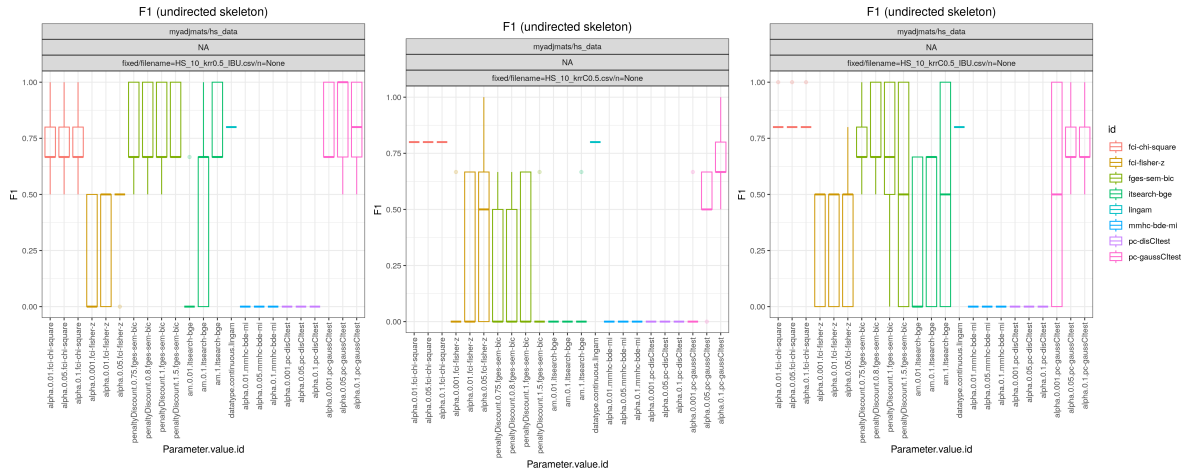


Figure 60: Human Stature data, k -RR C-wise IBU mechanism, max probability 0.5.

Figure 61: Human Stature data, k -RR Comb mechanism, max probability 0.5.

Figure 62: Human Stature data, k -RR Comb IBU mechanism, max probability 0.5.

D.4. SHD Score results Human Stature data set

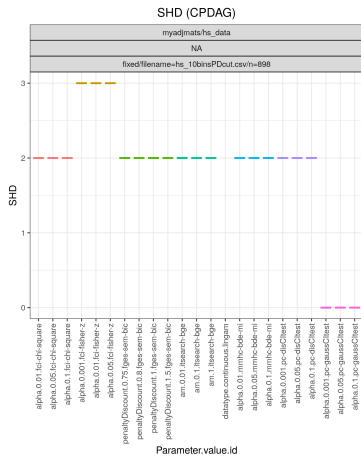


Figure 63: SHD Scores on the Human Stature data set. Discretized, no noise.

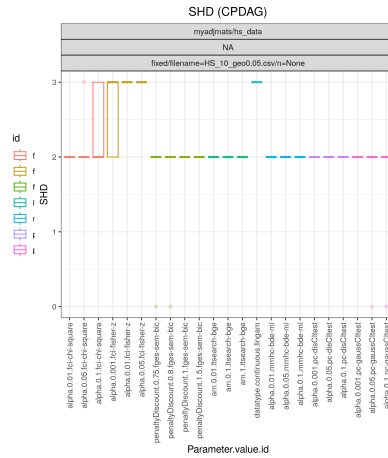


Figure 64: Human Stature data, Geo C-wise mechanism, max probability 0.05.

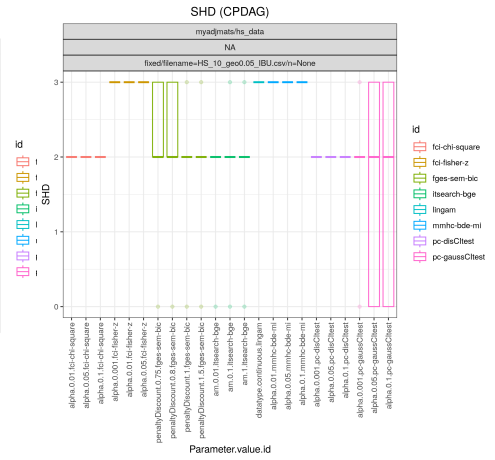


Figure 65: Human Stature data, Geo C-wise IBU mechanism, max probability 0.05.

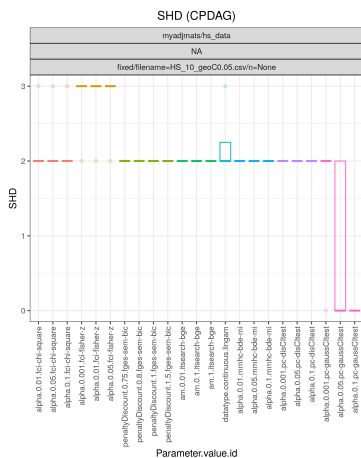


Figure 66: Human Stature data, Geo Comb mechanism, max probability 0.05.

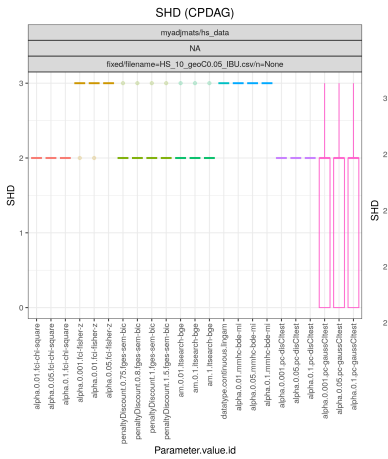


Figure 67: Human Stature data, Geo Comb IBU mechanism, max probability 0.05.

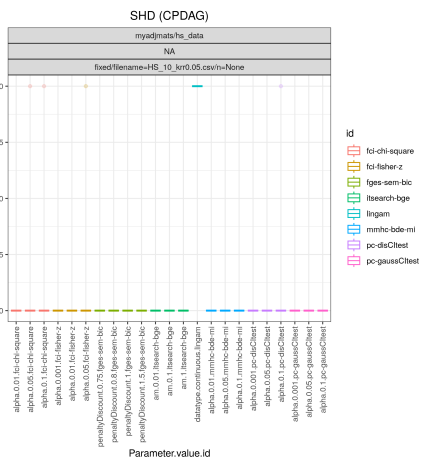


Figure 68: Human Stature data, k-RR C-wise mechanism, max probability 0.05.

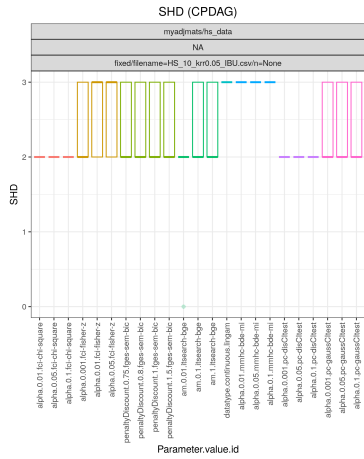


Figure 69: Human Stature data, k -RR C-wise IBU mechanism, max probability 0.05.

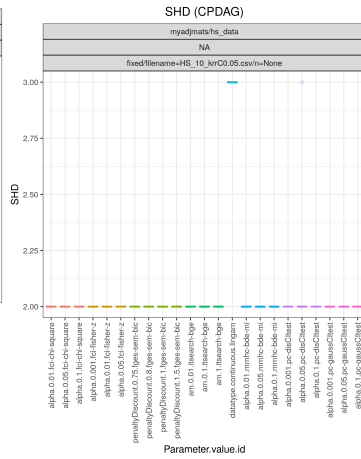


Figure 70: Human Stature data, k -RR Comb mechanism, max probability 0.05.

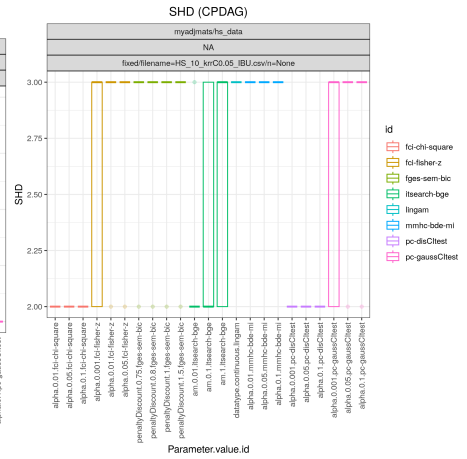


Figure 71: Human Stature data, k -RR Comb IBU mechanism, max probability 0.05.

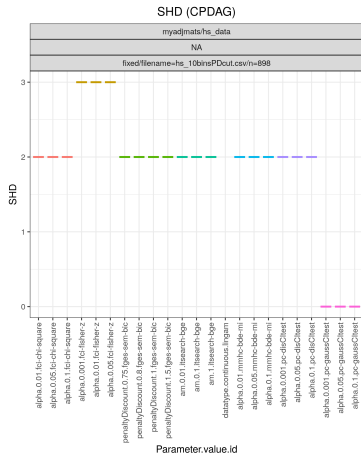


Figure 72: F1 Scores on the Figure 73: Human Stature data set. Discretized, no noise.

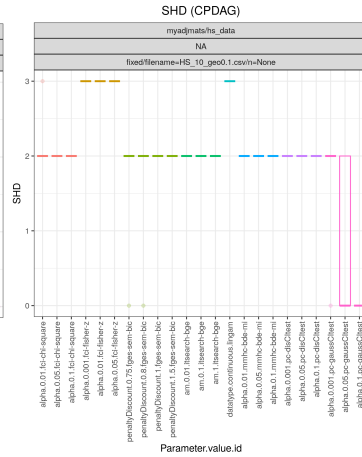


Figure 73: Human Stature data, Geo C-wise mechanism, max probability 0.1.

