

Relative Comparison Kernel Learning with Auxiliary Kernels

Eric Heim¹, Hamed Valizadegan², and Milos Hauskrecht¹

¹ University of Pittsburgh, Department of Computer Science
{eth13, milos}@cs.pitt.edu

² NASA Ames Research Center
hamed.valizadegan@nasa.gov

Abstract. In this work we consider the problem of learning a positive semidefinite kernel matrix from relative comparisons of the form: “object A is more similar to object B than it is to C ”, where comparisons are given by humans. Existing solutions to this problem assume many comparisons are provided to learn a high quality kernel. However, this can be considered unrealistic for many real-world tasks since a large amount of human input is often costly or difficult to obtain. Because of this, only a limited number of these comparisons may be provided. We propose a new kernel learning approach that supplements the few relative comparisons with “auxiliary” kernels built from more easily extractable features in order to learn a kernel that more completely models the notion of similarity gained from human feedback. Our proposed formulation is a convex optimization problem that adds only minor overhead to methods that use no auxiliary information. Empirical results show that in the presence of few training relative comparisons, our method can learn kernels that generalize to more out-of-sample comparisons than methods that do not utilize auxiliary information, as well as similar methods that learn metrics over objects.

1 Introduction

The effectiveness of many kernel methods for unsupervised [24,6], semi-supervised [30,29,27], and supervised [18] learning is highly dependent how meaningful the input kernel is for modeling similarity among objects for a given task. In practice, kernels are often built by using a standard kernel function on features extracted from data. For example, when building a kernel over clothing items, features can be extracted for each item regarding attributes like size, style, and color. Then, a predefined kernel function (e.g. the Gaussian kernel function) can be applied to this feature representation to build a kernel over clothing. However, for certain tasks, objects may not be represented well by extracted features alone. Consider a product recommendation system for suggesting replacements for out of stock clothing items. Such a system requires a model of similarity based on how humans perceive clothing, which may not be captured entirely by features. For this, it is likely that human input is necessary to construct a meaningful kernel.

In general, obtaining reliable information from humans can be challenging, but retrieving *relative comparisons* of the form “object A is more similar to object B than it is to object C ” has several attractive characteristics. First, relative comparison questions are less mentally fatiguing to humans compared to other forms (e.g. quantitative comparison questions: “On a scale from 1 to 10, how similar are objects A and B ?”) [10]. Second, there is no need to reconcile individual humans’ personal scales of similarity. Finally, relative comparison feedback can be drawn implicitly through certain human-computer interactions, such as mouse clicks. In this work, we consider the specific problem of learning a kernel from relative comparisons; A problem we will refer to as *relative comparison kernel learning* (RCKL).

Current RCKL methods [1,22,14] assume that all necessary human feedback to build a useful kernel is provided. This is often not the case in real-world scenarios. A large amount of feedback is needed to build a kernel that represents a meaningful notion of how a human views the relationships among objects, and obtaining feedback from humans is often difficult or costly. Hence, it is a realistic assumption that only a limited amount of feedback can be obtained.

In order to learn a meaningful kernel from only a limited number of relative comparisons we propose a novel RCKL method that learns a combination of *auxiliary kernels* built from extracted features and a *learned kernel*. The intuition behind this approach is that while human feedback is necessary to construct an appropriate kernel, some aspects of how humans view similarity among objects are likely captured in easily extractable features. If “auxiliary” kernels are built from these features, then they can be used to reduce the need of many relative comparisons. To learn the aforementioned combination, we formulate a convex optimization that adds a only small amount of computational overhead to traditional RCKL methods. Experimentally, we show that by using auxiliary information, our method can learn a kernel that accurately models the relationship among objects from few relative comparisons, including relationships among objects not explicitly given. More specifically, our method is shown to generalize to more held out relative comparisons than traditional RCKL methods. In addition, we compare our method to similar, state-of-the-art methods in metric learning, and, again, we show that our method generalizes to more held out relative comparisons.

The remainder of the paper is organized as follows. Section 2 provides our formal definition of RCKL. Section 3 motivates our problem. Section 4 introduces a general framework for extending RCKL methods to use auxiliary information. Section 5 overviews related work. Section 6 presents an evaluation of our method. Section 7 concludes with future work.

2 Preliminaries

The RCKL problem considered in this work is defined by a set of n objects, $\mathcal{X} = \{x_1, \dots, x_n\} \subseteq \mathbb{X}$, where \mathbb{X} is the set of all possible objects. Similarity information among objects is given in the form of a set \mathcal{T} of triplets:

$$\mathcal{T} = \{(a, b, c) | x_a \text{ is more similar to } x_b \text{ than } x_c\} \quad (1)$$

The goal is to find a positive semidefinite (PSD) kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$ that satisfies the following constraints:

$$\begin{aligned} \forall_{(a,b,c) \in \mathcal{T}} : d_{\mathbf{K}}(x_a, x_b) < d_{\mathbf{K}}(x_a, x_c) \\ \text{Where } d_{\mathbf{K}}(x_a, x_b) = \mathbf{K}_{aa} + \mathbf{K}_{bb} - 2\mathbf{K}_{ab} \end{aligned} \quad (2)$$

Here, \mathbf{K}_{ab} is the element in the a th row and b th column of \mathbf{K} , representing the similarity between the a th and b th objects. The elements of \mathbf{K} can be interpreted as the inner products of the objects embedded in a Reproducing Kernel Hilbert Space (RKHS), $\mathcal{H}_{\mathbf{K}}$, endowed with a mapping $\Phi_{\mathbf{K}} : \mathbb{X} \rightarrow \mathcal{H}_{\mathbf{K}}$. With this interpretation $\mathbf{K}_{aa} + \mathbf{K}_{bb} - 2\mathbf{K}_{ab} = \|\Phi_{\mathbf{K}}(x_a) - \Phi_{\mathbf{K}}(x_b)\|_2^2$. Thus, learning a kernel matrix \mathbf{K} that satisfies the constraints in (2) is equivalent to embedding the objects in a space, such that for all triplets $(a, b, c) \in \mathcal{T}$, x_a is closer to x_b than it is to x_c without explicitly learning the mapping $\Phi_{\mathbf{K}}$. We say that a triplet (a, b, c) is *satisfied* if the corresponding constraint in (2) is satisfied.

One interpretation of (2) is that triplets define a binary relation over distances among objects. For example, if $\mathcal{X} = \{x_1, x_2, x_3\}$ and $\mathcal{T} = \{(1, 2, 3), (2, 1, 3)\}$, then \mathcal{T} defines a relation, $R_{\mathcal{T}}$ over $\mathcal{S}_{\mathcal{X}} = \{(d_{\mathbf{K}}(x_1, x_2), d_{\mathbf{K}}(x_1, x_3), d_{\mathbf{K}}(x_2, x_3))\}$, such that $R_{\mathcal{T}} = \{(d_{\mathbf{K}}(x_1, x_2), d_{\mathbf{K}}(x_1, x_3)), (d_{\mathbf{K}}(x_1, x_2), d_{\mathbf{K}}(x_2, x_3))\}$. With this in mind, we continue onto the next section where we discuss the RCKL problem in more depth.

3 The Impact of Few Triplets

To help motivate this work, we provide some insight into why it can be assumed, in practice, that only a limited number of triplets can be obtained from humans, and the potential impact it has on learning an accurate kernel. First, we begin by defining some properties of sets of triplets:

Definition 1. Given a set of triplets \mathcal{T} , let \mathcal{T}^{∞} be the transitive closure of \mathcal{T}

Definition 2. Given a set of triplets \mathcal{T} , let $\mathcal{T}^{trans} = \mathcal{T}^{\infty} \setminus \mathcal{T}$

Definition 2 simply defines \mathcal{T}^{trans} as the set of triplets that can be inferred by transitivity of triplets in \mathcal{T} . For example, if $\mathcal{T} = \{(a, b, c), (c, a, b)\}$ then $\mathcal{T}^{trans} = \{(b, a, c)\}$.

Definition 3. A set \mathcal{T} of triplets is conflicting if $\exists a, b, c : (a, b, c) \in \mathcal{T}^{\infty} \wedge (a, c, b) \in \mathcal{T}^{\infty}$

A set of conflicting triplets given by a source can be seen as inconsistent or contradictory in terms of how the source is comparing objects. In practice, this can be handled by prompting the source of triplets to resolve this conflict or by using simplifying methods such as in [16]. We defer to these methods in terms

of how conflicts can be dealt with and consider the *non-conflicting* case. Let \mathcal{T}^{total} be the set of all non-conflicting triplets that would be given by a source, if prompted with every relative comparison question over n objects. We begin by stating the following:

Theorem 1. For n objects, $|\mathcal{T}^{total}| = \frac{1}{2}(n^3 - 3n^2 + 2n)$

Theorem 1 is proven in Sec. A.1 of the appendix. For even a small number of objects, obtaining most of \mathcal{T}^{total} from humans would be too difficult or costly in many practical scenarios, especially if feedback is gained through natural use of a system, such as an online store, or if feedback requires an expert’s opinion, such as in the medical domain. Let $\mathcal{T} \subseteq \mathcal{T}^{total}$ be the set of triplets actually *obtained* from a source. We say that a triplet t is *unobtained* if $t \in \mathcal{T}^{total} \setminus \mathcal{T}$. To build a model that accurately reflects the true notion of similarity given by a source of triplets, an RCKL method should learn a kernel \mathbf{K} that not only satisfies the obtained triplets, but also many of the triplets in \mathcal{T}^{total} , including those that were unobtained. This means that given small number of obtained triplets, an RCKL method should somehow *infer* a portion of the unobtained triplets in order to build an accurate model of similarity. In the remainder of this section we consider two possible scenarios where unobtained triplets could potentially be inferred.

