



HAL
open science

Robustness and Generalization for Metric Learning

Aurélien Bellet, Amaury Habrard

► **To cite this version:**

Aurélien Bellet, Amaury Habrard. Robustness and Generalization for Metric Learning. Neurocomputing, 2015, pp.16. 10.1016/j.neucom.2014.09.044 . hal-01075370

HAL Id: hal-01075370

<https://hal.science/hal-01075370>

Submitted on 21 Oct 2014

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.

Robustness and Generalization for Metric Learning[☆]

Aurélien Bellet^{a,1,*}, Amaury Habrard^b

^a*Department of Computer Science, University of Southern California, Los Angeles, CA 90089, USA*

^b*Laboratoire Hubert Curien UMR 5516, Université de Saint-Etienne, 18 rue Benoit Lauras, 42000 St-Etienne, France*

Abstract

Metric learning has attracted a lot of interest over the last decade, but the generalization ability of such methods has not been thoroughly studied. In this paper, we introduce an adaptation of the notion of algorithmic robustness (previously introduced by Xu and Mannor) that can be used to derive generalization bounds for metric learning. We further show that a weak notion of robustness is in fact a necessary and sufficient condition for a metric learning algorithm to generalize. To illustrate the applicability of the proposed framework, we derive generalization results for a large family of existing metric learning algorithms, including some sparse formulations that are not covered by previous results.

Keywords: Metric learning, Algorithmic robustness, Generalization bounds

1. Introduction

Metric learning consists in automatically adjusting a distance or similarity function using training examples. The resulting metric is tailored to the problem of interest and can lead to dramatic improvement in classification, clustering or ranking performance. For this reason, metric learning has attracted a lot of interest for the past decade (see [1, 2] for recent surveys).

*Corresponding author

Email addresses: bellet@usc.edu (Aurélien Bellet),
amaury.habrard@univ-st-etienne.fr (Amaury Habrard)

¹Most of the work in this paper was carried out while the author was affiliated with Laboratoire Hubert Curien UMR 5516, Université de Saint-Etienne, France.

Existing approaches rely on the principle that pairs of examples with the same (resp. different) labels should be close to each other (resp. far away) under a good metric. Learning thus generally consists in finding the best parameters of the metric function given a set of labeled pairs.² Many methods focus on learning a Mahalanobis distance, which is parameterized by a positive semi-definite (PSD) matrix and can be seen as finding a linear projection of the data to a space where the Euclidean distance performs well on the training pairs (see for instance [3, 4, 5, 6, 7, 8, 9]). More flexible metrics have also been considered, such as similarity functions without PSD constraint [10, 11, 12]. The resulting distance or similarity is used to improve the performance of a metric-based algorithm such as k -nearest neighbors [5, 7], linear separators [12, 13], K -Means clustering [3] or ranking [9].

Despite the practical success of metric learning, little work has gone into a formal analysis of the generalization ability of the resulting metrics on unseen data. The main reason for this lack of results is that metric learning violates the common assumption of independent and identically distributed (IID) data. Indeed, the training pairs are generally given by an expert and/or extracted from a sample of individual instances, by considering all possible pairs or only a subset based for instance on the nearest or farthest neighbors of each example, some criterion of diversity [14] or a random sample. Online learning algorithms [15, 6, 10] can still offer some guarantees in this setting, but only in the form of regret bounds assessing the deviation between the cumulative loss suffered by the online algorithm and the loss induced by the best hypothesis that can be chosen in hindsight. These may be converted into proper generalization bounds under restrictive assumptions [16]. Apart from these results on online metric learning, very few papers have looked at the generalization ability of batch methods. The approach of Bian and Tao [17, 18] uses a statistical analysis to give generalization guarantees for loss minimization approaches, but their results rely on restrictive assumptions on the distribution of the examples and do not take into account any regularization on the metric. Jin et al. [19] adapted the framework of uniform stability [20] to regularized metric learning. However, their approach is based on a Frobenius norm regularizer and cannot be applied to other type of regularization, in particular sparsity-inducing norms [21] that are used in many recent

²Some methods use triplets (x, y, z) such that x should be closer to y than to z , where x and y share the same label, but not z .

metric learning approaches [22, 8, 23, 9]. Independently and in parallel to our work,³ Cao et al. [25] proposed a framework based on Rademacher analysis, which is general but rather complex and limited to pair constraints.

In this paper, we propose to study the generalization ability of metric learning algorithms according to a notion of *algorithmic robustness*. This framework, introduced by Xu et al. [26, 27], allows one to derive generalization bounds when the variation in the loss associated with two “close” training and testing examples is bounded. The notion of closeness relies on a partition of the input space into different regions such that two examples in the same region are considered close. Robustness has been successfully used to derive generalization bounds in the classic supervised learning setting, with results for SVM, LASSO, etc. We propose here to adapt algorithmic robustness to metric learning. We show that, in this context, the problem of non-IIDness of the training pairs/triplets can be worked around by simply assuming that they are built from an IID sample of labeled examples. Moreover, following [27], we provide a notion of weak robustness that is necessary and sufficient for metric learning algorithms to generalize well, confirming that robustness is a fundamental property. We illustrate the applicability of the proposed framework by deriving generalization bounds, using very few approach-specific arguments, for a family of problems that is larger than what is considered in previous work [19, 17, 18, 25]. In particular, results apply to a vast choice of regularizers, without any assumption on the distribution of the examples and using a simple proof technique.

The rest of the paper is organized as follows. We introduce some preliminaries and notations in Section 2. Our notion of algorithmic robustness for metric learning is presented in Section 3. The necessity and sufficiency of weak robustness is shown in Section 4. Section 5 illustrates the wide applicability of our framework by deriving bounds for existing metric learning formulations. Section 6 discusses the merits and limitations of the proposed analysis compared to related work, and we conclude in Section 7.

2. Preliminaries

2.1. Notations

Let X be the instance space, Y be a finite label set and let $\mathcal{Z} = X \times Y$. In the following, $z = (x, y) \in \mathcal{Z}$ means $x \in X$ and $y \in Y$. Let μ be an

³We posted a preliminary version of the present work on arXiv in 2012 [24].

unknown probability distribution over \mathcal{Z} . We assume that X is a compact convex metric space w.r.t. a norm $\|\cdot\|$ such that $X \subset \mathbb{R}^d$, thus there exists a constant R such that $\forall x \in X, \|x\| \leq R$. A similarity or distance function is a pairwise function $f : X \times X \rightarrow \mathbb{R}$. In the following, we use the generic term *metric* to refer to either a similarity or a distance function. We denote by \mathbf{s} a labeled training sample consisting of n training instances (s_1, \dots, s_n) drawn IID from μ . The sample of all possible pairs built from \mathbf{s} is denoted by $p_{\mathbf{s}}$ such that $p_{\mathbf{s}} = \{(s_1, s_1), \dots, (s_1, s_n), \dots, (s_n, s_n)\}$. A metric learning algorithm \mathcal{A} takes as input a finite set of pairs from $(\mathcal{Z} \times \mathcal{Z})^n$ and outputs a metric. We denote by $\mathcal{A}_{p_{\mathbf{s}}}$ the metric learned by an algorithm \mathcal{A} from a sample $p_{\mathbf{s}}$ of pairs. For any pair of labeled examples (z, z') and any metric f , we associate a loss function $l(f, z, z')$ which depends on the examples and their labels. This loss is assumed to be nonnegative and uniformly bounded by a constant B . We define the generalization loss (or true loss) over μ as

$$\mathcal{L}(f) = \mathbb{E}_{z, z' \sim \mu} l(f, z, z'),$$

and the empirical loss over the sample $p_{\mathbf{s}}$ as

$$l_{emp}(f) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n l(f, s_i, s_j) = \frac{1}{n^2} \sum_{(s_i, s_j) \in p_{\mathbf{s}}} l(f, s_i, s_j).$$

We are interested in bounding the deviation between $l_{emp}(f)$ and $\mathcal{L}(f)$.

2.2. Algorithmic Robustness in Classic Supervised Learning

The notion of algorithmic robustness, introduced by Xu and Mannor [26, 27] in the context of classic supervised learning, is based on the deviation between the loss associated with two training and testing instances that are “close”. Formally, an algorithm is said $(K, \epsilon(\mathbf{s}))$ -robust if there exists a partition of the space $\mathcal{Z} = X \times Y$ into K disjoint subsets such that for every training and testing instances belonging to the same region of the partition, the variation in their associated loss is bounded by a term $\epsilon(\mathbf{s})$. From this definition, the authors have proved a bound for the difference between the empirical loss and the true loss that has the form

$$\epsilon(\mathbf{s}) + B \sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}}, \quad (1)$$

with probability $1 - \delta$. This bound depends on K and $\epsilon(\mathbf{s})$. The latter should tend to zero as K increases to ensure that (1) also goes to zero when $n \rightarrow \infty$.⁴ When considering metric spaces, the partition of \mathcal{Z} can be obtained by the notion of covering number [28].

