



HAL
open science

Sparse GEMINI 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 GEMINI for joint discriminative clustering and feature selection. 2023. hal-04195208

HAL Id: hal-04195208

<https://hal.science/hal-04195208>

Preprint submitted on 4 Sep 2023

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.

Sparse GEMINI

for joint discriminative clustering and feature selection

Louis Ohl^{1,2,3,4}

Pierre-Alexandre Mattei^{3,5}

Charles Bouveyron^{1,3,5}

Mickaël Leclercq⁴

Arnaud Droit⁴

Frédéric Precioso^{1,2,3}

1. Université Côte d'Azur 2. I3S, CNRS 3. Inria, Maasai team 4. CHU de Québec Research Centre, Laval University
5. LJAD, CNRS

{louis.ohl, pierre-alexandre.mattei}@inria.fr

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 $p(\mathbf{x})$, 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 clustering model $p_\theta(y|\mathbf{x})$. We demonstrate the performances of Sparse GEMINI on synthetic datasets as well as 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.

look relevant yet 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. Solutions may thus be hardly scalable to high-dimensional data (Raftery & Dean, 2006) or to the number of samples (Witten & Tibshirani, 2010) when the selection process is part of the model.

Therefore reducing the number of variables on which to learn to a relevant few is of interest, notably in terms of interpretation (Fop & 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 & Dean, 2006) or number of samples (Witten & 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) for a scalable discriminative clustering with penalised feature selection. The contributions of Sparse GEMINI are:

1. Introduction

It is common that clustering algorithms as well as supervised models rely on all available features for the best performance. Yet, as datasets become high-dimensional, clustering algorithms tend to break under the curse of dimensionality (Bouveyron & Brunet-Saumard, 2014b). 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 whereas irrelevant variables do not bring any new knowledge regarding the cluster distribution (Tadesse et al., 2005) and redundant variables

- A simple novel algorithm efficiently combining feature selection and discriminative clustering.
- A scalable feature selection and clustering model compatible with deep learning architectures.
- Demonstrations of performances on multiple synthetic and real datasets as well as a large-scale transcriptomics dataset.

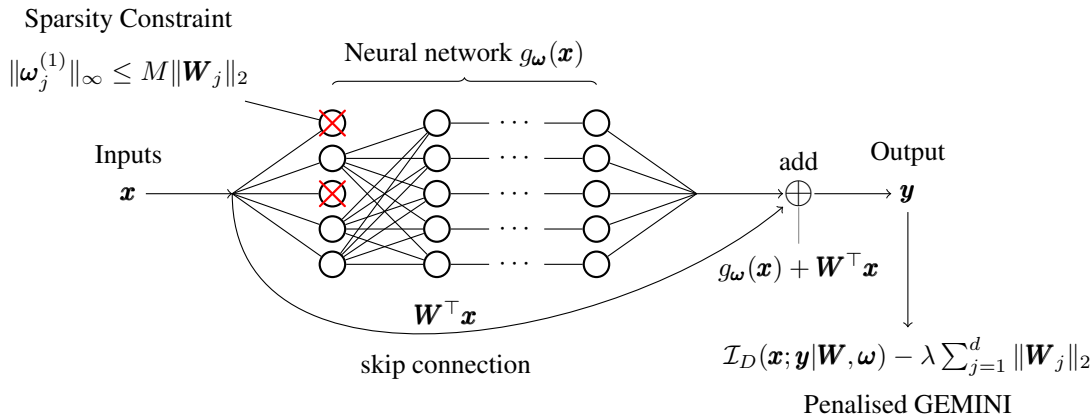


Figure 1. Description of the complete Sparse GEMINI model. Through a proximal gradient, clusters learned by GEMINI drop irrelevant features both in a skip connection and an MLP.

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 & Liu, 2007). Those methods are thus easily scalable and quick despite bearing the challenge of defining unsupervised feature interestingness (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 & Langley, 1997) as compromises between the two first categories.

While the definition of the 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 & Murphy, 2018). Often, the terms relevant variables, irrelevant variables (Tadesse et al., 2005) for the notion of conveying information are used. Others may consider as well redundant variables as those that bring already available information (Maugis et al., 2009). A key difference in models would then be to consider whether the informative variables are independent given the cluster assignment (local independence) or dependent (global independence from the uninformative variables), yet the latter hardly accounts for redundant variables (Fop & Murphy, 2018).

