Stockholms universitet

Pruning of complex networks in psychiatric symptomatology

Vera Andersson

Masteruppsats i matematisk statistik
Master Thesis in Mathematical Statistics

Masteruppsats 2023:9
Matematisk statistik
Juni 2023

www.math.su.se

Matematisk statistik
Matematiska institutionen
Stockholms universitet
106 91 Stockholm

Mathematical Statistics
Stockholm University
Master Thesis **2023:9**
http://www.math.su.se

Pruning of complex networks in psychiatric symptomatology

Vera Andersson*

June 2023

Abstract

Psychiatric symptoms can be conceptualized as nodes in a directed network with edges representing the perceived causality between symptoms, and the visualizations can be useful for identification of interventions and diagnosis of symptoms and mental disorders. However, these networks are often large with complex structures, and we therefore aim to develop simplification methods that remove superfluous edges while maintaining significant structures. We propose three centrality based pruning methods that iteratively remove edges with the lowest edge betweenness, and a generalization of the "brute force" approach in Zhou, Malher Toivonen [23] for directed networks that allows for disconnected components, and this method prunes the edge retaining the connectivity the most in each step. All methods incorporate the node and edge weights, and the second and third centrality based methods compute the PageRank and updated PageRank centrality of all nodes, respectively. We evaluate the methods by comparing the simplifications of pairs of similar, empirical networks from a self-reported survey completed twice by Swedish teenagers having screened positive for depression, and by using synthetic networks with added noise to test reliability. We also let two sets of psychologists who know and do not know the patients visually evaluate the pruned results to select the most useful visualizations. The findings suggest that all methods have trouble with the empirical network pairs, but the basic edge betweenness approach and the non-updating PageRank approach maintain more similarity in general. However, all methods significantly increase the similarity between the synthetic and pruned noisy networks, especially when adding many edges with smaller weights. The psychologists overall prefer the method using connectivity, but those who know the patients prefer the simple edge betweenness approach and the PageRank approach. Therefore, selecting the best method is difficult and might need to be based on the tradeoff between practical usefulness, accuracy, and time complexity.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden. E-mail: vera.c.andersson@gmail.com. Supervisor: Tom Britton.

Sammanfattning

Psykiatriska symtom kan konceptualiseras som noder i ett riktat nätverk med kanter som representerar den upplevda kausaliteten mellan symtom och visualiseringarna kan vara användbara för identifiering av interventioner för och diagnosering av symtom och psykiska störningar. Dessa nätverk är dock ofta stora med komplexa strukturer och vi ämnar därför att utveckla förenklingsmetoder som tar bort överflödiga kanter samtidigt som betydande strukturer bibehålls. Vi föreslår tre centralitetsbaserade *pruning*-metoder som iterativt tar bort kanter med lägst edge betweenness, och en generalisering av "brute force"-metoden i Zhou et al. [23] för riktade nätverk som tillåter isolerade komponenter och denna metod avlägsnar kanten som behåller sammanhängningen mest i varje steg. Alla metoder inkluderar nod- och kantvikterna och den andra och tredje centralitetsbaserade metoderna beräknar PageRank respektive uppdaterad PageRank-centralitet för alla noder. Vi utvärderar metoderna genom att jämföra förenklingarna av par av liknande, empiriska nätverk från en självrapporterad undersökning som genomförts två gånger av svenska tonåringar som screenats positivt för depression, samt genom att använda syntetiska nätverk med tillagd störning för att testa robustheten. Vi låter också två uppsättningar psykologer som känner och inte känner patienterna visuellt utvärdera de förenklade resultaten för att välja de mest användbara visualiseringarna. Resultaten tyder på att alla metoder har problem med de empiriska nätverksparen, men den enklaste edge betweenness-metoden och den icke-uppdaterande PageRank-metoden bibehåller mer likhet mellan nätversparen i allmänhet. Alla metoder ökar dock avsevärt likheten mellan de syntetiska och förenklade störda nätverken, särskilt när vi lägger till många kanter med mindre vikter. Psykologerna föredrar metoden som generaliserar "brute force"-metoden, men de som känner patienterna föredrar den enkla edge betweenness-metoden och PageRank-metoden. Därför är det svårt att välja en metod och valet kan behöva baseras på avvägningen mellan praktisk användbarhet, noggrannhet och tidskomplexitet.

Acknowledgments

I would like to express my sincere gratitude to my supervisor Tom Britton at Stockholm University and my external supervisor Lars Klintwall at Karolinska Institutet for their invaluable support, expert guidance, and helpful feedback throughout the entire writing process. I am immensely grateful for their encouragement and enthusiasm towards this thesis.

A very special thanks to the psychologists who participated in the human expert evaluation survey, and to Lars for helping me organize it.

Contents

1 Introduction

In psychiatric symptomatology, the conceptualization of the interaction of psychiatric symptoms through networks has been developed as a way to use statistical models to aid the diagnosis and identification of possible treatment measures of mental disorders. The networks are referred to as *symptom networks*, where the nodes represent the symptoms and the directed edges represent the causal interactions between them, and the idea is that mental disorders form due to these interactions in the network, according to Borsboom [3]. By studying the visualizations of the symptom networks, psychologists can identify influential nodes and structures, such as cycles where symptoms continue triggering each other, and determine where to optimally introduce interventions to treat mental disorders and prevent them from arising. These networks can be individualized for patients, and the visualizations can then also be used to help the patients understand their own symptomatological processes and to discuss possible treatment options. Although, as with all models, symptom networks are not a perfect depiction of the true interactions between symptoms and the underlying processes that cause them, as these can be more complicated in reality and can be affected by multiple additional factors not quantifiable or representable by the available network structures. Nevertheless, symptom networks can still be useful for psychologists by providing a glimpse into the complex processes involved in the formation of mental disorders.

In this thesis, we consider symptom networks based on symptoms of depression, which is defined by the World Health Organization (WHO) [20] as a mental disorder characterized by experiencing a prolonged depressed mood, or increased disinterest in or indifference towards activities. The networks are created using self-reported data collected from surveys completed by teenagers in Sweden screening positive for this mental disorder, which is globally common in adolescents as the WHO reported in late 2021 [21] that depression was estimated to be experienced by 1.1% of 10-14 year-olds and 2.8% of 15-19 year-olds. These seemingly small percentages represent large numbers of people, as the WHO also stated that one in six people were 10-19-year-olds at the time. Hence, symptom networks can be of great utility for diagnosis and treatment related to these age groups.

However, when using these statistical models in practice, clinicians have discovered that it is not always straightforward to visually interpret the resulting networks created from empirical data, as the networks are often too complex with an overwhelming number of edges that obscure any significant structures a psychologist would be keen to detect. An example of such an empirical symptom network that is included in the data used in this thesis can be found in Figure 1. The nodes represent symptoms experienced by an individual and are labelled accordingly, while the node sizes and edge widths represent the perceived painfulness of the symptoms and the experienced frequency of causality, respectively. Based on this network, it is difficult to identify influential nodes to target for treatment, since the underlying structure is not clearly visible due to the entanglement of numerous overlapping edges which complicates the visual interpretability. It would thus also be difficult to explain the network to the individual whose symptoms it represents. Therefore, to make the symptom networks more practical, clinicians are in need of a systematic approach that can simplify the intricate structures to reveal the most important causal relationships between symptoms, while maintaining the main features in the original network.

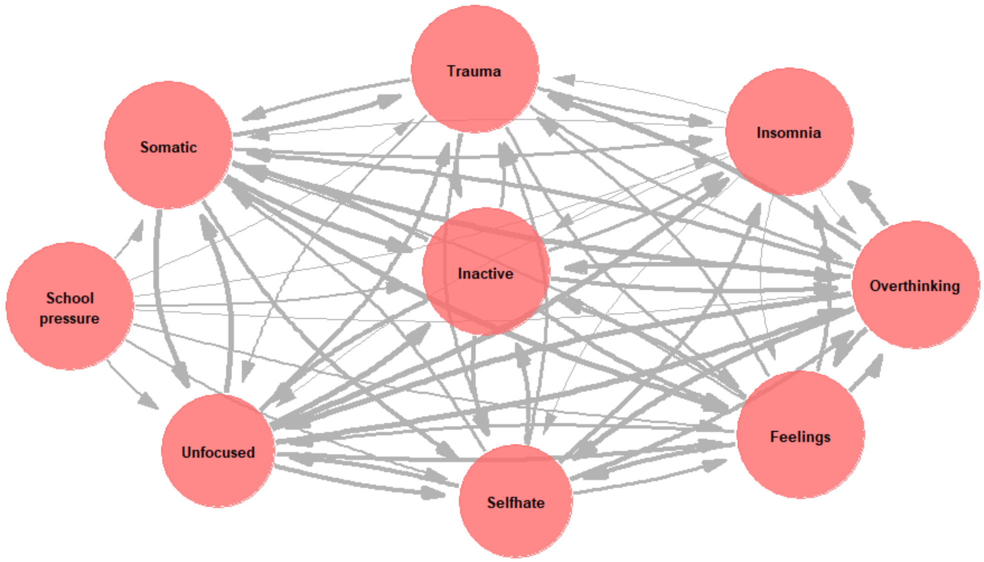

Figure 1: An example of a complex, empirical symptom network showing the importance of establishing pruning methods to improve interpretability. Each node is labeled according to the symptom it represents, and the arrows illustrate the direction of the perceived causal effect by the source node on the end node. The sizes of the nodes represent the painfulness of each symptom and the edge widths represent the frequency of the perceived causal effect. The larger the diameter of the nodes, the larger the node weight, and similarly, the thicker the edges, the larger the edge weight. This network is included in the data used in this thesis (as further described in Section 3) and, in particular, it is the network reported by patient 292 on day 2. The adjacency matrix and the node weights corresponding to this graph can be accessed through the link provided in Appendix A.

This motivation leads us to the aim of this thesis, which is to develop such mathematical methods that simplify the structures of the visualized symptom networks while preserving the key components, by removing superfluous edges that are less influential when it comes to the perceived frequency of interactions between painful symptoms. We refer to the process of systematically removing edges from networks as *pruning*, and variations of simplifying networks have been developed in a wide range of areas. Jiang & Claramunt [15] use thresholds of degree, closeness, and betweenness centralities to simplify graphs by removing the nodes, rather than edges, in a network where each node represents a named street and the edges represent the street intersections. They find in a case study applied to a small street network of Gävle, Sweden that the most important structures are maintained with the filtering. Girvan & Newman [12] extend the notion of betweenness from nodes to edges and use this for community detection by removing the edges with the highest betweenness. Furthermore, Elibol & Chong [8] also use edge betweenness to limit the number of overlapping image pairs in image blending for networks where images are represented by the nodes, and the pairs having an overlapping area are represented by the edges. A different simplification approach involving random graphs is proposed by Dianati [7], who uses a local and a global measure of significance of each edge based on a null model, where nodes of high strength (sum of in-coming weights) are more likely to be connected by an edge.

The main focus of the methods considered in this thesis is on assessing the importance of the edges in the symptom networks to identify possible candidates for pruning. We propose, test, and evaluate four different methods for edge pruning where the first three use edge betweenness centrality to rate the interconnection of each edge, and the last computes a ratio defined in terms of the connectivity kept through maintained quality of the paths present in the network after pruning an edge. The former three methods are distinguishable by their difference in incorporation of the node weights, which represent the perceived painfulness of each symptom while the edge weights represent the perceived frequency of causality between symptoms. The first method, also referred to as the "Edge betweenness" approach, incorporates the end node weight of each edge into their edge weights through normalized multiplication to account for the severity of the symptom affected by the edge. These new edge weights represent the connection strengths between the nodes, and larger weights represent shorter paths in the edge betweenness computations. As edge betweenness ranks edges as more important if they frequently occur on the shortest paths in the network, we iteratively prune the edge with the lowest betweenness until a threshold based on the original network size is reached.

The next method, referred to as the "PageRank" approach, is an extension of the former which additionally incorporates the weighted PageRank centrality to transform the node weights and rate them according to the importance of the symptoms with direct connections to this node as a way to include the network structure, rather than relying solely on the reported painfulness of each symptom. The edge weights are then transformed in the same way as in the first method but with these new PageRank node weights, and the pruning is conducted in the same way using edge betweenness. Similarly, the third method uses the same approach as the second method, but also updates the PageRank scores in each iteration to account for the updated network structure in each step. This method is therefore referred to as the "Updated PageRank" method.

Finally, the fourth and last method is inspired by Zhou et al. [23], who present a pruning approach which they call the "brute force" algorithm, where a ratio of the connectivity kept after removing an edge is computed in each iteration for all remaining edges, and the edge with the smallest effect on the overall connectivity upon removal is eventually pruned. This connectivity is computed in terms of maximizing a path quality function, and we choose to use the reciprocal length of the shortest paths between symptoms to represent the quality of a path. The authors present several other pruning approaches as simplifications of the brute force method as it can be costly due to the iterated computations of the ratio of connectivity kept in each iteration for each edge, but our networks are relatively small, and we therefore select this more accurate but slow method. However, we also modify the method to fit our networks and pruning purposes, by allowing for disconnected components and directed networks, and we refer to this modification as the "Connectivity kept" approach.

The precision of these pruning methods is first evaluated when comparing the results for the networks in pairs of two, possibly quite different, empirical networks reported by the same patient within an interval of a few days to determine if the methods are able to detect common important structures in both networks. The similarities are evaluated using the correlation between the resulting directed and weighted adjacency matrices, the number of common edges in the pruned networks, as well as the number of network pairs for which the pruning methods have identified the same node as having the highest out-degree centrality and thus affecting the greatest number of other symptoms. The results indicate a loss of similarity for all methods when considering the correlations and the number of edges the pruned empirical networks have in common, but the PageRank approach and Edge betweenness approach are able to maintain the most similarity. These methods also increase the similarity the most often post pruning, even if this only occurs for a small number of network pairs, and identify the same most central node for more than half of the network pairs, as originally a third of the empirical network pairs have the same most central symptom.

Moreover, we also evaluate the robustness of the methods by adding varying levels of noise to the nodes, existing edges, and many new edges of small weight in synthetic networks of different sizes chosen to resemble desired pruned results. This is done to ensure that a true, underlying simple network is present in the noisy version of the network, which we do not have access to in the previous evaluation method using empirical networks, and we compare the similarity between the pruned results and the corresponding true networks. The results indicate that all methods are quite robust to low levels of noise when many new edges of low weight are added, and that they perform comparably, with the Connectivity kept approach having slightly higher similarity after pruning. They still perform quite well even when the new edges have higher weights, but there is a decrease in similarity compared to the true synthetic networks. Furthermore, the methods seem to be more robust for larger networks with many more edges, and the Edge betweenness and Connectivity kept approaches identify the same node as the most central more frequently than the other methods.

Nevertheless, the main objective of the pruning methods is to achieve reasonable simplifications of the original complex networks to improve visual interpretability, and we therefore let human experts evaluate the resulting pruned visualizations. They could also be thought of as having access to a true, but perhaps subjective, underlying network through their expertise in psychology. To do this, we let three psychologists evaluate the visualized pruned networks of three of the patients from the data, and we also ask two more psychologists to evaluate the pruned networks of three additional patients not included in the original data who are known to the psychologists. In this way, the latter two psychologists would have an impression of the individual patients and could evaluate how reasonably the networks represent their personal symptomatological processes. We find that the psychologists who do not know the patients predominantly select the Connectivity kept pruning results, while the two psychologists who know their patients favor the Edge betweenness method and the PageRank method. Overall, the Connectivity kept approach is selected the most often, followed by the Edge betweenness approach, and it seems like the selected pruned results usually contained more intricate structures such as cycles involving more symptoms.

Based on the similar evaluation results, it is difficult to choose a method that outperforms the rest, but the Edge betweenness and PageRank approaches give the best results for the empirical networks, handle extra added edges of low weight well, and these methods are favored by the psychologists who know the patients. However, the Connectivity kept approach does well for the synthetic data, and it is favored by the psychologists who do not know the patients. Thus, the best methods seem to be either the Edge betweenness and PageRank methods, as they demonstrate the ability to produce more similar pruning results both for empirical and synthetic data, or the Connectivity kept method as it is favored by the psychologists overall. As a compromise, perhaps the Edge betweenness approach triumphs when considering all evaluation results, as it is both robust and produces more useful visualizations. Another reason to favor this method is that is it the least complex and costly.

2 Theoretical aspects

In this section we introduce notation and basic network theory, and present the centrality measures used in the pruning methods established in Section 4. We also describe the use of networks in psychiatric symptomatology to illustrate this connection.

2.1 Basic network theory

Networks, or graphs, have a wide range of practical applications and can be used to model for instance the spread of infectious diseases, information flow, social relationships, and webpages joined by links on the World Wide Web, among many other real-life situations. Following the terminology presented in Bondy & Murty [2, p.12-14] and the notation in Zhang, Wang & Yan [22, p.2], we define the notion of graphs and some related properties. Specifically, we let $G = (V, E)$ be a graph where V is the set of its vertices, or nodes, and E is the set of its edges, or links. The number of vertices $|V|$ and the number of edges $|E|$ are considered to be finite in this thesis. Furthermore, we consider directed, simple, and weighted graphs, meaning that each edge has a direction, that there are only at most two edges in opposite directions between two vertices, and that there are no nodes with self-loops, i.e. there are no edges with both ends connected to the same node. In this thesis, all pruned graphs are subgraphs $H \subset G$ whose edge set is a subset of E and whose vertex set is equal to V, such that H has the same vertices as G but fewer edges.

Moreover, the nodes and edges are each associated with non-negative numerical weights, and we denote the weight of a directed edge e, or $\{ij\}$, from vertex $i \in V$ to vertex $j \in V$ by w_{ij}, and the weight of a vertex i by v_i. The graph G is thus defined by its weighted adjacency matrix $W := (w_{ij})$ where $w_{ij} = 0$ if there is no directed edge from i to j. If the network is unweighted and thus has no edge weights, the standard adjacency matrix $A = (a_{ij})$ is used instead, where $a_{ij} \in \{0, 1\}$. Furthermore, we let $d_i^{(\text{out})} := \sum_{j \in V} a_{ij}$ and $d_i^{(\text{in})} := \sum_{j \in V} a_{ji}$ denote the out-degree and in-degree of node i, respectively. These definitions correspond to the number of outgoing and incident directed edges of node i. Similarly, for the weighted case we have that $s_i^{(\text{out})} := \sum_{j \in V} w_{ij}$ and $s_i^{(\text{in})} := \sum_{j \in V} w_{ji}$ correspond to the out-strength and in-strength of node i, where *strength* refers to the sum of the edge weights.

Additionally, we let a (directed) *path* between vertices i and j, where $i \neq j$, be a sequence of directed edges that directly lead from vertex i to vertex j without loops such that no vertex is visited twice. If a path P exists from vertex i to j, we say that they are *connected* and denote this by $i \overset{P}{\rightsquigarrow} j$. The length of such a path in a directed, weighted network is determined by the sum of the edge weights along the edges comprising the path, and the *shortest path* is that which has the smallest sum of edge weights. One can for instance use Dijkstra's algorithm, which is described in Section 2.2.3, to find the shortest paths in a network.

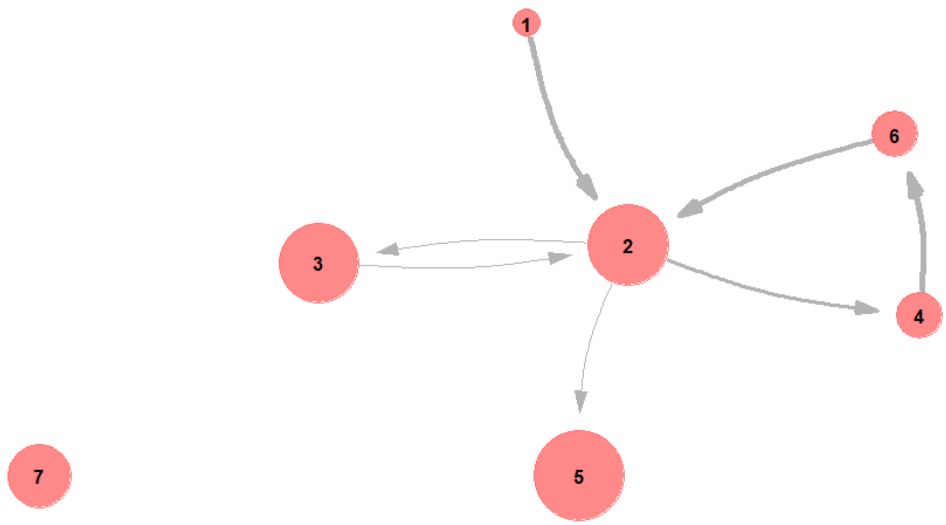

Figure 2: An example of a directed, weighted network where the sizes of the nodes illustrate the node weights and the edge widths visualize the edge weights. The larger the diameter of the nodes, the larger the node weight, and similarly, the thicker the edges, the larger the edge weight. Each arrow points in the direction of the directed edge and the different nodes are labeled from 1 to 7. The adjacency matrix and the node weights corresponding to this graph can be found in Appendix B.

An example of a simple, directed, weighted graph is illustrated in Figure 2, where we have constructed a graph of 7 nodes and 7 directed edges. The sizes of the nodes represent their node weights and, similarly, the widths of the edges represent the corresponding edge weights, where larger sizes and widths indicate larger weights in both cases. Each edge has a direction, as illustrated by the arrows, and the nodes are labeled from 1 to 7. The layout of the nodes is determined using stress majorization, which aims to represent the graph-theoretic distance between nodes by the Euclidean distance in the illustration. This layout method is used for all networks in this thesis, and Gasner et al. [11] describe the algorithm in more detail.

Furthermore, a *cycle*, or a *loop* as we also call it in this thesis, is a path with the same start and end nodes but with distinct nodes in between. In Figure 2, nodes 2 and 3 form loops involving two directed edges, and nodes 2, 4, and 6 also form loops, but with three directed edges. Since node 1 does not have any incoming edges, we call this a *start node* in this thesis, and conversely, we call nodes like node 5 with only incoming edges *sink nodes*. The disconnected vertex 7 forms its own component of size 1 and has no direct relation through any edges to vertices 1 to 6. Additionally, we note that nodes 1 to 6 make up the largest component, which is *weakly connected* as the underlying undirected component is connected, according to Kolaczyk & Csárdi [18, p.23-24]. This does not necessarily mean that each node can be reached by the others by following a directed edge, but merely that these nodes form a separate component if the directions are disregarded. Had it been the case that each vertex in this component could be reached by a directed edge, the component would have been *strongly connected*.

Nevertheless, the networks described by the theory above are oftentimes complicated and difficult to interpret, evaluate, and compare. Their complexity can for example be affected by the sizes of the networks, as it quickly becomes difficult to visually assess the illustrated relationships when the graph resembles a "hairball" with numerous nodes and edges concealing any principal structures. This is especially the case for our application of using networks to describe perceived relationships between psychiatric symptoms, and an example of a complex symptom network that is included in the data used in this thesis (as further described in Section 3) was presented in the introduction in Figure 1. The many edges are overwhelming, and it is difficult to make out any important structures or to identify the most influential symptoms. We are hence in need of various techniques and measures to collectively use to produce a deeper analysis of the networks, in particular for the purpose of comparisons between different networks. This leads us to the use of centrality measures to identify the most intrinsic structures in the networks, and some selected measures are described below in Section 2.2. We first describe out-degree centrality for nodes, which ranks the nodes in decreasing order of out-degree and this measures how central the nodes are in terms of how many other nodes are in turn affected by the selected node. Next, we move on to betweenness centrality for nodes and edges, respectively, where the nodes or edges are ranked based on how often they appear on the shortest paths in the network and this gives an indication of how central they are to the fastest and most frequent interactions between nodes. Moreover, we also introduce the PageRank centrality for nodes, which takes into consideration the importance of the other nodes leading to the current node, in contrast to the aforementioned node centralities.

2.2 Centrality measures

When analyzing networks, a crucial notion is the importance of nodes (and edges), which can be defined in various ways. As pointed out by Estrada [9, p.121], the purpose of centrality measures is to compute a score for each vertex (or edge, as we can extend the definitions to edges in some cases as described in Section 2.2.2 for betweenness centrality) in a network that ranks their importance or position in the graph. In this way, one can make inference based on these values to determine how interconnected each vertex or edge is in the network structure. For instance, one could use node centrality to determine the most popular person in a community where nodes represent different people and the edges represent friendships, or one could use edge centrality to find the most important physical interactions in a network of people in a community where a disease could spread. In this thesis, we use centrality measures to evaluate the impact of symptoms on the causality of other symptoms in networks, as well as to assess the usefulness of the edges to decide which edges to keep in our simplified networks. The various methods used are described in the coming subsections.

