Multiclass Rotations in Epistemic Network Analysis

Mariah A. Knowles[1](✉), Amanda Barany[2], Zhiqiang Cai[1], and David Williamson Shaffer[1]

[1] University of Wisconsin–Madison, Madison, WI 53711, USA
{mariah.knowles,zhiqiang.cai}@wisc.edu, dws@education.wisc.edu
[2] Drexel University, Philadelphia, PA 19104, USA
amb595@drexel.edu

Abstract. The task of succinctly and insightfully discussing themes in the differences between several (three or more) groups in naturalistic, ethnographic research faces a number of constraints. The number of all possible pairs is a quadratic function of the number of groups, and prior order and stand-out subsets may not exist to narrow that number down. We define and compare methods for guiding this task during Epistemic Network Analysis.

Keywords: Epistemic Network Analysis · Means Rotation · Linear Discriminant Analysis · Multiclass Rotations · Singular Value Decomposition

1 Introduction

It is a common task in naturalistic, ethnographic research to model and discuss the differences between multiple groups. Our focus in this paper is on the case where one has three or more (ie, $g \geq 3$) groups, as this presents a number of challenges when (i) the number of groups continues to increase, (ii) there is no meaningful prior order in which to guide one's comparison, and (iii) there is no clear subset of the data one can justify giving narrowed attention to. Generally, this task amounts to identifying themes of difference: imagine considering what it is that makes any two groups different from one another, then succinctly summing up what you find. Actually approaching the task exhaustively like this quickly becomes too burdensome without some way to guide one's analytic focus. For example, to compare 15 groups this way one would need to consider 105 distinct pairs.

To get at this task, we first summarize existing approaches to structuring themes of difference throughout the past three years of ICQE. Second, we define and compare a number of dimensionality reduction techniques usable in Epistemic Network Analysis (ENA), namely Singular Value Decomposition (SVD), Linear Discriminant Analysis (LDA), and a method we define here, Multi-Class Means Rotation (MCMR). And finally we illustrate our approach using a wellknown dataset in our community, Nephrotex, showing how one might choose among these methods and arrive at a story structured around a succinct number of themes of difference.

C. Damşa and A. Barany (Eds.): ICQE 2022, CCIS 1785, pp. 58–70, 2023.
https://doi.org/10.1007/978-3-031-31726-2_5

One of the strengths of multiclass methods is that they provide a reduced number of axes around which one can discuss the differences of their groups: to compare 15 groups, one would only need 4 (at least) to 14 (at most) axes. Axes provide themes of difference in terms of spectra, and structuring one's telling of the story around these spectra may help alleviate the complexity inherent in telling stories that move over multiple group difference. However, as we show, the existing approach (SVD) fails to identify trends that actually discriminate between multiple groups; LDA and MCMR both overcome this, balancing between discrimination and ENA goodness of fit scores differently.

2 Theory and Prior Literature

2.1 Multiclass Comparisons at ICQE

In the past three years of ICQE, data structured into multiple groups (that is, $g \geq 3$) has been approached in numerous ways. Generally, these boil down to comparing each group to the collective rest, comparing all possible pairs, justifying some focus, or discussing general trends instead. In each case, these strategies impose limitations when the number of groups continues to grow, and these limitations differ when one's groups have vs. don't have a pre-existing sense of order.

In some cases, groups in the data have pre-existing ordinality that often aligns with the passage of time, such as weeks in a course or stages in an intervention, and so there may be better reason to talk about them in one order or another [1–14]. In other cases, the groups in the data have no sense of ordinality, such as schools or countries, and so the order in which one ought to discuss and compare them depends on one's storytelling substance, constraints, goals, and commitments [10–22].

Researchers approached these cases in one of six ways (Fig. 1):

1. *Punt the Ball*—One can forego discussing group differences and instead describe each group in its own right without direct or inferred comparisons [15].
2. *One Against the Rest*—One can describe how each individual group in turn compared to all other groups together, perhaps after interpreting the grand mean of all groups [10, 11, 16–19].
3. *All Pairs*—One can describe each possible difference in each possible pairing of two groups [20]. In the ordinal case, one can also describe the differences of each adjacent pair of groups [6, 14].
4. *General Trend*—One can interpret possible meanings of the four plotted quadrants and use those to discuss overall trends in differences [14]. In the ordinal case, one can also fit or justify an overall temporal trend, then describe features of that trend [5, 7, 8, 12, 13].
5. *Justified Focus*—One can describe only a subset of possible pairings and provide a rationale for one's focus on that subset [11–13, 21, 22].
6. *Play the Tape*—One can, in the ordinal case only, describe the empirical qualities of individual groups in early-to-late order, perhaps running through this order multiple times to focus on the changes in particular qualities [1–5, 10]

With the exception of a few that used sequential pairwise means rotations (which only considered two groups at a time), most used an SVD projection to guide and/or illustrate these descriptions.

