Article 25fa pilot End User Agreement

This publication is distributed under the terms of Article 25fa of the Dutch Copyright Act (Auteurswet) with explicit consent by the author. Dutch law entitles the maker of a short scientific work funded either wholly or partially by Dutch public funds to make that work publicly available for no consideration following a reasonable period of time after the work was first published, provided that clear reference is made to the source of the first publication of the work.

This publication is distributed under The Association of Universities in the Netherlands (VSNU)'Article 25fa implementation' pilot project. In this pilot research outputs of researchers employed by Dutch Universities that comply with the legal requirements of Article 25fa of the Dutch Copyright Act are distributed online and free of cost or other barriers in institutional repositories. Research outputs are distributed six months after their first online publication in the original published version and with proper attribution to the source of the original publication.

You are permitted to download and use the publication for personal purposes. Please note that you are not allowed to share this article on other platforms, but can link to it. All rights remain with the author(s) and/or copyrights owner(s) of this work. Any use of the publication or parts of it other than authorised under this licence or copyright law is prohibited. Neither Radboud University nor the authors of this publication are liable for any damage resulting from your (re)use of this publication.

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please contact the Library through email: copyright@ubn.ru.nl, or send a letter to:

University Library
Radboud University
Copyright Information Point
PO Box 9100
6500 HA Nijmegen

You will be contacted as soon as possible.

Causality on cross-sectional data: Stable specification search in constrained structural equation modeling

Ridho Rahmadi^{a,b,*}, Perry Groot^b, Marianne Heins^d, Hans Knoop^c, Tom Heskes^b, The OPTIMISTIC consortium¹

^a Department of Informatics, Universitas Islam Indonesia, Indonesia

^b Institute for Computing and Information Sciences, Radboud University Nijmegen, The Netherlands

^c Expert Centre for Chronic Fatigue, Radboud University Medical Centre, Nijmegen, The Netherlands

^d Netherlands Institute for Health Services Research, Utrecht, The Netherlands

ARTICLE INFO

Article history:

Received 13 December 2015

Received in revised form 8 July 2016

Accepted 4 October 2016

Available online 11 October 2016

Keywords:

Causal modeling

Structural equation model

Stability selection

Multi-objective evolutionary algorithm

NSGA-II

ABSTRACT

Causal modeling has long been an attractive topic for many researchers and in recent decades there has been a surge in theoretical development and discovery algorithms. Generally discovery algorithms can be divided into two approaches: constraint-based and score-based. The constraint-based approach is able to detect common causes of the observed variables but the use of independence tests makes it less reliable. The score-based approach produces a result that is easier to interpret as it also measures the reliability of the inferred causal relationships, but it is unable to detect common confounders of the observed variables. A drawback of both score-based and constrained-based approaches is the inherent instability in structure estimation. With finite samples small changes in the data can lead to completely different optimal structures. The present work introduces a new hypothesis-free score-based causal discovery algorithm, called stable specification search, that is robust for finite samples based on recent advances in stability selection using subsampling and selection algorithms. Structure search is performed over structural equation models. Our approach uses exploratory search but allows incorporation of prior background knowledge. We validated our approach on one simulated data set, which we compare to the known ground truth, and two real-world data sets for chronic fatigue syndrome and attention deficit hyperactivity disorder, which we compare to earlier medical studies. The results on the simulated data set show significant improvement over alternative approaches and the results on the real-world data sets show consistency with the hypothesis driven models constructed by medical experts.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Causal modeling has been an attractive topic for many researchers for decades. Especially since the 1990s there has been an enormous increase in theoretical development, partly because of advances in graphical modeling [2]. This has led to a variety of causal discovery algorithms in the literature. In general, causal discovery algorithms can be divided into two approaches: constraint-based and score-based. Constraint-based approaches work with conditional independence tests. First, they construct a skeleton graph starting with the complete graph and excluding

edges between variables that are conditionally independent. Second, edges are oriented to arrive at a causal graph. Examples of constraint-based approaches are the IC algorithm [3], PC-FCI [4], and TC [5]. Constraint-based approaches do not have to rely on the causal sufficiency assumption, and then can detect common causes of the observed variables [4]. A disadvantage of this approach is the use of independence tests on a large number of conditioning variables, making it less reliable [6]. Score-based approaches assign scores to particular graph structures based on the data fit and the complexity of the graph. Different scoring metrics that are often used are the Bayesian score [7] and the BIC score [8]. An example of a score-based method is *greedy equivalence search* (GES) [9]. The goal of the score-based approach is to find the graph structure with the highest score. An advantage of this approach is that it measures the reliability of the inferred causal relationships, which makes the result easy to interpret [10]. Score-based approaches typically do make the causal sufficiency assumption, and then cannot detect

* Corresponding author at: Faculty of Science, Toernooiveld 212, 6525EC Nijmegen, The Netherlands.

E-mail address: r.rahmadi@cs.ru.nl (R. Rahmadi).

¹ The members of OPTIMISTIC consortium are described in [1].

common confounders of the observed variables. Moreover, the involved optimization problem is usually NP-hard, so that different search heuristics are often used. The approach advocated in this paper is an example of a score-based approach.

Furthermore, in causal modeling based on observational data, the causal models are undetermined unless a preference for parsimonious models over more complex models is made [6]. In score-based approaches, such simplicity assumptions are typically implemented by adding a penalty for model complexity [9]. Constraint-based approaches often make the implicit assumption of so-called *causal faithfulness* [6], which states that there are no conditional independencies that hold in the density over a set of variables V , except those that are entailed by the causal structure. However, in practice faithfulness can be violated and better constrained-based approaches have been developed to handle this, such as CPC [11] and ACPC [12].

A drawback of both score-based and constrained-based approaches, however, is the inherent instability in structure estimation. With finite samples small changes in the data can lead to completely different optimal structures. Outcomes of borderline independence tests can be incorrect and can lead to multiple errors when propagated by the discovery algorithm [6].

The present work introduces a new score-based causal discovery algorithm, called *stable specification search*, that is robust for finite samples based on advances in stability selection using subsampling and selection algorithms. Structure search is performed over structural equation models (SEM), which is the most widely used language for causal discovery in various scientific disciplines. The method uses exploratory search, but allows incorporation of prior background knowledge. In order to show that our method can handle various kinds of data (continuous, discrete, and a combination of both) we evaluated our method on simulated and real-world data. The simulated data is used to compare our method with some advanced constrained-based approaches (PC-stable [13], CPC) and a score-based approach (GES). Specifically, we compare the robustness of each method in computing causal structure. The real-world data sets, about chronic fatigue syndrome and attention deficit hyperactivity disorder, are used to compare our results with some previous studies. The results show that our exploratory, hypothesis-free approach gives significant improvement over alternative approaches, and is able to obtain structure estimates that are consistent with the hypothesis driven models constructed by medical experts based on medical data and years of experience.

The rest of this paper is structured as follows. Section 2 describes all the background material obtained from the existing literature. Section 3 describes our robust score-based approach for causal discovery. Section 4 presents experimental result on one simulated and two real-world data sets. Section 5 gives conclusions and suggestions for future work.

