

Network Analysis of Clinical Trials on Depression: Implications for Comparative Effectiveness Research

Suresh K. Bhavnani^{1,2} PhD, Simona Carini³ MA, Jessica Ross⁴ MD MS, Ida Sim³ MD PhD

¹Center for Computational Medicine & Bioinformatics, ²MICHR, Univ. of Michigan, Ann Arbor, MI; ³Div. of General Internal Medicine, Univ. of California San Francisco, San Francisco, CA; ⁴San Francisco VA Health Center

Abstract

A comprehensive understanding of evidence related to treatments for a disease is critical for planning effective clinical care, and for designing future trials. However, it is often difficult to comprehend the available evidence because of the complex combination of interventions across trials, in addition to the limited search and retrieval tools available in databases such as ClinicalTrials.gov. Here we demonstrate the use of networks to visualize and quantitatively analyze the co-occurrence of drug interventions across trials on depression in ClinicalTrials.gov. The analysis identified general co-occurrence patterns of interventions across all depression trials, and specific co-occurrence patterns related to antidepressants and natural supplements. These results led to insights about the current state of depression trials, and to a graph-theoretic measure to categorize interventions for a disease. We conclude by discussing the opportunities and challenges of generalizing our approach to analyze comparative interventional studies for any disease.

Introduction

Researchers regularly search ClinicalTrials.gov to determine the state of human studies related to treatments for a particular disease [1]. These searches inform both comparative effectiveness research on drugs and other interventions, in addition to the design of future trials. However, the complexity of how interventions are tested with and against each other across trials makes it difficult for researchers to easily assess the breadth and depth of what has been studied [2]. For example, a search for trials that test selective serotonin reuptake inhibitors (SSRIs) does not reveal how they relate to trials that include natural supplements. However, a search for SSRIs and natural supplements together yields a long list of results that requires extensive processing to reveal which specific drugs were compared to each other.

To address such shortcomings, a few researchers have attempted to use network analyses to go beyond long lists of results. For example, Salanti et al. [3] used networks to visually represent intervention comparisons from published meta-analyses of 18 diseases. Each network showed interventions from a meta-analysis of a single disease. These networks consisted of nodes that represented interventions and

weighted edges between the nodes that represented the frequency with which two interventions co-occur across trials. The networks helped to identify comparison biases such as the absence of head-to-head comparisons between specific interventions, and gaps in evidence for each domain. However, the study was limited to trials reported in meta-analyses that focus on subsets of interventions in a domain, and therefore did not attempt to understand all the trials in a particular domain. Such analyses therefore can provide only a limited view of all the possible trials that have and are being conducted.

We therefore attempted to use networks to analyze patterns of intervention co-occurrence across an entire disease domain, namely depression. The advantage of using a single network to represent an entire domain include the potential to reveal (1) **general** patterns that have implications for the entire domain, and (2) **specific** patterns related to a subset of interventions that have implications for that subset *in context* to the rest of the domain (which is often not included in meta analyses due to its complexity).

We begin by describing how we extracted all drug and dietary supplement trials related to depression from ClinicalTrials.gov. We then describe why and how we used networks to visually and quantitatively analyze the data. These analyses revealed a new understanding of the general and specific topological properties related to depression trials. We conclude with the implications of the results for conducting similar analysis of interventions in any disease. This approach should be useful to comparative effectiveness researchers, and to network scientists interested in phenomena related to clinical trials.

Method

Our research began with the question: *How do drug treatments for depression co-occur across trials in ClinicalTrials.gov?* To address this research question, we made critical decisions regarding *data selection*, *data representation* and *data analysis*.

Data Collection. The following method was used to identify relevant depression trials from ClinicalTrials.gov. (1) We extracted 1081 interventional trials that met the criteria of “depression OR dysthymia OR dysphoric” as condition, and “drug OR dietary supplement” as intervention. (2) 87 trials were excluded because they

did not include either a drug or a dietary supplement. (3) 26 trials were excluded because they neither belonged to the domain, nor assessed depression as an outcome. This procedure resulted in 968 trials.