2.2.1 Out-degree centrality

One centrality measure for nodes is *degree* centrality, which is defined in Estrada [9, p.121-123] by using the degree of each node to rate their overall importance. For directed networks, we can distinguish between the in-degrees and out-degrees for the nodes, and in this thesis we consider the latter, as we are more interested in the causality of each symptom. The *out-degree* centrality measure simply assigns node i its out-degree $d_i^{(\text{out})} := \sum_{j \in V} a_{ij}$ as its centrality score. As an example, we note in Figure 2 that node 2 has the highest out-degree centrality. This could be useful in scenarios where the number of further connections of a vertex is of importance, for example if we are interested in how many other symptoms are caused by a given symptom represented by a vertex. This is similar to how infection spreads from one individual to the next in the example of a network representing a community mentioned above, where each edge demonstrates the physical interactions between the individuals represented by the nodes, and out-degree centrality would be an appropriate measure in this case as well to find the person who directly infects the most number of people. However, this is quite a simple centrality measure, and we can also find central objects by determining which agents in the network lie on a high number of the shortest communication paths between nodes. A centrality measure based on this is described in the subsequent section.

2.2.2 Edge betweenness centrality

Another centrality measure is *betweenness* centrality for nodes which is defined in Freeman [10], and this takes into account the ability of nodes to regulate the flow of information in the network by forwarding, halting, or altering the information transmitted from connected nodes. Vertices are considered central in the sense of being *between* other nodes in relation to the paths connecting them, and the level of centrality of a node is computed using the number of occurrences of the node on the shortest paths in the networks. Therefore, nodes have less control of the information flow if they appear on fewer of the shortest paths in the network. In particular, Freedman defines betweenness of node k as

$$B(k) := \sum_{i\neq j, i\neq k, j\neq k} \frac{\sigma_{ij}(k)}{\sigma_{ij}} \tag{1}$$

for distinct vertices i, j, and k, where σ_{ij} denotes the number of shortest paths between i and j, and $\sigma_{ij}(k)$ denotes the number of shortest paths between i and j where vertex k is included on the path. The author defines this centrality measure for an undirected, unweighted graph, but the definition holds for directed, weighted graphs if we let the definitions of a path and its length include directions and weights of the edges, as we do in this thesis. Moreover, if there is no path between vertices i and j, we let $B(k) = 0$ since there is no shortest path between the two vertices where vertex k would appear. One can think of Eq. (1) as the sum of partial betweenness values for pairs of nodes i and j, where $\frac{\sigma_{ij}(k)}{\sigma_{ij}}$ is the probability that node k lies on a randomly selected shortest path between nodes i and j (and there could be more than one shortest path if there are multiple paths of the shortest length). The overall partial betweenness of node k given in Eq. (1) thus increases by 1 if node k is included in the only shortest path between nodes i and j, but if there are multiple shortest paths, it increases proportionally to the number of times node k is included on those respective paths. In the case of weighted networks, it is highly unlikely that there are multiple shortest paths between nodes, and we can therefore simplify Eq. (1) by letting $\sigma_{ij} = 1$ to get that betweenness is defined by

$$B(k) := \sum_{i\neq j, i\neq k, j\neq k} \sigma_{ij}(k)$$

in this case.

To illustrate this centrality measure, we include an example of a directed network (that is unweighted for simplicity) in Figure 3, where the node with the highest betweenness is colored red. This highlighted center node is acting as a gatekeeper in the right direction between the nodes on each side of it, and thus has the highest betweenness as all the shortest paths from the nodes on the left side to the nodes on the right must pass through this center node. The two nodes directly connected to the red node on each side also have high betweenness, since they are in turn directly connected to several peripheral nodes whose shortest paths have to pass through these nodes. However, the red node in the middle has one more path going through it, which is the path from the directly connected node on the left to the one on the right. Intuitively, this network could for instance demonstrate how information can spread between people represented by the nodes who communicate through forwarding messages in the direction of the edges. In this case, the node with the highest betweenness would represent an influential person who controls the flow of information. If this person is removed from the network and hence does not receive the forwarded messages from the people on the left, then the people on the right will have no chance of receiving this information according to this structure.

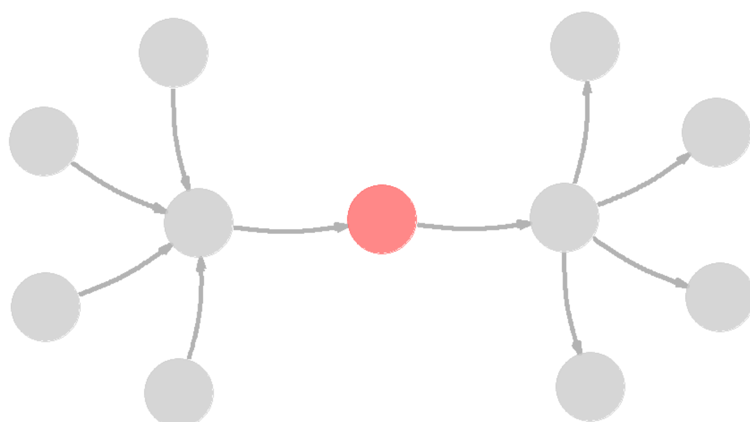

Figure 3: An example where the node with the highest betweenness in an unweighted, directed network is highlighted in red. This node has a betweenness score of 25, while the nodes it is directly connected to each have a betweenness score of 24, as only one additional path passes through the red node. The remaining nodes all have 0 betweenness. The direction of the edges in the network is represented by the direction of the arrows.

Correspondingly, this centrality measure can be extended to edges and we call this *edge betweenness*, which is similarly defined by Girvan & Newman [12] as the number of shortest paths between pairs of nodes containing the given edge. We can thus generalize Eq. (1) to the following

$$B(e) := \sum_{i \neq j} \frac{\sigma_{ij}(e)}{\sigma_{ij}} \tag{2}$$

for an edge e, where $\sigma_{ij}(e)$ now represents the number of shortest paths between distinct vertices i and j passing through the edge e. The definition in Eq. (2) again holds for directed, weighted networks and $B(e) = 0$ if there is no path between vertices i and j for similar reasons as above. The case of multiple shortest paths is handled in the same way as for node betweenness, and we can hence simplify Eq. (2) to similarly become

$$B(e) := \sum_{i \neq j} \sigma_{ij}(e)$$

for weighted networks. Furthermore, the definition suggests that edges that have high edge betweenness will be more central to the network flow, as for example a network with clusters of nodes where the clusters are only interconnected with a few edges will have all the shortest paths between these groups of nodes passing through the few links between the clusters. Hence, these edges will have high edge betweenness scores and examples of this are illustrated in Figure 4 for two unweighted, directed graphs, where the edges with the highest betweenness are colored in red. In Figure 4a, the case where two groups of nodes are interconnected by a single edge acting as a bridge is shown, and in Figure 4b, the case where two groups of nodes are connected with a loop of three edges in the middle is shown, and a darker shade of red indicates a higher edge betweenness in this figure. We also note that the direction of the edges in the right-hand side group is flipped compared to in Figure 4a. In addition, we observe that the peripheral edges only connected to a few nodes in the figures tend to have lower edge betweenness as they are involved in fewer of the shortest paths between vertices. However, it is also possible that such edges have high betweenness if for example most of the nodes in the network are reachable by following this initial edge, and there are no alternative routes that are shorter than starting with this edge.

Furthermore, connecting these networks to the example above involving people (nodes) forwarding messages along communication paths (edges), we would consider the bridge edge in Figure 4a the strongest controller of communication flow, and the information would not be able to reach the people grouped together on the right-hand side of this edge if we remove it by disallowing interaction between the people linked by the edge. Equivalently, the edges in the loop in Figure 4b also regulate the flow of information the most by circulating the messages between the two groups. Moreover, these example networks are unweighted for simplicity, but the edge weights also affect edge betweenness, since paths are defined in terms of their lengths which are computed using the edge weights. Thus, if edge weight represents the strength of the link, a path with edges of greater edge weight would be shorter than one with smaller edge weights. Conversely, if the edge weights represent the distance between the nodes at each end of the edge, a path would be considered longer if the edge weights are greater.

In this thesis, we make use of edge betweenness and the length of shortest paths in our network simplification methods presented in Section 4, and in particular we use the function `edge_betweenness` from the `igraph` package in R to compute the edge betweenness scores. According to the documentation, this function utilizes an algorithm similar to the one described by Brandes [4] in the computations, which requires $\mathcal{O}(|V| + |E|)$ space and runs in $\mathcal{O}(|V|(|E| + |V| \log |V|))$ time. This algorithm uses a modification of Dijkstra's algorithm to more efficiently compute betweenness scores for weighted graphs, and the shortest paths used in Method IV in Section 4.4 are computed using the function `distances` also from the `igraph` package in R, which also uses Dijkstra's algorithm. We therefore describe this algorithm in the next section for completeness.

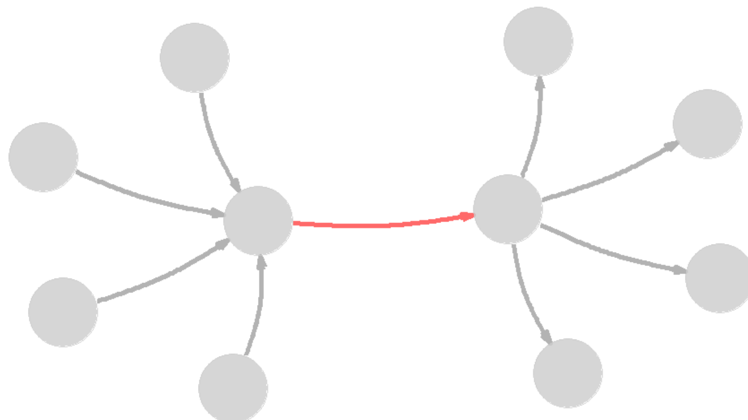

(a) Bridge with highest betweenness.

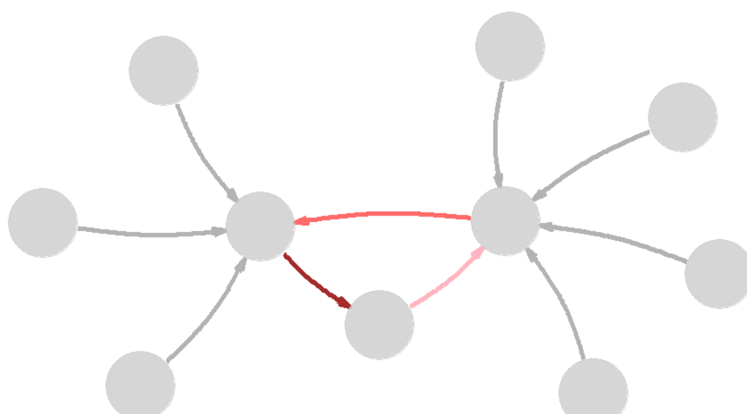

(b) Loop of edges with high betweenness.

Figure 4: Two examples where the edges with the highest betweenness in an unweighted, directed network are highlighted in shades of red. A darker color indicates a higher betweenness if multiple edges are colored red, and the arrows indicate the direction of the edges. In the top network, the edge acting as a bridge has a betweenness score of 25 and the other gray edges have a betweenness score of 6. In the bottom network, the edges within the loop have betweenness scores 13, 11, and 6, while the remaining gray edges each have a betweenness score of 3.

2.2.3 Dijkstra's algorithm

A way to find the shortest paths between all vertices in a network using the weights of the edges involved in the paths is through *Dijkstra's algorithm*, which is described in Grimaldi [13, p.631-634] and aims to find the shortest paths and their lengths from a source node to all other nodes in the graph. The algorithm is defined by the author for directed, weighted graphs and the weights are positive, real numbers for existing edges, and infinite for edges not present in the given graph. The length of the shortest path is the sum of the weights of the edges along this path and if no path exists between vertices i and j in V, we let this be equal to infinity. This is in contrast to our definitions of weighted graphs above, as we assign non-existing edges weight 0, but this is not an issue as we can adapt to the definition in our pruning methods by replacing the infinite values by 0. Furthermore, we also let the length of the shortest path from a node to itself be equal to 0. In this way, by letting $l : V \times V \to \mathbb{R}^+ \cup \{0, \infty\}$ be the function giving the length of the shortest path between vertices in V, we have that $l(i,i) = 0$, $l(i,j) = \infty$ for vertices $i, j \in V$ if they are not connected through a path from i to j, and $l(i,j) \in \mathbb{R}^+$ if such a path does exist.

Using this, the purpose of the algorithm is to find the shortest distances between a fixed vertex $v_0 \in V$ and all other vertices $i \in V$, as well as the directed paths from v_0 to i if the shortest distance is finite. We follow the description of the method by Grimaldi and hence, to do this, we first need some useful properties for the function l. We let $S \subset V$ be a set where the fixed vertex $v_0 \in S$ with $\bar{S} = V \setminus S$ being the set of the remaining vertices, and the distance from v_0 to the set $\bar{S}$ is then

$$l(v_0, \bar{S}) = \min_{i \in \bar{S}} l(v_0, i)$$

which means that the distance is given by the shortest directed distance to the closest vertex in the set $\bar{S}$, for finite values (asserting that such a vertex exists). It follows that

$$l(v_0, \bar{S}) = \min_{j \in S, k \in \bar{S}} \{l(v_0, j) + w_{jk}\}$$

since this means that the shortest path between v_0 and $\bar{S}$ is the minimum of the weight of the edge on the shortest single-edge path between j and $\bar{S}$ added to the shortest path between v_0 and j. Based on this, the algorithm first considers a single fixed vertex v_0 and starts with the sets $S_0 = \{v_0\}$ and $\bar{S}_0 = V \setminus S_0$ to compute

$$l(v_0, \bar{S}_0) = \min_{j \in S_0, k \in \bar{S}_0} \{l(v_0, j) + w_{jk}\} = \min_{k \in \bar{S}_0} w_{v_0 k}$$

where the last equality follows since $l(v_0, v_0) = 0$. For a vertex $v_1 \in \bar{S}_0$ with $l(v_0, v_1) = w_{v_0 v1}$ we now fix v_0 and v_1 such that $S_1 = S_0 \cup \{v_1\}$ and $\bar{S}_1 = V \setminus S_1$, and compute

$$l(v_0, \bar{S}_1) = \min_{j \in S_1, k \in \bar{S}_1} \{l(v_0, j) + w_{jk}\}.$$

In the next step, we fix $v_2 \in \bar{S}_1$ with $l(v_0, \bar{S}_1) = l(v_0, v_2)$ and continue this procedure until there are no additional vertices left to fix, either because all vertices have been visited or because there are no vertices left in V that are connected to v_0.

During the steps of the algorithm, the vertices are assigned labels where the final version of the label is of the form $(L(j), i)$ with $L(j) = l(v_0, j)$ being the shortest distance from v_0 to j, and i being the vertex located before j along this shortest path. This means that the edge from i to j is the last edge in this shortest path from v_0 to j. The label for the starting vertex is $(0, -)$ and the labels for the other vertices along the path are initially $(\infty, -)$ which are then updated, possibly multiple times, to the ultimate label $(L(j), i)$ for vertices connected to v_0.

The steps of the algorithm are summarized as

Step 1: Fix a vertex $v_0 \in V$ such that $S_0 = \{v_0\}$ and let the distance from v_0 to itself be 0, and the distances from v_0 to all other vertices in V be infinite. If there is only one vertex in the vertex set, the algorithm terminates here and otherwise it continues to the next step.

Step 2: For all other vertices j in V that are not fixed, update the distances using the label

$$L(j) = \min_{i \in S_i}\{L(j), L(i) + w_{ij}\}.$$

Step 3: If the fixed vertices cannot reach any of the other vertices, we are done. If not, there is at least one attainable vertex and we do the following:

1. Select a vertex v_{i+1} with a minimum label $L(v_{i+1})$. If there are multiple such vertices, we can choose either. This vertex is hence the closest to v_0 out of all the vertices that are yet to be visited.
2. Add the vertex v_{i+1} to the set of fixed vertices.
3. Continue until we have performed $|V| - 1$ iterations.

In summary, the algorithm thus starts at an initial source node, analyzes the graph by visiting all the nodes connected to this node and determining the currently shortest path between the source node and the remaining nodes. These values may be updated if a shorter path is determined in future steps, and the nodes are marked as visited once the shortest paths to these nodes have been established. The visited nodes are added to the current path and the process keeps going until all connected nodes have been visited. The result is that we have found the shortest paths from the source nodes to all other nodes in the network that this source node can reach.

As an example, we consider the network in Figure 5 which has 5 nodes labeled from 1 to 5, and labeled, weighted edges where the weight represents the distance between each node (i.e. a smaller weight indicates a shorter path). If we use node 1 as a source node, we observe that the shortest path to node 2 is by simply following the edge $1 \to 2$ to this node, which has weight 12, as this is the only possible way to reach this node. The current shortest path to node 3 is by following the edge $1 \to 3$ of weight 30. However, in the next step when fixing node 2, we note that the shortest path to node 3 is not the directed edge $1 \to 3$ from node 1 to 3, but a shorter path is actually to go from node 1 to 2, and then from node 2 to 3, as the length of the path $1 \to 2 \to 3$ is $10 + 12 = 22$ which is less than 30. These nodes are now marked as visited, as we cannot find shorter paths from node 1 to nodes 2 and 3. Additionally, when fixing node 2 in the next step, we find that the shortest path from node 2 to 4 is by first visiting node 3 and then following the directed edge to node 4, meaning that the path $2 \to 3 \to 4$ is shorter than $2 \to 4$ since the lengths of those paths are $10+30 = 40$ and 86, respectively. The shortest path from the source node to node 4 is thus $1 \to 2 \to 3 \to 4$ of length $12 + 30 + 10 = 52$. Finally, we fix node 3 to find that the shortest path from the source node to node 5 is by following the shortest path to node 3 and then using the edge $3 \to 5$, as this is the shortest distance between node 5 and the already visited nodes 2, 3, and 4. Hence, the last shortest path is $1 \to 2 \to 3 \to 5$ of length $12 + 10 + 15 = 37$.

This algorithm is used both to compute edge betweenness scores for the pruning methods described in Sections 4.1-4.3, and to find the shortest paths in the pruning method described in Section 4.4. We now shift our focus back to node centrality measures and describe another such measure that was inspired by internet webpages.

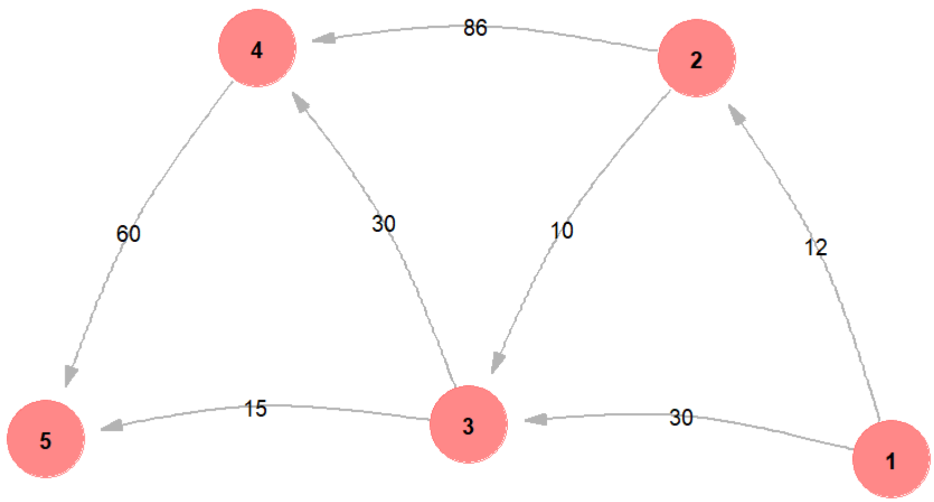

Figure 5: An example of a directed, weighted network used to illustrate Dijkstra's algorithm, where the labels of the edges illustrate the edge weights but we do not consider node weights for simplicity. Each arrow points in the direction of the directed edge and the different nodes are labeled from 1 to 5.

2.2.4 PageRank centrality

Another centrality measure is *PageRank*, which was introduced by Brin & Page [5, p.107-110] when they designed Google in 1998 to create an improved search engine with more accurate and useful search results. At the time, search engines were overwhelmed with the number of search results and were in need of a way to improve the quality of the first few displayed results by filtering and ranking their relevance. This, since users would only consider the first couple of results after entering a search query. To do this, they introduced PageRank, which ranks the quality and importance of the pages (or nodes in a network) by considering variations in link (or edge in a network) quality and by normalizing the values using the number of outgoing edges from a node. More formally, according to the abovementioned article, the PageRank of node i is recursively defined as

$$PR(i) = \frac{(1-\delta)}{N} + \delta \sum_{j \in V} \frac{a_{ji}}{d_j^{(\text{out})}} PR(j) \tag{3}$$

where $N = |V|$ is the total number of nodes, $d_j^{(\text{out})}$ is the number of edges going out from node j, a_{ji} is 0 if the directed edge $\{ji\}$ does not exist and 1 if it does, and $\delta \in [0, 1[$ is a damping factor. This damping factor is usually set to 0.85 and can be thought of as the probability that an edge between the current node and the next will be followed, while $1 - \delta$ is the probability of instead switching to another random node. The PageRanks correspond to the principal eigenvector of the normalized adjacency matrix and to the probability distribution of visiting each node, meaning that the sum of all PageRanks for a node is 1. Additionally, as the algorithm is recursively defined, the starting distribution of the nodes is uniform and iteratively updated by the PageRank values. However, the starting distribution can be personalized to be for instance a scaled version of the node weights that sums to 1 to incorporate these into the algorithm.

Furthermore, as mentioned in the article, PageRank can be interpreted as modeling the actions of a person using a search engine, where we assume that there is a random surfer who starts out at an arbitrary website, or node, and clicks on subsequent links until the user gets bored and decides to start from another random node. The damping factor δ represents the probability of the latter occurring, while the PageRank value of a page is the probability that the user visits this page. An advantage of using PageRank is that it marks a node as important if it has many incident edges from other nodes with high PageRanks, but also if there are only a few incident edges from other highly ranked nodes. This means that this centrality measure takes into consideration the importance of other nodes leading to the given node, rather than merely considering the number of incident links. Thus, a single link from another important node can increase the probability of reaching the node and hence give it a higher PageRank.

2.2.5 Weighted PageRank centrality

Zhang et al. [22, p.2] introduce *weighted PageRank* as an extension of the original PageRank algorithm involving weighted edges as well as weighted nodes. The weighted PageRank is defined recursively as

$$PR(i) = (1-\delta)\frac{\beta_i}{\sum_{i \in V} \beta_i} + \delta \sum_{j \in V} \left(\theta \frac{w_{ji}}{s_j^{(\text{out})}} + (1-\theta) \frac{a_{ji}}{d_j^{(\text{out})}} \right) PR(j) \tag{4}$$

where $\theta \in [0, 1]$ is an added parameter determining the relevance of the weights, which can be adjusted to preference. A motivation for including this parameter is that a balance between the number of links and the quality of the links may often be needed to be considered, and by adjusting this parameter one would alter the relationship in the algorithm. If for example $\theta = 0$, we only consider the number of incoming links as in the original version in Eq. (3), whereas we only consider the weights of the incoming links if $\theta = 1$. Additionally, the values of β_i for $i = 1, 2, ..., n$ are assumed to be independent of the edge weights and can be chosen to correspond to the relative importance of the nodes, using for example their node weights. By letting $\beta_i = 1$ for all nodes, we get that Eq. (4) is equivalent to Eq. (3) as we then let the nodes be chosen uniformly at random.