For the following analysis, we assume that triplets are obtained one at a time. Also, we assume that the order in which triplets are obtained is random. This could be a reasonable assumption in applications, such as search engines, where the goal of asking relative comparison questions that are most useful in the learning process comes secondary to providing the best search results, and as such, no assumptions can be made regarding which relative comparison questions are posed to a source. Thus, the worst-case in the following analysis is with adversarial choice of both \mathcal{T}^{total} and the order in which triplets are obtained. Let \mathcal{T}_i be the set of triplets given by an adversary after i triplets are given. Under these assumptions, we state the following theorem:

Theorem 2. In the worst-case, $\forall_{i=1, \dots, |\mathcal{T}^{total}|} : \mathcal{T}_i^{trans} \setminus \mathcal{T}_i = \emptyset$

Theorem 2 is proven in Sec. A.2 of the appendix. This states that in the worst case, no unobtained triplet can be inferred by transitive relationship among obtained triplets. As a result, it may fall on the RCKL methods themselves to infer triplets. Many RCKL methods attempt to do this by assuming the learned kernel \mathbf{K} has low rank. By limiting the rank of \mathbf{K} to $r < n$, an RCKL method may effectively infer unobtained triplets by eliminating those that cannot be satisfied by a rank r kernel. For instance, assume an RCKL method attempts to learn a rank r kernel from \mathcal{T} , and assume the triplets (a, b, c) and (a, c, b) are not in \mathcal{T} . If the set $\mathcal{T} \cup (a, c, b)$ cannot be satisfied by a rank r kernel, but $\mathcal{T} \cup (a, b, c)$ can, then the RCKL method can only learn a kernel in which (a, b, c) is satisfied. Let \mathcal{T}^{rank-r} be the set of all unobtained, not otherwise inferred, triplets that are inferred when an RCKL method enforces $\text{rank}(\mathbf{K}) \leq r$. For adversarial choice of \mathcal{T}^{total} we can state the following theorem:

Theorem 3. *In the worst case, $\forall_{t \in \mathcal{T}_{rank-r}} : t \notin \mathcal{T}^{total}$*

Theorem 3 is proven in Sec. A.3 of the appendix. This theorem shows that it could be the case that any triplet inferred by limiting the rank of \mathbf{K} is not a triplet a source would give. If a large portion of \mathcal{T}^{total} cannot be obtained or correctly inferred, then much of the information needed for an RCKL method to learn a kernel that reflects how a source views the relationship among objects is simply not available. The goal of this work is to use auxiliary information describing the objects to supplement obtained triplets in order to learn a kernel that can satisfy more unobtained triplets than traditional methods. In the following section we propose a novel RCKL method that extends traditional RCKL methods to use auxiliary information.

4 Learning a Kernel with Auxiliary Information

In this section we introduce a generalized framework for traditional RCKL methods. Then, we expand upon this to create two new frameworks: One that combines auxiliary kernels to satisfy triplets, and another that is a hybrid of the previous two.

4.1 Traditional RCKL

Many RCKL methods can be generalized by the following optimization problem:

$$\begin{aligned} \min_{\mathbf{K}} E(\mathbf{K}, \mathcal{T}) + \lambda \text{trace}(\mathbf{K}) \\ \text{s.t. } \mathbf{K} \succeq 0, \end{aligned} \tag{3}$$

The first term, $E(\mathbf{K}, \mathcal{T})$, is a function of the error the objective incurs for \mathbf{K} not satisfying triplets in \mathcal{T} . The second term regularizes \mathbf{K} by its trace. Here, the trace is used as a convex approximation of the non-convex rank function. The rank of \mathbf{K} directly reflects the dimensionality of the embedding of the objects in $\mathcal{H}_{\mathbf{K}}$. A low setting of the hyperparameter λ favors a more accurate embedding, while a high value prefers a lower rank kernel. The PSD constraint ensures that \mathbf{K} is a valid kernel matrix, and makes (3) a semidefinite program (SDP) over n^2 variables. For the remainder of this paper we will refer to (3) as Traditional Relative Comparison Kernel Learning (RCKL-T).

4.2 RCKL via Conic Combination

In general, if there are few triplets in \mathcal{T} relative to n , there are many different RCKL solutions. Without using information regarding how the objects relate other than \mathcal{T} , RCKL-T methods may not be able to prefer solutions that generalize well to the many unobtained triplets. However, objects can often be described by features drawn from data. From these features, A auxiliary kernels $\mathbf{K}_1, \dots, \mathbf{K}_A \in \mathbb{R}^{n \times n}$ can be constructed using standard kernel functions to

model the relationship among objects. If one or more auxiliary kernels satisfy many triplets in \mathcal{T} , they may represent factors that influence how some of the unobtained triplets would have been answered. For instance, if a user considers characteristics such as color and size to be important when comparing clothing items, then kernels built from color and size may model a trend in how the user answers triplets over clothing items. If these kernels do represent a trend, then they could not only satisfy a portion of triplets in \mathcal{T} , but also a portion of the unobtained triplets. We wish to identify which of the given auxiliary kernels model trends in given triplets and combine them to satisfy triplets in \mathcal{T} . An approach popularized by multiple kernel learning methods is to combine kernels by a weighted sum:

$$\mathbf{K}' = \sum_{a=1}^A \mu_a \mathbf{K}_a \quad \boldsymbol{\mu} \in \mathbb{R}_+^A \quad (4)$$

\mathbf{K}' is a conic combination of PSD kernels, so itself is a PSD kernel [18]. \mathbf{K}' induces the mapping $\Phi_{\mathbf{K}'} : \mathbb{X} \rightarrow \mathbb{R}^D$ [7]:

$$\Phi_{\mathbf{K}'}(x_i) = [\sqrt{\mu_1} \Phi_1(x_i), \dots, \sqrt{\mu_A} \Phi_A(x_i)] \quad (5)$$

Here $\Phi_j : \mathbb{X} \rightarrow \mathbb{R}^{d_j}$ is a mapping from an object into the RKHS defined by \mathbf{K}_j , and $D = \sum_{a=1}^A d_a$. In short, (4) induces a mapping of the objects into a feature space defined as the weighted concatenation of the individual kernels' feature spaces. Consider, then, the following optimization:

$$\begin{aligned} \min_{\boldsymbol{\mu}} E(\mathbf{K}', \mathcal{T}) + \lambda \|\boldsymbol{\mu}\|_1 \\ \text{s.t. } \boldsymbol{\mu} \geq 0 \end{aligned} \quad (6)$$

By learning the weight vector $\boldsymbol{\mu}$, (6) scales the individual concatenated feature spaces to emphasize those that reflect \mathcal{T} well, and reduce the influence of those that do not. Because of its relationship to multiple kernel learning, we call this formulation RCKL-MKL.

Since the auxiliary kernels are fixed, regularizing them by their traces has no effect on their rank nor the rank of \mathbf{K}' . Instead, we choose to regularize $\boldsymbol{\mu}$ by its ℓ_1 -norm, a technique first made popular for its use in regression [23]. For a proper setting of λ , this has the effect of eliminating the contribution of kernels that do not help in reducing the error by forcing their corresponding weights to be exactly zero. Note that RCKL-MKL does not learn the elements of a kernel directly, and as a result is a linear program over A variables.

By limiting the optimization to only a conic combination of the predefined auxiliary kernels, RCKL-MKL does not necessarily produce a kernel that satisfies any triplets in \mathcal{T} . To capture the potential generalization power of using auxiliary information while retaining the ability to satisfy triplets in \mathcal{T} , we propose to learn a combination of the auxiliary kernels and \mathbf{K}_0 , a kernel similar to the one in RCKL-T whose elements are learned directly. By doing this, we force RCKL-T to prefer solutions more similar to the auxiliary kernels, which could satisfy

unobtained triplets. We call this hybrid approach Relative Comparison Kernel Learning with Auxiliary Kernels (RCKL-AK).

4.3 RCKL-AK

RCKL-AK learns the following kernel combination:

$$\mathbf{K}'' = \mathbf{K}_0 + \sum_{a=1}^A \mu_a \mathbf{K}_a \quad \boldsymbol{\mu} \in \mathbb{R}_+^A, \mathbf{K}_0 \succeq 0 \quad (7)$$

(7) is a conic combination of kernel matrices that induces the mapping:

$$\Phi_{\mathbf{K}''}(x_i) = [\Phi_0(x_i), \sqrt{\mu_1} \Phi_1(x_i), \dots, \sqrt{\mu_A} \Phi_A(x_i)] \quad (8)$$

The intuition behind this combination is that auxiliary kernels that satisfy many triplets are emphasized by weighing them more, and \mathbf{K}_0 , which is learned directly, can satisfy the triplets that cannot be satisfied by the conic combination of the auxiliary kernels. Consider, again, the example of a person comparing clothing items on an online store. She may compare clothes by characteristics such as color, size, and material, which are features that can be extracted and used to build the auxiliary kernels. However, other factors may influence how she compares clothes, such as designer or pattern, which may be omitted from the auxiliary kernels. In addition, she may have a personal sense of style that is impossible to be gained from features alone. \mathbf{K}_0 , and thus features induced by the mapping Φ_0 , is learned to model factors she uses to compare clothes that are omitted from the auxiliary kernels or cannot be modeled by extracted features. Using (7), we propose the following optimization:

$$\begin{aligned} \min_{\mathbf{K}_0, \boldsymbol{\mu}} \quad & E(\mathbf{K}'', \mathcal{T}) + \lambda_1 \text{trace}(\mathbf{K}_0) + \lambda_2 \|\boldsymbol{\mu}\|_1 \\ \text{s.t.} \quad & \mathbf{K}_0 \succeq 0, \boldsymbol{\mu} \geq 0 \end{aligned} \quad (9)$$

This objective has two regularization terms: trace regulation on \mathbf{K}_0 , and ℓ_1 -norm regularization on $\boldsymbol{\mu}$. Increasing λ_1 limits the expressiveness of \mathbf{K}_0 by reducing its rank, while increasing λ_2 reduces the influence of the auxiliary kernels by forcing the elements of $\boldsymbol{\mu}$ towards zero. Thus, λ_1 and λ_2 represent a trade-off between finding a kernel that is more influenced by \mathbf{K}_0 and one more influenced by the auxiliary kernels. Like RCKL-T, RCKL-AK is an SDP, but with $n^2 + A$ optimization variables. For practical A , RCKL-AK can be solved with minimal additional computational overhead to RCKL-T.

