

cna: An R Package for Configurational Causal Inference and Modeling

Michael Baumgartner
University of Bergen, Norway

Mathias Ambuehl
Consult AG, Switzerland

Abstract

The R package **cna** provides comprehensive functionalities for causal inference and modeling with *Coincidence Analysis* (CNA), which is a configurational comparative method of causal data analysis. In this vignette, we first review the theoretical and methodological background of CNA. Second, we introduce the data types processable by CNA, the package’s core analytical functions with their arguments, and some auxiliary functions for data simulations. Third, CNA’s output along with relevant fit parameters and output attributes are discussed. Fourth, we provide guidance on how to interpret that output and, in particular, on how to proceed in case of model ambiguities. Finally, some considerations are offered on benchmarking the reliability of CNA.

Keywords: configurational comparative methods, set-theoretic methods, Coincidence Analysis, Qualitative Comparative Analysis, INUS causation, Boolean causation.

1. Introduction

Coincidence Analysis (CNA) is a configurational comparative method of causal data analysis that was introduced for crisp-set (i.e. binary) data in (Baumgartner 2009a; 2009b; 2015) and substantively extended, reworked, and generalized for multi-value and fuzzy-set data in (Baumgartner and Ambuehl 2020). The **cna** package reflects and implements the method’s latest stage of development.

CNA infers causal structures, as defined by the so-called *INUS theory*, from empirical data. The INUS theory was first developed by Mackie (1974) and later refined to the *MINUS theory* (Graßhoff and May 2001; Beirlaen *et al.* 2018; Baumgartner and Falk 2019). (M)INUS structures have two characteristic features that are frequently encountered in many disciplines: *conjunctivity*—complex causes only become operative when all of their components are co-instantiated, each of which, in isolation, is ineffective—and *disjunctivity*—effects can be brought about along alternative causal routes such that, when suppressing one route, the effect may still be produced via another one.¹ A viral disease is only contracted if exposure to the virus occurs while it is active and the host’s immune system has not yet acquired immunity (e.g. Rothman 1976). A variation in a phenotype only occurs if multiple single-nucleotide polymorphisms (SNPs) interact, and various such SNP interactions can independently of one another induce the same phenotype (e.g. Culverhouse *et al.* 2002). Or, a country only be-

¹Various other labels are used for these features in different disciplines: “component causation”, “conjunctural causation”, “alternative causation”, “equifinality”, etc.

comes a liberal democracy if it has a history of medieval constitutionalism combined with the absence of military revolution (in the early modern period), which, in turn, results from either a geography that deters invading armies, commercial wealth, foreign resource mobilization, or foreign alliances (Downing 1992).

Causal structures featuring conjunctivity or disjunctivity often do not exhibit additive or, more generally, linear dependencies between pairs of exogenous and endogenous factors, which gives rise to various problems when it comes to discovering them in data. Most importantly, there may not be any dependencies between an individual component X of a conjunctive cause and the corresponding effect Y as long as not all components are properly co-instantiated and alternative causal pathways to Y suppressed. It follows that X cannot be identified as a (M)INUS cause of Y by searching for suitable pairwise dependencies between X and Y but only by embedding X in a complex Boolean structure over many factors and fitting that structure as a whole to the data. But the space of Boolean functions over even a handful of factors is vast. So, a method for (M)INUS discovery must find ways to purposefully and efficiently navigate in that vast space of possibilities.

Methods of data analysis capable of handling (M)INUS structures have been developed in various disciplines—independently of one another. Apart from CNA, the methods with the most advanced software implementations are *Logic Regression* (LR; Ruczinski *et al.* 2003; Kooperberg and Ruczinski 2005) from biostatistics and *Qualitative Comparative Analysis* (QCA; Ragin 1987; Rihoux and Ragin 2009; Cronqvist and Berg-Schlusser 2009) from the social sciences.² LR, which is implemented in the R package **LogicReg** (Kooperberg and Ruczinski 2019)³ is mainly used to model higher-order interactions of SNPs in genetic association studies,⁴ to which end it fits Boolean functions to data by embedding them in a generalized regression framework and using the residual sum of squares to measure the model fit. It is optimized for data with high noise levels, large numbers of exogenous factors, and data-generating structures with no more than one endogenous factor (outcome). LR does not aim to build all data-fitting models—not even for data of low dimensionality—but uses search heuristics (standardly a simulated annealing algorithm) that zoom in on one or a small number of best fitting models. QCA, whose most powerful implementations are provided by the R packages **QCApro** (Thiem 2018) and **QCA** (Duşa 2007),⁵ has been widely applied in diverse areas in the social sciences and beyond (for an overview over corresponding publications see the bibliography section on the **COMPASSS** website). It is optimized for data with low noise levels, a maximum of around 15 exogenous factors, and, like LR, tracks data-generating structures with single outcomes only. Unlike LR, however, QCA does not rely on techniques from regression analysis but from Boolean algebra. In particular, it uses consistency and coverage (Ragin 2006a) as model fit parameters and builds its models using Quine-McCluskey optimization (QMC; McCluskey 1965; the QCA R packages use modern versions this optimization approach, see Duşa 2018), which is an algorithm from switching circuit theory that,

²In epidemiology, Rothman’s (1976) *sufficient-component cause model* with its various extensions (VanderWeele and Robins 2009; Wensink *et al.* 2014) is also designed for INUS discovery, but a proper software implementation is still lacking.

³Another package implementing a variation of LR is **logicFS** (Schwender and Tietz 2020).

⁴There is no methodological reason why the application of LR should be restricted to SNP studies. In fact, applying LR for (M)INUS discovery more generally would harness the potential for considerable synergies with CNA or QCA.

⁵Other useful QCA software include **QCAfalsePositive** (Braumoeller 2015) and **SetMethods** (Oana *et al.* 2020).

for data of low dimensionality, can be used to build all data-fitting models.

CNA has commonalities and differences with both QCA and LR. Like QCA, CNA is embedded in Boolean algebra and it also uses consistency and coverage as measures for data fit. But CNA does not build models by means of QMC, rather it comes with a search algorithm custom-built for causal data analysis. While QMC generates models from the top down by first building maximal Boolean dependency structures and then gradually eliminating redundant elements, CNA—on a par with LR—assembles causal models from the bottom up by gradually combining single factor values to complex dependency structures until the requested thresholds of model fit are met. One important upshot of this bottom-up approach is that redundant elements are not incorporated in models to begin with, meaning that models complying with fit thresholds are automatically redundancy-free. As a direct consequence, contrary to QCA and in line with LR, CNA straightforwardly handles data fragmentation (limited diversity) without resorting to counterfactual reasoning, it does not produce analogues to QCA’s so-called intermediate and conservative solutions, which are expressly allowed to comprise redundancies, and it is more successful at inferring redundancy-free sufficient and necessary conditions from noisy data than QCA’s so-called parsimonious search strategy (cf. Baumgartner and Ambuehl 2020), meaning—since redundancy-freeness is crucial for (M)INUS causation—that CNA is more successful at discovering (M)INUS causation in noisy data than any of QCA’s search approaches. Contrary to LR, CNA does not rely on a search heuristic but produces all data-fitting models within user-defined bounds for model complexity. Consequently, like QCA, CNA can currently only fully process data with a maximum of about 15 exogenous factors (in reasonable time). CNA and LR perform comparably when it comes to inferring (M)INUS structures with single outcomes from noisy data, with complementary strengths and weaknesses (for a detailed performance comparison see Baumgartner manuscript). Finally, unlike both QCA and LR, CNA allows for multiple effects (outcomes/endogenous factors), meaning it can analyze common-cause and causal chain structures; that is, it not only groups causes conjunctively and disjunctively but also sequentially.

Hence, CNA is the only currently available method for (M)INUS discovery and, correspondingly, **cna** the only currently available (M)INUS software that builds multi-outcome models. This vignette provides a detailed introduction to **cna**. We first exhibit **cna**’s theoretical and methodological background. Second, we discuss the main inputs of the package’s core function `cna()` along with numerous auxiliary functions for data review and simulation. Third, the working of the algorithm implemented in `cna()` is presented. Fourth, we explain `cna()`’s output along with relevant fit parameters and output attributes. Fifth, we provide some guidance on how to interpret that output and, in particular, on how to proceed in case of model ambiguities. Finally, some considerations are offered on benchmarking the reliability of `cna()`.

2. Background

The (M)INUS theory of causation belongs to the family of so-called *regularity theories*, which have roots as far back as (Hume 1999 (1748)). It is a type-level theory of causation (cf. Baumgartner 2020) that analyzes the dependence relation of causal relevance between factors/variables taking on specific values, as in “ $X=\chi$ is causally relevant to $Y=\gamma$ ”. It assumes that causation is ultimately a deterministic form of dependence, such that whenever the same complete cause occurs the same effect follows. This entails that indeterministic behavior pat-

terns in data result from insufficient control over background influences generating noise and not from the indeterministic nature of the underlying causal processes. For $X=\chi$ to be an (M)INUS cause of $Y=\gamma$, $X=\chi$ must be a difference-maker of $Y=\gamma$, meaning—roughly—that there exists a context in which other causes take constant values and a change from $X\neq\chi$ to $X=\chi$ is associated with a change from $Y\neq\gamma$ to $Y=\gamma$.

To further clarify that theory as well as the characteristics and requirements of inferring (M)INUS structures from empirical data a number of preliminaries are needed.

2.1. Factors and their values

Factors are the basic modeling devices of CNA. They are analogous to (random) variables in statistics, that is, they are functions from (measured) properties into a range of values (typically integers). They can be used to represent categorical properties that partition sets of units of observation (cases) either into two sets, in case of binary properties, or into more than two (but finitely many) sets, in case of multi-value properties. Factors representing binary properties can be *crisp-set* (*cs*) or *fuzzy-set* (*fs*); the former can take on 0 and 1 as possible values, whereas the latter can take on any (continuous) values from the unit interval $[0, 1]$. Factors representing multi-value properties are called *multi-value* (*mv*) *factors*; they can take on any of an open (but finite) number of non-negative integers.

Values of a *cs* or *fs* factor X can be interpreted as membership scores in the set of cases exhibiting the property represented by X . A case of type $X=1$ is a full member of that set, a case of type $X=0$ is a (full) non-member, and a case of type $X=\chi_i$, $0 < \chi_i < 1$, is a member to degree χ_i . An alternative interpretation, which lends itself particularly well for causal modeling, is that “ $X=1$ ” stands for the full presence of the property represented by X , “ $X=0$ ” for its full absence, and “ $X=\chi_i$ ” for its partial presence (to degree χ_i). By contrast, the values of an *mv* factor X designate the particular way in which the property represented by X is exemplified. For instance, if X represents the education of subjects, $X=2$ may stand for “high school”, with $X=1$ (“no completed primary schooling”) and $X=3$ (“university”) designating other possible property exemplifications. *Mv* factors taking on one of their possible values also define sets, but the values themselves must not be interpreted as membership scores; rather they denote the relevant property exemplification.

As the explicit “Factor=value” notation yields convoluted syntactic expressions with increasing model complexity, the **cna** package uses the following shorthand notation, which is conventional in Boolean algebra: membership in a set is expressed by italicized upper case and non-membership by lower case Roman letters. Hence, in case of *cs* and *fs* factors, we write “ X ” for $X=1$ and “ x ” for $X=0$. It must be emphasized that, while this notation significantly simplifies the syntax of causal models, it introduces a risk of misinterpretation, for it yields that the factor X and its taking on the value 1 are both expressed by “ X ”. Disambiguation must hence be facilitated by the concrete context in which “ X ” appears. Therefore, whenever we do not explicitly characterize italicized Roman letters as “factors” in this vignette, we use them in terms of the shorthand notation. In case of *mv* factors, value assignments to variables are not abbreviated but always written out, using the “Factor=value” notation.

2.2. Boolean operations

The (M)INUS theory defines causation using the Boolean operations of negation ($\neg X$, or x), conjunction ($X*Y$), disjunction ($X + Y$), implication ($X \rightarrow Y$), and equivalence ($X \leftrightarrow Y$).

Inputs		Outputs				
X	Y	$\neg X$	$X*Y$	$X + Y$	$X \rightarrow Y$	$X \leftrightarrow Y$
1	1	0	1	1	1	1
1	0	0	0	1	0	0
0	1	1	0	1	1	0
0	0	1	0	0	1	1

Table 1: Classical Boolean operations applied to *cs* factors.

Negation is a unary truth function, the other operations are binary truth functions. That is, they take one resp. two truth values as inputs and output a truth value. When applied to *cs* factors, both their input and output set is $\{0, 1\}$. Negation is typically translated by “not”, conjunction by “and”, disjunction by “or”, implication by “if ... then”, and equivalence by “if and only if (iff)”. Their classical definitions are given in Table 1.

These operations can be straightforwardly applied to *mv* factors as well, in which case they amount to functions from the *mv* factors’ domain of values into the set $\{0, 1\}$. To illustrate, assume that both X and Y are ternary factors with values from the domain $\{0, 1, 2\}$. The negation of $X=2$, *viz.* $\neg(X=2)$, then returns 1 iff X is not 2, meaning iff X is 0 or 1. $X=2*Y=0$ yields 1 iff X is 2 and Y is 0. $X=2 + Y=0$ returns 1 iff X is 2 or Y is 0. $X=2 \rightarrow Y=0$ yields 1 iff either X is not 2 or Y is 0. $X=2 \leftrightarrow Y=0$ issues 1 iff either X is 2 and Y is 0 or X is not 2 and Y is not 0.

For *fs* factors whose values are interpreted as membership scores in fuzzy sets, the classical Boolean operations must be translated into fuzzy logic. There exist numerous systems of fuzzy logic (for an overview cf. Hájek 1998), each of which comes with its own rendering of Boolean operations. In the context of CNA (and QCA), the following fuzzy-logic renderings are standard: negation $\neg X$ is translated in terms of $1 - X$, conjunction $X*Y$ in terms of the minimum membership score in X and Y , i.e., $\min(X, Y)$, disjunction $X + Y$ in terms of the maximum membership score in X and Y , i.e., $\max(X, Y)$, an implication $X \rightarrow Y$ is taken to express that the membership score in X is smaller or equal to Y ($X \leq Y$), and an equivalence $X \leftrightarrow Y$ that the membership scores in X and Y are equal ($X = Y$).

Based on the implication operator, the notions of *sufficiency* and *necessity* are defined, which are the two Boolean dependencies exploited by the (M)INUS theory:

Sufficiency X is sufficient for Y iff $X \rightarrow Y$ (or equivalently: $x + Y$; and colloquially: “if X is present, then Y is present”);

Necessity X is necessary for Y iff $Y \rightarrow X$ (or equivalently: $\neg X \rightarrow \neg Y$ or $y + X$; and colloquially: “if Y is present, then X is present”).

Analogously for more complex expressions:

- $X=3*Z=2$ is sufficient for $Y=4$ iff $X=3*Z=2 \rightarrow Y=4$;
- $X=3 + Z=2$ is necessary for $Y=4$ iff $Y=4 \rightarrow X=3 + Z=2$;
- $X=3 + Z=2$ is sufficient and necessary for $Y=4$ iff $X=3 + Z=2 \leftrightarrow Y=4$.