Each of these approaches impose limitations when g grows and one lacks a sense of ordinality. *All Pairs* approaches are prohibitively dense, as there are up to $(g^2 - g)/2$ possible pairs, which is too great a burden on one's page length and reader's attention to fully describe. *One Against the Rest* approaches are less dense as they require only g steps, and in some cases this may be appropriate, but in others this may lead to redundant descriptions of similar groups or not lead to clear insights about common patterns of difference among groups. *Justified Focus* approaches solve these issues, but only when one is fortunate enough to have data with immediately clear stand-out patterns. And while *General Trend* approaches exist in the ordinal case, authors have relied on SVD rotations to infer these trends in the non-ordinal case: as we show in our results below, SVD is ill-suited for this task, as it aims to maximize *overall* variance, not *between group* variance, and thus can fail to show differences that otherwise exist in the data. In theory, a *General Trend* approach could describe the themes of differences between all possible pairs of non-ordinal groups in as few as $\lceil \log_2 g \rceil$ axes, each axis dividing the groups in two in roughly orthogonal ways. At most, one would need $g - 1$ axes, which would amount to a *One Against the Rest* approach but dropping one axis, as it would be redundant with the rest. The method we propose achieves this lower bound in our worked example below, while capturing differences failed to be seen by SVD.

2.2 Singular Value Decomposition, Linear Discriminant Analysis, & Means Rotation

Let us consider two dimensionality reduction techniques commonly used in quantitative ethnography, Singular Value Decomposition (SVD) and Means Rotation (MR), as well as a related technique, Linear Discriminant Analysis (LDA). All three seek to find an axis of a high dimensional space that maximizes some aspect of variance: SVD maximizes overall variance, MR maximizes between-group variance of two groups, and LDA maximizes between-group variance while minimizing within-group variance (put another way, LDA maximizes effect size) [23, 24].

The calculations for SVD and LDA are closely related. Where X is one's high dimensional data, S_{cov} is the covariance matrix of X, $\overline{x}$ is the mean vector of X, S_b is the between-group scatter matrix of X given g groups, $\overline{x}^{(i)}$ is the mean vector of group i within X, and n_i is the sample size of the ith group, we first compute.

$$S_{cov} = \frac{1}{n-1}\left((X - \overline{x})^T (X - \overline{x})\right)$$

$$S_b = \sum_{i=1}^{g} n\left(\overline{x}^{(i)} - \overline{x}\right)\left(\overline{x}^{(i)} - \overline{x}\right)^T$$

then, in SVD, one finds the eigenvalues and eigenvectors of S_{cov} and uses those vectors with the highest eigenvalues to determine the axes of one's lower dimensional embedding; in LDA one does the same, instead finding the eigenvalues and eigenvectors of