One desirable property of (9) is that under certain conditions, it is a convex optimization problem:

Proposition 1. *If E is a convex function in both \mathbf{K}_0 and $\boldsymbol{\mu}$, then (9) is a convex optimization problem.*

Proposition 1 is proven in Sec. B.1 of the appendix. While Prop. 1 may seem simple, it allows us to leverage traditional RCKL methods that contain

error functions that are convex in \mathbf{K}_0 and $\boldsymbol{\mu}$ in order to solve (9) using convex optimization techniques. Two such error functions are discussed in the following subsections.

STE-AK Stochastic Triplet Embedding (STE) [14] proposes the following probability that a triplet is satisfied:

$$p_{abc}^{\mathbf{K}} = \frac{\exp(-d_{\mathbf{K}}(x_a, x_b))}{\exp(-d_{\mathbf{K}}(x_a, x_b)) + \exp(-d_{\mathbf{K}}(x_a, x_c))}$$

If this probability is high, then x_a is closer to x_b than it is to x_c . As such, we minimize the negative sum of the log-probabilities over all triplets.

$$E_{\text{STE}}(\mathbf{K}'', \mathcal{T}) = - \sum_{(a,b,c) \in \mathcal{T}} \log(p_{abc}^{\mathbf{K}''}) \quad (10)$$

With this error function we call our method STE-AK and can state the following proposition:

Proposition 2. (10) is convex in both \mathbf{K}_0 and $\boldsymbol{\mu}$

Proposition 2 is proven in Sec. B.2 of the appendix. By Props. 1 and 2, STE-AK is a convex optimization problem.

GNMDS-AK Another potential error function is one motivated by Generalized Non-Metric Multidimensional Scaling (GNMDS) [1] which uses hinge loss:

$$E_{\text{GNMDS}}(\mathbf{K}'', \mathcal{T}) = \sum_{(a,b,c) \in \mathcal{T}} \max(0, d_{\mathbf{K}''}(x_a, x_b) - d_{\mathbf{K}''}(x_a, x_c) + 1) \quad (11)$$

We call our method with this error function GNMDS-AK. The hinge loss ensures that only triplets that are unsatisfied by a margin of one increase the objective. GNMDS-AK is also a convex optimization problem, due to Prop. 1 and the following:

Proposition 3. (11) is convex in both \mathbf{K}_0 and $\boldsymbol{\mu}$.

Proposition 3 is proven in Sec. B.3 of the appendix. For a more rigorous comparison of RCKL methods see [14]. We propose to solve both STE-AK and GNMDS-AK via projected gradient descent methods; both of which are outlined in Sec. C of the appendix.

5 Related Work

RCKL-AK can be viewed as a combination of *multiple kernel learning* (MKL) and *non-metric multidimensional scaling* (NMDS). Learning a non-negative sum of kernels, as in (4), appears often in MKL literature, which is focused on finding

efficient methods for learning a combination of predefined kernels for a learning task. The most widely studied problem in MKL has been Support Vector Classification [13,17,26,9]. To our knowledge there has been no application of MKL techniques to the task of learning a kernel from relative comparisons.

The RCKL problem posed in Sec. 2 is a special case of the NMDS problem first formalized in [1], which in turn is a generalization of the Shepard-Kruskal NMDS problem [21]. GNMDS, STE, and Crowd Kernel Learning (CKL) [22] are all methods that can be applied to the RCKL problem. However, none of these methods consider inputs beyond relative comparisons. Our work creates a novel RCKL method that uses ideas popularized in MKL research to incorporate side information into the learning problem.

Relative comparisons have also been considered in *metric learning* [19,4,8]. In metric learning the focus is on learning a distance metric over objects that can be applied to out-of-sample *objects*. This work focuses specifically on finding a kernel over given objects that generalizes well to out-of-sample (unobtained) *triplets*. In this way, the goal of metric learning methods is somewhat different than the one in this work. Two recent works propose methods to learn a Mahalanobis distance metric with multiple kernels: Metric Learning with Multiple Kernels (ML-MKL) [28] and Multiple Kernel Partial Order Embedding (MKPOE) [16]; the latter focusing exclusively on relative distance constraints similar to those in this work. The kernel learned by RCKL-AK induces a mapping that is fundamentally different than those learned by these metric learning techniques. Consider the mapping induced by one of the metric learning methods proposed in both [16] (Section 4.2) and [28] (Equation 6):

$$\Phi_{\mu,\Omega}(x) = \Omega [\sqrt{\mu_1}\Phi_1(x), \dots, \sqrt{\mu_A}\Phi_A(x)] \quad (12)$$

The derivation of this mapping can be found in Sec. D of the appendix. Here $\Omega \in \mathbb{R}^{m \times D}$ produces a new feature space by transforming the feature spaces induced by the auxiliary kernels. Without Ω , (12) learns a mapping similar to (5). The matrix Ω plays a role similar to the one \mathbf{K}_0 plays in RCKL-AK: it is learned to satisfy triplets that the auxiliary kernels alone cannot. Instead of linearly transforming the auxiliary kernel feature spaces, RCKL-AK implicitly learns new features that are concatenated onto the concatenated auxiliary kernel feature spaces (see (8)).

In both works, the authors propose non-convex optimizations to solve for their metrics, and, in addition, different convex relaxations. A critical issue with the convex solutions is that they employ SDPs over $n^2 * A$ (MKPOE-Full) and $n^2 * A^2$ (NR-ML-MKL) optimization variables, respectively. For moderately sized problems these methods are impractical. To resolve this issue, [16] propose a method that imposes further diagonal structure on the learned metric, reducing the number of optimization variables to $n * A$ (MKPOE-Diag), but in the process, greatly limit the structure of the metric. RCKL-AK is a convex SDP with $n^2 + A$ optimization variables that does not impose strict structure on the learned kernel. Unfortunately, by learning the unique kernel \mathbf{K}_0 directly and not the mapping Φ_0 or a generating function of \mathbf{K}'' , our method cannot be applied to

out-of-sample objects. Data analysis that does not require the addition of out-of-sample objects can be used over that kernel. There are many unsupervised and semi-supervised techniques that fit this use case.

6 Experiments

In order to show that RCKL-AK can learn kernels from few triplets that generalize well to unobtained triplets, we perform two experiments: one using synthetic data, and one using real-world data. In both experiments, we compare STE and GNMDS variants of RCKL-T, RCKL-MKL, and RCKL-AK, as well as non-convex and convex variants of MKPOE. For the MKPOE methods, we consider a triplet (a, b, c) to be satisfied if $d_{\mathbf{M}}(x_a, x_b) < d_{\mathbf{M}}(x_a, x_c)$, where $d_{\mathbf{M}}$ is the distance function defined by the metric. The STE and GNMDS implementations used are from [14], which are made publicly available on the authors’ websites. The MKL and AK versions were extended from these implementations. MKPOE implementations were provided to us by their original authors. All auxiliary kernels are normalized to unit trace, and all hyperparameters were validated via line or grid search using validation sets.

6.1 Synthetic Data

To generate synthetic data we begun by randomly generating 100 points in seven, independent, two-dimensional feature spaces where both dimensions were over the interval $[0, 1]$. Then, we created seven linear kernels, $\mathbf{K}_0, \dots, \mathbf{K}_6$ from these seven spaces. We combined four of the seven kernels:

$$\mathbf{K}^* = \frac{1}{2}\mathbf{K}_0 + \frac{1}{4}\mathbf{K}_1 + \frac{1}{6}\mathbf{K}_2 + \frac{1}{12}\mathbf{K}_3 \quad (13)$$

We then used \mathbf{K}^* as the ground truth to answer all possible, non-redundant triplets. Following the experimental setup in [22], we divided these triplets into 100 triplet “rounds”. A round is a set of triplets where each object appears once as the head a being compared to randomly chosen objects b and c . From the pool of rounds, 20 were chosen to be the training set, 10 were chosen to be the validation set, and the remaining rounds were the test set. This was repeated ten times to create ten different trials.

Next, we took all seven feature spaces and perturbed each point with randomly generated Gaussian noise. From these new spaces we created seven new linear kernels $\hat{\mathbf{K}}_0, \dots, \hat{\mathbf{K}}_6$, of which $\hat{\mathbf{K}}_1, \dots, \hat{\mathbf{K}}_6$ were used as the input auxiliary kernels in the experiment. Here, $\hat{\mathbf{K}}_1, \dots, \hat{\mathbf{K}}_3$ are kernels that represent attributes that influence how the ground truth makes comparisons between objects. $\hat{\mathbf{K}}_4, \dots, \hat{\mathbf{K}}_6$ contain information that is not considered when making comparisons, and \mathbf{K}_0 represents intuition about the objects that was not or cannot be input as an auxiliary kernel.

