



HAL
open science

Sparse and geometry-aware generalisation of the mutual information for joint discriminative clustering and feature selection

Louis Ohl, Pierre-Alexandre Mattei, Charles Bouveyron, Mickaël Leclercq, Arnaud Droit, Frédéric Precioso

► To cite this version:

Louis Ohl, Pierre-Alexandre Mattei, Charles Bouveyron, Mickaël Leclercq, Arnaud Droit, et al.. Sparse and geometry-aware generalisation of the mutual information for joint discriminative clustering and feature selection. *Statistics and Computing*, 2024, 34 (5), pp.155. 10.1007/s11222-024-10467-9 . hal-04755942

HAL Id: hal-04755942

<https://hal.science/hal-04755942v1>

Submitted on 20 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Sparse and geometry-aware generalisation of the mutual information for joint discriminative clustering and feature selection

Louis Ohl^{1,2} · Pierre-Alexandre Mattei¹ · Charles Bouveyron¹ · Mickaël Leclercq² · Arnaud Droit² · Frédéric Precioso¹

Received: 21 December 2023 / Accepted: 9 July 2024 / Published online: 17 July 2024
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

Feature selection in clustering is a hard task which involves simultaneously the discovery of relevant clusters as well as relevant variables with respect to these clusters. While feature selection algorithms are often model-based through optimised model selection or strong assumptions on the data distribution, we introduce a discriminative clustering model trying to maximise a geometry-aware generalisation of the mutual information called GEMINI with a simple ℓ_1 penalty: the Sparse GEMINI. This algorithm avoids the burden of combinatorial feature subset exploration and is easily scalable to high-dimensional data and large amounts of samples while only designing a discriminative clustering model. We demonstrate the performances of Sparse GEMINI on synthetic datasets and large-scale datasets. Our results show that Sparse GEMINI is a competitive algorithm and has the ability to select relevant subsets of variables with respect to the clustering without using relevance criteria or prior hypotheses.

Keywords Discriminative clustering · neural networks · unsupervised learning · feature selection · group-lasso penalty

1 Introduction

It is common that clustering algorithms and supervised models rely on all available features for the best performance. However, as data sets become high-dimensional, clustering algorithms tend to break under the curse of dimensionality (Bouveyron and Brunet-Saumard 2014b), for instance in biological micro-array data where the number of variables outweighs the number of samples (McLachlan et al. 2002). To alleviate this burden, feature selection is a method of choice. Indeed, all features may not always be of interest: some variables can be perceived as relevant or not with respect to the clustering objective. Relevant variables bring information that is useful for the clustering operation, while irrelevant variables do not bring any new knowledge regarding the clus-

ter distribution (Tadesse et al. 2005) and redundant variables look relevant but do not bring beneficial knowledge (Maugis et al. 2009). The challenge of selecting the relevant variables often comes with the burden of combinatorial search in the variable space. Therefore, solutions may be hardly scalable to high-dimensional data (Raftery and Dean 2006) or to the number of samples (Witten and Tibshirani 2010) when the selection process is part of the model. Therefore reducing the number of variables for learning to a relevant few is of interest, notably in terms of interpretation (Fop and Murphy 2018). The necessity of variable selection notably met successful applications in genomics (Marbac et al. 2020), multi-omics (Meng et al. 2016; Ramazzotti et al. 2018; Shen et al. 2012).

Often, integrating the selection process as part of the model will lead to either not scaling well (Solorio-Fernández et al. 2020) in terms of number of features (Raftery and Dean 2006) or number of samples (Witten and Tibshirani 2010) or imposing too constrained decision boundaries due to the nature of strong parametric assumptions. To alleviate both problems, we present the Sparse GEMINI: a model that combines the LassoNet architecture (Lemhadri et al. 2021) and the discriminative clustering objective GEMINI (Ohl et al. 2022, 2023) for a scalable discriminative clustering

✉ Louis Ohl
louis.ohl@inria.fr

✉ Pierre-Alexandre Mattei
pierre-alexandre.mattei@inria.fr

¹ Inria, CNRS, Maasai Team, Université Côte d'Azur, Nice, France

² CHU de Québec Research Centre, Laval University, Quebec City, Canada

with penalised feature selection. The contributions of Sparse GEMINI are:

- A simple novel scalable algorithm efficiently combining feature selection and discriminative clustering compatible with models ranging from logistic regression to deep neural networks.
- Demonstrations of performances on multiple synthetic and real datasets including a large-scale transcriptomics dataset.
- An extension of the work on the generalised mutual information by proposing a package containing the methods of previous work and the Sparse GEMINI model thanks to explicit computations of GEMINI gradients.

2 Related works

Feature selection algorithms can be divided into 2 distinct categories (John et al. 1994; Dy 2007): filter methods and wrapper methods. Filter methods apply in an independent step feature selection using a relevance criterion to eliminate irrelevant features before performing clustering. This can be done, for example, using information theory (Cover 1999) with the SVD-Entropy (Varshavsky et al. 2006) or spectral analysis (von Luxburg 2007; He et al. 2005; Zhao and Liu 2007). Those methods are thus easily scalable and quick despite the challenge of defining unsupervised feature relevance (Dy 2007). Wrapper methods encompass the selection process within the model and exploit their clustering results to guide the feature selection (Solorio-Fernández et al. 2020). Other related works sometimes refer to a third category named hybrid model (Alelyani et al. 2018) or embedded models (Blum and Langley 1997) as a compromise between the first two categories.

Although the definition of relevance of a variable is more straightforward for supervised learning, its definition in unsupervised learning clearly impacts the choice of selection criterion for filter methods or distribution design in model-based methods (Fop and Murphy 2018). Often, the terms relevant variables, irrelevant variables (Tadesse et al. 2005) for the notion of conveying information are used. Others may consider redundant variables as those that bring information already available (Maugis et al. 2009). Another key difference in models would be to consider whether the informative variables are independent given the cluster membership (local independence) or dependent (global independence from the non-informative variables), although the latter hardly accounts for redundant variables (Fop and Murphy 2018).

Feature selection should not be mistaken with dimensionality reduction, sometimes called feature reduction, which is the process of finding a latent space of lower dimension

leveraging good manifolds for clustering, e.g. using matrix factorisation (Shen et al. 2012). In this sense, methods seeking a sparse subspace for spectral clustering (Peng et al. 2016) or for KMeans clustering through PCA (Long et al. 2021) are discriminative. However, the nature of such projection forces the clustering to be done according to linear boundaries due to the projections. Still, by enforcing the projection matrix to be sparse, feature selection can be recovered in the original space (Bouveyron and Brunet-Saumard 2014a). Similarly, subspace clustering seeks to find clusters in different subspaces of the data (Zografos et al. 2013; Chen et al. 2018) and is thus an extension of feature selection (Parsons et al. 2004), particularly with the motivation that several latent variables could explain the heterogeneity of the data (Vandewalle 2020). However, such problems usually incorporate a mechanism to merge clusters, which is challenging as well, while we are interested in a method that selects features while producing a single clustering output.

Finally, clustering models in feature selection are often model-based (Scrucca and Raftery 2018; Raftery and Dean 2006; Maugis et al. 2009), which implies that they assume a parametric mixture model that can explain the distribution of the data, including the distribution of irrelevant variables. To perform well, these methods need a good selection criterion to compare models with one another (Raftery and Dean 2006; Marbac et al. 2020; Maugis et al. 2009). To the best of our knowledge, there do not exist models for joint feature selection and clustering in the discriminative sense of Minka (2005) and Krause et al. (2010), i.e., models that only design $p_{\theta}(y|\mathbf{x})$ with end-to-end training. Finally, most of these generative wrapper methods hardly scale in both sample quantity and/or variable quantity.

3 The sparse GEMINI

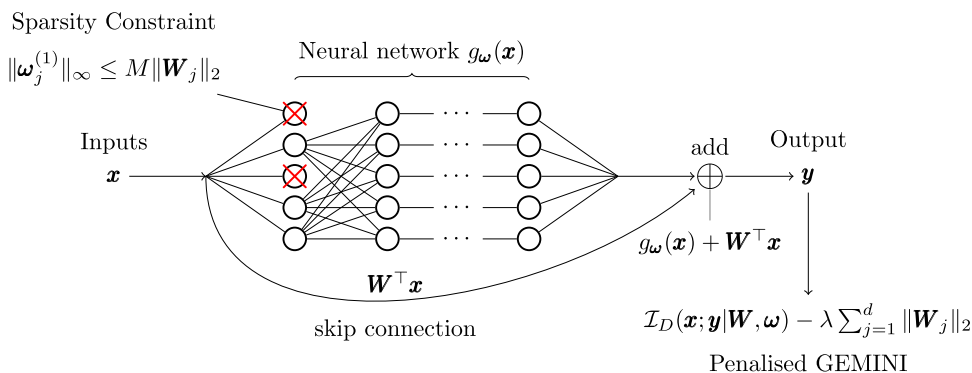
Sparse GEMINI is a combination of the generalised mutual information objective for discriminative clustering (Ohl et al. 2023) with the LassoNet framework for feature selection (Lemhadri et al. 2021) in neural networks, including the sparse logistic regression as a specific case. The model is summarised in Fig. 1.

3.1 The GEMINI objective

Let $\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^N \subset \mathcal{X}$ a dataset of N observations of dimension d . We note each feature $\mathbf{x}_{ij} \in \mathcal{X}_j$, thus: $\mathcal{X} = \prod_{j=1}^d \mathcal{X}_j$. We seek to cluster this dataset by learning a distribution $p_{\theta}(y|\mathbf{x})$ where y is a discrete variable taking K values. This distribution is defined by a softmax-ended function which ensures that the elements of the resulting vector add up to 1:

$$y|\mathbf{x} \sim \text{Categorical}(\text{SoftMax} \circ f_{\theta}(\mathbf{x})), \quad (1)$$

Fig. 1 Description of the complete Sparse GEMINI model. Through a proximal gradient, clusters learned by GEMINI drop irrelevant features both in a skip connection and an MLP. Setting $M = 0$ recovers a sparse unsupervised logistic regression



where $f_\theta : \mathcal{X} \mapsto \mathbb{R}^K$ has parameters θ . For example, setting f to an affine function recovers the multiclass logistic regression. In order to perform clustering with f as a discriminative distribution, we train the parameters θ using a generalised mutual information (GEMINI, Ohl et al., 2023). This objective was introduced to circumvent the need for parametric assumptions regarding $p(\mathbf{x})$ in clustering and thus leads to designing only a discriminative clustering model $p_\theta(y|\mathbf{x})$ (Minka 2005). With the help of Bayes theorem, this objective can be estimated without assumptions of the data distribution $p(\mathbf{x})$ using only the output of the clustering distribution $p_\theta(y|\mathbf{x})$. Note that, as we avoid parametric hypotheses on the data distribution, we do not have a relationship between the data and the parameters, hence the writing $p(\mathbf{x})$ without θ . Despite the absence of assumptions, we are able to sample from $p(\mathbf{x})$ owing to the dataset we have at hand. Overall, the GEMINI aims at separating according to a distance D the cluster distributions from either the data distribution (one-vs-all):

$$\mathcal{I}_D^{\text{ova}}(\mathbf{x}; y|\theta) = \mathbb{E}_{y \sim p_\theta(y)} [D(p_\theta(\mathbf{x}|y) \| p(\mathbf{x}))], \tag{2}$$

or other cluster distributions (one-vs-one):

$$\mathcal{I}_D^{\text{ovo}}(\mathbf{x}; y|\theta) = \mathbb{E}_{y_1, y_2 \sim p_\theta(y)} [D(p_\theta(\mathbf{x}|y_1) \| p(\mathbf{x}|y_2))]. \tag{3}$$

The novelty of GEMINI is to consider different types of distances D between distributions with a special focus on the maximum mean discrepancy (MMD, Gretton et al., 2012) or the Wasserstein distance (Peyré and Cuturi 2019) compared to former discriminative approaches using the standard mutual information (Bridle et al. 1992; Krause et al. 2010) which require regularisation to learn. The former corresponds to the distance between the expectations of the respective distributions projected into a Hilbert space and the latter is an optimal transport distance describing the minimum of energy necessary to reshape one distribution as the other. Both of them incorporate geometrical information on the data respectively through a kernel κ or a distance δ in the data space. Any neural network that is trainable through cross-entropy

loss can be switched to unsupervised learning at the cost of choosing a metric or kernel in the data space. While we cannot compute the true GEMINI values, we can estimate them using Bayes theorem to get an expression depending only on the predictions of the model. For instance, the Wasserstein GEMINI can be estimated using importance weights to estimate the densities of each cluster distribution. The MMD can be computed with sums of kernel terms weighted by the predictions of the model (Ohl et al. 2023, Table 1). GEMINI can therefore train any discriminative model of the form $p_\theta(y|\mathbf{x})$ as long as the distance D can be evaluated using only the model’s outputs. Thus, the written GEMINI objective becomes for a batch of size N :

$$\hat{\mathcal{I}}_D = \sum_{k=1}^K f_D(\{p_\theta(y = k|\mathbf{x} = \mathbf{x}_1), \dots, p_\theta(y = k|\mathbf{x} = \mathbf{x}_N)\}), \tag{4}$$

where f_D is a function depending on the chosen distance. For instance, if D is the KL distance, we can show that:

$$f_D(\{p_\theta(y = k|\mathbf{x} = \mathbf{x}_1), \dots, p_\theta(y = k|\mathbf{x} = \mathbf{x}_N)\}) = \sum_{i=1}^N p_\theta(y = k|\mathbf{x} = \mathbf{x}_i) \log \frac{p_\theta(y = k|\mathbf{x} = \mathbf{x}_i)}{p_\theta(y = k)}. \tag{5}$$

All GEMINIs from Ohl et al. (2022) are available in the GemClus package that we will present in Sect. 4.3.

