# Wasserstein *t*-SNE

Fynn Bachmann[1,2]✉, Philipp Hennig[2], and Dmitry Kobak[2]

[1] University of Hamburg, Germany
fsvbach@gmail.com
[2] University of Tübingen, Germany
{philipp.hennig, dmitry.kobak}@uni-tuebingen.de

**Abstract.** Scientific datasets often have hierarchical structure: for example, in surveys, individual participants (samples) might be grouped at a higher level (units) such as their geographical region. In these settings, the interest is often in exploring the structure on the unit level rather than on the sample level. Units can be compared based on the distance between their means, however this ignores the within-unit distribution of samples. Here we develop an approach for exploratory analysis of hierarchical datasets using the Wasserstein distance metric that takes into account the shapes of within-unit distributions. We use *t*-SNE to construct 2D embeddings of the units, based on the matrix of pairwise Wasserstein distances between them. The distance matrix can be efficiently computed by approximating each unit with a Gaussian distribution, but we also provide a scalable method to compute exact Wasserstein distances. We use synthetic data to demonstrate the effectiveness of our *Wasserstein t-SNE*, and apply it to data from the 2017 German parliamentary election, considering polling stations as samples and voting districts as units. The resulting embedding uncovers meaningful structure in the data.

**Keywords:** Wasserstein metric · *t*-SNE · Dimensionality reduction · Election data · Hierarchical data · Optimal transport

## 1    Introduction

We consider dimensionality reduction for the purpose of data visualization, for the situation in which each 'data point' is a probability distribution, or a set of samples from it. This situation naturally arises when the data have *hierarchical structure*, i.e. the individual samples can be grouped at a higher level. Throughout this work we will use the word 'unit' to refer to this grouping level; for each 'unit' there is a number of 'samples' in the data (Figure 1). For example, in a social science survey, participants can be seen as samples and their countries of origin can be seen as units. For exploratory analysis, the interest may often be in the relationships between units (countries), rather than samples (participants).

A common approach for data exploration is to visualize the dataset as a 2D embedding, using dimensionality reduction algorithms such as PCA, MDS, *t*-SNE [14] or UMAP [15]. These algorithms are designed to get vectors as input, and compute pairwise distances (e.g. Euclidean) between the input vectors.
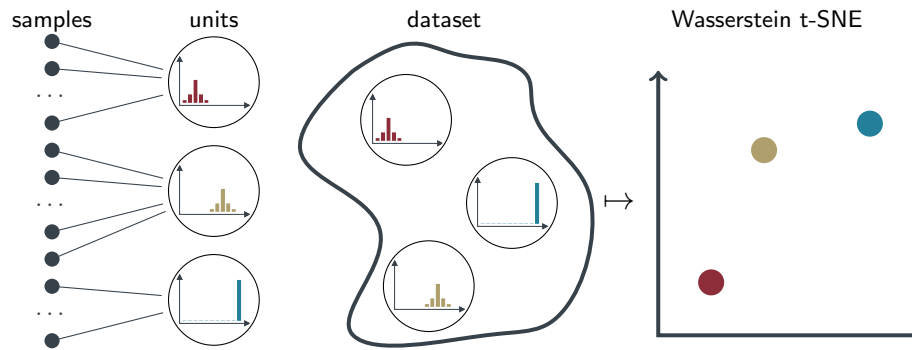
**Fig. 1.** Hierarchical data. Individual *samples* can be grouped into *units*. Each unit forms a probability distribution over its samples. In our *Wasserstein t-SNE* approach, units in the dataset are compared using the Wasserstein metric to construct a pairwise distance matrix, which is then embedded in two dimensions using the *t*-SNE algorithm. Units with similar probability distributions end up close together in the 2D embedding.

However, when analyzing units in a hierarchical dataset, each single unit forms an entire probability distribution over its samples, and cannot be represented by one vector. A simple approach would be to collapse all within-unit distributions to their means, and then apply any standard dimensionality reduction algorithm. However, this procedure can loose important information, particularly when some of the units share the same mean but have different shape.

Here we propose to use the Wasserstein metric [9] to compute pairwise distances between units. The Wasserstein distance has got recent attention in applications to Generative Adversarial Networks [1] or discriminant analysis [4] where it was used to compare probability densities with different support. The Wasserstein metric is convenient because there exists a closed-form solution for Gaussian distributions [2]. Using the Gaussian approximation, it is possible to efficiently construct the pairwise distance matrix between units in a hierarchical dataset. This distance matrix can then be used for downstream analysis, such as clustering or dimensionality reduction. Our focus here will be on *t*-SNE embeddings.

In the first part of this work we use simulated data to demonstrate the effectiveness of our *Wasserstein t-SNE*. In the second part we apply the same method to real-world data, in particular the data from the 2017 German parliamentary election. Here, samples correspond to polling stations while the units correspond to voting districts. We use the Gaussian approximation but also compute the exact Wasserstein distances, using an efficient linear programming approach.

The Python implementation of Wasserstein *t*-SNE is available on GitHub at fsvbach/WassersteinTSNE and as a package on PyPi WasserteinTSNE. The analysis code reproducing all figures in this paper can be found on GitHub at fsvbach/wassersteinTSNE-paper together with the analyzed data.

