

---

# Introduction: Psychological Networks

---

## 1.1 Introduction

There are over 7 billion people in the world, each with a different brain containing 15 to 33 billion neurons. These people are intelligent entities who develop and change over time and who interact with each other in complicated social structures. Consequently, human behavior is likely to be complex. In recent years, research on dynamical systems in psychology has emerged, which is analogous to other fields such as biology and physics. One popular and promising line of research involves the modeling of psychological systems as causal systems or networks of cellular automata (Van Der Maas et al., 2006; Borsboom, 2008; Cramer, Waldorp, van der Maas, & Borsboom, 2010; Borsboom, Cramer, Schmittmann, Epskamp, & Waldorp, 2011). The general hypothesis is that noticeable macroscopic behavior—the co-occurrence of aspects of psychology such as cognitive abilities, psychopathological symptoms, or behavior—is not due to the influence of unobserved common causes, such as general intelligence, psychopathological disorders, or personality traits, but rather to emergent behavior in a network of interacting psychological, sociological, biological, and other components. I will term such networks *psychological networks* to distinguish these models from other networks used in psychology, such as social networks and neural networks.

Figure 1.1 shows an example of such a psychological network, estimated on the `bfi` dataset from the `psych` R package (Revelle, 2010). This dataset contains the responses of 2,800 people on 25 items designed to measure the Big Five personality traits (McCrae & Costa, 1997). The network shows many meaningful connections, such as “make friends easily” being linked to “make people feel at ease,” “don’t talk a lot” being linked to “find it difficult to approach others,” and “carry the conversation to a higher level” being linked to “know how to captivate

---

Parts of this chapter have been adapted from: Epskamp, S., Borsboom, D., and Fried, E.I. (in press). Estimating Psychological Networks and their Accuracy: A Tutorial Paper. *Behavior Research Methods*.

people.” A more detailed description on how to interpret such networks is included below and is discussed further (with respect to personality) in Chapter 10. Psychological networks are strikingly different to network models typically used in complexity research. These are networks between variables (nodes can take one of multiple states) rather than between concrete entities, such as people, computers, cities, and so forth. Furthermore, we do not know the structure of psychological networks. Due to the nascent status of this field, we do not yet know basic properties of psychological networks such as connectivity and clustering, which forces researchers to estimate network structures from data.

This dissertation addresses the problem of how to estimate network models from psychological data, and how such models should consequently be analyzed and interpreted. This is the first dissertation fully devoted to estimating and interpreting psychological networks. Some of the methods discussed (most notably the *qgraph* package; Chapter 9) have already grown to be commonly used psychological research. As a result, the number of researchers working on the estimation of psychological networks has grown substantively during the course of this PhD project. We can now speak of a new field of research: *network psychometrics*. The goal of this dissertation is to start of researchers in this field on the right foot.

Because most chapters of this dissertation utilize a certain class of network models, pairwise Markov random fields (PMRF; Lauritzen, 1996; Murphy, 2012), this introduction is followed by a general introduction to interpreting such models (adapted from Epskamp, Borsboom, & Fried, 2016, which is further used in Chapter 3). The introduction will conclude with a general outline of the dissertation.

## 1.2 Psychological Networks

A psychological network is a model in which *nodes* represent observed psychological variables, usually psychometric test items such as responses to questions about whether a person suffered from insomnia or fatigue in past weeks. These nodes are connected by *edges* which indicate some statistical relationship between them. These models are conceptually different from commonly used reflective latent variable models that explain the co-occurrence among symptoms (e.g., the fact that individuals often suffer from sadness, insomnia, fatigue, and concentration problems at the same time) by invoking an underlying unobserved latent trait (e.g., depression) as the common cause of all the symptoms. Psychological networks offer a different conceptual interpretation of the data and explain such co-occurrences via direct relationships between symptoms; for example, someone who sleeps poorly will be tired, and someone who is tired will not concentrate well (Fried et al., 2015; Schmittmann et al., 2013). Such relationships can then be more easily interpreted when drawn as a network structure where edges indicate pathways on which nodes can affect each other. The edges can differ in strength of connection, also termed *edge weight*, indicating if a relationship is strong (commonly visualized with thick edges) or weak (thin, less saturated edges) and positive (green edges) or negative (red edges). After a network structure is estimated, the visualization of the graph itself tells the researcher a detailed story of the multivariate dependencies in the data. Additionally, many inference methods from




