

An ROC Approach to Estimating Interpersonal Networks

C. Deniz Yenigün *

Güneş Ertan †

Michael Siciliano ‡

Abstract

This paper focuses on estimating interpersonal networks from a random sample of individuals that belong to the network. A design known as Cognitive Social Structures (CSS) asks individuals to report on their own ties as well as their perceptions of the networks that surround them. Based on data collected using CSS design, first section of the paper dissects cognitive errors as errors of commission and omission, and displays the relationship between the two. Second section of the paper uses a Receiver Operating Characteristic (ROC) based approach to estimate networks from a small sample of randomly selected individuals, while controlling for both types of errors. Results from a numerical study are provided to illustrate the performance of the suggested method.

Key Words: Network estimation, cognitive social structures, network sampling

1. Introduction

Most individual level network oriented studies ask respondents to recall their own direct social ties. Cognitive Social Structure (CSS) based studies on the other hand, collect data not only on the actor's own relations, but also his/her perception of social ties among all other actors in the network. In complete CSS designs for networks of size N , each respondent is required to answer questions about $N^2 - N$ possible relations in the network. Formalized by David Krackhardt (1987), CSS scholarship mostly centers around determinants of accuracy of network cognition.

This paper deviates from traditional CSS studies and aims at (1) presenting a tool for CSS researchers to classify and analyze different errors of perception, and (2) providing a new network estimation method that is designed to minimize overall errors while controlling for different error types. The method is based on randomly selected CSS slices (actors' perception matrices) and therefore does not require high response rates to estimate the network structure. The proposed density weighted Receiver Operator Characteristic (ROC) curve based method is shown to produce network estimates with relatively low errors. Moreover, this new method provides a decision tool for researchers for determining the threshold k (the minimum required number of actors claiming the presence of a tie to assign a tie between two actors in the estimated network) that enables them to visually inspect the distribution of errors at every level of k . The recognition of error variation in data collection is likely to benefit network scholars in areas in which costs of making an error of omission or commission may be significantly different.

The rest of this paper is organized as follows: In Section 2 we introduce the five data sets we consider and present summaries of different error types in individuals' perceptions. In Section 3 we propose a network estimation method controlling these error types. Section 4 reports the results of a numerical study for evaluating the performance of the proposed estimation methodology, and Section 5 concludes.

2. Error Types in CSS

To date most CSS based studies conceptualized errors as absent or present, only very few studies looked at the specific type of perception errors, and to our knowledge none compared the relationship between error types. For example, based on an experimental study O'Connor and Gladstone (2015) found that experience of social exclusion leads to false positives, in other words, errors of commission. Flynn et al. (2010) showed that need for closure is also associated with false positives due to respondents' tendency to perceive transitive ties. In another experimental study Dessi et al.

*Istanbul Bilgi University, Department of Industrial Engineering, 34060, Istanbul, Turkey

†Koç University, Department of International Relations, 34450, Istanbul, Turkey

‡University of Illinois at Chicago, Department of Public Administration, 60607, Chicago, IL, USA

(2016) showed that individuals tended to underestimate mean degree of the networks that they were asked to recall during the experiments.

Unlike this line of research, we first conceptualize and assess cognitive errors as Type 1 (errors of commission, or false positives) and Type 2 (errors of omission, or false negatives), and examine the relationship between the two types of cognitive inaccuracy and their association with overall respondent reliability. To be more specific, we define Type 1 error as an instance of error when an individual perceives there to be a tie between two actors when in fact there is no tie, and Type 2 error as an instance when perception says there is no tie when in fact there is a tie.

We use the following five data sets to analyze the errors and test our estimation technique: (1) High Tech Managers: 21 managers in a machinery firm (Krackhardt, 1987); (2) Silicon Systems: 36 semiskilled production and service workers from a small entrepreneurial firm (Krackhardt, 1990); (3) Pacific Distributors: 33 key personnel from the headquarters of a logistics company (Krackhardt and Kilduff, 1990); (4) Government Office: 36 government employees at the federal level (Krackhardt and Kilduff, 1999); and (5) Italian University: 25 researchers across partner research centers at a university (Casciaro, 1998).