As stated by the authors, the random surfer model mentioned in Section 2.2.4 is an irreducible Markov chain when $\delta \neq 0$, since this means that each node can reach all the other nodes through the probability of switching to a random node, where each node is a state in the Markov chain. They show that this is also applicable to the weighted PageRank centrality, since the transition matrix $M := (m_{ij})$ would have entries

$$m_{ij} = \begin{cases} \theta \frac{w_{ji}}{s_j^{(\text{out})}} + (1-\theta)\frac{a_{ji}}{d_j^{(\text{out})}}, & \text{if } d_j^{(\text{out})} \neq 0 \\ \frac{\beta_i}{\sum_{i \in V} \beta_i}, & \text{if } d_j^{(\text{out})} = 0 \end{cases}$$

depending on whether you have reached a sink node or not, where

$$\sum_{i \in V} \left(\theta \frac{w_{ji}}{s_j^{(\text{out})}} + (1-\theta)\frac{a_{ji}}{d_j^{(\text{out})}} \right) = \theta + (1-\theta) = 1$$

and hence each column is a probability vector. If we let $R := (PR(1), PR(2), ..., PR(N))^T$ be the column vector of PageRank scores of each node, we can reformulate Eq. (4) as

$$R = \delta M R + (1-\delta)\hat{\beta} = (\delta M + (1-\delta)B)R =: \hat{M}R \tag{5}$$

for $\hat{\beta} = (\beta_1, \beta_2, ..., \beta_N)^T / \sum_{i \in V} \beta_i$ and B being an $N \times N$ matrix with the jth column given by $\beta_j / \sum_{i \in V} \beta_i \mathbf{1}$, since the entries in R sum to 1. Hence, the columns of B also sum to 1 and it follows that the columns of $\hat{M}$ must also sum to 1 as long as the entries of B are non-zero, and $\hat{M}$ is also a positive matrix. The authors thus note that the Perron-Frobenius theorem yields that the largest eigenvalue of $\hat{M}$ is 1, with eigenvector equal to the solution R of PageRank scores in Eq. (5). This, since the theorem states that non-negative square matrices have a unique principal eigenvalue with a corresponding eigenvector that can be chosen to be positive, and this eigenvalue is 1 for matrices with column sums of 1, according to Strang [19, p.314]. The normalized solution to Eq. (5) is thus a stationary distribution of the Markov chain with $\hat{M}$ as the probability transition matrix.

In this thesis, we use the weighted PageRank algorithm to compute centrality scores for the nodes in the pruning methods described in Sections 4.2-4.3. These scores are computed using the `page_rank` function from the `igraph` package in R, where the PageRank values are computed by solving the problem above using the PRPACK library, which, according to the documentation, is more efficient than the ARPACK library mentioned in the article as a solver of the eigenvalue problem in Eq. (5).

2.3 The role of networks in psychiatric symptomatology

In this thesis, the application of networks is in psychiatric symptomatology, and theoretical aspects of using network theory to conceptualize mental disorders have been developed by Borsboom [3] (used as a main source in this section), who describes the concept of using symptom networks to illustrate the dynamics of the interactions between different psychological symptoms. In the article, the term "mental disorder" is defined as a "syndromic constellation of symptoms that hang together empirically, often for unknown reasons", meaning that multiple symptoms cooperate and influence each other in different ways in mental disorders. An example of a problem mentioned in the article is social anxiety, which could affect a person's social life such that they feel lonely and sad. In this case, the symptom social anxiety leads to the symptoms loneliness and sadness, and a psychiatrist would be inclined to identify the source of the problem and possible treatment options. The author further explains how the direct causal links may be caused by

(a) basic biological processes (eg. insomnia causing fatigue),

(b) psychological processes (eg. losing interest could cause feelings of guilt),

(c) homeostatic couplings (eg. symptoms may be correlated, for instance appetite and sleep could be simultaneously affected as they are connected through biological processes), and

(d) societal norms (eg. engaging in illegal substance use could lead to juridical issues).

However, it is mentioned that finding certain symptoms that are common as the effects caused by the same underlying mental disorder may be an impossible task, and that the psychiatric symptoms might rather be causing each other. The example given is that delusion through believing that people can read your mind can lead to paranoia through suspicion, which may eventually lead to social isolation by abstaining from social interactions. This in turn could then further feed into the delusions from the lack of social contact and curation, which illustrates how a feedback loop is formed where a person is stuck in a scenario where the symptoms keep stimulating each other, causing a mental disorder to form.

This is where symptom networks come into play, as the interconnections between the symptoms can be visualized in a network, where each node represents a symptom and the links between the nodes represent the direct causal relationship between the symptoms at each end node. Using this representation, a mental disorder can then be conceptualized as the process of node activation, where the triggering of a symptom increases the likelihood of other connected symptoms to successively get triggered as well, and the mental disorder forms as a result of these activations being retained by feedback loops. These feedback loops are then represented by cycles present in the networks. Additionally, the networks of different mental disorders can be connected to form a single, joint network through symptoms common in both disorders, showing how multiple illnesses could interact. An example of this given by the author is how post-traumatic stress disorders can cause insomnia, which can lead to fatigue and impaired focus, and these symptoms are also connected to major depressive episode and generalized anxiety disorder.

Furthermore, there could also be factors outside the network in the *external field* that trigger the symptoms as, for example, a major life event such as losing your job could cause feelings of hopelessness and low self-esteem and this could further trigger depression in the person. Such an external factor could hence be a triggering cause of a mental disorder and this leads to the definition of *hysteresis*, which, according to the article, is when symptoms continuously stimulate each other in a strongly connected symptom network, despite the triggering event having passed. On the other hand, in a weakly connected network, hysteresis cannot take place as the symptoms are not independently active unless there is a feedback loop present and this is the case with for instance normal grief, according to the author. However, it could be possible for mental disorders to arise in weakly connected networks representing multiple disorders if there are bridge symptoms with strong connections and feedback loops involving the symptoms associated with a single mental illness.

Symptom networks give rise to the alternative definition of *mental health* as the equilibrium state of a weakly connected network, as stated by the author. This means that such a network can be disturbed with several symptoms being activated, but the system will eventually return to a state of mental health with the symptoms still present but dormant. Similarly, *mental disorder* can also be defined as the equilibrium state of a strongly connected network corresponding to the state of hysteresis. Based on this, diagnosis and treatment of mental disorders aim to identify such strongly connected systems and equilibrium states, and intervene to instead form weakly connected networks that can achieve the mentally healthy stable state. In particular, Borsboom states that diagnosis involves determining the existing symptoms and the network structures causing their activation, and one such network structure could be a cycle. Hence, clinicians are interested in identifying cycles in networks to find possible interventions that potentially break these cycles. Furthermore, treatment involves altering the network structures through

(a) symptom interventions (directly treat one or several nodes),

(b) interventions in the external field (remove triggering causes of nodes), and

(c) network interventions (target the edges between nodes to adjust the connections between symptoms).

An example of symptom intervention would be to prescribe antianxiety medication to a person suffering from anxiety caused by a build-up of stress from pressure at work, as this directly treats the anxiety, while an example of intervention in the external field would be to remove anxiety-inducing triggers by having the person take time off from work or shortening their working hours. Additionally, treatment in the form of network intervention could be to provide the person with the tools necessary to deal with the stress, such that this stops the person from experiencing anxiety.

Nevertheless, it is important to note that certain assumptions are made in the network representation of symptoms, and that it is not a perfect depiction of reality. For example, one assumes that the nodes and the edges between them are all known or possible to identify, but it could be more difficult to recognize these relationships in practice, especially for someone lacking a complete understanding of their own symptoms and the ability to express their interactions. Moreover, there could be certain mental disorders that do not fit the network structure, such as cyclic disorders or those developing slowly over time, and Borsboom mentions the example of autism where children avoiding eye contact could have difficulty learning social cues, reducing their ability to make social connections in the future. The author also states that this process could lead to complicated networks within networks through other processes of social interactions, and the framework may therefore not be suitable for this type of disorder.

Furthermore, Bringmann et al. [6] also outline procedural problems of using network theory, such as node selection and assessment, and managing modifications of the network structure over time. In this thesis, we consider symptoms as possible nodes, but the authors also put forward other possibilities such as theoretical causes not included in diagnosis criteria, and temporary states including for instance thoughts and feelings that are not considered symptoms. Furthermore, the chosen nodes should only contain those necessary to represent the chosen psychological process and exclude redundant nodes, while being dissimilar enough to be adequately represented by distinct nodes. This is related to how it should be possible to treat a single node without immediately directly treating other nodes simultaneously. However, it is important to limit the number of chosen nodes as to not create a network that is too cluttered, and one also has to limit the number of questions a participant is asked to reduce the participation load. In addition, Bringmann et al. specify that the psychological networks could change over time, and also that the edges, i.e. the connections between symptoms, could change in intensity based on if an individual is healthy or for instance experiencing depression. The example given is that lack of sleep might not cause sadness or anxiety in a healthy individual, but it could have more aggravating effects in someone close to experiencing a depressive episode.

Despite these problems and although the symptom networks are merely statistical models and not perfectly realistic, they can still be useful for clinicians and patients. A potential use of the symptom networks that is outlined by Bringmann et al. is to make predictions that can improve intervention methods for mental disorders. This can be done for individual patients by interpreting their personal network structures to understand the dynamics of their collection of symptoms, which is vital when seeking possible treatment measures. Nevertheless, the article notes that no study, as of the time of publishing the article (2022), has proven that existing network-based treatment measures are superior to the existing traditional measures and further research needs to be done in this area. Moreover, the authors also raise the issue of selecting the optimal edges or nodes to focus treatment on, and state that centrality measures can be used to identify these target points. However, it can be difficult to choose among the various centrality measures and the article states that there are some who argue that their usefulness is unclear in this application, and that focusing interventions on the most central nodes may not necessarily cause deactivation of connected nodes as desired. Furthermore, it could also be difficult to find possible treatment methods in practice for the identified node and the interventions could also end up not being effective enough to affect the network structure.

Nonetheless, the article suggests that one may expect that personalized networks are more useful than general networks in finding potential interventions, and it is crucial that the available networks are not too complex as this otherwise obstructs the visual and statistical analysis. Individual symptom networks could be more complicated compared to those for a general mental disorder, as an individual could suffer from more than one mental disorder at a time and hence experience several psychological symptoms. This motivates the aim of this thesis to find useful simplification methods that reduce the personalized networks into less complicated structures, while maintaining the most essential information. In this way, these methods aid clinicians and patients to interpret the individual networks to understand the processes and formalize possible treatments tailored to the individual's needs. These simplification methods are developed in Section 4.

3 Data

In this thesis, we use data comprising networks created based on the information reported by 21 individuals, where each person participated twice in a Swedish, self-reported survey within an interval of 2-7 days. The version of the survey used is new and unpublished at the time of writing this thesis, as it is a master's thesis by Kaariniemi & Bosund [16], but the approach is somewhat similar to the methods outlined and utilized in Klintwall, Bellander & Cervin [17]. In the version we use, participants are qualified to partake in the survey by being teenagers of both sexes having screened positive for depression. The survey asks the respondent to select at most 9 symptoms experienced within the past week from a list of possible symptoms. In this case, a symptom can be for example a behavior or emotion, such as overthinking or feelings of self-hatred. Each participant is able to choose between 17 symptoms, as well as add an 18th open item of their choice, which are listed in Swedish with their English translations in Table 12 in Appendix A. These symptoms always appear in this same order in the survey. Furthermore, the participants can also add a brief description of their experience with each chosen symptom.

Additionally, the participants are asked to rate the severity of each chosen symptom using values ranging from 0 to 100, where 0 means that the symptom is not painful at all and 100 corresponds to the highest painfulness. The respondents are then asked to report on the proportion of days each chosen symptom is experienced and how often each symptom is perceived to lead to the other available symptoms as a percentage. These interactions are represented by the possible directed edges that can be formed between the nodes representing the symptoms. The reported frequency related to each node is then multiplied by how often the interactions are experienced to occur. To illustrate, if a participant includes anxiety and insomnia in their list of experienced symptoms, and they report that anxiety is experienced 50% of the days of the week and that it further causes insomnia 10% out of those days, this means that the directed edge from the node "anxiety" to the node "insomnia" receives a weight of 5, since it is experienced 5% of the time.

The symptoms chosen by each participant in the survey are used to create a directed and weighted adjacency matrix $W := (w_{ij})$ for each individual, where each edge weight w_{ij} ranges from 0 to 100 and represents the reported frequency of interaction between the symptoms by the computation method described above, where 0 indicates no interaction and 100 indicates the highest frequency. Each column in the raw data matrix shows the symptoms experienced to be caused by the symptom represented by the column, and we hence transpose the matrix before applying our simplification algorithms such that the element w_{ij} instead displays how often symptom i is perceived to lead to symptom j, in accordance with the definitions in Section 2.1. The weights in the matrix are the edge weights in the corresponding network, and the painfulness scores of each symptom are similarly represented by the node weights in the resulting networks. Most networks have around 50 edges with positive weights (the mean number of edges is 51, and the median number of edges is 52), but it varies between as low as 4 up to 72, where all possible edges (excluding self-loops) have positive weights. Only two networks have a disconnected symptom included in the original versions, and most networks have the maximum number of nodes while only 5 networks have fewer than 9 nodes.

Moreover, we assess our pruning methods using human expert evaluations in Section 5.3 where five psychologists rate the pruning results. In this evaluation, we use three of the most reliable patients when considering the correlation between their full adjacency matrices from the first and second reported networks (we later decided to only use intersecting symptoms for the correlations computed in the coming evaluations, to increase similarity between the matrices), and those three patients are unknown to three of the psychologists who evaluate their results. We also include three additional patients not included in the 21 above whom the two other psychologists know, and who have completed the survey once. Only these two psychologists evaluate their results. The evaluation method is explained in more detail in Section 5.3.

The raw data for all 24 patients are accessible through the link in Appendix A and the corresponding networks are all visualized in the PDF files found by following the same link. The original networks are placed in the top left corner of each page, with the pruned results included as well, as described in the appendix.

4 Pruning methods

In this thesis, the main purpose of using edge pruning is to increase the interpretability of the networks by simplifying their structure, while maintaining as much as possible of the originally reported causation between symptoms by preserving the most important underlying structures in the networks. In the subsections below, we thus establish several methods to assess the quality of each edge, in terms of their importance for the overall network structure, to help us select which edges to prune. In particular, Sections 4.1 to 4.4 describe the pruning methods used, and Section 4.5 describes an abandoned attempt to add penalization of removing incoming edges to these methods. We later present three evaluation methods to assess and compare these established pruning methods in Section 5.

4.1 Method I: Edge betweenness

The first network simplification method evaluates the relevance of the edges using the centrality measure *edge betweenness* described in Section 2.2.2, to determine which edges are the least involved in the interactions between the nodes through their occurrence on the shortest paths in the network. Connecting this to the networks considered in this thesis, the edge betweenness scores rank each edge, or perceived relationship between symptoms, according to how often it is contained in the shortest paths between symptoms, which we define in terms of perceived frequency of triggering other symptoms and the painfulness associated with these symptoms. We thus let a high value of edge betweenness correspond to higher importance, as this means that the edge is prominent in the various paths between symptoms, and we prune the edges with the lowest edge betweenness scores iteratively until $n = \lceil \gamma(|E| - |V|) \rceil$ edges have been pruned, where $\gamma \in [0, 1]$ is a parameter that can be chosen to scale the level of pruning intensity. If $\gamma = 1$, the number of edges in the resulting network equals the number of nodes, while no edges are pruned if $\gamma = 0$. Initially, the intention was to prune edges until the number of remaining edges was the same as the number of nodes to hopefully have a connected pruned network without too many edges in the end result. However, it is not completely obvious that this is the most reasonable cutoff, and we hence include the parameter γ to allow for extra edges if this yields superior outcomes. For example, by setting $\gamma = 0.95$ instead of $\gamma = 1$ for a network with 10 nodes and 50 edges, we keep 12 edges instead of only 10 which allows for more information in the end results while maintaining a simple enough structure.

Moreover, to account for the given node weights and edge weights indicating the perceived painfulness of each symptom and how often they are experienced to influence each other, we transform the edge weights through multiplication with the node weight of the end node. This, since the symptom the patient experiences to be caused by the existence of the edge is considered more important than the symptom causing the edge, as one potential use of the networks is to determine candidate symptoms to focus treatment on. In that case, the symptoms that are subsequently triggered if the selected node is left untreated would be more relevant, as the psychologist may desire to treat as many of the most severe symptoms as possible by instead focusing on only a select few nodes. Furthermore, we normalize the respective weights through dividing by their sum prior to the multiplication to ensure that the node weights and edge weights are both in the range of $[0, 1]$, and to avoid that one of them outweighs the other. Thus, the new directed edge weights in the network are

$$w_{ij}^{(new)} = \left(\frac{v_j}{\sum_{k \in V} v_k} \right) \left(\frac{w_{ij}}{\sum_{k,l \in V, \{kl\} \in E, k \neq l} w_{kl}} \right), \tag{6}$$

where i and j are nodes in V and we only sum over the existing edge weights. It is important to keep in mind that this assumes that the relationship between the node weights and edge weights is linear, which might actually not be a realistic assumption if the painfulness and frequency of the symptoms are not considered to be of equal importance, and this relationship could be tuned to preference.

These modified edge weights are then used in the edge betweenness computations to determine which edges to prune. However, since the new edge weights in Eq. (6) indicate connection strengths between the symptoms through their perceived frequency of causality and experienced painfulness of the affected symptoms rather than distances, where large weights indicate strong connections but also long distances, we must use a transformed version of our edge weights in the computations to take this into account if we want short paths to be defined by strong connections. This, since the algorithm used in the edge betweenness computations for finding the shortest paths considers the edge weights as distances, and we would thus find the *longest* paths with the weakest connections when using our new weights $w_{ij}^{(new)}$. Therefore, we instead use

$$1 - w_{ij}^{(new)} = 1 - \left(\frac{v_j}{\sum_{k \in V} v_k} \right) \left(\frac{w_{ij}}{\sum_{k,l \in V, \{kl\} \in E, k \neq l} w_{kl}} \right) \tag{7}$$

as the edge weights when assigning edge betweenness scores. These transformed weights are now equivalent to our definition of shorter paths through connection strengths, as low values of the weights in Eq. (7) yield high values of $w_{ij}^{(new)}$, meaning that shorter paths are more likely to involve symptoms that frequently trigger other painful symptoms. This particular transformation is chosen to keep a linear relationship between the length of the shortest paths and the connection strength of the paths in terms of the new edge weights, but could also be adjusted to preference. Moreover, it is possible that several edges have the same betweenness scores, but since this means that the edges are of equal importance for the network flow in terms of edge betweenness, we simply select any of the edges to prune if there are multiple edges with the lowest betweenness scores. However, this might be too simple and one could generalize this to include another feature for edge selection and use this together with the edge betweenness scores to hopefully get unique scores for each edge.

In summary, the pruning method described in this section utilizes edge betweenness scores computed using a transformation of the combined node and edge weights of a weighted, directed graph $G = (V, E)$ to obtain a simplified subgraph $H = (V, F)$, such that $F \subset E$ with $|F| = |E| - n$. Since this method uses a version of Dijkstra's algorithm to find the shortest paths needed to compute edge betweenness for the edges, as described in Section 2.2.2, this means that the time complexity for this method is $\mathcal{O}(n|V|(|E|+|V|\log|V|))$ since we do this in each of the n iterations. This pruning method is summarized as pseudocode in Algorithm 1.

Nevertheless, this method is quite simple in its handling of the node weights, and it may be useful to be more considerate about the *node* centrality as well when deciding which edges to prune, as this is also indicative of the network structure. One may be especially interested in accounting for the importance of nodes leading to a given symptom when considering the structural importance of the symptom itself. Hence, we incorporate a version of this in the following method using *PageRank* centrality.

Algorithm 1 Pruning with edge betweenness

Input: A weighted, directed graph $G = (V, E)$, and the parameter $\gamma \in [0, 1]$.
Output: A pruned, weighted, and directed subgraph $H = (V, F)$ such that $F \subset E$.

$n \leftarrow \lceil \gamma(|E| - |V|) \rceil$
$F \leftarrow E$
$w_{ij}^{(new)} \leftarrow \left(\frac{v_j}{\sum_{k \in V} v_k}\right) \left(\frac{w_{ij}}{\sum_{k,l \in V, \{kl\} \in E, k \neq l} w_{kl}}\right)$ ▷ Set new edge weights of edges in F.
while $|F| > |E| - n$ **do**
    {Compute edge betweenness $B(e)$ with $1 - w_{ij}^{(new)}$ as edge weights.}
    {Rank edges in order of betweenness and find the edge e_min with the lowest score.}
    $F \leftarrow F \setminus e\_min$
end while
return $H = (V, F)$

4.2 Method II: PageRank

To include more of the structural information of the network in the edge weights instead of only using the perceptions of painfulness reported by the patient, we now add a feature to the node weights in the previous method. In this method, we take into consideration the importance of each symptom in relation to its connection to other symptoms by using the weighted centrality measure *PageRank* described in Section 2.2.5 for each of the nodes, prior to transforming the edge weights in the graph and using edge betweenness to rate the importance of each edge. In this way, we weight each node according to how important the other symptoms leading to this symptom are, and a symptom is considered more important if many other important symptoms can trigger it. We start by computing the PageRank scores for each node, with the normalized original node weights $v_i / \sum_{j \in V} v_j$ as the starting values $\beta_i / \sum_{j \in V} \beta_j$ to incorporate the painfulness of each symptom in the algorithm. This means that if we encounter a sink node, we are more likely to restart the PageRank algorithm at a more painful symptom. However, this could be considered unrealistic since it may not be true in general that more painful symptoms are more likely to occur, but in this way we ensure that these symptoms are given priority in the algorithm, which is crucial as symptoms with higher painfulness could be more severe and therefore more important to keep connected in the network to track the relationships to other nodes.

Additionally, the original edge weights w_{ij} are also included as a measure of connection strength, and we keep the damping factor $\delta = 0.85$ and let $\theta = 1$ to consider the edge weights when computing the weighted PageRank scores. As in Method I in Section 4.1, the edge weights in the network are transformed to incorporate the importance of the end nodes, but we now use the PageRank value of the end node of each edge instead of the original node weight. Since the PageRank scores represent a probability and are already in the interval $[0, 1]$, we do not need to normalize them and the scores are thus multiplied by the corresponding normalized edge weights to create the new edge weights in the network.

These are thus given by

$$w_{ij}^{(new)} = PR(j) \left(\frac{w_{ij}}{\sum_{k,l \in V, \{kl\} \in E, k \neq l} w_{kl}} \right)$$

where $PR(j)$ is the PageRank score of the end node j which is computed according to the weighted PageRank algorithm described in Section 2.2.5. Similar to in the first method, we use these new weights to compute the edge betweenness scores for each edge and rank them in decreasing order in order to prune the edge with the lowest betweenness, and the new edge weights are again first transformed to ensure that they are equivalent to lengths rather than connection strengths. The weights used in the edge betweenness computations are hence

$$1 - w_{ij}^{(new)} = 1 - PR(j) \left(\frac{w_{ij}}{\sum_{k,l \in V, \{kl\} \in E, k \neq l} w_{kl}} \right).$$

Moreover, we again prune the lowest scoring edges until we have removed $n = \lceil \gamma(|E| - |V|) \rceil$ edges in total from the directed, weighted graph $G = (V, E)$ for $\gamma \in [0, 1]$, and are left with a subgraph $H = (V, F)$ where $F \subset E$ and $|F| = |E| - n$. The method is summarized as pseudocode in Algorithm 2.

Algorithm 2 Pruning with edge betweenness and PageRank

Input: A weighted, directed graph $G = (V, E)$, and the parameter $\gamma \in [0, 1]$.
Output: A pruned, weighted, and directed subgraph $H = (V, F)$ such that $F \subset E$.

```
n ← ⌈γ(|E| − |V|)⌉
F ← E
θ ← 1                                   ▷ Parameter for PageRank to only use edge weights.
δ ← 0.85                                ▷ Damping factor.
β_i / Σ_{j∈V} β_j ← v_i / Σ_{j∈V} v_j   ▷ Original node weights as starting values for PageRank.
{Compute PageRank scores PR(i) using personalized β_i.}
w_ij^(new) ← PR(j) ( w_ij / Σ_{k,l∈V,{kl}∈E,k≠l} w_kl )   ▷ Set new edge weights of edges in F.
while |F| > |E| − n do
    {Compute edge betweenness B(e) with 1 − w_ij^(new) as edge weights.}
    {Rank edges in order of betweenness and find the edge e_min with the lowest score.}
    F ← F \ e_min
end while
return H = (V, F)
```

However, this pruning method only utilizes the PageRank scores of the original network, but as we remove edges, the PageRank scores could change depending on the transformed structure of the new network. Consequently, we establish the next method below, which incorporates the updated structure of the networks after pruning by computing new PageRank values in each step.

