

# A Tutorial on Regularized Partial Correlation Networks

Sacha Epskamp and Eiko I. Fried

University of Amsterdam: Department of Psychological Methods

<b>Abstract</b> . . . . .	<b>1</b>
<b>Introduction</b> . . . . .	<b>1</b>
<b>Partial Correlation Networks</b> . . . . .	<b>3</b>
<b>LASSO regularization</b> . . . . .	<b>5</b>
<b>Non-normal data</b> . . . . .	<b>7</b>
<b>Example</b> . . . . .	<b>8</b>
<b>Sample size selection and replicability</b> . . . . .	<b>11</b>
A priori sample size analysis . . . . .	11
Post-hoc stability analysis . . . . .	12
<b>Common Problems and Questions</b> . . . . .	<b>13</b>
<b>Conclusion</b> . . . . .	<b>16</b>
<b>References</b> . . . . .	<b>16</b>

## Abstract

Recent years have seen an emergence of network modeling applied to moods, attitudes, and problems in the realm of psychology. In this framework, psychological variables are understood to directly affect each other rather than being caused by an unobserved latent entity. In this tutorial, we introduce the reader to estimating the most popular network model for psychological data: the partial correlation network. We describe how regularization techniques can be used to efficiently estimate a parsimonious and interpretable network structure in psychological data. We show how to perform these analyses in R and demonstrate the method in an empirical example on post-traumatic stress disorder data. In addition, we discuss the effect of the hyperparameter that needs to be manually set by the researcher, how to handle non-normal data, how to determine the required sample size for a network analysis, and provide a checklist with potential solutions for problems that can arise when estimating regularized partial correlation networks.

## Introduction

Recent years have seen increasing use of network modeling for exploratory studies of psychological behavior as

an alternative to latent variable modeling (Borsboom and Cramer, 2013; Schmittmann et al., 2013). In these so-called *psychological networks* (Epskamp et al., 2017a), nodes represent psychological variables such as mood states, symptoms, or attitudes, while edges (links connecting two nodes) represent unknown statistical relationships that need to be estimated. As a result, this class of network models is strikingly different from social networks, in which edges are known (Wasserman and Faust, 1994), posing novel problems for statistical inference. A great body of technical literature exists on the estimation of network models (e.g., Meinshausen and Bühlmann 2006; Friedman et al. 2008; Hastie et al. 2001, 2015; Foygel and Drton 2010). However, this line of literature often requires a more technical background and does not focus on the unique problems that come with analyzing psychological data, such as the handling of ordinal data, how a limited sample size affects the results, and the correspondence between network models and latent variable models.

Currently, the most common model used to estimate psychological networks based on continuous data is the *partial correlation network*. Partial correlation networks are usually estimated using *regularization* techniques originating from the field of machine learning. These techniques have been shown to perform well in retrieving the true network struc-

ture (Meinshausen and Bühlmann, 2006; Friedman et al., 2008; Foygel and Drton, 2010). Regularization involves estimating a statistical model with an extra penalty for model complexity. Doing so leads to a model to be estimated that is *sparse*: many parameters are estimated to be exactly zero. When estimating networks, this means that edges that are likely to be spurious are removed from the model, leading to networks that are simpler to interpret. Regularization therefore jointly performs model-selection and parameter estimation. Regularization techniques have grown prominent in many analytic methods, ranging from regression analysis to principal component analysis (Hastie et al., 2015). In this tutorial, we will only discuss regularization in the context of network estimation. For an overview of such methods applied more broadly in psychological methods, we refer the reader to Chapman et al. (2016).

Regularized network estimation has already been used in a substantive number of publications in diverse fields, such as clinical psychology (e.g., Fried et al. 2016; Knefel et al. 2016; Levine and Leucht 2016; Jaya et al. 2016; Deserno et al. 2016; van Borkulo et al. 2015; Deserno et al. 2016), psychiatry (e.g., Isvoranu et al. 2016b,a; McNally 2016), personality research (e.g., Costantini et al. 2015a,b), and health sciences (e.g., Kossakowski et al. 2015; Langley et al. 2015). What these papers have in common is that they assume observed variables to causally influence one another, leading to network models consisting of nodes such as psychopathology symptoms (e.g. sad mood, fatigue and insomnia), items of personality domains like conscientiousness (e.g. impulse-control, orderliness and industriousness), or health behaviors (e.g. feeling full of energy, getting sick easily, having difficulties performing daily tasks). From this network perspective, correlations among items stem from mutual associations among variables, which differs from the traditional perspective where latent variables are thought to explain the correlation among variables (Schmittmann et al., 2013). Psychological networks thus offer a different view of item clusters: syndromes such as depression or anxiety disorder in the realm of mental disorders (Cramer et al., 2010; Borsboom, 2017; Fried et al., 2017b), personality facets or domains such as extraversion or neuroticism in personality research (Möttus and Allerhand, 2017; Cramer et al., 2012), health domains like physical or social functioning in health research (Kossakowski et al., 2015), and the *g*-factor in intelligence research (Van Der Maas et al., 2006, 2017). Important to note is that one does not have to adhere to this network perspective (i.e., network theory) in order to use the methods described in this tutorial (i.e. network methodology): psychological networks can be powerful tools to explore multicollinearity and predictive mediation, and can even be used to highlight the presence of latent variables.

We are not aware of concise and clear introductions aimed at empirical researchers that explain regularization. The goal

of this paper is thus (1) to provide an introduction to regularized partial correlation networks, (2) to outline the commands used in R to estimate these models, and (3) to address the most common problems and questions arising from analyzing regularized partial correlation networks. The methodology introduced in this tutorial comes with the assumption that the cases (the rows of the spreadsheet) in the data set are independent, which is usually the case in cross-sectional data. Applying these methods to time-series data does not take temporal dependence between consecutive cases into account. We refer the reader to Epskamp et al. (2017d) to a discussion of extending this framework to such temporally ordered datasets. While this tutorial is primarily aimed at empirical researchers in psychology, the methodology can readily be applied to other fields of research as well.

This tutorial builds on the work of two prior tutorials: Costantini et al. (2015a) focused on psychological networks in the domain of personality research, described different types of networks ranging from correlation networks to adaptive LASSO networks (Zou, 2006; Krämer et al., 2009), and introduced basic concepts such as centrality estimation in R. Epskamp et al. (2017a) introduced several tests that allow researchers to investigate the accuracy and stability of psychological networks and derived graph-theoretical measures such as centrality, tackling the topics of generalizability and replicability. The present tutorial goes beyond these papers in the following ways:

- Costantini et al. (2015a) estimated the network structure using a different form of regularization (adaptive LASSO; Zou 2006), a different method for estimating the parameters (node-wise regressions; Meinshausen and Bühlmann 2006), and a different method for selecting the regularization tuning parameter (cross-validation; Krämer et al. 2009). While an acceptable method for estimating regularized partial correlation networks, this procedure can lead to unstable results due to differences in the cross-validation sample selection (see section 2.5.6 of Costantini et al. 2015a) and is not capable of handling ordinal data. We estimate regularized partial correlation networks via the EBIC graphical LASSO (Foygel and Drton, 2010), using polychoric correlations as input when data are ordinal. We detail advantages of this methodology, an important one being that it can be used with ordinal variables that are very common in psychological research.
- We offer a detailed explanation of partial correlations and how these should be interpreted. Especially since the work of Costantini et al. (2015a), researchers have gained a better understanding of the interpretation of partial correlation networks and their correspondence to multicollinearity and latent variable modeling. We

summarize the most recent insights in these topics.

- We provide a state-of-the-art FAQ addressing issues that researchers regularly struggle with—including power analysis and sample size recommendations that have been called for repeatedly (Epskamp et al., 2017a; Fried and Cramer, 2017)—and offer novel solutions to these challenges.

The following sections are structured as follows. First, we introduce partial correlation networks and their estimation, providing detailed information on how these networks can be interpreted. Second, we explain regularization, an integral step in the estimation of partial correlation networks to avoid spurious relationships among items. Third, we explain how to best deal with non-normal (e.g., ordinal) data when estimating partial correlation networks. Fourth, we show researchers how to estimate partial correlation networks in R using an empirical example dataset. The fifth section covers replicability and power analysis for partial correlation networks. In this section, we present the simulation tool `netSimulator`, which allows researchers to determine the sample size that would be required to successfully examine a specific network structure. We also summarize post-hoc stability and accuracy analyses that are described in detail elsewhere (Epskamp et al., 2017a). Finally, we conclude with solutions to the most commonly encountered problems when estimating network models and cover supplemental topics such as comparing networks given unequal sample sizes, unexpected negative relationships among items, very strong positive relationships, or empty networks without any edges.

### Partial Correlation Networks

The most commonly used framework for constructing a psychological network on data that can be assumed to be multivariate normal is to estimate a network of *partial correlation coefficients* (McNally et al., 2015; Borsboom and Cramer, 2013). These coefficients range from  $-1$  to  $1$  and encode the remaining association between two nodes after controlling for all other information possible, also known as conditional independence associations. Partial correlation networks have also been called *concentration graphs* (Cox and Wermuth, 1994) or *Gaussian graphical models* (Lauritzen, 1996), and are part of a more general class of statistical models termed *pairwise Markov random fields* (see e.g., Koller and Friedman 2009 and Murphy 2012 for an extensive description of pairwise Markov random fields). The interpretation of partial correlation networks has recently been described in the psychological literature (e.g., conceptual guides are included in Costantini et al. 2015a and the online supplementary materials of Epskamp et al. 2017a; an extensive technical introduction is included in Epskamp et al. 2017d). To keep the present tutorial self-contained, we suc-

cinctly summarize the interpretation of partial correlations below.

**Drawing partial correlation networks.** After partial correlations have been estimated, they can be visualized in a weighted network structure. Each node represents a variable and each edge represents that two variables are not independent after conditioning on all variables in the dataset. These edges have a weight, *edge weights*, which are the partial correlation coefficients described below. Whenever the partial correlation is exactly zero, no edge is drawn between two nodes, indicating that two variables are independent after controlling for all other variables in the network. Several different software packages can be used to visualize the network. For example, one could use the freely available software packages *cytoscape* (Shannon et al., 2003), *gephi* (Bastian et al., 2009), or R packages *qgraph* (Epskamp et al., 2012), *igraph* (Csardi and Nepusz, 2006), *Rgraphviz* (Gentry et al., 2011) or *ggraph* (Pedersen, 2017). The *qgraph* package has commonly been used in psychological literature as it automates many steps for drawing weighted networks and includes the estimation methods discussed in this paper. When drawing a network model, the color and weight of an edges indicates its magnitude and direction. Using *qgraph*, red lines indicate negative partial correlations, green (using the classic theme) or blue (using the colorblind theme) lines indicate positive partial correlations, with wider and more saturated lines indicating stronger partial correlations (Epskamp et al., 2012).

