

Exploring dual information in distance metric learning for clustering

R. Randel, D. Aloise, A. Hertz

G-2021-32

May 2021

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Citation suggérée : R. Randel, D. Aloise, A. Hertz (Mai 2021). Exploring dual information in distance metric learning for clustering, Rapport technique, Les Cahiers du GERAD G- 2021-32, GERAD, HEC Montréal, Canada.

Suggested citation: R. Randel, D. Aloise, A. Hertz (May 2021). Exploring dual information in distance metric learning for clustering, Technical report, Les Cahiers du GERAD G-2021- 32, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2021-32>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2021-32>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2021
– Bibliothèque et Archives Canada, 2021

Legal deposit – Bibliothèque et Archives nationales du Québec, 2021
– Library and Archives Canada, 2021

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Télec. : 514 340-5665
info@gerad.ca
www.gerad.ca

Exploring dual information in distance metric learning for clustering

Rodrigo Randel ^a

Daniel Aloise ^a

Alain Hertz ^b

^a GERAD & Département de Génie Informatique et Génie Logiciel, Polytechnique Montréal, Montréal (Québec), Canada H3C 3A7

^b GERAD & Département de Mathématiques et de Génie Industriel, Polytechnique Montréal, Montréal (Québec), Canada H3C 3A7

rodrigo.randel@polymtl.ca

daniel.aloise@polymtl.ca

alain.hertz@polymtl.ca

May 2021

Les Cahiers du GERAD

G–2021–32

Copyright © 2021 GERAD, Randel, Aloise, Hertz

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : Distance metric learning algorithms aim to appropriately measure similarities and distances between data points. In the context of clustering, metric learning is typically applied with the assist of side-information provided by experts, most commonly expressed in the form of cannot-link and must-link constraints. In this setting, distance metric learning algorithms move closer pairs of data points involved in must-link constraints, while pairs of points involved in cannot-link constraints are moved away from each other. For these algorithms to be effective, it is important to use a distance metric that matches the expert knowledge, beliefs, and expectations, and the transformations made to stick to the side-information should preserve geometrical properties of the dataset. Also, it is interesting to filter the constraints provided by the experts to keep only the most useful and reject those that can harm the clustering process. To address these issues, we propose to exploit the dual information associated with the pairwise constraints of the semi-supervised clustering problem. Experiments clearly show that distance metric learning algorithms benefit from integrating this dual information.

Keywords: Clustering, distance metric learning, pairwise constraints, Lagrangian relaxation, duality theory

1 Introduction

A large collection of data mining techniques strongly rely on the use of similarity or dissimilarity measures among data objects to successfully perform their associated tasks (Aggarwal, 2015). Consequently, defining a metric capable of recovering the underlying data structure is a critical component for achieving good results in many pattern recognition problems (Chang and Yeung, 2006). The effects of an inappropriate dissimilarity metric are notably severe under the unsupervised learning paradigm, mainly due to the lack of background information to evaluate classification performance (Xing et al., 2003). This is the case, for example, for the *clustering* problem, one of the most popular data mining tasks, which is aimed to identify hidden homogeneous subgroups, named clusters, from the data (Hansen and Jaumard, 1997).

Whereas defining a good fitting dissimilarity metric is highly problem-specific (Xiang et al., 2008), the user of a clustering model has often no evidence or external information to assess the quality of the dissimilarity metric adopted nor of the clusters obtained. To mitigate this difficulty, *distance metric learning* techniques emerged as a mechanism to automatically learn how to appropriately measure similarities and distances through the use of external knowledge about the data (Bellet et al., 2015). In clustering, distance learning frequently occurs under the semi-supervised paradigm, where domain experts are allowed to provide additional side-information regarding the data distribution. Therefore, semi-supervised clustering methods aim to find solutions that are more in line with the expert knowledge, beliefs, and expectations.

Typically, side-information is formulated by means of pairwise constraints, for which the user provides information regarding the relationship between a pair of data objects. In this sense, a *must-link* constraint informs that two data points are similar and, therefore, must be assigned to the same cluster. Likewise, a *cannot-link* constraint ensures that a pair of points must be assigned to different clusters. Pairwise constraints arise naturally in many applications, e.g., image retrieval (Bar-Hillel et al., 2005), and, in many circumstances, are more practical to be obtained than class-labels (Zhang et al., 2008).

Accordingly, under the semi-supervised clustering paradigm with pairwise constraints, distance learning methods are designed to find transformations to the data features, so as to bring closer pairs of data objects involved in must-link constraints and to move away pairs of data objects involved in cannot-link constraints for the metric space under consideration. Thus, distance metric learning can mitigate the effect of adopting a less appropriate clustering model to capture the underlying structure of the data.

Establishing a re-featured dataset through distance metric learning has its own challenges. For example, the choice of the metric to be used initially in order to define the notion of dissimilarity is crucial and has a significant impact on the resulting clustering solution. For illustration, an example is shown in Figure 1, where clustering solutions for the popular *Wine* dataset (Dua and Graff, 2017) are obtained with the distance metric learning algorithm MPCK-Means (Bilenko et al., 2004), using two different metrics and 20 randomly generated pairwise constraints. The data points are displayed with the two principal components. We observe that the Euclidean distance is a more suitable metric for finding the ground-truth partition than the Cosine distance, even when the same set of pairwise constraints is used. This example illustrates the need to first find a suitable metric so that distance metric learning algorithms can actually benefit from the background knowledge.

Another issue with distance metric learning methods is that the transformations that are applied to the data points do not necessarily preserve the geometric properties of the dataset. This can be problematic since these transformations are based on pairwise constraints collected from different sources than the ones that generated the dataset. This occurs because domain experts provide pairwise constraints based on their beliefs and expectations, and these may not be consistent with the distribution of the data. Such a defect is particularly apparent when the pairwise constraints fail to deal well with critical regions of the space (e.g., overlapping of classes) (Le Capitaine, 2018). Therefore, a method capable of keeping the modified space as close as possible to the original one is of great interest. This

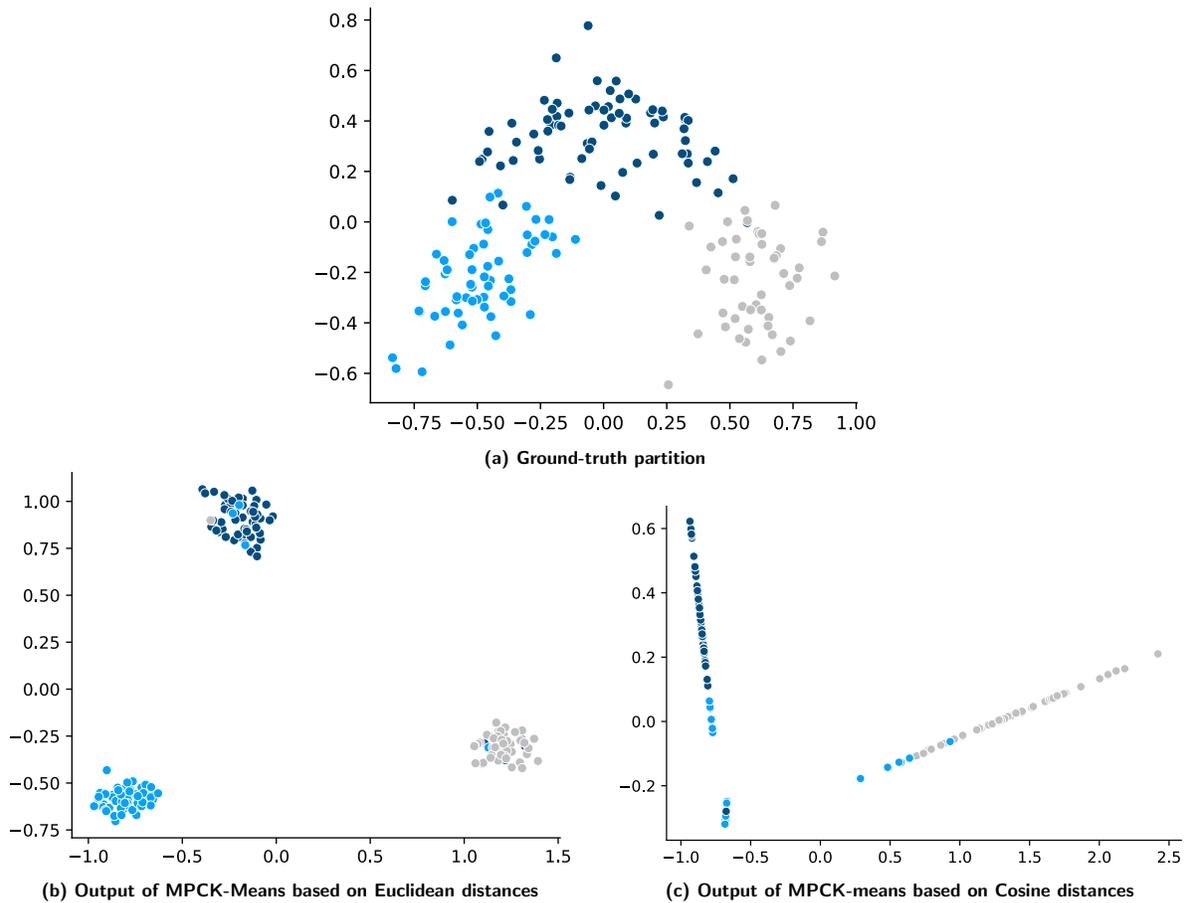


Figure 1: Illustration of the use of MPCK-Means on the Wine dataset, with two metrics and 20 pairwise constraints.

not only makes the interpretation of the clustering more reliable, but it also contributes to a more accurate view of the transformed dataset when represented using the learned metric.

It has been shown that some constraints can negatively affect the clustering quality (Davidson et al., 2006). Nevertheless, little is known about which constraints are best able to guide a semi-supervised clustering algorithm. Users of clustering algorithms are thus left with the sole option of using all the available constraints, which corresponds to a hit-or-miss game (Zhang et al., 2021). A mechanism to assess the importance of pairwise constraints, to keep only those that have a positive effect, would therefore be of great benefit to users.

