Proceedings of the 89th Annual International Meeting of the Psychometric Society, Prague, Czech Republic (2024).

Exploratory subgroup detection of psychological networks: Assessing the impact of ordinal and skewed data

T. Kvetnaya,* K. J. Nehler, and M. Schultze

Department of Psychology, Goethe University Frankfurt, 60478 Frankfurt am Main, Germany
*Corresponding author. Email: kvetnaya@psych.uni-frankfurt.de

Abstract

Exploratory subgroup identification can be a valuable tool for psychological network science, e.g., to identify patient subgroups with distinct symptom constellations in mental disorders. Gaussian mixture modeling (GMM) – a popular method for investigating heterogeneity in multivariate data – offers a promising avenue to achieve this. GMM approaches allow participants to be clustered into subgroups based on their subgroup-specific network structures, rather than symptom profiles or sumscores. Recent advancements in graphical GMM approaches were extended to explicitly consider the structure of associations among variables within each cluster (e.g., Fop et al., 2019). By introducing a graph structure search step into the expectation–maximization (EM) algorithm, it allows for not only optimizing parameters but also graph edge sets. However, this approach assumes continuous, normally distributed data, whereas real-world psychological data is often ordinal and/or skewed in nature. In this study, we seek to explore how effectively the structural EM algorithm is able to recover underlying subgroups in data under conditions frequently encountered in psychological data. To this end, we generate cross-sectional data stemming from 3 subgroups with different degrees of network sparsity, echoing findings from previous network analyses of psychological disorders. By varying the cluster proportions, the number of ordinal answer categories, and variable skewness in the simulated datasets, we evaluate the performance of graphical GMM in terms of clustering and structure recovery. Classification goodness, as well as recovery of the true cluster proportions, edge sets, and weight estimates are used as performance indicators.

Keywords: model-based clustering, ordinal data, network analysis, subgroup detection, Gaussian graphical models

1. Introduction

An intriguing question in network psychometrics is whether networks of psychological variables can be used to detect subgroups in an exploratory, data-driven manner. Model-based clustering, or Gaussian mixture modeling (GMM), is a popular method to uncover subgroups, which assumes that observations arise from a mixture of distributions, each representing one cluster or subgroup (Fraley, 1998). It is common in fields handling high-dimensional data like bioinformatics (Gao et al., 2016; McNicholas, 2010), though its use in psychology remains relatively scarce (i.e. Paul et al., 2019). For continuous, normally distributed data, GMMs reliably identify the correct number of clusters (Keribin, 2000; Scrucca et al., 2016). However, this performance is only guaranteed if the data-generating mechanism is correctly specified, with a common specification error occurring when ordinal data are treated as continuous (Haslbeck et al., 2023). The present simulation study explores the impact of ordinal and skewed data on a clustering method for mixtures of Gaussian graphical models (*mixggm,* Fop et al., 2019) – which is of particular interest for psychological applications.

1.1 Sample heterogeneity in network psychometrics

Sample heterogeneity refers to the possibility that observed data arise from several latent subgroups (Fried & Nesse, 2015; Hoekstra et al., 2022). For instance, Fried and Nesse (2015) demonstrated marked heterogeneity using questionnaire data from patients with depression, in which they identified 1,030 unique symptom profiles. Indeed, a number of psychological hypotheses suggest the existence of subpopulations characterized by different network structures. For example, network *connectivity* or *density* has been proposed to vary between groups with and without psychological disorders, with more severe symptomatology being associated with more densely connected symptom networks (Cramer et al., 2016; van Borkulo et al., 2015). While such hypotheses have been predominantly investigated by comparing a priori-defined groups, exploratory, data-driven approaches to identify possible subgroups based on within-cluster associations could be particularly valuable.

Untangling sample heterogeneity based on empirical network structure has proven challenging, as existing GMM frameworks rarely incorporate within-cluster associations as the guiding property for clustering (Fop et al., 2019). Some either assume local independence, i.e. diagonal within-cluster matrices (e.g. *pgmm*, McNicholas, 2010) or are computationally expensive and risk overparameterization (e.g. *mclust*, Haslbeck et al., 2023; Scrucca et al., 2016). While methods have been proposed to model within-cluster associations for continuous data (Gao et al., 2016; Ren et al., 2022), preliminary work showed a variety of convergence and estimation issues (Kvetnaya et al., 2023).