3.2 Sparse models

3.2.1 Unsupervised logistic regression architecture

We start with the simplest discriminative model for variable selection: logistic regression. This corresponds to the case where our distribution $p_\theta(y|\mathbf{x})$ is characterised by the set of linear functions:

$$\mathcal{F} = \{f_\theta : \mathbf{x} \mapsto \theta^\top \mathbf{x}\}, \tag{6}$$

with $\theta \in \mathbb{R}^{d \times K}$ for d features and K clusters. Notice the absence of bias as this linear model is a sub-case of the neural network model covered in the next section. To properly ensure that vector weights are eliminated at once, a group-lasso penalty is preferred (Hastie et al. 2015, Section 4.3) also known as ℓ_1/ℓ_2 penalty (Bach et al. 2012). We consider a user-defined partition of the input features into $G \leq d$ groups, each with associated parameter subset θ_j . Note that the dimensions of θ_j vary depending on the number of features within the j -th group. For example, a categorical variable taking M values transformed into a one-hot-encoded vector of dimension M can be associated to a single group. Note that all clusters use the same subset of selected groups of variables. Thus, the optimal parameters should satisfy:

$$\hat{\theta} = \arg \max_{\theta} \mathcal{I}_D(\mathbf{x}; y|\theta) - \lambda \sum_{j=1}^G \|\theta_j\|_2. \tag{7}$$

This is exactly the same objective formulation as the supervised multi-class Lasso if we replace the GEMINI by the maximum likelihood or any other supervised loss. Notice that λ is positive because we seek to simultaneously maximise the GEMINI and minimise the ℓ_1/ℓ_2 penalty. During training, the sparse linear parameter will progressively remove variables by setting all grouped parameters to 0. If we set $G = d$, then each variable can be removed on its own as there are no groups of variables. A similar objective without group-lasso and using the standard mutual information can be found in (Kong et al. 2015, Eq. 4), although specific initialisation strategies were required to circumvent the unspecificity of mutual information local maxima, as described by Ohl et al. (2023).

3.2.2 The LassoNet architecture

We extend this procedure to neural network by adapting the LassoNet (Lemhadri et al. 2021) framework with GEMINIs. The neural network $f_{\theta} : \mathcal{X} \mapsto \mathbb{R}^K$ is taken from a family of architectures \mathcal{F} consisting of one multi-layered perceptron (MLP) and a linear skip connection:

$$\mathcal{F} = \{f_{\theta} : \mathbf{x} \mapsto g_{\omega}(\mathbf{x}) + \mathbf{W}^T \mathbf{x}\}, \tag{8}$$

with $\theta = \{\omega, \mathbf{W}\}$ including ω the parameters of the MLP g_{ω} and $\mathbf{W} \in \mathbb{R}^{d \times K}$ the weights of a linear skip connection penalised by group-lasso. This leads to the same optimisation objective as previously with a focus on the skip connection parameters:

$$\hat{\theta} = \arg \max_{\theta} \mathcal{I}_D(\mathbf{x}; y|\theta) - \lambda \sum_{j=1}^G \|\mathbf{W}_j\|_2, \tag{9}$$

with \mathbf{W}_j , the weights of the j -th group of features from \mathbf{W} . It is a matrix of dimension K times the group size. As the sparse skip connection \mathbf{W} loses some feature subset, we must force the MLP to drop this same subset of features as well. Therefore the weights of the first layer $\omega^{(1)}$ are constrained such that:

$$\|\omega_j^{(1)}\|_{\infty} \leq M \|\mathbf{W}_j\|_2, \forall j \leq G. \tag{10}$$

where M is called the hierarchy coefficient. Thus, when a feature j is eliminated, all weights starting from this feature in the MLP will be equal to 0 as well. When $M = 0$, the method is equivalent to the penalised logistic regression from the previous section because all entry weights of the MLP are equal to zero, hence passing no information.

Interestingly, while the constraints are designed to specifically select features, dimension reduction can be performed as well by extracting representations from lower-dimension layers in the network g_{ω} . However, this intermediate representation would not be complete as it misses the information from the skip connection.

4 Optimisation

4.1 Training and model selection

We follow Lemhadri et al. (2021) in proposing a *dense-to-sparse* training strategy for the penalty coefficient. Training is carried along a path where the ℓ_1 penalty parameter λ is geometrically increased: $\lambda = \lambda_0 \rho^t$ ($\rho > 1$) at time step t after an initial step without ℓ_1 penalty. We stop when the number of remaining features used by the model is below a user-defined threshold $0 < F_{\text{thres}} < d$ which can be thought of as the minimum number of useful variables required. Each time the number of features decreases during training, we save its associate intermediate model. It is thus possible to restore any intermediate model to get a clustering on different subsets of variables.

Once the training is finished, we look again at all GEMINI scores during the feature decrease and select the model with the minimum of features that managed to remain in the arbitrary range of 90% of the best GEMINI value. This best value is most of the time the loss evaluated with the model exploiting all features. Still, automatic model selection may not be necessary and in this case, we can look at all intermediate models produced during the path as a set of possible solutions for clustering with different number of selected variables.

As GEMINI depends on a metric defined in the data space, the metric will still be computed using all features despite selection by the model. Consequently, we propose a less grounded yet efficient training mode in appendix A taking into account variable selection in the metric computation.

However, this training mode yields incomparable GEMINI scores due to the change of metric definition, rendering the selection strategy inapplicable.

Interestingly, as the MMD-GEMINI maximisation is equivalent to a kernel KMeans objective, as shown by França et al. (2020), a subset of the method resembles to a sparse kernel KMeans. However, unlike related works (França et al. 2020; Witten and Tibshirani 2010), the selection process is done through gradient descent, i.e. without indicative variables, and without the definition of explicit centroids, thus being less strict regarding the number of clusters to find, i.e. GEMINI models can find fewer clusters than asked.

4.2 Gradient considerations

We offer here more detailed insights on the gradients of the model for ensuring true variable elimination and how the GEMINI behaves.

4.2.1 Proximal gradients

To ensure the convergence of the parameters to 0 upon elimination, we adopt a proximal gradient strategy. In the case of sparse logistic regression, the gradient ascent hence follows the two classical steps:

$$\beta = \theta^t + \eta^t \nabla_{\theta} \mathcal{I}_D(\mathbf{x}; y|\theta), \tag{11}$$

$$\theta_j^{t+1} = \mathcal{S}_{\rho^t \lambda, \|\beta_j\|_2}(\beta_j), \forall j \leq d, \tag{12}$$

where η^t is the learning at timestep t . The soft-thresholding operation \mathcal{S}_{α} :

$$\mathcal{S}_{\alpha}(x) = \text{sign}(x) \max\{0, |x| - \alpha\}, \tag{13}$$

is the closed-form solution of the proximal operator to project the parameters on the constrained space due to the group-lasso penalty (Hastie et al. 2015, Section 5.3.3). Consequently, we are sure to obtain true zeroes in the linear weights of the logistic regression or the weights of the skip connection for the neural network. For the case of the complete neural network model, Lemhadri et al. (2021) gracefully provide a proximal gradient operation to satisfy inequality constraints during training time which guarantees true zeros in the first MLP layer as well.

4.2.2 GEMINI gradients

We extend our initial work on gradients (Ohl et al. 2023) here by providing explicit gradients for the GEMINI functions and some explanations. Let us consider for example the derivative of the OvA MMD-GEMINI with respect to the probability output of some $p_{\theta}(y = k|\mathbf{x} = \mathbf{x}_i)$ for a sample \mathbf{x}_i in a

batch containing N samples. Adapting Eq. (2) to this specific distance, the objective to maximise is:

$$\mathcal{I} = \mathbb{E}_{y \sim p_{\theta}(y)} [\text{MMD}(p_{\theta}(\mathbf{x}|y) \| p_{\text{data}}(\mathbf{x}))]. \tag{14}$$

We can show that the gradient of an estimated OvA MMD-GEMINI $\hat{\mathcal{I}}$ is:

$$\begin{aligned} \frac{\partial \hat{\mathcal{I}}}{\partial p_{\theta}(y = k|\mathbf{x} = \mathbf{x}_i)} &= \frac{1}{\text{MMD}_k} \left[\frac{1}{N^3} \sum_{j,l}^N \kappa(\mathbf{x}_j, \mathbf{x}_l) \right. \\ &\quad \left. - \frac{2}{N^3} \sum_j^N \sum_{l=1}^N \kappa(\mathbf{x}_j, \mathbf{x}_l) \frac{p_{\theta}(y = k|\mathbf{x}_l)}{p_{\theta}(y = k)} \right. \\ &\quad \left. + \sum_j^N \kappa(\mathbf{x}_i, \mathbf{x}_j) \left(\frac{p_{\theta}(y = k|\mathbf{x} = \mathbf{x}_j)}{p_{\theta}(y = k)} - 1 \right) \right], \end{aligned} \tag{15}$$

where MMD_k stands for the MMD between the k -th cluster and the data distribution and κ is the kernel of choice. Regarding the interpretation of the gradient, we can observe that for all clusters k , the gradient is divided by the respective MMD distance. In other words, the smaller the MMD the stronger the gradient: we need to increase the MMD. All clusters k get their own common gradient throughout samples. This is due to the interdependence nature of the OvA MMD that requires batches of samples to be evaluated and cannot be performed on one sample at a time. Moreover, this common term comprises a constant which is the estimated participation of the data to the MMD. The second term is the cross-contribution of $p_{\theta}(y|\mathbf{x})$ and $p(\mathbf{x})$ in the evaluation of the distance. The stronger this cross-contribution is (and so the smaller the MMD), the smaller the gradient. Finally, each individual sample receives finally its own gradient which is its weighted average kernel strength by $\frac{p_{\theta}(y=k|\mathbf{x})}{p(y=k)} - 1$. Thus, the greater the conditional distribution of a sample compared to the cluster distribution, the stronger the gradient. Indeed, the more representative a sample of the cluster distribution, the greater its contribution to pulling the gradient to this representation.

We give further details on other gradients as well in Appendix B,C and D. The latter is dedicated to the Wasserstein GEMINI.

4.3 Implementations

Owing to the exact computations of the gradients, we developed a Python package encompassing all sparse GEMINI methods as well as original non-sparse GEMINI methods (Ohl et al. 2023) named *GemClus*.¹ Overall, the package

¹ The GemClus package can be found at <https://gemini-clustering.github.io/>.

Table 1 Brief description of datasets involved in experiments

Name	Samples	Features	#Classes
US-Congress	435	16	2
Heart-statlog	270	13	2
MNIST	12,000	784	10
MNIST-BR	12,000	784	10
Prostate-BCR	171	25,904	2

is designed for small datasets as it encompasses the precise and optimised computations of the gradients with lightweight dependencies, rather than using the automatic differentiation processes of heavy packages like PyTorch. The code of this paper was written using the latter.² Results may consequently slightly differ. We give detailed examples of code snippets in App. E, including how to reproduce the numerical experiments from Sect. 5.3.

5 Experiments

A brief summary of the datasets used in these experiments can be found in Table 1.

5.1 Metrics

Depending on the experiments for comparison purposes, we report 3 different metrics. The adjusted rand index (ARI, Hubert and Arabie, 1985) describes how close the clustering is to the classes, with a correction to random guesses. The variable selection error rate (VSER), for instance used by Celeux et al. (2014), describes the percentage of variables that the model erroneously omitted or accepted, therefore the lower the better. We finally report the correct variable rate (CVR) which describes how many of the expected variables were selected: higher is better. For example, a model selecting all variables of a dataset with d variables and d' good variables will get a CVR of 100% and a VSER of $1 - \frac{d'}{d}$.

5.2 Default hyperparameters

We set the hierarchy coefficient to $M = 10$, as Lemhadri et al. (2021) report that this value seems to “work well for a variety of datasets”. We also report the performances for the logistic regression mode when $M = 0$. The optimiser for the initial training step with $\lambda = 0$ is Adam (Kingma and Ba 2014) with a learning rate of 10^{-3} while other steps are done with SGD with momentum 0.9 and the same learning rate. Most of our experiments are done with 100 epochs per

step with early stopping as soon as the global objective does not improve by 1% for 10 consecutive epochs. The early stopping criterion is evaluated on the same training set since we do not seek to separate the dataset in train and validation sets in clustering. All activation functions are ReLUs. The default starting penalty is $\lambda_0 = 1$ with a 5% increase per step. We keep the linear kernel and the Euclidean distance respectively in conjunction with the MMD and Wasserstein distances when evaluating the GEMINI. Finally, we evaluate in most experiments the method with the exact same number of clusters as the number of known (supervised) labels.

