

Population-Based Hierarchical Non-negative Matrix Factorization for Survey Data

Xiaofu Ding

Math Dept., UCLA

Los Angeles, United States
jding62@g.ucla.edu

Xinyu Dong

Math Dept., UCLA

Los Angeles, United States
xinyudong@g.ucla.edu

Olivia McGough

Math Dept., Reed College

Portland, United States
mcgougho@reed.edu

Chenxin Shen

Math Dept., UCLA

Los Angeles, United States
amyshen@g.ucla.edu

Annie Ulichney

Appl. Math Dept., Yale University

New Haven, United States
annie.ulichney@yale.edu

Ruiyao Xu

Math Dept., UCLA

Los Angeles, United States
ruiyaox@g.ucla.edu

William Swartworth

Math Dept., UCLA

Los Angeles, United States
wswartworth@gmail.com

Jocelyn T. Chi

Math Dept., UCLA

Los Angeles, United States
jtchi@math.ucla.edu

Deanna Needell

Math Dept., UCLA

Los Angeles, United States
deanna@math.ucla.edu

Abstract—Motivated by the problem of identifying potential hierarchical population structure on modern survey data containing a wide range of complex data types, we introduce population-based hierarchical non-negative matrix factorization (PHNMF). PHNMF is a variant of hierarchical non-negative matrix factorization based on feature similarity. As such, it enables an automatic and interpretable approach for identifying and understanding hierarchical structure in a data matrix constructed from a wide range of data types. Our numerical experiments on synthetic and real survey data demonstrate that PHNMF can recover latent hierarchical population structure in complex data with high accuracy. Moreover, the recovered subpopulation structure is meaningful and can be useful for improving downstream inference.

Index Terms—Non-negative matrix factorization, hierarchical clustering, survey data, latent classes, population structure

I. INTRODUCTION

Modern survey data are rapidly growing larger and more complex. Technological improvements in recent years and the ubiquity of social media have dramatically lowered barriers to collecting and storing a wide range of data types on numerous respondents and variables. Indeed, modern survey data now frequently contain not only traditional multiple choice questions but also a wide variety of questions enabling complex, open-form short responses.

A key issue in many studies involving survey data is the identification of latent hierarchical groups, or subpopulations, within the observations. However, traditional statistical methods for identifying subpopulations in survey data were developed prior to modern technologies that enable the ready incorporation of complex textual responses. Consequently, these methods may be limited in the types of data they can

This work was partially supported by NSF grants DMS-2011140, DMS-2108479, and DMS-2103093.

analyze as well as the numbers of variables that they can practically handle.

Motivated by the problem of identifying potential hierarchical population structure on modern survey data containing a wide range of complex data types, we introduce population-based hierarchical non-negative matrix factorization (PHNMF). PHNMF is a variant of hierarchical non-negative matrix factorization based on feature similarity. As such, it enables an automatic and interpretable approach for identifying and understanding hierarchical structure in a data matrix constructed from a wide range of data types. Our numerical experiments in Sections IV and V demonstrate that PHNMF can recover latent hierarchical population structure with high accuracy. Moreover, the recovered subpopulation structure is interpretable and can be useful for improving downstream inference.

In the rest of this paper, we first discuss some related works and highlight our contributions in Section I-A. We then set some notation and describe background on non-negative matrix factorization (NMF), its employment in topic modeling, and its hierarchical variant in Section II. We introduce PHNMF and highlight its differences from hierarchical NMF in Section III, describe the design of our numerical experiments in Section IV, and present our results in Section V. Finally, we conclude with a brief discussion and some directions for future work in Section VI.

A. Related Work and Contributions

Traditional statistical methods for discovering latent subgroups from observed data include latent class analysis (LCA) and latent profile analysis (LPA). LCA is a probabilistic method for discovering latent groupings of observed categorical variables and assumes that variables within each group

are independent ([1], [2]). LPA is the continuous data version of LCA. It is an instance of multivariate mixture estimation for continuous variables that likewise assumes that variables within each group are independent ([3], [4]). One can therefore employ these methods to cluster members of a population by assigning members to the most probable latent class. While the models obtained from LCA and LPA depict a flat clustering structure, one can also employ the methods recursively within groups to obtain a hierarchical clustering tree.

PHNMF makes several improvements over these methods. First, PHNMF makes no distributional assumptions on the observed data and does not require independence on subsets of the variables. Nonetheless, our numerical experiments in Section IV demonstrate that PHNMF can recover subpopulation structure with high accuracy. Therefore, PHNMF may better accommodate modern survey data, which can include complex data with non-trivial correlations.

A second advantage of PHNMF is improved interpretability of the resulting clustering structure. While LCA and LPA return the probabilities that the members belong to particular latent classes, they treat all variables within a class as equally important. By contrast, PHNMF returns the relative importance of each variable to a member's responses. This enables understanding of how the variables contribute to each member's clustering assignments.

A third advantage of PHNMF is that it is designed specifically for discovering hierarchical structure in complex data. Our numerical results in Section V illustrate how LCA and LPA may struggle to identify meaningful structure in complex data with a large number of variables. By contrast, PHNMF readily discovers interpretable hierarchical population structure that aligns with existing literature.

II. BACKGROUND

We first set some notation for the rest of this paper. We then present a brief overview of non-negative matrix factorization, its applications to topic modeling, and hierarchical non-negative matrix factorization.

We assume that all vectors are column vectors. Given a matrix \mathbf{X} , we denote its $(i, j)^{th}$ entry by $\mathbf{X}_{i,j}$, its i^{th} row by \mathbf{x}_i^T , and its Frobenius norm by $\|\mathbf{X}\|_F$. We employ the notation $\mathbf{X} \geq 0$ to indicate that it is non-negative so that all its entries are greater than or equal to 0.

A. Non-Negative Matrix Factorization (NMF)

Non-negative matrix factorization (NMF) [5] is a linear dimension reduction technique for interpretable analysis of non-negative data. Given a non-negative matrix $\mathbf{X} \in \mathbb{R}^{n \times m}$ and a rank $k \leq \min(n, m)$, NMF finds $\mathbf{W} \in \mathbb{R}^{n \times k}$ and $\mathbf{H} \in \mathbb{R}^{k \times m}$ that minimize

$$(\mathbf{W}, \mathbf{H}) = \arg \min_{\mathbf{W} \geq 0, \mathbf{H} \geq 0} \frac{1}{2} \|\mathbf{X} - \mathbf{WH}\|_F^2.$$

Therefore, NMF approximates \mathbf{X} as the product of two lower-dimensional factor matrices, \mathbf{W} and \mathbf{H} . NMF employs

the following multiplicative update (MU) rules from [5] and [6] to compute the factor matrices

$$\mathbf{W}_{i,j} \leftarrow \mathbf{W}_{i,j} \frac{(\mathbf{X}(\mathbf{H})^T)_{i,j}}{(\mathbf{WH}(\mathbf{H})^T)_{i,j}}, \quad \text{and} \quad (1)$$

$$\mathbf{H}_{i,j} \leftarrow \mathbf{H}_{i,j} \frac{((\mathbf{W})^T \mathbf{X})_{i,j}}{((\mathbf{W})^T \mathbf{WH})_{i,j}}. \quad (2)$$

Alg. 1 presents pseudocode for computing an NMF factorization. Compared with other matrix factorization techniques that can produce factor matrices containing negative components, the non-negativity constraint of NMF enables straight-forward interpretation. Due to its interpretability, NMF is popular in a wide variety of applications including image processing [7] and hyperspectral unmixing [8].

