

Learning Multi-modal Similarity

Brian McFee

*Department of Computer Science and Engineering
University of California
San Diego, CA 92093-0404, USA*

BMCFEE@CS.UCSD.EDU

Gert Lanckriet

*Department of Electrical and Computer Engineering
University of California
San Diego, CA 92093-0407, USA*

GERT@ECE.UCSD.EDU

Editor: Tony Jebara

Abstract

In many applications involving multi-media data, the definition of similarity between items is integral to several key tasks, including nearest-neighbor retrieval, classification, and recommendation. Data in such regimes typically exhibits multiple modalities, such as acoustic and visual content of video. Integrating such heterogeneous data to form a holistic similarity space is therefore a key challenge to be overcome in many real-world applications.

We present a novel multiple kernel learning technique for integrating heterogeneous data into a single, unified similarity space. Our algorithm learns an optimal ensemble of kernel transformations which conform to measurements of human perceptual similarity, as expressed by relative comparisons. To cope with the ubiquitous problems of subjectivity and inconsistency in multi-media similarity, we develop graph-based techniques to filter similarity measurements, resulting in a simplified and robust training procedure.

Keywords: multiple kernel learning, metric learning, similarity

1. Introduction

In applications such as content-based recommendation systems, the definition of a proper similarity measure between items is crucial to many tasks, including nearest-neighbor retrieval and classification. In some cases, a natural notion of similarity may emerge from domain knowledge, for example, cosine similarity for bag-of-words models of text. However, in more complex, multi-media domains, there is often no obvious choice of similarity measure. Rather, viewing different aspects of the data may lead to several different, and apparently equally valid notions of similarity. For example, if the corpus consists of musical data, each song or artist may be represented simultaneously by acoustic features (such as rhythm and timbre), semantic features (tags, lyrics), or social features (collaborative filtering, artist reviews and biographies, etc). Although domain knowledge may be incorporated to endow each representation with an intrinsic geometry—and, therefore, a sense of similarity—the different notions of similarity may not be mutually consistent. In such cases, there is generally no obvious way to combine representations to form a unified similarity space which optimally integrates heterogeneous data.

Without extra information to guide the construction of a similarity measure, the situation seems hopeless. However, if some side-information is available, for example, as provided by human labelers, it can be used to formulate a learning algorithm to optimize the similarity measure.

This idea of using side-information to optimize a similarity function has received a great deal of attention in recent years. Typically, the notion of similarity is captured by a distance metric over a vector space (e.g., Euclidean distance in \mathbb{R}^d), and the problem of optimizing similarity reduces to finding a suitable embedding of the data under a specific choice of the distance metric. *Metric learning* methods, as they are known in the machine learning literature, can be informed by various types of side-information, including class labels (Xing et al., 2003; Goldberger et al., 2005; Globerson and Roweis, 2006; Weinberger et al., 2006), or binary *similar/dissimilar* pairwise labels (Wagstaff et al., 2001; Shental et al., 2002; Bilenko et al., 2004; Globerson and Roweis, 2007; Davis et al., 2007). Alternatively, multidimensional scaling (MDS) techniques are typically formulated in terms of quantitative (dis)similarity measurements (Torgerson, 1952; Kruskal, 1964; Cox and Cox, 1994; Borg and Groenen, 2005). In these settings, the representation of data is optimized so that distance (typically Euclidean) conforms to side-information. Once a suitable metric has been learned, similarity to new, unseen data can be computed either directly (if the metric takes a certain parametric form, for example, a linear projection matrix), or via out-of-sample extensions (Bengio et al., 2004).

To guide the construction of a similarity space for multi-modal data, we adopt the idea of using similarity measurements, provided by human labelers, as side-information. However, it has to be noted that, especially in heterogeneous, multi-media domains, similarity may itself be a highly subjective concept and vary from one labeler to the next (Ellis et al., 2002). Moreover, a single labeler may not be able to consistently decide if or to what extent two objects are similar, but she may still be able to reliably produce a rank-ordering of similarity over pairs (Kendall and Gibbons, 1990). Thus, rather than rely on quantitative similarity or hard binary labels of pairwise similarity, it is now becoming increasingly common to collect similarity information in the form of triadic or *relative* comparisons (Schultz and Joachims, 2004; Agarwal et al., 2007), in which human labelers answer questions of the form:

“Is x more similar to y or z ?”

Although this form of similarity measurement has been observed to be more stable than quantitative similarity (Kendall and Gibbons, 1990), and clearly provides a richer representation than binary pairwise similarities, it is still subject to problems of consistency and inter-labeler agreement. It is therefore imperative that great care be taken to ensure some sense of robustness when working with perceptual similarity measurements.

In the present work, our goal is to develop a framework for integrating multi-modal data so as to optimally conform to perceptual similarity encoded by relative comparisons. In particular, we follow three guiding principles in the development of our framework:

1. The algorithm should be robust against subjectivity and inter-labeler disagreement.
2. The algorithm must be able to integrate multi-modal data in an optimal way, that is, the distances between embedded points should conform to perceptual similarity measurements.
3. It must be possible to compute distances to new, unseen data as it becomes available.

We formulate this problem of heterogeneous feature integration as a learning problem: given a data set, and a collection of relative comparisons between pairs, we learn a representation of

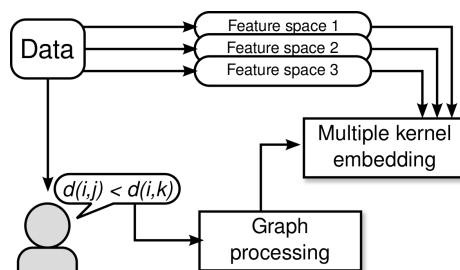


Figure 1: An overview of our proposed framework for multi-modal feature integration. Data is represented in multiple feature spaces (each encoded by a kernel function). Humans supply perceptual similarity measurements in the form of relative pairwise comparisons, which are in turn filtered by graph processing algorithms, and then used as constraints to optimize the multiple kernel embedding.

the data that optimally reproduces the similarity measurements. This type of embedding problem has been previously studied by Agarwal et al. (2007) and Schultz and Joachims (2004). However, Agarwal et al. (2007) provide no out-of-sample extension, and neither support heterogeneous feature integration, nor do they address the problem of noisy similarity measurements.

A common approach to optimally integrate heterogeneous data is based on *multiple kernel learning*, where each kernel encodes a different modality of the data. Heterogeneous feature integration via multiple kernel learning has been addressed by previous authors in a variety of contexts, including classification (Lanckriet et al., 2004; Zien and Ong, 2007; Kloft et al., 2009; Jagarlapudi et al., 2009), regression (Sonnenburg et al., 2006; Bach, 2008; Cortes et al., 2009), and dimensionality reduction (Lin et al., 2009). However, none of these methods specifically address the problem of learning a unified data representation which conforms to perceptual similarity measurements.

1.1 Contributions

Our contributions in this work are two-fold. First, we develop the *partial order embedding* (POE) framework (McFee and Lanckriet, 2009b), which allows us to use graph-theoretic algorithms to filter a collection of subjective similarity measurements for consistency and redundancy. We then formulate a novel multiple kernel learning (MKL) algorithm which learns an ensemble of feature space projections to produce a unified similarity space. Our method is able to produce non-linear embedding functions which generalize to unseen, out-of-sample data. Figure 1 provides a high-level overview of the proposed methods.

The remainder of this paper is structured as follows. In Section 2, we develop a graphical framework for interpreting and manipulating subjective similarity measurements. In Section 3, we derive an embedding algorithm which learns an optimal transformation of a single feature space. In Section 4, we develop a novel multiple-kernel learning formulation for embedding problems, and derive an algorithm to learn an optimal space from heterogeneous data. Section 5 provides experimental results illustrating the effects of graph-processing on noisy similarity data, and the effectiveness of the multiple-kernel embedding algorithm on a music similarity task with human

perception measurements. Finally, we prove hardness of dimensionality reduction in this setting in Section 6, and conclude in Section 7.

1.2 Preliminaries

A (*strict*) *partial order* is a binary relation R over a set Z ($R \subseteq Z^2$) which satisfies the following properties:¹

- Irreflexivity: $(a, a) \notin R$,
- Transitivity: $(a, b) \in R \wedge (b, c) \in R \Rightarrow (a, c) \in R$,
- Anti-symmetry: $(a, b) \in R \Rightarrow (b, a) \notin R$.

Every partial order can be equivalently represented as a directed acyclic graph (DAG), where each vertex is an element of Z and an edge is drawn from a to b if $(a, b) \in R$. For any partial order, R may refer to either the set of ordered tuples $\{(a, b)\}$ or the graph (DAG) representation of the partial order; the use will be clear from context.

For a directed graph G , we denote by G^∞ its *transitive closure*, that is, G^∞ contains an edge (i, j) if and only if there exists a path from i to j in G . Similarly, the *transitive reduction* (denoted G^{\min}) is the minimal graph with equivalent transitivity to G , that is, the graph with the fewest edges such that $(G^{\min})^\infty = G^\infty$.

Let $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$ denote the training set of n items. A *Euclidean embedding* is a function $g : \mathcal{X} \rightarrow \mathbb{R}^d$ which maps \mathcal{X} into a d -dimensional space equipped with the Euclidean (ℓ_2) metric:

$$\|x - y\|_2 = \sqrt{(x - y)^\top (x - y)}.$$

For any matrix B , let B_i denote its i^{th} column vector. A symmetric matrix $A \in \mathbb{R}^{n \times n}$ has a spectral decomposition $A = V\Lambda V^\top$, where $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$ is a diagonal matrix containing the eigenvalues of A , and V contains the eigenvectors of A . We adopt the convention that eigenvalues (and corresponding eigenvectors) are sorted in descending order. A is *positive semi-definite* (PSD), denoted by $A \succeq 0$, if each eigenvalue is non-negative: $\lambda_i \geq 0$, $i = 1, \dots, n$. Finally, a PSD matrix A gives rise to the Mahalanobis distance function

$$\|x - y\|_A = \sqrt{(x - y)^\top A (x - y)}.$$

2. A Graphical View of Similarity

Before we can construct an embedding algorithm for multi-modal data, we must first establish the form of side-information that will drive the algorithm, that is, the similarity measurements that will be collected from human labelers. There is an extensive body of work on the topic of constructing a geometric representation of data to fit perceptual similarity measurements. Primarily, this work falls under the umbrella of multi-dimensional scaling (MDS), in which perceptual similarity is modeled by numerical responses corresponding to the perceived “distance” between a pair of items, for

1. The standard definition of a (non-strict) partial order also includes the *reflexive* property: $\forall a, (a, a) \in \mathbb{R}$. For reasons that will become clear in Section 2, we take the *strict* definition here, and omit the reflexive property.

example, on a similarity scale of 1–10. (See Cox and Cox 1994 and Borg and Groenen 2005 for comprehensive overviews of MDS techniques.)