This research proposes to address the above issues by defining a methodology for measuring and analyzing the effects of pairwise constraints in a clustering model. To this end, we resort to a Lagrangian formulation of the clustering problem, where pairwise constraints are relaxed and transferred to the clustering objective function. This process introduces Lagrange multipliers, called *dual variables*, which indicate the price to be paid by the model in order to violate a constraint. In other words, by associating each pairwise constraint with a dual variable, we can measure the impact of each constraint in the clustering solution. This idea was introduced in our previous work (Randel et al., 2021) with the objective of identifying the presence of erroneous/contradictory pairwise constraints. Motivated by the wealth of information that can be acquired from dual variables, we propose in this paper to further exploit this knowledge by developing tools to help manage the problems described above in distance metric learning for clustering.

The rest of the paper is organized as follows. In Section 2, we summarize prior research on distance metric learning for clustering, with supervision based on pairwise constraints. We then explain in

Section 3 how to generate and exploit dual information obtained from pairwise constraints within clustering optimization problems. The next three sections illustrate the use of the proposed tools: we show in Section 4 how to identify the most appropriate metric in clustering models; in Section 5, we explain how to design distance metric learning methods so that the transformations that are applied to the data points preserve as much as possible the geometric properties of the dataset; we describe in Section 6 a procedure that filters the most useful constraints and demonstrate how it can be integrated with a deep learning framework recently proposed in the literature. Concluding remarks are given in the last section.

2 Related works

In this paper, we focus on distance metric learning algorithms for clustering, with supervision based on pairwise constraints. The literature on metric learning is however much broader, and the reader is referred to Bellet et al. (2015) for a thorough review on the subject.

The distance metric learning task is often performed by assuming that dissimilarities are expressed as Mahalanobis distances (Mahalanobis, 1936). This distance has generated a lot of interest due to its flexibility and nice interpretation in terms of a linear projection (Bellet et al., 2015; Nguyen and Baets, 2019). Given a dataset $O = \{o_1, \dots, o_n\}$ of n points in the d -dimensional space, the objective of the Mahalanobis distance learning is to determine a positive semi-definite $d \times d$ matrix \mathcal{A} , such that the Mahalanobis distance between two points o_i and o_j is increased if a constraint imposes that they must belong to the same cluster, and is decreased if the points are to be in different clusters. The Mahalanobis distance between points o_i and o_j in \mathbb{R}^d is defined as:

$$d_{\mathcal{A}}(o_i, o_j) = \sqrt{(o_i - o_j)^T \mathcal{A} (o_i - o_j)}. \quad (1)$$

The learned matrix \mathcal{A} can be used to transform the dataset O into a set O' by defining $o'_i = \mathcal{A}^{1/2} o_i$ for $i = 1, \dots, n$ (where $\mathcal{A}^{1/2}$ is the unique matrix \mathcal{B} that is positive semidefinite and such that $\mathcal{B}\mathcal{B} = \mathcal{B}^T \mathcal{B} = \mathcal{A}$). The standard Euclidean distance can then be used on O' by classical data mining algorithms. The seminal work exploring this idea was done by Xing et al. (2003). Formally, given a set \mathcal{ML} of must-link constraints and a set \mathcal{CL} of cannot-link constraints, the *Mahalanobis distance metric learning problem* is expressed as follows, where \mathbb{S} is the set of positive semidefinite $d \times d$ matrices :

$$\begin{aligned} & \max_{\mathcal{A} \in \mathbb{S}} \sum_{(o_i, o_j) \in \mathcal{CL}} d_{\mathcal{A}}(o_i, o_j) \\ & \text{subject to} \sum_{(o_i, o_j) \in \mathcal{ML}} d_{\mathcal{A}}(o_i, o_j) \leq c \end{aligned} \quad (2)$$

In words, the sum of the distances between dissimilar points is maximized, while for similar points, this sum must be less than a constant c which is typically set equal to 1. One could also choose to minimize the sum of the distances between similar points, while keeping the sum of distances between dissimilar points greater than c . Problem (2) can be solved to optimality using a projected gradient descent algorithm. However, if \mathcal{A} is a full matrix, the time complexity of determining the eigenvalues of \mathcal{A} is $O(d^2)$, which is computationally too expensive for high-dimensional data (Bilenko et al., 2004; Wu et al., 2012). To reduce this complexity, one can impose that \mathcal{A} be a diagonal matrix, which is equivalent to performing a feature weighting on the dataset.

A classical distance metric learning technique for clustering with pairwise constraints is the MPCK-Means algorithm (Bilenko et al., 2004). Instead of computing a single matrix \mathcal{A} , MPCK-Means determines one matrix per cluster. The method can thus define local transformations which allow the clusters to have different shapes. This feature can be particularly useful to enforce specific shapes. For example, squared Euclidean distances used by the classical k -means heuristic are suitable for spheroidal clusters, but fail to represent dissimilarities when the clusters have other shapes.

We should also mention that the learning of nonlinear shapes can be achieved by using Bregman distances (Lei et al., 2009; Wu et al., 2012) and kernel functions (Baghshah and Shouraki, 2009; Kalintha et al., 2017; Nguyen and Baets, 2019), both being more suitable than Mahalanobis distances for high-dimensional data, or to work with datasets with nonlinear structures.

Methods based on deep learning have also been developed (Hsu and Kira, 2015; Fogel et al., 2019; Zhang et al., 2021). For example, Zhang et al. (2021) recently proposed to use the Deep Embedded Clustering (DEC) algorithm (Xie et al., 2016) followed by a training procedure that incorporates the violation of secondary constraints in the loss function. The DEC method consists of first learning an embedded representation of the data points using an autoencoder network, and then executing a self-training routine to learn a clustering partition. To accomplish that, DEC uses the embedded points to define a *soft* membership distribution (points can be partially assigned to more than one cluster), and then approximates it to a target distribution that resembles a *hard* clustering membership (where data points are assigned to exactly one cluster) using the Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951) loss function. The work of Zhang et al. (2021), extends the loss function by adding terms that accounts for diverse types of supervision, including pairwise constraints. It has been shown that the proposed algorithm outperforms previous semi-supervised clustering algorithms in terms of clustering accuracy.

3 Dual information from the pairwise constraints

A k -partition of a set is a grouping of its elements into k disjoint non-empty subsets called clusters. Given a set O of n data points and an integer k , we write $\mathcal{P}(O, k)$ for the set of all k -partitions of O . Let $f : \mathcal{P}(O, k) \rightarrow \mathbb{R}$, be a function, usually called *clustering criterion*, that assigns a value to every k -partition in $\mathcal{P}(O, k)$. The clustering problem is to determine a k -partition $P \in \mathcal{P}(O, k)$ with minimum (or maximum) value $f(P)$.

A widely used clustering criterion is the minimum sum of squared Euclidean distances from each data point to the centroid of its cluster, or minimum sum-of-squares clustering (MSSC) for short, which expresses both homogeneity and separation (Aloise et al., 2012). Let \mathcal{ML} and \mathcal{CL} be the sets of pairs (o_i, o_j) of data points involved in must-link and cannot-link constraints, respectively, where $(o_i, o_j) \in \mathcal{ML} \cup \mathcal{CL}$ only if $i < j$. The semi-supervised version of the clustering problem with MSSC can be expressed by means of the following optimization problem:

$$\min \sum_{i=1}^n \sum_{c=1}^k x_i^c \|o_i - y_c\|^2 \quad (3)$$

$$\text{subject to } \sum_{c=1}^k x_i^c = 1 \quad \forall i = 1 \dots n \quad (4)$$

$$x_i^c + x_j^c \leq 1 \quad \forall (o_i, o_j) \in \mathcal{CL}, \quad \forall c = 1, \dots, k \quad (5)$$

$$x_i^c - x_j^c = 0 \quad \forall (o_i, o_j) \in \mathcal{ML}, \quad \forall c = 1, \dots, k \quad (6)$$

$$x_i^c \in \{0, 1\} \quad \forall i = 1, \dots, n, \quad \forall c = 1, \dots, k. \quad (7)$$

The feasible solutions to this problem correspond to partitions $\{C_1, \dots, C_k\} \in \mathcal{P}(O, k)$ where every binary decision variable x_i^c indicates whether data point o_i is assigned to cluster C_c , and y_c is the centroid of cluster C_c . Constraints (4) ensure that each data point is assigned to exactly one cluster. The cannot-link constraints are expressed by equations (5) and the must-link constraints by equations (6). To avoid situations where these constraints are satisfied with equality, we can replace them by the following equivalent constraints where ϵ is any real number in $]0, 1[$:

$$x_i^c + x_j^c \leq 1 + \epsilon \quad \forall (o_i, o_j) \in \mathcal{CL}, \quad \forall c = 1, \dots, k \quad (5')$$

$$x_i^c - x_j^c \leq \epsilon \quad \forall (o_i, o_j) \in \mathcal{ML}, \quad \forall c = 1, \dots, k \quad (6')$$

$$x_j^c - x_i^c \leq \epsilon \quad \forall (o_i, o_j) \in \mathcal{ML}, \quad \forall c = 1, \dots, k. \quad (6'')$$

We aim to investigate how the pairwise constraints affect the clustering objective as expressed in (3). We believe that the cost of complying with a constraint may provide information on its importance. We thus attempt to answer the following questions:

- what are the most difficult constraints to satisfy?
- is the adopted metric in agreement with the imposed constraints?
- what are the most useful constraints for obtaining a low cost partition?

This constitutes a form of sensitivity analysis (Pichery, 2014) whose goal is to measure the impact of modifications of the input variables on the result of a clustering model. To obtain such information, we have recourse to the Lagrangian duality theory, as explained in Randel et al. (2021)

The Lagrangian function for the optimization problem (3)–(7) can be obtained by introducing dual variables η_{ij}^c , λ_{ij}^c and γ_{ij}^c to penalize violations of inequality constraints (5'), (6') and (6''). Specifically, the Lagrangian function $L(\eta, \lambda, \gamma)$ is defined as follows:

$$\begin{aligned}
L(\eta, \lambda, \gamma) = \min \quad & \sum_{i=1}^n \sum_{c=1}^k x_i^c \|o_i - y_c\|^2 \\
& + \sum_{(o_i, o_j) \in \mathcal{CL}} \sum_{c=1}^k \eta_{ij}^c (1 + \epsilon - x_i^c - x_j^c) \\
& + \sum_{(o_i, o_j) \in \mathcal{ML}} \sum_{c=1}^k \lambda_{ij}^c (\epsilon + x_i^c - x_j^c) \\
& + \sum_{(o_i, o_j) \in \mathcal{ML}} \sum_{c=1}^k \gamma_{ij}^c (\epsilon + x_j^c - x_i^c) \tag{8}
\end{aligned}$$

$$\text{subject to} \quad \sum_{c=1}^k x_i^c = 1 \quad \forall i = 1 \dots n \tag{9}$$

$$x_i^c \in \{0, 1\} \quad \forall i = 1, \dots, n; \quad \forall c = 1, \dots, k \tag{10}$$

and the dual of the integer program (3)–(7) can be expressed as follows:

$$L_D = \max_{\eta, \lambda, \gamma \leq 0} L(\eta, \lambda, \gamma). \tag{11}$$

The dual variables η_{ij}^c , λ_{ij}^c and γ_{ij}^c provide an estimation of the cost paid by the clustering optimization model for respecting the associated constraints. If a pairwise constraint is inherently satisfied by (3)–(7), its associated dual variable should be zero. On the other hand, if a constraint is necessary to obtain an optimal solution to (3)–(7), a penalty must be paid for its violation, which implies that the associated dual variable should be strictly negative.

