Racing for Conditional Independence Inference

Remco R. Bouckaert[1] and Milan Studený[2⋆]

1. Computer Science Department, University of Waikato & Xtal Mountain
Information Technology, New Zealand (`remco@cs.waikato.ac.nz, rrb@xm.co.nz`)
2. Institute of Information Theory and Automation, Academy of Sciences of the
Czech Republic, Prague, Czech Republic (`studeny@utia.cas.cz`)

Abstract. In this article, we consider the computational aspects of deciding whether a conditional independence statement t is implied by a list of conditional independence statements L using the implication related to the method of structural imsets. We present two methods which have the interesting complementary properties that one method performs well to prove that t is implied by L, while the other performs well to prove that t is not implied by L. However, both methods do not perform well the opposite. This gives rise to a parallel algorithm in which both methods race against each other in order to determine effectively whether t is or is not implied.

Some empirical evidence is provided that suggest this racing algorithms method performs a lot better than an existing method based on so-called skeletal characterization of the respective implication. Furthermore, the method is able to handle more than five variables.

1 Introduction

Conditional independence (CI) is a crucial notion in many calculi for dealing with knowledge and uncertainty in artificial intelligence [2, 3]. A powerful formalism for describing probabilistic CI structures is provided by the method of structural imsets [7]. In this algebraic approach, CI structures are described by certain vectors whose components are integers, called *structural imsets*. An important question is to decide whether a CI statement is implied by a set of CI statements. The method of structural imsets offers a sufficient condition for the probabilistic implication of CI statements. The offered inference mechanism is based on linear algebraic operations with imsets. The basic idea is that every CI statement can be translated into a simple imset and the respective algebraic relation between imsets, called *independence implication*, forces the probabilistic implication of CI statements. Techniques were developed in [5] to test the independence implication through systematic calculation when there are up to five variables involved.

For reasoning about CI statements with more than five variables one may resort to making severe assumptions. For example, one can assume that the CI

---

structure is graph isomorphic for a class of graphs such as directed acyclic graphs (DAG) [3, 8], undirected graphs (UG) [2], chain graphs (CG) [1], etc. Then CI inference from a set of CI statements of a special form, a so-called *input list*, can be made as follows. The list is used to construct a graph and CI statements are read from the graph through the respective graphical separation criterion. However, the assumption that the CI structure is graph isomorphic may be too strong in many cases and only special input lists can be processed anyway. Using the method of structural imsets, many more CI structures can be described than with DAGs, UGs or CGs. However, the computational effort required when more than five variables are involved is not clear at present.

Fortunately, structural imsets have some properties that we can exploit. First, a relatively easy sufficient condition for independence implication is that the respective linear combination of imsets can be decomposed into so-called *elementary imsets*. The existence of this decomposition can be found relatively quickly. On the other hand, to prove that the decomposition does not exist requires trying all decompositions, which often takes a long time. Second, there exists a method to show that the independence implication does not hold. It suffices to find a certain vector, called *supermodular function*, such that its inner product with the respective combination of structural imsets is negative. These supermodular functions can be generated randomly. This only allows us to disprove independence implication of imsets, not to disprove probabilistic implication of respective CI statements. However, if the obtained supermodular function is a multiple of a multiinformation function of a probability distribution [7] then it also allows us to disprove probabilistic implication of respective CI statements. Thus, we have one method that allows us to find a proof that a statement is implied, and one method to find a proof that a statement is not implied. However, both methods perform poorly in proving their opposite outcome. This gives rise to a race: both methods are started at the same time and the method that returns first also returns a proof whether the statement of interest is implied or not.

The following section introduces formal terminology and the fundamentals of CI inference using imsets. The racing algorithms are described in Section 3 where many more smaller optimizations are described as well. Section 4 presents experiments that were performed to get an impression of the run-times of various variants of inference algorithms. We conclude with some final comments and directions for further research.

2   Terminology

Let N be a set of variables $\{x_1, \ldots, x_n\}$ ($n \geq 1$), as will be assumed throughout the paper. Let X and Y be subsets of N. We use XY to denote the union of X and Y and $X \setminus Y$ to denote the set of variables that are in X but not in Y. Further, let x be a variable in N, then x will also denote the singleton $\{x\}$.

