Comparing Apples and Oranges Through Partial Orders: An Empirical Approach

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 this paper, we try to understand what people mean when they say that two objects are "similar." This is an important question in the area of human-robot interactions, where robots must interpret human movements in order to act in a "similar" manner. Specifically, we assume that we are given a collection of empirically generated pairwise comparisons between a subset of so-called alternatives (members of a given set), which produces a partial order over the set of alternatives. Based on this partial order, an inverse optimization problem is solved, producing a cost associated with each alternative that is consistent with the partial order. This cost is, moreover, assumed to be generative in that it can be used to select the globally best alternative. An experimental study involving the comparison of apples and oranges is presented to highlight the operation of the proposed approach.**

I. INTRODUCTION

As the saying goes, one can not compare apples and oranges. But why not? It is clear that some apples look more like oranges than others. One can thus ask the question "What makes apple X look more like an orange than apple Y?" Or, more interestingly (yet also more absurdly), "If apple X is in fact a robot apple, how should it act in order to make it more like an orange?" These questions may seem like nonsense, but this is exactly what mobile robots are asked to do in some areas of human-robot interactions, e.g. Programming by Demonstration, where humans (oranges) ask robots (apples) to behave "like" them, e.g. [1],[2], [3], [4].

The basic idea behind the Programming by Demonstration paradigm is that the human operator should be able to instruct robots (typically manipulator arms) how to act without having to write code or use any other type of formal interface language. Instead, the operator should be able to "show" the robot how to act. But, since the robot typically has a completely different set of dynamical constraints, degrees of freedom, and even spatial scales, it is not at all clear what it should be doing in response to the human operator. For more on these issues, see for example [5],[6],[7],[8],[9].

In this paper, we address this seemingly ill-posed problem by formulating a version of it in such a way that it is amenable to analytical solution, while still being based on subjective judgments of "similarity." The reason for insisting on the subjective element is that, at the end of the day, "similarity" is to be understood as "what people consider to be similar." We should note, already at this point, that a related idea is pursued in [10], [11]. But, in those references, the basic premise is not that one is given a collection of comparisons, but rather a metrically ranked collection of examples. What *is* inherently similar between those works and the work presented in this paper is that the empirical data is to be used to find an underlying cost function that it is postulated people use when making judgments about the similarity of objects and motions.

It should also be noted that in econometrics, the problem of inferring people's *utility functions* from data is well-studied. What is different with this problem is that the objective typically is to either understand game theoretic decision making strategies in a market or stochastic setting (e.g., [12], [13], [14]), or to locate clusters of consumers with similar preferences (e.g., [15], [16]). This paper is really about finding deterministic "similarity" measures that fit the empirical data, and, as such, it has an entirely different focus and objective.

II. PROBLEM FORMULATION

A. Experimentally Generated Comparisons

Let $\mathcal{A}$ be the set of all possible objects ("apples," or robot actions) that we would like to compare relative to some ideal ("the orange," or a human action) – most generally, we will refer to these as *alternatives*. To begin, we will assume very little about $\mathcal{A}$; it may not even be countable. If we are interested in programming a particular robot by demonstration, $\mathcal{A}$ might be the set of all possible motions that the robot can perform. If we are comparing actual apples to oranges in order to find the most "orangelike" apple, then $\mathcal{A}$ would be the set of all conceivable apples. To capture the subjectivity of comparing elements of this set to the ideal, suppose that we can conduct experiments, in each of which we ask human observers to perform a pairwise comparison of elements of $\mathcal{A}$; specifically, assume that we can pose questions of the form, "Which of these apples, X or Y, is more orangelike?" This particular form of experiment, it should be noted, has the advantage over other forms (e.g., rankings on a scale from 1-10) that it is less prone to *batch effects*, a psychological phenomenon in which people's rankings are only accurate among objects compared at around the same time [15], [17].

We will assume for our purposes that the human observers we ask to compare alternatives are "similar," in that they tend to have the same opinions about which alternatives are more like the ideal, so we can treat their responses meaningfully in an aggregate sense (i.e., we will not be concerned with Condorcet's paradox).