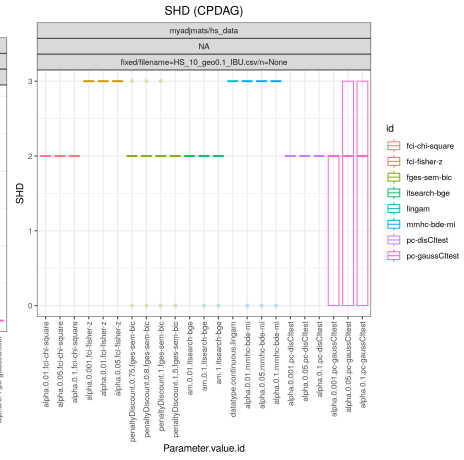


Figure 74: Human Stature data, Geo C-wise IBU mechanism, max probability 0.1.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

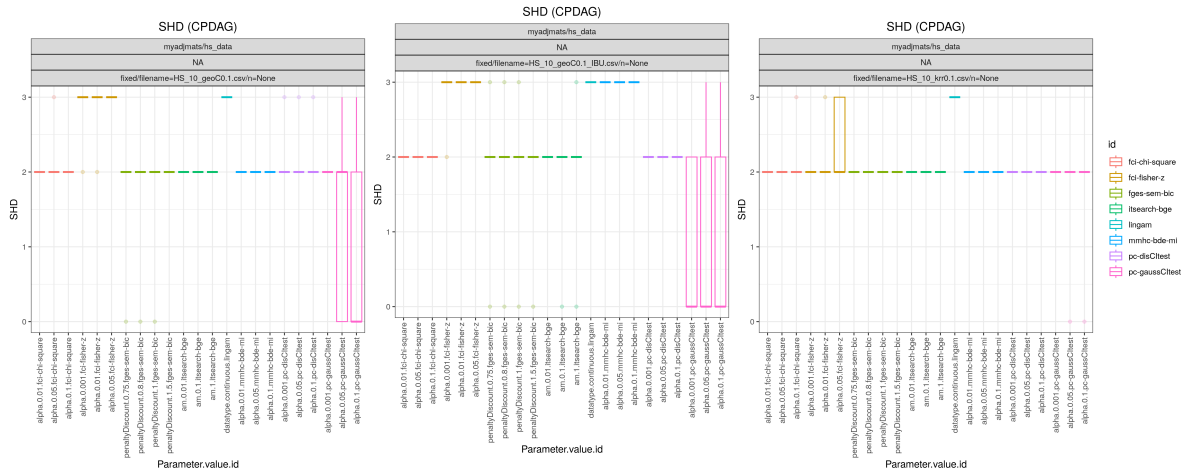


Figure 75: Human Stature data, Geo Comb mechanism, max probability 0.1.

Figure 76: Human Stature data, Geo Comb IBU mechanism, max probability 0.1.

Figure 77: Human Stature data, k -RR C-wise mechanism, max probability 0.1.

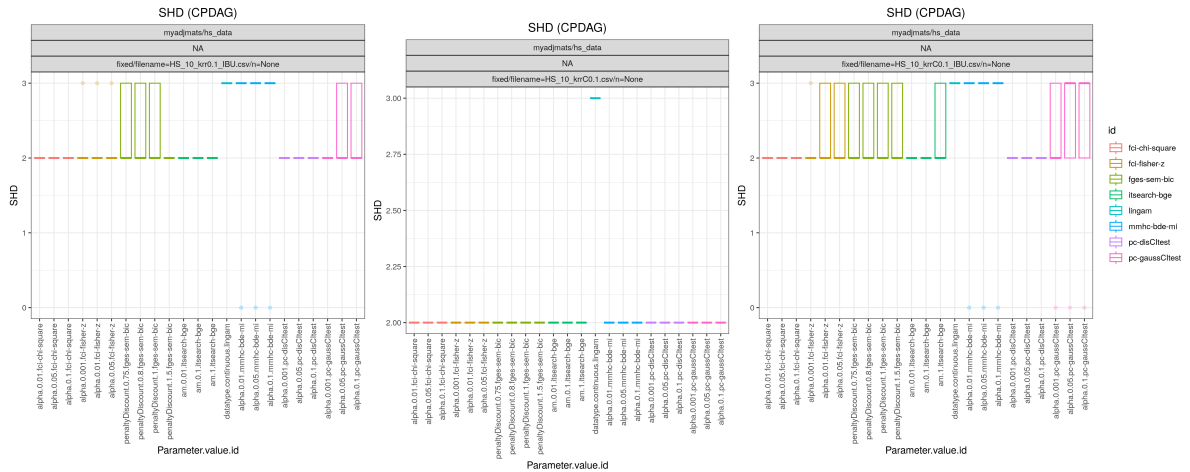


Figure 78: Human Stature data, k -RR C-wise IBU mechanism, max probability 0.1.

Figure 79: Human Stature data, k -RR Comb mechanism, max probability 0.1.

Figure 80: Human Stature data, k -RR Comb IBU mechanism, max probability 0.1.

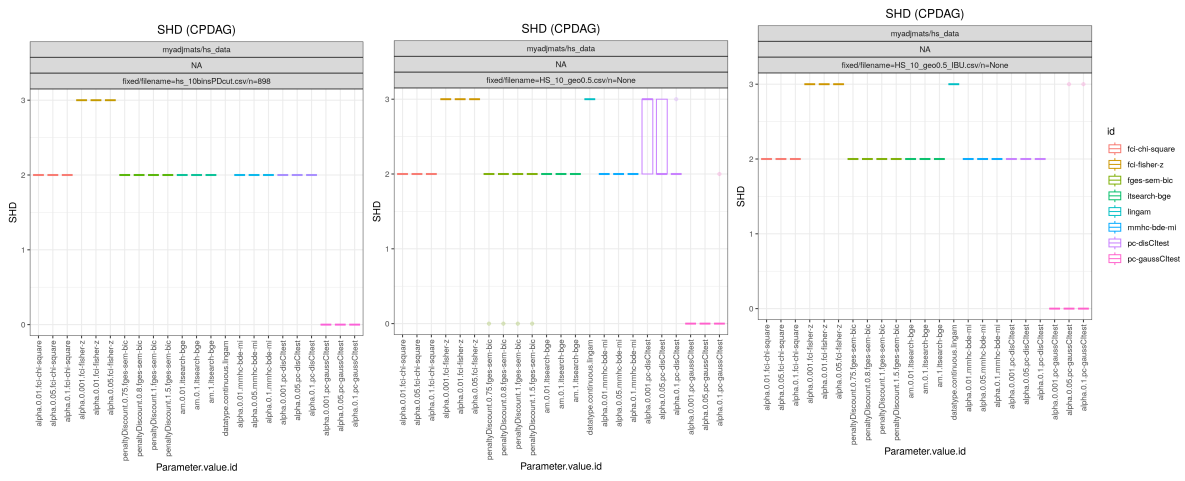


Figure 81: SHD Scores on the Human Stature data set. Discretized, no noise.

Figure 82: Human Stature data, Geo C-wise mechanism, max probability 0.5.

Figure 83: Human Stature data, Geo C-wise IBU mechanism, max probability 0.5.

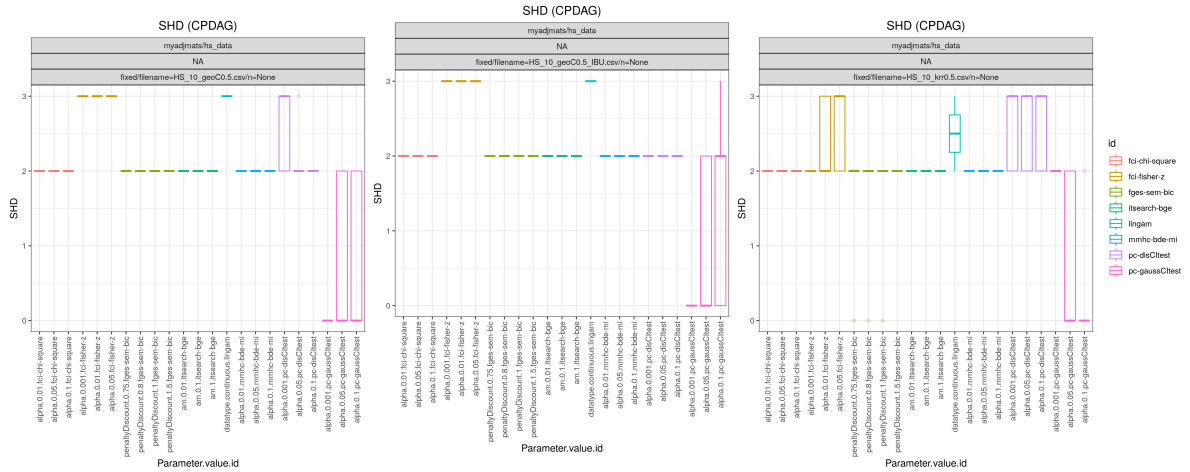


Figure 84: Human Stature data, Geo Comb mechanism, max probability 0.5.

Figure 85: Human Stature data, Geo Comb IBU mechanism, max probability 0.5.

Figure 86: Human Stature data, k -RR C-wise mechanism, max probability 0.5.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

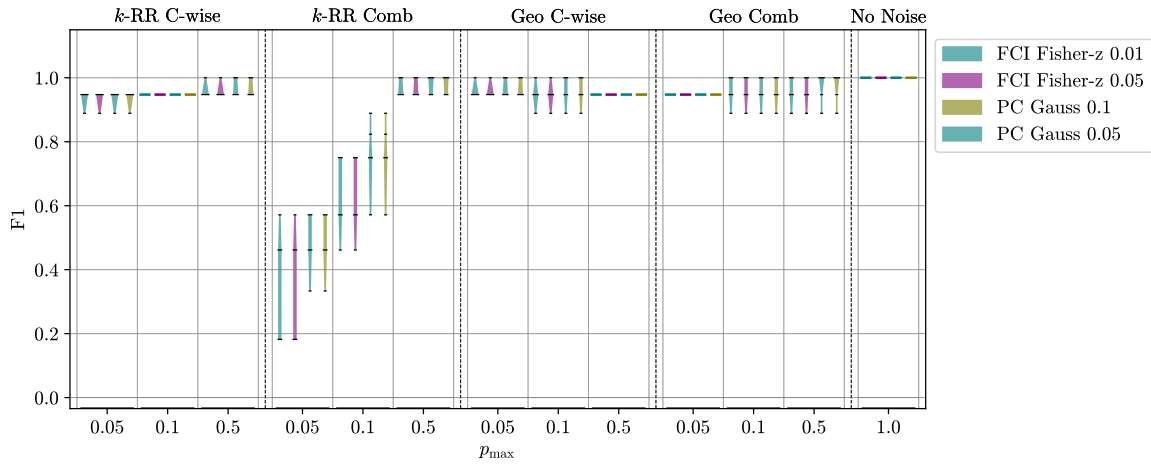


Figure 90: Synthetic data, 5 nodes, F1.

