

Potential Conditional Mutual Information: Estimators and Properties

Arman Rahimzamani and Sreeram Kannan

Department of Electrical Engineering, University of Washington, Seattle, WA
 Email: {armanrz, ksreeram}@uw.edu

Abstract—The conditional mutual information $I(X; Y|Z)$ measures the average information that X and Y contain about each other given Z . This is an important primitive in many learning problems including conditional independence testing, graphical model inference, causal strength estimation and time-series problems. In several applications, it is desirable to have a functional purely of the conditional distribution $p_{Y|X,Z}$ rather than of the joint distribution $p_{X,Y,Z}$. We define the potential conditional mutual information as the conditional mutual information calculated with a modified joint distribution $p_{Y|X,Z}q_{X,Z}$, where $q_{X,Z}$ is a potential distribution, fixed airport. We develop K nearest neighbor based estimators for this functional, employing importance sampling, and a coupling trick, and prove the finite k consistency of such an estimator. We demonstrate that the estimator has excellent practical performance and show an application in dynamical system inference.

I. INTRODUCTION

Given three random variables X, Y, Z , the conditional mutual information $I(X; Y|Z)$ (CMI) is the expected value of the mutual information between X and Y given Z , and can be expressed as follows [1]:

$$\begin{aligned} CMI_{X \leftrightarrow Y|Z}(p_{X,Y,Z}) &:= I(X; Y|Z) \\ &= D(p_{X,Y,Z} || p_Z p_{X|Z} p_{Y|Z}) \end{aligned} \quad (1)$$

Thus CMI is a functional of the joint distribution $p_{X,Y,Z}$. A basic property of CMI, and a key application, is the following: $I(X; Y|Z) = 0$ iff X is independent of Y given Z . This measure depends on the joint distribution between the three variables $p_{X,Y,Z}$. There are certain circumstances where such a dependence on the entire joint distribution is not favorable, and a measure that depends purely on the conditional distribution $p_{Y|X,Z}$ is more useful. This is because, in a way, conditional independence can be well defined purely in terms of the conditional distribution and the measure $p_{X,Z}$ is extraneous. This motivates the direction that we explore in this paper: we define potential conditional mutual information as a function purely of $p_{Y|X,Z}$ evaluated with a distribution $q_{X,Z}$ that is fixed a-priori.

An Example: Consider the following causal graph where $X \rightarrow Y$, $Z \rightarrow X$ and $Z \rightarrow Y$ shown in Figure 1a. Let

us say $p_{Y|X,Z}$ has a strong dependence on both X and Z , say defined by the structural equation $Y = X + Z$. We would like to measure the strength of the edge $X \rightarrow Y$ in this causal graphical model. One natural measure in this context is $I(X; Y|Z)$. However, if we use $I(X; Y|Z)$ as the strength, the strength goes to zero when $Z \approx X$ and this is undesirable. In such a case, Janzing et al [2] pointed out that a better strength of causal influence is given by the following:

$$C(X \rightarrow Y) := D(p_{X,Y,Z} || p_Z p_{X|Z} p_{Y|Z}). \quad (2)$$

This causal measure satisfies certain axioms laid out in that paper and is nonzero in the aforesaid example. However, in the case that the distribution p_X approaches a deterministic distribution (X is approximately a constant), this measure becomes zero, irrespective of the fact that the relationship from X and Z to Y remains unaltered. We would like to define a *potential dependence measure* that is dependent purely on $p_{Y|X,Z}$ and which has no dependence on the observed $p_{X,Z}$. We note that such a measure should give a (strong) non-zero result if $Y = X + Z$.

The Measure: We define potential conditional information measure as the conditional mutual information evaluated under a predefined distribution $q_{X,Z}$, and denote it as $qCMI(X \leftrightarrow Y|Z)$, and express it as follows.

$$qCMI_{X \leftrightarrow Y|Z}(p_{Y|X,Z}) := CMI_{X \leftrightarrow Y|Z}(q_{X,Z} p_{Y|X,Z}). \quad (3)$$

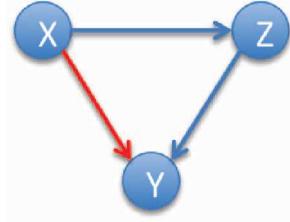
A Simple Property: The main question here is how to choose $q_{X,Z}$. A simple property that maybe of interest is the following, which can be easily stated in case that all three variables X, Y, Z are discrete. In such a case, it will be useful if we can have that $qCMI_{X \leftrightarrow Y|Z}(p_{Y|X,Z}) = 0$ if and only if $p_{Y|X,Z}$ depends purely only on Z . Such a property will be true for qCMI as long as $q_{X,Z}$ is non-zero for every value of X, Z . In case that all three variables are real valued, a similar statement can be asserted when $q_{X,Z}$ is a positive everywhere density, under the assumption that $p_{X,Y,Z}$ induces a joint density.

Instantiations: We will propose some instantiations of the potential CMI by giving examples of the distribution $q_{X,Z}$.

