Causal Relationships in Longitudinal Observational Data: An Integrative Modelling Approach

Claudinei E. Biazoli Jr.^{1,2}, João R. Sato², Michael Pluess¹

1-Department of Biological and Experimental Psychology, Queen Mary University of London, UK.

2- Center of Mathematics, Computing and Cognition, Universidade Federal do ABC, Brazil.

Authors' Note: CEB and MP have received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement #896021, CARINGTrauma project. JRS is supported by São Paulo Research Foundation - FAPESP Grant #2018/21934-5

Abstract

Much research in psychology relies on data from observational studies that traditionally do not allow for causal interpretation. However, a range of approaches in statistics and computational sciences have been developed to infer causality from correlational data. Based on conceptual and theoretical considerations on the integration of interventional and time-restraint notions of causality, we set out to design and empirically test a new approach in order to identify potential causal factors in longitudinal correlational data. A principled and representative set of simulations and an illustrative application to identify early-life determinants of cognitive development in a large cohort study are presented. The simulation results illustrate the potential but also the limitations for discovering causal factors from observational data. In the illustrative application, plausible and reasonably well-established early life determinants of cognitive abilities in 5 year old children were identified. Based on these results, we discuss the possibilities of using exploratory causal discovery in psychological research but also highlight its limits and potential misuses and misinterpretations.

Keywords: causal inference; causal discovery; cognitive development; machine learning; Granger causality.

1. Introduction

Many highly relevant questions in psychological research, such as determining the long-term effects of childhood environment on cognitive development, can only be addressed by answering questions about complex causal relations (Hamaker et al., 2020). However, non-experimental fields in psychology that heavily rely on large cohort studies have only recently begun to explicitly incorporate causal inference language and tools (Kline, 2015; Rohrer, 2018). Inferring causality from observational data is not often applied in psychology (Grosz et al., 2020), mainly due to legitimate worries about the perils of over-interpreting correlational evidence as indicative of causal effects. This scenario is rapidly changing with the increasing adoption of more contemporary ways to conceptualize and quantify causality, based mainly on recent methodological advances in statistics and computational sciences (Pearl, 2009a; Pearl & Mackenzie, 2018; D. B. Rubin, 2010a; Spirtes & Zhang, 2016; Hofman et al., 2021).

Establishing causal relations is a challenging task and the contemporary landscape of approaches for causal inquiry includes divergent and contentious theoretical frameworks. In a recent perspective on the emerging field of computational social science, Hofman et al. (2021) identify a line of tension between the epistemic values and scientific goals of constructing accurate *predictive* data-driven models and that of formulating *explanatory* theory-driven accounts. This tension arises between methodological advances in what have been called the "computational revolution" (Lazer et al., 2020; Watts, 2007) in predictive power and the "causal revolution" (Pearson et al., 2018) in structural modeling. Arguing for taking both predictive and causal inquiries seriously, Hofman et al. (2021) propose the development of *integrative models*.

Here, we sought to design and empirically test an example of such an integrative method. For this, we build on recent theoretical and methodological advances in both the structural causal models (SCM; Kline, 2015; Pearl, 2009a) framework and predictive modelling with machine-learning (Eronen, 2020), as well as formal results on the possibility

of integrating model-free and model-dependent notions of causality (Eichler, 2012; Eichler & Didelez, 2010).

The more general account for causal inquiry, Structural Causal Modeling (SCM) relies on diagrammatic representations of one's causal assumptions using tools from graph theory. One of the many advantages of this representation is that it allows to clearly differentiate two tasks of causal research: *causal discovery* and *causal inference or identification*. Causal discovery is the task of determining the causal structure of a problem or, in other words, finding the causal diagram that best describes the process generating the data. Causal inference or identification is the task of calculating the strength of the causal effects for a known or presumed causal structure. As will be described in more detail below, the task of causal identification can be performed, as long as certain conditions are met, in a fully automated way and tools for doing so are available in common statistical packages. But finding the causal structure is a much harder task and depends on expertise, domain knowledge and ultimately on untestable assumptions.

Specifically, we set out to construct and test an exploratory discovery method useful to search for the most relevant putative causes in large cohort datasets, arguably reflecting complex generative processes where multiple causal factors are expected. In other words, we aimed to expand researchers' toolbox for the hard causal discovery task by adding a tool to search for *potential direct and relevant causal factors*. The set of candidate causal factors discovered using this tool might then be prioritized for further causal investigations in the same dataset or using further data and lines of evidence.

Additionally, we aim to echo previous efforts on introducing more clear causal reasoning in psychological inquiry (Rohrer, 2018). For these purposes, we first present a brief overview of SCM and an introduction to the representation of causal structures using Directed Acyclic Graphs (DAGs). Then, a short presentation of currently available causal discovery algorithms and the foundations for integrating predictive frameworks to discover causal links in a data-driven way is provided. Results of representative simulated scenarios applying the new tool and a concrete illustrative application of searching for the early-life determinants of cognitive development are presented. Finally, we discuss some possible ways

of using exploratory causal discovery in psychological research but also highlight its limits and potential misuses and misinterpretations.

1.1 Structural Causal Modelling

Randomized controlled trials (RCTs) are the gold standard procedure to directly test specific causal hypotheses. In a proper RCT, a correlation might indeed be a precise measure of the causal effect of a treatment on an outcome. However, conducting an RCT is often impossible, impractical or unethical for a wide range of fundamental and interesting causal questions in psychology. Treatments, exposures or factors that are protracted or that causally influence outcomes occurring much later in development, for instance, usually cannot be studied using RCTs (Ness et al., 2011; Richmond et al., 2014). Moreover, psychological interventions are usually fat-handed, i.e., tend to change several variables at once, further difficulting determining causal effects with RCTs (Eronen, 2020). RCTs might also be of limited use to uncover complex causal structures as they are tailored to determine the strength of only one or a few causal effects (Grosz et al., 2020).

A conventional approach to infer causal effects in non-experimental settings consists of attempting to statistically simulate RCTs in observational data (McCaffrey et al., 2004; Rosenbaum & Rubin, 1983; D. B. Rubin, 2010b). This approach is based on two main ideas: (1) to consider the causal inference problem analogous to a missing data problem in which a counterfactual potential outcome is not observed (Donald B. Rubin, 1978), and (2) ascribing a probability for an individual to receive treatment from the observed variables (D. B. Rubin, 2010b). For this, a model for the assignment mechanism, i.e. for how individuals end up to receive one treatment or another, must be in place. With such a model, it is possible to calculate the probability of a given individual to receive a treatment, the so-called propensity score (Harder et al., 2010; McCaffrey et al., 2004; Rosenbaum & Rubin, 1983). The data is then treated as coming from a randomized experiment but with missing outcomes. Though useful for a range of causal inference problems, this Potential Outcomes (PO) framework might not be optimal to investigate complex multicausal structures. Limitations of matching using propensity scores have been recently highlighted (King & Nielsen, 2019), particularly

its susceptibility to bias due to omitting confounders or including covariates that should not be included in a given model (Sjölander, 2009).

Causal inference problems can also be tackled using Structural Equation Modelling (SEM) (Kline, 2015; D. B. Rubin, 2010b). In SEM, causal assumptions can be thought of as being reflected in the choices of which structural coefficients or covariates are set to zero or allowed to assume non-zero values and then estimated from data (Bollen & Pearl, 2013). Expanding and integrating PO and SEM frameworks, Pearl proposed and developed the Structural Causal Modelling (SCM) framework (Pearl, 2000, 2009a, 2009b, 2014).

Applied in computational sciences and artificial intelligence over the last few decades, SCM provides a principled and general framework for causal inference from observational data. The use of SCM has gained traction in many fields in the last decade and it is arguably the current framework of choice when inferring causality from non-experimental designs in epidemiology (Hammerton & Munafò, 2021; Lawlor et al., 2016; Tennant et al., 2021). However, SCM frameworks and tools are still largely underappreciated in psychological research, despite increasing recent interest (Austin et al., 2019; Grosz et al., 2020; Hamaker et al., 2020; Rohrer, 2018).

1.2 SCM and the language of Directed Acyclic Graphs (DAGs)

The mathematical language and tools of graph theory provide a clear framework to construct, test and communicate structural causal models (Pearl, 2000; Spirtes et al., 2001). Many excellent general accounts and introductory texts on the mathematical bases and on how to use DAGs to investigate causality are currently available. Of worth mentioning are the introductory texts from Judea Pearl (Pearl, 2000; Pearl & Mackenzie, 2018), the main proponent of these approaches. Julia Rohrer has recently contextualized the potential implications of causal inference approaches to psychological research, including an excellent introduction to DAGs (Rohrer, 2018).