All five data sets are collected according to the CSS method, where each actor not only reports his or her self-ties, but also answers questions on all possible dyads in the network. A CSS for a network involving N individuals is usually represented by a three dimensional array $R_{i,j,m}$ ($i, j, m = 1, \dots, N$), where i is the sender, j is the receiver, and m is the perceiver of the relationship (Krackhardt, 1987). As mentioned above, an actor's network perception matrix is referred to as a CSS slice. Based on the work of Krackhardt (1990), for a given CSS array we construct the true network using the locally aggregate structures (LAS) intersection rule, according to which a tie is considered to be present in the true network if and only if both parties in a dyad report the tie. Note that these ties considered in the networks are directional, and thus for the i - j tie to exist, i must claim to send a tie to j , and j must claim to receive a tie from i . After constructing the true network in this fashion, we explore the different types of errors in perceptions.

Data Set	N	Type I Error Rate		Type II Error Rate		Type I & II Correlation
		Mean	St. Dev.	Mean	St. Dev.	
Italian University	24	0.087	0.090	0.539	0.179	-0.68
High Tech Managers	21	0.052	0.049	0.636	0.174	-0.77
Silicon Systems	36	0.136	0.140	0.641	0.204	-0.82
Government Office	36	0.052	0.073	0.723	0.243	-0.94
Pacific Distributors	48	0.028	0.029	0.701	0.176	-0.86

Table 1: Summary statistics for error rates across five data sets.

For all five data sets, Table 1 displays the mean and standard deviation of both error rates, as well as the correlation between them. Error rates are calculated based on the ratio of the frequency of the particular error type for each actor to the possible number of each error type. Type 1 error rates are considerably low in comparison to Type 2 rates. Average Type 1 error rates range between 0.028 and 0.136 (with standard deviations ranging between 0.029 and 0.140) while average Type 2 error rates range between 0.539 and 0.723 (with standard deviations ranging between 0.174 and 0.243). This difference can be explained by the low true network densities, all of which are less than 0.2. In a sparse network even a small number of omission errors lead to a large rate of Type 2 error since overall number of "1's" in the true adjacency matrix is very small. However, committing errors of commission in similar amounts will translate into lower Type 1 error rates since there is a large number of "0's" in the adjacency matrix. Since majority of social networks tend to have low density (Anderson et al. 1999), these patterns are likely to be observed in many other contexts as well.

In sum, scholarly studies on network cognition so far mostly focused on determinants of cognitive accuracy and ignored different types of cognitive errors. A closer inspection of errors across five commonly used CSS data sets reveals that there is a strong negative relationship between errors of omission and commission, and the probability of making these errors may be associated with the density of the true network. Investigation of the determinants of heterogeneity of perception errors may be beneficial for advancing CSS studies. In the next section we build on divergence of errors in cognition of social relations and introduce a network estimation technique that enables researchers to assess the differential costs of Type 1 and Type 2 errors using a ROC curve based approach.

3. Network Estimation with CSS

In this section, given a random sample of CSS slices, our goal is to estimate the whole network by properly aggregating the observed slices. A straightforward approach for this aggregation is to set a threshold, say k , add the matrices representing the CSS slices, and if k perceptions have accumulated for a possible tie classify it as an existing tie, otherwise classify as a non-existing tie. Although this seems to be a simple and intuitive approach, there are two important issues. The first issue is the origin of a reported tie, in other words, whether it is a self report or it is a perception on others ties. The second issue is the proper choice of threshold k . In this section we first give some details from Siciliano et al. (2012) which considers network estimation by aggregating the randomly selected CSS slices, controlling the Type 1 error rate. We then discuss the main contribution of this paper which is a similar estimation procedure, but controls both Type 1 and Type 2 error rates, therefore considered to be more powerful as well as more flexible.