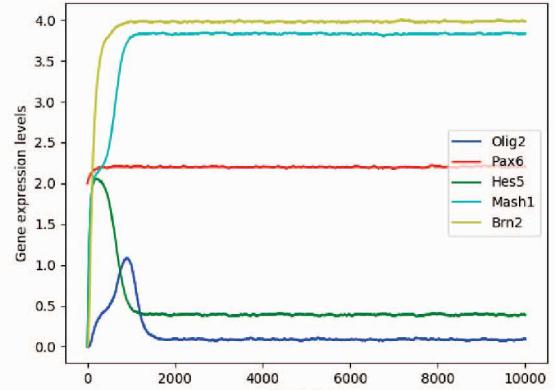
- CMI: $q_{X,Z} = p_{X,Z}$. Here the $q_{X,Z}$ is the factual measure and hence qCMI devolves to pure CMI.
- uCMI: $q_{X,Z} = u_{X,Z}$, where $u_{X,Z}$ is the product uniform distribution on X, Z . This is well defined when X, Z is either discrete or has a joint density with a bounded support.
- nCMI: $q_{X,Z} = n_{X,Z}$, where $n_{X,Z}$ is the i.i.d. Gaussian distribution on X, Z . This is well defined when X, Z are real-valued (whether scalar or vector).
- maxCMI = $\max_{q_{X,Z}}$ CMI($q_{X,Z} p_{Y|X,Z}$) is defined as an analog of the Shannon capacity in the conditional case, where we maximize the CMI over all possible distributions on X, Z . This is akin to tuning the input distribution to maximize the signal in the graph. Note that uCMI or nCMI is not invariant to invertible functional transformations on X, Z , whereas maxCMI is indeed invariant to such functional transformations.
- iCMI: $q_{X,Z} = p_X p_Z$ is the CMI evaluated not under the true joint distribution of X, Z but under the product distribution on X, Z . This measure is related to the causal strength measure proposed in [2], though not identical.

Note that uCMI, nCMI, maxCMI all satisfy the property that they are zero if and only if $p_{Y|X,Z}$ has no dependence on X , whereas CMI and iCMI do not.

Applications: A key application of the potential information measures is in testing graphical models, where conditional independence tests are the basic primitive by which models are built [3], [4], [5]. To give a concrete example of the setting, which motivated us to pursue this line of study, consider the following problem, which can model gene regulatory network inference from time-series data. We observe a set of n time series, $X_i(t)$ for $t = 1, 2, \dots, T$ with $i = 1, 2, \dots, n$ and wish to infer the graph of the dynamical system. The underlying model assumption is that \vec{X} is a markov chain with $X_i(t)$ depending only on $X_j(t-1)$ for $j \in Pa(i)$ and the goal is to determine $Pa(i)$, the set of parents of a given node. This was originally studied in the setting when the variables are jointly Gaussian and hence the dependence is linear (see [6] for the original treatment, and [7], [8] for versions with latent variables). This problem was generalized to the setting with arbitrary probability distributions and temporal dependences in [9] and studied further in [10], for one-step markov chains in [11] and deterministic relationships in [12]. From these works, under



(a)



(b)

Fig. 1: (a): A causal graph, where the interest is in determining the strength of X to Y . (b) A gene expression trace as a function of time for a few example genes.

some technical condition, we can assert that the following method is guaranteed to be consistent,

$$x_i \rightarrow x_j \iff I\{X_i(t-1); X_j(t)|X_{i^c}(t-1)\} > 0. \quad (4)$$

Thus to solve this problem, we estimate the CMI between the aforesaid variables. However, we observed while experimenting with gene regulatory network data (from [13]), that there is a strange phenomenon; the performance of the inference worsens as we collect more data: the number of data points *increases*.

An example of a gene expression time series for a few genes is shown in Figure 1b. It is clear that as the number of time points increases, the system is moving into an equilibrium with very little change in gene expression values. This induces a distribution on any $X_i(t)$ which looks more and more like a deterministic distribution.

In such a case, an information measure such as CMI which depends on the “input” distribution $p_{x_i(t-1)}$ will converge to zero and thus its performance will deteriorate as the number of samples increases. However a measure that depends on the conditional distribution $p_{x_j(t|x_{i^c}(t-1), x_i)}$ need not deteriorate with increasing number of samples. Thus

qCMI is more appropriate in this context (see Sec. IV-C for performance of qCMI on this problem).

Related work: In the case that there is a pair of random variables X, Y , recent work [14] explored conditional dependence measures which depend only on $p_{Y|X}$. Again in the two-variable case, a measure that had weak dependence on p_X was studied in [15]. The proposal there was to use the strong data processing constant and hypercontractivity [16], [17] to infer causal strength; this has strong relationships to information bottleneck [18]. In this paper, we extend [19] to conditional independence (rather than independence studied there). In a related but different direction, Shannon capacity, which is a potential dependence metric, was proposed in [20] to infer causality from observational data [21].

Main Contributions: In this paper, we make the following main contributions:

- 1) We propose *potential conditional mutual information* as a way of quantifying conditional independence, but depending only on the conditional distribution $p_{Y|X,Z}$.
- 2) We propose *new estimators* in the real-valued case that combine ideas from importance sampling, a coupling trick and k -nearest neighbors estimation to estimate potential CMI.
- 3) We prove that the proposed estimator is *consistent* for a fixed k , which does not depend on the number of samples N .
- 4) We demonstrate by simulation studies that the proposed estimator has excellent performance when there are a finite number of samples, as well as an application in gene network inference, where we show that qCMI can solve the non-monotonicity problem.

