Causal Inference in Multi-Agent Causal Models

Sam Maes, Stijn Meganck, Bernard Manderick

Computational Modeling Lab,
Vrije Universiteit Brussel,
Pleinlaan 2 - 1050 Brussel,
sammaes, smeganck, bmanderi@vub.ac.be,
http ://como.vub.ac.be

1 Introduction

This paper treats the calculation of the effect of an intervention (also called causal effect) on a variable from a combination of observational data and some theoretical assumptions. Observational data implies that the modeler has no way to do experiments to assess the effect of one variable on some others, instead he possesses data collected by observing variables in the domain he is investigating.

The theoretical assumptions are represented by a semi-Markovian causal model (SMCM), containing both arrows and bi-directed arcs. An arrow indicates a direct causal relationship between the corresponding variables from cause to effect, meaning that in the underlying domain there is a stochastic process *P(effect|cause)* specifying how the effect is determined by its cause. Furthermore this stochastic process must be autonomous, i.e., changes or interventions in *P(effect|cause)* may not influence the assignment of other stochastic processes in the domain. A bi-directed arc represents a spurious dependency between two variables due to an unmeasured common cause (Tian and Pearl, 2002), this is also called a confounding factor between the corresponding variables.

Deciding if a causal effect is identifiable (i.e. can be computed) in a SMCM amounts to assessing whether the assumptions of a diagram are sufficient to calculate the effect of the desired intervention from observational data. When all variables of a domain can be observed, all causal effects are identifiable. In the presence of unmeasured confounders, identifiability becomes an issue (e.g. the causal effect of X on Y is not identifiable in the causal diagram of Figure 1, since we can not distinguish causal influence from X to Y form the influence via the unobserved confounder (Pearl, 2000).

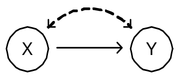

FIG. 1 – *The causal effect of X on Y is not identifiable in this SMCM.*

In this paper we introduce an algorithm for the identification of causal effects in a context where no agent has complete access to the overall domain. Instead we consider a multi-agent approach where several agents each observe only a subset of the variables. The main advantages of the multi-agent solution is that the identification of causal effects can be assessed in cases where information of a local agent cannot be disclosed to other agents for reasons of sensitivity or limited time. Our approach allows to perform causal inference in situations where parts of models are kept confidential by their distributors.

Imagine for example the bi-agent causal model of Figure 2 (based on (Nadkarni and Shenoy, 2001)), where a company wants to assess the influence of its *product pricing* strategy on *product decision*, this is whether a new product is launched or not. Furthermore this company consists of two divisions, one modeled by *agent1*, roughly responsible for external issues such as the market situation, and one modeled by *agent2* pertaining to internal issues such as research & development. With the techniques introduced in this paper the 2 divisions could use the information stored in both their models to calculate the wanted effect, while only communicating over shared variables.

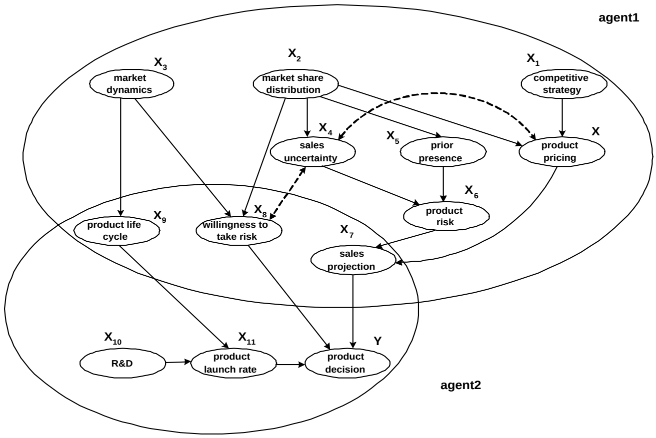

FIG. 2 – *Example of a multi-agent causal model of a product decision model.*

2 Notations and Definitions

In this work uppercase letters V are used to represent variables or sets of variables, while corresponding lowercase letters v are used to represent their values. $Pa(V)$ and $Ch(V)$ are used to represent the observed parents and children of variable V respectively. Furthermore, $Pa(v)$ represents the values of the parents of V.

Let a path entirely composed of bi-directed edges be called a *bi-directed path*. In a semi-Markovian model the set of variables can be partitioned into disjoint groups by assigning two variables to the same group if and only if they are connected by a path only consisting of bi-directed links. We call such a group a *c-component* (from "confounded component") (Tian and Pearl, 2002).

In multi-agent causal models (MACM) (Maes et al., 2003) we no longer assume that there is a central agent having access to all the variables, instead a collection of agents has access to non-disjoint subsets V_i of all observable domain variables V.

A *MACM* consists of n agents M_i, each of which is represented by :
$M_i = \langle V_{M_i}, G_{M_i}, P(V_{M_i}), K_i \rangle$ for $i \in \{1, 2\}$.