To address the inconsistent use of the fields provided by ClinicalTrials.gov (e.g., many trials did not have their arms clearly defined), we manually inspected and modified the interventions for the 968 trials, as follows. (1) Separated all interventions that were listed together (e.g., Escitalopram + Ramelteon) into individual interventions. (2) Excluded all text that did not pertain to the active ingredient (e.g., dosage information, and extended release). (3) Added placebo as an intervention to trials that were described as placebo-controlled, but had no placebo entered in their intervention fields. Similarly, we added placebo as an intervention to trials that were identified as islands (nodes that were disconnected from the main network) in our initial network analysis, but on inspection mentioned a placebo in other fields. Because of the above difficulties related to determining which single or multiple interventions were tested against each other, we report here only the *co-occurrence* of depression interventions in ClinicalTrials.gov trials, not their comparisons.

As the focus of our analysis was on drug and dietary supplement interventions, we classified depression interventions as follows. (1) Drugs (n=267) were classified into 9 antidepressant classes (e.g., SSRI and MAOI), and 11 non-antidepressant drug classes (e.g., Stimulant, Natural Supplement). (2) The remaining interventions were classified as Other (n=6) (e.g., Behavioral) and Usual Care (n=6) (e.g., Treatment as Usual). These classifications were independently checked by two reviewers. The above method resulted in 968 trials, and 279 unique interventions grouped in 22 intervention classes.

Data Representation. Networks are increasingly being used to analyze a wide range of phenomena, such as how diseases relate to genes [4]. A network is a graph consisting of nodes and edges; nodes represent one or more types of entities (e.g., trials or interventions), and edges between the nodes represent a specific relationship between the entities (e.g., a trial has an intervention). Figure 1 shows a bipartite network (where edges exist only between two different types of entities) of trials (black nodes) and their interventions (colored nodes). The size of each node is proportional to the number of edges (referred to the node's degree) incident to that node. Therefore large intervention nodes occur in many trials, whereas small intervention nodes occur in few.

Networks have two advantages for analyzing complex relationships. (1) They represent a particular

relationship between different nodes and therefore can reveal, for example, regularities in how specific trials are connected to specific interventions. (2) They can be rapidly visualized and analyzed using a toolbox of network algorithms to reveal global patterns in the relationships. For example, Figure 1 shows how the *Fruchterman-Reingold* layout algorithm [5], which is particularly suited for analyzing large networks, helps to visualize trials and interventions. The algorithm pulls together nodes that have common neighbors, and pushes apart nodes that do not. The result is that trials that have similar interventions are placed close to each other, and close to their interventions. All the networks were created using *Pajek* (version 1.24).

Data Analysis. We used the following visual and quantitative methods to understand general and specific patterns in the data. To identify the **general patterns** of intervention co-occurrences across trials, we first visually analyzed the bipartite network. The network revealed a pattern related to the distance that interventions and trials were from the placebo node. To quantitatively analyze this finding, we used the *k-neighbors* algorithm [6] in *Pajek* to calculate the shortest distance (least number of connected edges) of each node in the network to the placebo node. We refer to this measure as the *Placebo Distance*, which was used to mark nodes the same color if they shared the same distance. Finally, we plotted a distribution of the Placebo Distances for interventions and trials to further analyze the topological properties of the network.

To identify **specific patterns** of how a subset of interventions co-occurred across trials, we analyzed the co-occurrence of antidepressants and natural supplements across trials. This analysis was done by transforming the bipartite network of the above subset using a method called a *one-mode projection* [4]. As shown in Figure 3, all trial nodes were removed, and an edge was placed between two interventions if they co-occurred in one or more trials. This network therefore showed how frequently (based on the thickness of the edge) pairs of interventions of both classes co-occurred across trials.