Feature selection is to be not 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, f.e. using matrix factorisation (Shen et al., 2012). Moreover, by enforcing the projection matrix to be sparse, feature selection can be recovered in the original space (Bouveyron & 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), notably 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 clusterings which is challenging as well while we are interested in a method that selects features while producing a single clustering output.

Finally, models for clustering in feature selection are often model-based (Scrucca & Raftery, 2018; Raftery & Dean, 2006; Maugis et al., 2009), implying that they assume a parametric mixture model that can either explain the distribution of the data, as well as the distribution of the irrelevant variables. To perform well, these methods need a good selection criterion to compare models with one another (Raftery & Dean, 2006; Marbac et al., 2020; Maugis et al., 2009). To the best of our knowledge, there does 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 $p_\theta(y|\mathbf{x})$. Finally, most of these generative wrapper methods hardly scale both in 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., 2022) with the LassoNet framework for feature selection (Lemhadri et al., 2021) in neural networks. The model

is summarised in Figure 1.

3.1. The GEMINI objective

Let $\mathcal{D} = \{\mathbf{x}_{i=1}^N\}_{j=1}^N \subset \mathcal{X}$ a dataset of N observations, each of dimension d . We note each feature $\mathbf{x}^j \in \mathcal{X}_j$, 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:

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

where $f_\theta : \mathcal{X} \mapsto \mathbb{R}^K$ has parameters θ . 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., 2022). This objective was introduced to circumvent the need for parametric assumptions regarding $p(x)$ in clustering and thus leads to designing only a discriminative clustering model $p_\theta(y|\mathbf{x})$. With the help of Bayes theorem, this objective can be estimated without knowledge of the data distribution $p(\mathbf{x})$ using only the output of the clustering distribution $p_\theta(y|\mathbf{x})$. Overall, the GEMINI aims at separating according to a distance D the cluster distributions from either the data distribution (one-vs-all):

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

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

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

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 & Cuturi, 2019). The former corresponds to the distance between the expectations of the respective distributions projected into a Hilbert space and the second 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. Thus, 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.

3.2. The LassoNet architecture

To perform variable selection inside the neural network, we chose to adapt 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:

$$\mathcal{F} = \{f_\theta : \mathbf{x} \mapsto g_\omega(\mathbf{x}) + \mathbf{W}^\top \mathbf{x}\}, \quad (4)$$

with $\theta = \{\omega, \mathbf{W}\}$ including ω the parameters of the MLP and $\mathbf{W} \in \mathbb{R}^{K \times d}$ the weights of a linear skip connection penalised by ℓ_1 , similarly to the Lasso (Tibshirani, 1996). However, to properly ensure that an entire vector weights is eliminated at once, a group-lasso penalty is preferred (Hastie et al., 2015, Section 3.3.3) also known as ℓ_1/ℓ_2 penalty (Bach et al., 2012). Thus, the optimal parameters should satisfy:

$$\hat{\theta} = \operatorname{argmax}_\theta \mathcal{I}_D(\theta) - \lambda \sum_{j=1}^d \|\mathbf{W}_j\|_2, \quad (5)$$

with $\mathbf{W}_j \in \mathbb{R}^K$, the j -th column of \mathbf{W} . Notice that λ is positive because we seek to simultaneously maximise the GEMINI and minimise the ℓ_1/ℓ_2 penalty. During training, the sparsity-induced linear parameter \mathbf{W} will remove some feature subset I . In order to force the MLP to drop this same subset of features as well, the weights of the first layer $\omega^{(1)}$ are constrained such that:

$$\|\omega_j^{(1)}\|_\infty \leq M \|\mathbf{W}_j\|_2, \forall j \leq d. \quad (6)$$

where M is called the hierarchy coefficient. When $M = 0$, the method is equivalent to a penalised logistic regression. Thus, when a feature j is eliminated, all weights starting from this feature in the MLP will be equal to 0 as well. Lemhadri et al. (2021) gracefully provide a proximal gradient operation to satisfy this constraint during training time which guarantees true zeros in the first MLP layer and the skip connection.