Definition 1. For a metric space (X, ρ) , and $T \subset X$, we say that $\hat{T} \subset T$ is a γ -cover of T , if $\forall t \in T, \exists \hat{t} \in \hat{T}$ such that $\rho(t, \hat{t}) \leq \gamma$. The γ -covering number of T is

$$\mathcal{N}(\gamma, T, \rho) = \min\{|\hat{T}| : \hat{T} \text{ is a } \gamma\text{-cover of } T\}.$$

When X is a compact convex space, for any $\gamma > 0$, the quantity $\mathcal{N}(\gamma, X, \rho)$ is finite leading to a finite cover. If we consider the space \mathcal{Z} , note that the label set can be partitioned into $|Y|$ sets. Thus, \mathcal{Z} can be partitioned into $|Y|\mathcal{N}(\gamma, X, \rho)$ subsets such that if two instances $z_1 = (x_1, y_1)$, $z_2 = (x_2, y_2)$ belong to the same subset, then $y_1 = y_2$ and $\rho(x_1, x_2) \leq \gamma$.

3. Robustness and Generalization for Metric Learning

We present here our adaptation of robustness to metric learning. The idea is to use the partition of \mathcal{Z} at the pair level: if a new test pair of examples is close to a training pair, then the loss value for each pair must be close. Two pairs are close when each instance of the first pair fall into the same subset of the partition of \mathcal{Z} as the corresponding instance of the other pair, as shown in Figure 1. A metric learning algorithm with this property is said robust. This notion is formalized as follows.

Definition 2. An algorithm \mathcal{A} is $(K, \epsilon(\cdot))$ robust for $K \in \mathbb{N}$ and $\epsilon(\cdot) : (\mathcal{Z} \times \mathcal{Z})^n \rightarrow \mathbb{R}$ if \mathcal{Z} can be partitioned into K disjoint sets, denoted by $\{C_i\}_{i=1}^K$, such that for all sample $\mathbf{s} \in \mathcal{Z}^n$ and the pair set $p(\mathbf{s})$ associated to this sample, the following holds:

$\forall (s_1, s_2) \in p(\mathbf{s}), \forall z_1, z_2 \in \mathcal{Z}, \forall i, j = 1, \dots, K : \text{if } s_1, z_1 \in C_i \text{ and } s_2, z_2 \in C_j$
then

$$|l(\mathcal{A}_{p_{\mathbf{s}}}, s_1, s_2) - l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2)| \leq \epsilon(p_{\mathbf{s}}). \quad (2)$$

⁴This point will be made clear by the examples provided in Section 5.




Figure 1: Illustration of the robustness property in the classic and metric learning settings. In this example, we use a cover based on the L_1 norm. In the classic definition, if any example z' falls in the same region C_i as a training example z , then the deviation between their loss must be bounded. In the metric learning definition proposed in this work, for any pair (z, z') and a training pair (z_1, z_2) , if z, z_1 belong to some region C_i and z', z_2 to some region C_j , then the deviation between the loss of these two pairs must be bounded.

K and $\epsilon(\cdot)$ quantify the robustness of the algorithm and depend on the training sample. The property of robustness is required for every training pair of the sample; we will later see that this property can be relaxed.

Note that this definition of robustness can be easily extended to triplet based metric learning algorithms. Instead of considering all the pairs $p_{\mathbf{s}}$ from an IID sample \mathbf{s} , we take the admissible triplet set $trip_{\mathbf{s}}$ of \mathbf{s} such that $(s_1, s_2, s_3) \in trip_{\mathbf{s}}$ means s_1 and s_2 share the same label while s_1 and s_3 have different ones, with the interpretation that s_1 must be more similar to s_2 than to s_3 . The robustness property can then be expressed by: $\forall (s_1, s_2, s_3) \in trip_{\mathbf{s}}, \forall z_1, z_2, z_3 \in \mathcal{Z}, \forall i, j, k = 1, \dots, K$: if $s_1, z_1 \in C_i$, $s_2, z_2 \in C_j$ and $s_3, z_3 \in C_k$ then

$$|l(\mathcal{A}_{trip_{\mathbf{s}}}, s_1, s_2, s_3) - l(\mathcal{A}_{trip_{\mathbf{s}}}, z_1, z_2, z_3)| \leq \epsilon(trip_{\mathbf{s}}). \quad (3)$$

3.1. Generalization of robust algorithms

We now give a PAC generalization bound for metric learning algorithms fulfilling the property of robustness (Definition 2). We first begin by presenting a concentration inequality that will help us to derive the bound.

Proposition 1 ([29]). *Let $(|N_1|, \dots, |N_K|)$ an IID multinomial random variable with parameters n and $(\mu(C_1), \dots, \mu(C_K))$. By the Bretaganolle-Huber-Carol inequality we have: $\Pr \left\{ \sum_{i=1}^K \left| \frac{|N_i|}{n} - \mu(C_i) \right| \geq \lambda \right\} \leq 2^K \exp \left(\frac{-n\lambda^2}{2} \right)$, hence with probability at least $1 - \delta$,*

$$\sum_{i=1}^K \left| \frac{N_i}{n} - \mu(C_i) \right| \leq \sqrt{\frac{2K \ln 2 + 2 \ln(1/\delta)}{n}}. \quad (4)$$

We now give our first result on the generalization of metric learning algorithms.

Theorem 1. *If a learning algorithm \mathcal{A} is $(K, \epsilon(\cdot))$ -robust and the training sample is made of the pairs $p_{\mathbf{s}}$ obtained from a sample \mathbf{s} generated by n IID draws from μ , then for any $\delta > 0$, with probability at least $1 - \delta$ we have:*

$$|\mathcal{L}(\mathcal{A}_{p_{\mathbf{s}}}) - l_{emp}(\mathcal{A}_{p_{\mathbf{s}}})| \leq \epsilon(p_{\mathbf{s}}) + 2B \sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}}.$$

Proof Let N_i be the set of index of points of \mathbf{s} that fall into the C_i . $(|N_1|, \dots, |N_K|)$ is a IID random variable with parameters n and $(\mu(C_1), \dots, \mu(C_K))$. We have:

$$\begin{aligned} & |\mathcal{L}(\mathcal{A}_{p_{\mathbf{s}}}) - l_{emp}(\mathcal{A}_{p_{\mathbf{s}}})| \\ &= \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \mu(C_i) \mu(C_j) - \frac{1}{n^2} \sum_{i,j=1}^n l(\mathcal{A}_{p_{\mathbf{s}}}, s_i, s_j) \right| \\ &\stackrel{(a)}{\leq} \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \mu(C_i) \mu(C_j) - \right. \\ &\quad \left. \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \mu(C_i) \frac{|N_j|}{n} \right| + \\ &\quad \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \mu(C_i) \frac{|N_j|}{n} - \frac{1}{n^2} \sum_{i,j=1}^n l(\mathcal{A}_{p_{\mathbf{s}}}, s_i, s_j) \right| \end{aligned}$$

$$\begin{aligned}
&\stackrel{(b)}{\leq} \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \mu(C_i) (\mu(C_j) - \frac{|N_j|}{n}) \right| + \\
&\left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \mu(C_i) \frac{|N_j|}{n} - \right. \\
&\quad \left. \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \frac{|N_i||N_j|}{n} \right| + \\
&\left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \frac{|N_i||N_j|}{n} - \frac{1}{n^2} \sum_{i,j=1}^n l(\mathcal{A}_{p_{\mathbf{s}}}, s_i, s_j) \right| \\
&\stackrel{(c)}{\leq} B \left(\left| \sum_{j=1}^K \mu(C_j) - \frac{|N_j|}{n} \right| + \left| \sum_{i=1}^K \mu(C_i) - \frac{|N_i|}{n} \right| \right) + \\
&\left| \frac{1}{n^2} \sum_{i,j=1}^K \sum_{s_o \in N_i} \sum_{s_l \in N_j} \max_{z \in C_i} \max_{z' \in C_j} |l(\mathcal{A}_{p_{\mathbf{s}}}, z, z') - l(\mathcal{A}_{p_{\mathbf{s}}}, s_o, s_l)| \right| \\
&\stackrel{(d)}{\leq} \epsilon(p_{\mathbf{s}}) + 2B \sum_{i=1}^K \left| \frac{|N_i|}{n} - \mu(C_i) \right| \stackrel{(e)}{\leq} \epsilon(p_{\mathbf{s}}) + 2B \sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}}.
\end{aligned}$$

Inequalities (a) and (b) are due to the triangle inequality, (c) uses the fact that l is bounded by B , that $\sum_{i=1}^K \mu(C_i) = 1$ by definition of a multinomial random variable and that $\sum_{j=1}^K \frac{|N_j|}{n} = 1$ by definition of the N_j . Lastly, (d) is due to the hypothesis of robustness (Equation 2) and (e) to the application of Proposition 1. \square

The previous bound depends on K which is given by the cover chosen for \mathcal{Z} . If for any K , the associated $\epsilon(\cdot)$ is a constant (i.e. $\epsilon_K(\mathbf{s}) = \epsilon_K$) for any \mathbf{s} , we can obtain a bound that holds uniformly for all K :

$$|\mathcal{L}(\mathcal{A}_{p_{\mathbf{s}}}) - l_{emp}(\mathcal{A}_{p_{\mathbf{s}}})| \leq \inf_{K \geq 1} \left[\epsilon_K + 2B \sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}} \right].$$

For triplet based metric learning algorithms, by following the definition of robustness given by Equation 3 and adapting straightforwardly the losses to triplets such that they output zero for non admissible ones, Theorem 1

can be easily extended to obtain the following generalization bound:

$$|\mathcal{L}(\mathcal{A}_{trip_{p_{\mathbf{s}}}}) - l_{emp}(\mathcal{A}_{trip_{p_{\mathbf{s}}})}| \leq \epsilon(trip_{p_{\mathbf{s}}}) + 3B\sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}}. \quad (5)$$

3.2. Pseudo-robustness

The previous study requires the robustness property to be satisfied for every training pair. In this section, we show that it is possible to relax the robustness such that it must hold only for a subset of the possible pairs, while still providing generalization guarantees.