A graph is formed by nodes and edges connecting those nodes. In general, these objects can be classified as directed or undirected, weighted or unweighted, and cyclic or

acyclic. The edges can be thought of as channels that convey information between nodes. A directed graph is thus one in which the direction of information flow is specified and the edges can be represented by arrows. If it is impossible to get back to any node following the paths formed by the arrows, the graph is said to be acyclic. For representing causal structures, directed acyclic graphs require the usually reasonable assumption that an effect cannot be the cause of any of its causes. In other words, the effect "listens" to the causes but no information is conveyed from the effect to the causes. Finally, a graph is unweighted when the strength of the connections is binary, i.e., can only assume the values 0, in which the arrow is absent, or 1, when it is present. The task of *causal discovery* is thus that of finding the unweighted DAG that best represents the causal structure of a problem (Figure 1). *Causal inference or identification*, on the other hand, can be thought of as the task to assign the weights of the relations between a cause and an effect, given a known causal structure represented by an unweighted DAG.

In sum, causal structures can be represented by DAGs in which one's causal assumptions are clearly and explicitly conveyed. Importantly, the use of DAGs to represent causal structure greatly clarifies the concept of confounder (Pearl & Mackenzie, 2018b). Such a clear conceptualization associated with formalized procedures allows for a principled choice of the variables that should or should not be included in a model in order to make unbiased causal inference. Three possible combinations of three variables (represented as edges) and the links between them (represented by arrows) provide the basic structural motives of a DAG. These three structural motives are called the path, the fork and the collider. Minimal DAGs with these structures are depicted in Figure 1 and are more comprehensively explored and illustrated in the simulation scenarios.

Figure 1: The three basic structural motives of DAGs from which the criteria for selecting a minimal set of covariates necessary to infer causality can be determined. In the pipe or causal path structure, the middle variable (CV2) is a direct cause of the outcome (OUT) and is also called a mediator of the indirect causal relation between the first variable (CV1) and the outcome. Directed arrows between two variables represent direct or *contrived* causal relations. Note that in the pipe structure, there is no contrived or direct causal relation between cv1 and out, but only an *indirect* causal relation through the mediator CV2. In the fork structure, the middle variable is the cause of two effects: the outcome of interest (out) and a third variable (CV1). In this case, CV2 is called a confounder since not taking this variable into account would induce a spurious relation between the third variable cv1 and the outcome of interest. Finally, in the collider structure, a variable (CV2) is caused by two other independent variables (CV1 and CV3).

A classical confounder is represented by the fork structure. If a confounder variable is omitted in a regression model, for instance, it will induce a spurious correlation between a third variable and the outcome (Figure 3-B). For the other two fundamental structures, the opposite is true: including either a collider or a mediator in a regression model will bias the inference of causal effects (Figure 3-A/C). These consequences of not including confounders or erroneously including mediators and colliders in a regression model are further illustrated in the simulations of the causal discovery method below.

But a lot of relevant problems in psychological research are likely to have a rather complex causal structure. Complex DAGs with multiple interacting causes and different relations between observed and unobserved variables would represent such complex structures. Despite the complexity of the DAG and importantly for their practical application, a minimal set of covariates, when such a set exists, can be automatically determined and tools to do so are available in the DAGitty (Knüppel & Stang, 2010; Textor et al., 2011) and the DAG Program (Knüppel & Stang, 2010) packages. The pcalg R package provides a wide range of tools for causal inference using DAGs (Kalisch et al., 2012). Including the minimal

set of covariate in a conventional regression model, for instance, yields unbiased estimates of the causal effect of a given variable in the outcome.

For some causal structures, there is no such a minimal set of covariates or it cannot be algorithmically determined using the so-called backdoor criterion (Knüppel & Stang, 2010; Pearl, 2014) (Textor et al., 2011). In other words, some causal effects cannot be inferred from non-experimental data. These causal effects are said to be *non-identifiable* and cannot be inferred from the observed variables independently of some unknown, unmeasured or not measurable variables. Critically, structural causal model framework also provides principled ways to evaluate the identifiability of causal effects from non-experimental data based on the DAG representation of the problem (Pearl, 2014).

1.3 - Causal discovery

As mentioned, DAGs language allows one to clearly differentiate two tasks on causal research: causal discovery and causal inference or identification. Causal discovery is the task of finding the DAG that describes the process of generating the data while causal inference or identification is the task of calculating the strength of specific DAG edges. In other words, the aim of a causal discovery algorithm is to recover the causal structure of a problem in a data-driven way (Balzer & Petersen, 2021; Robins & Hernan, 2008), and accurately detect the arrows of the underlying DAG. Thus, the powerful tools of causal inference can be applied to estimate the strength of the known or presumed causal relations.

For many causal problems, determining the causal structure can be quite challenging if possible at all. In general, no single automatic procedure suffices to uniquely discover causal structures but recent advances in Structural Equation Modeling (SEM) have led to the development of many causal discovery data-driven approaches (Spirtes & Zhang, 2016). Among these SEM approaches proposed to *fully* determine causal structure from data, the linear, non-Gaussian and acyclic model (LiNGAM, Shimizu et al., 2006) was a pioneering effort. Non-linear (Hoyer et al., 2009) and post-nonlinear (Zhang & Hyvärinen, 2009) extensions of LiNGAM have also been developed and tested. Though uniquely determining the *full* causal structure of a given problem might not be feasible or possible for all causal

problems, the increasing availability of large high-quality cohorts and advances in computational Machine Learning (ML) algorithms fostered the development of new causal discovery tools (Spirtes & Zhang, 2016).

Machine learning (ML) algorithms are the state-of-the-art approach to solve prediction problems in many fields (Yarkoni & Westfall, 2017). Although prediction is not causation to the same extent that correlation is not causation, predictive and causal inquiries can and should be integrated (Hofman et al., 2021; Spirtes & Zhang, 2016). Here, we argue that such integration is theoretically feasible based on results of convergence between temporal precedence and interventional notions of causality (Eichler, 2012).

Being able to better predict an outcome from a set of different variables does not necessarily inform us which of these variables would be amenable to be modified to effectively change the outcome. In other words, knowing that one variable is relevant for predicting an outcome does not inform the extent to which a *manipulation* or intervention on this variable would affect that outcome. In the structural account of causality, a relationship between a variable and an outcome is defined as causal if a *manipulation* (often referred to as intervention) of the variable is mapped in the outcome.

More generally, the interventional approach to causality builds upon counterfactual reasoning, i.e., hypothesizing what might have happened to the outcome under different scenarios. Counterfactuals are also the base of the Potential Outcome Model presented earlier and this model provides a good example for how counterfactual reasoning might be used to infer cause. In POM, causal inference is performed by considering what would be the outcome given an intervention on a variable. As the intervention, or treatment, is counterfactual, the causal effect is analogous to solving a missing data problem. The Structural Causal Model further advances this principled approach to counterfactuals (Pearl, 2009). As Rohrer et al. (2020) noted, DAGs and counterfactual reasoning are not radical novelties in psychology since these tools are in line with Campbell's approach for identifying and dealing with threats to internal validity (West & Thoemmes, 2010).

In an independent line of development to the interventional notion of causality underlying SCM, the aspect of temporal precedence of causes in relation to effects is emphasized. This alternative theoretical and practical account for inferring causality from observational data was developed and is widely applied in econometrics, where time-series data are common (Eichler, 2012; Eichler & Didelez, 2010; C. Granger, 1986; C. W. J. Granger, 1969). Proposed by Granger, this notion of causality is explicitly defined in terms of temporal and informational relations: being everything else equal in the universe, a variable is said to Granger-cause an outcome if the past of the variable time-series uniquely increases the *prediction* of the future outcome. Granger's definition of causality thus relies on two assumptions: (1) the cause precedes the effect, and (2) the cause necessarily contains unique information that helps to *predict* the effect. Critically, and in contrast to the interventional accounts to causality, determining Granger's causality depends only on prediction. Hence, Granger causality can be straightforwardly integrated in data-driven approaches.