**Obtaining partial correlation networks.** While multiple ways exist to compute partial correlation coefficients (Cohen et al., 2003), we focus on two commonly-used methods that have been shown to obtain the partial correlations quickly. First, the partial correlations can be directly obtained from the inverse of a variance–covariance matrix. Let  $\mathbf{y}$  represent a set of item responses, which we can assume without loss of generality to be centered. Let  $\mathbf{\Sigma}$  (sigma) denote a variance–covariance matrix. Then, the following states that we assume  $\mathbf{y}$  to have a multivariate normal distribution:

$$\mathbf{y} \sim N(\mathbf{0}, \mathbf{\Sigma}).$$

Let  $\mathbf{K}$  (kappa) denote the inverse of  $\mathbf{\Sigma}$ , also termed the *precision matrix*:

$$\mathbf{K} = \mathbf{\Sigma}^{-1},$$

then, element  $\kappa_{ij}$  (row  $i$ , column  $j$  of  $\mathbf{K}$ ) can be standardized to obtain the partial correlation coefficient between variable  $y_i$  and variable  $y_j$ , after conditioning on all other variables in  $\mathbf{y}$ ,  $\mathbf{y}_{-(i,j)}$  (Lauritzen, 1996):

$$\text{Cor}(y_i, y_j | \mathbf{y}_{-(i,j)}) = -\frac{\kappa_{ij}}{\sqrt{\kappa_{ii}} \sqrt{\kappa_{jj}}}.$$

An alternative way to obtain the partial correlation coefficients is by using *node-wise regressions* (Meinshausen and

Bühlmann, 2006). If one was to perform a multiple regression in which  $y_1$  is predicted from all other variables:

$$y_1 = \beta_{10} + \beta_{12}y_2 + \beta_{13}y_3 + \cdots + \varepsilon_1,$$

followed by a similar regression model for  $y_2$ :

$$y_2 = \beta_{20} + \beta_{21}y_1 + \beta_{23}y_3 + \cdots + \varepsilon_2,$$

and similarly for  $y_3$ ,  $y_4$ , etcetera, then, the same partial correlation coefficient between  $y_i$  and  $y_j$  is proportional to either the regression slope predicting  $y_i$  from  $y_j$  or the regression slope predicting  $y_j$  from  $y_i$  (Pourahmadi, 2011):

$$\text{Cor}(y_i, y_j | \mathbf{y}_{-(i,j)}) = \frac{\beta_{ij}\text{SD}(\varepsilon_j)}{\text{SD}(\varepsilon_i)} = \frac{\beta_{ji}\text{SD}(\varepsilon_i)}{\text{SD}(\varepsilon_j)},$$

in which SD stands for the standard-deviation. Obtaining a partial correlation coefficient by standardizing the precision matrix or performing node-wise regressions will lead to the exact same estimate.

**Interpreting partial correlation networks.** Partial correlation networks allow for several powerful inferences. These points are a summary of a more detailed and technical introduction by Epskamp et al. (2017d):

- *Partial correlation networks allow one to model unique interactions between variables.* If  $A$  correlates with  $B$ , and  $B$  correlates with  $C$ , then we would naturally expect  $A$  to correlate with  $C$ . An unconditional correlation of zero between  $A$  and  $C$  would be unexpected as only few causal structures would lead to such a correlational pattern.<sup>1</sup> If the data are normal, partial correlations can be interpreted as *pairwise interactions*<sup>2</sup>, of which we only need two to model the correlational pattern: an interaction between  $A$  and  $B$  and an interaction between  $B$  and  $C$ . This model will contain one degree of freedom and thus leads to a testable hypothesis (Epskamp et al., 2017c). Such a point of view is akin to loglinear modeling of categorical data (Agresti, 1990; Wickens, 1989), which is structurally comparable to the partial correlation network (Epskamp et al., 2016).
- *The partial correlation network maps out multicollinearity and predictive mediation.* As shown above, partial correlations are closely related to coefficients obtained in multiple regression models: when an independent variable does not predict the dependent variable, we would not expect an edge in the network. The strength of the partial correlation is furthermore directly related to the strength of the regression coefficient. The edges connected to a single node therefore show the researcher the expected result of a multiple regression analysis. Unlike what can be seen from a

multiple regression analysis of a single dependent variable, however, the network also shows which variables would predict the independent variables. By linking separate multiple regression models, partial correlation networks allow for mapping out linear prediction and multicollinearity among all variables. This allows for insight into predictive mediation: a network in which two variables are not directly connected but are indirectly connected (e.g.,  $A - B - C$ ) indicates that  $A$  and  $C$  may be correlated, but any predictive effect from  $A$  to  $C$  (or vice versa) is mediated by  $B$ .

- *Partial correlations can be indicative of potential causal pathways.* Conditional independence relationships, such as those encoded by partial correlation coefficients, play a crucial role in causal inference (Pearl, 2000). When all relevant variables are assumed to be observed (i.e., no latent variables), a partial correlation between variables  $A$  and  $B$  would only be expected to be non-zero if  $A$  causes  $B$ ,  $B$  causes  $A$ , there is a reciprocal relationship between  $A$  and  $B$ , or both  $A$  and  $B$  cause a third variable in the network (Pearl, 2000; Koller and Friedman, 2009). To this end, partial correlation networks are thought of as highly exploratory hypothesis-generating structures, indicative of potential causal effects. While exploratory algorithms exist that aim to discover directed (causal) networks, they rely on strong assumptions such as acyclicity (a variable may not eventually cause itself (e.g.,  $A \rightarrow B \rightarrow C \rightarrow A$ )), and are more strongly influenced by latent variables causing covariation (latent variables would induce directed edges between observed variables implying a strong causal hypothesis). Additionally, these models are not easily identified or parameterized: many equivalent directed models can fit the data equally well, all differently parameterized. Partial correlation networks, on the other hand, are well identified (no equivalent models) and easily parameterized using partial correlation coefficients. As such, exploratively estimating undirected networks offer an attractive alternative to exploratively estimating directed networks, without the troublesome and poorly identified direction of effect<sup>3</sup>.

<sup>1</sup>Two possible options are if  $B$  is a common effect of  $A$  and  $C$  or if two orthogonal latent variables cause covariation between  $A$  and  $B$  and between  $B$  and  $C$ .

<sup>2</sup>Not to be confused with interaction effects of two variables on an outcome variable.

<sup>3</sup>A partial correlation network should not be interpreted to equate the skeleton of a causal model (a directed network with arrowheads removed), as conditioning on a common effect can induce an edge in the partial correlation network. In addition, latent variables can induce edges in both directed and undirected networks. We discuss both common effects and latent variables in detail be-

- *Clusters in the network may highlight latent variables.* While partial correlations aim to highlight unique variance between two variables, they retain shared variance due to outside sources that cannot fully be partialled out by controlling for other variables in the network. As a result, if a latent variable causes covariation between two or more variables in the network, it is expected that all these variables will be connected in the network, forming a cluster (Golino and Epskamp, 2017). Such clusters can thus be indicative of latent variables (Epskamp et al., 2017d). We discuss the relationship between networks and latent variable models in more detail at the end of this paper.

### LASSO regularization

**Limiting spurious edges.** As shown above, partial correlations can readily be estimated by inverting the sample variance–covariance matrix or by performing sequential multiple regressions and standardizing the obtained coefficients. Estimating parameters from data, however, always comes with sampling variation, leading to estimates that are never exactly zero. Even when two variables are conditionally independent, we still obtain nonzero (although typically small) partial correlations that will be represented as very weak edges in the network. These edges are called *spurious* or *false positives* (Costantini et al., 2015a). In order to prevent over-interpretation and failures to replicate estimated network structures, an important goal in network estimation is to limit the number of spurious connections. One way to do so is to test all partial correlations for statistical significance and remove all edges that fail to reach significance (Drton and Perlman, 2004). However, this poses a problem of multiple testing, and correcting for this problem (e.g., by using a Bonferroni correction) results in a loss of power (Costantini et al., 2015a)<sup>4</sup>.

**The LASSO.** An increasingly popular method for limiting the number of spurious edges—as well as for obtaining more interpretable networks that better extrapolate to new samples—is to use statistical *regularization* techniques. An especially prominent method of regularization is the ‘least absolute shrinkage and selection operator’ (LASSO; Tibshirani 1996), which, unlike other regularization techniques, can lead to parameter estimates of exactly zero. In essence, the LASSO limits the sum of absolute partial correlation coefficients; as a result, all estimates shrink, and some become exactly zero. More technically, if  $\mathbf{S}$  represents the sample variance–covariance matrix, LASSO aims to estimate  $\mathbf{K}$  by maximizing the *penalized* likelihood function (Friedman et al., 2008):

$$\log \det(\mathbf{K}) - \text{trace}(\mathbf{S}\mathbf{K}) - \lambda \sum_{\langle i,j \rangle} |\kappa_{ij}|$$

Alternatively, LASSO regularization can be applied on the individual regression models if a network is estimated using node-wise regressions (Meinshausen and Bühlmann, 2006)<sup>5</sup>. Using the LASSO results in a *sparse* network in which likely spurious edges are removed (Epskamp et al., 2017b). The LASSO utilizes a tuning parameter  $\lambda$  (lambda) that controls the level of sparsity. As can be seen above,  $\lambda$  directly controls how much the likelihood is penalized for the sum of absolute parameter values. When the tuning parameter is low, only a few edges are removed, likely resulting in the retention of spurious edges. When the tuning parameter is high, many edges are removed, likely resulting in the removal of true edges in addition to the removal of spurious edges. The tuning parameter therefore needs to be carefully selected to create a network structure that minimizes the number of spurious edges while maximizing the number of true edges (Foygel Barber and Drton, 2015; Foygel and Drton, 2010).