2.3. (M)INUS causation

Boolean dependencies of sufficiency and necessity amount to mere patterns of co-occurrence

of factor values; as such, they carry no causal connotations whatsoever. In fact, most Boolean dependencies do not reflect causal dependencies. To mention just two well-rehearsed examples: the sinking of a properly functioning barometer in combination with high temperatures and blue skies is sufficient for weather changes, but it does not cause the weather; or whenever it rains, the street gets wet, hence, wetness of the street is necessary for rainfall but certainly not causally relevant for it. At the same time, some dependencies of sufficiency and necessity are in fact due to underlying causal dependencies: rainfall is sufficient for wet streets and also a cause thereof, or the presence of oxygen is necessary for fires and also a cause thereof.

That means the crucial problem to be solved by the (M)INUS theory is to filter out those Boolean dependencies that are due to underlying causal dependencies and are, hence, amenable to a causal interpretation. The main reason why most sufficiency and necessity relations do not reflect causation is that they either contain redundancies or are themselves redundant to account for the behavior of the outcome, whereas causal conditions do not feature redundant elements and are themselves indispensable to account for the outcome in at least one context. Accordingly, to filter out the causally interpretable Boolean dependencies, they need to be freed of redundancies. In Mackie's (1974, 62) words, causes are *Insufficient but Non-redundant parts of Unnecessary but Sufficient conditions* (thus the acronym INUS).

While Mackie's INUS theory only requires that sufficient conditions be freed of redundancies, he himself formulates a problem for that theory, *viz.* the *Manchester Factory Hooters* problem (Mackie 1974, 81-87), which Graßhoff and May (2001) solve by eliminating redundancies also from necessary conditions. Accordingly, modern versions of the INUS theory stipulate that whatever can be removed from sufficient or necessary conditions without affecting their sufficiency and necessity is not a difference-maker and, hence, not a cause. The causally interesting sufficient and necessary conditions are *minimal* in the following sense:

Minimal sufficiency A conjunction Φ of coincidentally instantiated factor values (e.g., $X_1 * \dots * X_n$) is a minimally sufficient condition of Y iff $\Phi \rightarrow Y$ and there does not exist a proper part Φ' of Φ such that $\Phi' \rightarrow Y$, where a proper part Φ' of Φ is the result of eliminating one or more conjuncts from Φ .

Minimal necessity A disjunction Ψ of minimally sufficient conditions (e.g., $\Phi_1 + \dots + \Phi_n$) is a minimally necessary condition of Y iff $Y \rightarrow \Psi$ and there does not exist a proper part Ψ' of Ψ such that $Y \rightarrow \Psi'$, where a proper part Ψ' of Ψ is the result of eliminating one or more disjuncts from Ψ .

Minimally sufficient and minimally necessary conditions can be combined to so-called *atomic MINUS-formulas* (Beirlaen *et al.* 2018; or, equivalently, *minimal theories*, Graßhoff and May 2001), which, in turn, can be combined to *complex MINUS-formulas*:⁶

Atomic MINUS-formula An atomic MINUS-formula of an outcome Y is an expression $\Psi \leftrightarrow Y$, where Ψ is a minimally necessary disjunction of minimally sufficient conditions of Y , in disjunctive normal form.⁷

⁶We provide suitably simplified definitions that suffice for our purposes here. For complete definitions see (Baumgartner and Falk 2019).

⁷An expression is in disjunctive normal form iff it is a disjunction of one or more conjunctions of one or more literals (i.e. factor values; Lemmon 1965, 190).

Complex MINUS-formula A complex MINUS-formula of outcomes Y_1, \dots, Y_n is a conjunction $(\Psi_1 \leftrightarrow Y_1) * \dots * (\Psi_n \leftrightarrow Y_n)$ of atomic MINUS-formulas that is itself redundancy-free, meaning it does not contain a logically equivalent proper part.⁸

MINUS-formulas connect Boolean dependencies to causal dependencies: only those Boolean dependencies are causally interpretable that appear in MINUS-formulas. To make this concrete, consider the following atomic MINUS-formula:

$$A*e + C*d \leftrightarrow B \quad (1)$$

(1) being a MINUS-formula of B entails that $A*e$ and $C*d$, but neither A , e , C , nor d alone, are sufficient for B and that $A*e + C*d$, but neither $A*e$ nor $C*d$ alone, are necessary for B . If this holds, it follows that for each factor value in (1) there exists a *difference-making pair*, meaning a pair of cases such that a change in that factor value alone accounts for a change in the outcome (Baumgartner and Falk 2019, 9). For example, A being part of the MINUS-formula (1) entails that there are two cases σ_i and σ_j such that e is given and $C*d$ is not given in both σ_i and σ_j while A and B are present in σ_i and absent in σ_j . Only if such a difference-making pair $\langle \sigma_i, \sigma_j \rangle$ exists is A indispensable to account for B . Analogously, (1) being a MINUS-formula entails that there exist difference-making pairs for all other factor values in (1).

To define causation in terms of Boolean difference-making, an additional constraint is needed because not all MINUS-formulas faithfully represent causation. Complete redundancy elimination is relative to the set of analyzed factors \mathbf{F} , meaning that factor values contained in MINUS-formulas relative to some \mathbf{F} may fail to be part of a MINUS-formulas relative to supersets of \mathbf{F} (Baumgartner 2013). In other words, by adding further factors to the analysis, factor values that originally appeared to be non-redundant to account for an outcome can turn out to be redundant after all. Hence, a *permanence* constraint needs to be imposed: only factor values that are permanently non-redundant, meaning that cannot be rendered redundant by expanding factor sets, are causally relevant.

These considerations yield the following definition of causation:

Causal Relevance (MINUS) X is causally relevant to Y if, and only if, (I) X is part of a MINUS-formula of Y relative to some factor set \mathbf{F} and (II) X remains part of a MINUS-formula of Y across all expansions of \mathbf{F} .

MINUS-formulas whose elements satisfy the permanence constraint not only identify causally relevant factor values but also place a Boolean ordering over these causes, such that conjunctivity, disjunctivity, and outcome-interlinking can be directly read off their syntax. Take the following example:

$$(A*b + a*B \leftrightarrow C) * (C*f + D \leftrightarrow E) \quad (2)$$

If (2) complies with (MINUS.I) and (MINUS.II), it entails these causal claims:

1. the factor values listed on the left-hand sides of “ \leftrightarrow ” are causally relevant for the factor values on the right-hand sides;

⁸The purpose of the redundancy-freeness constraint imposed on complex MINUS-formulas is to avoid structural and partial structural redundancies; see section 5.5 below.

2. A and b are jointly relevant to C and located on a causal path that differs from the path on which the jointly relevant a and B are located; C and f are jointly relevant to E and located on a path that differs from D 's path;
3. there is a causal chain from $A*b$ and $a*B$ via C to E .

2.4. Inferring MINUS causation from data

Inferring MINUS causation from data faces various challenges. First, as anticipated in section 1, causal structures for which conjunctivity and disjunctivity hold cannot be uncovered by scanning data for dependencies between pairs of factor values and suitably combining dependent pairs. Instead, discovering MINUS causation requires searching for dependencies between complex Boolean functions of exogenous factors and outcomes. But for n *cs* factors there exist 2^{2^n} possible Boolean functions; and in case of mv factors that space of possibilities explodes even more. So, a method for MINUS discovery must find ways to efficiently navigate in that vast search space.

Second, condition (MINUS.II) is not exhaustively testable. Once a MINUS-formula of an outcome Y comprising a factor value X has been inferred from data δ , the question arises whether the non-redundancy of X in accounting for Y is an artefact of δ , due, for example, to the uncontrolled variation of confounders, or whether it is genuine and persists when further factors are taken into consideration. But expanding the set of factors is only feasible within narrow confines in scientific practice. To make up for the impossibility to test (MINUS.II), the analyzed data δ should be collected in such a way that Boolean dependencies in δ are not induced by an uncontrolled variation of latent causes but by actual causal dependencies among the measured factors. If the dependencies in δ are not artefacts of latent causes, they cannot be neutralized by factor set expansions, meaning they are permanent and, hence, causal. It follows that in order for it to be certain that causal inferences drawn from δ are fallacy-free, δ must meet very high quality standards. In particular, the uncontrolled causal background of δ must be *homogeneous* (Baumgartner and Thiem 2020, 286):

Homogeneity The unmeasured causal background of data δ is homogeneous if, and only if, latent causes not connected to the outcome(s) in δ on causal paths via the measured exogenous factors (so-called *off-path* causes) take constant values (i.e. do not vary) in the cases recorded in δ .

However, third, real-life data often do not meet very high quality standards. Rather, they tend to be *fragmented* to the effect that not all empirically possible configurations of analyzed factors are actually observed. Moreover, real-life data typically feature *noise*, that is, configurations incompatible with data-generating causal structures. Noise is induced, for instance, by measurement error or limited control over latent causes, i.e. confounding. In the presence of noise there may be no strict Boolean sufficiency or necessity relations in the data, meaning that methods of MINUS discovery can only approximate strict MINUS structures by fitting their models more or less closely to the data using suitable parameters and thresholds of model fit. Moreover, noise stemming from the uncontrolled variation of latent causes gives rise to homogeneity violations, which yield that inferences to MINUS causation are not guaranteed to be fallacy-free. In order to nonetheless distill causal information from noisy data, strategies for avoiding over- and underfitting and estimating the fallacy-risk are needed.

Fourth, according to the MINUS theory, the inference to *causal irrelevance* is much more demanding than the inference to causal relevance. Establishing that X is a MINUS cause of Y requires demonstrating the existence of at least one context with a constant background in which a difference in X is associated with a difference in Y , whereas establishing that X is *not* a MINUS cause of Y requires demonstrating the *non-existence* of such a context, which is impossible on the basis of the non-exhaustive data samples that are typically analyzed in real-life studies. Correspondingly, the fact that, say, G does not appear in (2) does not imply that G is causally irrelevant to C or E . The non-inclusion of G simply means that the data from which (2) has been derived do not contain evidence for the causal relevance of G . However, future research having access to additional data might reveal the existence of a difference-making context for G and, hence, entail the causal relevance of G to C or E after all.

Finally, on a related note, as a result of the common fragmentation of real-life data δ MINUS-formulas inferred from δ cannot be expected to completely reflect the causal structure generating δ . That is, MINUS-formulas inferred from δ are inevitably going to be *incomplete*. They only detail those causally relevant factor values along with those conjunctive, disjunctive, and sequential groupings for which δ contain difference-making evidence. What difference-making evidence is contained in δ does not only depend on the cases recorded in δ but, when δ is noisy, also on the tuning thresholds imposed to approximate strict Boolean dependency structures; relative to some such tuning settings an association between X and Y may pass as a sufficiency or necessity relation whereas relative to another setting it will not. Hence, the inference to MINUS causation is sensitive to the chosen tuning settings, to the effect that choosing different settings is often going to be associated with changes in inferred MINUS-formulas.

Variance in inferred MINUS-formulas may be entirely unproblematic. Two different MINUS-formulas \mathbf{m}_i and \mathbf{m}_j derived from δ using different tuning settings are in no disagreement if the causal claims entailed by \mathbf{m}_i and \mathbf{m}_j stand in a subset relation; that is, if one formula is a submodel of the other:

Submodel relation A MINUS-formula \mathbf{m}_i is a *submodel* of another MINUS-formula \mathbf{m}_j if, and only if,

- (i) all factor values causally relevant according to \mathbf{m}_i are also causally relevant according to \mathbf{m}_j ,
- (ii) all factor values contained in two different disjuncts in \mathbf{m}_i are also contained in two different disjuncts in \mathbf{m}_j ,
- (iii) all factor values contained in the same conjunct in \mathbf{m}_i are also contained in the same conjunct in \mathbf{m}_j , and
- (iv) if \mathbf{m}_i and \mathbf{m}_j are complex MINUS-formulas, all atomic components \mathbf{m}_i^k of \mathbf{m}_i have a counterpart \mathbf{m}_j^k in \mathbf{m}_j such (i) to (iii) are satisfied for \mathbf{m}_i^k and \mathbf{m}_j^k .

If \mathbf{m}_i is a submodel of \mathbf{m}_j , \mathbf{m}_j is a *supermodel* of \mathbf{m}_i . All of \mathbf{m}_i 's causal ascriptions are contained in its supermodels' ascriptions, and \mathbf{m}_i contains the causal ascriptions of its own submodels. The submodel relation is reflexive: every model is a submodel (and supermodel) of itself; or differently, if \mathbf{m}_i and \mathbf{m}_j are submodels of one another, then \mathbf{m}_i and \mathbf{m}_j are identical. Most importantly, if two MINUS-formulas related by the submodel relation are

not identical, they can be interpreted as describing the same causal structure with varying granularity.

Before we turn to the details of MINUS discovery using the **cna** package, a terminological note is required. In the literature on configurational comparative methods it has become customary to refer to the models produced by the methods as *solution formulas*. To mirror that convention, the **cna** package refers to atomic MINUS-formulas inferred from data by CNA as *atomic solution formulas*, **asf**, for short, and to complex MINUS-formulas inferred from data as *complex solution formulas*, **csf**. For brevity, we will henceforth mainly use the shorthands *asf* and *csf*.

3. The input of CNA

The goal of CNA is to output all *asf* and *csf* within provided bounds of model complexity that fit an input data set relative to provided tuning settings, in particular, fit thresholds. The algorithm performing this task in the **cna** package is implemented in the function `cna()`. Its most important arguments are:

```
cna(x, type, ordering = NULL, strict = FALSE, con = 1, cov = 1, con.msc = con,
    notcols = NULL, maxstep = c(3, 4, 10), inus.only = TRUE, suff.only = FALSE,
    what = if (suff.only) "m" else "ac", details = FALSE, acyclic.only = FALSE)
```

This section explains most of these inputs and introduces some auxiliary functions. The arguments `inus.only`, `what`, `details`, and `acyclic.only` will be discussed in section 5.

3.1. Data

Data δ processed by CNA have the form of $m \times k$ matrices, where m is the number of units of observation (cases) and k is the number of measured factors. δ can either be of type “crisp-set” (*cs*), “multi-value” (*mv*) or “fuzzy-set” (*fs*). Data that feature *cs* factors only are *cs*. If the data contain at least one *mv* factor, they count as *mv*. Data featuring at least one *fs* factor are *fs*.⁹ Examples of each data type are given in Table 2. Raw data collected in a study typically need to be suitably calibrated before they can be fed to `cna()`. We do not address the calibration problem here because it is the same for CNA as for QCA, in which context it has been extensively discussed, for example, by Thiem and Duşa (2013) or Schneider and Wagemann (2012). The R packages **QCApro**, **QCA**, and **SetMethods** provide all tools necessary for automating data calibration.

Data is given to the `cna()` function via the argument `x`, which is a data frame or an object of class “configTable” as output by the `configTable()` function (see section 3.1.1 below). The **cna** package contains a number of exemplary data sets from published studies: `d.autonomy`, `d.educate`, `d.irrigate`, `d.jobsecurity`, `d.minaret`, `d.pacts`, `d.pban`, `d.performance`, `d.volatile`, `d.women`. For details on their contents and sources, see the **cna reference manual**. After having loaded the **cna** package, all of them are directly (i.e. without separate loading) available for processing:

⁹Mixing *mv* and *fs* factors in one analysis is (currently) not supported.

