G Model
SON-615; No. of Pages 13
ARTICLE IN PRESS
Social Networks xxx (2009) xxx–xxx
Contents lists available at ScienceDirect
Social Networks
journal homepage: www.elsevier.com/locate/socnet

Computing continuous core/periphery structures for social relations data with MINRES/SVD

John P. Boyd[a,*], William J. Fitzgerald[b], Matthew C. Mahutga[c], David A. Smith[d]

[a] Institute for Mathematical Behavioral Sciences, University of California, Irvine, CA 92697-5100, USA
[b] Orange County Department of Education, Costa Mesa, CA 92628, USA
[c] Department of Sociology, University of California, Riverside, CA 92521, USA
[d] Department of Sociology, University of California, Irvine, CA 92697-5100, USA

ARTICLE INFO

Keywords:
Core/periphery
SVD
MINRES
Least squares
Trade
Citation

ABSTRACT

When diagonal values are missing or excluded, MINRES is a natural continuous model for the core/periphery structure of a symmetric social network matrix. Symmetric models, however, are not so useful when dealing with asymmetric data. Singular value decomposition (SVD) is a natural choice to model asymmetry, but this method also requires the presence of diagonal values. In this paper we offer an alternative, more general, approach to continuous core/periphery structures, the minimum residual singular value decomposition (MINRES/SVD), where each node in the network receives two indices, an "in-coreness" and an "out-coreness." The algorithm for computing these coreness vectors is a least squares computation similar to, but distinct from the SVD, again because of the missing diagonal values. And in contrast to the standard, symmetric MINRES algorithm, we can more accurately model asymmetric matrices. This allows us to distinguish, for example, countries in the world economy that are more in the exporting core than they are in the importing core. We propose two nested PRE (proportional reduction of error) measures of fit: (1) the PRE from the MINRES vector with respect to the data and (2) the PRE of the product of the two MINRES/SVD vectors. Applying the resulting method to citations between journals and to international trade in clothing, we illustrate insights gained from being able to model asymmetrical flow patterns. Finally, two permutation tests are introduced to test independently for the MINRES and MINRES/SVD results.

1. Introduction

Core/periphery concepts and network structures are of great interest across a range of academic disciplines and research areas, from interpersonal networks to global systems, from the local transfer of pathogens to the international trade of commodities. Some examples include epidemiology (Jolly et al., 2001; Christley et al., 2005), small groups (Beck et al., 2003; Cummings and Cross, 2003), interpersonal networks (Bourgeois and Friedkin, 2001), linguistics (Dodsworth, 2005), groups in isolated or extreme environments (Johnson et al., 2003), networks of creative artists (Uzzi and Spiro, 2005), PhD exchange networks (Burris, 2004; Fowler et al., 2007), knowledge communities of firms (Giuliani and Bell, 2005), biology (Bosch et al., 2009; Luo et al., 2009), and regional studies and globalization (Alderson and Beckfield, 2004; Gray, 2005; Grbic, 2007; Lee et al., 2007; Mahutga, 2006; Schott, 1986). This plethora of substantive studies has prompted further attempts at conceptual refinement (Everett and Borgatti, 1999, 2005) and work to improve upon the existing computational methods (Boyd et al., 2006; Garcia Muniz and Ramos Carvajal, 2006).

In some substantive applications, "core" and "periphery" are seen as discrete subgroups. This discrete model partitions actors into core and periphery subgroups such that core actors are maximally connected to each other, peripheral actors are minimally connected to each other, and connections between subgroups are unconstrained (Borgatti and Everett, 1999; Boyd et al., 2006). Table 1 illustrates such a discrete core/periphery structure for an imaginary group of 10 actors. The five members of the core appear first and are connected by 1 s, which in practice means a "strong" connection, while the five members of the periphery are connected by 0 s, meaning little or no connections. Each connection between the core and the periphery, however, is indicated by an asterisk (*), meaning that this value is indeterminate and can take on any value. In addition, each diagonal entry has a dash (–), indicating that this tie is undefined. An actual network with a proposed discrete core can be evaluated by correlating its connections against a size-adjusted Table 1, omitting both the asterisks and dashes (Borgatti and Everett, 1999).

This concept can be expanded to several groups in between the core and the periphery, such as "semi-core" and a "semi-periphery."

* Corresponding author. Tel.: +1 949 824 6336; fax: +1 949 824 3733.
E-mail address: jpboyd@uci.edu (J.P. Boyd).

0378-8733/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.socnet.2009.09.003

Table 1
A discrete core/periphery matrix.

	A	B	C	D	E	F	G	H	I	J
A	–	1	1	1	1	*	*	*	*	*
B	1	–	1	1	1	*	*	*	*	*
C	1	1	–	1	1	*	*	*	*	*
D	1	1	1	–	1	*	*	*	*	*
E	1	1	1	1	–	*	*	*	*	*
F	*	*	*	*	*	–	0	0	0	0
G	*	*	*	*	*	0	–	0	0	0
H	*	*	*	*	*	0	0	–	0	0
I	*	*	*	*	*	0	0	0	–	0
J	*	*	*	*	*	0	0	0	0	–

The natural limit, as presented in this paper, is to consider the core and periphery as opposite ends of a continuum. In this view, actors with high values of coreness tend to be highly connected with each other, while those with low values, i.e., peripheral actors, tend not to be connected with each other.

Whether a researcher analyzes empirical data using a discrete or continuous core/periphery approach, a somewhat standard data analysis issue in the analysis of square matrices must first be addressed: how to treat the diagonal entries. The usual practice in social networks is to treat the diagonals as undefined (Wasserman and Faust, 1994). One notable exception is given by Batagelj et al. (2005), who analyze the Baker (1992) citation data among social work journals, where an author citing an article in the same journal is entered in the diagonal.

Diagonal data may possess at least five properties: it can be (1) missing, (2) incompatible, (3) irrelevant, (4) biased, or (5) meaningful. The first case, missing diagonals, is the most common one in social networks, and here one has no choice but to treat them as undefined. The second case, where diagonal entries are present but incompatible with the off-diagonals, is also common. For example, in a study of e-mail communications between individuals, Boyd et al. (2006: 176) decided to treat diagonal values as missing, since e-mails to oneself are clearly a different type of behavior (e.g., e-mail testing or the use of the "ReplyAll" option) than e-mails to others (communication about homework). Similarly, although it is possible to elicit a response to a reflexive network question, such as "liking yourself," the meaning of this response is clearly different from "liking" another person. The third case, irrelevancy, depends on the analysis. E.g., for estimating a distance function on a set of objects, the distance between an object and itself is zero by definition, independent of the data. The fourth case, bias, may be handled either by treating diagonals as missing data, analogous to eliminating outliers, or by constructing a more complex model that estimates and adjusts for this bias. An example of this kind of data is citation networks, where there is an obvious bias toward citing oneself. The fifth and final case, with meaningful diagonals, is best illustrated by symmetric matrices derived from two-mode data. For example, if B is a two-mode data matrix, then the "sums of squares and cross-products" matrix $A = BB^T$ has diagonals (the "squares") that are just as meaningful as are the off-diagonal entries (the "cross-products"). However, these matrices are positive semi-definite, as opposed to the general non-negative square matrices considered in this paper. While not denying the possibility of valid theory and empirical data for diagonal values, for the remainder of this paper we focus on the more common cases, where diagonal values are either absent or excluded.

The most widely used measures of continuous core/periphery structures in networks are due to Borgatti and Everett (1999), implemented as part of the UCINET package (Borgatti et al., 2002, Version 6.153). One of the methods they use to calculate coreness for each actor is MINRES, or minimum residual, a procedure developed by Harman (1967) and Comrey (1962) for approximating correlation matrices while ignoring the "communalities" on the diagonal. This model approximates a square matrix by finding a vector such that the vector times its transpose minimizes the sums of squares of the off-diagonal elements of the residual matrix (see Section 2.1).

In the context of social networks, the MINRES procedure, like eigenvector centrality (Bonacich, 1972, 2007), produces a vector w that takes account of indirect effects. The difference is that eigenvector centrality requires zeros on the diagonals, and then computes the principal eigenvector. By comparison, MINRES finds a vector w that is computed without using the diagonals, optimizing the fit to the off-diagonal entries Comrey (1962: 86), whether or not the data is symmetric.

Most empirical networks are not symmetric: the strength of a tie from i to j may well differ from the tie from j back to i. Yet the MINRES vector produces a symmetric structure matrix, making it impossible to capture empirical asymmetries. One choice for capturing empirical asymmetries, even for square matrices (one-mode data) as an important special case, is to employ the singular value decomposition (SVD), but this requires diagonal values (see Section 2.3). This paper simultaneously addresses the problems of missing diagonals and asymmetry by approximating data matrices with an expression analogous to, but distinct from, the SVD of a matrix. This formulation involves two vectors, representing outgoing and incoming tendencies for each node, but like MINRES, does not use the diagonal. For friendship networks, these vectors might be interpreted as "expansiveness" and "popularity," respectively; for international trade, the interpretation might be exporting and importing tendencies, respectively.

In the following section, we will first briefly review the MINRES continuous core/periphery model. Next, using mock and empirical datasets of differing size, we will compare results obtained from the MINRES algorithm to those from alternative models based upon variations of SVD. After considering the results of these comparisons, and addressing computational issues raised by large sparse matrices, we conclude with an illustrative substantive application of our proposed method to international commodity trade in clothing (Mahutga, 2008) for the year 2000.

2. The core/periphery continuum and MINRES/SVD

2.1. The base model: MINRES

The MINRES method (Harman, 1967; Comrey, 1962) seeks a (column) vector w such that the square n by n data matrix A is approximated by the *structure matrix* ww^T, in the sense that it minimizes the off-diagonal sums of squared differences, or residuals, $SS(A - ww^T) \equiv \sum_i \sum_{j \neq i} (A_{ij} - w_i w_j)^2$. In practice, UCINET normalizes the coreness measures, reporting $w/\sqrt{\sum w_i^2}$ instead of w itself. Each of these n non-negative real numbers w_i measure the *coreness* for actor i. MINRES is most suited to symmetric matrices because the structure matrix ww^T is itself symmetric. Note that if it were not for the exclusion of the diagonal elements of A, the optimal one-dimensional approximation would be given by $v\lambda v^T$, where λ is the first eigenvalue of the symmetrized matrix $(A + A^T)/2$ and v is its eigenvector.

