# Graphical Models and Sequential Decisions

Vanessa Didelez
*University College London, Department of Statistical Science*
*Gower Street*
*WC1E 6BT London, United Kingdom*
*vanessa@stats.ucl.ac.uk*

## 1. Motivation

This paper deals with the question of when and how we can identify the effect of a sequential decision strategy from the available data and how we can use suitable graphical models to facilitate this task. To illustrate such sequential decisions we might, for instance, consider dynamic treatment regimes (cf. Murphy, 2003), where decision rules for how the dosage level and type should vary with time are specified before the beginning of treatment and these rules are based on time–varying measurements of subject–specific need.

The problem that I am concerned with arises when the data available to evaluate a strategy for sequential decisions were not actually collected under this strategy and not even necessarily under an experimental design, but for instance were collected in an observational study or under a different strategy. This problem is essentially the same as the one labelled as 'causal inference' (Pearl, 2000, Pearl and Robins, 1995, Dawid et al., 2003), the underlying idea being that a variable is a 'cause' if setting this variable to some specific value (by intervention) changes the distribution of the response. In order to predict such changes, certain conditions have to be satisfied regarding the data and how it was collected, as well as the type of intervention and how these interventions are carried out. As opposed to many causal models that are based on counterfactual argument, the approach taken here is purely decision theoretic.

Further, the dynamic case is somewhat special and different from for instance the causal models considered by Pearl (2000). Sequential decisions will typically take the available information about the past into account, especially the effect of past interventions. The actual decisions are not known in advance, only the decision rules.

## 2. Decisions and Interventions

In a dynamic setting, decisions or interventions can take place either at predetermined points in time or at any time within a certain interval. For simplicity, let us consider the discrete time situation first. In this case, we consider variables $A_1, \ldots, A_T$ which are suitable for intervention, e.g. the administration of a drug, where $A_t$ indicates which action is taken (or treatment received) at time $t$. Further, some covariates $X_1, \ldots, X_T$ might be available, where we assume that $X_t$ is measured before $A_t$, $t = 1, \ldots, T$, i.e. $X_t$ can be used to make a decision about $A_t$. Note that $X_1, \ldots, X_T$ are meant to contain the *available* information — this can be more or less information than required to identify a strategy, as shown below. The final outcome, the effect on which is of interest, is denoted by $Y$. To simplify notation we use $\bar{A} = (A_1, \ldots, A_T)$ for the whole vector, $\bar{A}_t = (A_1, \ldots, A_t)$ for the past, $\bar{A}^t = (A_t, \ldots, A_T)$ for the future, with analogous notations for the other variables.

The two situations to be contrasted are: (a) The action variables $A_1, \ldots, A_T$ are 'under our control', i.e. we know the law that gives rise to them (called the *interventional regime*), or (b) they are not 'under our control'. In the latter case they might arise by nature or they might be controlled by someone else who follows a strategy unknown to us and who might even use information that is not contained in $\bar{X}_t$ and $\bar{A}_{t-1}$ when making a decision about $A_t$. This will be called the *observational regime.* In order to formalise and link these two situations we introduce intervention variables $\sigma_t$, $t = 1, \ldots T$, which are essentially indicators of whether at time $t$ a specific strategy $s \in \mathcal{S}$ is applied or not (cf. Dawid, 2002). Here, $\mathcal{S}$ represents the set of strategies that are to be compared or evaluated. Formally we define

$$\sigma_t = \begin{cases} o & \text{if no intervention takes place (observational regime)} \\ s & \text{if we intervene according to strategy } s \in \mathcal{S}. \end{cases}$$

Let $H_t = (\bar{A}_{t-1}, \bar{L}_t)$ be the history containing the available information about covariates and actions before deciding on $A_t$. As $\sigma_t$ is a decision variable, any probability statement about the remaining variables has to be conditional on one of the possible values of $\sigma_t$. Consequently, $p(A_t|H_t, \sigma_t = o)$ stands for the (typically unknown) distribution of $A_t$ under the observational regime. In contrast, $p(A_t|H_t, \sigma_t = s)$ describes the interventional regime and typically implies that $A_t$ is equal to some predetermined fixed value $a = s(H_t)$, i.e.

$$p(A_t = a|H_t, \sigma_t = s) = \begin{cases} 1 & \text{if } a = s(H_t) \\ 0 & \text{otherwise.} \end{cases}$$

However, one might also think of a random strategy $s$ where $p(A_t|H_t, \sigma_t = s)$ is a non–degenerate distribution.