2.1 Conditional Independence

Let P be a discrete probability distribution over N and X, Y, Z pairwise disjoint subsets of N. We say that X is *conditionally independent* of Y given Z if $P(\mathbf{x}|\mathbf{yz}) = P(\mathbf{x}|\mathbf{z})$ for all configurations $\mathbf{x},\mathbf{y},\mathbf{z}$ of values for X, Y, Z with $P(\mathbf{yz}) > 0$. We write then $X \perp\!\!\!\perp Y \,|\, Z \,[P]$ or just $X \perp\!\!\!\perp Y \,|\, Z$, and call it a *CI statement.* It is well-known that CI follows some simple rules, known as the *semi-graphoid axioms* defined as follows ($X, Y, Z, W \subseteq N$ are pairwise disjoint):

Symmetry	$X \perp\!\!\!\perp Y \,\vert\, Z$	$\Rightarrow$	$Y \perp\!\!\!\perp X \,\vert\, Z,$
Decomposition	$X \perp\!\!\!\perp WY \,\vert\, Z$	$\Rightarrow$	$X \perp\!\!\!\perp Y \,\vert\, Z,$
Weak union	$X \perp\!\!\!\perp WY \,\vert\, Z$	$\Rightarrow$	$X \perp\!\!\!\perp W \,\vert\, YZ,$
Contraction	$X \perp\!\!\!\perp W \,\vert\, YZ \ \& \ X \perp\!\!\!\perp Y \,\vert\, Z$	$\Rightarrow$	$X \perp\!\!\!\perp WY \,\vert\, Z.$

The problem we address in this paper is the following *inference problem.* Let L be a set of CI statements, called an *input list* and t is a CI statement $X \perp\!\!\!\perp Y \,|\, Z$. Does L imply t? More formally, is it true that for any discrete distribution P for which all statements in L hold necessarily t holds as well? This is *probabilistic implication* of those CI statements. The semi-graphoid axioms do not cover this implication. For example,

$$X \perp\!\!\!\perp Y \,|\, WZ \ \& \ W \perp\!\!\!\perp Z \,|\, X \ \& \ W \perp\!\!\!\perp Z \,|\, Y \ \& \ X \perp\!\!\!\perp Y \,|\, \emptyset \quad \Leftrightarrow$$
$$\Leftrightarrow \quad W \perp\!\!\!\perp Z \,|\, XY \ \& \ X \perp\!\!\!\perp Y \,|\, Z \ \& \ X \perp\!\!\!\perp Y \,|\, W \ \& \ W \perp\!\!\!\perp Z \,|\, \emptyset$$

is also a valid rule [7]. In fact, there is no complete finite set of rules of this kind describing relationships between probabilistic CI statements [4]. A more powerful formalism to describe the properties of CI is provided by the method of structural imsets.

2.2 Imsets

An *imset* over N (abbreviation for **i**nteger-valued **m**ulti**set**) is an integer-valued function on the power set of N. It can be viewed as a vector whose components, indexed by subsets of N, are integers. Given $X \subseteq N$, we use δ_X to denote the identifier imset, that is, $\delta_X(X) = 1$ and $\delta_X(Y) = 0$ for all $Y \subseteq N$, $Y \neq X$. An imset associated with a CI statement $X \perp\!\!\!\perp Y \,|\, Z$ is $u_{\langle X,Y|Z \rangle} = \delta_{XYZ} + \delta_Z - \delta_{XZ} - \delta_{YZ}$. The imset associated with an input list L is then $u_L = \sum_{t \in L} u_t$.

The basic technique for inference of a statement t from an input list L using the method of structural imsets is based on the following property. If $n \cdot u_L$ (for some natural number $n \in \mathbb{N}$) can be written as u_t plus the sum of some imsets associated with CI statements then t is implied by L. This can be derived from results of [7]. For example, if L consists of a single statement $X \perp\!\!\!\perp WY \,|\, Z$ and t is $X \perp\!\!\!\perp Y \,|\, Z$, we have (with $n = 1$)

$$\begin{aligned}
n \cdot u_L &= \delta_{WXYZ} + \delta_Z - \delta_{XZ} - \delta_{WYZ} \\
&= (\delta_{XYZ} + \delta_Z - \delta_{XZ} - \delta_{YZ}) + (\delta_{WXYZ} + \delta_{YZ} - \delta_{XYZ} - \delta_{WYZ}) \\
&= u_t + u_{\langle X,W|YZ \rangle}.
\end{aligned}$$

Thus, $X \perp\!\!\!\perp WY \mid Z$ implies t and we have derived the decomposition rule of the semi-graphoid axioms. Realize that any statement in the decomposition on the right-hand side can be swapped with t, so those statements are implied too. This means that above we have derived weak union as well.