3.1 Network Estimation Controlling Type 1 Error

The first study considering the estimation of a network from randomly sampled individuals' perceptions is Siciliano et al. (2012), where an aggregation method for the observed CSS slices is proposed, so that the Type 1 error rate in the estimated network is kept below a pre-defined tolerable level. This method is shown to have a good estimation power and it forms a basis of our current work, therefore we present the main ideas here. As we have noted above, the first issue regarding the aggregation method is the origin of a reported tie. This is particularly important because as described in Section 2, the true network is derived through *intersection* of all CSS slices, and an estimation procedure is expected to converge to this true network as the sample size increases. Recall that by intersection we mean that a tie between two actors is considered to exist if and only if both actors agree on its existence. Given this, there are three main scenarios that can arise when estimating a certain tie from an observed sample. In the first scenario, both actors of the tie are sampled, therefore existence or nonexistence of this tie will be estimated by the intersection of the self reports from both actors. No perception information will change this. These type of ties will constitute a region in the estimated network matrix referred to as the *knowledge region*. In the second scenario, none of the actors of the tie are sampled, therefore, the estimation of this tie only relies on others perceptions. If k perceptions have accumulated in the sample regarding this tie, it is estimated to exist, otherwise, not to exist. In the third scenario, only one of the actors of the tie are sampled, therefore the intersection method cannot be used for estimation. In this scenario, we treat the data from the sampled actor as a perception, and estimate this tie to exist only if $k - 1$ additional perception ties are reported by others. Then the estimated ties coming from scenarios 2 and 3 compliment the knowledge region in the estimated network matrix, which is referred to as the *perception region*. Now suppose the CSS slices of a random sample of size n are observed from an unknown network. For a given k , this aggregation method can be summarized in the following algorithm. In what follows, we will refer to this method as the *fixed threshold method* (FTM).

FTM Algorithm: CSS Aggregation for a Fixed k

1. Estimate the knowledge region of the network using the intersection of self reports (Scenario 1).
2. Estimate the perception region of the network by adding the perceptions and unverified self reports (Scenarios 2 and 3). If there are k reports on a tie, it is estimated to exist.
3. Combine the knowledge and perception parts of the network to get the final estimate.

The second issue regarding the aggregation of CSS slices for network estimation is the selection of a proper threshold k . Siciliano et al. (2012) set k such that the estimated Type 1 error rate is controlled below a pre-defined tolerable level α . For a given k and an observed sample of CSS slices, one may estimate the Type 1 error rate in the network by

$$\hat{\alpha}_k = \frac{\left\{ \begin{array}{l} \text{Number of Type 1 errors committed over} \\ \text{the knowledge region with threshold } k \end{array} \right\}}{\left\{ \begin{array}{l} \text{Number of possible Type 1 errors} \\ \text{(or the number of zeroes in the knowledge region)} \end{array} \right\}}. \quad (1)$$

Note that this is the same method as we calculate the Type 1 error rate in a complete CSS structure ($n = N$), where the knowledge region is the whole true network observed by intersection method. Since k is inversely proportional to $\hat{\alpha}_k$, the smallest k satisfying $\hat{\alpha}_k < \alpha$ can be used as a proper threshold value for aggregating CSS slices. Combining this approach with FTM Algorithm, the network estimation method of Siciliano et al. (2012) can be summarized in the following algorithm. As in the original paper, we will refer to this method as the *adaptive threshold method* (ATM).

ATM Algorithm: Network Estimation from CSS Slices, Controlling Type 1 Error

1. Set a tolerable Type 1 error rate α . Typical values are 0.05, 0.10, 0.15.
2. Draw a random sample of size n and observe the CSS slices.
3. Find the smallest k such that $\hat{\alpha}_k < \alpha$, and denote it by k^* .
4. Compute the estimated network by aggregating the CSS slices using FTM with threshold k^* .

3.2 Network Estimation Controlling Both Type 1 and Type 2 Errors: The ROC Approach

The estimation methodology described above is somewhat limited in that, it only controls for the Type 1 error rate. As described in Section 2, another important error type is the Type 2 error, and it is of interest to control both errors when estimating a network from a random sample of CSS slices. Therefore, in this section we propose a new estimation methodology which seeks a balance between Type 1 and Type 2 errors. We base our methodology on a commonly used tool in classification, the receiver operating characteristic (ROC) curve. We first give a brief insight on ROC curves, then we give the specifics of our methodology. We begin by noting that for a given k , similar to estimating the Type 1 error rate in the network, it is possible to estimate the Type 2 error rate by

$$\hat{\beta}_k = \frac{\left\{ \begin{array}{l} \text{Number of Type 2 errors committed over} \\ \text{the knowledge region with threshold } k \end{array} \right\}}{\left\{ \begin{array}{l} \text{Number of possible Type 2 errors} \\ \text{(or the number of ones in the knowledge region)} \end{array} \right\}}. \quad (2)$$