Figure 1.1: Example of a network model estimated the BFI personality dataset from the psych package in R. Nodes represent variables (in this case, personality inventory items) and links between the nodes (also termed *edges*) represent partial correlation coefficients. Green edges indicate positive partial correlations, red edges indicate negative partial correlations, and the width and saturation of an edge corresponds to the absolute value of the partial correlation. Estimation technique as outlined in Chapter 2 was used.

graph theory can be used to assess which nodes are the most important in the network, termed the *most central nodes*.

## Directed and Undirected Networks

In general, there are two types of edges that can be present in a network: an edge can be directed, in which case one head of the edge has an arrowhead indicating a one-way effect, or an edge can be undirected, indicating some mutual relationship. A network that contains only directed edges is termed a *directed network*, whereas a network that contains only undirected edges is termed an *undirected network* (Newman, 2010). Many fields of science consider directed networks interesting because they can be used to encode causal structures (Pearl, 2000). For example, the edge insomnia  $\rightarrow$  fatigue can be taken to indicate that insomnia causes fatigue. The work of Pearl describes that such causal structures can be tested using only observational cross-sectional data and can even be estimated to a certain extent (Kalisch, Mächler, Colombo, Maathuis, & Bühlmann, 2012; Scutari,

2010). However, when temporal information is lacking, there is only limited information present in cross-sectional observational data. Such estimation methods typically only work under two very strict assumptions (a) that all entities which play a causal role have been measured and (b) that the causal chain of cause and effect is not cyclic (i.e., a variable cannot cause itself via any path). Both assumptions are not very plausible in psychological systems. Furthermore, such directed networks suffer from the problem that many equivalent models can exist that feature the same relationships found in the data (MacCallum, Wegener, Uchino, & Fabrigar, 1993); this makes the interpretation of structures difficult. For example, the structure  $\text{insomnia} \rightarrow \text{fatigue} \rightarrow \text{concentration}$  is statistically equivalent to the structure  $\text{insomnia} \leftarrow \text{fatigue} \rightarrow \text{concentration}$  as well as the structure  $\text{insomnia} \leftarrow \text{fatigue} \leftarrow \text{concentration}$ : All three only indicate that insomnia and concentration problems are conditionally independent after controlling for fatigue.

For the reasons outlined above, psychological networks estimated on cross-sectional data are typically undirected networks. The current state-of-the-art method for estimating undirected psychological network structures involves the estimation of PMRFs. A PMRF is a network model in which edges indicate the full conditional association between two nodes after conditioning on all other nodes in the network. This means when two nodes are connected, there is a relationship between these two nodes that cannot be explained by any other node in the network. Simplified, it can be understood as a partial correlation controlling for all other connections. The absence of an edge between two nodes indicates that these nodes are conditionally independent of each other given the other nodes in the network. Thus, a completely equivalent undirected structure (compared to the structures described above) would be  $\text{insomnia} - \text{fatigue} - \text{concentration}$ , indicating that insomnia and concentration problems are conditionally independent after controlling for fatigue.

