2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

Motion Preference Learning

(Invited Paper)

Peter Kingston and Magnus Egerstedt
kingston@gatech.edu, magnus@ece.gatech.edu
Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

Abstract— **In order to control systems to meet subjective criteria, one would like to construct objective functions that accurately represent human preferences. To do this, we develop robust estimators based on convex optimization that, given empirical, pairwise comparisons between motions, produce both (1) objective functions that are compatible with the expressed preferences, and (2) global optimizers (i.e., "best motions") for these functions. The approach is demonstrated with an example in which human and synthetic motions are compared.**

I. Introduction

In control design, one typically first produces desired performance specifications and then tries to design the control laws in order to meet these specifications. However, if no such specifications are given, or even worse, if they are subjective, it is not clear what is meant by a successful control design.

One specific motivating application in which this issue arises is *robotic puppetry* (see e.g. [1], [2]), where marionettes with actuated strings are asked to perform expressive and aesthetically pleasing motions in the context of puppet plays. Given a particular human motion, how should the vastly-more-limited marionette move to not just mimic it but also communicate the same emotional intent? Similar questions arise whenever one wishes to control a system to meet aesthetic, artistic, or subjective goals – whether it be a puppet, an electronic musical instrument, or an acrobatic aircraft (to imagine just a few examples). Other examples include [3], [4], [5], [6], [7], [8], [9], and [10].

The role of human preferences in these problems is unavoidable, in that a motion is aesthetically pleasing only if we think it is pleasing. In this work we address techniques both for using empirical measurements to learn cost functions that are consistent with humans' aesthetic preferences, and for generalizing from these preference measurements to determine a globally best alternative. For the example of the marionette asked to mimic a human, this would mean finding the one marionette motion that best captures the subjective "essence" of a given human motion.

The idea of learning cost or rating functions from expressed preferences has been studied in the machine learning community, where it is known as *preference learning*. Here, it is assumed that one is given a collection of *pairwise comparisons* or *two-alternative forced choices* (i.e., for our case, answers to questions of the form "Which of these two options is better?"), and one seeks to generalize to other hypothetical comparisons, often by constructing a rating function from the choices. Pairwise comparison data has the advantage over numerical scores (e.g., a scale of 1 to 10) of being less prone to *batch effects*, a psychological phenomenon in which people's rankings are only accurate among objects compared at around the same time [11]. Other work sharing this philosophy includes [12], [13], [14], [15], [16], and [17].

Of the work on preference learning, ours is most closely related to [11], [12], and [13], which find rating functions that are *large margin* in a particular sense, by training a standard support-vector machine (SVM) classifier on a related set of instances. The method we propose – a Chebyshev estimation scheme – is similar in that it also employs a constrained optimization approach and in a particular limiting case can be reduced to an equivalent SVM classification problem. Our proposed approach differs from [11], [12], and [13] in that it aims (1) to find not only a rating function but also a globally best alternative, and (2) to solve only computationally-efficient convex programs; which in turn motivates different (and sometimes more efficient) problem formulations and solutions.

In the remainder of this paper, we give a precise problem formulation, and introduce two related optimization problems in detail in both *direct* and *instance vector expansion* forms. Then we show how in a particular limiting case a natural generalization is equivalent to a certain SVM classification problem, before demonstrating the approach with an example involving the comparison of human and synthetic motions.

II. Problem Formulation

At the core of preference learning is a collection of empirical, pairwise comparisons. The underlying assumption is that these comparisons reflect an underlying rating function. Hence, given a sequence of pairwise comparisons between points in an inner product space, we wish to find (1) a real-valued rating function that is consistent with those preferences, and (2) a global optimum to this function – the best point in the metric space. By solving these two problems

This work was supported by the U.S. National Science Foundation through Creative IT Grant #0757317. The human-study was conducted within the Georgia Institute of Technology, Institute Review Board Protocol H08162 - "Perceived Similarity Study."

we would have recovered what the underlying source for the comparisons is.

Formally, let $(X, \langle\cdot,\cdot\rangle)$ be the inner product space, and $S = \{(x_i^1, x_i^2)\}_{i=1}^N = \{s_1, \cdots, s_N\} \subset X \times X$ the sequence of comparisons; a pair (x_i^1, x_i^2) appears in the sequence S if and only if x_i^1 is preferred to x_i^2. The *preference graph* $\mathcal{G} = (V, S)$ corresponding to the comparison sequence S is the directed graph whose vertex set $V = \{x_1^1, x_1^2, \cdots, x_N^1, x_N^2\} \subset X$ is the collection of all unique points that have been compared, and whose edge set is S. We will index the vertices as $V = \{x_1, \cdots, x_M\}$, where $M \leq 2N$ is the cardinality of V.