3.2.1 ROC Curves

A ROC curve is a graphical plot illustrating the performance of a binary classifier for its varying discrimination threshold parameter. It plots the true positive rate (TPR = *Positives correctly classified / Total positives*) against the false positive rate (FPR = *Negatives incorrectly classified / Total negatives*) for various levels of the threshold parameter. In classical ROC analysis, typically TPR is plotted on the vertical axis and FPR on the horizontal axis, therefore, the threshold value closest to the top left corner of the plot is said to produce the best classification as it is the value that seeks for larger TPR and smaller FPR. For a detailed treatment of ROC analysis, see, for example, Fawcett (2004).

3.2.2 ROC Approach for Network Estimation

Estimating a network with FTM Algorithm is essentially a binary classification problem, where each entry of the network matrix is classified as an existing or non-existing tie, based on the threshold parameter k . Here, a Type 1 error instance corresponds to a false positive, and a Type 2 error instance corresponds to a false negative. Similarly, *Type 1 error rate* corresponds to FPR, and *1 - Type 2 error rate* corresponds to TPR. Then following the classical ROC approach, the threshold value k closest to the top left corner of the ROC plot should be the optimal threshold. In other words, when employed along with FTM, the threshold k minimizing the distance

$$\delta = \sqrt{\hat{\alpha}_k^2 + \hat{\beta}_k^2} \quad (3)$$

gives a network estimate which seeks a balance between Type 1 and Type 2 errors. Note that the threshold value k obtained by minimizing δ in (3) treats both error types as equally important. However, in practice, one error type may be more crucial than the other, and it is of interest to

minimize a weighted distance measure in order to handle these unbalanced cases. In the remaining of this section we first illustrate the ROC based network estimation using the unweighted distance δ on a real data set and point out potential problems. We then discuss how these problems can be addressed by a ROC approach employing a weighted distance δ_w .

Numerical Illustration

Consider the High Tech Managers data set introduced in Section 2, which contains the CSS slices of $N = 21$ managers of a machinery firm. Suppose that a random sample of size $n = 10$ is drawn and the network is to be estimated from the sampled CSS slices. For this illustration, suppose the individuals 2, 4, 5, 8, 9, 10, 11, 14, 18 and 19 are sampled. We consider FTM for aggregating the slices, and k will be determined by the ROC approach described above, which minimizes the distance δ in (3). For a given k , the true positive rate and the false positive rate in the network may be estimated over the exact region of the network estimate by $1 - \hat{\beta}_k$ and $\hat{\alpha}_k$, respectively. For the observed sample, estimated true positive rates and false positive rates for all possible k are given in Table 2, along with the calculated distance δ . Note that δ is minimized for $k = 1$, therefore FTM with $k = 1$ is considered to give a good estimate of the network. However, taking a closer look at the actual numbers of committed Type 1 and Type 2 errors reveal a potential flaw of this approach.

k	TPR ($1 - \hat{\beta}_k$)	FPR ($\hat{\alpha}_k$)	δ	δ_w
0	1.000	1.000	1.000	10.606
1	0.917	0.295	0.307	3.135
2	0.667	0.148	0.365	1.602
3	0.583	0.080	0.424	0.941
4	0.333	0.034	0.668	0.758
5	0.250	0.011	0.750	0.760
6	0.083	0.011	0.917	0.925
7	0.083	0.000	0.917	0.917
8	0.000	0.000	1.000	1.000

Table 2: TPR, FPR, and two distances considered for the numerical illustration.

