Available online at www.sciencedirect.com

SCIENCE DIRECT®

Social Networks 28 (2006) 275–282

SOCIAL NETWORKS

www.elsevier.com/locate/socnet

Book review

P. Doreian, V. Batagelj, A. Ferligoj, Generalized Blockmodeling, Cambridge University Press, New York, 2005 (xv + 384 pp).

Blockmodels hold a very prominent place among approaches to network analysis. Scott (2000, p. 33) refers to them as a central element of a "breakthrough... that firmly established social network analysis as a method of structural analysis." Likewise, Freeman (2004, chapter 8) features them in his discussion of a "renaissance" that catalyzed the development of the present interdisciplinary social network community. A bibliometric analysis of scholarship published in *Social Networks* (Hummon and Carley, 1993) found that two foundational articles on blockmodels (Lorrain and White (1971) and White et al. (1976)) were the most-cited works from outside the journal. It also identified a set of interrelated articles on blockmodels and role analysis as the most significant of six primary lines of social network research then underway.

So the appearance of an extended statement on blockmodeling written by three distinguished contributors to this literature is a signal publication event for the field. *Generalized Blockmodeling* is a major contribution to both social network analysis and mathematical social science. While much of this material previously appeared in article or chapter form, its sustained and extended treatment here amply rewards the reader's attention. This book will necessarily be a touchstone reference for further methodological developments in this area, and will enhance the quality of future applications of blockmodel analysis.

1. Accomplishments

Integrating results produced during more than fifteen years of collaboration, Doreian, Bategelj and Ferligoj (DBF) devote three chapters to mathematical foundations for blockmodeling, and one to a general discussion of methods of cluster analysis. I concentrate here, however, on their innovations regarding the practice of blockmodeling per se:

- selection of blockmodels using a "direct" approach that optimizes a fitting criterion;
- pre-specification of blockmodels through provision of ex ante information about the image of a relation or mapping of nodes to positions;
- generalization of the familiar criteria of structural and regular equivalence, thereby providing new bases for grouping nodes in a network into classes.

DBF treat the basic blockmodeling problem – of mapping a set of units to a smaller set of positions on the basis of commonalities in relational patterns – as an instance of a more general clustering problem. They refer to positions as "clusters", reserving the term "block" for an array

0378-8733/$ – see front matter
doi:10.1016/j.socnet.2006.02.001

of relations between units assigned to clusters, thus avoiding a long-standing ambiguity in the usage of "block" (White et al., 1976, p. 739).

"Conventional" blockmodeling methods approach this clustering problem "indirectly," focusing on dissimilarity measures calculated from relational data. Such methods consider a series of (usually nested) partitions that group units resembling one another according to some standard. DBF critique this widely-applied strategy on the grounds that:

(1) the dissimilarity measures used are not always "compatible" with the equivalence concepts (e.g., structural or regular equivalence) on which analyses rest; to be compatible with a given equivalence, zero dissimilarity must imply that a pair of units is equivalent, and vice versa;
(2) the linkage between equivalence concepts and goodness-of-fit criteria used to assess proposed partitions is often loose;
(3) the process of partitioning units into clusters is not guided by explicitly stated fitting criteria.

DBF begin by observing that a given definition of equivalence implies not only that compatible dissimilarity measures must be zero for units within a cluster, but also that blocks must take particular "ideal" forms. With the usual binary (0/1) network data, structural equivalence implies that all blocks are either filled with 0 s (null) or with 1 s (complete). Other equivalences imply ideal blocks that take different forms. For example, under regular equivalence, all blocks must be either null or "1-covered,"—that is, having at least one 1 in each row and column. This meets the requirement that regularly equivalent units must have (or lack) relationships to (and from) regularly equivalent units.

The "direct" or optimization approach to blockmodeling centers on the correspondence between a set of ideal block types and the empirical blocks obtained when a partition of nodes into clusters is imposed on network data. It does not interpose any dissimilarity measure. For a given equivalence definition, one may specify a compatible criterion or loss function measuring the extent of this correspondence.

DBF propose to select blockmodel partitions by minimizing such a loss function, doing so in practice via a local optimization algorithm. Notably, this often leads to multiple partitions that fit the data equally well.[1] In several examples they present, the value of a loss function optimized via the direct approach is lower – often appreciably so – than that obtained using a partition selected by way of the conventional approach.