Definition 3. An algorithm \mathcal{A} is $(K, \epsilon(\cdot), \hat{p}_n(\cdot))$ pseudo-robust for $K \in \mathbb{N}$, $\epsilon(\cdot) : (\mathcal{Z} \times \mathcal{Z})^n \rightarrow \mathbb{R}$ and $\hat{p}_n(\cdot) : (\mathcal{Z} \times \mathcal{Z})^n \rightarrow \{1, \dots, n^2\}$, if \mathcal{Z} can be partitioned into K disjoint sets, denoted by $\{C_i\}_{i=1}^K$, such that for all $\mathbf{s} \in \mathcal{Z}^n$ IID from μ , there exists a subset of training pairs samples $\hat{p}_{\mathbf{s}} \subseteq p_{\mathbf{s}}$, with $|\hat{p}_{\mathbf{s}}| = \hat{p}_n(p_{\mathbf{s}})$, such that the following holds:

$\forall (s_1, s_2) \in \hat{p}_{\mathbf{s}}, \forall z_1, z_2 \in \mathcal{Z}, \forall i, j = 1, \dots, K$: if $s_1, z_1 \in C_i$ and $s_2, z_2 \in C_j$ then

$$|l(\mathcal{A}_{p_{\mathbf{s}}}, s_1, s_2) - l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2)| \leq \epsilon(p_{\mathbf{s}}). \quad (6)$$

We can easily observe that $(K, \epsilon(\cdot))$ -robust is equivalent to $(K, \epsilon(\cdot), n^2)$ pseudo-robust. The following theorem gives the generalization guarantees associated with the pseudo-robustness property.

Theorem 2. If a learning algorithm \mathcal{A} is $(K, \epsilon(\cdot), \hat{p}_n(\cdot))$ pseudo-robust, the training pairs $p_{\mathbf{s}}$ come from a sample generated by n IID draws from μ , then for any $\delta > 0$, with probability at least $1 - \delta$ we have:

$$|\mathcal{L}(\mathcal{A}_{p_{\mathbf{s}}}) - l_{emp}(\mathcal{A}_{p_{\mathbf{s}}})| \leq \frac{\hat{p}_n(p_{\mathbf{s}})}{n^2} \epsilon(p_{\mathbf{s}}) + B \left(\frac{n^2 - \hat{p}_n(p_{\mathbf{s}})}{n^2} + 2\sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}} \right).$$

Proof It is similar to that of Theorem 1 and is given in Appendix A. \square

This notion of pseudo-robustness is very relevant to metric learning. Indeed, it is often difficult and potentially damaging to optimize the metric with respect to all possible pairs, and it has been observed in practice that focusing on a subset of carefully-selected pairs (e.g., defined according to nearest-neighbors) gives much better generalization performance [7, 12]. Theorem 2 confirms that this principle is well-founded: as long as the robustness property is fulfilled for a (large enough) subset of the pairs, the resulting metric has generalization guarantees. Note that this notion of pseudo-robustness can be also easily adapted to triplet based metric learning.

4. Necessity of Robustness

We prove here that a notion of weak robustness is actually necessary and sufficient to generalize in a metric learning setup. This result is based on an asymptotic analysis following the work of Xu and Mannor [27]. We consider pairs of instances coming from an increasing sample of training instances $\mathbf{s} = (s_1, s_2, \dots)$ and from a sample of test instances $\mathbf{t} = (t_1, t_2, \dots)$ such that both samples are assumed to be drawn IID from a distribution μ . We use $\mathbf{s}(n)$ and $\mathbf{t}(n)$ to denote the first n examples of the two samples respectively, while \mathbf{s}^* denotes a fixed sequence of examples.

We use $L(f, p_{\mathbf{t}(n)}) = \frac{1}{n^2} \sum_{(s_i, s_j) \in p_{\mathbf{t}(n)}} l(f, s_i, s_j)$ to refer to the average loss given a set of pairs for any learned metric f , and $\mathcal{L}(f) = \mathbb{E}_{z, z' \sim \mu} l(f, z, z')$ for the expected loss.

We first define a notion of generalizability for metric learning.

Definition 4. *Given a training pair set $p_{\mathbf{s}^*}$ coming from a sequence of examples \mathbf{s}^* , a metric learning method \mathcal{A} generalizes w.r.t. $p_{\mathbf{s}^*}$ if*

$$\lim_n \left| \mathcal{L}(\mathcal{A}_{p_{\mathbf{s}^*(n)}}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)}) \right| = 0.$$

Furthermore, a learning method \mathcal{A} generalizes with probability 1 if it generalizes with respect to the pairs $p_{\mathbf{s}}$ of almost all samples \mathbf{s} IID from μ .

Note this notion of generalizability implies convergence in mean. We then introduce the notion of weak robustness for metric learning.

Definition 5. *Given a set of training pairs $p_{\mathbf{s}^*}$ coming from a sequence of examples \mathbf{s}^* , a metric learning method \mathcal{A} is weakly robust with respect to $p_{\mathbf{s}^*}$ if there exists a sequence of $\{\mathcal{D}_n \subseteq \mathcal{Z}^n\}$ such that $\Pr(\mathbf{t}(n) \in \mathcal{D}_n) \rightarrow 1$ and*

$$\lim_n \left\{ \max_{\hat{\mathbf{s}}(n) \in \mathcal{D}_n} \left| L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\hat{\mathbf{s}}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)}) \right| \right\} = 0.$$

Furthermore, a learning method \mathcal{A} is almost surely weakly robust if it is robust with respect to almost all \mathbf{s} .

Recall that the definition of robustness requires the labeled sample space to be partitioned into disjoint subsets such that if some instances of pairs of train/test examples belong to the same partition, then they have similar loss. Weak robustness is a generalization of this notion where we consider the

average loss of testing and training pairs: if for a large (in the probabilistic sense) subset of data, the testing loss is close to the training loss, then the algorithm is weakly robust. From Proposition 1, we can see that if for any fixed $\epsilon > 0$ there exists K such that an algorithm \mathcal{A} is (K, ϵ) robust, then \mathcal{A} is weakly robust. We now give the main result of this section about the necessity of robustness.

Theorem 3. *Given a fixed sequence of training examples \mathbf{s}^* , a metric learning method \mathcal{A} generalizes with respect to $p_{\mathbf{s}^*}$ if and only if it is weakly robust with respect to $p_{\mathbf{s}^*}$.*

Proof Following [27], the sufficiency is obtained by the fact that the testing pairs are obtained from a sample $\mathbf{t}(n)$ constituted of n IID instances. We give the proof in Appendix B.

For the necessity, we need the following lemma which is a direct adaptation of a result introduced in [27] (Lemma 2). We provide the proof in Appendix C for the sake of completeness.

Lemma 1. *Given \mathbf{s}^* , if a learning method is not weakly robust w.r.t. $p_{\mathbf{s}^*}$, there exists $\epsilon^*, \delta^* > 0$ such that the following holds for infinitely many n :*

$$\Pr(|L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{t}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \geq \epsilon^*) \geq \delta^*. \quad (7)$$

Now, recall that l is positive and uniformly bounded by B , thus by the McDiarmid inequality (recalled in Appendix D) we have that for any $\epsilon, \delta > 0$ there exists an index n^* such that for any $n > n^*$, with probability at least $1 - \delta$, we have $|\frac{1}{n^2} \sum_{(t_i, t_j) \in p_{\mathbf{t}(n)}} l(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, t_i, t_j) - \mathcal{L}(\mathcal{A}_{p_{\mathbf{s}^*(n)}})| \leq \epsilon$. This implies the convergence $\mathcal{L}(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{t}(n)}) - \mathcal{L}(\mathcal{A}_{p_{\mathbf{s}^*(n)}}) \xrightarrow{Pr} 0$, and thus from a given index:

$$|\mathcal{L}(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{t}(n)}) - \mathcal{L}(\mathcal{A}_{p_{\mathbf{s}^*(n)}})| \leq \frac{\epsilon^*}{2}. \quad (8)$$

Now, by contradiction, suppose algorithm \mathcal{A} is not weakly robust, Lemma 1 implies Equation 7 holds for infinitely many n . This combined with Equation 8 implies that for infinitely many n :

$$|\mathcal{L}(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{t}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \geq \frac{\epsilon^*}{2}$$

which means \mathcal{A} does not generalize, thus the necessity of weak robustness is established. \square

The following corollary follows immediately from Theorem 3.