Despite the interventional notion of causality being stronger than the Granger-causality notion (Granger, 1980), recent theoretical results support that Granger-causality can be fruitfully embedded in the more general SCM framework (Eichler, 2012; Eichler & Didelez, 2010). The notion of Granger-causality and the consequent temporal constraint have been incorporated in forms of SEM to improve discovery of causal relations. In fact, the temporal constraint is useful to reduce the search space for causal structure in causal discovery procedures such as LiNGAM. A comprehensive account of the formalized integration between Granger causality and the effects of interventions in time series can be found in Eichler (Eichler, 2012). Particularly, Eichler and Didelez (2010) have shown that the backdoor criterion for DAGs are compatible with the conditions under which Granger-causality, or its lack thereof, allows to compute the effect, or absence of effect, of an intervention in time series data. Given that result, Granger-causality can be used, at least in principle, to *discover direct causal relations*. Hence, we sought to empirically test a data-driven procedure to discover direct causal relations based on Granger-causality. Importantly, the idea was to discover direct causal factors rather than search for the full causal structure. For searching full structures, a series of algorithms are currently available and practical guidelines can be found in (Glymour et al., 2019)

2 - Methods

2.1 - Rationale:

Despite being developed for time-series data, the *idea* of Granger-causality can be expanded to other data structures (Eichler & Didelez, 2010). Critically and as presented in the previous session, theoretical and empirical results support integrating Granger-causality concept in causal discovery frameworks (Eichler & Didelez, 2010; C. Granger, 1986; C. W. J. Granger, 1969). Given these previous formal results (Eichler & Didelez, 2010; C. Granger, 1986; C. W. J. Granger, 1969), we applied the temporal constraint underlying Granger-causality notion to the special case of longitudinal panel data with two data-points in order to explore potential or candidate causal factors. The main ideas for this application were to develop a method suitable to be applied to large-scale cohort studies common in developmental psychology and to illustrate how causal thinking can be embedded in a practical analysis.

Thus, we developed and tested a procedure using the idea of Granger causality for the particular data structures where potential causes are measured at baseline and necessarily precedes effects or outcomes measured at the follow-up. We empirically tested this procedure by simulating data sets constructed in a principled way, including the basic structural motifs (pipe, fork and collider) and complex identifiable and non-identifiable causal structures. In all simulations, we evaluated to which extent direct causal effects and hence known direct causal factors could be retrieved. The same algorithm was then empirically tested in a concrete application of identifying early-life determinants of cognitive abilities at age five in the Millennium Cohort Study (Connelly & Platt, 2014). For the evaluation of the procedure in this application, we compared the candidate causal factors found with the current state-of-art evidence in cognitive development.

2.2 - Description of the algorithm

Fully annotated R codes (R version 4.1.0) using XGBoost as the ML algorithm for both the simulations and application are available in Supplementary Material and might be readily adapted to other empirical problems. In general, we propose an algorithm specified by the following steps (also described in Figure 2):

1. Train a model in part of the sample using a supervised Machine-Learning algorithm to predict an outcome occurring after a set of predictors;
2. Use the trained model to predict the outcome for each individual in the remaining sample;
3. Calculate the prediction errors for each individual prediction (e.g., the absolute difference between the individual predicted and actual outcome);
4. Run a k-fold cross-validation procedure repeating steps 1-3;
5. Perform steps 1-4 including the whole set of independent variables;
6. Perform steps 1-4 with the whole set of predictors except one at each iteration;
7. Use a non-parametric bootstrap to null hypothesis testing, the alternative hypothesis being that the mean of individual predictions using all the independent variables is greater than the mean of individual predictions excluding a given variable.
8. A putative causal link between a given variable and the outcome will be concluded if the null hypothesis is rejected (i.e., if the variable significantly and uniquely increases the prediction accuracy).

Figure 2 - Schematic representation of the causal discovery algorithm. Specifically for ensemble ML methods, like the XGBoost, feature selection and cross-validation procedures are already embedded in the training model phase. Additionally, a five-fold cross-validation was used to out-of-sample prediction, from which the hypotheses on the existence of the causal links were tested with a nonparametric procedure.

As a first step, a procedure that optimizes prediction performance at the individual level is required. For optimizing prediction, supervised learning algorithms are applied. Such algorithms estimate the expected value of an outcome given a set of variables. More recently, ensemble methods have been developed as an extension of supervised methods (Balzer & Petersen, 2021; Naimi & Balzer, 2018). Rather than using a specified supervised ML algorithm, ensemble methods use cross-validation procedures to build optimized combinations of different candidate algorithms. Broadly, cross-validation consists in splitting the data and then training the ML models in the first part of the data before evaluating their prediction performance in the rest of data (i.e., determining the models out-of-sample performance). Out-of-sample performance can then be used to compare and select a model or combination of models. Among ensemble algorithms, the eXtreme Gradient Boost (XGBoost) has been gaining increasing attention in the ML community, winning prediction

competitions for a wide range of problems. Other ensemble methods include Super Learner and random forest (Balzer & Petersen, 2021).

For the presentation of the causal discovery approach, we used XGBoost based on its flexibility and consistent higher prediction power in a wide range of applications. However, it is important to emphasize that any suitable ML algorithm can be used and the goal is to achieve the best prediction performance for the specific problem at hand. Actually, if the best ML algorithm is already known for a particular application, the researcher should apply that one for the four first steps. The last four steps use the *idea* of Granger-causality to explicitly test whether a causal link between any variable and the outcome is probable.

The identification causality here assumes the temporal precedence of the causal factors in relation to the outcome. Hence, in a two wave data set, the features to be explored would be collected at baseline and should be used to predict the outcomes collected at the follow-up. To decide whether a variable measured at the baseline uniquely and significantly increases the prediction of an outcome, hypothesis testing is necessary. Thus a bootstrap test was implemented under the null hypothesis that the mean prediction error for all variables and for all variables except the variable at consideration at each time is the same. The alternative hypothesis for each variable was that the inclusion of a particular variable *increases the mean prediction accuracy* (i.e., decreases the mean error of individual prediction using the ML method). A potential direct causal relation is said to be present if the null hypothesis is rejected. As is the case for choosing the ML algorithms, the hypothesis testing was implemented to illustrate the concept of the procedure in a clear and intuitive way. However, other approaches can be readily implemented and tested using the same framework. In particular, Bayesian extensions can be straightforwardly applied.

3 - Simulations

3.1 - Methods

In general, as in any detection problem, the evaluation of the proposed procedure relies on its ability to recover true presence or absence of relation (true positives and true

negatives), while minimizing falsely discovering absent relations (false positives) or missing actual relations (false negatives). In the causal discovery language, the procedure should be able to recover at least part of the underlying causal structure represented by a DAG (i.e., correctly identify the presence and absence of the edges or directed arrows of a DAG).

Here, we proposed a method to recover the *direct or contrived* relations in a data-driven way, based on previous formal accounts of the relation between interventional and Granger causality frameworks (Eichler & Didelez, 2010). To test the extent to which the method is able to recover direct causal factors, a representative and principled series of simulated scenarios was used. The three fundamental DAG structures and the concept of *identifiability* of a DAG provided the principles for constructing the simulation scenarios. In order to both test the potential and limits of our approach and to illustrate the advantages of causal reasoning in general, we analyzed the same simulated dataset using multiple regression analyses in which the causal structure is not explicitly taken into account (Westreich & Greenland, 2013 ; Westfall & Yarkoni, 2016).

As a first step for all the simulations, toy datasets were generated by independently sampling from four continuous distributions: one for the simulated outcome (out); three for continuous independent variables (CV1, CV2, CV3); and two discrete distributions to generate extra independent variables (DV1 and DV2). Then, linear dependencies between the outcome and the explanatory variables were included according to the chosen causal structure. We report only on continuous outcomes, but the approach can also be used with discrete outcomes. Continuous outcomes were chosen to test the machine-learning algorithm in a more challenging scenario for individual prediction than more usual classification procedures.

To quantify the discovery properties, each simulated scenario was run 500 times and the accuracy (defined as the mean of true positive and true negative discoveries), false positive and false negative discovery ratios for causal relations were calculated. Each simulated dataset comprised 15000 subjects, with 10000 used for training the model and 5000 for prediction. The sample size was chosen to match the Millennium Cohort Study, which was used to illustrate application. To examine the extent to which sample size impacts the

discovery performance, a representative scenario was analyzed for a varying number of subjects (from 300 to 15000, Figure 4-B).