An *elementary imset* is an imset associated with an elementary CI statement $x \perp\!\!\!\perp y \mid Z$, namely $u_{\langle x,y|Z \rangle} = \delta_{xyZ} + \delta_Z - \delta_{xZ} - \delta_{yZ}$. It is convenient to denote the set of elementary imsets over N by $\mathcal{E}(N)$ or simply $\mathcal{E}$. A *structural imset* is an imset u that can be decomposed into elementary imsets when multiplied by positive natural number, that is,

$$n \cdot u = \sum_{v \in \mathcal{E}} k_v \cdot v$$

for some $n \in \mathbb{N}$ and $k_v \in \mathbb{Z}^+$. Note that every structural imset induces a whole CI structure through an algebraic criterion, which is omitted here. The attraction of the method of structural imsets is that every discrete probabilistic CI structure can be described in this way [7].

Let u, v be structural imsets over N. We say that u *independence implies* v and write $u \rightharpoonup v$ if there exists $k \in \mathbb{N}$ such that $k \cdot u - v$ is a structural imset. This terminology is motivated by the fact that $u \rightharpoonup v$ actually means that u encodes more CI statements than v – see Lemma 6.1 in [7]. If $v \in \mathcal{E}$ then the constant $k \in \mathbb{N}$ can be supposed lesser than a limit $k_{\max}$ depending on the number of variables $|N|$ – see Lemma 4 in [6]. However, the value of the exact limit $k_{\max}$ for $|N| \geq 6$ is not known. It follows from results of [5] that $k_{\max} = 1$ if $|N| \leq 4$ and $k_{\max} = 7$ if $|N| = 5$. In our computer programs for $|N| \geq 6$ we need a limit for k. Instead of the unknown exact theoretical limit $k_{\max}$ we use the number $2^{|N|}$. Although we have not a proof of that we believe that $k_{\max} \leq 2^{|N|}$. Now, we can reformulate our inference problem. Given an elementary CI statement t and an input list (of elementary CI statements) L we are going to test whether $u_L \rightharpoonup u_t$. This is a sufficient condition for probabilistic implication of t by L. However, in general, it is not a necessary condition for it.

3 Algorithms

This section introduces algorithms for testing the implication $u_L \rightharpoonup u_t$. In Section 3.1, we revisit a method based on skeletal characterization of structural imsets from [7] and optimize the method. In Section 3.2, an algorithm for verification of $u_L \rightharpoonup u_t$ is presented based on searching a decomposition of $k \cdot u_L - u_t$ into elementary imsets. Section 3.3 concentrates on a method of disproving $u_L \rightharpoonup u_t$ by exploiting properties of supermodular functions. Section 3.4 combines the two previous methods by letting them race against each other and the one that returns its outcome first has a proof whether $u_L \rightharpoonup u_t$ or not.

3.1 Skeletal characterization of independence implication

We will only consider the implementation details here. Technical details and motivation of this approach can be found in §6.2.2 of [7]. This skeletal charac-

terization is based on a particular set of imsets called the ℓ-skeleton, denoted as $\mathcal{K}_\ell^\diamond(N)$. It follows from Lemma 6.2 in [7] that, for this particular set of imsets, we have $u_L \rightharpoonup u_t$ iff

$$\text{for all } m \in \mathcal{K}_\ell^\diamond(N) \text{ if } \langle m, u_t \rangle > 0 \text{ then } \langle m, u_L \rangle > 0. \tag{1}$$

Recall that the inner product $\langle m, u \rangle$ of a function $m : \mathcal{P}(N) \to \mathbb{R}$ and an imset u is defined by $\sum_{S \subseteq N} m(S) \cdot u(S)$. Thus, to conclude $u_L \rightharpoonup u_t$, we just need to check the conditions in (1) for all imsets in the ℓ-skeleton.[1] It can be used to check which elementary imsets over five variables are implied in this sense by a user defining the input list.

