Indirect Elicitation of NIN-AND Trees in Causal Model Acquisition

Yang Xiang, Minh Truong, Jingyu Zhu, David Stanley, and Blair Nonnecke

University of Guelph, Canada
yxiang@socs.uoguelph.ca

Abstract. To specify a Bayes net, a conditional probability table, often of an effect conditioned on its n causes, needs to be assessed for each node. Its complexity is generally exponential in n and hence how to scale up is important to knowledge engineering. The non-impeding noisy-AND (NIN-AND) tree causal model reduces the complexity to linear while explicitly expressing both reinforcing and undermining interactions among causes. The key challenge to acquisition of such a model from an expert is the elicitation of the NIN-AND tree topology. In this work, we propose and empirically evaluate two methods that indirectly acquire the tree topology through a small subset of elicited multi-causal probabilities. We demonstrate the effectiveness of the methods in both human-based experiments and simulation-based studies.

1 Introduction

To specify a Bayes net (BN), a conditional probability table (CPT), needs to be assessed for each non-root node. A BN is often constructed in the causal direction, where a CPT is about an effect conditioned on its n causes. In general, specifying a CPT has the complexity exponential in n. Noisy-OR [Pearl(1988)] and a number of extensions, e.g., [Heckerman and Breese(1996), Galan and Diez(2000), Lemmer and Gossink(2004)] reduce the complexity to linear, but are limited to the reinforcing causal interaction.

The NIN-AND tree [Xiang and Jia(2007)] causal model, as well as its special case [Maaskant and Druzdzel(2008)], extends noisy-OR and explicitly encodes reinforcing and undermining causal interactions, as well as their mixture. Its specification consists of a linear (in n) number of probability parameters and a linear sized tree topology. Its default independence assumptions may be flexibly relaxed to trade efficiency for expressiveness. That is, by relaxing the assumptions incrementally and specifying more parameters, any CPT can be encoded.

The key challenge to specifying a NIN-AND tree causal model is the acquisition of the tree topology, which encodes types of causal interactions among causes. Elicitation of the tree topology requires nontrivial training of a domain expert on the syntax and semantics of NIN-AND tree causal models, and demands nontrivial mental exercise by the expert to articulate the partial order of causal interactions among causes. Usability of NIN-AND tree causal modeling will be enhanced if such training and mental exercise can be avoided during model acquisition.

We accomplish this by proposing two model acquisition methods that bypass direct elicitation of the NIN-AND tree topology. Instead, a small subset of causal probabilities

S. Benferhat and J. Grant (Eds.): SUM 2011, LNAI 6929, pp. 261–274, 2011.

in the order of $O(n^2)$ or $O(n^3)$ are elicited, from which a NIN-AND tree topology is generated. From these probabilities and the tree topology, a NIN-AND tree causal model is defined and the corresponding CPT can be constructed. We show that the acquired CPT is a good approximation of the underlying true CPT.

The remainder of the paper is organized as follows: Background on NIN-AND tree causal models is covered in Sect. 2. The task of NIN-AND tree acquisition and the assumption underlying this work are presented in Sect. 3. In Sect. 4 and 5, we propose two novel techniques for the task. Setup of human-based experiments for evaluation is described in Sect. 6 and results are presented in Sect. 7. They are followed in Sect. 8 by simulation-based studies. Sect. 9 draws the conclusion.

2 NIN-AND Tree Causal Models

An uncertain cause is a cause that can produce an effect but does not always do so. We denote a binary effect variable by $e \in \{e^+, e^-\}$, where e^+ denotes $e = true$, and a set of binary cause variables of e by $X = \{c_1, ..., c_n\}$, where $c_i \in \{c_i^+, c_i^-\}$ $(i = 1, ..., n)$.

A single-causal success is an event where c_i caused e to occur successfully when all other causes are absent. We denote the event by $e^+ \leftarrow c_i^+$ and its probability by $P(e^+ \leftarrow c_i^+)$. For instance, smoking causing lung cancer is denoted by $lc^+ \leftarrow smk^+$. A single-causal failure, where e is false when c_i is true and all other causes of e are false, is denoted by $e^+ \not\leftarrow c_i^+$. A multi-causal success is an event where a set $X = \{c_1, ..., c_n\}$ $(n > 1)$ of causes caused e, and is denoted by $e^+ \leftarrow c_1^+, ..., c_n^+$ or $e^+ \leftarrow \underline{x}^+$. Denote the set of all causes of e by C.

CPT $P(e|C)$ relates to probabilities of causal events as follows: If $C = \{c_1, c_2, c_3\}$, then $P(e^+|c_1^+, c_2^-, c_3^+) = P(e^+ \leftarrow c_1^+, c_3^+)$. C is assumed to include a leaky variable (if any) to capture causes not represented explicitly, and hence $P(e^+|c_1^-, c_2^-, c_3^-) = 0$.

Causes reinforce each other if collectively they are at least as effective as when some are active. For example, radiotherapy and chemotherapy are reinforcing causes for curing cancer. If collectively causes are less effective, they undermine each other. Living with mother and living with wife are undermining causes for the happiness of a man, as often observed. If $C = \{c_1, c_2\}$, and c_1 and c_2 undermine each other, the following hold: $P(e^+|c_1^-, c_2^-) = 0$, $P(e^+|c_1^+, c_2^-) > 0$, $P(e^+|c_1^-, c_2^+) > 0$,

$$P(e^+|c_1^+, c_2^+) < min(P(e^+|c_1^+, c_2^-), P(e^+|c_1^-, c_2^+)).$$

The following Def.1 defines the two types of causal interactions generally.