3.2 - Results

Building upon the DAGs language for causal structure, we constructed the three minimal causal structural motives for problems with three variables and two complex graphs, with and without identifiable causal relations, respectively. While Figure 4-A presents the results for all evaluated scenarios with 500 simulations each, Figure 3 shows the results of a single simulation for the three illustrative scenarios focusing on the fundamental structures of the pipe, the fork and the collider.

The causal discovery (GC-XGBoost) and Multiple Regression with the full set of covariates (MR) results are shown for a DAG with a pipe structure ($CV1 \rightarrow CV2 \rightarrow OUT$) and another independent cause ($DV1 \rightarrow OUT$) when all the variables are included in the models (Figure 3-A1) and when the direct causal variable CV2 is not measured or observed (Figure 3-A2). Using the discovery method, all true direct causal relations are correctly identified, which does not occur for the MR approach. Two properties of MR results on a pipe structure are represented in Figure 3A. First, the regression coefficient on scenario A1 recovers the real *strength* of the simulated data. Second, when the variable CV2 is not observed, an association between CV1 and the outcome is observed and its strength is the product of the actual causal relations ($CV1 \rightarrow CV2 * CV2 \rightarrow OUT$: $0.3 * 0.3 = 0.09$). This simple illustration shows how MR can produce accurate causal estimates for a given causal structure and how including a mediator (CV2) biases the inference of an indirect causal relation ($CV1 \rightarrow CV2 \rightarrow OUT$).

A1 - Pipe

Variable	GC-XGBoost		MR		
	p-value	discovery	p-value	discovery	Coeff
CV1	0.93		0.25		0
CV2	<0.001	→	<0.001	→	0.3
CV3	0.72		0.12		0
DV1	<0.001	→	0.46	✗	0
DV2	0.71		0.27		0

A2 - Pipe (mediator not included)

Variable	GC-XGBoost		MR		
	p-value	discovery	p-value	discovery	Coeff
CV1	0.62		<0.001	→	0.09
CV2					
CV3	0.81		0.13		0
DV1	<0.001	→	<0.001	→	0.3
DV2	0.72		0.19		0

B1 - Fork

Variable	GC-XGBoost		MR		
	p-value	discovery	p-value	discovery	Coeff
CV1	0.99		0.18		0
CV2	<0.001	→	<0.001	→	0.3
CV3	0.99		0.12		0
DV1	<0.001	→	<0.001	→	0.3
DV2	1		0.27		0

B2 - Fork (confounder not included)

Variable	GC-XGBoost		MR		
	p-value	discovery	p-value	discovery	Coeff
CV1	0.25		<0.001	→	0.17
CV2					
CV3	0.99		0.08		0
DV1	<0.001	→	<0.001	→	0.3
DV2	0.31		0.24		0

C2 - Collider

Variable	GC-XGBoost		MR		
	p-value	discovery	p-value	discovery	Coeff
CV1	<0.001	→	<0.001	→	0.19
CV2	0.24		<0.001	→	-0.05
CV3	0.65		0.08		0
DV1	0.42		0.67		0
DV2	0.30		0.35		0

C2 - Collider (collider not included)

Variable	GC-XGBoost		MR		
	p-value	discovery	p-value	discovery	Coeff
CV1					
CV2	0.90		0.45		0
CV3	1		0.12		0
DV1	0.78		0.46		0
DV2	0.96		0.27		0

Figure 3 - Example of simulated scenarios for each DAG's fundamental structures (A - pipe; B - fork; and C - collider). Green arrows represent true discoveries of causal relations, red arrows represent false discoveries (false positives) and the cross represent a failure to discover an actual causal relation (false negative). Note that in these examples, the GC-XGBoost causal discovery approach consistently found the true causal relations, except for the directionality on the collider structure (C1). This is in contrast with the results from Multiple Regression (MR) analyses in which all variables all included with no consideration of the underlying causal structure. Besides illustrating how the casual discovery approach deals with the fundamental causal structures, these examples show the effects of a mediator in the inference of the causal effect (A), the inducement of an spurious association when a confounder is not included in the model (B2 - dashed arrow) and how including a collider in the model leads to an spurious negative association with a third variable.

Figure 4 - Simulation results for all scenarios (A) and dependence on the sample size (B) for a complex graph with identifiable causal relations. FDR - False Discovery Ratio; MR - Multiple Regression; GC-XGBoost - Granger-causality from XGBoost results.

3.3 - Discussion

The fork structure provides a precise definition of a classic confounder (Figure 3-B1). If not taken into account in a regression model, a confounder variable will induce a spurious association between an unrelated variable and the outcome (Figure 3-B2). Importantly, as

illustrated in Figure 3-B and in results of simulations (Figure 4-A), the causal discovery algorithm was more robust than MR against falsely discovering a causal relation when a confounder is not included in the model.

In contrast with the classical confounder, including a collider variable into a model induces a negative spurious correlation between an unrelated third variable and the outcome. As emphasized in many accounts on the perils of "put-it-all" multiple regression approaches (Westreich & Greenland, 2013), colliders are one important example on how *including* a covariate may actually dampen causal inference (Figure 3C1 and Figure 4, scenario 4a). Note that, despite not falsely introducing the negative correlation of collider bias, the Granger-causality discovery approach is still susceptible to false discoveries induced by a collider structure (Figure 3C1, Figure 4 scenario 4a). We attempted to minimize the risks for collider bias by a design choice: the outcomes of interest must necessarily not affect the candidate causal factors, as is further illustrated in the example application.

Besides the scenarios illustrated in Figure 3, two scenarios with DAGs comprising five causally connected variables and combinations of the fundamental structures were also simulated (scenarios 5 and 6). To better illustrate the applicability and the limits of data-driven causal discovery approach, one of these graphs included non-identifiable causal effects (scenario 6). For non-identifiable DAGs, if the full set of variables is included in the model, an unbiased causal effect can be inferred (scenario 6a). However, an unobserved variable makes the causal discovery task undecidable (scenarios 6b and 6c).

In sum, the simulations illustrate the potential of combining Granger-causality in a predictive procedure to improve the discovery of direct causal relations from observational data. They also illustrate the current, and probably inescapable, limits for the use of a single algorithmic method in assigning causal relations, namely, the existence of non-identifiable causal structures. In the next section, we further test and illustrate these strengths and limitations of our approach by analyzing and interpreting the results of a concrete causal discovery problem of relevance to developmental psychology.

4 - An illustrative application: discovering early-life determinants of cognitive development

4.1 - Rationale and dataset

Knowing which psycho-social, environmental and biological factors in early life are more relevant and causality-related to cognitive abilities at school age is of paramount importance to design more effective interventions to foster cognitive development and educational achievement. Critically to illustrate the integrative approach proposed here, an extensive literature and at least some potential causal determinants of cognitive development are reasonably well-established. Moreover, large scale birth-cohort studies with valid and reliable measures of cognitive outcomes are available.

The UK Millennium Cohort Study (MSC) is an on-going nationally representative observational study of children born between September 2000 and January 2002 (the full cohort profile can be found in (Connelly & Platt, 2014) (Joshi & Fitzsimons, 2016). Crucial for our illustrative purpose, MSC aimed to collect a wide range of information, in contrast to the focus of previous cohorts on health measures. The multidimensional approach of the MSC was based on an ecological model for child development (Bronfenbrenner, 1979) that explicitly takes the putative influence of family and social contexts into account. This rationale behind the construction of the MSC makes this dataset particularly suitable to explore complex and multiple causes for cognitive development.

The first comprehensive survey of the MSC included data provided by the primary caregiver when the children were around 9 months of age and follow-ups at ages 3, 5, 7, 11, 14 and 17 are currently available. Here, we used data collected from the main caregiver at the first wave to construct the set of predictor variables for cognitive outcomes assessed when children were 5 years old (Wave 3). With this choice, we aimed to discover the most relevant early-life causal factors for cognitive development at the beginning of the school life. Besides applying our causal approach, results from conventional multiple regression models and the

relevance of each feature for the ML prediction were considered for comparison. We expected to find plausible causal links for each outcome that would not be captured by either a causally uninformed multiple regression model or simply by the features' relevance for ML algorithm prediction. In total, 15,460 children were assessed at MSC in wave 3. Of those, 14 170 were singletons with data on cognitive development at age 5 years and were thus included in the analysis.