	A	B	C	D		A	B	C	D		A	B	C	D	E
c_1	0	0	0	0	c_1	1	3	3	1	c_1	0.37	0.30	0.16	0.06	0.25
c_2	0	1	0	0	c_2	2	2	1	2	c_2	0.89	0.39	0.64	0.09	0.03
c_3	1	1	0	0	c_3	2	1	2	2	c_3	0.06	0.61	0.92	0.37	0.15
c_4	0	0	1	0	c_4	2	2	2	2	c_4	0.65	0.93	0.92	0.18	0.93
c_5	1	0	0	1	c_5	3	3	3	2	c_5	0.08	0.08	0.12	0.86	0.91
c_6	1	0	1	1	c_6	2	4	3	2	c_6	0.70	0.02	0.85	0.91	0.97
c_7	0	1	1	1	c_7	1	3	3	3	c_7	0.04	0.72	0.76	0.90	0.68
c_8	1	1	1	1	c_8	1	4	3	3	c_8	0.81	0.96	0.89	0.72	0.82

(a) *cs* data(b) *mv* data(c) *fs* data

Table 2: Data types processable by CNA.

```
R> library(cna)
R> cna(d.educate)
R> cna(d.women)
```

If the data are not of type *cs*, `cna()` must be told explicitly what type of data `x` contains using the `type` argument, which takes the values "mv" for *mv* data and "fs" for *fs* data. The functions `mvna(x, ...)` and `fsna(x, ...)` are available as shorthands for `cna(x, type = "mv", ...)` and `cna(x, type = "fs", ...)`, respectively.

```
R> cna(d.jobsecurity, type = "fs")
R> fsna(d.jobsecurity)
R> cna(d.pban, type = "mv")
R> mvna(d.pban)
```

Configuration tables

To facilitate the reviewing of data, the `configTable()` function assembles cases with identical configurations in a table called a *configuration table*. In previous versions of the `cna` package, these tables were called “truth tables”, which however led to confusion with the QCA terminology, where a very different type of object is also referred to by that label. While a QCA truth table indicates for every configuration of all exogenous factors (i.e. for every minterm) whether it is sufficient for the outcome, a CNA configuration table does not express relations of sufficiency but simply amounts to a compact representation of the data that lists all configurations exactly once and adds a column indicating how many instances (cases) of each configuration are contained in the data.

```
configTable(x, type = c("cs", "mv", "fs"), case.cutoff = 0)
```

The first input `x` is a data frame or matrix. The function then merges multiple rows of `x` featuring the same configuration into one row, such that each row of the resulting table corresponds to one determinate configuration of the factors in `x`. The number of occurrences of a configuration and an enumeration of the cases instantiating it are saved as attributes “n” and “cases”, respectively. When not applied to *cs* data, the data type must be specified with

the `type` argument. Alternatively, the shorthand functions `csct(x)`, `mvct(x)` and `fsct(x)` are available.

```
R> configTable(d.women)
R> mvct(d.pban)
```

`configTable()` provides a numeric argument called `case.cutoff`, which allows for setting a minimum frequency cutoff determining that configurations with less instances in the data are not included in the configuration table and the ensuing analysis. For instance, `configTable(x, case.cutoff = 3)` entails that configurations that are instantiated in less than 3 cases are excluded.

Configuration tables produced by `configTable()` can be directly passed to `cna()`. Moreover, as configuration tables generated by `configTable` are objects that are very particular to the **cna** package, the function `ct2df()` is available to transform configuration tables back into ordinary R data frames.

```
R> pact.ct <- configTable(d.pacts, type = "fs", case.cutoff = 2)
R> ct2df(pact.ct)
```

Data simulations

The **cna** package provides extensive functionalities for data simulations—which, in turn, are essential for inverse search trials that benchmark CNA’s output (see section 7). In a nutshell, the functions `allCombs()` and `full.ct()` generate the space of all logically possible configurations over a given set of factors, `selectCases()` selects, from this space, the configurations that are compatible with a data-generating causal structure, which, in turn, can be randomly drawn by `randomAsf()` and `randomCsf()`, `makeFuzzy()` fuzzifies that data, and `some()` randomly selects cases, for instance, to produce data fragmentation.

More specifically, `allCombs(x)` takes an integer vector `x` as input and generates a data frame of all possible value configurations of `length(x)` factors, the first factor having `x[1]` values, the second `x[2]` values etc. The factors are labeled using capital letters in alphabetical order. Analogously, but more flexibly, `full.ct(x)` generates a configuration table with all logically possible value configurations of the factors defined in the input `x`, which can be a configuration table, a data frame, an integer, a list specifying the factors’ value ranges, or a character vector featuring all admissible factor values.

```
R> allCombs(c(2, 2, 2)) - 1
R> allCombs(c(3, 4, 5))
R> full.ct("A + B*c")
R> full.ct(6)
R> full.ct(list(A = 1:2, B = 0:1, C = 1:4))
```

The input of `selectCases(cond, x)` is a character string `cond` specifying a Boolean function, which typically (but not necessarily) expresses a data-generating MINUS structure, as well as, optionally, a data frame or configuration table `x`. If `x` is specified, the function selects the cases that are compatible with `cond` from `x`; if `x` is not specified, it selects from `full.ct(cond)`. It is possible to randomly draw `cond` using `randomAsf(x)` or `randomCsf(x)`, which generate

random atomic and complex solution (i.e. MINUS-) formulas, respectively, from a data frame or configuration table `x`.

```
R> dat1 <- allCombs(c(2, 2, 2)) - 1
R> selectCases("A + B <-> C", dat1)
R> selectCases("(h*F + B*C*k + T*r <-> G)*(A*b + H*I*K <-> E)")
R> target <- randomCsf(full.ct(6))
R> selectCases(target)
```

The closely related function `selectCases1(cond, x, con = 1, cov = 1)` additionally allows for providing consistency (`con`) and coverage (`cov`) thresholds (see section 3.2), such that some cases that are incompatible with `cond` are also selected, as long as `cond` still meets `con` and `cov` in the resulting data. Thereby, measurement error or noise can be simulated in a manner that allows for controlling the degree of case incompatibilities.

```
R> dat2 <- full.ct(list(EN = 0:2, TE = 0:4, RU = 1:4))
R> selectCases1("EN=1*TE=3 + EN=2*TE=0 <-> RU=2", dat2, con = .75, cov = .75)
```

`makeFuzzy(x, fuzzvalues = c(0, 0.05, 0.1))` simulates fuzzy-set data by transforming a data frame or configuration table `x` consisting of *cs* factors into an *fs* configuration table. To this end, the function adds values selected at random from the argument `fuzzvalues` to the 0's and subtracts them from the 1's in `x`. `fuzzvalues` is a numeric vector of values from the interval $[0,1]$.

```
R> makeFuzzy(selectCases("Hunger + Heat <-> Run"),
+           fuzzvalues = seq(0, 0.4, 0.05))
```

Finally, `some(x, n = 10, replace = TRUE)` randomly selects `n` cases from a data frame or configuration table `x`, with or without replacement. If `x` features all configurations that are compatible with a data-generating structure and `n < nrow(x)`, the data frame or configuration table issued by `some()` is *fragmented*, meaning it does not contain all empirically possible configurations. If `n > nrow(x)`, data of large sample sizes can be generated featuring multiple instances of the empirically possible configurations.

```
R> dat3 <- allCombs(c(3, 4, 5))
R> dat4 <- selectCases("A=1*B=3 + A=3 <-> C=2", mvct(dat3))
R> some(dat4, n = 10, replace = FALSE)
R> some(dat4, n = 1000, replace = TRUE)
```

3.2. Consistency and coverage

As real-life data tend to feature measurement error or noise induced by variations in latent causes, strictly sufficient or necessary conditions for an outcome often do not exist. In order to nonetheless distill some causal information from such data, methods for MINUS discovery have to suitably fit their models to the data. To this end, Ragin (2006b) imported the so-called *consistency* and *coverage* measures (with values from the interval $[0,1]$) into the QCA protocol. Both of these measures are also serviceable for the purposes of CNA. Informally put,

consistency reflects the degree to which the behavior of an outcome obeys a corresponding sufficiency or necessity relationship or a whole model, whereas coverage reflects the degree to which a sufficiency or necessity relationship or a whole model accounts for the behavior of the corresponding outcome. As the implication operator underlying the notions of sufficiency and necessity is defined differently in classical and in fuzzy logic, the two measures are defined differently for crisp-set and multi-value data (which both have a classical footing), on the one hand, and fuzzy-set data, on the other. *Cs-consistency* (con^{cs}) and *cs-coverage* (cov^{cs}) are defined as follows:

$$con^{cs}(X \rightarrow Y) = \frac{|X*Y|_{\delta}}{|X|_{\delta}} \quad cov^{cs}(X \rightarrow Y) = \frac{|X*Y|_{\delta}}{|Y|_{\delta}}$$

where X and Y are individual factors or Boolean functions of factors in the analyzed data δ and $|\dots|_{\delta}$ stands for the cardinality of the set of cases in δ instantiating the enclosed expression. *Fs-consistency* (con^{fs}) and *fs-coverage* (cov^{fs}) of $X \rightarrow Y$ are defined as follows, where n is the number of cases in the data:

$$con^{fs}(X \rightarrow Y) = \frac{\sum_{i=1}^n \min(X_i, Y_i)}{\sum_{i=1}^n X_i} \quad cov^{fs}(X \rightarrow Y) = \frac{\sum_{i=1}^n \min(X_i, Y_i)}{\sum_{i=1}^n Y_i}$$

Whenever the values of X and Y are restricted to 1 and 0 in the crisp-set measures, con^{cs} and cov^{cs} are equivalent to con^{fs} and cov^{fs} , but for binary factors with values other than 1 and 0 and for multi-value factors that equivalence does not hold. Nonetheless, we will in the following not explicitly distinguish between the *cs* and *fs* measures because our discussion will make it sufficiently clear which of them is at issue.

Consistency and coverage thresholds can be given to the `cna()` function using the arguments `con.msc`, `con`, and `cov` that take values from the interval $[0, 1]$. `con.msc` sets the consistency threshold for minimally sufficient conditions (*msc*), `con` does the same for *asf* and *csf*, while `cov` sets the coverage threshold for *asf* and *csf* (no coverage threshold is imposed on *msc*). As illustrated on p. 17 of the [cna reference manual](#), setting different consistency thresholds for *msc* and *asf/csf* can enhance the informativeness of `cna()`'s output in certain cases but is non-standard. The standard setting is `con = con.msc`.

The default numeric value for all thresholds is 1, i.e. perfect consistency and coverage. Contrary to QCA, which often returns solutions that do not comply with the chosen consistency threshold and which does not impose a coverage threshold at all, CNA uses consistency and coverage as authoritative model building criteria such that, if they are not met, CNA abstains from issuing solutions. That means, if the default thresholds are used, `cna()` will only output perfectly consistent *msc*, *asf*, and *csf* and only perfectly covering *asf* and *csf*.

If the data are noisy, the default thresholds will typically not yield any solution formulas. In such cases, `con` and `cov` may be suitably lowered. By lowering `con` to, say, 0.75 in a *cs* analysis, `cna()` is given permission to treat X as sufficient for Y , even though in 25% of the cases X is not associated with Y . Or by lowering `cov` to 0.75 in an *fs* analysis, `cna()` is allowed to treat X as necessary for Y , even though the sum of the membership scores in Y over all cases in the data exceeds the sum of the membership scores in $\min(X, Y)$ by 25%.

Determining to which value `con` and `cov` are best lowered in a concrete discovery context is a delicate task. On the one hand, CNA faces a severe overfitting risk when the data contain configurations incompatible with the data-generating structure, meaning that `con` and `cov` must not be set too high (i.e. too close to 1). On the other hand, the lower `con` and `cov` are

set, the less complex and informative CNA's output will be, that is, the more CNA's purpose of uncovering causal complexity will be undermined. To find a suitable balance between over- and underfitting, Parkkinen and Baumgartner (forthcoming) systematically re-analyze the data at all `con` and `cov` settings in the interval $[0.7, 1]$, collect all solutions resulting from such a re-analysis series in a set \mathbf{M} , and select the solution formula(s) with the most sub- and supermodels in \mathbf{M} . These are the solutions with the highest *overlap in causal ascriptions* with the other solutions in \mathbf{M} . They are the most *robust* solutions inferrable from the data. This approach to robustness scoring has been automatized in the `frscore()` function obtainable from the [replication material](#) of (Parkkinen and Baumgartner forthcoming).

If the analyst does not want to conduct a whole robustness analysis, reasonable non-perfect consistency and coverage settings are `con = cov = 0.8` or `0.75`. To illustrate, `cna()` does not build solutions for the `d.jobsecurity` data at the following `con` and `cov` thresholds (the arguments `ordering` and `strict` are explained in section 3.3 below):

```
R> fscna(d.jobsecurity, ordering = list("JSR"), strict = TRUE,
+       con = 1, cov = 1)
R> fscna(d.jobsecurity, ordering = list("JSR"), strict = TRUE,
+       con = .9, cov = .9)
```

But if `con` and `cov` are set to `0.75`, 20 solutions are returned:

```
R> fscna(d.jobsecurity, ordering = list("JSR"), strict = TRUE,
+       con = .75, cov = .75)
```

In the presence of noise, it is generally advisable to vary the `con` and `cov` settings to some degree in order to get a sense for how sensitive the model space reacts to changes in tuning settings and for the overlap in causal ascriptions between different solutions. Less complex solutions with more overlap with the rest of the returned solutions are generally preferable over more complex solutions with less overlap. If the consistency and coverage scores of resulting solutions can be increased by raising the `con` and `cov` settings without, at the same time, disproportionately increasing the solutions' complexity, solutions with higher fit are preferable over solutions with lower fit. But if an increase in fit comes with a substantive increase in model complexity, less complex models with lower fit are to be preferred (to avoid overfitting).

3.3. Ordering

In principle, the `cna()` function does not need to be told which factors in the data \mathbf{x} are endogenous (i.e. effects) and which ones are exogenous (i.e. causes). It attempts to infer that from \mathbf{x} . But if prior causal knowledge is available as to which factors can figure as effects and which ones cannot, this information can be given to `cna()` via a *causal ordering*. A causal ordering is a relation $X_i \prec X_j$ defined on the factors in \mathbf{x} entailing that X_j cannot be a cause of X_i (e.g. because X_i is instantiated temporally before X_j). That is, an ordering excludes certain causal dependencies but does not stipulate any. If an ordering is provided, `cna()` only searches for MINUS-formulas in accordance with the ordering; if no ordering is provided, `cna()` treats all values of the factors in \mathbf{x} as potential outcomes and explores whether a MINUS-formula for them can be inferred.