## 2   Methods

### 2.1   $t$-SNE

T-distributed stochastic neighbor embedding ($t$-SNE) [14] is a dimensionality reduction algorithm used in many scientific fields to find structure in datasets. The main idea of $t$-SNE is to arrange points in a low-dimensional space, such that their pairwise similarities (affinities) are similar to those in the high-dimensional space. In particular, if $P$ and $Q$ are the affinity matrices of the data and embedding respectively, $t$-SNE minimizes their Kullback-Leibler divergence

$$\mathrm{KL}(P\|Q) := \sum_{ij} P_{ij} \log \frac{P_{ij}}{Q_{ij}}.$$

The affinity matrix $P$ is constructed from the pairwise distances $d_{ij}$ by Gaussian kernels with bandwidth $\sigma_i$

$$P_{j|i} = \frac{\exp(-d_{ij}^2/2\sigma_i^2)}{\sum_{k\neq i}\exp(-d_{ik}^2/2\sigma_i^2)}$$

such that all perplexities of the conditional distributions equal some predefined value. In most $t$-SNE implementations this parameter defaults to 30 which we leave untouched in our experiments. As a reminder, if $p(x)$ is a discrete probability density function, the perplexity of $p$ is given by

$$\mathcal{P}(p) := 2^{H(p)} = \prod_x p(x)^{-p(x)}.$$

The affinity matrix $P$ is then symmetrized by

$$P_{ij} = \frac{P_{j|i} + P_{i|j}}{2n}.$$

In the low-dimensional space, the affinity matrix $Q$ is based on the pairwise distances between the embedding vectors $\mathbf{y}_i$, using the $t$-distribution kernel:

$$Q_{ij} = \frac{\left(1 + \|\mathbf{y}_i - \mathbf{y}_j\|^2\right)^{-1}}{\sum_{k\neq l}\left(1 + \|\mathbf{y}_k - \mathbf{y}_l\|^2\right)^{-1}}.$$

The $t$-SNE algorithm minimizes the Kullback-Leibler divergence $\mathrm{KL}(P\|Q)$ with respect to the coordinates $\mathbf{y}_i$. The embedding is initialized randomly, or using another algorithm such as PCA [11]. The optimization is done with gradient descent, i.e. the points move along the gradient until convergence. This results in a local minimum where no point can be moved without yielding a worse embedding.

When interpreting $t$-SNE embeddings it is important to keep in mind that the algorithm puts emphasis on close points, i.e., similar data points are embedded close to each other. The opposite does not hold: points that are embedded far from each other are not necessarily far from each other in the original space.

In this work we use the implementation of openTSNE [17], and keep all parameters at their default values.

## 2.2   Wasserstein metric

The Wasserstein metric [9] is a natural choice to compare probability distributions. It can be used to compare densities which do not have the same support, as long as a distance measure of their support is given. The downside of the Wasserstein distance is its computational complexity, which is linked to optimal transport [20,16].

**Definition 1** *Let (M,d) be a metric space. The p-Wasserstein distance of two distributions $\mu$ and $\nu$ is defined as*

$$W_p(\mu, \nu) := \left( \inf_{\gamma \in \Gamma(\mu,\nu)} \int_{M \times M} d(x,y)^p \mathrm{d}\gamma(x,y) \right)^{\frac{1}{p}}$$

*where $\Gamma$ is the set of all couplings of $\mu$ and $\nu$.*

In computer science the 1-Wasserstein metric is also known as *Earth Mover's Distance*, because if one imagines the probability distributions as piles of earth, then $W_p(\mu, \nu)$ represents the minimal amount of work necessary to transfer this mass from $\mu$ to $\nu$. This intuition also explains why the probability distributions must be defined on a metric space $M$, because we have to measure how far two points are away from each other, i.e. how far the mass has to be transported.

In general, the $p$-Wasserstein distance for continuous distributions is hard to compute [20]. But there exists a closed-form solution for the 2-Wasserstein metric for multivariate Gaussian distributions [2] (also known as Fréchet Inception Distance [8]). If $\mu, \nu$ are two Gaussian distributions $\mathcal{N}_i(m_i, C_i)$ with means $m_i$ and covariance matrices $C_i$, the 2-Wasserstein distance between them is given by

$$W_2(\mu, \nu)^2 = \|m_1 - m_2\|_2^2 + \mathrm{tr}\left( C_1 + C_2 - 2 \left( C_2^{1/2} C_1 C_2^{1/2} \right)^{1/2} \right)$$
$$= \|m_1 - m_2\|_2^2 + \mathrm{tr}\left( C_1 + C_2 - 2 \left( C_2 C_1 \right)^{1/2} \right).$$

The first term here is the Euclidean distance between the means, while the second term defines a metric on the space of covariance matrices [2]. By introducing a hyperparameter $\lambda \in [0, 1]$ we can put emphasis either on the means or on the covariances. We therefore propose a convex generalization of the 2-Wasserstein distance for Gaussians:

$$\tilde{W}(\mu, \nu)^2 = (1 - \lambda) \cdot \|m_1 - m_2\|_2^2 + \lambda \cdot \mathrm{tr}\left( C_1 + C_2 - 2 \left( C_2 C_1 \right)^{1/2} \right). \quad (\star)$$

This reduces to the Euclidean distance between the means for $\lambda = 0$ and to the distance between covariance matrices for $\lambda = 1$. The 2-Wasserstein distance corresponds to $\lambda = 0.5$ (up to a scaling factor).

In closed form, the 2-Wasserstein distance between two Gaussians can be computed in polynomial time. The matrix multiplication and the eigenvalue decomposition (for taking the square root) have $\mathcal{O}(d^3)$ complexity, where $d$ is the number of features. If there are $n$ units in the dataset, the $n \times n$ pairwise distance matrix can be computed in $\mathcal{O}(n^2 d^3)$ time.
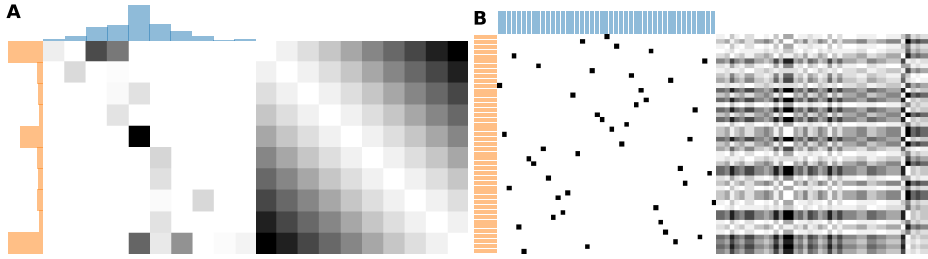
**Fig. 2.** Wasserstein distance as a linear program. (**A**) The optimal transport map $\gamma$ of two probability distributions $\nu$ (orange) and $\mu$ (blue) is shown. The heatmap represents the cost matrix $C$. (**B**) The same distributions can be visualized as a collection of samples, which have different support. The distance between samples $\nu_i$ and $\mu_j$ is given in the cost matrix entry $C_{ij}$. The size of the optimization variable $\gamma$ is then upper bounded by the product of the sample sizes.