The first item we seek, given some assumptions about its parametric form, is a function $f : X \to \mathbb{R}$ such that

$$f(x^1) \leq f(x^2) \Leftrightarrow (x^1, x^2) \in S \ . \tag{1}$$

That is, we adopt the convention that lower scores are better; hence we will refer to f as a *cost function*. Moreover, we would like f to optimize a particular smoothness criterion which we assume will make it generalize well to other points in X besides the ones we have seen in S.

The second item we seek is a global minimizer to f,

$$\bar{x} \triangleq \operatorname*{argmin}_x f(x) \tag{2}$$

which represents the best possible point in the inner product space.

Crucially, we would like to be able to determine f and $\bar{x}$ entirely by convex optimization – both so that the resulting problems are computationally efficient, and to ensure that any minima we find are in fact global optima. Although the SVM methodology employed in e.g. [13] and [11] finds f as the solution to a convex program, its use of the kernel trick introduces nonlinearities that prevent the determination of $\bar{x}$ by convex programming. Yet, without the kernel trick, and using the SVM approach, one arrives at linear cost functions that have no unique minima at all. What we will present in this paper is instead a set of convex programs that provide a useful compromise between these extremes, and which only reduce to an SVM classification problem in a particular limiting case. These will allow us to entertain the idea of a unique "best" point in X, and at the same time determine what it is by convex programming.

III. Metric Costs

Colloquially, when comparing various alternatives, we often speak of options as being "closer to what we would like," or of being "far from perfect." Motivated by this everyday use of geometric language, in [18] we considered *metric costs*, which have the form,

$$f(x) = ||x - \bar{x}||^2 \ . \tag{3}$$

In short, it is assumed that there exists some single best point $\bar{x}$ in X, and one alternative is preferred over another if and only if it is closer to that point.

What does an individual response (x^1, x^2) tell us about the location of $\bar{x}$? Simply, the following are equivalent:

1) $(x_i^1, x_i^2) \in S$
2) $f(x^1) \leq f(x^2)$
3) $\langle x_i^2 - x_i^1, \bar{x}\rangle - \frac{1}{2}\langle x_i^2 - x_i^1, x_i^2 + x_i^1\rangle < 0.$

In words, each comparison constrains $\bar{x}$ to lie within a particular halfspace of X. Defining,

$$d_i \triangleq x_i^2 - x_i^1 \tag{4}$$

$$\mu_i \triangleq \frac{1}{2}\left(x_i^1 + x_i^2\right) \tag{5}$$

$$b_i \triangleq \langle d_i, \mu_i\rangle \ , \tag{6}$$

the totality of what we know, then, about where $\bar{x}$ might lie is summarized by the inclusion over all the comparison halfspaces,

$$\bar{x} \in P \triangleq \bigcap_{i=1}^N \{x \mid \langle d_i, x\rangle - b_i < 0\} \ . \tag{7}$$

The set P, if it is bounded, is a polytope in X. In [18], we stated this system of inequalities and gave an asymptotic observer that converges to $\bar{x}$ under certain assumptions. Here, we ask another question: Out of all the points in this polytope, which is "best?"

When P is bounded, we propose to select $\bar{x}$ as the *incenter* or *Chebyshev center* of the polytope,

$$\bar{x} = \operatorname*{argmin}_x \max_i \frac{1}{||d_i||}\left(\langle d_i, x\rangle - b_i\right), \tag{8}$$

which is the point that is maximally far away from the closest constraint plane, as illustrated by Figure 1. In other words, when P is nonempty, $\bar{x}$ is the point that can be perturbed as much as possible without contradicting any of the preferences expressed in S; and when P is empty, it is the "compromise" point whose worst constraint violation is minimal.