Algorithm 1: NMF

1 **Input:** Data matrix $\mathbf{X} \geq 0$, $k \leq \min(n, m)$
2: initialize $\mathbf{W} \geq 0 \in \mathbb{R}^{n \times k}$, $\mathbf{H} \geq 0 \in \mathbb{R}^{k \times m}$
2: **while** not converged **do**
3: Compute \mathbf{WH}
4: Update \mathbf{H} by MU rules in (2)
5: Update \mathbf{W} by MU rules in (1)
6: **end while**
7: **return** $\mathbf{W} \geq 0 \in \mathbb{R}^{n \times k}$, $\mathbf{H} \geq 0 \in \mathbb{R}^{k \times m}$

B. NMF for Topic Modeling

Let \mathbf{X} represent a corpus of documents with the documents on the rows and the words in the documents on the columns. Then NMF recovers topics within the corpus and classifies the topical composition of each document [9]. The resulting NMF factorization produces \mathbf{W} and \mathbf{H} such that \mathbf{W}_{ij} represents the weight of the j^{th} topic in the i^{th} document, and \mathbf{H}_{ij} represents the weight of the j^{th} word in the i^{th} topic. This factorization enables simultaneous identification of latent topics and understanding of the topical composition of each document.

C. Hierarchical NMF (HNMF)

Hierarchical NMF (HNMF) applies NMF recursively to reveal latent hierarchical topic structure in the data. Therefore, HNMF can offer a richer understanding than classical NMF, which assumes a flat latent structure.

There are multiple approaches to performing HNMF. We consider top-down HNMF, which first discovers high-granularity topics and then recursively discovers subtopics within broader ones [10]. It does this by recursively applying Alg. 1. Concretely, let $k^{(0)}$ denote the rank parameter at the first iteration. Then top-down HNMF first factorizes \mathbf{X} as

$$\mathbf{X} \approx \mathbf{W}_0 \mathbf{H}_0,$$

where $\mathbf{H}_0 \in \mathbb{R}^{k^{(0)} \times m}$ describes the composition of the broadest supertopics and $\mathbf{W}_0 \in \mathbb{R}^{n \times k^{(0)}}$ encodes the corresponding representations of each document. At the i^{th} iteration, top-down HNMF splits documents into submatrices

$\mathbf{X}_1^{(i)}, \mathbf{X}_2^{(i)}, \dots, \mathbf{X}_{k^{(i)}}^{(i)}$, where $\mathbf{X}_j^{(i)}$ for $1 \leq j \leq k^{(i)}$ contains only the documents that discuss the j^{th} topic as determined by the coefficients in \mathbf{H}_i and a minimum discussion threshold α . Top-down HNMF then applies NMF to each of these submatrices to discover the subtopics that form the j^{th} supertopic. The algorithm continues until subtopics no longer contain some minimum number of documents t . Alg. 2 presents pseudocode for these procedures.

Algorithm 2: Top-down HNMF

1 **Input:** Data matrix $\mathbf{X} \in \mathbf{R}^{n \times m}$ with $\mathbf{X} \geq 0$, minimum topic discussion threshold α , minimum number of documents in a topic t

- 1: Initialize $i \leftarrow 0$
- 2: Determine $k^{(i)}$, number of latent topics in \mathbf{X}
- 3: Compute $(\mathbf{W}_i, \mathbf{H}_i) = \text{NMF}(\mathbf{X}, k^{(i)})$
- 4: **for** each document $1 \leq l \leq n$ in \mathbf{X} **do**
- 5: Assign document l to topic submatrix $\mathbf{X}_j^{(i)}$ if $(\mathbf{W}_i)_{l,j} > \alpha$
- 6: **end for**
- 7: **for** each topic submatrix $\mathbf{X}_j^{(i)}$ **do**
- 8: **while** number of documents in topic $j > m$ **do**
- 9: Update $i \leftarrow i + 1$
- 10: Determine $k^{(i)}$, number of latent topics in $\mathbf{X}_j^{(i-1)}$
- 11: Compute $(\mathbf{W}_i, \mathbf{H}_i) = \text{NMF}(\mathbf{X}_j^{(i-1)}, k^{(i)})$
- 12: Assign documents in $\mathbf{X}_j^{(i-1)}$ to topic submatrices based on whether corresponding coefficients in $\mathbf{W}_i > \alpha$
- 13: **end while**
- 14: **end for**

Lines 2 and 10 in Alg. 2 involve determining an input rank k in Alg. 1. While Alg. 1 accommodates any method for choosing the rank, we highlight the method in [11, Algorithm 2]. This method computes NMF factorizations with each of $k \in [2, 9]$ topics and produces a score between 0 and 1 for each k based on the cohesiveness of the resulting topics. Higher scores indicate better model fit.

III. POPULATION-BASED HNMF (PHNMF)

We now present population-based hierarchical NMF (PHNMF), a variation on HNMF that automatically discovers latent population hierarchical structure based on feature similarity, where similarity is based on the pairwise cosine similarity as defined in [11, Section 5.1]. Suppose that \mathbf{X} is a data matrix with individuals on the rows and features derived from responses to survey questions on the columns. At each iteration, PHNMF divides the population into disjoint subpopulations by whether a respondent's coefficients for a subpopulation exceeds a threshold quantity α . The coefficients in the \mathbf{W} and \mathbf{H} factor matrices at each iteration determine population assignments and relative importance of the discovered features, respectively. Rather than terminating when the number of observations in a subpopulation falls below

a threshold quantity, however, PHNMF terminates when the feature similarity falls below a threshold amount.

At a high level, feature similarity measures the pairwise cosine similarity between the rows of the \mathbf{H} factor matrices obtained from multiple NMF runs on the same subpopulation with a pre-determined optimal rank and different initializations. To compute feature similarity, we employ a modified version of [11, Algorithm 2]. Rather than iterating over multiple ranks k , we employ only the pre-determined optimal rank. Additionally, rather than returning the maximum of the median value of the least similarity seeds, we return the minimum of the minimum of the least similarity seeds.

Alg. 3 presents pseudocode for PHNMF. PHNMF differs from HNMF in two key ways. First, PHNMF performs “hard” splits so that each level produces disjoint subpopulations. This ensures that each respondent belongs to at most one subpopulation at each level. Second, PHNMF employs feature similarity as the stopping criteria rather than terminating when each subpopulation acquires a minimum number of respondents. This feature similarity criteria simultaneously avoids arbitrary termination and produces more interpretable subpopulation discovery.