**Selecting the LASSO tuning parameter.** Typically, several networks are estimated under different values of  $\lambda$  (Zhao and Yu, 2006). The different  $\lambda$  values can be chosen from a logarithmically spaced range between a maximal  $\lambda$  value for which no edge is retained (when  $\lambda$  equals the largest absolute correlation; Zhao et al. 2015), and some scalar times this maximal  $\lambda$  value<sup>6</sup>. Thus, the LASSO is commonly used to estimate a *collection* of networks rather than a single network, ranging from a fully connected network to a fully disconnected network. Next, one needs to select the best network out of this collection of networks. This selection can be done by optimizing the fit of the network to the data by minimizing some information criterion. Minimizing the Extended Bayesian Information Criterion (EBIC; Chen and Chen 2008) has been shown to work particularly well in retrieving the true network structure (Foygel Barber and Drton, 2015; Foygel and Drton, 2010; van Borkulo et al., 2014), especially when the generating network is sparse (i.e., does not contain many edges). LASSO regularization with EBIC model selection has been shown to feature high specificity all-around (i.e., not estimating edges that are not in the true network) but a varying sensitivity (i.e., estimating edges that are in the true network) based on the true network structure and sample size. For example, sensitivity typically is less

low.

<sup>4</sup>Unregularized partial correlations can also be seen to already reduce spurious edges in a network comprised of marginal correlation coefficients (Costantini et al., 2015a).

<sup>5</sup>In regularized node-wise regressions, partial correlations obtained from the regression model for one node might slightly differ from partial correlations obtained from the regression model for another node. A single estimate can then be obtained by averaging the two estimated partial correlations.

<sup>6</sup>Current *qgraph* package version 1.4.4 uses 0.01 as scalar and estimates 100 networks by default.

when the true network is dense (contains many edges) or features some nodes with many edges (hubs).

**Choosing the EBIC hyperparameter.** The EBIC uses a hyperparameter<sup>7</sup>  $\gamma$  (gamma) that controls how much the EBIC prefers simpler models (fewer edges; Chen and Chen 2008; Foygel and Drton 2010):

$$\text{EBIC} = -2L + E \log(N) + 4\gamma E \log(P),$$

in which  $L$  denotes the log-likelihood,  $N$  the sample size,  $E$  the number of non-zero edges and  $P$  the number of nodes. This hyperparameter  $\gamma$  should not be confused with the LASSO tuning parameter  $\lambda$ , and needs to be set manually. It typically is set between 0 and 0.5 (Foygel and Drton, 2010), with higher values indicating that simpler models (more parsimonious models with fewer edges) are preferred. Setting the hyperparameter to 0 errs on the side of discovery: more edges are estimated, including possible spurious ones (the network has a higher sensitivity). Setting the hyperparameter to 0.5, as suggested by Foygel and Drton (2010), errs on the side of caution or parsimony: fewer edges are obtained, avoiding most spurious edges but possibly missing some edges (i.e., the network has a higher specificity). Even when setting the hyperparameter to 0, the network will still be sparser compared to a partial correlation network that does not employ any form of regularization; setting  $\gamma$  to 0 indicates that the EBIC reduces to the standard BIC, which still prefers simple models.

Many variants of the LASSO have been implemented in open-source software (e.g., Krämer et al. 2009; Zhao et al. 2015). We suggest the variant termed the ‘graphical LASSO’ (glasso; Friedman et al. 2008), which is specifically aimed at estimating partial correlation networks by inverting the sample variance–covariance matrix. The glasso algorithm has been implemented in the *glasso* package (Friedman et al., 2014) for the statistical programming language R (R Core Team, 2016). A function that uses this package in combination with EBIC model selection as described by Foygel and Drton (2010) has been implemented in the R package *qgraph* (Epskamp et al., 2012), and can be called via the *bootnet* package (Epskamp et al., 2017a). The glasso algorithm directly penalizes elements of the variance–covariance matrix, which differs from other LASSO network estimation methods which typically aim to estimate a network structure by penalizing regression coefficients in a series of multiple regression models (Meinshausen and Bühlmann, 2006). We suggest using this routine because it can be engaged using simple input commands and because it only requires an estimate of the covariance matrix and not the raw data, allowing one to use polychoric correlation matrices when the data are ordinal (discussed below).

To exemplify the above-described method of selecting a best-fitting regularized partial correlation network, we simulated a dataset of 100 people and 8 nodes (variables) based




Figure 1. True network structure used in simulation example. The network represents a *partial correlation network*: nodes represent observed variables and edges represent partial correlations between two variables after conditioning on all other variables. The simulated structure is a *chain graph* in which all absolute partial correlation coefficients were drawn randomly between 0.3 and 0.4.

on the *chain graph* shown in Figure 1. Such graphs are particularly suitable for our example because the true network (the one we want to recover with our statistical analysis) only features edges among neighboring nodes visualized in a circle. This makes spurious edges—any edge that connects non-neighboring nodes—easy to identify visually. We used the *qgraph* package to estimate 100 different network structures, based on different values for  $\lambda$ , and computed the EBIC under different values of  $\gamma$ . Figure 2 depicts a representative sample of 10 of these networks. Networks 1 through 7 feature spurious edges and err on the side of discovery, while networks 9 and 10 recover too few edges and err on the side of caution. For each network, we computed the EBIC based on  $\gamma$  of 0, 0.25 and 0.5 (the hyperparameter the researchers needs to set manually). The boldface values show the best fitting models, indicating which models would be selected using a certain value of  $\gamma$ . When  $\gamma = 0$  was used, network 7 was selected that featured three weak spurious edges. When  $\gamma$  was set to 0.25 or 0.5 (the latter being the default in *qgraph*) respectively, network 8 was selected, which has the same structure as the true network shown in Figure 1. These results show that in our case, varying  $\gamma$  changed the results only slightly. Importantly, this simulation does not imply that  $\gamma = 0.5$  always leads to the true model; simulation

<sup>7</sup>A hyperparameter is a parameter that controls other parameters, and usually needs to be set manually.

work has shown that 0.5 is fairly conservative and may result in omitting true edges from the network (Foygel and Drton, 2010). In sum, the choice of the hyperparameter is somewhat arbitrary and up to the researcher, and depends on the relative importance assigned to caution or discovery (Dziak et al., 2012). Which of these  $\gamma$  values work best is a complex function of the (usually unknown) true network structure.

**A note on sparsity.** It is important to note that although LASSO regularization<sup>8</sup> will lead to edges being removed from the network, it does not present evidence that these edges are, in fact, zero (Epskamp et al., 2017b). This is because LASSO seeks to maximize *specificity*; that is, it aims to include as few *false positives* (edges that are not in the true model) as possible. As a result, observing an estimated network that is sparse (containing missing edges), or even observing an empty network, is in no way evidence that there are, in fact, missing edges. LASSO estimation may result in many *false negatives*, edges that are not present in the estimated network but are present in the true network. This is related to a well-known problem of null hypothesis testing: *Not rejecting the null-hypothesis is not evidence that the null hypothesis is true* (Wagenmakers, 2007). We might not include an edge either because the data are too noisy or because the null hypothesis is true; LASSO regularization, like classical significance testing, cannot differentiate between these two reasons. Quantifying evidence for edge weights being zero is still a topic of future research (Epskamp, 2017; Wetzels and Wagenmakers, 2012).

### Non-normal data

Common challenges to estimating partial correlation networks relate to the assumption of multivariate normality. The estimation of partial correlation networks is closely related to structural equation modeling (Epskamp et al., 2017c), and, as such, also requires multivariate normal distributions. Not only does this mean that the marginal distributions must be normal, all relationships between variables must also be linear. But what do we do with non-normal (e.g. ordered categorical) data that are common in psychological data? Several solutions proposed in the structural equation modeling literature may offer solutions to network modeling as well.

The assumption of normality can be relaxed by assuming that the observed data are a transformation of a latent multivariate normally distributed system (Liu et al., 2009). Figure 3, Panel (a) shows an example of such a model. In this model, squares indicate observed variables, circles indicate normally distributed latent variables and directed arrows indicate monotone (every value is transformed into one unique value, keeping ordinality intact; higher values in the original scale are also higher on the transformed scale) transformation functions. Note that we do not assume measurement error, which could be included by having multiple indicators per latent variable (Epskamp et al., 2017c). Here, we

assume every observed variable indicates one latent variable (Muthén, 1984).

The most common two scenarios are that the observed variables are continuous, or that they consist of ordered categories. When observed variables are continuous, but not normally distributed, the variables can be transformed to have a marginal normal distribution. A powerful method that has been used in network estimation is to apply a *nonparanormal transformation* (Liu et al., 2009). This transformation uses the cumulative distributions (encoding the probability that a variable is below some level) to transform the distribution of the observed variable to that of the latent normally distributed variable. Figure 3, Panel (b) shows a simplified example on how two distributions can be linked by their cumulative distribution. Suppose  $X$  is normally distributed, and  $Y$  is gamma distributed (potentially skewed). Then, values of  $X$  can be mapped to the cumulative distribution by using the probability function (in R: `pnorm`). These cumulative probabilities can then be mapped to values of the gamma distribution by using the quantile function (in R: `qgamma`). In practice, however, the distribution of  $Y$  (top right panel) is not known. The density and cumulative density of  $X$  (left panels), on the other hand, are known, and the cumulative distribution of  $Y$  can be estimated by computing the empirical cumulative distribution function (in R: `ecdf`). Thus, to map values of  $Y$  to values of the normally distributed variable  $X$ , one needs to estimate a smooth transformation function between the bottom two panels. This is the core of the nonparanormal transformation, which aims to map every unique outcome of a variable (e.g., 1, 2 or 3) to one unique outcome of a standard normal variable (e.g., -1.96, 0, 1.65). The `huge.npn` function from the `huge` package (Zhao et al., 2015) can be used to this end. Important to note is that this transformation assumes smoothly increasing cumulative distributions, and will therefore not work when, only a few possible answering options are present (such as in Likert scales). When the data are binary, the transformed data will still be binary, just using different labels than 0 and 1.

When only few item categories are available and the answer options can be assumed to be ordinal (Stevens, 1946), one can make use of threshold functions (Muthén, 1984) as the data transforming functions. Now, the observed score is again assumed to be reflective of a latent normally distributed score, but correlations between items can directly be estimated without having to transform the data. An example of such a threshold function is shown in Figure 3, Panel (c). In this panel, three thresholds are estimated to accommodate four answering categories (0, 1, 2 or 3). The normal distribution corresponds to the latent item score and vertical bars correspond to the thresholds; a person with a latent score below

<sup>8</sup>These arguments apply for other frequentist model selection methods as well, such as removing edges based on statistical significance.