In more detail, what Figure 1 portrays are two examples for the case when $X = \mathbb{R}^2$. Shades of gray indicate the number of violated constraints (points in darker regions violate more constraints), and discontinuities in the derivative of the piecewise-linear function $x \mapsto \max_i \frac{1}{||d_i||}\left(\langle d_i, x\rangle - b_i\right)$ are indicated by dashed lines. In the first example (top), $P \neq \emptyset$ (white region), and $\bar{x}$ is its incenter, the point maximally far away from the closest of the constraint surfaces (thin, solid lines) - i.e., it is the center of the largest inscribed sphere (thick, solid curve). In the second example (bottom), $P = \emptyset$, and the resulting optimum, $\bar{x}$, is the point whose worst constraint violation is minimal.

Note that with the definition (8), if the constraints are feasible (i.e., if $P \neq \emptyset$), then $\bar{x} \in P$. This can be viewed as minimizing the ∞-norm of the vector of constraints. Additionally, $\bar{x} \in \operatorname{aff}\left\{x_1^1, x_1^2, \cdots, x_N^1, x_N^2\right\}$ and hence we need only solve for the coefficients of an expansion in terms of this basis (see Theorem III.1). Furthermore, this minimization problem has a sensible solution even when P is empty; it is the point whose worst (i.e., largest) constraint violation is as small as possible.

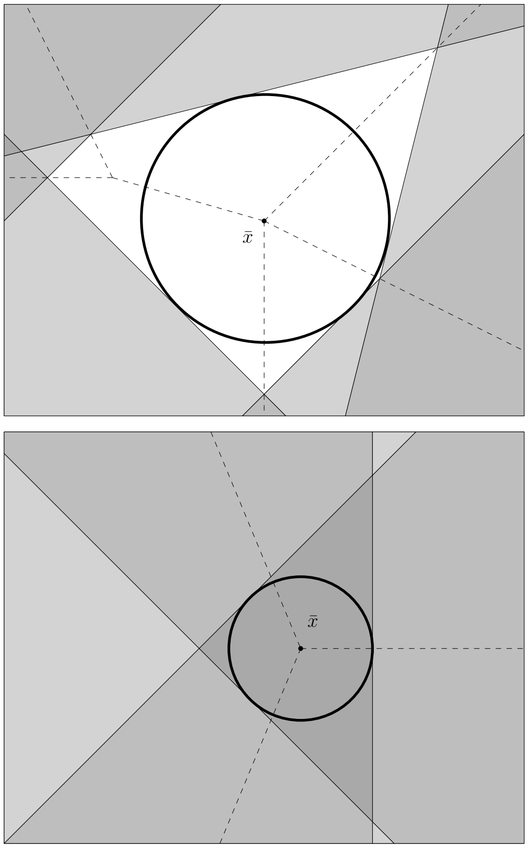

Fig. 1. The Chebyshev center of a polytope (top), and its natural generalization in the infeasible case (bottom).

The minimization problem (8) can be rewritten in epigraph form as,

$$\begin{aligned}(\bar{z}, \bar{x}) = &\operatorname*{argmin}_{(z,x)} z \\ \text{s.t.} \quad &||d_i|| z \geq \langle d_i, x\rangle - b_i\end{aligned} \tag{9}$$

which is always feasible (but possibly unbounded), and satisfies $\bar{z} > 0 \iff P = \emptyset$. This is a linear program, which, if $\dim(X)$ is not too large, can be solved directly by general-purpose codes. In order to work with very-large or infinite-dimensional vectors like motions, however, we will require the instance vector expansion form described in section III-A.

Theorem III.1 *If (8) has a global minimizer, then it has a global minimizer in* $\operatorname{aff}\left\{x_1^1, x_1^2, \cdots, x_N^1, x_N^2\right\}$.

Proof : Let x be a global minimum to (8), and $\bar{x}$ be the projection of x onto $\operatorname{aff}\left\{x_1^1, x_1^2, \cdots, x_N^1, x_N^2\right\}$; i.e., $\bar{x} = x + \delta$ with $\delta \perp \operatorname{span}\{d_1, \cdots, d_N\}$. Then for all $i \in \{1, \cdots, N\}$, since $\langle d_i, \delta\rangle = 0$ and by linearity of the inner product, $\frac{1}{||d_i||}\langle d_i, \bar{x}\rangle - b_i = \frac{1}{||d_i||}\langle d_i, x\rangle - b_i$, and hence the value of the objective function in (8) is the same at either x or $\bar{x}$. ■