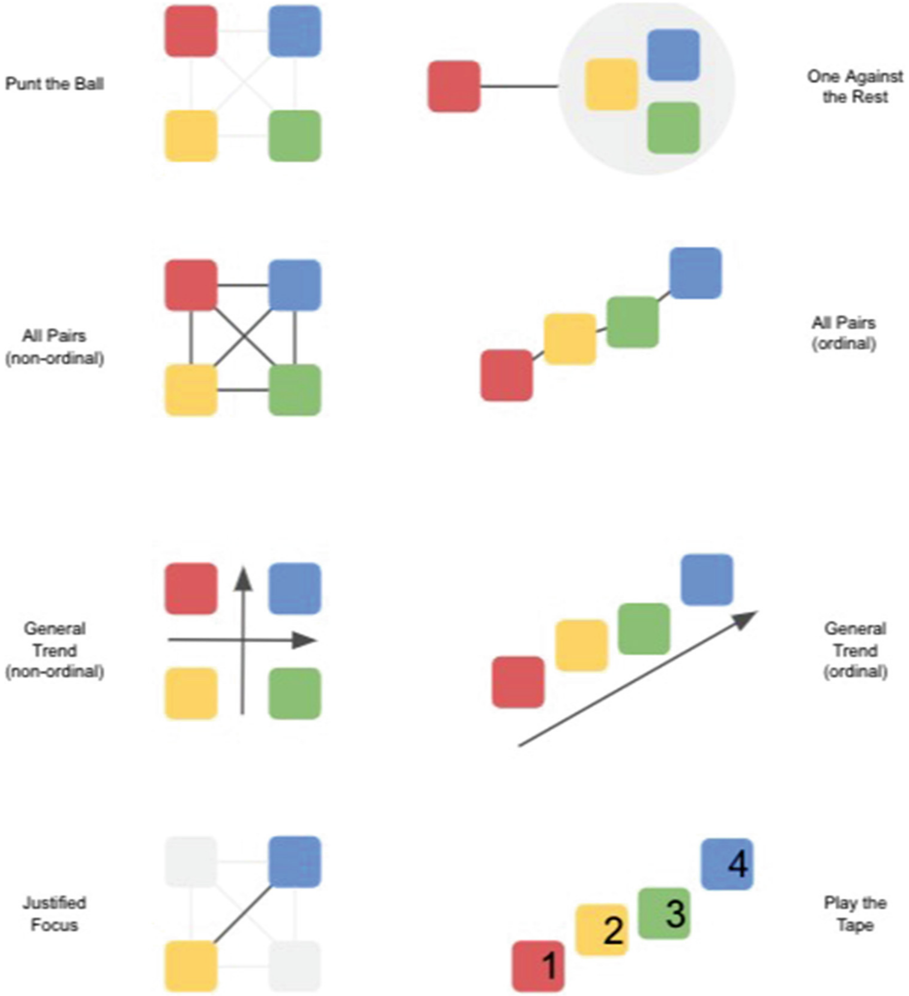

Fig. 1. Illustrations of the approaches taken in the past three years of ICQE for exploring themes of differences among three or more groups

${S_{cov}}^{-1}S_b$. Note, SVD is guaranteed to find orthogonal axes in all cases, while LDA only guarantees this when S_{cov} is symmetric. A number of approaches have been proposed to address this and related limitations of LDA [24–35]. For the sake of demonstration, we consider a simple approach, discussed in the section below.

For MR, given two groups j and k, one instead takes $\overline{x}^{(k)} - \overline{x}^{(j)}$ as the x-axis and the first dimension of an SVD of the remaining dimensions as the y-axis.

In essence, each technique highlights different features of the data. SVD finds the dimensions that highlight the greatest overall differences between units in one's higher dimensional embedding, a useful task when one seeks to understand the major turns of one's global structure quickly. LDA maximizes the discrimination (effect size) between groups in one's data, a useful task when one seeks to design automated classifiers. And MR gives an easily interpretable x-axis for non-technical readers: the x-axis (in most cases) runs through the two group means. Moreover, MR can be generalized through a regression framework, which allows one to moderate this projection for possible conflations or hierarchical effects often seen in nested data (*eg.*, students within classes within schools within districts) [36]. However, MR is limited to the $g = 2$ case, and so unlike SVD and LDA, MR (as it currently stands) is not appropriate for modeling $g \geq 3$ groups simultaneously. As we show in the proposed method below, MR can be reformulated using the same framework as SVD and LDA, allowing it to be generalized to the multiclass case.

2.3 Epistemic Network Analysis Rotations

In this paper, we assume familiarity with Epistemic Network Analysis (ENA) [23, 37–40]. Still, some ground clearing is worthwhile about how ENA rotates high dimensional data.

The general process of ENA involves three steps: we construct a high dimensional model of the connections between qualitative codes; we reduce the dimensionality of that space while highlighting features of interest; and we project a network into that space as a way to illustrate its dynamics [11, 23, 38, 40]. Let X represent this high dimensional space, where X_{ij} corresponds to the ith unit's connection strength between the jth pair of qualitative codes.

In whatever rotation method one chooses in the ENA tool (rENA or WebENA [38–40]), the rotation amounts to reducing the dimensionality of X by finding a pair of vectors, v_x and v_y, such that Xv_x and Xv_y are the dimensions that most highlight one's features of interest. Because these dimensions are taken as the x- and y-axis of the ENA plot and the distances between plotted points must be uniformly interpretable (as in a rigid body rotation), we have the further requirements that v_x and v_y be orthogonal to one another and have equal length. In a case where one's underlying dimensionality reduction technique does not produce orthogonal axes (as with LDA), we can instead take as our y-axis an approximation found by rejecting v_y from v_x and re-normalizing [36]. Put another way, when v_x and v_y are not exactly orthogonal, we identify the *plane* they exist in, rotate that plane such that v_x aligns with our x-axis, and plot the result.