4.3 Method III: Updated PageRank

Now, to take into consideration how the relative importance of the nodes may change after removing edges, we alter Method II in Section 4.2 above such that the PageRank scores for each edge are updated in each step after an edge is removed. The starting values are given by the normalized original node weights in the first iteration as in the previous section, while the subsequent iterations use the previous PageRank scores as the current starting values. As all these values sum to 1, they can be considered probabilities of the algorithm starting over at each node when a sink node is encountered. Hence, we redefine this probability according to the importance established in the previous iteration for each node, meaning that we define the probability of activation of a different symptom when getting stuck in a sink node according to the importance of each symptom. As in the previous method, the PageRank values are then multiplied with the normalized original edge weights to get $w_{ij}^{(new)}$, and the weights $1 - w_{ij}^{(new)}$ are then used to compute the edge betweenness scores in each step. The only difference is thus that the PageRank scores are updated in each iteration using the previous PageRank values as the prior information β_i, and the final edge weights in the pruned subgraph are the values of $w_{ij}^{(new)}$ obtained in the last iteration. The method is otherwise the same as Method II, such that we prune a directed, weighted graph $G = (V, E)$ until $n = \lceil \gamma(|E| - |V|) \rceil$ edges are removed with $\gamma \in [0, 1]$, and we have obtained a simplified subgraph $H = (V, F)$ with $F \subset E$ and $|F| = |E| - n$. This updated version is summarized as pseudocode in Algorithm 3.

Algorithm 3 Pruning with edge betweenness and updated PageRank

Input: A weighted, directed graph $G = (V, E)$, and the parameter $\gamma \in [0, 1]$.
Output: A pruned, weighted, and directed subgraph $H = (V, F)$ such that $F \subset E$.

$n \leftarrow \lceil \gamma(|E| - |V|) \rceil$
$F \leftarrow E$
$\theta \leftarrow 1$ ▷ Parameter for PageRank to only use edge weights.
$\delta \leftarrow 0.85$ ▷ Damping factor.
$\frac{\beta_i}{\sum_{j \in V} \beta_j} \leftarrow \frac{v_i}{\sum_{j \in V} v_j}$ ▷ Original node weights as starting values for first PageRanks.
$w_{ij} \leftarrow \frac{w_{ij}}{\sum_{k,l \in V, \{kl\} \in E, k \neq l} w_{kl}}$ ▷ Normalize original edge weights.
while $|F| > |E| - n$ **do**
  {Compute PageRank scores $PR(i)$ using personalized and updated β_i.}
  $w_{ij}^{(new)} \leftarrow PR(j) w_{ij}$ ▷ Update edge weights.
  {Compute edge betweenness $B(e)$ with $1 - w_{ij}^{(new)}$ as edge weights.}
  {Rank edges in order of betweenness and find the edge e_min with the lowest score.}
  $F \leftarrow F \setminus e\_min$
  $\beta_i \leftarrow PR(i)$ ▷ Use previous PageRank scores as starting values for next iteration.
end while
$w_{ij}^{(new)} \leftarrow PR(j) w_{ij}$ ▷ Set new edge weights of edges in F using final PageRanks.
return $H = (V, F)$

4.4 Method IV: Connectivity kept

A different way to monitor the connectedness of the network and the quality of its paths, without relying on the centrality measures used in Methods I-III, is to consider the relative change of the network quality in each pruning step. Accordingly, we now present another method which is inspired by the "brute force" edge pruning approach in Zhou, et al. [23], where a quantity referred to as the *ratio of connectivity kept*, or rk, is computed in each step and the edge producing the highest ratio of connectivity kept upon removal is pruned. The original "brute force" method from the article is described first, and then we modify it to fit our networks and pruning objectives in the subsequent section.

4.4.1 Original version

In the article, Zhou et al. [23] consider undirected, weighted networks, in contrast to the directed networks in our data. The authors define a positively real-valued *path quality function* $q(P)$ for a path P which is used to compute the connectivity of the graph. The choice of this function depends on the related networks and the pruning objectives, and the authors mention that for instance the probability of the existence of a path may be used in random graphs, while the length of a path might be used for the shortest paths in weighted networks. Furthermore, we can examine the *connectivity between two* ***vertices*** i and j to evaluate how strongly they are connected by an edge or a path, which is given by

$$C(i, j; E) = \begin{cases} \max_{P \subset E: i \overset{P}{\rightsquigarrow} j} q(P), & \text{if } P \text{ exists} \\ -\infty, & \text{otherwise} \end{cases}$$

for the paths P between i and j. In this way, we let the connectivity be the value of the path quality function for the best path, if there is such a path. From this, we can define the *connectivity of a* ***graph*** $G = (V, E)$ as

$$C(V, E) = \frac{2}{|V|(|V| - 1)} \sum_{i,j \in V, i \neq j} C(i, j; E), \tag{8}$$

which is the average connectivity for all possible edges formed by selecting pairs of nodes, where self-loops and multiple edges are not allowed, since there are $\frac{|V|(|V|-1)}{2}$ possible undirected pairs. In the article, the graph is assumed to be connected to ensure that $C(V, E) > 0$, and if not, each connected component is pruned separately and hence this assumption can be made without limiting the method to connected input networks. Furthermore, the rk value is defined as

$$rk(V, E, E_R) = \frac{C(V, E \setminus E_R)}{C(V, E)} \tag{9}$$

when removing a set E_R of edges. The ratio represents the relative change in connectivity of a graph after pruning the edges in E_R, where $rk = 1$ indicates no change in connectivity, $0 < rk < 1$ indicates a decrease in connectivity, and $rk = -\infty$ means that the graph has become disconnected.

Using the above, the purpose of the approach is to prune edges while minimizing the decrease in connectivity indicated by the value of rk, and we hence want to keep this ratio as close to 1 as possible to try to maintain the connectivity even after removing links between nodes. Since the ratio reduces to $-\infty$ if the graph becomes disconnected, the authors only prune those edges that do not cut the graph by storing these in a set M and only computing rk for edges not in this set. Hence, the number of potential edges subject to pruning ranges from $|V| - 1$ to $|E|$, as $|V| - 1$ is the minimum number of edges necessary for the graph to be connected, and the authors thus specify that $n = \lceil \gamma(|E| - (|V| - 1)) \rceil$

edges are to be removed by the algorithm, where the parameter $\gamma \in [0,1]$ is set to allow for different pruning intensities. If $\gamma = 1$, the graph is maximally pruned and if $\gamma = 0$, no pruning occurs. Consequently, the approach prunes a weighted, undirected graph $G = (V, E)$ into a simpler subgraph $H = (V, F)$, where $F \subset E$ and the number of pruned edges is $n = \lceil \gamma(|E| - (|V| - 1)) \rceil$, by maximizing $rk(V, E, E \setminus F)$ given a path quality function q and the pruning intensity parameter γ.

Specifically, the approach first checks that the graph is connected and then calculates $rk(V, F, e)$ for every edge e in $F \subset E$. The value for and identity of the edge whose $rk(V, F, e)$ value is currently the highest are saved in each iteration. After having gone through the entire list of edges, the edge whose ratio of connectivity kept is the highest is pruned, unless the graph has become disconnected. In that case, the edge is added to the set M and is ignored in the following iterations. To compute the ratio of connectivity kept, we need to compute the path quality function in each step, which can be time-consuming for large graphs and this is the reason why the authors named this the "brute force" algorithm. In particular, they mention that the cost of computing the best paths is $\mathcal{O}(|V|(|E| + |V|)\log|V|)$, meaning that the total time complexity is $\mathcal{O}(n|E||V|(|E| + |V|)\log|V|)$ inside the loop where the ratio of connectivity kept is computed n times for each edge. This requires more time than computing the edge betweenness scores using the modified version of Dijkstra's algorithm in each iteration in Methods I-III, as this has time complexity $\mathcal{O}(n|V|(|E| + |V|\log|V|))$. Hence, they also present alternatives that are faster but also less accurate in evaluating the quality of the networks after pruning in each step. However, since the networks in our data are quite small, we choose to implement this less efficient but more accurate method rather than its alternative approaches, as we value accuracy over a (in this case) smaller increase in algorithm execution speed. This original version is summarized as pseudocode in Algorithm 4, and is slightly modified in the section below to fit our needs and purposes.

4.4.2 Modified version

The version of the original "brute force" approach in Zhou et al. [23] described in the previous section that is used in this thesis is a modified version that is more compatible with our networks. In particular, instead of undirected graphs, we consider directed, weighted graphs and hence a path P exists between vertices i and j if and only if a sequence of directed edges lead from i to j. This means that the connectivity of a directed graph $G = (V, E)$ is

$$C(V, E) = \frac{1}{|V|(|V| - 1)} \sum_{i,j \in V, i \neq j} C(i, j; E),$$

as there are $|V|(|V| - 1)$ possible directed edges to be formed, if no self-loops or multiple edges are allowed. Hence, the sole difference between this and the connectivity of an undirected graph in Eq. (8) is a factor of 2. Moreover, we want our modified version to allow for disconnected components to be less restrictive than the original version and to be comparable to Methods I-III. Therefore, we define the connectivity for vertices as

$$C(i, j; E) = \begin{cases} \max_{P \subset E : i \overset{P}{\rightsquigarrow} j} q(P), & \text{if } P \text{ exists} \\ 0, & \text{otherwise} \end{cases}$$

since in that way, the value of the ratio of connectivity kept cannot drop to $-\infty$ if there is no path between i and j in that direction, or in any direction (which corresponds to the graph being disconnected). This will result in lower connectivity overall if the graph

Algorithm 4 Pruning with connectivity kept (original "brute force")

Input: A weighted, undirected graph $G = (V, E)$, a positively real-valued path quality function q, and a parameter $\gamma \in [0, 1]$

Output: A pruned, weighted, and directed subgraph $H = (V, F)$ such that $F \subset E$

$n \leftarrow \lceil \gamma(|E| - (|V| - 1)) \rceil$
$F \leftarrow E$
$M \leftarrow \emptyset$ ▷ Edges known to disconnect graph upon removal.
while $|F| > |E| - n$ **do**
    $rk\_largest \leftarrow -\infty$
    $e\_largest \leftarrow null$
    for edge $e \in F$ such that $e \notin M$ **do**
        if subgraph $(V, F \setminus \{e\})$ connected **then**
            {Compute $rk(V, F, e)$.}
            if $rk(V, F, e) > rk\_largest$ **then**
                $rk\_largest \leftarrow rk(V, F, e)$
                $e\_largest \leftarrow e$
            end if
        else
            $M \leftarrow M + e$
        end if
    end for
    {Prune edge of H with the highest rk value.}
    $F \leftarrow F \setminus \{e\_largest\}$
end while
return $H = (V, F)$

ends up getting disconnected, but this is hopefully avoided when trying to keep the ratio of connectivity kept as high as possible, which is defined in the same way as in Eq. (9).

Additionally, we let the path quality function $q(P)$ be 1 divided by the length of the directed path P between i and j, if such a path exists, as a high value indicates a better and shorter path. For the same reasons as in Method I, the lengths of the paths are computed as the sum of 1 minus the new edge weights of the edges comprising the path using Eq. (7), where the new edge weights are the normalized original edge weights multiplied with the normalized original node weights. Hence, this method aims to maintain the overall connection strength in the network when removing an edge in each step, meaning that the lengths of the shortest paths between symptoms should be minimally affected when removing a link. Furthermore, Dijkstra's algorithm, described in Section 2.2.3, is used to find the shortest paths between each node, since the length of the shortest path will maximize the path quality function. Thus, this has the same time complexity as the original version. The difference between removing edges according to their ratios of connectivity kept and doing so according to their edge betweenness scores, is that we for each edge $e \in E$ in each step in the first case consider the relative change in the sum of the reciprocal lengths of the shortest paths between each symptom upon removing e, while we in the latter case consider the number of times the edges occur on these shortest paths. This, since

$$\begin{aligned} C(V, E \setminus \{e\}) &= \frac{1}{|V|(|V|-1)} \sum_{i,j \in V, i \neq j} C(i,j; E \setminus \{e\}) \\ &= \frac{1}{|V|(|V|-1)} \sum_{i,j \in V, i \neq j} \max_{P \subset E \setminus \{e\}: i \overset{P}{\rightsquigarrow} j} \frac{1}{\text{length}(P)} \\ &= \frac{1}{|V|(|V|-1)} \sum_{i,j \in V, \{ij\} \in E \setminus \{e\}, i \neq j} \frac{1}{L(i,j)} \end{aligned}$$

for the shortest path $L(i,j)$ between nodes i and j, and in the ratio of connectivity kept, this becomes

$$\begin{aligned} rk(V, E, E \setminus \{e\}) &= \frac{C(V, E \setminus \{e\})}{C(V, E)} \\ &= \frac{\sum_{i,j \in V, i \neq j} C(i,j; E \setminus \{e\})}{\sum_{k,l \in V, k \neq l} C(k,l; E)} \\ &= \frac{\sum_{i,j \in V, \{ij\} \in E \setminus \{e\}, i \neq j} 1/L(i,j)}{\sum_{k,l \in V, \{kl\} \in E, k \neq l} 1/L(k,l)}. \end{aligned}$$

Therefore, when choosing the edge to remove in this method, we consider the overall effect on the network structure through how the removal of an edge affects the summed quality of all paths using the ratio of connectivity kept. When using edge betweenness for pruning, we instead consider the frequency of the inclusion of the individual edges on the shortest paths in the network. Hence, this method minimizes the loss in connection strength between symptoms, while the approaches based on edge betweenness maximize the centrality of the existing edges.

Finally, we also let $n = \lceil \gamma(|E| - |V|) \rceil$ as in Methods I-III, with $\gamma \in [0, 1]$ such that this simplification method also prunes a directed, weighted graph $G = (V, E)$ to produce a subgraph $H = (V, F)$ with $F \subset E$ and $|F| = |E| - n$. The algorithm is summarized as pseudocode in Algorithm 5, and the method is renamed for clarity.

Algorithm 5 Pruning using connectivity kept (modified "brute force")

Input: A weighted, directed graph $G = (V, E)$, a positively real-valued path quality function q, and a parameter $\gamma \in [0, 1]$

Output: A pruned, weighted, and directed subgraph $H = (V, F)$ such that $F \subset E$.

$n \leftarrow \lceil \gamma(|E| - |V|) \rceil$
$F \leftarrow E$
$w_{ij}^{(new)} \leftarrow \left(\frac{v_j}{\sum_{k \in V} v_k} \right) \left(\frac{w_{ij}}{\sum_{k,l \in V, \{kl\} \in E, k \neq l} w_{kl}} \right)$ ▷ Set new edge weights of edges in F.
while $|F| > |E| - n$ **do**
  $rk\_largest \leftarrow -\infty$
  $e\_largest \leftarrow null$
  for edge $e \in F$ **do**
    {Compute $rk(V, F \setminus \{e\})$ with $1 - w_{ij}^{(new)}$ as edge weights and $q(P) = \frac{1}{\text{length}(P)}$.}
    if $rk(V, F \setminus \{e\}) > rk\_largest$ **then**
      $rk\_largest \leftarrow rk(V, F \setminus \{e\})$
      $e\_largest \leftarrow e$
    end if
  end for
  {Prune edge of H with the highest rk value.}
  $F \leftarrow F \setminus \{e\_largest\}$
end while
return $H = (V, F)$

4.5 Attempt to penalize start nodes

A potential issue with the above established pruning methods is that they allow for the resulting networks to create new start nodes, which is undesirable as this means that the symptoms represented by these nodes are not shown to be affected by the remaining symptoms, even if this were the case in the original network. This hinders the determination of intervention methods based on the pruned networks that target these start nodes through the treatment of adjacent nodes, if one is unable to directly treat these symptoms. For this reason, we looked into ways to penalize the creation of new start nodes and instead prioritize certain edges to ensure we do not lose the incoming connections from other vertices. The first way we did this was by completely forbidding the emergence of new start nodes by checking in each iteration of the pruning algorithms if a new start node has been created after the removal of an edge, and if so, the edge was reattached to the graph and stored in a set of edges to avoid pruning in the next steps, and we started over with the next potential edge to remove. If no new start node was created in that iteration we simply removed the edge according to the algorithms and moved on to the next possible edge to prune.

However, strictly forbidding the appearance of new start nodes was deemed too harsh as there are situations where start nodes may be reasonable in a symptom network, for example if a symptom is truly not significantly affected by the other symptoms in the network or if it is affected by external causes that are not represented in the network structure. It could thus be useful to keep some such nodes, if the removed edges causing the creation of the start nodes are not central to the overall network structure and would otherwise only be kept to avoid an unwanted start node. This could also mean that a more important edge that affects more severe symptoms, or that plays an integral part in the activation of other structurally important nodes, is removed for the sake of avoiding an unwanted structure merely because it makes it difficult to treat that symptom alone. However, the start node could represent a serious symptom that is crucial but difficult to treat directly, and we therefore might want to avoid the removal of incoming edges to such a symptom even if these edges are less central to the structure itself.

Therefore, we wanted to relax the restriction to a penalization rather than a strict prohibition, where the creation of new start nodes is tolerated, but discouraged. In particular, one may for example want to do this in proportion to the severity of the potential start node or to the edge weights of the removed incoming edges. One way we did this was by checking in each iteration if the removed edge created a new start node, and adding weight to this particular edge if this were the case to increase its importance and hopefully not prune it if we add it back and reiterate the step. This, since we in Methods I-III use edge betweenness to decide which edges to prune and edges with larger weights are interpreted as shorter paths when computing the edge betweenness scores. In Method IV, we also use the shortest paths to decide which edges to prune as we in each iteration remove the edge whose impact on the ratio of connectivity kept is the smallest. However, this method of increasing the weights proved to be ineffective, since the resulting methods still created new unwanted start nodes. This could be due to how the edges being removed as a result of having low edge betweenness scores or low effect on the ratio of connectivity kept are not highly interconnected within the rest of the network, as they are not found on a high number of the shortest paths. They are more likely to be found in the outermost structures of the graphs and, hence, increasing their weights does not increase their edge betweenness scores or effect on the connectivity even if the paths the edges are present on would become shorter.

Since the above methods were not useful, we decided against using penalization. Furthermore, another reason why we decided not to use penalization is that adding a cost to the creation of start nodes somehow abuses the data provided by the patient, as this disregards the originally reported relations between symptoms. It means that we give more importance to certain relations for the sake of keeping some incoming edges, rather than solely focusing on the interconnectedness of each edge in terms of their relevance for the network structure as evaluated by for example the centrality measures. Therefore, we decided that it may be unreasonable to penalize the creation of new start nodes, as it may very well be the case that the originally reported incoming edges that end up getting removed are not as vital for the overall network structure. This, despite the new start nodes complicating the choice of treatment methods. However, the treatability of nodes could be incorporated in different ways, as mentioned in the discussion in Section 8.

5 Evaluation methods

In order to compare and assess the quality of the above pruning methods, we need to test their ability to retain the key components, and their practical usefulness. Hence, we use three different evaluation approaches that are based on

(a) comparing the outcomes for two empirical networks reported by the same patient on different days through **repeated trials**,

(b) the stability from **added noise** to synthetic networks, and

(c) the visual quality of the resulting networks from empirical data, as reported by psychologists through **human expert evaluation**.

These assessment criteria test the robustness and usefulness of the methods, which are important features since the main goal of the pruning methods is to reduce the networks to reasonable simplifications that are visually more straightforward to interpret and analyze. The details of the evaluation techniques are outlined in the sections below.

5.1 Repeated trials

A way to assess the precision of the pruning methods is to analyze the similarity between pruning results for two empirical networks from the same patient. In the data, patients have completed the survey twice with a 2-7-day interval, and we are interested in how similar the pruned versions of these networks are. Since the participants are able to choose different symptoms when repeating the survey, we restrict this analysis to the coinciding chosen nodes for the first and second trial. The similarity is analyzed through the correlation of the directed, weighted adjacency matrices, through the number of edges the networks have in common, as well as by if the out-degree centrality rates the same symptom as the most central in both networks. To evaluate the correlation, we use the definitions in Alm & Britton [1, p.398-399] to compute the *Pearson* and *Spearman correlation coefficients* for the respective pruned networks for all simplification methods, and for the original networks to determine whether networks are more similar after pruning. The Pearson correlation coefficient is defined for the adjacency matrices as

$$r(X,Y) := \frac{\sum_{i,j=1,i\neq j}^{n}(x_{ij}-\bar{x})(y_{ij}-\bar{y})}{\sqrt{\sum_{i,j=1,i\neq j}^{n}(x_{ij}-\bar{x})^2}\sqrt{\sum_{i,j=1,i\neq j}^{n}(y_{ij}-\bar{y})^2}},$$

where $X = (x_{ij})_{i,j=1,2,\dots,n}$ and $Y = (y_{ij})_{i,j=1,2,\dots,n}$ are the directed, weighted adjacency matrices of size $n \times n$ for the networks corresponding to the first and second trial, with means $\bar{x}$ and $\bar{y}$, respectively. In the computations, we do not sum over $i = j$, as we do not have any self-loops. Similarly, the Spearman correlation coefficient is computed in the same way, but with X and Y replaced by the ranks of the entries in X and Y from 1 to $n \times n$, such that $s_{ij} := \text{rank}(x_{ij})$ and $t_{ij} := \text{rank}(y_{ij})$. Hence, this coefficient is given by

$$r(S,T) := \frac{\sum_{i=1,i\neq j}^{n}(s_{ij}-\bar{s})(t_{ij}-\bar{t})}{\sqrt{\sum_{i=1,i\neq j}^{n}(s_{ij}-\bar{s})^2}\sqrt{\sum_{i=1,i\neq j}^{n}(t_{ij}-\bar{t})^2}}$$

where $S = (s_{ij})_{i,j=1,2,\dots,n}$ and $T = (t_{ij})_{i,j=1,2,\dots,n}$ are the matrices containing the ranks of the entries of the adjacency matrices, and $\bar{s}$ and $\bar{t}$ are the means of S and T, respectively. This definition is applicable even if we have tied events. Moreover, to evaluate how many edges the network pairs have in common, we use Eq. (7.4) in Kolaczyk & Csárdi [18] to compute the *Jaccard coefficient* which is defined as

$$J(E_1, E_2) := \frac{|E_1 \cap E_2|}{|E_1 \cup E_2|}$$

for the sets of edges E_1 and E_2 of the networks reported in the first and second trial, respectively.

The first two coefficients give a measure of how correlated the networks are in terms of their existing edges and their corresponding edge weights, whereas the Jaccard coefficient gives a ratio of the number of edges the pruned networks have in common compared to how many distinct edges they have altogether. In particular, the Pearson correlation coefficient assesses how linear the relationship between the two matrices is, while the Spearman correlation coefficient measures their monotonic relationship, since we only consider the ranks of the entries and not their actual values.

Furthermore, the out-degree centrality measure described in Section 2.2.1 is utilized to find the most central nodes in each of the network pairs, out of the nodes chosen in both pairs. We use this to compare how many of the network pairs have the same most central node. This is again done for the pruned networks and the original networks to determine whether pruning yields more similar networks.

Using these measures, we hence compare the outcomes for empirical networks observed over time where the underlying network structures are not known beforehand, and we therefore do not have a known network to compare the pruned results to. The goal is thus to compare how similar the pruning results are for two noisy networks whose connection is that they are reported by the same patient, but the two reported networks might not necessarily have a similar underlying structure if the respondent answers differently between trials. Therefore, we also investigate the robustness of the pruning methods with synthetic data in the next section, where we know the simple version of the network beforehand and add noise to it prior to applying the pruning algorithms.

5.2 Added noise

Since it could be the case that patients report significantly different results after a couple of days when they have had a chance to think more about their symptoms, how they are experienced to interact, and their perceived painfulness, it may also be of interest to assess the accuracy of the results when the true network is known beforehand in contrast to the previous evaluation method. In accordance with this, we now create two directed and weighted synthetic networks similar to those in the given data, one smaller and one larger, which are chosen to have structures considered desirable post pruning with weights similar to that of the empirical networks used in the previous section. We then add random noise to these synthetic networks and apply the pruning methods on these noisy counterparts to examine their ability to identify and extract the underlying network structures hidden underneath the noise.

