

Special Issue

Analysis of Spatial Networks From Bipartite Projections Using the R Backbone Package

Zachary P. Neal , Rachel Domagalski , Bruce Sagan

Michigan State University, East Lansing, MI USA

Bipartite projections have become a common way to measure spatial networks. They are now used in many subfields of geography, and are among the most common ways to measure the world city network, where intercity links are inferred from firm co-location patterns. Bipartite projections are attractive because a network can be indirectly inferred from readily available data. However, spatial bipartite projections are difficult to analyze because the links in these networks are weighted, and larger weights do not necessarily indicate stronger or more important connections. Methods for extracting the backbone of bipartite projections offer a solution by using statistical models for identifying the links that have statistically significant weights. In this article, we introduce the open-source `backbone` R package, which implements several backbone models, and demonstrate its key features by using it to measure a world city network.

Introduction

Spatial analysis and quantitative geography have a long history of using network analysis (e.g., Haggett and Chorley 1969; Neal 2013b; Smith and Timberlake 1995; Ter Wal and Boschma 2009). Although there are many ways to measure spatial networks, bipartite projections have emerged as one of the most widely used approaches. A bipartite projection defines a network among a set of nodes (e.g., cities, countries) in which the strength of the connections between them is measured using their number of shared attributes (e.g., the number of firms located in two cities, the number of treaties signed by two counties). This approach has become a *de facto* method for not only measuring the world city networks (Taylor 2001; Taylor and Derudder 2016), but is also used in other areas of geography at multiple geographic scales: at the macroscale bipartite projections measure networks of international relations (e.g., Hafner-Burton, Kahler, and Montgomery 2009), at the microscale they measure neighborhood social networks (e.g., Browning et al. 2017), and at a meta-scale they have been used to study the structure of schools of thought in geography (e.g., Peris, Meijers, and van Ham, 2018). Despite their widespread adoption, using bipartite projections to measure spatial networks is not always straightforward. In this article, we introduce and demonstrate the open-source `backbone` R package,

Correspondence: Zachary P. Neal, Michigan State University, 316 Physics Rd., Rm 262, East Lansing, MI e-mail: zpneal@msu.edu

Submitted: July 29, 2020; Revised version accepted: December 16, 2020

which is a general-purpose set of commands for constructing bipartite projections, focusing on its applications for spatial networks.

The article is organized into four sections. In the first section, we provide a brief introduction to bipartite projections, reviewing their use in spatial analysis, noting key methodological challenges, and describing backbones as a solution. In the second section, we introduce the `backbone` package, providing an overview of its syntax and functions. In the third section, we provide a replicable demonstration of the `backbone` package in the context of spatial analysis, using it to examine the world city network and identify the most central cities. Finally, we conclude in the fourth section by providing recommendations for using bipartite projections to measure spatial networks.

Background

Introduction to bipartite networks and projections

A (unipartite) network is a collection of objects, called *nodes*, and connections, called *edges*, between pairs of nodes. It can be represented visually as a graph or sociogram, where shapes represent nodes, which are connected by lines representing edges. It can also be represented mathematically as a square matrix, where the rows and columns represent nodes, and the cells indicate whether (or how strong) the edge is connecting the respective row and column nodes.

A bipartite network is a type of network composed of two sets of nodes, which following Neal (2014a) we call *agents* and *artifacts*, in which an edge can only connect an agent to an artifact. Bipartite networks are often also called two-mode networks because they contain two types of nodes, or affiliation networks because they describe how agents are affiliated with artifacts. They can be represented visually as a graph or mathematically as a rectangular matrix \mathbf{B} , where the rows represent agents, the columns represent artifacts, and cell $B_{ij} = 1$ if agent i is connected to artifact j , and otherwise is 0.

A bipartite network can be transformed into a unipartite network via projection. The projection of a bipartite network is computed as $\mathbf{P} = \mathbf{B}\mathbf{B}'$, where \mathbf{B}' indicates the transpose of \mathbf{B} . So \mathbf{P} is a symmetric square matrix, where the rows and columns represent the agents in \mathbf{B} and cell P_{ij} contains the number of artifacts shared by agents i and j for $i \neq j$. Cell P_{ii} contains the number of artifacts associated with agent i , but in practice is ignored in analysis.¹

Fig. 1 illustrates a simple bipartite network (left) and its projection (right), each represented as both a graph (top) and matrix (bottom). The bipartite graph shows five agents (squares) and their connections to four artifacts (circles), while the bipartite matrix shows the pattern of agent–artifact connections using 0s and 1s. The row and column sums of the bipartite matrix capture the total number of connections of each agent and artifact, respectively. The bipartite projection graph shows these five agents connected to each other, to the extent that they share artifacts. For example, A and B are connected because they share artifact 1, while C and E are connected twice (i.e., with an edge of weight 2) because they share both artifact 3 and artifact 4. Notably, the fact that C is associated with artifact 2 plays no role in the projection because no other agent is associated with this artifact. The bipartite projection matrix shows the number of shared artifacts for each pair of agents. Notably, the diagonal cells in the projection matrix indicate each agent’s total number of artifacts (e.g., C is associated with 4 artifacts in total), but are not represented in the projection graph and are ignored in subsequent network analysis.

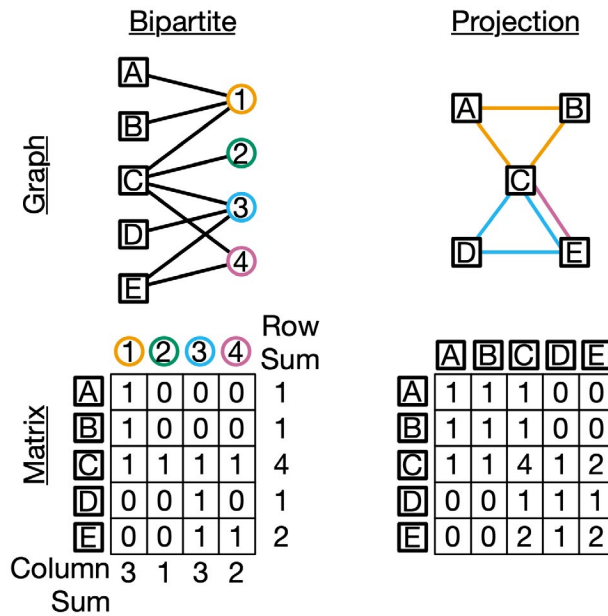


Figure 1. Bipartite and bipartite projection networks

Bipartite projection networks in spatial analysis

Bipartite projections appear in many contexts (Vasques Filho and O’Neale 2020), including spatial analysis, where they can take two distinct forms depending on whether the agents or artifacts are spatial entities (i.e., locations). In the *locations-as-agents* approach, a spatial bipartite projection is a network of locations, such that a pair of locations is connected to the extent that they share artifacts. Calling it the “interlocking world city network model,” this is the approach that Taylor (2001) proposed and which launched a wave of research on world city networks: major cities (the agents, which are locations) are connected to the extent that they house branch offices of the same advanced producer services firms (e.g., finance, accounting, consulting; the artifacts). It rests on the logic that offices of the same firm must communicate and interact with one another, and therefore that when two cities have an office of the same firm, there is likely interaction between them. Spatial networks adopting the *locations-as-agents* approach to measurement via bipartite projection are quite common at multiple spatial scales, and have been used to measure networks among urban locations connected by twitter users (Poorthuis 2018), bus routes (Liu and Duan 2020), networks among cities connected by patents (Balland and Rigby 2017), banking syndicates (Pažitka, Wójcik, and Knight 2021), networks among countries connected by treaties (Hafner-Burton et al. 2009), trade (Straka, Caldarelli, and Saracco 2017), and corporate executives (Heemskerck, Fennema, and Carroll 2016).