Figure 2. Ten different partial correlation networks estimated using LASSO regularization. Setting the LASSO tuningparameter  $\lambda$  that controls sparsity leads to networks ranging from densely connected to fully unconnected. Data were simulated under the network represented in Figure 1. The fit of every network was assessed using the EBIC, using hyperparameter  $\gamma$  set to 0, 0.25 or 0.5. The bold-faced EBIC value is the best, indicating the network which would be selected and returned using that  $\gamma$  value.

the first would score a 0, a person with a latent score between the first and second threshold would score a 1, etc. After the thresholds are estimated, the correlations between latent variables can be estimated pairwise. These are termed *polychoric* correlations when both variables are ordinal (Olsson, 1979), or *polyserial* correlations when only one of the two variables is ordinal (Olsson et al., 1982). The *lavCor* function from the *lavaan* package (Rosseel, 2012) can be used to compute polychoric and polyserial correlations, which can subsequently be used as input to the *glasso* algorithm (Epskamp, 2016). Regularized partial correlations using *glasso* with EBIC model selection based on polychoric correlations has become standard when estimating psychopathology networks due to the high prevalence of ordered-categorical data. An important limitation is that these methods rely on an assumption that the latent variables underlying the observed ordinal variables are normally distributed, which might not be plausible. For example, some psychopathological symptoms, such as suicidal ideation, might plausibly have a real “zero” point — the absence of a symptom. Properly handling such variables is still a topic of future research (Epskamp, 2017).

When data are binary, one could also use tetrachoric and biserial correlations (special cases of polychoric and polyserial correlations respectively). However, these data would not be best handled using partial correlation networks because of the underlying assumption of normality. When all variables are binary, one can estimate the Ising Model using the *IsingFit* R package (van Borkulo and Epskamp, 2014). The resulting network has a similar interpretation as partial correlation networks, and is also estimated using LASSO

with EBIC model selection (van Borkulo et al., 2014). When the data consist of both categorical and continuous variables, the state-of-the-art network model is termed the Mixed Graphical Model, which is implemented in the *mgm* package (Haslbeck and Waldorp, 2016), also making use of LASSO estimation with EBIC model selection.


### Example

In this section, we estimate a network based on the data of 221 people with a sub-threshold post-traumatic stress disorder (PTSD) diagnosis. The network features 20 PTSD symptoms. A detailed description of the dataset can be found elsewhere (Armour et al., 2017), and the full R code for this analysis can be found in the supplementary materials.<sup>9</sup>


The following R code performs regularized estimation of a partial correlation network using EBIC selection (Foygel and Drton, 2010). This methodology has been implemented in the *EBICglasso* function from the *qgraph* package (Epskamp et al., 2012), which in turn utilizes the *glasso* package for the *glasso* algorithm (Friedman et al., 2014). A convenient wrapper around this (and several other network estimation methodologies such as the Ising Model and the Mixed Graphical Model) is implemented in the *bootnet* package (see Epskamp et al. 2017a for an extensive description), which we use here in order to perform (a) model estimation, (b) a priori sample size analysis, and (c) post-hoc accuracy

<sup>9</sup>We performed these analyses using R version 3.5.0, *bootnet* version 1.0.1 and *qgraph* version 1.4.4, using OSX version 10.11.6.






(a) Observed variables (squares) assumed to be transformations of multivariate normal latent variables (circles).



(b) Visualization on how marginal distributions can be used to transform a variable to have a normal marginal distribution.



(c) Visualization of a threshold model, used in polychoric and polyserial correlations.

Figure 3. Methods for relaxing the assumption of multivariate normality.

and stability analysis. This code assumes the data is present in R under the object name `data`.

```
library("bootnet")
results <- estimateNetwork(
  data,
  default = "EBICglasso",
  corMethod = "cor_auto",
  tuning = 0.5)
```

In this code, `library("bootnet")` loads the package into R, and the `default = "EBICglasso"` specifies that the `EBICglasso` function from *qgraph* is used. The `corMethod = "cor_auto"` argument specifies that the `cor_auto` function from *qgraph* is used to obtain the necessary correlations. This function automatically detects ordinal variables (variables with up to 7 unique integer values) and uses the *lavaan* package (Rosseel, 2012) to estimate polychoric, polyserial and Pearson correlations. Finally, the `tuning = 0.5` argument sets the EBIC hyperparameter,  $\gamma$ ,

to 0.5. After estimation, the network structure can be obtained using the code:

```
results$graph
```

and the network can be plotted using the plot method of *bootnet* using the code:

```
plot(results)
```

This function uses the `qgraph` function from the *qgraph* package to draw the network (Epskamp et al., 2012).<sup>10</sup> By default, edges are drawn using a colorblind-friendly theme (blue edges indicate positive partial correlations and red edges indicate negative partial correlations). Nodes are placed using an modified version of the Fruchterman-Reingold algorithm (Fruchterman and Reingold, 1991) for weighted networks (Epskamp et al., 2012). This algorithm aims to place nodes in an informative way by positioning connected nodes close to each other. A downside of the Fruchterman-Reingold algorithm is that it can behave chaotically: every input will lead to one exact output, but small differences in the input (e.g., a difference of 0.01 in an edge weight or using a different computer architecture) can lead to an entirely different placement of nodes (nodes will likely be placed about the same distance from one-another, but might be placed on a different side of the plotting area). Thus, the eventual placement cannot be interpreted in any substantial way, and might differ substantially between two networks even when there are only very small differences in the network structures. To compare two networks, one should constrain the layout to be equal for both networks. One way to do so is by using `averageLayout` from the *qgraph* package, which was used in drawing Figure 4<sup>11</sup>.

Figure 4 shows the resulting network estimated under three different values of  $\gamma$ : 0, 0.25, and 0.5. Table 1 shows the description of the nodes. As expected, the network with the largest hyperparameter has the fewest edges: the networks feature 105 edges with  $\gamma = 0$ , 95 edges with  $\gamma = 0.25$ , and 87 edges with  $\gamma = 0.5$ .

We can further investigate how important nodes are in the network using measures called centrality indices. These indices can be obtained as followed:

```
centrality(results)
```

<sup>10</sup>Any argument used in this plot method is used in the underlying call to `qgraph`. The *bootnet* plot method has three different default arguments than *qgraph*: (1) the `cut` argument is set to zero, the `layout` argument is set to "spring", and the `theme` argument is set to "colorblind". For more details on these arguments and other ways in which *qgraph* visualizes networks we refer the reader to Epskamp et al. (2012) and the online documentation at <https://CRAN.R-project.org/package=qgraph>.

<sup>11</sup>See online supplementary materials for exact R codes.

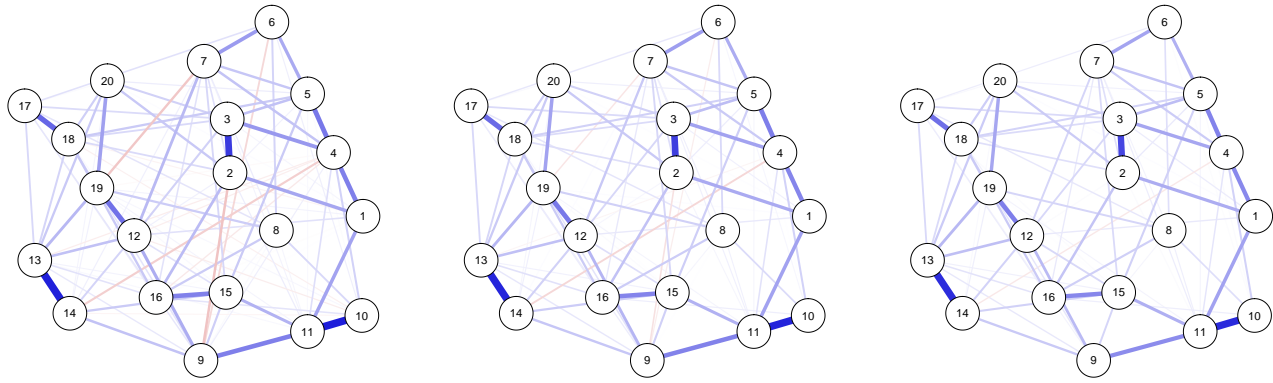
(a)  $\gamma = 0$ (b)  $\gamma = 0.25$ (c)  $\gamma = 0.5$ 

Figure 4. Partial correlation networks estimated on responses of 221 subjects on 20 PTSD symptoms, with increasing levels of the LASSO hyperparameter  $\gamma$  (from left to right: Panel (a) = 0, Panel (b) = 0.25, Panel (c) = 0.5).

Table 1

Description of nodes shown in Figure 4

Node	Description
1	Intrusive Thoughts
2	Nightmares
3	Flashbacks
4	Emotional cue reactivity
5	Psychological cue reactivity
6	Avoidance of thoughts
7	Avoidance of reminders
8	Trauma-related amnesia
9	Negative beliefs
10	Blame of self or others
11	Negative trauma-related emotions
12	Loss of interest
13	Detachment
14	Restricted affect
15	Irritability/anger
16	Self-destructive/reckless behavior
17	Hypervigilance
18	Exaggerated startle response
19	Difficulty concentrating
20	Sleep disturbance

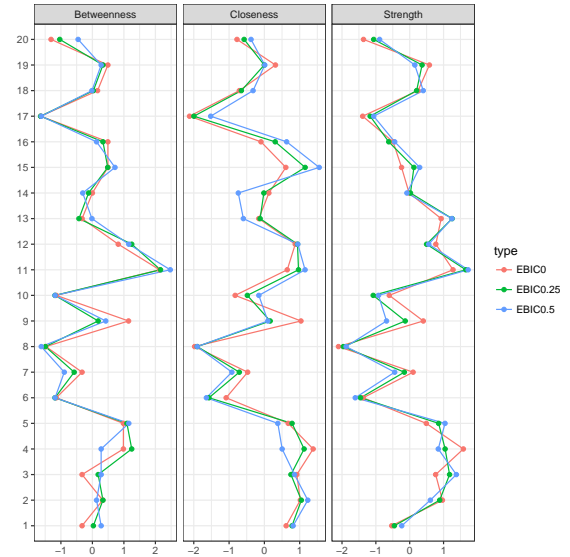


Figure 5. Closeness, betweenness, and degree centrality of the three networks described in Figure 4 with increasing levels of the LASSO hyperparameter  $\gamma$ . Centrality indices are plotted using standardized z-scores in order to facilitate interpretation.