Specifically, we create a small graph of 5 nodes and 5 edges, and a larger graph of 9 nodes and 11 edges, which are both connected and contain loops. These structures are chosen as all symptoms are in a single component and are thus not isolated, and as psychologists may be interested in identifying loops where symptoms keep affecting and triggering each other. Such networks can thus be used to determine points of intervention and treatment in the structure. Furthermore, we again let $\gamma = 0.95$ for consistency, although this could be set to other values as well, and this means that we cannot have more than 5 edges in the synthetic small network, since we prune $\lceil 0.95(|E| - 5) \rceil$ edges in each method, where $|E|$ is the total number of edges in the network. We would therefore have $|E| - \lceil 0.95(|E| - 5) \rceil$ edges left in the network after pruning, but the maximum total number of edges that can be present in this network without creating self-loops is 20 which yields $20 - \lceil 0.95(20 - 5) \rceil = 5$ as the maximum number of edges left after pruning. Hence, we let this network have 5 edges as we want to be able to prune the noisy networks and get back to this original network, and as we want there to be at least as many edges as nodes. Moreover, to further ensure similarity between the synthetic and empirical networks, we want to evaluate the case where the large network has roughly the same number of edges as the mean number of 51 edges in the given data, and by letting the number of edges be 50 we would have $50 - \lceil 0.95(50-9) \rceil = 11$ edges left after pruning such a network using our methods. Hence, this is the number of edges selected for the larger network. The directed, weighted adjacency matrices and tables of the node weights can be found in Appendix B for both synthetic graphs, and a visualization of the large graph is included in Figure 15 which is also found there (as we visualize the small one below).

Now, the synthetic graphs above are used to add random noise to a percentage of the nodes and the existing edges proportionally to their weights, and to add random edges with weights proportional to the mean weight of the existing edges to ensure that the weights are not too small or large in relation to the existing edges. The nodes and edges to add noise to are randomly sampled, and when adding new edges, we do not add self-loops or multiple edges to keep the networks simple. Furthermore, if the set V contains the existing nodes, the set E the existing edges, and the set $\hat{E}$ the potential new edges that can be added (excluding self-edges and multiple edges), we let N_V denote the set of vertices to add noise to, N_E the set of existing edges to add noise to, and $N_{\hat{E}}$ the set of new noisy edges to add to the network. We also let the number of noisy nodes $|N_V|$, the number of noisy existing edges $|N_E|$, and the number of new edges $|N_{\hat{E}}|$ be equal to

- $|N_V| = \lceil \mu_V |V| \rceil$,
- $|N_E| = \lceil \mu_E |E| \rceil$,
- $|N_{\hat{E}}| = \lceil \mu_{\hat{E}}(|V|(|V|-1) - |E|) \rceil$

for predetermined percentages $\mu_V, \mu_E, \mu_{\hat{E}} > 0$ of the nodes, existing edges, and potential edges to add noise to, respectively. The number of new edges in the third point follows since there are $|V|^2 - |V| - |E|$ possible new edges to be formed in a network with $|V|$ vertices, no self-loops, and $|E|$ existing edges.

Using this notation, the entries of the new noisy vector of node weights are thus

$$v_i^{(noisy)} = v_i(1 + \lambda_V \varepsilon_i \mathbb{1}\{i \in N_V\})$$

and the entries of the new noisy adjacency matrix are

$$w_{ij}^{(noisy)} = \begin{cases} w_{ij}(1 + \lambda_E \varepsilon_{ij} \mathbb{1}\{e \in N_E\}), & \text{if } w_{ij} > 0 \\ \lambda_{\hat{E}} \bar{w}(1 + \lambda_E \hat{\varepsilon}_{ij}) \mathbb{1}\{e \in N_{\hat{E}}\}, & \text{if } w_{ij} = 0, i \neq j \end{cases} \tag{10}$$

where $\varepsilon_i, \varepsilon_{ij}, \hat{\varepsilon}_{ij} \sim N(0,1)$ are random variables used to create noise for the original node weights v_i and edge weights w_{ij}, with edges e from node i to j in V. The mean existing edge weight $\bar{w}$ in the network is given by $\bar{w} = \frac{1}{|E|} \sum_{k,l \in V, k \neq l} w_{kl}$, and the edge weight $w_{ij} = 0$ if there is no edge between nodes i and j. The parameters $\lambda_V, \lambda_E, \lambda_{\hat{E}} > 0$ can be set to vary the amount of noise added to the node weights, existing edge weights, and new edge weights, and $\lambda_{\hat{E}}$ is included to ensure that the new added edges have a smaller weight on average compared to the original existing edges. If the new noisy weights $v_i^{(noisy)}$ and $w_{ij}^{(noisy)}$ are negative, which could potentially happen if the generated noise is negative and larger in magnitude than the original weights, we consider this as no noise being added and instead keep the original weights. In this way, we do not consider removing existing edges since the pruning methods cannot add back those edges. The node weights are handled in the same way for consistency, even though they could potentially be set to 0 instead without interfering with the abilities of the pruning methods. As in the empirical data, all weights are in the interval $[0, 100]$ and if the new noisy weights are greater than 100, which is the maximum weight of both the nodes and the edges, we let the new noisy weights be this maximum value to keep all weights in the same range. Due to this, when selecting the node and edge weights for our synthetic networks, we randomly sample weights from a $N(70, 10)$-distribution to ensure that the weights are large with some variance, but hopefully not too close to 100. The weights are also rounded to 2 digits to be comparable to the weights in the empirical data.

Using the above technique of adding noise, we do so in different combinations of adding noise to each component (nodes, existing edges, and potential edges) with varying intensity before pruning the noisy networks with the methods described in Section 4. Since the synthetic networks are examples of desired pruning results, they already have the desired number of edges, and it is therefore ineffective to use the pruning methods on noisy networks where only the nodes and existing edges have added interference. We therefore must consider combinations of noise added through new edges, together with potential noise added to existing node and edge weights. We evaluate the results using the Pearson and Spearman correlation coefficients, and the Jaccard coefficient to compare the noisy networks and the pruned noisy networks with the original synthetic networks. This is done to determine how successful the methods are in returning to the true networks after obscuring the structures with interference. We do this for a small amount and a larger amount of noise added to the nodes and edges in the networks, as further described in more detail in Section 6.2.

To emphasize the effects of adding noise, we visualize the small synthetic graph in Figure 6 together with an illustrated example of a version with noise added to the nodes and edges. All nodes and existing edges have been selected to have 20% ($\lambda_V = 0.2$ and $\lambda_E = 0.2$) of noise added to them, respectively, and five new edges are added with 20% ($\lambda_{\hat{E}} = 0.2$) noise. The resulting noisy adjacency matrix and node weights are included in Appendix B. We notice that some of the original edges are thicker after adding noise and that they are thicker than the new edges, as well as that node 2 has shrunk in size. The new edges have also created several new and more involved loops.

This and the previous evaluation method focus on the reliability of the pruning algorithms, but we are also in need of a method that judges the usefulness of the pruning results visually, since the main objective of the algorithms is to create visually simpler network structures that can be used in practice to interpret psychological processes. A natural way to do this is by asking potential users of the algorithms to rate the outcomes, and we therefore let psychologists choose the most reasonable pruning result for a selection of patients in the next evaluation method, which is outlined below.

5.3 Human expert evaluation

As the main objective is to produce networks that are simpler to understand yet have a reasonable structure that contains useful information, it is essential to evaluate the visual quality of the resulting networks to get a sense of the usefulness of the methods. We therefore plot the pruned versions of the networks provided by three additional patients with patient IDs A, B, and C who also participated in the data collection survey, and these are separate from the 21 patients who completed the survey twice. Additionally, we select the three

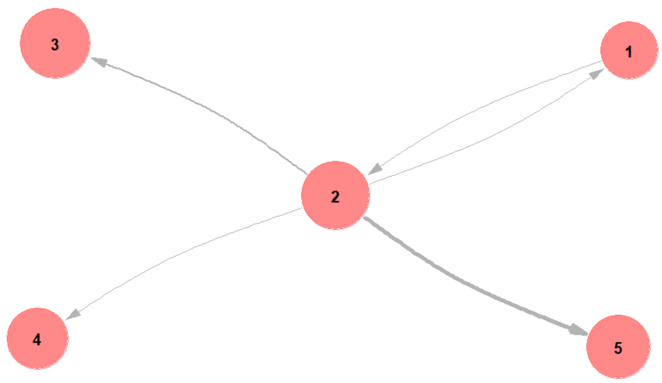

(a) No noise.

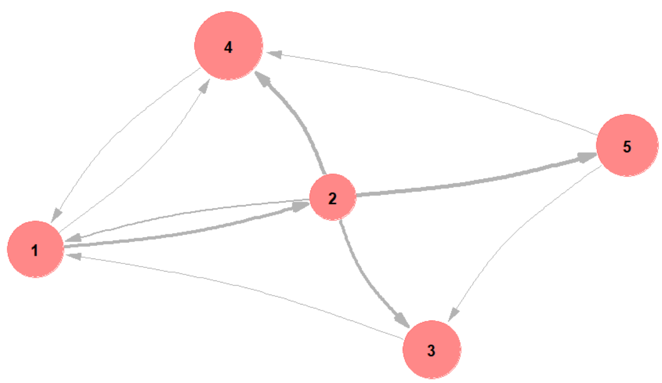

(b) Effect of noise.

Figure 6: An example of the effect of adding noise to the nodes and edges of the small example graph. In particular, all nodes and existing edges are selected to have 20% and 20% of added noise, respectively, while five new edges are added with 20% noise. The original version without noise is in Figure 6a and the noisy version is in Figure 6b. The edge widths represent the frequency of the interactions between symptoms and the node sizes represent the painfulness of each symptom. The direction of the edges is illustrated by the arrows.

most reliable patient IDs (247, 276, and 292) from the original 21 patients with the highest correlation between adjacency matrices and also plot their pruning results for their second reported network, as this second network is hopefully more reliable than the first, given that the participant is familiar with the survey and has had a chance to consider their responses at the time of the second trial.

The resulting networks for these six patients are presented in a new survey to five licensed psychologists, where two of them are the new patients' real psychologists who know their psychological state well enough to give their opinion on how well the networks illustrate the individual's symptomatological interactions. In particular, the patients with IDs A and B are known by the same psychologist, while patient C is known by another psychologist and the networks for these patients are evaluated by their respective psychologist. In this way, we could consider these psychologists as having access to a subjective, underlying "true" network through their knowledge and experience of their patients' mental state. Conversely, the other three psychologists do *not* know the patients selected from the original data and their psychiatric history, and are only asked to evaluate the networks for these patients. This is of course a small sample of psychologists, limited by the available patient data, and one may prefer to collect more data in the future by asking more psychologist-patient pairs to participate.

The survey given to the psychologists is conducted online in Swedish and only the two who know the patients are aware of their identity. In addition, the patient ID (an anonymous code), the original network provided by the patient, as well as the pruning results from Methods I-IV in Section 4 are presented for each patient, where the pruned networks are unlabeled to avoid influencing the participant's perception of the pruning methods. In all networks, the node sizes illustrate the painfulness of the symptoms as reported by the patient, and the original network also visualizes the reported frequency of the experienced interactions through the edge widths. Each node is labeled with the symptom name and a short description of the patients' experience with the symptom. The psychologists are asked to choose which pruning result they believe is the most reasonable for the given patient, and to provide comments about if there is any information considered missing in the pruned networks. The psychologists who know the patients can thus select the networks based on how well they believe the networks represent their perception of the patients' individual psychological processes, whereas those who do not know the patients can make their selections based on how realistic they believe the pruning results are in general. The networks included in the survey can be found by following the link in Appendix A.

6 Results

We now prune the networks available in the data and evaluate them according to the methods described in Section 5. When choosing the value of γ to decide the stopping criteria, the initial intention was to fix the number of edges to prune such that the number of edges in the resulting network equals the number of nodes to make it possible for the pruned networks to still be connected, while avoiding extra edges obscuring the visual interpretability. Nevertheless, after trying out different values of γ, we decided to let $\gamma = 0.95$ to allow for a few (in proportion to the number of edges in the original network) extra edges to be present in the final pruned versions of the networks, as it was found not to disturb the visualizations but instead preserve useful information. We also tried using a smaller value of γ, such as 0.90, which was found to make the pruned networks too difficult to interpret, and a higher value, such as 0.97, which was deemed less efficient than $\gamma = 0.95$ as the resulting networks could sometimes visually benefit from having additional edges, in terms of for example displaying interesting loops and avoiding disconnections.

6.1 Repeated trials

We first evaluate the results using the empirical data provided by the patients having participated in the survey twice within an interval of a few days. In this case, we do not have access to any true, simplified versions of the networks, if such exist, but only to what the patients have reported in their survey answers. Therefore, we do not compare the outcomes to a theoretical underlying structure, but instead we compare the reported networks and their pruning results from the first and second trials, as they hopefully have similar underlying structures that are detectable by the pruning methods. Hence, the goal is to evaluate how well the pruning methods detect the most important coinciding structures in similar but noisy networks, and if the networks become more similar post pruning. For this reason, we also restrict ourselves to the coinciding nodes in the networks, as the excess nodes are redundant when comparing the similarity before and after pruning. To do this, we compute the Pearson and Spearman correlation coefficients, the Jaccard coefficients, and the out-degree node centralities for the weighted and directed adjacency matrices corresponding to the first and second trials. This evaluation method is described in more detail in Section 5.1. However, when doing this, we observe that the Updated PageRank pruning result for patient 193 from trial 1 has no edges left when considering the coinciding symptoms of trial day 1 and 2, and the correlations cannot be computed, as the standard deviation is zero in the adjacency matrix of zeros for trial 1. We hence let these networks have correlation 0 instead, as they have no edges in common. Although, we note that this could be handled differently as well, as the two adjacency matrices only differ in the places where the second network has non-zero values.

To get an overview of the distributions of the similarity coefficients, we have illustrated these as boxplots in Figures 7-9. In Figure 7, we see the distributions of the Pearson coefficients measuring the linear relationships between the edge weights of the networks from the first and second trial. We immediately note that the box part is in general much higher up on the scale for the original networks than for the pruned results, with some overlapping for the PageRank approach and some minor overlapping for the Edge betweenness and Updated PageRank approaches, indicating that most networks do not become more similar post pruning for any of the algorithms. This, since most, or almost 75%, of the original correlation coefficients are higher than most of the correlation coefficients after pruning and in particular, half of the original networks have correlations between about $0.55 - 0.80$ whereas 75% of the respective data points are below 0.60 for all but the PageRank pruning method, whose upper quartile is below 0.70. In addition, the original boxplot is completely overlapped by the boxplots for both the Edge betweenness and PageRank methods, but only partially overlaps with the other two methods.

The whiskers are also in general longer for the simplified networks, and the boxes and whiskers span almost the entire interval from 0 to 1 for the Edge betweenness method and the PageRank method, and includes even some negative values. The algorithms using updated PageRank values and the ratio of connectivity kept have slightly smaller spans, but the maximum values are lower as they only produce correlations of at most about 0.80 and just under 0.75, respectively. Meanwhile, the maximum correlation produced by the Edge betweenness algorithm is greater than 0.90, and that of the PageRank algorithm is just under 1.0. This is comparable to the original networks, as the highest value is just greater than 0.90. In contrast to the original networks, all methods produce some graphs with a slight negative correlation, which indicates that the methods sometimes result in networks with slightly inversely proportional edge weights, that are close to having zero correlation. The Edge betweenness approach produces the most negative correlations overall of around -0.15, but all methods have some slightly negatively correlated results. This is a significant decrease compared to the reported networks, as the lowest value for the original networks when disregarding the outlier is around 0.30. However, most of the correlations are still positive which means that most networks maintain some similarity after the simplification, even if the similarity seems to be lower than before.

Moreover, it is also clear that the correlations between the pruned networks are much more varied compared to the original versions, as the interquartile ranges (IQR) are larger compared to that of the original networks, as seen by the size of the boxes. The Updated PageRank approach has the largest IQR and the Connectivity kept method seems to have the shortest, and these also have the lowest quartiles that are just under or close to the dashed red line dividing the positive and negative side of the scale, meaning that the interval containing half of the data points includes slightly negatively correlated networks for the Updated PageRank approach and almost does so for the Connectivity kept approach. Since the IQR for the latter method is the shortest, this also means that most networks are barely or weakly correlated as 75% of the networks have a lower correlation than 0.45. In addition, the median is the highest at about 0.35 for the PageRank method and lowest at around 0.15 for the Updated PageRank approach, and all methods have medians lower than that of the original networks which is close to 0.65 and demonstrates relatively strong linear relationships for the original networks. However, the results indicate that these relationships are in general not well-kept, as the correlations are in general much lower for the pruned networks, as all pruning methods have medians much lower than the lower quartile for the original networks and most are lower than the minimum value when discounting the outlier.

Similar results are noticeable in Figure 8 when considering the Spearman coefficients measuring the monotone relationships using ranks, since we again have that the original networks are in general more correlated than the pruned results as shown by how the boxes are higher for the original networks. We again have overlapping boxes for the PageRank approach and some minor overlapping for the Edge betweenness and Updated PageRank approaches, while the box for Connectivity kept is much lower than the box for the non-simplified graphs, but the Spearman coefficients also suggest that the correlations are in general lower after pruning. The original Spearman correlations have an almost identical boxplot compared to the Pearson correlations, with the main difference being the absence of

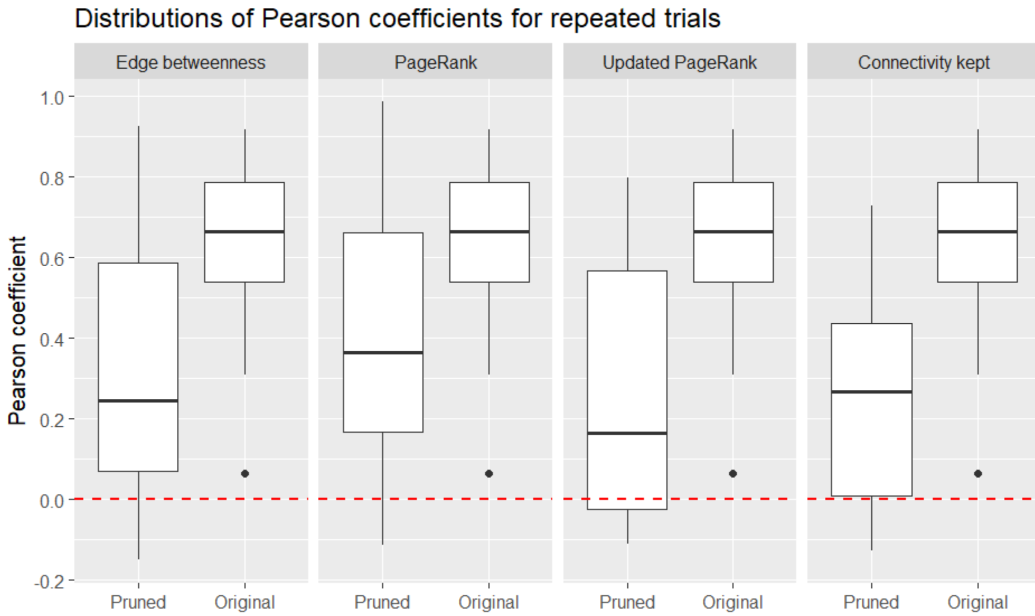

Figure 7: The distributions of the Pearson correlation coefficients for the original and pruned graphs of the networks reported by patients on day 1 and day 2. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV in Section 4.

an outlier which could be explained by how the Pearson coefficient measures linear relationships while the Spearman coefficient measures monotone relationships and may therefore rate the outlier as more correlated when considering ranks rather than the edge weights themselves. Other than that, the quartiles and median are slightly lower for the Spearman coefficient, but still extremely similar. As for the Pearson correlation, this full boxplot is also completely overlapped by the boxplots for the Edge betweenness and PageRank methods, but not by the other two methods.

Furthermore, we also see that the whiskers are again longer in general for the simplified networks with large spans overall, and the Edge betweenness and PageRank methods have the largest difference between the minimum and maximum Spearman correlations, with minimums around -0.17 and -0.21, and maximums around 0.91 and 1.0, respectively. The Updated PageRank and Connectivity kept methods have minimums around -0.20 and -0.05, and maximums around 0.77 and 0.79, respectively. This again shows that the former two methods have maximum correlations closer to that of the original networks, but they also have much lower minimum values. All pruning methods thus similarly have results that are slightly negatively correlated or uncorrelated, but the majority of the networks have a positive Spearman correlation. Although, they are again shown to have more varied correlations than the original networks, as they all have an IQR larger than that of the reported networks. However, in contrast to the Pearson correlation, it appears that the Edge betweenness approach has the largest box, but the Connectivity kept method still has the smallest one. Though, in this case, the lower quartile is now above the dashed red line.

The lower quartile is still below the dashed line for the Updated PageRank method, but it is a bit lower now for the PageRank approach compared to for the Pearson coefficient. The upper quartile is also lower than in Figure 7, and hence the median is a bit lower too at close to 0.28, making the median for the Connectivity kept method the highest out of the four pruning methods at around 0.33. The medians for the remaining two methods are about the same. However, since the medians are all again much lower than the median for the original networks and also closer to the minimum value of the original correlations, this suggests that the monotonic relationships of the networks are not well-preserved after applying the simplification techniques.

Thus, the Pearson correlation coefficients seem to suggest that similarity decreases after pruning, but that the PageRank and Edge betweenness methods are better suited to extract similarities in the most important underlying structures, as they produce higher correlations on average. They also seem better at producing simplifications that are more similar than before, but it is difficult to tell from the boxplots other than from the maximum values being slightly larger. However, all methods generally produce results that are less correlated than before pruning. Similarly, the Spearman correlation coefficients also seem to favor the Edge betweenness and PageRank methods as these boxes are closer to the box for the original method and these also produce the highest maximum coefficients compared to the remaining methods, but it is more difficult to select a method based on these results. The Connectivity kept method has the highest median and a bit less varied results, but higher correlations seem rarer for this method than for the Edge betweenness and PageRank approaches.

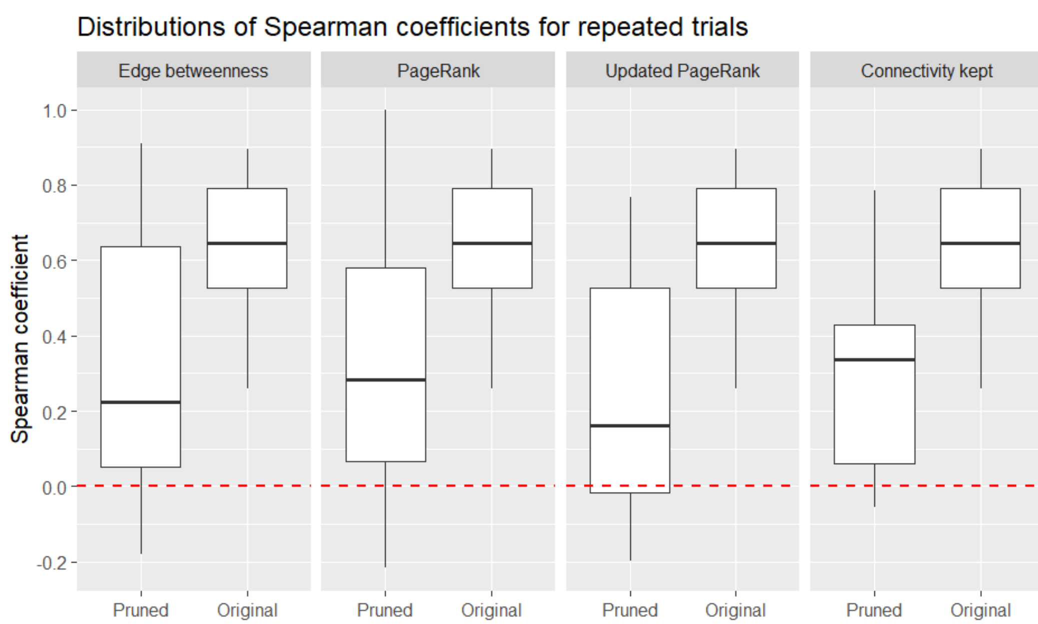

Figure 8: The distributions of the Spearman correlation coefficients for the original and pruned graphs of the networks reported by patients on day 1 and day 2. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV in Section 4.

Moreover, as the correlation methods above take into account the edge weights, we also analyze how the number of specific edges in common for both networks is changed after pruning, and the distributions of the Jaccard coefficients are displayed in Figure 9. As one might expect based on the correlation results, we again note that the original networks are much more similar than the pruned networks as their boxplot is much higher than the boxplots for the pruning methods, showing that the original networks have in general more edges in common. In fact, most Jaccard scores are less than the minimum for the original networks for all pruning methods, suggesting that most simplified networks are less similar between trials after pruning. Furthermore, the Jaccard coefficients are much less varied than the correlation coefficients, which might be due to the varying edge weights compared to the consistent binary possibility of the edge being present or not. We note that the original networks have at most the exact same edges chosen for the first and second trial, and have at least more than half of the combined edges in common, as the minimum value is about 0.55 if we disregard the outlier at about 0.25. The median is close to 0.80 and half of the networks have a Jaccard coefficient between about 0.73 and 0.88.