2. Background

2.1. Directed acyclic graph

We first describe some graphical notation and terminology used in the remaining sections. A graph is a pair (V, E) with V a set of nodes and E a set of edges. A *directed graph* has all edges in E directed (arc); a single arrowhead on every edge, e.g., $A \rightarrow B$. Directed cycles represent feedback or reciprocal relationships, e.g., $A \rightarrow B \rightarrow A$. A graph with no directed cycles is called *acyclic*. A graph which is both directed and acyclic is called a *directed acyclic graph* (DAG) [2]. Fig. 1 depicts a DAG of four variables.

The *skeleton* of a DAG is the undirected graph that results from removing the directionality of every edge. A *v-structure* in a DAG G

Fig. 1. A DAG of four variables.

is an ordered triple (x, y, z) such that G contains the directed edges $x \rightarrow y$, $z \rightarrow y$, and x and z are not adjacent in G [14].

2.2. Causal modeling in SEM

In this study, we focus on causal models with no reciprocal or feedback relationships, and no latent variables. Generally, there are two common ways of representing a model in SEM: by stating all relations in the set as equations, which is called a *causal model*, or by drawing them as a *causal diagram* (graph). The general form of the equations is

$$x_i = f_i(\text{pa}_i, \varepsilon_i), \quad i = 1, \dots, n. \quad (1)$$

where pa_i denotes the *parents* which represent the set of variables considered to be direct causes of X_i and ε_i represents errors on account of omitted factors that are assumed to be mutually independent [2].

2.3. Specification search in SEM

Typically, a SEM is used as follows: (1) set a hypothesis as the prior model, (2) fit the model to the data, (3) evaluate the model, and (4) modify the model to improve the parsimony and score [15]. The last step is called *specification search* [16,17]. This typical model refinement approach is hypothesis-driven. It works by adding or deleting some arcs between variables from the initial model. Typically only a few models are evaluated, making it difficult to derive causal relationships.

An alternative approach is exploratory search in which no prior hypothesis is specified. Typical approaches in the literature for addressing the exponential search space include tabu search [18], genetic algorithms [19,20], ant colony optimization [21], and others [22–24].

2.4. Multi-objective optimization

Following the principle of Occam's razor, we should prefer models that are simple and fit the data well. These two objectives, however, are often conflicting as a well-fit model is likely to be a complex model. In this paper, we propose to make use of multi-objective optimization to explicitly optimize both objectives.

In multi-objective optimization, optimal solutions are defined in terms of *domination*. A model $\mathbf{x}_1$ is said to dominate model $\mathbf{x}_2$, if the following conditions are satisfied [25]:

$$\mathbf{x}_1 \leq \mathbf{x}_2 \quad \text{iff} \quad \begin{cases} \forall i \in \{1, \dots, M\} & f_i(\mathbf{x}_1) \leq f_i(\mathbf{x}_2) \\ \exists j \in \{1, \dots, M\} & f_j(\mathbf{x}_1) < f_j(\mathbf{x}_2) \end{cases} \quad (2)$$

The first condition states that the model $\mathbf{x}_1$ is no worse than $\mathbf{x}_2$ in all objectives f_i . The second condition states that the model $\mathbf{x}_1$ is strictly better than $\mathbf{x}_2$ in at least one objective. By using this concept, given the population of models P , we can partition P into n sets called *fronts* $F_1, \dots, F_n$, such that F_k dominates F_l where $1 \leq k < l \leq n$ and the models within the same front do not dominate each other. The

Fig. 2. Example of a population P partitioned into fronts $F_1, \dots, F_n$ when minimizing objectives f_1 and f_2 . F_1 is the Pareto front not dominated by any member of P .

so-called *Pareto Front* or *non-dominated set* F_1 includes models that are not dominated by any member of P . Essentially, using multi-objective optimization we efficiently find the best fitting models over a whole range of model complexities using a single coherent optimization approach. Fig. 2 provides a sketch.

2.4.1. NSGA-II

Non-dominated sorting genetic algorithm II or NSGA-II [26] is a well-known multi-objective evolutionary algorithm (MOEA), still widely applied in various fields, such as image retrieval [27], reactive power planning [28], building design [29], and robot grippers [30]. A characteristic feature is *fast non-dominated sorting* which sorts models based on the concept of domination. With M the number of objectives and N the size of population, the time complexity has order $\mathcal{O}(MN^2)$, which is better than a naïve approach with $\mathcal{O}(MN^3)$. Another characteristic feature is *crowding distance sorting* which is implemented to preserve the diversity among the solutions in the Pareto front. This feature sorts models based on the distance metric which explains the proximity of a model to other models.

The iterative procedure of NSGA-II shown in Fig. 3 is a sequence of steps started by generating a population of solutions P of size N . P is then manipulated by genetic operators such as selection, crossover, and mutation, forming a new population Q of size N . P and Q are then combined into population R with size $2N$. After that R is sorted using fast non-dominated sorting, yielding a set of fronts F . In the next iteration each front in F is sorted using the crowding distance sorting and the first N members are used to generate a new population P . At $t=0$, P is formed by creating N random solutions sorted with fast non-dominated sorting.

2.5. Stability selection

Structure estimation is a notoriously difficult problem, both because of computational aspects (finding the optimal structure can be NP hard) and because of instability (small changes in the data can lead to completely different optimal structures). In this section we describe the method of [31] for robust estimation of model structure based on subsampling in combination with selection algorithms. The method has been shown to yield finite sample family wise error control and improved structure estimates.

Let β be a sparse p -dimensional vector which generally represents, for example, the coefficient vector in linear regression or the edges in a graph. In structure estimation the goal is to infer the set $S = \{k : \beta_k \neq 0\}$ of non-zero components from noisy observations. Many methods tackle this problem by minimizing some loss

function augmented with a regularization term to avoid overfitting. Usually the regularization term is parameterized by $\lambda \in \Lambda \subseteq \mathbb{R}^+$ and each λ leads to an estimated structure $\hat{S}^\lambda \subseteq \{1, \dots, p\}$. The objective is to determine λ such that $\hat{S}^\lambda$ is identical to S with high probability. To this end, [31] introduces the concepts of *selection probabilities* and *stability paths*.

Definition 1 (Selection probabilities). Let I be a subset of $\{1, \dots, n\}$ of size $\lfloor n/2 \rfloor$ randomly drawn without replacement, $K \subseteq \{1, \dots, p\}$, and $\hat{S}^\lambda(I)$ be the selected set $\hat{S}^\lambda$ for subsample I . The probability of K being in set $\hat{S}^\lambda(I)$ is

$$\hat{\Pi}_K^\lambda = P(K \subseteq \hat{S}^\lambda(I))$$

where the probability being is with respect to the random subsampling and possibly the construction of $\hat{S}^\lambda(I)$.

Definition 2 (Stability path). For each variable $k = 1, \dots, p$ the *stability path* is given by the selection probabilities $\{\hat{\Pi}_k^\lambda : \lambda \in \Lambda\}$.

Furthermore, in stability selection we do not select a single element from the set of models $\{\hat{S}^\lambda : \lambda \in \Lambda\}$ as traditional methods do, but perturb the data many times and select structures that occur in a large fraction of selected sets. To this end, [31] introduces the concept of *stable variables*.

Definition 3 (Stable variables). The set of stable variables is defined as

$$\hat{S}^{\text{stable}} = \{k : \max_{\lambda \in \Lambda} \hat{\Pi}_k^\lambda \geq \pi_{\text{thr}}\}$$

where π_{thr} is a cutoff with $0 < \pi_{\text{thr}} < 1$.