1.2 Ordinal and skewed data in psychology

A further challenge is that psychological data are frequently ordinal (Liddell & Kruschke, 2018) and/or not normally distributed in the population (e.g., suicidality scores), necessitating either robustness checks or dedicated modeling approaches. For example, many established diagnostic instruments (e.g., Beck et al., 1996; Zigmond & Snaith, 1983) use ordinal scales with four response categories, and Likert-type scales with 5 to 7 points are common (Liddell & Kruschke, 2018). This raises the longstanding question: When is it appropriate to treat ordinal data as continuous? Previous simulation research found that network estimation performance improved when ordinal categories exceed $c > 5$ and sample size N is sufficiently large ($N > 3{,}000$), even under moderate to strong skewness (Johal & Rhemtulla, 2021). Similarly, Haslbeck et al. (2023) found that high clustering accuracy with *mclust* (Scrucca et al., 2016) can be achieved for $c > 5$ categories when $N \geq 1{,}000$ and clusters are well-separated. However, especially with freely estimated covariance matrices, accuracy decreased as parameters grew with rising number of variables p, making model recovery difficult without a very large N ($\geq 10{,}000$).

While existing approaches like *clustMD* (McParland & Gormley, 2016) can accommodate ordinal data, they too assume local independence, as does latent class analysis (*LCA*, Visser & Depaoli, 2022). Another promising method for ordinal graphical models (Lee et al., 2022) was not readily available at the time of this study. We therefore proceed with the method by Fop et al. (2019), which explicitly considers within-cluster associations by optimizing for sparse precision matrices, thus avoiding overparameterization problems. In the presence of within-cluster associations, this method outperformed *mclust* (Fop et al., 2019). Therefore, the key question is whether existing findings about treating ordinal data as continuous apply to this method, which we will now introduce.

1.3 Model-based clustering for sparse precision matrices (Fop et al., 2019)

The present method assumes that observations x_i are generated by a mixture of K distributions. The probability density function for each x_i is given by:

$$f(\mathbf{x}_i \mid \boldsymbol{\Theta}, \mathbb{G}) = \sum_{k=1}^{K} \tau_k \, \phi\left(\mathbf{x}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Omega}_k, \mathcal{G}_k\right) \tag{1}$$

where the parameter vector $\boldsymbol{\Theta}$ comprises cluster proportions τ_k, means $\boldsymbol{\mu}_k$, and precision matrices $\boldsymbol{\Omega}_k$, with ϕ as the multivariate Gaussian density. $\mathbb{G}$ denotes the collection of graph structures $\mathcal{G}_k$, encoded by adjacency matrices $\mathbf{A}_k$.

Typically, for a fixed number of K, the classical EM algorithm (Dempster et al., 1977) alternates between the *expectation* (E) and the *maximization* (M) steps. Crucially, Fop et al. (2019) supplement the M step with a *structure search* (S) step, which determines the optimal $\mathcal{G}_k$.

In the E step, the conditional expectation of the log-likelihood ℓ is computed given the observed data and $\boldsymbol{\Theta}$ from the previous iteration. In the M step, the algorithm maximizes the regularized ℓ_R with respect to $\boldsymbol{\Theta}$ and $\mathbb{G}$, adding a penalty term Q to introduce sparsity in $\mathbf{A}_k$:

$$\ell_R = \sum_{i=1}^{N} \log \left\{ \sum_{k=1}^{K} \tau_k \, \phi \left(\mathbf{x}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Omega}_k, \mathcal{G}_k \right) \right\} + \sum_{k=1}^{K} \log p(\boldsymbol{\Omega}_k) - \sum_{k=1}^{K} Q(\mathbf{A}_k). \tag{2}$$