Algorithm 3: Population-Based HNMF (PHNMF)

1 **Input:** Data matrix $\mathbf{X} \in \mathbf{R}^{n \times m}$ with $\mathbf{X} \geq 0$, minimum subpopulation threshold α , minimum feature similarity threshold β

- 1: Initialize $i \leftarrow 0$
- 2: Determine $k^{(i)}$, number of latent subpopulations in \mathbf{X}
- 3: Compute $(\mathbf{W}_i, \mathbf{H}_i) = \text{NMF}(\mathbf{X}, k^{(i)})$
- 4: **for** each respondent $1 \leq l \leq n$ in \mathbf{X} **do**
- 5: Compute $j = \max_{1 \leq f \leq m} (\mathbf{W}_i)_{l,f}$
- 6: Assign respondent l to population submatrix $\mathbf{X}_j^{(i)}$ if $(\mathbf{W}_i)_{l,j} > \alpha$
- 7: **end for**
- 8: **for** each population submatrix $\mathbf{X}_j^{(i)}$ **do**
- 9: Compute feature similarity for $\mathbf{X}_j^{(i)}$
- 10: **while** feature similarity $> \beta$ **do**
- 11: Update $i \leftarrow i + 1$
- 12: Determine $k^{(i)}$, number of latent subpopulations in $\mathbf{X}_j^{(i-1)}$
- 13: Compute $(\mathbf{W}_i, \mathbf{H}_i) = \text{NMF}(\mathbf{X}_j^{(i-1)}, k^{(i)})$
- 14: Assign respondents in $\mathbf{X}_j^{(i-1)}$ to population submatrices based on whether corresponding coefficients in $\mathbf{W}_i > \alpha$
- 15: **end while**
- 16: **end for**

IV. EXPERIMENTAL SETUP

We demonstrate PHNMF on both synthetic and real data. Our synthetic data mimic responses from distinct subpopulations to both multiple choice and open-form questions. We also employ real data from two survey datasets.

A. Synthetic Data

We construct two synthetic datasets with hierarchical population structure containing continuous and categorical variables, respectively. These data mimic the numerical encoding of responses to multiple choice and open-form survey questions. Each dataset consists of a pair of synthetic explanatory variables and response vector. We describe our procedures for constructing each below.

1) *Synthetic Continuous Data*: We construct synthetic explanatory variables $\mathbf{X}^I \in \mathbb{R}^{1600 \times 120}$ containing 1600 observations on 120 continuous features with hierarchical population structure as described in Fig. 1. In the context of topic modeling for survey data, these data mimic a population containing eight groups, each containing 200 respondents. Each respondent discusses some combination of four topics and each topic consists of 30 words. While we describe these data in the context of topic modeling, our approach also applies to continuous variables in general.

We construct \mathbf{X}^I as the product of two rank-4 matrices with

$$\mathbf{X}^I = \mathbf{W}^I \mathbf{H}^I,$$

where $\mathbf{W}^I \in \mathbb{R}^{1600 \times 4}$ is a person-topic matrix and $\mathbf{H}^I \in \mathbb{R}^{4 \times 120}$ is a topic-word matrix. Therefore, $\mathbf{W}^I_{i,j}$ reflects the amount that the j^{th} topic is discussed by the i^{th} person, and $\mathbf{H}^I_{i,j}$ reflects the importance of the j^{th} word in the i^{th} topic.

We construct \mathbf{W}^I according to the hierarchical structure depicted in Fig. 1. At the first split, Groups 1 and 2 discuss words in Topic 1 according to zero-truncated $\mathcal{N}(64, 9)$ and $\mathcal{N}(3, 9)$ distributions, respectively. At the second split, Groups 1a and 2a and Groups 1b and 2b discuss words in Topic 2 according to zero-truncated $\mathcal{N}(45, 9)$ and $\mathcal{N}(3, 9)$ distributions, respectively. At the third split, Groups 1a1, 2a1, 1b1, and 2b1 and Groups 1a2, 2a2, 1b2, and 2b2 discuss words in Topic 3 according to zero-truncated $\mathcal{N}(3, 9)$ and $\mathcal{N}(50, 9)$ distributions, respectively. Finally, the entire population discusses words in Topic 4 according to a zero-truncated $\mathcal{N}(50, 9)$ distribution.

We construct \mathbf{H}^I as follows. Suppose that each of the four topics consists of 30 words. For simplicity, we assume that the words in each topic appear together in the columns of \mathbf{H}^I . For each topic, we construct 30 vectors of length four by sampling from a multinomial distribution so that the words in each topic are four times more likely to appear in that topic than in the others. This simulates a more disjoint vocabulary; words are less likely to appear in multiple topics simultaneously. Fig. 2 depicts the structure in the resulting synthetic explanatory variables matrix \mathbf{X}^I .

We construct a synthetic coefficient vector $\boldsymbol{\theta}_g \in \mathbb{R}^4$ with $1 \leq g \leq 8$ for the subgroups. We obtain a synthetic response $\mathbf{y}^I \in \mathbb{R}^{1600}$ by multiplying the corresponding rows of \mathbf{W}^I with $\boldsymbol{\theta}_g$ for $1 \leq g \leq 8$. This simulates a response that is determined by different combinations of the topics for each subgroup.

2) *Synthetic Categorical Data*: We construct synthetic explanatory variables $\mathbf{X}^{II} \in \mathbb{R}^{1600 \times 120}$ containing 1600 observations on 120 dichotomous categorical features with hierarchi-

cal population structure as described in Fig. 1. We construct \mathbf{X}^{II} according to

$$\mathbf{X}^{II} = \text{threshold}(\mathbf{W}^I \mathbf{H}^{II}),$$

where we employ the same \mathbf{W}^I and construct \mathbf{H}^{II} similarly to \mathbf{H}^I . We threshold the entries in the product $\mathbf{W}^I \mathbf{H}^{II}$ so that \mathbf{X}^{II} contains only 1's and 0's.

We highlight some differences in the procedures for constructing \mathbf{H}^{II} . Rather than employing 30 words in each topic as in \mathbf{H}^I , we vary the number of words in each topic to ensure ordered topical importance. To see this, notice that a topic's relative importance is determined by the magnitudes of its entries and its size, or the number of words it contains. For example, equal numbers of entries containing 1's and 0's and equal topic sizes produces equally important topics. Therefore, we vary the number of words in each topic to induce an importance hierarchy. Topic 1 consists of 65 words, Topic 2 consists of 30 words, Topic 3 consists of 20 words, and Topic 4 consists of 5 words. We again assume that words appear together by topic.

To ensure that \mathbf{X}^{II} contains only dichotomous categorical variables, we threshold the entries in the product $\mathbf{W}^I \mathbf{H}^{II}$ as follows. Suppose that the j^{th} word belongs to the l^{th} topic and let med_l denote the median value of the words in Topic l . Then we obtain $\mathbf{X}^{II}_{i,j}$ with

$$\mathbf{X}^{II}_{i,j} = \begin{cases} 1 & \text{if } (\mathbf{W}^I \mathbf{H}^{II})_{i,j} > \text{med}_l, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Fig. 3 depicts the structure in \mathbf{X}^{II} . We employ the same response variable \mathbf{y}^I since the hierarchical population structure determined by \mathbf{W}^I remains unchanged.

B. Real Data

1) *City of Austin Satisfaction Survey*: We employ real data containing survey responses from residents of Austin, Texas on their satisfaction with various city services [12]. Our dataset consists of the 3,723 completed surveys obtained since 2018 that do not contain any empty columns.

Approximately 200 columns correspond to multiple choice questions such as the following.