A sub-gradient optimization algorithm can be used to solve (11) and thus determine optimal values of the dual variables. The weak duality theorem (Bertsimas and Tsitsiklis, 1997; Fisher, 2004) asserts that L_D is the best possible lower bound on the optimal value of (3)–(7). Consequently, the optimal values of the dual variables are reliable estimates of the price to be paid for complying with the constraints associated with them. For more details, the reader is referred to Randel et al. (2021).

Formulation (3)–(7) is based on the MSSC clustering criterion. Nevertheless, the proposed methodology is not limited to this particular objective. In fact, the strategy is flexible to be used in conjunction with several other clustering criteria and models represented in mathematical programming language (Hansen and Jaumard, 1997). To illustrate that, we use the k-medoids model (Kaufman and Rousseeuw, 2009). More precisely, let D be a matrix, where each entry d_{ij} indicates the dissimilarity between data points o_i and o_j . The medoid of a cluster is defined as the data point in the cluster

whose average dissimilarity to all the data points in the cluster is minimal. The k -medoids problem is to determine a partition in $\mathcal{P}(O, k)$ that minimizes the sum of the distances of the data points to the medoid of their cluster. The semi-supervised version of this clustering problem can be expressed by means of the following optimization problem:

$$\min \sum_{i=1}^n \sum_{c=1}^n x_i^c d_{ic} \quad (12)$$

$$\text{subject to } \sum_{c=1}^n x_i^c = 1 \quad \forall i = 1, \dots, n \quad (13)$$

$$x_i^c \leq y_c \quad \forall i = 1, \dots, n, \quad \forall c = 1, \dots, n \quad (14)$$

$$\sum_{c=1}^n y_c = k \quad (15)$$

$$x_i^c + x_j^c \leq 1 \quad \forall (o_i, o_j) \in \mathcal{CL}, \quad \forall c = 1, \dots, n \quad (16)$$

$$x_i^c - x_j^c = 0 \quad \forall (o_i, o_j) \in \mathcal{ML}, \quad \forall c = 1, \dots, n \quad (17)$$

$$x_i^c \in \{0, 1\} \quad \forall i = 1, \dots, n, \quad \forall c = 1, \dots, n, \quad (18)$$

$$y_c \in \{0, 1\} \quad \forall c = 1, \dots, n, \quad (19)$$

where every binary variable y_c indicates whether the data point o_c is one of the k medoids, and every binary variable x_i^c indicates whether data point o_i is assigned to a cluster having o_c as medoid. The Lagrangian function $L(\eta, \lambda, \gamma)$ associated with this optimization problem is then defined as in the previous case by replacing constraints (16) and (17) by penalty terms in the objective function and the same sub-gradient optimization technique as in (Randel et al., 2021) can be used to determine $L_D = \max_{\eta, \lambda, \gamma \leq 0} L(\eta, \lambda, \gamma)$.

4 Identifying an appropriate dissimilarity measure

Although the choice of the most appropriate dissimilarity metric is problem-specific, users of an algorithm or a clustering model often have very little information to support their choice which can have a significant impact on the analysis of the quality of a solution. For example, the homogeneity of a cluster depends on the similarity measure between the objects which is defined using the chosen metric.

Distance learning methods are also highly dependent on the chosen distance metric. For example, the most commonly used dissimilarity measure for quantitative data is the Euclidean distance (Wu et al., 2012). It assumes that each data feature is independent and equally important, which is not necessarily the case for the considered dataset. An alternative is to weight the relative importance of the data features. For instance, in the weighted Euclidean Distance-Based approach (Rao, 2013), data features are weighted according to their discriminating effect. These weights are generally set based on user preferences and interpretation, which is a subjective task prone to error. One can also choose to use secondary information such as pairwise constraints to improve the initial dissimilarities calculated from Euclidean distances, but this might not compensate for the fact that a more suitable metric exists for the underlying clustering optimization model.

To tackle this issue, we propose to use a *fitting score* which measures the adequacy of a metric with a given clustering model. This score is therefore a tool that helps experts to identify the metric that most closely matches their knowledge and beliefs. The fitness score $F(m)$ of a metric m is defined as

$$F(m) = \sum_{(o_i, o_j) \in \mathcal{CL}} \sum_{c=1}^k \delta(\eta_{ij}^c = 0) + \sum_{(o_i, o_j) \in \mathcal{ML}} \sum_{c=1}^k \delta(\lambda_{ij}^c = \gamma_{ij}^c = 0), \quad (20)$$

where η_{ij}^c , λ_{ij}^c and γ_{ij}^c are the optimal values of the dual variables resulting from the optimization of L_D in (11), and $\delta(e) = 1$ if expression e is true. The metric with the highest fitness score should be preferred for the given clustering model, as more pairwise constraints are inherently respected.

In the following subsections, we validate the ability of the fitness score to detect the most adequate metric for a clustering problem. Experiments are conducted on synthetic and real-world datasets.

4.1 Validation of the fitness score for synthetic datasets

In order to assess the usefulness of the proposed fitness score, we selected four of the most commonly used distance metrics in clustering problems, namely the Euclidean, Manhattan, Chebyshev, and Mahalanobis distances.

For each of the four metrics, we generated a two-dimensional dataset with 200 points and three balanced clusters C_1, C_2, C_3 . This is how we proceeded.

- For the Euclidean, Manhattan and Chebyshev distances, we first generated 1000 pairs of points with coordinates drawn from a normal distribution $\mathcal{N}(10, 1)$, and determined the maximum distance Δ (using the chosen metric) between the two points of a pair. We then repeatedly generated three points c_1, c_2, c_3 with coordinates drawn from a normal distribution $\mathcal{N}(10, 1)$, until $\min\{d(c_1, c_2), d(c_1, c_3), d(c_2, c_3)\} \geq \Delta/3$, where d is the considered distance metric. The three resulting points were defined as the centers of the three clusters. Starting from $C_1=C_2=C_3=\emptyset$, we then iteratively generated points p with coordinates drawn from a normal distribution $\mathcal{N}(10, 1)$, determined the index i such that $d(p, c_i) = \min\{d(p, c_1), d(p, c_2), d(p, c_3)\}$, and added p to C_i if $|C_i| < 67$. The process stopped when $|C_1| + |C_2| + |C_3| = 200$.
- For the Mahalanobis distance, we have defined $c_1=(50, 0)$, $c_2=(50, 5)$ and $c_3=(50, 10)$ as centers of the three clusters. The 200 data points were then generated so that every x coordinate is a random number chosen from the uniform distribution on the interval $[0, 100]$, and the y coordinate of every data point in C_i ($i = 1, 2, 3$) is a random number generated from the uniform distribution on the interval $[5(i-1)-0.1, 5(i-1)+0.1]$.

The four generated datasets and their ground-truth partitions are shown on Figure 2. Pairwise constraints were obtained using the *boosting aggregating framework* (Breiman, 1996). For each of the 4 datasets, we generated 500 sets E_1, \dots, E_{500} of constraints, using the ground-truth partition. Each set E_i contains n_i constraints, with n_i chosen randomly according to a uniform distribution in $\{1, \dots, 100\}$.

We take advantage of the flexibility of the k -medoids model to solve the unsupervised clustering problem associated with each of the four datasets and each of the four metrics. More precisely, we set d_{ij} equal to the distance between o_i and o_j , using the considered metrics, and use CPLEX 12.8 to solve model (12)–(19) without constraints (16) and (17). The clustering accuracy of the resulting solutions is obtained using the Adjusted Random Index (ARI) (Hubert and Arabie, 1985), which measures the similarity between the obtained clustering solutions and the ground-truth partition.

For each dataset, each set E_i of pairwise constraints, and each metric m , we ran the sub-gradient algorithm of Randel et al. (2021) to solve (11) and so produce optimal dual variables which give a fitness score as defined in (20). A time limit of 10 seconds has been allocated to each execution of the sub-gradient algorithm. We report in Figure 3, for each dataset and each of the four metrics, the average fitness score calculated over the 500 sets of constraints. We also indicate the ARI values of the solutions produced by the unsupervised clustering algorithm based on the k -medoids model.

We observe that the fitness score is able to identify the most suitable metric for all datasets. The fluctuation of the fitness score agrees with that of the ARI and the best score is always obtained with the appropriate metric. Although the experiments were carried out on synthetic data, the results demonstrate that the optimal dual values of (11) are useful for exploiting the information provided by the pairwise constraints, in order to suggest the metric which seems to adhere the most to the data.