We wish to evaluate the performance of each method as the number of triplets increases. With more triplets each method should be able to build models that

satisfy more unobtained triplets. To show this we performed the following experiment. For each trial, the 20 training rounds and 10 validation rounds are divided into ten subsets, each containing two training rounds and one validation round. Starting with one of the subsets, each model is trained, setting the hyperparameters through cross-validation on the validation set, and evaluated on the test set. Then, another subset is added to the training and validation sets. We repeat this process until all ten subsets are included. We evaluate the methods by the total number of unsatisfied triplets in the test set divided by the total number of triplets in the test set (test error). Here, the test set represents unobtained triplets. For all of the following figures, error bars represent a 95% confidence interval.

Discussion: Figure 1 shows the mean test error over the ten trials as a function of the number of triplets in the training set. Both RCKL-MKL methods improve performance initially, but achieve their approximate peak performance early and fail to improve as triplets are added. This supports the claim that RCKL-MKL is limited by only being able to combine auxiliary kernels through a conic combination. Both RCKL-T methods perform much worse than either of the methods that use auxiliary kernels. Without the side information provided by the auxiliary kernels, RCKL-T cannot generalize to test triplets with few training triplets.




Fig. 1: Mean test error over ten trials of the synthetic data set

We believe this experiment demonstrates the utility of both \mathbf{K}_0 and the auxiliary kernels in RCKL-AK. With very few training triplets, the RCKL-AK methods relied on the auxiliary kernels, thus the performance is similar to the RCKL-MKL methods. As triplets are added, the RCKL-AK methods used \mathbf{K}_0 to satisfy the triplets that a conic combination of the auxiliary kernels could not. Further evidence for this is shown by the fact that the rank of \mathbf{K}_0 increased as the number of training triplets increased. For example, for STE-AK, the mean rank of \mathbf{K}_0 was 85.6, 94.2, and 96.2 for 200, 400, and 600 triplets in the training

set, respectively. In other words, the optimal settings of λ_1 and λ_2 made \mathbf{K}_0 more expressive as the number of triplets increased.

Ideally, the RCKL-AK methods should eliminate $\hat{\mathbf{K}}_4$, $\hat{\mathbf{K}}_5$, and $\hat{\mathbf{K}}_6$ from the model by reducing their corresponding weights μ_4 , μ_5 , and μ_6 to exactly zero. Figure 2 shows the values of the μ parameter for STE-AK and GNMDS-AK as the number of triplets increase. Both RCKL-AK methods correctly identify the three auxiliary kernels from which the ground truth kernel was created by setting their corresponding weight parameters to be non-zero. In addition, they assigned weights to the kernels roughly proportional to the ground truth. The three noise kernels were assigned very low, and often zero weights. The RCKL-MKL methods learned similar values for the elements of μ than those in Fig. 2. Since RCKL-MKL learned the relative importance of the auxiliary kernels with only few triplets, it had achieved approximately its peak performance and could not improve further with the addition of more triplets.




Fig. 2: Mean values of μ on synthetic data

Figure 3a shows the same STE-AK, GNMDS-AK, and GNMDS-MKL error plots as Fig. 1, but also includes three variations of MKPOE: A non-convex formulation (MKPOE-NC), and two convex formulations (MKPOE-Full and MKPOE-Diag). All metric learning methods perform very similarly, yet worse than RCKL-MKL and RCKL-AK. We believe that the MKPOE methods must transform the auxiliary kernel space drastically to satisfy the few triplets. By doing this they lose much of the information in the auxiliary kernels that allows RCKL-MKL and RCKL-AK methods to form more general solutions.

6.2 Music Artist Data

We also performed an experiment using comparisons among popular music artists. The *aset400* dataset [5] contains 16,385 relative comparisons of 412 music artists gathered from a web survey, and [15] provides five kernels built from various features describing each artist and their music. Two of the kernels were built from

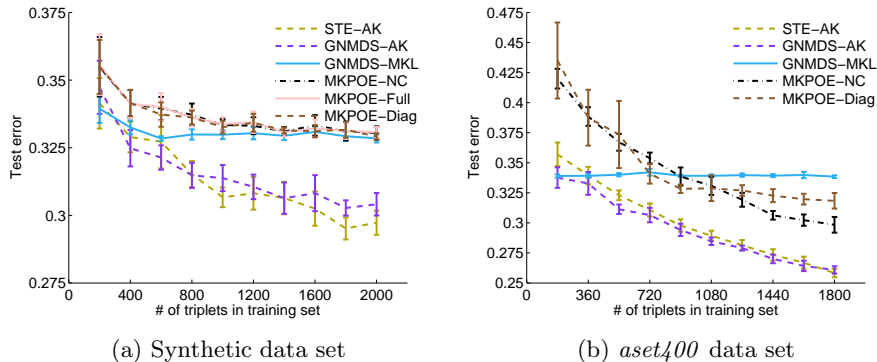


Fig. 3: Mean test error over ten trials

text descriptions of the artists, and three were built by extracting acoustic features from songs by each artist.

The *aset400* dataset provides a challenge absent in the synthetic data: not all artists appear in the same number of triplets. In fact, some artists never appear as the head of a triplet at all. As a result, this dataset represents a setting where feedback was gathered non-uniformly amongst the objects. In light of this, instead of training the models in rounds of triplets, we randomly chose 2000 triplets as the development set; the rest were used as the test set. Like before, we broke the development set into ten subsets, and progressively added subsets to the working set, training and testing each iteration. Ten percent of the working set was used for validation and 90 percent was used for training. The experiment was performed ten times on different randomly chosen train/validation/test splits.

Discussion: The results, shown in Fig. 4, are similar to those for the synthetic data with a few key differences. The RCKL-MKL methods did not perform as well relative to the RCKL-T methods. This could be attributed to the fact that the auxiliary kernels here did not reflect the triplets as well as those in the synthetic data. Only one kernel was consistently used in every iteration (the kernel built from artist tags). The rest were either given little weight or completely removed from the model. As with the synthetic data, with 200 and 400 training triplets the RCKL-AK methods performed as well as their respective RCKL-MKL counterparts, but as more triplets were added to the training set, the RCKL-AK methods began to perform much better. In this experiment, the RCKL-T methods became more competitive, but were outperformed significantly by RCKL-AK much of the time. This, again, could be because the auxiliary kernels were less useful than with the synthetic data.

Figure 3b compares the performance of MKPOE-NC and MKPOE-Diag on the *aset400* dataset to the RCKL-AK methods as well as GNMDS-MKL. MKPOE-Full could not be included in this experiment due to its impractically long run-time for an experiment of this size. Both MKPOE methods perform similarly, and seem to suffer greatly from the lack of meaningful auxiliary kernels,




Fig. 4: Mean test error over ten trials of the *aset400* data set

but do improve as the number of triplets increases. Still, over all experiments, the MKPOE methods have statistically significantly higher test error than both RCKL-AK methods.

7 Conclusions and Future Work

In this work we propose a method for learning a kernel from relative comparisons called Relative Comparison Kernel Learning with Auxiliary Kernels (RCKL-AK) that supplements given relative comparisons with auxiliary information. RCKL-AK is a convex SDP that can be solved by adding slight computational overhead to traditional methods and more efficiently than many metric learning alternatives. Experimentally, we show that RCKL-AK learns kernels that generalize to more out of sample relative comparisons than the aforementioned traditional and metric learning methods.

There are three main directions of future work. First, we wish to create a more efficient formulation of RCKL-AK. While common in solving SDPs, the most time-consuming step in RCKL-AK is performing eigendecomposition in order to project the kernel onto the PSD cone after each gradient step. In [12] the authors create a kernel learning method that eliminates the need to project onto the PSD cone. We will investigate extending their method to use auxiliary kernels. Second, we would like to extend our method to out-of-sample objects. In [2], generating functions are learned for kernels that were learnt through various popular techniques (LLE, Isomap, etc.). Similarly, we will attempt to learn a generating function for our kernel. Finally, we will explore practical applications of our method, specifically, the use of RCKL-AK for product recommendation.

A Proof of Theorems

A.1 Theorem 1

Below we prove Thm. 1 from Section 3

Proof. For n objects there are $\binom{n}{2}$ pair-wise distances between objects (we do not consider distances between objects and themselves). A triplet is a comparison between two pair-wise distances with a common object. Consider a single pairwise distance, and without loss of generality, let this distance be $d(x_a, x_b)$. There exists exactly $n - 2$ other pairwise distances that contain x_a (each with a different second object that is not x_a or x_b), and exactly $n - 2$ other pairwise distances that contain x_b . Thus, each pair-wise distance can be compared to $2n - 4$ other distances. As a result there are exactly, $\frac{1}{2}\binom{n}{2} * (2n - 4) = \frac{1}{2}(n^3 - 3n^2 + 2n)$ triplets (the $\frac{1}{2}$ comes from the fact that a single triplet (a, b, c) counts as an answer to the relative comparison “is x_a more similar to x_b than x_c ” and “is x_a more similar to x_c than x_b ”). \square

A.2 Theorem 2

To prove Thm. 2 from Section 3 we use the directed acyclic graph representation of relative comparisons given in Section 2 of [16]. A set of triplets \mathcal{T} is represented by a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Here, a vertex represents a distance between two objects, and an edge represents a relative comparison. Let $v_{\{1,2\}} \in \mathcal{V}$ be a vertex representing the distance between objects x_1 and x_2 indexed by the unordered pair $\{1, 2\}$. Note the slight notational change from the main body of the paper: Instead of indexing objects by alphabetic characters, we choose to index them by natural numbers for convenience. As an example, the triplet $(1, 2, 3)$ can be represented by a directed edge $e = (v_{\{1,2\}} \rightarrow v_{\{1,3\}}) \in \mathcal{E}$ from vertex $v_{\{1,2\}}$ to vertex $v_{\{1,3\}}$. Because this work considers triplets, vertices need to have a common object between them for there to be an edge between them. As an example, vertices $v_{\{1,2\}}$ and $v_{\{1,3\}}$ can have an edge between them, but $v_{\{1,2\}}$ and $v_{\{3,4\}}$ cannot. Any cycle in \mathcal{G} constitutes a conflict in relative comparisons. For this work we assume that there exists no conflicts or that conflicts can be resolved either algorithmically, such as in [16] or by querying the source of triplets again. From this point forward we say an edge is “valid” if it does not create a conflict and it connects two vertices that have a common object between them. An edge, and thus the corresponding triplet, can be inferred if there exists a path of length greater than 1. For example, if there is an edge from $v_{\{1,2\}}$ to $v_{\{1,3\}}$, and an edge from $v_{\{1,3\}}$ to $v_{\{1,4\}}$ then the edge from $v_{\{1,2\}}$ to $v_{\{1,4\}}$ can be inferred.