DBF next show that their representation of the blockmodeling problem permits analysts to take both exploratory and more confirmatory stances. Conventional blockmodeling is maximally exploratory, requiring only that the equivalence criterion and the number of clusters be designated in advance. Such analyses can be conducted via DBF's direct approach by requiring that all ideal blocks be of types compatible with the specified equivalence. But the criteria for partitions may also be further pre-specified by designating an ideal form for blocks in specific locations within a model. For example, an analysis could require that some or all clusters be cliques, by providing that within-cluster ideal blocks should be complete.[2] Attribute data too may be used for pre-specification, by requiring that units within a cluster share some non-relational property.

[1] An early recognition of the possibility of multiple partitions is found in White et al.'s (1976) paper, p. 745. See also Heil and White's (1976) identification of "floater" units that can be assigned to more than one cluster within a blockmodel.

[2] Pre-specification of the entire array of blocks – a blockmodel hypothesis – was required by the Heil and White (1976) algorithm. That algorithm reported only clusterings that exactly satisfied its fitting criterion, yielding no result when no such clustering exists.

Optimization and pre-specification arguably enhance and extend blockmodeling based on concepts of structural or regular equivalence. DBF's generalization of blockmodeling begins by defining new criteria for treating units as equivalent. These involve new block types beyond complete, 1-covered, and null. For example, in a "row-regular" block, every row unit must have a relationship to at least one column unit (i.e., the rows must be 1-covered), but no corresponding requirement is imposed on the column units. A special case of row-regularity is a "row-functional" block in which each row unit has a relationship to one and only one column unit. In a "row-dominant" block, at least one row unit is linked to all column units. DBF suggest other block types based on density and degree, and invite others to propose still more types.

Combining this variety in block types with pre-specification, DBF introduce a further generalization by permitting the ideal block types defining equivalence to differ across pairs of clusters. For example, one might require that units within a cluster be structurally equivalent in their relationships to one another but relax this to regular equivalence for their relationships to elements in other clusters. Such a generalized definition of equivalence comes at some cost: The criteria for grouping units into clusters can no longer be concisely stated, but are instead encoded in the array of ideal block types. Indeed, the properties that define units as "equivalent" may differ across clusters and – within clusters – across relationships to "alter" clusters. Notwithstanding this additional descriptive complexity, the generalization offers much flexibility in the definition of blockmodels.

DBF's focus on block types rather than dissimilarity measures facilitates development of several applications and extensions of the framework. One involves blockmodels that simultaneously cluster the row and column units in two-mode network data. DBF apply the same analytic scheme to directed one-mode data, obtaining different clusterings of units in their capacities as sources and as destinations of relations. Secondly, they consider blockmodels for signed network data that record positive (1), neutral (0) and negative (-1) relationships. Here, theories of structural balance and clusterability imply the ideal block types: positive relations within clusters and negative ones between clusters. A third application is to systems of ranked clusters. Relaxing the very stringent requirement of within-cluster completeness (i.e., graph theoretic cliques) in such models, DBF introduce the symmetric block as a new block type. Ranking is enforced by a pre-specification that between-cluster relationships contain no cycles, implying that at least one of the two blocks linking each pair of clusters must be null. The latter two applications highlight the subgroup-level configurations in models of balance, clusterability, and ranking. Those structural images imply regularities in triadic microstructures (Davis and Leinhardt, 1972) that statistical analyses of local network structure examine. Rarely, however, do local analyses return to the more macroscopic level, where these applications of the generalized blockmodeling framework focus.

2. Discussion

I raise several issues in reflecting on DBF's book, most of which the authors address in one way or another themselves, especially in their thoughtful concluding chapter. These include interpretation of the new block types, definition of the loss function and its role in model selection, choosing among multiple partitions, the handling of multi-relational data, and constructing blockmodels using data from network samples.

2.1. Interpretation

At the beginning of the book (pp. 2–3), DBF distinguish two distinct rationales for blockmodeling: delineating role structures and discerning structural regularities in networks. They emphasize the latter, more general, of these. Lorrain and White (1971, p. 69) stressed that for the outcome of a blockmodel analysis to be interpreted as a role structure, there must be a clear correspondence between network ties and expectations held by actors in a population. DBF give limited explicit consideration to the actor standpoint, usually taking the position of an observer.

The new block types allow analysts to consider a great diversity of possible blockmodels. Given so many possibilities, the question of how and why to include each type in a model arises, and this in turn begs the question how to understand each of the types in substantive settings. The interpretation to be assigned to new block types involved in the models of balance/clusterability or ranked clusters appears self-evident, since these types derive from relatively well-specified models of group structure resting on structural and social psychological theories.