Because “distances” supplied by test subjects may not satisfy metric properties—in particular, they may not correspond to Euclidean distances—alternative *non-metric* MDS (NMDS) techniques have been proposed (Kruskal, 1964). Unlike classical or metric MDS techniques, which seek to preserve quantitative distances, NDMS seeks an embedding in which the rank-ordering of distances is preserved.

Since NMDS only needs the rank-ordering of distances, and not the distances themselves, the task of collecting similarity measurements can be simplified by asking test subjects to order pairs of points by similarity:

“Are i and j more similar than k and ℓ ?”

or, as a special case, the “triadic comparison”

“Is i more similar to j or ℓ ?”

Based on this kind of *relative comparison* data, the embedding problem can be formulated as follows. Given is a set of objects \mathcal{X} , and a set of similarity measurements $\mathcal{C} = \{(i, j, k, \ell)\} \subseteq \mathcal{X}^4$, where a tuple (i, j, k, ℓ) is interpreted as “ i and j are more similar than k and ℓ .” (This formulation subsumes the triadic comparisons model when $i = k$.) The goal is to find an embedding function $g : \mathcal{X} \rightarrow \mathbb{R}^d$ such that

$$\forall (i, j, k, \ell) \in \mathcal{C} : \|g(i) - g(j)\|^2 + 1 < \|g(k) - g(\ell)\|^2. \quad (1)$$

The unit margin is forced between the constrained distances for numerical stability.

Agarwal et al. (2007) work with this kind of relative comparison data and describe a generalized NMDS algorithm (GNMDS), which formulates the embedding problem as a semi-definite program. Schultz and Joachims (2004) derive a similar algorithm which solves a quadratic program to learn a linear, axis-aligned transformation of data to fit relative comparisons.

Previous work on relative comparison data often treats each measurement $(i, j, k, \ell) \in \mathcal{C}$ as effectively independent (Schultz and Joachims, 2004; Agarwal et al., 2007). However, due to their semantic interpretation as encoding pairwise similarity comparisons, and the fact that a pair (i, j) may participate in several comparisons with other pairs, there may be some *global* structure to \mathcal{C} which these previous methods are unable to exploit.

In Section 2.1, we develop a graphical framework to infer and interpret the global structure exhibited by the constraints of the embedding problem. Graph-theoretic algorithms presented in Section 2.2 then exploit this representation to filter this collection of noisy similarity measurements for consistency and redundancy. The final, reduced set of relative comparison constraints defines a partial order, making for a more robust and efficient embedding problem.

2.1 Similarity Graphs

To gain more insight into the underlying structure of a collection of comparisons \mathcal{C} , we can represent \mathcal{C} as a directed graph over \mathcal{X}^2 . Each vertex in the graph corresponds to a pair $(i, j) \in \mathcal{X}^2$, and an edge from (i, j) to (k, ℓ) corresponds to a similarity measurement (i, j, k, ℓ) (see Figure 2). Interpreting \mathcal{C} as a graph will allow us to infer properties of *global* (graphical) structure of \mathcal{C} . In particular, two facts become immediately apparent:




Figure 2: The graph representation (left) of a set of relative comparisons (right).

1. If C contains cycles, then there exists no embedding which can satisfy C .
2. If C is acyclic, any embedding that satisfies the transitive reduction C^{\min} also satisfies C .

The first fact implies that no algorithm can produce an embedding which satisfies all measurements if the graph is cyclic. In fact, the converse of this statement is also true: if C is acyclic, then an embedding exists in which all similarity measurements are preserved (see Appendix A). If C is cyclic, however, by analyzing the graph, it is possible to identify an “unlearnable” subset of C which must be violated by any embedding.

Similarly, the second fact exploits the transitive nature of distance comparisons. In the example depicted in Figure 2, any g that satisfies (j, k, j, l) and (j, l, i, k) must also satisfy (j, k, i, k) . In effect, the constraint (j, k, i, k) is redundant, and may also be safely omitted from C .

These two observations allude to two desirable properties in C for embedding methods: *transitivity* and *anti-symmetry*. Together with irreflexivity, these fit the defining characteristics of a *partial order*. Due to subjectivity and inter-labeler disagreement, however, most collections of relative comparisons will not define a partial order. Some graph processing, presented next, based on an approximate maximum acyclic subgraph algorithm, can reduce them to a partial order.

2.2 Graph Simplification

Because a set of similarity measurements C containing cycles cannot be embedded in any Euclidean space, C is inherently inconsistent. Cycles in C therefore constitute a form of *label noise*. As noted by Angelova (2004), label noise can have adverse effects on both model complexity and generalization. This problem can be mitigated by detecting and pruning noisy (confusing) examples, and training on a reduced, but certifiably “clean” set (Angelova et al., 2005; Vezhnevets and Barinova, 2007).

Unlike most settings, where the noise process affects each label independently—for example, random classification noise (Angluin and Laird, 1988)—the graphical structure of interrelated relative comparisons can be exploited to detect and prune inconsistent measurements. By eliminating similarity measurements which cannot be realized by any embedding, the optimization procedure can be carried out more efficiently and reliably on a reduced constraint set.

Ideally, when eliminating edges from the graph, we would like to retain as much information as possible. Unfortunately, this is equivalent to the *maximum acyclic subgraph* problem, which is NP-Complete (Garey and Johnson, 1979). A $1/2$ -approximate solution can be achieved by a simple greedy algorithm (Algorithm 1) (Berger and Shor, 1990).

Once a consistent subset of similarity measurements has been produced, it can be simplified further by pruning redundancies. In the graph view of similarity measurements, redundancies can be easily removed by computing the transitive reduction of the graph (Aho et al., 1972).

Algorithm 1 Approximate maximum acyclic subgraph (Aho et al., 1972)

Input: Directed graph $G = (V, E)$
Output: Acyclic graph G'
 $E' \leftarrow \emptyset$
for each $(u, v) \in E$ in random order **do**
 if $E' \cup \{(u, v)\}$ is acyclic **then**
 $E' \leftarrow E' \cup \{(u, v)\}$
 end if
end for
 $G' \leftarrow (V, E')$

By filtering the constraint set for consistency, we ensure that embedding algorithms are not learning from spurious information. Additionally, pruning the constraint set by transitive reduction focuses embedding algorithms on the most important core set of constraints while reducing overhead due to redundant information.

3. Partial Order Embedding

Now that we have developed a language for expressing similarity between items, we are ready to formulate the embedding problem. In this section, we develop an algorithm that learns a representation of data consistent with a collection of relative similarity measurements, and allows to map unseen data into the learned similarity space after learning. In order to accomplish this, we will assume a feature representation for \mathcal{X} . By parameterizing the embedding function g in terms of the feature representation, we will be able to apply g to any point in the feature space, thereby generalizing to data outside of the training set.

3.1 Linear Projection

To start, we assume that the data originally lies in some Euclidean space, that is, $\mathcal{X} \subset \mathbb{R}^D$. There are of course many ways to define an embedding function $g : \mathbb{R}^D \rightarrow \mathbb{R}^d$. Here, we will restrict attention to embeddings parameterized by a linear projection matrix M , so that for a vector $x \in \mathbb{R}^D$,

$$g(x) \doteq Mx.$$

Collecting the vector representations of the training set as columns of a matrix $X \in \mathbb{R}^{D \times n}$, the inner product matrix of the embedded points can be characterized as

$$A = X^T M^T M X.$$

Now, for a relative comparison (i, j, k, ℓ) , we can express the distance constraint (1) between embedded points as follows:

$$(X_i - X_j)^T M^T M (X_i - X_j) + 1 \leq (X_k - X_\ell)^T M^T M (X_k - X_\ell).$$

These inequalities can then be used to form the constraint set of an optimization problem to solve for M . Because, in general, C may not be satisfiable by a linear projection of \mathcal{X} , we soften the

constraints by introducing a slack variable $\xi_{ijkl} \geq 0$ for each constraint, and minimize the empirical hinge loss over constraint violations $1/|C| \sum_C \xi_{ijkl}$. This choice of loss function can be interpreted as a convex approximation to a generalization of the area under an ROC curve (see Appendix C).

To avoid over-fitting, we introduce a regularization term $\text{tr}(M^T M)$, and a trade-off parameter $\beta > 0$ to control the balance between regularization and loss minimization. This leads to a regularized risk minimization objective:

$$\begin{aligned} \min_{M, \xi \geq 0} \quad & \text{tr}(M^T M) + \frac{\beta}{|C|} \sum_C \xi_{ijkl} & (2) \\ \text{s.t.} \quad & (X_i - X_j)^T M^T M (X_i - X_j) + 1 \leq (X_k - X_\ell)^T M^T M (X_k - X_\ell) + \xi_{ijkl}, \\ & \forall (i, j, k, \ell) \in C. \end{aligned}$$

After learning M by solving this optimization problem, the embedding can be extended to out-of-sample points x' by applying the projection: $x' \mapsto Mx'$.

Note that the distance constraints in (2) involve differences of quadratic terms, and are therefore not convex. However, since M only appears in the form $M^T M$ in (2), the optimization problem can be expressed in terms of a positive semi-definite matrix $W \doteq M^T M$. This change of variables results in Algorithm 2, a (convex) semi-definite programming (SDP) problem (Boyd and Vandenberghe, 2004), since objective and constraints are linear in W , including the linear matrix inequality $W \succeq 0$. The corresponding inner product matrix is

$$A = X^T W X.$$

Finally, after the optimal W is found, the embedding function $g : \mathbb{R}^D \rightarrow \mathbb{R}^D$ can be recovered from the spectral decomposition of W :

$$W = V \Lambda V^T \quad \Rightarrow \quad g(x) = \Lambda^{1/2} V^T x,$$

and a d -dimensional approximation can be recovered by taking the leading d eigenvectors of W .

Algorithm 2 Linear partial order embedding (LPOE)

Input: n objects \mathcal{X} ,
 partial order \mathcal{C} ,
 data matrix $X \in \mathbb{R}^{D \times n}$,
 $\beta > 0$

Output: mapping $g : \mathcal{X} \rightarrow \mathbb{R}^d$

$$\begin{aligned} \min_{W, \xi} \quad & \text{tr}(W) + \frac{\beta}{|C|} \sum_C \xi_{ijkl} \\ & d(x_i, x_j) \doteq (X_i - X_j)^T W (X_i - X_j) \\ & d(x_i, x_j) + 1 \leq d(x_k, x_\ell) + \xi_{ijkl} \\ & \xi_{ijkl} \geq 0 & \forall (i, j, k, \ell) \in C \\ & W \succeq 0 \end{aligned}$$

3.2 Non-linear Projection via Kernels

The formulation in Algorithm 2 can be generalized to support non-linear embeddings by the use of kernels, following the method of Globerson and Roweis (2007): we first map the data into a reproducing kernel Hilbert space (RKHS) \mathcal{H} via a feature map ϕ with corresponding kernel function $k(x, y) = \langle \phi(x), \phi(y) \rangle_{\mathcal{H}}$; then, the data is mapped to \mathbb{R}^d by a linear projection $M : \mathcal{H} \rightarrow \mathbb{R}^d$. The embedding function $g : \mathcal{X} \rightarrow \mathbb{R}^d$ is therefore the composition of the projection M with ϕ :

$$g(x) = M(\phi(x)).$$

Because ϕ may be non-linear, this allows us to learn a non-linear embedding g .