Here, the term $\log p(\boldsymbol{\Omega}_k)$ introduces Bayesian regularization to prevent singularity in $\boldsymbol{\Omega}_k$. In this work, we set the penalty Q to correspond to the Bayesian Information Criterion (BIC):

$$Q_{BIC}(\mathbf{A}_k) = \frac{1}{2} E_k \log N, \tag{3}$$

where E_k denotes the number of edges in the candidate graph $\mathcal{G}_k$. In the S step, to efficiently explore the space of possible graph configurations, a stepwise search or a probabilistic genetic algorithm are available. Once an optimal $\mathbb{G}$ is identified, the method estimates a positive-definite $\boldsymbol{\Omega}_k$ with fixed zero entries given $\mathcal{G}_k$. BIC is then used to select the model with the optimal number of clusters K_{opt}. Finally, individual observations are assigned to clusters.

1.4 Research question

While *mixggm* offers a promising way to tackle the first challenge and identify underlying subgroups in psychological data based on their cluster-specific associations, the second challenge of handling ordinality and skew remains unresolved. Our central question therefore is how tolerant *mixggm* is to violations of distributional assumptions. We conducted a simulation to evaluate the method's performance under different degrees of assumption violations. The outcome measures we focus on here are the recovery rate of the correct number of clusters K, the classification quality, and the recovery of edge structures within the clusters.

2. Method

All analyses were conducted in R (Version 4.3.3, R Core Team, 2024). Model-based clustering was implemented using the *mixggm* package (Version 1.0, Fop et al., 2019). All code used in the simulation study is available on OSF.io: https://osf.io/6n9tq.

2.1 Population networks

First, we generated cross-sectional partial correlation matrices for $K = 3$ clusters and $p = 10$ variables each to serve as true population networks for data generation in the simulation study. To reflect hypotheses about differential network density in psychological research (Cramer et al., 2016), we varied the proportion of non-zero edges across population networks (90%, 50%, and 10%, respectively). To define edge weights for each population network, we first randomly sampled values from an exponential distribution, with smaller values occurring more frequently. This choice was informed by characteristics reported in recent network analysis studies of psychological disorders (Wysocki & Rhemtulla, 2021). The sampled values were randomly assigned to the off-diagonal entries of each 10×10 matrix according to the predefined density level. Diagonal entries were

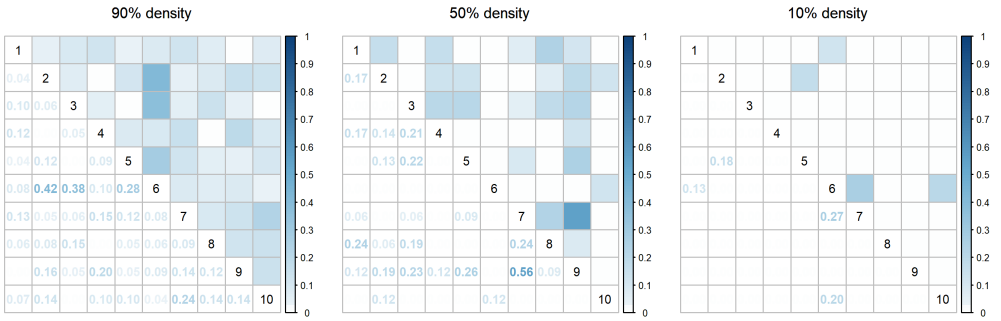

Figure 1. Partial correlation matrices representing the true population networks used to generate data in each simulation replication. Network density and edge weight placement vary by matrix, representing distinct correlational structures for each subgroup. Numbers on the main diagonal refer to variable indices.

iteratively increased by a small constant until positive definiteness was achieved, while preserving the original association structure. After this adjustment, the final partial correlation values ranged from 0.04 to 0.56. All three population matrices are illustrated in Figure 1.

2.2 Data generation

Multivariate normal data were generated based on the specified population matrices in each simulation replication using the *MASS* package, (Version 7.3-60.0.1., Venables & Ripley, 2002), preserving the correlational structure within each cluster. Cluster proportions τ were manipulated to be either balanced (1 : 1 : 1) or unbalanced (1 : 2 : 3). In the balanced condition, each cluster comprised 33.3% of the sample. In the unbalanced condition, the clusters contained n_1 = 16.7%, n_2 = 33.3%, and n_3 = 50% of the total sample, respectively. Importantly, in this condition, the smallest cluster (n_1) had the highest network density, while the largest cluster (n_3) had the lowest. This design reflects hypotheses such as the connectivity hypothesis of mental disorders, which suggests that a smaller subset of the population may exhibit higher symptom connectivity (Cramer et al., 2016).

