Information-Theoretic Implications of Classical and Quantum Causal Structures (Extended Abstract of [5] and [6])

Rafael Chaves, Christian Majenz, L. Luft, T. O. Maciel, D. Janzing, B. Schölkopf, David Gross

Inferring causal relationships from empirical data is one of the main goals of science. To that aim a simple but crucial observation is that the correlations that can be observed between a set of variables depend on the causal structure underpinning them. Within that context, during the past year we have developed and formalized a new research program for the study of causal relations between classical as well as quantum variables. What we propose is an information-theoretic framework for computing constraints that a causal structure can give rise to. We illustrate the power and generality of our method by applying it to a variety of classical and quantum scenarios: i) inferring the direction of causation from marginal observations, ii) the quantification of direct causal influence in classical models, iii) the derivation of constraints implied by the topology of distributed quantum architectures and iv) a strengthened version of the information causality principle.

The root of this research program lies in previous work on Bell's theorem [1–4], that we have adapted and generalized for its use in classical causal inference problems [5] and that in our most recent work [6] has been tailored to deal with quantum causal inference problems as well. The undergoing evolution of this program highlights a fruitful interplay between the causal inference literature and problems in quantum information, in particular nonlocality, a connection which is increasingly appreciated among quantum physicists.

It is an oft-repeated piece of advice that "correlation does not imply causation". In spite of this cautious remark, the emerging field of causal inference aims exactly that: to identify cause-effect relationships given some observed correlations. The fact that empirical data *can* contain information about causation rather than mere correlation had been the subject of a fiercely debate until the seminal works of Pearl [7]. Since then it has become popular to use directed acyclic graphs (DAGs) to represent *causal structures*, with random variables as nodes and arrows meaning direct causal influence [7, 8]. Such a causal model offers a means of *explaining* dependencies between *classical variables*, by specifying the process that gave rise to them. More formally, variables $X_1, \ldots, X_n$ form a *Bayesian network* with respect to a DAG, if every variable X_i depends only on its graph-theoretic parents $\mathrm{pa_i}$. This is the case [7, 8] if and only if the distribution factorizes as in

$$p(x_1, \ldots, x_n) = \prod_{i=1}^{n} p(x_i | x_{\mathrm{PA}_i}). \tag{1}$$

One can then ask the following fundamental question: *Given a subset of classical variables, which correlations between them are compatible with a given causal structure?* Arguably, this problem have wide applications in the classical treatment of causal models, such as statistics and machine learning, and via these, in genetics, economics and other disciplines wherein causal discovery plays a prominent role.

Aside its fundamental role in classical causal inference, this problem also appears in many related quantum mechanical problems. To illustrate, we highlight that the reasoning employed above is close to a core question in the foundations of quantum mechanics: what is the set of quantum mechanical correlations compatible with a presumed causal structure? Since Bell's theorem [9], we know that our classical conceptions of causal relations must be taken with care, as they fail, for example, to commit with the results obtained in some quantum experiments, the phenomenon known as quantum *nonlocality*. In this case one must deal with quantum generalizations of causal structures, where nodes are now allowed to represent either quantum or classical systems, and edges may represent quantum operations. An important conceptual difference to the purely classical setup is rooted in the fact that quantum operations disturb their input. Put differently, quantum mechanics does not assign a joint state to the input and the output of an operation. Therefore, there is in general no analogue to (1), i.e., a global density operator for all nodes in a quantum causal structure