For the sake of demonstration, this is the technique we use for ensuring our proposed methods conform to ENA's rigid body requirements.

3 Methods

3.1 Proposed Method: Multiclass Rotations

To date, the only linear projection used (that we are aware of) for simultaneously comparing $g \geq 3$ groups in an ENA context is SVD. And by default, this is the behavior of WebENA except when $g = 2$ exactly, where MR is used instead [39].

We consider two alternatives to those methods, LDA and a multiclass generalization of MR (MCMR), which together we think of as members of a more general class of possible multiclass rotations: rotations of an ENA space designed to highlight differences among $g \geq 3$ groups when ordinality is not guaranteed. Moreover, the process of these two approaches is identical to SVD rotations, except LDA considers the eigenvalues and eigenvectors of ${S_{cov}}^{-1} S_b$ and MCMR considers that of just S_b. That is, SVD maximizes overall variance, LDA maximizes between-group variance while minimizing within-group variance, MCMR only maximizes between-group variance, and none of the these approaches is more or less conceptually complex than the other.

We claim that MCMR generalizes MR: the two are identical along the x-axis when $g = 2$. Let us sketch a proof: Let j and k be our two groups and $S_{cov}^{(j)}$ and $S_{cov}^{(k)}$ be their covariance matrices such that $S_{cov}^{(j)} + S_{cov}^{(k)} = S_{cov}$. MCMR's generalized eigenvalue problem is $S_b v = \lambda v$, where λ is an eigenvalue and v is an eigenvector. This is equivalent to an LDA eigenvector problem $S_b v = \lambda S_{cov} v$ when the covariance matrix is proportional to the identity matrix, *ie.* when the columns of X are exactly independent. In such a case, it is known that the solution of LDA is proportional to the vector $\left(S_{cov}^{(j)} + S_{cov}^{(k)}\right)^{-1}\left(\overline{x}^{(k)} - \overline{x}^{(j)}\right) \propto x^{-(k)} - x^{-(j)}$. That is, MCMR is a special case of LDA which, when $g = 2$, reduces exactly to the definition of MR.

Whereas a generalization of MR based on a regression framework adds the ability to control one's projection in any way one can a regression [36], this generalization of MR based on an eigenvector framework adds the ability to explore differences between $g \geq 3$ groups even when ordinality is not guaranteed.

The question is, which of these two multiclass rotations is better (and when), what are the features of that difference, and what can these features tell us about telling stories around themes of difference between non-ordinal groups?

3.2 Data

To illustrate the task of telling a story of multiclass difference, we turn to Nephrotex [41]. We choose this dataset because (i) the ICQE community is familiar with it and (ii) it has a manageable number of multiple groups. Nephrotex was implemented across $g = 5$ schools (Iowa, KSU, Pitt, Rowan, and UW) during 2014 and 2015, and Nephrotex outcomes have been reported related to professional thinking [41], entrepreneurial mindsets [42], and complex collaborative thinking [43].

Nephrotex is a virtual internship designed to synchronously guide student groups through authentic biomedical engineering experiences. This provides students an educational task in which they can come to practice and understand the roles as engineers. This task has been designed with deliberate difficulties— problems to be overcome—to help guide students' learning and help them develop the skills necessary to achieve their goals. Nephrotex is also a collaborative environment where students are expected and encouraged to work together, and participants often ask each other for help in response to the data they encounter. As we show in the results below, on average, the discourse between the five sites differed in how students talked about these facets of the internship experience.

3.3 Evaluation

For our task of using one or more ENA plots to illustrate themes of difference, such plots need to be useful in a number of senses: they need to illustrate discrimination between one's groups (where there are differences), and the network embedding needs to aid interpretation of the space in trustworthy ways. So, to compare the usefulness of SVD, LDA, and MCMR, we will consider the discrimination between sites (Kruskal-Wallis H), the variance explained along the relevant axes (R^2), and the co-registration Pearson correlation of the network embedding along the relevant axes (r). Finally, we will use the best of these approaches to demonstrate how one might use it when closing the interpretive loop.

4 Results

4.1 Comparing Approaches

Table 1 summarizes the evaluation results for each method. SVD is designed to maximize variance explained, so naturally it outperforms the others along this metric. Moreover, it is worth noting that LDA explains only a small amount of variance. This suggests to us that LDA is too eager to minimize the variance within groups, and we see the results of this in the coregistration metrics: LDA is the only one that does not have a near-perfect score, scoring 10 percentage points lower than SVD and MCMR. Finally, SVD underperforms on discrimination between groups; LDA and MCMR have an H score more than 8 times that of SVD.