A natural measure of fit in this context is the *proportional reduction of error*, $\mathrm{PRE}(ww^T|\bar{A}) \equiv 1 - \mathrm{SS}(A - ww^T)/SS(A - \bar{A})$, where $\mathrm{SS}(A - \bar{A})$ is the sum of squared deviations of the off-diagonal elements of A from the global mean $\bar{A}$. Obviously, minimizing $\mathrm{SS}(A - ww^T)$ is equivalent to maximizing the PRE. Note that this PRE is the same as the proportional reduction of variance.

If the vector w is chosen badly, the PRE can be zero or even negative. For example, let A be an extreme form of a core/periphery structure, viz., Table 1 but with all the *s replaced by 0 s. Now

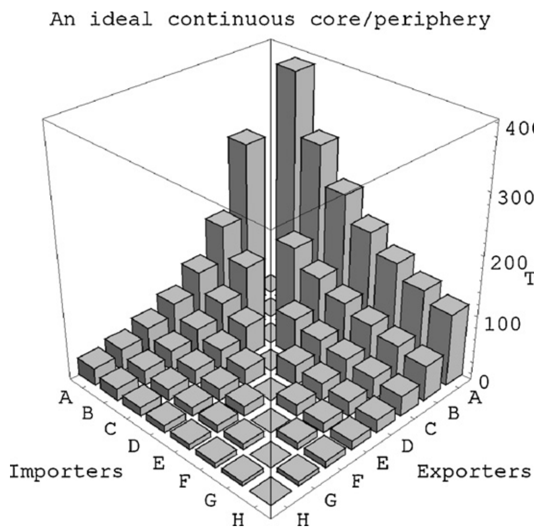

Fig. 1. An ideal, continuous, non-symmetric core/periphery structure.

the matrix has 70 0 s and 20 1 s, so the mean is 2/9, resulting in $SS(A-\bar{A}) = 20(7/9)^2 + 70(2/9)^2 = 1260/81 = 15.(5)$, where the last pair of parentheses represents a repeated sequence of 5 s. The best choice for w would be five 1 s followed by five 0 s, giving a perfect fit with a PRE of 1. On the other hand, the worst choice for w, while holding its inner product w^Tw constant at 5, would be to switch the 0 s and 1 s from the previous w, giving a $SS(A - ww^T)$ of 40, and a PRE of 1–40/15·(5) = −11/7. For any A, however, the PRE for the optimal w always lies between 0 and 1, since w can always be chosen to be the constant vector, the constant being the square root of the mean, so that the entries in the structure matrix ww^T will equal the global mean.

Later we will introduce PREs with respect to the structure matrix generated by the MINRES/SVD model. If X, Y, and Z are three structure matrices with increasing sums of squared deviations from A, denoted by x, y, and z, respectively, then PRE($X|Y$), read as "the PRE of X given Y," is defined as $1 - x/y$. The natural question to ask is how can the largest PRE, X given Z, be expressed in terms of the other two. The correct expression is

$$\mathrm{PRE}(X|Z) = \mathrm{PRE}(X|Y) + \mathrm{PRE}(Y|Z) - \mathrm{PRE}(X|Y)\mathrm{PRE}(Y|Z), \quad (1)$$

which can be easily verified. Obviously, the rule of combining PREs could not be simple addition, since this would lead to PREs greater than one.

A data matrix A satisfies the continuous core/periphery model if its measure of fit, PRE($ww^T|\bar{A}$), is both "large," which depending on the data might mean "larger than 0.5," and "significant," as measured by the permutation tests of Section 3.3. The improved model introduced in this paper, MINRES/SVD, is conditional on the older MINRES being true enough to say that we are in fact looking at a core/periphery model.

2.2. The problems of symmetry and measurement

When empirical data is not symmetric, forcing symmetry on the predicted structure matrix has significant consequences for the theoretical and substantive interpretation of results based upon a symmetric core/periphery modeling approach. To illustrate these consequences, ideally, one would like to model data like that depicted in Fig. 1 and presented in tabular form in Table 2. One can see that coreness drops sharply as one goes from actor A to actor E, following a power law often found in income distributions for either individuals or nations. However, there are also notable asymmetries, most obviously, the larger flow from A to B than from B to A.

Table 2
A non-symmetric, continuous core/periphery matrix.

	A	B	C	D	E	F	G	H
A	–	396	296	234	191	161	138	120
B	281	–	144	114	93	78	67	59
C	161	110	–	65	53	45	39	34
D	102	70	52	–	34	28	24	21
E	69	48	36	28	–	19	17	14
F	50	34	26	20	16	–	12	10
G	37	25	19	15	12	10	–	8
H	28	20	15	12	9	8	7	–

Fig. 1 has several interesting limiting cases obtained by emphasizing symmetry, asymmetry, or the formation of two discrete groups, the latter case being inconsistent with a core/periphery structure. An uninteresting special case would arise if the response were to be independent of either the row or the column variables. For example, it would be possible for each row to be constant.

Any algorithm that forces symmetry in the predicted structure matrix, such as MINRES and other algorithms available in UCINET, is restricted to modeling core/periphery structures corresponding to the symmetric limit of the discrete or continuous structure found in Tables 1 and 2, respectively. Observed data containing significant asymmetry will result in a poor fit to these same core/periphery algorithms.

It certainly would be a mistake to symmetrize all asymmetric observed data before conducting any analysis because the researcher must make an arbitrary decision regarding the method to do this (e.g., take the maximum value of A_{ij} and A_{ji}, or the minimum, the arithmetic mean, the geometric mean, etc.). Furthermore, forcing symmetry may completely remove or fundamentally alter theoretically or substantively important patterns that existed in the original directed data. What is needed, then, is a method that is able to detect continuous core/periphery structures, where they exist, in both symmetric data and asymmetric data.

2.3. SVD solves the one problem, forced symmetry, but now fails to exclude the diagonal

The singular value decomposition is defined for rectangular matrices, for which symmetry is not an issue, but it can also handle square matrices, whether symmetric or asymmetric, as a special case. However, it does require the presence of the diagonal elements of the matrix. More formally, the *singular value decomposition* (SVD) (Schmidt, 1907; Stewart, 1993; Ben-Israel and Greville, 2003) of a real m by n matrix A of rank r is a triple of matrices (U, D, V) such that

$$A = UDV^T, \quad (2)$$

where U is an m by r real matrix with orthonormal columns, D is an r by r diagonal matrix, and V is an n by r real matrix with orthonormal columns. Recall that the *rank* of a matrix is its maximum number of linearly independent rows (or, equivalently, columns). Also, a set of vectors is *orthonormal* if each vector is normal (its sum of squares equals 1) and each pair of distinct vectors is orthogonal (the sum of their cross-products equals 0). In terms of matrix equations, U and V are orthonormal if and only if:

$$U^TU = I_r \quad \text{and} \quad V^TV = I_r, \quad (3)$$

where I_r is the identity matrix (r by r diagonal matrix with 1 s on the diagonal). The columns of U and V are called *singular vectors*. The diagonal elements d_i of D are called *singular values* and are ordered as follows: $d_1 \geq \cdots \geq d_r > 0$. Finally, note that the Us and the Vs, can each be determined from the other, since

$$U = AVD^{-1} \quad \text{and} \quad V = AUD^{-1}. \quad (4)$$

The critical property possessed by the SVD of A is that, for any $k \leq r$, it gives the best least squares rank-k approximation of A, denoted by $A_{(k)}$, by the formula:

$$A_{(k)} = U_{(k)} D_{(k)} V_{(k)}^T, \tag{5}$$

where $U_{(k)}$ and $V_{(k)}$ are the first k columns of U and V, respectively, and $D_{(k)}$ is the diagonal matrix formed from the first k singular values (Schmidt, 1907). This approximation for A is unique if and only if the kth and $(k+1)$st singular values of A are not equal. The squared approximation error is given by

$$\sum_i^m \sum_j^n \left(A_{ij} - \sum_h^k U_{ih} d_h V_{jh} \right)^2 = \sum_{i=k+1}^r d_i^2. \tag{6}$$

Here we are employing the usual convention that the lower summation indices start from 1, unless otherwise noted. The number of real numbers involved in a rank-k approximation is $k(m+n+1)$, counting both sets of singular vectors plus the singular values. However, subtracting the orthonormal constraints in Eq. (3), we are left with just $k(m+n-k)$ free variables to be estimated (degrees of freedom). In the special case of a square n by n matrix of full rank, this formula gives the degrees of freedom as n^2, which agrees with the n^2 free choices for entries in the matrix.

In many applications of SVD, the matrix A is preprocessed before doing the decomposition. In correspondence analysis, for example, each element of A is divided by the square root of the product of its row and column sum (Greenacre, 1984). This is especially useful in two-mode data or when the rows or columns are not measured in the same units. For example, in a person by attribute matrix, the attributes might include incommensurable variables such as height, weight, and religion. However, in the one-mode data sets discussed here, citations between journals and trade between countries of specific commodities, both the rows and columns are from the same relatively homogeneous set of variables (i.e., they are all journals or they are all countries) and the entries are in uniform units, the number of citations and the U.S. dollar value of trade, respectively.

This idea is illustrated in Fig. 1, which is a bar graph of a matrix of a fictional trade network with eight countries, lettered A–H. The height of the (i, j)th bar indicates the amount that country i exports to country j and is generated by the udv^T formula. The singular vectors are $u = \{0.85, 0.42, 0.24, 0.15, 0.10, 0.07, 0.05, 0.04\}$ and $v = \{0.68, 0.46, 0.35, 0.27, 0.22, 0.19, 0.16, 0.14\}$. The singular value d was chosen to be 1000 just so that the trade values would be in the hundreds. The singular value is a scaling factor that gives the optimal least squares fit to the data and can adjust for a change in units, say from dollars to euros. For example, the value of the exports from country A to country B in the model is $u_1 d v_2 = 0.85 \times 1000 \times 0.46 = 400$. The ability of the expression udv^T to capture asymmetry is illustrated here by the reciprocal exports from country B to country A, $u_2 d v_1 = 0.42 \times 1000 \times 0.68 = 286$.