4.1.1 - Cognitive development outcomes

Cognitive development at age 5 years was assessed in children using the British Ability Scales (BAS; Hill, 2005), including the subscales of *naming vocabulary*, *picture similarity* and *pattern construction*. These subscales were used to assess the cognitive functions of *verbal ability*, *pictorial reasoning* and *spatial abilities*, respectively. Data was collected by trained interviewers directly from children in the home. All children enrolled in the MSC cohort were assessed, except those with severe learning disability or severe behavioural problems or children who did not have the required level of either English or Welsh.

Scores were rescaled by dividing raw scores by the maximum punctuation for each of the BAS subscales. Moderate correlations ($r = 0.34 - 0.37$) between the subscales were found, suggesting that each comprise an independent construct of information-processing skill. Importantly, BAS subscales were shown to be robust and individually interpretable, allowing for a proper evaluation of each of these three main domains of information-processing (Hill, 2005).

Given that the cognitive skills assessed at 5 years are not present at the first assessment (though some precursor of cognitive development might be so), it is reasonable to suppose that no relevant reverse causality is present. In other words, not only the time precedence of the predictor variables but the implausibility of a *direct* reversal causal influence of a cognitive skill developing later in life on the variables measured at 9 months make the assumption of no post-treatment effects with a collider structure reasonable. To give a more concrete example, it is highly unlikely that the linguistic abilities of the child at 5

years old would have any influence on environmental factors or behaviour at 9 months, since such skills would not have developed yet. Hence, we assume that arrows on the DAG would necessarily point from variables at 9 months to the later cognitive outcomes not only because of the temporal precedence but for the additional reason that skills not yet developed would not plausibly influence other factors.

4.1.2 - Potential causal predictors

Potential causal predictors were based exclusively on self-report data from the main caregiver at MSC wave 1, when children were around 9 months of age. The main caregiver was usually the biological mother, except for 18 cases where the father responded to the questionnaire and 7 cases where the questionnaire was not completed by any biological parent. A rich range of multidimensional variables were included, taking advantage of the MSC multidisciplinary approach. A first screen identified 194 individual questions as potential predictors, excluding only duplicates and variables related to quality control or evidently not relevant to the analysis (e.g., information on twins). Single questions conveying relevant demographic information were included (referred to as covariates in Table 1). A second screening identified highly redundant questions by examining bivariate correlations and questions that conveying unique information were selected when correlations were high ($r > 0.75$). These variables comprising a single question are marked in italics in Table 1 and include, for example, smoking, alcohol consumption, and breastfeeding. Finally, an exploratory factor analysis (EFA) was conducted for the remaining 167 single questions resulting in the final multiple dimension predictors described in Table 1. Dimensionality reduction using EFA was necessary to improve the interpretability of the results and minimize potential multicollinearity issues (McElreath, 2018).

Table 1 - Predictors variables after selection and dimensionality reduction using EFA.

Dimension	Factor	# of items	Description
Demographic	<i>Age (days)</i>	1	Age in days at the moment of the cognitive evaluation
	<i>Sex</i>	1	Child sex
	<i>Weight at birth</i>	1	Child weight at birth
	<i>Mother age</i>	1	Maternal age at the time of birth in years
	<i>Country</i>	1	Country of residence in the UK
	<i>Religion</i>	1	Reported maternal religion
	<i>Marital status</i>	1	Reported marital status of the mother at wave 1
Child Individual level	Developmental Stage	10	Summary of developmental hallmarks (sits up, waves, gives toys)
	Developmental problems	6	Delayed hallmarks (smiling, grabbing objects, solid food)
	<i>Hearing Problems</i>	1	Any reported child hearing problem
	Routine	4	Regular timing of sleep and feeding
	Well-being	8	Baby pleasant during routine activities (calm, happy sounds)
	Anxiety	4	Fretful in new places, wary of strangers, shiness
	Sleep problems	4	Fuzziness around sleep time and frequency that wakes up at night
Caregiver	Maternal Health	11	Maternal health problems, including depression
	Control	3	Maternal perceived control over own life
	Cognitive difficulties	4	Maternal difficulties with routine activities, reading or numerical ability
	Negative self-evaluation	10	Maternal negative self-evaluation (thinks is no good, feels a failure)
	Positive self-evaluation	8	Maternal positive attitudes towards self (includes fathers involvement)
	<i>Smoking</i>	1	Whether the mother have ever smoked
	<i>Alcohol</i>	1	Maternal frequency of alcohol consumption per week
	Feeling when away	2	Maternal thoughts and feelings about baby when apart
	Impatient with baby	5	Feelings of irritation when caring for baby
	<i>Maternal education</i>	1	Maternal age when leave formal education
Family level (relationships)	<i>Breastfeeding</i>	1	Age when last had breast milk, including currently having
	Response to crying	3	Picking up baby when crying
	Positive parenting	4	Importance of child stimulation, including talking and cuddling
	Interest in baby	3	Expressed interest in baby
	Planned pregnancy	3	Planned pregnancy, feelings when became pregnant
	Relationship	7	Quality of the partners relationship, including domestic violence
Social level (environment)	<i>SES</i>	1	Socioeconomic status as measure by long version NS SEC
	<i>Language at home</i>	1	Language spoken at home
	<i>Mother at full time job</i>	1	Whether the mother is at a full time job
	<i>Partner at full time job</i>	1	Whether the father is at a full time job
	<i>Father alive</i>	1	Whether the father is alive at wave 1
	<i>Mother alive</i>	1	Whether the mother is alive at wave 1
	<i>Frequency sees mother</i>	1	The frequency the child is with the mother
	<i>Frequency sees father</i>	1	The frequency the child is with the father
	<i>Parents separated</i>	1	If the parent have ever been separated

	Paid Childcare	1	Hours of paid childcare per week
	Childcare	3	Usage of childcare in general
	Interest in politics	8	General interest in politics and whether receiving any benefits
	Moral attitudes	3	Beliefs towards having children without being married, interracial marriage
	Religion intolerance	2	Religion based insults or attacks
	Home responsibility	15	Level of shared responsibility for home chores and baby routine
	Workload	8	Maternal working load, including hours working per day
	Work pregnancy	7	Maternal working load during pregnancy
	Flexible work	2	Usage of flexible work arrangements and amount of time with baby
	Home satisfaction	9	Satisfaction with home and with home area, including access to garden
	Perceived environment	5	Neighbourhood level nuisance (noise, rubbish, pollution, insults)

We used parallel analysis to determine the number of factors. For robustness, an unsupervised clustering analysis was also conducted, which indicated a similar number of factors to the parallel analysis. Then, EFA with oblique rotation was conducted reducing the 167 original measures to 29 factors. Each factor comprised meaningful and interpretable measures which are summarized in Table 1 and presented in detail in supplementary materials.

4.2 - Statistical analysis

As a first step, Granger-causality inference from XGBoost results (GC-XGBoost) was obtained using the same algorithm described and applied to the simulated scenarios. After initial tuning, the XGBoost hyperparameters were fixed (max.depth=6, eta=0.3, nrounds=100). To evaluate the machine-learning algorithm performance, the distributions of the absolute error of prediction for each individual were obtained and plotted (Figure 5).

Once the individual predictions were shown to be reasonably accurate, we proceed to the hypothesis testing. For comparison and to illustrate how the algorithm can be integrated in a causal inference strategy, single and multiple regression models (MR) were also implemented. First, multiple regression models with the full set of predictors were run for each of the three outcomes. Note that these models are not informed by the potential causal structure of the problem. A second set of MR models was run but including only the potential causal factors obtained by the GC-XGBoost method. Also, single regressions were

implemented for each of these factors independently to obtain uncorrected estimates. If the factors are causally related to the outcomes and independent of each other (see simulation scenario 1), the uncorrected and corrected regression coefficients from single and multiple regressions must be similar. Hence, if convergent, the results for these regression models would reinforce the causal discovery results and can be used to infer the strength of the candidate causal relations.

For comparison and to illustrate the differences between a causally uninformed regression approach and the causal discovery framework, the regression coefficients from the multiple regression model that included all variables are shown (Figure 6, first column). Also for comparison, we report the relevance of each independent variable for the prediction using the XGBoost ML procedure (Figure 6, middle column). Such measure of relevance of a variable or feature to the prediction model is often used as an indicator of the feature's explanatory power in ML studies. Finally, regression analyses informed by the causal discovery method were implemented. Single and multiple regressions coefficients and confidence intervals were calculated only for the set of candidate causal factors. If the causal variables are independent from each other, the coefficients from the single and multiple regressions are expected to be similar. This is illustrated in the first simulated scenario, in which the true value of the regression coefficient is obtained for each of the independent causal factors in the graph (Figure 4A, Supplementary Figure 1).