A. Instance Vector Expansion

Since $\bar{x} \in \operatorname{aff}\left\{x_1^1, x_1^2, \cdots, x_N^1, x_N^2\right\}$, the optimization problem (8) can be solved as a finite-dimensional problem even when X is not finite-dimensional, by expanding $\bar{x}$ in terms of a finite-dimensional basis, as described by the following theorem:

Theorem III.2 *The point*

$$\bar{x} = \sum_{k=1}^{N} \bar{c}_k d_k + x^* \tag{10}$$

solves the optimization problem (8), where

$$x^* = \operatorname*{argmin}_x \left\{||x||^2 \mid x \in \operatorname{aff}\left\{x_1^1, x_1^2, \cdots, x_N^1, x_N^2\right\}\right\} , \tag{11}$$

and $\bar{c}$ is found by solving

$$\begin{aligned}(\bar{z}, \bar{c}) = &\operatorname*{argmin}_{(z,c)} z \\ \text{s.t.} \quad &G^{dd} c - Dz \leq \beta ,\end{aligned} \tag{12}$$

with $D = (||d_1||, \cdots, ||d_N||)$, $\beta \in \mathbb{R}^N$ defined by

$$\beta_i \triangleq \langle d_i, \mu_i\rangle \tag{13}$$

and $G^{dd} \in \mathbb{R}^{N\times N}$ being the Gramian,

$$G^{dd} \triangleq \begin{bmatrix} \langle d_1, d_1\rangle & \cdots & \langle d_1, d_N\rangle \\ \vdots & \ddots & \vdots \\ \langle d_N, d_1\rangle & \cdots & \langle d_N, d_N\rangle \end{bmatrix} . \tag{14}$$

Proof : Defining x^* by (11), one can write any x in the affine span of the data in the form (10). Substituting the expansion (10) into (9) and noting that by Hilbert's Projection Theorem $x^* \perp d_i$ for all $i \in \{1, \cdots, N\}$, one obtains (12). ■

Remark III.1 We also note at this point that (10) can be written,

$$x = \sum_{k=1}^{M} \left(\operatorname{indeg}_c(x_k) - \operatorname{outdeg}_c(x_k)\right) x_k + x^* \tag{15}$$

$$\triangleq \sum_{k=1}^{M} \xi_k x_k + x^* \tag{16}$$

by treating c as a vector of edge weights to the preference graph, and denoting the weighted in- and out-degrees of a given node x_k by $\operatorname{indeg}_c(x_k)$ and $\operatorname{outdeg}_c(x_k)$ respectively. Precisely,

$$\operatorname{indeg}_c(x_k) \triangleq \sum_{i \mid x_i^2 = x_k} c_i \tag{17}$$

$$\operatorname{outdeg}_c(x_k) \triangleq \sum_{i \mid x_i^1 = x_k} c_i . \tag{18}$$

Remark III.2 Moreover, β can be written,

$$\beta_i = e_i^T G^{\mu d} e_i , \tag{19}$$

where $G^{\mu d} \in \mathbb{R}^{N\times N}$ is the cross-Gramian

$$G^{\mu d} \triangleq \begin{bmatrix} \langle d_1, \mu_1\rangle & \cdots & \langle d_1, \mu_N\rangle \\ \vdots & \ddots & \vdots \\ \langle d_N, \mu_1\rangle & \cdots & \langle d_N, \mu_N\rangle \end{bmatrix} \tag{20}$$

and e_i denotes the i-th element of the natural basis.

Remark III.3 Note that this problem depends only on inner products of the various d_i and u_i vectors, and hence the problem can be solved even when X is infinite-dimensional. Precisely, $\frac{N(N+1)}{2} + N^2 \sim O(N^2)$ inner products must be computed to build the matrices G^{dd} and $G^{\mu d}$, where N is the number of comparisons. Alternatively, the relevant matrices can also be produced directly from inner products of elements of S, as

$$G^{dd} = K^{22} - K^{21} - K^{12} + K^{11} \tag{21}$$

$$G^{\mu d} = \frac{1}{2}(K^{22} + K^{21} - K^{12} - K^{11}) \tag{22}$$

where each matrix $K^{lm} \in \mathbb{R}^{N\times N}$ is defined by

$$K^{lm}_{ij} = \langle x^l_i, x^m_j\rangle \tag{23}$$

and can be built by indexing into the single Gramian (or kernel) matrix $K \in \mathbb{R}^{M\times M}$ defined by

$$K_{ij} = \langle x_i, x_j\rangle \ . \tag{24}$$

Moreover, D in equation (12) can now be expressed as $D = (\sqrt{G^{dd}_{11}}, \sqrt{G^{dd}_{22}}, \sqrt{G^{dd}_{33}}, \cdots, \sqrt{G^{dd}_{NN}})$.