This code provides three commonly used centrality indices: *node strength*, which takes the sum of absolute edge weights connected to each node, *closeness*, which takes the inverse of the sum of distances from one node to all other nodes in the network, and *betweenness*, which quantifies how often one node is in the shortest paths between other nodes. A more extensive overview of these measures and their interpretation is described elsewhere (Costantini et al., 2015a; Epskamp et al., 2017a; Opsahl et al., 2010). All measures

indicate how important nodes are in a network, with higher values indicating that nodes are more important. Figure 5 is the result of the function `centralityPlot` and shows the centrality of all three networks shown in Figure 4. For a substantive interpretation of the network model obtained from this dataset we refer the reader to Armour et al. (2017).

### Sample size selection and replicability

An increasingly important topic in psychological research is the replicability of results (Open Science Collaboration, 2015). High-dimensional exploratory network estimation, as presented in this tutorial paper, lends itself to generating many different measures (e.g., edge weights, network structures, centrality indices) that may or may not replicate or generalize across samples. Recent work has put the importance of replicability in network modeling of psychological data in the spotlight (Epskamp et al. 2017a; Fried et al. 2017b; Fried and Cramer 2017; Fried et al. 2017a; Forbes et al. in press, but see also Borsboom et al. in press). However, it is not easy to determine the replicability of an estimated network. Many factors can influence the stability and accuracy of results, such as the sample size, the true network structure and other characteristics of the data<sup>12</sup>. Even when a network is estimated stably, measures derived from the network structure (e.g., graph theoretical measures such as centrality metrics) might still not be interpretable. For example, all nodes in the true network shown in Figure 1 have exactly the same betweenness (0, all shortest paths do not go via third nodes). Thus, any differences in betweenness in estimated networks are due to chance, regardless of sample size.

We therefore recommend sample size analyses both before and after collecting the data for analysis. A priori sample size analyses let researchers know if the sample size is appropriate for the expected network structure, and post-hoc stability analyses provide researchers with information about the stability of their results. We describe a priori sample size analysis in detail in the next section, which has not been done before in the psychological network literature, and then summarize post-hoc stability analyses that are explicated in detail elsewhere (Epskamp et al., 2017a).

### A priori sample size analysis

An important consideration for any statistical analysis is the sample size required for an analysis, which is often referred to as *power analysis* (Cohen, 1977). To perform such an analysis, one needs to have a prior expectation of the effect size—the expected strength of the true effect. In network modeling, the analogy to an expected effect size is the expected weighted network: a high-dimensional interplay of the network structure (zero and non-zero edges) and the strength of edges (the weight of the non-zero edges). For a partial correlation network of  $P$  nodes, one needs to have a prior expectation on  $P(P - 1)/2$  parameters (edges) to estimate how well edges or any descriptive statistics derived from the network structure, such as centrality indices, can be estimated stably given a certain sample size.<sup>13</sup>

When estimating a network structure, three properties are of primary interest (van Borkulo et al., 2014):

- *Sensitivity*: Also termed the true-positive rate, the pro-

portion of edges present in the true network that were detected in the estimated network.

- *Specificity*: Also termed the true-negative rate, the proportion of missing edges in the true network that were also detected correctly to be absent edges in the estimated network.
- The *correlation* between edge weights of the true network and edge weights of the estimated network, or between centrality estimates based on the true network and centrality estimates based on the estimated network.

A researcher wants sensitivity to increase with sample size and preferably to be high (although a moderate sensitivity can be acceptable as that at least indicates the strongest edges are discovered). When specificity is low, the estimation procedure mistakingly detects many edges that are not present in the true network (false positives). As a result, we argue that researchers always want high specificity. Finally, the correlation indicates how well the true network structure and the estimated network structure mimic one-another. Especially when a researcher is interested in analyzing the network structure as a whole (e.g., for shortest paths analyses), the researcher wants this to be high. In addition to this correlation, the correlation between between centrality indices of the true network and the estimated network might also be of interest, which can be low even though the edge weights are estimated accurately (e.g., when centrality does not differ in the true network, such as betweenness in Figure 1).

Simulation studies have shown that LASSO regularized network estimation generally results in a high specificity, while sensitivity and correlation increases with sample size (Epskamp, 2016; van Borkulo et al., 2014; Foygel and Drton, 2010). This means that whenever LASSO regularization is used, one can interpret edges that are discovered by the method as likely to represent edges in the true network, but should take into account that the method might not discover some true edges. Unfortunately, the precise values of sensitivity, specificity and different correlations are strongly influenced by the expected network structure, similar to how the expected effect size influences a power analysis. As a result, judging the required sample size is far from trivial, but has been called for multiple times in the recent literature (Epskamp et al., 2017a; Fried and Cramer, 2017).

We recommend three ways forward on this issue: (1) more research estimating network models from psychological data

<sup>12</sup>For example, Borsboom et al. (in press) show how data-imputation strategies can lead to unstable edge parameters even at large sample size.

<sup>13</sup>Other network models, such as the Ising model, also require a prior expectation for the  $P$  intercepts. The partial correlation network does not require intercepts as data can be assumed centered.

will make clear what one could expect as a true network structure, especially if researchers make the statistical parameters of their network models publicly available; (2) researchers should simulate network models under a wide variety of potential true network structures, using different estimation methods; (3) researchers should simulate data under an expected network structure to gain some insight in the required sample size. To aid researchers in (2) and (3), we have implemented the `netSimulator` function in the *bootnet* package, which can be used to flexibly set up simulation studies assessing sample size and estimation methods given an expected network structure.

The `netSimulator` function can simulate data under a given network model and expected network structure. Since partial correlation networks feature many parameters, and the field of estimating these models is still young, researchers cannot be expected to have strong theoretical expectations on the network structure. One option is to simulate data under the parameters of a previously published network model, which can be obtained by re-analyzing the data of the original authors or, if the data are not available, asking the original authors to send the adjacency matrix encoding the edge weights. Below, we will conduct such a simulation study by using the estimated network structure in Figure 4, Panel (c) as the simulation baseline. Simulating data under LASSO regularized parameters, however, poses a problem in that these parameters will be biased towards zero due to shrinkage, and therefore might imply a weaker effect than can be expected. To accommodate this, we can first fit the model by using LASSO to obtain a network structure (i.e. which edges are present), and then *refit* a model with only those edges without LASSO regularization (see also Epskamp et al. 2017c on confirmatory partial correlation network analysis). This can be done by using the `refit` argument in `estimateNetwork`:

```
network <- estimateNetwork(
  data,
  default = "EBICglasso",
  corMethod = "cor_auto",
  tuning = 0.5,
  refit = TRUE)
```

Next, a simulation study can be performed using the following R code:

```
simRes <- netSimulator(network$graph,
  dataGenerator = ggmGenerator(
    ordinal = TRUE, nLevels = 5),
  default = "EBICglasso",
  nCases = c(100, 250, 500, 1000, 2500),
  tuning = 0.5,
  nReps = 100,
  nCores = 8
)
```

The `netSimulator` can use any argument of `estimateNetwork`, with a vector of options describing multiple conditions are estimated (e.g., `tuning = c(0.25, 0.5)` would vary the tuning parameter). The first argument is a weights matrix encoding an expected network (or a list with a weights matrix and intercepts vector for the Ising Model which is not needed for partial correlation networks), the `dataGenerator` argument specifies the data generating process (can be ignored for non-ordinal data), `nCases` encodes the sample size conditions, `nReps` the number of repetitions per condition, and `nCores` the number of computer cores to use. Next, results can be printed:

```
simRes
```

or plotted:

```
plot(simRes)
plot(simRes,
  yvar = c("strength", "closeness", "betweenness"))
```

Figure 6 shows the corresponding plots. These plots may be used to gain a rough insight into the required sample size, based on the requirements of the researcher. For example,  $N = 250$  achieves a correlation between the 'true' and estimated networks above 0.8 for edge weights and strength, and above 0.7 for sensitivity. Noteworthy is that specificity is moderate, but not as high as in other studies (Epskamp, 2016; van Borkulo et al., 2014; Foygel and Drton, 2010), possibly a result of the true network structure used being very sparse (54% of the edges were zero in the generating network).

### Post-hoc stability analysis

After estimating a network, bootstrapping methods (Chernick, 2011; Efron, 1979) can be used to gain insight into the accuracy and stability of the network parameters and descriptive statistics based on the estimated network structure (e.g., centrality indices). These are extensively discussed by Epskamp et al. (2017a), including a tutorial on how to perform these analyses using the *bootnet* package. In short, *bootnet* can be used to perform several types of bootstraps using the original data and the estimation method. The two most important methods are:

```
boot1 <- bootnet(results, nCores = 8,
  nBoots = 1000, type = "nonparametric")
boot2 <- bootnet(results, nCores = 8,
  nBoots = 1000, type = "case")
```

The first bootstrap is a non-parametric bootstrap (using re-sampled data with replacement), which can be used to construct confidence intervals around the regularized edge weights (Hastie et al., 2015) and perform significance tests




Figure 6. Simulation results using the estimated refitted PTSD network as true network structure. The top panel shows the sensitivity (true positive rate), specificity (true negative rate) and correlation between true and estimated networks, and the bottom panel shows the correlation between true and estimated centrality indices.

on the difference between different edge weights (e.g. comparing edge  $A - B$  with edge  $A - C$ ) and different centrality indices (e.g., comparing node strength centrality of node  $A$  vs. node  $B$ ). Confidence intervals can *not* be constructed for centrality indices (see the supplementary materials of Epskamp et al. 2017a). To assess the stability of centrality indices, one can perform a case-dropping bootstrap (subsampling without replacement). Based on these bootstraps, the steps from Epskamp et al. (2017a) can be followed to create several plots, which we include for the network in Figure 4, Panel (c) in the supplementary files to this paper. The plots show sizable sampling variation around the edge weights and a poor stability for closeness and betweenness. Strength was more stable, although not many nodes differed from each other significantly in strength. The results of the case-dropping bootstrap can also be summarized in a coefficient, the  $CS$ -coefficient (correlation stability), which quantifies the proportion of data that can be dropped to retain with 95% certainty a correlation of at least 0.7 with the original centrality coefficients. Ideally this coefficient should be above 0.5, and should be at least above 0.25. Strength was shown to be stable ( $CS(\text{cor} = 0.7) \approx 0.516$ ) while closeness ( $CS(\text{cor} = 0.7) \approx 0.204$ ) and betweenness ( $CS(\text{cor} = 0.7) \approx 0.05$ ) were not. Thus, the post-hoc analysis shows that the estimated network structure and derived centrality indices should be interpreted with some care for our example network of PTSD symptoms.