As the ground-truth causal relations are obviously unknown for this real-world dataset, we sought to test whether plausible causal relations were discovered. Particularly, we expected to find determinants of cognitive development that are well-established in the literature by converging lines of evidence.

4.3 - Results

4.3.1 - Prediction performance

Reasonable performance for individual predictions were obtained for the three cognitive outcomes using the XGBoost algorithm. The distributions of individual errors for

the out-of-sample predictions are shown in Figure 5. All outcomes were re-scaled to the percentage of the possible maximum score to facilitate interpretation and comparison. The median (sd) individual prediction error was 0.08 (+/- 0.077) for the naming task, 0.06 (+/- 0.053) for the pattern construction task and 0.09 (+/- 0.1) for the picture recognition task. As can be seen in Figure 6, the vast majority of individual prediction errors did not exceed 20% for all tasks and the predictive performance of the algorithm was fairly similar for all outcomes.

Figure 5 - Prediction performance of the XGBoost algorithm for each of the three cognitive outcomes. Individual prediction errors were defined as the absolute difference between the out-of-sample individual prediction and the actual individual punctuation. Note that the absolute error in individual prediction for all outcomes were similar, with the most errors not exceeding 20%.

4.3.2 - Candidate causal factors

For verbal ability, the candidate direct causal factors identified were: (1) the language spoken at home; (2) the maternal working status during pregnancy; (3) hours of paid childcare per week at 9 months; (4) hearing problems; (5) the maternal educational

achievement measured as the age she left full-time education; the (6) maternal self-reported cognitive difficulties in daily routine (including functional illiteracy); (7) mother's age; (8) children age in days at the moment of the cognitive assessment; and (9) how frequently child and mother interacts. For spatial reasoning ability, the candidate early-life direct causal factors found were: (1) the child gender; (2) the perceived quality of the neighbourhood; (3) the working load of the mother (4) and her partner; (5) the maternal educational achievement measured as the age she left full-time education; and (6) religion. For pictorial reasoning ability, the candidate direct causal factors were: (1) the perceived quality of the neighbourhood; (2) maternal self-reported cognitive difficulties in daily routine; (3) delayed developmental hallmarks at 9 months; (4) maternal negative self-evaluation; (5) mother's age; and (6) child weight at birth.

Figure 6 - Discovered early-life causal links for each cognitive ability outcome. For comparison, the first column depicts the regression coefficients in a "put them all" MR strategy. The second column presents the relevance of each feature for the predictive ML model. The last column shows the increased out-of-sample prediction for each variable. The vertical dashed lines mark the threshold for significance calculated using bootstrap ($p < .05$).

4.4 - Discussion

As expected, the causal candidates found were mostly plausible or reasonably well-established determinants of the assessed specific cognitive abilities. Hence, applying our approach based on prediction and time constraint to a large cohort dataset yielded plausible causal candidate factors for cognitive development. Discovered factors should be interpreted as indicative of causal relations that must be confirmed using causal inference tools, ideally in independent large samples. In a general causal inquiry, other causal discovery tools can also be applied and a better informed DAG can hence be drawn and tested. To further illustrate the practical analysis, we conducted a series of MR models including all the causal candidate factors as predictors. As expected, we observed similar strengths and patterns of relation between the factors and outcomes for single regression models and a multiple regression model including all causal candidates at once, indicating that most factors might constitute independent causes. Quantifying the associations also indicated that some candidates were indeed weakly associated with the outcomes and might not be considered in further investigations.

Factors directly related to the linguistic environment experienced by the child at 9 months were the most relevant among the causal candidate factors for determining verbal ability at 5 years old (Melhuish et al., 2008). Moreover, the putative causal factors can quite reasonably affect language development in the expected direction. The “speaking language at home being other than English (or Welsh)” was related to lower verbal ability. “Maternal cognitive difficulties impairing daily life tasks” were also associated with children's lower verbal ability. On the other hand, “higher maternal educational achievement” and “working status during pregnancy” and hours spent with paid childcare” were associated with higher verbal ability. Unexpectedly, the presence of any hearing problem detected up to 9 months

were associated with higher verbal ability at 5 years old. Out of the nine causal candidates, three (mother's age, children's age in days at cognitive assessment and the frequency of child-mother interactions) present negligible associations in the regression models. Finally, though not superposed, the coefficients calculated from single and multiple regression using only the candidate causal factors had the same direction, order of magnitude and pattern, further suggesting that the remaining six factors appear to constitute the most relevant and independent causal factors of verbal ability.

Regarding spatial ability, the child's sex stood out as the most relevant causal candidate, with higher scores for boys. Sex was the strongest factor and with the exact same coefficient for both the single and multiple regression models. These findings are in line with highly consistent evidence of sex difference in spatial ability (Peters et al., 2006; Voyer et al., 1995). Maternal educational achievement and mother's and her partner's work status when the child was 9 months old were also associated with higher spatial ability at 5 years old. A higher quality of the neighborhood environment at early age was related to higher spatial ability scores. However, differentiating between the effects of the general environment and shared genetic variation is not possible with the data used here. Though a causal candidate, maternal religion was not associated with the outcome in further single and multiple regression analyses. Similar to what was observed for verbal ability, most of the coefficients calculated from single and multiple regression using only the candidate causal factors were not superposed. However, all coefficients had the same direction, order of magnitude and pattern, also further suggesting that the factors could constitute independent causal factors.

Finally, pictorial reasoning was found to be importantly affected by more general child development problems and by maternal self-reported cognitive difficulties. The quality of the neighborhood environment, on the other hand, was associated with higher pictorial reasoning. Unexpectedly, the same was observed for maternal negative self-evaluation. Maternal age and child's weight at birth were putative causal factors but with no significant association in the causal informed regression analyses. Moreover, all the coefficients calculated from single and multiple regressions including only the independent causal candidates were highly similar.

Figure 7 - Causally informed single and multiple regression models results. Regression coefficients with 95% confidence intervals are presented for single regressions of each candidate causal factors (smaller balls) and for a multiple regression model including all direct causal factors candidates (bigger balls). Note that, if the predictors are indeed independent causal factors, the regression coefficients should be similar for both single and multiple regression models.

4 - General Discussion

Accurately identifying and quantifying causal relations in dynamic systems with multiple causes is of paramount importance in psychological research, policy advising and practice. Here we sought to develop and test a new tool for the toolbox of causal inquiry. The feasibility and utility of the approach proposed here relies on recent theoretical and technical advances on causal modeling and machine-learning.

The increasing availability of large and high-quality prospective longitudinal observational studies with a broad range of measures and assessments allows refining discovery and testing of causal relations for complex multicausal problems. Technical advances in machine-learning and in popular programming languages like R and python now allow the broad use of high-performance data-driven predictive tools. Last but not least, the so-called "causal revolution" led by the development do-calculus, DAGs and the

comprehensive structural causal model (SCM) framework provide a principled way to infer and quantify causal relations from non or quasi-experimental data.

Based on previous formal results on the partial convergence between a measure of causality based only on prediction, i.e. Granger-causality, and the structural causal modeling framework, we proposed a data-driven approach to search for relevant direct causal factors. A principled and comprehensive set of simulations was implemented supporting that the approach robustly recovers known causal structures from a reasonable set of representative scenarios. An important and expected exception is causal structures with non-identifiable causal relations (i.e., causal relations that cannot be inferred from observational data [Imai et al., 2014; Pearl, 2014]).

This is an important limitation to emphasize as it reflects a more general account on the limits of causal discovery research. A fully automated algorithm for discovering causal relations from observational data is likely not an achievable goal. However, this impossibility of a general algorithm does not necessarily imply that data-driven tools to causal discovery cannot be useful. Here we sought to develop such a tool and illustrate how it can be embedded in causal inquiry. It is also important to emphasize that other causal discovery methods are currently available and should be considered to the specific causal question one is attempting to address (Glymour et al., 2019). In contrast to these methods, we aimed to design an approach to search for the most relevant, direct and independent candidate factors rather than automatically finding a full causal structure.