Definition 1. *Let $R = \{W_1, W_2, ...\}$ be a partition of a set X of causes, $R' \subset R$ be any proper subset of R, and $Y = \cup_{W_i \in R'} W_i$. Sets of causes in R* **reinforce** *each other, iff*

$$\forall R' \; P(e^+ \leftarrow \underline{y}^+) \leq P(e^+ \leftarrow \underline{x}^+).$$

Sets of causes in R **undermine** *each other, iff* $\forall R' \; P(e^+ \leftarrow \underline{y}^+) > P(e^+ \leftarrow \underline{x}^+)$.

Reinforcement and undermining occur between individual causes as well as sets of them. When the interaction is between individual causes, each W_i is a singleton. Otherwise, each W_i can be a generic set. For instance, consider $X = \{c_1, c_2, c_3, c_4\}$, $W_1 =$

$\{c_1, c_2\}$, $W_2 = \{c_3, c_4\}$, $R = \{W_1, W_2\}$, where c_1 and c_2 reinforce each other, and so do c_3 and c_4. But sets W_1 and W_2 can undermine each other.

Disjoint sets of causes $W_1, ..., W_m$ satisfy failure conjunction iff

$$(e^+ \not\leftarrow \underline{w}_1^+, ..., \underline{w}_m^+) = (e^+ \not\leftarrow \underline{w}_1^+) \wedge ... \wedge (e^+ \not\leftarrow \underline{w}_m^+).$$

That is, when causes collectively fail to produce the effect, each must have failed to do so. They also satisfy failure independence iff

$$P((e^+ \not\leftarrow \underline{w}_1^+) \wedge ... \wedge (e^+ \not\leftarrow \underline{w}_m^+)) = P(e^+ \not\leftarrow \underline{w}_1^+) \dots P(e^+ \not\leftarrow \underline{w}_m^+). \quad (1)$$

Disjoint sets of causes $W_1, ..., W_m$ satisfy success conjunction iff

$$(e^+ \leftarrow \underline{w}_1^+, ..., \underline{w}_m^+) = (e^+ \leftarrow \underline{w}_1^+) \wedge ... \wedge (e^+ \leftarrow \underline{w}_m^+).$$

That is, collective success requires individual effectiveness. They also satisfy success independence iff

$$P((e^+ \leftarrow \underline{w}_1^+) \wedge ... \wedge (e^+ \leftarrow \underline{w}_m^+)) = P(e^+ \leftarrow \underline{w}_1^+) \dots P(e^+ \leftarrow \underline{w}_m^+). \quad (2)$$

It has been shown [Xiang and Jia(2007)] that causes are undermining when they satisfy success conjunction and independence. Hence, undermining can be modeled by a direct NIN-AND gate (Fig. 1, left). Its root nodes (top) are single-causal successes, and its leaf node (bottom) is the multi-causal success in question. Success conjunction is expressed by AND gate, and success independence by disconnection of root nodes other than through the gate. The probability of the leaf event can be computed by Eqn. (2). Similarly, causes are reinforcing when they satisfy failure conjunction and independence. Hence, reinforcement can be modeled by a dual NIN-AND gate (Fig. 1, middle). The leaf event probability is obtained by Eqn. (1).

By organizing multiple direct and dual NIN-AND gates in a tree, both reinforcement and undermining, as well as their mixture at multiple levels can be expressed in a NIN-AND tree model. A simple example is given below and more can be found in [Xiang and Jia(2007)]. Consider $C = \{c_1, c_2, c_3\}$, where c_1 and c_3 undermine each other, but collectively they reinforce c_2. Assuming event conjunction and independence, their causal interaction (a two-level mixture of reinforcement and undermining) relative to the event $e^+ \leftarrow c_1^+, c_2^+, c_3^+$ can be expressed by the NIN-AND tree in Fig. 1 (right). The top gate is direct and the bottom gate (the leaf gate) is dual. The link downward from node $e^+ \leftarrow c_1^+, c_3^+$ has a white oval end (a negation link) and

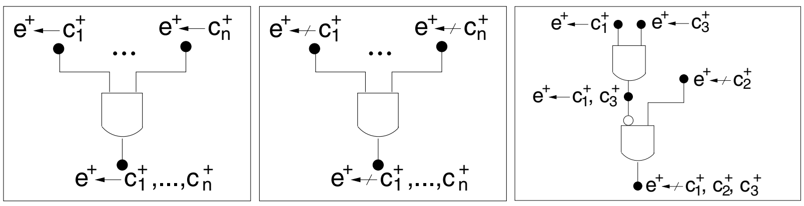

Fig. 1. Direct (left), dual (middle) NIN-AND gates, and a NIN-AND tree (right)

negates the event. All other links are forward links. Probability of the leaf event can be computed by Eqn. (1) and (2). For instance, from single-causal probabilities for root events, $P(e^+ \leftarrow c_1^+) = 0.85,\ P(e^+ \leftarrow c_2^+) = 0.8,\ P(e^+ \leftarrow c_3^+) = 0.7$, probability $P(e^+ \not\leftarrow c_1^+, c_2^+, c_3^+)$ is derived:

$$\begin{aligned} P(e^+ \leftarrow c_1^+, c_3^+) &= P(e^+ \leftarrow c_1^+)P(e^+ \leftarrow c_3^+) = 0.595 \\ P(e^+ \not\leftarrow c_1^+, c_2^+, c_3^+) &= P(e^+ \not\leftarrow c_1^+, c_3^+)P(e^+ \not\leftarrow c_2^+) \\ &= (1 - P(e^+ \leftarrow c_1^+, c_3^+))(1 - P(e^+ \leftarrow c_2^+)) = 0.081 \end{aligned}$$

Furthermore, using a more sophisticated algorithm [Xiang(2010a)], the CPT in Table 1 can be obtained from the NIN-AND tree and these parameters.