Theorem 2 considers the case where triplets are given one at a time. Let t_i , be the i th triplet to be given, and $\mathcal{T}_i = \{t_j | j \leq i\}$. Let \mathcal{T}_i^{trans} be the triplets that can be inferred by \mathcal{T}_i . For this analysis we study the directed acyclic graph analogues of these sets of triplets: \mathcal{G}_i and \mathcal{G}_i^{trans} . Let $e_i \in \mathcal{E}_i$ be the edge given at time i . We assume that an edge given at time i has not been given at time $j < i$. Next, we introduce the following recursive adversarial strategy for giving edges to this graph (square brackets (\square) to denote an ordered list of elements with subscripts indicating each element’s position in the list):

The following lines describe the process in which edges are given by Alg. 1 in more detail:

Algorithm 1 Adversarial Strategy

```
1: function ADVERSARY( $n \in \mathbb{N}_{>2}$ )
2:   if  $n == 3$  then
3:      $e_1 \leftarrow (v_{\{1,3\}} \rightarrow v_{\{1,2\}})$ 
4:      $e_2 \leftarrow (v_{\{2,3\}} \rightarrow v_{\{1,2\}})$ 
5:      $e_3 \leftarrow (v_{\{2,3\}} \rightarrow v_{\{1,3\}})$ 
6:      $idx \leftarrow 3$ 
7:   else
8:      $idx \leftarrow \frac{1}{2}((n-1)^3 - 3(n-1)^2 + 2(n-1))$ 
9:      $[e_1, \dots, e_{idx}] \leftarrow \text{ADVERSARY}(n-1)$ 
10:     $\mathcal{V}^{new} \leftarrow \text{GETNEWVERTICES}(n)$ 
11:     $[v_1^{old}, \dots, v_{\binom{n-1}{2}}^{old}] \leftarrow \text{FINDOLDVERTEXORDER}([e_1, \dots, e_{idx}])$ 
12:    for  $j = 1$  to  $\binom{n-1}{2}$  do
13:       $[e_{idx+1}, e_{idx+2}] \leftarrow \text{ADDEREMAININGEDGES}(v_j^{old}, \mathcal{V}^{new})$ 
14:       $idx \leftarrow idx + 2$ 
15:    end for
16:    for all vertices in  $\mathcal{V}^{new}$  do
17:       $v_{rnd} \leftarrow \text{SECTRANDOMVERTEX}(\mathcal{V}^{new})$ 
18:       $[e_{idx+1}, \dots, e_{idx'}], idx' \leftarrow \text{ADDEREMAININGEDGES}(v_{rnd}, \mathcal{V}^{new})$ 
19:       $idx \leftarrow idx'$ 
20:       $\mathcal{V}^{new} \leftarrow \mathcal{V}^{new} \setminus v_{rnd}$ 
21:    end for
22:  end if
23:  return  $[e_1, \dots, e_{idx}]$ 
24: end function
```

1. **Lines 2-6:** If $n=3$ (the base case) give the three edges on lines 3-5.
2. **Line 9:** Get the adversarial solution for $n-1$ objects.
3. **Line 10:** Call the procedure “GetNewVertices” which returns the set of vertices in the graph for n objects that are not in the graph for $n-1$ objects. This set is all vertices with n their indexing pair. We will call these vertices “new vertices” and all edges that connect to at least one new vertex “new edges”. We will call the vertices in the solution that are not new vertices “old vertices”, and the edges that connect two old vertices “old edges”. Note, that e_1, \dots, e_{idx} on line 9 contains every old edge and no new edges.
4. **Line 11:** Call the procedure “FindOldVertexOrder”, which takes the ordered set of old edges given as input. This procedure orders all old vertices in the order in which they appear in the incoming side of an edge. If an old vertex has no incoming edges it is added to the end of the list. For instance, if the input is the solution for the base case, the output would be the list $[v_{\{1,2\}}, v_{\{1,3\}}, v_{\{2,3\}}]$.
5. **Line 12:** Loop for each old vertex.
6. **Line 13:** Call the procedure “AddRemainingEdges” which exhaustively gives all remaining edges involving v_i^{old} in any order and points them towards v_i^{old} , assigning them to e_{idx+1} and e_{idx+2} . After this, all valid edges to or from v_i^{old} are given. Note that these new edges given here point from new vertices to v_i^{old} .
7. **Line 14:** The index for the given edge numbers is incremented.
8. **Line 16:** Loop for each new vertex
9. **Line 17:** Call procedure “SelectRandomVertex” which randomly selects a vertex from \mathcal{V}^{new} and assigns it to v_{rnd}
10. **Line 18:** Call the procedure “AddRemainingEdges” which gives all remaining valid edges $e_{idx+1}, \dots, e_{idx'}$ involving v_{rnd} and have them point toward v_{rnd} . Note that these edges only connect new vertices together.
11. **Lines 19 and 20:** Update the current number of edges given and remove v_{rnd} from \mathcal{V}^{new} so it does not get chosen in line 21 next iteration.
12. **Line 23:** Return the given edges, their order indicated by their indices.

Next, we state the following lemmas:

Lemma 1. *Algorithm 1 exhaustively gives all valid edges for a given n .*

Proof. The loop beginning on line 12 loops over all old vertices and exhaustively gives all edges involving these vertices. The loop beginning at line 16 does the same for all new vertices. Since the sets of old and new vertices together include every vertex, valid edges are given by line 23 and are returned. \square

In the subsequent proofs we will refer to any valid edge $e \notin \mathcal{E}_i$ to be “ungiven” at time i . A fact used to prove the following two lemmas is that when the adversarial strategy gives an edge at time i , a vertex or group of vertices can never be used to infer an ungiven edge if two conditions hold:

1. If all valid edges have been given involving a vertex or a group of vertices at time $i-1$

2. If there does not exist an edge leaving a vertex or group of vertices to a vertex with at least one ungiven edge at time $i - 1$

The first condition is somewhat obvious in that if all edges are given involving a vertex or group of vertices, then no ungiven edges involving this vertex can be inferred, because there are no ungiven edges involving this vertex. Thus, the only way for a vertex or group of vertices can be used to infer an ungiven edge is that a path through them is used to infer an ungiven edge between two external vertices. The second condition ensures that no such path exists. As a result, regardless of what edge the adversary gives at time i , the ungiven inferred edges of graph \mathcal{G}_i is always the same as that on \mathcal{G}_i with all vertices or groups of vertices removed that satisfy these two conditions. With this we can prove the following two lemmas:

Lemma 2. *Line 13 of Alg. 1 never gives edges that infer ungiven edges.*

Proof. Consider line 11 of Alg. 1. It orders the old vertices in the order in which they appear as the vertex being pointed to in the solution for n objects. Any vertex in this order only points to those that appear before it in the list, because when a vertex is chosen by this algorithm all valid edges involving this vertex are given pointing inward. Consider the first iteration of the for loop starting on line 12. The first vertex in this list has only incoming edges (namely, $v_{1,2}$ due to the base case). Adding all remaining ungiven edges inward to this node cannot create any inferred edges because all old edges are going inward due to it being first in the list, thus Lem. 2 holds. In addition, after line 13 this node satisfies the first condition above. It also satisfies the second condition as there exists no outgoing edges, thus no paths can go through it. By these two conditions, it can be eliminated from the graph. By doing so, all edges connecting it to other vertices are also eliminated. As a result, the second vertex effectively becomes the first vertex in the list, making it only have incoming edges. In the next iteration the second vertex is chosen, and, again, all remaining edges are pointed inward. Adding of these edges also satisfies Lem. 2, the same conditions hold, and it can effectively be removed. The loop repeats this for all old vertices. As such, for all iterations of the loop starting on line 12, line 13 never gives edges that infer ungiven edges. \square

Lemma 3. *Line 18 of Alg. 1 never gives edges that infer ungiven edges.*

Proof. By line 16 all edges connected to or from any old vertex has already been given via the loop starting on line 12. Thus, the first condition above is satisfied for all old vertices. Line 13 only adds edges from new vertices to old vertices, no edge leaves the subgraph of old vertices to the subgraph of new vertices. This satisfies the second condition. As a result, all old vertices can be eliminated from consideration. What is left is just the new vertices with no edges connecting them to any other vertex. Consider the first iteration of the loop starting on line 16. Selecting a vertex $v_{r,n}$ and giving all edges pointing inward cannot infer any edges, because it has only incoming edges. Thus, this operation does not violate