To illustrate one possible use of a data-driven causal discovery method based on the Granger-causality idea, we set out to discover the early-life determinants of cognitive abilities at the beginning of school-life using data from the Millennium Cohort Study. Of importance for our illustration, cognitive abilities were measured directly in children when they were 5 years old using a well validated tool. Thus, the factors predicting the well-defined and measured outcomes of verbal, spatial, and pictorial reasoning abilities could be investigated. From a range of 50 independent factors measured by an extensive survey with the main caregiver, usually the mother, when the children were 9 months, several sets of highly plausible causal candidates were discovered. Counterfactual reasoning on such factors is

useful to clarify their plausibility. For instance, the language spoken at home being the same as that assessed by the naming task or the child having hearing problems corrected or being exposed to a more challenging linguistic environment in childcare or the mother having access to full-time education for longer would all quite possibly change the level of verbal ability when the child is 5 years old.

Causal investigation requires exploring different methods, approaches and lines of evidence. Following the steps proposed by Rohrer (2018), the discovered putative causes using the approach described here might be used to further articulate specific causal questions, draw the corresponding set of plausible DAGs and then estimate the causal effects using causal inference techniques (Tennant et al., 2021). Finally, robustness checks (Duncan et al., 2014), triangulation (Hammerton & Munafò, 2021), or sensitivity analyses (Imai et al., 2010) must be performed to determine the impact of violations of the causal assumptions on the estimates. Most definitive and accurate causal answers in psychological research will ultimately depend on accumulated evidence from different studies designs and lines of evidence (Eronen, 2020).

Following the conceptual framework recently proposed by Hofman et al. (2021), we illustratively explored the possibility of constructing and evaluating an *integrative* tool. Hofman et al. (2021) proposes a taxonomy of modelling comprising descriptive, explanatory, predictive and integrative approaches. While causal inference corresponds to explanatory modelling, predictive modelling, particularly supervised machine learning, aims to directly predict an outcome of interest without explicitly taking causal relations into consideration. Integrative modelling seeks a compromise between predictive power and assessing true causal rather than only associational relations. Integrative models might then potentially generate accurate predictions even in a changing world (Hofman et al., 2021). But in contrast with this framework, we propose that useful integrative approaches might not need to be fully explanatory, in the sense of establishing mechanisms, to achieve integrative research goals.

A more general and often underappreciated differentiation that can be clarified using the interventional approach to causality regards the role of detailed mechanistic accounts in causal inquiries. Unravelling a detailed causal mechanism is not necessary nor sufficient to

develop and test an effective intervention. Actually, discovering the causal factors that are prone to be changed by an intervention might be a more directly effective way to guide clinicians and policymakers. In the classical example of the association between smoking and lung cancer, establishing that the association is causal and estimating the *total* causal effect suffices to guide policies aiming to reduce smoking in the population. In this case, discovering the cellular and molecular mechanisms by which smoking leads to lung cancer was historically useful as an additional line of evidence for the plausibility of the causality link. But such mechanistic evidence would not be absolutely *necessary nor sufficient* to determine the causal link. Furthermore, knowing the detailed molecular mechanisms would not necessarily lead to better estimates of the effect of quitting smoking on the population incidence of lung cancer.

In sum, discovering causes and describing mechanisms are different scientific tasks, with complementary but different potential uses and implications. Not clearly differentiating these two tasks can actually introduce unnecessary confusion and, in some cases, hampers our ability to accurately determine a causal effect. This danger is clearly demonstrated in DAGs approach, where conditioning on a mediator biases the estimation of the *total* effect of a causal factor. More generally, Hoel et al., (2013) have used SCM and information theory to demonstrate that, for certain causal structures, causal relations not presented in micro-scales emerge in systems' macro spatio-temporal levels. Such phenomena of causal emergence in macro systems provides further support for the idea of focusing causal investigation on the "differences that make a difference" (Tononi, 2012). Though SCM and PO have been applied to causal mediation analysis (Imai et al., 2010, 2014; Pearl, 2014), we here focused on adding to the researcher toolbox for causal inference on macro systems in expanded temporal scales rather than mechanistic investigations.

In conclusion, we empirically tested a procedure to identify causal factors in complex causal structures. More generally, we attempted to illustrate the potential of clearly stating causal questions in psychology (Rohrer, 2018), using SCM and DAGs repertoire (Pearl, 2009a) and integrating up-to-date predictive tools in causal inquiries (Eichler, 2012; Hofman et al., 2021). The particular procedure tested here might be useful in prioritizing further

causal investigation of factors by assessing which plausible factors would more significantly modify a given outcome.

References

- Austin, A. E., Desrosiers, T. A., & Shanahan, M. E. (2019). Directed acyclic graphs: An under-utilized tool for child maltreatment research. In *Child Abuse & Neglect* (Vol. 91, pp. 78–87). <https://doi.org/10.1016/j.chiabu.2019.02.011>
- Balzer, L. B., & Petersen, M. L. (2021). Machine Learning in Causal Inference: How do I love thee? Let me count the ways. *American Journal of Epidemiology*. <https://doi.org/10.1093/aje/kwab048>
- Bollen, K. A., & Pearl, J. (2013). Eight Myths About Causality and Structural Equation Models. In *Handbooks of Sociology and Social Research* (pp. 301–328). https://doi.org/10.1007/978-94-007-6094-3_15
- Bronfenbrenner, U. (1979). *The Ecology of Human Development: Experiments by Nature and Design*. Cambridge, Mass. : Harvard University Press.
- Connelly, R., & Platt, L. (2014). Cohort Profile: UK Millennium Cohort Study (MCS). In *International Journal of Epidemiology* (Vol. 43, Issue 6, pp. 1719–1725). <https://doi.org/10.1093/ije/dyu001>
- Duncan, G. J., Engel, M., Claessens, A., & Dowsett, C. J. (2014). Replication and robustness in developmental research. *Developmental Psychology*, 50(11), 2417–2425.
- Eichler, M. (2012). Causal Inference in Time Series Analysis. In *Causality* (pp. 327–354). <https://doi.org/10.1002/9781119945710.ch22>
- Eichler, M., & Didelez, V. (2010). On Granger causality and the effect of interventions in time series. *Lifetime Data Analysis*, 16(1), 3–32.
- Eronen, M. I. (2020). Causal discovery and the problem of psychological interventions. In *New Ideas in Psychology* (Vol. 59, p. 100785). <https://doi.org/10.1016/j.newideapsych.2020.100785>

- Glymour, C., Zhang, K., & Spirtes, P. (2019). Review of Causal Discovery Methods Based on Graphical Models. In *Frontiers in Genetics* (Vol. 10).
<https://doi.org/10.3389/fgene.2019.00524>
- Granger, C. (1986). Statistics and Causal Inference: Comment. In *Journal of the American Statistical Association* (Vol. 81, Issue 396, p. 967). <https://doi.org/10.2307/2289068>
- Granger, C. W. J. (n.d.). Some Recent Developments in a Concept of Causality. In *Essays in Econometrics vol II: Collected Papers of Clive W. J. Granger* (pp. 71–83).
<https://doi.org/10.1017/ccol052179207x.004>
- Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. In *Econometrica* (Vol. 37, Issue 3, p. 424).
<https://doi.org/10.2307/1912791>
- Grosz, M. P., Rohrer, J. M., & Thoemmes, F. (2020). The Taboo Against Explicit Causal Inference in Nonexperimental Psychology. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 15(5), 1243–1255.
- Hamaker, E. L., Mulder, J. D., & van IJzendoorn, M. H. (2020). Description, prediction and causation: Methodological challenges of studying child and adolescent development. *Developmental Cognitive Neuroscience*, 46, 100867.
- Hammerton, G., & Munafò, M. R. (2021). Causal inference with observational data: the need for triangulation of evidence. *Psychological Medicine*, 51(4), 563–578.
- Harder, V. S., Stuart, E. A., & Anthony, J. C. (2010). Propensity score techniques and the assessment of measured covariate balance to test causal associations in psychological research. In *Psychological Methods* (Vol. 15, Issue 3, pp. 234–249).
<https://doi.org/10.1037/a0019623>
- Hill, V. (2005). Through the Past Darkly: A Review of the British Ability Scales Second