II. ESTIMATOR

In most real settings, we do not have access to either the joint distribution $p_{X,Y,Z}$ or the conditional distribution $p_{Y|X,Z}$, but need to estimate the requisite information functionals from observed samples. We are given N independent identically distributed samples $\{(x_i, y_i, z_i)_{i=1,2,\dots,N}\}$ from $p_{X,Y,Z}$. In the case of qCMI, the estimator is also given as input the modified distribution $q_{X,Z}$. The estimator needs to estimate qCMI from samples.

In the case of discrete valued distributions, it is possible to empirically estimate $p_{X,Y,Z}$ from samples and calculate the qCMI from this distribution. We focus our attention here on the case of continuous valued alphabet, where each variable takes on values in a bounded subset of \mathbb{R}^d . We assume that X, Y, Z are of dimensions d_x, d_y, d_z respectively, and let $f_{X,Y,Z}$ denote the joint density of the three variables (we assume that it exists). In such a case, it is possible to estimate $f_{X,Y,Z}$ using kernel density estimators [22], [23] and then warp the estimate using the potential measure $q_{X,Z}$. However, it is known that k -nearest neighbors based

estimators perform better even in the simpler case of mutual information estimation and are widely used in practice [24], [25]. Therefore in this work, we develop KNN based estimators for qCMI estimation.

A. Entropy estimation

Consider first the estimation of the differential entropy of a random variable X with density f_X and observed samples x_1, \dots, x_N . A simple method to estimate the differential entropy is to use the re-substitution estimator, where we calculate $\hat{h}(X) := \frac{1}{N} \sum_{i=1}^N \log(\hat{f}_X(x_i))$, where \hat{f}_X is an estimate of the density of X . We can estimate the density using a KNN based estimator. To do so, we fix k a-priori, and for each sample x_i , find the distance $\rho_{k,i}$ to the nearest neighbor.

$$\hat{f}_X(x_i) c_d \rho_{k,i}^d \approx \frac{k}{N}. \quad (5)$$

This estimator is not consistent when k is fixed, and it was shown in a remarkable result by Kozhachenko and Loenenko [26] that the bias is independent of the distribution and can be computed apriori. Thus the following estimator was shown in [26] to be consistent for differential entropy.

$$\hat{h}_{\text{KL}}(X) = \frac{1}{N} \sum_{i=1}^N \log \frac{N \rho_{k,i}^d c_d}{k} + \log k - \psi(k).$$

While it is possible to have estimators which fix an ϵ apriori and then find the number of nearest neighbors to plug into the formula, such estimators do not adapt to the density (some regions will have many more points inside an ϵ neighborhood than others) and do not have a consistency proof as well. We mention this as fixed ϵ estimators are used for a sub-problem in our estimator.

B. Coupling trick

The conditional mutual information can be written as a sum of 4 differential entropies, and one can estimate these differential entropies independently using KNN estimators and sum them.

$$I(X; Y|Z) = -h(X, Y, Z) - h(Z) + h(X, Z) + h(Y, Z).$$

However, even in the case of mutual information, the estimation can be improved by an inspired coupling trick, in what is called the KSG estimator [24]. We note that the original KSG estimator did not have a proof of consistency and its consistency and convergence rates were analyzed in a recent paper [27]. Also of interest is the fact that the coupling trick has been shown to be quite useful in problems where X, Y or both have a mixture of discrete and continuous distributions or components [28].

This trick was applied in the context of conditional mutual information estimation in [29]. However, we note that this estimator of CMI does not have a proof of consistency to the

Algorithm 1: qCMI algorithm

Data: Data Samples (x_i, y_i, z_i) for $i = 1, \dots, N$ and $q_{X,Z}$

Result: $q\hat{CMI}$ an estimate of $qCMI$

Step 1: Calculate weights ω_i

for $i = 1, \dots, N$ **do**

Estimate $\hat{f}_{XZ}(x_i, z_i)$ using a Kernel density estimator [23], [22].
 $\omega_i := \frac{q_{XZ}(x_i, z_i)}{\hat{f}_{XZ}(x_i, z_i)}$, the importance sampling estimate of sample i .

end

Step 2: Calculate information samples I_i

for $i = 1, \dots, N$ **do**

$\rho_{k,i} :=$ Distance of k -th nearest neighbor of (x_i, y_i, z_i) .
 $n_{xz,i} := \sum_{j \neq i: \|(x_i, z_i) - (x_j, z_j)\| < \rho_{k,i}} 1$, the number of neighbors of (x_i, z_i) within distance $\rho_{k,i}$.
 $n_{yz,i} := \sum_{j \neq i: \|(y_i, z_i) - (y_j, z_j)\| < \rho_{k,i}} \omega_j$, the weighted number of neighbors of (y_i, z_i) within distance $\rho_{k,i}$.
 $n_{z,i} := \sum_{j \neq i: \|z_i - z_j\| < \rho_{k,i}} \omega_i$, the weighted number of neighbors of z_i within distance $\rho_{k,i}$.
 $I_i := \psi(k) - \log(n_{xz,i}) - \log(n_{yz,i}) + \log(n_{z,i})$.

end

Return $q\hat{CMI} = \frac{1}{N} \sum_{i=1}^N I_i + \log \left(\frac{c_{d_x+d_z} c_{d_y+d_z}}{c_{d_x+d_y+d_z} c_{d_z}} \right)$.