Interestingly, while the constraints are designed to specifically select features in the dataset, 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.

3.3. 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 an arbitrary threshold $0 < F_{\text{thres}} < d$ which can be thought as the minimum number of useful variables required. Each time the number of features decrease during training, we save its associate intermediate model

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 range of 90% of the best GEMINI value. This best value is most of the time the loss evaluated with the model exploiting

Table 1. Brief description of datasets involved in experiments

Name	Samples	Features	#Classes
US-Congress	435	16	2
Heart-statlog	270	13	2
MNIST	12000	784	10
MNIST-BR	12000	784	10
Prostate-BCR	171	24508	2

all features. We propose as well a less grounded yet efficient training mode in appendix A.

4. Experiments

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

4.1. Metrics

Depending on the experiments for comparison purposes, we report 3 different metrics. The adjusted rand index (ARI, [Hubert & 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, thus 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 $1 - \frac{d'}{d}$. All metrics are written in percentage form.

4.2. 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”. The optimiser for the initial training step with $\lambda = 0$ is Adam ([Kingma & 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 like [Lemhadri et al. \(2021\)](#). 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.

4.3. Numerical experiments

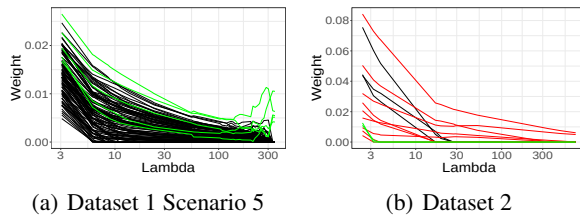


Figure 2. Example of convergence of the norm of the weights of the skip connection for every feature during training for the Wasserstein OVA objective. Green lines are the informative variables, black lines are the noise and red are the correlated variables. In the case of noisy variables, Sparse GEMINI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

We experimented Sparse GEMINI on two synthetic datasets proposed by [Celeux et al. \(2014\)](#) and also used by ([Bouveyron & 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 described along results in Table 2.

The second dataset consists of $n = 2000$ samples of 14 variables, 2 of them being informative and most others being 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:

$$\mathbf{x}^{3-11} = [0, 0, 0.4, 0.8, 1.2, 1.6, 2.0, 2.4, 2.8]^\top + \mathbf{x}^{1-2^\top} \begin{bmatrix} 0.5 & 2 & 0 & -1 & 2 & 0.5 & 4 & 3 & 2 \\ 1 & 0 & 3 & 2 & -4 & 0 & 0.5 & 0 & 1 \end{bmatrix} + \epsilon, \quad (7)$$

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

Table 2. Performances of Sparse GEMINI using on synthetic datasets after 20 runs. 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 parentheses