The ℓ-skeleton for five variables consists of 117978 imsets, which break into 1319 permutational types with each involving at most 120 imsets. So, checking whether $u_L \rightharpoonup u_t$ requires at most 117978 operations [5]. However, if t is not implied by L, we might find out far earlier that (1) does not hold for a particular imset in $\mathcal{K}_\ell^\diamond(N)$. By ordering skeletal imsets such that imsets that are more likely to cause violation in (1) are tried earlier, the required time can be minimized. The likelihood of violating (1) by $m \in \mathcal{K}_\ell^\diamond(N)$ grows with the number of zeros in $\{\langle m, v \rangle \, ; \; v \in \mathcal{E}\}$. Thus, sorting skeletal imsets on basis of this criterion helps to speed up the inference. The second auxiliary criterion is the number of sets $S \subseteq N$ with $u(S) = 0$.

Unfortunately, the skeletal characterization approach is hard to extend to more than five variables. First, because finding all elements of the ℓ-skeleton for more than five variables is computationally infeasible. Second, because it appears that the size of the ℓ-skeleton grows extremely fast with a growing number of variables. Therefore, we will consider different approaches to perform the inference in the rest of the paper.

3.2 Verification algorithm

If an imset u is a combination of elementary imsets $u = \sum_{v \in \mathcal{E}} k_v \cdot v, \; k_v \in \mathbb{Z}^+$ then we say that it is a *combinatorial imset*. This is a sufficient condition for an imset to be structural and it is an open question if it is also a necessary condition [7]. The method to verify $u_L \rightharpoonup u_t$ presented in this section is based on testing whether $u \equiv k \cdot u_L - u_t$ is a combinatorial imset for some $k \in \mathbb{N}$.

Testing whether u is combinatorial can be done recursively, by checking, for each $v \in \mathcal{E}$, whether $u - v$ is combinatorial. Obviously, this naive approach is computationally demanding and it requires some guidance and extra tests in order to reduce the search space.

There are a number of sanity checks we can apply, before starting the search. First of all, let t be $X \perp\!\!\!\perp Y \mid Z$, then $u_L \rightharpoonup u_t$ implies there exists $W \supseteq XYZ$ with $u_L(W) > 0$. This can be shown by Proposition 4.4 from [7] where we use $m^{A\uparrow}$ with $A = XYZ$. Another sanity check is as follows. Whenever u is a

---

[1] An applet at http://www.utia.cas.cz/user_data/studeny/VerifyView.html uses this method.

structural imset and $S \subseteq N$ a maximal set with respect to inclusion satisfying $u(S) \neq 0$ then $u(S) > 0$. Likewise, $u(S) > 0$ for any minimal set satisfying $u(S) \neq 0$ – see Lemma 6.5 in [7].

To guide the search, for each elementary imset $v \in \mathcal{E}$, we define the *deviance* of v from a non-zero imset u as follows. Let $maxcard(u)$ be the cardinality of the largest set $S \subseteq N$ for which $u(S) \neq 0$. It follows from the notes above that if u is structural then $u(S) \geq 0$ whenever $|S| = maxcard(u)$. Then, with $v = u_{\langle x,y|Z \rangle}$,

$$dev(v|u) = \begin{cases} \infty & |xyZ| < maxcard(u) \text{ or } u(xyZ) \leq 0, \\ \sum_{S \subseteq N} |v(S) - u(S)| & \text{otherwise.} \end{cases}$$

Thus, the deviance of v from a combinatorial imset u is finite only if δ_{xyZ} has a positive coefficient in u and no set larger than $|xyZ|$ has a positive coefficient in u. We pick the elementary imset with the lowest deviance first. Observe that if u is a non-zero combinatorial imset then $v \in \mathcal{E}$ with finite $dev(v|u)$ exists.

The deviance is defined in such a way that the elementary imsets that cancel as many of the coefficients in u as possible are tried before the imsets that cancel out fewer of the coefficients. For example, let $u = u_{\langle x,wy|z \rangle} + u_{\langle x,y|z \rangle} = \delta_{xywz} + 2\delta_z - 2\delta_{xz} - \delta_{wyz} + \delta_{xyz} - \delta_{yz}$ and $v_1 = u_{\langle x,w|yz \rangle} = \delta_{xywz} + \delta_{yz} - \delta_{xyz} - \delta_{wyz}$ then $dev(v_1|u) = 8$ while $v_2 = u_{\langle w,z|xy \rangle} = \delta_{xywz} + \delta_{xy} - \delta_{wxy} - \delta_{xyz}$ has the deviance $dev(v_2|u) = 10$. Furthermore $v_3 = u_{\langle x,y|z \rangle}$ has infinite deviance since $|xyz| = 3$ while $maxcard(u) = 4$. Finally, $v_4 = u_{\langle w,y|rz \rangle}$ has infinite deviance as $u(rwyz) = 0$. Therefore, v_1 will be tried before v_2, while v_3 and v_4 will not be tried at all in this cycle.