Results

The analysis revealed general and specific co-occurrence patterns related to interventions for depression.

General Patterns in Intervention Co-occurrence

As shown in Figure 1, the bipartite network visually represents the explicit relationships between the 968

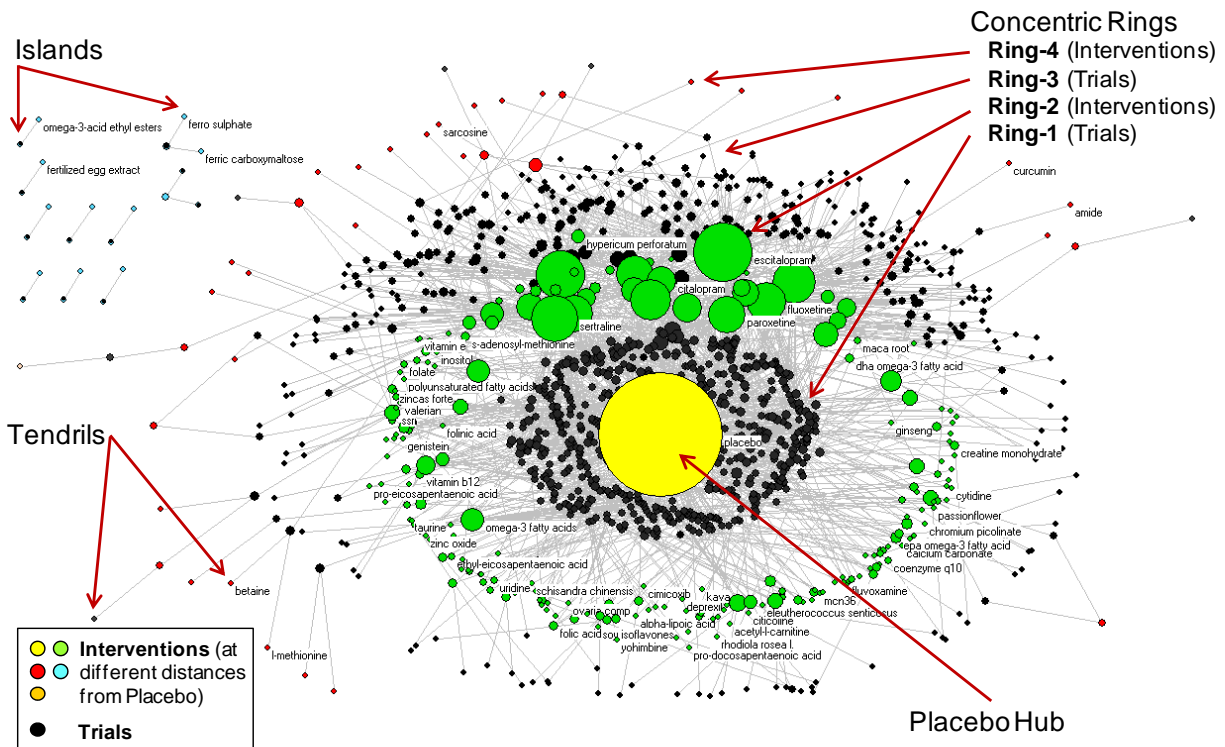


Figure 1. A bipartite network (automatically generated by the *Fruchterman-Reingold* algorithm [5]) shows how 279 interventions (colored nodes) co-occur across 968 trials (black nodes). The size of the nodes is proportional to the edges that connect to them. Therefore common interventions have large nodes, whereas rare interventions have smaller nodes. Colors represent different distances from the placebo as determined by the *k*-neighbors algorithm. To enable readability of the network, only the placebo, SSRI (a class of antidepressants), and natural supplement nodes have been labeled.

depression trials and 279 interventions. Our analysis revealed four distinct general patterns in the network:

1. High Degree Placebo Hub. The placebo node is the largest node (connected to 559 or 57% of the trials), and centrally located. In addition to the placebo node, a few other intervention nodes have a high degree (e.g., Escitalopram is connected to 126 trials). However the majority of the intervention nodes have a low degree, resulting in a right-skewed intervention degree distribution ($y = -10.11 \ln(x) + 40.464$).