Variables with a high selection probability are kept whereas those with low selection probabilities are disregarded. The threshold π_{thr} is a tuning parameter but its influence is small and sensible values (e.g., $\pi_{\text{thr}} \in (0.6, 0.9)$) tend to give similar results.

2.6. Model equivalence

There is one further subtlety that makes our approach for finding stable models (or sub-models) slightly more complicated than that in [31]. If we find a particular model, we have to account for the fact that there may be different models that are observationally indistinguishable. Causal models represented by DAGs have their corresponding model equivalent classes, called *completed partially directed acyclic graph* (CPDAG). This means that every probability distribution derived from a model in a particular CPDAG, can also be derived by models belonging to the same CPDAG. In SEMs, these models are called covariance equivalent [2].

The characterization of equivalent structures is given by the following theorem [32].

Theorem 1. (Verma and Pearl, 1990) Two DAGs are equivalent if and only if they have the same skeletons and the same v-structures.

Furthermore, a directed edge $x \rightarrow y$ is compelled in $\mathcal{G}$ if for every DAG $\mathcal{G}'$ equivalent to $\mathcal{G}$, $x \rightarrow y$ exists in $\mathcal{G}'$. For any edge e in $\mathcal{G}$, if e is not compelled in $\mathcal{G}$, then e is reversible. A CPDAG can be represented by a directed edge (arc) for every compelled edge and an undirected edge for every reversible edge [14].

Converting a model into a CPDAG allows one to observe the relations that hold among the variables. Arcs in a CPDAG indicate a cause–effect relation among variables since the same arc occurs in all members of the CPDAG. Undirected edges $A - B$ in a CPDAG indicate that some members of the CPDAG contain an arc $A \rightarrow B$ whereas other members contain an arc $B \rightarrow A$.

Fig. 3. Adopted from [26]. P is the current population with size N and is manipulated to make a new population Q . Both are combined, forming R , which will be sorted using fast non-dominated yielding a set of fronts F . Every member of front $F_n \in F$ will be assigned a so-called crowding distance in order to sort F_k . The first N members of F will be selected to be the next population P .

Fig. 4. The proposed method consists of two phases: *search* and *visualization*. The search phase is an iterative process using an *outer loop* and *inner loop* that combines SEM, NSGA-II, and stability selection, which outputs all relevant edges and causal paths between two variables. The visualization phase displays the relevant relationships as a causal model.

3. Proposed method

3.1. The general idea

Our proposed method can be divided into two phases. The first phase is *search* and the second phase is *visualization*. In the search phase SEM and NSGA-II are synergically combined for exploratory search of the model space. As portrayed in Fig. 4, the *inner loop* is an iterative process, searching over the model space and returns a Pareto front of models. The *outer loop* is an iterative process that samples a different subset of the data in each iteration and at the end returns a number of Pareto fronts coming from those subsets.

Each model returned by the outer loop is transformed into a CPDAG which are then used to compute the *edge stability graph* and the *causal path stability graph*.

Definition 4. (Stability graphs) Let A and B be two variables and G a multiset (or bag) of CPDAGS. Let G_c be the submultiset of G containing all CPDAGS with complexity c . The edge stability for A and B at complexity c is the number of models in G_c for which there exists an edge between A and B (i.e., $A \rightarrow B$, $B \rightarrow A$, or $A - B$) divided by the total number of models in G_c . The causal path stability for A

to B at complexity c is the number of models in G_c for which there is a directed path from A to B (of any length) divided by the total number of models in G_c . The terms edge stability graph and causal path stability graph are used to denote the corresponding measures for all variable pairs and all complexity levels.

On top of the stability graphs we perform stability selection. In [31], stability selection is defined in terms of a regularization parameter λ . In our approach we do not have a regularization parameter and instead use model complexity (defined in Section 4.2) which is one of the objectives in our multi-objective optimization approach. We therefore define two thresholds. The first threshold is the boundary of selection probability π_{sel} and corresponds to π_{thr} in [31]. For example, setting $\pi_{\text{sel}} = 0.6$ (as described in [31]) means that all causal relationships with edge stability or causal path stability (Fig. 5) above this threshold are considered *stable*. The second threshold is the boundary of complexity π_{bic} , which is used to control overfitting and corresponds to minimal λ in [31]. We set π_{bic} to the level of model complexity at which the minimum average Bayesian Information Criterion (BIC) score is found. For example, $\pi_{\text{bic}} = 8$ means that all causal relationships with an edge stability or a causal path stability lower than this threshold (Fig. 5) are considered *parsimonious*. Causal relationships that intersect with the top-left region are considered both stable and parsimonious and called *relevant*.

In the visualization phase we combine the stability graphs into a graph with nodes and edges. This is done by adding the relevant edges and orienting them using background knowledge (if any, see Section 3.2) and the relevant causal paths. In addition we annotate each edge with the highest selection probability it has across different model complexities in the top-left region of the edge stability. This visualization eases interpretation but the stability graphs are considered to be the main outcome of our approach.

3.2. Constrained SEM

In practise, one often has prior knowledge about the domain, for example, that A does not cause B directly, denoted by $A \nrightarrow B$. The method proposed here can include such prior knowledge,

Fig. 5. Example stability graphs from an artificial data set of 400 instances with six continuous variables, without prior knowledge. (a) Edge stability graph. (b) Causal path stability graph. Each line in (a) represents an edge between a pair of variables and each line in (b) represents a causal path with any length from a variable to another variable. The threshold of selection probability, π_{sel} , is set to 0.6 and the threshold for model complexity, π_{bic} , is chosen to minimize the average BIC score. See the main text for more details.

extending previous work [33], since this translates to a DAG with no directed edge from A to B .²

Model specifications should comply with any prior knowledge when performing specification search and when measuring the edge and causal path stability. When DAGs are converted into CPDAGs in the outer loop, a constraint $A \rightarrow B$ may be violated since arcs $B \rightarrow A$ in the DAG may be converted into undirected (reversible) edges $A - B$ in the CPDAG. In order to preserve constraints we therefore extended the efficient DAG-TO-CPDAG algorithm of [14] which runs in time $\mathcal{O}(|E|)$ given a DAG $G = (V, E)$.

Fig. 6 provides pseudocode for the constrained DAG to CPDAG algorithm. Line 2 produces a total ordering E' over the edges in DAG G . Lines 3–6 impose an arc upon the edges that match the constraints. Finally, Line 7 uses [14] to label the remaining edges $E \setminus E'$ in G with “compelled” or “reversible” and returns the constrained CPDAG G' .

A DAG without edges will always be transformed into a CPDAG without edges. A fully connected DAG without constraints will be transformed into a CPDAG with only undirected edges. However, if background knowledge is added, a fully connected DAG

will be transformed into a CPDAG in which the edges corresponding to the background knowledge are directed. From these observations it follows that in the edge stability graph all paths start with a selection probability of 0 and end up in a selection probability of 1. In the causal path stability graph when no prior knowledge has been added all paths start with a selection probability of 0 and end up in a selection probability of 0. However, when prior knowledge is added some of the paths may end up in a selection probability of 1 because of the added constraints.