More precisely, we consider M as being comprised of d elements of \mathcal{H} , that is, $\{\omega_1, \omega_2, \dots, \omega_d\} \subseteq \mathcal{H}$. The embedding g can thus be expressed as

$$g(x) = (\langle \omega_p, \phi(x) \rangle_{\mathcal{H}})_{p=1}^d,$$

where $(\cdot)_{p=1}^d$ denotes concatenation.

Note that in general, \mathcal{H} may be infinite-dimensional, so directly optimizing M may not be feasible. However, by appropriately regularizing M , we may invoke the generalized representer theorem (Schölkopf et al., 2001). Our choice of regularization is the Hilbert-Schmidt norm of M , which, in this case, reduces to

$$\|M\|_{\text{HS}}^2 = \sum_{p=1}^d \langle \omega_p, \omega_p \rangle_{\mathcal{H}}.$$

With this choice of regularization, it follows from the generalized representer theorem that at an optimum, each ω_p must lie in the span of the training data, that is,

$$\omega_p = \sum_{i=1}^n N_{pi} \phi(x_i), \quad p = 1, \dots, d,$$

for some real-valued matrix $N \in \mathbb{R}^{d \times n}$. If Φ is a matrix representation of \mathcal{X} in \mathcal{H} (i.e., $\Phi_i = \phi(x_i)$ for $x_i \in \mathcal{X}$), then the projection operator M can be expressed as

$$M = N\Phi^{\top}. \tag{3}$$

We can now reformulate the embedding problem as an optimization over N rather than M . Using (3), the regularization term can be expressed as

$$\|M\|_{\text{HS}}^2 = \text{tr}(\Phi N^{\top} N \Phi^{\top}) = \text{tr}(N^{\top} N \Phi^{\top} \Phi) = \text{tr}(N^{\top} N K),$$

where K is the kernel matrix over \mathcal{X} :

$$K = \Phi^{\top} \Phi, \quad \text{with } K_{ij} = \langle \phi(x_i), \phi(x_j) \rangle_{\mathcal{H}} = k(x_i, x_j).$$

To formulate the distance constraints in terms of N , we first express the embedding g in terms of N and the kernel function:

$$g(x) = M(\phi(x)) = N\Phi^{\top}(\phi(x)) = N(\langle \Phi_i, \phi(x) \rangle_{\mathcal{H}})_{i=1}^n = N(k(x_i, x))_{i=1}^n = NK_x,$$

where K_x is the column vector formed by evaluating the kernel function k at x against the training set. The inner product matrix of embedded points can therefore be expressed as

$$A = KN^T NK,$$

which allows to express the distance constraints in terms of N and the kernel matrix K :

$$(K_i - K_j)^T N^T N (K_i - K_j) + 1 \leq (K_k - K_\ell)^T N^T N (K_k - K_\ell).$$

The embedding problem thus amounts to solving the following optimization problem in N and ξ :

$$\begin{aligned} \min_{N, \xi \geq 0} \quad & \text{tr}(N^T NK) + \frac{\beta}{|\mathcal{C}|} \sum_{\mathcal{C}} \xi_{ijk\ell} \\ \text{s.t.} \quad & (K_i - K_j)^T N^T N (K_i - K_j) + 1 \leq (K_k - K_\ell)^T N^T N (K_k - K_\ell) + \xi_{ijk\ell}, \\ & \forall (i, j, k, \ell) \in \mathcal{C}. \end{aligned} \tag{4}$$

Again, the distance constraints in (4) are non-convex due to the differences of quadratic terms. And, as in the previous section, N only appears in the form of inner products $N^T N$ in (4)—both in the constraints, and in the regularization term—so we can again derive a convex optimization problem by changing variables to $W \doteq N^T N \succeq 0$. The resulting embedding problem is listed as Algorithm 3, again a semi-definite programming problem (SDP), with an objective function and constraints that are linear in W .

After solving for W , the matrix N can be recovered by computing the spectral decomposition $W = V\Lambda V^T$, and defining $N = \Lambda^{1/2}V^T$. The resulting embedding function takes the form:

$$g(x) = \Lambda^{1/2}V^T K_x.$$

As in Schultz and Joachims (2004), this formulation can be interpreted as learning a Mahalanobis distance metric $\Phi W \Phi^T$ over \mathcal{H} . More generally, we can view this as a form of kernel learning, where the kernel matrix A is restricted to the set

$$A \in \{KWK : W \succeq 0\}. \tag{5}$$

3.3 Connection to GNMDS

We conclude this section by drawing a connection between Algorithm 3 and the generalized non-metric MDS (GNMDS) algorithm of Agarwal et al. (2007).

First, we observe that the i -th column, K_i , of the kernel matrix K can be expressed in terms of K and the i^{th} standard basis vector e_i :

$$K_i = Ke_i.$$

From this, it follows that distance computations in Algorithm 3 can be equivalently expressed as

$$\begin{aligned} d(x_i, x_j) &= (K_i - K_j)^T W (K_i - K_j) \\ &= (K(e_i - e_j))^T W (K(e_i - e_j)) \\ &= (e_i - e_j)^T K^T W K (e_i - e_j). \end{aligned} \tag{6}$$

Algorithm 3 Kernel partial order embedding (KPOE)

Input: n objects \mathcal{X} ,
 partial order \mathcal{C} ,
 kernel matrix K ,
 $\beta > 0$

Output: mapping $g : \mathcal{X} \rightarrow \mathbb{R}^n$

$$\begin{aligned} \min_{W, \xi} \quad & \text{tr}(WK) + \frac{\beta}{|\mathcal{C}|} \sum_{\mathcal{C}} \xi_{ijkl} \\ d(x_i, x_j) \doteq & (K_i - K_j)^\top W (K_i - K_j) \\ d(x_i, x_j) + 1 \leq & d(x_k, x_\ell) + \xi_{ijkl} \\ \xi_{ijkl} \geq 0 & \quad \forall (i, j, k, \ell) \in \mathcal{C} \\ W \succeq 0 & \end{aligned}$$

If we consider the extremal case where $K = I$, that is, we have no prior feature-based knowledge of similarity between points, then Equation 6 simplifies to

$$d(x_i, x_j) = (e_i - e_j)^\top I W I (e_i - e_j) = W_{ii} + W_{jj} - W_{ij} - W_{ji}.$$

Therefore, in this setting, rather than defining a feature transformation, W directly encodes the inner products between embedded training points. Similarly, the regularization term becomes

$$\text{tr}(WK) = \text{tr}(WI) = \text{tr}(W).$$

Minimizing the regularization term can be interpreted as minimizing a convex upper bound on the rank of W (Boyd and Vandenberghe, 2004), which expresses a preference for low-dimensional embeddings. Thus, by setting $K = I$ in Algorithm 3, we directly recover the GNMDS algorithm.

Note that directly learning inner products between embedded training data points rather than a feature transformation does not allow a meaningful out-of-sample extension, to embed unseen data points. On the other hand, by Equation 5, it is clear that the algorithm optimizes over the entire cone of PSD matrices. Thus, if \mathcal{C} defines a DAG, we could exploit the fact that a partial order over distances always allows an embedding which satisfies all constraints in \mathcal{C} (see Appendix A) to eliminate the slack variables from the program entirely.

4. Multiple Kernel Embedding

In the previous section, we derived an algorithm to learn an optimal projection from a kernel space \mathcal{H} to \mathbb{R}^d such that Euclidean distance between embedded points conforms to perceptual similarity. If, however, the data is heterogeneous in nature, it may not be realistic to assume that a single feature representation can sufficiently capture the inherent structure in the data. For example, if the objects in question are images, it may be natural to encode texture information by one set of features, and color in another, and it is not immediately clear how to reconcile these two disparate sources of information into a single kernel space.

However, by encoding each source of information independently by separate feature spaces $\mathcal{H}^1, \mathcal{H}^2, \dots$ —equivalently, kernel matrices K^1, K^2, \dots —we can formulate a multiple kernel learning algorithm to optimally combine all feature spaces into a single, unified embedding space. In this section, we will derive a novel, projection-based approach to multiple-kernel learning and extend Algorithm 3 to support heterogeneous data in a principled way.

4.1 Unweighted Combination

Let K^1, K^2, \dots, K^m be a set of kernel matrices, each with a corresponding feature map ϕ^p and RKHS \mathcal{H}^p , for $p \in 1, \dots, m$. One natural way to combine the kernels is to look at the product space, which is formed by concatenating the feature maps:

$$\phi(x_i) = (\phi^1(x_i), \phi^2(x_i), \dots, \phi^m(x_i)) = (\phi^p(x_i))_{p=1}^m.$$

Inner products can be computed in this space by summing across each feature map:

$$\langle \phi(x_i), \phi(x_j) \rangle = \sum_{p=1}^m \langle \phi^p(x_i), \phi^p(x_j) \rangle_{\mathcal{H}^p}.$$

resulting in the *sum-kernel*—also known as the *average kernel* or *product space kernel*. The corresponding kernel matrix can be conveniently represented as the unweighted sum of the base kernel matrices:

$$\widehat{K} = \sum_{p=1}^m K^p. \tag{7}$$

Since \widehat{K} is a valid kernel matrix itself, we could use \widehat{K} as input for Algorithm 3. As a result, the algorithm would learn a kernel from the family

$$\begin{aligned} \mathcal{K}_1 &= \left\{ \left(\sum_{p=1}^m K^p \right) W \left(\sum_{p=1}^m K^p \right) : W \succeq 0 \right\} \\ &= \left\{ \sum_{p,q=1}^m K^p W K^q : W \succeq 0 \right\}. \end{aligned}$$

4.2 Weighted Combination

Note that \mathcal{K}_1 treats each kernel equally; it is therefore impossible to distinguish *good* features (i.e., those which can be transformed to best fit \mathcal{C}) from *bad* features, and as a result, the quality of the resulting embedding may be degraded. To combat this phenomenon, it is common to learn a scheme for weighting the kernels in a way which is optimal for a particular task. The most common approach to combining the base kernels is to take a positive-weighted sum

$$\sum_{p=1}^m \mu_p K^p \quad (\mu_p \geq 0),$$

where the weights μ_p are learned in conjunction with a predictor (Lanckriet et al., 2004; Sonnenburg et al., 2006; Bach, 2008; Cortes et al., 2009). Equivalently, this can be viewed as learning a feature map

$$\phi(x_i) = (\sqrt{\mu_p} \phi^p(x_i))_{p=1}^m,$$

where each base feature map has been scaled by the corresponding weight $\sqrt{\mu_p}$.