Ordinal datasets were obtained by thresholding the generated continuous data following Rhemtulla et al. (2012). We applied symmetric, moderately, and heavily asymmetric thresholds to discretize data into 4, 5, or 7 ordinal categories with varied skews. Symmetric thresholds were evenly spaced around 0, covering –2.5 to 2.5 standard deviations (e.g., in four categories: –1.25, 0, 1.25). For moderate asymmetry, the peak shifted left of center, while for heavy asymmetry, the lowest category contained the most observations, with progressively fewer observations in higher categories.

2.3 Simulation settings

In total, the resulting the simulation conditions were as follows:

- Number of ordinal categories c: 4, 5, 7, continuous
- Skewness of ordinal data: symmetric, moderate, heavy
- Cluster proportions τ: balanced (1 : 1 : 1), unbalanced (1 : 2 : 3)
- Sample size N: 600, 1200, 2400, 4800, 9400

This resulted in 100 unique conditions. Conditions were fully crossed, with the exception that no condition involved continuous, skewed data due to the selected ordinalization method. For each condition, 150 datasets were simulated, on which clustering with *mixggm* was performed for $K = [1, 5]$. A forward stepwise search algorithm was chosen for the S step.

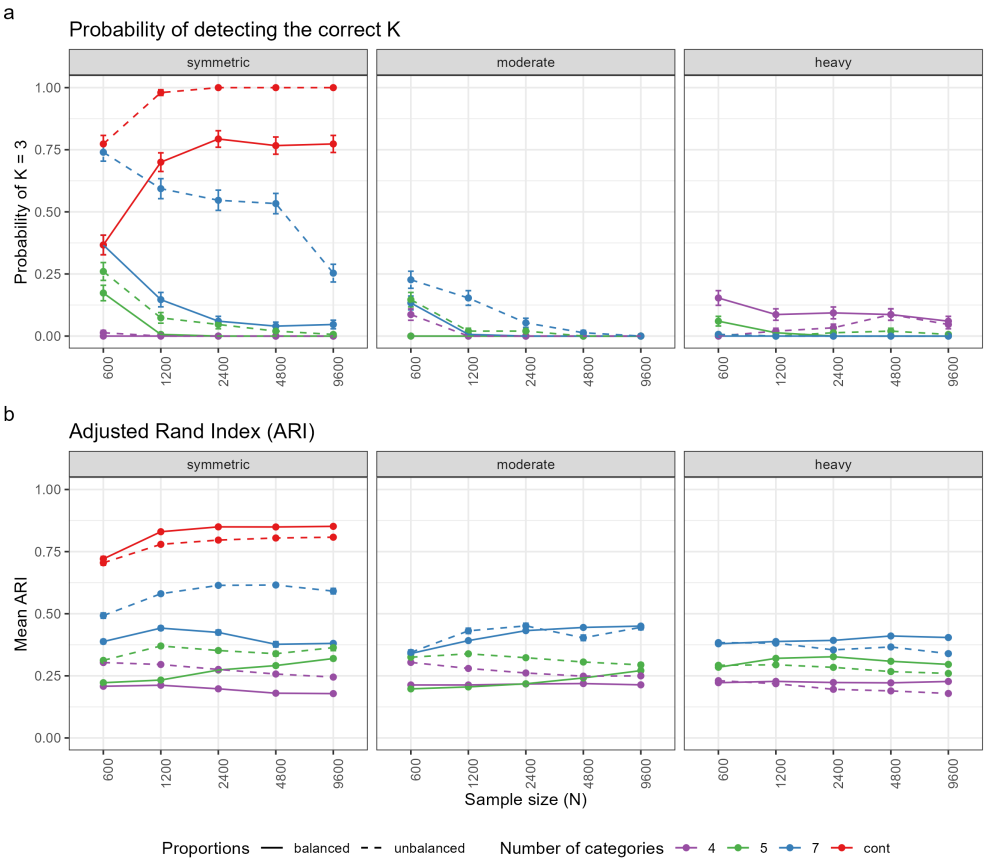