Table 1. CPT of the example NIN-AND tree model

$P(e^+\|c_1^-, c_2^-, c_3^-)$	0	$P(e^+\|c_1^-, c_2^+, c_3^-)$	0.8	$P(e^+\|c_1^+, c_2^-, c_3^+)$	0.595	$P(e^+\|c_1^+, c_2^+, c_3^-)$	0.97
$P(e^+\|c_1^+, c_2^-, c_3^-)$	0.85	$P(e^+\|c_1^-, c_2^-, c_3^+)$	0.7	$P(e^+\|c_1^-, c_2^+, c_3^+)$	0.94	$P(e^+\|c_1^+, c_2^+, c_3^+)$	0.919

Variables in a NIN-AND tree model can generally be multi-valued [Xiang(2010b)]. Assumptions on event conjunction and independence can also be relaxed, in which case some root events will be multi-causal. In this work, we focus on binary effect and causes, and on models whose root events are single-causal.

3 Acquisition of NIN-AND Tree Models

As illustrated above, a NIN-AND tree model over e and C consists of its tree topology as well as a single-causal probability for each $c_i \in C$. In general, a NIN-AND tree causal model M is a tuple $M = (e, C, T, PS)$, where e is the effect, C is the set of all causes of e, T is a NIN-AND tree, and PS is the set of single causal probabilities one for each cause in C. From M, a CPT $P(e|C)$ can be uniquely constructed. M and $P(e|C)$ are said to be *consistent*.

Furthermore, NIN-AND tree causal models $M = (e, C, T, PS)$ and $M' = (e, C, T', PS')$ are said to be *structurally consistent* if T and T' are isomorphic. M and M' are said to be *consistent* if they are consistent with the same CPT.

To acquire M, its tree topology T may be elicited directly from the expert. To complete such a task, the expert must have a thorough understanding of the syntax and semantics of NIN-AND tree models, in order to assess and articulate the partial order of causal interactions among causes and cause groups. This demands an nontrivial amount of training of the domain expert before elicitation and nontrivial mental exercise of the expert during elicitation.

To ease these burdens for model acquisition, we investigate the idea to bypass direct tree elicitation. Instead, we elicit a small number of multi-causal probabilities (in addition to the single-causal probabilities PS), and generate T from elicited probabilities. Our work is based on the following assumption:

Assumption 1. *Let $P_t(e|C)$ be the (true) CPT that characterizes the probabilistic relation over an effect e and its causes C, such that the following hold:*

1. *There exists a NIN-AND tree causal model $M_t = (e,C,T,PS)$ that is consistent with $P_t(e|C)$.*
2. *A domain expert is able to approximately assess all single-causal probabilities and some multi-causal probabilities relative to $P_t(e|C)$.*

The first condition is justified by the observation that reinforcement and undermining capture intuitive patterns of causal interaction, and reinforcement based causal models, such as Noisy-OR, have been widely applied. The second condition is justified by knowledge engineering practice in building BNs. Note that the condition does not require the expert to assess all multi-causal probabilities, nor to assess them accurately.

In the following, we investigate two alternative techniques to generate tree topology based on structure elimination (SE) and pairwise causal interaction (PCI).

4 Generate NIN-AND Tree by Structure Elimination

The SE technique builds on minimal NIN-AND tree models [Xiang et al(2009a)] and their enumeration [Xiang et al(2009b)]. Models $M = (e,C,T,PS)$ and $M' = (e,C,T',PS)$ may be consistent even though they are not structurally consistent. By limiting T and T' within the space of minimal NIN-AND trees, model consistency implies structure consistency in general. This means that a unique minimal tree exists for each pattern of causal interactions among a set of causes.

Definition 2. *Let T be a NIN-AND tree. If T contains a gate t that outputs to a gate g of the same type (direct or dual), delete t and connect its inputs to g. Apply such deletion until no longer possible. The resultant NIN-AND tree is minimal.*

The uniqueness of minimal NIN-AND trees allows them to be enumerated explicitly, e.g., using the two-step enumeration algorithm in [Xiang et al(2009b)]. For binary effect and causes, if $|C| = 4$, there are 52 minimal NIN-AND trees. For $|C| = 5,6,7$, the number is 472, 5504, 78416, respectively.

We **propose** the SE technique as follows. Denote $n = |C|$. First, a set PS_e of n single-causal probabilities, e.g., $P_e(e^+|c_i^+)$, are elicited from the expert, where subscript e denotes 'elicited'. Then the set TM of minimal NIN-AND trees over C are enumerated. Combining each $T \in TM$ with PS_e, a set NM_e of NIN-AND tree models is obtained. In general, a unique CPT over e and C can be constructed from each model in NM_e. A set CPT_e of CPTs is thus defined. Note that there is a one-to-one mapping between TM and NM_e, and generally also between NM_e and CPT_e.