Figure 1.2 shows a PMRF similar to the example described above. In this network, there is a positive relationship between insomnia and fatigue and a negative relationship between fatigue and concentration. The positive edge is thicker and more saturated than the negative edge, indicating that this interaction effect is stronger than that of the negative edge. This network shows that insomnia and concentration do not directly interact with each other in any way other than through their common connection with fatigue. Therefore, fatigue is the most important node in this network—a concept we will later quantify as *centrality*. These edges can be interpreted in several different ways.<sup>1</sup> First, as shown above, the model is in line with causal interpretations of associations among the symptoms. Second, this model implies that insomnia and fatigue predict each other after controlling for concentration; even when we know someone is concentrating poorly, that person is more likely to suffer from insomnia when we observe that person suffering from fatigue. Similarly, fatigue and concentration predict each other after controlling for insomnia. After controlling for fatigue, there is no longer any predictive quality between insomnia and concentration, even though these variables are correlated; fatigue now mediates the prediction between these

---

<sup>1</sup>A more detailed description of the interpretation of such models can be found in Chapter 6 and Chapter 8.




Figure 1.2: Example of a pairwise Markov random field. Node  $A$  positively interacts with node  $B$ , and node  $B$  negatively interacts with node  $C$ . Nodes  $A$  and  $C$  are *conditionally independent* given node  $B$ .

two symptoms. Finally, these edges can represent genuine symmetric causal interactions between symptoms (e.g., in statistical physics, a PRMF called the Ising model is used to model particles that cause neighboring particles to be aligned).

### Network Inference

In the first step of network analysis, the obtained network is typically presented graphically to show the structure of the data. Afterwards, inference methods derived from graph theory can be applied to the network structure. The estimated PRMF is always a weighted network, which means that we not only look at the structure of the network (e.g., are two nodes connected or not) but also at the strength of connection between pairs of nodes. Because of this, many typical inference methods that concern the global structure of the network (e.g., small-worldness, density, and global clustering; Kolaczyk, 2009; Newman, 2010; Watts & Strogatz, 1998) are less useful in the context of psychological networks because they only take into account whether nodes are connected or not and not the strength of association among nodes. Because the global inference methods for weighted networks and PRMFs are still in development and no consensus has been reached, the network inference section focuses on local network properties: How are two nodes related, and what is the influence of a single node?

**Relationship between two nodes.** The relationship between two nodes can be assessed in two ways. First, we can directly assess the edge weight. This is

always a number that is nonzero because an edge weight of zero would indicate there is no edge. The sign of the edge weight (positive or negative) indicates the type of interaction, and the absolute value of the edge weight indicates the strength of the effect. For example, a positive edge weight of 0.5 is equal in strength to a negative edge weight of  $-0.5$  and both are stronger than an edge weight of 0.2. Two strongly connected nodes influence each other more easily than two weakly connected nodes. This is similar to how two persons standing closer to each other can communicate more easily (via talking) than two people standing far away from each other (via shouting)—two strongly connected nodes are closer to each other. As such, the length of an edge is defined as the inverse of the edge strength. Finally, the distance between two nodes is equal to the sum of the lengths of all edges on the shortest path between two nodes (Newman, 2010).

**Node centrality.** The importance of individual nodes in the network can be assessed by investigating the node centrality. A visualization of a network, such as the one shown in Figure 1.2, is an abstract rendition of a high-dimensional space in two dimensions. Although visualizations of network models often aim to place highly connected nodes into the center of the graph, for instance using the Fruchterman-Reingold algorithm (Fruchterman & Reingold, 1991), the two-dimensional visualization cannot properly reflect the true space of the model. Thus, the metric distance between the placement of nodes in the two-dimensional space has no direct interpretation as it has in multidimensional scaling, for instance. Therefore, graph theory has developed several methods to more objectively quantify which node is most central in a network. Three such centrality measures have appropriate weighted generalizations that can be used with psychological networks (Opsahl, Agneessens, & Skvoretz, 2010). First, node strength, also called degree in unweighted networks (Newman, 2010), simply adds up the strength of all connected edges to a node; if the network is made up of partial correlation coefficients, the node strength equals the sum of absolute partial correlation coefficients between a node and all other nodes. Second, closeness takes the inverse of the sum of all the shortest paths between one node and all other nodes in the network. Thus, where node strength investigates how strongly a node is directly connected to other nodes in the network, closeness investigates how strongly a node is indirectly connected to other nodes in the network. Finally, betweenness looks at how many of the shortest paths between two nodes go through the node in question; the higher the betweenness, the more important a node is in connecting other nodes.