Method		Sparse KMeans	Clustvarsel	vscc	SFEM	MMD		Wasserstein	
						OvA	OvO	OvA	OvO
S1	ARI	25 (5.9)	9.4 (0)	-0.8 (0)	17 (1.5)	23 (5.9)	22 (6.8)	8.8 (7.1)	8.8 (8.4)
$N = 30$	VSER	30 (20)	28 (0)	80 (0)	27 (3.6)	29 (11)	32 (11)	54 (9.8)	47 (8.4)
$\alpha = 0.6$	CVR	59 (28)	0 (0)	100 (0)	40 (0)	63 (15)	68 (15)	79 (22)	76 (20)
$p = 20$	# Var	8.4 (7.9)	2.0 (0)	25 (0)	5.8 (0.9)	8.5 (3.1)	9.8 (3.0)	16.4 (2.0)	14.2 (3.2)
S2	ARI	82 (0)	9.4 (0)	-0.8 (0)	90 (0)	49 (4.1)	89 (8.7)	54 (18)	56 (14)
$N = 30$	VSER	80 (0)	28 (0)	80 (0)	40 (0)	8.2 (6.0)	0 (0)	18 (12)	14 (9.9)
$\alpha = 1.7$	CVR	100 (0)	0(0)	100 (0)	20 (0)	78 (16)	100 (0)	83 (18)	81 (18)
$p = 20$	# Var	25 (0)	2 (0)	25 (0)	7 (0)	4.8 (0.37)	5 (0)	7.8 (2.9)	6.5 (2.0)
S3	ARI	9.1 (0.1)	0.5 (0)	24 (0)	19 (0)	21 (2.0)	20 (2.4)	9.1 (5.2)	11 (4.7)
$N = 300$	VSER	80 (0)	24 (0)	80 (0)	18 (0)	21 (7.9)	20 (8.8)	67 (9.5)	49 (18)
$\alpha = 0.6$	CVR	100 (0)	20 (0)	100 (0)	33 (9.8)	99 (4.5)	100 (0)	96 (11)	85 (19)
$p = 20$	# Var	25 (0)	3.0 (0)	25 (0)	2.8 (0.64)	10.2 (2.2)	9.9 (2.2)	21.4 (2.8)	15.7 (5.3)
S4	ARI	86 (0)	87 (0)	50 (0)	86 (0)	50 (5.77)	86 (0.6)	80 (12)	81 (12)
$N = 300$	VSER	80 (0)	4 (0)	80 (0)	24 (0)	0 (0)	0 (0)	0.8 (2.5)	0.4 (1.2)
$\alpha = 1.7$	CVR	100 (0)	100 (0)	100 (0)	60 (0)	100 (0)	100 (0)	100 (0)	100 (0)
$p = 20$	# Var	25 (0)	6 (0)	25 (0)	7 (0)	5 (0)	5 (0)	5.2 (0.6)	5.1 (0.3)
S5	ARI	86 (0)	87 (0)	0 (0)	86 (0)	77 (7.63)	86 (0.5)	58 (19)	74 (16)
$N = 300$	VSER	95 (0)	1 (0)	95 (0)	12 (0)	0 (0)	0 (0)	5.6 (6.4)	0.8 (2.0)
$\alpha = 1.7$	CVR	100 (0)	100 (0)	100 (0)	60 (0)	1 (0)	100 (0)	95 (11)	97 (7.3)
$p = 95$	# Var	100 (0)	6 (0)	100 (0)	13 (0)	5 (0)	5 (0)	10.1 (5.9)	5.5 (1.4)
D2	ARI	30 (0)	60 (0)	54 (0)	57 (0)	56 (4.2)	55 (2.8)	56 (2.2)	55 (3.0)
	VSER	86 (0)	0 (0)	71 (0)	50 (0)	29 (0)	29 (0)	30 (3.7)	29 (0)
	CVR	100 (0)	100 (0)	100 (0)	100 (0)	0 (0)	0 (0)	0 (0)	0 (0)
	# Var	14 (0)	2 (0)	12 (0)	9 (0)	2.1 (0.3)	2.1 (0.3)	2.2 (0.5)	2 (0)

$$\Omega = \text{diag} \left(\mathbf{I}_3, 0.5\mathbf{I}_2, \text{diag}([1, 3])\text{Rot}\left(\frac{\pi}{3}\right), \text{diag}[2, 6]\text{Rot}\left(\frac{\pi}{6}\right) \right). \quad (8)$$

Finally, the last 3 variables are independently sampled 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. We compare our results against our own runs of other methods using their R package: SparseKMeans (Witten et al., 2013), ClustVarSel (Scrucca & Raftery, 2018), vscc (Andrews & McNicholas, 2013; 2014) and SparseFisherEM (Bouveyron & Brunet, 2012).

It appears that the Sparse GEMINI is efficient in selecting the relevant variables when several others are noisy, especially with the MMD-OvO objective. Moreover, while we do not systematically get the best ARI, our performances never fall far behind the most competitive method. We can observe as well that the MMD objective learns well despite the presence of few samples in scenarios 2 and 3. 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. It also appears that we performed poorly at selecting the correct variables in 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 top-selected variables by our models were the 6th and the 8th variables. We focus on this difference

of convergence in Figure 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. Overall, Clustvarsel (Scrucca & Raftery, 2018) performed better on this type of synthetic dataset in terms of variable selections because it explicitly assumes linear dependency between relevant variables and others.