Subsequently, the expert is asked to assess some multi-causal probabilities. Let $P_e(e^+|c_i^+,c_j^+,c_k^+)$ be elicited from an expert, and $P'(e^+|c_i^+,c_j^+,c_k^+)$ be from a CPT $P'(e|C) \in CPT_e$. If $P'(e^+|c_i^+,c_j^+,c_k^+)$ differs significantly from $P_e(e^+|c_i^+,c_j^+,c_k^+)$, $P'(e|C)$ is deemed to be *inconsistent* with the true CPT, and the NIN-AND tree model corresponding to $P'(e|C)$ is eliminated from the candidate set NM_e. Based on such comparison of CPTs in CPT_e and elicited multi-causal probabilities, all models in NM_e except one, $M_e = (e,C,T_e,PS_e)$, will be eliminated. M_e is returned as the indirectly elicited model and T_e is the indirectly elicited NIN-AND tree. Below, we investigate several variations for elicitation and elimination procedures:

[Threshold based sequential elimination] Since elicitation from an expert is sequential, it is natural to interleave model elimination with elicitation. Elicitation and elimination proceed in rounds. Each round starts with elicitation of a multi-causal probability, followed by elimination of one or more inconsistent NIN-AND tree models. The process continues until a single model in NM_e remains in the last round.

The elimination operation requires a threshold s. Only when difference $\delta = |P_e(e^+|c_i^+, c_j^+, c_k^+) - P'(e^+|c_i^+, c_j^+, c_k^+)| > s$, $P'(e|C)$ is deemed inconsistent with the true CPT. However, choosing the adequate threshold value is difficult in practice for the reason below.

By Assumption 1, the expert assessment of single-causal probabilities PS_e is approximate. Hence, none of the models in NM_e is consistent with the true model M_t. Furthermore, by assumption, an elicited multi-causal probability may also differ from the corresponding true probability. Hence, δ above contain elicitation errors. If s is set too low, even if a model $M \in NM_e$ is structurally consistent with the true model M_t, it may still be eliminated because δ exceeds s. On the other hand, if s is set too high, multiple models structurally inconsistent with the true model M_t may pass each round, and no single model can be selected in the last round.

[Bounded sequential elimination] Elicitation and elimination proceed in K rounds, where K is the number of multi-causal probabilities to be elicited, is predetermined, and can be varied based on expert availability. In each round, after elicitation of a multi-causal probability, its difference δ from each CPT in CPT_e is calculated, a given number of models in NM_e with the minimum δ values are retained, and the other models are eliminated. The number of models retained in each round decreases over consecutive rounds, and it is one for the Kth round.

The threshold is no longer needed, and its drawback is avoided. Instead, a set of K bounds is used, one for the number of retained models in each round. For example, if $K = 4$, numbers of models retained in succeeding rounds can be 16, 8, 4, and 1.

One limitation is that the model returned may depend on the order in which the K multi-causal probabilities are elicited. The NIN-AND tree model $M \in NM_e$ that is structurally consistent with M_t (such M is unique whenever single-causal probabilities by $P_t(e|C)$ are distinct) may be eliminated in an earlier round. This occurs when the probability elicited in the current round is not distinguishing, and too many models in NM_e have similar, small δ values: If the bound for the current round is m, the model M may be eliminated because its δ value is slightly larger than that of the model ranked m. Whereas if multi-causal probabilities were elicited in another order, M may be retained in each round and returned in the end.

[Simultaneous elimination] Only one round of elicitation and elimination is conducted. A set PM_e of K multi-causal probabilities are first elicited. Its root-mean-square (rms) distance from the corresponding set PM' of multi-causal probabilities determined by each CPT in CPT_e is calculated:

$$d(PM_e, PM') = \sqrt{\frac{1}{K}\sum_{i=1}^{K}(P_e(e^+|\underline{x}_i^+) - P'(e^+|\underline{x}_i^+))^2} \quad (3)$$

The model in NM_e with the minimum distance will be returned.

The method overcomes the limitation on threshold or elicitation order by the two alternative procedures. It is thus used in the further investigation of the SE technique. Although any multi-causal probabilities may be used with the SE technique, in the remainder of the paper, we assume that they are triple-causal.

5 Generate NIN-AND Tree by Pairwise Causal Interaction

The PCI technique builds on the pairwise causal interaction function defined by a NIN-AND tree [Xiang et al(2009a)].

Proposition 1. *Let T be a minimal NIN-AND tree for effect e and its causes C. Then T defines a function pci from pairs of distinct causes $\{c_i, c_j\} \subset C$, where $i \neq j$, to the set $\{rif, udm\}$, where rif stands for reinforcing and udm for undermining.*

The pci function signifies explicitly the causal interaction between each pair of causes. For instance, the NIN-AND tree in Fig. 1 (right) defines the function: $pci(c_1, c_2) = rif, pci(c_1, c_3) = udm, pci(c_2, c_3) = rif$.

Let TM be the set of all minimal NIN-AND trees over n causes. Then each NIN-AND tree $T \in TM$ has a distinct pci function (exhaustively confirmed for $n = 3, ..., 10$). Hence, a NIN-AND tree can be identified from a given pci function.

Based on this idea, we **propose** the PCI technique for generating a NIN-AND tree as follows: First, elicit a set PS_e of single-causal probabilities from the expert, and enumerate the set TM, as done in the SE technique. From TM, a set $PCIF$ of pci functions, one for each NIN-AND tree $T \in TM$ is defined. Then, a set PD_e of all double-causal probabilities (a total of $n\ (n-1)/2$ values) are elicited from the expert.

From PS_e and PD_e, a pci function $pci_e()$ can be determined according to Def. 1. For example, suppose the CPT in Table 1 is the true CPT, elicited single-causal probabilities include $P_e(e^+ \leftarrow c_2^+) = 0.82$, $P_e(e^+ \leftarrow c_3^+) = 0.67$, and elicited double-causal probabilities include $P_e(e^+ \leftarrow c_2^+, c_3^+) = 0.91$. From $P_e(e^+ \leftarrow c_2^+, c_3^+) > P_e(e^+ \leftarrow c_2^+)$ and $P_e(e^+ \leftarrow c_2^+, c_3^+) > P_e(e^+ \leftarrow c_3^+)$, the function value $pci(c_2, c_3) = rif$ can be determined.