### 2.3 Linear programming

We are also interested in computing exact Wasserstein distances without relying on the Gaussian approximation. Since real-life datasets contain discrete samples, this is possible by discretizing Definition 1 to

$$W_p(\mu, \nu)^p := \min_{\gamma \in \Gamma(\mu,\nu)} \sum_{M \times M} d(m_i, m_j)^p \cdot \gamma(m_i, m_j),$$

which is equivalent to the following linear program [16]:

$$
\begin{array}{c|c}
\textbf{primal form}: & \textbf{dual form}: \\
\begin{aligned}
\text{minimize} \quad z &= \mathbf{c}^T\mathbf{x}, \\
\text{so that} \quad \mathbf{Ax} &= \mathbf{b} \\
\text{and} \quad \mathbf{x} &\geq \mathbf{0}
\end{aligned}
&
\begin{aligned}
\text{maximize} \quad \tilde{z} &= \mathbf{b}^T\mathbf{y}, \\
\text{so that} \quad \mathbf{A}^T\mathbf{y} &\leq \mathbf{c}.
\end{aligned}
\end{array}
\quad (\star\star)
$$

The vectorized matrix $\mathbf{c}$ defines the transport cost, i.e., $c_{ij} = d(m_i, m_j)^p$ represents the $L_p$-distance of the points $m_i$ and $m_j$, where $M$ is the discrete metric space on which the probability distributions are defined. The optimization variable $\mathbf{x}$ represents the vectorized transport plan $\gamma$ as in Figure 2A. Each entry of $\mathbf{x}$ must be non-negative. The constraint $\mathbf{Ax} = \mathbf{b}$ is set up such that it is satisfied if the marginals of $\gamma$ equal the densities $\mu, \nu$. The primal form in ($\star\star$) yields an explicit transport plan while the dual form has less variables and is faster. Due to the strong duality of a linear program the resulting solution $z$ is the same. In practice we therefore use the dual form to compute exact Wasserstein distances.

Simplex algorithms and interior-point methods can solve real-world linear programs with a unique solution in polynomial time [7]. However, the exact complexity depends on the constraint matrix in the problem formulation. In general, the runtime of a linear program depends on the size of the optimization variable. In our case this is given by the product of the support sizes of the two discrete probability distributions. Figure 2A provides an example of two one-dimensional distributions, defined on the same support of size 10 (which could
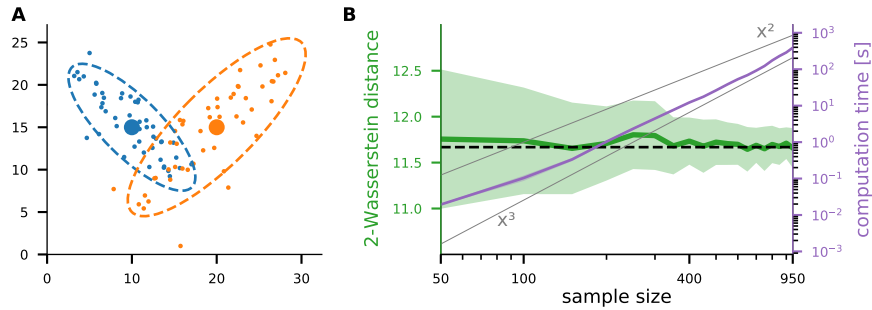
**Fig. 3.** Computation time and accuracy. (**A**) Two multivariate Gaussian distributions with 50 samples each. (**B**) The Wasserstein distance between the two probability distributions is computed using a different number of samples. The ground-truth distance is obtained by the closed-form solution and is shown with the dashed black line. The Wasserstein distance estimates using our linear program approach are shown in green (mean and standard deviation over 50 repetitions). The purple line shows the average runtime.

e.g. be a ten-point rating scale from 1 to 10). Both probability mass functions have 10 values so the resulting optimization variable $\gamma$ has $10{\times}10 = 100$ entries. While this linear program is easily solvable, the problem becomes computationally hard if we add additional feature dimensions (for example, if we add another ten-point feature, each probability density will become a two-dimensional mass function over 100 values, so then $\gamma$ has length 10,000). The number of variables in the transport map therefore grows exponentially with the number of features, thus this approach is intractable for datasets with many features.

Instead, we reduce the probability densities to the subspace where samples have actually been observed, rather than comparing distributions on the complete space $M$. That is, we consider both distributions uniformly distributed over their samples (Figure 2B). The marginal distributions $\mu$ and $\nu$ in Figure 2B therefore become uniform distributions with supports of size $n$ and $m$ respectively, where $n$ and $m$ are the two sample sizes. The size of the optimization variable $\gamma$ now becomes upper bounded by $nm$ regardless of the number of features. The cost matrix is given by the pairwise $L_p$-distance between all samples, which can be computed efficiently. We are not aware of a prior use of this shortcut, which however is only applicable when the number of samples is not large.

A way to see that both approaches are equivalent is to consider the constraints $\mathbf{x} \geq 0$ and $\mathbf{A}\mathbf{x} = \mathbf{b}$. When the number of features is large while the sample size is small, the sample density at most support values will have zero probability mass, because no sample has been observed at that point. Since the rows and the columns in the transport plan must sum to the marginal distributions, each entry with zero probability mass forces the corresponding row or column to be empty. Therefore these entries can be left out in the problem formulation and the size of the optimization variable is upper bounded by $nm$.

One consequence of this approach is that the samples no longer need to come from a discrete distribution, and indeed we can use the same approach to compute the exact Wasserstein distance between the samples coming from two Gaussian distributions (Figure 3). To demonstrate that, we chose a pair of two-dimensional Gaussian distributions with Wasserstein distance $d_W = 11.7$, where the Euclidean distance between the means is $d_E = 10.0$ and the distance between the covariances is $d_C = 6.0$ (Figure 3A). As the sample size grows from 50 to 1000, the solution of the linear program converges to the ground truth (Figure 3B, green line), while the runtime increases approximately as $\mathcal{O}(m^3)$ (purple line). However, for larger sample sizes the complexity will likely grow faster, as it is known that integer linear programs have exponential complexity [10]. Note that the dimensionality of the feature space (in this example, it is two-dimensional) does not strongly influence runtime.