Corollary 1. *A metric learning method \mathcal{A} generalizes with probability 1 if and only if it is almost surely weakly robust.*

5. Examples of Robust Metric Learning Algorithms

We first restrict our attention to Mahalanobis distance learning algorithms of the following form:

$$\min_{\mathbf{M} \succeq 0} c\|\mathbf{M}\| + \frac{1}{n^2} \sum_{(s_i, s_j) \in p_s} g(y_{ij}[1 - f(\mathbf{M}, x_i, x_j)]), \quad (9)$$

where $s_i = (x_i, y_i)$, $s_j = (x_j, y_j)$, $y_{ij} = 1$ if $y_i = y_j$ and -1 otherwise, $f(\mathbf{M}, x_i, x_j) = (x_i - x_j)^T \mathbf{M} (x_i - x_j)$ is the Mahalanobis distance parameterized by the $d \times d$ PSD matrix \mathbf{M} , $\|\cdot\|$ some matrix norm and c a regularization parameter. The loss function $l(f, s_i, s_j) = g(y_{ij}[1 - f(\mathbf{M}, x_i, x_j)])$ outputs a small value when its input is large positive and a large value when it is large negative. We assume g to be nonnegative and Lipschitz continuous with Lipschitz constant U . Lastly, $g_0 = \sup_{s_i, s_j} g(y_{ij}[1 - f(\mathbf{0}, x_i, x_j)])$ is the largest loss when \mathbf{M} is $\mathbf{0}$. The general form (9) encompasses many existing metric learning formulations. For instance, in the case of the hinge loss and Frobenius norm regularization, we recover [19], while the family of formulations studied in [30] corresponds to a trace norm regularizer.

To prove the robustness of (9), we will use the following theorem, which is based on the geometric intuition behind robustness. It essentially says that if a metric learning algorithm achieves approximately the same testing loss for testing pairs that are close to each other, then it is robust.

Theorem 4. *Fix $\gamma > 0$ and a metric ρ of \mathcal{Z} . Suppose \mathcal{A} satisfies:*

$$\forall z_1, z_2, z'_1, z'_2 : z_1, z_2 \in \mathbf{s}, \rho(z_1, z'_1) \leq \gamma, \rho(z_2, z'_2) \leq \gamma,$$

$$|l(\mathcal{A}_{p_s}, z_1, z_2) - l(\mathcal{A}_{p_s}, z'_1, z'_2)| \leq \epsilon(p_s)$$

and $\mathcal{N}(\gamma/2, \mathcal{Z}, \rho) < \infty$. Then \mathcal{A} is $(\mathcal{N}(\gamma/2, \mathcal{Z}, \rho), \epsilon(p_s))$ -robust.

Proof By definition of covering number, we can partition X in $\mathcal{N}(\gamma/2, X, \rho)$ subsets such that each subset has a diameter less or equal to γ . Furthermore, since Y is a finite set, we can partition \mathcal{Z} into $|Y|\mathcal{N}(\gamma/2, X, \rho)$ subsets $\{C_i\}$

such that $z_1, z'_1 \in C_i \Rightarrow \rho(z_1, z'_1) \leq \gamma$. Therefore, $\forall z_1, z_2, z'_1, z'_2 : z_1, z_2 \in \mathbf{s}, \rho(z_1, z'_1) \leq \gamma, \rho(z_2, z'_2) \leq \gamma$,

$$|l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) - l(\mathcal{A}_{p_{\mathbf{s}}}, z'_1, z'_2)| \leq \epsilon(p_{\mathbf{s}}),$$

this implies $z_1, z_2 \in \mathbf{s}, z_1, z'_1 \in C_i, z_2, z'_2 \in C_j \Rightarrow |l(\mathcal{A}_{p_{\mathbf{s}}}, z_1, z_2) - l(\mathcal{A}_{p_{\mathbf{s}}}, z'_1, z'_2)| \leq \epsilon(p_{\mathbf{s}})$, which establishes the theorem. \square

We provide a similar theorem for triplet-based approaches in Appendix E. This theorem provides a roadmap for deriving generalization guarantees based on the robustness framework. Indeed, given a partition of the input space, one must bound the deviation between the loss for any pair of examples with corresponding elements belonging to the same partitions. This bound is generally a constant that depends on the problem to solve and the thinness of the partition defined by γ . This bound tends to zero as $\gamma \rightarrow 0$, which ensures the consistency of the approach. While this framework is rather general, the price to pay is the relative looseness of the bounds, as discussed in Section 6.

Recall that we assume that $\forall x \in X, \|x\| \leq R$ for some convenient norm $\|\cdot\|$. Following Theorem 4, we now prove the robustness of (9) when $\|\mathbf{M}\|$ is the Frobenius norm.

Example 1 (Frobenius norm). *Algorithm (9) with the Frobenius norm $\|\mathbf{M}\| = \|\mathbf{M}\|_{\mathcal{F}} = \sqrt{\sum_{i=1}^d \sum_{j=1}^d m_{ij}^2}$ is $(|Y|\mathcal{N}(\gamma/2, X, \|\cdot\|_2), \frac{8UR\gamma g_0}{c})$ -robust.*

Proof Let \mathbf{M}^* be the solution given training data $p_{\mathbf{s}}$. Thus, due to optimality of \mathbf{M}^* , we have

$$\begin{aligned} c\|\mathbf{M}^*\|_{\mathcal{F}} + \frac{1}{n^2} \sum_{(s_i, s_j) \in p_{\mathbf{s}}} g(y_{ij}[1 - f(\mathbf{M}, x_i, x_j)]) &\leq \\ c\|\mathbf{0}\|_{\mathcal{F}} + \frac{1}{n^2} \sum_{(s_i, s_j) \in p_{\mathbf{s}}} g(y_{ij}[1 - f(\mathbf{0}, x_i, x_j)]) &= g_0 \end{aligned}$$

and thus $\|\mathbf{M}^*\|_{\mathcal{F}} \leq g_0/c$. We can partition \mathcal{Z} as $|Y|\mathcal{N}(\gamma/2, X, \|\cdot\|_2)$ sets, such that if z and z' belong to the same set, then $y = y'$ and $\|x - x'\|_2 \leq \gamma$. Now, for $z_1, z_2, z'_1, z'_2 \in \mathcal{Z}$, if $y_1 = y'_1, \|x_1 - x'_1\|_2 \leq \gamma, y_2 = y'_2$ and $\|x_2 - x'_2\|_2 \leq \gamma$,

then:

$$\begin{aligned}
& |g(y_{12}[1 - f(\mathbf{M}^*, x_1, x_2)]) - g(y'_{12}[1 - f(\mathbf{M}^*, x'_1, x'_2)])| \\
& \leq U|(x_1 - x_2)^T \mathbf{M}^*(x_1 - x_2) - (x'_1 - x'_2)^T \mathbf{M}^*(x'_1 - x'_2)| \\
& = U|(x_1 - x_2)^T \mathbf{M}^*(x_1 - x_2) - (x_1 - x_2)^T \mathbf{M}^*(x'_1 - x'_2) \\
& \quad + (x_1 - x_2)^T \mathbf{M}^*(x'_1 - x'_2) - (x'_1 - x'_2)^T \mathbf{M}^*(x'_1 - x'_2)| \\
& = U|(x_1 - x_2)^T \mathbf{M}^*(x_1 - x_2 - (x'_1 + x'_2)) + \\
& \quad (x_1 - x_2 - (x'_1 + x'_2))^T \mathbf{M}^*(x'_1 + x'_2)| \\
& \leq U(|(x_1 - x_2)^T \mathbf{M}^*(x_1 - x'_1)| + |(x_1 - x_2)^T \mathbf{M}^*(x'_2 - x_2)| \\
& \quad + |(x_1 - x'_1)^T \mathbf{M}^*(x'_1 + x'_2)| + |(x'_2 - x_2)^T \mathbf{M}^*(x'_1 + x'_2)|) \\
& \leq U(\|x_1 - x_2\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x_1 - x'_1\|_2 + \|x_1 - x_2\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x'_2 - x_2\|_2 \\
& \quad + \|x_1 - x'_1\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x'_1 - x'_2\|_2 + \|x'_2 - x_2\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x'_1 - x'_2\|_2) \\
& \leq \frac{8UR\gamma g_0}{c}
\end{aligned}$$

Hence, the example holds by Theorem 4. \square

Note that for the special case of Example 1, a generalization bound (with same order of convergence rate) based on uniform stability was derived in [19]. However, it is known that sparse algorithms are not stable [21], and thus stability-based analysis fails to assess the generalization ability of recent sparse metric learning approaches [22, 23, 8, 9, 30]. The key advantage of robustness over stability is that we can obtain bounds similar to the Frobenius case for arbitrary p -norms (or even any regularizer which is bounded below by some p -norm) using equivalence of norms arguments. To illustrate this, we show the robustness when $\|\mathbf{M}\|$ is the ℓ_1 norm (used in [22, 23]) which promotes sparsity at the entry level, the $\ell_{2,1}$ norm (used e.g. in [8]) which induces sparsity at the column/row level, and the trace norm (used e.g. in [9, 30]) which favors low-rank matrices.⁵ The proofs are reminiscent of that of Example 1 and can be found in Appendix F and Appendix G, respectively.