3.3. Stable specification search algorithm

Fig. 7 provides pseudocode for our approach (cf. Fig. 4). Lines 3–18 represent the outer loop, Lines 6–16 represent the inner loops, Lines 19–21 compute stability graphs.

An inner loop (Lines 6–16) starts by forming a population P of size N , initially at random, or else from a previous population using crowding distance sorting (Lines 7–12). Models are represented with a binary vector $\mathbf{y}$ with $y_i \in \{0, 1\}$ denoting the existence of some arc $A \rightarrow B$. Line 13 forms a new population Q by manipulating P using binary tournament selection, one-point crossover, and one-bit flip mutation, which are compatible with a binary representation. The selection scheme selects N times two models from P and places the best model (i.e., lowest front or else smallest crowding distance) in a mating pool M_{pool} . One-point crossover takes two models from M_{pool} and swaps the data after the crossover point (the middle). One-bit flip mutation flips each bit according to a predetermined rate. Line 14 combines P and Q and sorts them using fast non-dominated sorting. Line 15 updates the Pareto front in F_1 .

An outer loop (Lines 3–18) randomly samples a subset T from D with size $\lfloor |D|/2 \rfloor$ (Line 4), runs the inner loop I times to obtain a Pareto front (Lines 6–16), and stores it in H (Line 17). After J iterations, H contains J Pareto fronts.

Lines 19–21 convert the J Pareto fronts in H from DAGs into CPDAGs using the algorithm in Fig. 7 and then computes the edge and causal path stability graphs. The stability graphs are considered to be the main outcome of our approach, but can also be visualized as a graph with nodes and edges.

4. Experimental study

4.1. Implementation

We implemented the stable specification search as an R package named *stablespec*. The package is publicly available at the Comprehensive R Archive Network (CRAN),³ so it can be installed directly, e.g., from R console by typing `install.packages("stablespec")` or from RStudio by using feature to install package. We also included a package documentation as a brief tutorial of using the functions. All experiments were run on an Intel Xeon E7-4870 v2 Processor 2.3 GHz, 15 Core, 96 of 32 GB LRDIMM.

4.2. Parameter settings

For all experiments, we employed the same set of NSGA-II parameters and stability thresholds. We had 100 iterations in the outer loop, and in each iteration we drew a subsample with size $\lfloor |D|/2 \rfloor$. We did not do a comprehensive parameter tuning for NSGA-II, instead, we followed guidelines provided in [34]. The parameters

² This still allows for directed edges from B to A or indirect relations from A to B .

³ <https://cran.r-project.org/web/packages/stablespec/index.html>.

```

1: function consDag2Cpdag(DAG  $G$ , constraint  $\mathcal{C}$ )
2:    $E' \leftarrow \text{orderEdges}(G)$ 
3:   for every constraint  $c \in \mathcal{C}$  do
4:     get  $e \in E'$  that matches  $c$ 
5:     label  $e$  with “compelled” in the direction consistent with  $c$ 
6:   end for
7:   return  $G' \leftarrow \text{labelEdges}(G, E')$     ▷ label remaining edges using [14]
8: end function

```

Fig. 6. The constrained DAG-TO-CPDAG algorithm returns a CPDAG which is consistent with the added prior knowledge and extends [14]. The algorithm first labels the edges that match the constraints with “compelled” and then labels the remaining edges with “reversible” or “compelled” using [14].

```

1: procedure stableSpecificationSearch(data set  $D$ , constraint  $\mathcal{C}$ )
2:    $H \leftarrow ()$     ▷ initialize
3:   for  $j \leftarrow 0, \dots, J-1$  do    ▷  $J$  is number of outer loop iterations
4:      $T \leftarrow$  subset of  $D$  with size  $\lfloor |D|/2 \rfloor$  without replacement
5:      $F_1 \leftarrow ()$     ▷ initialize Pareto fronts to empty list
6:     for  $i \leftarrow 0, \dots, I-1$  do    ▷  $I$  is number of inner loop iterations
7:       if  $i = 0$  then
8:          $P \leftarrow N$  random DAGs consistent with  $\mathcal{C}$ 
9:          $P \leftarrow \text{fastNonDominatedSort}(P)$ 
10:      else
11:         $P \leftarrow \text{crowdingDistanceSort}(F)$     ▷ draw the first  $N$  models
12:      end if
13:       $Q \leftarrow$  make population from  $P$ 
14:       $F \leftarrow \text{fastNonDominatedSort}(P \cup Q)$ 
15:       $F_1 \leftarrow$  pareto front of  $F$  and  $F_1$ 
16:    end for
17:     $H \leftarrow H \cup F_1$     ▷ concatenation
18:  end for
19:   $G \leftarrow \text{consDag2Cpdag}(H, \mathcal{C})$ 
20:  edges  $\leftarrow$  edge stability of  $G$ 
21:  paths  $\leftarrow$  path stability of  $G$ 
22: end procedure

```

Fig. 7. Stable specification search consists of an outer and an inner loop. The outer loop samples a subset of the data, and for every subset, the inner loop searches for the Pareto front by applying NSGA-II. The Pareto fronts are converted into constrained CPDAGs which are then used to compute the edge and causal path stability graph.

were set as follows: the number of generations (inner loop) was 20, the size of the population P was 100, the crossover rate was 0.85, the mutation rate was 0.075 and with binary tournament selection.

We score models using the *chi-square* χ^2 and the *model complexity*. The χ^2 is considered the original fit index in SEM and measures whether the model-implied covariance matrix is close enough to the sample covariance matrix [35].

The model complexity represents how many predicted parameters the model contains. Assuming that variances of parameters are always predicted, the maximum model complexity with n variables is given by $n(n-1)/2$.

When using multi-objective optimization we minimize both the χ^2 and model complexity objectives. These two objectives are, however, conflicting with each other. For example, minimizing the model complexity typically means compromising the data fit.

Fig. 8. The Waste Incinerator network. Rectangular nodes represent discrete variables, oval nodes represent continuous variables, and arcs represent direct causal relations.

4.3. Application to simulated data

4.3.1. Data generation

In this experiment we generated data using the Waste Incinerator network in Fig. 8, which is a model of waste emissions from an incinerator plant [36]. This model contains both discrete and continuous random variables, with B the waste burning regimen, W the compositional differences in incoming waste, C the concentration of CO_2 , F the filter state, E the filter efficiency, L the light penetrability, D the emission of dust, M_{in} the metals in waste, and M_{out} the metals emission.

We generated 10 data sets containing 400 samples from this network using the BNT toolbox with the default parameter setting as described in [37].

4.3.2. Performance measure

We compared the stable specification search with GES (score-based method), PC-stable, and CPC (both constrained-based methods). Our method intrinsically subdivides the data in a number of subsets, here 50 of size 200 samples, and then runs the multi-objective optimization to obtain 50 Pareto fronts (see Section 2.4). For a fair comparison, for each algorithm we consider subsampling (e.g., [38]), giving each method 50 subsets. For every subset, each algorithm returns a CPDAG from which we can derive the edges and causal paths.