## 2.4   Data

**Simulated data** To demonstrate and validate our method, we simulated hierarchical datasets, i.e. we defined the *hierarchical Gaussian mixture model* (HGMM). Similar to a Gaussian mixture model, a HGMM has multiple classes from which units are drawn. But here, each unit defines a Gaussian distribution with a unit-specific mean and covariance matrix. In each class, the unit means come from a class-specific Gaussian distribution, while the unit covariance matrices come from a class-specific Wishart distribution.

**Definition 2** *Let $\mathcal{N}$ and $\mathcal{W}$ denote Gaussian and Wishart distributions respectively. A hierarchical Gaussian mixture model is then defined by the number of classes ($K$), the number of units per class ($N_i$ for $i = 1 \ldots K$), the number of samples per unit ($M_j$ for $j = 1 \ldots \sum_{i=1}^{K} N_i$) and their feature dimensionality $F$, where*

- *each class i is characterized by a Gaussian distribution $\mathcal{N}(\mu_i, \Sigma_i)$ with $\mu_i \in \mathbb{R}^F$ and $\Sigma_i \in \mathbb{R}^{F \times F}$ and a Wishart distribution $\mathcal{W}(n_i, \Lambda_i)$ with $n_i \geq F$ and $\Lambda_i \in \mathbb{R}^{F \times F}$;*
- *each unit $X_j$ belonging to a class $i$ is characterized by a Gaussian distribution $\mathcal{N}(\nu_j, \Gamma_j)$. Unit means are samples from the class-specific Gaussian $\nu_j \sim \mathcal{N}(\mu_i, \Sigma_i)$ and unit covariance matrices are samples from the class-specific Wishart distribution $\Gamma_j \sim \mathcal{W}(n_i, \Lambda_i)$;*
- *the samples $S_k$ of each unit $j$ are iid distributed as $S_k \sim \mathcal{N}(\nu_j, \Gamma_j)$.*

A HGMM is specified by the set of class-specific parameters $\{\mu_i, \Sigma_i, n_i, \Lambda_i\}$. For example, Figure 4 shows a two-dimensional ($F = 2$) HGMM with $K = 4$ classes, $N = 100$ units in each class, and $M = 15$ samples in each unit. The class-specific Gaussian distributions (dashed black lines) are far away from each other because their means $\mu_i$ are chosen to be sufficiently different. The Within-class similarity of the unit means can be adjusted by $\Sigma_i$ (defining the shape of the dashed contours). Each class has its own Wishart scale matrix $\Lambda_i$; in this
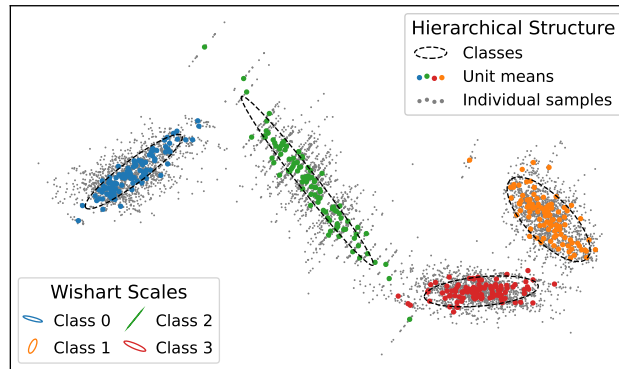
**Fig. 4.** Hierarchical Gaussian mixture model (HGMM). This two-dimensional example dataset has $K = 4$ classes with $N = 100$ units each. The gray points show the samples from all units ($M = 15$ samples per unit). Note that some of the units in the red and green classes have similar means, but their covariances are very different.

example, the peculiar Wishart scale of the green class makes green units easily distinguishable from the red units even when their means come close to the red class. Note, that the unit covariance matrices $\Gamma_j$ in any given class are not all the same. Their sampling process (from a Wishart distribution) is equivalent to drawing $n_i$ samples from a zero-centered Gaussian distribution with the Wishart scale as covariance matrix and estimating the sample covariance matrix. The larger the $n_i$, the closer all $\Gamma_j$ are to the respective $\Lambda_i$. We used $n_i = 4$ in Figure 4.

**German election data** The German parliamentary election was held in September 2017 with six major parties making it to the parliament. Germany is divided into 299 voting districts (*Wahlkreise*). In each voting district, multiple polling stations (*Wahlbezirke*) are set up. In our analysis we consider each voting district to be a unit with its polling stations being its samples.

The election data were directly downloaded from the *Bundeswahlleiter* website (https://tinyurl.com/mpevp355). We removed results of all minor parties and normalized each polling station so that the percentages of the six major parties — CDU (including the Bavarian-only CSU), SPD, AfD, FDP, Grüne and Linke — sum to 1 (the feature dimension of each sample is therefore six). Voting by mail was counted in separate mail-only polling stations of the respective voting district.