In the *locations-as-artifacts* approach, a spatial bipartite projection is a network of agents (often people or other social actors), such that a pair of agents is connected to the extent that they share locations. The *locations-as-artifacts* approach is less common in geography because the spatial units play only an instrumental role in the network, forging the links between agents, but

do not appear in the bipartite projection network itself. However, it is common in sociological research, where the focus is on social networks emerging from spatial interactions. For example, Browning et al. (2017) and Xi, Calder, and Browning (2020) use this approach to measure and study the social network among households in Los Angeles: households (the agents) are connected to the extent that they visit the same routine activity locations (e.g., school, work; the artifacts). This rests on the logic that places offer opportunities for casual encounters which lead to the formation of social bonds, and therefore when two households frequent the same places, they are more likely to interact with each other (Jacobs 1961). Hidalgo et al. (2007) adopted a similar locations-as-artifacts approach to derive a ‘product space’ in which export products were connected to the extent that they were exported by the same countries. This follows the logic that “if [the production of] two goods...require similar institutions, infrastructure, physical factors, technology, or some combination thereof, they will tend to be produced [in the same location],” and therefore, the spatial co-production of products indirectly captures their production technology similarity (Hidalgo et al. 2007, p. 484).

There is an important link between these two approaches. When \mathbf{B} is a bipartite network where the rows represent locations, then \mathbf{BB}' will yield a locations-as-agents bipartite projection, while $\mathbf{B}'\mathbf{B}$ will yield a locations-as-artifacts bipartite projection. Therefore, a single bipartite network can be studied from both perspectives. For example, although the world cities literature usually focuses on cities linked by sharing firms, some have simultaneously examined a network of firms linked by their co-location in cities (e.g., Neal 2008; Van Meeteren, Neal, and Derudder 2016). Similarly, Straka et al. (2017) examined not only a network of countries linked by trading the same products, but also a network of products that are traded by the same countries.

The key advantage to measuring spatial networks using bipartite projections lies in the relative ease of data collection. For example, data about economic exchanges between cities may not be available from official government sources, and collecting such data directly is often impractical. However, data about where firms’ offices are located is readily available, usually on the firms’ own websites. Accordingly, bipartite projections offer a practical way for researchers to indirectly approximate a city-level economic network. Similarly, because social network analysis requires data from a population (not a sample) and is sensitive to missingness, it is often impractical to collect data on the social network among residents of a large city. However, data about the places residents visit or tweet about can be collected using routine surveys, remote sensing, and digital trace measures. Accordingly, bipartite projections also offer a practical way for researchers to indirectly approximate social networks in large geographic areas.

Challenges with bipartite projections

Although bipartite projections are promising for measuring spatial networks, they also present some significant challenges. Some of these challenges are conceptual or theoretical. For example, bipartite projection allows *any* rectangular matrix of 0s and 1s to be transformed into a symmetric square matrix that *resembles* a network, but this does not necessarily mean it can be interpreted and analyzed as a network (Derudder 2020; Neal 2021; Nordlund 2004). The suitability of a bipartite projection as an indirect approximation of a spatial network depends on the researcher’s ability to articulate a theory about why the sharing of artifacts suggests a connection between two locations, or about why the sharing of locations suggests a connection between two agents, and a description of the type of connection such a phenomenon represents. It is important to emphasize that *in the absence of such a theory, bipartite projections are not appropriate for*

measuring spatial networks. As a theoretical challenge, it is not resolvable through the use of open-source software or indeed by any methodological tools.

Determining whether or not using a bipartite projection is an appropriate way to measure a spatial network is rarely straightforward (Derudder 2020; Neal 2014b, 2021; Pažitka et al. 2021). However, for the sake of clarity, consider two contrasting cases. In the first case, a researcher collects data in the form of a binary rectangular matrix where the rows are countries, and the columns are colors. A cell in this matrix contains a 1 if the country’s flag contains the respective color (e.g., $B_{\text{USA, red}} = 1$ and $B_{\text{USA, green}} = 0$). The researcher then constructs a bipartite projection from these data and analyzes it as a network. Although this exercise is mathematically possible, it is unlikely that a network in which countries are connected by shared flag colors has any real meaning; this type of analysis should be avoided. In the second case, Taylor (2001) collects data on firm locations in cities, then constructs a bipartite projection and analyzes it as a network in which cities are connected to the extent that they share firms. To justify this measurement approach he explicitly articulates a theory, drawing on Sassen (1991), that multinational firms share information through their global office networks and therefore intra-firm office networks provide information about flows between cities. Although other researchers may disagree with this rationale (e.g., Pažitka et al. 2021), an explicit theory about the meaning of the network exists and can be evaluated; this type of analysis is the essence of science and should be pursued...with caution.

When a bipartite projection is a theoretically sound approach to measuring a spatial network, the researcher must then confront several methodological challenges. A bipartite projection “transforms the problem of analysing a bipartite structure into the problem of analysing a weighted one, which is not easier” (Latapy, Magnien, and Del Vecchio 2008, p. 34–35). As Fig. 1 illustrates, all bipartite projections are weighted networks, where the weights capture the number of artifacts shared by two agents. Although the analysis of any weighted network can be complex (Newman 2004), the analysis of a weighted bipartite projection is particularly difficult because larger edge weights do not necessarily indicate stronger or more important connections.

A standard solution to the challenge of analyzing a weighted bipartite projection has been to transform it into an unweighted network by applying a universal threshold: edges with weights above the threshold are kept, while weaker edges are discarded. However, this solution can distort the structure of the network. First, it ensures that “nodes with small [degree centrality] are systematically overlooked,” yielding a network focused only on the agents that are most well connected in the original bipartite network (Serrano, Boguná\xAD, and Vespignani 2009, p. 6484). This helps explain why many studies of world city networks focus on cities with strong connections such as New York and London, while cities with weaker connections are “off the map” (Robinson 2002). Second, it ensures that “even a random bipartite network – one that has no particular structure built into it at all – will be highly clustered” (Watts 2008, p. 128). This helps explain why world city networks almost always contain clusters or ‘cliques’ of cities (Derudder and Taylor 2005). Finally, some network structures, such as open triads (e.g., a trade circuit) and stars (e.g., a hub-and-spoke transportation arrangement) are not observable. This helps explain why trade brokerage is rarely observed in city networks measured using bipartite projection (Neal 2012), but is readily observable in city networks measured using other methods (Martinus et al. 2021).

Backbones of bipartite projections

To overcome these challenges, it is necessary to extract the backbone of the weighted bipartite projection by using a statistical test to identify the most important (i.e., statistically significant) edges, which are preserved in an unweighted *backbone* network. The statistical tests used by different backbone models all aim to answer the same question: “Is the weight of the edge between two agents stronger than would be expected at random?” Answering this question involves comparing an edge’s observed weight to the distribution of weights it would have if some features of the original network were preserved, but the network was otherwise random.

A large class of such backbone models already exist for extracting the backbone of weighted networks that are *not* the product of a bipartite projection (e.g., Dianati 2016; Serrano et al. 2009). Such natively unipartite weighted networks arise frequently in spatial analysis, for example, in the form of transportation networks where the weights of edges directly capture flows from one location to another, and not (as they would in a bipartite projection) the number of shared artifacts. However, these models should not be used for extracting the backbone of bipartite projections because stronger edges in a bipartite projection are not necessarily more important. Instead, it is necessary to use backbone extraction models developed specifically for bipartite projections.

Three models for extracting the backbone of bipartite projections are implemented in the *backbone* package we introduce below: the hypergeometric model (HM), the fixed degree sequence model (FDSM), and the stochastic degree sequence model (SDSM). The mathematical details of these models are described by Domagalski, Neal, and Sagan (2021), however they differ solely in how they define “at random” when asking “Is the weight of the edge between two agents stronger than would be expected at random?” Here, we briefly sketch their definition of random and its implications for their scope of application.