Because our goal is to tell a story about group differences, and because this task demands the ability to discriminate between groups, it is clear that LDA and MCMR are more appropriate models than SVD. However, the choice between LDA and MCMR depends on one's commitments: if one values fit of the network embedding higher, then MCMR wins out; if one values discrimination higher, then it's LDA; and if one values discrimination so long as network fit does not fall below some threshold, then it depends on where that threshold is set. Because, on inspection, MCMR and LDA were both able to discriminate between any pair of groups within their first three axes, and because we value network fit highly, we chose to explore and compare patterns in the Nephrotex dataset using MCMR.

Table 1. Statistical Evaluations

Model	R^2	r	H
SVD	.2830	.9964	8.201
LDA	.0685	.8921	74.22
MCMR	.1178	.9946	69.76

4.2 Quantitative Results

With $g = 5$ groups, the MCMR algorithm may produce up to $g - 1 = 4$ axes.

However, we focus on just the first three axes of the rotation (Fig. 3), *ie.* those that highlight the most between group variance. We do this because, altogether, these suffice to show how any one school was different from any other school. Along all three axes, there were significant differences between at least one pair of groups ($p_1 < .0001, H_1 > 69, p_2 < .0001, H_2 > 26, p_3 < .0001, H_3 > 24, g = 5$) and each had a high co-registration Pearson correlation which suggests strong goodnesses of fit between the visualizations and original models ($r_1 > .99$, $r_2 > .96$, $r_3 > .96$).

Figure 3 illustrates these axes. At a glance, the first MCMR axis discriminates between Pitt vs. Rowan vs. the rest in terms of client requests vs. technical constraints. The second discriminates between UW vs. Pitt and Iowa in terms of talk demonstrating the work of engineers vs. collaboration with one another within the virtual internship. And the third discriminates between KSU vs. the rest in terms of data-driven design vs. the affordances of the virtual platform. This suffices to show the differences between any pair of groups, and it achieves the theoretical lower bound of $\lceil \log_2 g \rceil = 3$ axes. These features of the data amount to a minimum number of themes of difference along which we might organize our qualitative account.

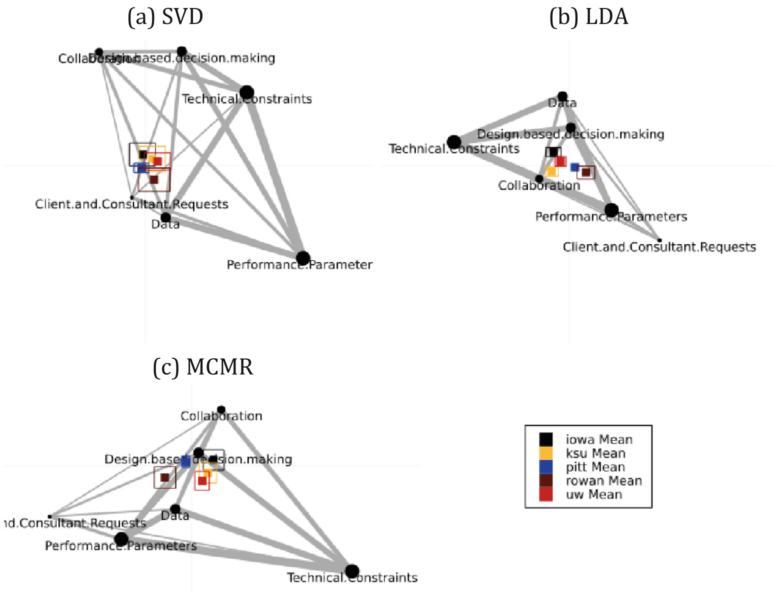

Fig. 2. ENA plots for all three models, showing the grand mean of connection strengths and confidence intervals for each school