## 3. Conditions for the identifiability

The problem of identifiability can now be posed as to how to infer the post–interventional distribution $p(Y|\bar{\sigma} = s)$ from the observational $p(Y|\bar{\sigma} = o)$. Note that, for the distribution $p(Y|\bar{\sigma} = s)$ to be meaningful and related to a 'real world' quantity, it is important that the strategy $s$ can actually be carried out, i.e. that one can think of interventions that allow to implement this strategy. In the following, conditions on the interventions will be considered that ensure identifiability — however, if such interventions cannot actually be carried out it is obviously impossible to verify such conditions. In general, identifiability can be an impossible or very difficult task as there is, in principle, no reason why the distribution under intervention should be in any way related to the distribution without intervention. The intervention could be of such a kind that all 'natural' relations among the variables are altered or destroyed. Also, it might be the case that some strategies are identifiable but others are not.

The joint distribution of all the variables involved given the strategy $\sigma$ can naturally be decomposed into the univariate conditional distribution according to the chronological development as

$$(1) \qquad p(Y, \bar{A}, \bar{X}|\bar{\sigma}) = p(Y|\bar{A}, \bar{X}, \bar{\sigma}) \prod_{t=1}^{T} p(X_t|\bar{A}_{t-1}, \bar{X}_{t-1}, \bar{\sigma}_{t-1}) p(A_t|\bar{A}_{t-1}, \bar{X}_t, \bar{\sigma}_t).$$

For a strategy to be sensible, we assume that it is decided before seeing the data whether a strategy is followed or not, i.e. we assume that $\sigma_t \perp\!\!\!\perp \bar{A}_{t-1}, \bar{X}_t$. Note that this does not mean that the actual value assigned to $A_t$ is independent of $\bar{A}_{t-1}, \bar{X}_t$, in particular when $\sigma_t = s$. If, in addition,

$$(2) \qquad Y \perp\!\!\!\perp \bar{\sigma} \mid \bar{A}, \bar{X} \text{ and}$$
$$(3) \qquad X_t \perp\!\!\!\perp \bar{\sigma}_{t-1} \mid \bar{A}_{t-1}, \bar{X}_{t-1}, \qquad t = 2, \ldots, T,$$

it is obvious from the factorisation (1) that the distribution $p(Y, \bar{A}, \bar{X}|\bar{\sigma} = s)$ can be inferred from $p(Y, \bar{A}, \bar{X}|\bar{\sigma} = o)$, as $p(A_t|\bar{A}_{t-1}, \bar{X}_t, \bar{\sigma}_t = s)$ is known. Conditions (2) and (3) are satisfied if the observational regime corresponds to *sequential randomisation,* but this is not a necessary prerequisite. Further, it can be shown that $p(Y|\bar{\sigma} = s)$ is identifiable under weaker conditions (Dawid et al., 2003). In fact, under these weaker conditions we can calculate the post–interventional distribution as

$$p(Y|\bar{\sigma} = s) = \int \cdots \int p(Y|\bar{A}, \bar{X}, \sigma = o) \prod_{t=1}^{T} p(X_t|\bar{A}_{t-1}, \bar{X}_{t-1}, \bar{\sigma}_{t-1} = o)$$
$$p(A_t|\bar{A}_{t-1}, \bar{X}_t, \sigma_t = s)\, dA_T\, dX_T \cdots dA_1\, dX_1,$$

where all involved conditional distributions can either be estimated from the observational data or are the known (possibly degenerate) distributions $p(A_t|\bar{A}_{t-1}, \bar{X}_t, \sigma_t = s)$. The above is also known as the $G$–formula (Pearl and Robins, 1995).

## 4. Graphical representation

The aim of a graphical representation is to make it easier to think about and to verify whether the rather abstract conditions (2) and (3) are satisfied. In fact, these conditional independencies can be depicted by an influence diagram as in Figure 1 (for the case $T = 2$). Missing directed edges represent specific conditional independencies that can be checked by applying one of the usual separation criteria for directed acyclic graphs, e.g. moralisation.