2. Concentric Rings of Nodes. There are four distinct concentric rings of nodes around the placebo node. Each ring is visible in the network layout, and quantitatively identified by the *k*-neighbors algorithm with colors based on their shortest distance to the placebo. These rings alternate between trials (black nodes) and interventions (colored nodes). The **trial** nodes in Ring-1 (black nodes) are connected directly to a placebo, and therefore represent placebo-controlled trials. The **intervention** nodes in Ring-2 (colored green) are connected to trial nodes in the first ring and therefore co-occur with a placebo; many of these nodes are also connected to trials in the third ring. The **trial** nodes in Ring-3 (colored black) represent trials that do not include a placebo,

but have at least one intervention that has been tested in a placebo-controlled trial (green nodes in Ring-2). Finally, the **intervention** nodes in Ring-4 (colored red) are included in the trials in Ring-3, but have far fewer interventions compared to those in Ring-2.

3. Tendrils. There are tendrils (connected sequences of nodes with decreasing degree and terminating in a one degree node) that emanate from the network. These contain rare interventions that are a long distance from the placebo node, and have been pushed out to the periphery of the network. For example, Betaine is tested in only one trial, and is five steps removed from the placebo node.

4. Islands. There are 10 islands that are disconnected from the giant main network. These trial-intervention sets are disconnected from the rest of the trials as they include interventions that have neither a direct, nor an indirect connection, to a placebo.

Because distance from the placebo node appeared to be the main measure underlying the above network topologies, we calculated the *Placebo Distance* for each node, and plotted their distribution. The goal of generating the distribution was to relate this measure to the observed network topologies (rings, tendrils, and islands).

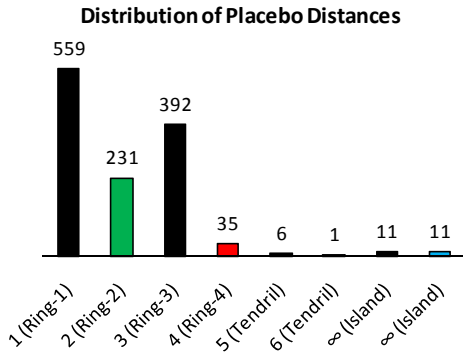


Figure 2. The distribution of 1246 Placebo Distances (excluding the distance of the placebo to itself), and their relationship to the trial (black) and intervention (colored) topologies identified in the network.

As shown in Figure 2, Placebo Distances 1-4 corresponds to Ring-1 to Ring-4, Placebo Distances 5-6 correspond to tendrils, and Placebo Distances= ∞ correspond to the trials and interventions in the islands. The distribution shows that there are 559 trials at Placebo Distance= ∞ , and the remaining non-placebo trials have different profiles based on their complex relationship with interventions at different distances from the placebo. The Placebo Distance therefore provides a richer understanding of the complexities in no-placebo trials. This understanding could enable researchers to make sense of global trends in an entire domain, and identify specific categories of interventions and trials to target for close inspection. For example, tendrils in the network could be caused when interventions are abandoned by results from earlier trials, and replaced by new ones.

Islands might exist because the interventions they contain have been tested against a placebo in another domain (e.g., a cardiac trial for impact on cardiovascular outcomes) but not tested against a placebo in the depression domain.

Specific Patterns in Intervention Co-occurrence

While the general patterns revealed how all depression interventions co-occurred across depression trials, the bipartite network in Figure 1 also revealed that all SSRI nodes (one of the 9 classes of antidepressants in the network) were all in Ring-3, whereas there were several nutritional supplements in Ring-5 and in the islands. We therefore analyzed how the specific subset of antidepressants and nutritional supplements co-occurred across trials.