Figure 1-a illustrates the classical ROC curve produced by varying k from 0 to 8, along with the value of k at each step. We may see that when $k = 1$, the curve is closest to top left corner, i.e., δ is minimized at this point. Figure 1-b indicates the actual number of Type 1 and Type 2 errors committed for each k . We may see that when $k = 1$ we commit 26 Type 1 errors out of 88 possible instances ($\hat{\alpha}_k = 26/88 = 0.295$), and one Type 2 error out of 12 possible instances ($\hat{\beta}_k = 1/12 = 0.083$). The potential drawback here is that the classical ROC approach minimizing δ assigns equal importance to both error types, and it gives the optimal k based on error rates, not the actual counts. However, in low density networks reducing Type 1 error may be more critical compared to Type 2 error rate since the true network contains much more zeroes than ones, creating much more potential Type 1 error instances. This can be easily seen in Figure 1-b, where increasing k from one to two reduces the Type 1 error count by 13 at the expense of only three Type 2 error counts, and similarly for the next increment of k . Therefore, we consider a weighted distance δ_w to be minimized in the ROC analysis,

$$\delta_w = \sqrt{(w\hat{\alpha}_k)^2 + \hat{\beta}_k^2}, \quad (4)$$

where a choice of $w > 1$ gives more emphasis on Type 1 error, a choice of $w < 1$ gives more emphasis on Type 2 error, and the choice of $w = 1$ corresponds to classical ROC distance δ . Based on our observation on the effect of network density on the potential error rates, we propose that $w = 1/\bar{d}$ is a reasonable choice of w for the low density networks, where \bar{d} is the average density of all CSS slices in the random sample. In the illustrative example $\bar{d} = 0.094$, so we set $w = 1/0.094 = 10.606$. As marked in Figure 1-a (as well as in Table 2), when $w = 10.606$ the minimum δ_w is attained for $k = 4$, which produces a Type 1 error rate of $\hat{\alpha}_k = 3/88 = 0.034$ and Type 2 error rate of $\hat{\beta}_k = 8/12 = 0.667$. Then employing FTM with $k = 4$ is considered to give a better estimate of the unknown network based on the sample information. As an evidence of the improvement in estimation, we note that when δ is minimized ($k = 1$) the correlation between the estimated network and the true network is 0.644, however, when δ_w is minimized ($k = 4$) the correlation is 0.749. Here we also note that the correlation measure we use in this study is the similarity index S_{14} suggested by Krackhardt (1990).

Proposed Estimation Method

Based on our observations on the above illustrative example, as well as our investigation on other data sets, we observed that the ROC approach is a powerful one for estimating networks from a random sample of CSS slices since it seeks for a balance between the committed Type 1 and Type 2 errors. We also observed that giving more emphasis on Type 1 error rate improves the estimation in sparse networks. Therefore, we control the level of this emphasis to be inversely proportional to the average density of the observed CSS slices. Our estimation method is summarized in the following algorithm. In what follows, we will refer to this method as the *ROC based threshold method* (RTM).

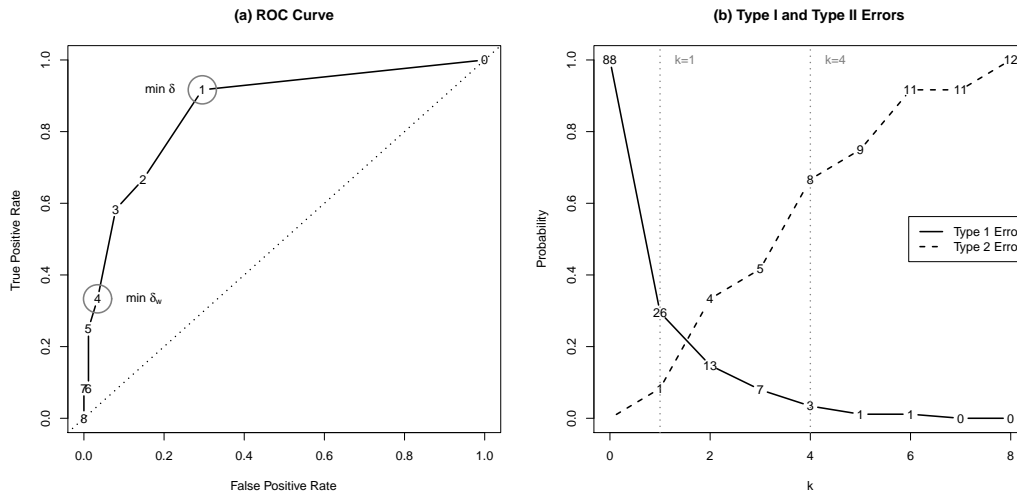


Figure 1: Plots for the numerical illustration.