An ordering is given to `cna()` via the argument `ordering`, which takes as value a list of character vectors specifying the causal ordering of the factors in `x`. For example, `ordering = list(c("A", "B"), "C")` determines that C is causally located after A and B (i.e. $A, B \prec C$), meaning that C is not a potential cause of A and B . The latter are located on the same level of the ordering, for A and B are unrelated by \prec , whereas C is located on a level that is downstream of the A, B -level. `cna()` then only checks whether values of A and B can be modeled as causes of values of C ; the test for a causal dependency in the upstream direction is skipped. If the argument `ordering` is not specified, `cna()` searches for dependencies between all factors in `x`. An ordering does not need to explicitly mention all factors in `x`. If only a subset of the factors are included in the ordering, the non-included factors are entailed to be causally before the included ones. Hence, `ordering = list("C")` means that C is causally located after all other factors in `x`.

Additionally, the logical argument `strict` is available. It determines whether the elements of one level in an ordering can be causally related or not. For example, if `ordering = list(c("A", "B"), "C")` and `strict = TRUE`, then A and B are excluded to be causally related and `cna()` skips corresponding tests. By contrast, if `ordering = list(c("A", "B"), "C")` and `strict = FALSE`, then `cna()` also searches for dependencies among A and B .

Let us illustrate with the data set `d.autonomy`. Relative to the following function call, which stipulates that AU cannot be a cause of EM, SP , and CO and that the latter factors are not mutually causally related, `cna()` infers that SP is causally relevant to AU (i.e. $SP \leftrightarrow AU$; the function `csf()` builds the `csf` from a `cna()` solution object; see section 5.1):

```
R> dat.aut.1 <- d.autonomy[15:30, c("AU", "EM", "SP", "CO")]
R> ana.aut.1 <- fscna(dat.aut.1, ordering = list(c("EM", "SP", "CO"), "AU"),
+   strict = TRUE, con = .91, cov = .91)
R> printCols <- c("condition", "consistency", "coverage")
R> csf(ana.aut.1)[printCols]
```

```
   condition consistency coverage
1 SP <-> AU           0.935     0.915
```

If we set `strict` to `FALSE` and, thereby, allow for causal dependencies among EM, SP , and CO , it turns out that SP not only causes AU , but, on another causal path, also makes a difference to EM :

```
R> ana.aut.2 <- fscna(dat.aut.1, ordering = list(c("EM", "SP", "CO"), "AU"),
+   strict = FALSE, con = .91, cov = .91)
R> csf(ana.aut.2)[printCols]
```

```
   condition                                consistency coverage
1 (SP <-> AU)*(SP + CO <-> EM)           0.912     0.915
```

3.4. Maxstep

As anticipated in section 1 and as will be exhibited in more detail in section 4, `cna()` builds minimally necessary disjunctions of minimally sufficient conditions (i.e. `asf`) from the bottom

up by gradually permutating and testing conjunctions and disjunctions of increasing complexity for sufficiency and necessity. The combinatorial search space that this algorithm has to scan depends on a variety of different aspects, for instance, on the number of factors in \mathbf{x} , on the number of values these factors can take, on the number and length of the *msc* recovered in the first computational phase, etc. As the search space may be too large to be exhaustively scanned in reasonable time, the argument `maxstep` allows for setting an upper bound for the complexity of the generated *asf*. `maxstep` takes a vector of three integers $c(i, j, k)$ as input, entailing that the generated *asf* have maximally j disjuncts with maximally i conjuncts each and a total of maximally k factors. The default is `maxstep = c(3,4,10)`. The user can set it to any complexity level if computational time is not an issue.

The `maxstep` argument is particularly relevant for the analysis of data featuring severe model ambiguities. A telling case in point is the data set `d.volatile`. At the default `maxstep`, `cna()` recovers 416 *csf*. By increasing `maxstep`, this number increases rapidly. (In the standard print method of `cna()`, the `n.init` parameter in `csf()` is set to 1000; to get all *csf* from the second call below, this parameter needs to be increased. See section 5.1 for details.)

```
R> cna(d.volatile, ordering = list("V02"), maxstep = c(3,4,10))
R> vol1 <- cna(d.volatile, ordering = list("V02"), maxstep = c(4,4,10))
R> csf(vol1, n.init = 3000)
```

If the values of `maxstep` are further increased, the analysis will quickly fail to terminate in reasonable time. When a complete analysis cannot be completed, `cna()` can be told to only search for minimally sufficient conditions (*msc*) by setting the argument `suff.only` to its non-default value `TRUE`. As the search for *msc* is the part of a CNA analysis that is least computationally demanding, it will typically terminate quickly and, thus, shed some light on the dependencies among the factors in \mathbf{x} even when a complete analysis is infeasible.

```
R> cna(d.volatile, ordering = list("V02"), maxstep = c(8,10,40),
+     suff.only = TRUE)
```

While the `maxstep` argument is valuable for controlling the search space in case of large and ambiguous data sets, it also comes with a pitfall: it may happen that `cna()` fails to find a model because of a `maxstep` that is too low. An example is `d.jobsecurity`. At the default `maxstep`, `cna()` does not build a solution, but if `maxstep` is increased, two solutions are found.

```
R> ana.jsc.1 <- fscna(d.jobsecurity, ordering = list("JSR"), con=.9, cov=.85)
R> csf(ana.jsc.1)[printCols]
```

```
[1] condition consistency coverage
<0 rows> (or 0-length row.names)
```

```
R> ana.jsc.2 <- fscna(d.jobsecurity, ordering = list("JSR"), con=.9, cov=.85,
+     maxstep = c(3,5,12))
R> csf(ana.jsc.2)[printCols]
```

condition	consistency coverage	
1 S*V + C*1 + L*P + R*V + S*c*R <-> JSR	0.906	0.859
2 C*1 + R*V + P*v + S*c*R + S*C*P <-> JSR	0.912	0.853

In sum, there are two possible reasons for why `cna()` fails to build a solution: (i) the chosen `maxstep` is too low; (ii) the chosen `con` and/or `cov` values are too high, meaning the processed data `x` are too noisy. Accordingly, in case of a null result, two paths should be explored (in that order): (i) gradually increase `maxstep`; (ii) lower `con` and `cov`, as described in section 3.2 above.¹⁰

3.5. Notcols

In classical Boolean logic, the law of Contraposition ensures that an expression of type $\Psi \leftrightarrow Y$ is equivalent to the expression that results from negating both sides of the double arrow: $\neg\Psi \leftrightarrow \neg Y$. Applied to the context of configurational causal modeling that entails that an *asf* for Y can be transformed into an *asf* for the negation of Y , *viz.* y , based on logical principles alone, i.e. without a separate data analysis. However, that transformability only holds for *asf* with perfect consistency and coverage (`con` = `cov` = 1) that are inferred from exhaustive (non-fragmented) data (see section 5.3 for details on exhaustiveness). If an *asf* of an outcome Y does not reach perfect consistency or coverage or is inferred from fragmented data, identifying the causes of y requires a separate application of `cna()` explicitly targeting the causes of the negated outcome.

To this end, the argument `notcols` allows for negating the values of factors in *cs* and *fs* data (in case of *mv* data, `cna()` automatically searches for models of all possible values of endogenous factors, thereby rendering `notcols` redundant). If `notcols` = "all", all factors are negated, i.e. their values i are replaced by $1 - i$. If `notcols` is given a character vector of factors in the data, only the factors in that vector are negated. For example, `notcols` = `c("A", "B")` determines that only factors A and B are negated.

When processing *cs* or *fs* data, CNA should first be applied to recover solutions for the positive outcomes. If resulting *asf* and *csf* do not reach perfect consistency, coverage, and exhaustiveness scores, a second CNA should be run negating the values of all factors that have been modeled as outcomes in the first CNA. To illustrate, we revisit our analysis of `d.autonomy` from section 3.3, which identified AU and EM as outcomes.

```
R> ana.aut.3 <- fscna(dat.aut.1, ordering = list(c("EM", "SP", "CO"), "AU"),
+   strict = FALSE, con = .88, cov = .82, notcols = c("AU", "EM"))
R> csf(ana.aut.3)[printCols]
```

condition	consistency	coverage
1 (sp <-> au)*(sp*co <-> em)	0.882	0.821

4. The CNA algorithm

This section explains the working of the algorithm implemented in the `cna()` function. We first provide an informal summary and then a detailed outline in four stages. The aim of `cna()` is to find all *msc*, *asf*, and *csf* that meet `con.msc`, `con` and `cov` in the input data `x` in

¹⁰For crisp-set and multi-value data the R package `cnaOpt` (Ambuehl and Baumgartner 2019) provides functions that automatically find the optimal consistency and coverage scores obtainable from a given data set (see also Baumgartner and Ambuehl manuscript).

accordance with **ordering** and **maxstep**. The algorithm starts with single factor values and tests whether they meet **con.msc**; if that is not the case, it proceeds to test conjunctions of two factor values, then to conjunctions of three, and so on. Whenever a conjunction meets **con.msc** (and no proper part of it has previously been identified to meet **con.msc**), it is automatically a minimally sufficient condition *msc*, and supersets of it do not need to be tested any more. Then, it tests whether single *msc* meet **con** and **cov**; if not, it proceeds to disjunctions of two, then to disjunctions of three, and so on. Whenever a disjunction meets **con** and **cov** (and no proper part of it has previously been identified to meet **con** and **cov**), it is automatically a minimally necessary disjunction of *msc*, and supersets of it do not need to be tested any more. All and only those disjunctions of *msc* that meet both **con** and **cov** are then issued as *asf*. Finally, recovered *asf* are conjunctively concatenated to *csf* while eliminating structural redundancies and deleting tautologous and contradictory solutions as well as solutions with partial structural redundancies and constant factors.

The **cna()** algorithm can be more specifically broken down into four stages.

Stage 1 On the basis of **ordering**, **cna()** first builds a set of potential outcomes $\mathbf{O} = \{O_h=\omega_f, \dots, O_m=\omega_g\}$ from the set of factors $\mathbf{F} = \{O_1, \dots, O_n\}$ in \mathbf{x} ,¹¹ where $1 \leq h \leq m \leq n$, and second assigns a set of potential cause factors \mathbf{C}_{O_i} from $\mathbf{F} \setminus \{O_i\}$ to every element $O_i=\omega_k$ of \mathbf{O} . If no **ordering** is provided, all value assignments to all elements of \mathbf{F} are treated as possible outcomes in case of *mv* data, whereas in case of *cs* and *fs* data \mathbf{O} is set to $\{O_1=1, \dots, O_n=1\}$.

Stage 2 **cna()** attempts to build a set $\mathbf{msc}_{O_i=\omega_k}$ of minimally sufficient conditions that meet **con.msc** for each $O_i=\omega_k \in \mathbf{O}$. To this end, it first checks for each value assignment $X_h=\chi_j$ of each element of \mathbf{C}_{O_i} , such that $X_h=\chi_j$ has a membership score above 0.5 in at least one case in \mathbf{x} , whether the consistency of $X_h=\chi_j \rightarrow O_i=\omega_k$ meets **con.msc**, i.e. whether $\text{con}(X_h=\chi_j \rightarrow O_i=\omega_k) \geq \text{con.msc}$. If, and only if, that is the case, $X_h=\chi_j$ is put into the set $\mathbf{msc}_{O_i=\omega_k}$. Next, **cna()** checks for each conjunction of two factor values $X_m=\chi_j * X_n=\chi_l$ from \mathbf{C}_{O_i} , such that $X_m=\chi_j * X_n=\chi_l$ has a membership score above 0.5 in at least one case in \mathbf{x} and no part of $X_m=\chi_j * X_n=\chi_l$ is already contained in $\mathbf{msc}_{O_i=\omega_k}$, whether $\text{con}(X_m=\chi_j * X_n=\chi_l \rightarrow O_i=\omega_k) \geq \text{con.msc}$. If, and only if, that is the case, $X_m=\chi_j * X_n=\chi_l$ is put into the set $\mathbf{msc}_{O_i=\omega_k}$. Next, conjunctions of three factor values with no parts already contained in $\mathbf{msc}_{O_i=\omega_k}$ are tested, then conjunctions of four factor values, etc., until either all logically possible conjunctions of the elements of \mathbf{C}_{O_i} have been tested or **maxstep** is reached. Every non-empty $\mathbf{msc}_{O_i=\omega_k}$ is passed on to the third stage.

Stage 3 **cna()** attempts to build a set $\mathbf{asf}_{O_i=\omega_k}$ of atomic solution formulas for every $O_i=\omega_k \in \mathbf{O}$, which has a non-empty $\mathbf{msc}_{O_i=\omega_k}$, by disjunctively concatenating the elements of $\mathbf{msc}_{O_i=\omega_k}$ to minimally necessary conditions of $O_i=\omega_k$ that meet **con** and **cov**. To this end, it first checks for each single condition $\Phi_h \in \mathbf{msc}_{O_i=\omega_k}$ whether $\text{con}(\Phi_h \rightarrow O_i=\omega_k) \geq \text{con}$ and $\text{cov}(\Phi_h \rightarrow O_i=\omega_k) \geq \text{cov}$. If, and only if, that is the case, Φ_h is put into the set $\mathbf{asf}_{O_i=\omega_k}$. Next, **cna()** checks for each disjunction of two conditions $\Phi_m + \Phi_n$ from $\mathbf{msc}_{O_i=\omega_k}$, such that no part of $\Phi_m + \Phi_n$ is already contained in $\mathbf{asf}_{O_i=\omega_k}$, whether $\text{con}(\Phi_m + \Phi_n \rightarrow O_i=\omega_k) \geq \text{con}$ and $\text{cov}(\Phi_m + \Phi_n \rightarrow O_i=\omega_k) \geq \text{cov}$. If, and

¹¹Note that if \mathbf{x} is a data frame, **cna()** first transforms \mathbf{x} into a configuration table using **configTable(x)**, thereby passing the argument **type** (and the two additional arguments **rm.dup.factors** and **rm.const.factors**) to the **configTable()** function.

only if, that is the case, $\Phi_m + \Phi_n$ is put into the set $\mathbf{asf}_{O_i=\omega_k}$. Next, disjunctions of three conditions from $\mathbf{msc}_{O_i=\omega_k}$ with no parts already contained in $\mathbf{asf}_{O_i=\omega_k}$ are tested, then disjunctions of four conditions, etc., until either all logically possible disjunctions of the elements of $\mathbf{msc}_{O_i=\omega_k}$ have been tested or `maxstep` is reached. Every non-empty $\mathbf{asf}_{O_i=\omega_k}$ is passed on to the fourth stage.

Stage 4 `cna()` calls the function `csf()`, which builds a set \mathbf{csf}_O of complex solution formulas. This is done in a stepwise manner as follows. First, all logically possible conjunctions of exactly one element from every non-empty $\mathbf{asf}_{O_i=\omega_k}$ are constructed. Second, if `inus.only = TRUE` (see section 5.2 below), the conjunctions resulting from step 1 are freed of structural redundancies (cf. Baumgartner and Falk 2019), and tautologous and contradictory solutions as well as solutions with partial structural redundancies and constant factors are eliminated. Third, if `acyclic.only = TRUE`, solutions with cyclic substructures are eliminated. Fourth, for those solutions that were modified in the previous steps, consistency and coverage are re-calculated and solutions that no longer reach `con` or `cov` are eliminated. The remaining solutions are returned as \mathbf{csf}_O . If there is only one non-empty set $\mathbf{asf}_{O_i=\omega_k}$, \mathbf{csf}_O is identical to $\mathbf{asf}_{O_i=\omega_k}$.