Thus, the deviance leads our search in a direction where we can hope to find a proper decomposition. Obviously, if t is not implied by L, the verification algorithm can spend a long time searching through the complete space of possible partial decompositions.

3.3 Falsification algorithm

Falsification is based on supermodular functions. A *supermodular function* is a function $m : \mathcal{P}(N) \to \mathbb{R}$ such that, for all $X, Y \subseteq N$,

$$m(XY) + m(X \cap Y) - m(X) - m(Y) \geq 0\,.$$

Note that an equivalent definition is that $\langle m, v \rangle \geq 0$ for every $v \in \mathcal{E}$. For example, δ_N is a supermodular function. By a *supermodular imset* we understand an imset which is a supermodular function.

Theorem 1. *An imset u is structural iff $\langle m, u \rangle \geq 0$ for any supermodular function m and $\sum_{S, S \supseteq K} u(S) = 0$ for any $K \subseteq N$ with $|K| \leq 1$.*

Proof. The necessity of the conditions is easy for they both hold for elementary imsets and can be extended to structural imsets. The sufficiency follows from Theorem 5.1 in [7] which claims that the same holds for a finite subset of the class of supermodular functions, namely the ℓ-skeleton $\mathcal{K}_\ell^\diamond(N)$.

Thus, we can exploit Theorem 1 to disprove $u_L \rightharpoonup u_t$ by constructing non-negative supermodular imsets randomly and taking their inner products with $k \cdot u_L - u_t$. If u_t is elementary and, for all $1 \leq k \leq k_{\max}$, the inner product is negative then we can conclude that $\neg(u_L \rightharpoonup u_t)$. A random supermodular imset m can be generated by first generating a 'base' imset m_{base} and then by modifying it to ensure the resulting imset is supermodular. We randomly select the size n of the base, then randomly select n different subsets $S_1, \ldots, S_n$ of N and assign $m_{base} = \sum_{S \in \{S_1, \ldots, S_n\}} k_S \cdot \delta_S$ where k_S are randomly selected integers in the range from 1 to $2^{|N|}$. Selecting larger values of the coefficients k_S would not make difference. On the other hand, they also would not help.

Now, m_{base} needs to be modified to ensure that the obtained function m is supermodular. We perform the following operation on m_{base}. Let $S_1, \ldots, S_{2^{|N|}}$ be an ordering of the subsets of N with $S_j \subseteq S_i \Rightarrow j \leq i$. For $i = 1, \ldots, 2^{|N|}$ define $m(S_i)$ to be the maximum of $m_{base}(S_i)$ and $m(S_i \setminus x) + m(S_i \setminus y) - m(S_i \setminus xy)$ for all $x, y \in S_i$. This ensures that $\langle m, v \rangle \geq 0$ for all $v \in \mathcal{E}$ and we have constructed an imset m which is supermodular.

Note that this technique can be used to disprove $u_L \rightharpoonup u_t$ but it cannot be used to prove it.

3.4 Racing algorithms for a proof

Typically, the verification algorithm from Section 3.2 can quickly find a decomposition of $k \cdot u_L - u_t$ into $\sum_{v \in \mathcal{E}} k_v \cdot v$, which proves that t is implied by L. Nevertheless, if $\neg(u_L \rightharpoonup u_t)$, the verification algorithm may spend a long time before it exhausts the whole space of possible decompositions of $k \cdot u_L - u_t$. However, the falsification algorithm from Section 3.3 can find a supermodular imset m with $\langle m, k \cdot u_L - u_t \rangle < 0$, which proves u_t is not implied by u_L. On the other hand, it will not be able to prove that $u_L \rightharpoonup u_t$.