RTM Algorithm: Network Estimation from CSS Slices, Controlling both Type 1 and Type 2 Error Rates

1. Draw a random sample of size n and observe the CSS slices.
2. Calculate the average density \bar{d} of all CSS slices in the sample, and set $w = 1/\bar{d}$.
3. Perform the weighted ROC analysis to find the k such that δ_w is minimized, and denote it by k^* .
4. Compute the estimated network by aggregating the CSS slices using FTM with threshold k^* .

As will be illustrated in the following section, our numerical study across five data sets indicate that setting the weight to $w = 1/\bar{d}$ is a reasonable choice. However, for various reasons, a researcher might want to use a different weight, and our ROC based methodology may be employed for any choice of w . This may be needed for certain social networks such as terrorist networks in which costs of Type 2 errors may be considered to be significantly larger than Type 1 errors. We finalize this section by noting that FTM, ATM, and RTM algorithms may be implemented by the functions `ftm`, `atm` and `rtm`, respectively, in the `cssTools` package for the statistical software R. The illustrative High Tech Managers data set is also included in the package, along with functions for producing tables and figures similar to Table 2 and Figure 1, and some useful tools for the analysis of CSS data.

4. Numerical Study

In this section we present the results of a numerical study aiming to illustrate the performance of the proposed estimation methodology. Since the adaptive threshold method (ATM) of Siciliano et al. (2012) for estimating networks from a random sample of CSS slices is known to outperform

traditional roster and ego network methods, we only compare the proposed ROC based adaptive threshold methodology (RTM) with ATM. For the Italian University data set introduced in Section 2, we generate random samples of sizes 4 to N , estimate the network based on the sample information only, and compute the correlation (S_{14}) between the estimated network and the true network obtained by the intersection method. Our results are displayed in Figure 2, which displays the boxplots of the observed correlations for different estimation methods. In this plot, the horizontal axis represents the sample size, the vertical axis represents the correlation, and the boxplots illustrate the distribution of the observed correlations based on 1000 simulations. In other words, for each sample size the vertical boxes represent the middle 50%, and the vertical lines represent the lower and upper 25% of the distribution of the observed correlations between the true and estimated networks. The colors indicate the method used for estimation.




Figure 2: Italian University data. Correlations between true and estimated networks.

We may see from the figure that the overall performances of both methods considered seem to be satisfactory. As expected, the correlation between the estimated and true networks increase as the sample size increases, and when all individuals are included in the sample the correlation equals one. The performances of the two methods considered seem to be comparable, however, ATM turns out to be sensitive to the selection of tolerable Type 1 error rate α . We ran similar simulations for the remaining four data sets introduced in Section 2 and obtained similar results. When using ATM, setting α to a larger value such as 0.15 generally seem to be a poor selection, which is most visible for Silicon Systems and Pacific Distributors data sets. This is an expected result as too much Type 1 error is allowed in the estimation procedure. On the other extreme, setting α to a smaller value such as 0.01 seems to produce better results across most data sets, however, as we saw that in Pacific Distributors data for small to moderate sample sizes, choice of 0.15 may outperform choice of 0.01. This may be because of the fact that forcing to reduce Type 1 error to a great extent may trigger more Type 2 errors and thus result in a larger overall error rate. Across the five data sets considered, setting α to 0.05 or 0.10 seem to be safer, but this feature cannot be generalized to any data set. In short, a researcher to employ ATM may be advised to set α to 0.05 or 0.10, but, they must also be advised that these are the preferable levels for the five data sets investigated, and it is not guaranteed that they will be the best choices for any data set.

On the other hand, the proposed RTM performs as good as ATM with a good choice of α , and it provides consistent results across all data sets. In this sense it is a robust method for estimating networks, and the fact that it doesn't require a pre-determined decision threshold such as α is appealing. RTM has an automatic way of balancing Type 1 and Type 2 errors, and our numerical study shows that its estimation performance is comparable with ATM, even when the ATM user is assumed to choose the best α level.