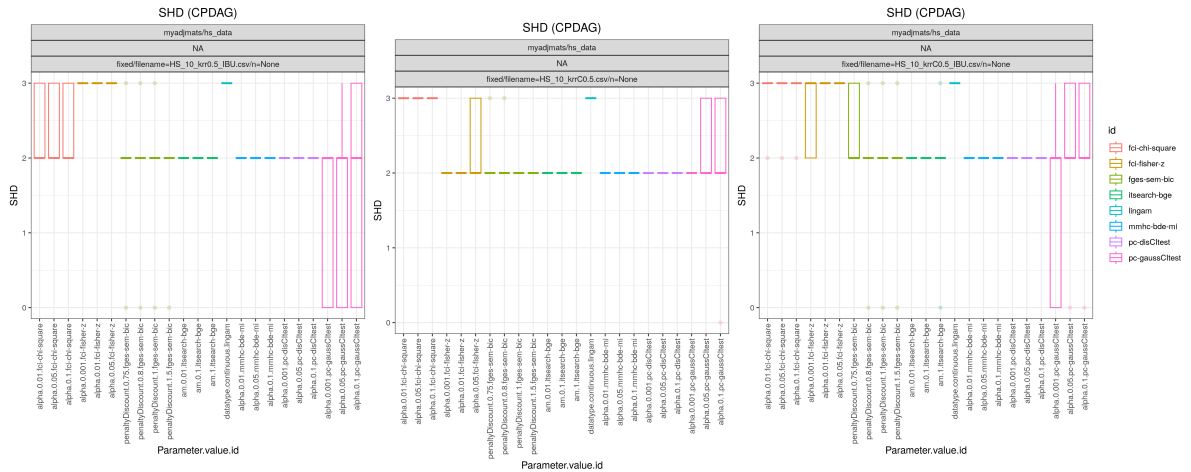


Figure 87: Human Stature data, k -RR C-wise IBU mechanism, max probability 0.5.

Figure 88: Human Stature data, k -RR Comb mechanism, max probability 0.5.

Figure 89: Human Stature data, k -RR Comb IBU mechanism, max probability 0.5.

D.5. F1 Score results Synthetic 5 nodes data set

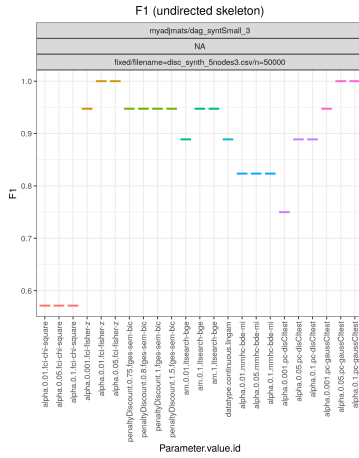


Figure 91: F1 Scores on the Synthetic 5 nodes data set. Discretized, no noise.

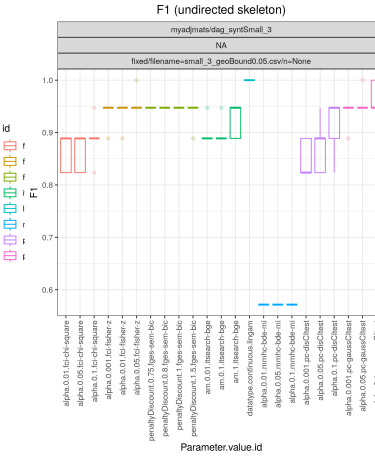


Figure 92: Synthetic 5 nodes data, Geo C-wise mechanism, max probability 0.05.

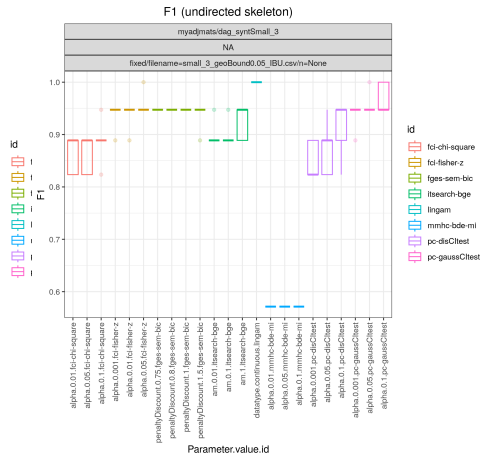


Figure 93: Synthetic 5 nodes data, Geo C-wise IBU mechanism, max probability 0.05.

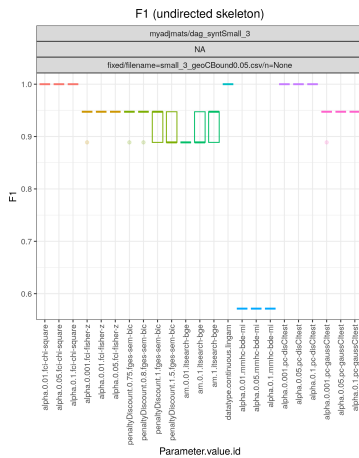


Figure 94: Synthetic 5 nodes data, Geo Comb mechanism, max probability 0.05.

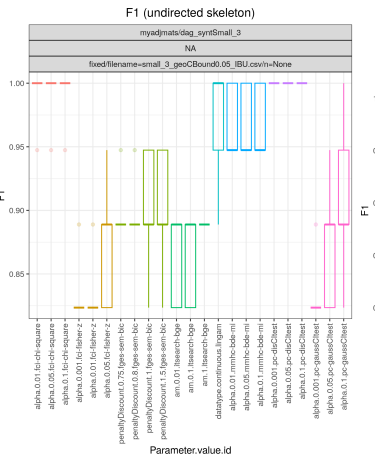


Figure 95: Synthetic 5 nodes data, Geo Comb IBU mechanism, max probability 0.05.

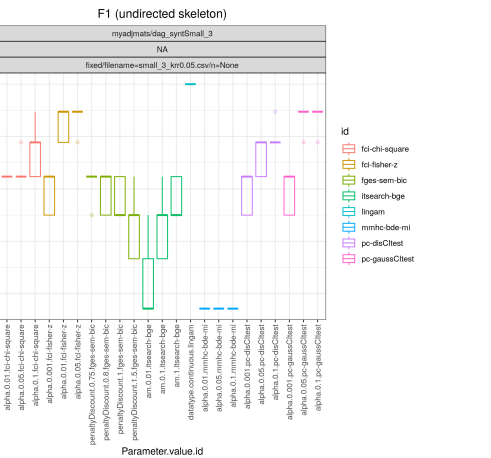


Figure 96: Synthetic 5 nodes data, k-RR C-wise mechanism, max probability 0.05.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

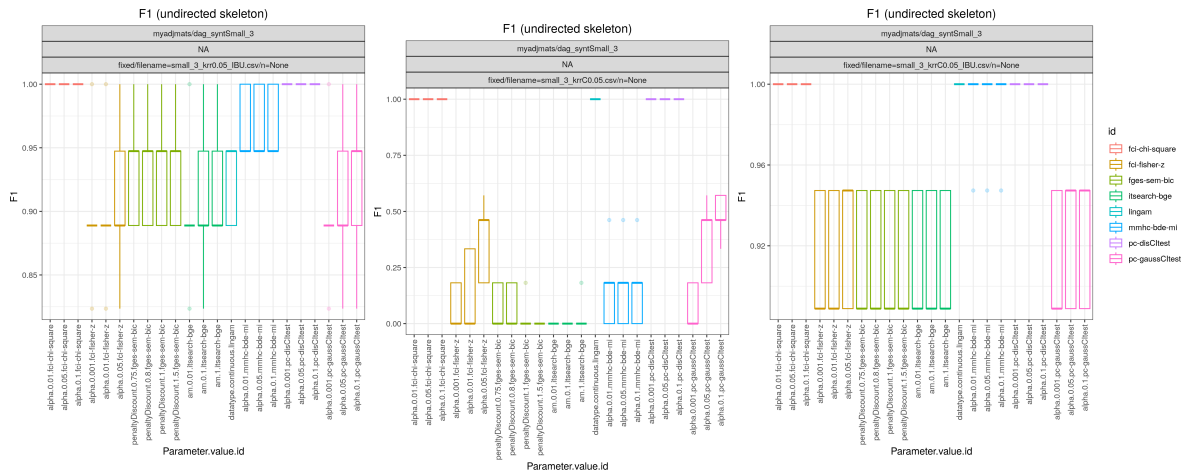


Figure 97: Synthetic 5 nodes data, k -RR C-wise IBU mechanism, max probability 0.05.

Figure 98: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.05.

Figure 99: Synthetic 5 nodes data, k -RR Comb IBU mechanism, max probability 0.05.

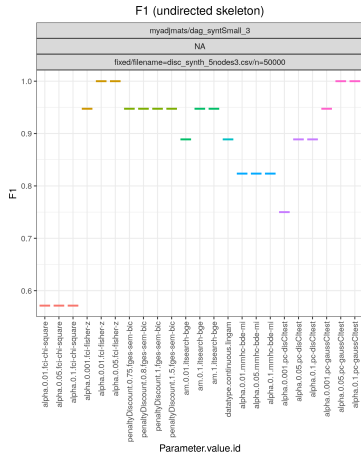


Figure 100: F1 Scores on the Synthetic 5 nodes data set. Discretized, no noise.

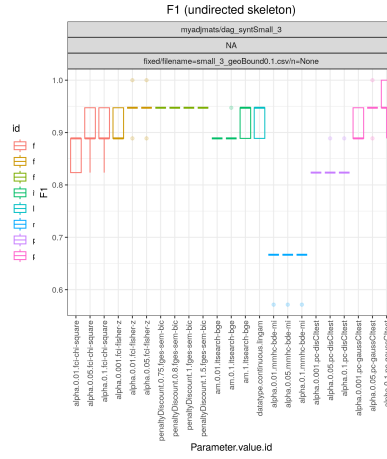


Figure 101: Synthetic 5 nodes data, Geo C-wise mechanism, max probability 0.1.

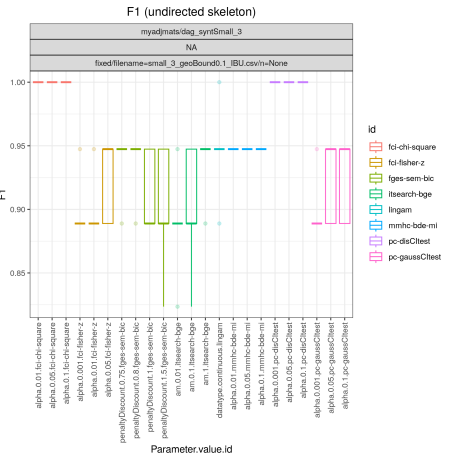


Figure 102: Synthetic 5 nodes data, Geo C-wise IBU mechanism, max probability 0.1.