best of our knowledge. The CMI estimator essentially fixes a k for the (X, Y, Z) vector and calculates for each sample (x_i, y_i, z_i) , the distance $\rho_{k,i}$ to the k -th nearest neighbor. The estimator fixes this $\rho_{k,i}$ as the distance and calculates the number of nearest neighbors within $\rho_{k,i}$ in the Z , (X, Z) and (Y, Z) dimensions as $n_{z,i}, n_{xz,i}, n_{yz,i}$ respectively. The CMI estimator is then given by,

$$q\hat{CMI} := \frac{1}{N} \sum_{i=1}^N (-\log(n_{xz,i}) - \log(n_{yz,i}) + \log(n_{z,i})) + \psi(k) + \log \left(\frac{c_{d_x+d_z} c_{d_y+d_z}}{c_{d_x+d_y+d_z} c_{d_z}} \right) \quad (6)$$

C. qCMI estimator

Here, we adapt this estimator to calculate the qCMI for a given potential distribution $q_{X,Z}$. The major difference is the utilization of an importance sampling estimator to get the importance of each sample i estimated as follows,

$$\omega_i := \frac{q_{XZ}(x_i, z_i)}{\hat{f}_{XZ}(x_i, z_i)}. \quad (7)$$

However, importance sampling based reweighting alone is insufficient to handle qCMI estimation, since there is a logarithm term which depends on the density also. We handle this effect by appropriately re-weighting the number of nearest neighbors for the (y, z) and z terms carefully using the importance sampling estimators. The estimation algorithm is described in detail in Algorithm 1.

III. PROPERTIES

Our main technical result is the consistency of the proposed potential conditional mutual information estimator. This proof requires combining several elements from importance sampling, and accounting for the correlation induced by the coupling trick, in addition to handling the fact that the k is fixed and hence introduces a bias into estimation.

Assumption 1. We make the following assumptions.

- a) $\int f_{XYZ}(x, y, z) (\log f_{XYZ}(x, y, z))^2 dx dy dz < \infty$.
- b) All the probability density functions (PDF) are absolutely integrable, i.e. for all $A, B \subset \{X, Y, Z\}$, $\int |f_{A|B}(a|b)| da < \infty$ and $\int |f_{A|B}^q(a|b)| da < \infty$.
- c) There exists a finite constant C such that the hessian matrices of f_{XYZ} and f_{XYZ}^q exist and it's true that $\max\{\|h(f_{XYZ})\|_2, \|h(f_{XYZ}^q)\|_2\} < C$ almost everywhere.
- d) All the PDFs are upper-bounded, i.e. there exists a positive constant C' such that for all $A, B \subset \{X, Y, Z\}$, $f_{A|B} < C'$ and $f_{A|B}^q < C'$ almost everywhere.
- e) f_{XZ} is upper and lower-bounded, i.e. there exist positive constants C_1 and C_2 such that $C_1 q_{XZ}(x, z) < f_{XZ}(x, z) < C_2 q_{XZ}(x, z)$ almost everywhere.
- f) There bandwidth h_N of kernel density estimator is chosen as $h_N = \frac{1}{2} N^{-1/(2d_x+2d_z+3)}$.
- g) The k for the KNN estimator is chosen satisfying $k > \max\{\frac{d_z}{d_x+d_y}, \frac{d_x+d_y}{d_z}, \frac{d_x+d_z}{d_y}\}$

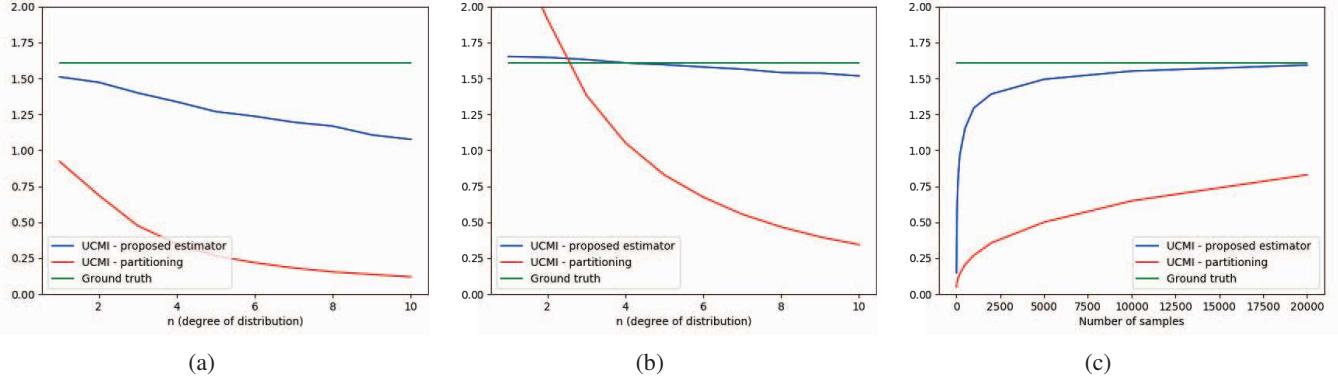


Fig. 2: The qCMI values calculated for $u^n(0, 1)$ distributions for X and Z , and uniform noise: (a) $N=1000$ samples and (b) $N=20000$ samples. (c) Degree $n = 5$.