Example 2 (ℓ_1 norm). *Algorithm (9) with $\|\mathbf{M}\| = \|\mathbf{M}\|_1$ is $(|Y|\mathcal{N}(\gamma, \mathcal{X}, \|\cdot\|_1), \frac{8UR\gamma g_0}{c})$ -robust.*

⁵In the last two cases, the linear projection space of the data induced by the learned Mahalanobis distance is of lower dimension than the original space, allowing more efficient computations and reduced memory usage.

Example 3 ($\ell_{2,1}$ norm and trace norm). Consider Algorithm (9) with $\|\mathbf{M}\| = \|\mathbf{M}\|_{2,1} = \sum_{i=1}^d \|m^i\|_2$, where m^i is the i -th column of \mathbf{M} . This algorithm is $(|Y|\mathcal{N}(\gamma, \mathcal{X}, \|\cdot\|_2), \frac{8UR\gamma g_0}{c})$ -robust. The same holds for the trace norm $\|\mathbf{M}\|_*$, which is the sum of the singular values of \mathbf{M} .

Some metric learning algorithms have kernelized versions, for instance [4, 5]. In the following example we show robustness for a kernelized formulation.

Example 4 (Kernelization). Consider the kernelized version of (9):

$$\min_{\mathbf{M} \succeq 0} c\|\mathbf{M}\|_{\mathbb{H}} + \frac{1}{n^2} \sum_{(s_i, s_j) \in p_s} g(y_{ij}[1 - f(\mathbf{M}, \phi(x_i), \phi(x_j))]), \quad (10)$$

where $\phi(\cdot)$ is a feature mapping to a kernel space \mathbb{H} , $\|\cdot\|_{\mathbb{H}}$ the norm function of \mathbb{H} and $k(\cdot, \cdot)$ the kernel function. Consider a cover of X by $\|\cdot\|_2$ (X being compact) and let $f_{\mathbb{H}}(\gamma) \triangleq \max_{a, b \in X, \|a-b\|_2 \leq \gamma} (k(a, a) + k(b, b) - 2k(a, b))$ and $B_{\gamma} = \max_{x \in X} \sqrt{k(x, x)}$. If the kernel function is continuous, B_{γ} and $f_{\mathbb{H}}$ are finite for any $\gamma > 0$ and thus Algorithm 10 is $(|Y|\mathcal{N}(\gamma, X, \|\cdot\|_2), \frac{8UB_{\gamma}\sqrt{f_{\mathbb{H}}g_0}}{c})$ -robust.

Proof Given in Appendix H. □

Finally, we illustrate the flexibility of our framework by deriving bounds for another form of metric as well as for formulations based on triplet constraints.

Example 5. Consider Algorithm (9) with the bilinear similarity $f(\mathbf{M}, x_i, x_j) = x_i^T \mathbf{M} x_j$ instead of the Mahalanobis distance, as studied in [10, 11, 12]. For the regularizers considered in Examples 1 – 3, we can improve the robustness to $2UR\gamma g_0/c$. For sake of completeness, the proof is given in Appendix I.

Example 6. Using triplet-based robustness (Equation 3), we can show the robustness of two popular triplet-based metric learning approaches [4, 8] for which no generalization guarantees were known (to the best of our knowledge). These algorithms have the following form:

$$\min_{\mathbf{M} \succeq 0} c\|\mathbf{M}\| + \frac{1}{|\text{trip}_s|} \sum_{(s_i, s_j, s_k) \in \text{trip}_s} [1 - (x_i - x_k)^T \mathbf{M} (x_i - x_k) + (x_i - x_j)^T \mathbf{M} (x_i - x_j)]_+,$$

where $\|\mathbf{M}\| = \|\mathbf{M}\|_{\mathcal{F}}$ in [4] or $\|\mathbf{M}\| = \|\mathbf{M}\|_{1,2}$ in [8]. These methods are $(\mathcal{N}(\gamma, \mathcal{Z}, \|\cdot\|_2), \frac{16UR\gamma g_0}{c})$ -robust. The proof is given in Appendix J, the additional factor 2 comes from the use of triplets instead of pairs. This example illustrates that robustness results for triplet-based approaches can be derived from constraint-based methods.

6. Discussion

This section discusses the bounds derived from the proposed framework and put them into perspective with other approaches.

As seen in the previous section, our approach is rather general and allows one to derive generalization bounds for many metric learning methods. The counterpart of this generality is the relative looseness of the resulting bounds: although the $O(1/\sqrt{n})$ convergence rate is the same as in the alternative frameworks presented below, the covering number constants are difficult to estimate and can be large. Therefore, these bounds are useful to establish the consistency of a metric learning approach but do not provide sharp estimates of the generalization loss. This is in accordance with the original robustness bounds introduced in [26, 27].

The guarantees proposed in [17, 18] can be tighter but hold only under strong assumptions on the distribution of examples. Moreover, these results only apply to a specific metric learning formulation and it is not clear how they can be adapted to more general forms. Bounds based on uniform stability [19] are also tighter and can deal with various loss functions, but fail to address sparsity-inducing regularizers. This is known to be a general limitation of stability-based analysis [21].

More recently, independently and in parallel to our work, generalization bounds for metric learning based on Rademacher analysis have been proposed [25, 13]. These bounds are tighter than the ones obtained with robustness and can tackle some sparsity-inducing regularizers. Their derivation is however more involved as it requires to compute Rademacher average estimates related to the matrix dual norm. For this reason, their analysis is limited to matrix norm regularization, while our framework can essentially accommodate any regularizer that is bounded below by a matrix norm (following the same proof technique as in Section 5). Furthermore, robustness is flexible enough to tackle other settings (such as triplet-based constraints), as illustrated in Section 5.

We conclude this discussion by noting that the proposed framework can be used to obtain generalization bounds for linear classifiers that use the learned metrics, following the work of [12, 13].

7. Conclusion

We proposed a new theoretical framework for evaluating the generalization ability of metric learning based on the notion of algorithm robustness originally introduced in [27]. We showed that a weak notion of robustness characterizes the generalizability of metric learning algorithms, justifying that robustness is fundamental for such algorithms. The proposed framework has an intuitive geometric meaning and allows us to derive generalization bounds for a large class of algorithms with different regularizations (such as sparsity inducing norms), showing that it has a wider applicability than existing frameworks. Moreover, few algorithm-specific arguments are needed. The price to pay is the relative looseness of the resulting bounds.

A perspective of this work is to take advantage of the generality and flexibility of the robustness framework to tackle more complex metric learning settings, for instance other regularizers (such as the LogDet divergence used in [5, 6]), methods that learn multiple metrics (e.g., [31, 32]), and metric learning for domain adaptation [33, 34]. It is also promising to investigate whether robustness could be used to derive guarantees for online algorithms such as [15, 6, 10].

Another exciting direction for future work is to investigate new metric learning algorithms based on the robustness property. For instance, given a partition of the labeled input space and for any two regions, such an algorithm could minimize the maximum loss over pairs of examples belonging to each region. This is reminiscent of concepts from robust optimization [35] and could be useful to deal with noisy settings.

Appendix A. Proof of Theorem 2 (pseudo-robustness)

Proof From the proof of Theorem 1, we can easily deduce that:

$$|\mathcal{L}(\mathcal{A}_{p_s}) - l_{emp}(\mathcal{A}_{p_s})| \leq 2B \sum_{i=1}^K \left| \frac{|N_i|}{n} - \mu(C_i) \right| + \left| \sum_{i,j=1}^K \mathbb{E}_{z_1, z_2 \sim \mu} (l(\mathcal{A}_{p_s}, z_1, z_2) | z_1 \in C_i, z_2 \in C_j) \frac{|N_i||N_j|}{n} - \frac{1}{n^2} \sum_{i,j=1}^n l(\mathcal{A}_{p_s}, s_i, s_j) \right|.$$

Then, we have

$$\begin{aligned}
&\leq 2B \sum_{i=1}^K \left| \frac{|N_i|}{n} - \mu(C_i) \right| + \\
&\quad \left| \frac{1}{n^2} \sum_{i,j=1}^K \sum_{(s_o, s_l) \in \hat{p}(\mathbf{s})} \sum_{s_o \in N_i} \sum_{s_l \in N_j} \max_{z \in C_i} \max_{z' \in C_j} |l(\mathcal{A}_{p_{\mathbf{s}}}, z, z') - l(\mathcal{A}_{p_{\mathbf{s}}}, s_o, s_l)| \right| + \\
&\quad \left| \frac{1}{n^2} \sum_{i,j=1}^K \sum_{(s_o, s_l) \notin \hat{p}(\mathbf{s})} \sum_{s_o \in N_i} \sum_{s_l \in N_j} \max_{z \in C_i} \max_{z' \in C_j} |l(\mathcal{A}_{p_{\mathbf{s}}}, z, z') - l(\mathcal{A}_{p_{\mathbf{s}}}, s_o, s_l)| \right| \\
&\leq \frac{\hat{p}_n(p_{\mathbf{s}})}{n^2} \epsilon(p_{\mathbf{s}}) + B \left(\frac{n^2 - \hat{p}_n(p_{\mathbf{s}})}{n^2} + 2\sqrt{\frac{2K \ln 2 + 2 \ln 1/\delta}{n}} \right).
\end{aligned}$$