- 1) Overall quality of public safety services (i.e. police, fire and ambulance) (*Choices: Very Dissatisfied, Dissatisfied, Neutral, Satisfied, Very Satisfied, Don't Know*)
- 2) I feel safe in my neighborhood during the day (*Choices: Strongly Disagree, Disagree, Neutral, Agree, Strongly Agree, Don't Know*)
- 3) Have you had contact with the City of Austin Municipal Court? (*Choices: Yes, No, Don't Know*)
- 4) Which two items listed in Quality of Life do you think are most important for the City to provide? (*Choices: 'Feel Welcome', 'Overall quality of life', 'Place to Live', 'Place to Work', 'Raise Children', 'Retire', 'None'*)

In addition to the multiple choice questions, we include eight columns corresponding to open-form response questions on respondent satisfaction levels. We exclude open-form questions that closely resemble the multiple choice questions.

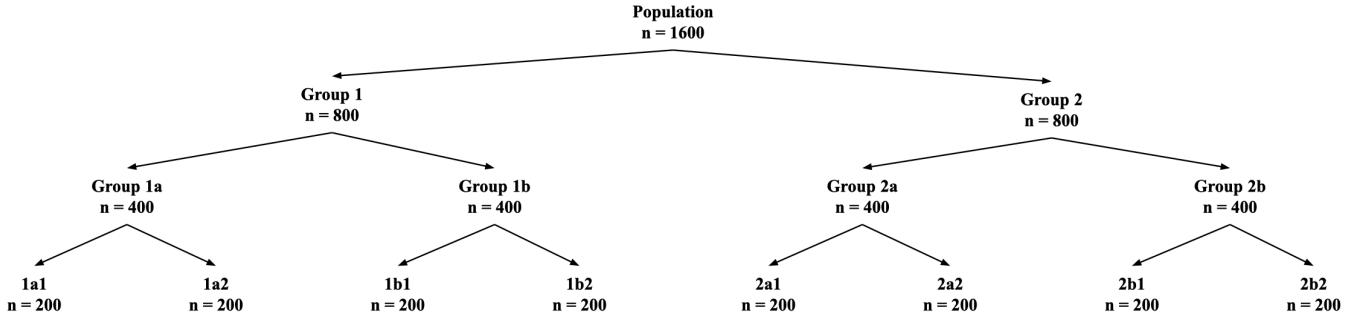


Fig. 1. Hierarchical population structure of synthetic datasets.

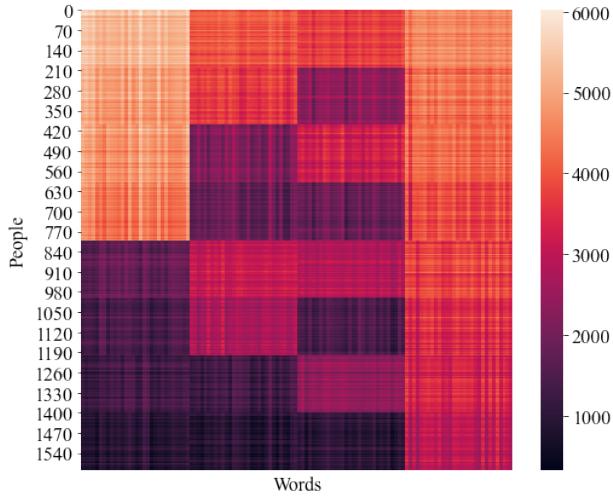


Fig. 2. Heatmap depicting structure of synthetic continuous data matrix \mathbf{X}^I .

We consolidate multiple choice responses on satisfaction to positive, neutral, and negative responses. We construct indicator variables for positive and negative responses and retain these as input variables. For each open-form question, we employ `tfidfvectorizer` from `scikit-learn` in Python [13] to convert text responses to a respondent-word matrix and perform NMF to identify the primary response topics. For each topic, we construct a new variable whose entries are the corresponding NMF coefficients indicating how much each respondent discusses each topic, and set the largest coefficient to 1. The resulting dataset contains 3,723 observations on 250 continuous and categorical variables. We exclude demographic variables in performing PHNMF and employ them only to ascertain the meaningfulness of the PHNMF subgroups.

2) *Facebook Climate Change Survey*: We employ real data containing multiple choice questions from a Facebook survey measuring public opinion on climate change from 2,432 respondents [14]. We emphasize that these data and results do not reflect our personal views on climate change. We apply PHNMF to the following seven multiple choice questions.

- 1) Do you think that global warming is happening? (Choices: Yes, No, Don't know)

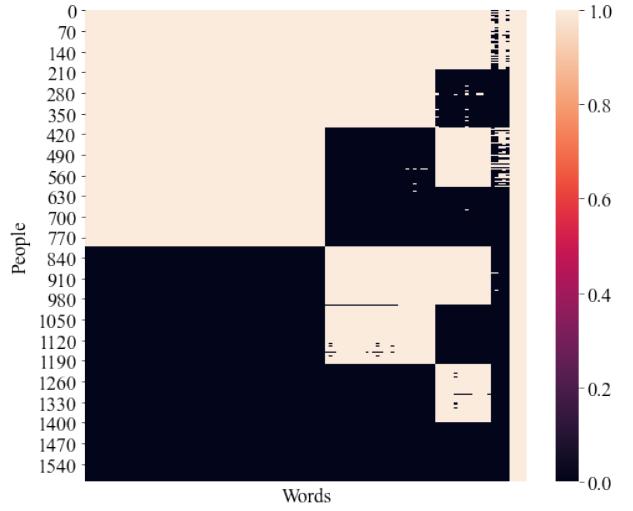


Fig. 3. Heatmap depicting structure of synthetic categorical data matrix \mathbf{X}^{II} .

- 2) Assuming global warming is happening, do you think it is mostly... (Choices: *Caused by human activities*, *Caused by natural environmental changes*, *Other*)
- 3) Which comes closest to your own view? (Choices: *Most scientists think global warming is not happening*, *There is much disagreement among scientists about whether or not global warming is happening*, *Most scientists think global warming is happening*, *Don't know*)
- 4) How much do you think global warming will harm you personally? (Choices: *A great deal*, *A moderate amount*, *Only a little*, *Not at all*, *Don't know*)
- 5) How much do you think global warming will harm future generations of people? (Choices: *A great deal*, *A moderate amount*, *Only a little*, *Not at all*, *Don't know*)
- 6) How much do you think global warming will harm plant and animal species? (Choices: *A great deal*, *A moderate amount*, *Only a little*, *Not at all*, *Don't know*)
- 7) How much do you support or oppose funding more research into renewable energy sources, such as solar and wind power. (Choices: *Strongly support*, *Somewhat support*, *Somewhat oppose*, *Strongly oppose*)

For each question, we construct indicator variables for each

multiple choice response. The resulting dataset contains 2,427 observations on 35 categorical variables. Again, we exclude demographic variables when performing PHNMF and employ them only to investigate the meaningfulness of the PHNMF subgroups.

C. Experimental Setup

We perform two sets of experiments to demonstrate the performance of our proposed method. 1) First, we perform experiments to test PHNMF’s ability to recover latent hierarchical clustering structure on both continuous and categorical data. 2) Second, we perform experiments to illustrate the potential usefulness of the recovered clustering structure for improved inference with linear and ridge regression as appropriate.

We perform 1,000 replicates with PHNMF for the continuous and categorical synthetic data described in Sec. IV-A. For each replicate, we construct a new pair of synthetic explanatory and response variables and record the classification accuracy.