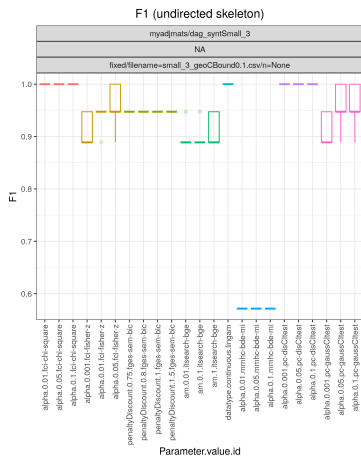


Figure 103: Synthetic 5 nodes data, Geo Comb mechanism, max probability 0.1.

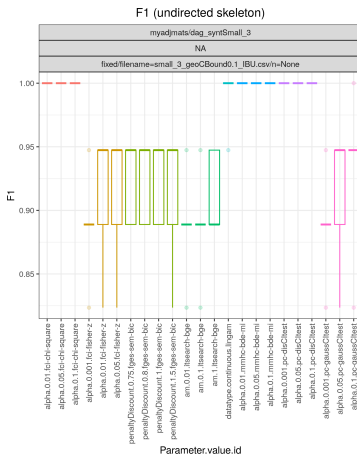


Figure 104: Synthetic 5 nodes data, Geo Comb IBU mechanism, max probability 0.1.

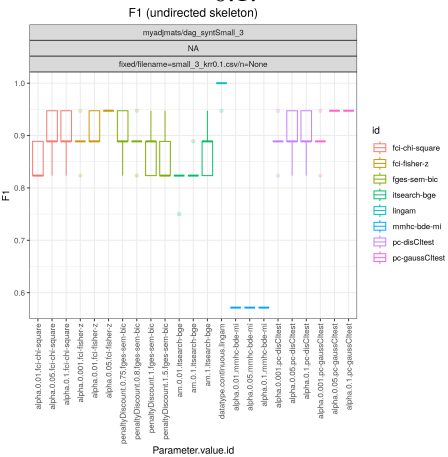


Figure 105: Synthetic 5 nodes data, k -RR C-wise mechanism, max probability 0.1.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

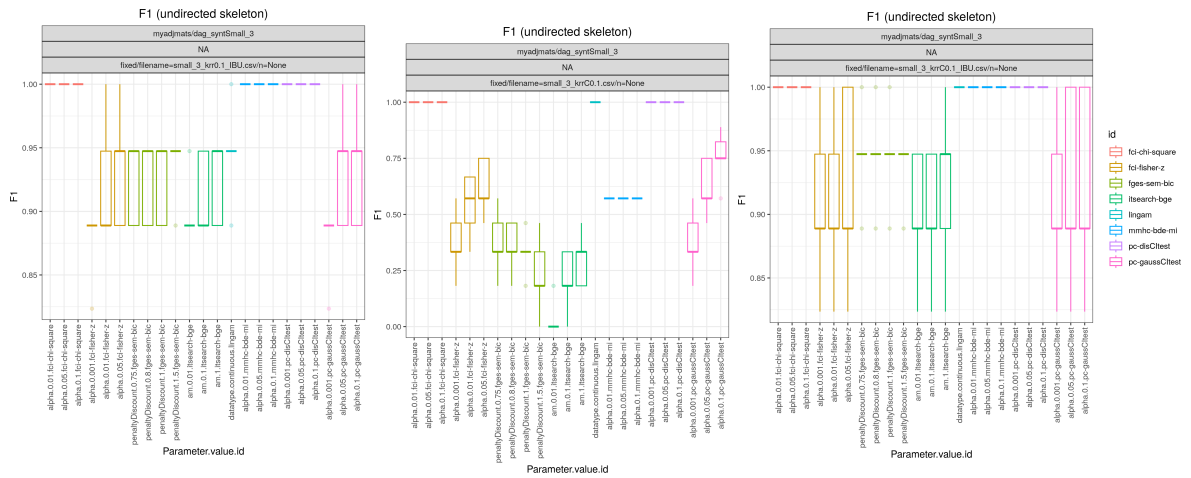


Figure 106: Synthetic 5 nodes data, k -RR C-wise IBU mechanism, max probability 0.1.

Figure 107: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.1.

Figure 108: Synthetic 5 nodes data, k -RR Comb IBU mechanism, max probability 0.1.

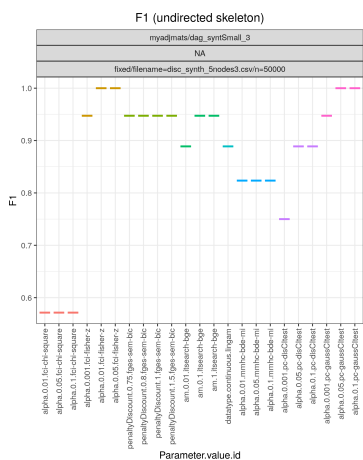


Figure 109: F1 Scores on the Synthetic 5 nodes data set. Discretized, no noise.

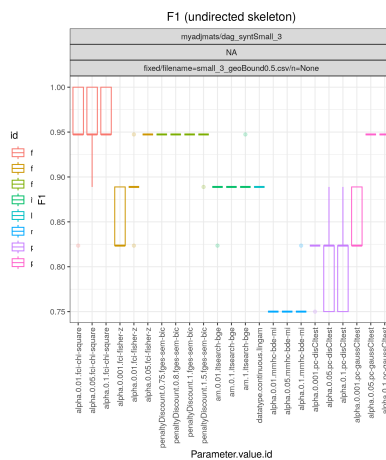


Figure 110: Synthetic 5 nodes data, Geo C-wise mechanism, max probability 0.5.

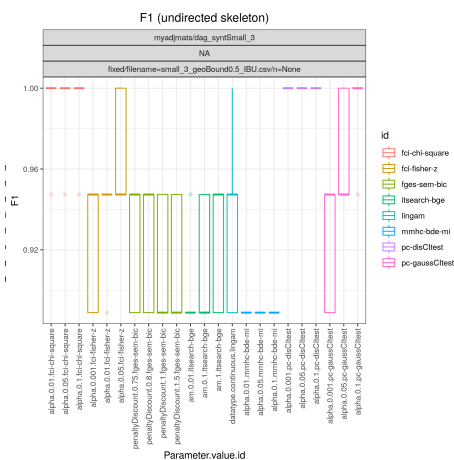


Figure 111: Synthetic 5 nodes data, Geo C-wise IBU mechanism, max probability 0.5.

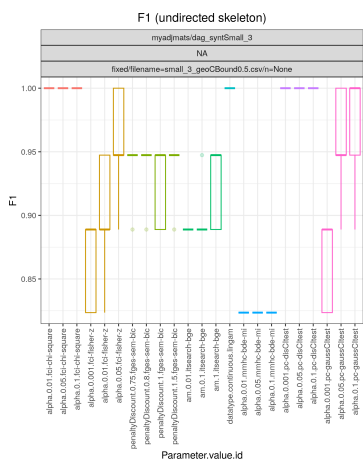


Figure 112: Synthetic 5 nodes data, Geo Comb mechanism, max probability 0.5.

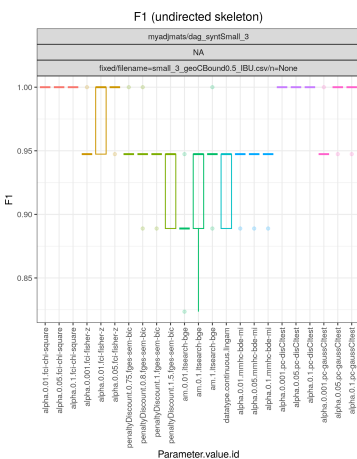


Figure 113: Synthetic 5 nodes data, Geo Comb IBU mechanism, max probability 0.5.

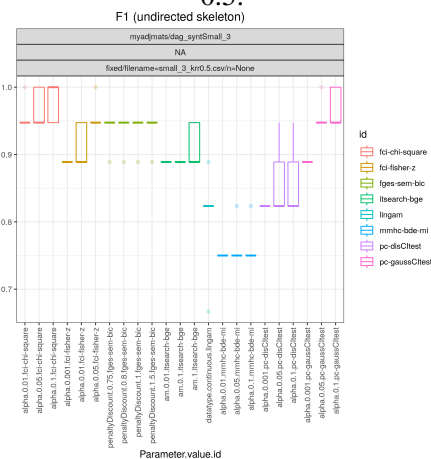


Figure 114: Synthetic 5 nodes data, k -RR C-wise mechanism, max probability 0.5.

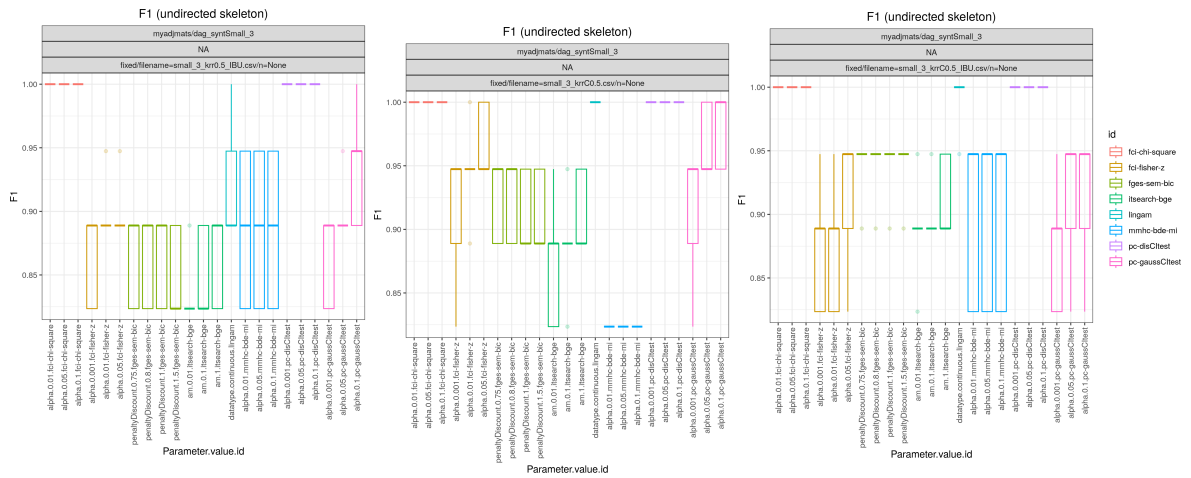


Figure 115: Synthetic 5 nodes data, k -RR C-wise IBU mechanism, max probability 0.5.