The statistical test used by the HM to determine when an edge weight is statistically significant controls for the row sums of \mathbf{B} (i.e., the number of artifacts associated with each agent), but not for the column sums of \mathbf{B} (i.e., the number of agents associated with each artifact). It is most suitable for application to cases where the column sums are (nearly) equal, or are unimportant, and therefore do not need to be controlled. In practice, this is likely to be rare in spatial data. For example, there is substantial variation in the number of cities (agents) in which different firms (artifacts) are located; some firms are big and maintain locations in many cities, while other firms are small and maintain locations in just a few cities. This variation likely matters for making inferences about which cities have economic interactions.

The statistical test used by the FDSM is more restrictive, controlling for *both* the row and column sums of \mathbf{B} (Zweig and Kaufmann 2011). The FDSM is more appropriate than the HM for most spatial data because it is able to control for variation in both these features of the data. However, this additional control comes at a high computational cost; the FDSM relies on a numerical simulation that can require a significant and sometimes impractical amount of time to extract the backbone of a large bipartite projection.

Finally, the statistical test used by the SDSM *approximately* controls for both the row and column sums of \mathbf{B} (Neal 2014a). By approximately controlling for both features of the data, the SDSM yield backbones that are similar to those generated by FDSM, but does so more efficiently. Therefore, the SDSM is often a reasonable choice for the extraction of backbones from most bipartite projections, particularly when the FDSM is computationally impractical.

Each of these backbone models has previously been used to study spatial networks. Neal (2013a) used the HM to study linkages formed in the world city network by a process through which firms are sorted into cities. In contrast to a conventional world city network dominated

by global financial capitals, he described a network structured by national institutions such as the US Federal Reserve banking system. Van Meeteren et al. (2016) used the SDSM to study agglomeration patterns, finding that advanced producer service firms agglomerate intranationally and pursue sector-specific global location strategies. Finally, Poorthuis (2018) used the FDSM to identify neighborhoods as clusters of locations tagged by twitter users. Although backbone models are increasingly widely used to approximate spatial networks from bipartite projections, progress and transparency have been limited by the lack of software that implements these models.

The backbone package

The `backbone` package is an open-source collection of commands for R that facilitates the analysis of bipartite projections (Domagalski et al. 2021; R Core Team 2018). It is freely available from the Comprehensive R Archive Network (CRAN). To install, load, and verify the version of the package, type:

```
> install.packages("backbone") #install
> library(backbone) #load
> sessionInfo() #verify
R version 4.0.2 (2020-06-22)
other attached packages:
[1] backbone_1.2.2
```

This article describes `backbone v1.2.2` running in `R v4.0.2`, and the example is intended for use with these or newer versions. We present only the package's primary functions, with a focus on their application for spatial networks. Details about these functions' formal mathematical specification are described by Domagalski et al. (2021) and complete documentation of all commands available in the package is available by typing:

```
> ?backbone    #for documentation
> vignette("backbone")    #for an example
```

The `backbone` package is composed of three types of functions. First, the `universal()` function constructs conventional weighted bipartite projections, as well as simple backbones using a universal threshold. Second, the `hyperg()`, `sds()`, and `fds()` functions derive probability distributions that can be used to test the statistical significance of edges in a weighted bipartite projection using the hypergeometric, SDSM, and FDSM models, respectively. Finally, the `backbone.extract()` function constructs a backbone network that contains only the

statistically significant edges. Each of these functions offers several optional parameters to customize their output; we illustrate the most commonly used options in the next section.

Using *backbone* to examine the World City Network

The *backbone* package is a general-purpose set of commands designed to facilitate the analysis of bipartite projections. In the context of spatial analysis, it can be used for research adopting a locations-as-agents approach, to infer the spatial network among a set of locations from data on their shared characteristics. However, it can also be used for research adopting a locations-as-artifacts approach, to infer a social network among a set of actors from data on their shared locations. To illustrate *backbone*'s application in one specific spatial analytic context, in this section, we demonstrate its use to examine the world city network and identify the most central cities in it. We selected this context for illustration for two reasons. First, the topic of world city networks has been the subject of many recent *Geographical Analysis* articles (e.g., Derudder 2020; Liu and Derudder 2012; Neal 2012, 2020; Taylor 2001), some of which are among the journal's most highly cited (Franklin 2021). Second, the analyses can be easily replicated by readers because one widely studied bipartite data set concerning world cities is publicly available. We intend the example analyses presented below to serve as an illustration of the *backbone* package, and not necessarily to make novel contributions to the substantive literature on world city networks. These analyses can be replicated by pasting the code below into R, however the complete replication R script and data are also available at <https://osf.io/r2evn/>.

Data

The Globalization and World Cities (GaWC) "Data Set 11" was originally collected in 2000, and records the extent of 100 advanced producer services firms' presence in each of 315 large cities (Taylor, Catalano, and Walker 2002). These data served as the foundation for one of the earliest and most comprehensive empirical studies of the world city network (Taylor 2004), and as a template for a substantial body of empirical research conducted by those associated with the GaWC research network.

Formally, the data set takes the form of a rectangular 315×100 bipartite matrix \mathbf{B} , in which B_{ij} contains the "service value" of firm j 's presence in city i . The service values are an ordinal scale intended to capture the importance or extent of a firm's presence in a city, and ranged from 0 (no presence) to 5 (global headquarters), with a value of 2 representing an presence that provides "the 'normal' or 'typical' service level of the given firm in a city" (Taylor et al. 2002, p. 2370). These publicly available data can be loaded into R directly from the GaWC website (as of 5 November 2020) and converted to matrix form:

```
> B <- read.csv(file="https://www.lboro.ac.
                uk/gawc/datasets/da11.csv",
                header = TRUE,
                row.names = 1)
> B <- as.matrix(B)
```


The `backbone` package is designed for use with binary bipartite data, so for this illustration we transform the original ordinal \mathbf{B} into a binary \mathbf{B}' such that

$$B'_{ij} = \begin{cases} 1 & \text{if } B_{ij} \geq 3 \\ 0 & \text{if } B_{ij} \leq 2 \end{cases}.$$

This transformation can be achieved, and the cities that contain no firms with a larger-than-typical presence can be excluded, by typing:

```
> B[B <= 2] <- 0
> B[B >= 3] <- 1
> B <- B[rowSums(B) != 0,]
```

This transformation allows us to focus only on firms that maintain a larger-than-typical presence in a given city, and only on the 196 cities that contain at least one such firm.² For convenience, we use \mathbf{B} to refer to this binary matrix in the remainder of this section.

Once the bipartite data has been loaded and transformed, it is possible to examine some of its features. For example, it is possible to look at the pattern of firms' presence in cities.

```
> B[114:117,8:11]
           Horwath KPMG Summit...Baker RSM
MELBOURNE      0    1           0    1
MEXICO CITY     0    1           0    0
MIAMI           1    1           0    1
MILAN           0    0           0    1
```

This command shows the portion of \mathbf{B} that includes the 114th to 117th cities, and 8th to 11th firms. The output shows that while the accounting firms of KPMG and RSM maintained offices in several of these cities, Horwath and Summit International+Baker Tilley did not.

Two key characteristics of any bipartite data are the row sums and column sums. In these data, the row sums indicate the number of firms located in a city, while the column sums indicate the number of cities in which a firm maintains a presence.

```
> rowSums(B) ["AMSTERDAM"]
AMSTERDAM
      29

> rowSums(B) ["NEW YORK"]
NEW YORK
      74

> colSums(B) ["KPMG"]
KPMG
      76

> colSums(B) ["HSBC"]
HSBC
      43
```

For example, there are 74 firms that maintain a larger-than-typical presence in New York, but only 29 firms that maintain a larger-than-typical presence in Amsterdam. Likewise, KPMG maintains a larger-than-typical presence in 76 cities, while HSBC maintains a larger-than-typical presence in only 43 cities. Fig. 2 illustrates these values for all cities and firms in these data. Specifically, Fig. 2A shows that while most cities contain fewer than 20 firms, some cities contain many more firms. Similarly, Fig. 2B shows that while most firms maintain a presence fewer than 40 cities, some firms maintain a presence of many more cities.