To illustrate, the following code chunk, first, simulates the data in Table 2c, p. 11, and second, runs `cna()` (and `csf()`) on that data with `con = .8` and `cov = .8`, with default `maxstep`, and without `ordering`.

```
R> dat5 <- allCombs(c(2, 2, 2, 2, 2)) -1
R> dat6 <- selectCases("(A + B <-> C)*(A*B + D <-> E)", dat5)
R> set.seed(3)
R> tab2c <- makeFuzzy(dat6, fuzzvalues = seq(0, 0.4, 0.01))
R> fscna(tab2c, con = .8, cov = .8, what = "mac")
```

Table 2c contains data of type *fs*, meaning that the values in the data matrix are interpreted as membership scores in fuzzy sets. As is customary for this data type, we use uppercase letters for membership in a set and lowercase letters for non-membership. In the absence of an ordering, the set of potential outcomes is determined to be $O = \{A, B, C, D, E\}$ in stage 1, that is, the presence of each factor in Table 2c is treated as a potential outcome. Moreover, all other factors are potential cause factors of every element of O , hence, $C_A = \{B, C, D, E\}$, $C_B = \{A, C, D, E\}$, $C_C = \{A, B, D, E\}$, $C_D = \{A, B, C, E\}$, and $C_E = \{A, B, C, D\}$.

In stage 2, `cna()` succeeds in building non-empty sets of minimally sufficient conditions in compliance with `con` for all elements of O : $\mathbf{msc}_A = \{B*d*E\}$, $\mathbf{msc}_B = \{C*d, d*E, a*C*D, a*C*E, a*C*e\}$, $\mathbf{msc}_C = \{A, B, d*E\}$, $\mathbf{msc}_D = \{b*E, a*E, c*E\}$, $\mathbf{msc}_E = \{D, A*B, A*C\}$. But only the elements of \mathbf{msc}_C and \mathbf{msc}_E can be disjunctively combined to atomic solution formulas that meet `cov` in stage 3: $\mathbf{asf}_C = \{A + B \leftrightarrow C\}$ and $\mathbf{asf}_E = \{D + A*B \leftrightarrow E, D + A*C \leftrightarrow E\}$. For the other three factors in O , the coverage threshold of 0.8 cannot be satisfied. `cna()` therefore abstains from issuing *asf* for A , B and D .

Finally, in stage 4 one redundancy-free *csf* is built from the inventory of *asf* in \mathbf{asf}_C and \mathbf{asf}_E , which constitutes `cna()`'s final output for Table 2c:

$$(A + B \leftrightarrow C) * (D + A*B \leftrightarrow E) \quad con = 0.836; cov = 0.897 \quad (3)$$

5. The output of CNA

5.1. Customizing the output

The default output of `cna()` first lists the provided ordering, second, the *asf* that were recovered in accordance with the ordering, and third, the *csf*. *Asf* and *csf* are ordered by complexity and the product of consistency and coverage. For *asf* and *csf*, four attributes are standardly computed: `consistency`, `coverage`, `complexity`, and `inus`. `consistency` and `coverage` correspond to a solution's consistency and coverage scores, which measures have been explained in section 3.2 above. The `complexity` score amounts to the number of factor values on the left-hand sides of " \rightarrow " or " \leftrightarrow " in *asf* and *csf*; and the `inus` attribute indicates whether a solution has the form of a well-formed MINUS structure (for more on the `inus` attribute cf. section 5.2 below).

`cna()` can compute additional solution attributes, all of which will be explained below: `exhaustiveness`, and `faithfulness` for both *asf* and *csf*, as well as `coherence`, `redundant`, and `cyclic` for *csf*. These attributes are accessible via the `details` argument, which can be given the values `TRUE/FALSE`, for computing all/none of the additional attributes, or a character vector specifying the specific attributes to be computed: for example, `details = c("faithfulness", "exhaustiveness")`—the strings can also be abbreviated, e.g. "f" for "faithfulness", "e" for "exhaustiveness", etc.

```
R> cna(d.educate, details = TRUE)
R> cna(d.educate, details = c("co", "cy"))
```

The output of `cna()` can be further customized through the argument `what` that controls which solution items to print. It can be given a character string specifying the requested solution items: "t" stands for the configuration table, "m" for minimally sufficient conditions (*msc*), "a" for *asf*, "c" for *csf*, and "all" for all solution items.

```
R> cna(d.educate, what = "tm")
R> cna(d.educate, what = "mac")
R> cna(d.educate, what = "all")
```

As shown in section 3.4, it can happen that many *asf* and *csf* fit the data equally well. The standard output of `cna()` only features 5 solution items of each type. To recover all *msc* and *asf* the functions `msc(x)` and `asf(x)` are available, where `x` is a solution object of `cna()`.

```
R> vol2 <- cna(d.volatile, ordering = list("V02"))
R> msc(vol2)
R> asf(vol2)
```

While `msc()` and `asf()` simply access the complete sets of *msc* and *asf* stored in `x`, the *csf* are not stored in `x`. The construction of *csf* in the fourth stage of the CNA algorithm is not conducted by the `cna()` function itself, rather, it is outsourced to the function `csf()` with the these main arguments:

```
csf(x, n.init = 1000, inus.only = x$inus.only,
    minimalizeCsf = inus.only, cyclic.only = x$cyclic.only,
    cycle.type = x$cycle.type, verbose = FALSE)
```

The arguments `inus.only`, `minimalizeCsf`, `cyclic.only`, and `cycle.type` will be further discussed in the following sections, `n.init` and `verbose` are explained in the remainder of this one. It can happen that the set $\mathbf{asf}_{O_i=\omega_k}$ contains too many *asf* to construct all *csf* in reasonable time. The argument `n.init` therefore allows for controlling how many conjunctions of *asf* are initially built in the first step of *csf* construction (see stage 4 of the CNA algorithm); it defaults to 1000. Increasing or lowering that default results in more or less *csf* being built and in longer or shorter computing times, respectively.

```
R> csf(vol2, n.init = 2000)
R> csf(vol2, n.init = 100)
```

Setting the argument `verbose` to its non-default value `TRUE` prints some information about the *csf* construction process to the console, e.g. how many structural redundancies or cyclic substructures have been eliminated along the way.

```
R> csf(vol2, verbose = TRUE)
```

5.2. INUS vs. non-INUS solutions

The (M)INUS-theory of causation (cf. section 2.3) has been developed for strictly Boolean discovery contexts, meaning for deterministic (i.e. noise-free) data that feature perfectly sufficient and necessary conditions. In such contexts, some Boolean expressions can be identified as non-minimal (i.e. as featuring redundant elements) on mere *logical grounds*, that is, independently of data. For instance, in an expression as

$$A + a*B \leftrightarrow C \tag{4}$$

a in the second disjunct is redundant, for (4) is logically equivalent to $A + B \leftrightarrow C$. These two formulas state exactly the same. Under no conceivable circumstances could *a* as contained in (4) ever make a difference to *C*. To see this, note that a necessary condition for *a*B* to be a complex cause of *C* is that there exists a context \mathcal{F} such that *C* is only instantiated when *both* *a* and *B* are given. That means that, in \mathcal{F} , *C* is not instantiated if *B* is given but *a* is not, which, in turn, means that *C* is not instantiated if *B* is given and *A* (*viz.* not-*a*) is given. But such an \mathcal{F} cannot possibly exist, for *A* itself is sufficient for *C* according to (4). It follows that in every context where *B* is instantiated, a change from *A* to *a* is not associated with a change in *C* (which takes the value 1 throughout the change in the factor *A*), meaning that *a* cannot possibly make a difference to *C* and, hence, cannot be a cause of *C*. That is, (4) can be identified as non-minimal independently of all data. (4) is not a well-formed causal model. It is not a MINUS-formula, or not an *INUS solution*.

Correspondingly, a solution as (4) or any other non-INUS expression cannot ever be inferred from strictly deterministic data. By contrast, it is possible that a necessary disjunction of sufficient conditions that does *not* have INUS form is redundancy-free (or minimal) relative to indeterministic data. Hence, the solution attribute `inus` makes explicit whether an *asf* or *csf* is an INUS solution. Moreover, the `cna()` and the `csf()` functions both have an argument `inus.only` controlling whether only INUS solutions (MINUS-formulas) shall be built or whether non-INUS solutions shall be issued as well (if they are inferrable from data). As `inus.only` defaults to `TRUE`, standard calls of `cna()` and `csf()` will never yield non-INUS

solutions, even if such solutions might be inferable from data. But if the analyst is interested in non-(M)INUS structures as well, she can request corresponding solutions by `inus.only = FALSE`.

To illustrate this, we first show that CNA never infers non-(M)INUS solutions from deterministic data—even if `inus.only = FALSE`. For this purpose, we simulate ideal data from the non-INUS solution in (4); `cna(..., inus.only = FALSE)` will always, i.e. upon an open number of re-runs of the following code chunk, return $A + B \leftrightarrow C$, regardless of the fact that we select cases based on (4) with `selectCases1("A + a*B <-> C", ...)`.

```
R> dat.inu.1 <- allCombs(c(2, 2, 2)) -1
R> dat.inu.2 <- some(dat.inu.1, 40, replace = TRUE)
R> dat.inu.3 <- selectCases1("A + a*B <-> C", con = 1, cov = 1, dat.inu.2)
R> asf(cna(dat.inu.3, con = 1, cov = 1, inus.only = FALSE))
```

	outcome	condition	consistency	coverage	complexity	inus
1	C	A + B <-> C	1	1	2	TRUE

But in real-life discovery contexts, especially in observational studies, deterministic dependencies are the exception rather than the norm. Ordinary (observational) data are indeterministic, meaning that causes tend to be combined both with the presence and the absence of outcomes. In such discovery contexts, which can be simulated by lowering `con` and `cov` in `selectCases1()`, non-INUS expressions may count as minimally necessary disjunctions of minimally sufficient conditions. Correspondingly, if `inus.only` is set to `FALSE`, `cna()` may infer non-INUS solutions:

```
R> set.seed(26)
R> dat.inu.4 <- some(dat.inu.1, 40, replace = TRUE)
R> dat.inu.5 <- selectCases1("A + a*B <-> C", con = .8, cov = .8, dat.inu.4)
R> asf(cna(dat.inu.5, con = .8, cov = .8, inus.only = FALSE))
```

	outcome	condition	consistency	coverage	complexity	inus
1	C	A + a*B <-> C	0.826	0.826	3	FALSE

In indeterministic data, it can happen that a is needed to lift the consistency of B above the chosen `con` threshold. In such a case, a can be argued to make a difference to C : only in conjunction with a does B reach `con`; and this holds notwithstanding the fact that A itself also meets `con`. Or put differently, when the consistency of $A \rightarrow C$ is below 1, there exist cases where A is instantiated and C is not, which, in turn, yields that it becomes possible for a change from A to a , while B is constantly instantiated, to be associated with a change in C , meaning that a can turn out to be a difference-maker for C .

If a necessary disjunction of sufficient conditions turns out to be redundancy-free relative to indeterministic data even though it does not have (M)INUS form, the crucial follow-up question is whether the indeterminism in the data is due to insufficient control of background influences (i.e. to noise, measurement error, etc.) or to the inherent indeterministic nature of the physical processes themselves (as can e.g. be found in the domain of quantum mechanics, cf. Albert 1992). If the latter is the case, the difference-making relations stipulated by non-INUS solutions should be taken seriously and, hence, the argument `inus.only` should be set to

its non-default value `FALSE`. By contrast, if the former is the case (i.e. the indeterminism is due to noise), the difference-making relations entailed by non-INUS solutions are mere artifacts of the noise in the data, meaning that they would disappear if the corresponding causal structure were investigated under more ideal discovery circumstances. In that case, which obtains in the macro domains to which CNA is typically applied, the argument `inus.only` should be left at its default value `TRUE` both in `cna()` and `csf()`. If `inus.only` is switched to `TRUE` in the `cna()`-call from the previous code chunk, no solution is returned any more, that is, there does not exist a MINUS-formula satisfying `con` and `cov` for `dat.inu.5`:

```
R> asf(cna(dat.inu.5, con = .8, cov = .8, inus.only = TRUE))

[1] outcome      condition      consistency coverage      complexity
[6] inus
<0 rows> (or 0-length row.names)
```

The function behind the solution attribute `inus` and the argument `inus.only` is also available as stand-alone function `is.inus()`. Logical redundancies as contained in non-INUS solutions can be eliminated by means of the function `minimalize()` (see the [cna reference manual](#) for details).

5.3. Exhaustiveness and faithfulness

Exhaustiveness and faithfulness are two measures of model fit that quantify the degree of correspondence between the configurations that are, in principle, compatible with a solution and the configurations actually contained in the data from which that solution is derived. Exhaustiveness is high when *all* or most configurations *compatible* with a solution are in the data. More specifically, it amounts to the ratio of the number of configurations in the data that are compatible with a solution to the number of configurations in total that are compatible with a solution. To illustrate, consider `d.educate`, which contains all configurations that are compatible with the two `csf` issued by `cna()` and `csf()`:

```
R> printCols <- c("condition", "consistency", "coverage", "exhaustiveness")
R> csf(cna(d.educate, details = "exhaust"))[printCols]
```

	condition	consistency	coverage	exhaustiveness
1	(L + G <-> E)*(U + D <-> L)	1	1	1
2	(U + D + G <-> E)*(U + D <-> L)	1	1	1

If, say, the first configuration in `d.educate` (*viz.* $U*D*L*G*E$) is not observed or removed—as in `d.educate[-1,]`—, CNA still builds the same solutions (with perfect consistency and coverage). In that case, however, the resulting `csf` are not exhaustively represented in the data, for one configuration that is compatible with both `csf` is not contained therein.

```
R> csf(cna(d.educate[-1,], details = "exhaust"))[printCols]
```

	condition	consistency	coverage	exhaustiveness
1	(L + G <-> E)*(U + D <-> L)	1	1	0.875
2	(U + D + G <-> E)*(U + D <-> L)	1	1	0.875

In a sense, faithfulness is the complement of exhaustiveness. It is high when *no* or only few configurations that are *incompatible* with a solution are in the data. More specifically, faithfulness amounts to the ratio of the number of configurations in the data that are compatible with a solution to the total number of configurations in the data. The two *csf* resulting from `d.educate` also reach perfect faithfulness:

```
R> printCols <- c("condition", "consistency", "coverage", "faithfulness")
R> csf(cna(d.educate, details = "faithful"))[printCols]
```

	condition	consistency	coverage	faithfulness
1	(L + G <-> E)*(U + D <-> L)	1	1	1
2	(U + D + G <-> E)*(U + D <-> L)	1	1	1

If we add a configuration that is not compatible with these *csf*, say, $U*D*I*G*e$ and lower the consistency threshold, the same solutions along with two others result—this time, however, with non-perfect faithfulness scores.

```
R> csf(cna(rbind(d.educate,c(1,1,0,1,0)), con = .8, details = "f"))[printCols]
```

	condition	consistency	coverage	faithfulness
1	(L + G <-> E)*(E <-> L)	0.857	1	0.778
2	(L + G <-> E)*(U + D <-> L)	0.857	1	0.889
3	(U + D + G <-> E)*(E <-> L)	0.857	1	0.778
4	(U + D + G <-> E)*(U + D <-> L)	0.857	1	0.889

If both exhaustiveness and faithfulness are high, the configurations in the data are all and only the configurations that are compatible with the solution. Low exhaustiveness and/or faithfulness, by contrast, means that the data do not contain many configurations compatible with the solution and/or the data contain many configurations not compatible with the solution. In general, solutions with higher exhaustiveness and faithfulness scores are preferable over solutions with lower scores.