The second inequality is obtained by the triangle inequality, the last one is obtained by the application of Proposition 1, the hypothesis of pseudo-robustness and the fact that l is positive and bounded by B , thus we have $|l(\mathcal{A}_{p_{\mathbf{s}}}, z, z') - l(\mathcal{A}_{p_{\mathbf{s}}}, s_o, s_l)| \leq B$. \square

Appendix B. Proof of sufficiency of Theorem 3

Proof The proof of sufficiency closely follows the first part of the proof of Theorem 8 in [27]. When \mathcal{A} is weakly robust, there exists a sequence $\{D_n\}$ such that for any $\delta, \epsilon > 0$ there exists $N(\delta, \epsilon)$ such that for all $n > N(\delta, \epsilon)$, $\Pr(\mathbf{t}(n) \in D_n) > 1 - \delta$ and

$$\max_{\hat{\mathbf{s}}(n) \in D_n} \left| L(\mathcal{A}_{p_{\hat{\mathbf{s}}(n)}}, p_{\hat{\mathbf{s}}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)}) \right| < \epsilon. \quad (\text{B.1})$$

Therefore for any $n > N(\delta, \epsilon)$,

$$\begin{aligned}
&|\mathcal{L}(\mathcal{A}_{p_{\hat{\mathbf{s}}(n)}}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \\
&= |\mathbb{E}_{\mathbf{t}(n)}(L(\mathcal{A}_{p_{\hat{\mathbf{s}}(n)}}, p_{\mathbf{t}(n)})) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \\
&= |\Pr(\mathbf{t}(n) \notin D_n) \mathbb{E}(L(\mathcal{A}_{p_{\hat{\mathbf{s}}(n)}}, p_{\mathbf{t}(n)}) | \mathbf{t}(n) \notin D_n) \\
&\quad + \Pr(\mathbf{t}(n) \in D_n) \mathbb{E}(L(\mathcal{A}_{p_{\hat{\mathbf{s}}(n)}}, p_{\mathbf{t}(n)}) | \mathbf{t}(n) \in D_n) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \\
&\leq \Pr(\mathbf{t}(n) \notin D_n) |\mathbb{E}(L(\mathcal{A}_{p_{\hat{\mathbf{s}}(n)}}, p_{\mathbf{t}(n)}) | \mathbf{t}(n) \notin D_n) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \\
&\quad + \Pr(\mathbf{t}(n) \in D_n) |\mathbb{E}(L(\mathcal{A}_{p_{\hat{\mathbf{s}}(n)}}, p_{\mathbf{t}(n)}) | \mathbf{t}(n) \in D_n) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \\
&\leq \delta B + \max_{\hat{\mathbf{s}}(n) \in D_n} |L(\mathcal{A}_{p_{\hat{\mathbf{s}}(n)}}, p_{\hat{\mathbf{s}}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \\
&\leq \delta B + \epsilon.
\end{aligned}$$

The first inequality holds because the testing samples $\mathbf{t}(n)$ consist of n instances IID from μ . The second equality is obtained by conditional expectation. The next inequality uses the positiveness and the upper bound B of the loss function. Finally, we apply Equation B.1. We thus conclude that \mathcal{A} generalizes for $p_{\mathbf{s}^*}$ because ϵ and δ can be chosen arbitrary. \square

Appendix C. Proof of Lemma 1

Proof This proof follows the same principle as the proof of Lemma 2 from [27]. By contradiction, assume ϵ^* and δ^* do not exist. Let $\epsilon_v = \delta_v = 1/v$ for $v = 1, 2, \dots$, then there exists a non decreasing sequence $\{N(v)\}_{v=1}^{\infty}$ such that for all v , if $n \geq N(v)$ then $\Pr(|L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{t}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \geq \epsilon_v) < \delta_v$. For each n we define

$$D_n^v \triangleq \{\hat{\mathbf{s}}(n) | L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\hat{\mathbf{s}}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| < \epsilon_v\}.$$

For each $n \geq N(v)$ we have

$$\Pr(\mathbf{t}(n) \in D_n^v) = 1 - \Pr(|L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{t}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \geq \epsilon_v) > 1 - \delta_v.$$

For $n \geq N(1)$, define $D_n \triangleq D_n^{v(n)}$, where $v(n) = \max\{v | N(v) \leq n; v \leq n\}$. Thus for all, $n \geq N(1)$ we have $\Pr(\mathbf{t}(n) \in D_n) > 1 - \delta_{v(n)}$ and

$$\sup_{\hat{\mathbf{s}}(n) \in D_n} |L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\hat{\mathbf{s}}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| < \epsilon_{v(n)}.$$

Note that $v(n)$ tends to infinity, it follows that $\delta_{v(n)} \rightarrow 0$ and $\epsilon_{v(n)} \rightarrow 0$. Therefore, $\Pr(\mathbf{t}(n) \in D_n) \rightarrow 1$ and

$$\lim_{n \rightarrow \infty} \left\{ \sup_{\hat{\mathbf{s}}(n) \in D_n} |L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\hat{\mathbf{s}}(n)}) - L(\mathcal{A}_{p_{\mathbf{s}^*(n)}}, p_{\mathbf{s}^*(n)})| \right\} = 0.$$

That is \mathcal{A} is weakly robust. w.r.t. $p_{\mathbf{s}}$ which is a desired contradiction. \square

Appendix D. Mc Diarmid inequality

Let X_1, \dots, X_n be n independent random variables taking values in X and let $Z = f(X_1, \dots, X_n)$. If for each $1 \leq i \leq n$, there exists a constant c_i such that

$$\sup_{x_1, \dots, x_n, x'_i \in \mathcal{X}} |f(x_1, \dots, x_i, \dots, x_n) - f(x_1, \dots, x'_i, \dots, x_n)| \leq c_i, \forall 1 \leq i \leq n,$$

$$\text{then for any } \epsilon > 0, \quad \Pr[|Z - \mathbb{E}[Z]| \geq \epsilon] \leq 2 \exp\left(\frac{-2\epsilon^2}{\sum_{i=1}^n c_i^2}\right).$$

Appendix E. Robustness Theorem for Triplet-based Approaches

We give here an adaptation of Theorem 4 for triplet-based approaches. The proof follows the same principle as the one of Theorem 4.

Theorem 5. *Fix $\gamma > 0$ and a metric ρ of \mathcal{Z} . Suppose \mathcal{A} satisfies:*

$$\forall z_1, z_2, z_3, z'_1, z'_2, z'_3 : z_1, z_2, z_3 \in \mathbf{s}, \rho(z_1, z'_1) \leq \gamma, \rho(z_2, z'_2) \leq \gamma, \rho(z_3, z'_3) \leq \gamma,$$

$$|l(\mathcal{A}_{trip_{\mathbf{s}}}, z_1, z_2, z_3) - l(\mathcal{A}_{trip_{\mathbf{s}}}, z'_1, z'_2, z'_3)| \leq \epsilon(trip_{\mathbf{s}})$$

and $\mathcal{N}(\gamma/2, \mathcal{Z}, \rho) < \infty$. Then \mathcal{A} is $(\mathcal{N}(\gamma/2, \mathcal{Z}, \rho), \epsilon(trip_{\mathbf{s}}))$ -robust.

Appendix F. Proof of Example 2 (ℓ_1 norm)

Proof Let \mathbf{M}^* be the solution given training data $p_{\mathbf{s}}$. Due to optimality of \mathbf{M}^* , we have $\|\mathbf{M}^*\|_1 \leq g_0/c$. We can partition \mathcal{Z} as $|Y|\mathcal{N}(\gamma/2, X, \|\cdot\|_1)$ sets, such that if z and z' belong to the same set, then $y = y'$ and $\|x - x'\|_1 \leq \gamma$. Now, for $z_1, z_2, z'_1, z'_2 \in \mathcal{Z}$, if $y_1 = y'_1$, $\|x_1 - x'_1\|_1 \leq \gamma$, $y_2 = y'_2$ and $\|x_2 - x'_2\|_1 \leq \gamma$, then:

$$\begin{aligned} & |g(y_{12}[1 - f(\mathbf{M}^*, x_1, x_2)]) - g(y'_{12}[1 - f(\mathbf{M}^*, x'_1, x'_2)])| \\ & \leq U(|(x_1 - x_2)^T \mathbf{M}^*(x_1 - x'_1)| + |(x_1 - x_2)^T \mathbf{M}^*(x'_2 - x_2)| \\ & \quad + |(x_1 - x'_1)^T \mathbf{M}^*(x'_1 + x'_2)| + |(x'_2 - x_2)^T \mathbf{M}^*(x'_1 + x'_2)|) \\ & \leq U(\|x_1 - x_2\|_{\infty} \|\mathbf{M}^*\|_1 \|x_1 - x'_1\|_1 + \|x_1 - x_2\|_{\infty} \|\mathbf{M}^*\|_1 \|x'_2 - x_2\|_1 \\ & \quad + \|x_1 - x'_1\|_1 \|\mathbf{M}^*\|_1 \|x'_1 - x'_2\|_{\infty} + \|x'_2 - x_2\|_1 \|\mathbf{M}^*\|_1 \|x'_1 - x'_2\|_{\infty}) \\ & \leq \frac{8UR\gamma g_0}{c}. \end{aligned}$$