4.4. Examples on MNIST and variations

We demonstrate as well performances 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 use as well 5% of dropout inside an MLP with 2 hidden layers of 1200 dimensions each (Hinton et al., 2012). We report in Figure 3 the selected features by the clustering algorithms and the evolution of the ARI. We extended this experiment as well to the variations of MNIST (Larochelle et al., 2007) by showing the performances on the MNIST-BR dataset¹. The former consists in samples of MNIST with the black background being replaced by uniform noise hence displaying conditional noise on the data whereas the latter replaces that background by real images. 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 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 figures 3(b) and 3(d), the respective optimal 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., 2022) which contributed here to lowering the ARI when evaluating with the true digits targets.

4.5. Real datasets

4.5.1. 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

¹Datasets were available at <https://web.archive.org/web/20180519112150/http://www.iro.umontreal.ca/~lisa/twiki/bin/view.cgi/Public/MnistVariations>

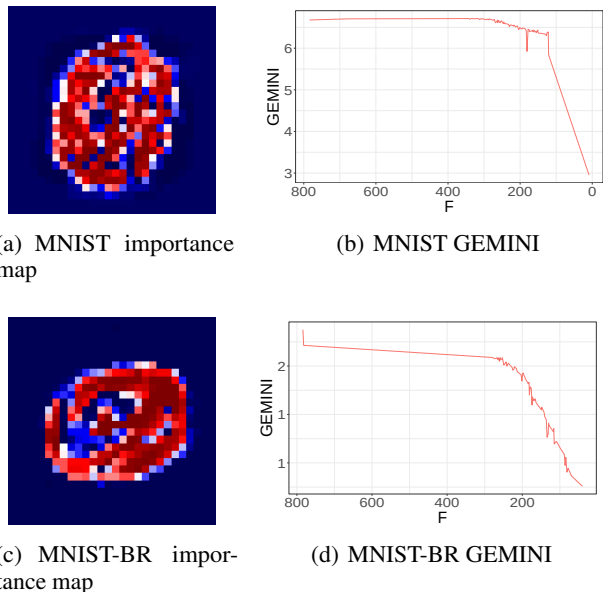



Figure 3. Relative importance of MNIST features after dynamic training of Sparse GEMINI with a log-scale color map. Blue features were eliminated at the first steps of λ and red features were eliminated last. On the right: evolution of the GEMINI depending on λ . F stands for the number of selected features.


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 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 Figure 5 for the US Congress dataset and in Figure 4 for Heart-statlog. Table 4 contains the performances for the US Congress dataset and Table 3 those of the Heart-statlog dataset. Both reports the average number of selected variables over 20 runs according to our postprocessing selection criterion. We added as well the performances of competitors from the previous section. However, we only managed to run Sparse Fisher EM on the Heart-statlog dataset, hence its presence only in Table 3. For comparison purposes, the best unsupervised accuracy reported on the Heart-statlog dataset in 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 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.



(a) GEMINI depending on selected feature number (b) Selected features depending on λ

Figure 4. Average training curves of Sparse GEMINI on the Heart Statlog dataset over 20 runs. Blue lines are Wasserstein, red lines are MMD.



(a) GEMINI Loss depending on selected feature number (b) Selected features depending on λ

Figure 5. Average training curves of Sparse GEMINI on the US Congress dataset over 50 runs. Blue lines are Wasserstein, red lines are MMD.

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 & Murphy, 2018). The MMD objective only added the Physician fee freeze vote to this subset. Regarding the heart 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.

4.5.2. PROSTATE-BCR DATASET

To show the scalability of Sparse GEMINI, we demonstrate its performance as well on the Prostate-BCR dataset, taken from (Vittrant et al., 2020) and publicly avail-

Table 3. ARI of Sparse GEMINI on the Heart-statlog dataset with the average number of selected features. Standard deviation in parentheses. Scores with an asterisk are the initial performances when using all features.

	ARI	# Variables
SparseKMeans	18.1 (0)	13 (0)
Clustvarsel	2.8 (0)	13 (0)
vsc	27 (0)	1 (0)
Sparse Fisher EM	19 (0)	1 (0)
MMD	32 (1.4)	8 (0)
Wasserstein	32 (8.8)	8.4 (2.7)
MMD*	37 (2.0)	13 (-)
Wasserstein*	33 (9.1)	13 (-)

Table 4. ARI of Sparse GEMINI on the US Congress dataset with the average number of selected features. Standard deviation in parentheses. Scores with an asterisk are the initial performances when using all features.