For this dataset, we a priori defined four classes: *Cities* (all voting districts with population density of at least 1000 people per square kilometer; population densities (https://tinyurl.com/3262nf8b) also obtained from the *Bundeswahlleiter* website), *Southern Germany* (all districts in Bavaria and Baden-Würtemberg, excluding previously defined cities), *Eastern Germany* (former DDR, excluding cities) and *Western Germany* (the rest).
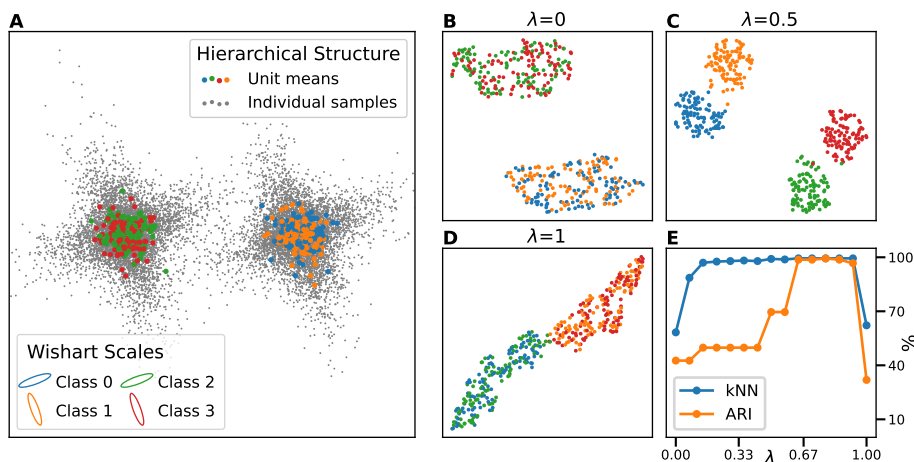
**Fig. 5.** Wasserstein *t*-SNE. (**A**) This two-dimensional ($F = 2$) HGMM was generated using $K = 4$ classes with $N = 100$ units each ($M = 30$ samples per unit). Two pairs of classes have the same distribution of unit means, while two other pairs of classes have the same distribution of unit covariance matrices. (**B**) The mean-based embedding ($\lambda = 0$) is not able to separate some of the classes. (**C**) The Wasserstein embedding ($\lambda = 0.5$) separates all four classes. (**D**) The covariance-based embedding ($\lambda = 1$) is not able to separate some of the classes. (**E**) The performance at different values of $\lambda$ was assessed using the kNN accuracy ($k = 5$) in the 2D embedding and the adjusted Rand index (ARI) obtained from Leiden clustering of the original distance matrix (kNN graph with $k = 5$, resolution parameter $\gamma = 0.08$).

## 3   Results

### 3.1   Wasserstein *t*-SNE on simulated data

To perform *Wasserstein t-SNE*, we first compute the pairwise distance matrix between units in a dataset where each unit is considered to be a probability distribution over its samples. We then embed these units in 2D using the *t*-SNE algorithm.

Figure 5A shows a two-dimensional ($F = 2$) toy dataset that consists of $K = 4$ classes, with $N = 100$ units per class, and $M = 30$ samples per unit. The red and green classes have the same distribution of unit means; the same is true for the blue and the orange classes. Likewise, the red and orange classes have the same distribution of unit covariances; the same is true for the blue and the green classes. For a more detailed description see Section 2.4.

Depending on the value of $\lambda$ in ($\star$), the resulting *t*-SNE embeddings show different structure. The mean-based embedding ($\lambda = 0$) only separates the dataset into two clusters (Figure 5B). Since it only takes the means into account, the orange class cannot be separated from the blue one, and the red class cannot be separated from the green one. Similarly, the covariance-based embedding ($\lambda = 1$) only finds two clusters as well, mixing up blue with green and orange with red
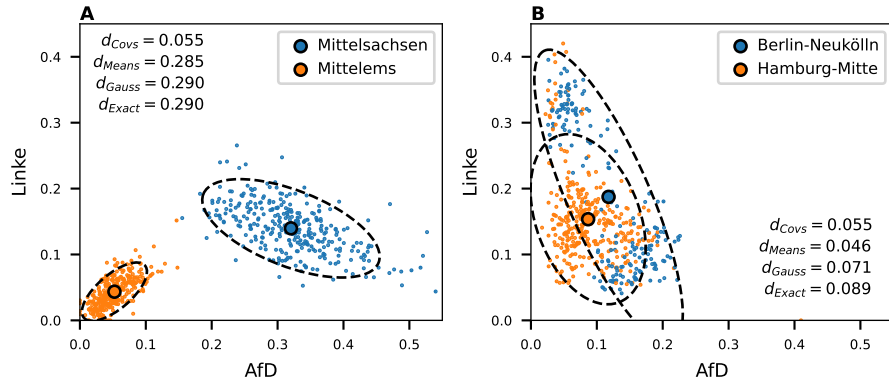
**Fig. 6.** Example voting districts in the 2017 German parliamentary election. Four voting districts (units) are shown together with their respective polling stations (samples). (**A**) Mittelems, which is located in the rural western Germany, exhibits positive correlation between the votes for AfD and for Linke, whereas Mittelsachsen in eastern Germany shows negative correlation. (**B**) The politically diverse district of Berlin-Neukölln has bimodal structure in the within-unit distribution of AfD and Linke votes, whereas Hamburg-Mitte does not show such bimodality.

(Figure 5D). In contrast, the Wasserstein embedding ($\lambda = 0.5$) successfully separates all four classes from each other (Figure 5C).

To measure the performance of the different $\lambda$ values, we used two different metrics. One metric is the k-nearest-neighbor (kNN) classification accuracy that measures the probability that a unit is labeled correctly by the majority vote of its $k = 5$ nearest neighbors in the embedding (we used the sklearn implementation [5]). The second metric is the adjusted Rand index (ARI) [18] that evaluates the agreement between the ground truth classes and the results of unsupervised clustering. We used the Leiden clustering algorithm [19], applied to the Wasserstein distance matrix (here and below we used the leidenalg implementation [19] with resolution parameter $\gamma = 0.08$ on the kNN graph built with $k = 5$). Note that unlike the kNN accuracy, the ARI metric is independent of $t$-SNE.

Both metrics, kNN accuracy and ARI, peaked at $\lambda \in [0.7, 0.8]$ and showed markedly worse performance at both $\lambda = 0$ and $\lambda = 1$. Moreover, while the kNN accuracy was close to the peak at already $\lambda = 0.5$, the ARI achieved higher values only for $\lambda > 0.7$. The result indicates that putting more emphasis on the covariance structure helps the algorithm to cluster the data correctly. This shows the power of our generalized Wasserstein distance for Gaussian distributions, as in this case it outperforms the exact Wasserstein distance (corresponding to $\lambda = 0.5$).