Figure 116: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.5.

Figure 117: Synthetic 5 nodes data, k -RR Comb IBU mechanism, max probability 0.5.

D.6. SHD Score results Synthetic 5 nodes data set

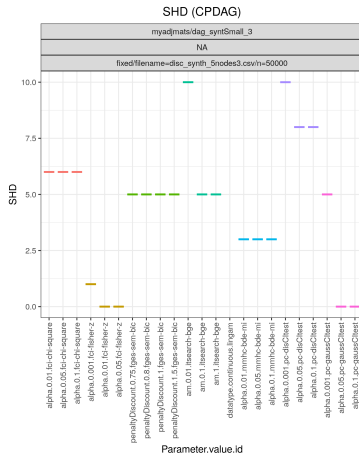


Figure 118: SHD Scores on the Synthetic 5 nodes data set. Discretized, no noise.

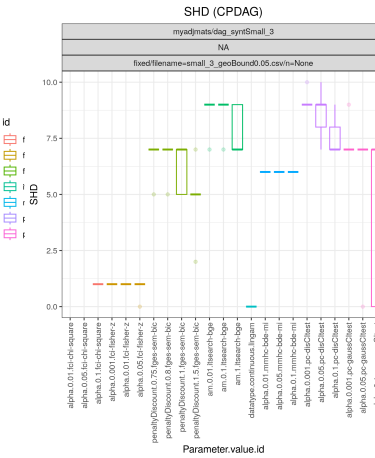


Figure 119: Synthetic 5 nodes data, Geo C-wise mechanism, max probability 0.05.

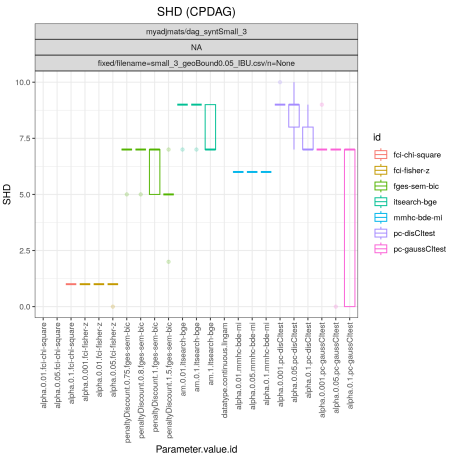


Figure 120: Synthetic 5 nodes data, Geo C-wise IBU mechanism, max probability 0.05.

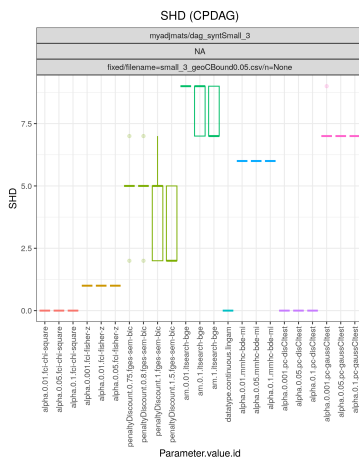


Figure 121: Synthetic 5 nodes data, Geo Comb mechanism, max probability 0.05.

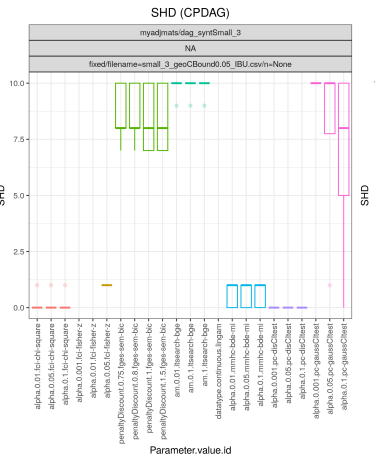


Figure 122: Synthetic 5 nodes data, Geo Comb IBU mechanism, max probability 0.05.

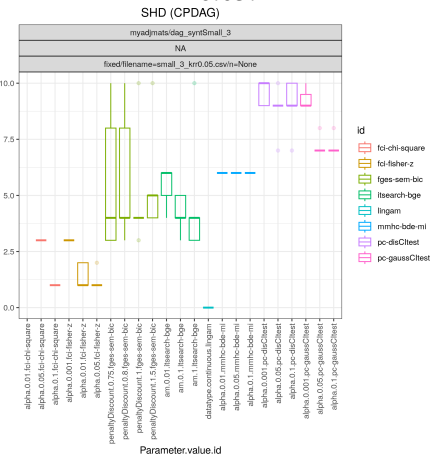


Figure 123: Synthetic 5 nodes data, k-RR C-wise mechanism, max probability 0.05.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

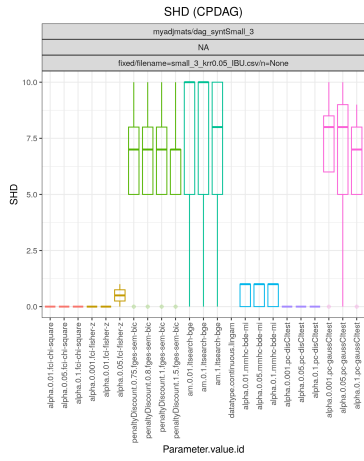


Figure 124: Synthetic 5 nodes data, k -RR C-wise IBU mechanism, max probability 0.05.

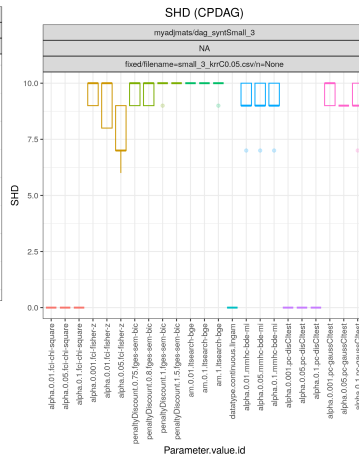


Figure 125: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.05.

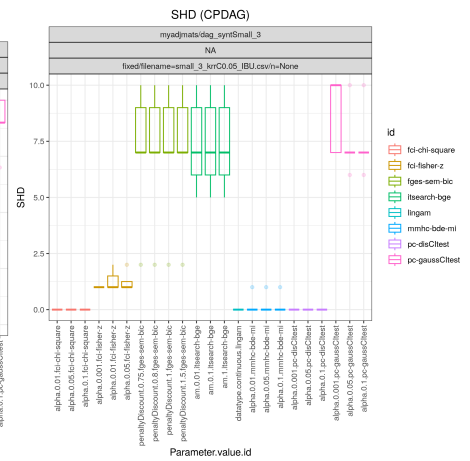


Figure 126: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.05.

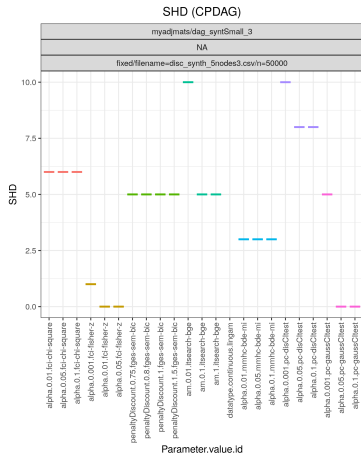


Figure 127: SHD Scores on the Synthetic 5 nodes data set. Discretized, no noise.

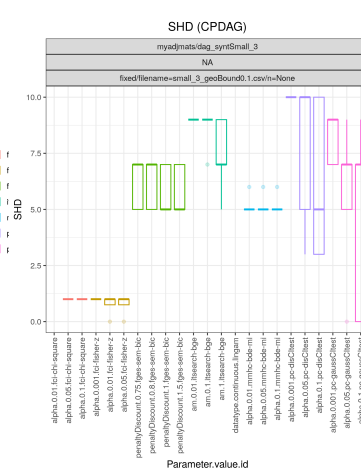


Figure 128: Synthetic 5 nodes data, Geo C-wise mechanism, max probability 0.1.

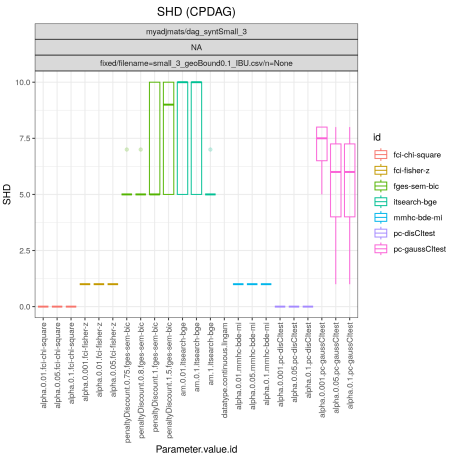


Figure 129: Synthetic 5 nodes data, Geo C-wise IBU mechanism, max probability 0.1.

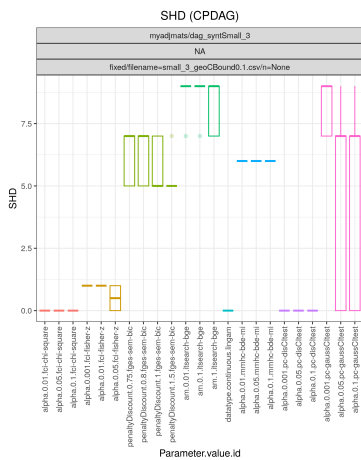


Figure 130: Synthetic 5 nodes data, Geo Comb mechanism, max probability 0.1.

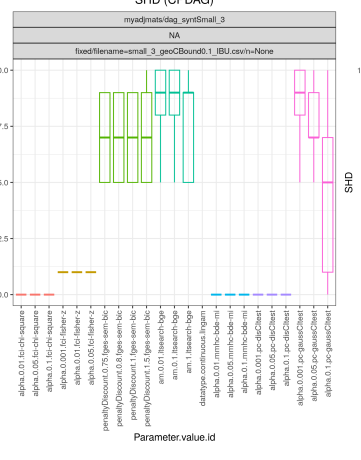


Figure 131: Synthetic 5 nodes data, Geo Comb IBU mechanism, max probability 0.1.