We can combine the two algorithms by starting two threads, one with the verification algorithm and one with the falsification algorithm. The one that finds a proof first, returns its outcome and stops the other thread. Figure 1 illustrates the algorithm.

```
Algorithm: Racing for inference with structural imsets
Input: Input list L, CI statement t
1: thread1 = new RaceThread(Verify(L, t, proof))
2: thread2 = new RaceThread(Falsify(L, t, proof), thread1)
4: thread1.start(); thread2.start()
5: thread1.join() // wait for thread1 to stop
// if thread2 finished first, it will stop thread1
6: thread2.stop()
return proof
```

Fig. 1. Racing algorithm.

4 Experiments

We would like to judge the algorithms above on computational speed. However, it is hard to get a general impression of the performance of the algorithms, because it depends on the distribution of inference problems, which is unknown.

Still, we think we can get a representative impression of the relative performance of the algorithms by generating inference problems randomly and measuring the computation speed. We generated inference problems over five variables so that we can compare the performance of the skeleton-based algorithm from Section 3.1 with the others. A thousand input lists each were generated by randomly selecting 3,4 up to 10 elementary CI statements, giving a total of 8000 input lists. The algorithms described in Section 3 were applied to this class of lists with each of the elementary CI statements that were not in the list. This gave 1000×77 inference problems for input lists with 3 statements, 1000×76 inference problems for input lists with 4 statements, etc. In total, this created $1000 \times ([80 - 3] + [80 - 4] + \ldots + [80 - 10]) = 588.000$ inference problems over five variables.

Figure 2 shows the total number of elementary CI statements that are implied (labeled by Accept) and not implied (labeled by Reject) grouped by the number of elementary CI statements (3, 4 up to 10) in the input list. Naturally, the number of implied statements increases with increased input list size.

Fig. 2. Total number of rejects and accepts per experiment over 5 variables for various input list sizes. The size of the input list is shown on the x-axis. The number of rejects, accepts and total of unknown elementary statements is shown on the y-axis.


Figure 3 shows the total run-times for running the experiments comparing skeleton-based testing with sorted skeleton-based testing. We distinguish between run-time for accepts, rejects and total because the run-time for accepts is not influenced by the order of skeletal imsets as all of them need to be inspected. Indeed, run-times for accepts hardly differed (run-times only slightly differ due

to the fact that at random intervals garbage collection and other processes were performed). Run-times for rejects are reduced by about one order of magnitude so that total run-times are about halved. Thus, sorting the skeleton indeed helps significantly.

Fig. 3. Original skeleton-based testing compared with sorted skeleton-based testing. Sequences marked with asterisk are results for the sorted testing.


Figure 4 shows the striking difference in reject times for the racing algorithms method from Section 3.4 and the skeleton-based method from Section 3.1, which clearly favors the new method. Only input lists of size 10 are shown, but the shapes for input lists of other size are the same.

Fig. 4. Distribution of reject times of sorted skeleton-based method and racing algorithms method for input lists of size 10. The x-axis shows time, and the y-axis the number of elementary statements rejected in that time.

Unfortunately, the distribution of accept times shows a different picture, as illustrated in Figure 5. The graph for skeleton-based method shows just one

peak around 6 seconds per elementary CI statement, because that is how long it approximately takes to visit all skeletal imsets. The graph for the racing algorithms[2] shows a peak close to 10 milliseconds, that drops off pretty quickly. Shapes for input lists of other size look very similar, though the tail gets thinner with decreasing size of input lists.

Fig. 5. Distribution of accept times of the sorted skeleton-based method and the racing algorithms method for input lists of size 10. The x-axis shows time, and the y-axis the number of elementary statements accepted in that time.

An alternative approach is to only run the falsification algorithm and run it long enough that the complete space of elementary statements is covered. Table 1 shows the number of fails[3] of the falsification algorithm.

Table 1. Number of fails of the falsification algorithm with two different methods of generating random base imsets and various input list sizes (times $1000 \times k_{\max}$).

| $|L|$	Rnd 1	Rnd 2					
	1	1	2	3	4	5	20
3	1	0	0	0	0	0	0
4	19	2	0	0	0	0	0
5	57	18	3	6	2	3	1
6	147	50	37	24	18	16	5
7	243	92	61	39	46	42	21
8	429	189	144	124	109	95	48
9	423	195	138	112	97	92	46
10	547	299	239	201	192	193	110

---

[2] It is actually enlargement of the graph for the verification algorithm since the falsification thread cannot return acceptance.

[3] These are those elementary CI statements that are not implied by the input list but the algorithm did not succeed to identify them in a fixed time limit.