Figure 2. Probability to detect the correct number of clusters K (a) and classification quality (b) measured by ARI, as a function of skewness (symmetric, moderate, heavy), number of ordinal categories (4 - continuous), N, and balanced (solid line) vs. unbalanced (dashed line) cluster-size proportions in the population. Results for continuous data are presented only for the symmetric condition, as no skewed continuous data was generated in this study. Error bars indicate SE.

2.4 Evaluation criteria

We evaluated the performance of the model using the following metrics:

- Clustering performance: Probability of correctly identifying the number of clusters K
- Classification performance: Adjusted Rand Index (ARI, Hubert & Arabie, 1985), measuring the similarity between two data clustering assignments
- Sensitivity: True positive rate of edge recovery
- Specificity: True negative rate of edge recovery

Sensitivity and specificity were calculated only for replications in which estimated clusters could be matched to population clusters, which was achieved by minimizing SRMR between estimated and true covariance matrices.

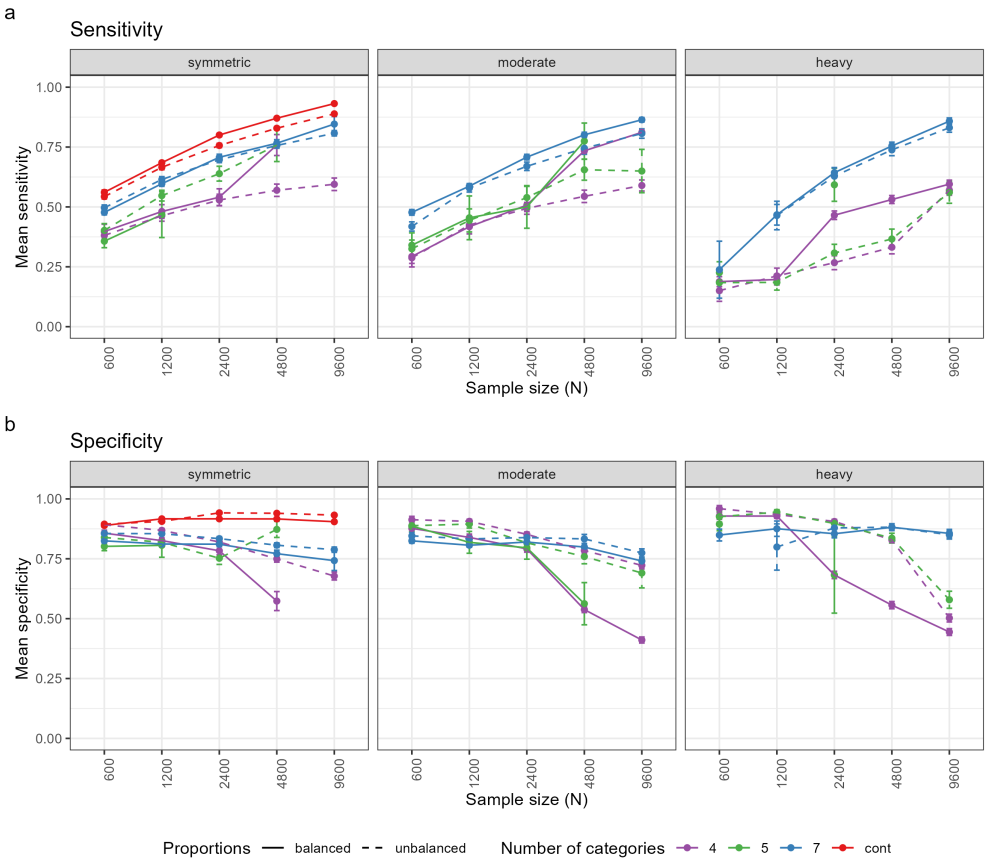

Figure 3. Sensitivity (a) and specificity (b) as a function of skewness (symmetric, moderate, heavy), number of categories (4 - continuous), N, and balanced (solid line) vs. unbalanced (dashed line) cluster-size proportions in the population. Results for continuous data are presented only for the symmetric condition, as no skewed continuous data was generated in this study. Error bars indicate SE.