Lem. 3. In addition, the first condition is satisfied when all edges are given, and since there are no paths through this vertex, condition 2 is satisfied. As a result, this vertex too can then be eliminated. For the next iteration we now have one less new vertex, but, again, no edges connecting the ones left. This loop repeats until all edges are given. \square

Given these three lemmas we proceed by proving the following proposition:

Proposition 4. *For a given $n \in \mathbb{N}_{>2}$, there exists an adversarial strategy for giving edges such that $\forall_{i \in \{1, \dots, |\mathcal{E}^{total}|\}}$, $\mathcal{E}_i^{trans} \setminus \mathcal{E}_i = \emptyset$*

Proof. Let $\mathcal{E}_i = \{e_1, \dots, e_i\}$ where $\{e_1, \dots, e_i\}$ is the first i triplets in the ordered set returned by Alg. 1. First, we need to ensure \mathcal{E}_i can be constructed for all $i \in \{1, \dots, |\mathcal{E}^{total}|\}$. Lemma 1 states that Alg. 1 gives all edges for a given n , thus proving this. Lastly, we need to show that Alg. 1 never adds an edge to the sequence of edges that will infer an ungiven edge. To prove this, we use induction. The base case is $n = 3$, which is the fewest number of objects for which triplets can be defined. The base case is defined on lines 2-6 of Alg. 1. The edges e_1 and e_2 do not infer the third and final edge e_3 , thus Prop. 4 is true for the case $n = 3$.

Line 9 in Alg. 1 returns the solution for $n - 1$. Thus, if Alg. 1 never gives an edge that can infer an ungiven edge after this point in the algorithm, the inductive step is proven. The only two lines after line 9 that add edges are lines 13 and 18. By Lem. 2, line 13 never adds an edge that can infer an ungiven edge and by Lem. 3 neither does line 18. As a result the inductive step and thus Prop. 4 is proven. \square

Finally, with these propositions we can prove Thm. 2 from Section 3:

Proof. By Prop. 4, there exists an adversarial strategy such that, for all $i \in \{1, \dots, |\mathcal{E}^{total}|\}$, $\mathcal{E}_i^{trans} \setminus \mathcal{E}_i = \emptyset$. By construction of the directed acyclic graphs representing sets of triplets we can construct sets \mathcal{T}_i from \mathcal{E}_i , \mathcal{T}_i^{trans} from \mathcal{E}_i^{trans} , and \mathcal{T}^{total} from \mathcal{E}^{total} . Thus we can state, $\forall_{i=1, \dots, |\mathcal{T}^{total}|} : \mathcal{T}_i^{trans} \setminus \mathcal{T}_i = \emptyset$ for the worst case defined by Alg. 1. \square

A.3 Theorem 3

We begin by stating that the RCKL problem is equivalent to learning an embedding of objects in a space that satisfies distance constraints imposed by the triplets \mathcal{T} . Limiting the comparisons to triplets makes this embedding problem a special case of non-metric multidimensional scaling (NMDS). Also, the rank of the learned kernel \mathbf{K} is equivalent to the rank of the learned embedding (i.e. if $\mathbf{K} = \mathbf{A}\mathbf{A}^T$, then $\text{rank}(\mathbf{K}) = \text{rank}(\mathbf{A})$). As a result, enforcing $\text{rank}(\mathbf{K}) \leq r$ is equivalent to enforcing that the objects be embedded in \mathbb{R}^d where $d \leq r$. With this in mind we state the following proposition from Appendix A of [16]:

Proposition 5. *Any set of objects \mathcal{X} with a partial order of distances \mathcal{C} can be embedded in \mathbb{R}^{n-1}*

When \mathcal{C} is a total ordering over all pairs, the problem of embedding \mathcal{X} in a space that respects the distances in \mathcal{C} reduces to NMDS [11]. This implies that given any objects \mathcal{X} and any non-conflicting set of triplets \mathcal{T} , the objects in \mathcal{X} can be embedded in \mathbb{R}^{n-1} and satisfy all triplets in \mathcal{T} . Equivalently, given any set of objects \mathcal{X} and any non-conflicting set of triplets \mathcal{T} , all triplets in \mathcal{T} can be satisfied by a rank $n - 1$ kernel. With this in mind we can prove Thm. 3 from Section 3:

Proof. Assume an RCKL method enforces $\text{rank}(\mathbf{K}) = r$. By Prop. 5, r must be less than $n - 1$ to infer any triplets. Without loss of generality let $t = (a, b, c) \in \mathcal{T}^{\text{rank}-r}$. Because, $t \notin \mathcal{T}^{\text{trans}}$ by definition of $\mathcal{T}^{\text{rank}-r}$, $\mathcal{T} \cup (a, c, b)$ does not cause a conflict. An adversary can construct $\mathcal{T}^{\text{total}}$ with $\text{rank } r^{\text{total}} > r$, such that $(a, c, b) \in \mathcal{T}^{\text{total}}$, because any non-conflicting set of triplets $\mathcal{T}^{\text{total}}$ can be satisfied by some choice of $r^{\text{total}} \leq n - 1$ from Prop. 5. By virtue of the fact that $\mathcal{T}^{\text{total}}$ contains no conflicting triplets, we can then deduce that $(a, b, c) \notin \mathcal{T}^{\text{total}}$. Trivially, this can be said of any rank less than r , as well. Thus, any triplet $t \in \mathcal{T}^{\text{rank}-r}$ is not an element in $\mathcal{T}^{\text{total}}$ through adversarial choice of $\mathcal{T}^{\text{total}}$, proving Thm. 3. \square

B Proofs of Propositions

The strategy employed throughout this section to prove the stated functions are convex is to build each using convex combinations of convex functions. In order to use this strategy, we need to establish the following Lems. (Above each Lemma is a reference to a source for each lemma, respectively).

Section 3.2.1 of [3]:

Lemma 4. *If f and g are both convex functions, then so is their sum $f + g$.*

Section 2.3.2 of [3]:

Lemma 5. *Affine functions of the form $f(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{x} \in \mathbb{R}^n$, and $\mathbf{b} \in \mathbb{R}^m$ are convex in \mathbf{x}*

Section 3.2.3 of [3]:

Lemma 6. *If f and g are convex functions, then their point-wise maximum, $\max(f(x), g(x))$, is also convex.*

In addition, we will use the concept of *logarithmic convexity*:

Definition 4. *A function f is logarithmically convex (log-convex) if $f(x) > 0$ for all $x \in \text{dom}f$ and $\log f$ is convex.*

Which we then use to state the following Lemma.

Section 3.5.2 of [3]:

Lemma 7. *If f and g are both log-convex functions, then so is their sum $f + g$.*

Finally, for the sake of notational brevity, let us define the following:

$$\begin{aligned} d_{\mathbf{K}}^{ab} &= d_{\mathbf{K}}(x_a, x_b) \\ D_{\mathbf{K}}^{abc} &= d_{\mathbf{K}}(x_a, x_b) - d_{\mathbf{K}}(x_a, x_c) \end{aligned}$$

These short-hand versions of our established notation will be used throughout this section.

B.1 Proposition 1

Proof of Prop. 1:

Proof.

1. In order for (9) to be a convex optimization problem, it's objective and constraints must be convex in the optimization variables.
2. By Lem. 4, if $E(\mathbf{K}'', \mathcal{T})$, $\lambda_1 \text{trace}(\mathbf{K}_0)$, and $\lambda_2 \|\boldsymbol{\mu}\|_1$ are convex, then the objective in (9) is convex.
3. It is an assumption of the proposition that $E(\mathbf{K}'', \mathcal{T})$ is convex.
4. $\lambda_1 \text{trace}(\mathbf{K}_0)$ is defined as a constant times the sum of the diagonal elements of the matrix \mathbf{K}_0 , which is a sum of convex functions (Lem. 4).
5. $\lambda_2 \|\boldsymbol{\mu}\|_1$ is defined as a constant times the the sum of the absolute values of the elements of $\boldsymbol{\mu}$, which is a sum of convex functions (Lem. 4)
6. By lines 2, 3, 4, and 5, the objective in (9) is convex.
7. The positivity constraint on $\boldsymbol{\mu}$ is trivially convex.
8. The positive semidefinite constraint is known to be convex [25].
9. By lines 7 and 8, both constraints of (9) are convex.
10. By lines 1, 6, and 9, (9) is a convex optimization problem.

□

B.2 Proposition 2

Proof of Prop. 2:

Proof.