	ARI	# Variables
SparseKMeans	54 (0)	16 (0)
Clustvarsel	0.4 (0)	2 (0)
vsc	40 (0)	11 (0)
MMD	48 (0.2)	3.1 (0.08)
Wasserstein	47 (0)	2.0 (0)
MMD*	55 (0.7)	16 (-)
Wasserstein*	55 (1.7)	16 (-)

able at https://github.com/ArnaudDroitLab/prostate_BCR_prediction. This dataset is a combination of transcriptomics data from 3 different sources. Those are the Cancer Genom atlas (Abeshouse et al., 2015), the GSE54460 dataset from the NCBI website, and the PR-JEB6530 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 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 converging to 400 features or less, following (Vittrant et al., 2020). We trained Sparse GEMINI 5 times to find either 2 clusters or 3 clusters in order to break down possible substructures among the supervised targets. For the evaluation of the 3 clusters case, we binarised the results by mapping each cluster to the class in which it had the most samples

Table 5. ARI scores of the Prostate BCR dataset for various numbers of clusters depending on the chosen type of targets. We either use the expected targets (BCR) regarding cancer prediction, or data source targets that identify the data origin of each sample.

Objective		#Var	BCR targets ARI	Data source targets ARI
MMD	2 clusters	385 (11)	-0.5 (0)	79 (0)
	3 clusters	8293 (11308)	8.2 (0.6)	98 (2)
Wasserstein	2 clusters	381 (16)	-0.3 (0.2)	70 (6)
	3 clusters	10598 (13971)	5.3 (4.9)	84 (12)

Interestingly, 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 100% ARI in the case of the MMD GEMINI. Thus, the unsupervised algorithm was able to find sufficient differences in distribution between each source of data to discriminate them. We report these scores in Figure 5. Additionally, consistent subsets of features were always selected as the final subset on all 5 runs depending on the GEMINI. This implies that even without the best GEMINI within a range for feature selection, several runs can lead to identifying subsets of relevant data as well.

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. It has been demonstrated that batch effects in microarray experiments originated from multiple causes, including variations in the labelling and hybridization protocols used, which led to differences in the intensity of gene expression signals (Luo et al., 2010).

To minimise batch effects, it is important to control for variables such as reagents, protocols, and equipment used, and to use appropriate normalisation and data analysis methods to account for these variations. There are several approaches that 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. Discussion

Our first observation from Table 2 is that the Sparse GEMINI algorithm can reach performances 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 distribution. Specifically, the algorithm is good at discarding noisy variables, but less competitive regarding 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 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 figures 5(b) and 4(b) 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. Owing to the nature of multi-layered perceptrons, Sparse GEMINI is easily scalable to high-dimensional data and provides thus 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.

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.