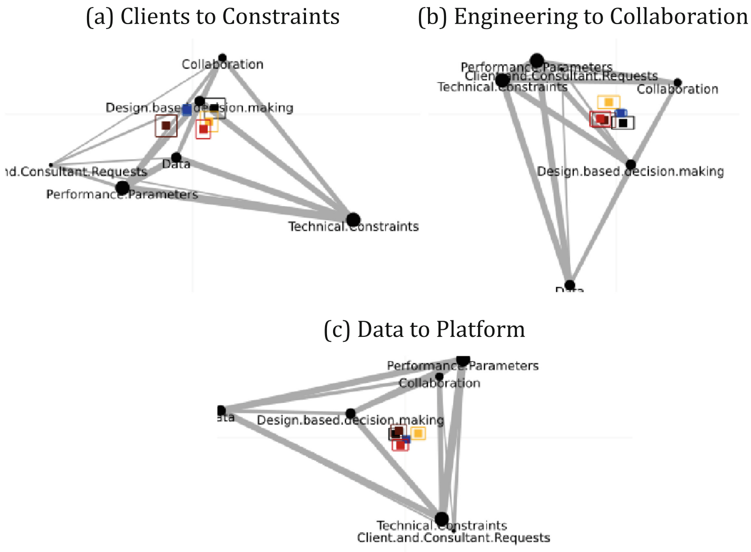

Fig. 3. First three axes of the MCMR rotation

4.3 Qualitative Description

One of the goals of Nephrotex as a virtual internship is for students to practice meeting stakeholder needs *as an engineer* (Fig. 3b). Throughout the internship's activities, students work together and with mentors, and as they do, they verbalize their understanding of the relationships between stakeholder needs and various engineering decisions. Along these lines, groups at UW, more often than Iowa or Pitt, discussed (1) hitting the design requirements of the virtual internship (*eg*., "I have submitted my surfactant data to Alex twice and both times he has told me that some of my data is incorrect") and (2) using performance data to inform their design choices (*eg*., "I agree with [student] in saying that steric hindering was the best option. It provided the most categories scoring in the higher ranges.").

Notably, the internship purposefully presented the students with tensions between stakeholder needs and constraints on the design space, and teams engaged with this *balancing act* in different ways (Fig. 3a). At Rowan, more so than other sites, this discourse centered around the requests of the internship's stakeholders as presented to students, as well as how students imagined future stakeholders' needs (*eg*., "I found our reliability at least meets the required and preferred standard of both consultants"). On the other hand, at Iowa, KSU, or UW, groups talked more about the burdens of technical constraints themselves (*eg*., "Cost was also a factor in my previous decision, otherwise the steric hindering surfactant would have been my top choice [goes on to list specific prices for choices in Nephrotex]"). Pitt, having much more variance in its implementation than the other schools, spanned this spectrum.

Finally, while setting parameter thresholds for their design in order to achieve their design goals, *in response to data*, and within the hard boundaries set by the internship tool, students often asked one another for help (Fig. 3c). This occurred least at KSU, where conversations favored more general discussions of the affordances of the Nephrotex platform (*eg*., “The biological surfactant could be a good option if we could lower its cost or improve its reliability”).

5 Discussion

In this work, we explore the use of three rotations for simultaneously comparing $g \geq 3$ groups, SVD, LDA, and MCMR, seeking to understand the contexts in which each approach might best serve a research project based on data structure and goals. Using data from Nephrotex, which consists of a number of comparable groups, we applied all three approaches, choosing MCMR as the best fit for the further exploration. We then shared the visualizations of the dataset with the MCMR rotation applied across three axes, and illustrated how these visualizations can be used to tell stories of non-ordinal themes of differences among multiple groups simultaneously.

We see two main takeaways for this discussion: MCMR and LDA’s improvement over SVD, and the role of multiclass rotations in illustrating themes of difference.

First, this work illustrated how MCMR and LDA approaches improved upon SVD in terms of discrimination between groups in the data. Yes, all three can be used to produce a set of axes that could guide an approach to storytelling organized around general trends in the data, but SVD may fail to identify trends that actually discriminate between groups. MCMR and LDA overcome this. The choice between these two depends on one’s commitments. When one prioritizes network embedding, they should choose MCMR. And when one prioritizes discrimination between groups, and lower coregistration Pearson correlations are acceptable, LDA may be more appropriate.

And second, this work showcased how models generated using a multiclass rotation can help to tell an ethnographic story of differences between several schools’ use of Nephrotex. We considered the first three axes of the MCMR rotation: this allowed us to illustrate the differences between any pair of schools in the fewest number of axes. Moreover, these axes illustrate the structure of one’s themes of difference by providing a set of spectra identifying different aspects of the data. This modeling process helps to alleviate the complexity inherent in telling stories that move over multiple groups. Exhaustively exploring all possible pairs of groups, exploring all possible ways to compare one group against the rest, and being fortunate enough to see readily clear stand-out patterns—these are unreasonable asks of researchers as the number of groups grows. Instead, a well-chosen ENA rotation can more directly illustrate a minimum number of spectra around which one can structure their qualitative account.