Finally, $\bar{x}$ can be reconstructed using (10) and

$$x^* = \sum_{i=1}^{M} \alpha_i x_i \tag{25}$$

$$\alpha = \frac{1}{\mathbf{1}^T K^\dagger \mathbf{1}} K^\dagger \mathbf{1} \tag{26}$$

where $K^\dagger$ denotes the Moore-Penrose pseudoinverse of K, and $\mathbf{1} = (1, 1, \cdots, 1) \in \mathbb{R}^M$.

In particular, the costs of the presented instances can be reconstructed as,

$$f(x_k) = (e_k - \xi - \alpha)^T K (e_k - \xi - \alpha) \tag{27}$$

where ξ is related to c by (15), (17), and (18).

B. Unbounded Case: The minimax-rate problem

When P is nonempty but unbounded, the minimization problem (8) is ill-posed; one can choose x with arbitrarily large norm, and in the process make $\max_i \frac{1}{||d_i||}(\langle d_i, x\rangle - b_i)$ arbitrarily small. Practically, this corresponds to the situation in which people would prefer some aspect of a motion not in a particular amount, but rather to the largest (or smallest) degree possible. Hence in the case of unbounded P we ask a slightly different question: What is the "point at infinity," or *direction*, that is best? More precisely, what we seek in this case is a unit vector

$$\bar{v} = \operatorname*{argmin}_{v\in X | ||v||=1} \lim_{t\to\infty} \frac{1}{t}\left[\max_i \frac{1}{||d_i||}\left(\langle d_i, tv\rangle - b_i\right)\right] \tag{28}$$

$$= \operatorname*{argmin}_{v\in X | ||v||=1} \max_i \frac{1}{||d_i||}\langle d_i, v\rangle \tag{29}$$

or equivalently,

$$(\bar{p}, \bar{v}) = \operatorname*{argmin}_{v\in X, p\in\mathbb{R}} p \tag{30}$$
$$\text{s.t.} \quad \left\{ \begin{array}{lcl} ||d_i|| p & \geq & \langle d_i, v\rangle \quad \forall i \in \{1, \cdots, N\} \\ ||v||^2 & \leq & 1 \end{array} \right. .$$

The interpretation is that, although there exists no best *point* $\bar{x} \in X$ because constraints can be satisfied arbitrarily well, there does exist a *direction* $v \in X$ along which constraint satisfaction improves most rapidly.

As before, an instance vector expansion is possible:

Theorem III.3 *Letting $v = \sum_{k=1}^{N} c_k d_k$, the optimization problem (30) is equivalent to*

$$(\bar{p}, \bar{c}) = \operatorname*{argmin}_{(p,c)} p \tag{31}$$
$$\text{s.t.} \quad \left\{ \begin{array}{rcl} G^{dd}c - Dp & \leq & 0 \\ c^T G^{dd} c & \leq & 1 \end{array} \right.$$

with the matrices G^{dd} and D as defined in the previous subsection.

Proof : The proof takes the form of Theorem III.2's. ∎

The problem (31) is a finite-dimensional second-order cone program (SOCP), which can be solved efficiently.

The cost function for the unbounded case arises from a similar limit process to (29), as

$$f(x) = \lim_{t\to\infty} \left(\frac{1}{t}||x - vt||^2 - t\right) \tag{32}$$

$$= \lim_{t\to\infty} \left[\frac{1}{t}\left(||x||^2 - 2\langle x, vt\rangle + ||vt||^2\right) - t\right] \tag{33}$$

$$= -2\langle x, v\rangle \tag{34}$$

which can be evaluated at the instances as,

$$f(x_k) = -2 e_k^T K \xi \ . \tag{35}$$

1) QP Form and Relation to SVMs: When $\operatorname{int} P$ is nonempty - as is almost always true in the unbounded case - the minimization problem (29) can be rewritten as an equivalent quadratic program (QP), which will make the relationship to the usual SVM approach very clear. In fact, (29) is equivalent to a particular SVM classification problem (which differs from but is related to that studied in e.g. [11] and [13]).