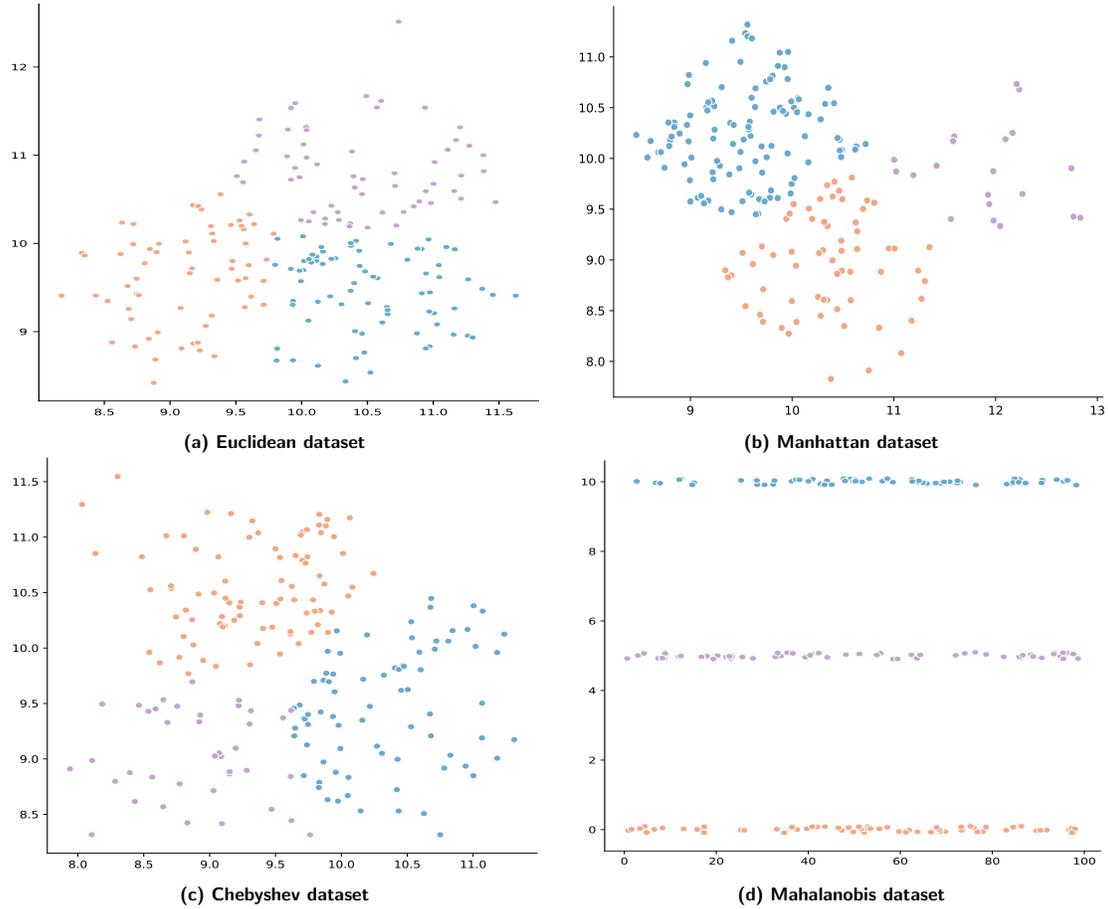


Figure 2: Four synthetic datasets for four distance metrics, with their ground-truth partitions.

4.2 Validation of the fitness score for five real-world datasets

We now consider five real-world datasets, the first four being available on the UCI repository (Dua and Graff, 2017), and the fifth being described in Rodrigues et al. (2020). Their characteristics are summarized in Table 1.

Table 1: Real-data Applications for Evaluating the Score.

Dataset	Samples	Classes	Features	Type
Iris	150	3	4	Quantitative
Wine	178	3	13	Quantitative
Control	600	6	60	Quantitative; time series
Twenty newsgroup	600	4	-	Text
Eclipse bug report	460	5	-	Discrete sequence; text

4.2.1 Quantitative datasets

For the quantitative datasets *Iris* and *Wine*, we used the same distance metrics as those carried out for the synthetic datasets of Section 4.1, that is, Euclidean, Manhattan, Chebyshev, and Mahalanobis. For the *Control* dataset, each data point represents a time series composed of 60 values. Each time series was decomposed into six segments of 10 consecutive values, and we have set $o_i = (o_i^1, \dots, o_i^6)$, where o_i^k ($k = 1, \dots, 6$) are the segments of o_i . The distance between two data points o_i and o_j was then defined as $\sum_{k=1}^6 d(o_i^k, o_j^k)$, where d is the considered metric.

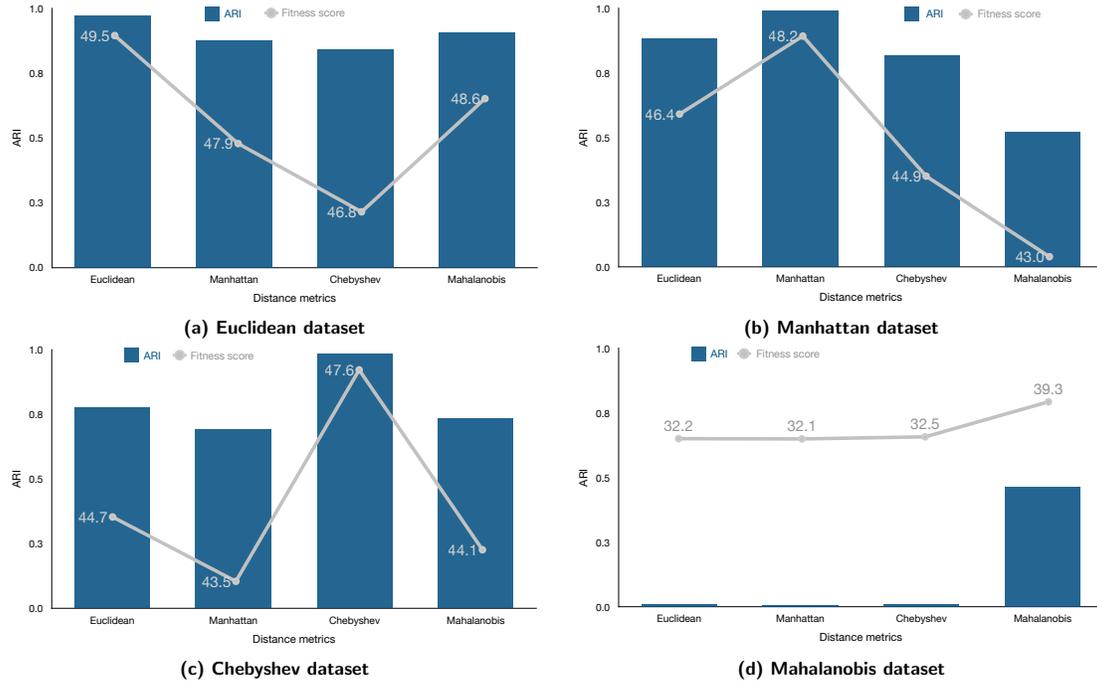


Figure 3: Fitness scores and ARI for synthetic datasets.

For each dataset, we generated 500 constraints sets from the available ground-truth partitions, using the *boosting aggregating framework* (Breiman, 1996), as was done for the synthetic dataset. The time limit for the sub-gradient algorithm has been set at 10 seconds for the *Iris* and *Wine* instances, and at 30 seconds for the *Control* dataset. The results are shown in Figure 4, alongside the representation of the datasets with the two principal components and their ground-truth partitions.

We observe that the L_p norm-based distances (i.e., the Euclidean, Manhattan and Chebyshev distances) give similar results, both in terms of clustering accuracy and fitness score. This is not surprising given the popularity of these types of metrics for quantitative data clustering.

4.2.2 Textual data

The datasets *Twenty newsgroup* and *Eclipse bug report* are composed of text samples. As typically done in document classification tasks, the texts were first converted to quantitative data through the use of a bag of words. For both datasets, we defined 10,000 features, which are originated from the most frequent terms in the dataset vocabulary. The direct use of L_p norm-based distances on raw textual data is not recommended due to data sparsity and high-dimensionality. Besides, they are prone to be inconsistent if not normalized because the distance between two long documents is very often larger than that between two short documents even if the two long documents are similar, and the short documents are unrelated (Aggarwal, 2015). In contrast, the *cosine distance* works by computing the angle between data samples, and is a popular choice for varied length-sized document classification.

In addition, instead of working with the raw frequency of the terms, we enhanced the features by calculating the *term frequency-inverse document frequency* (TF-IDF) (e.g. Manning and Schütze, 1999) to use global statistical measures to improve the dissimilarity computation. TF-IDF is based on the principle that documents matching with respect to rare terms are more likely to be similar than documents sharing many common terms (e.g. *the, like*). In our experiments, we considered every data point as a vector in $\mathbb{R}^{10,000}$, where each coordinate indicates the TF-IDF of the corresponding term.

The *Damerau-Levenshtein* distance is commonly used to establish the dissimilarity between discrete sequences. It measures the minimum number of operations necessary to change one sequence into