In future work, the authors hope to explore the pros and cons of MCMR and LDA approaches across more diverse datasets and context, with the goal of offering a roadmap for future QE scholars for well-reasoned choice between available rotations.

References

1. Shah, M., Foster, A., Talafian, H., Barany, A.: Examining the impact of virtual city planning on high school students’ identity exploration. In: Eagan, B., Misfeldt, M., Siebert-Evenstone, A. (eds.) Advances in Quantitative Ethnography. ICQE 2019. CCIS, vol. 1112. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-33232-7_17

2. Espino, D.P., et al.: Reflections of health care workers on their in-hospital experiences during the onset of COVID-19. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_17
3. Bressler, D.M.: Understanding off-topic utterances: do off-topic comments serve a purpose in collaborative learning? In First International Conference on Quantitative Ethnography: Conference Proceedings Supplement (2019)
4. Ha, S.Y., Lin, T.-J.L.: Development of epistemic cognition aboutsocial knowledge through collaborative small-group discussions. In: First International Conference on Quantitative Ethnography: Conference Proceedings Supplement (2019)
5. Brohinsky, J., Marquart, C., Wang, J., Ruis, A.R., Shaffer, D.W.: Trajectories in epistemic network analysis. In: Ruis, A.R., Lee, S.B. (eds.) ICQE 2021. CCIS, vol. 1312, pp. 106–121. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-67788-6_8
6. Wakimoto, T., et al.: Student teachers' discourse during puppetry-based microteaching. In: Eagan, B., Misfeldt, M., Siebert-Evenstone, A. (eds.) Advances in Quantitative Ethnography. ICQE 2019. CCIS, vol. 1112. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-33232-7_20
7. Wright, T., Oliveira, L., Espino, D.P., Lee, S.B., Hamilton, E.: Getting there together: examining patterns of a long-term collaboration in a virtual STEM makerspace. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_22
8. Barany, A., Philips, M., Kawakubo, A.J.T., Oshima, J.: Choosing units of analysis in temporal discourse. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_6
9. Mochizuki, T., et al.: Effects of perspective-taking through tangible puppetry in microteaching and reflection on the role-play with 3d animation. In: Eagan, B., Misfeldt, M., Siebert-Evenstone, A. (eds.) ICQE 2019. CCIS, vol. 1112, pp. 315–325. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-33232-7_28
10. Espino, D.P., et al.: News media communication of risk and mitigation factors during early stages of the covid-19 pandemic. In: Second International Conference on Quantitative Ethnography: Conference Proceedings Supplement, p. 23 (2021)
11. Carmona, G., Galarza-Tohen, B., Martinez-Medina, G.: Exploring interactions between computational and critical thinking in model-eliciting activities through epistemic network analysis. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_23
12. Knowles, M.A.: Telling stories of transitions: a demonstration of nonlinear epistemic network analysis. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_8
13. Mohammadhassan, N., Mitrovic, A.: Discovering differences in learning behaviours during active video watching using epistemic network analysis. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_24
14. Benna, A.M., Reynolds, K.: Teachers' beliefs shift across year-long professional development: ENA graphs transformation of privately held beliefs over time. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_13
15. Bressler, D.M.: Differences in group communication between game and nongame collaborations. In: First International Conference on Quantitative Ethnography: Conference Proceedings Supplement (2019)