### Common Problems and Questions

**Difficulties in interpreting networks.** Regularized networks can sometimes lead to network structures that are hard to interpret. Here, we list several common problems and questions encountered when estimating and interpreting these models, and try to provide potential ways forward.

1. The estimated network has no or very few edges. This can occur in the unlikely case when variables of interest do not exhibit (partial) correlations. More likely, it

occurs when the sample size is too low for the number of nodes in the network. The EBIC penalizes edge weights based on sample size to avoid false positive associations, which means that with increasing sample size, the partial correlation network will be more and more similar to the regularized partial correlation network. With smaller  $N$  fewer edges will be retained. Figure 7, Panel (a) shows a network estimated on the same data as Figure 4, but this time with only 50 instead of the 221 participants: it is devoid of any edges. A way to remediate this problem is by setting the hyperparameter lower (e.g., 0; see Figure 7, Panel (b)), but note that this increases the likelihood that the network will contain spurious edges. An alternative solution is to reduce the variables of interest and estimate a network based on a subset of variables, because fewer nodes mean that fewer parameters are estimated. However, doing so would lead one to not use all the available data, and might lead to failing to condition on relevant nodes.

2. The network is densely connected (i.e., many edges) including many unexpected negative edges and many implausibly high partial correlations (e.g., higher than 0.8). As the LASSO aims to remove edges and returns a relatively sparse network, we would not expect densely connected networks in any data that are not extremely large. In addition, we would not expect many partial correlations to be so high, as (partial) correlations above 0.8 indicate near-perfect collinearity between variables. These structures can occur when the correlation matrix used as input is not *positive definite*, which can occur when a sample is too small, or when estimating polychoric correlations. Just as a variance has to be positive, a variance-covariance matrix has to be positive-definite (all eigenvalues higher than zero) or at least positive semi-definite (all eigenvalues at least zero). When a covariance matrix is estimated pairwise, however, the resulting matrix is not guaran-

teed to be positive-definite or positive-semi-definite. Polychoric correlation matrices are estimated in such a pairwise manner. In case of a non-positive definite correlation matrix, `cor_auto` will warn the user when it estimates a non-positive definite correlation matrix and attempt to correct for this by searching for a nearest positive definite matrix. This matrix, however, can still lead to very unstable results. When the network looks very strongly connected with few (if any) missing edges and partial correlations near 1 and  $-1$ , the network structure is likely resulting from such a problem and should not be interpreted. We suggest that researchers always compare networks based on polychoric correlations with networks based on Spearman correlations (they should look somewhat similar) to determine if the estimation of polychoric correlations is the source of this problem.

3. While in general the graph looks as expected (i.e., relatively sparse), some edges are extremely high and/or unexpectedly extremely negative. This problem is related to the previous problem. The estimation of polychoric correlations relies on the pairwise crossing of variables in the dataset. When the sample size is relatively low, some cells in the item by item frequency table can be low or even zero (e.g., nobody was observed that scored a 2 on one item and a 1 on another item). The estimation of polychoric correlations is based on these frequency tables and is biased whenever an expected frequency is too small (i.e., below 10; Olsson 1979). Low frequencies can thus lead to biased polychoric correlations, which can compound to large biases in the estimated partial correlations. Another situation in which one might obtain low frequencies is when the scores are highly skewed (Rigdon and Ferguson Jr, 1991), which unfortunately often is the case in psychopathology data. Again, the network based on polychoric correlations should be compared to a network based on Spearman correlations. Obtaining very different networks indicates that the estimation of the polychoric correlations may not be trustworthy.
4. A network has negative edges where the researcher would expect positive ones. This can occur when one conditions on a *common effect* (Pearl, 2000). Suppose one measures three variables: psychology students' grades on a recent statistics test, their motivation to pass the test, and the easiness of the test (Koller and Friedman, 2009). The grade is likely positively influenced by both test easiness and student motivation, and we do not expect any correlation between motivation and easiness: knowing a student is motivated does not help us predict how difficult a professor makes a test. However, we can artificially induce a negative

partial correlation between motivation and easiness by conditioning on a common effect: if we know an unmotivated student obtained an A, we now *can* expect that the test must have been very easy. These negative relationships can occur when common effect relationships are present, and unexpected negative relationships might indicate common effect structures.

Another way these unexpected negative relationships can occur is if the network is based on a subsample, defined by a function on the observed variables. This is because taking a subsample based on a function of the observed variables is the same as conditioning on a common effect (Muthén, 1989). For example, a function of the observed variables might be the sum-score. When using this sum-score to select people to include in the analysis (e.g., to investigate the network structure of subjects with severe symptoms compared to subjects with less severe symptoms), then that subsample is derived by conditioning on the sumscore (e.g., only people with a sumscore above 10 are included). This will lead to spurious negative edges in the expected network structure (Muthén, 1989). Results based on such subsamples should be interpreted with care. In general, this poses a somewhat curious problem: on the one hand, we want to include as many variables as possible; on the other hand, we want to avoid controlling for (i.e. condition on) common effects. Important to note is that one would *not* expect negative partial correlations to occur if the common cause model is true and all variables are scored such that factor loadings are positive (Holland and Rosenbaum, 1986), as such negative relationships where one would expect positive ones can be of particular interest to the researcher.

**Comparing networks.** Another common question is if one can compare two different groups of people (e.g., clinical patients and healthy controls) regarding the connectivity or density of their networks (i.e. the number of edges)? The answer depends on the differences in sample size. As mentioned before, the EBIC is a function of the sample size: the lower the sample size, the more parsimonious the network structure. This means that comparing the connectivity of two networks is meaningful if they were estimated on roughly the same sample size, but that differences should not be compared if this assumption is not met (e.g., see Rhemtulla et al. 2016). One option is to perform a *permutation test* (Van Borkulo et al., 2017). A permutation test is a data-driven method in which all data are first pooled and then randomly assigned to two groups, resulting in two estimated networks. Repeating this process a number of times (e.g., 1000) leads to a distribution of differences between networks given that the two groups come from the same population. This distribution can subsequently be used to perform sta-




Figure 7. Network of 20 PTSD symptoms. Instead of the full data like in Figure 4 (221 subjects), only 50 subjects were used. Panel (a): LASSO hyperparameter  $\gamma$  set to the default of 0.5; Panel (b):  $\gamma$  set to 0 for discovery.

tistical tests on differences of the network structure between the two groups. The permutation test is implemented in the R package *NetworkComparisonTest*.

**Network models versus latent variables.** A final common question relates to how much network modeling and latent variable modeling overlap. Network modeling has been proposed as an alternative to latent variable modeling. As such, researchers might wonder if fitting a network model can provide evidence that the data are indeed generated by a system of variables causally influencing each-other, and not from a common cause model where the covariance between variables is explained by one or more underlying latent variables (Schmittmann et al., 2013)? The short answer is no. While psychological networks have been introduced as an alternative modeling framework to latent variable modeling, and are capable of strongly changing the point of focus from the common shared variance to unique variance between variables (Costantini et al., 2015a), they cannot disprove the latent variable model. This is because there is a direct equivalence between network models and latent variable models (Epskamp et al., 2017c, 2016; Van Der Maas et al., 2006; Kruis and Maris, 2016; Marsman et al., 2015). As discussed above, a latent variable causing covariation on multiple items should lead to a fully connected cluster of items if they are modeled as a network.

While the presence of a latent variable results in a fully connected cluster in the network, this does not mean that when the *estimated* network does not contain fully connected clusters, the latent variable model must be false. As ex-

plained above, the LASSO retaining an edge can provide evidence that an edge is present, but not retaining an edge does not provide evidence that the edge is not present because an edge could simply not be estimated due to a lack of power. We refer the reader to Epskamp et al. (2017b) for a more detailed discussion on this topic and to Epskamp et al. (2017c) for methodology on statistically comparing the fit of a network model and to that of a latent variable model. Finally, just because two models are equivalent does not mean that they are equally plausible. For example, a lattice shaped network structure (nodes ordered on a grid and connected only to neighbors) is equivalent to some latent variable model, but the latent variable model is complicated and very implausible (many latent variables would be needed to explain the data; Marsman et al. 2015).

Even when one expects a network model to largely explain the data, it may be implausible to assume that no latent variables cause any covariation in the network model (Fried and Cramer, 2017; Epskamp et al., 2017c; Chandrasekaran et al., 2012). To this end, estimating causal networks can lead to faulty causal hypotheses in the presence of latent variables. This issue is less problematic when estimating (undirected) partial correlation networks, as no direction of effect is coupled to the estimated edges. Methodologies to combine latent variable modeling and network modeling are currently being developed, which would allow researchers to use strengths from one framework to overcome weaknesses of the other framework. To overcome induced edges from latent variables, one can estimate a network structure after taking co-



variation due to one or more common causes into account (termed a *residual* network; Pan et al. in press; Epskamp et al. 2017c; Chen et al. 2016; Chandrasekaran et al. 2012). Another way of combining network models with latent variable models is to use latent variables as nodes in a network (termed a *latent* network; Epskamp et al. 2017c). Doing so can cope with potential measurement error in the observed variables, allowing for powerful exploratory model search on the structural effects between latent variables (Guyon et al., 2017). Finally, statistical tests to distinguish sparse networks from latent variable models are currently being developed (Van Bork, 2015).

### Conclusion

This paper contains a tutorial on how to estimate psychological networks using a popular estimation technique: LASSO regularization with the EBIC model selection. This method provides a network of partial correlation coefficients with a limited number of spurious edges and can be based on either continuous or ordered-categorical data. This methodology has grown prominent in the past years and is featured in an increasing number of publications throughout various fields of psychological research. In addition, this paper (a) discusses in detail what partial correlations and partial correlation networks are and how these should be interpreted, (b) shows how researchers can estimate these network models in psychological datasets, (c) introduces a new simulation tool to perform power analysis for psychological networks, (d) summarizes post-hoc stability and accuracy analyses, and (e) describes how to deal with most commonly encountered issues when estimating and interpreting regularized partial correlation networks.