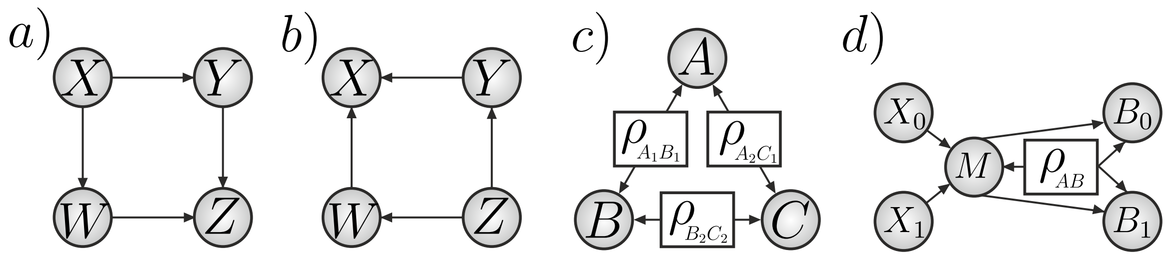

Figure 1: **(a)** and **(b)** DAGs with opposite causal directions and that can be distinguished even if only pairwise information is available. **(c)** An example of distributed architecture involving bipartite entangled states. Each of the underlying quantum states can connect at most two of the observable variables, what implies a non-trivial monogamy of correlations, for example, that $I(A:B)+I(A:C)\leq H(A)$. **d** The quantum causal structure associated with the information causality principle.

cannot be defined. However, if we pick a set of nodes that do coexist (e.g. because they are classical, or because they are created at the same instance of time), then we can again ask: *Which correlations, between coexisting nodes, can result from a given quantum causal structure?*

To begin with, we note that most often and for a variety of reasons, one is interested in testing causal structures with latent ("hidden") variables, that is, joint observations might be constrained. For instance, in quantum non-locality these reasons are physical: random variables corresponding to non-commuting observables cannot always be jointly measured. In general, there might also be practical reasons, for example: we may have no access to the variable describing the genetic disposition of a patient to become both a smoker and to develop lung cancer (not the least because we do not know whether such a genetic influence exists). In that case, we are faced with the problem of characterizing the set of *marginal distributions* a given causal structure can give rise to. If an observed distribution lies outside the set of marginals of a candidate model, then that model can be rejected as an explanation of the data. Unfortunately, even in the purely classical case, it is widely appreciated that causal structures involving latent variables impose highly non-trivial constraints on the distributions compatible with it [10–12].

The technical difficulties stem from the fact that the the causal relations implied by a given causal structure are encoded in *conditional independencies* (CI), that amount to non-trivial algebraic conditions on probabilities. More precisely, the marginal regions are semi-algebraic sets that can, in principle, be characterized by a finite number of polynomial equalities and inequalities [10]. However, it is safe to say that in practice, algebraic statistics is still limited to very simple models.

In order to circumvent this problem, we proposed in [5] an information-theoretic approach for causal inference, describing a systematic and practical algorithm for answering the classical question. The basic insight is that the *algebraic* condition $p(x,y)=p_1(x)p_2(y)$ for independence becomes a *linear* relation $H(X,Y)=H(X)+H(Y)$ on the level of entropies. This opens up the possibility of using computational tools such as linear programming to find marginal constraints – which contrasts pleasantly with the complexity of algebraic methods that would otherwise be necessary. In the subsequent work [6] we have generalized the classical framework in order to deal with quantum generalizations of causal structures.

Our main message is that a significant amount of information about causation is contained in the entropies of observable variables and that there are relatively simple and systematic ways of unlocking that information, both in classical [5] and quantum realms [6]. We will make that case by discussing a great variety of classical and quantum applications, which we briefly summarize here.

Distinguishing causal directions.— If the full joint distribution of (X,Y,W,Z) is accessible, then the two networks in Fig. 1 a) and b) can be distinguished, thus revealing the "direction of causation". Here we show that the same is possible even if only two variables are jointly accessible at any time. Fundamentally, inferring the direction of causation between two variables is a notoriously thorny issue, hence it is far from trivial that it can be done from information about several pairwise distributions. On the practical side one can cite Mendelian randomization as a good example where the joint distribution on all variables is often unavailable [13]. In general, the variables in Fig. 1 a) and b) could represent properties of short-lived or fragile specimen that decay before several tests can be completed, or that are degraded by multiple invasive probes.