Two methods of generating random 'base' imsets are compared. The first method draws weights from the interval 1 to 32 for randomly selected subsets, while the second always selects 1. The second method appears far more effective in identifying rejections as one can judge from the number of fails in the columns labeled 1 in Table 1. We also looked at the impact of the number of randomly selected supermodular imsets on the number of fails. Increasing this number decreases the failure rate, but the rate only drops very slowly. Even when generating the same number of supermodular functions as the number of skeletal imsets in the skeleton-based method, not all statements are correctly classified.

Figure 6 shows run-times of the racing algorithms method compared with pure falsification algorithm (without the verification part). While reject times are about a third on average for pure falsification, non-reject times are about four times larger than the accept times of the combined algorithm.

Fig. 6. Racing algorithms vs. sole falsification algorithm. Sequences marked with asterisk are results for the falsification.

The same experiments as for five variables were performed with six variables, but obviously the skeleton-based algorithm was not applied on these problems. Apart from longer run-times of the algorithms, all observation as for five variables were confirmed.

5   Conclusions

We considered the computational aspects of performing CI inference using the method of structural imset, that is, deciding whether a CI statement t follows from an input list L of CI statements in that sense. The existing skeleton-based algorithm [5] that allows inference with up to five variables was improved. We presented an algorithm for creating a constructive proof that t follows from L. Unfortunately, this method does not perform well if t is not implied by L. Fortunately, we can prove t is not implied by L by randomly generating supermodular

functions and testing whether the inner product based on L and t is negative. But this method cannot be used to give a conclusive proof that t is implied by L. Together, these methods can race against each other on the same problem.[4]

Empirical evidence suggests the mode of the run-time of the racing algorithms method is an order of magnitude less than the skeleton-based method. Furthermore, the new method also works well for problems with more than five variables, unlike the old one. An analysis of accept times of the new method indicates that the verification algorithm sometimes cannot find the decomposition efficiently. This suggests that it can benefit from further guidance.

Some questions remain open, in particular finding an upper estimate on $k_{\max}$ (see Section 2.2) for six and more variables. A good upper estimate can decrease the computational effort in proving t is not implied by L.

Though the falsification algorithm cannot give a conclusive proof that an statement t is implied by L, we found that it was often very good at finding all elementary CI statements that are not implied by L in our experiments. This suggests that one can have some confidence that the falsification algorithm can identify statements that are implied by L. Deriving theoretical bounds on the probability that the falsification algorithm actually correctly identifies such statements would be interesting, since this would allow us to quantify our confidence.

References

1. R.R. Bouckaert and M. Studený, Chain graphs: semantics and expressiveness, in Symbolic and Quantitative Approaches to Reasoning and Uncertainty (C. Froidevaux, J. Kohlas eds.), Lecture Notes in AI 946, Springer-Verlag 1995, 67-76.
2. R.G. Cowell, S.L. Lauritzen, A.P. Dawid, D.J. Spiegelhalter, Probabilistic Networks and Expert Systems, Springer-Verlag, New York, 1999.
3. J. Pearl, Probabilistic Reasoning in Intelligent Systems, Morgan Kaufmann, San Mateo, 1988.
4. M. Studený, Conditional indpendence relations have no finite complete characterization, in Information Theory, Statistical Decision Functions and Random Processes vol. B (S. Kubík, J.Á. Víšek eds.), Kluwer, Dordrecht, 1999, 377-396.
5. M. Studený, R.R. Bouckaert, T. Kočka, Extreme supermodular set functions over five variables, research report n. 1977, Institute of Information Theory and Automation, Prague, January 2000.
6. M. Studený, Structural imsets: an algebraic method for describing conditional independence structures, in Proceedings of IPMU 2004 (B. Bouchon-Meunier, G. Coletti, R.R. Yager eds.), 1323-1330.
7. M. Studený, Probabilistic Conditional Independence Structures, Springer-Verlag, London, 2005.
8. T. Verma and J. Pearl, Causal networks: semantics and expressiveness, in Uncertainty in Artificial Intelligence 4 (R.D. Shachter, T.S. Lewitt, L.N. Kanal, J.F. Lemmer eds.), North-Holland, Amsterdam, 1990, 69-76.

---

[4] An applet is available at http://www.cs.waikato.ac.nz/~remco/ci/index.html