5.4. Coherence

Coherence is a measure for model fit that is custom-built for *csf*. It measures the degree to which the *asf* combined in a *csf* cohere, that is, are instantiated together in the data rather than independently of one another. Coherence is intended to capture the following intuition. Suppose a *csf* entails that A is a sufficient cause of B , which, in turn, is entailed to be a sufficient cause of C . Corresponding data δ should be such that the $A - B$ link of that causal chain and the $B - C$ link are either both instantiated or both not instantiated in the cases recorded in δ . By contrast, a case in δ such that, say, only the $A - B$ link is instantiated but the $B - C$ link is not, pulls down the coherence of that *csf*. The more such non-cohering cases are contained in δ , the lower the overall coherence score of the *csf*.

Coherence is more specifically defined as the ratio of the number of cases satisfying all *asf* contained in a *csf* to the number of cases satisfying at least one *asf* in the *csf*. More formally, let a *csf* contain $asf_1, asf_2, \dots, asf_n$, coherence then amounts to (where $|\dots|_\delta$ represents the

cardinality of the set of cases in δ instantiating the corresponding expression):

$$\frac{|asf_1 * asf_2 * \dots * asf_n|_\delta}{|asf_1 + asf_2 + \dots + asf_n|_\delta}$$

To illustrate, we add a case of type $U*d*L*g*e$ to `d.educate`. When applied to the resulting data (`d.edu.exp1`), `cna()` and `csf()` issue three *csf*.

```
R> d.edu.exp1 <- rbind(d.educate, c(1,0,1,0,0))
R> printCols <- c("condition", "consistency", "coverage", "coherence")
R> csf(cna(d.edu.exp1, con = .8, details = "cohere"))[printCols]
```

condition	consistency	coverage	coherence
1 (L + G <-> E)*(U + D <-> L)	0.875	1	0.889
2 (L + G <-> E)*(U + E <-> L)	0.875	1	0.778
3 (U + D + G <-> E)*(U + D <-> L)	0.875	1	0.889

In the added case, none of these three *csf* cohere, as only one of their component *asf* is instantiated. Moreover, for the second *csf* there is yet another non-cohering case in `d.edu.exp1` (case #7).

Coherence is an additional parameter of model fit that allows for selecting among multiple solutions: the higher the coherence score of a *csf*, the better the overall model fit. In `d.edu.exp1`, thus, the first and the third *csf* are preferable over the second solution subject to their superior coherence.

5.5. Structural redundancies and partial structural redundancies

It is not only possible that Boolean expressions describing the behavior of single outcomes contain redundant proper parts, but such expressions can themselves—as a whole—be redundant in superordinate structures. For instance, when three *asf* are conjunctively concatenated to $asf_1 * asf_2 * asf_3$ in the first step of stage 4 of the CNA algorithm (see section 4), it can happen that $asf_1 * asf_2 * asf_3$ is logically equivalent to $asf_1 * asf_2$, meaning that asf_3 makes no difference to accounting for the behavior of the outcomes in that structure and is, thus, redundant. We speak of a *structural redundancy* in that context (for a detailed discussion see Baumgartner and Falk 2019).

This type of redundancy is best introduced with a concrete example. Consider the following complex MINUS-formula:

$$(A*B + C \leftrightarrow D) * (a + c \leftrightarrow E) \tag{5}$$

(5) represents a causal structure such that $A*B$ and C are two alternative causes of D and a and c are two alternative causes of E . That is, factors A and C are positively relevant to D and negatively relevant to E . A possible interpretation of these factors might be the following. Suppose a city has two power stations: a wind farm and a nuclear plant. Let A express that the wind farm is operational and C that the nuclear plant is operational and let operationality be sufficient for a nuclear plant to produce electricity, while a wind farm produces electricity provided it is operational and there is wind (B). Hence, the wind farm

being operational while it is windy or the nuclear plant being operational ($A*B + C$) are two alternative causes of the city being power supplied (D). Whereas the wind farm or the nuclear plant not being operational ($a + c$) are two alternative causes of an alarm being triggered (E). The following data (`dat.redun`) comprise all and only the configurations that are compatible with (5):

```
R> (dat.redun <- ct2df(selectCases("(A*B + C <-> D)*(a + c <-> E)")))
```

```

  A B C D E
2  0 1 1 1 1
4  0 0 1 1 1
5  1 1 0 1 1
14 0 1 0 0 1
15 1 0 0 0 1
16 0 0 0 0 1
17 1 1 1 1 0
19 1 0 1 1 0

```

The problem now is that `dat.redun` does not only entail the two *asf* contained in (5), *viz.* (6) and (7), but also a third one, *viz.* (8):

$$A*B + C \leftrightarrow D \quad (6)$$

$$a + c \leftrightarrow E \quad (7)$$

$$a*D + e \leftrightarrow C \quad (8)$$

That means the behavior of factor C , which is exogenous in the data-generating structure (5), can be expressed as a redundancy-free Boolean function of its two effects D and E . (8), hence, amounts to an upstream (or backtracking) *asf*, which, obviously, must not be causally interpreted. Indeed, when (8) is embedded in the superordinate dependency structure (9) that results from a conjunctive concatenation of all *asf* that follow from `dat.redun`, it turns out that (8) is redundant. The reason is that (9) has a proper part which is logically equivalent to (9), namely (5).

$$(A*B + C \leftrightarrow D) * (a + c \leftrightarrow E) * (a*D + e \leftrightarrow C) \quad (9)$$

(9) and (5) state exactly the same about the behavior of the factors in `dat.redun`, meaning that (8) makes no difference to that behavior over and above (6) and (7). By contrast, neither (6) nor (7) can be eliminated from (9) such that the remaining expression is logically equivalent to (9). Both of these downstream *asf* make their own distinctive difference to the behavior of the factors in `dat.redun`. The upstream *asf* (8), however, is a *structural redundancy* in (9). (9) must not be causally interpreted because it is not a complex MINUS-formula (see pp. 6-7 above).

Accordingly, in its default setting, the `csf()` function, which performs stage 4 of the CNA algorithm, removes all structurally redundant *asf* from conjunctions of *asf*; that is, when applied to `dat.redun` it returns (5), not (9):

```
R> printCols <- c("condition", "consistency", "coverage", "inus", "redundant")
R> csf(cna(dat.redun, details = "r"))[printCols]
```

```

condition                                consistency coverage inus redundant
1 (C + A*B <-> D)*(a + c <-> E)          1          1 TRUE     FALSE

```

In previous versions (< 3.0) of the **cna** package, structurally redundant *asf* were not automatically removed but only marked by means of the solution attribute `redundant`. The solutions with `redundant = TRUE` then had to be further processed by the function `minimalizeCsf()`. To reproduce that old behavior, `csf()` now has an additional argument `minimalizeCsf`, which defaults to `TRUE`. If set to `FALSE`, structural redundancies are not automatically eliminated. Accordingly, the following call returns (9) and marks it as containing a structural redundancy—and thus as not having (M)INUS form:

```

R> csf(cna(dat.redun, details = "r"), inus.only = FALSE,
+      minimalizeCsf = FALSE)[printCols]

```

```

condition                                consistency coverage
1 (e + a*D <-> C)*(C + A*B <-> D)*(a + c <-> E)          1          1
  inus redundant
1 FALSE          TRUE

```

As *csf* with `redundant = TRUE` must never be causally interpreted, the setting `minimalizeCsf = FALSE` is deprecated. It is mainly kept in the package for backwards compatibility and developing purposes. Correspondingly, the solution attribute `redundant` is no longer relevant, as `cna()` and `csf()` no longer output *csf* with structural redundancies in the first place.

While structural redundancies of whole *asf* can occur in both deterministic and indeterministic data, the latter type of data may induce yet another, but related, type of redundancy. When data do not feature strict Boolean dependencies, building *csf* from the inventory of *asf* recovered in stage 3 of the CNA algorithm may lead to the redundancy of proper parts of *asf*—parts which are not redundant when those *asf* are considered in isolation. That is, a complex structure can entail that one of its *asf* has a redundant proper part, which redundancy, however, is not visible in the data. We call this a *partial structural redundancy*. Again, a concrete example helps to clarify the problem. Consider solution #6 resulting from the following analysis of data `d.autonomy`:

```

R> printCols <- c("condition", "consistency", "coverage", "inus")
R> csf(fscna(d.autonomy, ordering = list("AU"), con = .9, cov = .94,
+      maxstep = c(2, 2, 8), inus.only = FALSE))[printCols]

```

```

condition                                consistency coverage inus
1 (SP*RE + CO*cn <-> EM)*(ci + EM*co <-> SP)          0.938    0.947 TRUE
2 (SP*co + CO*cn <-> EM)*(ci + EM*co <-> SP)          0.938    0.947 TRUE
3 (SP*RE + ci*cn <-> EM)*(ci + EM*co <-> SP)          0.938    0.947 TRUE
4 (SP*RE + CO*cn <-> EM)*(ci + EM*RE <-> SP)          0.928    0.947 TRUE
5 (SP*co + CO*cn <-> EM)*(ci + EM*RE <-> SP)          0.928    0.947 TRUE
6 (SP*RE + ci*cn <-> EM)*(ci + EM*RE <-> SP)          0.928    0.947 FALSE

```

Solution #6 is logically equivalent to (10):

$$(SP*RE + ci*cn \leftrightarrow EM)*(ci + EM \leftrightarrow SP) \quad (10)$$

That is, if the behavior of EM is regulated by the first asf in (10), RE in solution #6—for pure logical reasons—cannot make a difference to SP and, hence, is redundant. That partial structural redundancy, however, is not visible in the data `d.autonomy` where EM alone (i.e. without RE) is not sufficient for SP with consistency 0.9, which is the threshold chosen for the above analysis:

```
R> condTbl("EM -> SP", fsct(d.autonomy))

  outcome condition consistency coverage complexity
1 SP          EM -> SP      0.891    0.863          1
```

Hence, the data suggest that RE makes a difference to SP , to the effect that meeting $con = 0.9$ for all msc requires $EM*RE$ (and not EM alone) to be treated as cause of SP . At the same time, the equivalence of solution #6 and (10) entails it to be logically excluded that RE is a difference-maker of SP in the context of solution #6. The result is a contradiction: the data call for including RE as cause of SP , whereas the structure inferred from that data entails not to include RE . That cannot be resolved by modifying solution #6; rather, that solution is not a well-formed MINUS-formula and cannot be turned into one that would meet con . It must be eliminated from the output. This is exactly what happens if `cna()` and `csf()` are run with the default `inus.only = TRUE`:

```
R> csf(fscna(d.autonomy, ordering = list("AU"), con = .9, cov = .94,
+       maxstep = c(2, 2, 8), inus.only = TRUE))[printCols]

  condition                                consistency coverage inus
1 (SP*RE + CO*cn <-> EM)*(ci + EM*co <-> SP)      0.938    0.947 TRUE
2 (SP*co + CO*cn <-> EM)*(ci + EM*co <-> SP)      0.938    0.947 TRUE
3 (SP*RE + ci*cn <-> EM)*(ci + EM*co <-> SP)      0.938    0.947 TRUE
4 (SP*RE + CO*cn <-> EM)*(ci + EM*RE <-> SP)      0.928    0.947 TRUE
5 (SP*co + CO*cn <-> EM)*(ci + EM*RE <-> SP)      0.928    0.947 TRUE
```

In sum, as of version 3.0 of the `cna` package, both structural and partial structural redundancies are automatically resolved and eliminated. The functions `cna()` and `csf()` now exclusively output MINUS-formulas (i.e. INUS solutions).

5.6. Cycles

Detecting causal cycles is one of the most challenging tasks in causal data analysis—in all methodological traditions. In a nutshell, the reason is that factors in a cyclic structure are so highly interdependent that, even under optimal discovery conditions, the diversity of (observational) data tends to be too limited to draw informative conclusions about the data-generating structure. In consequence, various methods (most notably, Bayes nets methods, cf. [Spirtes et al. 2000](#)) assume that analyzed data-generating structures are acyclic.

`cna()` and `csf()` output cyclic *csf* if they fit the data, and the optional solution attribute `cyclic` identifies those *csf* that contain a cyclic substructure. A causal structure has a cyclic substructure if, and only if, it contains a directed causal path from at least one cause back to itself. The MINUS theory spells this criterion out more explicitly as follows:

Cycle A complex MINUS-formula \mathbf{m} has a cyclic substructure if, and only if, \mathbf{m} contains a sequence $\langle Z_1, Z_2, \dots, Z_n \rangle$ such that every Z_i is contained in an atomic MINUS-formula of Z_{i+1} and $Z_1 = Z_n$ in \mathbf{m} .

To illustrate with a simple example, (11) contains the cyclic sequence $\langle A, B, A \rangle$ and, thus, represents a causal cycle:

$$(A*c + D*E \leftrightarrow B)*(B*F + G*h \leftrightarrow A) \quad (11)$$

Typically, when cyclic models fit the data, the output of `cna()` and `csf()` is massively ambiguous. Therefore, if there are independent reasons to assume that the data are not generated by a cyclic structure, both `cna()` and `csf()` have the argument `acyclic.only`, which, if set to its non-default value `TRUE`, prevents solutions with cycles from being returned and, thereby, reduces model ambiguities. For example, by switching `acyclic.only` from `FALSE` to `TRUE` in the following analysis, the solution space is reduced from 4 to 1:

```
R> printCols <- c("condition", "cyclic")
R> csf(fscna(d.pacts, con = .8, cov = .8, acyclic.only = FALSE,
+         details = "cy"))[printCols]

      condition                                cyclic
1 (E + p <-> PACT)*(E*p + PACT*1*P <-> W)      FALSE
2 (w + p <-> E)*(W + l + P <-> PACT)*(e + PACT*L*p <-> W)    TRUE
3 (W + PACT*L <-> E)*(E + p <-> PACT)*(E*p + PACT*1*P <-> W)    TRUE
4 (PACT*L + PACT*P <-> E)*(E + p <-> PACT)*(E*p + PACT*1*P <-> W)    TRUE

R> csf(fscna(d.pacts, con = .8, cov = .8, acyclic.only = TRUE,
+         details = "cy"))[printCols]

      condition                                cyclic
1 (E + p <-> PACT)*(E*p + PACT*1*P <-> W)      FALSE
```

The `cycle.type` argument—also available in both `cna()` and `csf()`—controls whether a cyclic sequence $\langle Z_1, Z_2, \dots, Z_n \rangle$ is composed of factors (`cycle.type = "factor"`), which is the default, or factor values (`cycle.type = "value"`). To illustrate, if `cycle.type = "factor"`, (12) counts as cyclic:

$$(A + B \leftrightarrow C) * (c + D \leftrightarrow A) \quad (12)$$

The factor A (with value 1) appears in an *asf* of C (with value 1), and the factor C (with value 0) appears in an *asf* of A (with value 1). But if `cycle.type = "value"`, (12) does not pass as cyclic. Although the factor value 1 of A appears in an *asf* of the factor value 1 of C ,

that same value of C does not appear in an *asf* of A ; rather, the value 0 of C appears in the *asf* of A .