Applying this reasoning to learning an embedding that conforms to perceptual similarity, one might consider a two-stage approach to parameterizing the embedding (Figure 3(a)): first construct a weighted kernel combination, and then project from the combined kernel space. Lin et al. (2009) formulate a dimensionality reduction algorithm in this way. In the present setting, this would be achieved by simultaneously optimizing W and μ_p to choose an inner product matrix A from the set

$$\begin{aligned} \mathcal{K}_2 &= \left\{ \left(\sum_{p=1}^m \mu_p K^p \right) W \left(\sum_{p=1}^m \mu_p K^p \right) : W \succeq 0, \forall p, \mu_p \geq 0 \right\} \\ &= \left\{ \sum_{p,q=1}^m \mu_p K^p W \mu_q K^q : W \succeq 0, \forall p, \mu_p \geq 0 \right\}. \end{aligned}$$

The corresponding distance constraints, however, contain differences of terms cubic in the optimization variables W and μ_p :

$$\sum_{p,q} (K_i^p - K_j^p)^\top \mu_p W \mu_q (K_i^q - K_j^q) + 1 \leq \sum_{p,q} (K_k^p - K_\ell^p)^\top \mu_p W \mu_q (K_k^q - K_\ell^q),$$

and are therefore non-convex and difficult to optimize. Even simplifying the class by removing cross-terms, that is, restricting A to the form

$$\mathcal{K}_3 = \left\{ \sum_{p=1}^m \mu_p^2 K^p W K^p : W \succeq 0, \forall p, \mu_p \geq 0 \right\},$$

still leads to a non-convex problem, due to the difference of positive quadratic terms introduced by distance calculations:

$$\sum_{p=1}^m (K_i^p - K_j^p)^\top \mu_p^2 W (\mu_p K_i^p - K_j^p) + 1 \leq \sum_{p=1}^m (K_k^p - K_\ell^p)^\top \mu_p^2 W (\mu_p K_k^p - K_\ell^p).$$

However, a more subtle problem with this formulation lies in the assumption that a single weight can characterize the contribution of a kernel to the optimal embedding. In general, different kernels may be more or less informative on different subsets of \mathcal{X} or different regions of the corresponding feature space. Constraining the embedding to a single metric W with a single weight μ_p for each kernel may be too restrictive to take advantage of this phenomenon.

4.3 Concatenated Projection

We now return to the original intuition behind Equation 7. The sum-kernel represents the inner product between points in the space formed by concatenating the base feature maps ϕ^p . The sets \mathcal{K}_2 and \mathcal{K}_3 characterize projections of the weighted combination space, and turn out to not be amenable to efficient optimization (Figure 3(a)). This can be seen as a consequence of prematurely combining kernels prior to projection.

Rather than projecting the (weighted) concatenation of $\phi^p(\cdot)$, we could alternatively concatenate learned projections $M^p(\phi^p(\cdot))$, as illustrated by Figure 3(b). Intuitively, by defining the embedding as the concatenation of m different projections, we allow the algorithm to learn an ensemble of

projections, each tailored to its corresponding domain space and jointly optimized to produce an optimal space. By contrast, the previously discussed formulations apply essentially the same projection to each (weighted) feature space, and are thus much less flexible than our proposed approach. Mathematically, an embedding function of this form can be expressed as the concatenation

$$g(x) = (M^p (\phi^p(x)))_{p=1}^m.$$

Now, given this characterization of the embedding function, we can adapt Algorithm 3 to optimize over multiple kernels. As in the single-kernel case, we introduce regularization terms for each projection operator M^p

$$\sum_{p=1}^m \|M^p\|_{\text{HS}}^2$$

to the objective function. Again, by invoking the representer theorem for each M^p , it follows that

$$M^p = N^p (\Phi^p)^\top,$$

for some matrix N^p , which allows to reformulate the embedding problem as a joint optimization over N^p , $p = 1, \dots, m$ rather than M^p , $p = 1, \dots, m$. Indeed, the regularization terms can be expressed as

$$\sum_{p=1}^m \|M^p\|_{\text{HS}}^2 = \sum_{p=1}^m \text{tr} \left((N^p)^\top (N^p) K^p \right). \quad (8)$$

The embedding function can now be rewritten as

$$g(x) = (M^p (\phi^p(x)))_{p=1}^m = (N^p K_x^p)_{p=1}^m, \quad (9)$$

and the inner products between embedded points take the form:

$$\begin{aligned} A_{ij} = \langle g(x_i), g(x_j) \rangle &= \sum_{p=1}^m (N^p K_i^p)^\top (N^p K_j^p) \\ &= \sum_{p=1}^m (K_i^p)^\top (N^p)^\top (N^p) (K_j^p). \end{aligned}$$

Similarly, squared Euclidean distance also decomposes by kernel:

$$\|g(x_i) - g(x_j)\|^2 = \sum_{p=1}^m (K_i^p - K_j^p)^\top (N^p)^\top (N^p) (K_i^p - K_j^p). \quad (10)$$

Finally, since the matrices N^p , $p = 1, \dots, m$ only appear in the form of inner products in (8) and (10), we may instead optimize over PSD matrices $W^p = (N^p)^\top (N^p)$. This renders the regularization terms (8) and distances (10) linear in the optimization variables W^p . Extending Algorithm 3 to this parameterization of $g(\cdot)$ therefore results in an SDP, which is listed as Algorithm 4. To solve the SDP, we implemented a gradient descent solver, which is described in Appendix B.

The class of kernels over which Algorithm 4 optimizes can be expressed simply as the set

$$\mathcal{K}_4 = \left\{ \sum_{p=1}^m K^p W^p K^p : \forall p, W^p \succeq 0 \right\}$$

Kernel class	Learned kernel matrix
$\mathcal{K}_1 = \{\sum_{p,q} K^p W K^q\}$	$[K^1 + K^2 + \dots + K^m] [W] [K^1 + K^2 + \dots + K^m]$
$\mathcal{K}_2 = \{\sum_{p,q} \mu_p \mu_q K^p W K^q\}$	$\begin{bmatrix} K^1 \\ K^2 \\ \vdots \\ K^m \end{bmatrix}^\top \begin{bmatrix} \mu_1^2 W & \mu_1 \mu_2 W & \dots & \mu_1 \mu_m W \\ \mu_2 \mu_1 W & \mu_2^2 W & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ \mu_m \mu_1 W & & & \mu_m^2 W \end{bmatrix} \begin{bmatrix} K^1 \\ K^2 \\ \vdots \\ K^m \end{bmatrix}$
$\mathcal{K}_3 = \{\sum_p \mu_p^2 K^p W K^p\}$	$\begin{bmatrix} K^1 \\ K^2 \\ \vdots \\ K^m \end{bmatrix}^\top \begin{bmatrix} \mu_1^2 W & 0 & \dots & 0 \\ 0 & \mu_2^2 W & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ 0 & & & \mu_m^2 W \end{bmatrix} \begin{bmatrix} K^1 \\ K^2 \\ \vdots \\ K^m \end{bmatrix}$
$\mathcal{K}_4 = \{\sum_p K^p W^p K^p\}$	$\begin{bmatrix} K^1 \\ K^2 \\ \vdots \\ K^m \end{bmatrix}^\top \begin{bmatrix} W^1 & 0 & \dots & 0 \\ 0 & W^2 & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ 0 & & & W^m \end{bmatrix} \begin{bmatrix} K^1 \\ K^2 \\ \vdots \\ K^m \end{bmatrix}$

Table 1: Block-matrix formulations of metric learning for multiple-kernel formulations (\mathcal{K}_1 – \mathcal{K}_4). Each W^p is taken to be positive semi-definite. Note that all sets are equal when there is only one base kernel.

Note that \mathcal{K}_4 contains \mathcal{K}_3 as a special case when all W^p are positive scalar multiples of each-other. However, \mathcal{K}_4 leads to a convex optimization problem, where \mathcal{K}_3 does not.

Table 1 lists the block-matrix formulations of each of the kernel combination rules described in this section. It is worth noting that it is certainly valid to first form the unweighted combination kernel \widehat{K} and then use \mathcal{K}_1 (Algorithm 3) to learn an optimal projection of the product space. However, as we will demonstrate in Section 5, our proposed multiple-kernel formulation (\mathcal{K}_4) outperforms the simple unweighted combination rule in practice.

4.4 Diagonal Learning

The MKPOE optimization is formulated as a semi-definite program over m different $n \times n$ matrices W^p —or, as shown in Table 1, a single $mn \times mn$ PSD matrix with a block-diagonal sparsity structure. Scaling this approach to large data sets can become problematic, as they require optimizing over multiple high-dimensional PSD matrices.

To cope with larger problems, the optimization problem can be refined to constrain each W^p to the set of diagonal matrices. If W^p are all diagonal, positive semi-definiteness is equivalent to non-negativity of the diagonal values (since they are also the eigenvalues of the matrix). This allows the constraints $W^p \succeq 0$ to be replaced by linear constraints $W_{ii}^p \geq 0$, and the resulting optimization problem is a linear program (LP), rather than an SDP. This modification reduces the flexibility of the model, but leads to a much more efficient optimization procedure.




Figure 3: Two variants of multiple-kernel embedding. (a) A data point $x \in \mathcal{X}$ is mapped into m feature spaces via $\phi^1, \phi^2, \dots, \phi^m$, which are then scaled by $\mu_1, \mu_2, \dots, \mu_m$ to form a weighted feature space \mathcal{H}^* , which is subsequently projected to the embedding space via M . (b) x is first mapped into each kernel’s feature space and then its image in each space is directly projected into a Euclidean space via the corresponding projections M^p . The projections are jointly optimized to produce the embedding space.

Algorithm 4 Multiple kernel partial order embedding (MKPOE)

Input: n objects \mathcal{X} ,
 partial order \mathcal{C} ,
 m kernel matrices K^1, K^2, \dots, K^m ,
 $\beta > 0$
Output: mapping $g : \mathcal{X} \rightarrow \mathbb{R}^{mn}$

$$\min_{W^p, \xi} \sum_{p=1}^m \text{tr}(W^p K^p) + \frac{\beta}{|\mathcal{C}|} \sum_{\mathcal{C}} \xi_{ijkl}$$

$$d(x_i, x_j) \doteq \sum_{p=1}^m (K_i^p - K_j^p)^\top W^p (K_i^p - K_j^p)$$

$$d(x_i, x_j) + 1 \leq d(x_k, x_\ell) + \xi_{ijkl}$$

$$\xi_{ijkl} \geq 0 \quad \forall (i, j, k, \ell) \in \mathcal{C}$$

$$W^p \succeq 0 \quad p = 1, 2, \dots, m$$

More specifically, our implementation of Algorithm 4 operates by alternating sub-gradient descent on W^p and projection onto the feasible set $W^p \succeq 0$ (see Appendix B for details). For full matrices, this projection is accomplished by computing the spectral decomposition of each W^p , and thresholding the eigenvalues at 0. For diagonal matrices, this projection is accomplished simply by

$$W_{ii}^p \mapsto \max \{0, W_{ii}^p\},$$

which can be computed in $O(mn)$ time, compared to the $O(mn^3)$ time required to compute m spectral decompositions.