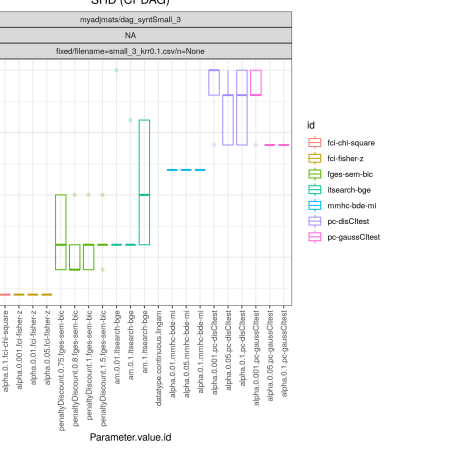


Figure 132: Synthetic 5 nodes data, k -RR C-wise mechanism, max probability 0.1.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

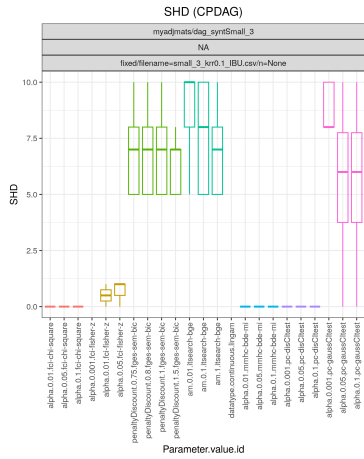


Figure 133: Synthetic 5 nodes data, k -RR C-wise IBU mechanism, max probability 0.1.

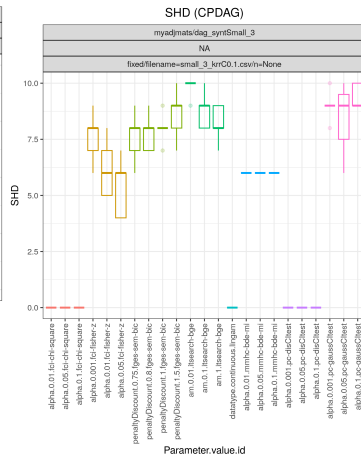


Figure 134: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.1.

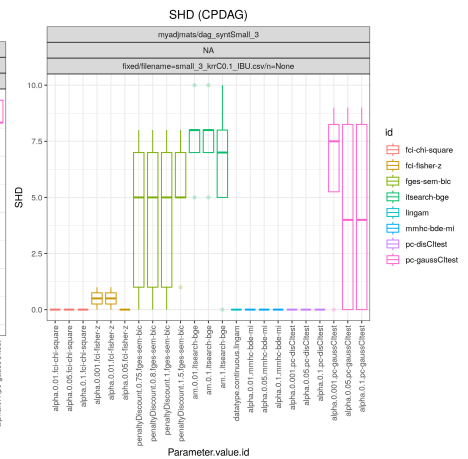


Figure 135: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.1.

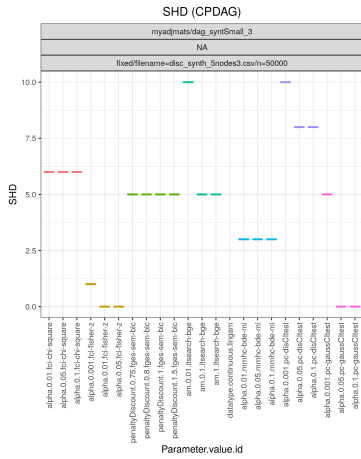


Figure 136: SHD Scores on the Synthetic 5 nodes data set. Discretized, no noise.

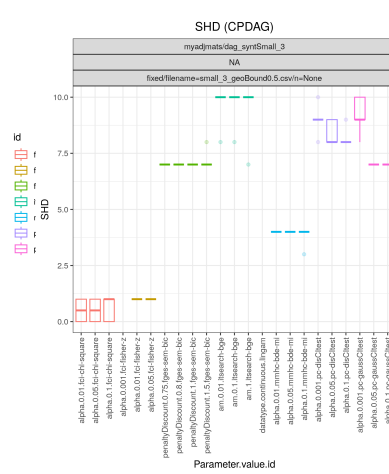


Figure 137: Synthetic 5 nodes data, Geo C-wise mechanism, max probability 0.5.

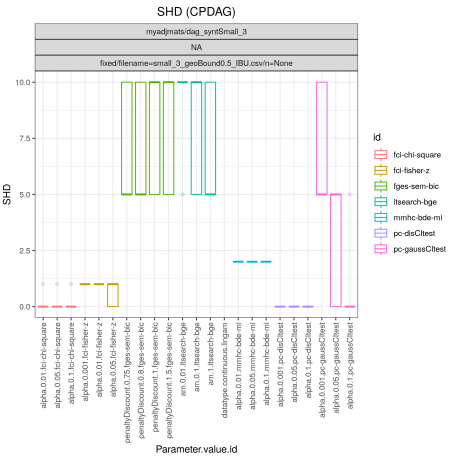


Figure 138: Synthetic 5 nodes data, Geo C-wise IBU mechanism, max probability 0.5.

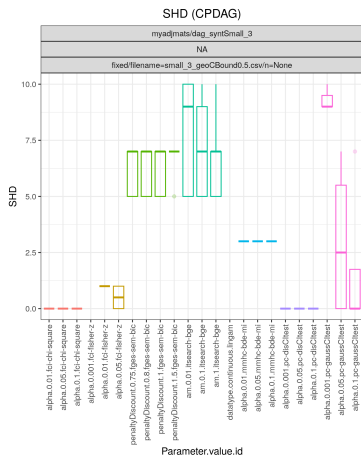


Figure 139: Synthetic 5 nodes data, Geo Comb mechanism, max probability 0.5.

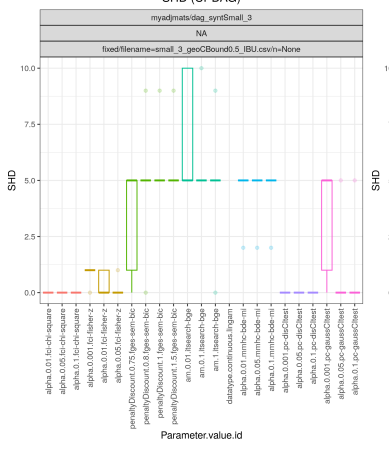


Figure 140: Synthetic 5 nodes data, Geo Comb IBU mechanism, max probability 0.5.

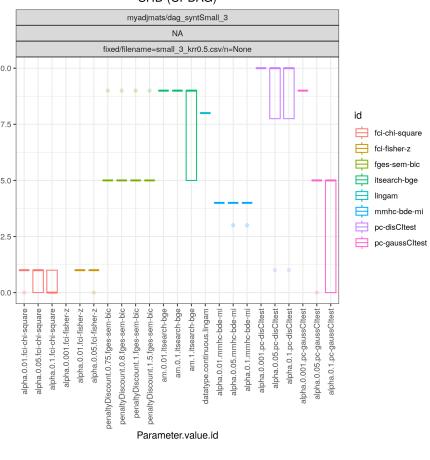


Figure 141: Synthetic 5 nodes data, k-RR C-wise mechanism, max probability 0.5.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

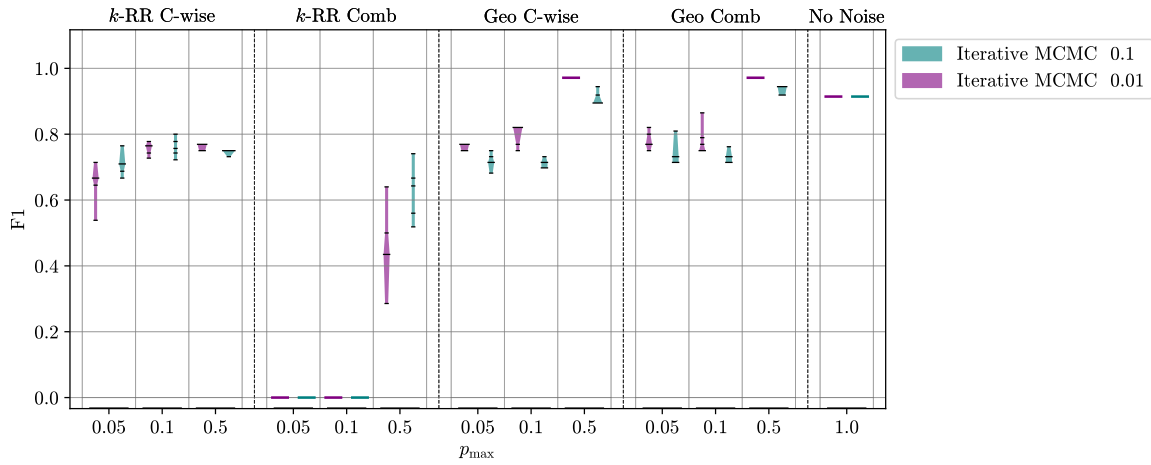


Figure 145: Synthetic data, 10 nodes, F1.

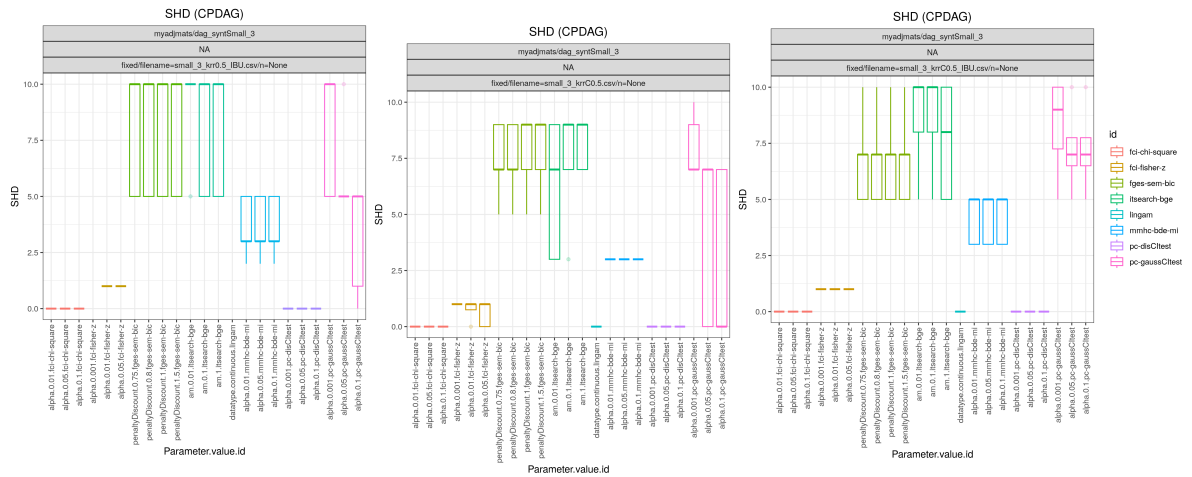


Figure 142: Synthetic 5 nodes data, k -RR C-wise mechanism, max probability 0.5.