Interpretation does not seem as straightforward for some of the other new types, but some conjectures can be offered. For example, in a row-regular block, it is obligatory that each row unit send a relationship toward some column unit, but each column unit need not receive one. From a role analysis standpoint, this might indicate some one-sided rather than interlocking expectation: having the relationship under study is a binding role-requirement from one end, but an option from the other. One might, for example, be an ordinary member of a religious group without receiving pastoral counseling, but leaders of such groups might have to counsel at least some members.

Or consider a row-dominant block, in which at least one of the row units is structurally distinguishable from the others by virtue of its linkage to all column units (as DBF acknowledge, p. 222).[3] Such a dominant row unit might be seen as a "representative" of the row cluster to the column cluster. Row dominance, however, imposes no requirements on relationships between the non-dominant row units and the column units. Such blocks might be informative when used together with requirements elsewhere in a model that provide a basis for placing certain non-dominant units with the dominant one(s). A "representative" or "liaison" interpretation of row-dominance might be plausible, for example, if the block linking the row units to one another were to require some degree of cohesion (e.g., completeness or symmetry).

I highlight questions of interpretation to underline points that DBF emphasize (e.g., p. 244). Some of their new block types are weak in that they place few restrictions on the data. Hence, nominally satisfactory blockmodels (in the sense that a loss function takes a low value) may readily be obtained for most data sets when the ideal block types specified are sufficiently weak. Specification of block types must be informed by substantive and/or theoretical considerations. I expect that the appeal of the new types will grow as applications connect them with engaging interpretations.

2.2. Loss functions

The loss function is a pivotal construction within the direct approach to blockmodeling, serving as the standard against which partitions of units into clusters are judged. DBF warn against mechanical reliance on the loss function as a guide to selecting blockmodels. They stress that its value will vary with the ideal block types specified under a model; hence, numerical comparisons

[3] This is also true of the column units in a row-regular block. A row-dominant block can be partitioned into a complete sub-block involving the dominant row unit(s) and all column units, and a non-complete sub-block involving the non-dominant row units and all column units. Likewise, a row-regular block can be partitioned into a regular sub-block involving all row units and some column units, and a null sub-block involving all row units and the remaining column units. Use of the row-dominant or row-regular blocks as ideals implies that these within-cluster distinctions should not prevent the units involved from being judged "equivalent" to one another.

across models involving more and less restrictive ideal block types can be misleading. DBF also allow (p. 208) that "optimal" partitions in terms of the loss function may not always be preferred to others that take into account substantive data. They do contend that such optimal partitions merit consideration as viable alternative interpretations of a data set, though, because such clusterings best satisfy the specified relational standard of equivalence.

Judgments about how the partitioning of units into clusters should be sensitive to different features of a blockmodel become explicit when a loss function is defined. There is room for debate over how DBF measure the deviations of empirical blocks from some ideals. Such details are important because they can affect both the selection of optimal partitions and (in exploratory analyses) the determination of a blockmodel image. Consider this 4-row, 8-column block (labels are arbitrary, 0 s shown as blanks), evaluated against a standard of regular equivalence (which requires that blocks be either 1-covered or null):

	A	B	C	D	E	F	G	H
J								
K			1	1				1
L								
M			1	1	1	1		

Two rows (J, L) are not 1-covered, as are three columns (A, B, G). It would seem that as few as three changes of 0 into 1 s (in, for example, the [J,A], [J,B], and [L,G] cells) would suffice to convert this into a 1-covered block. DBF's deviation measure (p. 224), however, counts each 0 in a null row, and each 0 in a null column, as a departure from 1-coverage. That measure yields 28 inconsistencies between the above block and a regular ideal. Since this block includes only seven 1 s, this means that it is evaluated as being closer to null than to regular. The deviation measure given by DBF seems a rather severe standard for departure from 1-coverage; indeed, since the above block has only 25 0 s, it would be judged as closer to complete than to 1-covered using their deviation measures.[4]

The principal way in which DBF shape loss functions is by weighting some inconsistencies between model and data more heavily than others. In one of their illustrations, for example, they treat between-cluster cycles as a much more important departure from the ranked clusters model than within-cluster asymmetries, thereby making the optimization algorithm especially prone to arrive at acyclic partitions. The "lean fit" criterion used by Heil and White (1976) is, in essence, a weighted loss function: it attends only to the presence of 1 s in blocks pre-specified as null, taking no account of the data-model correspondence in other blocks.