Restricting W^p to be diagonal not only simplifies the problem to linear programming, but carries the added interpretation of weighting the contribution of each (kernel, training point) pair in the construction of the embedding. A large value at W_{ii}^p corresponds to point i being a landmark for the

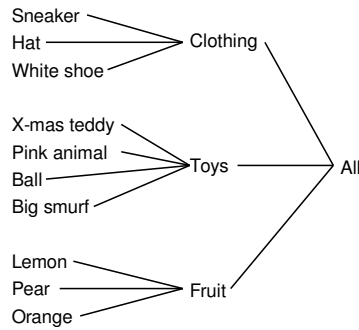


Figure 4: The label taxonomy for the experiment in Section 5.1.

features encoded in K^p . Note that each of the formulations listed in Table 1 has a corresponding diagonal variant, however, as in the full matrix case, only \mathcal{K}_1 and \mathcal{K}_4 lead to convex optimization problems.

5. Experiments

To evaluate our framework for learning multi-modal similarity, we first test the multiple kernel learning formulation on a simple toy taxonomy data set, and then on a real-world data set of musical perceptual similarity measurements.

5.1 Toy Experiment: Taxonomy Embedding

For our first experiment, we generated a toy data set from the Amsterdam Library of Object Images (ALOI) data set (Geusebroek et al., 2005). ALOI consists of RGB images of 1000 classes of objects against a black background. Each class corresponds to a single object, and examples are provided of the object under varying degrees of out-of-plane rotation.

In our experiment, we first selected 10 object classes, and from each class, sampled 20 examples. We then constructed an artificial taxonomy over the label set, as depicted in Figure 4. Using the taxonomy, we synthesized relative comparisons to span subtrees via their least common ancestor. For example,

$$\begin{aligned}
 & (Lemon\#1, Lemon\#2, Lemon\#1, Pear\#1), \\
 & (Lemon\#1, Pear\#, 1, Lemon\#1, Sneaker\#1),
 \end{aligned}$$

and so on. These comparisons are consistent and therefore can be represented as a directed acyclic graph. They are generated so as to avoid redundant, transitive edges in the graph.

For features, we generated five kernel matrices. The first is a simple linear kernel over the grayscale intensity values of the images, which, roughly speaking, compares objects by shape. The other four are Gaussian kernels over histograms in the (background-subtracted) red, green, blue, and intensity channels, and these kernels compare objects based on their color or intensity distributions.

We augment this set of kernels with five “noise” kernels, each of which was generated by sampling random points from the unit sphere in \mathbb{R}^3 and applying the linear kernel.

The data was partitioned into five 80/20 training and test set splits. To tune β , we further split the training set for 5-fold cross-validation, and swept over $\beta \in \{10^{-2}, 10^{-1}, \dots, 10^6\}$. For

each fold, we learned a diagonally-constrained embedding with Algorithm 4, using the subset of relative comparisons (i, j, k, ℓ) with i, j, k and ℓ restricted to the training set. After learning the embedding, the held out data (validation or test) was mapped into the space, and the accuracy of the embedding was determined by counting the fraction of correctly predicted relative comparisons. In the validation and test sets, comparisons were processed to only include comparisons of the form (i, j, i, k) where i belongs to the validation (or test) set, and j and k belong to the training set.

We repeat this experiment for each base kernel individually (that is, optimizing over \mathcal{K}_1 with a single base kernel), as well as the unweighted sum kernel (\mathcal{K}_1 with all base kernels), and finally MKPOE (\mathcal{K}_4 with all base kernels). The results are averaged over all training/test splits, and collected in Figure 5. For comparison purposes, we include the prediction accuracy achieved by computing distances in each kernel’s native space before learning. In each case, the optimized space indeed achieves higher accuracy than the corresponding native space. (Of course, the random noise kernels still predict randomly after optimization.)

As illustrated in Figure 5, taking the unweighted combination of kernels significantly degrades performance (relative to the best kernel) both in the native space (0.718 accuracy versus 0.862 for the linear kernel) and the optimized sum-kernel space (0.861 accuracy for the sum versus 0.951 for the linear kernel), that is, the unweighted sum kernel optimized by Algorithm 3. However, MKPOE (\mathcal{K}_4) correctly identifies and omits the random noise kernels by assigning them negligible weight, and achieves higher accuracy (0.984) than any of the single kernels (0.951 for the linear kernel, after learning).

5.2 Musical Artist Similarity

To test our framework on a real data set, we applied the MKPOE algorithm to the task of learning a similarity function between musical artists. The artist similarity problem is motivated by several real-world applications, including recommendation and playlist-generation for online radio. Because artists may be represented by a wide variety of different features (e.g., tags, acoustic features, social data), such applications can benefit greatly from an optimally integrated similarity metric.

The training data is derived from the *aset400* corpus of Ellis et al. (2002), which consists of 412 popular musicians, and 16385 relative comparisons of the form (i, j, i, k) . Relative comparisons were acquired from human test subjects through a web survey; subjects were presented with a query artist (i), and asked to choose what they believe to be the most similar artist (j) from a list of 10 candidates. From each single response, 9 relative comparisons are synthesized, indicating that j is more similar to i than the remaining 9 artists (k) which were not chosen.

Our experiments here replicate and extend previous work on this data set (McFee and Lanckriet, 2009a). In the remainder of this section, we will first give an overview of the various types of features used to characterize each artist in Section 5.2.1. We will then discuss the experimental procedure in more detail in Section 5.2.2. The MKL embedding results are presented in Section 5.2.3, and are followed by an experiment detailing the efficacy of our constraint graph processing approach in Section 5.2.4.

5.2.1 FEATURES

We construct five base kernels over the data, incorporating acoustic, semantic, and social views of the artists.




Figure 5: Mean test set accuracy for the experiment of Section 5.1. Error bars correspond to one standard deviation across folds. Accuracy is computed by counting the fraction of correctly predicted relative comparisons in the native space of each base kernel, and then in the optimized space produced by KPOE (\mathcal{K}_1 with a single base kernel). The unweighted combination of kernels (*Sum*) significantly degrades performance in both the native and optimized spaces. MKPOE (*MKL*, \mathcal{K}_4) correctly rejects the random kernels, and significantly outperforms the unweighted combination and the single best kernel.

- **MFCC**: for each artist, we collected between 1 and 10 songs (mean 4). For each song, we extracted a short clip consisting of 10000 half-overlapping 23ms windows. For each window, we computed the first 13 Mel Frequency Cepstral Coefficients (MFCCs) (Davis and Mermelstein, 1990), as well as their first and second instantaneous derivatives. This results in a sequence of 39-dimensional vectors (delta-MFCCs) for each song. Each artist i was then summarized by a Gaussian mixture model (GMM) p_i over delta-MFCCs extracted from the corresponding songs. Each GMM has 8 components and diagonal covariance matrices. Finally, the kernel between artists i and j is the probability product kernel (Jebara et al., 2004) between their corresponding delta-MFCC distributions p_i, p_j :

$$K_{ij}^{\text{mfcc}} = \int \sqrt{p_i(x)p_j(x)} dx.$$

- **Auto-tags (AT)**: Using the MFCC features described above, we applied the automatic tagging algorithm of Turnbull et al. (2008), which for each song yields a multinomial distribution over a set T of 149 musically-relevant tag words (*auto-tags*). Artist-level tag distributions q_i were formed by averaging model parameters (i.e., tag probabilities) across all of the songs of artist i . The kernel between artists i and j for auto-tags is a radial basis function applied to the

χ^2 -distance between the multinomial distributions q_i and q_j :

$$K_{ij}^{\text{at}} = \exp \left(-\sigma \sum_{t \in T} \frac{(q_i(t) - q_j(t))^2}{q_i(t) + q_j(t)} \right).$$

In these experiments, we fixed $\sigma = 256$.

- **Social tags (ST):** For each artist, we collected the top 100 most frequently used tag words from Last.fm,² a social music website which allows users to label songs or artists with arbitrary tag words or *social tags*. After stemming and stop-word removal, this results in a vocabulary of 7737 tag words. Each artist is then represented by a bag-of-words vector in \mathbb{R}^{7737} , and processed by TF-IDF. The kernel between artists for social tags is the cosine similarity (linear kernel) between TF-IDF vectors.
- **Biography (Bio):** Last.fm also provides textual descriptions of artists in the form of user-contributed biographies. We collected biographies for each artist in the *aset400* data set, and after stemming and stop-word removal, we arrived at a vocabulary of 16753 biography words. As with social tags, the kernel between artists is the cosine similarity between TF-IDF bag-of-words vectors.
- **Collaborative filtering (CF):** Celma (2008) collected collaborative filtering data from Last.fm in the form of a bipartite graph over users and artists, where each user is associated with the artists in her listening history. We filtered this data down to include only the *aset400* artists, of which all but 5 were found in the collaborative filtering graph. The resulting graph has 336527 users and 407 artists, and is equivalently represented by a binary matrix where each row i corresponds to an artist, and each column j corresponds to a user. The ij entry of this matrix is 1 if we observe a user-artist association, and 0 otherwise. The kernel between artists in this view is the cosine of the angle between corresponding rows in the matrix, which can be interpreted as counting the amount of overlap between the sets of users listening to each artist and normalizing for overall artist popularity. For the 5 artists not found in the graph, we fill in the corresponding rows and columns of the kernel matrix with the identity matrix.

5.2.2 EXPERIMENTAL PROCEDURE

The data was randomly partitioned into ten 90/10 training/test splits. Given the inherent ambiguity in the task, and format of the survey, there is a great deal of conflicting information in the survey responses. To obtain a more accurate and internally consistent set of training comparisons, directly contradictory comparisons (e.g., (i, j, i, k) and (i, k, i, j)) were removed from both the training and test sets. Each training set was further cleaned by finding an acyclic subset of comparisons and taking its transitive reduction, resulting in a minimal partial order. (No further processing was performed on test comparisons.)

After training, test artists were mapped into the learned space (by Equation 9), and accuracy was measured by counting the number of measurements (i, j, i, k) correctly predicted by distance in the learned space, where i belongs to the test set, and j, k belong to the training set.

For each experiment, β is chosen from $\{10^{-2}, 10^{-1}, \dots, 10^7\}$ by holding out 30% of the training constraints for validation. (Validation splits are generated from the unprocessed training set, and the

2. Last.fm can be found at <http://last.fm>.




Figure 6: aset400 embedding results for each of the base kernels. Accuracy is computed in each kernel’s native feature space, as well as the space produced by applying Algorithm 3 (i.e., optimizing over \mathcal{K}_1 with a single kernel) with either the diagonal or full-matrix formulation. Error bars correspond to one standard deviation across training/test splits.

remaining training constraints are processed as described above.) After finding the best-performing β , the embedding is trained on the full (processed) training set.