Quantifying (classical) causal influences.— Various measures of (classical) causal influence have been studied in the literature. Of particular relevance to us is the one recently introduced in [14]. The main idea is that the causal strength $\mathcal{C}_{X\to Y}$ between a variable X on another variable Y should measure the impact of an intervention that removes the arrow between them. It follows that

$$\mathcal{C}_{X\to Y} \geq I(X:Y|\mathrm{PA}_Y^X), \tag{2}$$

where PA_Y^X stands for the parents of variable Y other than X. Because the quantity $I(X:Y|\mathrm{PA}_Y^X)$ appears naturally in our description, it readily allows us to bound the causal strength $\mathcal{C}_{X\to Y}$ in arbitrary classical causal structures. We then go on to present two corollaries of this result: First, it follows that the degree of violation of an entropic inequality often carries an operational meaning. Second, under some assumptions, the finding will allow us to introduce a novel way of distinguishing dependence created through common ancestors from direct causal influence. That is, in these cases we can safely state, opposed to the common folklore, that "(some kinds of)correlations do imply causation".

Distributed quantum architectures.—Consider a scenario where, in a first step, several few-body quantum states are distributed among a number of parties. In a second step, each party processes those parts of the states it has access to (e.g. by performing a coherent operation or a joint measurement). Such setups are studied e.g. in distributed quantum computing [15, 16], quantum networks [17], quantum non-locality [18], and quantum repeaters [19]. Which limits on the resulting correlations are implied by the network topology alone? Our framework can be used to compute these systematically. We prove, e.g., certain *monogamy relations* between the correlations that can result from measurements on distributed quantum states (see Fig. 1).

Information causality.—The "no-signalling principle" alone is insufficient to explain the "degree of non-locality" exhibited by quantum mechanics [20]. This has motivated the search for stronger, operationally motivated principles, that may single out quantum mechanical correlations [21–27]. One of these is *information causality* (IC) [22] which posits that an m bit message from Alice to Bob must not allow Bob to learn more than m bits about a string held by Alice. A precise formulation of the protocol involves a relatively complicated quantum causal structure (Fig. 1b). It implies an information-theoretic bound on the mutual information between bits X_i held by Alice and guesses Y_i of these by Bob [22], namely: $\sum_{i=1}^n I(X_i:Y_i) \leq H(M)$. Here, we note that the IC setup falls into our framework and we put the machinery to use to generalize and strengthen it. We will show that by taking additional information into account, our strengthened IC principle given by

$$\sum_{i=1}^{n} I(X_i:Y_i,M) + \sum_{i=2}^{n} I(X_1:X_i|Y_i,M) \leq H(M), \tag{3}$$

can identify super-quantum correlations that could not have been detected in the original formulation.

These applications aside, we believe that the main contribution of this research program is to highlight the power of systematically analyzing entropic marginals. A number of future directions for research immediately suggest themselves. In the classical realm it would be fruitful to further explore our method to test highly non-trivial models (inducing semi-algebraic sets) such as in the "homophily versus contagion" problem arising in the study of social networks [12]. In the quantum domain, it will likely be fruitful to consider multi-partite versions of information causality or other information theoretical principles, to further look into the operational meaning of entropy inequality violations and possibly to look for applications in many body problems (for example, with nearest neighbours interactions).

[1] R. Chaves and T. Fritz. Entropic approach to local realism and noncontextuality. *Phys. Rev. A*, 85:032113, Mar 2012.

[2] T. Fritz and R. Chaves. Entropic inequalities and marginal problems. *IEEE Trans. Inform. Theory*, 59:803, Feb 2013.

[3] R. Chaves. Entropic inequalities as a necessary and sufficient condition to noncontextuality and locality. *Phys. Rev. A*, 87:022102, Feb 2013.

[4] R. Chaves, L. Luft, and D. Gross. Causal structures from entropic information: Geometry and novel scenarios. *New J. Phys.*, 16:043001, 2014.