- Edition. *Child and Adolescent Mental Health*, 10(2), 87–98.
- Hoel, E. P., Albantakis, L., & Tononi, G. (2013). Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences of the United States of America*, 110(49), 19790–19795.
- Hofman, J. M., Watts, D. J., Athey, S., Garip, F., Griffiths, T. L., Kleinberg, J., Margetts, H., Mullainathan, S., Salganik, M. J., Vazire, S., Vespignani, A., & Yarkoni, T. (2021). Integrating explanation and prediction in computational social science. *Nature*, 595(7866), 181–188.
- Imai, K., Keele, L., & Tingley, D. (2010). A general approach to causal mediation analysis. *Psychological Methods*, 15(4), 309–334.
- Imai, K., Keele, L., Tingley, D., & Yamamoto, T. (2014). Comment on Pearl: Practical implications of theoretical results for causal mediation analysis. In *Psychological Methods* (Vol. 19, Issue 4, pp. 482–487). <https://doi.org/10.1037/met0000021>
- Joshi, H., & Fitzsimons, E. (2016). The Millennium Cohort Study: the making of a multi-purpose resource for social science and policy. In *Longitudinal and Life Course Studies* (Vol. 7, Issue 4). <https://doi.org/10.14301/llcs.v7i4.410>
- Kalisch, M., Mächler, M., Colombo, D., Maathuis, M. H., & Bühlmann, P. (2012). Causal Inference Using Graphical Models with the R Package pcalg. In *Journal of Statistical Software* (Vol. 47, Issue 11). <https://doi.org/10.18637/jss.v047.i11>
- King, G., & Nielsen, R. (2019). Why Propensity Scores Should Not Be Used for Matching. In *Political Analysis* (Vol. 27, Issue 4, pp. 435–454). <https://doi.org/10.1017/pan.2019.11>
- Kline, R. B. (2015). *Principles and Practice of Structural Equation Modeling*, Fourth Edition. Guilford Publications.
- Knüppel, S., & Stang, A. (2010). DAG program: identifying minimal sufficient adjustment

- sets. *Epidemiology*, 21(1), 159.
- Lawlor, D. A., Tilling, K., & Davey Smith, G. (2016). Triangulation in aetiological epidemiology. *International Journal of Epidemiology*, 45(6), 1866–1886.
- Lazer, D. M. J., Pentland, A., Watts, D. J., Aral, S., Athey, S., Contractor, N., Freelon, D., Gonzalez-Bailon, S., King, G., Margetts, H., Nelson, A., Salganik, M. J., Strohmaier, M., Vespignani, A., & Wagner, C. (2020). Computational social science: Obstacles and opportunities. *Science*, 369(6507), 1060–1062.
- McCaffrey, D. F., Ridgeway, G., & Morral, A. R. (2004). Propensity score estimation with boosted regression for evaluating causal effects in observational studies. *Psychological Methods*, 9(4), 403–425.
- McElreath, R. (2018). *Statistical Rethinking*. <https://doi.org/10.1201/9781315372495>
- Melhuish, E. C., Phan, M. B., Sylva, K., Sammons, P., Siraj-Blatchford, I., & Taggart, B. (2008). Effects of the Home Learning Environment and Preschool Center Experience upon Literacy and Numeracy Development in Early Primary School. In *Journal of Social Issues* (Vol. 64, Issue 1, pp. 95–114). <https://doi.org/10.1111/j.1540-4560.2008.00550.x>
- Naimi, A. I., & Balzer, L. B. (2018). Stacked generalization: an introduction to super learning. *European Journal of Epidemiology*, 33(5), 459–464.
- Ness, A. R., Griffiths, A. E., Howe, L. D., & Leary, S. D. (2011). Drawing causal inferences in epidemiologic studies of early life influences. *The American Journal of Clinical Nutrition*, 94(6 Suppl), 1959S – 1963S.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- Pearl, J. (2009a). Causal inference in statistics: An overview. In *Statistics Surveys* (Vol. 3, Issue none). <https://doi.org/10.1214/09-ss057>

- Pearl, J. (2009b). *Causality*. <https://doi.org/10.1017/cbo9780511803161>
- Pearl, J. (2014). Interpretation and identification of causal mediation. *Psychological Methods*, 19(4), 459–481.
- Pearl, J., & Mackenzie, D. (2018). *The Book of Why: The New Science of Cause and Effect*. Penguin UK.
- Peters, M., Lehmann, W., Takahira, S., Takeuchi, Y., & Jordan, K. (2006). Mental rotation test performance in four cross-cultural samples (n = 3367): overall sex differences and the role of academic program in performance. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 42(7), 1005–1014.
- Richmond, R. C., Al-Amin, A., Smith, G. D., & Relton, C. L. (2014). Approaches for drawing causal inferences from epidemiological birth cohorts: a review. *Early Human Development*, 90(11), 769–780.
- Robins, J., & Hernan, M. (2008). Estimation of the causal effects of time-varying exposures. In *Chapman & Hall/CRC Handbooks of Modern Statistical Methods* (pp. 553–599). <https://doi.org/10.1201/9781420011579.ch23>
- Rohrer, J. M. (2018). *Thinking Clearly About Correlations and Causation: Graphical Causal Models for Observational Data*. <https://doi.org/10.31234/osf.io/t3qub>
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. In *Biometrika* (Vol. 70, Issue 1, pp. 41–55). <https://doi.org/10.1093/biomet/70.1.41>
- Rubin, D. B. (1978). Bayesian Inference for Causal Effects: The Role of Randomization. In *The Annals of Statistics* (Vol. 6, Issue 1). <https://doi.org/10.1214/aos/1176344064>
- Rubin, D. B. (2010a). Causal Inference. In *International Encyclopedia of Education* (pp.

- 66–71). <https://doi.org/10.1016/b978-0-08-044894-7.01313-0>
- Rubin, D. B. (2010b). Causal Inference. In *International Encyclopedia of Education* (pp. 66–71). <https://doi.org/10.1016/b978-0-08-044894-7.01313-0>
- Sjölander, A. (2009). Propensity scores and M-structures [Review of *Propensity scores and M-structures*]. *Statistics in Medicine*, 28(9), 1416–1420; author reply 1420–1423.
- Spirtes, P., Glymour, C., & Scheines, R. (2001). *Causation, Prediction, and Search*. <https://doi.org/10.7551/mitpress/1754.001.0001>
- Spirtes, P., & Zhang, K. (2016). Causal discovery and inference: concepts and recent methodological advances. *Applied Informatics*, 3, 3.
- Tennant, P. W. G., Murray, E. J., Arnold, K. F., Berrie, L., Fox, M. P., Gadd, S. C., Harrison, W. J., Keeble, C., Ranker, L. R., Textor, J., Tomova, G. D., Gilthorpe, M. S., & Ellison, G. T. H. (2021). Use of directed acyclic graphs (DAGs) to identify confounders in applied health research: review and recommendations. In *International Journal of Epidemiology* (Vol. 50, Issue 2, pp. 620–632). <https://doi.org/10.1093/ije/dyaa213>
- Textor, J., Hardt, J., & Knüppel, S. (2011). DAGitty: a graphical tool for analyzing causal diagrams. *Epidemiology*, 22(5), 745.
- Tononi, G. (2012). Integrated information theory of consciousness: an updated account. *Archives Italiennes de Biologie*, 150(4), 293–329.
- Voyer, D., Voyer, S., & Bryden, M. P. (1995). Magnitude of sex differences in spatial abilities: a meta-analysis and consideration of critical variables. *Psychological Bulletin*, 117(2), 250–270.
- Watts, D. J. (2007). A twenty-first century science. *Nature*, 445(7127), 489.
- Westfall, J., & Yarkoni, T. (2016). Statistically Controlling for Confounding Constructs Is Harder than You Think. *PloS One*, 11(3), e0152719.

- Westreich, D., & Greenland, S. (2013). The table 2 fallacy: presenting and interpreting confounder and modifier coefficients. *American Journal of Epidemiology*, 177(4), 292–298.
- West, S. G., & Thoemmes, F. (2010). Campbell's and Rubin's perspectives on causal inference. In *Psychological Methods* (Vol. 15, Issue 1, pp. 18–37).
<https://doi.org/10.1037/a0015917>
- Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 12(6), 1100–1122.
- Zhang, K., & Hyvärinen, A. (2009). Causality Discovery with Additive Disturbances: An Information-Theoretical Perspective. In *Machine Learning and Knowledge Discovery in Databases* (pp. 570–585). https://doi.org/10.1007/978-3-642-04174-7_37
- Hoyer, P. O., Janzing, D., Mooij, J. M., Peters, J., & Schölkopf, B. (2008, December). Nonlinear causal discovery with additive noise models. In *NIPS* (Vol. 21, pp. 689-696)
- Shimizu, S., Hoyer, P. O., Hyvärinen, A., Kerminen, A., & Jordan, M. (2006). A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10).