Theorem 1. Under the Assumption 1, the qCMI estimator expressed in Algorithm 1 converges to the true value qCMI.

$$\hat{qCMI} \xrightarrow{p} qCMI \quad (8)$$

Proof: Please see [30] for the proof. ■

IV. SIMULATION STUDY

In this section, we describe some simulated experiments we did to test the qCMI algorithm. The reader should notice that all the tests we have done are taking q_{XZ} as u_{XZ} , i.e. all the tests are done for the special case of uCMI. So by exploiting the qCMI notations we mean uCMI everywhere.

A. qCMI consistency

The first numerical experiment we do is to test the consistency of our qCMI estimation algorithm. We set up a system of three variables X , Y and Z . The variables X and Z are independent taken from $u^n(0, 1)$ distribution, i.e. X and Z are taken from a uniform distribution and then raised to a power of n . When $n = 1$, the variables X and Z are already uniform. When n is large, the $u^n(0, 1)$ distribution skews more towards 0. For simplicity we apply identical n for both X and Z here. Then Y is generated as $Y = (X + Z + W) \bmod 1$ in which the noise term W is sampled from $u(0, 0.2)$. From elementary information theory calculation, we can deduce that $I(X; Y|Z) = \log(\frac{1}{2}) = 1.609$ if $n = 1$. Thus $I^q(X; Y|Z) = 1.609$ for all n . We plot the estimated value against the ground truth.

As the first part of the experiment, we keep the number of samples constant at 1000 and 20000, and change the degree n from 1 to 10. We compare the results of our KSG-based method with the simple partitioning method, and the theoretical value of qCMI. For the partitioning method, the number of partitions at each dimension is determined by

$\sqrt[3]{100N}$, so that we observe on average 100 samples inside each quantization bin.

The results are shown in Figure 2a and Figure 2b. Our expectation is that qCMI remains constant as n (degree of distribution) changes. We see that with relatively high number of samples, the accuracy of proposed qCMI is satisfactorily high.

As the second part of the experiment, we do the same experiment as the first part, but this time we keep $n = 5$ and change the number of samples. The result is shown in Figure 2c. We can see convergence of KSG-based qCMI estimator to the true value and how it outperforms the partitioning-based qCMI method.

As the third part of the experiment, we repeated the process for the first part, but replaced the $u^n(0, 1)$ distributions with $\beta(1.5, 1.5)$ and the noise distribution with $N(0, \sigma^2)$ and repeated the experiment for $\sigma = 0.3, 1.0$. For this part we kept the number of partitions at 25 for each dimension. The results of calculated qCMI values are shown in Figure 3a and Figure 3b.

B. Dealing with discrete components

As we discussed before, the qCMI algorithm replaces the observed distribution f_{XZ} distribution with a distribution q_{XZ} . This property comes in handy when we want to remove the bias caused by repeated samples. For example, as discussed earlier, suppose that we want to measure the mutual information of two coupled variables in a dynamical system evolving through time. Such systems usually start from an initial state, go through a transient state and eventually reach a steady state. If one takes samples of the system's state at a constant rate to study the interaction of two variables, they might end up taking too many samples from the initial and steady states while the transient phase which usually happens in a relatively short time might be

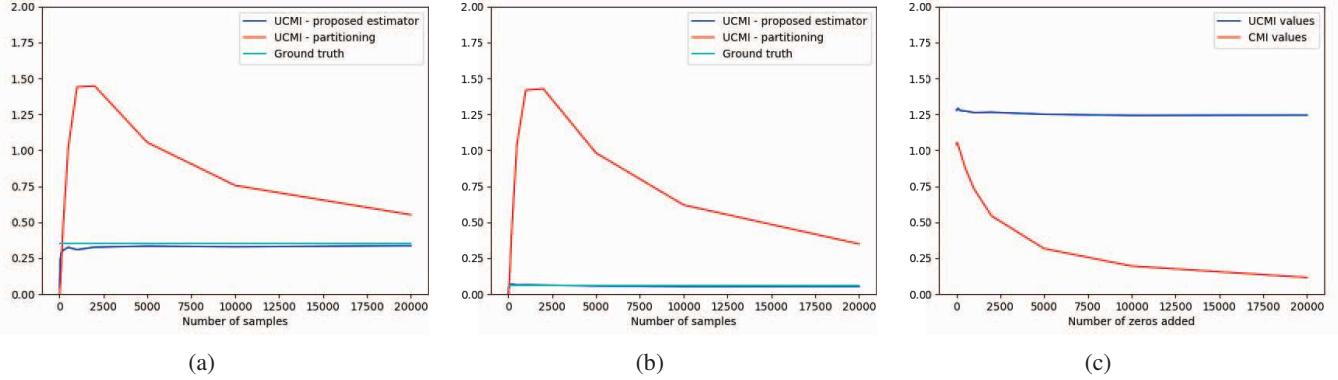


Fig. 3: The qCMI values calculated for a system with beta distribution for X and Z and Gaussian additive noise: (a) $\sigma = 0.3$, and (b) $\sigma = 1.0$. (c) The qCMI and CMI values versus the number of zeros added.

more informative. The conditional mutual information is not able to deal with this undesirable bias caused by the initial and steady states, while qCMI inherently deals with the effect by compensating for the samples which are less likely to happen.