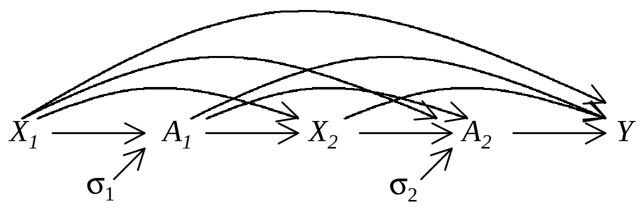


***Figure 1: An influence diagram for*** $T = 2$.

On the variables $\bar{A}, \bar{X}, Y$, Figure 1 is a complete graph. The only restrictions refer to the intervention variables. There are no directed edges into $\sigma_t$ implying that whether or not we intervene is not affected by past observations — this is true if the strategy is fixed in advance, as mentioned above. Also, $\sigma_t$ is pointing only at $A_t$, implying that whether or not we intervene affects future variables *only* through $A_t$, i.e. once the value of $A_t$ is known $\sigma_t$ carries no further information about the future variables $\bar{A}^{t+1}$, $\bar{X}^{t+1}$ and the response $Y$. This condition can be violated, for instance, by what is often called unobserved confounders (Parner and Arjas, 1999). In such a case, the value of $\sigma_t$, observation or intervention, might be informative for predicting $\bar{X}^{t+1}$ even when $A_t$ is given. Directed acyclic graphs, like the one in Figure 1, can be used to verify the conditions for identifiability by including all available knowledge of potential unobserved confounders or other possibly relevant variables and the way they affect the observable variables in the graph. Obviously, this requires sound subject matter knowledge in order to be reliable. Using graph separation, conditional independencies can then be checked, in particular the conditions (2) and (3).

## 5. Generalisations

The above approach to identifying and computing the effect of sequential strategies can be generalised to other dynamic data situations using suitable graphical representations.

*Time series.* Let $\{X_t\}$ and $\{A_t\}$, $t \in \mathcal{T}$, be time series with a similar interpretation as above, i.e. $\{A_t\}$ has to be suitable for intervention. Instead of a response variable $Y$ the outcome variable in the time series setting will typically be a 'later' observation $Y = X_{t_N}$ when contemplating interventions at some (not necessarily successive) points in time $t_1, \ldots, t_{N-1}$. As this is a discrete time situation it can in principle be reduced to the situation of longitudinal data as described earlier. However, for time series a more succinct graphical representation of the dependence structure is available which is based on so–called Granger–causality (Eichler, 2000) and makes use of the repetitive nature of dependencies in stationary time series. For the corresponding 'causality graphs', conditions for identifiability similar to the ones above can be derived to verify whether identifiability is given (Eichler and Didelez, 2003).

*Marked point processes (continuous time).* The situation of continuous time is inherently different as it is not possible anymore to use conditional independence graphs to represent dependencies among the processes except in very special cases. Instead, a dependence concept similar to Granger–causality can be used, called *local independence.* For marked point processes, local independence means that the intensity of an event of a specific type is independent of the occurrence of certain prior events. Such local independencies (or dependencies) can be depicted by *local independence graphs* (Didelez, 2000). Interventions are again included in the graph as additional $\sigma$–nodes and conditions on the independence structure can be checked by applying the so–called $\delta$–separation. However, we have to think carefully about what interventions in a continuous time marked point process might mean. It requires that we are able to induce an event at a given time or to prevent its occurence. An alternative approach to the continuous time situation can be found in Lok (2001).

## REFERENCES


Dawid, A.P., Didelez, V. and Murphy, S. (2003, in preparation). On the conditions underlying causal inference from longitudinal data.

Dawid, A.P. (2002). Influence diagrams for causal modelling and inference. International Statistical Review 70, 161-189.

Didelez, V. (2000). Graphical models for event history data based on local independence. Logos.

Eichler, M. (2000). Causality graphs for stationary time series. Research report. University of Heidelberg.

Eichler, M. and Didelez, V. (2003, in preparation). Causal reasoning in graphical time series models.

Lok, J. (2001). Statistical modelling of causal effects in time. PhD–thesis. University of Amsterdam.

Murphy, S. (2003). Optimal dynamic treatment regimes. Journal of the Royal Statistical Society B (to be published).

Parner, J. and Arjas, E. (1999). Causal Reasoning from longitudinal data. Research report A27 of Rolf Nevanlinna Institute, University of Helsinki.

Pearl, J. and Robins, J. (1995). Probabilistic evaluation of sequential plans with hidden variables. Uncertainty in Artifical Intelligence, Volume 11.

Pearl, J. (2000). Causality: models, reasoning and inference. Cambridge University press.


## RÉSUMÉ


*Cet article traite de la question de savoir quand et comment on peut identifier l'effet d'une stratégy séquentielle en n'utilisant que les données disponibles. Ceci peut poser un problème si les données ont eté collectionnées en utilisant des stratégies différentes ou simplement observées sans suivre de stratégie. Les conditions necessaires pour l'identification dépendent non seulement de la structure des données mais aussi de la stratégie ce qui est souvent oublié. Les modèles graphiques peuvent aider à vérifier si ces conditions sont valables en visualisant la structure des données en ajoutant des noeuds de decision. Les méthodes graphiques habituelles pour vérifier les propriétés de séparation peuvent de même être utilisées pour ces diagrammes de décisions.*