Weighted bipartite projections

The conventional “specification of the world city network” used in GaWC research involves computing a weighted bipartite projection \mathbf{P} from the original bipartite data \mathbf{B} Taylor (2001).

```
> P <- B %*% t(B)
```

Following this specification, the cities are treated as agents and the firms are treated as artifacts. The resulting square matrix \mathbf{P} is treated as a weighted world city network in which the strength of the connection between a pair of cities is measured by their number of co-located firms. For example, examining the matrix cell corresponding to the connection between Amsterdam and

```
> P["AMSTERDAM", "NEW YORK"]
[1] 26
```

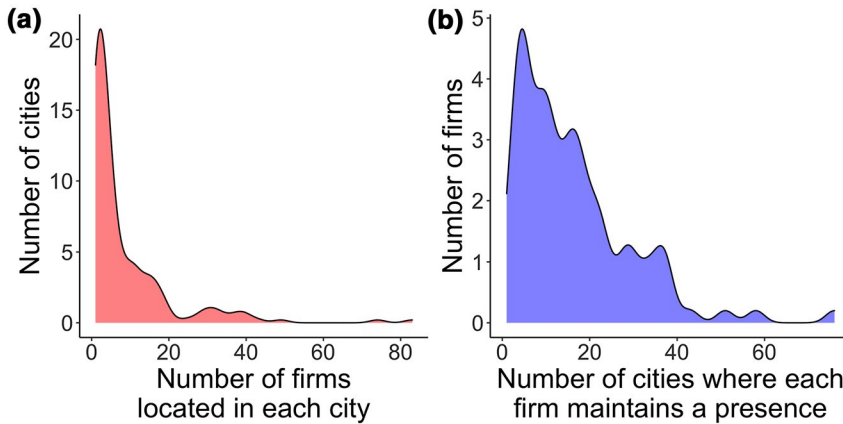


Figure 2. The distribution of (A) row sums and (B) column sums in the GaWC Data set 11.

New York indicates that 26 firms maintain a presence in both cities, and might be interpreted as evidence that they interact economically.

Many analyses of the world city network focus on cities' degree centrality, or what is sometimes called a city's "global network centrality" (GNC). This value measures a city's total number or strength of connections in the network, and is interpreted as an indicator of a city's status or importance in the network.

```
> sort(rowSums(P), decreasing = TRUE)[1:5]
LONDON NEW YORK PARIS HONG KONG SINGAPORE
1496    1403  1043    1032    913
```

In these data, London and New York have the greatest centrality, occupying the top tier of the urban hierarchy as what GaWC research calls *Alpha++* cities (Beaverstock, Smith, and Taylor 1999). They are followed by a second tier of *Alpha+* cities that include Paris, Hong Kong, and Singapore. This approach appears to successfully identify what nearly any scholar of globalization would regard as the cities "used by global capital as basing points in the spatial organization and articulation of production and markets" (Friedmann 1986, p. 71).

However, these values and this weighted spatial network are less informative than they might seem. The centrality values derived from this network are almost perfectly correlated with the number of firms located in each city (i.e. the row sums of **B**).

```
> cor(rowSums(P), rowSums(B))
[1] 0.9767704
```

The high correlation indicates that this approach to identifying central cities in a world city network is actually just identifying cities that contain many firms. This occurs because measuring a world city network using a weighted bipartite projection of firm locations guarantees that cities with many firms will have stronger connections and larger centrality values (Neal 2012). If world city researchers were simply interested in finding cities with many firms, there are much simpler ways achieve this (e.g., counting a city's number of firms).

The backbone of the world city network

In practice, world city researchers are interested in something more nuanced: studying cities that are central in a network of economic interactions. The challenge is that although firm co-location may provide information about which cities interact economically, firm co-location is not the same as economic interaction. The `backbone` package can be used to make inferences about which cities are engaged in economic interaction based on firm co-location patterns. Specifically, it can be used to estimate whether the number of firms co-located in two cities is large enough to warrant concluding that the two cities are engaged in meaningful economic interaction. The *backbone* of the world city network is a binary network in which pairs of cities are connected only if their number of co-located firms suggests they are engaged in meaningful economic interaction, and therefore provides a simplified and potentially more focused depiction of the world city network. The `backbone` package offers four ways to make such inferences and extract this backbone.

Using universal thresholds

The `universal()` function offers the simplest approach to extracting the backbone of the world city network by applying a single researcher-specified threshold value to all city pairs. Given a threshold T , any pair of cities with more than T co-located firms is defined as connected in the network. For example, choosing $T = 0$ implies that *any* number of firm co-locations is interpreted as evidence of economic interaction between a pair of cities. Extracting the backbone using a threshold of 0 is achieved by typing:

```
> universal0 <- universal(B, upper = 0,  
                          bipartite = TRUE)
```

This command extracts the backbone from an input data set `B`, which is a bipartite matrix (i.e., `bipartite = TRUE`), by applying an upper threshold of 0 (i.e., `upper = 0`), and stores it in a new matrix `universal0`. Once extracted, it is possible to examine the features of this universal threshold backbone:

```

> table(universal0$backbone)
  0    1
21506 16910

> mean(universal0$backbone)

[1] 0.4401812

> sort(rowSums(universal0$backbone),
       decreasing = TRUE)[1:5]

LONDON NEW YORK PARIS HONG KONG LOS ANGELES
  191    185   175     171     171

> cor(rowSums(universal0$backbone),
      rowSums(B))

[1] 0.7407175

```

Unlike the weighted bipartite projection, the backbone is a binary network; pairs of cities either are ($N = 16910$) or are not ($N = 21506$) connected. A backbone extracted using $T = 0$ is quite dense (44% of possible intercity connections are present) because it treats even small numbers of firm co-locations as evidence of economic interaction between cities. As a result, the most central cities are still obviously large cities that contain many firms, and indeed, cities' centrality in this network remains highly correlated ($r = 0.74$) with their total number of firms.

A sparser network containing fewer intercity connections can be obtained using a higher (i.e., more stringent) threshold that retains only particularly strong connections (e.g., Derudder and Taylor 2005). For example, the `universal()` function can be used to extract a backbone where $T = 25$ and therefore only cities with more than 25 co-located firms are counted as connected:

```

> universal25 <- universal(B, upper = 25,
                          bipartite = TRUE)

> mean(universal25$backbone)

[1] 0.001665973

> sort(rowSums(universal25$backbone),
       decreasing = TRUE)[1:5]

LONDON  NEW YORK HONG KONG PARIS CHICAGO
   15      12      5      5      3

> cor(rowSums(universal25$backbone),
      rowSums(B))

[1] 0.8381523

```

This more stringent universal threshold is indeed much less dense (only 0.16% of possible edges are present). However, it still remains focused on the largest cities, whose centrality is highly correlated ($r = 0.84$) with the total number of firms.

Both of these approaches involve an arbitrarily selected threshold; however, the `universal()` function can also be used to apply a universal threshold that is based on characteristics of the weighted bipartite projection \mathbf{P} . For example, it is possible to extract a backbone in which cities are connected if they have more than two standard deviations above the average number of co-located firms.

```
> universal.meansd <- universal(B, upper =
      function(x)mean(x)+2*sd(x),
      bipartite = TRUE)
> mean(universal.meansd$backbone)
[1] 0.03092461
> sort(rowSums(universal.meansd$backbone),
      decreasing = TRUE)[1:5]
LONDON NEW YORK HONG KONG PARIS SINGAPORE
      64      61      51      49      42
> cor(rowSums(universal.meansd$backbone),
      rowSums(B))
[1] 0.9655334
```

This backbone is also lower density (3% of possible edges are present), but once again it focuses only on large cities, whose centrality is nearly identical to their total number of firms ($r = 0.97$).