To better observe the effect, we repeat the first experiment of the previous section, but this time we generate 1000 samples from the scenario, and then add zeros to the X , Y and Z to create a high probability of occurrence at $(0, 0, 0)$. The proof of consistency of the estimator holds only when there is a joint density, i.e., the joint measure is absolutely continuous with respect to the Lebesgue measure, and hence does not directly apply to this case. We refer the reader to [28] for an analysis of a similar coupled KNN estimator for mutual information in the discrete-continuous mixture case.

Changing the number of zero points added from 0 to 20000, we apply the conditional MI and qCMI to the data generated and compare the results. As we can see in figure 3c, with the number of zeros increasing, the value of conditional MI falls down to zero, unable to capture the inter-dependence of X and Y given Z , while qCMI value remains unchanged, properly discovering the inter-dependence from the transient values.

C. Non-linear Neuron Cells' Development Process

In this section, we apply the RDI and uRDI algorithms to neuron cells' development process simulated based on a model from [13] which can be modeled as a dynamical system. A dynamical System is described as a set of variables shown by a vector of \underline{x} which evolve through time starting from an initial state $\underline{x}(0)$. The evolution can be described as a vector function $\underline{g}(\cdot)$ such that $\underline{x}(t) = \underline{g}(\underline{x}(t-1))$. Note that \underline{g} can be a stochastic function in general, i.e. it may include random coefficients, additive noise and so on.

The dynamical system here describes the evolution of

13 genes through the development process. The non-linear equations governing the development process approximate a continuous development process, in which $\dot{\underline{x}}(t) = \underline{g}(\underline{x}(t-1))$. In other words, $\underline{x}(t) = \underline{x}(t-1) + dt \cdot \underline{g}(\underline{x}(t-1)) + \underline{n}(t)$ in which \underline{n} are independent Gaussian noises $\sim N(0, \sigma^2)$.

For this system, we want to infer the true network of causal inferences. In a dynamical system, we say x_i causes x_j if $x_j(t)$ is a function of $x_i(t-1)$. For this purpose, we first apply the RDI algorithm [12] to extract the pairwise directed causality between the variables by calculating $I(x_i(t-1), x_j(t) | x_j(t-1))$. Then we apply the uRDI algorithm, in which the conditional mutual information $I(X; Y | Z)$ in RDI is replaced with qCMI as $I^q(X; Y | Z)$ using $q_{X, Z}$ as a uniform distribution.

This system is a good example of a system in which the genes undergo a rather short transient state compared to the initial and steady states, and hence we expect an improvement in the performance of causal inference by applying uRDI (see Figure 1b for an example run of the system). The details of the dynamical system are given in [13].

We simulated the system for discretization $dt = 0.1$ and $\sigma = .001$, and changed the number of steps until which the system continues developing, and then applied the RDI and uRDI algorithms to evaluate the performance of each of the algorithms in terms of the area-under-the-ROC-Curve (AUC). The results are shown in Figure 4a. As we can see, with the number of steps increasing implying the number of samples captured in the steady state are increased, the uRDI algorithm outperforms RDI. In another test scenario, we fixed the number of steps at 200, but concatenated several runs of the same process. The results and the improvement of performance by uRDI can be seen in the Figure 4b.

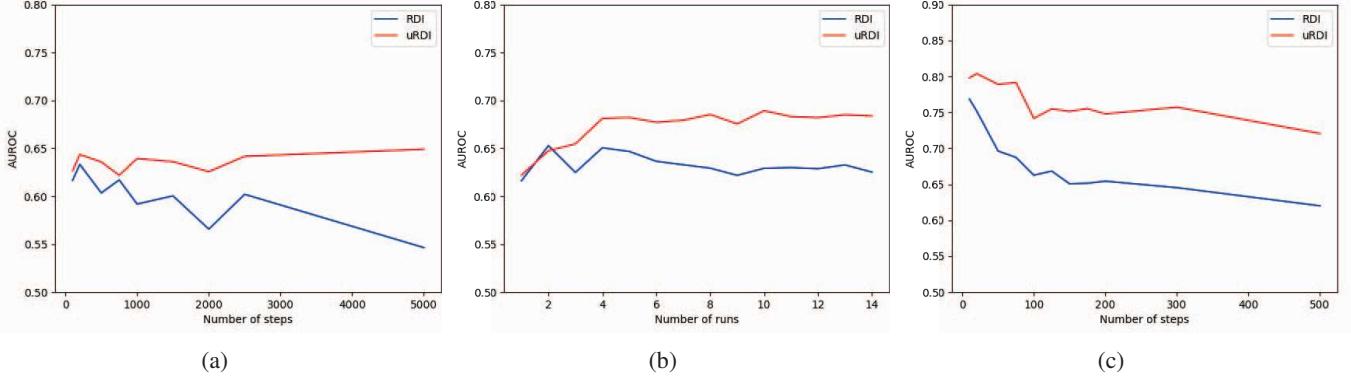


Fig. 4: AUC values for the neuron cells' development process: a) versus the number of steps. b) versus the number of runs. (c) for the decaying linear system

D. Decaying Linear Dynamical System

In this section, we simulate a linear decaying dynamical system. A dynamical system in the simple case of a deterministic linear system can be described as:

$$\underline{x}(t) = A\underline{x}(t-1) \quad (9)$$

In which A is a square matrix.