Since udv^T generates the matrix of Fig. 1 exactly, the value of $\mathrm{PRE}(udv^T|\bar{A})$ is 1.0, indicating a completely one-dimensional structure, while the correlation between u and v is 0.984 indicating high, but not complete, symmetry. By way of comparison, the ordinary MINRES has $\mathrm{PRE}(ww^T|\bar{A})$ of 0.849, while $\mathrm{PRE}(udv^T|ww^T)$, because of Eq. (1), also equals 1.0. Of course, if we had chosen u and v to be equal, then the data would be symmetric, and $\mathrm{PRE}(ww^T|\bar{A})$ would be 1.0, while $\mathrm{PRE}(udv^T|ww^T)$ would be 0.

For the reasons specified in Section 1, the diagonal entries, which are down the center of Fig. 1, are missing values. Unfortunately, these missing values mean that SVD cannot be directly applied to compute the first singular value d or the vectors u or v from the off-diagonal values.

2.4. MINRES/SVD solves both problems

To simultaneously address both the asymmetry and missing diagonal problems, we developed a technique called minimum residual singular value decomposition (MINRES/SVD), which has nice properties of both methods. In fact, our MINRES/SVD is the minimal generalization that captures the virtues of both MINRES and SVD. The purpose of MINRES/SVD is to find the best approximation to the n by n matrix A, but where the diagonal elements are excluded (as in MINRES), and where two vectors are allowed (as in an SVD of rank one). That is, we want to minimize the sum of the non-diagonal squared residuals,

$$f = \sum_i \sum_{j \neq i} (A_{ij} - u_i d v_j)^2. \tag{7}$$

The motivation for MINRES/SVD is that it takes desirable features from both SVD (separate u and v vectors to handle asymmetry) and from MINRES (not letting the undefined diagonal elements affect the results). By analogy to the symmetric case, where the w-vector values were interpreted as "coreness," in the remainder of this paper the u- and v-vectors will be referred to as "out-coreness" and "in-coreness," respectively.

2.5. Algorithms for MINRES and MINRES/SVD

One approach to finding the optimal MINRES/SVD solution is to minimize the function f in Eq. (7) with one of the many high level programming languages, such as MATLAB or *Mathematica* (Wolfram, 2003). We did just that using *Mathematica*'s NMinimize procedure, which successfully finds the global optimum in all cases we considered. However, this procedure is relatively slow, making it unsuited for the permutation tests considered later in this paper. A better approach, which follows the usual methods in calculating eigenvalues, is to differentiate Eq. (7), set the result equal to zero, and solve with another built in *Mathematica* function, FindRoot. The disadvantage of FindRoot is that in certain extreme cases it can get stuck in local optima, giving negative PREs. For example, if A is a 10 by 10 matrix consisting of four 5 by 5 blocks: the two off-diagonal blocks are all zeros, the upper left block is all 1 s (except for 0 s on the diagonal), while the lower right block is filled with the constant 0.603 (except for 0 s on the diagonal), then FindRoot gives a PRE of −0.261695, while NMinimze finds the right answer, a PRE of 0.541236. However, if the constant in the lower right block is lowered by 0.001, both procedures give the same answer. Note that this wild behavior occurs in a case totally unsuited to the core/periphery model: two non-communicating cliques. We have not yet encountered this problem on empirical data or with any of its permutations. Furthermore, since these types of structures can be detected via appropriate pre-analysis data screening, we shall proceed with FindRoot.

Before we do this, however, note that we can eliminate both the singular value d and the normality constraints on the singular vectors, u and v. The singular value d is a scale factor whose role is to enlarge or shrink the numbers in the matrix udv^T so that they best approximate the data matrix A. Since u and v are normalized, the numbers u_i and v_j are small (less than one) and their products are even smaller. So if the average number in A is large, say 1000, then d must be even larger to compensate. For example, if A is a constant 100 by 100 matrix such that $A_{ij} = 1000$, then by symmetry $u_i = v_j = 0.1$ (since 0.1 squared is 0.01, which sums to 1). In order to equal 1000, then d must equal 10,000. This gives a perfect fit: i.e., $A = udv^T$.

If we eliminate the normality constraint, then the singular value d can be eliminated by absorbing it into the unconstrained vectors u and v. In the example above we could have $u_i = 1$ and $v_j = 1000$. Of course, the solution is not unique. Nevertheless, without the d our approximating model is now simpler: $A \approx uv^T$. Now we replace Eq. (7) with the new minimization problem of the following equation:

$$f = \sum_i \sum_{j \neq i} (A_{ij} - u_i v_j)^2. \tag{8}$$

Next, we want to get back to the original u, d, and v for ease of interpretation, uniqueness, and comparability with SVD. Here's how: let $d = ||u||\,||v||$, where $||u|| = \sqrt{\sum_i^n u_i^2}$ is the *norm* of u, and then normalize (divide by their norms) u and v. To minimize f in Eq. (8), differentiate with respect to u_i and equate to zero, as in the following equation:

$$\frac{\partial f}{\partial u_i} = \sum_{j \neq i} 2u_i v_j^2 - \sum_{j \neq i} 2A_{ij} v_j = 0. \tag{9}$$

Eq. (9) can be simplified to Eq. (10). Note that the sums are now unrestricted on the left-hand side because the diagonal elements, A_{ii}, are assumed to be zero, and on the right-hand side because it is subtracted out by the term v_i^2.

$$\sum_j A_{ij} v_j = u_i \left(\sum_j v_j^2 - v_i^2 \right). \tag{10}$$

This can best be expressed as the vector Eq. (11), which also includes, in the right-hand equation, the partial derivatives with respect to v.

$$Av = u(v^T v - v^2) \quad \text{and} \quad u^T A = v^T (u^T u - u^2)^T. \tag{11}$$

This has the conventions that a vector squared (or two column vectors that are adjacent) means to square (or multiply) element-wise, and that a scalar, such as $v^T v$ or $u^T u$, plus a vector adds the scalar to each element in the vector. *Mathematica* insists on a period to indicate inner product (matrix–vector or vector–vector), so that the distinction between a vector and its transpose becomes redundant. Finally, the complete *Mathematica* code for solving Eq. (11), and hence finding MINRES/SVD, is the single line in Eq. (12), where u_0 and v_0 are initial guesses for u and v.

$$\text{FindRoot}[\{A.v = u(v.v - v^2),$$
$$u.A = v(u.u - u^2)\}, \ \{\{u, u_0\}, \ \{v, v_0\}\}]. \tag{12}$$

See Appendix A for the computation of good initial guesses for u_0 and v_0.

To most consistently compare the computed MINRES/SVD results with MINRES, we also used *Mathematica* to compute MINRES results. MINRES minimizes the residual Eq. (8), except that $u_i v_j$ is replaced by $w_i w_j$. Differentiating as before, we get Eq. (13), which is the analog of Eq. (11).

$$w^T A + Aw = w(w^T w - w^2). \tag{13}$$

Eq. (13) is solved by the even simpler *Mathematica* line found in the following equation:

$$\text{FindRoot}[A.w + w.A = 2w(w.w - w^2), \ \{w, w_0\}]. \tag{14}$$

3. Application to journal citations and international trade in clothing for 2000

3.1. Baker's 1992 citations among social work journals

In order to illustrate MINRES/SVD, we apply it to Baker's (1992) data, the number of citations from one journal to another among 20 social work journals over a 1-year period between 1985 and 1986. These same data, after first imposing symmetry by choosing the larger of A_{ij} and A_{ji}, were used by Borgatti and Everett (1999:386 Table 8) to demonstrate UCINET's continuous core/periphery algorithms. The raw Baker (1992) data is shown in Table 3, ordered by our MINRES vector w so that the higher core journals come first. The abbreviations in Table 3 are as follows: AMH (Administration in Mental Health), ASW (Administration in Social [henceforth, S] Work [W]), BJSW (British J of S W), CAN (Child[C] Abuse and Neglect), CCQ (C Care Quarterly), CW (C Welfare), CYSR (Children and Youth Services Rev), CSWJ (Clinical S W J), FR (Family Relations), IJSW (Indian J S W), JGSW (J Gerontological S W), JSP (J S Policy), JESW (J Education for S W), PW (Public Welfare), SCW (S Casework), SSR (S Service Rev), SW (S W), SWG (S W with Groups), SWHC (S W in Health Care), SWRA (S W Research and Abstracts),

Note the large asymmetries in the Baker (1992) data, such as that found between the first two journals SW and SCW: 124 versus 58. The diagonal entries show self-reference, although as previously discussed they will be ignored in the analysis that follows.

Both the Baker (1992) journal citation matrix and the world clothing (Mahutga, 2008) trade matrix for 2000, which we will also examine, have a distribution of cell values that is highly skewed to the right. This positive skew is very common with large data sets of non-negative numbers, such as the distribution of income or city size. Skewed data present a problem with any type of analysis that minimizes squared residuals (e.g., regression, SVD, or MINRES),

Table 3
Raw Baker (1992) data, ordered by *w*.

	SW	SCW	SSR	JSWE	CW	SWRA	ASW	SWG	SWHC	CYSR	CSWJ	FR	CAN	PW	JGSW	BJSW	CCQ	IJSW	AMH	JSP
SW	356	58	53	33	52	8	15	15	43	0	0	9	0	19	0	0	0	0	0	0
SCW	124	149	36	21	17	18	8	8	6	0	8	6	6	0	6	0	0	0	0	0
SSR	106	30	105	9	17	25	7	0	0	0	0	0	0	0	0	0	0	0	0	0
JSWE	58	18	16	104	0	16	9	0	7	0	0	0	0	0	0	0	0	0	0	0
CW	58	32	10	11	187	0	0	0	0	6	0	0	7	7	0	0	0	0	0	0
SWRA	44	8	39	24	8	40	0	0	0	0	0	0	0	0	0	0	0	0	0	0
ASW	73	0	21	18	0	7	70	0	0	0	0	0	0	13	0	0	0	0	0	0
SWG	40	9	7	9	0	0	0	41	9	0	0	0	0	0	0	2	0	0	0	0
SWHC	26	20	0	0	0	0	0	0	86	0	0	0	0	0	0	0	0	0	0	0
CYSR	28	8	14	0	70	5	0	0	0	26	0	4	12	6	0	0	5	0	0	0
CSWJ	45	47	20	0	0	0	0	0	0	0	40	0	0	0	0	0	0	0	0	0
FR	9	18	0	0	0	0	0	0	0	0	0	205	0	0	0	0	0	0	0	0
CAN	8	6	0	0	9	0	0	0	0	0	0	0	109	0	0	0	0	0	0	0
PW	0	0	0	0	4	0	0	0	0	0	0	0	7	9	0	0	0	0	0	0
JGSW	18	16	0	0	0	0	0	0	0	0	0	0	0	0	9	0	0	0	0	0
BJSW	19	0	0	13	0	0	0	0	0	0	0	0	0	0	0	95	0	0	0	0
CCQ	0	3	0	0	12	0	0	0	0	0	0	0	0	0	0	0	92	0	0	0
IJSW	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	6	0	0
AMH	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	13	0
JSP	0	0	7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	35

Table 4
Results from MINRES, MINRES/SVD, and the Independence Model for the logged Baker Data.