Despite the original networks having many edges in common between trials, we find that all simplified networks are less similar in terms of the Jaccard scores than before pruning, as none of the boxes overlap with the box for the original networks, and only the whiskers overlap for all methods except the Connectivity kept approach, which only has an outlier at about 0.72 and the maximum value is otherwise 0.50. This outlier also slightly increases the median to about 0.28, and without the outlier the median is otherwise around 0.23. Additionally, this is the only method whose maximum value is below 0.50, but all other methods have at least 75% of their Jaccard coefficients below 0.50. The PageRank method is the only method to have fully coinciding edges in a network after pruning, but this method also produces networks with no edges in common, as does all other methods except for the Connectivity kept approach.

When comparing the methods, we note that the Connectivity kept method seems more consistent as the boxplot is smaller than the others if we forget the outlier, but this method and the Updated PageRank approach have overall lower Jaccard values than the PageRank and Edge betweenness approaches, even if the boxes mostly overlap. The latter two simplification techniques have similar boxplots, with the significant difference that the Edge betweenness method does not produce scores above 0.85. Their medians are close to about 0.23 and 0.20, respectively, which is quite low and also comparable to the median of the Connectivity kept approach without the outlier, which otherwise has the highest median out of the four methods. Based on this, and that the Edge betweenness and PageRank approaches have higher upper quartiles and maximum values, we determine that these methods seem to produce networks with more edges in common based on the boxplots, with the latter being slightly more successful than the former.

Nevertheless, it is difficult to interpret the boxplots in terms of how well the methods increase the similarity after pruning, and we hence analyze this further by directly counting how many of the 21 network pairs have higher similarity coefficients post pruning compared to before. These results are displayed in Table 1 where we observe that most of the resulting networks are less similar than their original versions, since the highest number of networks having the same or a higher similarity coefficient after pruning is only 6 out of 21 for the Pearson correlation for the PageRank method. This method seems to increase the similarity the most overall, as it has the most increased similarity coefficients out of the four methods, even if only 6, 5, and 3 networks have increased Pearson, Spearman, and Jaccard coefficients after applying the approach. The second-best method is the Edge betweenness method

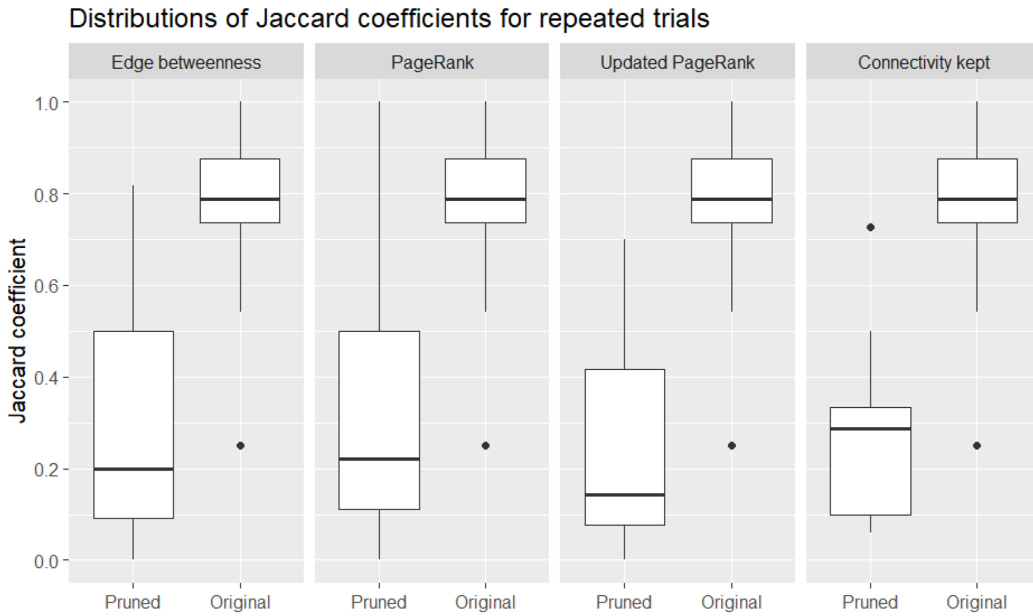

Figure 9: The distributions of Jaccard coefficients for the original and pruned graphs of the networks reported by patients on day 1 and day 2. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV in Section 4.

which increases the Pearson and Spearman correlations for 5 of the networks, respectively, and increases the Jaccard coefficient for 2 of them. The Connectivity kept method scores the lowest with just under 3 networks being more correlated after pruning according to the Pearson coefficient. From the table, the PageRank approach still seems to be the most favorable out of the four methods as it is able to increase the similarity for a little less than a third of the network pairs after pruning, but this is still quite low and rather shows that most networks are less similar after simplification and hence the methods are not very successful in finding similar features in the two networks reported by the same patient.

Furthermore, a reason why the correlations and number of edges in common in general decrease after pruning could be that the edges that the original networks have in common have lower edge weights and are connected to less painful end nodes. This could make the edges less central to the network structure according to the pruning methods, and so they are hence removed to favor other edges that the two networks do not have in common. In this way, perhaps the coinciding structures might not be the most important for the shortest paths in the network, which decreases the edge betweenness scores and increases the ratios of connectivity kept for the edges in common after pruning, leading to those edges being removed.

Table 1: The total number of the Pearson and Spearman correlation coefficients, and Jaccard coefficients of the pruned results that are larger than that of or equal to their corresponding original networks for each of the pruning methods, for the repeated trials. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV described in Section 4. There are 21 network pairs in total.

	# network pairs (/21)		
Method	Pearson	Spearman	Jaccard
Edge betweenness	5	5	2
PageRank	6	5	3
Updated PageRank	4	3	1
Connectivity kept	2	0	0

Moreover, since the above does not give an intuition about whether networks need to be highly correlated in order to also be correlated after pruning, we investigate the direct relationships between the similarity coefficients for the original network pairs compared to the pruned network pairs. The results are shown in scatter plots with fitted regression lines for the Pearson coefficient in Figure 10. There seems to be a somewhat increasing relationship overall, but it is weak and most points fall far from the regression line for all methods and especially for the Edge betweenness and Updated PageRank methods, where the regression line has almost no slope and the R^2 value is close to 0. On the other hand, the regression line shows a stronger positive relationship between the original Pearson correlation and the pruned correlations for the PageRank and Connectivity kept methods, with significantly higher R^2 values of 0.11 and 0.18 each, which are nonetheless still low values. This suggests that there is a faint increasing trend, but it does not explain a lot of the variance in the pruned Pearson coefficients.

Furthermore, the first point to the left for the Edge betweenness, PageRank, and Updated PageRank approaches seems to be an outlier, as this patient (patient ID: 208) reported a small enough network on day 1 that it did not get pruned at all while the network on day 2 had significantly more edges. This could explain why the correlation is so high after pruning for this particular patient even though the original networks were not highly correlated, since the resulting networks would become more similar as long as the second network prunes more of the edges it does not have in common with the first network. Additionally, this could also explain why it is so low for the Connectivity kept approach, since the second network only needs to have the few edges they both have in common removed to decrease their similarity drastically, which was low to begin with. Hence, it is reasonable that the correlation could vary more for smaller networks paired with larger networks.

Similarly, we also show the relationships for the Spearman correlation coefficients in the scatter plots in Figure 11 and observe the same trend, where we have a weakly positive relationship between the original correlations and the pruned correlations overall. The weakest relationships are again detected for the Edge betweenness and Updated PageRank approaches, where the trend looks almost constant with R^2 values close to 0 as most points are not close to the linear regression line. The linear relationships are a bit clearer for the PageRank and Connectivity kept methods, with R^2 values of 0.10 and 0.27 and the latter is noticeably larger than for the Pearson coefficients for the same method.

Thus, it could be the case that higher original correlation yields more correlated pruned

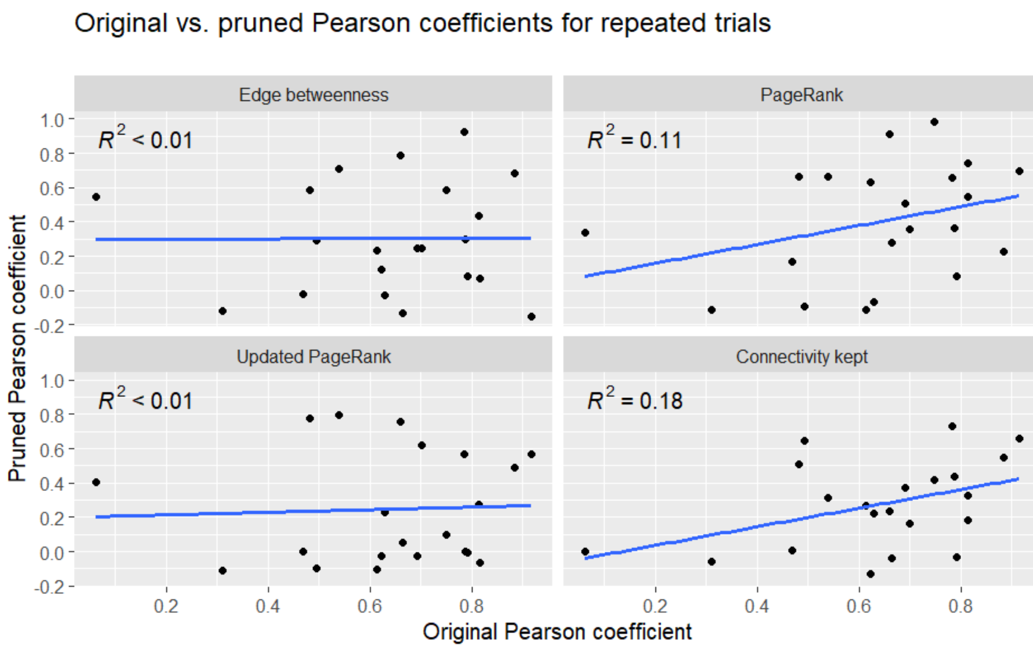

Figure 10: The distributions of Pearson correlation coefficients for the pruned graphs of the networks reported by patients on day 1 and day 2, when comparing the values to the Pearson correlation coefficients of the corresponding original networks. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV in Section 4.

networks to some degree, which is reasonable since networks that are highly similar or identical should get similar pruning results, and likewise, networks that have close to no structures in common should get different simplification results as well. However, it does not seem to be the only variable impacting the resulting correlation, and it does seem feasible that the trend is weak since we systematically remove structures from the networks. It is possible to have two large networks with many edges that are highly correlated, but have differences in their underlying structures in a couple of edges that are connected to painful nodes with frequent interactions with other nodes. In this case, since the pruning methods remove a percentage of the difference between the number of edges and nodes, a lot of edges would get removed from both networks, and it would therefore be reasonable that the underlying distinguishing structures are kept if they are more frequently involved with other, more painful symptoms. In this way, the edges in this structure would have high betweenness and low ratios of connectivity kept upon removal, leading to other, less important edges being removed instead. This means that the originally highly correlated networks have lost a significant amount of similarity after pruning.

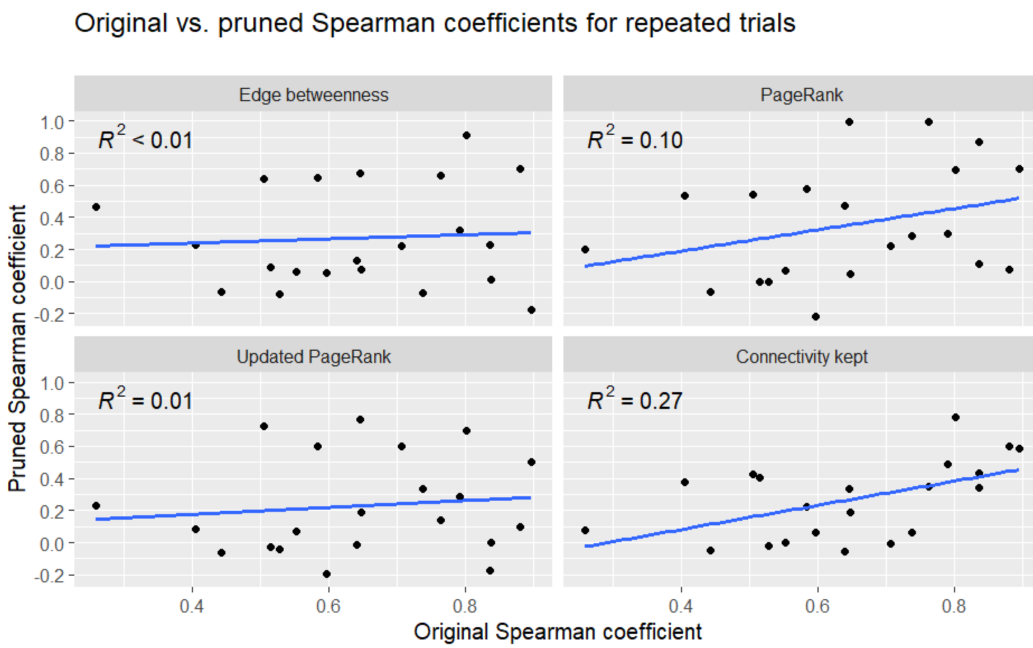

Figure 11: The distributions of Spearman correlation coefficients for the pruned graphs of the networks reported by patients on day 1 and day 2, when comparing the values to the Spearman correlation coefficients of the corresponding original networks. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV in Section 4.

Furthermore, in Figure 12 we plot the Jaccard coefficients for the pruned graphs compared to the original coefficients and note somewhat corresponding trends, but these are much weaker. Only the PageRank method has a linear regression line with a notable positive slope, but the R^2 value is only 0.06 and we observe that many of the data points are scattered far from the line. Meanwhile, the remaining three methods have regression lines with almost no slope and the R^2 values are closer to 0. This suggests that the original graphs having a high number of edges in common does not necessarily mean that they will keep most of those edges in their simplified networks. Thus, it seems like the initial correlation does not have a large effect on the pruning result in general, but could have a small positive effect. However, it could also be possible that more data are needed to be able to see clearer trends, or that the weights provided by the patients differ even though the networks contain a similar set of edges. This could influence the PageRank scores, edge betweenness scores and ratios of connectivity kept in the pruning methods, yielding different resulting networks even if the original networks had similar structures.

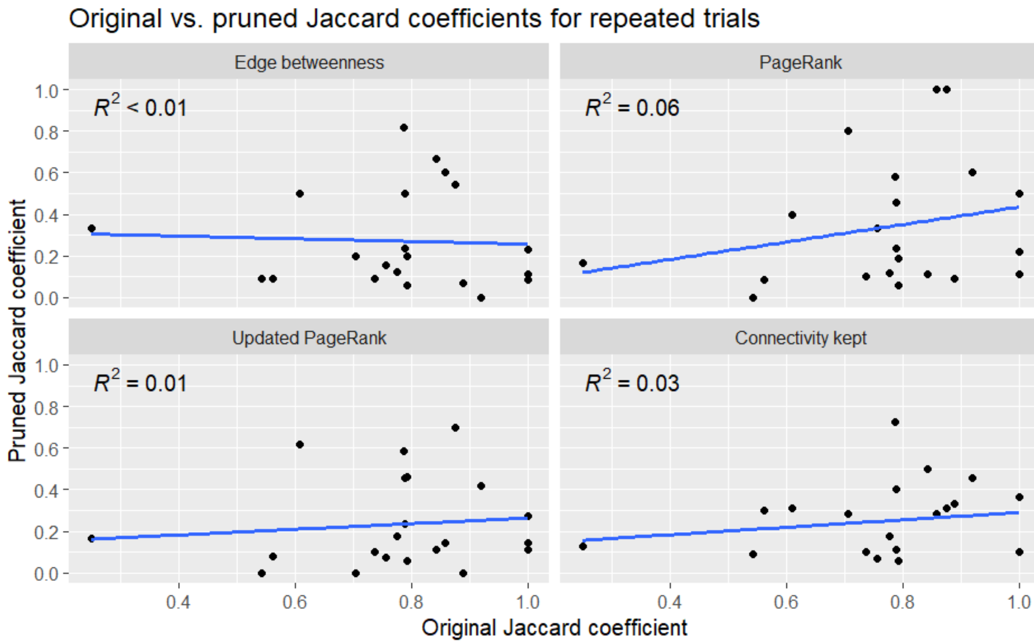

Figure 12: The distributions of Jaccard coefficients for the pruned graphs of the networks reported by patients on day 1 and day 2, when comparing the values to the Jaccard coefficients of the corresponding original networks. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV in Section 4.

Finally, we are also interested in how similar the most central structures in the networks are, and we thus also compute the out-degree centrality scores for all networks and compare the nodes selected as the most central for the first and second network reported by the patients. In Table 9, we can see the total number of network pairs deeming the same node as the most central, and we observe that only the Updated PageRank method has fewer coinciding most central nodes than in the original network pairs. The Edge betweenness and PageRank methods both have 13 network pairs with coinciding central nodes, which is an improvement compared to the initial networks only having 8 network pairs with the same most important node. The Connectivity kept approach performs as well as the original networks in this sense. Hence, in terms of out-degree centrality, the Edge betweenness and PageRank methods seem to be the most successful in detecting the most important, coinciding node structures.

Table 2: The total number of network pairs from the test-retest results having the same node ranked highest by out-degree centrality for the original versions and for each of the pruning methods described in Section 4, where Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV. There are 21 network pairs in total.

Method	# network pairs (/21)
Original	8
Edge betweenness	13
PageRank	13
Updated PageRank	2
Connectivity kept	8

6.2 Added noise

We now investigate the robustness of the pruning results by evaluating the effect of adding noise to synthetic networks with structures chosen to resemble possible simplified symptom networks with a central node and feedback loops. The noisy original networks and the pruned versions of the noisy networks are compared to the corresponding original synthetic networks, and in this way, we know the structure of the true, simplified network which is hidden among the noise. This is in contrast to the previous evaluation method where we use empirical data from two different trials and do not have access to the true simplified networks which we in that case assume exist, and there we also assume that the important, underlying structures of the two networks are similar in order to evaluate if their common most important structures are detectable by the methods. In this evaluation method, we instead know that such significant, underlying structures exist in the synthetic networks and their noisy counterparts as we only add a small amount of noise to ensure that the structures that the original and noisy networks have in common correspond to the most important structures, since higher weights correspond to shorter distances. Thus, the purpose is to investigate how well the pruning methods detect and obtain this true, underlying network in their resulting networks. This evaluation method and the synthetic graphs used are described in more detail in Section 5.2.

In Table 3, we have summarized the average node and edge weights of the original synthetic networks and these weights were chosen to be large enough to increase their importance, but not be too close to 100 since this would increase the likelihood of noisy weights exceeding 100 resulting in less noise added. Similarly, we do not want the weights to be too close to 0 as this increases the likelihood of the noisy weights being negative. These weights are each sampled from a normal distribution with mean 70 and standard deviation 10 and can be found in Appendix B.

The idea is now to add interference to the node weights and the existing edge weights, as well as add many new edges with noisy weights that are smaller than the original weights. However, since both synthetic networks are chosen to have structures desirable post pruning with an acceptable number of edges, the case of applying the pruning methods when adding noise to the existing edges and nodes alone without also adding new noisy edges is not very useful, and we want the noisy networks to resemble the large and complex empirical networks used in practice. We thus compare the outcomes for four different combinations

Table 3: The means of the original node weights and edge weights of the small and large synthetic networks. These weights are sampled from a normal distribution with mean 70 and standard deviation of 10 to ensure large enough weights that do not vary too much.

Synthetic network	Mean node weight	Mean edge weight
Small	76.2	71.2
Large	77.6	74.3

of adding noise to

(i) the edge weights of **potential new edges**,

(ii) the edge weights of **potential new edges** and to the **node weights**,

(iii) the edge weights of **potential new edges** and to the **existing edge weights**,

(iv) the edge weights of **potential new edges**, the **existing edge weights**, and the **node weights**

where we vary the level of noise added through changing the values of λ_V, λ_E and $\lambda_{\hat{E}}$ simultaneously for simplicity, but these could also be varied separately. This is done in steps without changing the previously given interference between cases, as it is otherwise difficult to see the changes due to the new noise being added. In addition, to be aware of how much noise is added in general, we compute the mean absolute value of the noise added to these weights, and the mean absolute value of the noise added as weights of the new edges (i.e. the resulting weight of the new edges). These values are presented in Table 4 for when we let $\lambda_V, \lambda_E, \lambda_{\hat{E}} = 0.1, 0.2$ and in the first case we thus add about 10% noise to the nodes and existing edges, and add 10% of the noisy mean of the existing edge weights, which also has 10% noise added to it according to Eq. (10). The values in the table show that these levels of interference do not change the node or edge weights too much, and that the new edges have much lower weights than the existing edges. The noise sampling is repeated to create five different noisy networks that are reused for each case to account for the randomness in the results, and the computed values are the averages for these repeated trials. The unweighted resulting networks are visualized for both synthetic graphs in Appendix B for completeness.

In Table 5, we present the effects of adding noise according to the four cases to all nodes, all existing edges, and 10 or 39 new edges where applicable in the small and large synthetic networks, and we let $\lambda_V, \lambda_E, \lambda_{\hat{E}} = 0.1$. By adding 39 new edges to the large network, we obtain 50 edges in total, which is close to the mean number of edges in the empirical data. Based on the table, we first note that the Pearson correlation coefficient is almost equal to 1 when comparing the synthetic networks to their noisy versions for both the small and large networks, which does not seem like a reasonable correlation as we have added many edges to the networks. The values are overall higher than for the Spearman coefficient and this suggests that the Pearson correlation is not very suitable for this data. This could be due to there being many zeros in the original adjacency matrices and the new edges having small weights such that the sets of edge weights are skewed with large differences between the existing and new edge weights. Furthermore, the Jaccard coefficient suggests that the intended edges are indeed added to the networks, since the (rounded) ratios indicate that

Table 4: The mean absolute values of the interference for the nodes and existing edges, and the mean absolute values of the noisy edge weights of the new edges for the original small and large synthetic graphs for the different ways of adding noise. These are the average values when repeating the addition of noise five times. The large network has 50 edges in the noisy versions (39 new edges added) and the small network has 15 edges in the noisy versions (10 new edges added).

	Percentage of noise	Noise nodes	Noise existing edges	Noise new edges
Small network	0.1	5.14	4.45	7.02
	0.2	9.92	8.44	13.85
Large network	0.1	4.83	12.08	7.51
	0.2	9.25	23.44	15.18

the small synthetic graph has its 5 original edges in common with its noisy version and that the large synthetic graph has its 11 edges in common with the noisy graph for all four cases.

Moreover, we observe no change in the Jaccard coefficient for the small graph for any of the centrality based pruning methods, while these values have increased for the large graph for all methods, suggesting that the noisy graph has more edges in common with the original synthetic graph after the simplification for all pruning methods. The Connectivity kept method seems to maintain this similarity slightly more than the remaining methods for all cases of possible noise addition, and it also manages to increase the similarity for the smaller graph. However, all Spearman correlations are lower than that of the comparison between the original graph and its noisy version, suggesting that the weights of the edges may not be maintained well. Hence, for low levels of interference, the similarity for this small graph is not well-kept, and less so for the centrality based methods. However, the similarity is increased for the larger graph, both in terms of the number of edges the networks have in common and in terms of the Spearman correlation for the edge weights, which is reasonable since we give more weight to the original structure and add many edges that are more likely to get pruned due to their low weights. For the small graph, there are not as many edges to choose from, and removing an edge would have a larger impact on the overall structure than when removing an edge from a large, dense network. Moreover, we do not observe a major difference between methods and the types (i)-(iv) of noise addition, except that the Jaccard coefficient seems to decrease with added noise for the Connectivity kept method.

Table 5: *All types, 10% noise*: The average Pearson and Spearman correlation coefficients, and Jaccard coefficients of the pruned results and the corresponding original networks for each of the pruning methods described in Section 4 for the five repeated results when adding noise according to the four noise types. For both networks, 10% of noise is added to the possible weights for each of the four methods of adding noise. The large network has 50 edges in the noisy versions (39 new edges added) and the small network has 15 edges in the noisy versions (10 new edges added).