The loss functions specified by DBF do not depend directly on the number of clusters in a model. For many models, the loss function declines monotonically as the number of clusters rises, eventually reaching 0 as that number grows large enough.[5] For some purposes – particularly when the number of clusters itself is at issue – it may be worthwhile to develop loss functions that treat a given level of data-model inconsistency as less serious in simpler (fewer clusters) models. Wheat (2005) suggests a model-based criterion that does this, seeking to balance the "accuracy" of a blockmodel relative to data against the "complexity" required to achieve that level of fit.

[4] DBF do acknowledge that less stringent measures exist (p. 187). The illustration given here is based on the deviation measures presented when they introduce the generalized blockmodeling framework (p. 224). Those deviation measures are used in the current Pajek implementation of generalized blockmodeling (De Nooy et al., 2005).

[5] This does not necessarily hold, however, for balance/clusterability models. There, adding clusters may create more between-cluster positive ties than the number of within-cluster negative ties it eliminates; the loss function for such models always has a unique lowest value. See DBF, p. 305.

Descriptive indices of goodness of fit could be useful adjuncts to loss functions, though they would not constitute alternative fitting criteria. Late in the book (p. 351), DBF mention a proportional reduction in error (PRE) measure of the extent to which a partition obtained via the direct approach improves over one obtained by indirect methods. PRE measures might also be calculated against other baselines to give an indication of how closely a blockmodel meets a pre-specification. One useful baseline might be the maximal inconsistency that could have arisen under a given partition.[6] Other indices might take the number of dyads in a network into account, supporting comparisons of blockmodel fit across networks differing in size.

2.3. Local optimization and multiple partitions

The iterative "relocation" algorithm that DBF invoke in fitting blockmodels searches for a partition that best satisfies a loss function, but at each step it considers only alternative clusterings that lie relatively close to a current provisional partition. There is no assurance that such a search will lead to a partition that has a globally minimal value of the loss function. To protect against this risk, the authors advise that such searches be repeated many times, beginning with different initial partitions, and that the number of repetitions should increase with network size (p. 190).

With many repetitions comes the prospect that two or more partitions will fit the data equally well. By identifying such multiple partitions, the direct approach obligates the analyst to consider each and select among them. In contrast, the conventional approach – in the absence of an explicit fitting criterion – ordinarily examines only one partition for any given number of clusters. DBF counsel that alternative optimal clusterings should be compared: commonalities across them would represent stable "cores" of clusters. Should alternative partitions differ vastly, however, analysts may have to resort to other, perhaps non-relational, data about a group or a network to choose among clusterings assessed – in terms of the loss function – as being equally good.

2.4. Multiple relations

In one respect, DBF's treatment of blockmodels is more restricted than that found in much prior work. Throughout their discussion, and in all of their applications, they examine single-relational networks only. This is ironic in that one of the initial appeals of blockmodels was their capacity to model multi-relational data (White et al., 1976).

It is not that DBF ignore the issue (see pp. 356–358). They observe that there is usually little reason to anticipate that each of several relationships would satisfy a generalized blockmodel involving the same ideal block types or the same pre-specified pattern of block types (i.e., a particular blockmodel image). However, prior blockmodel work using multi-relational data has rarely insisted that multiple relations should have the same blockmodel image. Indeed, such models are most interesting when the images for different relations contrast with one another.

More arguable than the question of whether each of several relations should have the same image is that of whether the same partition of units into clusters should apply to each of several relations. DBF do indicate that their approach could be extended to accommodate multiple

[6] In an analysis based on structural equivalence, for example, blocks must be complete or null, so the "maximal" inconsistency would be half the number of entries in a block (assuming that the two possible types of inconsistency are equally weighted). For the stronger pre-specification that a block should be complete, the maximum possible inconsistency would instead be the number of entries itself. Comparing the total, across all blocks, of such maximal inconsistencies with the realized number of inconsistencies would give one descriptive indication of blockmodel fit.

relationships. This would necessitate development of a suitable loss function. An extension to multi-relational data would place considerable demands on analysts, especially in a confirmatory mode that would call for pre-specification of a distinct image pattern of block types for each relation in a network.