5.3 Numerical experiments

We tested Sparse GEMINI on two synthetic datasets proposed by Celeux et al. (2014) and also used by Bouveyron and Brunet-Saumard (2014a) to first highlight some properties of the algorithm and compare it with competitors.

The first synthetic dataset consists of a few informative variables amidst noisy independent variables. The first 5 variables are informative and drawn from an equiprobable multivariate Gaussian mixture distribution of 3 components. All covariances are set to the identity matrix. The means are $\mu_1 = -\mu_2 = \alpha \mathbf{1}$ and $\mu_3 = \mathbf{0}$. All remaining p variables follow independent noisy centred Gaussian distributions. The number of samples N , the mean proximity α and the number of non-informative variables p vary over 5 scenarios. For the 2 first scenarios, we use $N = 30$ samples and $N = 300$ for others. The scenarios 1 and 3 present the challenge of close Gaussian distributions with $\alpha = 0.6$ while others use $\alpha = 1.7$. Finally, we add $p = 20$ noisy variables, except for the fifth scenario which takes up to $p = 95$ uninformative variables.

The second dataset consists of $n = 2000$ samples of 14 variables, 2 of them informative and most others linearly dependent on the former. The Gaussian mixture is equiprobable with 4 Gaussian distributions of means $[0, 0]$, $[4, 0]$, $[0, 2]$ and $[4, 2]$ with identity covariances. The 9 following variables are sampled as follows:

$$\mathbf{x}^{3-11} = [0, 0, 0.4, 0.8, 1.2, 1.6, 2.0, 2.4, 2.8]^\top + \mathbf{x}^{1-2^\top} \begin{bmatrix} 0.5 & 2 & 0 & -1 & 2 & 0.5 & 4 & 3 & 2 \\ 1 & 0 & 3 & 2 & -4 & 0 & 0.5 & 0 & 1 \end{bmatrix} + \boldsymbol{\epsilon}, \quad (16)$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega})$ with the covariance:

$$\boldsymbol{\Omega} = \text{diag} \left(\mathbf{I}_3, 0.5 \mathbf{I}_2, \text{diag}([1, 3]) \text{Rot} \left(\frac{\pi}{3} \right), \text{diag}[2, 6] \text{Rot} \left(\frac{\pi}{6} \right) \right). \quad (17)$$

Finally, the last 3 variables are sampled independently from $\mathcal{N}([3.2, 3.6, 4], \mathbf{I}_3)$.

For all synthetic datasets, we asked training to stop with F_{thres} set to the expected quantity of variables. We report the results of Sparse GEMINI in Table 2 after 20 runs. For

² The code for this paper can be found at <https://github.com/oshillou/SparseGEMINI>.

Table 2 Performances of Sparse GEMINI (OvO only) on synthetic datasets after 20 runs


	Sparse KMeans	ClustVarSel	VSCC	SFEM	MMD-GEMINI		Wasserstein-GEMINI	
					Logistic	MLP	Logistic	MLP
<i>(a) ARI scores (greater is better)</i>								
S1	0.09 _{0.08}	0.05 _{0.07}	0.00 _{0.02}	0.13 _{0.11}	0.14 _{0.11}	0.08 _{0.07}	0.09 _{0.12}	0.04 _{0.07}
S2	0.80 _{0.15}	0.20 _{0.22}	0.02 _{0.03}	0.72 _{0.17}	0.59 _{0.17}	0.52 _{0.23}	0.44 _{0.18}	0.43 _{0.15}
S3	0.11 _{0.03}	0.04 _{0.10}	0.15 _{0.10}	0.22 _{0.05}	0.21 _{0.05}	0.21 _{0.04}	0.14 _{0.05}	0.10 _{0.05}
S4	0.87 _{0.04}	0.88 _{0.04}	0.84 _{0.12}	0.87 _{0.04}	0.74 _{0.07}	0.86 _{0.05}	0.73 _{0.19}	0.82 _{0.11}
S5	0.87 _{0.03}	0.65 _{0.38}	0.00 _{0.00}	0.83 _{0.03}	0.76 _{0.05}	0.86 _{0.03}	0.59 _{0.20}	0.67 _{0.20}
D2	0.31 _{0.03}	0.60 _{0.02}	0.58 _{0.02}	0.58 _{0.01}	0.57 _{0.01}	0.54 _{0.03}	0.57 _{0.01}	0.55 _{0.02}
<i>(b) VSER scores (lower is better)</i>								
S1	0.31 _{0.22}	0.28 _{0.06}	0.72 _{0.15}	0.24 _{0.07}	0.44 _{0.15}	0.43 _{0.11}	0.51 _{0.11}	0.56 _{0.13}
S2	0.75 _{0.18}	0.29 _{0.07}	0.73 _{0.09}	0.27 _{0.08}	0.06 _{0.05}	0.11 _{0.07}	0.15 _{0.10}	0.24 _{0.12}
S3	0.47 _{0.34}	0.25 _{0.07}	0.65 _{0.22}	0.20 _{0.04}	0.07 _{0.05}	0.20 _{0.11}	0.20 _{0.12}	0.56 _{0.14}
S4	0.80 _{0.00}	0.01 _{0.03}	0.64 _{0.29}	0.23 _{0.08}	0.00 _{0.00}	0.00 _{0.00}	0.01 _{0.03}	0.00 _{0.02}
S5	0.95 _{0.00}	0.05 _{0.03}	0.95 _{0.00}	0.10 _{0.02}	0.00 _{0.00}	0.00 _{0.00}	0.01 _{0.01}	0.01 _{0.01}
D2	0.84 _{0.06}	0.00 _{0.00}	0.74 _{0.13}	0.52 _{0.08}	0.29 _{0.00}	0.31 _{0.04}	0.29 _{0.00}	0.29 _{0.00}
<i>(c) CVR scores (greater is better)</i>								
S1	0.53 _{0.31}	0.11 _{0.10}	0.87 _{0.23}	0.28 _{0.18}	0.60 _{0.26}	0.64 _{0.19}	0.63 _{0.23}	0.59 _{0.17}
S2	1.00 _{0.00}	0.14 _{0.17}	0.66 _{0.39}	0.39 _{0.17}	0.93 _{0.10}	0.82 _{0.22}	0.83 _{0.13}	0.78 _{0.14}
S3	1.00 _{0.00}	0.19 _{0.30}	0.97 _{0.13}	0.27 _{0.15}	0.95 _{0.09}	0.99 _{0.04}	0.68 _{0.25}	0.92 _{0.12}
S4	1.00 _{0.00}	0.19 _{0.30}	0.97 _{0.13}	0.27 _{0.15}	1.00 _{0.00}	1.00 _{0.00}	0.99 _{0.04}	0.99 _{0.04}
S5	1.00 _{0.00}	0.75 _{0.44}	1.00 _{0.00}	0.66 _{0.18}	1.00 _{0.00}	1.00 _{0.00}	0.94 _{0.11}	0.96 _{0.10}
D2	0.98 _{0.11}	1.00 _{0.00}	1.00 _{0.00}	1.00 _{0.00}	0.00 _{0.00}	0.00 _{0.00}	0.00 _{0.00}	0.00 _{0.00}

We compare our performances against other methods. S stands for a scenario of the first synthetic dataset and D2 stands for the second synthetic dataset. Standard deviation is reported in subscript
 Bold value indicates the best results

Table 3 Average regret scores (lower is better) between Sparse GEMINI for OvO MMD GEMINI against the best performing method per dataset

Method Dataset	ARI		VSER		CVR	
	Linear	MLP	Linear	MLP	Linear	MLP
1s1	0.00 _{0.11}	0.06 _{0.07}	0.17 _{0.15}	0.16 _{0.11}	0.28 _{0.26}	0.24 _{0.19}
1s2	0.19 _{0.17}	0.26 _{0.23}	0.02 _{0.05}	0.07 _{0.07}	0.07 _{0.10}	0.18 _{0.22}
1s3	0.02 _{0.05}	0.02 _{0.04}	0.01 _{0.05}	0.14 _{0.11}	0.04 _{0.09}	0.00 _{0.04}
1s4	0.14 _{0.07}	0.02 _{0.05}	0.00 _{0.00}	0.00 _{0.00}	0.00 _{0.00}	0.00 _{0.00}
1s5	0.12 _{0.05}	0.02 _{0.03}	0.00 _{0.00}	0.00 _{0.00}	0.00 _{0.00}	0.00 _{0.00}
2s2	0.02 _{0.01}	0.05 _{0.03}	0.29 _{0.00}	0.31 _{0.04}	1.00 _{0.00}	1.00 _{0.00}

Fig. 2 Example of convergence of the norm of the weights of the skip connection for every feature during training for the OvA Wasserstein objective. Green lines are the informative variables, black lines are the noise and red are the correlated variables. **a** In the case of noisy variables, Sparse GEMINI can recover the informative variables. **b** In the presence of redundant variables, Sparse GEMINI eliminates informative variables to keep the redundant ones



detailed distributions' box plots, refer to Appendix G. We compare our results against our own runs of other methods using their R package: SparseKMeans (Witten et al. 2013), ClustVarSel (Scrucca and Raftery 2018), vscc (Andrews and McNicholas 2013, 2014) and SparseFisherEM (Bouveyron and Brunet 2012). Due to the lack of space, we only report the scores for the one-vs-one GEMINI in Table 2. Extensive results using the one-vs-all GEMINI with the two architectures can be found in Appendix F.

It appears that the Sparse GEMINI is efficient in selecting the relevant variables when several others are noisy, especially with the OvO MMD objective while maintaining a high ARI. Moreover, while we do not systematically get the best ARI, our performances never fall far behind the most competitive method. We report in Table 3 the performances of the Sparse GEMINI MMD OvO that we deemed best against the best method for both architectures with regret scores. Regret scores are computed as the average difference between the scores of our model and the best performing one. Most scores are close to 0 for the third, fourth and fifth scenarios of the first dataset, except on the ARI. Our worst CVR is for the second dataset where we did not select the correct variables at all. We can also observe in Table 2 that the MMD objective learns well despite the presence of few samples in scenarios 2 and 3 and that the usage of an MLP leads to a trade-off between ARI and VSER when we have enough samples. Additionally, the selection strategy often leads to selecting the correct number of variables for the MMD, except in scenarios 1 and 3 where the Gaussian distributions are close to each other which is hard given the large variance. For the Wasserstein objective, we notice that the performances in selection are improved with the presence of more samples. However, the clustering performances are worse than the MMD, which we can attribute to the contribution of the noisy variables to the computation of the distances between samples, thus troubling the holistic perspective of the Wasserstein distance on the cluster distribution. It also appears that we performed poorly at selecting the correct variables in the presence of redundancy in the second dataset. However, since all variables except 3 are correlated to the informative variables, we still managed to get a correct ARI on the dataset while using other variables. On average, the variables selected by our models were the 6th and the 8th variables. We focus on this difference of convergence in Fig. 2 where we plot the norm of the skip connection per feature W_j . In the case of noisy variables, we are able to recover them as the number of selected features decreases, whereas we eliminated the informative variable of the second dataset during the first steps. In general, Clustvarsel (Scrucca and Raftery 2018) performed better on this type of synthetic dataset in terms of variable selection because it explicitly assumes a linear dependency between relevant variables and others.

5.4 Examples on MNIST and variations

We also demonstrate performance of the Sparse GEMINI algorithm by running it on the MNIST dataset. The initial λ_0 was set to 40. Following Lemhadri et al. (2021), we chose to stop training after finding 50 features. We also use 5% of dropout inside an MLP with 2 hidden layers of 1200 dimensions each (Hinton et al. 2012). We report in Fig. 3 the selected features by the clustering algorithms and the evolution of the ARI. We also extended this experiment to the MNIST variations proposed by Larochelle et al. (2007) showing the performance on the MNIST-BR dataset,³ a challenging dataset for unsupervised variable selection (Mattei et al. 2016). This variation consists in samples of MNIST with the black background being replaced by uniform noise hence displaying conditional noise on the data. To be fair, we reduced MNIST to the first 12,000 samples of the training set in order to match the number of samples in MNIST-BR.

We observed in Fig. 3 that for both the default MNIST dataset and the MNIST-BR dataset despite the presence of noise, the feature map concentrates precisely on the good location of the digits in the picture. Following the GEMINI curves in the Fig. 3a, b, the respective selected numbers of features were 122 for MNIST and 243 for MNIST-BR. These chosen models also have a respective ARI of 0.34 for 7 clusters and 0.28 for 8 clusters. The presence of empty clusters is a possible outcome with GEMINI (Ohl et al. 2023) which contributed here to lowering the ARI when evaluating with the true digits targets.





5.5 Real datasets

5.5.1 OpenML datasets

We ran Sparse GEMINI on two OpenML datasets that are often shown in related works: the US Congress dataset (Almanac 1984) and the Heart-statlog dataset (Brown 2004). The US Congress dataset describes the choice of the 435 representatives on 16 key votes in 1984. The labels used for evaluation are the political affiliations: 164 Republican against 267 Democrats. We replaced the missing values with 0 and converted the yes/no answers to 1, -1. Thus, an unknown label is equidistant from both answers. The Heart-statlog dataset describes 13 clinical and heart-related features with labels describing the presence or absence of cardiac disease among patients. We preprocessed it with standard scaling. For the US Congress dataset, we used one hidden layer of 20 nodes and a batch size of 87 samples. For the Heart-statlog dataset, we used 10 nodes and 90 samples. As we

³ Datasets were available at <https://web.archive.org/web/2018051912150/http://www.iro.umontreal.ca/~lisa/twiki/bin/view.cgi/Public/MnistVariations>.