The function behind the solution attribute `cyclic` and the corresponding arguments is also available as stand-alone function `cyclic()` (see the [cna reference manual](#) for details).

6. Interpreting the output

The ultimate output of `cna()` and `csf()` is a set \mathbf{csf}_O of *csf*—which may be identical to *asf*, if the data comprise only one endogenous factor. The causal inferences that are warranted based on the data input \mathbf{x} relative to the chosen `con` and `cov` thresholds and the provided `ordering` and `maxstep` have to be read off that set \mathbf{csf}_O . This section explains this final interpretative step of a CNA analysis.

There are three possible types of outputs:

1. \mathbf{csf}_O contains no *csf* (and, correspondingly, no *asf*);
2. \mathbf{csf}_O contains exactly one *csf* (and, correspondingly, exactly one *asf* for each endogenous factor);
3. \mathbf{csf}_O contains more than one *csf* (and, correspondingly, more than one *asf* for at least one endogenous factor).

6.1. No solution

As indicated in section 3.4, a null result can have two sources: either the data are too noisy to render the chosen `con` and `cov` thresholds satisfiable or the selected `maxstep` is too low. If increasing `maxstep` does not yield solutions at the chosen `con` and `cov` thresholds, the latter may be lowered, preferably with a concomitant robustness analysis as described in section 3.2. If no solutions are recovered at `con = cov = .7`, the data are too noisy to warrant reliable causal inferences. Users are then advised to go back to the data and follow standard guidelines (known from other methodological frameworks) to improve data quality, e.g. by integrating further relevant factors into the analysis, enhancing the control of unmeasured causes, expanding the population of cases or disregarding inhomogeneous cases, correcting for measurement error, supplying missing values, etc.

It must be emphasized again (see section 2.4) that, under normal circumstances, an empty \mathbf{csf}_O does not warrant the conclusion that the factors contained in the data input \mathbf{x} are causally irrelevant to one another. The inference to causal irrelevance is much more demanding than the inference to causal relevance. A null result only furnishes evidence for causal irrelevance if there are independent reasons to assume that all potentially relevant factors are measured in \mathbf{x} and that \mathbf{x} exhausts the space of empirically possible configurations.

6.2. A unique solution

That \mathbf{csf}_O contains exactly one *csf* is the optimal completion of a CNA analysis. It means that the data input \mathbf{x} contains sufficient evidence for a determinate causal inference. The factor values on the left-hand sides of “ \leftrightarrow ” in the *asf* constituting that *csf* can be interpreted

as causes of the factor values on the right-hand sides. Moreover, their conjunctive, disjunctive, and sequential groupings reflect the actual properties of the data-generating causal structure. Plainly, as with any other method of causal inference, the reliability of CNA's causal conclusions essentially hinges on the quality of the processed data. If the data satisfy homogeneity (see section 2.4), a unique solution is guaranteed to correctly reflect the data-generating structure. With increasing data deficiencies (noise, measurement error, etc.), the (inductive) risk of committing causal fallacies inevitably increases as well. For details on the degree to which the reliability of CNA's causal conclusions decreases with increasing data deficiencies see Baumgartner and Ambuehl (2020) and Parkkinen and Baumgartner (forthcoming).

6.3. Multiple solutions

If $\mathbf{csf}_{\mathbf{O}}$ has more than one element, the processed data underdetermine their own causal modeling. That means the evidence contained in the data is insufficient to determine which of the solutions contained in $\mathbf{csf}_{\mathbf{O}}$ corresponds to the data-generating causal structure. An output set of multiple solutions $\{csf_1, csf_2, \dots, csf_n\}$ is to be interpreted *disjunctively*: the data-generating causal structure is

$$csf_1 \text{ OR } csf_2 \text{ OR } \dots \text{ OR } csf_n$$

but, based on the evidence contained in the data, it is ambiguous which disjunct is actually operative.

That empirical data underdetermine their own causal modeling is a very common phenomenon in all methodological traditions (Simon 1954; Spirtes *et al.* 2000, 59-72; Kalisch *et al.* 2012; Eberhardt 2013; Baumgartner and Thiem 2017). But while some methods are designed to automatically generate all fitting models, e.g. Bayes nets methods and configurational comparative methods, other methods rely on search heuristics that zoom in on one best fitting model only, e.g. logic regression or regression analytic methods, more generally. Whereas model ambiguities are a thoroughly investigated topic in certain traditions, e.g. Bayes nets methods, they are only beginning to be studied in the literature on configurational comparative methods, in particular on QCA. There is still an unfortunate practice of model-underreporting in QCA studies. In fact, most QCA software regularly fails to find all data-fitting models. The only currently available QCA program that recovers the whole model space by default is **QC Apro** (Thiem 2018).

CNA—on a par with any other method—cannot disambiguate what is empirically underdetermined. Rather, it draws those and only those causal conclusions for which the data *de facto* contain evidence. In cases of empirical underdetermination it therefore renders transparent all data-fitting models and leaves the disambiguation up to the analyst.

That `cna()` and `csf()` issue multiple solutions for some data input \mathbf{x} does not necessarily mean that \mathbf{x} is deficient. In fact, even data that is *ideal* by all quality standards of configurational causal modeling can give rise to model ambiguities. The following simulates a case in point:

```
R> dat7 <- selectCases("a*B + A*b + B*C <-> D")
R> printCols <- c("condition", "consistency", "coverage", "inus",
+ "exhaustiveness")
R> csf(cna(dat7, details = c("exhaust", "inus")))[printCols]
```

	condition	consistency	coverage	inus	exhaustiveness
1	$a*B + A*b + A*C \leftrightarrow D$	1	1	TRUE	1
2	$a*B + A*b + B*C \leftrightarrow D$	1	1	TRUE	1

`dat7` induces perfect consistency and coverage scores and is free of fragmentation; it contains all and only the configurations that are compatible with the target structure, which accordingly is exhaustively and faithfully reflected in `dat7`. Nonetheless, two *csf* can be inferred. The causal structures expressed by these two *csf* generate the exact same data, meaning they are *empirically indistinguishable*.

Although, a unique solution is more determinate and, thus, preferable to multiple solutions, the fact that `cna()` and `csf()` generate multiple equally data-fitting models is not generally an uninformative result. In the above example, both resulting *csf* feature $a*B + A*b$. That is, the data contain enough evidence to establish the joint relevance of $a*B$ and of $A*b$ for D (on alternative paths). What is more, it can be conclusively inferred that D has a further complex cause, *viz.* either $A*C$ or $B*C$. It is merely an open question which of these candidate causes is actually operative.

That different model candidates have some *msc* in common is a frequent phenomenon. Here's a real-life example, where two alternative causes, *viz.* $C=1 + F=2$, are present in all solutions:

```
R> csf(mvcna(d.pban, cov = .95, maxstep = c(3, 5, 10)))["condition"]
```

```
condition
1 C=1 + F=2 + C=0*F=1 + C=2*V=0 <-> PB=1
2 C=1 + F=2 + C=0*T=2 + C=2*V=0 <-> PB=1
3 C=1 + F=2 + C=2*F=0 + C=0*F=1 + F=1*V=0 <-> PB=1
4 C=1 + F=2 + C=2*F=0 + C=0*T=2 + F=1*V=0 <-> PB=1
5 C=1 + F=2 + C=0*F=1 + C=2*T=1 + T=2*V=0 <-> PB=1
6 C=1 + F=2 + C=0*F=1 + T=1*V=0 + T=2*V=0 <-> PB=1
7 C=1 + F=2 + C=0*T=2 + C=2*T=1 + T=2*V=0 <-> PB=1
8 C=1 + F=2 + C=0*T=2 + T=1*V=0 + T=2*V=0 <-> PB=1
```

Such commonalities can be reported as conclusive results.

Moreover, even though multiple solutions do not permit pinpointing the causal structure behind an outcome, they nonetheless allow for constraining the range of possibilities. In a context where the causes of some outcome are unknown it amounts to a significant gain of scientific insight when a study can show that the structure behind that outcome has one of a small number of possible forms, even if it cannot determine which one exactly.

However, the larger the amount of data-fitting solutions and the lower the amount of commonalities among them, the lower the overall informativeness of a CNA output. Indeed, the ambiguity ratio in configurational causal modeling can reach dimensions where nothing at all can be concluded about the data-generating structure any more. Hence, a highly ambiguous result is on a par with a null result. A telling example of this sort is `d.volatile` which was discussed in section 3.4 above (cf. also Baumgartner and Thiem 2017).

As the problem of model ambiguities is still underinvestigated in the CNA literature, there do not yet exist explicit guidelines for how to proceed in cases of ambiguities. The model fit scores and solution attributes reported in the output objects of `cna()` and `csf()` often provide

some leverage to narrow down the space of model candidates. For instance, if, in a particular discovery context, there is reason to assume that data have been collected as exhaustively as possible, to the effect that most configurations compatible with an investigated causal structure should be contained in the data, the model space may be restricted to *csf* with a high score on exhaustiveness. By way of example, for *d.pacts* a total of 10 *csf* are built at `con = cov = .8`:

```
R> ana.pact.1 <- fscna(d.pacts, ordering = list("PACT"), con = .8, cov = .8,
+   maxstep = c(4, 5, 15), details = TRUE)
R> csf.pact.1 <- csf(ana.pact.1, Inf)
R> length(csf.pact.1$condition)
```

```
[1] 10
```

If only *csf* with `exhaustiveness >= .85` are considered, the amount of candidate *csf* is reduced to 6:

```
R> csf.pact.1.ex <- subset(csf.pact.1, exhaustiveness >= .85)
R> length(csf.pact.1.ex$condition)
```

```
[1] 6
```

To also resolve this final ambiguity, coherence may be brought to bear. The higher the coherence of a *csf*, the higher its overall model fit. In the above example, if coherence is required to be as high as possible, only one *csf* remains:

```
R> csf.pact.1.ex.co <- subset(csf.pact.1.ex,
+   coherence == max(csf.pact.1.ex$coherence))
R> length(csf.pact.1.ex.co$condition)
```

```
[1] 1
```

Clearly though, the fit parameters and solution attributes provided by `cna()` and `csf()` will not always provide a basis for ambiguity reduction. The evidence contained in data may simply be insufficient to draw determinate causal conclusions. Maybe background theories or case knowledge can be brought to bear to select among the model candidates (see section 6.4). Nevertheless, the most important course of action in the face of ambiguities is to *render them transparent*. By default, readers of CNA publications should be presented with all data-fitting models and if space constraints do not permit so, readers must at least be informed about the degree of ambiguity. Full transparency with respect to model ambiguities, first, allows readers to determine for themselves how much confidence to have in the conclusions drawn in a study, and second, paves the way for follow-up studies that are purposefully designed to resolve previously encountered ambiguities.

6.4. “Back to the cases”

When CNA is applied to small- or intermediate-*N* data, researchers may be familiar with some or all of the cases in their data. For instance, they may know that in a particular case

certain causes of an outcome are operative while others are not. Or they may know why certain cases are outliers or why others feature an outcome but none of the potential causes. A proper interpretation of a CNA result may therefore require that the performance of the obtained models be assessed on the case level and against the background of the available case knowledge.

The function that facilitates the evaluation of recovered *msc*, *asf*, and *csf* on the case level is `condition(x, ct, type)`. Its first input is a character vector `x` specifying Boolean expressions—typically *asf* or *csf*—and its second input a configuration table `ct`. In case of *cs* or *mv* data, the output of `condition()` then highlights in which cases `x` is instantiated, whereas for *fs* data, the output lists relevant membership scores in exogenous and endogenous factors. Moreover, if `x` is an *asf* or *csf*, `condition()` issues their consistency and coverage scores.

Instead of a configuration table, it is also possible to give `condition()` a data frame as second input. In this case, the data type must be specified using the `type` argument. To abbreviate the specification of the data type, the functions `cscond(x, ct)`, `mvcond(x, ct)`, and `fscond(x, ct)` are available as shorthands.

To illustrate, we re-analyze `d.autonomy`:

```
R> dat.aut.2 <- d.autonomy[15:30, c("AU", "EM", "SP", "CO", "RE", "DE")]
R> ana.aut.3 <- fscna(dat.aut.2, con = .91, cov = .91,
+   ordering = list(c("RE", "DE", "SP", "CO"), "EM", "AU"),
+   strict = TRUE)
R> fscond(csf(ana.aut.3)$condition, dat.aut.2)
```

That function call returns a list of three tables, each corresponding to one of the three *csf* contained in `ana.aut.3` and breaking down the relevant *csf* to the case level by contrasting the membership scores in the left-hand and right-hand sides of the component *asf*. A case with a higher left-hand score is one that pulls down consistency, whereas a case with a higher right-hand score pulls down coverage. For each *csf*, `condition()` moreover returns overall consistency and coverage scores as well as consistency and coverage scores for the component *asf*.

The three *csf* in `ana.aut.3` differ only in regard to their component *asf* for outcome *AU*. The function `group.by.outcome(condlst)`, which takes an output object `condlst` of `condition()` as input, lets us more specifically compare these different *asf* with respect to how they fare on the case level.

```
R> group.by.outcome(fscond(asf(ana.aut.3)$condition, dat.aut.2))$AU
```

	SP	EM*RE+re*DE	EM*RE+CO*DE	AU		n.obs
ENacg1	1.0	1.0	1.0	1.0		1
ENacg2	0.6	0.4	0.4	0.4		1
ENacg3	0.8	0.6	0.9	0.8		1
ENacg4	0.6	1.0	1.0	1.0		1
ENacg5	0.4	0.4	0.4	0.4		1
ENacg6	0.6	0.7	0.7	0.6		1
ENacg7	1.0	0.8	0.8	1.0		1
ENacg8	1.0	1.0	1.0	1.0		1

ENacto1	0.4	0.4	0.6	0.4		1
ENacto2	0.4	0.4	0.4	0.4		1
ENacosa1	0.4	0.4	0.2	0.4		1
ENacosa2	0.4	0.4	0.4	0.2		1
ENacosa3	0.4	0.4	0.4	0.6		1
ENacat1	0.4	0.4	0.4	0.2		1
ENacat2	0.4	0.4	0.4	0.6		1
ENacat3	0.4	0.6	0.4	0.4		1

The first three columns of that table list the membership scores of each case in the left-hand sides of the *asf*, and the fourth column reports the membership scores in *AU*. The table shows that the first *asf* ($SP \leftrightarrow AU$) outperforms the other *asf* in cases ENacg3/6/7, ENacto1, ENacosa1, and ENacat3, while it is outperformed by another *asf* in cases ENacg2 and ENacg4. In all other cases, the three solution candidates fare equally. If prior knowledge is available about some of these cases, this information can help to choose among the candidates. For instance, if it is known that in case ENacg7, no other relevant factors are operative than the ones contained in `dat.aut.2`, it follows that ENacg7’s full membership in *AU* must be brought about by *SP*—which, in turn, disqualifies the other solutions. By contrast, if the absence of other relevant factors can be assumed for case ENacg4, the *asf* featuring *SP* as cause of *AU* is disqualified.