We compare our results to those obtained from applying LPA and LCA for continuous and categorical data, respectively. We apply LPA to continuous data with both the `mclust` and `tidyLPA` packages for R ([15], [16]). We apply LCA to categorical data with the `poLCA` package for R [17]. While Algorithm 3 and our synthetic data construction employ Python, the packages for LCA and LPA are in R. LCA and LPA also require substantial time to determine an optimal number of classes for each dataset. Therefore, we perform 300 replicates with LCA and LPA. We run 10 replicates of LPA for each of $1 \leq k \leq 10$ classes to account for randomness in the initialization and select the model with lowest BIC [18].

To illustrate the potential usefulness of the recovered clustering structure for improved inference with synthetic data, we perform regression with the explanatory and response variables for each subpopulation discovered by PHNMF and the population as a whole. We record the cosine similarity between the resulting coefficient vectors and both the known coefficient vectors in Sec. IV-A and the coefficient vector we obtain from regression with the entire population. We construct the continuous synthetic data so that all the subgroups have full column rank in the explanatory variables to avoid issues with identifiability. Therefore, we employ linear regression for the continuous data. Since the categorical synthetic data may lack full column rank in the subgroups, we instead employ ridge regression.

Since the real datasets lack known coefficient vectors, we select the Facebook Climate Change Survey to illustrate the potential usefulness of the recovered clustering for downstream inference. We employ an ordinal variable containing a self-identified political ideology score ranging from 1 to 5, where 1 is “very liberal” and 5 is “very conservative”. Since some subpopulations discovered by PHNMF may lack full column rank in the explanatory variables, we perform ordinal ridge regression [19] with the `mord` package in Python. We employ three iterations of 5-fold cross-validation

over $\lambda = [0.0001, 0.001, 0.01, 0.01, 1]$ to select the λ that minimizes the mean squared error for the ridge penalty.

V. RESULTS

We present the results of our numerical experiments on synthetic and real datasets, respectively.

A. Synthetic Data Results

Fig. 4 and Fig. 5 depict data matrices recovered from a single replicate of the synthetic continuous and categorical data with PHNMF, respectively. They are qualitatively similar to Fig. 2 and Fig. 3.

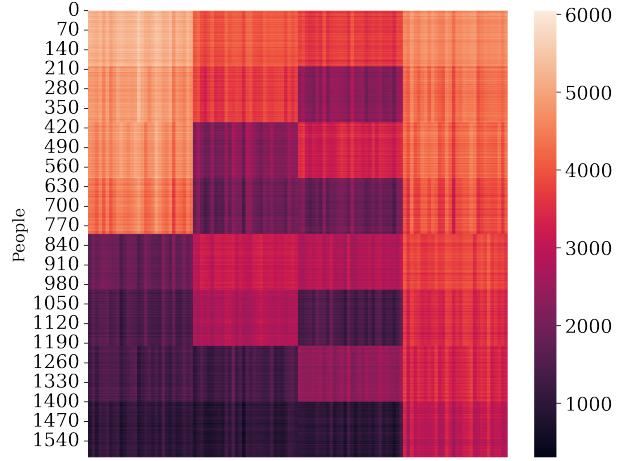


Fig. 4. Heatmap depicting structure of recovered and sorted synthetic continuous data matrix \mathbf{X}^I .

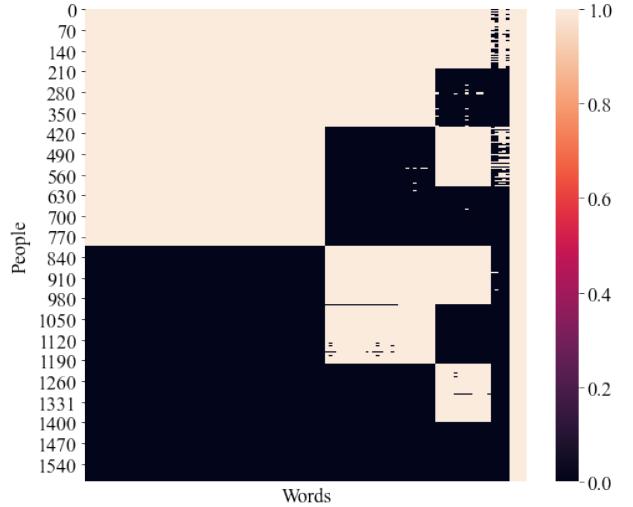


Fig. 5. Heatmap depicting structure of recovered and sorted synthetic categorical data matrix \mathbf{X}^{II} .

Tab. I depicts results of experiments testing PHNMF’s ability to recover latent hierarchical population structure on both continuous and categorical data. Notice that different

runs of our algorithm may return subgroups in a different ordering since a factorization $\mathbf{WH} = \mathbf{WP}^T \mathbf{PH}$, where \mathbf{P} is a permutation matrix. Therefore, we identify the eight subgroups by defining each group by the most frequently observed label for observations assigned to the group.

Tab. II shows that PHNMF recovers the latent hierarchical structure in both datasets with high accuracy. LCA performs nearly as well as PHNMF on the categorical data but does not handle continuous data. We attempt LPA with two R packages but neither produces functional results; `tidyLPA` assigns all observations to a single group while `mclust` does not run to completion. Since the packages can return results on subsets of the columns from the synthetic continuous datasets, it appears that the full datasets contain more variables than the LPA algorithms can practically handle.

TABLE I
CLUSTERING ACCURACY (MEAN AND STANDARD ERROR) OVER 1,000 REPLICATES OF PHNMF AND 300 REPLICATES OF LCA AND LPA ON SYNTHETIC DATASETS. WE DO NOT REPORT LPA RESULTS SINCE `TIDYLPA` ASSIGNS ALL OBSERVATIONS TO A SINGLE GROUP AND `MCLUST` DOES NOT RUN TO COMPLETION.

Method	Continuous Data		Categorical Data	
	Accuracy	Std. Err.	Accuracy	Std. Err.
PHNMF	98.5	0.04	99.97	0.0004
LCA	—	—	96.43	0.059
LPA	—	—	—	—

Tab. II presents results on the potential usefulness of the recovered PHNMF clustering structure for improved downstream inference from a single replicate of the continuous and categorical datasets, respectively. Since the cosine similarity computes the cosine of the angle between two vectors, it is 1 when two vectors are perfectly aligned and 0 when they are orthogonal. Columns 2 and 4 depict the cosine similarity between the known coefficient vectors and those recovered from regression on the subgroups for the synthetic continuous and categorical datasets, respectively. Columns 3 and 5 depict the cosine similarity between the coefficient vectors obtained from regression on the entire population and those recovered from regression on the subgroups for the synthetic continuous and categorical datasets, respectively.

Tab. II shows that coefficient vectors obtained on the PHNMF subgroups align more closely with the known coefficient vectors. This gives evidence that the subpopulation structure discovered by PHNMF may improve downstream inference.

B. Real Survey Data Results

1) *City of Austin Satisfaction Survey*: Fig. 6 depicts a heatmap of the population structure discovered with PHNMF on the City of Austin Satisfaction Survey. The heatmap suggests non-trivial structure on the columns and the rows.