Since the true model of the Waste Incinerator data is known, we can measure the performance of both methods by means of the receiver operating characteristic (ROC) curve [39]. The true positive rate (TPR) and the false positive rate (FPR) are computed with respect to the CPDAG of the true model. For example, in the case of causal path stability, a true positive means that a causal path with any length obtained through our approach or the other algorithms is actually present in the CPDAG of the true model. By increasing the threshold π_{sel} , we increase the TPR at the expense of the FPR. In addition, we conducted three significance tests to compare the ROC curves. The first test [40] compares the area under the curve (AUC) of the ROC curves based on the theory of U -statistics. The second test [41], a modification of [42], compares the AUC of ROC curves that are generated from bootstrap replicates. The third test [43] compares the actual ROC curves by evaluating the absolute difference. The null hypothesis is that the AUC of the ROC curves of our method and the other algorithm are equivalent.

We repeated the above procedure 10 times on different Waste Incinerator data sets and computed the ROC curves using two different schemes: averaging and individual. In the averaging scheme, the ROC curves are computed based on the average edge and causal path stability from different data sets. We conducted statistical significance tests on these average ROC curves. Conversely, in the individual scheme the ROC curves are computed directly from the edge and causal path stability on each data set. We conducted individual statistical significance tests on the ROC curves for each

(a)

(b)

Fig. 9. ROC curves for (a) the edge stability and (b) the causal path stability, for different values of π_{sel} in the range of [0, 1]. In (a), the AUCs are 0.96 (stable specification search), 0.89 (PC-stable), 0.88 (CPC), and 0.69 (GES). In (b), The AUCs are 0.98 (stable specification search), 0.85 (PC-stable), 0.88 (CPC), and 0.61 (GES).

data set and then used Fisher's method, as described in [44,45], to combine these tests into a single test statistic. Both schemes are intended to show empirically and comprehensively how robust the results of each algorithm are across changes in the data.

4.3.3. Discussion of Waste Incinerator result

Fig. 9 shows the ROC curves for (a) the edge stability and (b) the causal path stability from the averaging scheme. The corresponding AUCs for edge stability are 0.96 (stable specification search), 0.89 (PC-stable), 0.88 (CPC), and 0.69 (GES). The AUCs for causal path stability are 0.98 (stable specification search), 0.85 (PC-stable), 0.88 (CPC), and 0.61 (GES).

Table 1 lists the results of the significance tests for both the averaging and individual schemes. The ROC and AUC for the edge stability are comparable with PC-stable and CPC (p -value > 0.1), but always significant (p -value < 0.01) compared with GES. The ROC and AUC for the causal path stability compared with CPC are marginally significant (p -value < 0.1) using the averaging scheme, but significant using the individual scheme (p -value < 0.01); compared with PC-stable significant (p -value < 0.05) using the averaging scheme, but highly significant using the individual scheme (p -value $< 10^{-5}$); compared with GES highly significant using both schemes (p -value $< 10^{-5}$). To conclude, we show that the stable specification search

Table 1

Table of p -values from comparisons between stable specification search and alternative approaches. For each significance test, we compared the ROC of the edge (Edge) and causal path (Causal) stability on both averaging (Ave.) and individual (Ind.) schemes.

Significance test		GES		PC-stable		CPC	
		Ave.	Ind.	Ave.	Ind.	Ave.	Ind.
DeLong [40]	Edge	0.003	$<10^{-5}$	0.317	0.175	0.284	0.131
	Causal	$<10^{-5}$	$<10^{-5}$	0.027	$<10^{-5}$	0.073	$<10^{-5}$
Bootstrap [41]	Edge	0.003	$<10^{-5}$	0.296	0.135	0.261	0.098
	Causal	$<10^{-5}$	$<10^{-5}$	0.022	$<10^{-5}$	0.064	$<10^{-5}$
Venkatraman [43]	Edge	0.004	$<10^{-5}$	0.591	0.684	0.539	0.592
	Causal	$<10^{-5}$	$<10^{-5}$	0.023	$<10^{-5}$	0.096	0.005

obtains at least comparable performance as, but often significant improvement over alternative approaches, especially in obtaining the causal relations.

4.4. Application to real-world data

This section describes the results of applying our proposed method on two real-world data sets. Both of them are about particular diseases, for which the underlying causal relationships are often not clear. Revealing such causal relationships can lead to the development of (new) dedicated treatments and medications. Here, we consider data on *attention deficit hyperactivity disorder* (ADHD) and *chronic fatigue syndrome* (CFS).

4.4.1. Performance measure

Since the true model is unknown we measure the performance of our method using the edge stability and causal path stability graphs. We set the thresholds to $\pi_{\text{sel}} = 0.6$ and π_{bic} to the minimum average of BIC scores. The relevant causal relations are those which occur in the top-left region (see Fig. 5 as example). We compare the stability graphs to studies reported in the literature.

4.4.2. Application to CFS

In this experiment we consider a data set about *chronic fatigue syndrome* (CFS) of 183 subjects [46]. Originally the data comes from a longitudinal study with five time slices, but in this paper, we focus only on one time slice representing the subjects after the first treatment.

The data set contains six discrete variables; fatigue severity assessed with the subscale fatigue severity of the checklist individual strength (CIS), the sense of control over fatigue assessed with the *self-efficacy scale* (SES), focusing on symptoms measured with the *illness management questionnaire*, the objective activity of the patient measured using an *actometer* (oActivity), the subject's perceived activity measured with the subscale activity of the CIS (pActivity), and physical functioning measured with subscale physical functioning of the *medical outcomes survey* (SF36). We refer to the original paper [46], for a detailed description of the questionnaires used and the actometer. Missing values were imputed using an imputation method *expectation maximization* implemented in SPSS [47]. As all of the variables have large scales, e.g., in the range between 0 and 155, we treat them as continuous variables. We added prior knowledge that the variable fatigue does not cause any of the other variables directly.

The total computation time for one subset was around 5.5 min. Fig. 10 shows that eight relevant edges were found. These edges are between pActivity and fatigue, focusing and fatigue, functioning and fatigue, control and fatigue, pActivity and focusing, pActivity and oActivity, focusing and control, and functioning and control.

Fig. 10. The stability graphs for CFS together with π_{sel} and π_{bic} , yielding four regions. The top-left region is the area containing the relevant causal relations. (a) The edge stability graph showing eight relevant edges. (b) The causal path stability graph showing four relevant causal paths. See Tables A.1 and A.2 in Appendix A for more detail.

Fig. 10 shows that four relevant causal paths were found. These causal paths are: pActivity to fatigue, control to fatigue, functioning to fatigue, and focusing to fatigue.

The stability graphs can be combined into a model as follows. First, the nodes are connected according to the eight relevant edges obtained. Second, the edges are oriented according to the background knowledge added. The fact that the variable fatigue does not directly cause any other variable results in four directed edges, which, in this case, correspond exactly to the relevant causal paths obtained. The inferred model is shown in Fig. 11.

A (direct) causal path $X \rightarrow Y$ in Fig. 11 indicates that a change in variable X causes a change in variable Y . All variables except for objective activity were found to be direct causes for fatigue severity, which are corroborated by literature studies. In [48], changes in physical activity, sense of control, and focus on symptoms measured, were shown to result in changes in fatigue. In [49], changes in perceived activity, sense of control, and physical functioning were shown to result in changes in fatigue. In [46], an increase in sense of control, perceived activity, and self-reported physical functioning, as well as a decrease in focusing on symptoms resulted in a decrease of fatigue, whereas changes in objective activity did not result in any change in fatigue.