These examples illustrate that backbones extracted using a universal threshold and the `universal()` function will tend to focus on cities that contain many firms, which is not particularly illuminating. This occurs because the universal threshold approach to backbone extraction does not take into account variations in the number of firms located in each city. By not controlling for these variations (which are substantial in these data; see Fig. 2A) when deciding whether two cities are connected, it privileges cities that contain many firms. In these data, because there are large variations in the number of firms located in each city that must be controlled for, a universal threshold backbone is not appropriate. More generally, universal threshold backbones and the `universal()` function are appropriate only when there is limited variation in the row sums of \mathbf{B} .

Using the HM

In contrast to the universal threshold approach, the hypergeometric model *does* control for variations in the number of firms located in each city (i.e., the row sums of \mathbf{B}). It does so by using a unique threshold for each pair of cities in the network, rather than simply applying the same threshold to every pair. Extracting a backbone using the hypergeometric model involves two steps.

```

> hyper <- hyperg(B)
> hyperbb <- backbone.extract(hyper,
                             alpha = 0.1,
                             signed = FALSE)

```

The `hyperg()` function estimates an HM probability distribution like the one shown in Fig. 3, for each pair of cities in the network, storing the results in a backbone-class object called `hyper`³. The `backbone.extract()` function uses these distributions to identify statistically significant edges, storing the resulting backbone network in a matrix called `hyperbb`. The `signed = FALSE` option indicates that the backbone should only contain edges that are statistically significantly *stronger* than would be expected at random. The significance tests used by the backbone package are two-tailed, so for the backbones which focus only on strong edges (i.e., those in the upper tail of the distribution), the `alpha = 0.1` option ensures that the tests use the conventional $\alpha = 0.05$ (i.e., $0.1/2$) as the threshold for statistical significance.

Before examining the entire HM backbone, consider how the HM works for a single city-pair: Amsterdam and New York. We know that Amsterdam and New York have 26 co-located firms. The HM is designed to test whether this value is statistically significant *controlling for the number of firms in each city*. The green curve in Fig. 3 shows the number of firms that would be co-located in Amsterdam and New York if all firms located in cities randomly, but the number of firms in each city did not change. The 26 co-located firms actually observed in Amsterdam and New York is in the upper tail of this distribution, which indicates that it is much larger than would be expected at random (i.e., it is statistically significant). Therefore, the HM backbone includes a link between Amsterdam and New York.

Examining the backbone extracted using HM highlights how it differs from the weighted projection and universal threshold backbones in several ways.

```

> mean(hyperbb)
[1] 0.09225323
> sort(rowSums(hyperbb),
       decreasing = TRUE)[1:5]
INDIANAPOLIS PORTLAND MELBOURNE
           60          54          52
LYON AUCKLAND
           49          44
> cor(rowSums(hyperbb), rowSums(B))
[1] 0.3039028

```

First, it is less dense than the $T = 0$ backbone, but denser than the 25-threshold or mean-threshold backbones, containing 9.2% of possible edges. That is, this model does reduce the

complexity of the original network, but still preserves many intercity connections. Second, and perhaps more notably, because the HM controls for the number of firms in each city when deciding which intercity connections to keep, it does not simply focus on cities that are large and contain many firms. Indeed, while the most central cities are major financial centers, they are not the obvious ones typically highlighted in world cities research. Moreover, cities' centrality and total firm count are only modestly correlated ($r = 0.30$), indicating that cities' centrality in this network provides information that is unique from what could have been learned from simply counting their number of firms.

Although the HM does control for the number of firms in each city (i.e., the row sums of **B**), it does not control for the number of cities where each firm maintains a presence (i.e., the column sums of **B**). However, there is substantial variation in the number of cities where each firm maintains a presence (see Fig. 2B), and not controlling for this variation can distort decisions about whether a particular city pair's number of co-located firms is significant. For example, if Firm X maintains a presence in *every city*, then observing that it is co-located in Amsterdam and New York is trivial. In contrast, if Firm Y maintains a presence in *only two cities* then observing that it is co-located in Amsterdam and New York is quite noteworthy. Because these data contain not only large variations in the number of firms in each city (see Fig. 2A) but also large variations in the number of cities where each firm maintains a presence (see Fig. 2B), the HM is not appropriate. More generally, a HM backbone and the `hyperg()` function are appropriate only when there is variation in the row sums of **B**, but limited variation in the column sums of **B**.

Using the FDSM

In contrast to the hypergeometric model, the fixed degree sequence model controls for variations in *both* the number of firms located in each city (i.e., the row sums of **B**) and the number of cities where each firm maintains a presence (i.e., the column sums of **B**). Extracting a backbone using the fixed degree sequence model also involves two steps.

```
> set.seed(5) #optional
> fdsm <- fdsm(B, trials = 10000,
              progress = TRUE)
> fdsmbb <- backbone.extract(fdsm,
                             alpha = 0.1,
                             signed = FALSE)
```

The first line is not required, but will ensure that readers' FDSM results, which are generated via simulation, will match what is reported below. The `fdsm()` function estimates a FDSM probability distribution for each edge in the network, but unlike the `hyperg()` function above, allows some options. FDSM distributions cannot be computed exactly, and therefore must be derived via numerical simulation. The `trials = 10 000` option specifies the number of simulations to perform; more simulations will yield more precisely estimated distributions, but will also take longer. The `progress = TRUE` option displays a progress bar while the simulations

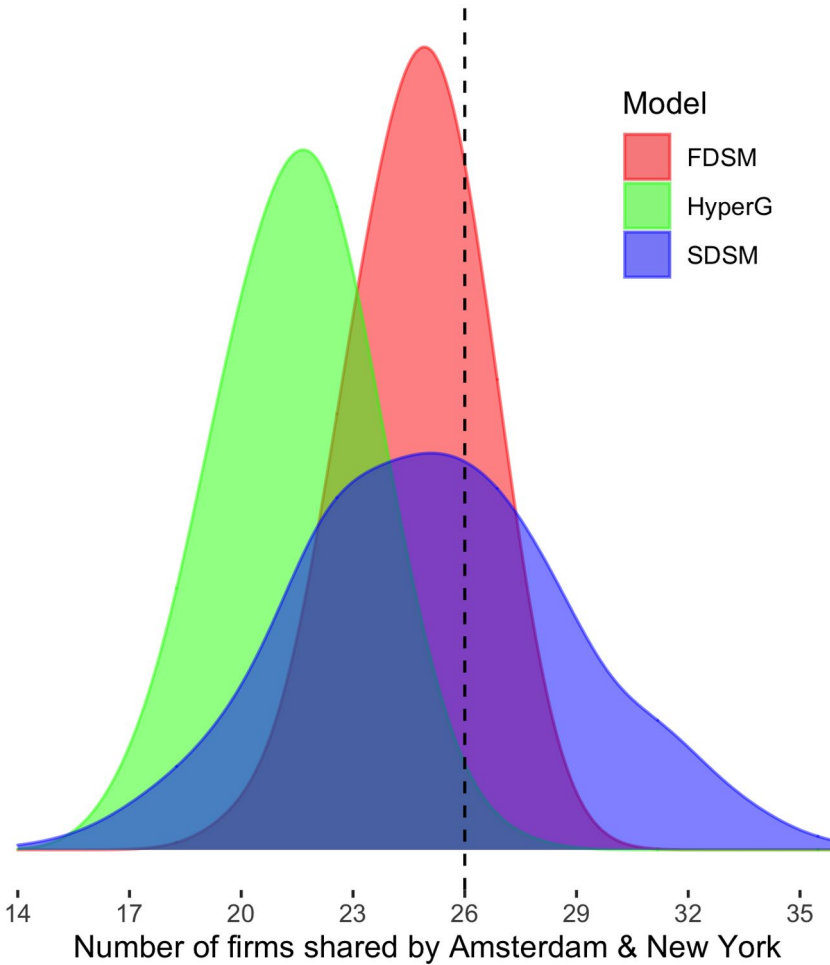


Figure 3. Null weight distributions generated using the `backbone` package on from the GaWC Data set 11.

run. The `backbone.extract()` function works similarly: it takes the resulting `fdsm` object and creates a backbone network called `fdsm_bbb` in which connections between cities are present if they are statistically significantly strong using a two-tailed $\alpha = 0.1$ test.