5. Conclusion

We first presented the common patterns of errors across five CSS data sets. Consistent in all data sets is the strong negative association between Type 1 and Type 2 errors. For the individuals with lower overall errors, most of the errors are of Type 2, whereas for the individuals with higher overall

errors most of the errors are of Type 1. Due to low density of the true networks considered, a common feature in many social networks, the uninformed perceivers tend to make much more Type 1 errors. These findings have especially important implications for measurement of network data, hence investigation of the determinants of tendency for different error types may be a useful next step in CSS research.

Next, we presented a novel estimation method, RTM, which requires a sample of CSS slices from the network of interest, and returns an estimate of the network while controlling for both Type 1 and Type 2 errors. RTM utilizes the ROC curve for balancing two error types, and it does not require a user decision on any cut-off values (such as α of ATM). Our comparative numerical study shows that the proposed method performs well. RTM gives weights to error types based on the network density, however, researchers may also choose to set their preferred weights depending on the emphasis they need to give to error types depending on their consequences.

Using RTM, it is possible to incorporate network data collection to random survey studies covering multiple organizations. Likewise more scholars are recognizing the need to study organization networks at multiple levels (Brass et al. 2004; Zappa and Lomi, 2015; Tranmer et al., 2016), however, data collection requirements are extremely demanding for studying inter-organizational and intra-organization networks simultaneously. RTM provides a convenient approach to meet data requirements for multi level organization research as well. One limitation of RTM is the network size. Due to obvious individual cognitive limitations, it may not be wise to apply RTM to networks involving more than 70 individuals. For large networks, a clustered version of RTM that is based on major departments or units in organizations can be considered.

Despite the promise of random sampling, gathering complete network data on the actors under study remains the best and recommended approach for collecting network data. We see these methods as useful when a researcher expects or obtains a low response rate or when the scope of the study prohibits the collection of network data from all participants, such as one looking at networks in a large number of organizations.

REFERENCES

- Anseel, F., Lievens, F., Schollaert, E., Choragwicka, B. (2010). Response Rates in Organizational Science, 1995-2008: A Meta-analytic Review and Guidelines for Survey Researchers. *Journal of Business and Psychology*, 25, 335-349.
- Baruch, Y., Holtom, B.C. (2008). Survey response rate levels and trends in organizational research. *Human Relations*, 61, 1139-1160.
- Brass, D. J., Galaskiewicz, J., Greve, H. R., Tsai, W. (2004). Taking stock of networks and organizations: A multilevel perspective. *Academy of management journal*, 47(6), 795-817.
- Casciaro, T. (1998). Seeing things clearly: Social structure, personality, and accuracy in social network perception. *Social Networks*, 20(4), 331-351.
- Dessi, R., Gallo, E., Goyal, S. (2016). Network cognition. *Journal of Economic Behavior and Organization*, 123, 78-96.
- Fawcett, T. (2004). ROC graphs: Notes and practical considerations for researchers. *Machine Learning*, 31(1), 1-38.
- Flynn, F. J., Reagans, R. E., Guillory, L. (2010). Do you two know each other? Transitivity, homophily, and the need for (network) closure. *Journal of Personality and Social Psychology*, 99(5), 855.
- Krackhardt, D. (1987). Cognitive social structures. *Social Networks*, 9(2), 109-134.
- Krackhardt, D. (1990). Assessing the political landscape: Structure, cognition, and power in organizations. *Administrative Science Quarterly*, 35, 342-369.
- Krackhardt, D., Kilduff, M. (1990). Friendship patterns and culture: The control of organizational diversity. *American Anthropologist*, 92(1), 142-154.
- Krackhardt, D., Kilduff, M. (1999). Whether close or far: Social distance effects on perceived balance in friendship networks. *Journal of Personality and Social Psychology*, 76(5), 770.
- Siciliano, M. D., Yenigun, D., Ertan, G. (2012). Estimating network structure via random sampling: Cognitive social structures and the adaptive threshold method. *Social Networks*, 34(4), 585-600.
- Tranmer, M., Pallotti, F., Lomi, A. (2016). The embeddedness of organizational performance: Multiple Membership Multiple Classification Models for the analysis of multilevel networks. *Social Networks*, 44, 269-280.
- Wasserman, S., Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge.
- Zappa, P., and Lomi, A. (2015). The Analysis of Multilevel Networks in Organizations Models and Empirical Tests. *Organizational Research Methods*, doi 10.1177/1094428115579225