References

- Abeshouse, A., Ahn, J., Akbani, R., Ally, A., Amin, S., Andry, C. D., Annala, M., Aprikian, A., Armenia, J., Arora, A., and others. The Molecular Taxonomy of Primary Prostate Cancer. *Cell*, 163(4):1011–1025, 2015.
- Alelyani, S., Tang, J., and Liu, H. Feature Selection for Clustering: A Review. *Data Clustering*, pp. 29–60, 2018.
- Almanac, C. Q. 98th Congress. *2nd session*, 40, 1984.
- Andrews, J. L. and McNicholas, P. D. Variable Selection for Clustering and Classification. *arXiv preprint*, 2013.
- Andrews, J. L. and McNicholas, P. D. Variable selection for clustering and classification. *Journal of Classification*, 31(2):136–153, 2014.
- Bach, F., Jenatton, R., Mairal, J., Obozinski, G., and others. Optimization with Sparsity-Inducing Penalties. *Foundations and Trends® in Machine Learning*, 4(1):1–106, 2012.
- Blum, A. L. and Langley, P. Selection of Relevant Features and Examples in Machine Learning. *Artificial intelligence*, 97(1-2):245–271, 1997.
- Bouveyron, C. and Brunet, C. Simultaneous Model-Based Clustering and Visualization in the Fisher Discriminative Subspace. *Statistics and Computing*, 22(1):301–324, 2012.
- Bouveyron, C. and Brunet-Saumard, C. Discriminative Variable Selection for Clustering with the Sparse Fisher-EM Algorithm. *Computational Statistics*, 29(3):489–513, June 2014a. ISSN 1613-9658. doi: 10.1007/s00180-013-0433-6.
- Bouveyron, C. and Brunet-Saumard, C. Model-Based Clustering of High-Dimensional Data: A Review. *Computational Statistics & Data Analysis*, 71:52–78, 2014b. ISSN 0167-9473. doi: <https://doi.org/10.1016/j.csda.2012.12.008>.
- Brown, G. *Diversity in Neural Network Ensembles*. PhD Thesis, University of Birmingham, 2004.
- Celeux, G., Martin-Magniette, M.-L., Maugis-Rabusseau, C., and 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., and He, R. Discriminative and Coherent Subspace Clustering. *Neurocomputing*, 284:177–186, 2018.
- Cover, T. M. *Elements of Information Theory*. John Wiley & Sons, 1999.
- Dy, J. G. Unsupervised Feature Selection. In *Computational methods of feature selection*, pp. 35–56. Chapman and Hall/CRC, 2007.
- Fop, M. and Murphy, T. B. Variable Selection Methods for Model-Based Clustering. *Statistics Surveys*, 12(none):18–65, 2018. doi: 10.1214/18-SS119.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A Kernel Two-Sample Test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Hastie, T., Tibshirani, R., and Wainwright, M. Statistical Learning with Sparsity. *Monographs on statistics and applied probability*, 143:143, 2015.
- He, X., Cai, D., and Niyogi, P. Laplacian Score for Feature Selection. *Advances in neural information processing systems*, 18, 2005.
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. R. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.
- Hubert, L. and Arabie, P. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985.
- John, G. H., Kohavi, R., and Pflieger, K. Irrelevant Features and the Subset Selection Problem. In *Machine learning proceedings 1994*, pp. 121–129. Elsevier, 1994.
- Kingma, D. P. and Ba, J. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- Krause, A., Perona, P., and Gomes, R. Discriminative Clustering by Regularized Information Maximization. In Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., and Culotta, A. (eds.), *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc., 2010.
- Larochelle, H., Erhan, D., Courville, A., Bergstra, J., and 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., Langmead, B., Johnson, W. E., Geman, D., Baggerly, K., and Irizarry, R. A. Tackling the Widespread and Critical Impact of Batch Effects in High-Throughput Data. *Nature Reviews Genetics*, 11(10):733–739, 2010.
- Lemhadri, I., Ruan, F., Abraham, L., and Tibshirani, R. LassoNet: A Neural Network with Feature Sparsity. *Journal of Machine Learning Research*, 22(127):1–29, 2021.
- Luo, J., Schumacher, M., Scherer, A., Sanoudou, D., Megherbi, D., Davison, T., Shi, T., Tong, W., Shi, L., Hong, H., and others. A Comparison of Batch Effect Removal Methods for Enhancement of Prediction Performance using MAQC-II Microarray Gene Expression Data. *The pharmacogenomics journal*, 10(4):278–291, 2010.
- Marbac, M., Sedki, M., and Patin, T. Variable Selection for Mixed Data Clustering: Application in Human Population Genomics. *Journal of Classification*, 37(1):124–142, 2020.
- Maugis, C., Celeux, G., and Martin-Magniette, M.-L. Variable Selection for Clustering with Gaussian Mixture Models. *Biometrics*, 65(3):701–709, 2009.
- Meng, C., Helm, D., Frejno, M., and Kuster, B. moCluster: Identifying Joint Patterns across Multiple Omics Data Sets. *Journal of proteome research*, 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., and Precioso, F. Generalised Mutual Information for Discriminative Clustering. In *Advances in Neural Information Processing Systems*, 2022.
- Parsons, L., Haque, E., and Liu, H. Subspace Clustering for High Dimensional Data: A Review. *Acm sigkdd explorations newsletter*, 6(1):90–105, 2004.
- Peyré, G. and Cuturi, M. Computational Optimal Transport: With Applications to Data Science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Raftery, A. E. and Dean, N. Variable Selection for Model-Based Clustering. *Journal of the American Statistical Association*, 101(473):168–178, 2006.
- Ramazzotti, D., Lal, A., Wang, B., Batzoglu, S., and Sidow, A. Multi-Omic Tumor Data Reveal Diversity of Molecular Mechanisms that Correlate with Survival. *Nature communications*, 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., and 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.
- Scrucca, L. and Raftery, A. E. clustvarsel: A Package Implementing Variable Selection for Gaussian Model-Based Clustering in R. *J. Stat. Soft.*, 84(1):1 – 28, April 2018. doi: 10.18637/jss.v084.i01.
- Shen, R., Mo, Q., Schultz, N., Seshan, V. E., Olshen, A. B., Huse, J., Ladanyi, M., and Sander, C. Integrative Subtype Discovery in Glioblastoma using iCluster. *PLoS one*, 7(4):e35236, 2012.
- Solorio-Fernández, S., Carrasco-Ochoa, J. A., and Martínez-Trinidad, J. F. A Review of Unsupervised Feature Selection Methods. *Artificial Intelligence Review*, 53(2): 907–948, 2020.
- Tadesse, M. G., Sha, N., and Vannucci, M. Bayesian Variable Selection in Clustering High-Dimensional Data. *Journal of the American Statistical Association*, 100(470): 602–617, 2005.
- Tibshirani, R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Vandewalle, V. Multi-Partitions Subspace Clustering. *Mathematics*, 8(4):597, 2020.
- Varshavsky, R., Gottlieb, A., Linial, M., and Horn, D. Novel Unsupervised Feature Filtering of Biological Data. *Bioinformatics*, 22(14):e507–e513, July 2006. ISSN 1367-4803. doi: 10.1093/bioinformatics/btl214.
- Vittrant, B., Leclercq, M., Martin-Magniette, M.-L., Collins, C., Bergeron, A., Fradet, Y., and Droit, A. Identification of a Transcriptomic Prognostic Signature by Machine Learning Using a Combination of Small Cohorts of Prostate Cancer. *Frontiers in Genetics*, 11, 2020. ISSN 1664-8021. doi: 10.3389/fgene.2020.550894.