### 1.3 Outline

The above describes, in short, the basis of the methodology introduced in this dissertation. Although the methods and software discussed in this dissertation are applicable to many fields both in and outside of psychology<sup>2</sup>, it is in three fields

---

<sup>2</sup>In recent literature, the methodology has been applied to diverse fields, such as attitude formation (Dalege et al., 2016), test validity (Ziegler, Booth, & Bensch, 2013), dog personality (Goold, Vas, Olsen, & Newberry, 2015), plant breeding (da Silva, Cecon, & Puiatti, 2015; Silva,

where the method has been most applied: clinical psychology, psychometrics, and personality research. As a result of this diversity, the audience of contributions in network psychometrics varies as well between empirical researchers without much programming knowledge, researchers with strong technical skills, and researchers that are familiar with programming in R. To this end, the dissertation is split in three parts: Part I is aimed at empirical researchers with an emphasis on clinical psychology, Part II is aimed at technical researchers with an emphasis on psychometrics, and Part III is aimed at R users with an emphasis on personality research.

Network models have gained extensive footing in clinical psychology for their ability to highlight the dynamics that may lead to someone developing or maintaining mental illness (Cramer et al., 2010; Borsboom & Cramer, 2013; Borsboom, in press). For this reason, Part I of this dissertation is aimed at empirical researchers, with an emphasis on clinical psychology. Chapter 2 will introduce the reader to the network estimation technique we now most often employ on cross-sectional data: model selection on a series of regularized networks. Chapter 3 continues on this topic and presents methods for assessing the accuracy of the network estimated as well as the stability of inferences made on the network structure after observing only subsets of people. Chapter 4 continuously discussing network estimation using regularization and provides a more critical note to argue that results from the methods described in this part are not without problems and certain conclusions cannot be drawn. Finally, Chapter 5 gives a conceptual introduction to time-series modeling on the data of a single patient in clinical practice.

Part II, on the other hand, is aimed at psychometricians and researchers with strong technical skills. Simply stated, these chapters contain equations. Chapter 6 will further outline the modeling of time-series data. This chapter will give the theoretical justification for the model described in Chapter 5 and extends this modeling framework to situations where time-series data is available from multiple people. Rather than contrasting network modeling to the latent variable modeling of classical psychometrics, Chapter 7 and Chapter 8 outline how these frameworks can be combined. Chapter 7 will introduce the network model for multivariate normal data as a formal psychometric model and will cast this model in the well-known structural equation modeling (SEM) framework. Combining network models into SEM allows for two new modeling frameworks, both of which show promise in psychometrics. Finally, Chapter 8 will outline the relationship between the best known network model for binary data, the Ising model, to well-known psychometric modeling frameworks such as logistic regression, loglinear modeling, and item-response theory.

Part III continues the combination of using networks to augment classical psychometrics in three tutorial papers aimed at empirical researchers familiar with programming in R. Two of these chapters are aimed at the third field in which network modeling has grown prominent: personality research (Cramer, Sluis, et al., 2012). These chapters are not technical and are all tutorials using R. Chapter 9 will outline how network-based visualizations can be used to gain insight

---

Rêgo, Pessoa, & Rêgo, 2016), and ecology (Wilkins, Shizuka, Joseph, Hubbard, & Safran, 2015; Gsell, Özkundakci, Hébert, & Adrian, 2016).

into high-dimensional correlational structures as well as large exploratory factor analysis results. Chapter 10 introduces the network perspective to personality researchers. Both Chapter 9 and Chapter 10 will show the use of these methods in personality inventory datasets. Finally, Chapter 11 will introduce an R package that can be used to draw path diagrams and to visualize correlational structures directly from SEM software output.