7. Benchmarking

Benchmarking the reliability of a method of causal inference is an essential element of method development and validation. In a nutshell, it amounts to testing to what degree the benchmarked method recovers the true data-generating structure Δ or proper substructures of Δ from data of varying quality. As Δ is not normally known in real-life discovery contexts, the reliability of a method cannot be assessed by applying it to real-life data. Instead, reliability benchmarking is done in so-called *inverse searches*, which reverse the order of causal discovery as it is commonly conducted in scientific practice. An inverse search comprises three steps:

- (1) a data-generating causal structure Δ is presupposed/drawn (as ground truth),
- (2) artificial data δ is simulated from Δ , possibly featuring various deficiencies (e.g. noise, fragmentation, measurement error etc.),
- (3) δ is processed by the tested method in order to check whether its output meets the tested reliability benchmark.

A benchmark test can measure various properties of a method’s output, for instance, whether it is fallacy-free, correct or complete, etc. As real-life data are often fragmented, methods for MINUS discovery typically do not infer the complete Δ from a real-life δ but only proper substructures thereof (see section 2.4). Thus, since completeness is not CNA’s primary aim, it should likewise not be the primary reliability benchmark for CNA; it is more important that its output scores high on *fallacy-freeness* and *correctness*.

CNA’s output, *viz.* the issued set \mathbf{csf}_O of *csf*, is fallacy-free iff it does not entail a causal claim that is false of the ground truth Δ (i.e. no false positive). That can be satisfied in two ways: either (i) \mathbf{csf}_O is empty, meaning no causal inferences are drawn, or (ii) \mathbf{csf}_O contains

at least one¹² solution \mathbf{m}_i that is correct of Δ , which is the case iff \mathbf{m}_i is a submodel of Δ (for details on the submodel relation see section 2.4). So, \mathbf{csf}_O satisfies the fallacy-freeness benchmark iff it satisfies conditions (i) or (ii). With increasing stringency, \mathbf{csf}_O can then be said to be correct of Δ iff condition (ii) is satisfied, meaning \mathbf{csf}_O actually contains at least one solution \mathbf{m}_i that is a submodel of Δ , and thus correct. Finally, completeness measures the informativeness of \mathbf{csf}_O , that is, the ratio of causal properties of Δ captured and revealed by the solutions in \mathbf{csf}_O .

The `cna` package provides all the functionalities necessary to conduct inverse searches that are tailor-made to benchmark the output of `cna()` and `csf()`. The functions `randomAsf()` and `randomCsf()` can be used to draw a data-generating structure Δ in step (1). `randomAsf(x)` generates a structure with a single outcome (i.e. a random *asf*) and `randomCsf(x)` an acyclic multi-outcome structure (i.e. a random *csf*), where x is a data frame or `configTable` defining the factors and their possible values from which the structures are drawn. The function `selectCases()`, which has already been discussed in section 3.1.2, can be employed to simulate data δ in the course of step (2). Finally, `is.submodel(x, y)` determines whether models are related by the submodel relation, which, in turn, helps in assessing whether \mathbf{csf}_O is true of Δ . `is.submodel()` takes a character vector x of *asf* or *csf* as first input and tests whether the elements of that vector are submodels of y , which, in turn, is a character string of length 1 representing the target *asf* or *csf* (i.e. Δ). Moreover, the function `identical.model(x, y)` is available to check whether x (which must have length 1) and y are identical.

Against that background, the following might be a core of a correctness benchmark test that simulates multi-value data with 20% missing observations and 10% random noise, and that runs `cna()` and `csf()` at `con = cov = 0.7` without giving the algorithm any prior causal information in an ordering.

```
R> # Draw the ground truth.
R> fullData <- mvct(allCombs(c(4,4,4,4,4)))
R> groundTruth <- randomCsf(fullData, n.asf = 2, compl = 2)
R> # Generate ideal data for groundTruth.
R> x <- ct2df(selectCases(groundTruth, fullData))
R> # Introduce 20% fragmentation.
R> x <- x[-sample(1:nrow(x), nrow(x)*0.2), ]
R> # Introduce 10% random noise.
R> x <- rbind(ct2df(fullData[sample(1:nrow(fullData), nrow(x)*0.1), ]), x)
R> # Run CNA without an ordering.
R> csfs <- csf(mvcna(x, con = .7, cov = .7, maxstep = c(3, 3, 12)))
R> # Check whether no causal fallacy (no false positive) is returned.
R> if(length(csfs$condition)==0) {
+   FALSE } else {any(is.submodel(csfs$condition, groundTruth))}
```

Every re-run of this code chunk generates a different ground truth and different data; in some runs CNA passes the test, in others it does not. To determine CNA's correctness ratio under these test conditions, the above core must be embedded in a suitable test loop. To estimate CNA's overall correctness ratio, the test conditions should be systematically varied

¹²Recall from section 6.3 that an output containing multiple solutions is to be interpreted disjunctively; and a disjunction of solutions is true iff at least one solution is true.

by, for instance, varying the complexity of the ground truth, the degree of fragmentation and noise, the consistency and coverage thresholds, or by drawing the noise with a bias or supplying CNA with more or less prior causal information via an ordering. Fallacy-freeness and completeness tests can be designed analogously, by suitably modifying the last line that evaluates the solution object `csfs`. For single-outcome structures (*asf*), benchmark tests with some of the above variations have been conducted in (Baumgartner and Ambuehl 2020) and (Baumgartner manuscript); corresponding tests for multi-outcome structures (*csf*) have been carried out in (Parkkinen and Baumgartner forthcoming).

8. Summary

This vignette introduced the theoretical foundations as well as the main functions of the **cna** R package for configurational causal inference and modeling with Coincidence Analysis (CNA). Moreover, we explained how to interpret the output of CNA, provided some guidance for how to use various model fit parameters for the purpose of ambiguity reduction, and supplied a benchmarking template.

CNA is currently the only method searching for (M)INUS causation in data that builds multi-outcome models and, hence, not only orders causes conjunctively and disjunctively but also sequentially. Moreover, it builds causal models on the basis of a bottom-up algorithm that is unique among configurational comparative methods and gives CNA an edge over QCA not only with respect to multi- but also single-outcome structures. Overall, CNA constitutes a powerful methodological alternative for researchers interested in (M)INUS structures and the Boolean dimensions of causality, more generally. The **cna** package makes that inferential power available to end-users.

Acknowledgments

We are grateful to Alrik Thiem and Martyna Swiatczak for helpful comments on earlier drafts of this vignette, and we thank the Toppforsk-program of the Bergen Research Foundation and the University of Bergen (grant nr. 811886) and the Swiss National Science Foundation (grant nr. PP00P1_144736/1) for generous support of this research.

References

- Albert DZ (1992). *Quantum Mechanics and Experience*. Harvard University Press, Cambridge.
- Ambuehl M, Baumgartner M (2019). **cnaOpt**: *Optimizing Consistency and Coverage in Configurational Causal Modeling*. R Package Version 0.1.1. URL <https://cran.r-project.org/package=cnaOpt>.
- Baumgartner M (2009a). “Inferring Causal Complexity.” *Sociological Methods & Research*, **38**, 71–101.
- Baumgartner M (2009b). “Uncovering Deterministic Causal Structures: A Boolean Approach.” *Synthese*, **170**, 71–96.

- Baumgartner M (2013). “A Regularity Theoretic Approach to Actual Causation.” *Erkenntnis*, **78**, 85–109.
- Baumgartner M (2015). “Parsimony and Causality.” *Quality & Quantity*, **49**, 839–856.
- Baumgartner M (2020). “Causation.” In D Berg-Schlosser, B Badie, L Morlino (eds.), *The SAGE Handbook of Political Science*, pp. 305–321. SAGE, London.
- Baumgartner M (manuscript). “Configurational Causal Modeling and Logic Regression.” URL https://people.uib.no/mba110/docs/CNA_LR.pdf.
- Baumgartner M, Ambuehl M (manuscript). “Optimizing Consistency and Coverage in Configurational Causal Modeling.” URL <https://people.uib.no/mba110/docs/ConCovOpt.pdf>.
- Baumgartner M, Ambuehl M (2020). “Causal Modeling with Multi-Value and Fuzzy-Set Coincidence Analysis.” *Political Science Research and Methods*, **8**, 526–542. doi:10.1017/psrm.2018.45.
- Baumgartner M, Falk C (2019). “Boolean Difference-Making: A Modern Regularity Theory of Causation.” *The British Journal for the Philosophy of Science*. doi:10.1093/bjps/axz047.
- Baumgartner M, Thiem A (2017). “Model Ambiguities in Configurational Comparative Research.” *Sociological Methods & Research*, **46**(4), 954–987.
- Baumgartner M, Thiem A (2020). “Often Trusted but Never (Properly) Tested: Evaluating Qualitative Comparative Analysis.” *Sociological Methods & Research*, **49**, 279–311. doi:10.1177/0049124117701487.
- Beirlaen M, Leuridan B, Van De Putte F (2018). “A Logic for the Discovery of Deterministic Causal Regularities.” *Synthese*, **195**(1), 367–399. doi:10.1007/s11229-016-1222-x.
- Braumoeller B (2015). *QCAfalsePositive: Tests for Type I Error in Qualitative Comparative Analysis (QCA)*. R package version 1.1.1, URL <https://CRAN.R-project.org/package=QCAfalsePositive>.
- Cronqvist L, Berg-Schlosser D (2009). “Multi-Value QCA (mvQCA).” In B Rihoux, CC Ragin (eds.), *Configurational Comparative Methods: Qualitative Comparative Analysis (QCA) and Related Techniques*, pp. 69–86. Sage Publications, London.
- Culverhouse R, Suarez BK, Lin J, Reich T (2002). “A Perspective on Epistasis: Limits of Models Displaying No Main Effect.” *The American Journal of Human Genetics*, **70**(2), 461–471. doi:<https://doi.org/10.1086/338759>.
- Downing BM (1992). *The Military Revolution and Political Change: Origins of Democracy and Autocracy in Early Modern Europe*. Princeton University Press, N.J.
- Duşa A (2007). “User Manual for the QCA(GUI) Package in R.” *Journal of Business Research*, **60**(5), 576–586.
- Duşa A (2018). “Consistency Cubes: A Fast, Efficient Method for Exact Boolean Minimization.” *The R Journal*, **10**(2), 357–370. doi:10.32614/RJ-2018-080.

- Eberhardt F (2013). “Experimental Indistinguishability of Causal Structures.” *Philosophy of Science*, **80**(5), 684–696.
- Graßhoff G, May M (2001). “Causal Regularities.” In W Spohn, M Ledwig, M Esfeld (eds.), *Current Issues in Causation*, pp. 85–114. Mentis, Paderborn.
- Hájek P (1998). *Metamatematics of Fuzzy Logic*. Kluwer, Dordrecht.
- Hume D (1999 (1748)). *An Enquiry Concerning Human Understanding*. Oxford University Press, Oxford.
- Kalisch M, Maechler M, Colombo D, Maathuis MH, Buehlmann P (2012). “Causal Inference Using Graphical Models with the R Package **pcalg**.” *Journal of Statistical Software*, **47**(11), 1–26.
- Kooperberg C, Ruczinski I (2005). “Identifying Interacting SNPs Using Monte Carlo Logic Regression.” *Genetic Epidemiology*, **28**(2), 157–170. doi:10.1002/gepi.20042.
- Kooperberg C, Ruczinski I (2019). **LogicReg**: *Logic Regression*. R package version 1.6.2. URL <https://CRAN.R-project.org/package=LogicReg>.
- Lemmon EJ (1965). *Beginning Logic*. Chapman & Hall, London.
- Mackie JL (1974). *The Cement of the Universe. A Study of Causation*. Clarendon Press, Oxford.
- McCluskey EJ (1965). *Introduction to the Theory of Switching Circuits*. Princeton University Press, Princeton.
- Oana IE, Medzihorsky J, Quaranta M, Schneider CQ (2020). *SetMethods: Functions for Set-Theoretic Multi-Method Research and Advanced QCA*. R package version 2.5, URL <https://CRAN.R-project.org/package=SetMethods>.
- Parkkinen VP, Baumgartner M (forthcoming). “Robustness and Model Selection in Configurational Causal Modeling.” *Sociological Methods & Research*. URL <https://people.uib.no/mba110/docs/robustness.pdf>.
- Ragin CC (1987). *The Comparative Method*. University of California Press, Berkeley.
- Ragin CC (2006a). “Set Relations in Social Research: Evaluating Their Consistency and Coverage.” *Political Analysis*, **14**, 291–310.
- Ragin CC (2006b). “Set Relations in Social Research: Evaluating Their Consistency and Coverage.” *Political Analysis*, **14**(3), 291–310.
- Rihoux B, Ragin CC (eds.) (2009). *Configurational Comparative Methods. Qualitative Comparative Analysis (QCA) and Related Techniques*. Sage, Thousand Oaks.
- Rothman KJ (1976). “Causes.” *American Journal of Epidemiology*, **104**(6), 587–592. doi:10.1093/oxfordjournals.aje.a112335.
- Ruczinski I, Kooperberg C, LeBlanc M (2003). “Logic Regression.” *Journal of Computational and Graphical Statistics*, **12**(3), 475–511. doi:10.1198/1061860032238.

- Schneider CQ, Wagemann C (2012). *Set-Theoretic Methods: A User's Guide for Qualitative Comparative Analysis (QCA) and Fuzzy-Sets in the Social Sciences*. Cambridge University Press, Cambridge.
- Schwender H, Tietz T (2020). **logicFS**: Identification of SNP Interactions. R package version 2.8.0. URL <https://bioconductor.org/packages/logicFS/>.
- Simon HA (1954). "Spurious Correlation: A Causal Interpretation." *Journal of the American Statistical Association*, **49**(267), 467–479.
- Spirtes P, Glymour C, Scheines R (2000). *Causation, Prediction, and Search*. 2 edition. MIT Press, Cambridge.
- Thiem A (2018). **QCApro**: Advanced Functionality for Performing and Evaluating Qualitative Comparative Analysis. R Package Version 1.1-2. URL <http://www.alrik-thiem.net/software/>.
- Thiem A, Duşa A (2013). *Qualitative Comparative Analysis with R: A User's Guide*. Springer, New York, NY.
- VanderWeele TJ, Robins JM (2009). "Minimal Sufficient Causation and Directed Acyclic Graphs." *Ann. Statist.*, **37**(3), 1437–1465. doi:10.1214/08-AOS613.
- Wensink M, Westendorp RG, Baudisch A (2014). "The Causal Pie Model: An Epidemiological Method Applied to Evolutionary Biology and Ecology." *Ecology and Evolution*, **4**, 1924–1930.

Affiliation:

Michael Baumgartner
University of Bergen
Department of Philosophy
Postboks 7805
5020 Bergen
Norway
E-mail: michael.baumgartner@uib.no
URL: <http://people.uib.no/mba110/>

Mathias Ambuehl
Consult AG Statistical Services
Tramstrasse 10
8050 Zürich
E-mail: mathias.ambuehl@consultag.ch