□

Appendix G. Proof of Example 3 ($\ell_{2,1}$ norm and trace norm)

Proof Let \mathbf{M}^* be the solution given training data $p_{\mathbf{s}}$. Due to optimality of \mathbf{M}^* , we have $\|\mathbf{M}^*\|_{2,1} \leq g_0/c$. We can partition \mathcal{Z} in the same way as in the proof of Example 1 and use the inequality $\|\mathbf{M}^*\|_{\mathcal{F}} \leq \|\mathbf{M}^*\|_{2,1}$ (from

Theorem 3 of [36, 37]) to derive the same bound:

$$\begin{aligned}
& |g(y_{12}[1 - f(\mathbf{M}^*, x_1, x_2)]) - g(y'_{12}[1 - f(\mathbf{M}^*, x'_1, x'_2)])| \\
& \leq U(\|x_1 - x_2\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x_1 - x'_1\|_2 + \|x_1 - x_2\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x'_2 - x_2\|_2 \\
& \quad + \|x_1 - x'_1\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x'_1 - x'_2\|_2 + \|x'_2 - x_2\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x'_1 - x'_2\|_2) \\
& \leq U(\|x_1 - x_2\|_2 \|\mathbf{M}^*\|_{2,1} \|x_1 - x'_1\|_2 + \|x_1 - x_2\|_2 \|\mathbf{M}^*\|_{2,1} \|x'_2 - x_2\|_2 \\
& \quad + \|x_1 - x'_1\|_2 \|\mathbf{M}^*\|_{2,1} \|x'_1 - x'_2\|_2 + \|x'_2 - x_2\|_2 \|\mathbf{M}^*\|_{2,1} \|x'_1 - x'_2\|_2) \\
& \leq \frac{8UR\gamma g_0}{c}.
\end{aligned}$$

For the trace norm, we use the classic result $\|\mathbf{M}^*\|_{\mathcal{F}} \leq \|\mathbf{M}\|_*$, which allows us to prove the same result by replacing $\|\cdot\|_{2,1}$ by $\|\cdot\|_*$ in the proof above.

□

Appendix H. Proof of Example 4 (Kernelization)

Proof We assume \mathbb{H} to be an Hilbert space with an inner product operator $\langle \cdot, \cdot \rangle$. The mapping $\phi(\cdot)$ is continuous from X to \mathbb{H} . The norm $\|\cdot\|_{\mathbb{H}} : \mathbb{H} \rightarrow \mathbb{R}$ is defined as $\|w\|_{\mathbb{H}} = \sqrt{\langle w, w \rangle}$ for all $w \in \mathbb{H}$, for matrices $\|\mathbf{M}\|_{\mathbb{H}}$ we take the entry wise norm by considering a matrix as a vector, corresponding to the Frobenius norm. The kernel function is defined as $k(x_1, x_2) = \langle \phi(x_1), \phi(x_2) \rangle$.

B_γ and $f_{\mathbb{H}}(\gamma)$ are finite by the compactness of X and continuity of $k(\cdot, \cdot)$. Let \mathbf{M}^* be the solution given training data p_s , by the optimality of \mathbf{M}^* and using the same trick as the other examples we have: $\|\mathbf{M}^*\|_{\mathbb{H}} \leq g_0/c$. Then, by considering a partition of \mathcal{Z} into $|Y|\mathcal{N}(\gamma/2, X, \|\cdot\|_2)$ disjoint subsets such that if (x_1, y_1) and (x_2, y_2) belong to the same set then $y_1 = y_2$ and $\|x_1 - x_2\|_2 \leq \gamma$. We have then,

$$\begin{aligned}
& |g(y_{ij}[1 - f(\mathbf{M}^*, \phi(x_1), \phi(x_2))]) - g(y_{ij}[1 - f(\mathbf{M}^*, \phi(x'_1), \phi(x'_2))])| \\
& \leq U(|(\phi(x_1) - \phi(x_2))^T \mathbf{M}^* (\phi(x_1) - \phi(x'_1))| + \\
& \quad |(\phi(x_1) - \phi(x_2))^T \mathbf{M}^* (\phi(x'_2) - \phi(x_2))| + \\
& \quad |(\phi(x_1) - \phi(x'_1))^T \mathbf{M}^* (\phi(x'_1) + \phi(x'_2))| + \\
& \quad |(\phi(x'_2) - \phi(x_2))^T \mathbf{M}^* (\phi(x'_1) + \phi(x'_2))|) \\
& \leq U(|\phi(x_1)^T \mathbf{M}^* (\phi(x_1) - \phi(x'_1))| + |\phi(x_2)^T \mathbf{M}^* (\phi(x_1) - \phi(x'_1))| + \\
& \quad |\phi(x_1)^T \mathbf{M}^* (\phi(x'_2) \phi(x_2))| + |\phi(x_2)^T \mathbf{M}^* (\phi(x'_2) - \phi(x_2))| + \\
& \quad |(\phi(x_1) - \phi(x'_1))^T \mathbf{M}^* \phi(x'_1)| + |(\phi(x_1) - \phi(x'_1))^T \mathbf{M}^* \phi(x'_2)| + \\
& \quad |(\phi(x'_2) - \phi(x_2))^T \mathbf{M}^* \phi(x'_1)| + |(\phi(x'_2) - \phi(x_2))^T \mathbf{M}^* \phi(x'_2)|). \quad (\text{H.1})
\end{aligned}$$

Then, note that

$$\begin{aligned}
& |\phi(x_1)^T \mathbf{M}^* (\phi(x_1) - \phi(x'_1))| \\
& \leq \sqrt{\langle \phi(x_1), \phi(x_1) \rangle} \|\mathbf{M}^*\|_{\mathbb{H}} \sqrt{\langle \phi(x'_1) - \phi(x'_2), \phi(x'_1) - \phi(x'_2) \rangle} \\
& \leq B_\gamma \frac{g_o}{c} \sqrt{f_{\mathbb{H}}(\gamma)}.
\end{aligned}$$

Thus, by applying the same principle to all the terms in the right part of inequality (H.1), we obtain:

$$|g(y_{ij}[1-f(\mathbf{M}^*, \phi(x_1), \phi(x_2))]) - g(y_{ij}[1-f(\mathbf{M}^*, \phi(x'_1), \phi(x'_2))])| \leq \frac{8UB_\gamma \sqrt{f_{\mathbb{H}}(\gamma)} g_o}{c}.$$

Appendix I. Proof of Example 5

Let \mathbf{M}^* be the solution given training data p_s , by the optimality of \mathbf{M}^* , we get $\|\mathbf{M}^*\| \leq g_o/c$ and we consider the same partition of \mathcal{Z} as in the proof of Example 1. We can then obtain easily:

$$\begin{aligned}
& |g(y_{12}[1 - f(\mathbf{M}^*, x_1, x_2)]) - g(y'_{12}[1 - f(\mathbf{M}^*, x'_1, x'_2)])| \\
& \leq U|x'_1 \mathbf{M}^* x'_2 - x_1 \mathbf{M}^* x_2| \\
& \leq U|x'_1 \mathbf{M}^* x'_2 - x_1 \mathbf{M}^* x'_2| + U|x_1 \mathbf{M}^* x'_2 - x_1 \mathbf{M}^* x_2| \\
& \leq U(\|x'_1 - x_1\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x'_2\|_2 + \|x_1\|_2 \|\mathbf{M}^*\|_{\mathcal{F}} \|x'_2 - x_2\|_2) \leq \frac{2UR\gamma g_o}{c}.
\end{aligned}$$

The proof is given for the Frobenius norm but can be easily adapted to the use of ℓ_1 and $\ell_{2,1}$ norms using similar arguments as in the proofs of Appendix F and Appendix G.

Appendix J. Proof of Example 6

We consider the following loss

$$\begin{aligned}
& g([1 - (x_i - x_k)^T \mathbf{M} (x_i - x_k) + (x_i - x_j)^T \mathbf{M} (x_i - x_j)]) \\
& = [1 - (x_i - x_k)^T \mathbf{M} (x_i - x_k) + (x_i - x_j)^T \mathbf{M} (x_i - x_j)]_+.
\end{aligned}$$

Using the Let \mathbf{M}^* be the solution given training data $trip_s$ formed of triplets. As usual, due to optimality of \mathbf{M}^* , using the same tricks as above, we get $\|\mathbf{M}^*\| \leq g_o/c$. Then, by considering a partition of \mathcal{Z} into $|Y|\mathcal{N}(\gamma/2, X, \|\cdot\|_2)$, three partitions C_1, C_2, C_3 and $z_1, z_2, z_3, z'_1, z'_2, z'_3 \in \mathcal{Z}$ such that $z_1, z'_1 \in C_1$,