DBF's reluctance to extend generalized blockmodeling to multiple relations raises the question of the conditions under which multi-relational data should be studied conjointly and separately. As usual, selection of an analytic approach cannot be divorced from conceptualization and subject matter understanding. In some substantive settings, it would seem that positions could be defined only by considering all networks at once. In world-system studies, for instance, the definition of core and peripheral positions may rest on the confluence of flows of raw materials and manufactured goods. Similarly, Merton's (1968, p. 423) notion of a role set ("that complement of role relationships which persons have by virtue of occupying a particular social status") too would seem to require that multiple relations be analyzed together: one partition of units, though not necessarily one blockmodel image, should apply to several networks. However, other frameworks allow for the prospect that actors may occupy multiple social statuses, each associated with one or more roles (e.g., Linton, 1936). Under such a conceptualization, statuses/positions might better be identified through the separate study of multiple networks, with the extent to which different statuses intersect or consolidate as an empirical question.

2.5. Blockmodels based on network samples

My discussion of DBF's book has not attended to the many mathematical foundations and results they present, but I will highlight one. In discussing properties of block types, they (p. 214) note that the defining properties for some block types also characterize sub-blocks involving subsets of units within a cluster (e.g., any sub-block of a complete block is itself complete). This result points to conditions under which analysis of a node sample from a network will identify a blockmodel image that also characterizes the whole network. Notably, the block types associated with structural equivalence (complete and null) seem most robust in that respect. This helps us to appreciate why White et al. (1976, p. 749) obtained similar blockmodels using three different 28-node samples drawn from a network of 107 scientists: the analyses used a criterion of structural equivalence. One should not necessarily anticipate such stability across node samples when using some of the weaker block types.[7]

3. Conclusion

Though *Generalized Blockmodeling* is a polished and substantial work, Doreian, Batagelj and Ferligoj present it in some measure as a progress report (p. xv). Their research program is ongoing: their session at the February 2005, International Sunbelt Social Network Conference covered, among other topics, extensions of their approach to incorporate valued network data and multiple networks. Implementations of most of their techniques are freely available within Pajek (De Nooy et al., 2005). Network analysts should appreciate the innovations and extensions outlined here, and explore them in substantive network analyses. The advances reported in *Generalized*

[7] Consider a row-regular block in a full network, in which each row unit is linked to at least one column unit. If a node sample happens to omit the column unit(s) to which a given sampled row unit is tied, then a blockmodel analysis based on the node sample will not correctly place the sampled row unit within a cluster involved in a row-regular block.

Blockmodeling offer foundations for many further developments in blockmodeling, by DBF and others alike, that promise to be very fruitful.

References

Davis, J.A., Leinhardt, S., 1972. The structure of positive interpersonal relations in small groups. In: Berger, J., Zelditch, M., Anderson, B. (Eds.), Sociological Theories in Progress, vol. 2. Houghton-Mifflin, Boston, pp. 218–251.

De Nooy, W., Mrvar, A., Batagelj, V., 2005. Exploratory Social Network Analysis with Pajek. Cambridge University Press, New York.

Freeman, L.C., 2004. The Development of Social Network Analysis: A Study in the Sociology of Science. Empirical Press, Vancouver, BC.

Heil, G.H., White, H.C., 1976. An algorithm for finding simultaneous homomorphic correspondences between graphs and their image graphs. Behavioral Science 21, 26–35.

Hummon, N.P., Carley, K., 1993. Social networks as normal science. Social Networks 15, 71–106.

Linton, R., 1936. The Study of Man: an Introduction. Appleton-Century, New York.

Lorrain, F., White, H.C., 1971. Structural equivalence of individuals in social networks. Journal of Mathematical Sociology 1, 49–80.

Merton, R.K., 1968. Continuities in the theory of reference groups and social structure. In: Merton, R.K. (Ed.), Social Theory and Social Structure, 1968, Enlarged ed. Free Press, New York, pp. 335–440.

Scott, J., 2000. Social Network Analysis: a Handbook, 2nd ed. Sage, Thousand Oaks, CA.

Wheat, C. 2005. Organizational Positions and the Social Structure of Exchange. Ph.D. dissertation, Program in Organizational Behavior, Harvard University.

White, H.C., Boorman, S.A., Breiger, R.L., 1976. Social structure from multiple networks. I. Blockmodels of roles and positions. American Journal of Sociology 81, 730–780.

Peter V. Marsden *
Department of Sociology, Faculty of Arts and Sciences,
Harvard University, 630 William James Hall,
Cambridge, MA 02138, USA

* Fax: +1 617 496 5794.
E-mail address: pvm@wjh.harvard.edu