Fig. 3 Relative importance of MNIST features after training of Sparse GEMINI with a log-scale color map. The blue features were eliminated at the first steps of λ , and the red features were eliminated last. On the right: evolution of the GEMINI depending on λ . F stands for the number of selected features



seek only two clusters, we only ran the one-vs-all versions of the GEMINI because it is strictly equal to the one-vs-one in binary clustering. Both datasets had a penalty increase of $\rho = 10\%$. We first show the number of selected features evolving with λ as well as the evolution of the GEMINI score as the number of features decreases respectively in Fig. 4 for the US Congress dataset and in Fig. 5 for Heart-statlog. Table 4 contains the performances for the two data sets, reporting the average number of variables selected over 20 runs according to our postprocessing selection criterion. We also added the performances of competitors from the previous section. However, we only managed to run Sparse Fisher EM on the Heart-statlog dataset, hence its absence for the US Congress scores. For comparison purposes, the best unsupervised accuracy reported on the Heart-statlog dataset by Solorio-Fernández et al. (2020) is 75.3%, while Sparse GEMINI achieves 79% with the MMD. The best score for all methods in the review (Solorio-Fernández et al. 2020) is 79.6%, but this encompasses filter methods which Sparse GEMINI is not. We also get similar results to the best performances of Marbac et al. (2020) who report 33% of ARI. Since most competitors retained all variables in the dataset, we chose to show as well the clustering performances without selection and hence with the greatest GEMINI score as well.


We averaged the number of times each feature was selected according to the model over the 20 runs and sorted them decreasingly. This post-process revealed that the Wasserstein objective consistently selected the El Salvador Aid and the Aid to Nicaraguan Contras votes as sufficient to perform clustering. Indeed, these two votes are among the most discriminating features between Republicans and Democrats and were often chosen by other model-based methods (Fop and Murphy 2018). The MMD objective only added the Physician fee freeze vote to this subset. Regarding the Heart-Statlog dataset, the MMD consistently picked a subset of 8 features out of 13, including, for example, age or chest pain type as relevant variables. Contrarily, the Wasserstein objective did not consistently choose the same subset of variables, yet its top variables that were selected more than 80% of the runs agree with the MMD selection as well.

5.5.2 Scalability example with the Prostate-BCR dataset

To show the scalability of Sparse GEMINI, we also demonstrate its performance on the Prostate-BCR dataset, taken from Vittrant et al. (2020).⁴ This dataset is a combination of transcriptomics data from 3 different sources: the Cancer Genom Atlas (Abeshouse et al. 2015), the GSE54460

⁴ The dataset is publicly available at https://github.com/ArnaudDroit/Lab/prostate_BCR_prediction.


Fig. 4 Average training curves of Sparse GEMINI on the US Congress dataset over 50 runs. Blue lines are Wasserstein, red lines are MMD



(a) GEMINI Loss depending on selected feature number

(b) Selected features depending on λ

Fig. 5 Average training curves of Sparse GEMINI on the Heart Statlog dataset over 20 runs. Blue lines are Wasserstein, red lines are MMD



(a) GEMINI depending on selected feature number

(b) Selected features depending on λ

Table 4 ARI of Sparse GEMINI (OvA) on the Heart-statlog and US Congress datasets with the average number of selected features

		Heart-statlog		US Congress	
		ARI	# Variables	ARI	# Variables
SparseKMeans		0.18 _{0.00}	13 _{0.00}	0.54 _{0.00}	16 _{0.0}
Clustvarsel		0.03 _{0.00}	2 _{0.00}	0.00 _{0.00}	2 _{0.00}
VSCC		0.27 _{0.00}	13 _{0.00}	0.40 _{0.00}	11 _{0.00}
Sparse Fisher EM		0.19 _{0.00}	1 _{0.00}	–	–
Logistic regression	MMD	0.37 _{0.03}	7.5 _{0.51}	0.53 _{0.02}	8.3 _{0.81}
	Wasserstein	0.33 _{0.08}	5.8 _{2.09}	0.48 _{0.00}	8.0 _{0.92}
MLP	MMD	0.32 _{0.01}	8.0 _{0.00}	0.48 _{0.00}	3.1 _{0.37}
	Wasserstein	0.32 _{0.09}	8.4 _{2.70}	0.47 _{0.00}	2.0 _{0.00}
MLP	MMD*	0.37 _{0.02}	13 _–	0.55 _{0.01}	16 _–
	Wasserstein*	0.33 _{0.09}	13 _–	0.55 _{0.02}	16 _–

Standard deviation in subscript. Scores with an asterisk are the initial performances when using all features
 Bold value indicates the best results

dataset from the NCBI website, and the PRJEB6530 project of the European Nucleotide Archive. The combined dataset contains 25,904 transcripts over 171 filtered patients with long-term follow-up, counting 52, 96 and 23 patients from the respective sources. The objective is to find biochemical recurrences (BCR) of prostate cancer through the transcriptomic signature, hence binary targets.

To carefully eliminate the variables, we increase λ gradually by 2%. We took a simple MLP with only one hidden layer of 100 neurons. We chose to run until we reached 400 features or less, following Vittrant et al. (2020). We trained Sparse GEMINI with OvA objectives 5 times to find either

2 or 3 clusters in order to break down possible substructures among the supervised targets.

Interestingly, we observed in Table 5 that the clustering results did not catch up with the actual BCR targets, with an ARI close to 0 most of the time. However, upon evaluation of the clusters with respect to the original source of each sample, we found scores close to 1 of ARI in the case of the MMD GEMINI. Thus, the unsupervised algorithm was able to find sufficient differences in distribution between each data source to discriminate them. Additionally, consistent subsets of features were always selected as the final subset on all 5 runs depending on the GEMINI. This implies that even

Table 5 ARI scores of the Prostate BCR dataset for various numbers of clusters depending on the chosen type of targets

Model	GEMINI	K	#Var	BCR targets ARI	Data source targets ARI
Logistic regression	MMD	2	810 ₅₉₀	-0.01 _{0.00}	0.79 _{0.01}
		3	1229 ₂₂₇₀	0.04 _{0.00}	1.00 _{0.01}
	Wasserstein	2	1334 ₂₅₆₁	0.01 _{0.02}	0.60 _{0.13}
		3	1430 ₃₁₂₇	0.04 _{0.01}	0.96 _{0.06}
MLP	MMD	2	4013 ₆₅₄₁	-0.01 _{0.00}	0.78 _{0.01}
		3	4287 ₆₅₉₀	0.03 _{0.02}	0.93 _{0.11}
	Wasserstein	2	4403 ₆₈₄₃	0.00 _{0.00}	0.65 _{0.04}
		3	4331 ₆₇₄₂	0.02 _{0.02}	0.80 _{0.20}

We either use the expected targets (BCR) regarding cancer prediction, or data source targets that identify the data origin of each sample. The indicated GEMINIs are in the one-vs-all setting
 Bold value indicates the best results

Table 6 ARI scores of the Prostate BCR dataset for various numbers of clusters depending on the chosen type of targets

Model	GEMINI	K	BCR targets ARI	Data source targets ARI
Logistic regression	MMD	2	-0.01 _{0.00}	0.80 _{0.01}
		3	0.03 _{0.01}	0.97 _{0.07}
	Wasserstein	2	0.00 _{0.00}	0.75 _{0.00}
		3	0.03 _{0.02}	0.90 _{0.13}
MLP	MMD	2	0.01 _{0.02}	0.56 _{0.11}
		3	0.04 _{0.01}	0.95 _{0.05}
	Wasserstein	2	0.00 _{0.00}	0.68 _{0.07}
		3	0.03 _{0.03}	0.86 _{0.16}

We either use the expected targets (BCR) regarding cancer prediction, or data source targets that identify the data origin of each sample. The indicated GEMINIs are in the one-vs-all setting

without the best GEMINI within a range for feature selection, several runs can lead to identifying subsets of relevant data as well. This example illustrates how even in the presence of potentially legitimate labels, there exist other valid cluster structures in the data (Hennig 2015). By comparing these results with the performances using all features in Table 6, we observe that in general the MMD objective gained in ARI with the decrease of the number of features, whereas Wasserstein GEMINI is stronger when observing all features.

These results can be viewed as discovering batch effect in the data. Batch effect, also known as batch variation, is a phenomenon that occurs in biological experiments where the results are affected by factors unrelated to the experimental variables being studied. These factors can include variations in sample processing, measurement conditions, people manipulating the samples, or equipment used. One common example of a batch effect is observed in microarray or RNA sequencing experiments, where the samples are processed in different batches and the results are affected by variations in the reagents or protocols used. Batch effects in microarray experiments have been shown to originate from multiple causes, including variations in the labelling and

hybridisation protocols used, leading to differences in the intensity of gene expression signals (Luo et al. 2010).

To minimise batch effects, it is important to control variables such as reagents, protocols, and equipment used, and to use appropriate normalisation and data analysis methods to account for these variations. Several approaches can be used to detect batch effects in RNA-seq experiments, including PCA (Reese et al. 2013) and clustering. For this latter, Hierarchical clustering is often used as a method that groups samples based on their similarity in gene expression patterns, and batch effects can be identified based on dendrogram analysis (Leek et al. 2010).

5.6 Discussion

Our first observation from Table 2 is that the Sparse GEMINI algorithm can achieve performance close to some competitors in terms of ARI while performing better in variable selection, especially for the one-vs-one MMD. The MMD is a distance computed between expectations, making it thus insensitive to small variations of the kernel, typically when noisy variables are introduced, contrary to the Wasserstein distance which takes a global point of view on the dis-

tribution. Specifically, the algorithm is good at discarding noisy variables, but less competitive with regard to redundant variables as illustrated with the second synthetic dataset. Nonetheless, the ARI remains competitive even though the model failed to give the correct ground for the clustering.

Additionally, the training path produces critical values of λ at which the features disappear. Thus, the algorithm produces an explicit unsupervised metric of the relevance of each feature according to the clustering. Typically, plateaus of the number of used variables like in Figs. 4b and 5b for the MMD shed light on different discriminating subsets. We also find that the empirical threshold of 90% of the maximal GEMINI to select fewer variables is an efficient criterion. In case of a too sudden collapse of variables, we encourage training over again models on iteratively selected subsets of features. Indeed, as λ increases during training, the collapse of the number of selected variables will often happen when the geometric increase is too strong which might lead to unstable selections.

6 Conclusion

We presented a novel algorithm named Sparse GEMINI that jointly performs clustering and feature selection by combining GEMINI for objective and an ℓ_1 penalised skip connection. The algorithm shows good performances in eliminating noisy irrelevant variables while maintaining relevant clustering. To eliminate redundant variables, our future works can focus on adding the correlation of selected variables in the penalty. Owing to the nature of multi-layered perceptrons, Sparse GEMINI is easily scalable to high-dimensional data and thus provides an unsupervised technique to get a projection of the data. However, the limits of the scalability are the number of clusters and samples per batch due to the complex nature of GEMINI. Thus, we believe that Sparse GEMINI is a relevant algorithm for multi-omics data where the number of samples is often little and the number of features large, especially when it is hard to design a good generative model for such data. As a concluding remark, we want to draw again the attention to the discriminative nature of the algorithm: Sparse GEMINI focuses on the design of a decision boundary instead of parametric assumptions.

Appendix A The dynamic training regime

As features get eliminated during the training, the notion of affinity (distance δ or kernel κ) and clustering with respect to GEMINI between two samples changes. Indeed, GEMINI aims at maximising a distance between two related distributions using an affinity computed between samples, yet removing features from the inference implies that we no

longer cluster the same original data space, but rather a subspace at step t : $\mathcal{X}_t = \prod_{j \in I_t} \mathcal{X}_j$. If we still compute our affinity function using all features from \mathcal{X} the extra removed features may bring noise compared to the affinity between the relevant features, and thus bring confusion with regards to the ideal decision boundary.

To respect the original notion of GEMINI in clustering, we introduce the dynamic training regime, where at each time step t , the affinity function is computed using only the subset of relevant features I_t . We call *static* regime the training with usage of all features in the affinity function as described in Sect. 3. The advantage of the dynamic training regime is that it respects the notion of GEMINI with regard to the decision boundary, while the static regime yields comparable values of GEMINI independently of the number of selected features. However, the dynamic regime is incompatible with the selection process described in Sect. 4.1 because any change of data space implies a change of values for kernels or distances and thus for GEMINI, making models incomparable. Moreover, we may have more theoretical guarantees of convergence for the usual static regime than in the dynamic regime which may seem unstable.

We experiment this approach again with the synthetic data sets and report the results in Table 7. For this experiment, we only evaluated the performances on the final subset of selected features. However, since Sparse GEMINI is trained until a user-defined number of features is reached, we avoid unfair comparisons with other variable selection methods and do not report the VSER and the CVR. Our main observation on the introduction of the dynamic regime is that it greatly improves the clustering performances of the Wasserstein-GEMINI while not affecting the MMD-GEMINI. This success can be explained by the removal of variables as the removal of noise in the distance computation which is crucial for the Wasserstein distance because it takes a global point of view on the complete distribution. In contrast, the MMD only considers the expectation, which helps in removing noisy variations of the distance around informative variables.