$z_2, z'_2 \in C_2$ and $z_3, z'_3 \in C_3$ with $y_1 = y'_1 = y_2 = y'_2, y_3 = y'_3, y_3 \neq y_1$, and $\|x_1 - x'_1\|_1 \leq \gamma, \|x_2 - x'_2\|_1 \leq \gamma, \|x_3 - x'_3\|_1 \leq \gamma$, we have:

$$\begin{aligned}
& |g([1 - (x_1 - x_3)^T \mathbf{M}^*(x_1 - x_3) + (x_1 - x_2)^T \mathbf{M}^*(x_1 - x_2)]) - \\
& \quad g([1 - (x'_1 - x'_3)^T \mathbf{M}^*(x'_1 - x'_3) + (x'_1 - x'_2)^T \mathbf{M}^*(x'_1 - x'_2)])| \\
& \leq U|(x'_1 - x'_3)^T \mathbf{M}^*(x'_1 - x'_3) - (x_1 - x_3)^T \mathbf{M}^*(x_1 - x_3) + \\
& \quad (x_1 - x_2)^T \mathbf{M}^*(x_1 - x_2) - (x'_1 - x'_2)^T \mathbf{M}^*(x'_1 - x'_2)| \\
& \leq U|(x'_1 - x'_3)^T \mathbf{M}^*(x'_1 - x'_3) - (x_1 - x_3)^T \mathbf{M}^*(x_1 - x_3)| + \\
& \quad U|(x_1 - x_2)^T \mathbf{M}^*(x_1 - x_2) - (x'_1 - x'_2)^T \mathbf{M}^*(x'_1 - x'_2)| \\
& \leq \frac{8UR\gamma g_0}{c} + \frac{8UR\gamma g_0}{c} = \frac{16UR\gamma g_0}{c}.
\end{aligned}$$

The first inequality is due to the U -lipschitz property of g , the second comes from the triangle inequality and the last one follows the same construction as in the proof of Example 1. Then, by Theorem 5, the example holds. \square

References

- [1] A. Bellet, A. Habrard, M. Sebban, A Survey on Metric Learning for Feature Vectors and Structured Data, Technical Report, arXiv:1306.6709, 2013.
- [2] B. Kulis, Metric Learning: A Survey, Foundations and Trends in Machine Learning (FTML) 5 (2012) 287–364.
- [3] E. P. Xing, A. Y. Ng, M. I. Jordan, S. J. Russell, Distance Metric Learning with Application to Clustering with Side-Information, in: Advances in Neural Information Processing Systems 15, 2002, pp. 505–512.
- [4] M. Schultz, T. Joachims, Learning a Distance Metric from Relative Comparisons, in: Advances in Neural Information Processing Systems (NIPS) 16, 2003.
- [5] J. V. Davis, B. Kulis, P. Jain, S. Sra, I. S. Dhillon, Information-theoretic metric learning, in: Proceedings of the 24th International Conference on Machine Learning (ICML), 2007, pp. 209–216.
- [6] P. Jain, B. Kulis, I. S. Dhillon, K. Grauman, Online Metric Learning and Fast Similarity Search, in: Advances in Neural Information Processing Systems (NIPS) 21, 2008, pp. 761–768.

- [7] K. Q. Weinberger, L. K. Saul, Distance Metric Learning for Large Margin Nearest Neighbor Classification, *Journal of Machine Learning Research (JMLR)* 10 (2009) 207–244.
- [8] Y. Ying, K. Huang, C. Campbell, Sparse Metric Learning via Smooth Optimization, in: *Advances in Neural Information Processing Systems (NIPS)* 22, 2009, pp. 2214–2222.
- [9] B. McFee, G. R. G. Lanckriet, Metric Learning to Rank, in: *Proceedings of the 27th International Conference on Machine Learning*, 2010, pp. 775–782.
- [10] G. Chechik, U. Shalit, V. Sharma, S. Bengio, An Online Algorithm for Large Scale Image Similarity Learning, in: *Advances in Neural Information Processing Systems (NIPS)* 22, 2009, pp. 306–314.
- [11] A. M. Qamar, Generalized Cosine and Similarity Metrics: A supervised learning approach based on nearest-neighbors, Ph.D. thesis, University of Grenoble, 2010.
- [12] A. Bellet, A. Habrard, M. Sebban, Similarity Learning for Provably Accurate Sparse Linear Classification, in: *Proceedings of the 29th International Conference on Machine Learning (ICML)*, 2012, pp. 1871–1878.
- [13] Z.-C. Guo, Y. Ying, Guaranteed Classification via Regularized Similarity Learning, *Neural Computation* 26 (2014) 497–522.
- [14] P. Kar, P. Jain, Similarity-based Learning via Data Driven Embeddings, in: *Advances in Neural Information Processing Systems (NIPS)* 24, 2011, pp. 1998–2006.
- [15] S. Shalev-Shwartz, Y. Singer, A. Y. Ng, Online and batch learning of pseudo-metrics, in: *Proceedings of the 21st International Conference on Machine Learning (ICML)*, 2004.
- [16] Y. Wang, R. Khardon, D. Pechyony, R. Jones, Generalization Bounds for Online Learning Algorithms with Pairwise Loss Functions, in: *Proceedings of the 25th Annual Conference on Learning Theory (COLT)*, 2012, pp. 13.1–13.22.

- [17] W. Bian, D. Tao, Learning a Distance Metric by Empirical Loss Minimization, in: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI), 2011, pp. 1186–1191.
- [18] W. Bian, D. Tao, Constrained Empirical Risk Minimization Framework for Distance Metric Learning, IEEE Transactions on Neural Networks and Learning Systems (TNNLS) 23 (2012) 1194–1205.
- [19] R. Jin, S. Wang, Y. Zhou, Regularized Distance Metric Learning: Theory and Algorithm, in: Advances in Neural Information Processing Systems (NIPS) 22, 2009, pp. 862–870.
- [20] O. Bousquet, A. Elisseeff, Stability and Generalization, Journal of Machine Learning Research (JMLR) 2 (2002) 499–526.
- [21] H. Xu, C. Caramanis, S. Mannor, Sparse Algorithms Are Not Stable: A No-Free-Lunch Theorem, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 34 (2012) 187–193.
- [22] R. Rosales, G. Fung, Learning Sparse Metrics via Linear Programming, in: Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2006, pp. 367–373.
- [23] G.-J. Qi, J. Tang, Z.-J. Zha, T.-S. Chua, H.-J. Zhang, An Efficient Sparse Metric Learning in High-Dimensional Space via l_1 -Penalized Log-Determinant Regularization, in: Proceedings of the 26th International Conference on Machine Learning (ICML), 2009.
- [24] A. Bellet, A. Habrard, Robustness and Generalization for Metric Learning, Technical Report, arXiv:1209.1086, 2012.
- [25] Q. Cao, Z.-C. Guo, Y. Ying, Generalization Bounds for Metric and Similarity Learning, Technical Report, University of Exeter, 2012. ArXiv:1207.5437.
- [26] H. Xu, S. Mannor, Robustness and Generalization, in: Proceedings of the 23rd Annual Conference on Learning Theory (COLT), 2010, pp. 503–515.
- [27] H. Xu, S. Mannor, Robustness and Generalization, Machine Learning 86 (2012) 391–423.

- [28] A. N. Kolmogorov, V. M. Tikhomirov, ϵ -entropy and ϵ -capacity of sets in functional spaces, *American Mathematical Society Translations* 2 (1961) 277–364.
- [29] A. W. van der Vaart, J. A. Wellner, *Weak convergence and empirical processes*, Springer, 2000.
- [30] G. Kunapuli, J. Shavlik, *Mirror Descent for Metric Learning: A Unified Approach*, in: *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Database (ECML/PKDD)*, 2012, pp. 859–874.
- [31] J. Wang, A. Woznica, A. Kalousis, *Parametric Local Metric Learning for Nearest Neighbor Classification*, in: *Advances in Neural Information Processing Systems (NIPS)* 25, 2012, pp. 1610–1618.
- [32] Y. Shi, A. Bellet, F. Sha, *Sparse Compositional Metric Learning*, in: *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, 2014.
- [33] B. Kulis, K. Saenko, T. Darrell, *What you saw is not what you get: Domain adaptation using asymmetric kernel transforms*, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 1785–1792.
- [34] B. Geng, D. Tao, C. Xu, *DAML: Domain Adaptation Metric Learning*, *IEEE Transactions on Image Processing (TIP)* 20 (2011) 2980–2989.
- [35] A. Ben-Tal, L. E. Ghaoui, A. Nemirovski, *Robust Optimization*, Princeton University Press, 2009.
- [36] B. Q. Feng, *Equivalence constants for certain matrix norms*, *Linear Algebra and Its Applications* 374 (2003) 247–253.
- [37] A.-L. Klaus, C.-K. Li, *Isometries for the vector (p,q) norm and the induced (p,q) norm*, *Linear & Multilinear Algebra* 38 (1995) 315–332.



Aurélien Bellet received his PhD in Machine Learning from the University of Saint-Etienne in 2012. His research focused on algorithmic and theoretical aspects of metric and similarity learning. He is currently a postdoctoral researcher at the University of Southern California, where he is involved in large-scale and distributed machine learning with applications to automatic speech recognition.



Amaury Habrard received a PhD in Machine Learning in 2004 from the University of Saint-Etienne. He was assistant professor at the Laboratoire d'Informatique Fondamentale of Aix-Marseille University until 2011, where he received an habilitation thesis in 2010. He is currently Professor in the Machine Learning group at the Laboratoire Hubert Curien of the University of Saint-Etienne. His research interests include metric learning, transfer learning, learning theory.