Defining,

$$w = \frac{1}{p} v \tag{36}$$

and restricting our attention to negative values for p (since when $\operatorname{int} P$ is nonempty, $p^* < 0$), we note that for $p < 0$

$$\operatorname{argmin} p = \operatorname{argmax} p^2 = \operatorname{argmin} \frac{1}{p^2} = \operatorname{argmin} ||w||^2 \ . \tag{37}$$

Additionally, the constraints in (30) can be replaced by,

$$\left\langle \frac{d_i}{||d_i||}, w \right\rangle \geq 1 \tag{38}$$

which results in the standard unbiased SVM problem,

$$\bar{w} = \underset{w}{\operatorname{argmin}} ||w||^2$$
$$\text{s.t.} \quad \left\langle \frac{d_i}{||d_i||}, w \right\rangle \geq 1 \quad \forall i \in \{1, \cdots, N\} \ . \tag{39}$$

This is equivalent to (30) in the unbounded case except when $\operatorname{int} P = \emptyset$; then, since $\bar{p} = 0$, $\bar{w}$ from (36) is undefined, but the solution to the SOCP problem (31) nevertheless exists.

The *minimax-rate* problem (39) differs from the SVM problem considered in e.g. [11] and [13] by the factor of $\frac{1}{||d_i||}$ included in each constraint. The difference is that whereas the standard SVM approach attempts to classify differences using a maximum-margin separating hyperplane, the minimax-rate approach finds the direction that maximizes the rate of constraint satisfaction.

This is illustrated in Figure 2. Here, a number of uniformly-randomly selected points in $[-1, 1] \times [-1, 1] \subset \mathbb{R}^2$ are compared according to a point at infinity (i.e., a linear cost function) (dotted), and both the traditional SVM (dashed) and the minimax-rate (solid) approaches are used to produce estimates of this direction from the comparisons. From the difference-classification point of view (top), one wishes to separate the vectors $\{d_i\}_{i=1}^N$ (displayed as "o"s) from the vectors $\{-d_i\}_{i=1}^N$ (displayed as "*"s). From the minimax-rate point of view (bottom), one wishes to find the direction that maximizes the rate of constraint satisfaction (the numbers of violated constraints are represented by shades of gray; the white region is feasible). The traditional SVM solution separates the positive from the negative differences with a larger margin (top), but the minimax-rate solution stays as far from the edge of the constraint cone as possible (bottom).

IV. Amoebas and Humans

To understand the comparison of higher-dimensional objects, and, in particular, *motions*, an experiment was performed in which an audience of 25 people was asked to perform pairwise comparisons of different motions of a computer-animated amoeba, relative to the motion-captured movement of a human who danced the bhangra. An example of one such question is illustrated in Figure 4. In this manner, a preference graph was generated as before, with 12 vertices (the amoeba motions) and 20 edges; this is shown in Figure 3.

Inner products between the various amoeba motions were computed by rasterizing the motions to binary videos, blurring each frame of the result, and computing the standard Euclidean inner product of these (extremely large) `[Frame Width]`×`[Frame Height]`×`[Number of Frames]`-dimensional vectors. We note that the sheer size of this representation highlights the advantage of the instance vector expansion described in Section III-A, without which the optimization problem simply could not be realistically solved.

The minimization problem (12) with the resulting data turns out to be unbounded and hence we again find an optimal direction via (10). We obtain the coefficient expansion