Journal	w	u	v	Row sum	Col sum
SW	0.571	0.509	0.634	0.482	0.648
SCW	0.420	0.389	0.411	0.473	0.447
SSR	0.349	0.319	0.372	0.278	0.360
JSWE	0.289	0.269	0.289	0.254	0.274
CW	0.272	0.265	0.266	0.276	0.277
SWRA	0.245	0.262	0.215	0.225	0.184
ASW	0.196	0.227	0.151	0.224	0.114
SWG	0.159	0.207	0.095	0.206	0.061
SWHC	0.152	0.138	0.153	0.094	0.123
CYSR	0.139	0.250	0.022	0.331	0.024
CSWJ	0.123	0.208	0.036	0.160	0.027
FR	0.104	0.108	0.095	0.078	0.071
CAN	0.097	0.114	0.082	0.096	0.106
PW	0.078	0.024	0.127	0.055	0.118
JGSW	0.078	0.122	0.031	0.086	0.024
BJSW	0.056	0.107	0.009	0.084	0.013
CCQ	0.032	0.050	0.018	0.059	0.022
IJSW	0.017	0.035	0.000	0.021	0.000
AMH	0.017	0.035	0.000	0.021	0.000
JSP	0.015	0.031	0.000	0.031	0.000
PRE	0.586	0.192		0.029	

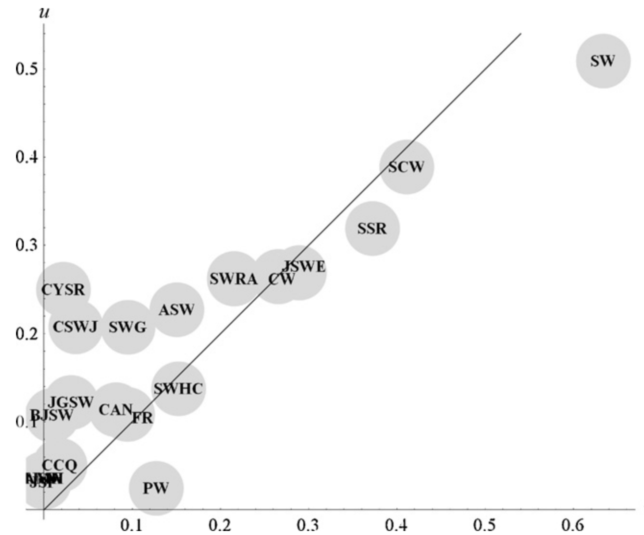

Fig. 2. The core/periphery structure for the logged Baker, 1992 Citations between social work journals.

since the outliers on the right-hand tail have undue influence, or leverage, on the result because squaring magnifies the effect of large deviations.

Skewness in data matrices can be reduced by a suitable transformation. For our purposes the most useful transformation is the $\log_{10}(x+1)$ transform, a special case of the Box–Cox family of transformations (Box and Cox, 1964). The "1" in the transformation is not arbitrary, but is required to produce a value of 0 at 0 with a slope of 1. For simplicity this transform or its application to data will be referred to as "the log transform" or "the logged data." The log transform also has three practical advantages: the base 10 interpretation in terms of approximate orders of magnitude, the implied model of multiplicative stochastic processes, and the reduction of skewness. When the log transform is applied to the Baker (1992), 2000 clothing data (Mahutga, 2008), the skewness is reduced from 4.647 and 38.62 to 1.664 and 1.044, respectively.

The results of MINRES and MINRES/SVD on the logged Baker (1992) data are show in Table 4.

A perusal of the raw data matrix in Table 3 explains why the journal SW (Social Work) receives by far the highest coreness scores; this journal cites and is cited by the most other journals with the highest frequencies. The column of Table 4 labeled w comes from MINRES applied to the logged data. PRE($ww^T|\bar{A}$) for this w is 0.586.[1]

The columns of Table 4 labeled u and v, come from MINRES/SVD applied to the logged data, resulting in PRE($udv^T|ww^T$) of 0.192. Recall that this PRE is over and above the error accounted for by MINRES alone; the PRE of udv^T given the mean is PRE($ww^T|\bar{A}$), which equals 0.665, which can be computed directly or from Eq. (1). The last two columns are merely the normalized row and column sums, also from the logged data. If we fit the data using only the row and column sums, an optimized choice of d results in a disappointing additional PRE, compared to MINRES, of only 0.029. However, the row and column sums can be the initial estimates for u and v, respectively, in the computation of MINRES/SVD. What is gained by computing u and v is a statistically significant increase in the PRE and in- and out-coreness scores that optimize both the direct and the indirect influences in the citation matrix. Just as college football ratings take into account not only wins and losses of a team, but also the wins and losses of its opponents, so too do the singular values. For example, let us examine why CW (Child Welfare) and SWRA (Social Work Research and Abstracts) can end up with almost the same u-scores, 0.265 and 0.262, while their normalized row sums are quite different, 0.276 and 0.225, respectively. The difference is accounted for by SWRA giving all of its citations to journals that are in the top four in u-scores, while CW cites these same journals but in addition uses 20 of its 131 citations on journals ranked in the bottom ten in u-scores, thus dissipating their contribution to its out-coreness. This can be justified more formally by noting that, as discussed in Appendix A, the row sums of the matrix AA^TA are an even better approximation to the u vector than are the row sums for the raw matrix A. The row sums in the matrix AA^TA for CW and SWRA are 679 and 698, which are more nearly equal, helping to explain why their u-scores are so close.

Fig. 2 plots v against u, illustrating the continuous core/periphery structure for the analysis of the log transformed Baker (1992) social work journal citation data (from Table 3). The peripheral journals cluster around the origin, while the core journals are scattered to the upper right, leading to the top journal, SW. Journals that cite more than they are cited, such as CYSR (Children and Youth Services Review), appear above the imaginary 45 degree line, while those that are cited more than they cited, such as PW (Public Welfare), are below this line.

So how does MINRES/SVD perform on much larger real world datasets? In the following section, we analyze one of the more commonly used data types in the context of core/periphery structures: international trade data.

3.2. International clothing trade in 2000

In this section, we analyze clothing trade in the year 2000. Indeed, network studies of international trade almost uniformly attempt to assess the extent to which trade networks exhibit a core/periphery structure. In turn, the main reason for the close association between studies of international trade and network analysis in the field of sociology is the world-system perspective, which sought to explain cross-national inequalities and "underdevelopment" by arguing that dominant "core" countries benefit from more powerful positions *vis-à-vis* "semi-peripheral" and "peripheral" countries (Wallerstein, 1974). Thus, a major task in the empirical world-systems literature has been assessing the extent to which the network of the world economy exhibited a

[1] Because of some inconsistent results we obtained from UCINET's MINRES calculations on asymmetric data (for which it was not designed), the w in Table 4 was computed as described in Section 2.5.

core/periphery interaction pattern; identifying core, peripheral and semi-peripheral countries; and establishing causal links between a country's structural location and level of development (e.g., Mahutga, 2006; Nemeth and Smith, 1985; Smith and White, 1992; Snyder and Kick, 1979; Van Rossem, 1996).

As these studies proliferated, so too did the methodological approaches to the measurement of core/periphery structures. Despite this proliferation and advancement, many previous approaches implicitly, if not explicitly, treat trade relationships as symmetric or discrete. For example, Mahutga (2006) and Smith and White (1992) allow asymmetry in the reduced data matrices, but do not have separate measures for export- and import-coreness because they reduce the multiple trade relationships to a regular equivalence matrix. Snyder and Kick (1979), Kick and Davis (2001), Kick et al. (2000), and Nemeth and Smith (1985) use CONCOR (Breiger et al., 1975), which gives discrete partitions based only on the correlations among the row (export) profiles, while ignoring the column (import) profiles. Kim and Shin (2002) symmetrize their data from the outset by taking the largest of A_{ij} and A_{ji}, and therefore analyze a matrix in which an important aspect of the data is explicitly removed prior to their analysis.

We contend that the asymmetry inherent in trade relationships, like many other kinds of relationships, is a valuable aspect of the data, and that our theoretical understanding of the role that various countries play in the world economy (as defined by country to country trade) is limited by forced symmetry in previous models. For example, a growing body of literature argues that the organization of some manufacturing industries changed over the course of economic globalization. One of the more common industries experiencing this reorganization is the garment industry, in which lead firms located in core countries design, market and retail finished garments but offshore the manufacturing activities to firms in poorer countries (e.g., Gereffi, 1994, 1999).

This shift in the organization of the global garment industry should manifest itself as asymmetry in patterns of international garment trade. For example, we expect to observe a tendency for historically "core" countries to have higher import-coreness than export-coreness scores in the garment industry, that a high ranking in in-coreness is a better predictor of a high ranking in symmetric coreness than is a high ranking in out-coreness, and that high ranking countries have less than expected bilateral trade in garments. In short, the sections that follow show how our proposed method not only verifies the findings of previous studies by fitting a core/periphery structure to international trade data, but also adds to them by differentiating between in-coreness (import-coreness) and out-coreness (export-coreness) in a manner consistent with the organization of the global garment industry.

The data for this analysis come from the United Nations' COMMTRADE database, which collects data on international trade classified by specific types of commodities. We analyze one type of bilateral trade in commodities: clothing (Mahutga, 2008) for the year 2000. Clothing is classified at the two-digit level according to the *Standard International Trade Classification (SITC) Revision 1* scheme (United Nations, 1963) as 84. Our data set is fairly comprehensive, as we were able to gather complete information for 94 countries for the year 2000. See Appendix B for details on how the data was gathered and for a list of included countries and their abbreviations.