Subsequently, the derived $pci_e()$ is compared against functions in $PCIF$. If $pci_e()$ matches $pci'() \in PCIF$, then the NIN-AND tree $T' \in TM$ that produces $pci'()$ will be returned.

The key operation of the PCI technique is the derivation of $pci_e()$ function from PS_e and PD_e. Below, we consider how to carry out the operation in practice. For any pair of causes c_i and c_j, $pci(c_i, c_j) \in \{rif, udm\}$. By Def. 1, $pci(c_i, c_j) = rif$ iff

$$P(e^+ \leftarrow c_i^+, c_j^+) \geq \max(P(e^+ \leftarrow c_i^+), P(e^+ \leftarrow c_j^+)), \tag{4}$$

and $pci(c_i, c_j) = udm$ iff

$$P(e^+ \leftarrow c_i^+, c_j^+) < \min(P(e^+ \leftarrow c_i^+), P(e^+ \leftarrow c_j^+)). \tag{5}$$

Therefore, in theory, it suffices to compare $P(e^+ \leftarrow c_i^+, c_j^+)$ and $P(e^+ \leftarrow c_i^+)$, and use the outcome to determine the value for $pci(c_i, c_j)$.

In practice, however, due to elicitation errors, it is possible that

$$P_e(e^+ \leftarrow c_i^+) < P_e(e^+ \leftarrow c_i^+, c_j^+) < P_e(e^+ \leftarrow c_j^+).$$

For example, if $P_t(e^+ \leftarrow c_i^+) = 0.6$, $P_t(e^+ \leftarrow c_j^+) = 0.9$, and c_i undermines c_j, we have $P_t(e^+ \leftarrow c_i^+, c_j^+) = 0.54$. Elicited values, however, may be

$$P_e(e^+ \leftarrow c_i^+) = 0.56 < P_e(e^+ \leftarrow c_i^+, c_j^+) = 0.59 < P_e(e^+ \leftarrow c_j^+) = 0.93$$

due to elicitation errors. Similarly, when c_i reinforces c_j, we have $P_t(e^+ \leftarrow c_i^+, c_j^+) = 1 - (0.4 * 0.1) = 0.96$, while elicited values may be

$$P_e(e^+ \leftarrow c_i^+) = 0.56 < P_e(e^+ \leftarrow c_i^+, c_j^+) = 0.91 < P_e(e^+ \leftarrow c_j^+) = 0.93.$$

When these happen, comparing $P_e(e^+ \leftarrow c_i^+, c_j^+)$ against one of $P_e(e^+ \leftarrow c_i^+)$ and $P_e(e^+ \leftarrow c_j^+)$ has a 0.5 chance to assign pci function value incorrectly. Comparing against both is not even feasible, because Eqn. (4) and (5) will both fail. To address this issue, we develop the following algorithm:

1. If Eqn. (4) holds for elicited probabilities, assign $pci(c_i, c_j) = rif$.
2. Else if Eqn. (5) holds for elicited probabilities, assign $pci(c_i, c_j) = udm$.
3. Else if
$$|P(e^+ \leftarrow c_i^+, c_j^+) - \min(P(e^+ \leftarrow c_i^+), P(e^+ \leftarrow c_j^+))|$$
$$< |P(e^+ \leftarrow c_i^+, c_j^+) - \max(P(e^+ \leftarrow c_i^+), P(e^+ \leftarrow c_j^+))|,$$
assign $pci(c_i, c_j) = udm$.
4. Else assign $pci(c_i, c_j) = rif$.

The algorithm handles normal cases (1 and 2) according to Eqn. (4) and (5). When elicitation errors fail these equations (cases 3 and 4), the pci function value is determined by assuming small errors. For the first example above, $pci(c_i, c_j) = udm$ will be assigned correctly due to case 3. For the second example, $pci(c_i, c_j) = rif$ will be assigned due to case 4.

It is possible that a derived function $pci_e() \notin PCIF$. That is, there exists no NIN-AND tree model that would produce the function $pci_e()$. The $pci_e()$ is said to be *invalid*. When this occurs, we apply a method in [Xiang(2010a)]: A valid pci function $pci_e^*()$ in *PCIF* which differs from $pci_e()$ the least will be selected, and its corresponding NIN-AND tree model will be returned as the indirectly elicited model.

6 Experimental Setup

To evaluate the effectiveness of SE and PCI techniques, human-based experiments are conducted, using an approach that extends that in [Zagorecki and Druzdzel(2004)]. A true causal model is simulated, from which a human is trained into an expert. A subset of causal probabilities are then elicited from the expert, from which a NIN-AND tree model is generated using the SE or PCI technique. The rms distance between the discovered model and the true model (similar to Eqn. (3)) is then measured to evaluate the effectiveness of these techniques. The experiment is organized into three stages elaborated below.

The first is *expert training*, during which each human participant is trained into an expert. A simulated NIN-AND tree model $M_t = (e, C, T, PS)$ is used as the true model, from which the true CPT $P_t(e|C)$ is constructed. Given the presence of a subset $X \subseteq C$ of active causes, an example $(e, \underline{x}^+)$, where $e \in \{e^+, e^-\}$, is generated by stochastic simulation from causal probability $P_t(e^+ \leftarrow \underline{x}^+)$. After seeing a sufficient number of examples for a sufficient number of distinctive $\underline{x}^+$ (detailed below), the participant is deemed to be an expert on model M_t.