Fig. 11. The inferred model of CFS by combining the edge stability and causal path stability graphs. Each edge has a reliability score which is the highest selection probability in the top-left region of the edge stability graph.

Fig. 12. The stability graphs for ADHD together with π_{sel} and π_{bic} , yielding four regions. The top-left region is the area containing the relevant causal relations. (a) The edge stability graph showing four relevant edges. (b) The causal path stability graph showing seven relevant causal paths. See Tables A.3 and A.4 in Appendix A for more detail.

4.4.3. Application to ADHD

In this experiment we consider a data set about *attention deficit hyperactivity disorder* (ADHD) of 245 subjects with 23 variables [50]. Following [51], we excluded instances with missing values and variables that either have insufficient instances or are considered irrelevant. The remaining data set consists of 221 instances and

Fig. 13. The inferred model of ADHD by combining the edge stability and causal path stability graphs. Each edge has a reliability score which is the highest selection probability in the top-left region of the edge stability graph.

six variables with gender the gender of subjects, AD the attention deficit measure, HI the assessment of hyperactivity/impulsivity symptoms, aggression the measure of aggressive behavior, medication the medication status of subjects, and handedness represents whether a subject uses the right and/or left hand. Following [52], we treat all discrete variables as continuous variables. We added prior knowledge that the variable gender does not cause any of the other variables directly.

The total computation time for one subset was around 4.9 min. Fig. 12 shows that there are four relevant edges, namely between gender and AD, AD and medication, AD and HI, and HI and aggression. Moreover, Fig. 12 shows that there are seven relevant causal paths; gender to AD, gender to HI, gender to medication, gender to aggression, AD to HI, AD to medication, and AD to aggression.

The stability graphs can be combined into a model as follows. First, the nodes are connected according to the four relevant edges obtained. Second, the edges are oriented according to the background knowledge added. The fact that the variable gender does not directly cause any other variable results in one directed edge $gender \rightarrow AD$. Third, the edges are oriented according to the relevant causal paths obtained. This results in two directed edges, $AD \rightarrow HI$ and $AD \rightarrow medication$. Since there is no relevant edge between AD and aggression and no relevant causal path from HI to aggression we cannot orient any other edges and therefore cannot represent two of the relevant causal paths in the model. We lose some information when converting the stability graphs into a model. The inferred model is shown in Fig. 13.

The causal relations obtained for ADHD are corroborated by studies reported in the literature. In [51], gender is shown to be a direct cause for attention deficit, attention deficit is shown to be a direct cause for both hyperactivity, medication, and aggression, and hyperactivity and aggression are related but neither variable is a direct cause for the other.

5. Conclusion and future work

In the last decades the field of causal modeling has seen a surge in theoretical development and the construction of various causal discovery algorithms. In general, causal discovery algorithms can be divided into two approaches: constraint-based and score-based. A disadvantage, however, of current causal discovery algorithms is the inherent instability of structure estimation. With finite samples small changes in the data can lead to completely different optimal structures.

The present work introduces a new hypothesis-free score-based causal discovery algorithm, stable specification search, that is robust for finite samples based on subsampling and selection algorithms. Our approach uses exploratory search to search over structural equation models and allows for the incorporation of prior background knowledge, without the need to specify the complete model structure in advance.

The comparison conducted on the simulated data shows that our method, the stable specification search, shows significant improvement over alternative approaches in obtaining the causal relations. The results on both real-world data sets, CFS and ADHD, are consistent with previous studies [48,49,46,51]. In general we may conclude that our causal discovery algorithm is able to robustly estimate the underlying causal structure.

Several issues have not yet been explored in our current approach that warrant further research, such as latent variables and longitudinal data. Taking into account the existence of latent variables can further improve our structure estimate by properly identifying dependencies between variables as an unmeasured

common cause acting on both variables. In longitudinal data several subjects are measured at different time slices which provides a richer structure that can be incorporated in the causal discovery algorithm. A first attempt in this direction can be found in [53].

Our approach can be viewed as a novel application of multi-objective optimization. The main idea of stability selection [31], is to increase the robustness of structure estimation by considering a whole range of model complexities. In the original work, this is done by varying a continuous regularization parameter. For causal discovery we have to explicitly consider different discrete model complexities. Furthermore, finding the optimal structure for each model complexity is a hard optimization problem. By

Table A.1
Edge stability of CFS.

Lines	Edges
□-----□-----□	fatigue and pActivity
○-----○-----○	fatigue and control
△-----△-----△	control and focusing
+-----+-----+	fatigue and functioning
×-----×-----×	control and functioning
◇-----◇-----◇	focusing and pActivity
□-----□-----□	fatigue and focusing
○-----○-----○	pActivity and oActivity
△-----△-----△	functioning and pActivity
+-----+-----+	control and pActivity
×-----×-----×	control and oActivity
◇-----◇-----◇	focusing and oActivity
□-----□-----□	fatigue and oActivity
○-----○-----○	functioning and oActivity
△-----△-----△	functioning and focusing

Table A.3
Edge stability of ADHD.

Lines	Edges
◇-----◇-----◇	AD and HI
□-----□-----□	AD and medication
×-----×-----×	HI and aggression
□-----□-----□	gender and AD
○-----○-----○	aggression and AD
+-----+-----+	AD and medication
△-----△-----△	handedness and aggression
□-----□-----□	medication and aggression
○-----○-----○	gender and HI
◇-----◇-----◇	HI and handedness
△-----△-----△	gender and medication
×-----×-----×	gender and handedness
△-----△-----△	AD and handedness
+-----+-----+	gender and aggression
○-----○-----○	medication and handedness

Table A.2
Causal path stability of CFS.

Lines	Causal Paths
□-----□-----□	pActivity to fatigue
○-----○-----○	control to fatigue
△-----△-----△	functioning to fatigue
+-----+-----+	focusing to fatigue
×-----×-----×	oActivity to fatigue
◇-----◇-----◇	focusing to pActivity
□-----□-----□	functioning to pActivity
○-----○-----○	oActivity to pActivity
△-----△-----△	control to pActivity
+-----+-----+	functioning to oActivity
×-----×-----×	focusing to oActivity
◇-----◇-----◇	focusing to control
□-----□-----□	control to oActivity
○-----○-----○	functioning to control
△-----△-----△	pActivity to oActivity
+-----+-----+	control to functioning
×-----×-----×	oActivity to functioning
◇-----◇-----◇	oActivity to control
□-----□-----□	oActivity to focusing
○-----○-----○	control to focusing
△-----△-----△	pActivity to functioning
+-----+-----+	focusing to functioning
×-----×-----×	functioning to focusing
◇-----◇-----◇	pActivity to control
□-----□-----□	pActivity to focusing
○-----○-----○	fatigue to pActivity
△-----△-----△	fatigue to oActivity
+-----+-----+	fatigue to focusing
×-----×-----×	fatigue to functioning
◇-----◇-----◇	fatigue to control

Table A.4
Causal path stability of ADHD.