As with the Baker (1992) data, we perform the log transformation on the 2000 clothing trade data. Then we apply MINRES and MINRES/SVD, which results in $\mathrm{PRE}(ww^T|\bar{A}) = 0.618$ and $\mathrm{PRE}(udv^T|ww^T) = 0.348$, respectively. The MINRES PRE is slightly higher than the 0.586 for the Baker (1992) data, while the additional PRE for the MINRES/SVD is substantially larger than the 0.192 for the Baker (1992) data. The overall PRE of the MINRES/SVD

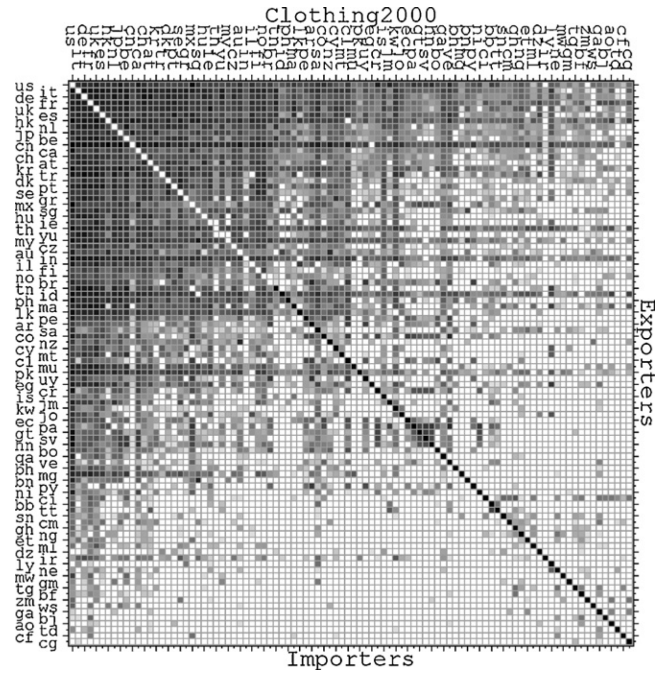

Fig. 3. Density matrix for clothing, 2000.

given the mean, $\mathrm{PRE}(udv^T|\bar{A}) = 0.751$, which again is larger than the 0.665 for the Baker (1992) data.

Fig. 3 shows the matrix density plot for the entire 94 countries in the data set, again ordered by w. The white squares on the diagonal indicate provisional discrete core countries. The overall pattern in this matrix reflects a classic continuous core/periphery structure, with interaction decreasing monotonically as you move down the continuum. However, there are notable asymmetries in the export-coreness and import-coreness scores of many countries, a pattern which is captured and reflected in the differing u and v scores, respectively, of these countries (see Table 5 below for some examples).

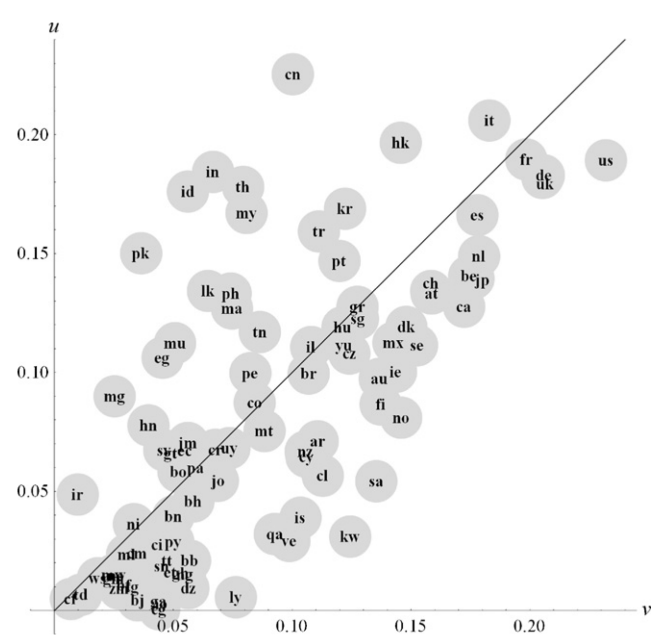

Fig. 4. Core/periphery plot for clothing, 2000.

Table 5
Top 15 Countries in 2000 Clothing Trade on Three Coreness Scores.

w (Symmetric)			*u* (Export-Coreness)			*v* (Import-Coreness)		
1	USA	0.220	1	China[a]	0.225	1	USA	0.232
2	Italy	0.204	2	Italy	0.206	2	Germany	0.206
3	Germany	0.204	3	Hong Kong[a]	0.197	3	UK	0.205
4	France	0.204	4	France	0.189	4	France	0.198
5	UK	0.203	5	USA	0.189	5	Italy	0.183
6	Spain	0.181	6	India[a]	0.184	6	Netherlands	0.179
7	Hong Kong[a]	0.178	7	Germany	0.183	7	Spain	0.178
8	Netherlands	0.172	8	UK	0.182	8	Japan	0.176
9	Japan	0.165	9	Thailand[a]	0.178	9	Belgium	0.174
10	Belgium	0.165	10	Indonesia[a]	0.176	10	Canada	0.172
11	China[a]	0.158	11	South Korea[a]	0.169	11	Austria	0.159
12	Canada	0.156	12	Malaysia[a]	0.167	12	Switzerland	0.158
13	Switzerland	0.153	13	Spain	0.166	13	Sweden	0.153
14	Austria	0.153	14	Turkey[a]	0.159	14	Denmark	0.148
15	South Korea[a]	0.151	15	Pakistan[a]	0.150	15	Norway	0.146

[a] Generally considered to be less developed countries in 2000.

Fig. 4 plots the export-coreness (vertical axis) and import-coreness (horizontal axis) of each country for clothing in 2000. Many countries do not fall on or relatively near the 45 degree line from bottom left to top right, which would be indicative of relatively similar import- and export-coreness. Instead, there are differences in the import-coreness and export-coreness scores for many countries, and some of these differences are relatively large. Consistent with our discussion of the organization of the garment industry above, the US is higher on import-coreness than export-coreness in the 2000 clothing industry. Indeed, this pattern is not limited to the US, but also holds for other widely accepted "core" countries, such as the UK, Germany, France, and Italy.

Table 5 lists the top 15 countries in the clothing industry in 2000 by their symmetric coreness w. The in-coreness v and the out-coreness u reveal notable differences between the countries in each list. Nine of the top 15 export-coreness countries are less developed. On the other hand, the top 15 import-coreness countries are made up entirely of more developed countries. Only three of the nine less developed countries among the top 15 on export-coreness also appear in the symmetric coreness measure in 2000, which is consistent with our expectation that a country's import-coreness score is much more important than its export-coreness score in terms of its relative symmetric coreness in the garment industry. China's placement is perhaps further indicative of the historical organizational changes in the garment industry. It is widely accepted that China was a major recipient of offshored garment manufacturing by 2000 and therefore a leading manufacturing exporter, and this idea is supported by its top ranking in out-coreness and comparatively low ranking on in-coreness. We emphasize again that these substantively important findings would have been obscured by analyses that employed symmetric models.

A final step in examining the fit of a model is to display its residuals. Although as proposed MINRES/SVD is the optimal one-dimensional solution, a review of the residuals for the clothing trade data in 2000 in Fig. 5 suggests the presence of higher dimensional patterns. For example, the top five core countries trade less with each other (as indicated by the blue tint) that is predicted under our model. Yet given the context of this dataset, this is precisely what one might intuitively expect. For example, not surprisingly, the residuals implicate geographical and political biases.

3.3. Two independent permutation tests

In order to assess the statistical significance of the PREs without assuming any particular underlying probability distribution, we turn to permutations tests (Good, 1999). As an illustration of a permutation test, suppose we want to test the association for bivariate data consisting of two vectors x and y, both of length n. Suppose further that we want to use the Pitman correlation statistic: $S = S(x, y) = \sum x_i y_i$, which is really just the inner (or dot) product. A permutation test for the significance of this Pittman association is to choose a large sample (say $N = 1000$) of permutations P of the set $\{1, \ldots, n\}$. Then for each π in P compute the corresponding Pitman correlation with one of the vectors, say y, permuted by π: $S_\pi = S_\pi(x, y) = \sum x_i y_{\pi(i)}$. The p-value is then reported as the fraction of permuted correlations greater than or equal to the original observed result. That is, $p = |\{\pi | S_\pi \geq S\}|/N$.

Note that in applying a permutation test for a given MINRES or MINRES/SVD, we have to permute the entries in the matrix and then re-optimize the PRE measure. This makes our tests more difficult than merely recalculating the sums of squares as in the Pitman correlation considered above.

The permutations should be chosen so as to "preserve" or "control for" as much of the structure as possible. For example, if in the example of the Pitman correlation the sample was divided into males and females, then one should only choose permutations π that preserve sex, meaning that person $\pi(i)$ should be of the same sex as person i.

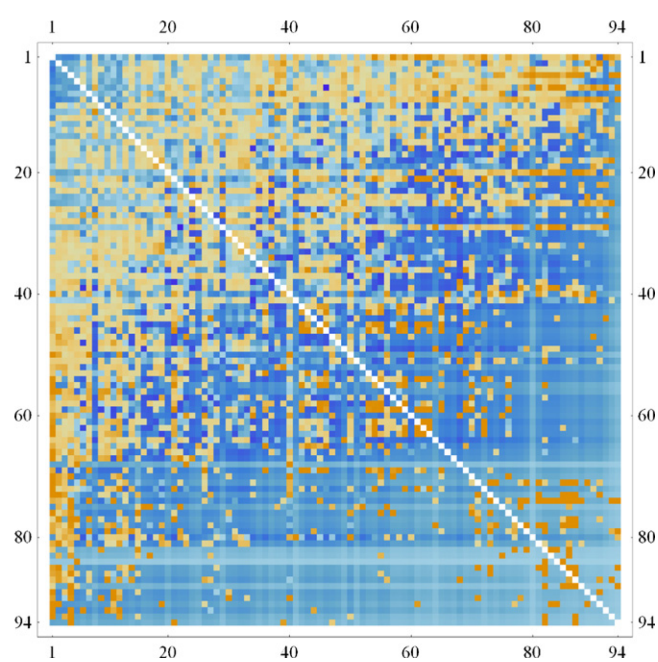

Fig. 5. Residual matrix for clothing, 2000.