### 3.2 German parliamentary election 2017

**Gaussian Wasserstein $t$-SNE** The dataset of the 2017 German parliamentary election dataset contains 299 voting districts (units), each having about
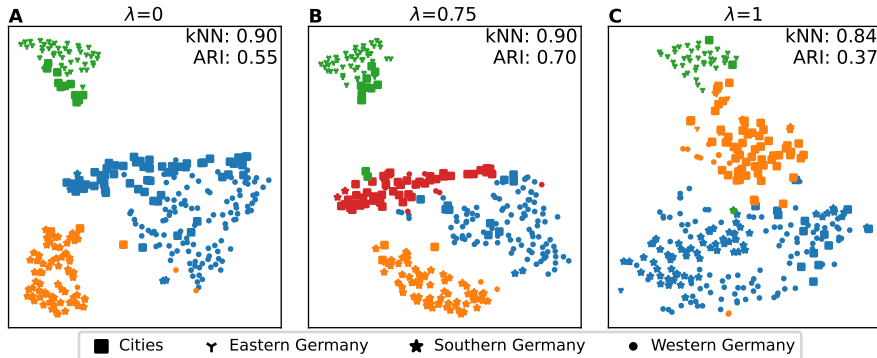
**Fig. 7.** Wasserstein $t$-SNE of the 2017 German parliamentary election. For clustering, we used the Leiden algorithm (on the original distance matrix) with $k = 5$ kNN graph and resolution parameter $\gamma = 0.08$. Colors correspond to the Leiden clusters; marker shape corresponds to the a priori classes. (**A**) The mean-based embedding with $\lambda = 0$ shows three clusters. (**B**) The Wasserstein embedding with $\lambda = 0.75$ shows four clusters. (**C**) The covariance-based embedding with $\lambda = 1$ shows three clusters.

150–850 polling stations (samples). The samples are represented as a points in six-dimensional space (corresponding to six political parties), as described in Section 2.4. For most units, the data could be reasonably well described by a multivariate Gaussian distribution (Figure 6). For example, the Gaussian Wasserstein distance and the exact Wasserstein distance between Mittelsachsen and Mittelems districts, both equaled 0.290 (Figure 6A). For some districts the approximation was less good: e.g. between Berlin-Neukölln and Hamburg-Mitte (Figure 6B), the Gaussian Wasserstein distance was 0.071 while the exact Wasserstein distance was 0.089. This can be explained by the polarization of Berlin-Neukölln, which had a bimodal structure that could not be captured by a multivariate normal distribution. However, we found that most districts were well approximated by a Gaussian.

We computed the pairwise Wasserstein distances between all pairs of units for different values of $\lambda$, and embedded the resulting distance matrices with the $t$-SNE algorithm (Figure 7). In parallel, we clustered each distance matrix by the Leiden algorithm (with resolution parameter $\gamma = 0.08$ and kNN graph with $k = 5$) and used the cluster assignments to color the $t$-SNE embeddings. The Wasserstein embedding with $\lambda = 0.75$ (Figure 7B) outperformed the mean-based ($\lambda = 0$ and the covariance-based ($\lambda = 1$) embeddings, and achieved a kNN accuracy of 0.90 and an ARI of 0.70. While the difference in the kNN accuracy was not very different for other values of $\lambda$ (0.90 for the mean-based and 0.82 for the covariance-based embeddings), the difference in the ARI was very pronounced (0.55 for the mean-based and 0.37 for the covariance-based embeddings).

The ARI is sensitive to the number of clusters, which is not pre-specified in the Leiden algorithm. While it automatically found three clusters with $\lambda = 0$
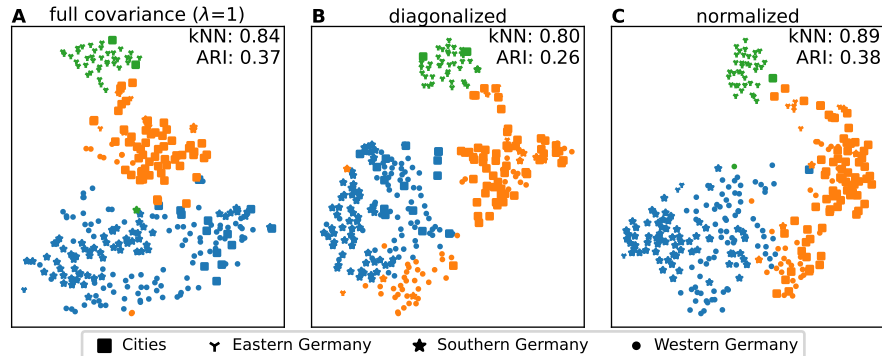
**Fig. 8.** Variants of the covariance-based ($\lambda = 1$) embedding. (**A**) The full covariance matrices were used to compute the pairwise distance matrix. The same embedding as in Figure 7C. (**B**) Only the diagonal entries of the covariance matrices were used, i.e. the marginal variances. (**C**) Each covariance matrix was normalized to become a correlation matrix.

(Figure 7A) and $\lambda = 1$ (Figure 7C), it found four clusters with $\lambda = 0.75$, and these clusters corresponded well to the four classes (Cities, Western Germany, Eastern Germany, Southern Germany) that we defined a priori. In contrast, the mean-based distance matrix merged Western Germany with the Cities, and the covariance-based embedding merged Western Germany with Southern Germany. This decreased the ARI for these embeddings, which is also visible in Figure 10C. While the exact ARI values depend on the choice of $k$ and $\gamma$ parameters, our results show that the Wasserstein distances with $0 < \lambda < 1$ can be an improvement compared to ignoring either the information about the means or about the covariance structure.

**Covariance and correlation structure** The covariance-based embedding clearly showed three clusters (Figure 7C), even though the means were completely ignored in this analysis. Where did this structure emerge from? Covariance is influenced by correlation and by marginal variances. To disentangle these two aspects of the data, we constructed a covariance-based embedding after zeroing out all off-diagonal values of all covariance matrices before computing the pairwise distance matrix (Figure 8B). We also constructed a covariance-based embedding after normalizing all covariance matrices to be correlation matrices (Figure 8C). Both embeddings were similar to the original covariance-based embedding (Figure 8A), suggesting that there was meaningful information in marginal variances as well as in pairwise correlations.