3. Results

3.1 Clustering performance

Figure 2, Panel a, summarizes the proportion of simulations $P(K = 3)$ in which the correct number of K was identified. Notably, performance was superior in the unbalanced condition for almost all data settings. For symmetric, continuous data, $P(K = 3)$ increased steadily with sample size. For unbalanced clusters, it approached 1 with increasing N, while $P(K_{correct})$ stabilized at around 0.75 in the balanced case. However, this pattern did not hold for ordinal conditions, where performance decreased with c and dropped to 0 in the least favorable, skewed scenarios. This might be due to the method often overestimating clusters when K_{opt} is chosen incorrectly. For example, mean K_{opt} was 3.10 ($SD = 0.45$) clusters for continuous data, but rose up to $M = 4.71$ ($SD = 0.51$) clusters for $c = 4$. Larger sample sizes did not mitigate this effect. Only for non-skewed data with $c = 7$ categories, the method markedly outperformed other ordinal categories in the unbalanced condition, but even moderate skew led to a drop in performance, such that $P(K_{correct})$ did not exceed 0.23 in any ordinal, skewed condition.

3.2 Classification performance

As illustrated by Figure 2, Panel b, ARI consistently exceeded 0.80 for continuous, symmetric data (except for $N = 600$) and showed slight improvement with increasing N, indicating good classification performance. By contrast, it remained below 0.50 for ordinal conditions, with the exception of $c = 7$ symmetric data in the unbalanced cluster-size condition.

3.3 Cluster network recovery

Sensitivity grew with higher N in all conditions (Figure 3, Panel a). It was highest for continuous data, and was negatively impacted by the extent of skew and ordinality. By contrast, specificity (Panel b) mostly decreased with rising sample size except in the continuous, symmetric condition, which remained above 0.88 throughout. However, these findings warrant cautious interpretation, as both metrics could only be calculated for replications in which estimated and true clusters could be matched: While for for continuous data, 95.5% of replicates could be matched, this number dropped to 45.3% for $c = 7$, 41.6% for $c = 4$, and only 7.4% for $c = 5$.

4. Discussion

The results indicate that both clustering performance and edge recovery are substantially affected by the presence of ordinality and skewness in the data. While *mixggm* performed well for continuous, symmetric and sparse network data – particularly when underlying networks were sparse – its effectiveness diminished markedly in all ordinal and skewed conditions. This performance decrease could not be compensated for by a larger N, as seen for ARI, and for the probability of detecting the correct K, an increasing N was actually detrimental. Edge recovery was impaired compared to continuous data, especially considering these metrics could only be calculated for the small subset of replications. Consistently, unbalanced cluster-size conditions yielded better performance than balanced conditions, aligning with *mixggm*'s optimization for sparse structures: When sparse subgroups represent a larger share of the sample, the method benefits. This indicates that classification performance is constrained not by power alone, but also the characteristics of the data.

Taken together, these findings suggest that prior results from the network estimation (Johal & Rhemtulla, 2021) and model-based clustering literature (Haslbeck et al., 2023) do not generalize to the approach evaluated here. A possible explanation is the compounded complexity of combining clustering and network recovery within a single framework. However, alternative reasons due to choices made in this study and limitations may also contribute.

4.1 Limitations and Future Directions

Firstly, we fixed key parameters including population structure, K, p, and penalty parameter Q, making it difficult to determine whether reduced performance stems from modeling deficiencies, or unsuitable properties of the population data structure.

The selection of appropriate data generation methods represents another critical decision point. Following Rhemtulla et al. (2012), we simulated multivariate normal data first and subsequently introduced skew to ordinal data through asymmetric thresholds. Due to this, a continuous skewed condition was absent in our study, posing a significant limitation. While prior research indicates it may be impossible to distinguish between a normally distributed latent variable with asymmetric thresholds and a skewed latent distribution with symmetric thresholds (Grønneberg & Foldnes, 2019), alternative ordinalizing approached might yield different results.