Appendix B Differentiation of the OvA MMD

B.1 Alternative computation of the forward pass

We consider the computations starting from a row-stochastic matrix $\boldsymbol{\tau} \in \mathbb{R}^{N \times K}$, typically the softmax output of a model. We focus here only on the computations of the objective function, the OvA MMD. First, we can compute the cluster proportions:

$$\boldsymbol{\pi} = \frac{1}{N} \mathbf{1}_N^\top \boldsymbol{\tau}. \quad (\text{B1})$$

Table 7 ARI scores on the synthetic datasets with the dynamic regime of training for the Sparse GEMINI using MLPs

Method	MMD		Wasserstein	
	OvA	OvO	OvA	OvO
Scenario 1	0.12 (0.13)	0.15 (0.12)	0.05 (0.09)	0.07 (0.06)
Scenario 2	0.48 (0.11)	0.63 (0.21)	0.37 (0.11)	0.30 (0.13)
Scenario 3	0.23 (0.04)	0.20 (0.03)	0.11 (0.05)	0.11 (0.06)
Scenario 4	0.45 (0.07)	0.88 (0.03)	0.82 (0.13)	0.85 (0.10)
Scenario 5	0.62 (0.09)	0.84 (0.05)	0.56 (0.20)	0.48 (0.17)
Dataset 2	0.54 (0.04)	0.54 (0.05)	0.50 (0.08)	0.56 (0.01)

Bold value indicates the best results

Our goal is to compute the vector $\Delta \in \mathbb{R}^K$ where the k -th component is the squared distance in the Hilbert space between one cluster distribution and the data distribution:

$$\Delta_k = \sum_{i,j}^{N,N} \tilde{\kappa}_{i,j} \left[\frac{\tau_{ki} \tau_{kj}}{\pi_k^2} + 1 - 2 \frac{\tau_{ki}}{\pi_k} \right]. \tag{B2}$$

To that end, we introduce the matrix $\alpha \in \mathbb{R}^{N \times K}$ which is the element-wise division of τ by the proportions of the matching cluster.

$$\alpha = \tau \oslash (\mathbf{1}_N \pi^\top) = \left[\frac{\tau_{ki}}{\pi_k} \right]. \tag{B3}$$

Individually, we can interpret the value of α_{ik} as the ratio $p(y = k | \mathbf{x}_i) / p(y = k)$ or $p(\mathbf{x}_i | y = k) / p(\mathbf{x}_i)$. This represents the relative strength of the sample in the cluster distribution. We can then deduce the writing of Δ :

To name the elements, \mathbf{a} is the cluster contribution: how the distribution of the cluster contributes to increase the distance, and the same goes for the constant \mathbf{c} which represents the agnostic data strength. Finally, \mathbf{b} is the agreement between the two distributions $p(y | = k | \mathbf{x})$ and $p(\mathbf{x})$ which diminishes the value of the MMD: the more the cluster distribution takes to the data (by having everything in the same cluster), the stronger \mathbf{b} is and the lower the MMD. Yet, to simplify the derivatives to compute later, we introduce two intermediary variables:

$$\gamma = \tilde{\kappa} \alpha, \tag{B4}$$

and

$$\omega = \alpha^\top \gamma, \tag{B5}$$

of respective shapes $N \times K$ and $K \times K$. Thus, we simply rewrite:

$$\Delta = \mathbf{a} + \mathbf{c} - 2\mathbf{b}, \tag{B6}$$

$$= \text{diag}(\omega) + \mathbf{1}_{K \times N} \tilde{\kappa} \mathbf{1}_N - 2\gamma^\top \mathbf{1}_N. \tag{B7}$$




Fig. 6 Summary of computations for the forward pass of the OvA MMD using matrices

Finally, assuming the square root is applied element-wise, we can write the final objective as:

$$\hat{\mathcal{I}}_{\text{MMD}}^{\text{ova}}(\mathbf{x}, y | \theta) = \pi^\top \sqrt{\Delta} \tag{B8}$$

The graph of computations is summarised in Fig. 6.

B.2 Backward pass

We can now compute the derivatives of each part of the graph with respect to the conditional probabilities: $\frac{\partial \hat{\mathcal{I}}}{\partial \tau}$. By reversing the graph, we get the list of the following sorted derivatives to compute: (1) $\frac{\partial \hat{\mathcal{I}}}{\partial \Delta}$ (2) $\frac{\partial \hat{\mathcal{I}}}{\partial \mathbf{a}}$ (3) $\frac{\partial \hat{\mathcal{I}}}{\partial \mathbf{b}}$ (4) $\frac{\partial \hat{\mathcal{I}}}{\partial \omega}$ (5) $\frac{\partial \hat{\mathcal{I}}}{\partial \gamma}$ (6) $\frac{\partial \hat{\mathcal{I}}}{\partial \alpha}$ (7) $\frac{\partial \hat{\mathcal{I}}}{\partial \pi}$ (8) $\frac{\partial \hat{\mathcal{I}}}{\partial \tau}$.

To compute these derivatives, we will follow an automatic differentiation procedure. All derivatives correspond to the gradient of a scalar with respect to a matrix or vector, hence all derivatives will keep the same shape as the denominator. Notice that because c and $\tilde{\kappa}$ do not depend on τ , they will not produce any gradient.

B.2.1 Deriving for Δ

We simply take the vector π for this derivative that summed the square root of Δ . Additionally, the element-wise square root is differentiated. Thus:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \Delta} = \frac{\pi}{2\sqrt{\Delta}}. \tag{B9}$$

B.2.2 Deriving for \mathbf{a} and \mathbf{b}

The contributions of \mathbf{a} and \mathbf{b} are simple element-wise sums of vectors. They have the same shape as the previous gradient. Therefore:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \mathbf{a}} = \frac{\partial \hat{\mathcal{I}}}{\partial \Delta} = \frac{\pi}{2\sqrt{\Delta}}, \tag{B10}$$

$$\frac{\partial \hat{\mathcal{I}}}{\partial \mathbf{b}} = -2 \frac{\partial \hat{\mathcal{I}}}{\partial \Delta} = -\frac{\pi}{\sqrt{\Delta}}. \tag{B11}$$

B.2.3 Deriving for ω

Since we took only the diagonal of ω for computation, the gradient with respect to ω will be a diagonal matrix, which diagonal is exactly the previously committed error:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \omega} = \text{diag} \left(\frac{\partial \hat{\mathcal{I}}}{\partial \mathbf{a}} \right) = \frac{1}{2} \text{diag} \left(\frac{\pi}{\sqrt{\Delta}} \right). \tag{B12}$$

B.2.4 Deriving for γ

The vector γ contributed two times in the computation graph: once to ω and another time for \mathbf{b} . Both cases involve simple matrix multiplications. We can sum the matrix gradient to both errors to get:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \gamma} = \frac{\partial \hat{\mathcal{I}}}{\partial \omega} \alpha^\top + \frac{\partial \hat{\mathcal{I}}}{\partial \mathbf{b}} \mathbf{1}_N^\top, \tag{B13}$$

$$= \frac{1}{2} \alpha \text{diag} \left(\frac{\pi}{\sqrt{\Delta}} \right) - \mathbf{1}_N \left(\frac{\pi}{\sqrt{\Delta}} \right)^\top. \tag{B14}$$

Notice that for the derivative from \mathbf{b} , we had to transpose the error since \mathbf{b} is computed using γ^\top . Here, we can remark that by unfolding the definition of α , and thanks to matrix product with the diagonal product, the values of π gets cancelled. Therefore:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \gamma} = \frac{1}{2} \frac{\tau}{\mathbf{1}_N \sqrt{\Delta}^\top} - \mathbf{1}_N \left(\frac{\pi}{\sqrt{\Delta}} \right)^\top. \tag{B15}$$

B.2.5 Deriving for α

As we did for γ , we need to sum here the gradient contributions of α to ω and γ . Both are matrix multiplications, however we add a transposition in the case of ω . Thus:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \alpha} = \left(\frac{\partial \hat{\mathcal{I}}}{\partial \omega} \gamma^\top \right)^\top + \tilde{\kappa}^\top \frac{\partial \hat{\mathcal{I}}}{\partial \gamma}, \tag{B16}$$

$$= \frac{1}{2} \gamma \text{diag} \left(\frac{\pi}{\sqrt{\Delta}} \right) + \tilde{\kappa} \left(\frac{1}{2} \frac{\tau}{\mathbf{1}_N \sqrt{\Delta}^\top} - \mathbf{1}_N \left(\frac{\pi}{\sqrt{\Delta}} \right)^\top \right). \tag{B17}$$

Here, we can unfold the definition of γ to make a common factor $\tilde{\kappa}$ appear on the left matrix multiplication. Thus:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \alpha} = \tilde{\kappa} \left[\frac{1}{2} \alpha \text{diag} \left(\frac{\pi}{\sqrt{\Delta}} \right) + \frac{1}{2} \frac{\tau}{\mathbf{1}_N \sqrt{\Delta}^\top} - \mathbf{1}_N \left(\frac{\pi}{\sqrt{\Delta}} \right)^\top \right]. \tag{B18}$$

We notice the exact same simplification on the left term between α and the diagonal matrix as we had for the gradient w.r.t. γ . Rewriting this term is exactly equal to the second, and thus:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \alpha} = \tilde{\kappa} \left[\frac{\tau}{\mathbf{1}_N \sqrt{\Delta}^\top} - \mathbf{1}_N \left(\frac{\pi}{\sqrt{\Delta}} \right)^\top \right]. \tag{B19}$$

B.2.6 Deriving for π

Same procedure for π by summing the contributions of the gradient from α and the dot product with $\sqrt{\Delta}$ in \mathcal{I} . For the derivative from α , we multiply the rows of the previous error by the squared inverse of π and β and sum them. Hence:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \pi} = \sqrt{\Delta} - \left[\left[\frac{\tau^\top}{\sqrt{\Delta} \mathbf{1}_N^\top} - \frac{\pi}{\sqrt{\Delta}} \mathbf{1}_N^\top \right] \tilde{\kappa} \odot \frac{\beta}{\pi^2 \mathbf{1}_N^\top} \right] \mathbf{1}_N, \tag{B20}$$

$$= \sqrt{\Delta} - \left[\left[\frac{\tau^\top}{\sqrt{\Delta} \mathbf{1}_N^\top} - \frac{\pi}{\sqrt{\Delta}} \mathbf{1}_N^\top \right] \tilde{\kappa} \odot \frac{\alpha}{\pi \mathbf{1}_N^\top} \right] \mathbf{1}_N. \tag{B21}$$

Since the inverse factor $1/\pi \mathbf{1}_N^\top$ is constant row-wise, we can incorporate it directly to the left term of the matrix multiplication. This simplifies again the notations:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \pi} = \sqrt{\Delta} + \left[\left(\frac{\alpha^\top \tilde{\kappa}}{\sqrt{\Delta} \mathbf{1}_N^\top} - \frac{\mathbf{1}_N^\top \tilde{\kappa}}{\sqrt{\Delta} \mathbf{1}_N^\top} \right) \odot \alpha \right] \mathbf{1}_N. \tag{B22}$$

Here, the combination of the element-wise multiplication by α followed by a sum over all samples is in fact equal to the respective distance terms \mathbf{a} and \mathbf{b} . First simplification yields:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \boldsymbol{\pi}} = \sqrt{\Delta} - \frac{\mathbf{a} - \mathbf{b}}{\sqrt{\Delta}}. \tag{B23}$$

To finally go further, we can use the definition of Δ to replace the left term by another. Indeed, since:

$$\mathbf{a} - \mathbf{b} = \Delta + \mathbf{b} - \mathbf{c}, \tag{B24}$$

and the denominator $\sqrt{\Delta}$ gets cancelled by Δ , we obtain:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \boldsymbol{\pi}} = \frac{\mathbf{c} - \mathbf{b}}{\sqrt{\Delta}}. \tag{B25}$$

B.2.7 Deriving for $\boldsymbol{\tau}$

Finally, the gradient for $\boldsymbol{\tau}$ sums contributions from both α and $\boldsymbol{\pi}$. In both cases, we just consider element-wise operations, so the global gradient will just be element-wise multiplication of the errors, with a specific repetition over all rows for the gradient from $\boldsymbol{\pi}$:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \boldsymbol{\tau}} = \frac{\mathbf{1}_N}{N} \frac{\partial \hat{\mathcal{I}}}{\partial \boldsymbol{\pi}} + \frac{\partial \hat{\mathcal{I}}}{\partial \boldsymbol{\alpha}} \odot \frac{1}{\mathbf{1}_N \boldsymbol{\pi}^\top}, \tag{B26}$$

$$= \mathbf{1}_N \frac{\mathbf{c} - \mathbf{b}^\top}{N \sqrt{\Delta}} + \tilde{\boldsymbol{\kappa}} \left[\frac{\boldsymbol{\tau}}{\mathbf{1}_N \sqrt{\Delta}^\top} - \mathbf{1}_N \left(\frac{\boldsymbol{\pi}}{\sqrt{\Delta}} \right)^\top \right] \odot \frac{1}{\mathbf{1}_N \boldsymbol{\pi}^\top}, \tag{B27}$$

$$= \mathbf{1}_N \frac{\mathbf{c} - \mathbf{b}^\top}{N \sqrt{\Delta}} + \tilde{\boldsymbol{\kappa}} \left[\frac{\boldsymbol{\alpha}}{\mathbf{1}_N \sqrt{\Delta}^\top} - \frac{1}{\mathbf{1}_N \sqrt{\Delta}^\top} \right]. \tag{B28}$$