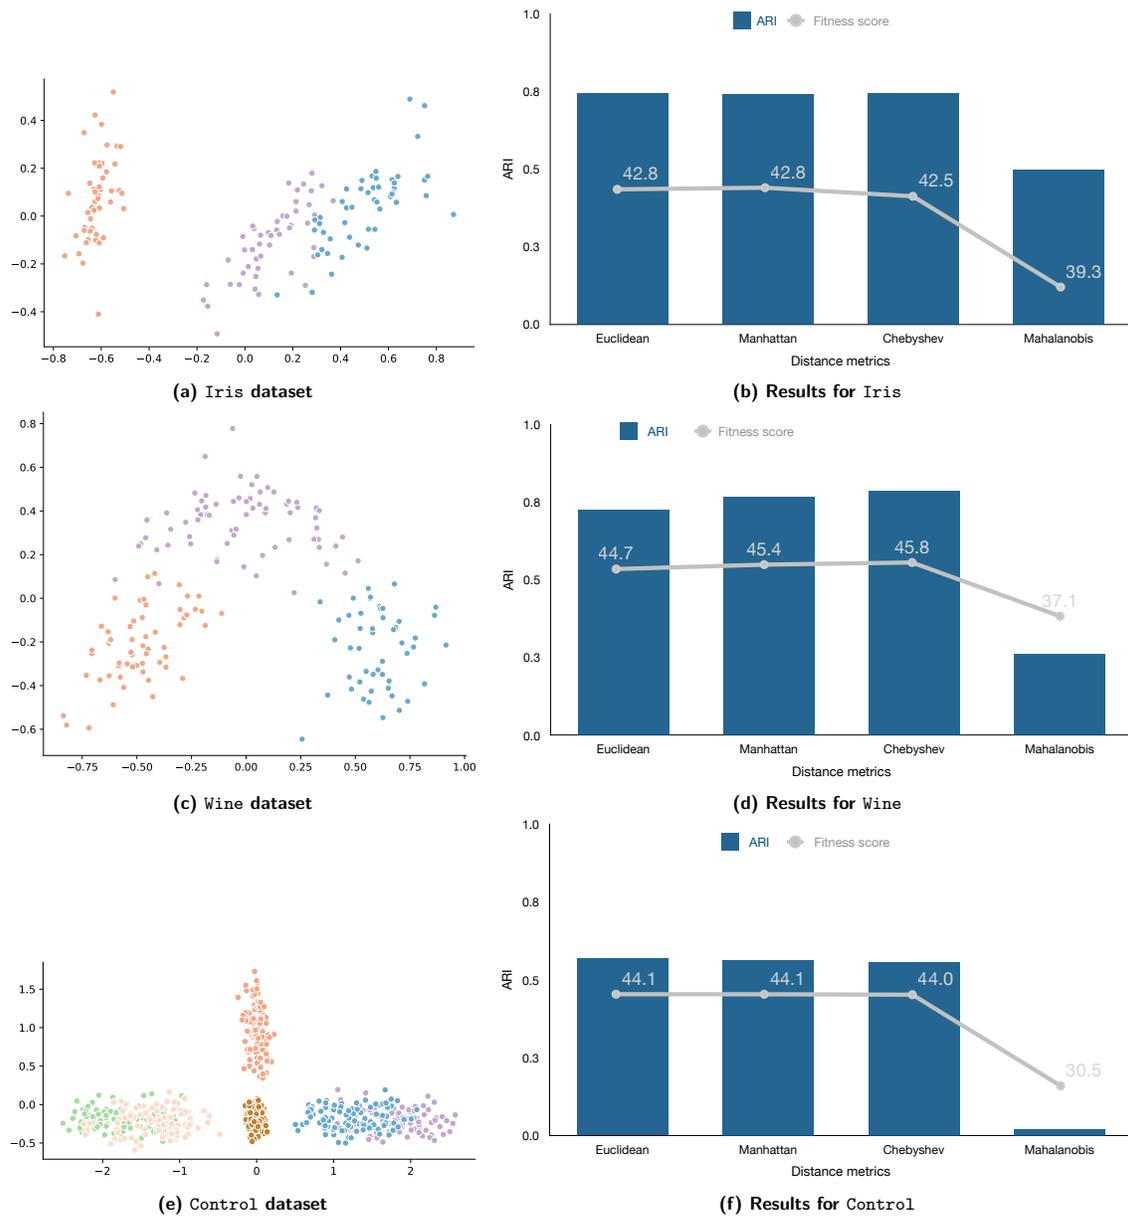


Figure 4: Ground-truth partition, fitness score and ARI for three quantitative datasets.

another. Since the Eclipse bug report dataset is structured as a discrete sequence of stack trace functions, we decided to include this other distance in our tests. A time limit of 30 seconds has been allocated to each execution of the sub-gradient algorithm.

The results are reported in Figure 5. We observe that our fitness score was able to identify the Cosine distance as an appropriate metric for clustering text documents. In addition, the Damerau-Levenshtein distance has given the highest fitness score for the Eclipse bug report instance. This demonstrates once again that the proposed fitness score is a valid tool to suggest a metric that adheres the most to the data.

5 Maintaining geometrical properties of the dataset

In this section, we demonstrate that dual information can be of great interest when experts are not entirely sure of the pairwise constraints they provide. In such circumstances, one may wish to perform

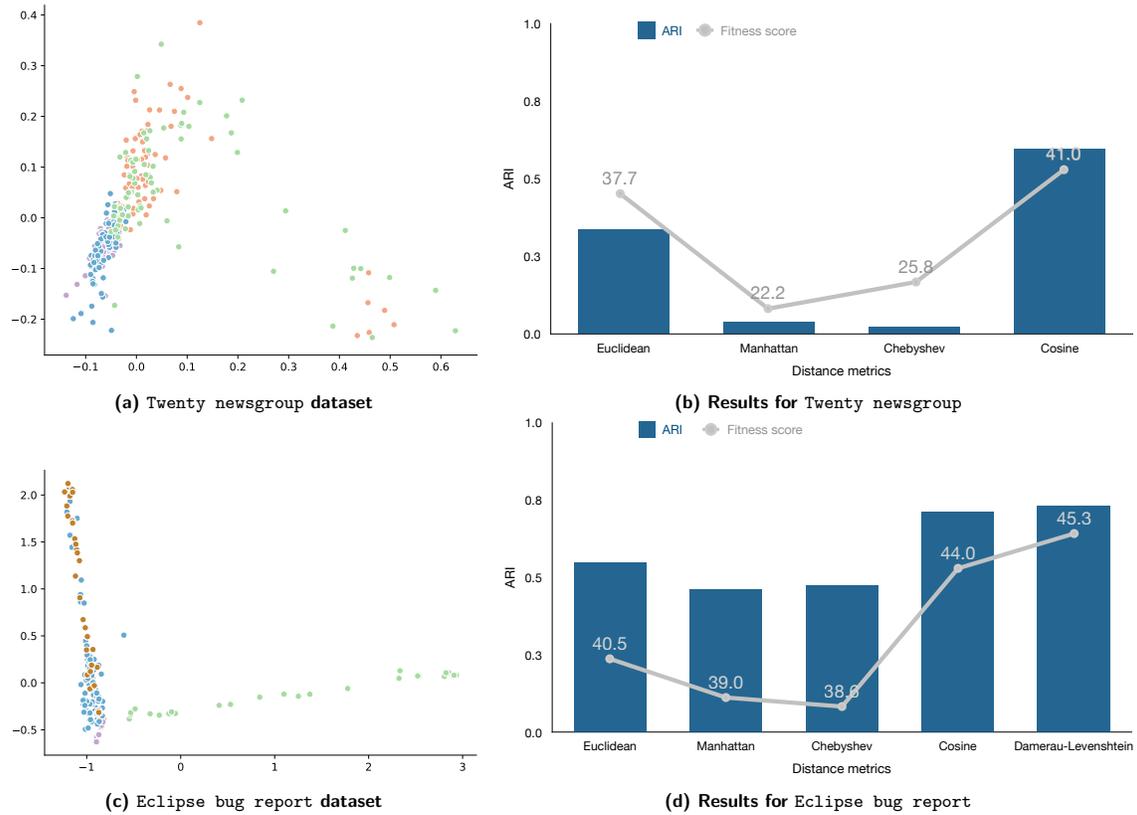


Figure 5: Ground-truth partition, fitness score and ARI for two text datasets.

distance metric learning more carefully, preserving as much as possible the characteristics of the original data distribution while incorporating the domain knowledge embedded in the pairwise constraints.

With that in mind, we propose to use dual information to establish an order in which pairwise constraints should be processed by metric learning algorithms. This sorting step aims to help these algorithms find the least impactful data transformations that can be applied, that is, those transformations that require minimal modifications to the original data. Thus, our strategy aims not only to integrate the information contained in the pairwise constraints, but also to maintain a reliable representation of the data.

With MSSC as underlying clustering model, we designed a very simple distance metric learning algorithm that iteratively selects a violated constraint, and then performs data transformations necessary to satisfy it. The method exploits the dual information in order to rank the pairwise constraints according to their estimated impact on the clustering objective function (3).

Algorithm 1 describes the steps of the proposed method. The algorithm begins by running an unsupervised clustering algorithm that generates an initial solution X to the clustering problem (3)–(7), without constraints (5) and (6) (Step 3). We then run the sub-gradient algorithm to solve (11) and thus produce optimal dual values η_{ij}^c , λ_{ij}^c and γ_{ij}^c (Step 4). An impact score \mathcal{I}_{ij} is then calculated for each pairwise constraint $(o_i, o_j) \in \mathcal{CL} \cup \mathcal{ML}$, this score being defined as the sum of the values of the dual variables which are associated with the constraint (Step 5). Next, we determine the pair $(o_i, o_j) \in \mathcal{CL} \cup \mathcal{ML}$ of maximum impact score \mathcal{I}_{ij} whose associated constraint is violated by X (Step 10). The reason behind this idea of maximizing the impact score is that the higher the value of a dual variable (i.e. the less negative it is), the lower will be the gain in the objective function if the associated constraint is relaxed. Thus, the data transformation required to satisfy such a constraint will presumably be small. We then run Algorithm 2 that moves o_i and/or o_j in order to try to satisfy

constraint (o_i, o_j) (Step 11). Details on how this is done are given below. This process is repeated until X satisfies all constraints.

Note that the pairwise constraints can be conflicting and Algorithm 1 would therefore never stop. In addition, the algorithm can cycle. Indeed, suppose that $k = 2$ and that there are only two cannot-link constraints (o_1, o_2) and (o_1, o_3) . If o_1 and o_2 belong to C_1 while o_3 belongs to C_2 , the algorithm can possibly move o_1 towards C_2 to satisfy constraint (o_1, o_2) . But (o_1, o_3) will then probably be violated, and the algorithm can possibly move o_1 back towards C_1 , and this cycle can be repeated indefinitely. To avoid such situations, when moving a data point $o_i \in C_j$ to a new position, we add the pair (i, j) in an initially empty list \mathcal{L} (Step 12), and we then forbid moving o_i back to C_j . More precisely, suppose constraint $(o_i, o_j) \in \mathcal{ML}$ is violated by X . At Step 6, we define M_{ij} as the set of indices c of clusters such that the move of o_i and o_j towards C_c is not forbidden (i.e., their move does not belong to \mathcal{L}). Similarly, for a violated constraint $(o_i, o_j) \in \mathcal{CL}$, we define M_{ij} at Step 7 as the set of indices c of clusters such that C_c does not contain o_i and o_j (i.e., $x_i^c = x_j^c = 0$), and the move of at least one of the two data points towards cluster C_c does not belong to \mathcal{L} . Let \mathcal{M} be equal to the union of the sets M_{ij} over all $(o_i, o_j) \in \mathcal{ML} \cup \mathcal{CL}$ (Step 8). The algorithm stops (Step 14) when $\mathcal{M} = \emptyset$, which means that X satisfies all the constraints, or all moves which would satisfy a violated constraint are in \mathcal{L} . The algorithm is finite since a data point o_i is moved towards a cluster C_c only if $(i, c) \notin \mathcal{L}$, while \mathcal{L} increases at each iteration.

Algorithm 1 Using dual information for satisfying all pairwise constraints.