5.2.3 EMBEDDING RESULTS

For each base kernel, we evaluate the test-set performance in the native space (i.e., by distances calculated directly from the entries of the kernel matrix), and by learned metrics, both diagonal and full (optimizing over \mathcal{K}_1 with a single base kernel). Figure 6 illustrates the results.

We then repeated the experiment by examining different groupings of base kernels: acoustic (MFCC and Auto-tags), semantic (Social tags and Bio), social (Collaborative filter), and combinations of the groups. The different sets of kernels were combined by Algorithm 4 (optimizing over \mathcal{K}_4). The results are listed in Figure 7.

In all cases, MKPOE improves over the unweighted combination of base kernels. Moreover, many combinations outperform the single best kernel (ST, 0.777 ± 0.02 after optimization), and the algorithm is generally robust in the presence of poorly-performing kernels (MFCC and AT). Note that the poor performance of MFCC and AT kernels may be expected, as they derive from song-level rather than artist-level features, whereas ST provides high-level semantic descriptions which are generally more homogeneous across the songs of an artist, and Bio and CF are directly constructed at the artist level. For comparison purposes, we trained metrics on the sum kernel with \mathcal{K}_1 (Algorithm 3), resulting in accuracies of 0.676 ± 0.05 (diagonal) and 0.765 ± 0.03 (full). The proposed approach (Algorithm 4) applied to all kernels results in 0.754 ± 0.03 (diagonal), and 0.795 ± 0.02 (full).

Figure 8 illustrates the weights learned by Algorithm 4 using all five kernels and diagonally-constrained W^p matrices. Note that the learned metrics are both sparse (many 0 weights) and non-uniform across different kernels. In particular, the (lowest-performing) MFCC kernel is eliminated




Figure 7: aset400 embedding results with multiple kernel learning: the learned metrics are optimized over \mathcal{X}_d by Algorithm 4. *Native* corresponds to distances calculated according to the unweighted sum of base kernels.

by the algorithm, and the majority of the weight is assigned to the (highest-performing) social tag (ST) kernel.

A t-SNE (van der Maaten and Hinton, 2008) visualization of the space produced by MKPOE is illustrated in Figure 9. The embedding captures a great deal of high-level genre structure: for example, the *classic rock* and *metal* genres lie at the opposite end of the space from *pop* and *hip-hop*.

5.2.4 GRAPH PROCESSING RESULTS

To evaluate the effects of processing the constraint set for consistency and redundancy, we repeat the experiment of the previous section with different levels of processing applied to \mathcal{C} . Here, we focus on the Biography kernel, since it exhibits the largest gap in performance between the native and learned spaces.

As a baseline, we first consider the full set of similarity measurements as provided by human judgements, including all inconsistencies. To first deal with what appear to be the most egregious inconsistencies, we prune all directly inconsistent training measurements; that is, whenever (i, j, i, k) and (i, k, i, j) both appear, both are removed.³ Finally, we consider the fully processed case by finding a maximal consistent subset (partial order) of \mathcal{C} and removing all redundancies. Table 2 lists the number of training constraints retained by each step of processing (averaged over the random splits).

Using each of these variants of the training set, we test the embedding algorithm with both diagonal and full-matrix formulations. The results are presented in Table 2. Each level of graph processing results in a significant reduction in the number of training comparisons (and, therefore,

3. A more sophisticated approach could be used here, for example, majority voting, provided there is sufficient over-sampling of comparisons. The aset400 data lacks sufficient over-sampling for majority voting, so we default to this relatively simple approach.




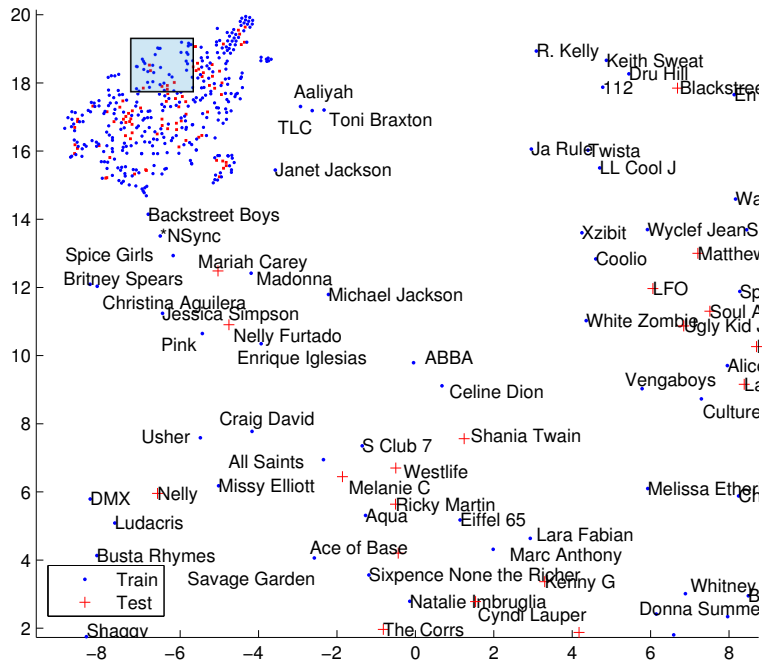
Figure 8: The weighting learned by Algorithm 4 using all five kernels and diagonal W^p . Each bar plot contains the diagonal of the corresponding kernel’s learned metric. The horizontal axis corresponds to the index in the training set, and the vertical axis corresponds to the learned weight in each kernel space.

\mathcal{C}	$ \mathcal{C} $ (Avg.)	Accuracy	
		Diagonal	Full
Full	8951.3	0.622±0.05	0.715±0.04
Length-2	6684.5	0.630±0.05	0.714±0.04
Processed	4814.5	0.628±0.05	0.716±0.04

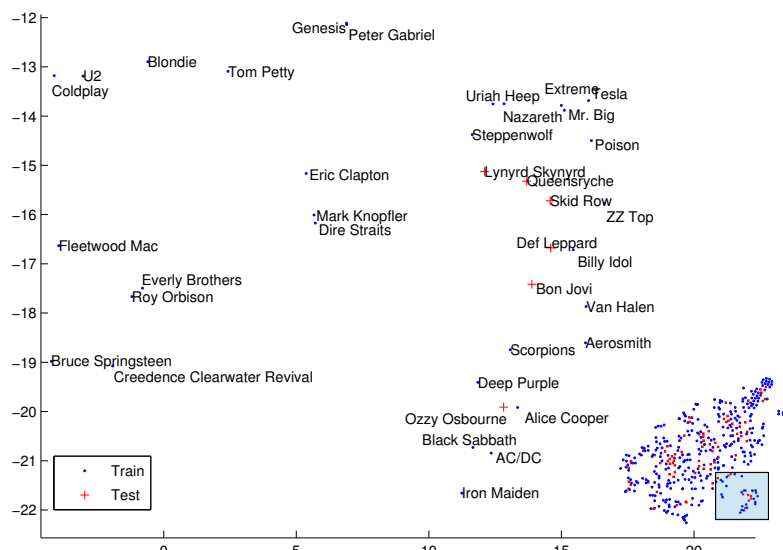
Table 2: aset400 embedding results (Biography kernel) for three possible refinements of the constraint set. *Full* includes all similarity measurements, with no pruning for consistency or redundancy. *Length-2* removes all length-2 cycles (i.e., (i, j, k, ℓ) and (k, ℓ, i, j)). *Processed* finds an approximate maximal consistent subset, and removes redundant constraints.

computational overhead of Algorithm 3), while not degrading the quality of the resulting embedding.

Finally, to test the sensitivity of the algorithm to randomness in the acyclic subgraph routine, we repeated the above experiment ten times, each with a different random maximal acyclic constraint set and the full matrix formulation of the algorithm. As depicted in Figure 10, the randomness in the constraint generation has little impact on the accuracy of the learned metric: the largest standard deviation is 0.007 (split #7).



(a)



(b)

Figure 9: t-SNE visualizations of an embedding of aset400 produced by MKPOE. The embedding is constructed by optimizing over \mathcal{K}_4 with all five base kernels. The two clusters shown roughly correspond to (a) pop/hip-hop, and (b) classic rock/metal genres. Out-of-sample points are indicated by a red +.

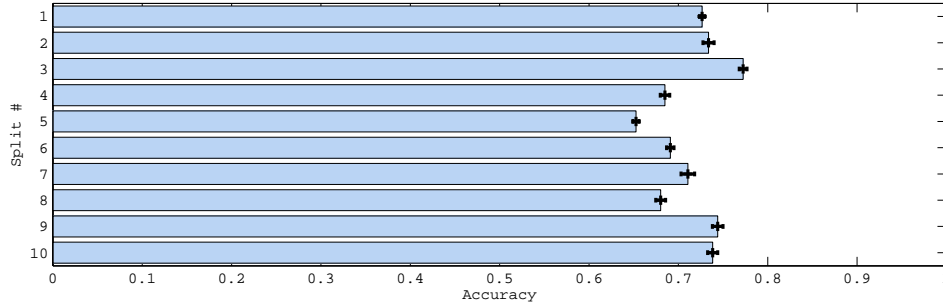


Figure 10: Accuracy of the learned embedding for each training/test split, averaged over ten trials with random maximal acyclic constraint subgraphs. Error bars correspond to one standard deviation.

6. Hardness of Dimensionality Reduction

The algorithms given in Sections 3 and 4 attempt to produce low-dimensional solutions by regularizing W , which can be seen as a convex approximation to the rank of the embedding. In general, because rank constraints are not convex, convex optimization techniques cannot efficiently minimize dimensionality. This does not necessarily imply other techniques could not work. So, it is natural to ask if exact solutions of minimal dimensionality can be found efficiently, particularly in the multidimensional scaling scenario, that is, when $K = I$ (Section 3.3).

As a special case, one may wonder if any instance $(\mathcal{X}, \mathcal{C})$ can be satisfied in \mathbb{R}^1 . As Figure 11 demonstrates, not all instances can be realized in one dimension. Even more, we show that it is NP-Complete to decide if a given \mathcal{C} can be satisfied in \mathbb{R}^1 . Given an embedding, it can be verified in polynomial time whether \mathcal{C} is satisfied or not by simply computing the distances between all pairs and checking each comparison in \mathcal{C} , so the decision problem is in NP. It remains to show that the \mathbb{R}^1 partial order embedding problem (hereafter referred to as *1-POE*) is NP-Hard. We reduce from the *Betweenness* problem (Opatrny, 1979), which is known to be NP-complete.