To conclude, we can factorise all terms by the common denominator:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \boldsymbol{\tau}} = \frac{1}{\mathbf{1}_N \sqrt{\Delta}^\top} \odot \left[\frac{\mathbf{1}_N}{N} (\mathbf{c} - \mathbf{b})^\top + \tilde{\boldsymbol{\kappa}} (\boldsymbol{\alpha} - \mathbf{1}_{N \times K}) \right]. \tag{B29}$$

For further simplification of the gradients, we can unfold again the definition of \mathbf{b} and \mathbf{c} as follows:

$$\frac{\mathbf{1}_N}{N} (\mathbf{c} - \mathbf{b})^\top = \frac{1}{N} [\mathbf{1}_{N \times N} \tilde{\boldsymbol{\kappa}} \mathbf{1}_{N \times K} - \mathbf{1}_{N \times N} \tilde{\boldsymbol{\kappa}} \boldsymbol{\alpha}], \tag{B30}$$

$$= \frac{1}{N} [\mathbf{1}_{N \times N} \tilde{\boldsymbol{\kappa}} (\mathbf{1}_{N \times K} - \boldsymbol{\alpha})]. \tag{B31}$$

Thus, we can conclude that the final equation for the gradient of the OvA MMD is:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \boldsymbol{\tau}} = \frac{1}{\mathbf{1}_N \sqrt{\Delta}^\top} \odot [(I_N - \mathbf{1}_{N \times N} / N) \tilde{\boldsymbol{\kappa}} (\boldsymbol{\alpha} - \mathbf{1}_{N \times K})]. \tag{B32}$$

To be more precise, we can even express the value for a component at position i, k :

$$\frac{\partial \hat{\mathcal{I}}}{\partial \boldsymbol{\tau}_{i,k}} = \left[\frac{1}{\sqrt{\Delta_k}} \sum_{j=1}^N \left(\tilde{\boldsymbol{\kappa}}_{ij} - \frac{1}{N} \sum_{l=1}^N \tilde{\boldsymbol{\kappa}}_{jl} \right) (\alpha_{jk} - 1) \right] \tag{B33}$$

$$= \left[\frac{1}{\sqrt{\Delta_k}} (c - b_k + \gamma_{ik} - \bar{\boldsymbol{\kappa}}_i) \right], \tag{B34}$$

with $\bar{\boldsymbol{\kappa}}_i = \sum_{j=1}^N \tilde{\boldsymbol{\kappa}}_{ij}$.

Appendix C Differentiation of the OvO MMD

C.1 Forward pass

We will proceed here to the exact same reasoning as in the OvA MMD. We first compute the distance Δ before summing them with $\boldsymbol{\pi}$. Contrary to the OvA MMD, Δ is now a matrix of shape $K \times K$ where each entry describes the distance between two clusters k and k' :

$$\hat{\mathcal{I}}_{\text{MMD}}^{\text{ovo}}(\mathbf{x}, y | \theta) = \boldsymbol{\pi}^\top \sqrt{\Delta} \boldsymbol{\pi}. \tag{C35}$$

As previously done, we can express the squared distance as the sum of two self-contributions minus a cross-contribution. These contributions will be here matrices of shape $K \times K$. Yet, we can notice that in the OvO MMD, the matrix Δ is symmetric. Simply put, the cross-contribution is symmetric, and the two self-contributions are the transposed of each other:

$$\Delta = \mathbf{A} + \mathbf{C} - 2\mathbf{B}, \tag{C36}$$

$$= \mathbf{A} + \mathbf{A}^\top - 2\mathbf{B}. \tag{C37}$$

We can here realise that the matrix \mathbf{A} is in fact a column-wise copy of the vector \mathbf{a} from the previous computations with OvA MMD. Similarly, \mathbf{B} is the entire matrix $\boldsymbol{\omega}$ while \mathbf{A} only consists in its diagonal. Therefore:

$$\Delta = \text{diag}(\boldsymbol{\omega}) \mathbf{1}_K^\top + \mathbf{1}_K \text{diag}(\boldsymbol{\omega})^\top - 2\boldsymbol{\omega}. \tag{C38}$$

The remaining of the definition of $\boldsymbol{\omega}$ strictly unfolds from the OvA MMD forward pass.

C.2 Backward pass

In the specific case of the OvO, we have square roots of values of Δ which can be equal to 0, hence undifferentiable. This is in fact not a burden since in principle, these 0 only happen when we evaluate the MMD between a cluster and itself. Thus, we can discard easily the null components of Δ during the final sum (expectation over π) and adopt locally the small convention that the derivative of \mathcal{I} w.r.t. Δ will be equal to 0 on the diagonal, despite the square root computation.

C.2.1 Deriving for Δ

We start simple, the derivative is simply a square matrix where all components are the cartesian product of the vector π :

$$\frac{\partial \hat{\mathcal{I}}}{\partial \Delta} = \frac{\pi \pi^\top}{2\sqrt{\Delta}}. \tag{C39}$$

From now on, we will arbitrarily say that $\left(\frac{\partial \hat{\mathcal{I}}}{\partial \Delta}\right)_{k,k} = 0$ because it was not summed at the end of the forward pass in the OvO MMD. Thus, we will write for clarity:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \Delta} = \frac{\pi \pi^\top}{2\sqrt{\Delta}} \odot (\mathbf{1}_{K \times K} - \mathbf{I}_K). \tag{C40}$$

C.2.2 Deriving for ω

For the gradient w.r.t. ω , we have two contributions to sum, one which comes from the diagonal of ω times 2, and another from the complete matrix ω :

$$\frac{\partial \hat{\mathcal{I}}}{\partial \omega} = \frac{\partial \hat{\mathcal{I}}}{\partial \Delta} \left(\frac{\partial \Delta}{\partial \text{diag}(\omega)} \frac{\partial \text{diag}(\omega)}{\partial \omega} + \frac{\partial \Delta}{\partial \text{diag}(\omega)^\top} \frac{\partial \text{diag}(\omega)^\top}{\partial \omega} \right) - 2 \times \frac{\partial \hat{\mathcal{I}}}{\partial \Delta} \frac{\partial \Delta}{\partial \omega}, \tag{C41}$$

$$= 2 \text{diag} \left(\frac{\partial \hat{\mathcal{I}}}{\partial \Delta} \mathbf{1}_K \right) - 2 \frac{\partial \hat{\mathcal{I}}}{\partial \Delta}. \tag{C42}$$

We can simplify the factor 2 to get:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \omega} = \text{diag} \left(\left[\frac{\pi \pi^\top}{\sqrt{\Delta}} \odot (\mathbf{1}_{K \times K} - \mathbf{I}_K) \right] \mathbf{1}_K \right) \tag{C43}$$

$$- \frac{\pi \pi^\top}{\sqrt{\Delta}} \odot (\mathbf{1}_{K \times K} - \mathbf{I}_K). \tag{C44}$$

This gradient says that on all parts of the matrix except the diagonal, we backpropagate the cross-contribution from ω , but sum all this contributions as well on the diagonal. To

ease later writings, we introduce Λ :

$$\Lambda = \frac{\pi \pi^\top}{\sqrt{\Delta}} \odot (\mathbf{1}_{K \times K} - \mathbf{I}_K). \tag{C45}$$

Thanks to Δ and the cross-product of the vector π , Λ is positive and symmetric.

C.2.3 Deriving for γ

Now, we can backpropagate as we did for the OvA MMD, except that γ only contributed once to the computation of ω . Thus:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \gamma} = \alpha \frac{\partial \hat{\mathcal{I}}}{\partial \omega}, \tag{C46}$$

$$= \alpha [\text{diag}(\Lambda \mathbf{1}_K) - \Lambda]. \tag{C47}$$

C.2.4 Deriving for α

The derivative w.r.t. α is a backpropagation through two matrix multiplications: one in γ and one in ω . Hence:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \alpha} = \tilde{\kappa} \frac{\partial \hat{\mathcal{I}}}{\partial \gamma} + \gamma \frac{\partial \hat{\mathcal{I}}}{\partial \omega}, \tag{C48}$$

$$= \tilde{\kappa} \alpha \frac{\partial \hat{\mathcal{I}}}{\partial \omega} + \gamma \frac{\partial \hat{\mathcal{I}}}{\partial \omega}, \tag{C49}$$

$$= 2\gamma \frac{\partial \hat{\mathcal{I}}}{\partial \omega}. \tag{C50}$$

Thus:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \alpha} = 2\gamma [\text{diag}(\Lambda \mathbf{1}_K) - \Lambda]. \tag{C51}$$

Note that we can write the first term differently because we multiply the matrix γ by a diagonal matrix. This is equivalent to doing an element-wise multiplication of γ by the vector inside the diag function repeated row-wise:

$$\frac{\partial \hat{\mathcal{I}}}{\partial \alpha} = 2\mathbf{1}_{N \times K} \Lambda \odot \gamma - 2\gamma \Lambda. \tag{C52}$$

C.2.5 Deriving for π

The proportions contributed in the final expectation with Δ and in the computations of α . This looks like what we had in the OvA MMD backpropagation. Therefore:

$$\begin{aligned} \frac{\partial \hat{\mathcal{I}}}{\partial \pi} &= \frac{\partial \pi^\top \sqrt{\Delta} \pi}{\partial \pi} - \mathbf{1}_N^\top \left[\frac{\alpha}{\mathbf{1}_N \pi^\top} \odot \frac{\partial \hat{\mathcal{I}}}{\partial \alpha} \right], \\ &= 2\pi^\top \sqrt{\Delta} \end{aligned} \tag{C53}$$

$$- 2 \times \mathbf{1}_N^\top \left[\frac{\boldsymbol{\alpha}}{\mathbf{1}_N \boldsymbol{\pi}^\top} \odot (\mathbf{1}_{N \times K} \boldsymbol{\Lambda} \odot \boldsymbol{\gamma} - \boldsymbol{\gamma} \boldsymbol{\Lambda}) \right]. \quad (C54)$$

By noticing that the matrix $\mathbf{1}_{N \times K} \boldsymbol{\Lambda}$ is constant row-wise, we can easily permute the element-wise operation with $\boldsymbol{\gamma}$ and perform the matrix multiplication with $\mathbf{1}_N^\top$ before thanks to factorisation. The element-wise product of $\boldsymbol{\alpha}$ with $\boldsymbol{\gamma}$ summed over all samples is equal to the diagonal of $\boldsymbol{\omega}$. We can rewrite:

$$\frac{\partial \hat{\mathcal{L}}}{\partial \boldsymbol{\pi}} = 2 \boldsymbol{\pi}^\top \sqrt{\Delta} - 2 \frac{\text{diag}(\boldsymbol{\omega})}{\boldsymbol{\pi}^\top} \odot \mathbf{1}_K \boldsymbol{\Lambda} \quad (C55)$$

$$+ \frac{2}{\boldsymbol{\pi}^\top} \odot \left(\mathbf{1}_N^\top [\boldsymbol{\alpha} \odot \boldsymbol{\gamma} \boldsymbol{\Lambda}] \right). \quad (C56)$$

C.2.6 Deriving for $\boldsymbol{\tau}$

We can finally draw a conclusion to this backpropagation by summing the gradient of the two contributions of $\boldsymbol{\tau}$: one from $\boldsymbol{\alpha}$ and another one from $\boldsymbol{\pi}$:

$$\frac{\partial \hat{\mathcal{L}}}{\partial \boldsymbol{\tau}} = \frac{1}{\mathbf{1}_N \boldsymbol{\pi}^\top} \odot \frac{\partial \hat{\mathcal{L}}}{\partial \boldsymbol{\alpha}} + \frac{1}{N} \mathbf{1}_N \frac{\partial \hat{\mathcal{L}}}{\partial \boldsymbol{\pi}}, \quad (C57)$$

$$\begin{aligned} &= \frac{2}{\mathbf{1}_N \boldsymbol{\pi}^\top} \odot [\boldsymbol{\gamma} \odot \mathbf{1}_{N \times K} \boldsymbol{\Lambda} - \boldsymbol{\gamma} \boldsymbol{\Lambda}] + \frac{2}{N} \mathbf{1}_N \boldsymbol{\pi}^\top \sqrt{\Delta} \\ &\quad - \frac{2}{N} \mathbf{1}_N \left[\frac{\text{diag}(\boldsymbol{\omega})}{\boldsymbol{\pi}^\top} \odot \mathbf{1}_K \boldsymbol{\Lambda} \right] \\ &\quad + \frac{2}{N} \mathbf{1}_N \left[\frac{1}{\boldsymbol{\pi}^\top} \odot \left(\mathbf{1}_N^\top [\boldsymbol{\alpha} \odot \boldsymbol{\gamma} \boldsymbol{\Lambda}] \right) \right], \quad (C58) \end{aligned}$$

$$\begin{aligned} &= \frac{2}{N} \mathbf{1}_N \boldsymbol{\pi}^\top \sqrt{\Delta} + \frac{2}{\mathbf{1}_N \boldsymbol{\pi}^\top} \odot [\boldsymbol{\gamma} \odot \mathbf{1}_{N \times K} \boldsymbol{\Lambda} - \boldsymbol{\gamma} \boldsymbol{\Lambda} \\ &\quad - \frac{1}{N} (\text{diag}(\boldsymbol{\omega}) \odot \mathbf{1}_K \boldsymbol{\Lambda}) + \frac{1_{N \times N}}{N} (\boldsymbol{\alpha} \odot \boldsymbol{\gamma} \boldsymbol{\Lambda})]. \quad (C59) \end{aligned}$$