Fig. 7 depicts the hierarchical population structure discovered with PHNMF on the same dataset. The population splits are interpretable and also provide an understanding of differences in survey responses among respondents. The first PHNMF split divides the 3,723 survey respondents into two

TABLE II
COSINE SIMILARITY BETWEEN RECOVERED COEFFICIENT VECTORS FROM REGRESSION ON THE SUBGROUPS COMPARED TO: 1) THE KNOWN SUBGROUP COEFFICIENT VECTORS AND 2) THE RECOVERED COEFFICIENT VECTORS FROM REGRESSION WITH THE ENTIRE POPULATION. COSINE SIMILARITY 1 INDICATES PERFECT DIRECTIONAL ALIGNMENT; 0 INDICATES ORTHOGONALITY.

Group Name	Continuous Data		Categorical Data	
	Subgroups	Population	Subgroups	Population
Group 1a1	0.8793	0.1020	0.9320	0.4773
Group 1a2	0.9730	0.2866	0.9157	0.5234
Group 1b1	0.9310	0.7022	0.9596	0.5008
Group 1b2	0.9271	0.0789	0.9117	0.5095
Group 2a1	0.9546	0.9229	0.9121	0.3572
Group 2a2	0.9642	0.9471	0.9959	0.3265
Group 2b1	0.9222	0.8786	0.9712	0.4728
Group 2b2	0.8023	0.3716	0.0100	0.7988

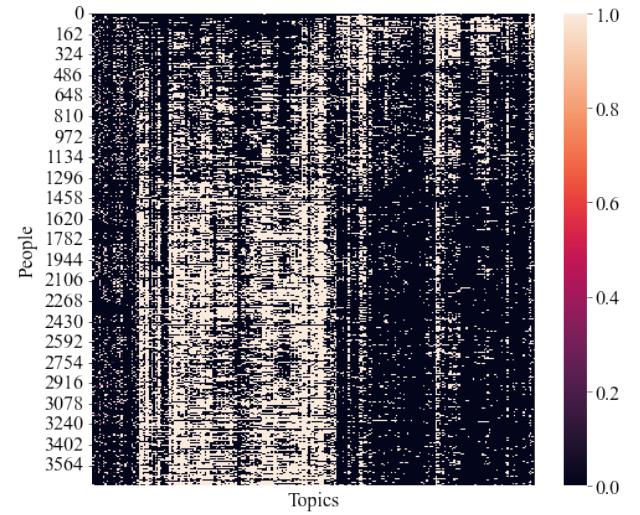


Fig. 6. Heatmap depicting structure discovered from PHNMF in City of Austin Satisfaction Survey dataset.

subgroups with 1,312 and 2,411 respondents, respectively. Group 1 primarily responds negatively to multiple choice questions on satisfaction with city services, while Group 2 primarily responds positively to the same questions.

At the second level, PHNMF divides Group 1 into two subgroups with 391 and 921 respondents, respectively. While Group 1a is distinguished by high levels of dissatisfaction, Group 1b reveals affirmative responses to feeling safe in Austin and satisfaction with the fire department and garbage collection services. We compare demographic statistics on the PHNMF subgroups and find that compared with the entire population, Group 1a reports lower average income, lower full-time employment rates, and lower percentages of owning their home compared to the entire population. Meanwhile, Group 1b reports comparable statistics to that of the entire population in these measures.

At the third level, PHNMF divides Group 1a into two groups with 145 and 246 respondents, respectively, and Group 1b into

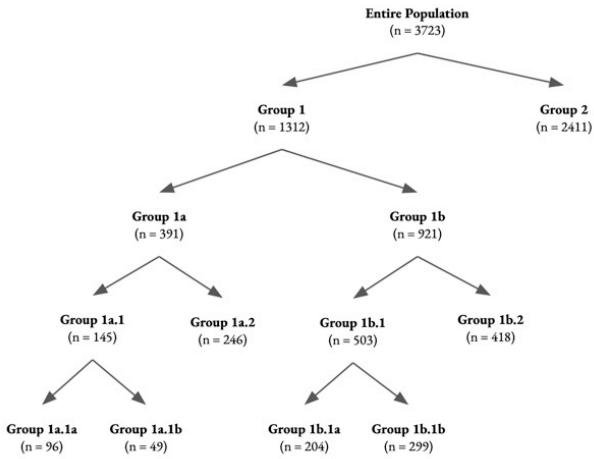


Fig. 7. Hierarchical population structure discovered with PHNMF in City of Austin Satisfaction Survey dataset.

two groups with 503 and 418 respondents, respectively. Group 1a.1 is distinguished by consistent dissatisfaction. Meanwhile, Group 1a.2 expresses dissatisfaction with traffic flow and city planning but also expresses satisfaction with Austin's emergency services, and affirms feeling safe in the city. Group 1b.1 is distinguished by satisfaction with city libraries and Austin as a place to live and feel welcome, and only express dissatisfaction with traffic flow. By contrast, Group 1b.2 has higher rates of dissatisfaction, but affirms feeling safe at home and in their neighborhood, and trusts emergency services.

At the fourth level, PHNMF splits Group 1a.1 into two groups with 96 and 51 respondents, respectively, and Group 1b.1 into two groups with 204 and 299 respondents, respectively. Group 1a.1a and Group 1a.1b consistently respond negatively but Group 1a.1a is more concerned with traffic and city planning. Meanwhile, Group 1a.1b primarily expresses dissatisfaction with the City of Austin's communications with its residents. By contrast, Group 1b.1a and Group 1b.1b mainly express satisfaction. Group 1b.1a is satisfied with the city library and learning centers, while Group 1b.1b is satisfied with Austin as a place to live and feel welcome, and report feeling safe in the city.

Since 27 of the 250 columns correspond to continuous variables, we convert these to categorical variables to employ LCA. However, LCA struggles to identify the best number of latent classes k in this dataset. When we enable k to range from 1 to 15, LCA reports insufficient degrees of freedom to select an optimal number of classes for $k \geq 13$. When we limit the number of classes to $k \leq 12$, however, LCA selects $k = 12$. Since LCA does not detail the importance of the variables in contributing to each class, the resulting 12 classes on the 250 columns are difficult to interpret. Since a small percentage of the variables in this dataset are continuous, we do not employ LPA.

2) *Facebook Climate Change Survey*: Fig. 8 depicts a heatmap of the population structure discovered with PHNMF

on the Facebook Climate Change Survey. The heatmap depicts non-trivial structure on the columns and the rows.

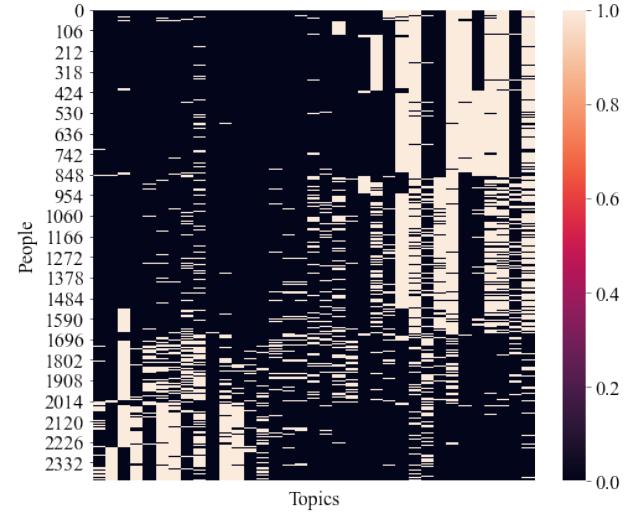


Fig. 8. Heatmap depicting structure discovered with PHNMF in Facebook Climate Survey dataset.