Again, before examining the entire FDSM backbone, consider how the FDSM determines whether the number of co-located firms is statistically significant for a single city-pair. The red curve in Fig. 3 shows the number of firms that would be co-located in Amsterdam and New York if all firms located in cities randomly, but the number of firms in each city did not change and the number of cities where each firm maintains a presence did not change. Notably the FDSM distribution is both narrower than, and to the right of, the HM distribution. These differences arise because HM and FDSM control for different characteristics of the data. The 26 co-located firms actually observed in Amsterdam and New York is in the middle of the FDSM distribution, which indicates that this value is about what might be expected even under random conditions (i.e., not statistically significant). Therefore, the FDSM backbone does not include a link between Amsterdam and New York.

The backbone extracted using FDSM is noticeably different from all the other networks.

```
> mean(fdsmbb)
[1] 0.02243857
> sort(rowSums(fdsmbb),
       decreasing = TRUE)[1:5]
KANSAS CITY CHARLOTTE INDIANAPOLIS
          24          21          20
RICHMOND GRENOBLE
          20          19
> cor(rowSums(fdsmbb), rowSums(B))
[1] -0.0291214
```

First, it has a very low density, containing only 2.2% of possible edges. Second, the cities with the highest centrality are medium-sized regional centers. Moreover, cities' centrality and total firm count are uncorrelated ($r = -0.03$), indicating that the FDSM backbone is detecting interaction patterns unrelated to a city's number of firms.

The original bipartite firm location data are known to contain substantial variation in both number of firms in each city (see Fig. 2A) but also large variations in the number of cities where each firm maintains a presence (see Fig. 2B). Because the FDSM controls for variation in these two characteristics, it is an appropriate model to use for backbone extraction in this case. Using it yields a world city network backbone that contains only those intercity links that are not simply the product of these characteristics. That is, the FDSM backbone allows world city researchers to look beyond these characteristics to identify pairs of cities with unexpectedly large numbers of firm co-locations, which are potentially indicative of unexpectedly strong economic interaction. More generally, the FDSM and `fdsmbb()` function are appropriate when there is variation in both the row sums of **B** and the column sums of **B**, which is likely to occur in most empirical bipartite data. However, although FDSM may often be the most suitable model for many empirical data, its simulation-based approach can be impractically slow when applied to bipartite data containing many agents and artifacts.

Using the SDSM

The stochastic degree sequence model offers a fast approximation of the FDSM by *approximately* controlling for variations in both the number of firms located in each city (i.e., the row sums of **B**) and the number of cities where each firm maintains a presence (i.e., the column sums of **B**). Extracting a backbone using the stochastic degree sequence model involves two steps.

```

> sdsmbb <- backbone.extract(sdsmbb,
                             alpha = 0.2,
                             signed = FALSE,
                             narrative = TRUE)

```

The `sdsmbb()` function estimates the SDSM probability distribution for each edge in the network. The `backbone.extract()` function is supplied the resulting `sdsmbb` object. In this example, we use $\alpha = 0.2$ rather than $\alpha = 0.1$ for reasons that we illustrate below. Finally, we also include the `narrative = TRUE` option, which can be used when extracting HM and FDSM backbones also. This option generates sample narrative text to be used in a manuscript's methods section:

```

From a bipartite graph containing 196 agents and 100 artifacts, we
obtained the weighted bipartite projection, then extracted its binary
backbone using the backbone package (Domagalski, Neal, & Sagan, 2020).
Edges were retained in the backbone if their weights were
statistically significant (alpha = 0.2) by comparison to a null
Stochastic Degree Sequence Model (Neal, 2014).

```

```

Domagalski, R., Neal, Z. P., and Sagan, B. (2020). backbone: An R
Package for Backbone Extraction of Weighted Graphs. arXiv:1912.12779
[cs.SI]

```

```

Neal, Z. P. (2014). The backbone of bipartite projections: Inferring
relationships from co-authorship, co-sponsorship, co-attendance and
other co-behaviors. Social Networks, 39, 84-97.
https://doi.org/10.1016/j.socnet.2014.06.001

```

Again, before examining the entire SDSM backbone, consider how it determines whether the number of co-located firms is statistically significant for a single city pair. The blue curve in Fig. 3 shows the number of firms that would be co-located in Amsterdam and New York if all firms located in cities randomly, but *on average* the number of firms in each city did not change and *on*

average the number of cities where each firm maintains a presence did not change. The SDSM distribution is wider and flatter than the FDSM distribution, but has nearly the same midpoint. These differences arise because the SDSM distribution is an approximation of the more targeted FDSM distribution. As an approximation with a wider distribution, the SDSM is less statistically powerful, therefore we use a more liberal threshold of statistical significance so that it will more closely mirror the FDSM. The 26 co-located firms actually observed in Amsterdam and New York is in the middle of the SDSM distribution, which indicates that this value is about what might be expected even under random conditions (i.e., not statistically significant). Therefore, the SDSM backbone does not include a link between Amsterdam and New York.

Because the SDSM backbone is an approximation of the FDSM backbone, the two share many features in common.

```
> mean(sdsmbb)
[1] 0.01973136

> sort(rowSums(sdsmbb),
       decreasing = TRUE)[1:5]
KANSAS CITY CHARLOTTE RICHMOND INDIANAPOLIS
           24         21         20          18
BORDEAUX
          17

> cor(rowSums(sdsmbb), rowSums(B))
[1] -0.1062661

> cor(as.vector(fdsmbb), as.vector(sdsmbb))
[1] 0.9301515
```

Like the FDSM, the SDSM backbone is a sparse network, in which medium-sized regional centers are the most central cities, and cities' centrality and total firm count are uncorrelated ($r = -0.11$). Importantly, the pattern of intercity links in the SDSM and FDSM backbones are highly correlated ($r = 0.93$).

These results highlight that the SDSM offers a close approximation of the FDSM. In different ways, both control for the number of firms located in each city (i.e., the row sums of **B**) and the number of cities where each firm maintains a presence (i.e., the column sums of **B**), however, as an approximation the SDSM does so more quickly. For example, extracting a FDSM backbone from these data on a 2.3 GHz processor requires approximately 4 minutes, while extracting an SDSM backbone requires less than one second. Therefore, the factors guiding a choice between SDSM and FDSM backbones are not methodological, but practical (how large is the data?) and theoretical (how strict should the controls be?). When the data are small and/or strict control is desired, FDSM is more appropriate, while when the data are large and/or less strict control is suitable, SDSM is more appropriate.

Discussion

Bipartite projections offer a way to indirectly measure spatial networks using data that is often relatively easy to obtain. For this reason, bipartite projections are now among the most common ways to measure the world city network (Taylor and Derudder 2016), and are frequently used to measure other spatial networks at the global (e.g., Hafner-Burton et al. 2009; Heemskerk et al. 2016; Straka et al. 2017) and local (Browning et al. 2017; Xi et al. 2020) scales, as well as to study the structure of geography as a discipline (Peris et al. 2018). It is often helpful to focus on the backbone of bipartite projections, which preserve only the most important connections between nodes. Multiple backbone models have already been used for spatial analysis (e.g., Neal 2013a; Poorthuis 2018; Van Meeteren et al. 2016); however, a lack of software implementing these models has limited their use. In this article, we have introduced the `backbone` package for R, which is an open-source set of commands for extracting the backbone of bipartite projections, and have demonstrated its use for spatial analysis by applying it to data on firms' locations in cities to understand the world city network. We conclude by offering some recommendations for using `backbone` for spatial analysis, commenting on its limitations, and identifying future directions for similar software development.