The methods described in this paper are only appropriate when the cases in the data can reasonably be assumed to be independent. As this is plausible in cross-sectional analysis, we have exemplified the methodology by analyzing such a dataset. Several authors criticize cross-sectional analysis for not being able to separate within- and between-person variation (Molenaar, 2004; Hamaker, 2012; Bos et al., 2017), and propose to study longitudinal data in order to capture within-person relationships (Bringmann et al., 2013). We refer the reader to Epskamp et al. (2017d) for discussion on this topic and simulation studies studying cross-sectional analysis, and to Weinberger (2015) for a discussion on the causal interpretation of relationships when within-subject variation is lacking. The methods discussed in this paper can readily be applied to within-person data to obtain network structures not confounded by between-subject effects (Epskamp et al., 2017d). For a recent tutorial on this methodology, we refer the reader to Costantini et al. (2017). A downside of this method is that temporal information is not taken into account when estimating network structures. One way to estimate partial correlation networks while taking temporal

information into account is by using the *graphical vector-autoregression* model (graphical VAR; Wild et al. 2010; Fisher et al. 2017; Epskamp et al. 2017d), for which LASSO regularization techniques have been worked out (Abegaz and Wit, 2013; Rothman et al., 2010). EBIC model selection using these routines has been implemented in the R packages *sparseTSCGM* (Abegaz and Wit 2015; aimed at estimating genetic networks) and *graphicalVAR* (Epskamp 2015; aimed at estimating  $n = 1$  psychological networks).

The use of network modeling in psychology is still a young field and is not without challenges. Several related topics were beyond the scope of this tutorial and are discussed elsewhere in the literature. For an overview of challenges and future directions in network modeling of psychological data we refer the reader to Fried and Cramer (2017) and Epskamp (2017). Psychological network analysis is a novel field that is rapidly changing and developing. We have not seen an accessible description of the most commonly used estimation procedure in the literature: LASSO regularization using EBIC model selection to estimate a sparse partial correlation network. This paper addresses this gap by providing an overview of this common and promising method.

### Acknowledgements

We would like to thank Jamie DeCoster for detailed feedback on earlier versions of this paper.

### References

- Abegaz, F. and Wit, E. (2013). Sparse time series chain graphical models for reconstructing genetic networks. *Biostatistics*, page kxt005.
- Abegaz, F. and Wit, E. (2015). *SparseTSCGM: Sparse Time Series Chain Graphical Models*. R package version 2.2.
- Agresti, A. (1990). *Categorical data analysis*. John Wiley & Sons Inc, New York, NY.
- Armour, C., Fried, E. I., Deserno, M. K., Tsai, J., and Pietrzak, R. H. (2017). A network analysis of dsm-5 post-traumatic stress disorder symptoms and correlates in us military veterans. *Journal of anxiety disorders*, 45:49–59.
- Bastian, M., Heymann, S., Jacomy, M., et al. (2009). Gephi: an open source software for exploring and manipulating networks. *Icwsm*, 8:361–362.
- Borsboom, D. (2017). A network theory of mental disorders. *World Psychiatry*, 16(1):5–13.
- Borsboom, D. and Cramer, A. O. J. (2013). Network analysis: an integrative approach to the structure of psychopathology. *Annual review of clinical psychology*, 9:91–121.
- Borsboom, D., Fried, E., Epskamp, S., Waldorp, L., van Borkulo, C., van der Maas, H., and Cramer, A. O. J. (in press). Replicability of psychopathology networks: the

- right question but the wrong answer. a comment on “evidence that psychopathology symptom networks have limited replicability” by forbes, wright, markon, and krueger. *Journal of Abnormal Psychology*.
- Bos, F. M., Snippe, E., de Vos, S., Hartmann, J. A., Simons, C. J. P., van der Krieke, L., de Jonge, P., and Wichers, M. (2017). Exploring the Idiographic Dynamics of Mood and Anxiety via Network Analysis.
- Bringmann, L. F., Vissers, N., Wichers, M., Geschwind, N., Kuppens, P., Peeters, F., Borsboom, D., and Tuerlinckx, F. (2013). A network approach to psychopathology: New insights into clinical longitudinal data. *PloS one*, 8(4):e60188.
- Chandrasekaran, V., Parrilo, P. A., and Willsky, A. S. (2012). Latent variable graphical model selection via convex optimization (with discussion). *The Annals of Statistics*, 40(4):1935–1967.
- Chapman, B. P., Weiss, A., and Duberstein, P. R. (2016). Statistical learning theory for high dimensional prediction: Application to criterion-keyed scale development. *Psychological methods*, 21(4):603–620.
- Chen, J. and Chen, Z. (2008). Extended bayesian information criteria for model selection with large model spaces. *Biometrika*, 95(3):759–771.
- Chen, Y., Li, X., Liu, J., and Ying, Z. (2016). A fused latent and graphical model for multivariate binary data. *arXiv preprint*, page arXiv:1606.08925.
- Chernick, M. R. (2011). *Bootstrap methods: A guide for practitioners and researchers*, volume 619. John Wiley & Sons, New York, NY, USA.
- Cohen, J. (1977). *Statistical power analysis for the behavioral sciences*. Academic Press, New York, NY, USA.
- Cohen, J., Cohen, P., West, S. G., and Aiken, L. S. (2003). *Applied multiple regression/correlation analysis for the behavioral sciences*. Erlbaum, Hillsdale, NJ.
- Costantini, G., Epskamp, S., Borsboom, D., Perugini, M., Möttus, R., Waldorp, L. J., and Cramer, A. O. J. (2015a). State of the aRt personality research: A tutorial on network analysis of personality data in R. *Journal of Research in Personality*, 54:13–29.
- Costantini, G., Richetin, J., Borsboom, D., Fried, E. I., Rhemtulla, M., and Perugini, M. (2015b). Development of indirect measures of conscientiousness: combining a facets approach and network analysis. *European Journal of Personality*, 29(5):548–567.
- Costantini, G., Richetin, J., Preti, E., Casini, E., Epskamp, S., and Perugini, M. (2017). Stability and variability of personality networks. a tutorial on recent developments in network psychometrics. *Personality and Individual Differences*.
- Cox, D. R. and Wermuth, N. (1994). A note on the quadratic exponential binary distribution. *Biometrika*, 81(2):403–408.
- Cramer, A. O. J., Sluis, S., Noordhof, A., Wichers, M., Geschwind, N., Aggen, S. H., Kendler, K. S., and Borsboom, D. (2012). Dimensions of normal personality as networks in search of equilibrium: You can’t like parties if you don’t like people. *European Journal of Personality*, 26(4):414–431.
- Cramer, A. O. J., Waldorp, L., van der Maas, H., and Borsboom, D. (2010). Comorbidity: A Network Perspective. *Behavioral and Brain Sciences*, 33(2-3):137–150.
- Csardi, G. and Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal, Complex Systems*:1695.
- Deserno, M. K., Borsboom, D., Begeer, S., and Geurts, H. M. (2016). Multicausal systems ask for multicausal approaches: A network perspective on subjective well-being in individuals with autism spectrum disorder. *Autism*.
- Drton, M. and Perlman, M. D. (2004). Model selection for gaussian concentration graphs. *Biometrika*, 91(3):591–602.
- Dziak, J. J., Coffman, D. L., Lanza, S. T., and Li, R. (2012). Sensitivity and specificity of information criteria. *The Methodology Center and Department of Statistics, Penn State, The Pennsylvania State University*.
- Efron, B. (1979). Bootstrap methods: another look at the jackknife. *The Annals of Statistics*, 7(1):1–26.
- Epskamp, S. (2015). *graphicalVAR: Graphical VAR for Experience Sampling Data*. R package version 0.1.3.
- Epskamp, S. (2016). Regularized Gaussian Psychological Networks: Brief Report on the Performance of Extended BIC Model Selection. *arXiv preprint*, page arXiv:1606.05771.
- Epskamp, S. (2017). Discussion: The road ahead. In Epskamp, S., editor, *Network Psychometrics*, chapter 12. PhD Dissertation.
- Epskamp, S., Borsboom, D., and Fried, E. I. (2017a). Estimating psychological networks and their accuracy: a tutorial paper. *Behavior Research Methods*.
- Epskamp, S., Cramer, A., Waldorp, L., Schmittmann, V. D., and Borsboom, D. (2012). qgraph: Network visualizations of relationships in psychometric data. *Journal of Statistical Software*, 48(1):1–18.
- Epskamp, S., Kruis, J., and Marsman, M. (2017b). Estimating psychopathological networks: be careful what you wish for. *PloSOne*, 12(6):e0179891.
- Epskamp, S., Maris, G., Waldorp, L., and Borsboom, D. (2016). Network psychometrics. In Irwing, P., Hughes, D., and Booth, T., editors, *Handbook of Psychometrics*. Wiley, New York, NY, USA.
- Epskamp, S., Rhemtulla, M., and Borsboom, D. (2017c). Generalized network psychometrics: Combining network and latent variable models. *Psychometrika*.
- Epskamp, S., Waldorp, L. J., Möttus, R., and Borsboom,