Please cite this article in press as: Boyd, J.P., et al., Computing continuous core/periphery structures for social relations data with MINRES/SVD. Soc. Netw. (2009), doi:10.1016/j.socnet.2009.09.003

G Model
SON-615; No. of Pages 13
ARTICLE IN PRESS

Table 6
Three examples of matrix permutations: a matrix isomorphism, MINRES, and SVD.

Starting configuration (upper left):

—	A	B	C	D
a	—	E	F	G
b	e	—	H	I
c	f	h	—	J
d	g	i	j	—

Isomorphism $1 \leftrightarrow 2$ (upper right):

—	**a**	**E**	**F**	**G**
A	—	**B**	**C**	**D**
e	**b**	—	H	I
f	**c**	h	—	J
g	**d**	i	j	—

MINRS: 1,2 $\updownarrow$ 3,5 $\Leftrightarrow$ 2,1 $\updownarrow$ 5,3 (lower left):

—	**I**	B	C	D
i	—	E	F	G
b	e	—	H	**A**
c	f	h	—	J
d	g	**a**		—

SVD: 1,2 $\leftrightarrow$ 2,1 (lower right):

—	**a**	B	C	D
A	—	E	F	G
b	e	—	H	I
c	f	h	—	J
d	g	i	j	—

In a permutation test for our two PRE statistics on an n by n matrix A, the set being permuted is the set I of $n(n-1)$ distinct ordered pairs that specify the off-diagonal elements of A. What is to be preserved in a permutation test for a MINRES vector w and its PRE measure, PRE($ww^T|\bar{A}$)? At least two natural properties come to mind: the transpose pairs and the order determined by w. A permutation on I *preserves transpose pairs* means that the permutation commutes with the transpose. That is, $\pi(i,j)=(p,q) \Leftrightarrow \pi(j,i)=(q,p)$ for all (i,j) in I. Next, this permutation is *order preserving* (with respect to the order determined by w) if $\pi(i,j)=(p,q)$ and $w_i \geq w_j$ implies $w_p \geq w_q$ for all (i,j) in I. We call a permutation that preserves both transpose pairs and order, a MINRES permutation, which is illustrated by the lower-left permutation in Table 6, with the starting configuration in the upper left. The set of all MINRES permutations is closed under composition so that it forms a finite group with $(n(n-1)/2)!$ elements.

Of course some permutations make no difference at all to the structure of the matrix. An isomorphism of the matrix A is determined by a permutation φ on the set of row (and column) indices, $\{1, \ldots, n\}$ by the rule that each (i,j) is mapped onto $(\varphi i, \varphi j)$. Such a permutation is illustrated by the upper right matrix in Table 6. The set of all matrix isomorphisms also form a finite group of order $n!$, with no permutation in common with the MINRES group except for the identity permutation. Both these disjoint groups are subgroups of the set of all $(n(n-1))!$ permutations on I. Despite the fact that isomorphic matrices are for most purposes the same, we will use the convention that when presenting matrices, we choose the isomorphism φ that puts the rows and columns in descending w-order.

In a permutation test for our PRE statistic for the MINES–SVD model, PRE($udv^T|ww^T$), we want a group that is disjoint from both the isomorphism and the MINRES groups. While we still want to preserve unordered pairs, we want to be able to investigate the effect of asymmetry by considering only those permutations that fix these unordered pairs. That is, MINRES/SVD permutations are of the form, $\pi(i,j)=(i,j)$ or (j,i). There are exactly $2^{n(n-1)/2}$ such permutations. See the lower right matrix in Table 6 for an example.

Table 7 gives the results of the permutation tests for the Baker (1992) journal citations and for the 2000 clothing trade data. The first row, labeled "Baker92 MINRES," gives the permutation test results for PRE($ww^T|\bar{A}$) of the w vector compared to the raw matrix minus the global mean. These permutations are of the MINRES type, which permute reciprocal pairs while keeping the orientation with respect to the main diagonal (e.g., lower-left in Table 6). The first column gives the observed (unpermuted) PRE of 0.586, meaning that about 59% of the variance is accounted for by the symmetric MINRES model, ww^T. The next column has the fraction, 0/1000, of times out of 1000 permutation that the PRE of the permuted matrices is greater than or equal to PRE_0. The empirical distribution of permuted PREs is then fitted to a Pearson type III distribution, giving a p-value of 9.8×10^{-15}. The final two columns give the z-score and coefficient of skewness: 13.89 and 0.58, respectively. The extreme values for both the p-value and the z-score should not be taken too literally, but they do aptly describe the low values and the low spread of the PREs of the permuted matrices compared to the initial PRE_0. While the skewness is relatively low, its effect is taken account of in the p-value derived from the Pearson type III distribution.

The second row in Table 7, labeled "Baker92 MINRES/SVD," gives the permutation test results for the PRE of the udv^T model compared to the PRE of the raw matrix minus the ww^T matrix, PRE($udv^T|ww^T$). That is, PRE of 0.192 gives the additional variance accounted for by MINRES/SVD over and above the 0.586 accounted for by MINRES alone. These permutations are of the MINRES/SVD type, which permute within reciprocal pairs, randomly changing their orientations with respect to the main diagonal (e.g., lower right in Table 6). Again, we find that 0 of 1000 PREs are greater than or equal to the observed PRE. However, the p-value and the z-score, 2.7×10^{-4} and 4.45, respectively, are less extreme than those for the MINRES PRE($ww^T|\bar{A}$). If we had computed a million PREs, then we would have expected to find 270 PREs that exceeded $PRE_0 = 0.192$.

The last two rows of Table 7 are for the 2000 clothing trade data. Here the MINRES PRE_0 is only slightly higher than for the Baker (1992) data (0.618 vs. 0.586), while its p-value and the z-score are astronomically higher (2.1×10^{-443} and 236.24). However, the larger sample size could explain these differences. Finally, note that in the last row the MINRES/SVD PRE_0 of 0.348 being almost double than for the corresponding PRE_0 for the Baker (1992) data indicates a greater asymmetry for the trade data.

Table 7
Baker and clothing permutation tests with Pearson III p-values.

	PRE_0	$n \geq PRE_0$	p_{III}-Value	z-Score	g_1
Baker92 MINRES	0.586	0/1000	9.8×10^{-15}	13.89	0.58
Baker92 MINRES/SVD	0.192	0/1000	2.7×10^{-4}	4.45	0.53
Cloth00 MINRES	0.618	0/1000	2.1×10^{-443}	236.24	0.43
Cloth00 MINES-SVD	0.348	0/1000	5.7×10^{-228}	128.17	0.44

Please cite this article in press as: Boyd, J.P., et al., Computing continuous core/periphery structures for social relations data with MINRES/SVD. Soc. Netw. (2009), doi:10.1016/j.socnet.2009.09.003

For both sets of data, the large and significant PRE($ww^T|\bar{A}$) suggests that the core/periphery model was appropriate to apply to these data. The large and significant PRE($udv^T|ww^T$) indicates that the MINRES/SVD model accounted for important empirical asymmetries in the data.

4. Conclusion and discussion

4.1. Implications for networks and global commodity trade

The journal citation data (Baker, 1992) demonstrate the usefulness of asymmetries in relational settings. In addition, a detailed analysis of some of the in- and out-coreness scores revealed the importance of second-order effects that are captured by the one-dimensional singular vectors of the MINRES/SVD model, as illustrated by the poor PRE (0.029) for the row and column marginals (Table 4). We suspect that both of these effects, asymmetry and higher-order interactions, are likely to be present in many other empirical relational contexts.

Our results on trade data also point to the theoretical and substantive ground gained by modeling asymmetry, where the out- and in-coreness is interpreted as export- and import-coreness, respectively. These two measures can be further compared with the MINRES-based symmetric coreness, yielding additional information. The discussion of the variation between import-coreness and export-coreness highlights the theoretical importance of modeling asymmetry where present, which carries over to a multitude of theoretical and empirical settings (Smith and Nemeth, 1988).

4.2. Refinements and generalizations

As mentioned in Section 2.3, SVD is so intrinsically multidimensional that the generalization of the continuous core/periphery model to more than one dimension is obvious. We conjecture that, due to the orthogonality of successive singular vectors of the SVD, that the core/periphery structure may well be replicated at many of the higher dimensions. For example, in the second dimension, one might find that the first core is itself split into a core/periphery structure, resulting in a "super-core" and a "semi-core." Similarly, the first periphery might split into a "semi-periphery" and an "extreme-periphery." The third dimension might display more localized core/periphery groupings, such as the Central American countries as a core with respect to relatively more peripheral, poorer nations to the south.

Another extension of the core/periphery concept is suggested by "multiple correspondence analysis" (Clausen and Sten-Erik, 1988; Greenarce and Blasius, 2006), which analyses 3-way or higher tables by means of "supplementary points," "stacking," "joint correspondence analysis," and other techniques. For example, multiple data sets on world trade (e.g., multiple commodities or years) could be represented by a single set of points in k-dimensional space, where each country is a single point. After some manipulations, some of which are optional, each of these datasets could be combined into a single matrix that is then subjected to an SVD.

It should be noted that neither Borgatti and Everett's (1999) symmetric continuous model nor our MINRES/SVD is a true generalization of the discrete core/periphery model in the following sense: there are matrices A that are perfectly fit by the discrete core/peripheral model, but which are not a perfect fit to either continuous model. For example, if all the *s in Table 1 were to be replaced by any constant, say 1 s, then the correlation between the diagonal blocks and the same diagonal blocks of Table 1 would be a perfect 1.0. On the other hand, MINRES produces the optimal vector w, whose first five entries are 1.092, followed by five entries of 0.693, resulting in an imperfect PRE of 0.466. Because this data is symmetric, MINRES/SVD produces the same answer, and therefore can explain no additional variance. That is, PRE($udv^T|ww^T$) = 0. The problem is that although there are only two values for the w_i, "large" and "small," there are three possible values in the structure matrix ww^T, "large" squared, "small" squared, and "large" times "small." The latter, intermediate value is the predicted value of the *s, instead of being indeterminate as in the discrete model.

One way to construct such a true generalization of the discrete core/periphery model would be to modify the least squares goodness of fit function in the style of Levine (2005). It might be reasonable to multiply each squared residual term, $(u_i dv_j - A_{ij})^2$, by a weighting function f of the squared difference, $(u_i - v_j)^2$, that gives less weight to those terms where u_i is very different from v_j. For example, we could choose $f(x) = \exp[-x^2/(2s^2)]$. The interpretation of this effect would be that for s less than 1, the nearly "off-diagonal" entries in the matrix would be almost "ignored" as in the discrete core/periphery model of Table 1. On the other hand, if s were large, we would get almost equal weight for each term, converging to the original MINRES/SVD. Finally, s itself could be estimated from the data.