To ensure that a participant's knowledge on M_t is obtained entirely from the training, and is not biased by outside experience, we presented M_t to be about phenomena from an imaginary planet. A software Environment Simulator (ES) is implemented accordingly to allow a participant to specify active causes $\underline{x}^+$ and observe simulated effects e. Note that this setup ensures condition 1 of Assumption 1.

The second stage is *elicitation*, during which a subset of causal probabilities $P_e(e^+ \leftarrow \underline{x}^+)$ are elicited from the expert. As stated in Assumption 1, generally, $P_e(e^+ \leftarrow \underline{x}^+) \neq P_t(e^+ \leftarrow \underline{x}^+)$. Their difference has so far been referred to as *elicitation error*, but in fact is the combination of two sources of errors.

1. Sampling error: Assuming $P_e(e^+ \leftarrow \underline{x}^+)$ is based on observed relative frequency $F(e^+ \leftarrow \underline{x}^+) = N(e^+ \leftarrow \underline{x}^+)/N(\underline{x}^+)$, where $N(e^+ \leftarrow \underline{x}^+)$ is the number of observations of example $(e^+, \underline{x}^+)$ and $N(\underline{x}^+)$ is the number of observations of $\underline{x}^+$, we have $F(e^+ \leftarrow \underline{x}^+) \neq P_t(e^+ \leftarrow \underline{x}^+)$ because $N(\underline{x}^+)$ is finite.
2. Retention-Articulation (RA) error: The participant may not be able to retain and articulate either $N(e^+ \leftarrow \underline{x}^+)$ and $N(\underline{x}^+)$, or $F(e^+ \leftarrow \underline{x}^+)$ accurately [Kahneman et al(1982)].

To ensure condition 2 of Assumption 1, both the sampling error and RA error need to be controlled. To control sampling error, we setup ES to enforce the requirement $N(\underline{x}^+) \geq 100$ for each $P_e(e^+ \leftarrow \underline{x}^+)$ to be elicited. That is, the participant must have a sufficient number of observations of $\underline{x}^+$ during training.

To control RA error, for each distinct $\underline{x}^+$, the frequency pair $F(e^+ \leftarrow \underline{x}^+)$ and $F(e^- \leftarrow \underline{x}^+)$ observed during the training stage is shown in a stacked bar graph (Fig. 2). The bar graph helps to reduce the RA error by providing a visual hint for the observed $F(e^+ \leftarrow \underline{x}^+)$. Yet, it does not eliminate RA error as it is visual, while $P_e(e^+ \leftarrow \underline{x}^+)$ is elicited numerically.

The final stage is *discovery*, during which the set of $P_e(e^+|\underline{x}^+)$ elicited is used to generate a NIN-AND tree model M_e.

Participants are recruited from university students (second year or above). Each participant is trained with a distinct true model $M_t = (e, C, T, PS)$. All models used have $|C| = 4$, but they differ in both T and PS.

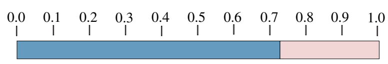

Fig. 2. A stacked bar graph where $F(e^+ \leftarrow \underline{x}^+) = 0.72$

Our objective is to evaluate the effectiveness of SE and PCI techniques. To facilitate the evaluation, we compare them against direct elicitation of each causal probability (all 15 parameters in $P_t(e|C)$). We refer to it as the *direct numerical* (DN) technique. For SE, we elicit 8 parameters (4 single-causal and 4 triple-causal). For PCI, we elicit 10 parameters (4 single-causal and 6 double-causal).

7 Experimental Results

Each data set consists of a number of causal probabilities elicited from one participant. A data set for evaluation of DN, SE, or PCI technique contains 15, 8 or 10 elicited probabilities, respectively, and the number of data sets collected are 23, 29, 29, respectively.

From the true CPT used to simulate training examples for a participant and probabilities elicited from the participant, the elicitation error (Section 6) of the participant is measured by the rms distance between the true CPT and elicited probabilities. The mean and standard deviation of elicitation errors over all participants are shown in Table 2 (column 4). Th elicitation error consists of sampling and RA errors (Section 6). From ES log of examples generated for training a participant and the true CPT used in example generation, the sampling error of training examples is measured by the rms distance between example frequencies and the true CPT. From the log of examples generated for training a participant and elicited probabilities, RA error of the participant is measured by rms distance between example frequencies and elicited probabilities. The means and standard deviations of sampling and RA errors over all participants are also shown in the table (columns 2 and 3). It can been seen that our elicitation aid by stacked bar graphs has effective control of the RA error. Hence, the elicitation error is composed mainly of the sampling error.

The DN technique directly elicits a CPT from the expert, which we refer to as the *CPT elicited with the DN technique*. On the other hand, for each data set collected for SE evaluation, the SE technique is applied to generate a NIN-AND tree model, from which a CPT is constructed. We refer to it as the *CPT elicited with the SE technique*. The *CPT elicited with the PCI technique* is similarly defined.

For each data set, the CPT elicited by the corresponding technique is compared against the true CPT used to drive expert training, and the rms distance between the two CPTs is calculated. For each of DN, SE, and PCI technique, the mean and standard deviation over the corresponding data sets are summarized in Table 3.

Results from all three techniques are comparable. Note that PCI technique depends on single and double-causal probabilities (10), SE technique depends on single and triple-causal probabilities (8), while DN technique depends on all causal probabilities (15). Hence, the results demonstrate that both SE and PCI techniques improve efficiency in CPT acquisition while maintaining comparable accuracy.