Here we simulate a system of 13 variables, all of them initialized from a $u(0.5, 2)$ distribution. The first 6 variables (x_1, \dots, x_6) are evolved through a linear deterministic process as in (9) in which A is a square 6×6 matrix initialized as:

$$A = \begin{bmatrix} a_{11} & 0 & 0 & 0 & a_{15} & 0 \\ a_{21} & a_{22} & 0 & 0 & 0 & a_{26} \\ 0 & a_{32} & a_{33} & 0 & 0 & 0 \\ 0 & a_{42} & 0 & a_{44} & 0 & 0 \\ 0 & 0 & a_{53} & a_{54} & a_{55} & 0 \\ a_{61} & a_{62} & 0 & 0 & 0 & a_{66} \end{bmatrix} \quad (10)$$

In which every non-zero a_{ij} is randomly taken from a distribution $u(0.75, 1.25)$. Then the matrix A is divided by $5 * \lambda_{\max}(A)$ in which $\lambda_{\max}(A)$ is the greatest eigenvalue of A . It's done to make sure that all the variables decay exponentially to 0. After initialization, the matrix A is kept constant throughout the development process, i.e. it doesn't change with time t .

The other 7 variables (x_7, \dots, x_{13}) are random independent Gaussian variables.

In this experiment, we simulate the system described above for various numbers of time-steps, keeping the standard deviation of the Gaussian variables at $\sigma = 0.1$, and applied both RDI and uRDI algorithms to infer the true causal inferences. Then we calculate the AUC values, the results are shown in Figure 4c. As we can see, the uRDI algorithm outperforms RDI by a margin of 0.1 in terms of AUC.

V. FUTURE DIRECTIONS

In this section, we will describe some promising directions for further investigation.

- 1) *Quantifying causal strength:* As pointed out earlier, potential conditional mutual information can be used as a metric for quantifying causal strength when the graph is a simple three node network (shown in Figure 1a). However, further work is needed in order to generalize the definition to deduce the causal strength of an edge or a set of edges in an arbitrary graph, akin to the formulation in [2] and to study the relative advantages and disadvantages of such a formulation.
- 2) *Discrete qCMI estimators:* It has been shown in recent work that such estimators are not optimal even for determining mutual information in the discrete alphabet case [31], [32], [33]. A very interesting question is how such minimax-rate optimal estimators can be developed in the potential measures problem.
- 3) *maxCMI estimation:* While we have developed efficient estimators for qCMI, in maxCMI, there is a further maximization over potential distributions q , which leads to some interesting interactions between estimation and optimization. Recent work has studied estimation of Shannon capacity on continuous alphabets, however, the formulation is not convex leading to possible local minima [14]. Further work is needed in order to find provably optimal estimators for maxCMI in the continuous case.
- 4) *Other conditional measures:* Recent work [15] has used strong data processing constants as a way for quantifying dependence between two variables, with relationships to information bottleneck. These measures depend *partially* on the factual measure p_X , and are implicitly regularized. One direction of future work is

to develop multi-variable versions of such estimators to estimate the strength of conditional independence, for example.

- 5) *Multivariable measures*: Develop estimators that can handle more general multi-variable information measures including total correlation [34] and multi-variate mutual information [35].
- 6) *Ensemble estimation*: Another approach exploiting k-nearest-neighbors for mutual information is the so-called ensemble estimation approach, where estimators for different k are combined together to get a stronger estimator, with fast convergence [36]. An interesting direction of research is to obtain ensemble estimators for potential measures.

VI. ACKNOWLEDGEMENT

The authors would like to thank Xiaojie Qiu and Cole Trapnell for discussions that triggered this work. This work was supported in part by NSF Career award (grant 1651236) and NIH award number R01HG008164.