1. Moving the negation into the sum, (10) becomes the sum of terms of the following form:

$$-\log(p_{abc}^{\mathbf{K}''}) \tag{14}$$

$$= -\log\left(\frac{\exp(-d_{\mathbf{K}''}^{ab})}{\exp(-d_{\mathbf{K}''}^{ac}) + \exp(-d_{\mathbf{K}''}^{ab})}\right)$$

$$= -\log\left(\frac{1}{1 + \exp(D_{\mathbf{K}''}^{abc})}\right)$$

$$= -\log(1) \tag{15}$$

$$+ \log(\exp(D_{\mathbf{K}''}^{abc}) + 1) \tag{16}$$

2. By Lem. 4, if (14) is convex for all triplets (a, b, c) , then (10) is convex.
3. By Lem. 4, (14) is convex if both (15) and (16) are convex.
4. (15) is a constant and trivially convex.
5. By Definition 4, if $\exp(D_{\mathbf{K}''}^{abc}) + 1$ is log-convex, then (16) is convex.
6. By Lem. 7 if $\exp(D_{\mathbf{K}''}^{abc})$ and 1 are both log-convex, then $\exp(D_{\mathbf{K}''}^{abc}) + 1$ is log-convex.
7. 1 is a constant and trivially log-convex.
8. The codomain of the exponential function is \mathbb{R}^+ , so $\exp(D_{\mathbf{K}''}^{abc}) > 0$ for all \mathbf{K}'' , which satisfies the first condition for log-convexity.
9. To show $\log(\exp(D_{\mathbf{K}''}^{abc}))$ is convex, thus satisfying the second condition of log-convexity, we start by stating the following equivalence by using the definition of \mathbf{K}'' (7):

$$\begin{aligned} \log(\exp(D_{\mathbf{K}''}^{abc})) &= D_{\mathbf{K}''}^{abc} \\ &= D_{\mathbf{K}_0}^{abc} + \sum_{i=1}^A \mu_i D_{\mathbf{K}_i}^{abc} \end{aligned} \quad (17)$$

10. By Lem. 4 if $D_{\mathbf{K}_0}^{abc}$ and $\sum_{i=1}^A \mu_i D_{\mathbf{K}_i}^{abc}$ are convex, then (17) is convex.
11. Let $\mathbf{k}_{abc} = (D_{\mathbf{K}_1}^{abc}, \dots, D_{\mathbf{K}_A}^{abc})$. Then, $\sum_{i=1}^A \mu_i D_{\mathbf{K}_i}^{abc} = \mathbf{k}_{abc}^T \boldsymbol{\mu}$.
12. $\mathbf{k}_{abc}^T \boldsymbol{\mu}$ is an affine function of $\boldsymbol{\mu}$ that has the form $f(\boldsymbol{\mu}) = \mathbf{A}\boldsymbol{\mu} + \mathbf{b}$ where $\mathbf{A} = \mathbf{k}_{abc}^T$ and \mathbf{b} is 0.
13. By Lem. 5 and the previous step, $\mathbf{k}_{abc}^T \boldsymbol{\mu}$ is convex in $\boldsymbol{\mu}$.
14. Using the definition of kernel distance from (2) (Note the slight change in notation: \mathbf{K}_0^{ab} refers to the a th column and b th row of \mathbf{K}_0):

$$D_{\mathbf{K}_0}^{abc} = \mathbf{K}_0^{bb} + 2\mathbf{K}_0^{ac} - \mathbf{K}_0^{cc} - 2\mathbf{K}_0^{ab} \quad (18)$$

15. By Lem. 4, if the individual terms of (18) are convex then (18) is convex.
16. The individual terms of (18) are simply elements of \mathbf{K}_0 multiplied by scalars, which are convex in \mathbf{K}_0 .
17. By lines 3-16, (14) is convex.
18. By lines 1, 2, and 17, (10) is convex.

□

B.3 Proposition 3

Proof of Prop. 3:

Proof.

1. By Lem. 4 if $\max(0, D_{\mathbf{K}''}^{abc} + 1)$ is convex for any triplet (a, b, c) , then (11) is convex.
2. By Lem. 6, if 0 and $D_{\mathbf{K}''}^{abc} + 1$ are convex, then $\max(0, D_{\mathbf{K}''}^{abc} + 1)$ is convex.
3. 0 is trivially convex.
4. By Lem. 4, if 1 and $D_{\mathbf{K}''}^{abc}$ are convex, then $D_{\mathbf{K}''}^{abc} + 1$ is convex.

5. 1 is trivially convex.
6. Steps 9-16 of Section B.2 showed $D_{\mathbf{K}''}^{abc}$ is convex in the optimization variables.
7. By lines 2-6, $\max(0, D_{\mathbf{K}''}^{abc} + 1)$ is convex for any triplet (a, b, c) .
8. By line 1 and 6, (11) is convex

□

C Auxiliary Kernel Algorithms

In this section we state and discuss the algorithms used to solve STE-AK and GNMDS-AK. They share many steps, so we begin stating STE-AK and discussing it in detail, and then state GNMDS-AK and highlight how it differs from STE-AK.

C.1 STE-AK

After initialization, Alg. 2 repeats the following steps until convergence:

1. **Line 6:** Take a gradient step for \mathbf{K}_0 (trace regularization included)
2. **Line 7:** Take a gradient step for $\boldsymbol{\mu}$ (ℓ_1 -norm regularization included)
3. **Line 8:** Project \mathbf{K}_0 onto the positive semidefinite cone
4. **Line 9:** Project the elements of $\boldsymbol{\mu}$ to be non-negative
5. **Line 10:** Update \mathbf{K}''

Algorithm 2 STE-AK Projected Gradient Descent

Input:

$$\begin{aligned} \mathcal{X} &= \{x_1, \dots, x_n\}, \\ \mathcal{T} &= \{(a, b, c) | x_a \text{ is more similar to } x_b \text{ than } x_c\}, \\ \mathbf{K}_1, \dots, \mathbf{K}_A &\in \mathbb{R}^{n \times n}, \lambda_1 \in \mathbb{R}^+, \lambda_2 \in \mathbb{R}^+, \eta \in \mathbb{R}^+ \end{aligned}$$

Output:

$$\mathbf{K}'' \in \mathbb{R}^{n \times n}$$

- 1: $t \leftarrow 0$
 - 2: $\mathbf{K}_0^0 \leftarrow \mathbf{I}^{n \times n}$
 - 3: $\mu_1^0, \dots, \mu_A^0 \leftarrow \frac{1}{A}$
 - 4: $\mathbf{K}'' \leftarrow \mathbf{K}_0^0 + \sum_{a=1}^A \mu_a^0 \mathbf{K}_a$
 - 5: **repeat**
 - 6: $\mathbf{K}_0^{t+1} \leftarrow \mathbf{K}^t - \eta * (\nabla_{\mathbf{K}^t} E_{\text{STE}}(\mathbf{K}'', \mathcal{T}) + \lambda_1 * \mathbf{I}^{n \times n})$
 - 7: $\boldsymbol{\mu}^{t+1} \leftarrow \boldsymbol{\mu}^t - \eta * (\nabla_{\boldsymbol{\mu}^t} E_{\text{STE}}(\mathbf{K}'', \mathcal{T}) + \lambda_2)$
 - 8: $\mathbf{K}_0^{t+1} \leftarrow \Pi_{PSD}(\mathbf{K}_0^{t+1})$
 - 9: $\boldsymbol{\mu}^{t+1} \leftarrow \Pi_+(\boldsymbol{\mu}^{t+1})$
 - 10: $\mathbf{K}'' \leftarrow \mathbf{K}_0^{t+1} + \sum_{a=1}^A \mu_a^{t+1} \mathbf{K}_a$
 - 11: $t \leftarrow t + 1$
 - 12: **until** convergence
-

Projection onto the positive semi-definite cone is done by performing eigen-decomposition of the matrix \mathbf{K}_0 , assigning all negative eigenvalues to zero, and then reassembling \mathbf{K}_0 from the original eigenvectors and the new eigenvalues [20]. Projection of the elements of $\boldsymbol{\mu}$ to be non-negative is simply done by assigning all negative elements to be zero. The ℓ_1 -norm regularization in Alg. 2 is performed by adding $\boldsymbol{\lambda}_2 = \lambda_2 * \mathbf{1}^A$ to the gradient (Line 7). Since $\boldsymbol{\mu}$ is constrained to the non-negative orthant, the subgradient of the ℓ_1 -norm function needs only to be over the non-negative orthant, thus $\boldsymbol{\lambda}_2$ is an acceptable subgradient. Moreover, since we then project the elements of $\boldsymbol{\mu}$ to be non-negative, we get the desired effect of the ℓ_1 -norm regularization: the reduction of some elements to be exactly zero.

C.2 GNMDS-AK

Algorithm 3 GNMDS-AK Projected Gradient Descent

Input:

$$\begin{aligned} \mathcal{X} &= \{x_1, \dots, x_n\}, \\ \mathcal{T} &= \{(a, b, c) | x_a \text{ is more similar to } x_b \text{ than } x_c\}, \\ \mathbf{K}_1, \dots, \mathbf{K}_A &\in \mathbb{R}^{n \times n}, \lambda_1 \in \mathbb{R}^+, \lambda_2 \in \mathbb{R}^+, \eta \in \mathbb{R}^+ \end{aligned}$$

Output:

$$\mathbf{K}'' \in \mathbb{R}^{n \times n}$$

```

1:  $t \leftarrow 0$ 
2:  $\mathbf{K}_0^0 \leftarrow \mathbf{I}^{n \times n}$ 
3:  $\mu_1^0, \dots, \mu_A^0 \leftarrow \frac{1}{A}$ 
4:  $\mathbf{K}'' \leftarrow \mathbf{K}_0^0 + \sum_{a=1}^A \mu_a^0 \mathbf{K}_a$ 
5: repeat
6:    $\mathcal{T}' \leftarrow \text{getActiveTriplets}(\mathcal{T}, \mathbf{K}'')$ 
7:    $\mathbf{K}_0^{t+1} \leftarrow \mathbf{K}^t - \eta * (\nabla_{\mathbf{K}^t} E'_{\text{GNMDS}}(\mathbf{K}'', \mathcal{T}') + \lambda_1 * \mathbf{I}^{n \times n})$ 
8:    $\boldsymbol{\mu}^{t+1} \leftarrow \boldsymbol{\mu}^t - \eta * (\nabla_{\boldsymbol{\mu}^t} E'_{\text{GNMDS}}(\mathbf{K}'', \mathcal{T}') + \boldsymbol{\lambda}_2)$ 
9:    $\mathbf{K}_0^{t+1} \leftarrow \Pi_{\text{PSD}}(\mathbf{K}_0^{t+1})$ 
10:   $\boldsymbol{\mu}^{t+1} \leftarrow \Pi_+(\boldsymbol{\mu}^{t+1})$ 
11:   $\mathbf{K}'' \leftarrow \mathbf{K}_0^{t+1} + \sum_{a=1}^A \mu_a^{t+1} \mathbf{K}_a$ 
12:   $t \leftarrow t + 1$ 
13: until convergence

```

The differences between Algs. 2 and 3 are in lines 6, 7, and 8.