Table 2. Mean (μ) and standard deviation (σ) of errors over all participants

	Sampling Errors	RA Errors	Elicitation Errors
μ	0.0293	0.0076	0.0301
σ	0.0096	0.0038	0.0099

Table 3. Mean (μ) and standard deviation (σ) of model distance by DN, SE and PCI techniques

	DN	SE	PCI
μ	0.0301	0.0356	0.0281
σ	0.0099	0.0343	0.0146

8 Simulation Study

Due to resource involved in human-based experiments, large numbers of participants and multiple setups are not feasible. To compensate this limitation, we enhanced human experiments with simulation-based studies.

For the DN technique, we simulated a true model $M_t = (e, C, T_t, PS_t)$ and constructed the true CPT $P_t(e|C)$ from M_t. For each subset $X \subseteq C$ of active causes, K examples $(e, \underline{x}^+)$ are stochastically generated from $P_t(e^+|\underline{x}^+)$. The elicited probability $P_e(e^+|\underline{x}^+)$ is simulated as the ratio between the number of examples $(e^+, \underline{x}^+)$ and K. This is justified by two observations. First, the elicitation errors in human experiments are made up mainly by sampling errors (Table 2). Second, as we decrease K, the elicitation error $|P_e(e^+|\underline{x}^+) - P_t(e^+|\underline{x}^+)|$ will increase. Hence, simulated elicitation errors can be well controlled through K.

After the elicited CPT $P_e(e|C)$ is thus simulated, we calculate the rms distance between $P_e(e|C)$ and $P_t(e|C)$. We repeat the above for W true models, and the effectiveness of the DN technique is evaluated by the mean distance from the W trials.

For the PCI technique, the true model $M_t = (e, C, T_t, PS_t)$ and true CPT $P_t(e|C)$ are simulated as above. A set $PS_e = \{P_e(e|c_i^+)\}$ of single-causal elicited probabilities and a set $PD_e = \{P_e(e|c_i^+, c_j^+)\}$ of double-causal elicited probabilities are simulated from $P_t(e|C)$. Applying the PCI technique to PS_e and PD_e, an indirectly elicited model $M_e = (e, C, T_e, PS_e)$ is generated.

From M_e, the elicited CPT $P_e(e|C)$ is constructed and the rms distance between $P_e(e|C)$ and $P_t(e|C)$ calculated. The effectiveness of the PCI technique is evaluated by repeating the above for W true models, and obtaining the mean distance.

For the SE technique, a set PS_e of single-causal elicited probabilities and a set $PT_e = \{P_e(e|c_i^+, c_j^+, c_k^+)\}$ of triple-causal elicited probabilities are simulated from $P_t(e|C)$. The set of all NIN-AND tree models $NM_e = \{(e, C, T, PS_e)\}$ are obtained by enumeration. Note that each model $M \in NM_e$ has a distinct NIN-AND tree topology T, but has the same PS_e. An indirectly elicited NIN-AND tree model M_e is then selected from NM_e if its corresponding CPT has the minimum distance from PT_e.

From M_e, CPT $P_e(e|C)$ is constructed and the rms distance between $P_e(e|C)$ and $P_t(e|C)$ is calculated. The SE technique is evaluated by the mean distance from simulation over W true models.

In simulation studies for the three techniques, we used $K = 100$ and $W = 1000$. $K = 100$ is chosen so that magnitudes of simulated elicitation errors are similar to those observed in the human-based study. $W = 1000$ is used as higher W values do not show significant difference in outcomes. For each technique, simulations are run for each of $n = |C| = 4, 5, 6, 7$. Table 4 shows the number of causal probabilities simulated for each technique and each n value.

Table 4. Number of simulated causal probabilities used by DN, SE and PCI studies

n	# CPT probs	# probs for DN	# probs for SE	# probs for PCI
4	16	15	8	10
5	32	31	15	15
6	64	63	26	21
7	128	127	42	28

The second column shows the number of independent probability parameters in $P(e|C)$, which is 2^n. The third column shows the number of elicited probabilities simulated by DN evaluation, which is $2^n - 1$, because NIN-AND tree models satisfy $P(e^+|\underline{c}^-) = 0$. The fourth column shows the count for SE evaluation, which is $n + C(n,3)$. The last column shows the count for PCI evaluation, which is $n + C(n,2)$.

Results from simulation-based studies are summarized in Table 8. Means and standard deviations of model distances for the three techniques are shown in columns 2, 3, 4, 5, 7, 8. Columns 6 and 9 show percentages of models indirectly elicited by SE and PCI that recover true tree topology T_t. The last column shows percentages of indirectly elicited *pci* functions that are invalid.

Table 5. Model distance by DN, SE and PCI techniques from simulation study

n	DN (μ)	DN (σ)	SE (μ)	SE (σ)	Rcv (%)	PCI (μ)	PCI (σ)	Rcv (%)	Ivad (%)
4	0.0363	0.0099	0.0470	0.0485	79.6	0.0352	0.0340	98.5	0.9
5	0.0368	0.0086	0.0352	0.0268	86.5	0.0369	0.0397	98.1	0.5
6	0.0364	0.0076	0.0317	0.0215	88.2	0.0338	0.0237	95.7	2.2
7	0.0356	0.0076	0.0311	0.0183	85.8	0.0344	0.0284	94.2	3.6

The mean distances for DN indicate the magnitudes of simulated elicitation errors in the studies of all three techniques, since the same $K = 100$ value is used. Note that the magnitudes are slightly higher than that observed in human-based experiments (Table 2).