Alternative methods designed specifically for ordinal data may hold promise. For example, while this method was not publicly accessible during this study, Lee et al. (2022) proposed an approach for clustering ordinal graphical models without assuming local independence, leveraging probit models to represent ordinal variables as discretizations of latent continuous variables. Tree-based methods such as SEM Trees and Network Trees (Grassi & Tarantino, 2023; Jones et al., 2020), which

recursively partition data based on directed or undirected edges, offer another avenue for settings where the structure of subgroups is complex or hierarchically organized.

5. Conclusion

In summary, this study demonstrates that violations of normality and sparsity assumptions substantially impact performance of *mixggm*. While we replicated its viability for network-based exploratory subgroup detection when assumptions hold, this study established that findings on treating ordinal data as continuous from previous literature do not directly carry over to this joint modeling approach. Careful consideration of data characteristics therefore remains essential when applying model-based clustering in psychological research.

References

Beck, A. T., Steer, R. A., & Brown, G. K. (1996). *Manual for the Beck Depression Inventory* (2nd edition). Psychological Corporation.

Cramer, A. O. J., Van Borkulo, C. D., Giltay, E. J., Van Der Maas, H. L. J., Kendler, K. S., Scheffer, M., & Borsboom, D. (2016). Major Depression as a Complex Dynamic System (I. Branchi, Ed.). *PLOS ONE*, *11*(12), e0167490. https://doi.org/10.1371/journal.pone.0167490

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum Likelihood from Incomplete Data Via the *EM* Algorithm. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *39*(1), 1–22. https://doi.org/10.1111/j.2517-6161.1977.tb01600.x

Fop, M., Murphy, T. B., & Scrucca, L. (2019). Model-based clustering with sparse covariance matrices. *Statistics and Computing*, *29*(4), 791–819. https://doi.org/10.1007/s11222-018-9838-y

Fraley, C. (1998). How Many Clusters? Which Clustering Method? Answers Via Model-Based Cluster Analysis. *The Computer Journal*, *41*(8), 578–588. https://doi.org/10.1093/comjnl/41.8.578

Fried, E. I., & Nesse, R. M. (2015). Depression sum-scores don't add up: Why analyzing specific depression symptoms is essential. *BMC Medicine*, *13*(1), 72. https://doi.org/10.1186/s12916-015-0325-4

Gao, C., Zhu, Y., Shen, X., & Pan, W. (2016). Estimation of multiple networks in Gaussian mixture models. *Electronic Journal of Statistics*, *10*(1), 1133–1154. https://doi.org/10.1214/16-EJS1135

Grassi, M., & Tarantino, B. (2023). SEMtree: Tree-based structure learning methods with structural equation models. *Bioinformatics*, *39*(6), btad377. https://doi.org/10.1093/bioinformatics/btad377

Grønneberg, S., & Foldnes, N. (2019). A Problem with Discretizing Vale–Maurelli in Simulation Studies. *Psychometrika*, *84*(2), 554–561. https://doi.org/10.1007/s11336-019-09663-8

Haslbeck, J. M. B., Vermunt, J. K., & Waldorp, L. J. (2023). The impact of ordinal scales on Gaussian mixture recovery. *Behavior Research Methods*, *55*(4), 2143–2156. https://doi.org/10.3758/s13428-022-01883-8

Hoekstra, R., Epskamp, S., & Borsboom, D. (2022). Heterogeneity in Individual Network Analysis: Reality or Illusion? *Multivariate Behavioral Research*, *58*, 1–25. https://doi.org/10.1080/00273171.2022.2128020

Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, *2*(1), 193–218. https://doi.org/10.1007/BF01908075

Johal, S. K., & Rhemtulla, M. (2021). Comparing estimation methods for psychometric networks with ordinal data. *Psychological Methods*, Advance online publication. https://doi.org/10.1037/met0000449

Jones, P. J., Mair, P., Simon, T., & Zeileis, A. (2020). Network Trees: A Method for Recursively Partitioning Covariance Structures. *Psychometrika*, *85*(4), 926–945. https://doi.org/10.1007/s11336-020-09731-4

Keribin, C. (2000). Consistent Estimation of the Order of Mixture Models. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)*, *62*(1), 49–66. Retrieved March 31, 2025, from https://www.jstor.org/stable/25051289

Kvetnaya, T., Nehler, K. J., & Schultze, M. (2023). Exploratory subgroup detection of psychological networks. An evaluation of current model based clustering approaches. *16th Conference of the DGPs Methods and Evaluation Section (FGME)*.