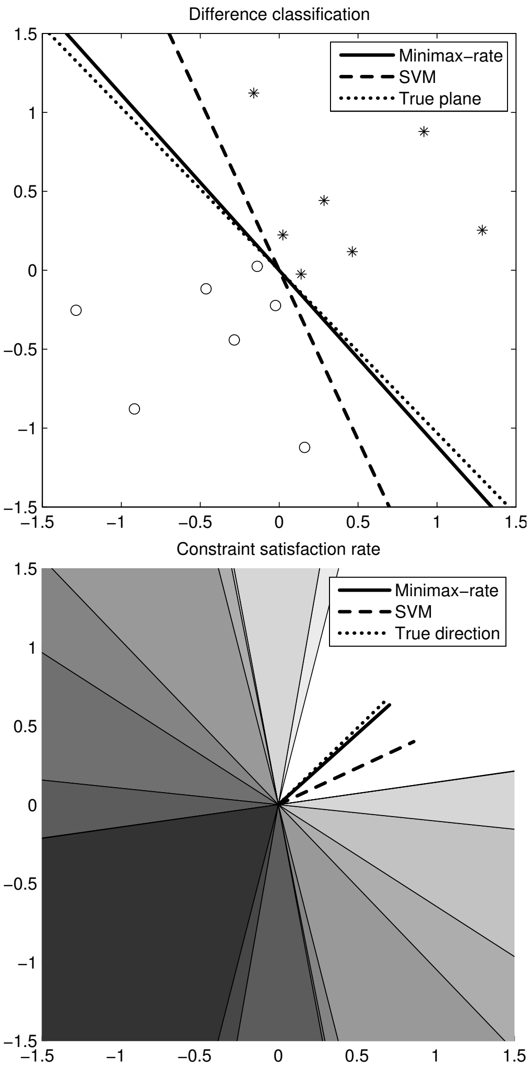

Fig. 2. Maximum-margin vs. minimax-rate cost functions.

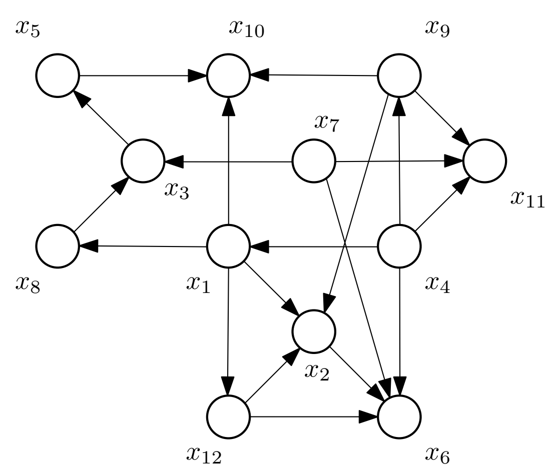

Fig. 3. The preference graph corresponding to the amoeba experiments.

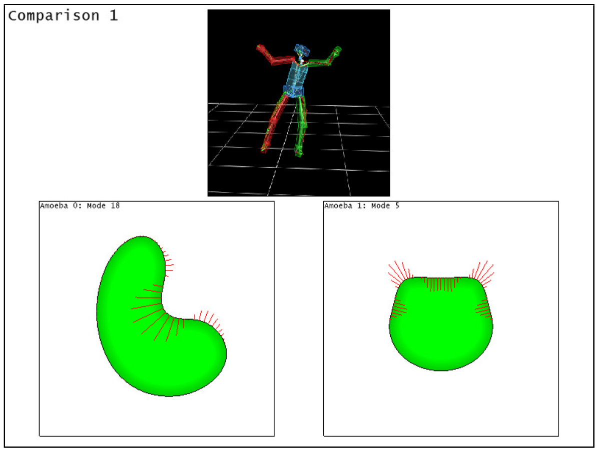

Fig. 4. Each question took the form, "Which of the two 'amoebas' (bottom) looks more like the [motion capture data from a] human dancer (top)?"

for the optimal direction,

$$\bar{v} = \sum_{k=1}^{M} \xi_k x_k \tag{40}$$

where

$$\begin{aligned}\xi = 10^3(\ & 1.4918, & -3.6556, && -0.1390, && 0.3113, \\ & -1.1243, & -0.1771, && 2.6335, && 0.5878, \\ & 1.8362, & -1.7319, && -0.2999, && 0.2672\ ) .\end{aligned}$$

What this means is that, in order to look as much like it is dancing the bhangra as possible, an amoeba should as its first priority aspire to be as much like amoeba 7 ($\xi_7 = 2.6335$) and as dissimilar from amoeba 2 ($\xi_2 = -3.6556$) as possible, and that it should to a lesser extent model itself after amoebas 1 and 9 ($\xi_1 = 1.4918, \xi_9 = 1.8362$) while avoiding the aesthetically unappealing moves of amoebas 5 and 10 ($\xi_5 = -1.1243, \xi_{10} = -1.7319$). Note here that, although this does not explain why, psychologically, e.g. amoeba 7 is preferred to amoeba 2 – i.e., we do not obtain a collection of characteristics that can be decoupled from the motions of the empirical amoeba moves – it does produce both a consistent cost structure, and an estimate for an amoeba motion that will be preferred to all others in the larger space of motions.