- V_{M_i} is the subset of variables *agent-i* can access.
- G_{M_i} is the causal diagram over variables V_{M_i}.
- $P(V_{M_i})$ is the joint probability distribution over V_{M_i}.
- K_i stores the intersections $V_{i,j}$ with other agents j, $\{V_{M_i} \cap V_{M_j}\}$. We assume that the agents agree on the structure and the distribution of their intersections.

An example MACM is given in Figure 2. In this case $V_{M_1} = \{X, X_1, \ldots, X_9\}$ and $V_{M_2} = \{X_7, \ldots, X_{11}, Y\}$, while $V_{1,2} = \{X_7, X_8, X_9\}$.

Our goal is to use MACMs for the identification of $P_x(y)$ in cases where variables X and Y are not modeled by the same agent.

3 Bi-Agent Identification

Without loss of generality we assume that *agent1* contains the intervention variable X and *agent2* contains the variable to be studied Y. Our bi-agent identification algorithm is based on a single agent algorithm due to Tian and Pearl (Tian and Pearl, 2002). It is based on a factorization of semi-Markovian models into distributions over c-components. The c-factor $Q[S_j]$ of c-component S_j is defined as (Tian and Pearl, 2002) :

$$Q[S_j] = \sum_{n_j} \prod_{\{i|V_i \in S_j\}} P(v_i|Pa(v_i), u^i)P(n_j) \tag{1}$$

In (Maes et al., 2005a), we introduce methods for calculating c-factors in a bi-agent causal model under the assumption that for each c-component S_j, the distribution of the variables $S_j \cup Pa(S_j)$ is visible in one of both agents. In the same article, we propose a method for deriving the c-factors of subsets of S^X, the c-component containing variable X, from $Q[S^X]$. These subsets are noted as D_j^X.

Our final formula for calculating $P_x(y)$ is as follows :

$$P_x(y) = \sum_{D_1 \cap (D_2 \setminus Y)} \tag{2}$$

$$\sum_{D_1 \setminus K_1} \left[\prod_j Q[D_j^X] \prod_i \sum_{S_i^1 \setminus D_1} Q[S_i^1] \right] \tag{3}$$

$$\sum_{D_2 \setminus (K_2 \cup Y)} \left[\prod_i \sum_{S_i^2 \setminus D_2} Q[S_i^2] \right] \tag{4}$$

The part on line (3) involves only variables of *agent1* and the part on line (4) involves only variables of *agent2*. The results of the summations on these two lines yield distributions over variables in the intersection and variables X and Y respectively. Finally, the summation over variables in the intersection on line (2), yields the distribution over X an Y.

So we have introduced an algorithm for the identification of $P_x(y)$ in a bi-agent causal model, where each agent combines confidential information stored in its local model with information concerning the intersection with the other agent and the variables being studied (in this case X and Y). In this algorithm no information concerning other variables than the intersection and the variables X and Y is being disclosed.

4 Future Work

We stress that in a MACM, agents are assumed to be honest and to cooperate to solve a problem, without disclosing their private information. Disclosure of confidential information via inference based on the combination of multiple non-confidential query results is a well known problem in statistical databases (Boyens et al., 2004). Investigating whether the solutions to this problem proposed there can be incorporated in our approach would be valuable future work.

Other possible future work would be to study identification algorithms with more than two agents. Preliminary results on this topic can be found in (Maes et al., 2005b). Another obvious extension of this research would be to devise a way to perform multi-agent causal inference when the assumption that every c-component and the parents of its constituents must belong to the same agent, does not hold. Finally, a method for calculating $P_t(s)$, where both T and S are sets, would also be an interesting extension of this work.

Références

Boyens, C., Gunther, O. et Lenz, H.-J. (2004). Handbook of Computational Statistics, chapter 9, pages 267–292. Springer-Verlag.

Maes, S., Meganck, S. et Manderick, B. (2005a). Causal inference in multi-agent causal models. In Proceedings of Modèles Graphiques Probabilistes pour la Modélisation des Connaissances, Atelier of EGC 05, pages 53–62.

Maes, S., Meganck, S. et Manderick, B. (2005b). Identification in chain multi-agent causal models. Accepted at the Special Track on Uncertain Reasoning of FLAIRS 2005.

Maes, S., Reumers, J. et Manderick, B. (2003). Identifiability of causal effects in a multi-agent causal model. In Proceedings of the 2003 IEEE/WIC International Conference on Intelligent Agent Technology (IAT).

Nadkarni, S. et Shenoy, P. P. (2001). A bayesian network approach to making inferences in causal maps. European Journal of Operational Research, 128 :479–498.

Pearl, J. (2000). Causality : Models, Reasoning and Inference. MIT Press.

Tian, J. et Pearl, J. (2002). On the identification of causal effects. Technical Report (R-290-L), UCLA C.S. Lab.