- D. (2017d). Discovering psychological dynamics in time-series data. *arXiv preprint*, page arXiv:1609.04156.
- Fisher, A. J., Reeves, J. W., Lawyer, G., Medaglia, J. D., and Rubel, J. A. (2017). Exploring the Idiographic Dynamics of Mood and Anxiety via Network Analysis.
- Forbes, M. K., Wright, A. G. C., Markon, K., and Krueger, R. (in press). Evidence that psychopathology symptom networks have limited replicability. *Journal of Abnormal Psychology*.
- Foygel, R. and Drton, M. (2010). Extended Bayesian information criteria for Gaussian graphical models. *Advances in Neural Information Processing Systems*, 23:2020–2028.
- Foygel Barber, R. and Drton, M. (2015). High-dimensional Ising model selection with bayesian information criteria. *Electronic Journal of Statistics*, 9(1):567–607.
- Fried, E. I. and Cramer, A. O. J. (2017). Moving forward: challenges and directions for psychopathological network theory and methodology. *Perspectives on Psychological Science*.
- Fried, E. I., Eidhof, M. B., Palic, S., Costantini, G., Huisman-van Diijk, H. M., Bockting, C. L. M., Engelhard, I., Armour, C., Nielsen, A. B. S., and Karstoft, K.-I. (2017a). Replicability and generalizability of ptsd networks: A cross-cultural multisite study of ptsd symptoms in four trauma patient samples.
- Fried, E. I., Epskamp, S., Nesse, R. M., Tuerlinckx, F., and Borsboom, D. (2016). What are ‘good’ depression symptoms? Comparing the centrality of DSM and non-DSM symptoms of depression in a network analysis. *Journal of Affective Disorders*, 189:314–320.
- Fried, E. I., van Borkulo, C. D., Cramer, A. O. J., Lynn, B., Schoevers, R. A., and Borsboom, D. (2017b). Mental disorders as networks of problems: a review of recent insights. *Social Psychiatry and Psychiatric Epidemiology*, 52:1–10.
- Friedman, J. H., Hastie, T., and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441.
- Friedman, J. H., Hastie, T., and Tibshirani, R. (2014). *glasso: Graphical lasso- estimation of Gaussian graphical models*. R package version 1.8.
- Fruchterman, T. and Reingold, E. (1991). Graph drawing by force-directed placement. *Software: Practice and Experience*, 21(11):1129–1164.
- Gentry, J., Long, L., Gentleman, R., Falcon, S., Hahne, F., Sarkar, D., and Hansen, K. (2011). *Rgraphviz: Provides Plotting Capabilities for R Graph Objects*. R package version 1.27.0.
- Golino, H. F. and Epskamp, S. (2017). Exploratory graph analysis: A new approach for estimating the number of dimensions in psychological research. *PloS one*, 12(6):e0174035.
- Guyon, H., Falissard, B., and Kop, J.-L. (2017). Modeling psychological attributes in psychology—an epistemological discussion: Network analysis vs. latent variables. *Frontiers in Psychology*, 8.
- Hamaker, E. L. (2012). Why researchers should think “within-person”: A paradigmatic rationale. *Handbook of research methods for studying daily life*, pages 43–61.
- Haslbeck, J. M. B. and Waldorp, L. J. (2016). mgm: Structure estimation for time-varying mixed graphical models in high-dimensional data. *arXiv preprint*, page arXiv:1510.06871.
- Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA.
- Hastie, T., Tibshirani, R., and Wainwright, M. (2015). *Statistical learning with sparsity: the lasso and generalizations*. CRC Press, Boca Raton, FL, USA.
- Holland, P. W. and Rosenbaum, P. R. (1986). Conditional association and unidimensionality in monotone latent variable models. *The Annals of Statistics*, 14:1523–1543.
- Isvoranu, A. M., Borsboom, D., van Os, J., and Guloksuz, S. (2016a). A Network Approach to Environmental Impact in Psychotic Disorders: Brief Theoretical Framework. *Schizophrenia Bulletin*, 42(4):870–873.
- Isvoranu, A. M., van Borkulo, C. D., Boyette, L., Wigman, J. T. W., Vinkers, C. H., Borsboom, D., and GROUP Investigators (2016b). A Network Approach to Psychosis: Pathways between Childhood Trauma and Psychotic Symptoms. *Schizophrenia Bulletin*. Advance Access published May 10, 2016.
- Jaya, E. S., Hillmann, T. E., Reininger, K. M., Gollwitzer, A., and Lincoln, T. M. (2016). Loneliness and psychotic symptoms: The mediating role of depression. *Cognitive Therapy and Research*.
- Knefel, M., Tran, U. S., and Lueger-Schuster, B. (2016). The association of posttraumatic stress disorder, complex posttraumatic stress disorder, and borderline personality disorder from a network analytical perspective. *Journal of Anxiety Disorders*, 43:70–78.
- Koller, D. and Friedman, N. (2009). *Probabilistic graphical models: principles and techniques*. MIT press, Cambridge, MA, USA.
- Kossakowski, J. J., Epskamp, S., Kieffer, J. M., van Borkulo, C. D., Rhemtulla, M., and Borsboom, D. (2015). The application of a network approach to health-related quality of life (HRQoL): Introducing a new method for assessing hrqol in healthy adults and cancer patient. *Quality of Life Research*, 25:781–92.
- Krämer, N., Schäfer, J., and Boulesteix, A.-L. (2009). Regularized estimation of large-scale gene association networks using graphical gaussian models. *BMC Bioinformatics*, 10(1):1–24.

- Kruis, J. and Maris, G. (2016). Three representations of the ising model. *Scientific Reports*, 6.
- Langley, D. J., Wijn, R., Epskamp, S., and Van Bork, R. (2015). Should I Get That Jab? Exploring Influence to Encourage Vaccination via Online Social Media. *ECIS 2015 Research-in-Progress Papers*, page Paper 64.
- Lauritzen, S. L. (1996). *Graphical models*. Clarendon Press, Oxford, UK.
- Levine, S. Z. and Leucht, S. (2016). Identifying a system of predominant negative symptoms: Network analysis of three randomized clinical trials. *Schizophrenia Research*.
- Liu, H., Lafferty, J. D., and Wasserman, L. (2009). The non-paranormal: Semiparametric estimation of high dimensional undirected graphs. *The Journal of Machine Learning Research*, 10:2295–2328.
- Marsman, M., Maris, G., Bechger, T., and Glas, C. (2015). Bayesian inference for low-rank ising networks. *Scientific reports*, 5(9050):1–7.
- McNally, R. J. (2016). Can network analysis transform psychopathology? *Behaviour Research and Therapy*.
- McNally, R. J., Robinaugh, D. J., Wu, G. W., Wang, L., Deserno, M. K., and Borsboom, D. (2015). Mental disorders as causal systems a network approach to posttraumatic stress disorder. *Clinical Psychological Science*, 3(6):836–849.
- Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *The annals of statistics*, pages 1436–1462.
- Molenaar, P. C. (2004). A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. *Measurement*, 2(4):201–218.
- Möttus, R. and Allerhand, M. (2017). Why do traits come together? the underlying trait and network approaches. *SAGE handbook of personality and individual differences*, 1:1–22.
- Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press, Cambridge, MA, USA.
- Muthén, B. (1984). A general structural equation model with dichotomous, ordered categorical, and continuous latent variable indicators. *Psychometrika*, 49(1):115–132.
- Muthén, B. O. (1989). Factor structure in groups selected on observed scores. *British journal of mathematical and statistical psychology*, 42(1):81–90.
- Olsson, U. (1979). Maximum likelihood estimation of the polychoric correlation coefficient. *Psychometrika*, 44(4):443–460.
- Olsson, U., Drasgow, F., and Dorans, N. J. (1982). The polyserial correlation coefficient. *Psychometrika*, 47(3):337–347.
- Open Science Collaboration (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251):aac4716–aac4716.
- Opsahl, T., Agneessens, F., and Skvoretz, J. (2010). Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, 32(3):245–251.
- Pan, J., Ip, E., and Dube, L. (In press). An alternative to post-hoc model modification in confirmatory factor analysis: The bayesian lasso. *Psychological Methods*.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge Univ Pr.
- Pedersen, T. L. (2017). *ggraph: An Implementation of Grammar of Graphics for Graphs and Networks*. R package version 1.0.0.
- Pourahmadi, M. (2011). Covariance estimation: The glm and regularization perspectives. *Statistical Science*, 26(3):369–387.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rhemtulla, M., Fried, E. I., Aggen, S. H., Tuerlinckx, F., Kendler, K. S., and Borsboom, D. (2016). Network analysis of substance abuse and dependence symptoms. *Drug and alcohol dependence*, 161:230–237.
- Rigdon, E. E. and Ferguson Jr, C. E. (1991). The performance of the polychoric correlation coefficient and selected fitting functions in confirmatory factor analysis with ordinal data. *Journal of marketing research*, pages 491–497.
- Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, 48(2):1–36.
- Rothman, A. J., Levina, E., and Zhu, J. (2010). Sparse multivariate regression with covariance estimation. *Journal of Computational and Graphical Statistics*, 19(4):947–962.
- Schmittmann, V. D., Cramer, A. O. J., Waldorp, L. J., Epskamp, S., Kievit, R. A., and Borsboom, D. (2013). Deconstructing the construct: A network perspective on psychological phenomena. *New Ideas in Psychology*, 31(1):43–53.
- Shannon, P., Markiel, A., Ozier, O., Baliga, N. S., Wang, J. T., Ramage, D., Amin, N., Schwikowski, B., and Ideker, T. (2003). Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome research*, 13(11):2498–2504.
- Stevens, S. S. (1946). On the theory of scales of measurement. *Science, New Series*, 103(2684):677–680.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58:267–288.
- Van Bork, R. (2015). Latent variable and network model implications for partial correlation structures. In *80th Annual Meeting of the Psychometric Society (IMPS)*.
- Van Borkulo, C., Boschloo, L., Kossakowski, J., Tio, P., Schoevers, R., Borsboom, D., and Waldorp, L. (2017). Comparing network structures on three aspects: A permutation test.

- van Borkulo, C. D., Borsboom, D., Epskamp, S., Blanken, T. F., Boschloo, L., Schoevers, R. A., and Waldorp, L. J. (2014). A new method for constructing networks from binary data. *Scientific reports*, 4(5918):1–10.
- van Borkulo, C. D., Boschloo, L., Borsboom, D., Penninx, B. W. J. H., Waldorp, L. J., and Schoevers, R. A. (2015). Association of Symptom Network Structure With the Course of Depression. *JAMA psychiatry*, 72(12):1219–1226.
- van Borkulo, C. D. and Epskamp, S. (2014). *IsingFit: Fitting Ising models using the eLasso method*. R package version 0.2.0.
- Van Der Maas, H., Kan, K.-J., Marsman, M., and Stevenson, C. E. (2017). Network models for cognitive development and intelligence.
- Van Der Maas, H. L., Dolan, C. V., Grasman, R. P., Wicherts, J. M., Huizenga, H. M., and Raijmakers, M. E. (2006). A dynamical model of general intelligence: the positive manifold of intelligence by mutualism. *Psychological review*, 113(4):842–861.
- Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic bulletin & review*, 14(5):779–804.
- Wasserman, S. and Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge University Press, Cambridge, UK.
- Weinberger, N. (2015). If intelligence is a cause, it is a within-subjects cause. *Theory & Psychology*, 25(3):346–361.
- Wetzels, R. and Wagenmakers, E.-J. (2012). A default bayesian hypothesis test for correlations and partial correlations. *Psychonomic bulletin & review*, 19(6):1057–1064.
- Wickens, T. D. (1989). *Multiway contingency tables analysis for the social sciences*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.
- Wild, B., Eichler, M., Friederich, H.-C., Hartmann, M., Zipfel, S., and Herzog, W. (2010). A graphical vector autoregressive modelling approach to the analysis of electronic diary data. *BMC medical research methodology*, 10(1):28.
- Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7:2541–2563.
- Zhao, T., Li, X., Liu, H., Roeder, K., Lafferty, J., and Wasserman, L. (2015). *huge: High-Dimensional Undirected Graph Estimation*. R package version 1.2.7.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429.