The most immediate and important impact of this paper would be for MINRES/SVD to be applied to other data. It is perhaps intuitively obvious to us that most social relations, whether between individuals, countries, or units in between, are not inherently or generally symmetric. If John likes Mary just as much as Mary likes John, it is not only boring, but it may also be due to a lack of precision in the measurements. Even if the measurements are intrinsically symmetric, such as the amount of time spent together for individuals, or the corporate interlock (Haunschild and Beckman, 1998) (number of board members in common) for businesses, asymmetry can result in two ways. First, if the interaction matrix were to undergo preprocessing, as in correspondence analysis (dividing each entry by the square root of the product of the row and columns sums), asymmetry can emerge if the marginals differ. Secondly, even a symmetric matrix can have different singular vectors after the first dimension. In both situations, it is important to not exclude the possibility of asymmetry by an unfortunate choice of method.

Given current software and hardware limitations, however, many researchers would find it difficult to apply our MINRES/SVD model to their data. While users of *Mathematica* could use our code on their machines, and larger datasets could be analyzed provided that sufficient memory and computational power were also available, we would like for this approach to be more widely distributable and accessible. As a final section in this paper, we propose a solution to this dilemma by suggesting an approximation that may be appropriate for such large datasets.

4.3. A MINRES/SVD approximation: SVD with imputed diagonals

As noted in Section 2.3, the SVD procedure, though very desirable both theoretically for its proven optimal minimization properties and practically for its computational speed, stability, and availability, cannot directly resolve the missing diagonal problem. If this problem can be solved, then much larger data sets can be analyzed. One approach to this problem is the seemingly *ad hoc* approach of imputing values to the missing entries and then carrying out the SVD on the result. Imputing missing values has a long history (Watson, 1956) and is widely practiced in log-linear analysis with the "quasi-independence" model (Agresti, 2002: 426). Our idea is to impute diagonal values such that they have little or no influence on the final result. Obviously, putting 0 s on the diagonal would be wrong since this value is an outlier and would have a great influence on the final result. After we impute the diagonal values and compute the SVD on the resulting matrix, we will use the results from MINRES/SVD as a benchmark to test the imputation. Note also that imputing diagonal values is consistent with our previous arguments against including empirical proxies for reflexive ties.

Before we proceed to the case of n missing diagonal values, let us begin by considering a single missing value. Greenacre (1984) suggests that a single missing value for the (i, j)th entry in a data matrix A be imputed by the formula obtained by assuming independence. If we let c_i, r_j, and s be the column, row, and total sums, respectively, where these sums do not include the unknown entry A_{ij}, then the equation for independence, adjusting for the absence of A_{ij} is $(r_i+A_{ij})(c_j+A_{ij})=A_{ij}(s+A_{ij})$. The solution for this equation for A_{ij} is shown in the following equation:

$$A_{ij} = \frac{r_i c_j}{s - r_i - c_j}. \quad (15)$$

If all of the diagonal elements were missing, as in most network data, one could use this formula to estimate each of the diagonal elements, A_{kk}. However, this neglects the contribution of the other $n-1$ diagonal elements to the total sum s. So a better approximation would be to estimate the sum of all the matrix elements by adding to s an estimate for the other $n-1$ diagonal elements, the average value of the off-diagonal elements, $s/(n(n-1))$. After cancelling the factor $n-1$, the independence model for estimating the diagonal elements appears in the following equation:

$$(r_k + A_{kk})(c_k + A_{kk}) = A_{kk}(s + s/n). \quad (16)$$

Solving Eq. (16) gives us the following equation:

$$A_{kk} = \frac{r_k c_k}{s + s/n - r_k - c_k}. \quad (17)$$

These imputed diagonals can be computed in both *Mathematica* and the free statistical program **R** by the very same expression, $(r^*c)/(s+s/n-r-c)$, using combined scalar and vector operations. For example, if we take the formula for generating the mock data of Fig. 1, then Eq. (17) reduces the root mean square error for estimating the diagonal by a factor of four over the estimate from Eq. (16). More refinements are possible, of course, but we suggest that Eq. (17) be used for imputing the diagonal values prior to an SVD.

4.4. Comparisons between SVD with imputed diagonals and MINRES/SVD

If we impute the diagonal according to Eq. (17), and then look at the first dimension of an SVD on the data of Fig. 1, then we get a PRE of 0.99972, which is very close to the 1.0 we get from MINRES/SVD. If, however, we use the naïve estimation from Eq. (16), the PRE is 0.99514, which is not quite as good. However, the PRE from setting the diagonals to 0 is 0.87898, indicating a more significant loss of information. For larger n, the choice of what to put on the diagonal has less effect, since the proportion of diagonal entries to total entries is $1/n$. For example, for $n=100$, and a formula similar to Fig. 1 (the exponents divided by 10), the SVD with diagonals estimated by either Eq. (15) or (17) gives a PRE of 1.0 to within machine epsilon. Even with zeros on the diagonal, the PRE is an acceptable 0.99983. We conclude that, for large n, imputing diagonals according to Eq. (17) followed by an SVD will produce fast results with acceptable accuracy.[2]

Another advantage, for large matrices, of imputing diagonal values followed by SVD is that anyone can compute the imputation of the diagonals found in Eq. (17), even within a standard spreadsheet application. Similarly, SVD is implemented in most statistical and linear algebra packages. Finally, SVD has a natural generalization to more than one dimension. The same is true of MINRES/SVD, except that the new singular vectors are not in general orthogonal to the old ones, as they are in SVD.

Acknowledgments

We would like to thank Katherine Faust and Donald Baker for many helpful comments.

Appendix A. Suggestions for initial in- and out-coreness, u_0 and v_0, in the MINRES/SVD computations

In the discussion of Eq. (12), which computes MINRES/SVD using *Mathematica*'s FindRoot function, the choice of the initial vectors, u_0 and v_0, was left for this appendix to specify. The choice of initial values is important in many numerical problems, since a poor choice can lead to a spurious result, such as returning the second eigenvector instead of the first. While this is unlikely in this case, as long as the initial vectors are positive, there is a concern that poor choices might increase the computation time. On the other hand, total computational time can be wasted if one spends too much time on estimating the initial vectors. Our experience is that the steps outlined below are a good compromise between simple computations and a good initial estimate.

Our strategy is to impute diagonals to the n by n data matrix A using Eq. (17), and then to use the first step in the "reciprocal averaging" method (Greenacre, 1984) for computing the SVD. This is justified by Section 4.3, which showed that MINES/SVD can be approximated by the imputation of diagonals followed by a one-dimensional SVD. The reciprocal averaging method is analogous to the power method for computing eigenvectors. It is not as efficient as other methods, but it does suggest a good way of estimating the initial vectors. Here's how reciprocal averaging works: choose initial vectors x_0 and y_0 to be $x_0(i)=y_0(i)=1$, for $i=1,\ldots,n$. Then the iterative formulas:

$$\begin{aligned} \tilde{x}_k &= Ay_{k-1}, \qquad & x_k &= \tilde{x}_k/||\tilde{x}_k||, \\ \tilde{y}_k &= x_{k-1}A, \qquad & y_k &= \tilde{y}_k/||\tilde{y}_k||, \end{aligned} \quad (18)$$

(where $||v||$ denotes the Euclidean norm of a vector v) give a sequence of vectors such that x_k and y_k converge to the first singular vectors u and v, respectively. Note that $\tilde{x}_1$ and $\tilde{y}_1$ are just the row and column sums (denoted by r and c) of the matrix A, including its imputed diagonals.

Therefore, x_2 and y_2 are just the vectors Ac and rA, normalized. If d is the first singular value, then a good approximation for udv^T, the optimal rank-one approximation for A, would be $x_2 d y_2{}^T$. However, we do not specify the singular value, but absorb it into the vectors u and v, which are now *not* normalized. It is better numerically to have the u and the v about the same size (measured by their norms), so our initial vectors, u_0 and v_0, could be $\sqrt{d}x_2$ and $\sqrt{d}y_2$, if only we had d. However, a good approximation for d is

$$d_0 = \frac{r}{||r||} A \frac{c^T}{||c||}. \quad (19)$$

Using these initial estimates on the logged Baker (1992) data, d_0 is 10.8, compared with the final value from FindRoot of 11.1. Similarly, the correlation between u_0 and the final u is 0.9964; and between v_0 and v, 0.9989.

Appendix B. Included countries and their ISO codes

Approximately half the data was purchased as an electronic file from the UN as reported imports by respective countries. However, many poor countries do not report their data every year. To collect data on missing countries, we simply collected reported export

[2] For example, a laptop with 2 GB of RAM operating at 2 GHz can do an SVD on a 2000 by 2000 matrix in 54 s.

data from countries that do report, to those that do not, and limited our inference to countries that appear in this time period. To fill in trade between missing countries, we collected data at the closest available time within 3 years from 2000. We also aggregated imports from the former Yugoslavian and Czechoslovakian republics to represent these. The included countries, ordered by their ISO abbreviations, are listed below:

ao: Angola, ar: Argentina, at: Austria, au: Australia, bb: Barbados, be: Belgium, bf: Burkina_Faso, bh: Bahrain, bj: Benin, bn: Brunei_Darussalam, bo: Bolivia, br: Brazil, ca: Canada, cf: Central_African_Republic, cg: Congo, ch: Switzerland, ci: Cote_Divoire, cl: Chile, cm: Cameroon, cn: China, co: Colombia, cr: Costa_Rica, cy: Cyprus, cz: Czechoslovakia, de: Germany, dk: Denmark, dz: Algeria, ec: Ecuador, eg: Egypt, es: Spain, et: Ethiopia, fi: Finland, fr: France, ga: Gabon, gh: Ghana, gm: Gambia, gr: Greece, gt: Guatemala, hk: Hong_Kong, hn: Honduras, hu: Hungary, id: Indonesia, ie: Ireland, il: Israel, in: India, ir: Iran, is: Iceland, it: Italy, jm: Jamaica, jo: Jordan, jp: Japan, kr: South_Korea, kw: Kuwait, lk: Sri_Lanka, ly: Libya, ma: Morocco, mg: Madagascar, ml: Mali, mt: Malta, mu: Mauritius, mw: Malawi, mx: Mexico, my: Malaysia, ne: Niger, ng: Nigeria, ni: Nicaragua, nl: Netherlands, no: Norway, nz: New_Zealand, pa: Panama, pe: Peru, ph: Philippines, pk: Pakistan, pt: Portugal, py: Paraguay, qa: Qatar, sa: Saudi_Arabia, se: Sweden, sg: Singapore, sn: Senegal, sv: El_Salvador, td: Chad, tg: Togo, th: Thailand, tn: Tunisia, tr: Turkey, tt: Trinidad/Tobago, uk: UK, us: USA, uy: Uruguay, ve: Venezuela, ws: Samoa, yu: Yugoslavia, zm: Zambia.