Figure 143: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.5.

Figure 144: Synthetic 5 nodes data, k -RR Comb mechanism, max probability 0.5.

D.7. F1 Score results Synthetic 10 nodes data set

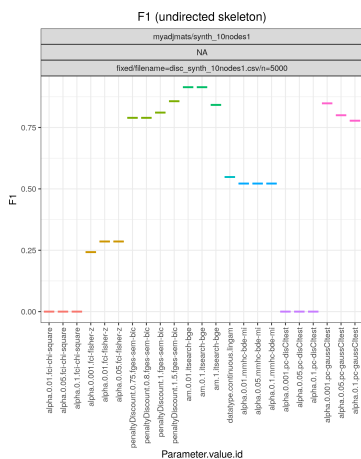


Figure 146: F1 Scores on the Synthetic 10 nodes data set. Discretized, no noise.

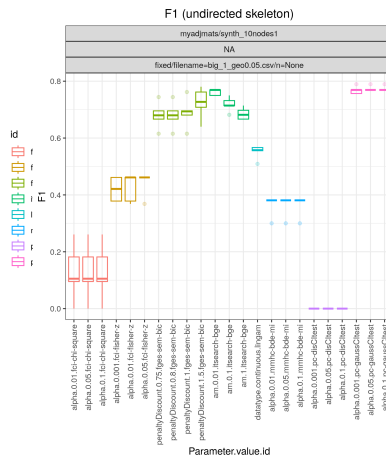


Figure 147: Synthetic 10 nodes data, Geo C-wise mechanism, max probability 0.05.

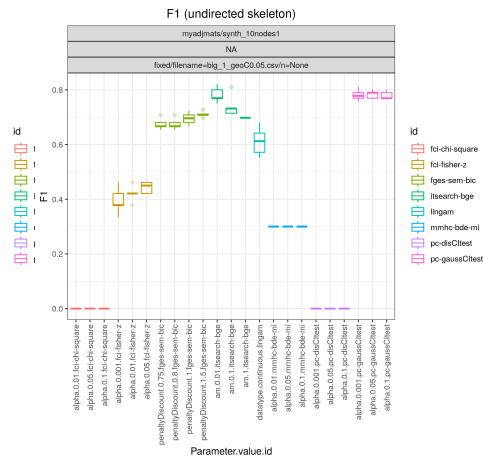


Figure 148: Synthetic 10 nodes data, Geo Comb mechanism, max probability 0.05.

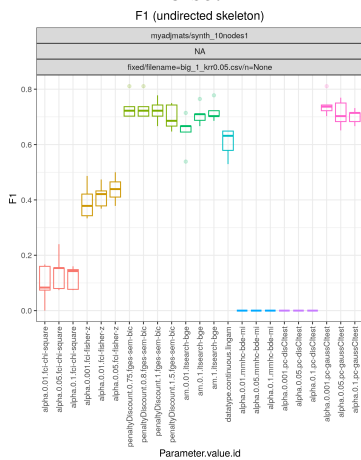


Figure 149: Synthetic 10 nodes data, k -RR C-wise mechanism, max probability 0.05.

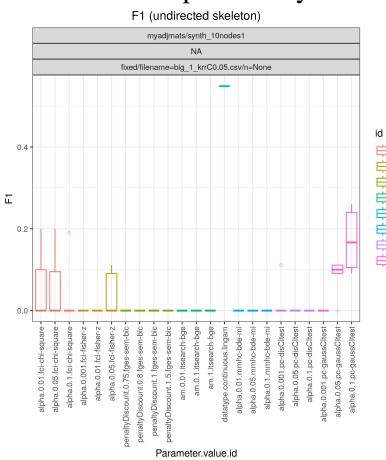


Figure 150: Synthetic 10 nodes data, k -RR Comb mechanism, max probability 0.05.

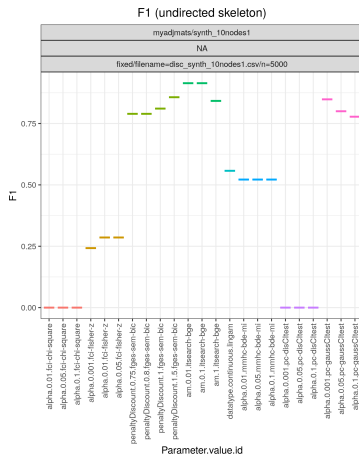


Figure 151: F1 Scores on the Synthetic 10 nodes data set. Discretized, no noise.

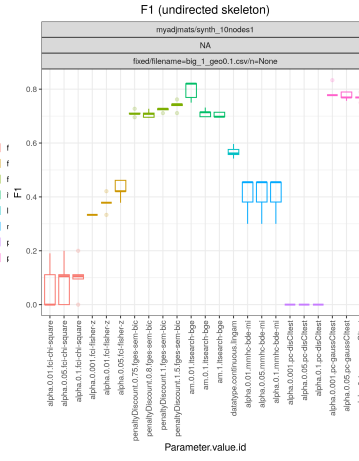


Figure 152: Synthetic 10 nodes data, Geometric mechanism, max probability 0.1.

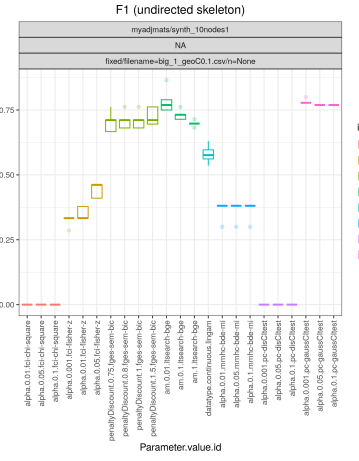


Figure 153: Synthetic 10 nodes data, Geometric mechanism, max probability 0.1.

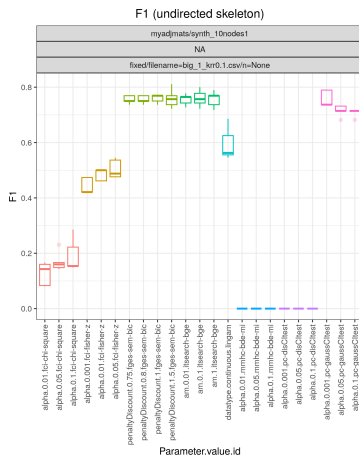


Figure 154: Synthetic 10 nodes data, k -RR mechanism, max probability 0.1.

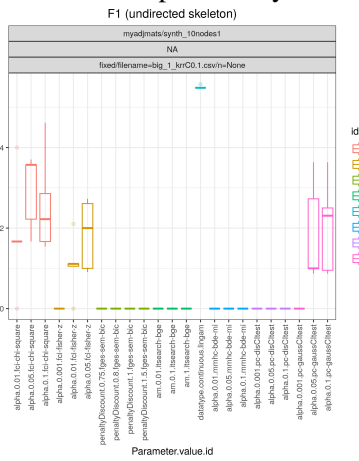


Figure 155: Synthetic 10 nodes data, Combinatorial mechanism, max probability 0.1.

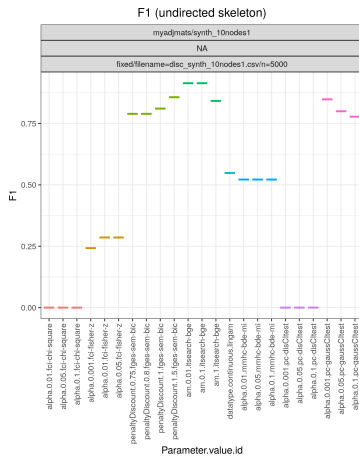


Figure 156: F1 Scores on the Synthetic 10 nodes data set. Discretized, no noise.

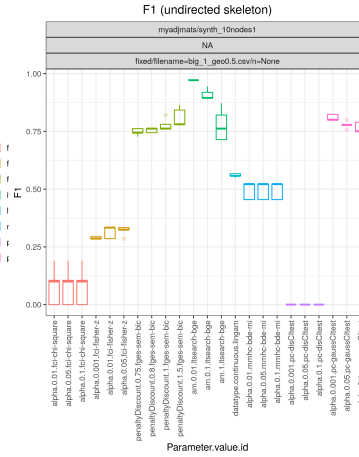


Figure 157: Synthetic 10 nodes data, Geometric mechanism, max probability 0.5.

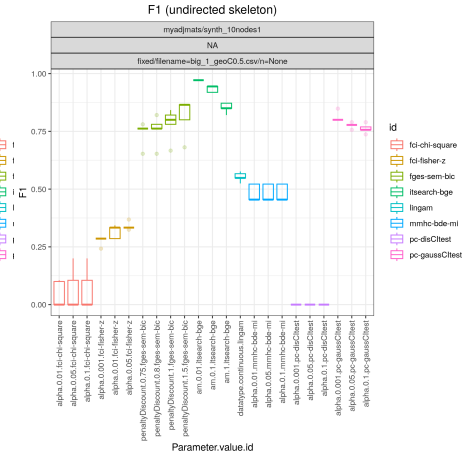


Figure 158: Synthetic 10 nodes data, Geometric mechanism, max probability 0.5.

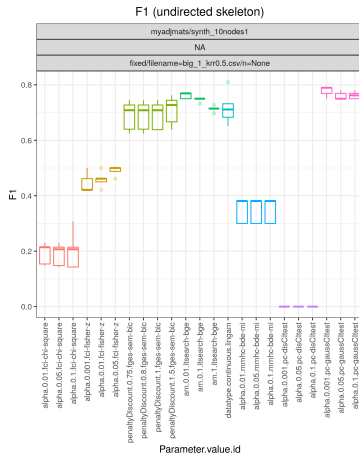


Figure 159: Synthetic 10 nodes data, k -RR mechanism, max probability 0.5.

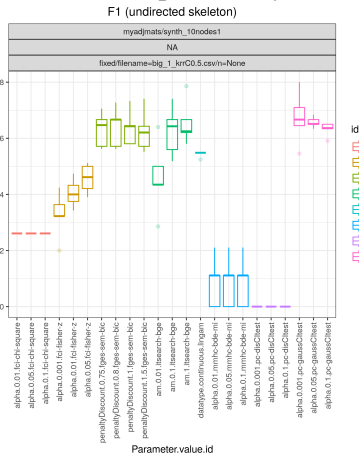


Figure 160: Synthetic 10 nodes data, k -RR mechanism, max probability 0.5.