Definition 1 (Betweenness) *Given a finite set Z and a collection T of ordered triples (a, b, c) of distinct elements from Z , is there a one-to-one function $f : Z \rightarrow \mathbb{R}$ such that for each $(a, b, c) \in T$, either $f(a) < f(b) < f(c)$ or $f(c) < f(b) < f(a)$?*

Theorem 1 *1-POE is NP-Hard.*

Proof Let (Z, T) be an instance of *Betweenness*. Let $\mathcal{X} = Z$, and for each $(a, b, c) \in T$, introduce constraints (a, b, a, c) and (b, c, a, c) to \mathcal{C} . Since Euclidean distance in \mathbb{R}^1 is simply line distance, these constraints force $g(b)$ to lie between $g(a)$ and $g(c)$. Therefore, the original instance $(Z, T) \in \text{Betweenness}$ if and only if the new instance $(\mathcal{X}, \mathcal{C}) \in \text{1-POE}$. Since *Betweenness* is NP-Hard, *1-POE* is NP-Hard as well. ■

Since *1-POE* can be reduced to the general optimization problem of finding an embedding of minimal dimensionality, we can conclude that dimensionality reduction subject to partial order constraints is also NP-Hard.




Figure 11: (a) The vertices of a square in \mathbb{R}^2 . (b) The partial order over distances induced by the square: each side is less than each diagonal. This constraint set cannot be satisfied in \mathbb{R}^1 .

7. Conclusion

We have demonstrated a novel method for optimally integrating heterogeneous data to conform to measurements of perceptual similarity. By interpreting a collection of relative similarity comparisons as a directed graph over pairs, we are able to apply graph-theoretic techniques to isolate and prune inconsistencies in the training set and reduce computational overhead by eliminating redundant constraints in the optimization procedure.

Our multiple-kernel formulation offers a principled way to integrate multiple feature modalities into a unified similarity space. Our formulation carries the intuitive geometric interpretation of concatenated projections, and results in a semidefinite program. By incorporating diagonal constraints as well, we are able to reduce the computational complexity of the algorithm, and learn a model which is both flexible—only using kernels in the portions of the space where they are informative—and interpretable—each diagonal weight corresponds to the contribution to the optimized space due to a single point within a single feature space. Table 1 provides a unified perspective of multiple kernel learning formulations for embedding problems, but it is clearly not complete. It will be the subject of future work to explore and compare alternative generalizations and restrictions of the formulations presented here.

Acknowledgments

The authors acknowledge support from NSF Grant DMS-MSPA 0625409 and eHarmony, Inc.

Appendix A. Embedding Partial Orders

In this appendix, we prove that any set \mathcal{X} with a partial order over distances \mathcal{C} can be embedded into \mathbb{R}^n while satisfying all distance comparisons.

In the special case where \mathcal{C} is a total ordering over all pairs (i.e., a chain graph), the problem reduces to non-metric multidimensional scaling (Kruskal, 1964), and a constraint-satisfying embedding can always be found by the constant-shift embedding algorithm of Roth et al. (2003). In

Algorithm 5 Naïve total order construction

Input: objects \mathcal{X} , partial order \mathcal{C}
Output: symmetric dissimilarity matrix $\Delta \in \mathbb{R}^{n \times n}$

```

for each  $i$  in  $1 \dots n$  do
     $\Delta_{ii} \leftarrow 0$ 
end for
for each  $(k, \ell)$  in topological order do
    if  $\text{in-degree}(k, \ell) = 0$  then
         $\Delta_{k\ell}, \Delta_{\ell k} \leftarrow 1$ 
    else
         $\Delta_{k\ell}, \Delta_{\ell k} \leftarrow \max_{(i,j,k,\ell) \in \mathcal{C}} \Delta_{ij} + 1$ 
    end if
end for
    
```

general, \mathcal{C} is not a total order, but a \mathcal{C} -respecting embedding can always be produced by reducing the partial order to a (weak) total order by topologically sorting the graph (see Algorithm 5).

Let Δ be the dissimilarity matrix produced by Algorithm 5 on an instance $(\mathcal{X}, \mathcal{C})$. An embedding can be found by first applying classical multidimensional scaling (MDS) (Cox and Cox, 1994) to Δ :

$$A = -\frac{1}{2}H\Delta H,$$

where $H = I - \frac{1}{n}\mathbf{1}\mathbf{1}^\top$ is the $n \times n$ centering matrix, and $\mathbf{1}$ is a vector of 1s. Shifting the spectrum of A yields

$$A - \lambda_n(A)I = \widehat{A} \succeq 0,$$

where $\lambda_n(A)$ is the minimum eigenvalue of A . The embedding g can be found by decomposing $\widehat{A} = V\widehat{\Lambda}V^\top$, so that $g(x_i)$ is the i^{th} column of $\widehat{\Lambda}^{1/2}V^\top$; this is the solution constructed by the constant-shift embedding non-metric MDS algorithm of Roth et al. (2003).

Applying this transformation to A affects distances by

$$\begin{aligned} \|g(x_i) - g(x_j)\|^2 &= \widehat{A}_{ii} + \widehat{A}_{jj} - 2\widehat{A}_{ij} = (A_{ii} - \lambda_n) + (A_{jj} - \lambda_n) - 2A_{ij} \\ &= A_{ii} + A_{jj} - 2A_{ij} - 2\lambda_n. \end{aligned}$$

Since adding a constant $(-2\lambda_n)$ preserves the ordering of distances, the total order (and hence \mathcal{C}) is preserved by this transformation. Thus, for any instance $(\mathcal{X}, \mathcal{C})$, an embedding can be found in \mathbb{R}^{n-1} .

Appendix B. Solver

Our implementation of Algorithm 4 is based on a simple projected (sub)gradient descent. To simplify exposition, we show the derivation of the single-kernel SDP version of the algorithm (Algorithm 3) with unit margins. (It is straightforward to extend the derivation to the multiple-kernel and LP settings.)

We first observe that a kernel matrix column K_i can be expressed as $K^\top e_i$ where e_i is the i^{th} standard basis vector. We can then denote the distance calculations in terms of Frobenius inner products:

$$\begin{aligned} d(x_i, x_j) &= (K_i - K_j)^\top W (K_i - K_j) \\ &= (e_i - e_j)^\top K W K (e_i - e_j) \\ &= \text{tr}(K W K (e_i - e_j)(e_i - e_j)^\top) = \text{tr}(W K E_{ij} K) \\ &= \langle W, K E_{ij} K \rangle_F, \end{aligned}$$

where $E_{ij} = (e_i - e_j)(e_i - e_j)^\top$.

A margin constraint (i, j, k, ℓ) can now be expressed as:

$$\begin{aligned} & d(x_i, x_j) + 1 \leq d(x_k, x_\ell) + \xi_{ijkl} \\ \Rightarrow & \langle W, K E_{ij} K \rangle_F + 1 \leq \langle W, K E_{k\ell} K \rangle_F + \xi_{ijkl} \\ \Rightarrow & \xi_{ijkl} \geq 1 + \langle W, K (E_{ij} - E_{k\ell}) K \rangle_F. \end{aligned}$$

The slack variables ξ_{ijkl} can be eliminated from the program by rewriting the objective in terms of the constraints:

$$\min_{W \geq 0} f(W) \text{ where } f(W) = \text{tr}(W K) + \frac{\beta}{|\mathcal{C}|} \sum_{\mathcal{C}} h(1 + \langle W, K (E_{ij} - E_{k\ell}) K \rangle_F),$$

where

$$h(x) = \begin{cases} 0 & x \leq 0 \\ x & x > 0 \end{cases}$$

is the hinge loss.

The gradient ∇f has two components: one due to regularization, and one due to the hinge loss. The gradient due to regularization is simply K . The loss term decomposes linearly, and for each $(i, j, k, \ell) \in \mathcal{C}$, a sub-gradient direction can be defined:

$$\frac{\partial}{\partial W} h(1 + d(x_i, x_j) - d(x_k, x_\ell)) = \begin{cases} 0 & d(x_i, x_j) + 1 \leq d(x_k, x_\ell) \\ K(E_{ij} - E_{k\ell})K & \text{otherwise.} \end{cases} \quad (11)$$

Rather than computing each gradient direction independently, we observe that each violated constraint contributes a matrix of the form $K(E_{ij} - E_{k\ell})K$. By linearity, we can collect all $(E_{ij} - E_{k\ell})$ terms and then pre- and post-multiply by K to obtain a more efficient calculation of ∇f :

$$\frac{\partial}{\partial W} f = K + K \left(\sum_{(i,j,k,\ell) \in \bar{\mathcal{C}}} E_{ij} - E_{k\ell} \right) K,$$

where $\bar{\mathcal{C}}$ is the set of all currently violated constraints.

After each gradient step $W \mapsto W - \alpha \nabla f$, the updated W is projected back onto the set of positive semidefinite matrices by computing its spectral decomposition and thresholding the eigenvalues by $\lambda_i \mapsto \max(0, \lambda_i)$.

To extend this derivation to the multiple-kernel case (Algorithm 4), we can define

$$d(x_i, x_j) \doteq \sum_{p=1}^m d^p(x_i, x_j),$$

and exploit linearity to compute each partial derivative $\partial/\partial W^p$ independently.

For the diagonally-constrained case, it suffices to substitute

$$K(E_{ij} - E_{kl})K \mapsto \text{diag}(K(E_{ij} - E_{kl})K)$$

in Equation 11. After each gradient step in the diagonal case, the PSD constraint on W can be enforced by the projection $W_{ii} \mapsto \max(0, W_{ii})$.

Appendix C. Relationship to AUC

In this appendix, we formalize the connection between partial orders over distances and query-by-example ranking. Recall that Algorithm 2 minimizes the loss $1/|C| \sum_C \xi_{ijkl}$, where each $\xi_{ijkl} \geq 0$ is a slack variable associated with a margin constraint

$$d(i, j) + 1 \leq d(k, \ell) + \xi_{ijkl}.$$

As noted by Schultz and Joachims (2004), the fraction of relative comparisons satisfied by an embedding g is closely related to the area under the receiver operating characteristic curve (AUC). To make this connection precise, consider the following information retrieval problem. For each point $x_i \in \mathcal{X}$, we are given a partition of $\mathcal{X} \setminus \{i\}$:

$$\begin{aligned} \mathcal{X}_i^+ &= \{x_j : x_j \in \mathcal{X} \text{ relevant for } x_i\}, \text{ and} \\ \mathcal{X}_i^- &= \{x_k : x_k \in \mathcal{X} \text{ irrelevant for } x_i\}. \end{aligned}$$

If we embed each $x_i \in \mathcal{X}$ into a Euclidean space, we can then rank the rest of the data $\mathcal{X} \setminus \{x_i\}$ by increasing distance from x_i . Truncating this ranked list at the top τ elements (i.e., closest τ points to x_i) will return a certain fraction of relevant points (true positives), and irrelevant points (false positives). Averaging over all values of τ defines the familiar AUC score, which can be compactly expressed as:

$$\text{AUC}(x_i|g) = \frac{1}{|\mathcal{X}_i^+| \cdot |\mathcal{X}_i^-|} \sum_{(x_j, x_k) \in \mathcal{X}_i^+ \times \mathcal{X}_i^-} \mathbb{1}[\|g(x_i) - g(x_j)\| < \|g(x_i) - g(x_k)\|].$$

Intuitively, AUC can be interpreted as an average over all pairs $(x_j, x_k) \in \mathcal{X}_i^+ \times \mathcal{X}_i^-$ of the number of times x_i was mapped closer to a relevant point x_j than an irrelevant point x_k . This in turn can be conveniently expressed by a set of relative comparisons for each $x_i \in \mathcal{X}$:

$$\forall (x_j, x_k) \in \mathcal{X}_i^+ \times \mathcal{X}_i^- : (i, j, i, k).$$

An embedding which satisfies a complete set of constraints of this form will receive an AUC score of 1, since every relevant point must be closer to x_i than every irrelevant point.