V. Concluding Remarks

In this work, we investigate the problem of motion preference learning under the assumption of an underlying *metric cost model*; here, the alternatives being compared are points in a metric space, and human judges are assumed to prefer one point to another if and only if it is closer to some fixed but unknown best alternative that they may not have been shown. This assumption appears to be a good one for the example considered and the features chosen, in that the feasible set P in this case is nonempty.

Based on the metric cost assumption, a Chebyshev estimator was given for the best point for the case when P is bounded, and a natural generalization, the minimax-rate estimator, was developed for when P is unbounded. In the first case, the solution was found, with an efficiency rivaling standard quadratic SVMs, as the solution to a linear program; and in the second case the problem was shown to in fact reduce to a particular SVM classification problem.

In order that the estimators for the bounded and unbounded cases be applicable to situations in which the compared alternatives inhabit high- or infinite- dimensional metric spaces – as is the case for motion signals – the optimization problems were additionally given in an instance vector expansion form, which results in optimization problems whose size is proportional not to the dimensionality of the metric space, but only to the number of comparisons available.

In all cases, optimal cost functions *and* points/directions were found efficiently by convex programming. The result is an efficient minimax estimator for the best possible alternative.

References

[1] M. Egerstedt, T. Murphey, and J. Ludwig, *Motion Programs for Puppet Choreography and Control*, pp. 190–202, Springer-Verlag, April 2007.

[2] Elliot Johnson and Todd Murphey, "Dynamic modeling and motion planning for marionettes: Rigid bodies articulated by massless strings," in *Robotics and Automation, 2007 IEEE International Conference on*, 2007, pp. 330–335.

[3] Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl, "Evaluating collaborative filtering recommender systems," *ACM Trans. Inf. Syst.*, vol. 22, no. 1, pp. 5–53, 2004.

[4] Gediminas Adomavicius and Alexander Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Trans. on Knowl. and Data Eng.*, vol. 17, no. 6, pp. 734–749, 2005.

[5] Miquel Montaner, Beatriz López, and Josep Lluís De La Rosa, "A taxonomy of recommender agents on the internet," *Artif. Intell. Rev.*, vol. 19, no. 4, pp. 285–330, 2003.

[6] Louis Leon Thurstone, *The measurement of values*, University of Chicago Press, 1959.

[7] Ralph Allan Bradley and Milton E. Terry, "Rank analysis of incomplete block designs: The method of paired comparisons," vol. 39, pp. 324–345, 1952.

[8] S. N. Afriat, "The construction of utility functions from expenditure data," vol. 8, no. 1, pp. 67–77, Feb 1967.

[9] H. R. Varian, "The nonparametric approach to demand analysis," vol. 50, no. 4, pp. 945–973, Jul 1982.

[10] A. Restificar and P. Haddawy, "Inferring implicit preferences from negotiation actions," in *Proc. Int'l Symposium on Artificial Intelligence and Mathematics*, Jan 2004.

[11] Jorge Díez, Juan José del Coz, Oscar Luaces, and Antonio Bahamonde, "Clustering people according to their preference criteria," *Expert Syst. Appl.*, vol. 34, no. 2, pp. 1274–1284, 2008.

[12] Ralf Herbrich, Thore Graepel, Peter Bollmann-Sdorra, and Klaus Obermayer, "Supervised learning of preference relations," in *Proceedings FGML-98, German National Workshop on Machine Learning*, 1998, pp. 43–47.

[13] Ralf Herbrich, Thore Graepel, and Klaus Obermayer, "Large margin rank boundaries for ordinal regression," in *Advances in large margin classifiers*, pp. 115–132. MIT Press, 2000.

[14] W. W. Cohen, R. E. Schapire, and Y. Singer, "Learning to order things," *Journal of Artificial Intelligence Research*, vol. 10, 1999.

[15] E. Hüllermeier, J. Frnkranz, W. Cheng, and K. Brinker, "Label ranking by learning pairwise preferences," *Artificial Intelligence*, vol. 172, pp. 1897–1917, 2008.

[16] Shivani Agarwal, "Ranking on graph data," in *Proceedings of the 23rd international conference on Machine learning*, 2006, ICML '06.

[17] Wei Chu and Zoubin Ghahramani, "Preference learning with gaussian processes," in *ICML '05: Proceedings of the 22nd international conference on Machine learning*, New York, NY, USA, 2005, pp. 137–144, ACM.

[18] Peter Kingston and Magnus Egerstedt, "Comparing apples and oranges through partial orders: An empirical approach," in *American Control Conference*, 2009.