Comparing columns 6 and 9, PCI technique performs better than SE in recovering true NIN-AND tree topology. On the other hand, although SE technique is less accurate in tree recovery, the mean model distance and standard deviation for $n = 5, 6, 7$ are slightly smaller than PCI. This observation shows that given the existence of elicitation errors, multiple NIN-AND tree models may generate similar CPTs, and the SE technique is robust under such condition. We attribute the reverse performance difference when $n = 4$, i.e., $SE(\mu) > PCI(\mu)$, to the number of elicited probabilities used (8 for SE and 10 for PCI).

Overall, SE and PCI techniques achieved the comparable model distance in comparison with DN technique, while requiring a much less number of elicited probabilities. In general, the number of probabilities to be elicited by the DN technique is $O(2^n)$. The number is $O(n^3)$ for SE and $O(n^2)$ for PCI. The performance of PCI technique makes it particularly attractive: It achieves about the same elicitation accuracy while requiring the smallest number of elicitations. For instance, when $n = 7$, DN requires 127, SE requires 42, while PCI requires only 28.

Finally, column 10 shows that although elicitation errors sometimes cause failure in constructing the pci function, our fault-tolerance method recovers from the failure well. Not only a valid NIN-AND tree model is returned under the failure condition, but the model is sufficiently close to the true model (shown by columns 7 and 8).

9 Conclusion

NIN-AND tree causal models provide an efficient tool for CPT acquisition in construction of Bayes nets. Direct elicitation of such a model involves elicitation of a number (linear in n) of single-causal probabilities, and a NIN-AND tree (of a size linear in n). The tree elicitation step requires nontrivial training of an expert on the syntax and semantics of these models, as well nontrivial mental exercise by the expert to identify correctly the partial order of interactions among causes.

In this work, we investigate the novel idea to substitute direct elicitation of a NIN-AND tree with elicitation of some multi-causal probabilities. The NIN-AND tree is then automatically generated based on elicited probabilities. We propose two alternative techniques that implement this idea with low-order multi-causal probabilities. Our human-based and simulation-based studies demonstrated the feasibility of the idea. These techniques eliminate above-mentioned expert training and demanding mental exercise, while remaining efficient. Numbers of probabilities to be elicited are $O(n^3)$ and $O(n^2)$ for (triple-causal based) SE and PCI, respectively.

The main assumption these techniques depend on is the expert's ability to approximately assess required causal probabilities. Elicitation error can be decomposed into sampling error and RA error. The RA error may be reduced through training and/or technical aids, although detailed investigation is beyond the scope of this work. Sampling error may be controlled by the number of examples observed for each causal combination (i.e., $\underline{x}^+$). Our experiments have shown that 100 examples per causal combination is sufficient for our techniques to work well.

Acknowledgements. Financial support from NSERC, Canada to the first author is acknowledged.

References

[Galan and Diez(2000)] Galan, S., Diez, F.: Modeling dynamic causal interaction with Bayesian networks: temporal noisy gates. In: Proc. 2nd Inter. Workshop on Causal Networks, pp. 1–5 (2000)

[Heckerman and Breese(1996)] Heckerman, D., Breese, J.: Causal independence for probabilistic assessment and inference using Bayesian networks. IEEE Trans. on System, Man and Cybernetics 26(6), 826–831 (1996)

[Kahneman et al(1982)] Kahneman, D., Slovic, P., Tversky, A. (eds.): Judgment under uncertainty: heuristics and biases. Cambridge University Press, Cambridge (1982)

[Lemmer and Gossink(2004)] Lemmer, J., Gossink, D.: Recursive noisy OR - a rule for estimating complex probabilistic interactions. IEEE Trans. on System, Man and Cybernetics, Part B 34(6), 2252–2261 (2004)

[Maaskant and Druzdzel(2008)] Maaskant, P., Druzdzel, M.: An independence of causal interactions model for opposing influences. In: Jaeger, M., Nielsen, T. (eds.) Proc. 4th European Workshop on Probabilistic Graphical Models, Hirtshals, Denmark, pp. 185–192 (2008)

[Pearl(1988)] Pearl, J.: Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann, San Francisco (1988)

[Xiang(2010a)] Xiang, Y.: Acquisition and computation issues with NIN-AND tree models. In: Myllymaki, P., Roos, T., Jaakkola, T. (eds.) Proc. 5th European Workshop on Probabilistic Graphical Models, Finland, pp. 281–289 (2010a)

[Xiang(2010b)] Xiang, Y.: Generalized non-impeding noisy-AND trees. In: Proc. 23th Inter. Florida Artificial Intelligence Research Society Conf., pp. 555–560 (2010b)

[Xiang and Jia(2007)] Xiang, Y., Jia, N.: Modeling causal reinforcement and undermining for efficient cpt elicitation. IEEE Trans. Knowledge and Data Engineering 19(12), 1708–1718 (2007)

[Xiang et al(2009a)] Xiang, Y., Li, Y., Zhu, J.: Towards effective elicitation of NIN-AND tree causal models. In: Godo, L., Pugliese, A. (eds.) SUM 2009. LNCS (LNAI), vol. 5785, pp. 282–296. Springer, Heidelberg (2009a)

[Xiang et al(2009b)] Xiang, Y., Zhu, J., Li, Y.: Enumerating unlabeled and root labeled trees for causal model acquisition. In: Gao, Y., Japkowicz, N. (eds.) AI 2009. LNCS (LNAI), vol. 5549, pp. 158–170. Springer, Heidelberg (2009b)

[Zagorecki and Druzdzel(2004)] Zagorecki, A., Druzdzel, M.: An empirical study of probability elicitation under Noisy-OR assumption. In: Proc. 17th Inter. Florida Artificial Intelligence Research Society Conf., pp. 880–885 (2004)