von Luxburg, U. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, December 2007. ISSN 1573-1375. doi: 10.1007/s11222-007-9033-z.

Witten, D. M. and Tibshirani, R. A Framework for Feature Selection in Clustering. *Journal of the American Statistical Association*, 105(490):713–726, 2010.

Witten, D. M., Tibshirani, R., and Witten, M. D. Package ‘sparcl’. 2013.

Zhao, Z. and 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., and Mester, R. Discriminative Subspace Clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2107–2114, 2013.

A. The dynamic training regime

As features get eliminated during the training, the notion of affinity 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 we do not cluster any longer 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 (distance δ or kernel κ) is computed using only the subset of relevant features I_t . We call *static* regime the training with usage of all features at all times despite the selection of some. 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 section 3.3 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 on the synthetic datasets again and report the results in Table 6. For this experiment, we only evaluated the performances on the final subset of selected features. However, since the Sparse GEMINI is trained until a user-defined number of features is reached, we avoid unfair comparisons with other variable selections 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. Contrarily, the MMD only considers expectation which helps getting rid of noisy variations of the distance around informative variables.

Table 6. Experiment on the synthetic datasets with the dynamic regime of training for Sparse GEMINI

Method	MMD		Wasserstein	
	OvA	OvO	OvA	OvO
Scenario 1	21 (6.5)	19 (6.4)	17 (9.8)	15 (7.2)
Scenario 2	49 (4.4)	81 (16)	61 (22)	61 (16)
Scenario 3	24 (2.4)	19 (2.2)	16 (3.3)	11 (4.9)
Scenario 4	50 (5.3)	86 (0.6)	84 (9.2)	80 (13)
Scenario 5	47 (2.5)	86 (0.4)	80 (9.9)	81 (4.1)
Dataset 2	57 (2.4)	53 (7.0)	51 (8.9)	55 (5.9)