REFERENCES

- [1] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [2] D. Janzing, D. Balduzzi, M. Grosse-Wentrup, and B. Schölkopf, “Quantifying causal influences,” *The Annals of Statistics*, pp. 2324–2358, 2013.
- [3] P. Spirtes, C. N. Glymour, and R. Scheines, *Causation, prediction, and search*. MIT press, 2000.
- [4] P. W. Holland, C. Glymour, and C. Granger, “Statistics and causal inference,” *ETS Research Report Series*, vol. 1985, no. 2, 1985.
- [5] A. P. Dawid, “Conditional independence in statistical theory,” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 1–31, 1979.
- [6] C. W. Granger, “Investigating causal relations by econometric models and cross-spectral methods,” *Econometrica: Journal of the Econometric Society*, pp. 424–438, 1969.
- [7] P. Geiger, K. Zhang, B. Schoelkopf, M. Gong, and D. Janzing, “Causal inference by identification of vector autoregressive processes with hidden components,” in *International Conference on Machine Learning*, pp. 1917–1925, 2015.
- [8] H. Hosseini, S. Kannan, B. Zhang, and R. Poovendran, “Learning temporal dependence from time-series data with latent variables,” in *Data Science and Advanced Analytics (DSAA), 2016 IEEE International Conference on*, pp. 253–262, IEEE, 2016.
- [9] M. Eichler, “Graphical modelling of multivariate time series,” *Probability Theory and Related Fields*, vol. 153, no. 1-2, pp. 233–268, 2012.
- [10] C. J. Quinn, N. Kiyavash, and T. P. Coleman, “Directed information graphs,” *IEEE Transactions on information theory*, vol. 61, no. 12, pp. 6887–6909, 2015.
- [11] J. Sun, D. Taylor, and E. M. Boltt, “Causal network inference by optimal causation entropy,” *SIAM Journal on Applied Dynamical Systems*, vol. 14, no. 1, pp. 73–106, 2015.
- [12] A. Rahimzamani and S. Kannan, “Network inference using directed information: The deterministic limit,” in *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on*, pp. 156–163, IEEE, 2016.
- [13] X. Qiu, S. Ding, and T. Shi, “From understanding the development landscape of the canonical fate-switch pair to constructing a dynamic landscape for two-step neural differentiation,” *PloS one*, vol. 7, no. 12, p. e49271, 2012.
- [14] W. Gao, S. Kannan, S. Oh, and P. Viswanath, “Causal strength via shannon capacity: Axioms, estimators and applications,” in *Proceedings of the 33rd International Conference on Machine Learning*, 2016.
- [15] H. Kim, W. Gao, S. Kanan, S. Oh, and P. Viswanath, “Discovering potential correlations via hypercontractivity,” *arXiv preprint arXiv:1709.04024*, 2017.
- [16] V. Anantharam, A. Gohari, S. Kamath, and C. Nair, “On maximal correlation, hypercontractivity, and the data processing inequality studied by erkip and cover,” *arXiv preprint arXiv:1304.6133*, 2013.
- [17] Y. Polyanskiy and Y. Wu, “Strong data-processing inequalities for channels and bayesian networks,” in *Convexity and Concentration*, pp. 211–249, Springer, 2017.
- [18] N. Tishby, F. C. Pereira, and W. Bialek, “The information bottleneck method,” *arXiv preprint physics/0004057*, 2000.
- [19] W. Gao, S. Kannan, S. Oh, and P. Viswanath, “Conditional dependence via shannon capacity: Axioms, estimators and applications,” *arXiv preprint arXiv:1602.03476*, 2016.
- [20] R. Kidambi and S. Kannan, “On shannon capacity and causal estimation,” in *Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on*, pp. 988–992, IEEE, 2015.
- [21] J. M. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf, “Distinguishing cause from effect using observational data: methods and benchmarks,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1103–1204, 2016.
- [22] L. Devroye and C. S. Penrod, “The consistency of automatic kernel density estimates,” *The Annals of Statistics*, pp. 1231–1249, 1984.
- [23] S. J. Sheather and M. C. Jones, “A reliable data-based bandwidth selection method for kernel density estimation,” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 683–690, 1991.
- [24] A. Kraskov, H. Stögbauer, and P. Grassberger, “Estimating mutual information,” *Physical review E*, vol. 69, no. 6, p. 066138, 2004.
- [25] S. Khan, S. Bandyopadhyay, A. R. Ganguly, S. Saigal, D. J. Erickson III, V. Protopopescu, and G. Ostroumov, “Relative performance of mutual information estimation methods for quantifying the dependence among short and noisy data,” *Physical Review E*, vol. 76, no. 2, p. 026209, 2007.
- [26] L. Kozachenko and N. N. Leonenko, “Sample estimate of the entropy of a random vector,” *Problemy Peredachi Informatsii*, vol. 23, no. 2, pp. 9–16, 1987.
- [27] W. Gao, S. Oh, and P. Viswanath, “Demystifying fixed k-nearest neighbor information estimators,” in *Information Theory (ISIT), 2017 IEEE International Symposium on*, pp. 1267–1271, IEEE, 2017.
- [28] W. Gao, S. Kannan, S. Oh, and P. Viswanath, “Estimating mutual information for discrete-continuous mixtures,” *arXiv preprint arXiv:1709.06212*, 2017.
- [29] S. Frenzel and B. Pompe, “Partial mutual information for coupling analysis of multivariate time series,” *Physical review letters*, vol. 99, no. 20, p. 204101, 2007.
- [30] A. Rahimzamani and S. Kannan, “Potential conditional mutual information: Estimators and properties,” *arXiv preprint*, 2017.
- [31] G. Valiant and P. Valiant, “Estimating the unseen: an $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new clts,” in *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 685–694, ACM, 2011.
- [32] J. Jiao, K. Venkat, Y. Han, and T. Weissman, “Minimax estimation of functionals of discrete distributions,” *IEEE Transactions on Information Theory*, vol. 61, no. 5, pp. 2835–2885, 2015.
- [33] Y. Wu and P. Yang, “Minimax rates of entropy estimation on large alphabets via best polynomial approximation,” *IEEE Transactions on Information Theory*, vol. 62, no. 6, pp. 3702–3720, 2016.
- [34] S. Watanabe, “Information theoretical analysis of multivariate correlation,” *IBM Journal of research and development*, vol. 4, no. 1, pp. 66–82, 1960.
- [35] C. Chan, A. Al-Bashabsheh, J. B. Ebrahimi, T. Kaced, and T. Liu, “Multivariate mutual information inspired by secret-key agreement,” *Proceedings of the IEEE*, vol. 103, no. 10, pp. 1883–1913, 2015.
- [36] K. R. Moon, K. Sricharan, and A. O. Hero III, “Ensemble estimation of mutual information,” *arXiv preprint arXiv:1701.08083*, 2017.