Figure 8C demonstrates that class information was present in the within-unit correlations, and indeed we saw earlier that parties' results can correlate differently in different voting districts (Figure 6). To visualize this effect, we overlayed pairwise correlation coefficients on the Gaussian Wasserstein embedding with $\lambda = 0.75$ (Figure 9A). For several pairs of parties, correlation strongly
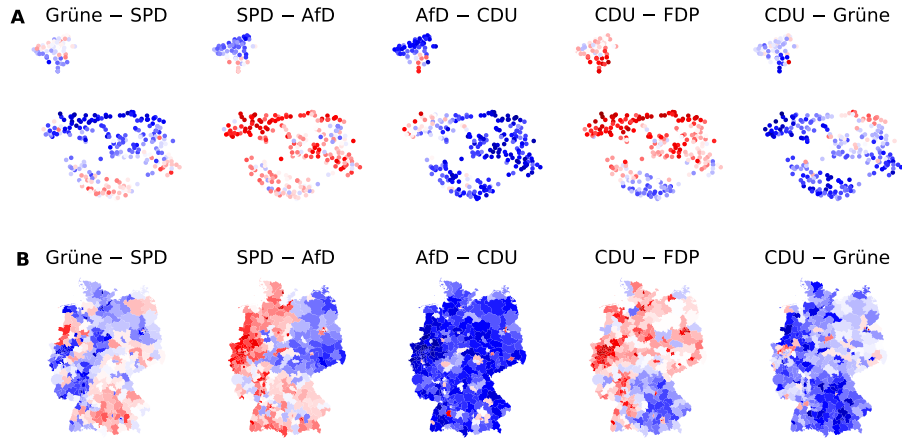
**Fig. 9.** Pairwise correlations between parties. Colors represent the Pearson correlation coefficient of the respective parties in each voting district, from blue ($-1$) to red ($1$). We chose five (out of 15) pairs of parties showing the most interesting patterns. (**A**) Party correlations overlayed over the Gaussian Wasserstein embedding with $\lambda = 0.75$ (as in Figure 7B). (**B**) Party correlations overlayed over the geographical map of Germany.

depended on the class, e.g. AfD and SPD were positively correlated in the Cities, but negatively correlated in Eastern Germany. This indicates that people who vote SPD in the cities tend to live in the same neighborhoods (i.e. same polling stations within a given district) as people who vote AfD, whereas in the rural east they tend to live in different neighborhoods. This suggests that these parties are perceived differently in different parts of the country, opposing each other in some of the regions but sharing sympathizers in others. Previous research has shown that many voters switched from SPD to AfD in the 2017 election [6]. Our analysis indicates that this effect may have happened mostly in the east.

**Exact Wasserstein Embedding** To verify that the Gaussian approximation used above did not strongly distort the embedding, we also did an embedding based on the exact Wasserstein distance. As explained in Section 2.3, we calculated the exact Wasserstein distances using linear programming. While this approach is more faithful to the data, it requires much longer computation time; calculating all $299 \cdot 298/2 = 44,551$ pairwise exact Wasserstein distances between units took 43 hours on a machine with 8 CPU cores at 3.0 GHz. The resulting embedding (Figure 10A) was very similar to the Gaussian embedding with $\lambda = 0.5$ (Figure 10B).

The Gaussian approximation has one important benefit beyond the faster runtime: namely, it allows to adjust the $\lambda$ parameter. As a function of $\lambda$, the kNN accuracy was nearly flat (Figure 10C) but the ARI of the Leiden clustering peaked for $\lambda > 0.6$, corresponding to the regime where the Leiden algorithm
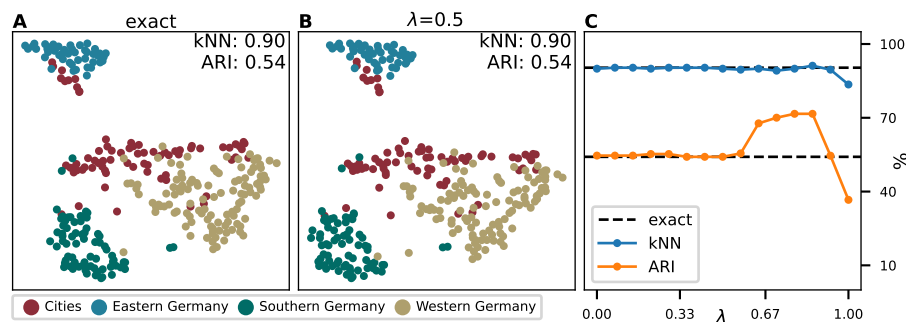
**Fig. 10.** Comparison of the Wasserstein $t$-SNE embeddings based on the Gaussian approximation and based on the exact Wasserstein distances. (**A**) The exact Wasserstein $t$-SNE embedding separates the classes. (**B**) The Gaussian Wasserstein embedding with $\lambda = 0.5$ shows similar structure. (**C**) The kNN classification accuracy ($k = 5$) and the ARI based on the Leiden clustering ($k = 5$ kNN graph, resolution parameter $\gamma = 0.08$) are shown for different values of $\lambda$. The black dashed lines show the kNN accuracy and the ARI of the exact Wasserstein embedding.

identifies four clusters in the data instead of three. At the same time, the performance of the exact Wasserstein embedding is similar to $\lambda = 0.5$ and hence is worse (Figure 10C, black dashed lines). This shows that, compared to the exact Wasserstein distances, the Gaussian approximation with the flexible $\lambda$ parameter can improve not only the runtime but also the final embedding.

## 4    Discussion

In this work we introduced *Wasserstein t-SNE* as a method to visualize hierarchical datasets. The main idea is to compute the pairwise distance matrix between the units of interest, using the Wasserstein metric to compare the distribution of their samples; then we use $t$-SNE to make a 2D embedding of the resulting distance matrix (Figure 1). Using a simulated and a real-life dataset, we showed that our approach can outperform standard $t$-SNE based on unit averages (Figures 5, 7).