- 1: Set $\mathcal{L} \leftarrow \emptyset$.
 - 2: **repeat**
 - 3: Run an unsupervised MSSC algorithm for solving model (3)–(7) without constraints (5) and (6) and let $X = (x_i^c)$ be the resulting solution.
 - 4: Run the sub-gradient algorithm to solve (11) and thus produce optimal dual values $\eta_{ij}^c, \lambda_{ij}^c$ and γ_{ij}^c .
 - 5: For all $(o_i, o_j) \in \mathcal{CL} \cup \mathcal{ML}$ do $\mathcal{I}_{ij} \leftarrow \begin{cases} \sum_{c=1}^k \eta_{ij}^c & \text{if } (o_i, o_j) \in \mathcal{CL} \\ \sum_{c=1}^k (\lambda_{ij}^c + \lambda'_{ij}^c) & \text{if } (o_i, o_j) \in \mathcal{ML}. \end{cases}$
 - 6: For all $(o_i, o_j) \in \mathcal{ML}$ set $M_{ij} \leftarrow \{c \in \{1, \dots, k\} \mid \{(i, c), (j, c)\} \cap \mathcal{L} = \emptyset\}$ if X violates (o_i, o_j) , and $M_{ij} \leftarrow \emptyset$ otherwise.
 - 7: For all $(o_i, o_j) \in \mathcal{CL}$ set $M_{ij} \leftarrow \{c \in \{1, \dots, k\} \mid x_i^c = x_j^c = 0 \text{ and } \{(i, c), (j, c)\} \not\subseteq \mathcal{L}\}$ if X violates (o_i, o_j) , and $M_{ij} \leftarrow \emptyset$ otherwise.
 - 8: Set \mathcal{M} equal to the union of the sets M_{ij} over all $(o_i, o_j) \in \mathcal{ML} \cup \mathcal{CL}$.
 - 9: **if** $\mathcal{M} \neq \emptyset$ **then**
 - 10: Determine the pair (i, j) that maximizes \mathcal{I}_{ij} among those with $M_{ij} \neq \emptyset$.
 - 11: Run Algorithm 2 to modify the coordinates of o_i and/or o_j .
 - 12: Let c_i and c_j be the indices of the clusters such that $x_i^{c_i} = x_j^{c_j} = 1$.
 if the coordinates of o_i have been changed at Step 11 **then** add (i, c_i) to \mathcal{L} .
 if the coordinates of o_j have been changed at Step 11 **then** add (j, c_j) to \mathcal{L} .
 - 13: **end if**
 - 14: **until** $\mathcal{M} = \emptyset$.
-

We now explain how Algorithm 2 moves o_i and/or o_j in an attempt to satisfy a violated constraint $(o_i, o_j) \in \mathcal{ML} \cup \mathcal{CL}$. Let us first define more precisely what we mean by moving a point $o_i \in C_c$ towards a new cluster C_{c^*} . Such a move is obtained by modifying the coordinates of o_i as follows. Let ℓ_1 be the line containing o_i and parallel to $\overrightarrow{y_c, y_{c^*}}$, let ℓ_2 be the perpendicular bisector of the segment that connects y_c with y_{c^*} , and let p be the point at the intersection of ℓ_1 and ℓ_2 : point o_i is moved to the new location p' such that $\overrightarrow{o_i p'} = \frac{101}{100} \overrightarrow{o_i p}$, which means that o_i is then slightly closer to y_{c^*} than to y_c .

As mentioned in Section 3, the value $f(X)$ of a solution X with MSSC as clustering criterion is

$$\sum_{i=1}^n \sum_{c=1}^k x_i^c \|o_i - y_c\|^2$$

where y_c is the centroid of cluster C_c . As shown in (Hansen and Mladenović, 2001), if a new solution X' is obtained from X by moving a data point $o_i \in C_c$ to a new cluster C_{c^*} with $c \neq c^*$ (i.e., $x_i^c = 1$ and $x_i^{c^*} = 0$ are replaced by $x_i^c = 0$ and $x_i^{c^*} = 1$), then the variation $\Delta(i, c, c^*) = f(X') - f(X)$ is equal to the following expression:

$$\Delta(i, c, c^*) = \frac{\sum_{i=1}^n x_i^{c^*}}{\sum_{i=1}^n x_i^{c^*} + 1} \|o_i - y_{c^*}\|^2 - \frac{\sum_{i=1}^n x_i^c}{\sum_{i=1}^n x_i^c - 1} \|o_i - y_c\|^2.$$

Since we do not want to move o_i towards C_{c^*} if $(i, c^*) \in \mathcal{L}$, and o_i does not move if $c = c^*$, we consider $\delta(o_i, c, c^*)$ defined as follows:

$$\delta(i, c, c^*) = \begin{cases} \Delta(i, c, c^*) & \text{if } (i, c^*) \notin \mathcal{L} \text{ and } c \neq c^* \\ \infty & \text{if } (i, c^*) \in \mathcal{L} \text{ and } c \neq c^* \\ 0 & \text{if } c = c^* \end{cases}$$

We are now ready to explain how Algorithm 2 works. Let (o_i, o_j) be a violated constraints $\in \mathcal{CL}$ and assume that o_i and o_j currently belong both to cluster C_c . Algorithm 2 determines the index $u \in \{i, j\}$ and the cluster index $c^* \neq c$ in M_{ij} with minimum value $\delta(u, c, c^*)$, and then moves o_u towards C_{c^*} , as explained above.

Similarly, let (o_i, o_j) be a violated constraint $\in \mathcal{ML}$ and let c_i and c_j be the indices of the clusters that contain o_i and o_j , respectively. Algorithm 2 determines the cluster index $c^* \in M_{ij}$ with minimum value $\delta(i, c_i, c^*) + \delta(j, c_j, c^*)$. If $o_i \notin C_{c^*}$ then o_i is moved towards C_{c^*} . Also, if $o_j \notin C_{c^*}$, then o_j is moved towards C_{c^*} .

Algorithm 2 Moving data points of violated constraints.

Input: a violated constraint $(o_i, o_j) \in \mathcal{ML} \cup \mathcal{CL}$.

- 1: **if** $(o_i, o_j) \in \mathcal{CL}$ **then**
 - 2: Determine $u \in \{i, j\}$ and the cluster index $c^* \in M_{ij}$ with minimum value $\delta(u, c, c^*)$.
 - 3: Move o_u towards C_{c^*} .
 - 4: **else**
 - 5: Let c_i and c_j be the indices of the clusters that contain o_i and o_j , respectively.
 - 6: Determine the cluster index $c^* \in M_{ij}$ with minimum value $\delta(i, c_i, c^*) + \delta(j, c_j, c^*)$.
 - 7: **If** $o_i \notin C_{c^*}$ **then** move o_i towards C_{c^*} .
 - 8: **If** $o_j \notin C_{c^*}$ **then** move o_j towards C_{c^*} .
 - 9: **end if**
-

We now demonstrate that algorithm 1 produces high-quality clustering solutions while maintaining geometrical properties of the data as much as possible. For this purpose, we compare it to a baseline procedure that iteratively selects, at random, one violated constraint, and then performs the move defined in Algorithm 2.

The algorithms are tested on three real-world datasets, namely the **Iris** dataset mentioned in Table 1, the **Seed** dataset that contains 210 samples with seven features, and the **Optical Digits** dataset composed of 3823 handwritten digits images of 8×8 pixels. These datasets were selected because (i) their ground-truth partitions are sufficiently different from the solutions obtained by optimizing the corresponding unsupervised MSSC model, and (ii) clustering them using the unsupervised MSSC model is still effective in terms of the obtained ARI indices, thereby reducing the number of necessary pairwise constraints to get close to the ground-truth partitions.

To get the solution X in step 3 of Algorithm 1, we executed 100 times the algorithm k -means, which is the most popular heuristic for solving unsupervised MSSC clustering problems, and we then set X equal to the solution with minimum MSSC value.

For each dataset, we generated 20 constraint sets. All these sets are composed of pairwise constraints, randomly chosen among those that are violated by the initial solution X . For the **Iris**, instance, these sets contain 15 constraints, while their number is 20 and 400 for the **Seed** and **Optical** datasets, respectively.

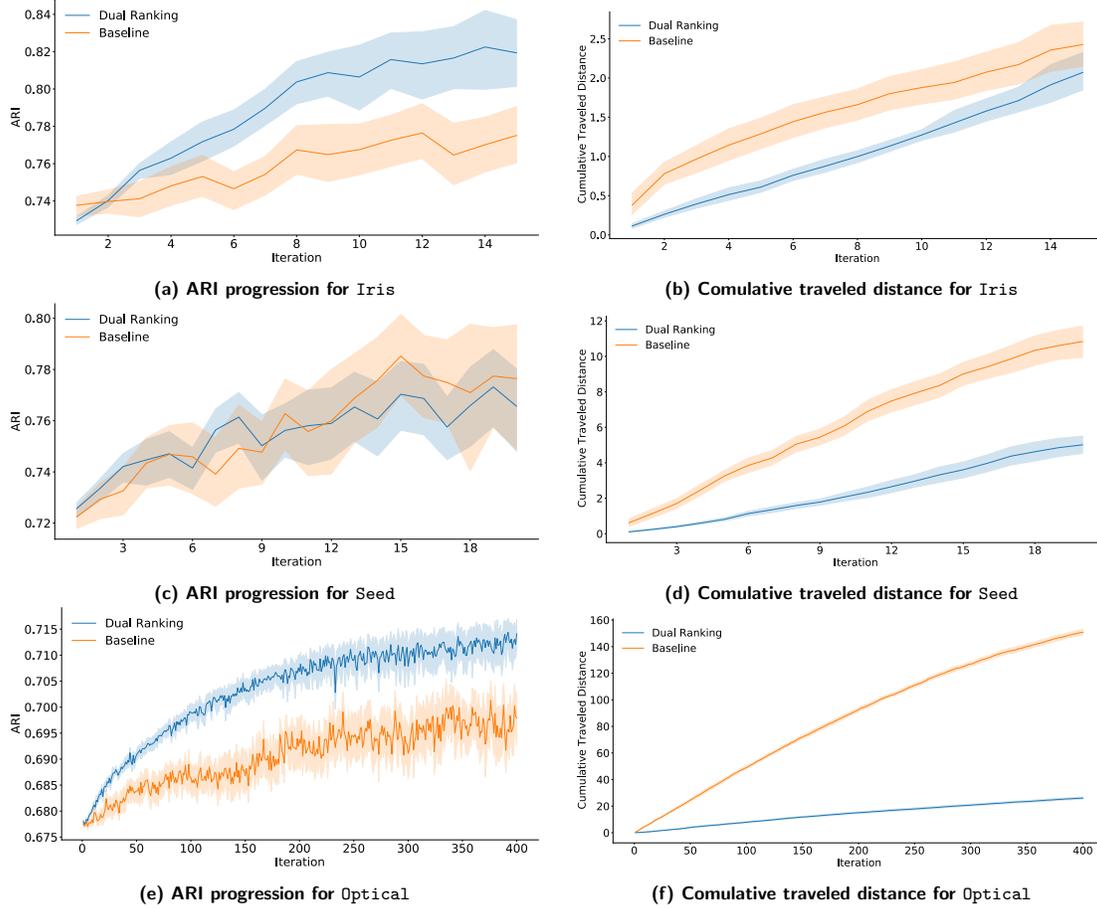


Figure 6: ARI progression and cumulative Euclidian traveled distance for Algorithm 1 and the baseline method applied to three real-world instances.