Figure 3 shows a one-mode projection which represents how 33 antidepressants (colored nodes) and 51 natural supplements (black nodes) co-occur across trials. As shown, there is a tightly connected collection of colored nodes in the center of the network with no intermingling black nodes. This means that while antidepressants frequently co-occur in trials, they only infrequently co-occur with natural supplements. In fact, only 7 out of 51 natural supplements co-occur in trials with antidepressants, and in addition tend to be tested singly in independent trials. Furthermore, despite the high media and patient interest in St. John's Wort (*Hypericum Perforatum*, pointed to by the right arrow) the network revealed that it has not been broadly tested for comparative efficacy.

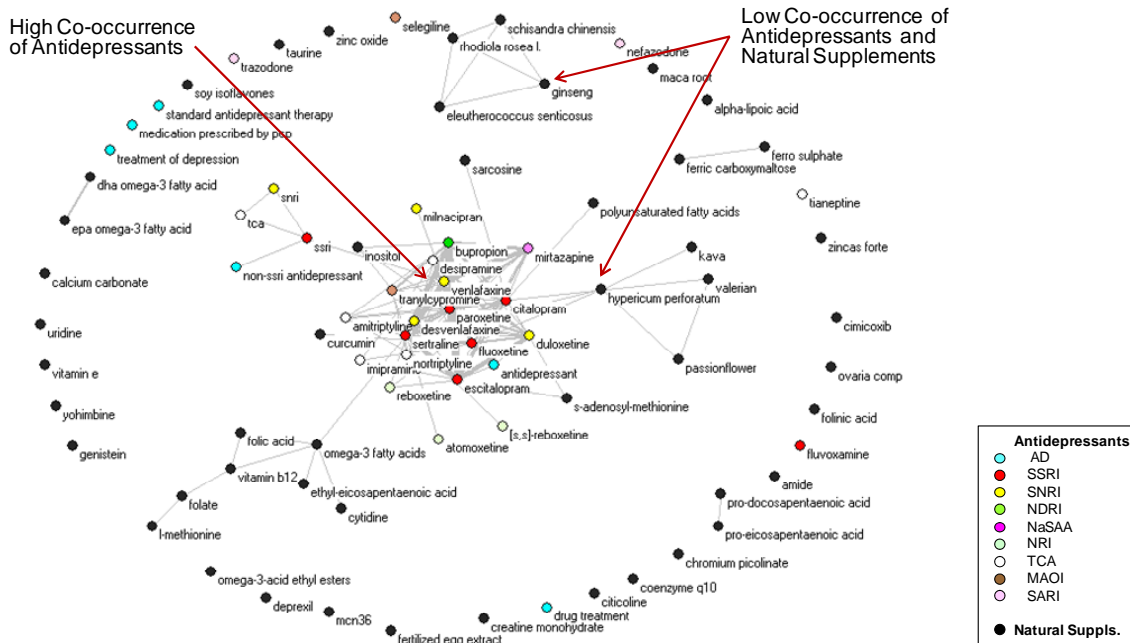


Figure 3. A one-mode network shows how antidepressants (colored nodes) co-occur with natural supplements (black nodes). The thickness of the edges is proportional to the frequency of co-occurrence. All nodes are sized equally and do not represent frequency of co-occurrence. Single disconnected nodes represent interventions that co-occurred with a placebo or with other classes of interventions (both of which were excluded to enable close inspection of the above co-occurrences).

The network also shows the relative low frequency of Tricyclic Antidepressants (TCAs) (white nodes), with Desipramine and Nortriptyline being the most commonly tested. This is probably because they are more clinically tolerable compared to other TCAs.