There are two different ways to use Wasserstein $t$-SNE. One way is to approximate each unit by a multivariate Gaussian distribution which is fast but not scalable to high feature dimensionality. The second way is to compute the exact Wasserstein distances which is more accurate but substantially slower if there are many samples. Both ways require to compute and store the pairwise distance matrix which gets unfeasible for very large datasets. A benefit of using the Gaussian approximation is that it allows to put emphasis on either the means or the covariances of the units by adjusting the $\lambda$ parameter in the distance definition ($\star$). We suggested this definition as a novel generalization of the closed-form solution for the 2-Wasserstein distance between Gaussian distributions. To compute the exact Wasserstein distances, we solve a linear program for each distance

calculation. We developed an approach that scales with the number of samples and allows to compute Wasserstein distances between samples from continuous distributions. Using this approach, it took us ∼0.1 s to compute the Wasserstein distance between two Gaussian samples of size $n = 100$ on a standard desktop computer, and ∼100 s for $n = 1000$ (Figure 3). The feature dimensionality of the Gaussians does not play a role here.

Electoral data provide ample opportunities for statistical analysis, such as e.g. statistical fraud detection [12,13]. Here we used the publicly available data from the 2017 German parliamentary election to demonstrate that Wasserstein $t$-SNE can be useful for analysis of real-life datasets. Our method produced a 2D visualization ('map') of the 299 German voting districts (Figure 7B). This visualization exhibited four clusters, that were in good agreement with the known sociopolitical division of Germany that we defined a priori. Moreover, the correlation coefficient between political parties varied smoothly over the embedding (Figure 9), and in some cases even changed the sign. This showed that the information about within-unit distributions can be valuable to provide concise visualizations of political landscape in an unsupervised way.

We are not aware of other methods specifically designed to visualize hierarchical data. The naive approach is to collapse units to their means. However, we showed that this can be suboptimal whenever there is meaningful information in the unit covariances (e.g. Figure 7). An alternative is to append some of the covariance-based features to the unit means. This approach was, e.g., used in the Wisconsin Breast Cancer dataset [21] (popularized by its role as a UCI benchmark) where samples are cells and units are the respective patients. Here the variance of each feature (such as cell radius or cell smoothness) was appended to the dataset as an additional separate feature. While this allows to use some of the covariance information, it removes all information about feature correlation. Finally, it is possible to base all the analysis on the sample level, instead of the unit level. Such a 'sample-based' $t$-SNE embedding would show many more points than a 'unit-based' $t$-SNE embedding. However, for datasets like the one shown in Figure 5 this would not result in a useful visualization, as it would yield only two clusters and not four (as there are two pairs of classes with strongly overlapping distribution of samples within the units).

In summary, we believe that Wasserstein $t$-SNE is a promising method to visualize hierarchical datasets. Our results on synthetic data and on the 2017 German election data demonstrate that Wasserstein $t$-SNE can outperform standard alternatives and uncover meaningful structure in the data. We hope that our method can be useful in various domains. For example, social science often deals with hierarchical datasets, such as the European Values Study [3] where geographical regions can be seen as units while individual participants of the survey can be seen as samples. We believe that Wasserstein $t$-SNE can also be useful beyond the social and political science, e.g. for biomedical data.

## References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: International Conference on Machine Learning. pp. 214–223. PMLR (2017)
2. Dowson, D., Landau, B.: The Fréchet distance between multivariate normal distributions. Journal of Multivariate Analysis **12**(3), 450–455 (1982)
3. European Values Study: Integrated Dataset (EVS 2017). GESIS Data Archive, Cologne. ZA7500 Data file Version 4.0.0 (2020)
4. Flamary, R., Cuturi, M., Courty, N., Rakotomamonjy, A.: Wasserstein discriminant analysis. Machine Learning **107**(12), 1923–1945 (2018)
5. Goldberger, J., Hinton, G.E., Roweis, S., Salakhutdinov, R.R.: Neighbourhood components analysis. Advances in Neural Information Processing Systems **17** (2004)
6. Görtz, A.: Sozialstrukturelle Ursachen für Wechselwahl: Warum wechseln WählerInnen von der SPD zur AfD? Social Science Open Access Repository (2020)
7. Huangfu, Q., Hall, J.J.: Parallelizing the dual revised simplex method. Mathematical Programming Computation **10**(1), 119–142 (2018)
8. Jung, S., Keuper, M.: Internalized Biases in Fréchet Inception Distance. In: NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications (2021)
9. Kantorovich, L.V.: The mathematical method of production planning and organization. Management Science **6**(4), 363–422 (1939)
10. Karp, R.M.: Reducibility among combinatorial problems. In: Complexity of Computer Computations, pp. 85–103. Springer (1972)
11. Kobak, D., Linderman, G.C.: Initialization is critical for preserving global data structure in both t-SNE and UMAP. Nature Biotechnology **39**(2), 156–157 (2021)
12. Kobak, D., Shpilkin, S., Pshenichnikov, M.S.: Integer percentages as electoral falsification fingerprints. The Annals of Applied Statistics **10**(1), 54 – 73 (2016)
13. Kobak, D., Shpilkin, S., Pshenichnikov, M.S.: Statistical fingerprints of electoral fraud? Significance **13**(4), 20–23 (2016)
14. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. Journal of Machine Learning Research **9**(11) (2008)
15. McInnes, L., John Healy, J.M.: UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. arXiv (2020)
16. Mémoli, F.: Gromov–Wasserstein distances and the metric approach to object matching. Foundations of Computational Mathematics **11**(4), 417–487 (2011)
17. Poličar, P.G., Stražar, M., Zupan, B.: openTSNE: a modular Python library for t-SNE dimensionality reduction and embedding. BioRxiv (2019)
18. Rand, W.M.: Objective criteria for the evaluation of clustering methods. Journal of the American Statistical Association **66**(336), 846–850 (1971)
19. Traag, V.A., Waltman, L., Van Eck, N.J.: From Louvain to Leiden: guaranteeing well-connected communities. Scientific Reports **9**(1), 1–12 (2019)
20. Villani, C.: The Wasserstein distances, pp. 93–111. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
21. Wolberg, W.: Breast Cancer Wisconsin (Original). UCI Machine Learning Repository (1992)