D.8. SHD Score results Synthetic 10 nodes data set

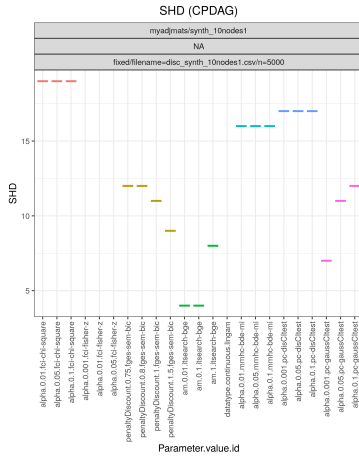


Figure 161: SHD Scores on the Synthetic 10 nodes data set. Discretized, no noise.

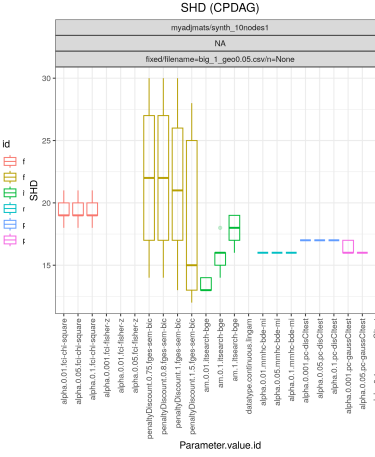


Figure 162: Synthetic 10 nodes data, Geo C-wise mechanism, max probability 0.05.

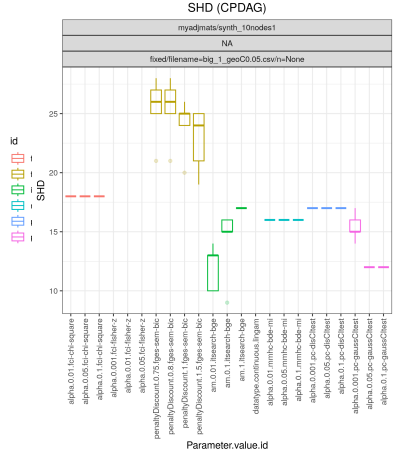


Figure 163: Synthetic 10 nodes data, Geo Comb mechanism, max probability 0.05.

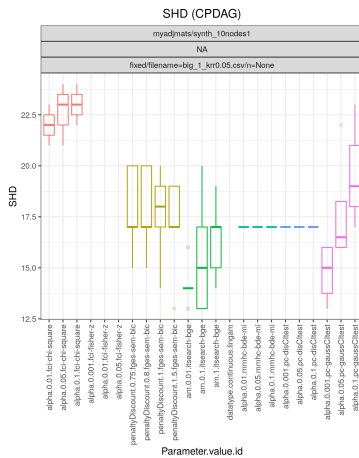


Figure 164: Synthetic 10 nodes data, k -RR C-wise mechanism, max probability 0.05.

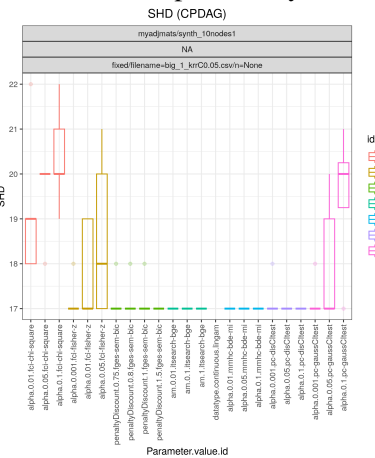


Figure 165: Synthetic 10 nodes data, k -RR Comb mechanism, max probability 0.05.

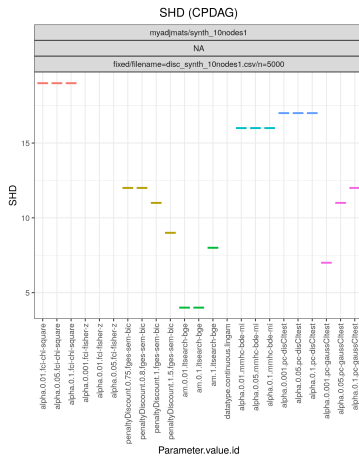


Figure 166: SHD Scores on the Synthetic 10 nodes data set. Discretized, no noise.

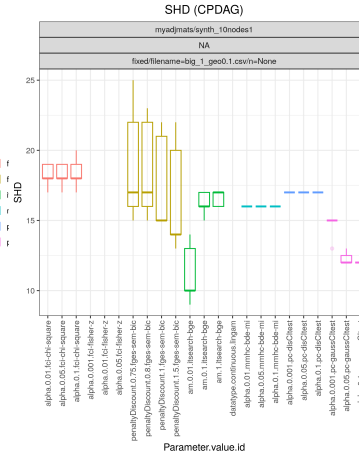


Figure 167: Synthetic 10 nodes data, Geometric mechanism, max probability 0.1.

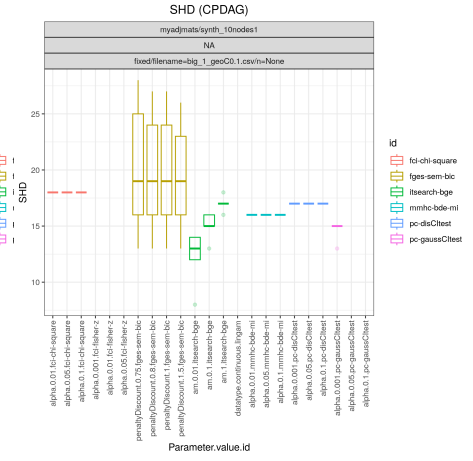


Figure 168: Synthetic 10 nodes data, Geometric mechanism, max probability 0.1.

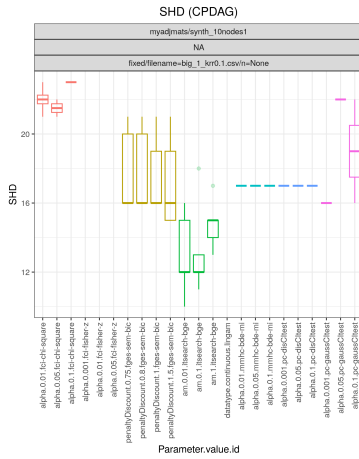


Figure 169: Synthetic 10 nodes data, k -RR mechanism, max probability 0.1.

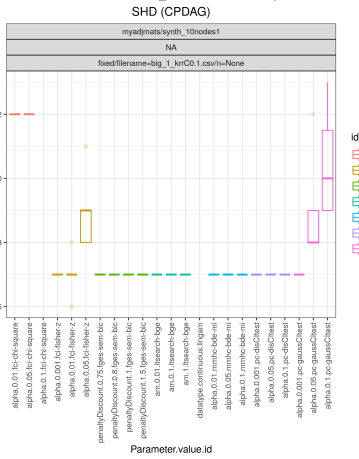


Figure 170: Synthetic 10 nodes data, k -RR mechanism, max probability 0.1.

CAUSAL DISCOVERY UNDER LOCAL PRIVACY

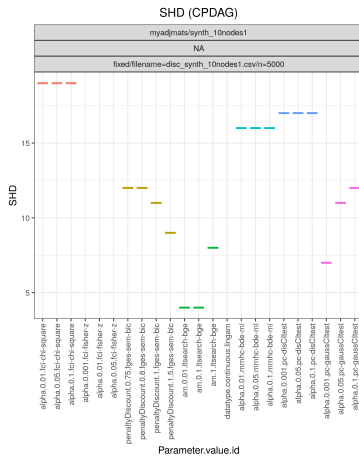


Figure 171: SHD Scores on the Synthetic 10 nodes data set. Discretized, no noise.

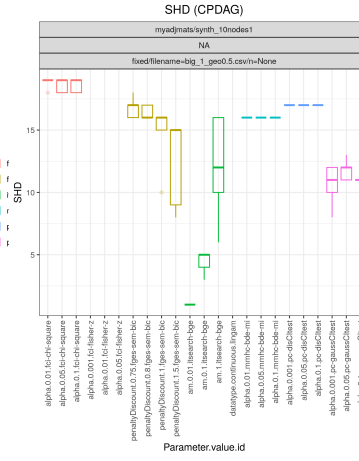


Figure 172: Synthetic 10 nodes data, Geometric mechanism, max probability 0.5.

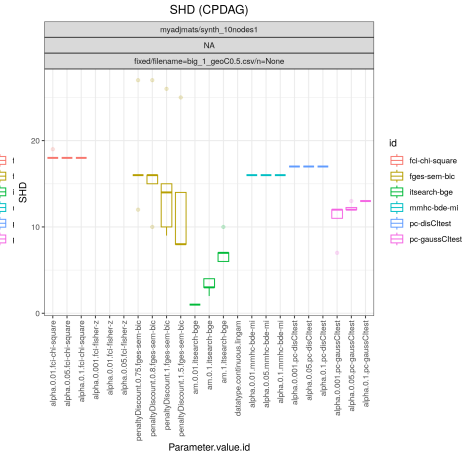


Figure 173: Synthetic 10 nodes data, Geometric mechanism, max probability 0.5.

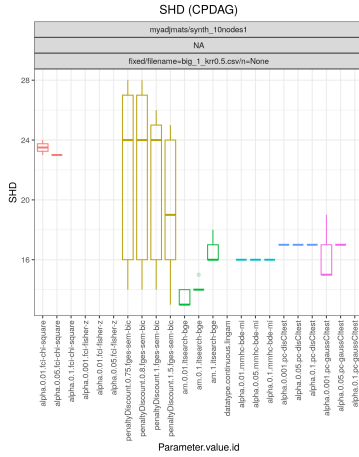


Figure 174: Synthetic 10 nodes data, k -RR mechanism, max probability 0.5.

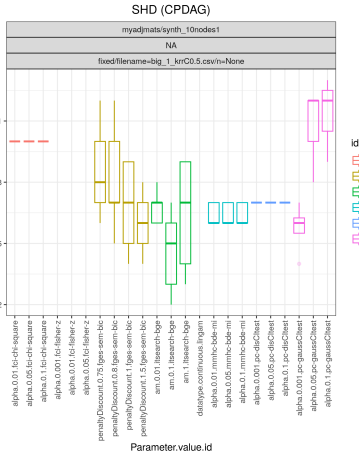


Figure 175: Synthetic 10 nodes data, k -RR mechanism, max probability 0.5.