Fig. 9 depicts the hierarchical population structure discovered with PHNMF on the same dataset. The population splits are interpretable and also align with existing findings on climate change. At the first split, PHNMF divides the 2,427 survey respondents into two subgroups with 1,665 and 762 respondents, respectively. Group 1 primarily answers affirmatively to whether global warming exists and expresses strong support for funding renewable energy. Meanwhile, Group 2 primarily indicates that global warming is caused by natural environmental changes and is unworried about these changes. We compare demographic statistics on the PHNMF subgroups and find that Group 1 is approximately 60% female while Group 2 is approximately 60% male. This aligns with findings that women are more likely to affirm and express greater concern over climate change ([20], [21]). Additionally, about 60% of Group 1 is under the age of 45 while about 60% of Group 2 is over the age of 45. This aligns with findings that younger generations are more likely to support a shift to renewable energy [22]. Finally, 44% of Group 1 identifies as Democratic and 7% as Republican whereas 4% of Group 2 identifies as Democratic 40% as Republican. This aligns with findings that Democrats tend to express greater concern over climate change [23].

At the second split, PHNMF divides Group 1 into subgroups Group 1.1 and Group 1.2 with 853 and 812 respondents, respectively. The distinguishing factor between these subgroups is their level of concern for global warming. Group 1.1 expresses a great deal of worry while Group 1.2 expresses being only somewhat worried. The percentage of women in Group 1.1, Group 1.2, and the population at large is 64%, 56%, 50%, respectively. This aligns with findings of greater levels of climate change concern among women [24]. Additionally, Group 1.1 consists of 51% Democrats and 4%

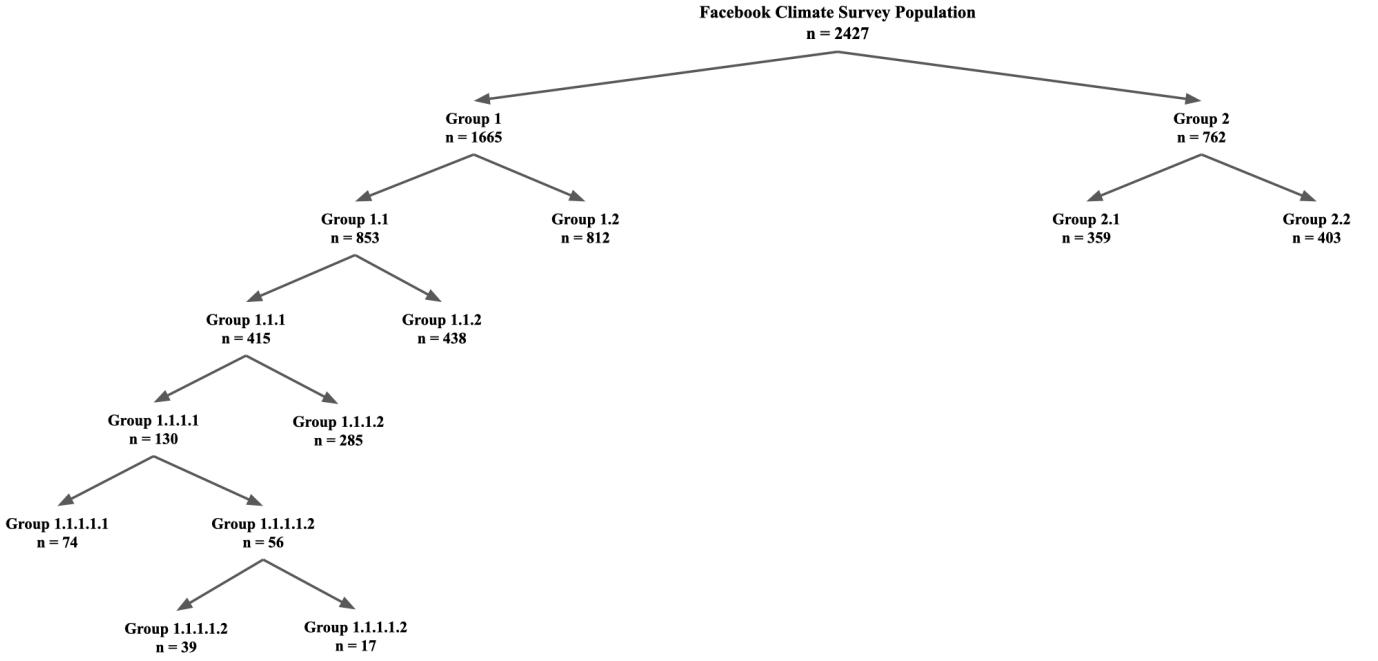


Fig. 9. Hierarchical population structure discovered with PHNMF in Facebook Climate Survey dataset.

Republicans while Group 1.2 consists of 37% Democrats and 10% Republicans. This aligns with findings that Democrats express greater concern over climate change [23]. At the second split, PHNMF also divides Group 2 into Groups 2.1 and 2.2 with 359 and 403 respondents, respectively. Group 2.1 acknowledges the existence of global warming to some extent while Group 2.2 denies global warming altogether.

The third split divides Group 1.1 into Group 1.1.1 and Group 1.1.2 with 415 and 438 respondents, respectively. Group 1.1.1 upholds that global warming is occurring and expresses concern, but only expresses moderate concern for the harm that global warming will cause them. Group 1.1.2 identifies that global warming will harm themselves, the planet, and future generations. Members of Group 1.1.2 are about 50% more likely to be under the age of 18 than those of Group 1.1.1. This aligns with findings that younger individuals are more likely to express concern that climate change will harm them [25]. Additional splits are likewise interpretable and appear to align similarly with existing findings on climate change views.

The corresponding LCA analysis finds 14 classes within the 35 variables on the first split. The large number of classes is difficult to interpret. When we enforce only two classes in the first LCA split, we obtain subgroups with 1,568 and 859 respondents, respectively. The distinguishing features between these two subgroups mirror the differences found between the PHNMF subgroups. However, when we run LCA again on these two subgroups, we obtain 6 and 7 subgroups for each group, respectively. This is again difficult to interpret. We do not employ LPA since it requires continuous variables.

Tab. III depicts the cosine similarity between the coefficient vectors obtained with ordinal ridge regression on the

discovered PHNMF subpopulations and those obtained from the same analysis on the entire population.

We observe greater alignment between coefficient vectors in subpopulations obtained from earlier PHNMF splits. However, the directions differ more as PHNMF gleans finer subpopulation structure with additional splits. While not depicted here, we also note that the most statistically significant regression variables identified by ordinal ridge regression agree with our interpretations of the PHNMF subpopulations. For example, for Group 1.1, the variable corresponding to “Yes, global warming is happening” is not significant since the entire group is distinguished by their agreement with this statement. However, within this subpopulation, the most significant variables relate to the extent that respondents are concerned about the impact of climate change on the future. This coincides with the distinguishing feature of the next PHNMF split.

TABLE III
COSINE SIMILARITY BETWEEN COEFFICIENT VECTORS OBTAINED FROM ORDINAL RIDGE REGRESSION ON THE SUBGROUPS COMPARED WITH ON THE POPULATION AS A WHOLE.