Implications for Comparative Effectiveness

The results have several implications for methods to improve comparative effectiveness research of treatments across all disease domains. (1) Global analysis of the trial-intervention topology of an entire domain could help make sense of the complex ways in which interventions co-occur across trials in critical domains. For example, such analyses done on all 100 priority Comparative Effectiveness Research (CER) topics listed by the Institute of Medicine [7] could provide an overview of the state of investigation in critical domains. (2) Similar network analyses could be conducted longitudinally to show for example the effect over time of funding policies on the trial-intervention topology.

However, the above methods are possible only if there are systematic attempts to address the unstructured and inconsistent nature of current clinical trials data. Therefore, while ClinicalTrials.gov was an important step to consolidate information about trials, projects such as the Human Studies Database Project (<http://hsdbwiki.org/>) should enable the capture of trial information in a well-modeled ontology of clinical research, to enable large scale visualization and analysis of human trials.

Conclusions and Future Research

While small-sized networks have been used to analyze results of meta-analyses [3], to the best of our knowledge this is the first attempt to use networks to analyze trials from an entire domain. Despite the complexity of how interventions co-occur across clinical trials (particularly in a heavily researched domain such as depression), the network quickly revealed key general and specific patterns related to interventions and trials.

At the general level, the analysis identified key network topologies (rings, tendrils, and islands) with a complex but understandable relationship to the placebo, based on the Placebo Distance. Therefore, while concepts such as placebo versus non-placebo trials, and indirect versus direct comparisons are well known, the analysis provided a deeper understanding of the relationships among trials and interventions. Future research should enable us to analyze whether this approach generalizes to trials in other domains, and whether the network results could be used to design future tools that categorize trials, and to detect patterns of comparisons among interventions.

At the specific level, the analysis revealed biases related to how antidepressants and nutritional supplements co-occur in trials, with implications for future trials. However, these biases might be caused by the selective registration of nutritional supplement trials, and therefore the current results should be combined with data from other sources.

The main limitation of this analysis stems largely from the unstructured and inconsistent nature of the data in ClinicalTrials.gov, which prevented us from analyzing which drugs were compared to each other. However, because arms' data is currently complex and time-consuming to extract for an entire domain, we believe that the co-occurrence networks demonstrated here provide a simpler first-cut understanding of interventions and trials. Such analyses could help to identify a subset of the trials for a targeted arms' based analysis.

Our future research will therefore attempt to use similar network analysis methods to analyze (1) how interventions were compared based on arms data, (2) how trial networks change over time, and (3) whether the concept of the Placebo Distance is helpful to analyze trials in other domains. Such analyses should enable a richer understanding of the evidence available from clinical trials, with the goal of enabling better treatment and design of future trials.

Acknowledgements

This research is funded in part by NIH grant # UL1RR024986, and R01-LM-06780. We thank R. Krishna, G. Vallabha, and A. Ganesan for feedback on the paper.

References

1. Zarin DA, Ide N, Tse T. et al. Issues in the Registration of Clinical Trials. *JAMA*, 2007;297:2112-2120.
2. Hernandez M, Falconer S, Storey M, Carini S. Synchronized Tag Clouds for Exploring Semi-Structured Clinical Trial Data. *Proceedings of CASCON*, 2008, 42-56.
3. Salanti G, Kavvoura FK, Ioannidis J. Exploring the Geometry of Treatment Networks, *Ann Intern Med*, April 1, 2008; 148(7): 544 - 553.
4. Newman M. The structure and function of complex networks. *SIAM Review* 2003; 45(2):167-256.
5. Fruchterman T, Reingold E. Graph drawing by force-directed placement. *Software: Practice and Experience* 1991; 21:1129-1164.
6. Nooy DW, Mrvar A, and Batagelj, V. *Exploratory Social Network Analysis with Pajek*, Cambridge, UK: Cambridge University Press.
7. <http://www.iom.edu/Reports/2009/ComparativeEffectivenessResearchPriorities.aspx>