Lines	Causal Paths
◇-----◇-----◇	gender to AD
×-----×-----×	gender to HI
+-----+-----+	gender to medication
×-----×-----×	AD to medication
◇-----◇-----◇	AD to HI
+-----+-----+	AD to aggression
△-----△-----△	gender to aggression
×-----×-----×	HI to aggression
◇-----◇-----◇	HI to medication
□-----□-----□	aggression to medication
○-----○-----○	aggression to HI
□-----□-----□	HI to AD
◇-----◇-----◇	medication to aggression
□-----□-----□	handedness to aggression
○-----○-----○	gender to handedness
◇-----◇-----◇	aggression to handedness
△-----△-----△	AD to handedness
+-----+-----+	HI to handedness
△-----△-----△	handedness to HI
○-----○-----○	handedness to medication
□-----□-----□	medication to HI
△-----△-----△	aggression to gender
×-----×-----×	medication to handedness
+-----+-----+	aggression to HI
○-----○-----○	medication to AD
□-----□-----□	AD to gender
○-----○-----○	HI to gender
△-----△-----△	medication to gender
+-----+-----+	aggression to gender
×-----×-----×	handedness to gender

rephrasing stability selection as a multi-objective optimization problem, we can jointly run over various model complexities and find the corresponding optimal structures for each model complexity. In this paper, we have used NSGA-II for multi-objective optimization, because of its popularity and availability, but realize that more recent multi-objective optimization approaches [54–57] may be even more efficient. This is beyond the scope of this work and left for future research. In the same spirit, one can easily combine freely available software packages, e.g., for scoring structural equation models, bootstrap sampling, and multi-objective optimization, to build one's own robust structural estimation approach.

Acknowledgement

The research leading to these results has received funding from the DGHE of Indonesia and the European Community's Seventh Framework Programme (FP7/2007–2013) under grant agreement no. 305697.

Appendix A. Table of stability

See Tables A.1–A.4.

References

- [1] B. van Engelen, The OPTIMISTIC consortium, Cognitive behaviour therapy plus aerobic exercise training to increase activity in patients with myotonic dystrophy type 1 (DM1) compared to usual care (OPTIMISTIC): study protocol for randomised controlled trial, *Trials* 16 (2015) 224.
- [2] J. Pearl, *Causality: Models, Reasoning and Inference*, vol. 29, Cambridge University Press, 2000.
- [3] J. Pearl, T. Verma, *A Theory of Inferred Causation*, Morgan Kaufmann, San Mateo, CA, 1991.
- [4] P. Spirtes, C.N. Glymour, R. Scheines, *Causation, Prediction, and Search*, MIT Press, 2000.
- [5] J.-P. Pellet, A. Elisseeff, Using Markov blankets for causal structure learning, *J. Mach. Learn. Res.* 9 (2008) 1295–1342.
- [6] P. Spirtes, Introduction to causal inference, *J. Mach. Learn. Res.* 11 (2010) 1643–1662.
- [7] A.P. Dawid, Present position and potential developments: some personal views: statistical theory: the prequential approach, *J. R. Stat. Soc. Ser. A Gen.* (1984) 278–292.
- [8] G. Schwarz, Estimating the dimension of a model, *Ann. Stat.* 6 (2) (1978) 461–464.
- [9] D.M. Chickering, Optimal structure identification with greedy search, *J. Mach. Learn. Res.* 3 (November) (2002) 507–554.
- [10] D. Heckerman, C. Meek, G. Cooper, A Bayesian approach to causal discovery, *Comput. Causation Discov.* 19 (1999) 141–166.
- [11] J. Ramsey, J. Zhang, P. Spirtes, Adjacency-faithfulness and conservative causal Inference, in: *Proceedings of the Twenty-second Annual Conference on Uncertainty in Artificial Intelligence*, 2006.
- [12] J. Lemeire, S. Meganck, F. Cartella, T. Liu, Conservative independence-based causal structure learning in absence of adjacency faithfulness, *Int. J. Approx. Reason.* 53 (9) (2012) 1305–1325.
- [13] D. Colombo, M.H. Maathuis, Order-independent constraint-based causal structure learning, *J. Mach. Learn. Res.* 15 (1) (2014) 3741–3782.
- [14] D.M. Chickering, Learning equivalence classes of Bayesian-network structures, *J. Mach. Learn. Res.* 2 (2002) 445–498.
- [15] R. MacCallum, Specification searches in covariance structure modeling, *Psychol. Bull.* 100 (1) (1986) 107–120.
- [16] E.E. Leamer, *Specification Searches: Ad Hoc Inference With Nonexperimental Data*, Wiley, New York, 1978.
- [17] J.S. Long, *Covariance Structure Models: An Introduction to LISREL*, Sage, 1983.
- [18] G.A. Marcoulides, Z. Drezner, R.E. Schumacker, Model specification searches in structural equation modeling using tabu search, *Struct. Equ. Model. Multidiscip. J.* 5 (4) (1998) 365–376.
- [19] G.A. Marcoulides, Z. Drezner, Specification searches in structural equation modeling with a genetic algorithm, *New Dev. Tech. Struct. Equ. Model.* (2001) 247–268.
- [20] H. Murohashi, H. Toyoda, Model specification search using a genetic algorithm with factor reordering for a simple structure factor analysis model, *Jpn. Psychol. Res.* 49 (3) (2007) 179–191.
- [21] G.A. Marcoulides, Z. Drezner, Model specification searches using ant colony optimization algorithms, *Struct. Equ. Model. Multidiscip. J.* 10 (1) (2003) 154–164.
- [22] J. Herting, H. Costner, Respecification in multiple indicator models, *Causal Models Soc. Sci.* (1985) 321–393.
- [23] P. Spirtes, R. Scheines, C. Glymour, Simulation studies of the reliability of computer-aided model specification using the TETRAD II, EQS, and LISREL programs, *Sociol. Methods Res.* 19 (1) (1990) 3–66.
- [24] W.E. Saris, A. Satorra, D. Sörbom, The detection and correction of specification errors in structural equation models, *Sociol. Methodol.* 17 (1987) 105–129.
- [25] K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms*, John Wiley & Sons, Chichester, 2001.
- [26] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Trans. Evol. Comput.* 6 (2) (2002) 182–197.
- [27] M. Arevalillo-Herráez, F.J. Ferri, S. Moreno-Picot, A hybrid multi-objective optimization algorithm for content based image retrieval, *Appl. Soft Comput.* 13 (11) (2013) 4358–4369.
- [28] F. Hajabdollahi, Z. Hajabdollahi, H. Hajabdollahi, Soft computing based multi-objective optimization of steam cycle power plant using NSGA-II and ANN, *Appl. Soft Comput.* 12 (11) (2012) 3648–3655.
- [29] A.E. Brownlee, J.A. Wright, Constrained, mixed-integer and multi-objective optimisation of building designs by NSGA-II with fitness approximation, *Appl. Soft Comput.* 33 (2015) 114–126.
- [30] R. Saravanan, S. Ramabalan, N.G.R. Ebenezer, C. Dharmaraja, Evolutionary multi criteria design optimization of robot grippers, *Appl. Soft Comput.* 9 (1) (2009) 159–172.
- [31] N. Meinshausen, P. Bühlmann, Stability selection, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 72 (4) (2010) 417–473.
- [32] T. Verma, J. Pearl, Equivalence and synthesis of causal models, in: *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*, Elsevier Science, Inc., 1990, pp. 255–270.
- [33] R. Rahmadi, P. Groot, T. Heskes, Stable specification searches in structural equation modeling using multi-objective evolutionary algorithm, in: *Proceedings of SNATI*, 2014.
- [34] J.J. Grefenstette, Optimization of control parameters for genetic algorithms, *IEEE Trans. Syst. Man Cybern.* 16 (1) (1986) 122–128.
- [35] R. Kline, *Principles and Practice of Structural Equation Modeling*, Methodology in the Social Sciences, Guilford Press, 2011.
- [36] S.L. Lauritzen, Propagation of probabilities, means, and variances in mixed graphical association models, *J. Am. Stat. Assoc.* 87 (420) (1992) 1098–1108.
- [37] K. Murphy, The Bayes net toolbox for Matlab, *Comput. Sci. Stat.* 33 (2) (2001) 1024–1034.
- [38] J. Ramsey, *Bootstrapping the PC and CPC Algorithms to Improve Search Accuracy*, 2010.
- [39] T. Fawcett, ROC graphs: notes and practical considerations for researchers, *Mach. Learn.* 31 (2004) 1–38.
- [40] E.R. DeLong, D.M. DeLong, D.L. Clarke-Pearson, Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach, *Biometrics* (1988) 837–845.
- [41] X. Robin, N. Turck, A. Hainard, N. Tiberti, F. Lisacek, J.-C. Sanchez, M. Müller, pROC: an open-source package for R and S+ to analyze and compare ROC curves, *BMC Bioinform.* 12 (1) (2011) 77.
- [42] J.A. Hanley, B.J. McNeil, A method of comparing the areas under receiver operating characteristic curves derived from the same cases, *Radiology* 148 (3) (1983) 839–843.
- [43] E. Venkatraman, C.B. Begg, A distribution-free procedure for comparing receiver operating characteristic curves from a paired experiment, *Biometrika* 83 (4) (1996) 835–848.
- [44] R.A. Fisher, *Statistical Methods for Research Workers*, Genesis Publishing Pvt Ltd, 1925.
- [45] R.A.F. Frederick Mosteller, Questions and answers, *Am. Stat.* 2 (5) (1948) 30–31 <http://www.jstor.org/stable/2681650>.
- [46] M.J. Heins, H. Knoop, W.J. Burk, G. Bleijenberg, The process of cognitive behaviour therapy for chronic fatigue syndrome: which changes in perpetuating cognitions and behaviour are related to a reduction in fatigue? *J. Psychosom. Res.* 75 (3) (2013) 235–241.
- [47] IBM SPSS Statistics for Windows, Version 19.0, IBM Corp, Armonk, NY, 2010.
- [48] J. Vercoulen, C. Swanink, J. Galama, J. Fennis, P. Jongen, O. Hommes, J. Van der Meer, G. Bleijenberg, The persistence of fatigue in chronic fatigue syndrome and multiple sclerosis: development of a model, *J. Psychosom. Res.* 45 (6) (1998) 507–517.
- [49] J.F. Wiborg, H. Knoop, L.E. Frank, G. Bleijenberg, Towards an evidence-based treatment model for cognitive behavioral interventions focusing on chronic fatigue syndrome, *J. Psychosom. Res.* 72 (5) (2012) 399–404.
- [50] Q. Cao, Y. Zang, L. Sun, M. Sui, X. Long, Q. Zou, Y. Wang, Abnormal neural activity in children with attention deficit hyperactivity disorder: a resting-state functional magnetic resonance imaging study, *Neuroreport* 17 (10) (2006) 1033–1036.
- [51] E. Sokolova, P. Groot, T. Claassen, T. Heskes, Causal discovery from databases with discrete and continuous variables, in: *Probabilistic Graphical Models*, Springer, 2014, pp. 442–457.
- [52] M. Rhemtulla, P.E. Brosseau-Liard, V. Savalei, When can categorical variables be treated as continuous? A comparison of robust continuous and categorical SEM estimation methods under suboptimal conditions, *Psychol. Methods* 17 (3) (2012) 354–373.