[5] R. Chaves, L. Luft, T. O. Maciel, D. Gross, D. Janzing, and B. Schölkopf. Inferring latent structures via information inequalities. 2014. Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence (UAI 2014), pp. 112 - 121, AUAI Press.

[6] Rafael Chaves, Christian Majenz, and David Gross. Information-theoretic implications of quantum causal structures. *arXiv:1407.3800 (Under review in Nature Communications)*, 2014.

[7] J. Pearl. *Causality*. Cambridge University Press, Cambridge, 2009.

[8] P. Spirtes, N. Glymour, and R. Scheienes. *Causation, Prediction, and Search, 2nd ed.* The MIT Press, 2001.

[9] J. S. Bell. On the Einstein–Podolsky–Rosen paradox. *Physics*, 1:195, 1964.

[10] D. Geiger and C. Meek. Graphical models and exponential families. In *Proceedings of the 14th conference on Uncertainty in Artificial Intelligence*, pages 156–165, 1998.

[11] J. Tian and J. Pearl. On the testable implications of causal models with hidden variables. In *Proceedings of the 28th conference on Uncertainty in Artificial Intelligence*, pages 519–527, 2002.

[12] G. Ver Steeg and A. Galstyan. A sequence of relaxations constraining hidden variable models. In *Proceedings of the 27th conference on Uncertainty in Artificial Intelligence*, 2011.

[13] Vanessa Didelez and Nuala Sheehan. Mendelian randomization as an instrumental variable approach to causal inference. *Stat Meth Med Res*, 16(4):309–330, 2007.

[14] D. Janzing, D. Balduzzi, M. Grosse-Wentrup, and B. Schölkopf. Quantifying causal influences. *The Annals of Statistics*, 41(5):2324–2358, 10 2013.

[15] R. Van Meter. *Quantum Networking*. Wiley, 2014.

[16] Harry Buhrman and Hein Röhrig. Distributed quantum computing. In *Mathematical Foundations of Computer Science 2003*. Springer, 2003.

[17] Antonio Acín, J Ignacio Cirac, and Maciej Lewenstein. Entanglement percolation in quantum networks. *Nature Physics*, 3(4):256–259, 2007.

[18] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. Bell nonlocality. *Rev. Mod. Phys.*, 86:419–478, Apr 2014.

[19] Nicolas Sangouard, Christoph Simon, Hugues de Riedmatten, and Nicolas Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Rev. Mod. Phys.*, 83:33–80, Mar 2011.

[20] Sandu Popescu and Daniel Rohrlich. Quantum nonlocality as an axiom. *Foundations of Physics*, 24(3):379–385, 1994.

[21] Wim van Dam. *Nonlocality & communication complexity*. PhD thesis, Faculty of Physical Sciences, University of Oxford, 1999.

[22] Marcin Pawłowski, Tomasz Paterek, Dagomir Kaszlikowski, Valerio Scarani, Andreas Winter, and Marek Żukowski. Information causality as a physical principle. *Nature*, 461(7267):1101–1104, 2009.

[23] David Gross, Markus Müller, Roger Colbeck, and Oscar CO Dahlsten. All reversible dynamics in maximally nonlocal theories are trivial. *Physical Review Letters*, 104(8), 2010.

[24] Gonzalo de la Torre, Lluís Masanes, Anthony J Short, and Markus P Müller. Deriving quantum theory from its local structure and reversibility. *Physical Review Letters*, 109(9), 2012.

[25] Miguel Navascués and Harald Wunderlich. A glance beyond the quantum model. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 466(2115):881–890, 2010.

[26] Tobias Fritz, Ana Belén Sainz, Remigiusz Augusiak, J Bohr Brask, Rafael Chaves, Anthony Leverrier, and A Acín. Local orthogonality as a multipartite principle for quantum correlations. *Nature communications*, 4, 2013.

[27] Miguel Navascués, Yelena Guryanova, Matty J Hoban, and Antonio Acín. Almost quantum correlations. 2014. arXiv:1403.4621.