Now, returning to the more general setting, we do not assume binary relevance scores or complete observations of relevance for all pairs of points. However, we can define the generalized

AUC score (GAUC) as simply the average number of correctly ordered pairs (equivalently, satisfied constraints) given a set of relative comparisons:

$$\text{GAUC}(g) = \frac{1}{|\mathcal{C}|} \sum_{(i,j,k,\ell) \in \mathcal{C}} \mathbb{1} [\|g(x_i) - g(x_j)\| < \|g(x_k) - g(x_\ell)\|].$$

Like AUC, GAUC is bounded between 0 and 1, and the two scores coincide exactly in the previously described ranking problem. A corresponding loss function can be defined by reversing the order of the inequality, that is,

$$L_{\text{GAUC}}(g) = \frac{1}{|\mathcal{C}|} \sum_{(i,j,k,\ell) \in \mathcal{C}} \mathbb{1} [\|g(x_i) - g(x_j)\| \geq \|g(x_k) - g(x_\ell)\|].$$

Note that L_{GAUC} takes the form of a sum over indicators, and can be interpreted as the average 0/1-loss over \mathcal{C} . This function is clearly not convex in g , and is therefore difficult to optimize. Algorithms 2, 3 and 4 instead optimize a convex upper bound on L_{GAUC} by replacing indicators by the hinge loss.

As in SVM, this is accomplished by introducing a unit margin and slack variable ξ_{ijkl} for each $(i, j, k, \ell) \in \mathcal{C}$, and minimizing $\frac{1}{|\mathcal{C}|} \sum_{\mathcal{C}} \xi_{ijkl}$.

References

- Sameer Agarwal, Joshua Wills, Lawrence Cayton, Gert Lanckriet, David Kriegman, and Serge Belongie. Generalized non-metric multi-dimensional scaling. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, 2007.
- Alfred V. Aho, Michael R. Garey, and Jeffrey D. Ullman. The transitive reduction of a directed graph. *SIAM Journal on Computing*, 1(2):131–137, 1972. doi: 10.1137/0201008. URL <http://link.aip.org/link/?SMJ/1/131/1>.
- Anelia Angelova. Data pruning. Master’s thesis, California Institute of Technology, 2004.
- Anelia Angelova, Yaser Abu-Mostafa, and Pietro Perona. Pruning training sets for learning of object categories. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 1: 494–501, 2005. ISSN 1063-6919. doi: <http://doi.ieeecomputersociety.org/10.1109/CVPR.2005.283>.
- Dana Angluin and Philip Laird. Learning from noisy examples. *Mach. Learn.*, 2(4):343–370, 1988. ISSN 0885-6125. doi: <http://dx.doi.org/10.1023/A:1022873112823>.
- Francis R. Bach. Consistency of the group lasso and multiple kernel learning. *J. Mach. Learn. Res.*, 9:1179–1225, 2008. ISSN 1532-4435.
- Yoshua Bengio, Jean-François Paiement, Pascal Vincent, Olivier Delalleau, Nicolas Le Roux, and Marie Ouimet. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. In Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.

- Bonnie Berger and Peter W. Shor. Approximation algorithms for the maximum acyclic subgraph problem. In *SODA '90: Proceedings of the First Annual ACM-SIAM Symposium on Discrete algorithms*, pages 236–243, Philadelphia, PA, USA, 1990. Society for Industrial and Applied Mathematics. ISBN 0-89871-251-3.
- Mikhail Bilenko, Sugato Basu, and Raymond J. Mooney. Integrating constraints and metric learning in semi-supervised clustering. In *Proceedings of the Twenty-first International Conference on Machine Learning*, pages 81–88, 2004.
- Ingwer Borg and Patrick J.F. Groenen. *Modern Multidimensional Scaling: Theory and Applications*. Springer-Verlag, second edition, 2005.
- Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- Oscar Celma. *Music Recommendation and Discovery in the Long Tail*. PhD thesis, Universitat Pompeu Fabra, Barcelona, Spain, 2008.
- Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Learning non-linear combinations of kernels. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 396–404. 2009.
- Trevor F. Cox and Michael A.A. Cox. *Multidimensional Scaling*. Chapman and Hall, 1994.
- Jason V. Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S. Dhillon. Information-theoretic metric learning. In *ICML '07: Proceedings of the 24th International Conference on Machine Learning*, pages 209–216, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-793-3. doi: <http://doi.acm.org/10.1145/1273496.1273523>.
- Steven B. Davis and Paul Mermelstein. Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. pages 65–74, 1990.
- Daniel P.W. Ellis, Brian Whitman, Adam Berenzweig, and Steve Lawrence. The quest for ground truth in musical artist similarity. In *Proceedings of the International Symposium on Music Information Retrieval (ISMIR2002)*, pages 170–177, October 2002.
- Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979. ISBN 0716710447.
- Jan-Mark Geusebroek, Gertjan J. Burghouts, and Arnold W. M. Smeulders. The Amsterdam library of object images. *Int. J. Comput. Vis.*, 61(1):103–112, 2005.
- Amir Globerson and Sam Roweis. Metric learning by collapsing classes. In Yair Weiss, Bernhard Schölkopf, and John Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 451–458, Cambridge, MA, 2006. MIT Press.
- Amir Globerson and Sam Roweis. Visualizing pairwise similarity via semidefinite embedding. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, 2007.

- Jacob Goldberger, Sam Roweis, Geoffrey Hinton, and Ruslan Salakhutdinov. Neighborhood components analysis. In Lawrence K. Saul, Yair Weiss, and Léon Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 513–520, Cambridge, MA, 2005. MIT Press.
- Saketha Nath Jagarlapudi, Dinesh G, Raman S, Chiranjib Bhattacharyya, Aharon Ben-Tal, and Ramakrishnan K.R. On the algorithmics and applications of a mixed-norm based kernel learning formulation. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 844–852. 2009.
- Tony Jebara, Risi Kondor, and Andrew Howard. Probability product kernels. *Journal of Machine Learning Research*, 5:819–844, 2004. ISSN 1533-7928.
- Maurice Kendall and Jean Dickinson Gibbons. *Rank correlation methods*. Oxford University Press, 1990.
- Marius Kloft, Ulf Brefeld, Soeren Sonnenburg, Pavel Laskov, Klaus-Robert Müller, and Alexander Zien. Efficient and accurate lp-norm multiple kernel learning. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 997–1005. 2009.
- Joseph B. Kruskal. Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29 (2), June 1964.
- Gert R. G. Lanckriet, Nello Cristianini, Peter Bartlett, Laurent El Ghaoui, and Michael I. Jordan. Learning the kernel matrix with semidefinite programming. *J. Mach. Learn. Res.*, 5:27–72, 2004. ISSN 1533-7928.
- Yen-Yu Lin, Tyng-Luh Liu, and Chiou-Shann Fuh. Dimensionality reduction for data in multiple feature representations. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 961–968. 2009.
- Brian McFee and Gert R. G. Lanckriet. Heterogeneous embedding for subjective artist similarity. In *Tenth International Symposium for Music Information Retrieval (ISMIR2009)*, October 2009a.
- Brian McFee and Gert R. G. Lanckriet. Partial order embedding with multiple kernels. In *Proceedings of the 26th International Conference on Machine Learning (ICML'09)*, pages 721–728, June 2009b.
- Jaroslav Opatrny. Total ordering problem. *SIAM J. Computing*, (8):111–114, 1979.
- Volker Roth, Julian Laub, Joachim M. Buhmann, and Klaus-Robert Müller. Going metric: denoising pairwise data. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 809–816, Cambridge, MA, 2003. MIT Press.
- Bernhard Schölkopf, Ralf Herbrich, Alex J. Smola, and Robert Williamson. A generalized representer theorem. In *Proceedings of the 14th Annual Conference on Computational Learning Theory*, pages 416–426, 2001.
- Matthew Schultz and Thorsten Joachims. Learning a distance metric from relative comparisons. In Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, Cambridge, MA, 2004. MIT Press.

- Noam Shental, Tomer Hertz, Daphna Weinshall, and Misha Pavel. Adjustment learning and relevant components analysis. In *European Conference on Computer Vision*, 2002.
- Sören Sonnenburg, Gunnar Rätsch, Christin Schäfer, and Bernhard Schölkopf. Large scale multiple kernel learning. *J. Mach. Learn. Res.*, 7:1531–1565, 2006. ISSN 1532-4435.
- Warren S. Torgerson. Multidimensional scaling: 1. theory and method. *Psychometrika*, 17:401–419, 1952.
- Douglas Turnbull, Luke Barrington, David Torres, and Gert Lanckriet. Semantic annotation and retrieval of music and sound effects. *IEEE Transactions on Audio, Speech and Language Processing*, 16(2):467–476, February 2008.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing high-dimensional data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- Alexander Vezhnevets and Olga Barinova. Avoiding boosting overfitting by removing confusing samples. In *ECML 2007*, pages 430–441, 2007.
- Kiri Wagstaff, Claire Cardie, Seth Rogers, and Stefan Schroedl. Constrained k-means clustering with background knowledge. In *Proceedings of the Eighteenth International Conference on Machine Learning*, pages 577–584, 2001.
- Kilian Q. Weinberger, John Blitzer, and Lawrence K. Saul. Distance metric learning for large margin nearest neighbor classification. In Yair Weiss, Bernhard Schölkopf, and John Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 451–458, Cambridge, MA, 2006. MIT Press.
- Eric P. Xing, Andrew Y. Ng, Michael I. Jordan, and Stuart Russell. Distance metric learning, with application to clustering with side-information. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 505–512, Cambridge, MA, 2003. MIT Press.
- Alexander Zien and Cheng Soon Ong. Multiclass multiple kernel learning. In *ICML '07: Proceedings of the 24th International Conference on Machine Learning*, pages 1191–1198, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-793-3. doi: <http://doi.acm.org/10.1145/1273496.1273646>.