When using bipartite projections to measure spatial networks, whether with the `backbone` package or with other tools, the most important requirement is to *have a theory*. The `backbone` package will transform almost any data into something that resembles a network, so it is essential that this transformation be grounded in a theory about why sharing artifacts (e.g., firms, treaties, activity spaces) provides information about interaction and specifically what kind of interaction it provides information about. The theory may be contested or may turn out to be wrong (after all, the purpose of science is to identify wrong theories), but it should at least be explicitly stated.

Even after offering an explicit theory about the suitability of a bipartite projection for network measurement, researchers have many degrees of freedom when using `backbone` to measure spatial networks with bipartite projections. Although methodological research on these topics is ongoing, Fig. 4 offers a preliminary guide to selecting among multiple backbone extraction models. Universal threshold backbones are appropriate when the bipartite data lacks any meaningful variation in the row and column sums, for example, if different cities *did not* contain different numbers of firms and different firms *did not* maintain a presence in different numbers of cities. When a universal threshold backbone is used, it is still necessary to choose the particular threshold value based on theory: how many shared artifacts (e.g., co-located firms) does theory suggest matters when it comes to agents (e.g., cities) interacting? Hypergeometric model (HM) backbones are suitable when there is variation in the row sums, but no meaningful variation in the column sums. Finally, fixed and stochastic degree sequence model (FDSM and SDSM) backbones are suitable when there is meaningful variation in both the row and column sums, with the SDSM offering a practical approximation when the data is large and computational time is a consideration.

There are a number of *future directions for research* on the extraction of backbones from bipartite projections. First, to date there have been limited attempts to formally validate these backbone models, that is, to determine which backbone model (if any) yields the “correct” network. Formal validation is challenging because it requires both spatial bipartite data to which `backbone` can be applied, and an independently measured “true” spatial network against which the resulting backbone can be compared. Preliminary work has attempted to validate bipartite projections as a measurement of world city networks by comparing them to airline

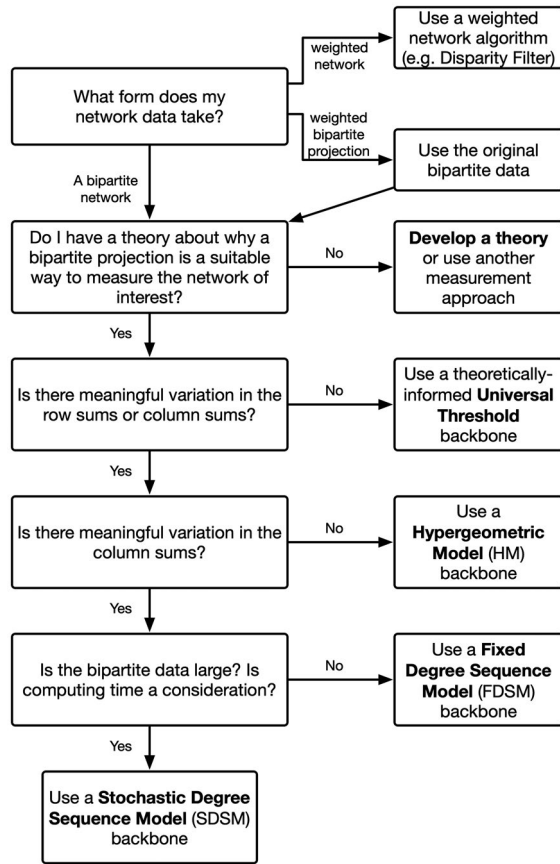


Figure 4. Decision tree for measuring spatial networks using backbones of bipartite projections.

traffic networks (Taylor, Derudder, and Witlox 2007), banking networks (Pažitka et al. 2021), and alternative backbone models (Neal 2014b), but a more general validation of spatial bipartite projections is needed. Second, models for the extraction of backbones from bipartite projections exist only for binary bipartite data, and for projections generated via matrix multiplication (i.e., $\mathbf{P} = \mathbf{B}\mathbf{B}'$). However, binary data are sometimes valued; for example, the original GaWC Data set 11 contains information not only about the presence or absence of firms in cities, but also the size of their presence on a 0–5 scale. Likewise, bipartite projections can be generated using mathematical functions including pairwise conditional probabilities (Hidalgo et al. 2007) and measures of interestingness (Zweig and Kaufmann 2011). The development of new backbone extraction models that can accommodate these cases is necessary (see Neal 2017). The `backbone` package offers a useful methodological tool for pursuing both lines of inquiry.

There are also a number of *future directions for software development* for the extraction of weighted network backbones. First, the `backbone` package currently only allows the extraction of backbones from weighted bipartite projections, but not from other types of weighted networks. However, geographers and spatial analysts often study weighted networks that are not bipartite projections, for example, transportation networks where edge weights convey capacity or volume. Therefore, future versions of `backbone` would benefit from implementing some of the already existing methods for extracting the spatial backbone from such non-projection weighted

networks (e.g., Dai, Derudder, and Liu 2018; Dianati 2016; Serrano et al. 2009). Second, although the `backbone` package implements several different backbone models, the selection of a particular model is left to the user. However, as Fig. 4 illustrates, model selection is driven by several features of the data itself. Therefore, future versions of `backbone` could automate the model selection process, thereby simplifying its use.

The measurement and analysis of spatial networks have become a core part of the spatial analysts' toolbox, alongside such other techniques as gravity models and GIS. Measuring spatial networks using bipartite projections has become increasingly common at both local and global spatial scales, and is a *de facto* approach for measuring the world city network. However, the analysis of bipartite projections requires special care because stronger links in these networks do not necessarily indicate more important connections. The `backbone` package for R is an open-source set of commands that facilitates the analysis of any bipartite projection. Although it is a general-purpose package that can be applied to any bipartite data, in this article, we have demonstrated its particular utility in the context of spatial networks, using the world city network as an illustrative example.

Notes

¹There are other ways to transform a data set into a matrix capturing the similarity among rows, including a Pearson correlation coefficient, pairwise conditional probabilities (e.g., Hidalgo et al. 2007), and measures of interestingness (e.g., Zweig and Kaufmann 2011). However, bipartite projection typically refers to the transformation described by $\mathbf{P} = \mathbf{BB}'$ (Breiger 1974). We restrict our focus to this bipartite projection function, which is the most common approach and the only one implemented in the `backbone` package.

²The transformation of the original valued bipartite data into binary bipartite data means that the analyses reported below may differ from those reported by researchers who do not apply such a transformation (e.g., Taylor 2004). Our goal in this section is to illustrate the functionality of the `backbone` package, and not necessarily to replicate any particular analysis. Neal (2017) described an extension of one of the models in the `backbone` package for ordinal bipartite data.

³These probability distributions are generated by `backbone` in the background as mathematical objects, but are not displayed in graphical form.

References

- Balland, P.-A., and D. Rigby (2017). "The Geography of Complex Knowledge." *Economic Geography* 93(1), 1–23.
- Beaverstock, J. V., R. G. Smith, and P. J. Taylor (1999). "A Roster of World Cities." *Cities* 16(6), 445–58.
- Breiger, R. L. (1974). "The Duality of Persons and Groups." *Social Forces* 53(2), 181–90.
- Browning, C. R., C. A. Calder, B. Soller, A. L. Jackson, and J. Dirlam (2017). "Ecological Networks and Neighborhood Social Organization." *American Journal of Sociology* 122(6), 1939–88.
- Dai, L., B. Derudder, and X. Liu (2018). "Transport Network Backbone Extraction: A Comparison of Techniques." *Journal of Transport Geography* 69, 271–81.
- Derudder, B. (2020). "Bipartite Network Projections of Multi-Locational Corporations: Realising the Potential." *Geographical Analysis*.
- Derudder, B., and P. Taylor (2005). "The Cliquishness of World Cities." *Global Networks* 5(1), 71–91.
- Dianati, N. (2016). "Unwinding the Hairball Graph: Pruning Algorithms for Weighted Complex Networks." *Physical Review E* 93, 012304. <http://doi.org/10.1103/PhysRevE.93.012304>.
- Domagalski R., Neal Z. P., Sagan B. (2021). Backbone: An R package for extracting the backbone of bipartite projections. *PLOS ONE*, 16(1), e0244363. <http://dx.doi.org/10.1371/journal.pone.0244363>
- Franklin, R. S. (2021). "Geographical Analysis at Midlife." *Geographical Analysis*.
- Friedmann, J. (1986). "The World City Hypothesis." *Development and Change* 17(1), 69–83.