- [53] R. Rahmadi, P. Groot, M. Heins, H. Knoop, T. Heskes, Causality on longitudinal data: stable specification search in constrained structural equation modeling, *Proceedings of AALTD 2015* (2015) 101.
- [54] Y. Qi, Z. Hou, M. Yin, H. Sun, J. Huang, An immune multi-objective optimization algorithm with differential evolution inspired recombination, *Appl. Soft Comput.* 29 (2015) 395–410.
- [55] S. Kukkonen, J. Lampinen, GDE3: the third evolution step of generalized differential evolution, in: *The 2005 IEEE Congress on Evolutionary Computation*, vol. 1, IEEE, 2005, pp. 443–450.
- [56] Q. Zhang, H. Li, MOEA/D: a multiobjective evolutionary algorithm based on decomposition, *IEEE Trans. Evol. Comput.* 11 (6) (2007) 712–731.
- [57] H. Taboada, J.F. Espiritu, D.W. Coit, MOMS-GA: a multi-objective multi-state genetic algorithm for system reliability optimization design problems, *IEEE Trans. Reliab.* 57 (1) (2008) 182–191.

Dr. Marianne Heins received a PhD degree from Radboud University Nijmegen Medical Centre, The Netherlands, in 2011. Currently she is a researcher at the Netherlands Institute for Health Service Research (NIVEL). Her main focus is on primary health care use in cancer survivors, using GP registry data and a patient panel. She is author of more than 20 publications in several scopes such as chronic fatigue syndrome and cancer.

Dr. Hans Knoop is a clinical psychologist and associate professor, and the head of the Expert Centre for Chronic Fatigue (ECCF) at the Radboud University Nijmegen Medical Centre, located in The Netherlands. His research focuses on possible determinants of chronic fatigue, development of cognitive behavioural intervention, and factors that contribute to outcomes. Besides author of almost 100 publications, he is also involved in several research projects.

Dr. Tom Heskes is a professor in artificial intelligence, and he leads the Data Science Group at the Institute for Computing and Information Sciences, Radboud University Nijmegen, The Netherlands. He is further affiliated Principal Investigator at the Donders Institute for Brain, Cognition and Behaviour. Prof. Heskes' research is on Bayesian inference, machine learning, and neural networks. He is involved in several projects that concern applications in, among others, neuroimaging, genomics, and bioinformatics. Prof. Heskes has published over 150 peer-reviewed research papers in the above areas. Prof. Heskes was the Editor-in-Chief of *Neurocomputing* between 2003 and 2015 and is currently associate editor of five other journals.

Ridho Rahmadi received master degrees in computer science (artificial intelligence and metaheuristics) in 2012 from Czech Technical University in Prague and Johannes Kepler University Linz, Austria. He is currently doing his PhD at the Data Science Group, Radboud University Nijmegen, and supervised by Tom Heskes and Perry Groot. His research interests are in the scope of causal discovery, machine learning, evolutionary algorithm, and submodularity.

Dr. Perry Groot received master degrees in computer science (artificial intelligence) and mathematics (topology) in 1998 and a PhD degree from the Department of Artificial Intelligence, Vrije University Amsterdam in 2003. Currently, he is a researcher at the Institute for Computing and Information Sciences at the Radboud University Nijmegen, The Netherlands. His research interests include Bayesian reasoning applications, causal modeling, particularly the use of Gaussian processes for preference learning and function optimization.