At the end of each iteration, we computed the ARI of the current solution X as well as the cumulative Euclidean distance traveled by the transformed data points, these moves resulting from running Algorithm 2. The results produced by Algorithm 1 and by the baseline method are reported in Figure 6. The line and the band refer to the mean and standard deviation with respect to the 20 runs of each method.

Our first observation is that for each set of m constraints ($m = 15, 20$ or 400), each execution of Algorithm 1 only required m iterations to satisfy them all, which means that each move resulting from a run of Algorithm 2 allowed to satisfy the violated constraint given in input.

Next, for all datasets, we observe an upward progression of the ARI both for the proposed algorithm that exploits the dual information and for the baseline method. This behavior asserts the effectiveness of the designed distance metric learning method to leverage information from the provided pairwise constraints. More interestingly, we highlight that Algorithm 1 improves the clustering quality faster (except for *Seed*) and with less significant transformations of the original dataset when compared to the baseline algorithm that does not exploit dual information.

6 Filtering useful pairwise constraints

Semi-supervised clustering techniques use constraints to guide an algorithm towards better clustering solutions. However, some constraints can have a negative effect on the clustering task (Davidson et al.,

2006), and it is not necessarily easy to identify which constraints are useful or harmful to the clustering process by basing this identification on a solution produced by an algorithm.

We propose to use the information associated with the dual variables to rank the pairwise constraints according to the likelihood that they contribute favorably to the semi-supervised clustering process. We then only use the best ranked constraints, this selection aiming to reduce the number of constraints to consider without compromising the quality of the solution produced. As in the previous section, the pairwise constraints $(o_i, o_j) \in \mathcal{ML} \cup \mathcal{CL}$ are ordered according to their associated impact \mathcal{I}_{ij} .

A particular family of methods that can greatly benefit from this ranking approach are those designed to handle massive datasets, where the time required to process the data itself and the set of constraints is crucial for their use in the practice. We demonstrate the effectiveness of our approach by integrating it into the semi-supervised deep learning framework recently proposed by Zhang et al. (2021). Their algorithm can be summarized as follows.

The algorithm first uses a deep neural network to determine an embedding from the data space to a lower-dimensional feature space Z . The algorithm then computes soft and hard allocation values q_{ic} and p_{ic} in $[0, 1]$ which, roughly speaking, can be interpreted as the probability of assigning sample o_i to cluster C_c . In the sequel, it defines a clustering loss function ℓ_C which is defined as

$$\ell_C = \sum_{i=1}^n \sum_{c=1}^k p_i^c \log \frac{p_i^c}{q_i^c},$$

as well as a loss $\ell_{\mathcal{ML}}$ for the must-link constraints, and a loss $\ell_{\mathcal{CL}}$ for the cannot-link constraints, where $\ell_{\mathcal{ML}}$ and $\ell_{\mathcal{CL}}$ are defined as

$$\ell_{\mathcal{ML}} = \sum_{(o_i, o_j) \in \mathcal{ML}} \log \sum_{c=1}^k q_i^c q_j^c \quad \text{and} \quad \ell_{\mathcal{CL}} = \sum_{(o_i, o_j) \in \mathcal{CL}} \log(1 - \sum_{c=1}^k q_i^c q_j^c).$$

The three loss functions ℓ_C , ℓ_{ML} and ℓ_{CL} are then used to update the parameters of the neural network and to obtain a new mapping of the data points in Z . This process is repeated at most m times, where m is a parameter, but may stop earlier if the ratio of changed cluster assignments between two consecutive iterations is smaller than 0.001, each data point o_i being assigned to the cluster c that maximizes q_i^c . For more details, the reader is referred to Zhang et al. (2021).

Instead of defining ℓ_{ML} and ℓ_{CL} by considering all the available pairwise constraints, we propose to use only a subset. More precisely, we run the sub-gradient algorithm to determine optimal dual values η_{ij}^c , λ_{ij}^c and γ_{ij}^c for the clustering problem based on the embedding of the points in Z . As in the previous section, these dual values are used to define the impact score \mathcal{I}_{ij} of each constraint $(o_i, o_j) \in \mathcal{ML} \cup \mathcal{CL}$, where \mathcal{I}_{ij} is defined as the sum of the values of the dual variables associated with that constraint. Let Ω be the subset of pairwise constraints having a strictly negative impact : we eliminate the $\lfloor \alpha \rfloor |\Omega|$ constraints of Ω which have the smallest impact score, where $\alpha \in [0, 1]$ is a parameter, and we denote by Ω^α the set of remaining constraints. By defining $\mathcal{ML}^\alpha = \mathcal{ML} \cap \Omega^\alpha$ and $\mathcal{CL}^\alpha = \mathcal{CL} \cap \Omega^\alpha$, we obtain the following new loss functions $\ell_{\mathcal{ML}}^\alpha$ and $\ell_{\mathcal{CL}}^\alpha$ given by :

$$\ell_{\mathcal{ML}}^\alpha = \sum_{(o_i, o_j) \in \mathcal{ML}^\alpha} \log \sum_{c=1}^k q_i^c q_j^c \quad \text{and} \quad \ell_{\mathcal{CL}}^\alpha = \sum_{(o_i, o_j) \in \mathcal{CL}^\alpha} \log(1 - \sum_{c=1}^k q_i^c q_j^c).$$

The neural network parameters are updated by using ℓ_{ML}^α and ℓ_{CL}^α instead of ℓ_{ML} and ℓ_{CL} . Algorithm 3 describes the main steps of this integration of our dual approach in the algorithm of Zhang et al. (2021).

Algorithm 3 is compared with the original algorithm of Zhang et al. (2021) on the same three datasets they used for their experiment: MNIST which contains 60,000 handwritten digits of 28×28 pixel

Algorithm 3 Integration of the dual approach in the deep learning framework of Zhang et al. (2021).

- 1: Train the neural network to obtain an embedding of the data points in a lower dimensional feature space Z . Initialize the iteration counter $iter$ to zero.
- 2: **repeat**
- 3: Compute the soft and hard allocation values q_i^c and p_i^c as well as the loss ℓ_C .
- 4: Run the sub-gradient algorithm to produce optimal dual values $\eta_{ij}^c, \lambda_{ij}^c$ and γ_{ij}^c for the clustering problem based on the current embedding of the data points in Z .
- 5: For all $(o_i, o_j) \in \mathcal{CL} \cup \mathcal{ML}$ do $\mathcal{I}_{ij} \leftarrow \begin{cases} \sum_{c=1}^k \eta_{ij}^c & \text{if } (o_i, o_j) \in \mathcal{CL} \\ \sum_{c=1}^k (\lambda_{ij}^c + \lambda'_{ij}{}^c) & \text{if } (o_i, o_j) \in \mathcal{ML}. \end{cases}$
- 6: Set $\Omega \leftarrow \{(o_i, o_j) \in \mathcal{ML} \cup \mathcal{CL} \mid \mathcal{I}_{ij} < 0\}$, sort these constraints in increasing order of their impact score \mathcal{I}_{ij} , and remove the $\lfloor \alpha \rfloor |\Omega|$ top-ranked constraints of Ω to obtain Ω^α .
- 7: Set $\mathcal{ML}^\alpha \leftarrow \mathcal{ML} \cap \Omega^\alpha$, $\mathcal{CL}^\alpha \leftarrow \mathcal{CL} \cap \Omega^\alpha$, and compute the losses ℓ_{ML}^α and ℓ_{CL}^α .
- 8: Set $iter \leftarrow iter + 1$, update the neural network parameters by using ℓ_C , ℓ_{CL}^α and ℓ_{ML}^α , and get a new embedding of the data points in Z .
- 9: **until** $iter = m$ or the ratio of changed cluster assignments between two consecutive iterations is below 10^{-3} , each o_i being assigned to the cluster c that maximizes q_i^c .

size which have to be grouped into $k = 10$ classes, **Fashion-MNIST** which contains 60,000 grayscale 28×28 images and $k = 10$ classes, and **Reuters-10K** which contains 10,000 document texts with 2,000 features and $k = 10$ classes. For each dataset, we generated 20 sets of pairwise constraints randomly chosen from the ground-truth partition. The number of constraints in each set is equal to 10% of the size of the dataset, which amounts to 6,000 for **MNIST** and **Fashion-MNIST**, and 1,000 for **Reuters-10K**. We ran Algorithm 3 for at most $m = 500$ iterations, with $\alpha = 0.95, 0.9, 0.8, 0.7, 0.6$ and 0.0 . Note that running our algorithm with $\alpha = 0.0$ is not totally equivalent to applying the original algorithm of Zhang et al. (2021) since we do not consider pairwise constraints (o_i, o_j) with impact score $\mathcal{I}_{ij} = 0$.

We indicate in Table 2 the average ARI (mean \pm standard deviation) over the 20 constraint sets for each value of α and for the three datasets. CPU times (in seconds) and the number of iterations required to reach convergence are shown in Table 3. We observe that our strategy that uses the dual information for filtering the constraints improves the unfiltered version of Zhang et al. (2021) for all three datasets. The best ARI is obtained by removing 70 to 80% of the constraints with the lowest impact score. Note also that in addition to improving the ARI, our filtering approach does not significantly modify the CPU times and the number of iterations necessary to reach convergence. More precisely, a reduction is obtained, both in terms of time and number of iterations, for the datasets **MNIST** and **Fashion-MNIST** when using our filtering algorithm with $\alpha = 70\%$. However, such gains are not observed for the smallest dataset **Reuters-10k**, where the unfiltered algorithm converges faster than Algorithm 3.

Table 2: Mean ARI and standard deviations for Algorithm 3 and for the unfiltered original algorithm of Zhang et al. (2021).