- Hafner-Burton, E. M., M. Kahler, and A. H. Montgomery (2009). "Network Analysis for International Relations." *International Organization* 559–592.
- Haggett, P., and R. J. Chorley (1969). *Network Analysis in Geography*. London, UK: Edward Arnold.
- Heemskerk, E. M., M. Fennema, and W. K. Carroll (2016). "The Global Corporate Elite After the Financial Crisis: Evidence from the Transnational Network of Interlocking Directorates." *Global Networks* 16(1), 68–88.
- Hidalgo, C. A., B. Klinger, A.-L. Barabási, and R. Hausmann (2007). "The Product Space Conditions the Development of Nations." *Science* 317(5837), 482–87.
- Jacobs, J. (1961). *The Death and Life of Great American Cities*. New York, NY: Random House.
- Latapy, M., C. Magnien, and N. Del Vecchio (2008). "Basic Notions for the Analysis of Large Two-Mode Networks." *Social Networks* 30(1), 31–48.
- Liu, C., and D. Duan (2020). "Spatial Inequality of Bus Transit Dependence on Urban Streets and its Relationships with Socioeconomic Intensities: A Tale of Two Megacities in China." *Journal of Transport Geography* 86, 102768.
- Liu, X., and B. Derudder (2012). "Two-Mode Networks and the Interlocking World City Network Model: A Reply to Neal." *Geographical Analysis* 44(2), 171–73.
- Martinus, K., T. Sigler, I. Iacopini, and B. Derudder (2021). "The Brokerage Role of Small States and Territories in Global Corporate Networks." *Growth and Change*.
- Neal, Z. P. (2008). "The Duality of World Cities and Firms: Comparing Networks, Hierarchies, and Inequalities in the Global Economy." *Global Networks* 8(1), 94–115.
- Neal, Z. P. (2012). "Structural Determinism in the Interlocking World City Network." *Geographical Analysis* 44(2), 162–170.
- Neal, Z. P. (2013a). "Brute Force and Sorting Processes: Two Perspectives on World City Network Formation." *Urban Studies* 50(6), 1277–91.
- Neal, Z. P. (2013b). *The Connected City: How Networks are Shaping the Modern Metropolis*. Routledge.
- Neal, Z. P. (2014a). "The Backbone of Bipartite Projections: Inferring Relationships from Co-Authorship, Co-Sponsorship, Co-Attendance and Other Co-Behaviors." *Social Networks* 39, 84–97. <http://doi.org/10.1016/j.socnet.2014.06.001>.
- Neal, Z. P. (2014b). "Validity in World City Network Measurements." *Tijdschrift voor economische en sociale geografie* 105(4), 427–43.
- Neal, Z. P. (2017). "Well Connected Compared to What? Rethinking Frames of Reference in World City Network Research." *Environment and Planning A: Economy and Space* 49(12), 2859–77.
- Neal, Z. P. (2021). "Fallacies in World City Network Measurement." *Geographical Analysis*.
- Newman, M. E. (2004). "Analysis of Weighted Networks." *Physical Review E* 70(5), 056131.
- Nordlund, C. (2004). "A Critical Comment on the Taylor Approach for Measuring World City Interlock Linkages." *Geographical Analysis* 36(3), 290–96.
- Pažitka, V., D. Wójcik, and E. Knight (2021). "Critiquing Construct Validity in World City Network Research: Moving from Office Location Networks to Inter-Organizational Projects in the Modeling of Intercity Business Flows." *Geographical Analysis*.
- Peris, A., E. Meijers, and M. van Ham (2018). "The Evolution of the Systems of Cities Literature Since 1995: Schools of Thought and Their Interaction." *Networks and Spatial Economics* 18(3), 533–54.
- Poorthuis, A. (2018). "How to Draw a Neighborhood? The Potential of Big Data, Regionalization, and Community Detection for Understanding the Heterogeneous Nature of Urban Neighborhoods." *Geographical Analysis* 50(2), 182–203.
- R Core Team. (2018). *R: A Language and Environment for Statistical Computing [Computer Software Manual]*. Vienna, Austria. Retrieved from <https://www.R-project.org>
- Robinson, J. (2002). "Global and World Cities: A View from Off the Map." *International Journal of Urban and Regional Research* 26(3), 531–54.
- Sassen, S. (1991). *The Global City: New York, London, and Tokyo*. Princeton University Press.
- Serrano, M. Á., M. Boguná, and A. Vespignani (2009). "Extracting the Multiscale Backbone of Complex Weighted Networks." *Proceedings of the National Academy of Sciences* 106(16), 6483–88.
- Smith, D. A., and M. Timberlake (1995). "Conceptualising and Mapping the Structure of the World System's City System." *Urban Studies* 32(2), 287–302.

- Zachary Neal, Rachel Domagalski, and Bruce Sagan **Analysis of Spatial Networks From Bipartite Projections Using the R Backbone Package**
- Straka, M. J., G. Caldarelli, and F. Saracco (2017). "Grand Canonical Validation of the Bipartite International Trade Network." *Physical Review E* 96(2), 022306.
- Taylor, P. J. (2001). "Specification of the World City Network." *Geographical Analysis* 33(2), 181–94.
- Taylor, P. J. (2004). *World City Network: A Global Urban Analysis*. New York, NY: Routledge.
- Taylor, P. J., G. Catalano, and D. R. Walker (2002). "Measurement of the World City Network." *Urban Studies* 39(13), 2367–76.
- Taylor, P. J., and B. Derudder (2016). *World City Network: A Global Urban Analysis*, 2nd ed. New York, NY: Routledge.
- Taylor, P. J., B. Derudder, and F. Witlox (2007). "Comparing airline Passenger Destinations with Global Service Connectivities: A Worldwide Empirical Study of 214 Cities." *Urban Geography* 28(3), 232–248.
- Ter Wal, A. L., and R. A. Boschma (2009). "Applying Social Network Analysis in Economic Geography: Framing Some Key Analytic Issues." *The Annals of Regional Science* 43(3), 739–56.
- Van Meeteren, M., Z. P. Neal, and B. Derudder (2016). "Disentangling Agglomeration and Network Externalities: A Conceptual Typology." *Papers in Regional Science* 95(1), 61–80.
- Vasques Filho, D., and D. R. J. O'Neale (2020). "Transitivity and Degree Assortativity Explained: The Bipartite Structure of Social Networks." *Physical Review E* 101, 052305. <http://doi.org/10.1103/PhysRevE.101.052305>
- Watts, D. J. (2008). *Six Degrees: The Science of a Connected Age*. New York, NY: W. W. Norton.
- Xi, W., C. A. Calder, and C. R. Browning (2020). "Beyond Activity Space: Detecting Communities in Ecological Networks." *Annals of the American Association of Geographers* 110, 1787–1806.
- Zweig, K. A., and M. Kaufmann (2011). "A systematic approach to the one-mode projection of bipartite graphs." *Social Network Analysis and Mining* 1(3), 187–218. <http://doi.org/10.1007/s13278-011-0021-0>.