16. Barany, A., Shah, M., Foster, A.: Connecting curricular design and student identity change: an epistemic network analysis. In: Ruis, A.R., Lee, S.B. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1312. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-67788-6_11
17. Phillips, M., Siebert-Evenstone, A., Kessler, A., Gasevic, D., Shaffer, D.W.: Professional decision making: reframing teachers' work using epistemic frame theory. In: Ruis, A.R., Lee, S.B. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1312. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-67788-6_18
18. Ma, L.: Using epistemic network analysis to explore emergent discourse dynamics of a grade 2 knowledge building community. In: First International Conference on Quantitative Ethnography: Conference Proceedings Supplement (2019)
19. Vachuska, K.: Using epistemic network analysis to measure and identify racialidentity development stages. In: First International Conference on Quantitative Ethnography: Conference Proceedings Supplement (2019)
20. Schnaider, K., Schiavetto, S., Meier, F., Wasson, B., Allsopp, B.B., Spikol, D.: Governmental response to the COVID-19 pandemic - a quantitative ethnographic comparison of public health authorities' communication in Denmark, Norway, and Sweden. In: Ruis, A.R., Lee, S.B. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1312. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-67788-6_28
21. Scianna, J., Kaliisa, R., Boisvenue, J.J., Zörgő, S.: Approaching structured debate with quantitative ethnography in mind. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_3
22. Hamilton, E.R., Lee, S.B., Charles, R., Molloy, J.: Peering a generation into the future: assessing workforce outcomes in the 2020s from an intervention in the 1990s. In: Wasson, B., Zörgő, S. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1522. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93859-8_11
23. Bowman, D., et al.: The mathematical foundations of epistemic network analysis. In: Ruis, A.R., Lee, S.B. (eds.) Advances in Quantitative Ethnography. ICQE 2021. CCIS, vol. 1312. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-67788-6_7
24. Van Loan, C.F., Golub, G.: Matrix computations (johns hopkins studies inmathematical sciences). Matrix Computations (1996)
25. Chu, D., Goh, S.R.: A new and fast orthogonal linear discriminant analysis on undersampled problems. SIAM J. Sci. Comput. **32**(4), 2274–2297 (2010)
26. Dai, D.-Q., Yuen, P.C.: Regularized discriminant analysis and its application to face recognition. Pattern Recogn. **36**(3), 845–847 (2003)
27. Friedman, J.H.: Regularized discriminant analysis. J. Am. Statist. Assoc. **84**(405), 165–175 (1989)
28. Chen, L.-F., Mark Liao, H.-Y., Ko, M.-T., Lin, J.-C., Yu, G.-J.: A new lda-based face recognition system which can solve the small sample size problem. Pattern Recogn. **33**(10), 1713–1726 (2000)
29. Howland, P., Jeon, M., Park, H.: Structure preserving dimensionreduction for clustered text data based on the generalized singular value decomposition. SIAM J. Matrix Anal. Appl. **25**(1), 165–179 (2003)
30. Howland, P., Park, H.: Generalizing discriminant analysis using the generalized singular value decomposition. IEEE Trans. Pattern Anal. Mach. Intell. **26**(8), 995–1006 (2004)
31. Huang, R., Liu, Q., Lu, H., Ma, S.: Solving the small samplesize problem of lda. In 2002 International Conference on Pattern Recognition, vol. 3, pp. 29–32. IEEE (2002)
32. Park, H., Drake, B.L., Lee, S., Park, C.H.: Fast linear discriminant analysis using QR decomposition and regularization. Technical report, Georgia Institute of Technology (2007)

33. Ye, J., Yu, B.: Characterization of a family of algorithms for generalized discriminant analysis on under sampled problems. J. Mach. Learn. Res. **6**(4) (2005)
34. Ye, J., Janardan, R., Park, C.H., Park, H.: An optimization criterion for generalized discriminant analysis on undersampled problems. IEEE Trans. Pattern Anal. Mach. Intell. **26**(8), 982–994 (2004)
35. Ye, J., Xiong, T., Madigan, D.: Computational and theoretical analysis of null space and orthogonal linear discriminant analysis. J. Mach. Learn. Res. **7**(7) (2006)
36. Knowles, M., Shaffer, D.W.: Hierarchical epistemic network analysis. In: Second International Conference on Quantitative Ethnography: Conference Proceedings Supplement. ICQE (2021)
37. Shaffer, D.W.: Quantitative ethnography. Lulu. com (2017)
38. Shaffer, D.W., Collier, W., Ruis, A.R.: A tutorial on epistemic network analysis: analyzing the structure of connections in cognitive, social, and interaction data. J. Learn. Anal. **3**(3):9–45 (2016)
39. Marquart, C.L., Hinojosa, C., Swiecki, Z., Eagan, B., Shaffer, D.W.: Epistemic networkanalysis (version 1.5. 2)[software] (2018)
40. Shaffer, D., Ruis, A.: Epistemic network analysis: a worked example of theory based learning analytics. Handbook of learning analytics (2017)
41. Arastoopour, G., et al.: Measuring first-year students' ways of professional thinking in a virtual internship. In: 2012 ASEE Annual Conference & Exposition, pp. 25–971 (2012)
42. Rogy, K.M., Bodnar, C.A., Clark, R.M.: Examining the entrepreneurial mindset of senior chemical engineering students as a result of exposure to the epistemic game "nephrotex". In: 2014 ASEE Annual Conference & Exposition, pp. 24–559 (2014)
43. Ruis, A.R., Siebert-Evenstone, A.L., Pozen, R., Eagan, B., Shaffer, D.W.: A method for determining the extent of recent temporal context in analyses of complex, collaborative thinking. In: 13th International Conference of the Learning Sciences (ICLS) 2018, vol. 3 (2018)