Appendix D Gradients for the Wasserstein GEMINI

We seek the expression of the gradient of the Wasserstein GEMINI for some output $\boldsymbol{\tau} \in \mathbb{R}^{N \times K}$ of some probabilistic model. The matrix $\boldsymbol{\tau}$ is therefore row-stochastic. The expression of the GEMINI in the one-vs-all context is then:

$$\mathcal{I}_{\mathcal{W}_\delta}^{\text{ova}}(\mathbf{x}; y|\theta) = \mathbb{E}_{y \sim p_\theta(y)} [\mathcal{W}_\delta(p_\theta(\mathbf{x}|y) \| p_{\text{data}}(x))], \quad (D60)$$

and the one-vs-one variant simply replaces the data distribution with another cluster distribution on which to perform the expectation as well. The distance between the samples is noted δ . During training, the model does not see continuous distribution and only gets batches of samples. Hence, the problem is discretised and the Wasserstein distance can be then evaluated using histogram vectors. We

Table 8 Performances of Sparse GEMINI using the one-vs-all objectives on the synthetic datasets

	MMD-GEMINI		Wasserstein-GEMINI	
	Logistic	MLP	Logistic	MLP
(a) ARI scores (greater is better)				
S1	0.08 _{0,08}	0.11 _{0.12}	0.03 _{0,05}	0.08 _{0,09}
S2	0.46 _{0.11}	0.47 _{0.11}	0.46 _{0.13}	0.41 _{0,15}
S3*	0.23 _{0.07}	0.22 _{0.05}	0.13 _{0,07}	0.08 _{0,06}
S4	0.43 _{0,05}	0.45 _{0,05}	0.56 _{0,17}	0.74 _{0.22}
S5	0.44 _{0,05}	0.58 _{0.11}	0.52 _{0,16}	0.47 _{0,18}
D2	0.53 _{0,06}	0.55 _{0,03}	0.57 _{0.02}	0.54 _{0,03}
(b) VSER scores (lower is better)				
S1	0.38 _{0.14}	0.37 _{0.11}	0.58 _{0,13}	0.55 _{0,10}
S2*	0.04 _{0.04}	0.10 _{0,06}	0.16 _{0,08}	0.25 _{0,11}
S3*	0.06 _{0.08}	0.19 _{0,11}	0.20 _{0,13}	0.68 _{0,16}
S4	0.00 _{0.00}	0.00 _{0.00}	0.03 _{0,03}	0.00 _{0.00}
S5	0.00 _{0.00}	0.00 _{0.00}	0.02 _{0,02}	0.08 _{0,07}
D2	0.29 _{0,00}	0.30 _{0,04}	0.29 _{0,00}	0.29 _{0,03}
(c) CVR scores (greater is better)				
S1	0.56 _{0,20}	0.63 _{0,31}	0.65 _{0,21}	0.71 _{0.17}
S2	0.93 _{0.10}	0.84 _{0,12}	0.91 _{0.10}	0.81 _{0,12}
S3	0.94 _{0,16}	0.98 _{0.06}	0.66 _{0,23}	0.96 _{0.08}
S4	1.00 _{0.00}	1.00 _{0.00}	0.95 _{0,09}	0.99 _{0,04}
S5	1.00 _{0.00}	1.00 _{0.00}	0.94 _{0,09}	0.95 _{0,09}
D2	0.00 _{0,00}	0.00 _{0,00}	0.00 _{0,00}	0.00 _{0,00}

Bold value indicates the best results

demonstrated (Ohl et al. 2023) that these histogram vectors consist in a cluster-wise normalisation of the predictions which arises from importance sampling. Thus, the discrete approximation of the Wasserstein GEMINI is:


$$\hat{\mathcal{I}}_{\mathcal{W}_\delta}^{\text{ova}}(\mathbf{x}; y|\theta) = \sum_{k=1}^K \pi_k \min_{\mathbf{P} \in U(\boldsymbol{\omega}_k, \mathbf{1}_N/N)} \sum_{i=1}^N \mathbf{P}_{i,j} \delta(\mathbf{x}_i, \mathbf{x}_j), \quad (D61)$$

where \mathbf{P} is constrained in a set that forces it to have rows summing to the values of $\boldsymbol{\omega}_k$ and columns summing to $\mathbf{1}_N/N$. The vector $\boldsymbol{\omega}_k = \boldsymbol{\tau}_{.k} / \sum_{i=1}^N \tau_{ik}$ is the normalised cluster predictions.


D.1 Gradient for the Wasserstein distance

The new formulation of the discrete Wasserstein distance corresponds to a linear program and is often referred to as the Kantorovich problem. This problem admits the following

Fig. 7 Box plots of one-vs-one methods against baseline on the dataset S1




(a) ARI




(b) VSER

Fig. 8 Box plots of one-vs-one methods against baseline on the dataset S2



(a) ARI




(b) VSER

dual (Peyré and Cuturi 2019):


$$\mathcal{W}_\delta(\omega_1 \| \omega_2) = \max_{\substack{(\mathbf{u}, \mathbf{v}) \in \mathbb{R}^N \times \mathbb{R}^N \\ \mathbf{u}_i + \mathbf{v}_j \leq \delta_{ij}, \forall i, j \leq N}} \langle \mathbf{u}, \omega_1 \rangle + \langle \mathbf{v}, \omega_2 \rangle, \quad (\text{D62})$$

thanks to the strong duality for linear programs (Bertsimas and Tsitsiklis 1997, p 148, Theorem 4.4). It immediately appears that once we found the optimal "Kantorovich potentials" \mathbf{u}^* and \mathbf{v}^* for each respective histogram vector ω_1 and ω_2 we can compute the gradient of the distance using these optimal values because we remove the max term. However,

Fig. 9 Box plots of one-vs-one methods against baseline on the dataset S3




(a) ARI




(b) V-SER

Fig. 10 Box plots of one-vs-one methods against baseline on the dataset S4




(a) ARI




(b) VSER

Fig. 11 Box plots of one-vs-one methods against baseline on the dataset S5



(a) ARI



(b) VSER

as we want to remain in the simplex, we need to recenter the mass of a gradient and thus subtract the mean of the dual variables:

$$\frac{\partial \mathcal{W}_\delta(\omega_1 \| \omega_2)}{\partial \omega_1} = \mathbf{u}^* - \sum_{i=1}^N \frac{\mathbf{u}_i^*}{N} = \bar{\mathbf{u}}, \tag{D63}$$

$$\frac{\partial \mathcal{W}_\delta(\omega_1 \| \omega_2)}{\partial \omega_2} = \mathbf{v}^* - \sum_{i=1}^N \frac{\mathbf{v}_i^*}{N} = \bar{\mathbf{v}}. \tag{D64}$$

D.2 Complete gradient for the OvA Wasserstein

We can now simply unfold the rules of derivation w.r.t. τ_{ik} between the product of the cluster proportions π_k and the Wasserstein distance. However, we must take into account that due to the self-normalisation of τ to produce the histogram vectors ω , we have to sum its derivative over all normalised samples. Thus:

$$\begin{aligned} \frac{\partial \hat{\mathcal{I}}}{\partial \tau_{ik}} &= \sum_{k'=1}^K \mathcal{W}_\delta(\omega_{\cdot k'} \| \mathbf{1}_N / N) \frac{\partial \pi_{k'}}{\partial \tau_{ik}} \\ &+ \sum_{j=1}^N \pi_{k'} \frac{\partial \mathcal{W}_\delta(\omega_{\cdot k'} \| \mathbf{1}_N / N)}{\partial \omega_{jk'}} \frac{\partial \omega_{jk'}}{\partial \tau_{ik}}, \end{aligned} \tag{D65}$$

$$\begin{aligned} &= \frac{\mathcal{W}_\delta(\omega_{\cdot k} \| \mathbf{1}_N)}{N} \\ &+ \pi_k \sum_{j=1}^N \bar{\mathbf{u}}_{jk} \left(\frac{\mathbb{1}[i == j]}{N \pi_k} - \frac{\tau_{jk}}{N^2 \pi_k^2} \right), \end{aligned} \tag{D66}$$

where $\mathbb{1}$ is the indicator function resulting from the derivative of the self-normalisation. After summing over all samples, we can conclude that the gradient of the one-vs-all Wasserstein GEMINI w.r.t. model predictions τ is:

$$\frac{\partial \hat{\mathcal{I}}_{\mathcal{W}_\delta}^{\text{ova}}}{\partial \tau_{ik}} = \frac{\mathcal{W}_\delta(\omega_{\cdot k} \| \mathbf{1}_N)}{N} + \frac{\bar{\mathbf{u}}_{ik}}{N} - \frac{\langle \bar{\mathbf{u}}_{\cdot k}, \boldsymbol{\tau}_k \rangle}{N^2 \pi_k}. \tag{D67}$$

D.3 Complete gradient for the OvO Wasserstein

The demonstration follows the same rules as before. We add as well the fact that the Wasserstein distance is symmetric, and hence its gradient is as well so we can permute the names $\bar{\mathbf{u}}$ and $\bar{\mathbf{v}}$ when changing $\mathcal{W}_\delta(\omega_1 \| \omega_2)$ for $\mathcal{W}_\delta(\omega_2 \| \omega_1)$. Therefore, we sum twice the gradients of the Wasserstein distances, as well as twice the gradients for the proportions due to the symmetric nature of this function. We can thus arrive to a final gradient that is very similar to the OvO scenario with an additional summing over adversarial proportions:

$$\frac{\partial \hat{\mathcal{I}}_{\mathcal{W}_\delta}^{\text{ovo}}}{\partial \tau_{ik}} = \sum_{k'=1}^K 2 \frac{\pi_k \mathcal{W}_\delta(\omega_{\cdot k} \| \omega_{\cdot k'})}{N} + 2 \frac{\pi_{k'} \bar{\mathbf{u}}_{j,k/k'}}{N}$$

Listing 1 An example of *gemclus* loading and data clustering

```

1 from gemclus.sparse import
   SparseLinearMMD
2 from sklearn.datasets import
   load_breast_cancer
3 # load data
4 X, _ = load_breast_cancer(return_X_y=
   True)
5 # Create simple logistic regression
   model and do clustering
6 y_pred = SparseLinearMMD(n_clusters
   =2).fit_predict(X)
    
```

$$-2 \frac{\langle \bar{\mathbf{u}}_{\cdot k/k'}, \boldsymbol{\tau}_k \rangle}{N^2 \pi_k}. \tag{D68}$$

Notice that we detailed in subscript for which Wasserstein evaluation a dual variable emerges using the notation k/k' . Since the one-vs-one GEMINI makes K^2 distance evaluation, we have K^2 dual variables as well when removing duplicate dual variables due to symmetry.

Appendix E Examples of code snippets with the package GemClus

GemClus is implemented to respect as much as possible the scikit-learn Pedregosa et al. (2011) naming conventions. For instance, Listing 1 shows how a the logistic regression can be trained with an MMD GEMINI on the breast cancer dataset. The Listing 2 shows how the numerical experiments from Sect. 5.3 can be run.

We also want to extend this package to other discriminative clustering methods for potentially small-scale datasets. Therefore, we include an implementation of the regularised mutual information (RIM) model by Krause et al. (2010) as shown in Listing 3 because we consider this model to be one of the very first proposed in the domain yet find few satisfying implementations. In this sense, we included as well (and as can be noted in Listing 2 functions for generating relevant synthetic datasets for clustering.

Appendix F Results on the synthetic datasets for the one-vs-all GEMINIs

We provide here a complement on the numerical experiments from Sect. 5.3 with the performances of Sparse GEMINI when using the one-vs-all GEMINIs in Table 8.

Listing 2 An example of sparse GEMINI model fitting the 5th scenario of the synthetic datasets

```

1 from gemclus.sparse import
   SparseMLPMMD
2 from gemclus.data import celeux_one
3
4 # Generate the data according to the
   5th scenario
5 X,y = celeux_one(n=300, p=95, mu=1.7)
6 # Prepare the model: MLP with the OvA
   MMD-GEMINI for 3 clusters
7 model = SparseMLPMMD(n_clusters=3)
8 # Progressively increase the penalty
   until all features are removed
9 # res contains the history of feature
   selection and best model weights
10 res = model.path(X)

```

Listing 3 The package *gemclus* incorporates as well the basic logistic regression with regularised mutual information by Krause et al. (2010)

```

1 from gemclus.linear import RIM
2 y_pred = RIM(n_clusters=3).
   fit_predict(X)

```

Appendix G Distribution of numerical experiments scores

We provide here additional details regarding the distribution of the ARI and VSER scores from Table 2 for the one-vs-one models in static training. These scores are displayed with box plots respectively in Figs. 7, 8, 9, 10 and 11 for the scenarios 1, 2, 3, 4, 5 of the first synthetic dataset.