α	Datasets		
	MNIST	Fashion-MNIST	Reuters-10K
0.95	0.9370 \pm 0.0136	0.4175 \pm 0.0244	0.5992 \pm 0.0140
0.90	0.9389 \pm 0.0083	0.4115 \pm 0.0193	0.6006 \pm 0.0157
0.80	0.9338 \pm 0.0132	0.4237 \pm 0.0250	0.6009 \pm 0.0153
0.70	0.9427 \pm 0.0077	0.4132 \pm 0.0127	0.5949 \pm 0.0154
0.60	0.9343 \pm 0.0244	0.4164 \pm 0.0101	0.5843 \pm 0.0153
0.00	0.9343 \pm 0.0161	0.4012 \pm 0.0167	0.5830 \pm 0.0120
Unfiltered	0.9314 \pm 0.0040	0.3916 \pm 0.0163	0.5995 \pm 0.0108

Table 3: CPU times (in seconds) and number of iterations to reach convergence for Algorithm 3 and for the unfiltered original algorithm of Zhang et al. (2021).

α	MNIST		Fashion-MNIST		Reuters-10K	
	CPU times	iterations	CPU times	iterations	CPU times	iterations
0.95	171.7 \pm 22.6	126.5 \pm 16.4	646.4 \pm 19.5	499.3 \pm 2.6	33.4 \pm 5.0	67.1 \pm 18.2
0.90	152.5 \pm 17.9	127.7 \pm 15.4	658.7 \pm 05.1	497.1 \pm 8.8	34.2 \pm 3.8	61.8 \pm 14.8
0.80	161.0 \pm 33.1	134.9 \pm 29.0	673.4 \pm 23.9	465.3 \pm 37.1	30.0 \pm 1.9	50.1 \pm 08.1
0.70	142.8 \pm 23.8	120.8 \pm 20.5	587.5 \pm 48.7	463.2 \pm 53.5	29.3 \pm 2.0	44.8 \pm 06.3
0.60	164.3 \pm 33.4	128.2 \pm 26.2	591.0 \pm 18.1	497.9 \pm 09.1	28.8 \pm 1.4	34.4 \pm 06.2
0.00	183.6 \pm 22.9	137.4 \pm 20.2	623.8 \pm 40.3	488.3 \pm 26.5	26.6 \pm 1.2	28.6 \pm 06.3
Unfiltered	187.2 \pm 27.5	149.7 \pm 14.8	775.3 \pm 27.4	500.0 \pm 00.0	13.2 \pm 0.4	10.7 \pm 01.9

7 Conclusion

Our goal was to explore how dual information can improve distance metric learning algorithms for clustering problems with must-link and cannot-link constraints. We have shown that the optimal dual values associated to these constraints in a clustering problem are an effective tool to estimate the impact of violating a constraint. These dual values are obtained by running a sub-gradient algorithm on a Lagrangian relaxation of the original clustering problem, where pairwise constraints are replaced by penalty terms in the objective function. We have illustrated this benefit for the following three tasks.

1. **Identification of an appropriate dissimilarity measure.** We have defined a fitness score based on optimal dual values for recommending to experts a dissimilarity measure that best agrees with their beliefs and expectations.
2. **Preserving geometrical properties of the dataset.** Distance metric learning algorithms aim to move closer pairs of data points involved in must-link constraints, and to move pairs of points involved in cannot-link constraints away from each other. These transformations are often applied without worrying too much about their magnitude. We have shown how the use of dual information makes it possible to determine transformations with little impact on the original space. Our methodology therefore offers a tool to data analysts to allow them to integrate their knowledge of the domain while controlling the way in which the data is modified.
3. **Filtering useful pairwise constraints.** While it is known that some pairwise constraints can have a negative effect on the clustering task, the identification of these harmful constraints is not an easy task. We have defined an impact score, based on optimal dual values, which helps to filter the constraints that appear to be the most useful. We have demonstrated the benefit of this filter by integrating it into the recently proposed deep learning framework of Zhang et al. (2021).

References

- Aggarwal CC (2015) Similarity and distances, Springer International Publishing, Cham, pp 63–91. DOI 10.1007/978-3-319-14142-8_3
- Aloise D, Hansen P, Liberti L (2012) An improved column generation algorithm for minimum sum-of-squares clustering. *Mathematical Programming* 131(1):195–220, DOI 10.1007/s10107-010-0349-7
- Baghshah MS, Shouraki SB (2009) Semi-supervised metric learning using pairwise constraints. In: *IJCAI International Joint Conference on Artificial Intelligence*, pp 1217–1222
- Bar-Hillel A, Hertz T, Shental N, Weinshall D (2005) Learning a Mahalanobis metric from equivalence constraints. *Journal of Machine Learning Research* 6(32):937–965
- Bellet A, Habrard A, Sebban M (2015) Metric learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 9(1):1–151
- Bertsimas D, Tsitsiklis J (1997) *Introduction to linear optimization*, 1st edn. Athena Scientific

- Bilenko M, Basu S, Mooney RJ (2004) Integrating constraints and metric learning in semi-supervised clustering. Proceedings, Twenty-First International Conference on Machine Learning, ICML 2004 pp 81–88, DOI 10.1145/1015330.1015360
- Breiman L (1996) Bagging predictors. *Machine Learning* 24(2):123–140, DOI 10.1007/BF00058655
- Chang H, Yeung DY (2006) Locally linear metric adaptation with application to semi-supervised clustering and image retrieval. *Pattern Recognition* 39(7):1253–1264, DOI 10.1016/j.patcog.2005.12.012
- Davidson I, Wagstaff KL, Basu S (2006) Measuring constraint-set utility for partitional clustering algorithms. In: Fürnkranz J, Scheffer T, Spiliopoulou M (eds) *Knowledge Discovery in Databases: PKDD 2006*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 115–126
- Dua D, Graff C (2017) UCI machine learning repository. URL <http://archive.ics.uci.edu/ml>
- Fisher ML (2004) The Lagrangian relaxation method for solving integer programming problems. *Management Science* 50(12_supplement):1861–1871, DOI 10.1287/mnsc.1040.0263
- Fogel S, Averbuch-Elor H, Cohen-Or D, Goldberger J (2019) Clustering-driven deep embedding with pairwise constraints. *IEEE Computer Graphics and Applications* 39(4):16–27, DOI 10.1109/MCG.2018.2881524
- Hansen P, Jaumard B (1997) Cluster analysis and mathematical programming. *Mathematical Programming* 79(1):191–215, DOI 10.1007/BF02614317
- Hansen P, Mladenović N (2001) J-means: a new local search heuristic for minimum sum of squares clustering. *Pattern Recognition* 34(2):405–413, DOI [https://doi.org/10.1016/S0031-3203\(99\)00216-2](https://doi.org/10.1016/S0031-3203(99)00216-2)
- Hsu YC, Kira Z (2015) Neural network-based clustering using pairwise constraints. *Arxiv.org/abs/1511.06321*, [1511.06321](https://arxiv.org/abs/1511.06321)
- Hubert L, Arabie P (1985) Comparing partitions. *Journal of Classification* 2(1):193–218, DOI 10.1007/BF01908075
- Kalintha W, Ono S, Numao M, Fukui KI (2017) Kernelized evolutionary distance metric learning for semi-supervised clustering. In: 31st AAAI Conference on Artificial Intelligence, AAAI 2017, vol 3, pp 4945–4946, DOI 10.3233/ida-184283
- Kaufman L, Rousseeuw PJ (2009) *Finding groups in data: an introduction to cluster analysis*, vol 344. John Wiley & Sons
- Kullback S, Leibler RA (1951) On information and sufficiency. *Ann Math Statist* 22(1):79–86, DOI 10.1214/aoms/1177729694
- Le Capitaine H (2018) Constraint selection in metric learning. *Knowledge-Based Systems* 146:91–103, DOI <https://doi.org/10.1016/j.knosys.2018.01.026>
- Lei W, Jin R, Hoit SC, Zhu J, Yu N (2009) Learning Bregman distance functions and its application for semi-supervised clustering. In: *Advances in Neural Information Processing Systems 22 — Proceedings of the 2009 Conference*, pp 2089–2097
- Mahalanobis PC (1936) On the generalized distance in statistics. *Proceedings of the National Institute of Sciences (Calcutta)* 2:49–55
- Manning C, Schütze H (1999) *Foundations of statistical natural language processing*. MIT press
- Nguyen B, Baets BD (2019) Kernel-based distance metric learning for supervised k -means clustering. *IEEE Transactions on Neural Networks and Learning Systems* 30(10):3084–3095, DOI 10.1109/TNNLS.2018.2890021
- Pichery C (2014) Sensitivity analysis. In: Wexler P (ed) *Encyclopedia of Toxicology (Third Edition)*, third edition edn, Academic Press, Oxford, pp 236–237, DOI <https://doi.org/10.1016/B978-0-12-386454-3.00431-0>
- Randel RA, Aloise D, Blanchard SJ, Hertz A (2021) A Lagrangian-based score for assessing the quality of pairwise constraints in semi-supervised clustering. *Data Mining and Knowledge Discovery*
- Rao RV (2013) *A novel weighted Euclidean distance-based approach*, Springer London, London, pp 159–191. DOI 10.1007/978-1-4471-4375-8_5
- Rodrigues I, Aloise D, Fernandes ER (2020) An effective sequence alignment method for duplicate crash report detection. In: 4th International Workshop on Machine Learning Techniques for Software Quality Evolution (MaLTeSQuE2020)
- Wu L, Hoi SC, Jin R, Zhu J, Yu N (2012) Learning Bregman distance functions for semi-supervised clustering. *IEEE Transactions on Knowledge and Data Engineering* 24(3):478–491, DOI 10.1109/TKDE.2010.215
- Xiang S, Nie F, Zhang C (2008) Learning a Mahalanobis distance metric for data clustering and classification. *Pattern Recognition* 41(12):3600–3612, DOI 10.1016/j.patcog.2008.05.018
- Xie J, Girshick R, Farhadi A (2016) Unsupervised deep embedding for clustering analysis. [1511.06335](https://arxiv.org/abs/1511.06335)

- Xing EP, Jordan MI, Russell SJ, Ng AY (2003) Distance metric learning with application to clustering with side-information. *Advances in Neural Information Processing Systems* 15 pp 521–528
- Zhang D, Chen S, Zhou ZH (2008) Constraint score: a new filter method for feature selection with pairwise constraints. *Pattern Recognition* 41(5):1440–1451, DOI <https://doi.org/10.1016/j.patcog.2007.10.009>
- Zhang H, Zhan T, Basu S, Davidson I (2021) A framework for deep constrained clustering. *Data Mining and Knowledge Discovery* 35(2):593–620, DOI 10.1007/s10618-020-00734-4