		Small network			Large network		
Noise type	Method	Pearson	Spearman	Jaccard	Pearson	Spearman	Jaccard
(i)	Original	0.99	0.72	0.33	0.99	0.61	0.22
	Edge betweenness	0.62	0.40	0.33	0.87	0.75	0.60
	PageRank	0.61	0.40	0.33	0.88	0.74	0.60
	Updated PageRank	0.58	0.38	0.33	0.89	0.74	0.60
	Connectivity kept	0.70	0.53	0.39	0.80	0.74	0.73
(ii)	Original	0.99	0.72	0.33	0.99	0.61	0.22
	Edge betweenness	0.62	0.40	0.33	0.87	0.75	0.60
	PageRank	0.61	0.40	0.33	0.88	0.74	0.60
	Updated PageRank	0.58	0.38	0.33	0.89	0.74	0.60
	Connectivity kept	0.70	0.53	0.39	0.84	0.77	0.70
(iii)	Original	0.99	0.72	0.33	0.99	0.61	0.22
	Edge betweenness	0.62	0.40	0.33	0.83	0.70	0.56
	PageRank	0.61	0.40	0.33	0.87	0.74	0.60
	Updated PageRank	0.58	0.38	0.33	0.89	0.74	0.60
	Connectivity kept	0.69	0.53	0.39	0.82	0.75	0.70
(iv)	Original	0.99	0.72	0.33	0.99	0.61	0.22
	Edge betweenness	0.62	0.40	0.33	0.87	0.75	0.60
	PageRank	0.61	0.40	0.33	0.87	0.74	0.60
	Updated PageRank	0.58	0.38	0.33	0.89	0.74	0.60
	Connectivity kept	0.70	0.53	0.39	0.85	0.74	0.65

Furthermore, we observe similar results in Table 6 when adding more noise to all the components in the data by letting $\lambda_V, \lambda_E, \lambda_{\hat{E}} = 0.2$, yielding about double the absolute mean of noise added previously. The Pearson coefficients between the original graph and its noisy version are again elevated, probably for the same reasons. These are lower for the pruned results, however, but still a bit higher than the Spearman correlations. Not much has changed for the small or the large graph in terms of their Spearman correlations and their Jaccard coefficients, except some small differences where for example the Updated PageRank method has a slightly lower Spearman correlation and Jaccard coefficient for when we add noise everywhere at once in case (iv). This might suggest that we need to add more noise to see a difference, as we observe in Table 4 that the mean absolute interference is quite small even for this increased value of λ_V, λ_E, and $\lambda_{\hat{E}}$.

Hence, we now try this for case (iv) in the former case when 20% of noise is added to now add even more noise in the new edge weights, and we vary $\lambda_{\hat{E}}$ by letting it equal as much as 0.7 and 0.9 to see what happens if we add edges with almost as large weights as the mean weight of the existing edges (however, the noise of the edges could also be varied more by also changing λ_E, as the weight of the new edges also depend on this parameter). In Table 7, we can see that the mean absolute noise is much higher for both networks, and

Table 6: *All types, 20% noise*: The average Pearson and Spearman correlation coefficients, and Jaccard coefficients of the pruned results and the corresponding original networks for each of the pruning methods described in Section 4 for the five repeated results when adding noise according to the four noise types. For both networks, 20% of noise is added to the possible weights for each of the four methods of adding noise. The large network has 50 edges in the noisy versions (39 new edges added) and the small network has 15 edges in the noisy versions (10 new edges added).

		Small network			Large network		
Noise type	Method	Pearson	Spearman	Jaccard	Pearson	Spearman	Jaccard
(i)	Original	0.97	0.72	0.33	0.95	0.61	0.22
	Edge betweenness	0.58	0.39	0.33	0.86	0.75	0.60
	PageRank	0.59	0.40	0.33	0.85	0.74	0.60
	Updated PageRank	0.56	0.38	0.33	0.88	0.73	0.60
	Connectivity kept	0.68	0.53	0.39	0.81	0.76	0.75
(ii)	Original	0.97	0.72	0.33	0.95	0.61	0.22
	Edge betweenness	0.58	0.40	0.33	0.86	0.75	0.60
	PageRank	0.59	0.40	0.33	0.85	0.74	0.60
	Updated PageRank	0.56	0.38	0.33	0.88	0.73	0.60
	Connectivity kept	0.67	0.53	0.39	0.82	0.77	0.75
(iii)	Original	0.96	0.72	0.33	0.94	0.61	0.22
	Edge betweenness	0.57	0.39	0.33	0.82	0.70	0.56
	PageRank	0.57	0.38	0.33	0.81	0.67	0.54
	Updated PageRank	0.55	0.38	0.33	0.87	0.76	0.60
	Connectivity kept	0.70	0.53	0.39	0.83	0.76	0.68
(iv)	Original	0.96	0.72	0.33	0.94	0.61	0.22
	Edge betweenness	0.64	0.47	0.36	0.86	0.73	0.57
	PageRank	0.58	0.40	0.33	0.81	0.67	0.54
	Updated PageRank	0.55	0.38	0.33	0.87	0.76	0.60
	Connectivity kept	0.67	0.53	0.39	0.83	0.74	0.65

closer to the respective means. The results in Table 8 suggest that all methods are still quite robust to this noise, but all similarity coefficients have decreased compared to in Table 6, suggesting that this increased noise does affect the pruning outcomes. Moreover, the Pearson correlation coefficient has decreased drastically, which could be because the new edge weights are more comparable to the original weights and are not considered outliers as much anymore, and we note that it is more comparable to the more robust Spearman correlation coefficient. The Spearman correlation has decreased for all methods and both graphs, and so has the Jaccard coefficient compared to the previous case when we added less noise, but these values indicate more similarity after pruning for the larger network for both values of $\lambda_{\hat{E}}$.

Table 7: *More noise to new edges*: The mean absolute value of the noisy edge weights of the new edges for the original small and large synthetic graphs for case (iv) of adding noise when $\lambda_V, \lambda_E = 0.2$ and when letting $\lambda_{\hat{E}}$ vary with higher values. These are the average such values when repeating this five times. The corresponding mean values of noise added to the other components are reported in Table 4.

$\lambda_{\hat{E}}$	Small network	Large network
0.7	48.48	52.08
0.9	62.34	66.96

Table 8: *More noise to new edges*: The average Pearson and Spearman correlation coefficients, and Jaccard coefficients of the pruned results and the corresponding original networks for each of the pruning methods described in Section 4 for the five repeated results when adding noise according to case (iv) with higher values of $\lambda_{\hat{E}}$. For both networks, 20% of noise is added to the node weights and existing edge weights. The large network has 50 edges in the noisy versions (39 new edges added) and the small network has 15 edges in the noisy versions (10 new edges added).

		Small network			Large network		
$\lambda_{\hat{E}}$	Method	Pearson	Spearman	Jaccard	Pearson	Spearman	Jaccard
0.7	Original	0.66	0.68	0.33	0.47	0.49	0.22
	Edge betweenness	0.41	0.35	0.29	0.67	0.62	0.51
	PageRank	0.38	0.35	0.29	0.55	0.58	0.54
	Updated PageRank	0.34	0.32	0.29	0.64	0.57	0.49
	Connectivity kept	0.44	0.41	0.33	0.68	0.66	0.59
0.9	Original	0.51	0.53	0.33	0.25	0.28	0.22
	Edge betweenness	0.31	0.29	0.26	0.48	0.47	0.43
	PageRank	0.21	0.20	0.23	0.33	0.41	0.38
	Updated PageRank	0.12	0.09	0.17	0.35	0.31	0.32
	Connectivity kept	0.38	0.38	0.33	0.17	0.15	0.20

Furthermore, in Table 9, we further investigate if the pruning methods can identify the same most central node in the five noisy graphs as in the original graphs by computing the out-degree centrality for when $\lambda_{\hat{E}} = 0.9$ in the previous case. We note that the methods do better for the large graph, and that the Edge betweenness and Connectivity kept approaches both find a coinciding most central node for 7 noisy networks in total, which means that they are the most successful overall.

The findings based on this evaluation method seem to suggest that the pruning methods are quite successful in identifying the most important structures and that they perform comparably, but the Connectivity kept approach might do slightly better overall. It also seems to be more efficient to use the pruning results than to simply not prune at all, as we do achieve more similarity after removing redundant edges. However, the practical usefulness of the results is better evaluated by the potential users of the algorithms who have a scientific background in the field of application. Hence, we finally let human experts visually evaluate some pruning results in the next section.

Table 9: The total number of noisy networks from the five trials having the same node ranked highest by out-degree centrality as for the original versions of the small and large synthetic graphs in the case of $\lambda_{\hat{E}} = 0.9$ in Table 7, for each of the pruning methods described in Section 4, where Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV. There are 5 noisy networks in total for the small and large graphs, respectively.

Method	# small networks (/5)	# large networks (/5)
Edge betweenness	2	5
PageRank	2	4
Updated PageRank	1	4
Connectivity kept	3	4

6.3 Human expert evaluation

We now examine the results from the human expert evaluations where psychologists visually assess the quality of the pruned networks based on their expertise in psychiatric symptomatology, and in some cases also on their personal experience with the individuals' mental health and this method is described in more detail in Section 5.3. We use empirical networks with self-reported causality between symptoms, and we do not know the true underlying simplified networks in this case, but are merely presented with the networks provided by the patients themselves. However, one could argue that the psychologists have some knowledge of these underlying networks, as they are experts in the possible and probable interactions between symptoms, and some psychologists have also worked closely with the selected patients and know their individual symptomatological profiles. This means that the input of the psychologists is valuable for determining the most efficient pruning method, not only because they can assess the quality of the resulting networks, but also because this helps determine the usefulness of the methods as tools in interpreting the interactions between symptoms and identifying possible interventions.

Table 10: The pruning methods selected by the two psychologists who know the three additional patients, whose patient IDs are given in the table. The first two patients are evaluated by the same psychologist who knows both patient A and B, and patient C is evaluated by a second psychologist. The results for patients B and C were the same for both the Edge betweenness and the PageRank approaches, as well as for the Edge betweenness and Connectivity kept approaches, respectively. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV in Section 4.

Patient ID	Method
A	PageRank
B	Edge betweenness/PageRank
C	Edge betweenness/Connectivity kept

The evaluation results by the two psychologists who know three of the patients are displayed in Table 10. As an example of what the results can look like, we include the original reported network by patient A and the pruned result selected by the corresponding psychologist in Figure 13. The pruned result in Figure 13b is much simpler to interpret, and the PageRank method has identified the symptom "Feelings" as the most central symptom, along with some feedback loops. These features were not clearly distinguishable in Figure 13a, and we note that most of the retained edges had a large original weight, but some original edges with lower weight are retained as well, as for example the original edge from "Unfocused" to "Overthinking" had a smaller weight but is not removed. This illustrates for instance how the edge weight is not the only deciding factor for this pruning method.

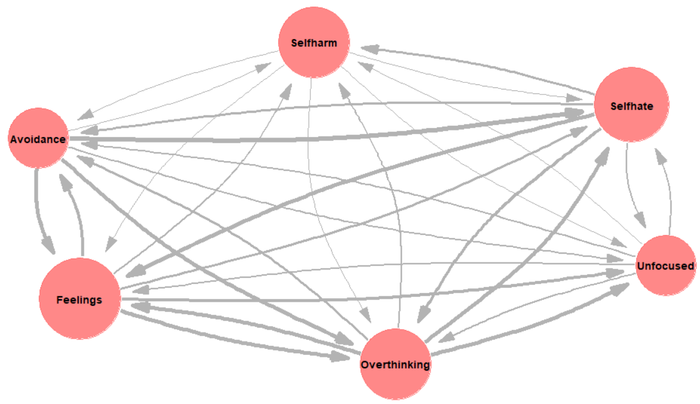

(a) Original reported network.

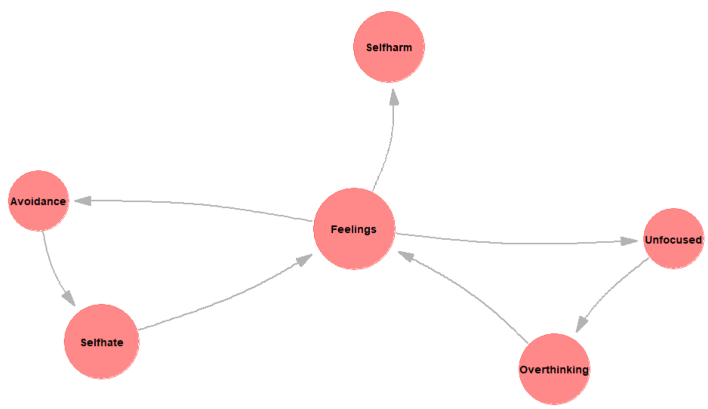

(b) Pruned result by the PageRank method.

Figure 13: An example of an empirical network and its pruned version by the PageRank method, as selected by a psychologist who knows the patient. This result is for patient A. The top network is weighted, while the bottom network only has node weights.

However, based on the table, the chosen pruning methods vary a bit with some overlapping network results and the Updated PageRank method is left out, suggesting that the approaches may be useful in different scenarios, but no method can be thought of as clearly outperforming the others based on this table. However, the Edge betweenness and PageRank approaches are chosen twice each due to overlapping results, and they yield the same network structure for patient B. Additionally, when considering the visualized networks linked in Appendix A, we note that the pruning results for patient A differ a bit between methods, but the selected result has two loops involving three symptoms, which also form a loop involving five out of six nodes, and this is also the only network with no start nodes.

The remaining results have less involved loops and one start node each, which suggests that the psychologist preferred more involved loops for this particular patient, which could be an intentional or a subconscious choice. Similarly, the chosen resulting network for patient B, which is identical to the result for the PageRank result, also has more involved loops and fewer start nodes than the remaining methods which further suggests that this psychologist prefers such network structures. However, the pruning method selected for the last patient, which yields an identical result to the pruning method only using edge betweenness centrality, only has one cycle involving two nodes whereas the remaining networks have a loop involving three nodes. This network also has two start nodes, while the other networks only have one each. Therefore, this psychologist does not favor these structures in their choice. They could instead be basing their decision on for instance the importance of the nodes involved in the loops. However, based on this, it seems like the PageRank approach and the method using only edge betweenness are more preferred by these two psychologists.

Table 11: The pruning methods selected by the three psychologists who do not know the patients from the original data, whose patient IDs are given in the table. Edge betweenness refers to Method I, PageRank to Method II. Updated PageRank to Method III, and Connectivity kept to Method IV in Section 4.

Patient ID	Psychologist 1	Psychologist 2	Psychologist 3
247	Edge betweenness	Connectivity kept	Connectivity kept
276	Edge betweenness	Connectivity kept	Updated PageRank
292	Connectivity kept	Connectivity kept	Connectivity kept

Moving on to the results for the psychologists who do not know the patients, which are displayed in Table 11, we note that the method labeled as Connectivity kept is chosen most often and Psychologist 2 selected this method throughout. For patient 247, we note that the network pruned by Connectivity kept has the most involved loops when considering the visualizations linked in Appendix A, as it contains two loops of three nodes each, that combined involve five out of the nine nodes, and only has two start nodes. The network pruned by the method labeled as Edge betweenness also has two such loops, but they only involve four nodes in their combined larger loop, and this network has one more start node. Additionally, the remaining pruned networks also have three start nodes but less complex loops. This could be a reason why these networks are not selected. Moreover, the networks for patient 276 have varied structures and the psychologists all preferred a different result, leaving out the PageRank result. We again note that the excluded result has more start nodes (three, compared to one and two in the remaining results), all results contain multiple loops of varying size. Moreover, for the last patient, 292, all psychologists agree that method using connectivity kept yields the best result. This network has the most loops and fewest start nodes, as it has four loops involving two or three nodes and only one start node, meanwhile the other networks have three loops of the same sizes and two start nodes each.

Based on the above, it seems like the psychologists who know the patients prefer the PageRank or Edge betweenness methods, while the psychologists who do not know the patients prefer the Connectivity kept approach. The latter method is also selected the most often, as it is selected more than half of the times (7 out of 12 times), and even though the Edge betweenness approach yields an identical result to the Connectivity kept result for patient 247, this means that this approach was only selected a third of the time. Therefore, it is difficult to choose a method based on these results, but it seems like the psychologists prefer the Connectivity kept and the Edge betweenness methods overall, but that the psychologists who know the patients prefer either the PageRank or the Edge betweenness methods. Since they also know the patients and have an idea of what their symptom networks should look like, this suggests that they believe these methods to be the most accurate to the subjective true underlying networks they have in mind when considering the given patients. Therefore, one may lean towards the latter approach, and it could also be considered a simpler method as it only computes the edge betweenness scores in each iteration, instead of computing the connectivity kept for each possible removal of an edge in each iteration before choosing to prune the edge with the highest value. Therefore, if one not only values accuracy but also simplicity of the pruning methods, especially if one chooses to use the algorithms for much larger networks with many more edges, then the Edge betweenness approach could be a better choice.

Moreover, the psychologists seem to prefer methods with more involved loops and fewer start nodes in general, and they commented that the networks could benefit from having more edges present in the end results. It could therefore be beneficial to decrease the parameter γ which is used in the methods to decide the number of edges to keep in the pruned versions, based on the number of edges and nodes in the networks. This value is currently set to 0.95, and could be set to an even smaller value in order to include more edges in the end results in the future even though we decided against it, since the psychologists are the potential users of the pruning algorithms and their opinions on the practical usefulness of the pruning results are more relevant.

7 Conclusions

The conclusions made from the repeated trials evaluations are that the similarity between the two networks reported by each patient seems to decrease quite significantly in general after pruning, when considering the Pearson and Spearman correlation coefficients, as well as the Jaccard correlation coefficient. However, the Edge betweenness and PageRank approaches seem more useful in detecting and extracting the similar structures in the networks as they have higher correlations than the other two methods, and they also manage to increase the similarity for more network pairs than the other methods, not only in terms of correlations and the number of distinct edges the results have in common, but also in terms of the common most central nodes. There is also weak support for a positive relationship between the original similarity between network pairs, with the strongest such trends for the PageRank approach, meaning that networks that already have some similarities might also be more similar post pruning, but it does not necessarily have to be the case. Therefore, the first evaluation method seems to favor the PageRank approach, and possibly the Edge betweenness approach, as the most precise in terms of keeping the similarity between the pruned networks from trial 1 and 2, and as the most successful in improving the similarity, even if all methods did quite poorly overall. The Connectivity kept approach is more precise in the sense of its consistency in similarity, but the similarity seems to be lower in general than for the other methods.

Nevertheless, even though the simplified empirical results are less similar after pruning, we can still find the resulting pruned networks useful as the main purpose is to simplify the networks to obtain a more interpretable structure that can be used to understand underlying structures and find treatment options. Therefore, even though we lose accuracy and precision in the results, as is somewhat expected since we actively remove information for the sake of simplicity, we still gain interpretability and this improves the understanding of the symptom interactions. This is more valuable in practice than maintaining the majority of the original information.

On the other hand, it seems like the methods do quite well to extract important structures when we use synthetic data, as we observe that all pruning methods manage to produce networks that are on average more similar after pruning than before, but the Connectivity kept approach sometimes does slightly better, and manages to more successfully find the same most central nodes together with the Edge betweenness approach. There seems to be almost no difference in robustness to noise for the former three methods, even when we vary the amount of noise added as well as which structures are affected by the noise in the original networks. All methods do better in general for the larger synthetic network, and especially when the edge weights in the original synthetic networks are much higher than the edge weights of the added edges, even if we add a large number of new edges to the network. This suggests that the pruning methods are indeed useful for finding important structures in the networks, at least in this synthetic case. However, one could also try out more combinations of levels of noise added to the different components to further investigate how the methods are affected by adding interference.

Furthermore, based on the human expert evaluations where psychologists evaluated the usefulness and reasonability of the pruned results for patients known and not known by them, it seems like those who know the patients prefer the Edge betweenness or PageRank approach, whereas those who do not know them prefer the Connectivity kept approach. Overall, the latter method is selected the most often, and therefore seems the most preferred, but the Edge betweenness method is also chosen relatively often and is one of the preferred methods by the psychologists who know the patients. Therefore, either approach might be suitable. However, the number of participating psychologists is small and one may want to even out the number of psychologists who know and do not know the patients to have more comparable data. Moreover, the psychologists seem to favor networks with more involved loops and fewer start nodes when making their choices, but we do not know if this is a conscious choice.

Hence, the Edge betweenness and PageRank approaches did better when comparing the empirical data, while all approaches seem to be robust to low levels of noise when we compare the outcomes to a true synthetic network, with the Edge betweenness and Connectivity kept approaches being able to most frequently find the same most central node as in the original data. Meanwhile, the psychologists prefer the Connectivity kept approach, together with the Edge betweenness and possibly PageRank approach. Therefore, if precision in empirical data is more important for the pruning results, one may prefer the Edge betweenness and PageRank methods, while one may be more inclined to choose the Connectivity kept or the Edge betweenness methods if the robustness to noise and quality of the visualized result are more important. However, the methods performed similarly overall, which is reasonable as three are related through extensions of one another, and it is thus difficult to choose a method that clearly outperforms the others. To choose a method, one needs to consider the tradeoff between the precision and accuracy, the quality of the visualizations, and perhaps also the complexity of the methods. It is clear that the Edge betweenness method is a bit simpler to understand and takes less time to run than the PageRank and Updated PageRank methods, as the latter two are extensions of the former, and this method is also simpler than the Connectivity kept method, which easily becomes slow as the sizes of the networks grow. This, since it computes the ratio of connectivity kept in each iteration for every edge, and this can take some time for a larger network with many edges. This is another reason to favor the Edge betweenness method, especially since a possible extension could be to gather more data by asking participants to include more symptoms and causal effects between them, which could result in many more edges being present in the empirical symptom networks.

8 Discussion

The findings suggest that the pruning methods established in this thesis are in general successful in simplifying the large, complex networks into smaller versions with fewer edges that still contain much of the essential information. However, it is difficult to choose between them as there is no method clearly out-performing the others, and there are possible ways to extend the methods for future research as there are certain assumptions made in our algorithms and evaluation methods which are potential areas of improvement.

When considering the validity of the results of the first evaluation method with repeated trials, the major drawback of this method is that we do not know what the true, underlying symptom network containing all processes looks like for each patient, or how much of its structures are present in each of the empirical networks reported by the patient. We assume that a superior simplification containing its most important features is hidden within each network for each network pair, and hope that the pruning methods are able to detect as much as possible of it in both cases. It is in a way a reasonable assumption as the data are self-reported, and the patients can perhaps more accurately describe their perceived symptoms than an external observer can identify the psychological processes experienced by another person. In this way, the individual has their own subjective idea of the underlying symptom network. The data are also collected within a few days, and the most important and most active symptomatological processes would reasonably reoccur or be maintained in this short period of time, and if not, the patient would probably still be affected enough by them to report these interactions in both survey answers. The differing nodes and/or edges would then be less important structures that perhaps do not occur as often or are not as painful if they were not included both times.

Nevertheless, it is not certain that the patient is aware of how all their symptoms interact, and it could be difficult to identify the causality between many, cross-interacting symptoms that could also be affected by external factors. Moreover, it is also possible that the second network is more accurate as the patient has had some more time to think about their symptoms and their interactions, and in this case they could choose different symptoms that replace previously chosen, less important nodes present in the first network since the survey answers are limited to selecting 9 symptoms. Another possibility is that the features common in both networks are not structurally important according to our algorithms, even if they are experienced to be painful and frequent and thus deemed important enough by the patient to report twice. This could occur if other nodes and edges are more interconnected, with those edges occurring more frequently on the shortest paths compared to the common nodes and edges. In this case, the common edges would get pruned in both cases and the similarity would decrease post pruning. These could be explanations of the decreased similarity observed by this evaluation method.

In contrast to this empirical evaluation method, the synthetic evaluation method with simulated noise has the advantage of the true simplification being given beforehand through our chosen networks. We can thus vary and affect the noise added to the network to ensure that the majority of our original structures are still more important, while we cannot do this with the empirical data. However, this is done by giving the original edges larger weight and the interfering edges a lower weight, and although edges with much higher weight correspond to much shorter paths, they could still have low edge betweenness through their position in the graph. Nevertheless, the findings show that we do improve the similarity between the synthetic network and its pruned version compared to the noisy version, which suggests that the methods successfully identify some strong structures and can improve the similarity to the true simplified network. Additionally, our synthetic networks are deterministically designed based on what we think a researcher would want to find, i.e. a central symptom with loops and out-going edges to all other symptoms, and these structures might not realistically represent the most important structures found in real symptom networks. One could also compare the outcomes for given simplifications of varied appearances with for example fewer loops, more clusters, more nodes, and more edges, to investigate the effect of certain structures on the pruning abilities. Moreover, one could add noise in other ways and could also try out larger simplified synthetic networks with more edges, as the psychologists suggested that more edges could be kept in the pruned results.