Lee, K. H., Chen, Q., DeSarbo, W. S., & Xue, L. (2022). Estimating Finite Mixtures of Ordinal Graphical Models. *Psychometrika*, *87*(1), 83–106. https://doi.org/10.1007/s11336-021-09781-2

Liddell, T. M., & Kruschke, J. K. (2018). Analyzing ordinal data with metric models: What could possibly go wrong? *Journal of Experimental Social Psychology*, *79*, 328–348. https://doi.org/10.1016/j.jesp.2018.08.009

McNicholas, P. D. (2010). Model-based classification using latent Gaussian mixture models. *Journal of Statistical Planning and Inference*, *140*(5), 1175–1181. https://doi.org/10.1016/j.jspi.2009.11.006

McParland, D., & Gormley, I. C. (2016). Model based clustering for mixed data: clustMD. *Advances in Data Analysis and Classification*, *10*(2), 155–169. https://doi.org/10.1007/s11634-016-0238-x

Paul, R., Andlauer, Till. F. M., Czamara, D., Hoehn, D., Lucae, S., Pütz, B., Lewis, C. M., Uher, R., Müller-Myhsok, B., Ising, M., & Sämann, P. G. (2019). Treatment response classes in major depressive disorder identified by model-based clustering and validated by clinical prediction models. *Translational Psychiatry*, *9*(1), 187. https://doi.org/10.1038/s41398-019-0524-4

R Core Team. (2024). *R: A Language and Environment for Statistical Computing* (Version 4.3.3.). Vienna, Austria. https://www.R-project.org/

Ren, M., Zhang, S., Zhang, Q., & Ma, S. (2022). Gaussian graphical model-based heterogeneity analysis via penalized fusion. *Biometrics*, *78*(2), 524–535. https://doi.org/10.1111/biom.13426

Rhemtulla, M., Brosseau-Liard, P. É., & Savalei, V. (2012). When can categorical variables be treated as continuous? A comparison of robust continuous and categorical SEM estimation methods under suboptimal conditions. *Psychological Methods*, *17*, 354–373. https://doi.org/10.1037/a0029315

Scrucca, L., Fop, M., Murphy, T. B., & Raftery, A. E. (2016). Mclust 5: Clustering, Classification and Density Estimation Using Gaussian Finite Mixture Models. *The R journal*, *8*(1), 289–317. Retrieved March 17, 2025, from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5096736/

van Borkulo, C., Boschloo, L., Borsboom, D., Penninx, B. W. J. H., Waldorp, L. J., & Schoevers, R. A. (2015). Association of Symptom Network Structure With the Course of Depression. *JAMA Psychiatry*, *72*(12), 1219–1226. https://doi.org/10.1001/jamapsychiatry.2015.2079

Venables, W. N., & Ripley, B. D. (2002). *Modern applied statistics with S* (4th ed). Springer.

Visser, M., & Depaoli, S. (2022). A Guide to Detecting and Modeling Local Dependence in Latent Class Analysis Models. *Structural Equation Modeling: A Multidisciplinary Journal*, *29*(6), 971–982. https://doi.org/10.1080/10705511.2022.2033622

Wysocki, A. C., & Rhemtulla, M. (2021). On Penalty Parameter Selection for Estimating Network Models. *Multivariate Behavioral Research*, *56*(2), 288–302. https://doi.org/10.1080/00273171.2019.1672516

Zigmond, A. S., & Snaith, R. P. (1983). The Hospital Anxiety and Depression Scale. *Acta Psychiatrica Scandinavica*, *67*(6), 361–370. https://doi.org/10.1111/j.1600-0447.1983.tb09716.x