1. **Line 6:** The function *getActiveTriplets* checks each triplet in \mathcal{T} to see if it violates the margin constraint $d_{\mathbf{K}''}(x_a, x_c) - d_{\mathbf{K}''}(x_a, x_b) < 1$. The set \mathcal{T}' is assigned to be the set containing only triplets that violate this constraint. What this effectively does is form a set of only the triplets that would contribute to the error E_{GNMDS} , because the point-wise maximum would be non-zero.

2. **Lines 7 and 8:** These two lines are similar to lines 6 and 7 in Alg. 3, respectively. The first difference is in the error function E'_{GNMDS} :

$$E'_{\text{GNMDS}}(\mathbf{K}, \mathcal{T}) = \sum_{(a,b,c) \in \mathcal{T}} d_{\mathbf{K}}(x_a, x_b) - d_{\mathbf{K}}(x_a, x_c) + 1$$

This is simply the sum over the triplets of the second argument in the point-wise maximum of E_{GNMDS} . Because the second argument of E'_{GNMDS} in these lines is the set \mathcal{T}' , $E'_{\text{GNMDS}}(\mathbf{K}'', \mathcal{T}')$ is equivalent to $E_{\text{GNMDS}}(\mathbf{K}'', \mathcal{T})$ for all possible triplets, but is differentiable everywhere.

D Derivation of Equation (12)

We begin by stating the definition of the distance metric learned in Equation (5) of [28]:

$$\begin{aligned} d_{\mathbf{A}, \boldsymbol{\mu}}^2(x_i, x_j) &= \sum_{l=1}^A \mu_l (\mathbf{K}_l^i - \mathbf{K}_l^j)^T \mathbf{A} \sum_{l=1}^A \mu_l (\mathbf{K}_l^i - \mathbf{K}_l^j) \\ &= (\boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_i) - \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_j))^T \boldsymbol{\Phi}_{\boldsymbol{\mu}}(\mathcal{X})^T \mathbf{A} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(\mathcal{X}) (\boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_i) - \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_j)) \end{aligned} \quad (19)$$

Here, \mathbf{K}_l^i is the i th row of the l th auxiliary kernel, and $\mathbf{A} \in \mathbb{R}^{n \times n}$. Also:

$$\boldsymbol{\Phi}_{\boldsymbol{\mu}}(x) = [\sqrt{\mu_1} \boldsymbol{\Phi}_1(x), \dots, \sqrt{\mu_A} \boldsymbol{\Phi}_A(x)] \quad (20)$$

Finally, $\boldsymbol{\Phi}_{\boldsymbol{\mu}}(\mathcal{X})$ is the matrix in which the rows are the mappings of the elements of \mathcal{X} to $\mathcal{H}_{\boldsymbol{\mu}}$ by $\boldsymbol{\Phi}_{\boldsymbol{\mu}}$. Note the slight change in notation from [28]. The domain of their mapping $\boldsymbol{\Phi}_{\boldsymbol{\mu}}$ is over vector representations of the objects (\mathbf{x}). In this work we assume that the mapping can exist over elements in a set (x) that do not necessarily need to be vectors. By factoring the matrix $\mathbf{A} = \mathbf{B}^T \mathbf{B}$ and defining $\boldsymbol{\Omega} = \mathbf{B} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(\mathcal{X})$ we can distribute $\boldsymbol{\Omega}$ through (19):

$$\begin{aligned} d_{\mathbf{A}, \boldsymbol{\mu}}^2(x_i, x_j) &= (\boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_i) - \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_j))^T \boldsymbol{\Omega}^T \boldsymbol{\Omega} (\boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_i) - \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_j)) \\ &= (\boldsymbol{\Omega} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_i) - \boldsymbol{\Omega} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_j))^T (\boldsymbol{\Omega} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_i) - \boldsymbol{\Omega} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_j)) \\ &= d^2(\boldsymbol{\Omega} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_i), \boldsymbol{\Omega} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x_j)) \end{aligned} \quad (21)$$

Where $d^2(\mathbf{x}_i, \mathbf{x}_j)$ is the squared Euclidean distance between \mathbf{x}_i and \mathbf{x}_j . In this form and by (20), we can see that learning the proposed Mahalanobis distance metric in Equation (5) of [28] is equivalent to learning a squared Euclidean distance of points transformed by the following linear transformation:

$$\begin{aligned} \boldsymbol{\Phi}_{\boldsymbol{\mu}, \boldsymbol{\Omega}}(x) &= \boldsymbol{\Omega} \boldsymbol{\Phi}_{\boldsymbol{\mu}}(x) \\ &= \boldsymbol{\Omega} [\sqrt{\mu_1} \boldsymbol{\Phi}_1(x), \dots, \sqrt{\mu_A} \boldsymbol{\Phi}_A(x)] \end{aligned}$$

Which is the mapping defined in (12) from Section 5.

References

1. Agarwal, S., Wills, J., Cayton, L., Lanckriet, G., Kriegman, D., Belongie, S.: Generalized non-metric multidimensional scaling. In: AISTATS (2007)
2. Bengio, Y., Paiement, J., Vincent, P., Delalleau, O., Roux, N.L., Ouimet, M.: Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. NIPS 16 (2004)
3. Boyd, S., Vandenberghe, L.: Convex optimization. Cambridge university press (2004)
4. Davis, J., Kulis, B., Jain, P., Sra, S., Dhillon, I.: Information-theoretic metric learning. In: ICML (2007)
5. Ellis, D., Whitman, B., Berenzweig, A., Lawrence, S.: The quest for ground truth in musical artist similarity. In: ISMIR (2002)
6. Filippone, M., Camastra, F., Masulli, F., Rovetta, S.: A survey of kernel and spectral methods for clustering. Pattern Recognition 41(1), 176–190 (2008)
7. Gönen, M., Alpaydm, E.: Multiple kernel learning algorithms. JMLR 12, 2211–2268 (2011)
8. Huang, K., Ying, Y., Campbell, C.: Generalized sparse metric learning with relative comparisons. KAIS 28(1), 25–45 (2011)
9. Jain, A., Vishwanathan, S., Varma, M.: Spg-gmkl: Generalized multiple kernel learning with a million kernels. In: SIGKDD (August 2012)
10. Kendall, M., Gibbons, J.: Rank Correlation Methods. Oxford University Press, fifth edn. (1990)
11. Kruskal, J.: Nonmetric multidimensional scaling: a numerical method. Psychometrika 29(2), 115–129 (1964)
12. Kulis, B., Sustik, M., Dhillon, I.: Learning low-rank kernel matrices. In: ICML (2006)
13. Lanckriet, G., Cristianini, N., Bartlett, P., Ghaoui, L., Jordan, M.: Learning the kernel matrix with semidefinite programming. JMLR 5, 27–72 (2004)
14. van der Maaten, L., Weinberger, K.: Stochastic triplet embedding. In: MLSP (2012)
15. McFee, B., Lanckriet, G.: Heterogeneous embedding for subjective artist similarity. In: ISMIR (2009)
16. McFee, B., Lanckriet, G.: Learning multi-modal similarity. JMLR 12, 491–523 (2011)
17. Rakotomamonjy, A., Bach, F., Canu, S., Grandvalet, Y.: Simplemkl. JMLR 9, 2491–2521 (2008)
18. Schölkopf, B., Smola, A.: Learning with kernels: support vector machines, regularization, optimization and beyond. the MIT Press (2002)
19. Schultz, M., Joachims, T.: Learning a distance metric from relative comparisons. NIPS (2004)
20. Schwertman, N., Allen, D.: Smoothing an indefinite variance-covariance matrix. JSCS 9(3), 183–194 (1979)
21. Shepard, R.: The analysis of proximities: Multidimensional scaling with an unknown distance function. i. Psychometrika 27(2), 125–140 (1962)
22. Tamuz, O., Liu, C., Belongie, S., Shamir, O., Kalai, A.: Adaptively learning the crowd kernel. ICML (2011)
23. Tibshirani, R.: Regression shrinkage and selection via the lasso. JRSS. Series B pp. 267–288 (1996)
24. Valizadegan, H., Jin, R.: Generalized maximum margin clustering and unsupervised kernel learning. In: NIPS (2006)

25. Vandenberghe, L., Boyd, S.: Semidefinite programming. *SIAM review* 38(1), 49–95 (1996)
26. Varma, M., Babu, B.: More generality in efficient multiple kernel learning. In: *ICML* (2009)
27. Wang, F., Zhang, C.: Label propagation through linear neighborhoods. *TKDE* 20(1), 55–67 (2008)
28. Wang, J., Do, H., Woznica, A., Kalousis, A.: Metric learning with multiple kernels. In: *NIPS* (2011)
29. Zhou, D., Bousquet, O., Lal, T., Weston, J., Schölkopf, B.: Learning with local and global consistency (2004)
30. Zhu, X., Ghahramani, Z.: Learning from labeled and unlabeled data with label propagation. Tech. rep., CMU-CALD-02-107, Carnegie Mellon University (2002)