Group Name	Cosine Similarity
Group 1.2	0.9521
Group 1.1.2	0.8545
Group 1.1.1.2	0.7226
Group 1.1.1.1	0.7758
Group 1.1.1.1.2	0.7951
Group 1.1.1.1.2.2	0.7966
Group 2.1	0.9983
Group 2.2	0.9533

VI. DISCUSSION AND FUTURE WORKS

Our results on synthetic data demonstrate empirically that PHNMF can identify hierarchical population structure with high accuracy. They also demonstrate that the discovered subpopulation structure can improve downstream inference with regression methods. Our results on real data highlight PHNMF's ability to identify interpretable and meaningful hierarchical population structure from complex data arising from real surveys. Moreover, when interpreted along with additional demographic statistics, the PHNMF splits align substantially with existing literature.

By contrast, LPA does not produce functional results for any of the synthetic or real datasets. Additionally, the LCA splits on the real datasets are inconclusive and uninterpretable. This empirical evidence suggests that these traditional methods for identifying latent class structure may be limited in data scenarios with complex hierarchical population structure and larger numbers of variables.

Future work may incorporate the inferential procedures within the PHNMF algorithm. For example, rather than first performing PHNMF and then applying regression methods, one might weave methods such as semi-supervised NMF into one or more layers within the PHNMF algorithm. Incorporation of semi-supervised NMF and other methods might provide increased flexibility for subpopulation discovery.

REFERENCES

- [1] D. Haughton, P. Legrand, and S. Woolford, "Review of three latent class cluster analysis packages: Latent gold, polca, and mclust," *The American Statistician*, vol. 63, pp. 81–91, 02 2009.
- [2] D. A. Linzer and J. B. Lewis, "polca: An r package for polytomous variable latent class analysis," *Journal of Statistical Software*, vol. 42, no. 10, p. 1–29, 2011. [Online]. Available: <https://www.jstatsoft.org/index.php/jss/article/view/v042i10>
- [3] J. Peugh and X. Fan, "Modeling unobserved heterogeneity using latent profile analysis: A monte carlo simulation," *Structural Equation Modeling: A Multidisciplinary Journal*, vol. 20, no. 4, pp. 616–639, 2013.
- [4] D. Oberski, *Mixture Models: Latent Profile and Latent Class Analysis*. Springer Nature, 03 2016, pp. 275–287.
- [5] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
- [6] D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Advances in Neural Information Processing Systems*, T. Leen, T. Dietterich, and V. Tresp, Eds., vol. 13. MIT Press, 2000.
- [7] W.-S. Chen, J. Liu, B. Pan, and B. Chen, "Face recognition using nonnegative matrix factorization with fractional power inner product kernel," *Neurocomputing*, vol. 348, pp. 40–53, 2019, advances in Data Representation and Learning for Pattern Analysis. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0925231218312815>
- [8] Y. Qian, S. Jia, J. Zhou, and A. Robles-Kelly, "Hyperspectral unmixing via $l_{1/2}$ sparsity-constrained nonnegative matrix factorization," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 11, pp. 4282–4297, 2011.
- [9] N. Gillis, "The why and how of nonnegative matrix factorization," 2014. [Online]. Available: <https://arxiv.org/abs/1401.5226>
- [10] D. Kuang and H. Park, "Fast rank-2 nonnegative matrix factorization for hierarchical document clustering," in *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 739–747. [Online]. Available: <https://doi.org/10.1145/2487575.2487606>
- [11] R. Grotheer, Y. Huang, P. Li, E. Rebrova, D. Needell, L. Huang, A. Kryshchenko, X. Li, K. Ha, and O. Kryshchenko, "COVID-19 literature topic-based search via hierarchical NMF," *CoRR*, vol. abs/2009.09074, 2020. [Online]. Available: <https://arxiv.org/abs/2009.09074>
- [12] City of Austin Office of Performance Management, "City of austin community survey," 2019. [Online]. Available: <https://data.austintexas.gov/City-Government/Community-Survey/s2py-ceb7/data>
- [13] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [14] B. Zhang, M. Mildenberger, P. D. Howe, J. Marlon, S. A. Rosenthal, and A. Leiserowitz, "Quota sampling using facebook advertisements," *Political Science Research and Methods*, vol. 8, no. 3, p. 558–564, 2020.
- [15] C. Fraley and A. Raftery, "Mclust: Software for model-based cluster and discriminant analysis," *Department of Statistics, University of Washington: Technical Report*, vol. 342, 1998.
- [16] J. M. Rosenberg, P. N. Beymer, D. J. Anderson, C. Van Lissa, and J. A. Schmidt, "tidyipa: An R package to easily carry out latent profile analysis (ipa) using open-source or commercial software," *Journal of Open Source Software*, vol. 3, no. 30, p. 978, 2019.
- [17] D. A. Linzer and J. B. Lewis, "poLCA: An R package for polytomous variable latent class analysis," *Journal of Statistical Software*, vol. 42, no. 10, pp. 1–29, 2011. [Online]. Available: <https://www.jstatsoft.org/v42/i10>
- [18] A. A. Neath and J. E. Cavanaugh, "The bayesian information criterion: background, derivation, and applications," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 4, no. 2, pp. 199–203, 2012.
- [19] F. M. Zahid and S. Ramzan, "Ordinal ridge regression with categorical predictors," *Journal of Applied Statistics*, vol. 39, no. 1, pp. 161–171, 2012. [Online]. Available: <https://doi.org/10.1080/02664763.2011.578622>
- [20] A. M. McCright, "The effects of gender on climate change knowledge and concern in the american public," *Population and Environment*, vol. 32, no. 1, pp. 66–87, 2010. [Online]. Available: <http://www.jstor.org/stable/40984168>
- [21] A. M. McCright and R. E. Dunlap, "Cool dudes: The denial of climate change among conservative white males in the united states," *Global Environmental Change*, vol. 21, no. 4, pp. 1163–1172, 2011. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S095937801100104X>
- [22] C. Funk, "Key findings: How Americans' attitudes about climate change differ by generation, party and other factors," Jun 2021. [Online]. Available: <https://www.pewresearch.org/fact-tank/2021/05/26/key-findings-how-americans-attitudes-about-climate-change-differ-by-generation-party/>
- [23] B. Kennedy and C. Johnson, "More Americans see climate change as a priority, but democrats are much more concerned than republicans," Aug 2020. [Online]. Available: <https://www.pewresearch.org/fact-tank/2020/02/28/more-americans-see-climate-change-as-a-priority-but-democrats-are-much-more-concerned-than-republicans/>
- [24] M. Ballew, J. Marlon, A. Leiserowitz, and E. Maibach, "Gender differences in public understanding of climate change," Dec 2018. [Online]. Available: <https://climatecommunication.yale.edu/publications/gender-differences-in-public-understanding-of-climate-change/>
- [25] J. Bell, J. Poushter, M. Fagan, and C. Huang, "In response to climate change, citizens in advanced economies are willing to alter how they live and work," Mar 2022. [Online]. Available: <https://www.pewresearch.org/global/2021/09/14/in-response-to-climate-change-citizens-in-advanced-economies-are-willing-to-alter-how-they-live-and-work/>