References

Alderson, A.S., Beckfield, J., 2004. Power and position in the world city system. American Journal of Sociology 109 (4), 811–851.
Agresti, A., 2002. Categorical Data Analysis. Wiley, New York.
Batagelj, V., Doreian, P., Ferligoj, A., 2005. Generalized Blockmodeling. Cambridge University Press, Cambridge.
Baker, D.R., 1992. A structural analysis of the social work journal network: 1985–1986. Journal of Social Service Research 15, 153–168.
Beck, R.J., Fitzgerald, W.J., Pauksztat, B., 2003. Individual behaviors and social structure in the development of communication networks of self-organizing online discussion groups. In: Wasson, B., Ludvigson, S., Hoppe, U. (Eds.), Designing for Change in Networked Learning Environments. Kluwer, Dordrecht, pp. 313–322.
Ben-Israel, A., Greville, T.N.E., 2003. Generalized Inverses: Theory and Applications, 2nd ed. Springer, New York.
Bonacich, P., 1972. Factoring weighting approaches to status scores clique identification. Journal of Mathematical Sociology 2, 113–120.
Bonacich, P., 2007. Some unique properties of eigenvector centrality. Social Networks 29, 555–564.
Borgatti, S.P., Everett, M.G., 1999. Models of core/periphery structures. Social Networks 21, 375–395.
Borgatti, S.P., Everett, M.G., Freeman, L.C., 2002. UCINET for Windows, Version 6. 153: Software for Social Network Analysis. Analytic Technologies, Harvard.
Bosch, J., Martin Gonzales, A., Rodrigo, A., Navarro, D., 2009. Plant-pollinator networks: adding the pollinator's perspective. Ecology Letters 12, 409–419.
Bourgeois, M., Friedkin, N.E., 2001. The distant core: social solidarity, social distance and interpersonal ties in core-periphery structures. Social Networks 24 (4), 245–260.
Box, G., Cox, D., 1964. The analysis of transformations (with discussion). Journal of the Royal Statistical Society B 26, 211–252.
Boyd, J.P., Fitzgerald, W.J., Beck, R.J., 2006. Computing core/periphery structures and permutation tests for social relations data. Social Networks 28 (2), 165–178.
Breiger, R.L., Boorman, S.A., Arabie, P., 1975. An algorithm for clustering relational data with applications to social network analysis and comparison with multidimensional scaling. Journal of Mathematical Psychology 12 (3), 328–383.
Burris, V., 2004. The academic caste system: prestige hierarchies in PhD exchange networks. American Sociological Review 69 (2), 239–264.
Christley, R.M., Pinchbeck, G.L., Bowers, R.G., Clancy, D., French, N.P., Bennett, R., Turner, J., 2005. Infection in social networks: using network analysis to identify high-risk individuals. American Journal of Epidemiology 162 (10), 1024–1031.
Clausen, Sten-Erik, 1988. Applied Correspondence Analysis: An Introduction. Sage University Paper Series on Quantitative Applications in the Social Sciences, Thousand Oaks, CA, 07-121.
Comrey, A.L., 1962. The minimum residual method of factor analysis. Psychological Reports 11, 15–18.
Cummings, J.N., Cross, R., 2003. Structural properties of work groups and their consequences for performance. Social Networks 25 (3), 197–210.
Dodsworth, R., 2005. Attribute networking: a technique for modeling social perceptions. Journal of Sociolinguistics 9 (2), 225–253.
Everett, M.G., Borgatti, S.P., 1999. Peripheries of cohesive subsets. Social Networks 21 (4), 397–407.
Everett, M.G., Borgatti, S.P., 2005. Extending centrality. In: Carrington, P.J., Scott, J., Wasserman, S. (Eds.), Models and Methods in Social Network Analysis. Cambridge University Press, New York, pp. 57–76.
Fowler, J.H., Grofman, B., Masuoka, N., 2007. Social networks in political science: hiring and placement of Ph.D.s, 1960–2002. Political Science and Politics 40 (4), 729–739.
Garcia Muniz, A.S., Ramos Carvajal, C., 2006. Core/periphery structure models: an alternative proposal. Social Networks 28 (4), 442–448.
Gereffi, G., 1999. International trade and industrial upgrading in the apparel commodity chain. Journal of International Economics 48 (1), 37–70.
Gereffi, G., 1994. The organization of buyer-driven global commodity chains: how U.S. retailers shape overseas production networks. In: Gereffi, G., Korczenewics, M. (Eds.), Commodity Chains and Global Capitalism. Praeger Press, Westport, CT, pp. 95–122.
Giuliani, E., Bell, M., 2005. The micro-determinants of meso-level learning and innovation: evidence from a Chilean wine cluster. Research Policy 34 (1), 47–68.
Good, P.I., 1999. Resampling Methods: A Practical Guide to Data Analysis. Birkhäuser, Boston.
Gray, D., 2005. An examination of regional interaction and super-regions in Britain: an error correction model approach. Regional Studies 39 (5), 619–632.
Grbic, D., 2007. The source, structure, and stability of control over Japan's financial sector. Social Science Research 36, 469–490.
Greenacre, M.J., 1984. Theory and Applications of Correspondence Analysis. Academic Press, Boca Raton, FL.
Greenarce, M.J., Blasius, J. (Eds.), 2006. Multiple Correspondence Analysis and Related Methods. Chapman and Hall/CRC, London.
Harman, H., 1967. Modern Factor Analysis, 2nd ed. University of Chicago Press, Chicago.
Haunschild, P., Beckman, C., 1998. When do interlocks matter? Alternate sources of information and interlock influence. Administrative Science Quarterly 43 (4), 815–844.
Johnson, J.C., Boster, J.S., Palinkas, L.A., 2003. Social roles and the evolution of networks in extreme and isolated environments. Journal of Mathematical Sociology 27, 89–121.
Jolly, A.M., Muth, S.Q., Wylie, J.L., Potterat, J.J., 2001. Sexual networks and sexually transmitted infections: a tale of two cities. Journal of Urban Health–Bulletin of the New York Academy of Medicine 78 (3), 433–445.
Kick, E.L., Davis, B.L., 2001. World-system structure and change: an analysis of global networks and economic growth across two time periods. American Behavioral Scientist 44 (10), 1567–1578.
Kick, E.L., Davis, B.L., Lehtinen, M., Burnes, T.J., 2000. World-system position, national political characteristics and economic development outcomes. Journal of Political and Military Sociology 28 (1), 131–155.
Kim, S, Shin, E., 2002. A longitudinal analysis of globalization and regionalization in international trade: a social network approach. Social Forces 81 (2), 445–471.
Lee, S., Monge, P., Bar, F., Matei, S.A., 2007. The emergence of clusters in the global telecommunications network. Journal of Communication 57, 415–434.
Levine, J., 2005. Extended correlation: not necessarily quadratic or quantitative. Sociological Methods & Research 34 (1), 31–75.
Luo, F., Li, B., Wan, X., Scheuermann, R., 2009. Core and periphery structures in protein interaction networks. BMC Bioinformatics 10 (Suppl. 4), S8.
Mahutga, M.C., 2006. The persistence of structural inequality? A network analysis of international trade, 1965–2000. Social Forces 84 (3), 1863–1889.
Mahutga, M.C., 2008. Embedded inequality in global production networks: the causes and consequences of production network formation. PhD Dissertation, Department of Sociology, University of California at Irvine.
Nemeth, R.J., Smith, D.A., 1985. International trade and world system structure: a multiple network analysis. Review (Fernand Braudel Center) 8 (4), 517–560.
Schmidt, E., 1907. Zur Theorie der linearen und nichlinearen Integralgleichungen. I. Entwicklung willküricher Funktionen nach Systemen vorgeschriebener. Math Annals 63, 433–476.
Schott, T., 1986. Models of dyadic and individual components of a social relation: applications to international trade. Journal of Mathematical Sociology 12, 225–249.
Smith, D.A., Nemeth, R.J., 1988. Research note: an empirical analysis of commodity exchange in the international economy: 1965–1980. International Studies Quarterly 32, 227–240.
Smith, D.A., White, D.R., 1992. Structure and dynamics of the global economy: network analysis of international trade 1965–1980. Social Forces 70 (4), 857–893.
Snyder, D., Kick, E.L., 1979. Structural position in the world system and economic growth, 1955–1970: a multiple-network analysis of transnational interactions. American Journal of Sociology 84 (5), 1096–1126.
Stewart, G.W., 1993. On the early history of the singular value decomposition. SIAM Review 35, 551–566.
United Nations, 1963. Commodity Indexes for the Standard International Trade Classification, Revised #1. Statistical Papers M:38(II).
Uzzi, B., Spiro, J., 2005. Collaboration and creativity: the small world problem. American Journal of Sociology 111 (2), 447–504.
Van Rossem, R., 1996. The world system paradigm as general theory of development: a cross-national test. American Sociological Review 61 (3), 508–527.

Wallerstein, I., 1974. The Modern World System I: Capitalist Agriculture and the Origins of the European World-Economy in the Sixteenth Century. Academic Press, New York.
Wasserman, S., Faust, K., 1994. Social Network Analysis: Methods and Applications. Cambridge University Press, Cambridge.
Watson, G.S., 1956. Missing and 'mixed up' frequencies in contingency tables. Biometrics 12, 47–50.
Wolfram, S., 2003. The Mathematica Book, 4th ed. Wolfram Media/Cambridge University Press, Cambridge.

Please cite this article in press as: Boyd, J.P., et al., Computing continuous core/periphery structures for social relations data with MINRES/SVD. Soc. Netw. (2009), doi:10.1016/j.socnet.2009.09.003