Acknowledgements This work has been supported by the French government, through the 3IA Côte d'Azur, Investment in the Future, project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002. We would also like to thank the France Canada Research Fund (FFCR) for their contribution to the project. This work was partly supported by EU Horizon 2020 project AI4Media, under contract no. 951911. Finally, it was also supported by the Health Data Hub.

Author Contributions L.O. wrote the main manuscript text. All other authors reviewed the manuscript and provided guidance for methods.

Funding Author L.Ohl received funding from the French government, through the 3IA Côte d'Azur, Investment in the Future, project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002. This work benefited as well from the EU Horizon 2020 project AI4Media, under contract no. 951911. This project was as well funded by the France Canada Research Fund (FFCR) to all authors of the paper except C. Bouveyron.

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no Conflict of interest.

References

- Abeshouse, A., Ahn, J., Akbani, R., Ally, A., Amin, S., Andry, C.D., Annala, M., Aprikian, A., Armenia, J., Arora, A., et al.: The molecular taxonomy of primary prostate cancer. *Cell* **163**(4), 1011–1025 (2015)
- Alelyani, S., Tang, J., Liu, H.: Feature selection for clustering: a review. *Data Clustering* 29–60 (2018)
- Andrews, J.L., McNicholas, P.D.: Variable Selection for Clustering and Classification. arXiv preprint (2013)
- Andrews, J.L., McNicholas, P.D.: Variable selection for clustering and classification. *J. Classif.* **31**(2), 136–153 (2014)
- Bach, F., Jenatton, R., Mairal, J., Obozinski, G., et al.: Optimization with sparsity-inducing penalties. *Found. Trends Mach. Learn.* **4**(1), 1–106 (2012)
- Bertsimas, D., Tsitsiklis, J.N.: *Introduction to Linear Optimization*, vol. 6. Athena Scientific, Belmont (1997)
- Blum, A.L., Langley, P.: Selection of relevant features and examples in machine learning. *Artif. Intell.* **97**(1–2), 245–271 (1997)
- Bouveyron, C., Brunet, C.: Simultaneous model-based clustering and visualization in the fisher discriminative subspace. *Stat. Comput.* **22**(1), 301–324 (2012)
- Bouveyron, C., Brunet-Saumard, C.: Discriminative variable selection for clustering with the sparse fisher-em algorithm. *Comput. Stat.* **29**(3), 489–513 (2014a). <https://doi.org/10.1007/s00180-013-0433-6>. (ISSN 1613-9658)
- Bouveyron, C., Brunet-Saumard, C.: Model-based clustering of high-dimensional data: a review. *Comput. Stat. Data Anal.* **71**, 52–78 (2014b). <https://doi.org/10.1016/j.csda.2012.12.008>. (ISSN 0167-9473)
- Bridle, J., Heading, A., MacKay, D.: Unsupervised Classifiers, Mutual Information and 'Phantom Targets'. In: Moody, J., Hanson, S., Lippmann, R.P. (eds.) *Advances in Neural Information Processing Systems*, vol. 4. Morgan-Kaufmann (1992)
- Brown, G.: Diversity in Neural Network Ensembles. PhD Thesis, University of Birmingham (2004)
- Celeux, G., Martin-Magniette, M.-L., Maugis-Rabusseau, C., Raftery, A.E.: Comparing model selection and regularization approaches to variable selection in model-based clustering. *Journal de la société française de statistique* **155**(2), 57–71 (2014)
- Chen, H., Wang, W., Feng, X., He, R.: Discriminative and coherent subspace clustering. *Neurocomputing* **284**, 177–186 (2018)
- Congressional Quarterly Almanac. 98th Congress. 2nd session, 40 (1984)
- Cover, T.M.: *Elements of Information Theory*. Wiley (1999)
- Dy, J.G.: Unsupervised feature selection. In: *Computational methods of feature selection*, pp. 35–56. Chapman and Hall/CRC (2007)
- Fop, M., Murphy, T.B.: Variable selection methods for model-based clustering. *Stat. Surv.* **12**(none), 18–65 (2018). <https://doi.org/10.1214/18-SS119>
- França, G., Rizzo, M.L., Vogelstein, J.T.: Kernel k-groups via Hartigan's method. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**(12), 4411–4425 (2020). (Publisher: IEEE)
- Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A.: A kernel two-sample test. *J. Mach. Learn. Res.* **13**(1), 723–773 (2012)

- Hastie, T., Tibshirani, R., Wainwright, M.: Statistical learning with sparsity. *Monographs Stat. Appl. Prob.* **143**, 143 (2015)
- He, X., Cai, D., Niyogi, P.: Laplacian score for feature selection. *Adv. Neural Inf. Process. Syst.* **18** (2005)
- Hennig, C.: What are the true clusters? *Pattern Recogn. Lett.* **64**, 53–62 (2015). <https://doi.org/10.1016/j.patrec.2015.04.009>
- Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors. (2012). arXiv preprint [arXiv:1207.0580](https://arxiv.org/abs/1207.0580)
- Hubert, L., Arabie, P.: Comparing partitions. *J. Classif.* **2**(1), 193–218 (1985)
- John, G.H., Kohavi, R., Pfleger, K.: Irrelevant Features and the Subset Selection Problem. In: *Machine Learning Proceedings 1994*, pp. 121–129. Elsevier (1994)
- Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization (2014). arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
- Kong, Y., Deng, Y., Dai, Q.: Discriminative clustering and feature selection for brain mri segmentation. *IEEE Signal Process. Lett.* **22**(5), 573–577 (2015). <https://doi.org/10.1109/LSP.2014.2364612>
- Krause, A., Perona, P., Gomes, R.: Discriminative clustering by regularized information maximization. In: Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., Culotta, A. (eds.) *Advances in Neural Information Processing Systems*, vol. 23. Curran Associates Inc. (2010)
- Larochelle, H., Erhan, D., Courville, A., Bergstra, J., Bengio, Y.: An empirical evaluation of deep architectures on problems with many factors of variation. In: *Proceedings of the 24th International Conference on Machine Learning*, pp. 473–480 (2007)
- Leek, J.T., Scharpf, R.B., Bravo, H.C., Simcha, D., Benjamin Langmead, W., Johnson, E., Geman, D., Baggerly, K., Irizarry, R.A.: Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat. Rev. Genet.* **11**(10), 733–739 (2010)
- Lemhadri, I., Ruan, F., Abraham, L., Tibshirani, R.: LassoNet: a neural network with feature sparsity. *J. Mach. Learn. Res.* **22**(127), 1–29 (2021)
- Long, Z.-Z., Xu, G., Du, J., Zhu, H., Yan, T., Yu, Y.-F.: Flexible subspace clustering: a joint feature selection and k-means clustering framework. *Big Data Res.* **23**, 100170 (2021). <https://doi.org/10.1016/j.bdr.2020.100170>
- Luo, J., Schumacher, M., Scherer, A., Sanoudou, D., Megherbi, D., Davison, T., Shi, T., Tong, W., Shi, L., Hong, H., et al.: A comparison of batch effect removal methods for enhancement of prediction performance using MAQC-II microarray gene expression data. *Pharmacogenomics J.* **10**(4), 278–291 (2010)
- Marbac, M., Sedki, M., Patin, T.: Variable selection for mixed data clustering: application in human population genomics. *J. Classif.* **37**(1), 124–142 (2020)
- Mattei, P.-A., Bouveyron, C., Latouche, P.: Globally sparse probabilistic pca. In: Gretton, A., Robert, C.C. (eds.) *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pp. 976–984, Cadiz, Spain, 09–11 May 2016. PMLR. <https://proceedings.mlr.press/v51/mattei16.html>
- Maugis, C., Celeux, G., Martin-Magniette, M.-L.: Variable selection for clustering with Gaussian mixture models. *Biometrics* **65**(3), 701–709 (2009)
- McLachlan, G.J., Bean, R.W., Peel, D.: A mixture model-based approach to the clustering of microarray expression data. *Bioinformatics* **18**(3), 413–422 (2002). <https://doi.org/10.1093/bioinformatics/18.3.413>
- Meng, C., Helm, D., Frejno, M., Kuster, B.: moCluster: identifying joint patterns across multiple omics data sets. *J. Proteome Res.* **15**(3), 755–765 (2016)
- Minka, T.: Discriminative Models, not Discriminative Training. Technical report, Technical Report MSR-TR-2005-144, Microsoft Research (2005)
- Ohl, L., Mattei, P.-A., Bouveyron, C., Harchaoui, W., Leclercq, M., Droit, A., Precioso, F.: Generalised mutual information for discriminative clustering. In: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A., editors, *Advances in Neural Information Processing Systems*, volume 35, pp. 3377–3390. Curran Associates, Inc., (2022). https://proceedings.neurips.cc/paper_files/paper/2022/file/16294049ed8de15830ac0b569b97f74a-Paper-Conference.pdf
- Ohl, L., Mattei, P.-A., Bouveyron, C., Harchaoui, W., Leclercq, M., Droit, A., Precioso, F.: Generalised mutual information: a framework for discriminative clustering. (2023). arXiv preprint [arXiv:2309.02858](https://arxiv.org/abs/2309.02858)
- Parsons, L., Haque, E., Liu, H.: Subspace clustering for high dimensional data: a review. *ACM SIGKDD Explor. Newsl.* **6**(1), 90–105 (2004)
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)
- Peng, C., Kang, Z., Yang, M., Cheng, Q.: Feature selection embedded subspace clustering. *IEEE Signal Process. Lett.* **23**(7), 1018–1022 (2016)
- Peyré, G., Cuturi, M.: Computational optimal transport: with applications to data science. *Found. Trends Mach. Learn.* **11**(5–6), 355–607 (2019)
- Raftery, A.E., Dean, N.: Variable selection for model-based clustering. *J. Am. Stat. Assoc.* **101**(473), 168–178 (2006)
- Ramazzotti, D., Lal, A., Wang, B., Batzoglu, S., Sidow, A.: Multi-omic tumor data reveal diversity of molecular mechanisms that correlate with survival. *Nat. Commun.* **9**(1), 1–14 (2018)
- Reese, S.E., Archer, K.J., Therneau, T.M., Atkinson, E.J., Vachon, C.M., De Andrade, M., Kocher, J.-P.A., Eckel-Passow, J.E.: A new statistic for identifying batch effects in high-throughput genomic data that uses guided principal component analysis. *Bioinformatics* **29**(22), 2877–2883 (2013)
- Saúl Solorio-Fernández, J., Carrasco-Ochoa, A., Martínez-Trinidad, J.F.: A review of unsupervised feature selection methods. *Artif. Intell. Rev.* **53**(2), 907–948 (2020)
- Scruella, L., Raftery, A.E.: clustvarsel: a package implementing variable selection for Gaussian model-based clustering in R. *J. Stat. Soft.* **84**(1), 1–28 (2018). <https://doi.org/10.18637/jss.v084.i01>
- Shen, R., Mo, Q., Schultz, N., Seshan, V.E., Olshen, A.B., Huse, J., Ladanyi, M., Sander, C.: Integrative subtype discovery in glioblastoma using iCluster. *PLoS ONE* **7**(4), e35236 (2012)
- Tadesse, M.G., Sha, N., Vannucci, M.: Bayesian variable selection in clustering high-dimensional data. *J. Am. Stat. Assoc.* **100**(470), 602–617 (2005)
- Vandewalle, V.: Multi-partitions subspace clustering. *Mathematics* **8**(4), 597 (2020)
- Varshavsky, R., Gottlieb, A., Linial, M., Horn, D.: Novel unsupervised feature filtering of biological data. *Bioinformatics* **22**(14), e507–e513 (2006). <https://doi.org/10.1093/bioinformatics/bt1214>. (ISSN 1367-4803)
- Vittrant, B., Leclercq, M., Martin-Magniette, M.-L., Collin, C., Bergeron, A., Fradet, Y., Droit, A.: Identification of a transcriptomic prognostic signature by machine learning using a combination of small cohorts of prostate cancer. *Front. Gene.* (2020). <https://doi.org/10.3389/fgene.2020.550894>
- von Luxburg, U.: A tutorial on spectral clustering. *Stat. Comput.* **17**(4), 395–416 (2007). <https://doi.org/10.1007/s11222-007-9033-z>. (ISSN 1573-1375)
- Witten, D.M., Tibshirani, R., Witten, M.D.: Package ‘sparcl’ (2013)
- Witten, D.M., Tibshirani, R.: A framework for feature selection in clustering. *J. Am. Stat. Assoc.* **105**(490), 713–726 (2010)

- Zhao, Z., Liu, H.: Spectral feature selection for supervised and unsupervised learning. In: Proceedings of the 24th International Conference on Machine Learning, pp. 1151–1157 (2007)
- Zografos, V., Ellis, L., Mester, R.: Discriminative subspace clustering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2107–2114 (2013)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.