The third evaluation method with psychologists as human experts has the main advantage of potential users of the algorithms visually evaluating the practical utility of the methods for comparisons. They can give an expert opinion on which of the results is the most theoretically reasonable, and those who know the patients can also base their decision on their perception of the symptom interactions and effects for this specific individual. Hence, these latter psychologists do not only have access to a theoretical and objective underlying network of how symptoms interact in general, but they also have their own subjective idea of what the true network looks like for the individuals and what a reasonable simplification could be. Nevertheless, the sample of psychologist-patient pairs is small, and it would be beneficial to include more participants to investigate if more people prefer a given method.

Furthermore, to evaluate the similarity between networks, we use Pearson and Spearman correlations as well as the Jaccard coefficient, and the former could be impacted by the many zeros in the adjacency matrices. We actually observe this for the Pearson correlation, which assumes that the data are normally distributed and this assumption is violated when we have many outliers in the form of empty slots for potential edges and many new edges with low weight, causing skewed distributions of data. The Spearman correlation is more robust in this sense, and the Jaccard coefficient is not affected by this in the same way since we only consider the existing edges for this similarity coefficient. Moreover, the out-degree centrality measure selected to compare the most central node finds that which directly affects the highest number of other nodes in networks, but the most central node could be the one that has the strongest indirect effect on the activation of other symptoms.

Additionally, there are limitations to the established pruning methods as well that should be acknowledged. For instance, we have selected the edge betweenness centrality method to rank the importance of the edges in pruning methods I-III, which favors the edges included in the highest number of shortest paths, and this number could be the same for several edges. It also might not necessarily be the case that the edge with the highest betweenness is the most important for the network structure, as the edge betweenness centrality could for instance favor an edge connected to a peripheral start node that is connected to many other nodes in a cluster. In this way, this edge could have a high betweenness score as all the shortest paths from this node to the other nodes in the cluster must pass through this edge, but it is not intuitively considered central to the network structure itself, as it is not included in the cluster. This might have impacted the attempt for penalization of pruning edges that create such start nodes described in Section 4.5.

Thus, it could be useful to redefine the notion of edge importance using a different centrality measure that considers the importance of other edges leading to this edge through a connected node. In this way, perhaps an extension of the PageRank algorithm could be applied to the edges as we did to the end nodes, or one could use the PageRank scores of the source node instead of the end node to use the indirect causal effect of the other symptoms leading to the edge. Another possibility is to adapt the *Eigenedge* centrality for edges, developed by Huand & Huang [14], which recursively defines the importance of an edge in an undirected network in terms of the in-degrees of its directly and indirectly connected nodes, where a high degree should yield a higher Eigenedge score. The Eigenedge scores of edges should therefore be proportional to the sum of that of their surrounding edges, which leads to these scores being computed as the principal eigenvector of the edge-by-edge adjacency matrix D due to the relation $e = De$, where e is the vector of Eigenedge scores. They find that this approach yields unique values to each edge, solving our problem with edge betweenness, and that it is a faster approach that performs comparably to using edge betweenness, however, they also find that using edge betweenness on much larger graphs is not appropriate.

Moreover, our transformation of the edge weights to incorporate the importance of the nodes may not be realistic, as the perceived frequency of causality and the painfulness associated with each symptom may not be of equal importance. One could also incorporate the *treatability* of each node in the definition of their importance, as a psychologist would be interested in finding treatable nodes to target for interventions, and it would therefore be helpful if one could prune the complex symptom networks to obtain a structure where the most influential edges can be disconnected through treatment of their source node. Using this, it may also be possible to avoid the creation of untreatable start nodes. The treatability of nodes could for instance replace the painfulness in the node weights, or it could be included similarly as the painfulness is currently included but for the source nodes to compute an edge weight representing the importance not only in terms of how severe the directly caused symptom is, but also in terms of the ability to prevent the causal effect by inactivating the source node in practice. In this way, it could be interesting to use a centrality measure, such as the PageRank algorithm, to update the importance of the source node such that it takes into account the treatability of the nodes that indirectly cause its activation, and to similarly use a corresponding algorithm that gives the end node a weight based on how severe it is and how often its activation causes further indirect activation of other severe nodes that in turn lead to more severe symptoms getting triggered. Hence, this could downplay the importance of edges leading to symptoms that indirectly get treated by treating other indirectly connected symptoms.

Another possible improvement could be to use the perceived frequency of each symptom as the prior information in the PageRank algorithms, to more realistically model how the symptoms are reactivated if the random surfer gets stuck in a sink node, which could represent the case where the other symptoms are no longer actively triggering each other, unless stuck in a feedback loop. Additionally, in the edge betweenness computations, we let the edge weights be $1 - w_{ij}^{(noisy)}$ to ensure that the shortest paths have the strongest connections between symptoms, but this relationship could also be achieved through other transformations, such as by using $\frac{1}{w_{ij}^{(noisy)}}$ which would shorten the paths between symptoms that are perceived as more painful or more frequently affected compared to if we maintain a linear relationship between the new edge weights $w_{ij}^{(noisy)}$ and the transformed edge weights used in the edge betweenness computations. However, these transformations likely have similar effects as they both favor paths involving more severe symptoms and/or frequent interactions. We could also change the values of some of the parameters in our methods, such as the parameter θ in the weighted PageRank algorithm which we set to 1 to only including edge weights, and in this way also include the degree in the computations.

The stopping criterion of the pruning methods could also be improved by decreasing γ to include more edges in the end results since the psychologists commented on this, and this might also help reduce start nodes, but it could also be changed such that the algorithm decides how many edges a pruned result should have independent of the network size and more dependent on how much of the original information is maintained. One way to do this could be by incorporating a connectivity criterion, by removing edges until the connectivity has reached a threshold. One could utilize the ratio of connectivity kept used in Method IV in this thesis, with perhaps a different path quality function, or look into another measure of global graph quality. The former ratio could for instance be used to stop the pruning if the loss of connectivity by removing the next edge is too great, but one would need to quantify what is meant by "too great". A different approach would be to save the resulting graphs in each step of the pruning process, or the last couple of steps, to provide the user with more options of when to stop pruning and the decision can be made visually and subjectively based on the utility of the results.

Furthermore, when using statistical models of any kind, it is important to remember that assumptions are made and that the model is merely a glimpse and not a reflection of reality. In the symptom networks, we assume that the symptoms are distinguishable enough to be represented by their own node and that the edges represent the true interactions between the given symptoms. However, it could be the case that several symptoms are affected by external factors and these are not represented by the symptom networks in this thesis. Moreover, we also assume that the reported networks represent the patients' experienced causal effects between symptoms, but information could be missing, and the given information may not be completely accurate. There is not really a way of knowing this, since the patients report these causal effects by themselves, but the psychologists could have some intuition of how reasonable the results are. However, this method of using self-reported surveys to create networks is used as it is a simple way to gather data, and to avoid the issue of symptoms interacting on different time scales which are usually unknown. Other data collection methods are based on time series data where patients report on their symptoms several times a day for a number of weeks, as stated by Klintwall et al. [17], and this can be costly as some symptoms interact within seconds which needs to be captured by the model. Therefore, fewer symptoms can be included in the data collection simultaneously, restricting the utility of the results.

Additionally, a possible extension could be to collect more data by having the participants provide more information about the networks in the survey by choosing more nodes and edges. This would make the networks larger and more complex, and allow the pruning methods to use even more information to decide which edges are the most important for the overall interactions between symptoms. However, as the networks grow in size, the time to run the algorithms increases as well, and it could therefore be better to use the Edge betweenness or PageRank approaches rather than the Connectivity kept method, as the latter takes longer to run and the time complexity grows fast for an increasing number of edges. Furthermore, another possible direction of future work related to increasing the size of the networks is to allow for extra edges to be added based on what the psychologists think is missing in the pruned end results. As they are the experts, they have insight on how symptoms interact theoretically and in practice, and it could be possible that the participants are unaware of how some of their own symptoms interact. Moreover, it could be useful to let the psychologists give predetermined scores for the symptoms based on the theoretically known severity which takes into account other effects, such as how for instance the symptom "suicidal" can further lead to death which is not a psychological symptom, since such effects might not necessarily be reflected in the self-reported data. These scores could further be combined with the experienced painfulness scores, and possibly treatability scores, to take both theoretical and individual importance into account, which is important since everyone may not experience all symptoms equivalently. The self-reported survey is thus biased in this way, as we cannot be sure to know the true interactions of all symptoms but have to rely on the patients to adequately express their perceived interactions. Furthermore, it could also be the case that some symptoms are missing in the reported networks due to the restrictions in the survey, and there could be external factors causing symptoms to get activated that are not represented in the given symptom networks. Therefore, one may be interested in also modeling the interactions of these factors somehow, or to investigate how multiple mental disorders affect the networks.

To conclude, the evaluation results are quite varied for the pruning methods established in this thesis, and there are several areas of possible improvement. However, the findings could still be useful for future research in the field of edge pruning, and in particular to further develop more helpful ways to simplify symptom networks into more useful tools for the conceptualization of the interactions of psychiatric symptoms.

Appendix

A Data and code

This appendix provides a link to the website containing the raw data, visualizations of the original networks and their pruning results, as well as the R code used to produce the results. The visualizations are explained below and a table of the English translations of the symptoms is also made available in Table 12 below.

The original networks reported by the patients and the R code used to compute the results are available at a public GitHub repository at https://github.com/veraandersson/master_thesis. The original data are given as directed, weighted adjacency matrices in CSV files for each anonymous patient with their original node weights in a column on the side in the same file. This repository also contains the visualizations of the empirical networks and their pruning results, which are used in the evaluation with repeated trials described in Section 5.1, and the human expert evaluation of the methods which is described in Section 5.3. Table 12 shows the corresponding names of the symptoms in Swedish and their English translations, since the networks are visualized in Swedish as the survey with participating psychologists is conducted in Swedish. The Swedish versions of the networks also contain more information, as the patients have added comments about their chosen symptoms in Swedish and these are included in the visualized networks. Each page shows the results for an individual patient and the top left network is the original network as reported by the patient in the survey, where the node sizes represent the reported painfulness and the edge widths represent the reported frequency of the interactions between symptoms. Moreover, the pruned networks in the last two rows are ordered from left to right as the edge betweenness approach, the PageRank approach, the Updated PageRank approach, and the Connectivity kept approach. These are not labeled, as they were unlabeled in the versions shown to the psychologists to avoid bias towards a particular method. The results for each of the three patients whose networks have been evaluated by psychologists who know them are in a separate PDF file, and the results for the remaining 21 patients are also divided into separate PDF files for the first and second trial. The patient IDs end with the number 1 or 2, which indicate if the results are from the first or the second trials, respectively. Only the networks from trial 1 are shown for patients A, B, and C (since they only completed the survey once), while the networks for both trial 1 and 2 are shown for the remaining 21 patients.

Table 12: The English and Swedish names of the possible symptoms to choose in the survey used to collect the data, as described in Section 3, where the Swedish versions of the symptoms are used since the survey is conducted in Swedish.

English	Swedish
Insomnia	Sömn
Avoidance	Undvikande
Feelings	Känslor
Overthinking	Övertänker
Inactive	Stillasittande
Selfharm	Självskada
Substances	Substanser
Eating	Ätande
Suicidal	Självmord
Selfhate	Självkritik
Unfocused	Fokusproblem
Somatic	Kroppsligt
School pressure	Skolkrav
Family situation	Familj
Peer problems	Jämnåriga
Trauma	Trauma
Dissociates	Dissocierar
Open item	Öppet

B Synthetic and example networks

This appendix contains the data used to create the synthetic networks utilized in the thesis for computations and examples. Specifically, the adjacency matrix and node weights for the example network used in Section 2.1 can be found in Table 13. Similarly, the adjacency matrices and node weights for the small synthetic graph before and after adding noise, as illustrated in Figure 6 in Section 5.2, can be found in Tables 14 and 15, respectively. Additionally, the adjacency matrix and node weights for the large synthetic graph are displayed in Table 16, and in Figures 14 and 15 we have illustrated these synthetic networks along with their five randomly sampled noisy networks where 10 and 39 edges are added each, to obtain new noisy graphs of 15 and 50 edges, respectively. The visualizations are included for completeness.

We recall that the node weights represent the perceived painfulness of each symptom as reported by the patient, where a higher score represents more painfulness, and that the directed edge weights correspond to the perceived frequency of the interactions between symptoms, where a higher level represents a higher frequency.

Table 13: The weighted and directed adjacency matrix for the example graph in Section 2.1. The rows show the edges originating from that node, and the right-most column contains the node weights.

	1	2	3	4	5	6	7	Node weight
1	0	0	0	0	0	0	0	30
2	100	0	20	0	0	86	0	90
3	0	10	0	0	0	0	0	88
4	0	70	0	0	0	0	0	50
5	0	40	0	0	0	0	0	100
6	0	0	0	100	0	0	0	50
7	0	0	0	0	0	0	0	70

Table 14: The weighted and directed adjacency matrix for the small graph used in the evaluation method where noise is added as described in Section 5.2. The rows show the edges originating from that node, and the right-most column contains the node weights.

	1	2	3	4	5	Node weight
1	0	69	0	0	0	99
2	67	0	72	65	83	77
3	0	0	0	0	0	87
4	0	0	0	0	0	67
5	0	0	0	0	0	60

Table 15: The weighted and directed adjacency matrix for the noisy small graph in Figure 6 in the evaluation method where noise is added as described in Section 5.2. The rows show the edges originating from that node, and the right-most column contains the node weights.

	1	2	3	4	5	Node weight
1	0.00	73.70	0.00	16.88	0.00	62.16
2	46.11	0.00	70.17	84.58	86.73	50.99
3	15.84	0.00	0.00	0.00	0.00	64.47
4	12.20	0.00	0.00	0.00	0.00	75.76
5	0.00	0.00	14.65	8.42	0.00	68.58

Table 16: The weighted and directed adjacency matrix for the large graph used in the evaluation method where noise is added as described in Section 5.2. The rows show the edges originating from that node, and the right-most column contains the node weights.

	1	2	3	4	5	6	7	8	9	Node weight
1	0	72	72	74	0	86	85	0	74	85
2	82	0	0	0	0	0	0	0	0	77
3	0	0	0	0	0	0	0	0	0	76
4	0	0	0	0	0	0	0	0	0	63
5	69	0	0	0	0	0	0	0	0	92
6	69	0	0	0	0	0	0	0	0	76
7	0	0	0	0	0	0	0	0	0	74
8	73	0	0	0	0	0	0	0	0	69
9	0	0	0	0	0	0	0	61	0	86

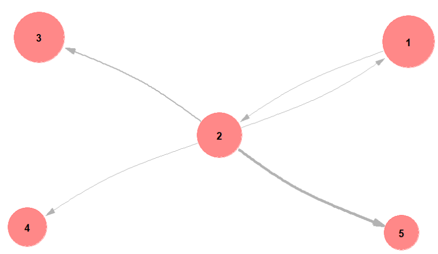

(a) Original small graph

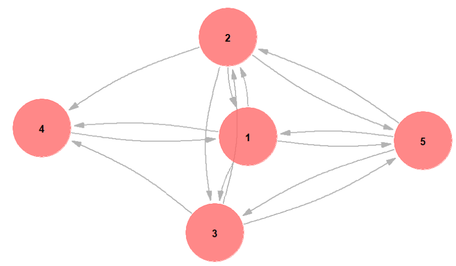

(b) Noisy 1

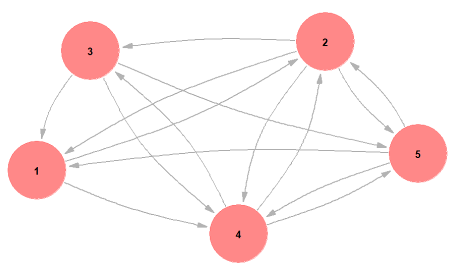

(c) Noisy 2

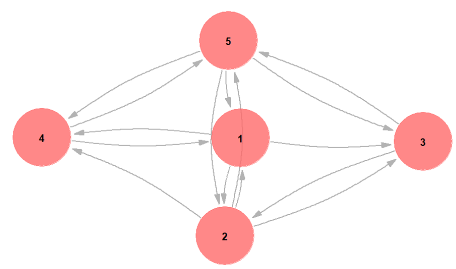

(d) Noisy 3

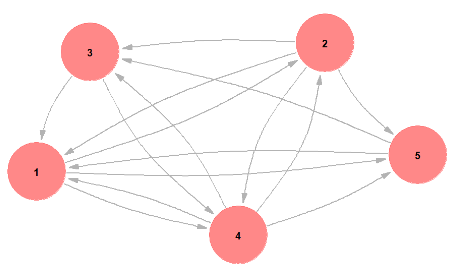

(e) Noisy 4

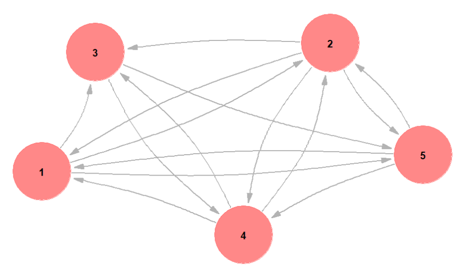

(f) Noisy 5

Figure 14: The original, weighted, small synthetic graph and its five resulting networks of when adding 10 random edges five times. The noisy networks are unweighted in this figure as the same structures were used to add random amounts of interference in several cases. The averages of their mean absolute noise and correlations with the original synthetic network are used in the analysis in Section 6.2.

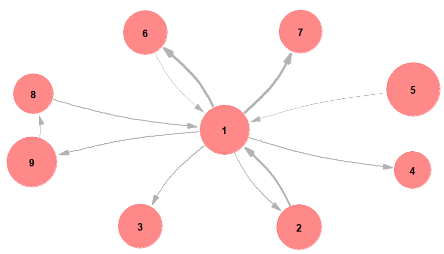

(a) Original large graph

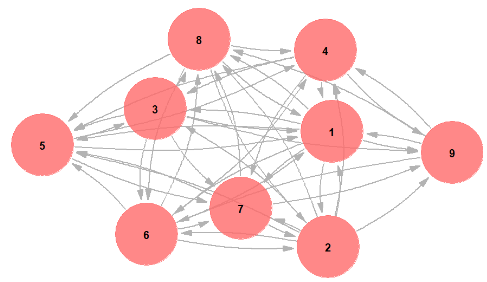

(b) Noisy 1

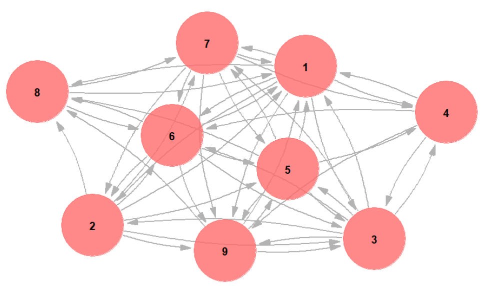

(c) Noisy 2

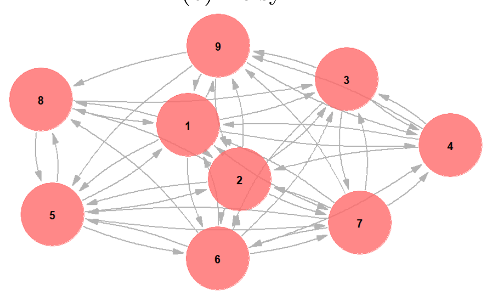

(d) Noisy 3

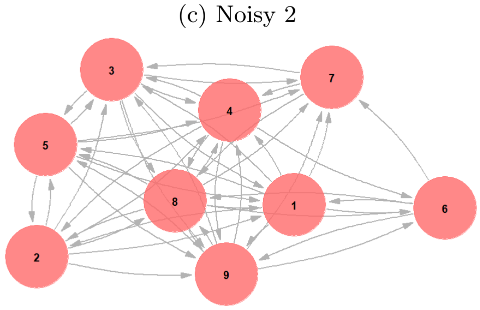

(e) Noisy 4

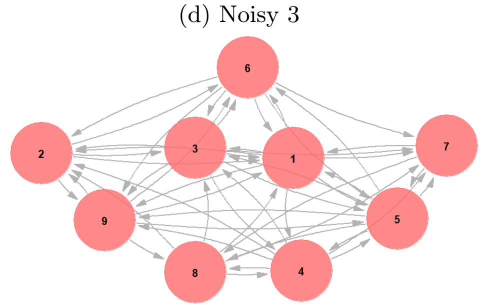

(f) Noisy 5

Figure 15: The original, weighted, large synthetic graph and its five resulting networks of when adding 39 random edges five times. The noisy networks are unweighted in this figure as the same structures were used to add random amounts of interference in several cases. The averages of their mean absolute noise and correlations with the original synthetic network are used in the analysis in Section 6.2.

9 References

[1] Alm, S. E., and Britton, T. *Stokastik: sannolikhetsteori och statistikteori med tillämpningar, [Stochastics: probability theory and statistical theory with applications].* Liber, 2008.

[2] Bondy, J. A., Murty, U. S. R., et al. *Graph theory with applications*, vol. 290. Macmillan London, 1976.

[3] Borsboom, D. A network theory of mental disorders. *World psychiatry 16*, 1 (2017), 5–13.

[4] Brandes, U. A faster algorithm for betweenness centrality. *Journal of mathematical sociology 25*, 2 (2001), 163–177.

[5] Brin, S., and Page, L. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems 30*, 1-7 (1998), 107–117.

[6] Bringmann, L. F., Albers, C., Bockting, C., Borsboom, D., Ceulemans, E., Cramer, A., Epskamp, S., Eronen, M. I., Hamaker, E., Kuppens, P., et al. Psychopathological networks: Theory, methods and practice. *Behaviour Research and Therapy 149* (2022), 104011.

[7] Dianati, N. Unwinding the hairball graph: Pruning algorithms for weighted complex networks. *Physical Review E 93*, 1 (2016), 012304.

[8] Elibol, A., and Chong, N.-Y. Topology graph pruning for optical mapping methods using edge betweenness centrality. In *2019 IEEE 4th International Conference on Advanced Robotics and Mechatronics (ICARM)* (2019), IEEE, pp. 954–959.

[9] Estrada, E. *The structure of complex networks: theory and applications.* Oxford University Press, 2012.

[10] Freeman, L. C. A set of measures of centrality based on betweenness. *Sociometry* (1977), 35–41.

[11] Gansner, E. R., Koren, Y., and North, S. Graph drawing by stress majorization. In *Graph Drawing: 12th International Symposium, GD 2004, New York, NY, USA, September 29-October 2, 2004, Revised Selected Papers 12* (2005), Springer, pp. 239–250.

[12] Girvan, M., and Newman, M. E. Community structure in social and biological networks. *Proceedings of the national academy of sciences 99*, 12 (2002), 7821–7826.

[13] Grimaldi, R. P. *Discrete and combinatorial mathematics: an applied introduction*, vol. 5. Pearson Addison-Wesley, 2004.

[14] Huang, X., and Huang, W. Eigenedge: A measure of edge centrality for big graph exploration. *Journal of Computer Languages 55* (2019), 100925.

[15] Jiang, B., and Claramunt, C. A structural approach to the model generalization of an urban street network. *GeoInformatica 8* (2004), 157–171.

[16] Kaariniemi, E., and Valter, B. Beyond the symptom list: Reliability and evaluation of utility of a structured interview to create personalized symptom networks for adolescents. Master's thesis, Karolinska Institutet, 2023.

[17] Klintwall, L., Bellander, M., and Cervin, M. Perceived causal problem networks: reliability, central problems, and clinical utility for depression. *Assessment 30*, 1 (2023), 73–83.

[18] Kolaczyk, E. D., and Csárdi, G. *Statistical Analysis of Network Data with R*, vol. 65. Springer, 2014.

[19] Strang, G., et al. *Linear algebra and learning from data*, vol. 4. Wellesley-Cambridge Press Cambridge, 2019.

[20] The World Health Organization (WHO). Depressive disorder (depression). https://www.who.int/news-room/fact-sheets/detail/depression. Last accessed on 2023-05-12.

[21] The World Health Organization (WHO). Mental health of adolescents. https://www.who.int/news-room/fact-sheets/detail/adolescent-mental-health. Last accessed on 2023-05-12.

[22] Zhang, P., Wang, T., and Yan, J. Pagerank centrality and algorithms for weighted, directed networks. *Physica A: Statistical Mechanics and its Applications 586* (2022), 126438.

[23] Zhou, F., Malher, S., and Toivonen, H. Network simplification with minimal loss of connectivity. In *2010 IEEE international conference on data mining* (2010), IEEE, pp. 659–668.