Given this assumption, one can think of the group of humans as a generalized *comparator*, or sorting function, which it will be our goal to "learn" from output measurements. More specifically, we assume that the human observers are presented with an unsorted pair of alternatives as "input," and that they "output" a corresponding sorted pair, ordered by similarity to the ideal. That is, the humans in these experiments behave as a map h from the set $\mathcal{U} = \mathcal{A} \times \mathcal{A}$ of questions, to the set $\mathcal{Y} = \mathcal{A} \times \mathcal{A}$ of answers, where h is defined,

$$h(a_a, a_b) = \begin{cases} (a_a, a_b) & a_a \succ a_b \\ (a_b, a_a) & a_b \succ a_a \end{cases} \tag{1}$$

and where we use "$a_a \succ a_b$" to denote "a_a is more 'orange-like' than a_b." From this point of view, an "experiment" – in which we present alternatives to human observers and obtain similarity judgments – is really just a function evaluation of h.

Now let $\mathcal{E} \subset \mathcal{A}$ be a finite, indexed subset of alternatives that we, the experimenters, will actually present to the human observers; we will refer to these as the *experimental alternatives*. The reason for defining this set is that, ideally, we would like to be able to present only a few alternatives (i.e., those in the set $\mathcal{E}$), in order to draw conclusions about preferences over *all* possible alternatives (everything in $\mathcal{A}$ – which may not even be countable). Additionally, let $\mathcal{P} = \{1, 2, \ldots, P\}$ be the index set associated with $\mathcal{E}$; we define this for convenience so that we may refer to "the i^{th} alternative," which we will abbreviate $\mathcal{E}_i \in \mathcal{E}$.

Next, suppose that we ask the human observers to perform a series of pairwise comparisons of alternatives. That is, we present elements of $\mathcal{E}$ two-at-a-time to human observers in an indexed set of questions $\mathcal{Q} = \{u^1, u^2, \ldots\} \subset \mathcal{E} \times \mathcal{E} \subset \mathcal{U}$ to obtain the indexed set of responses $\mathcal{R} = \{y^1, y^2, \ldots\} = \{h(u^1), h(u^2), \ldots\} \subset \mathcal{E} \times \mathcal{E} \subset \mathcal{Y}$. In other words, thinking of the humans again as a comparator, we input a sequence of unsorted pairs, and receive a corresponding sequence of sorted pairs as output; from this, we hope to deduce the inner workings of the comparator.

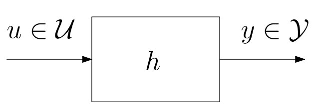

Fig. 1. Human as comparator – a memoryless nonlinear system

We will also assume that a vector of feature scores $\varphi(a) \in \mathcal{F} = \mathbb{R}^q$ associated with each alternative $a \in \mathcal{A}$ is available automatically from experimental data. If we are comparing actual apples to oranges, such features might include the average RGB colors of the apples, their dimensions along various directions, and similar measurements. (In fact, we do exactly this, as will be described later in the paper.) Moreover, for the special case when we are evaluating the features of an alternative in the experimental set $\mathcal{E}$, we will for compactness of notation write $\varphi(\mathcal{E}_k) = \varphi_k$.

Now, suppose that we can form a parametrized cost, $J : \mathbb{R}^N \times \mathcal{F} \to \mathbb{R}$ that, given some parameter $\rho \in \mathbb{R}^N$, maps the features (an element of $\mathcal{F}$) of any alternative in $\mathcal{A}$ to a real number, and that

$$J(\rho, \varphi(a_a)) < J(\rho, \varphi(a_b)) \Longleftrightarrow a_a \succ a_b. \tag{2}$$

In other words, we assume that we can define a cost such that "cheaper" alternatives are more like the ideal. What this means is that, given J, ρ, and the features corresponding to two "apples," we know exactly which of the two is more "orangelike." To make this ability to compare alternatives using J explicit, we define an associated output map $h_J : \mathcal{U} \to \mathcal{Y}$, that sorts pairs of alternatives in this fashion; i.e.,

$$h_J(a_a, a_b) = \begin{cases} (a_a, a_b) & J(\varphi(a_a)) > J(\varphi(a_b)) \\ (a_b, a_a) & J(\varphi(a_b)) > J(\varphi(a_a)) \end{cases} .$$

The significance of h_J is that it is the comparator function consistent with (2). Our goal, then, is to determine a ρ (for fixed J) and hence a cost function such that $h_J(u) = h(u)\ \forall u \in \mathcal{U}$.

In other words (and for compactness of notation defining $J(\rho, \varphi(\mathcal{E}_k)) = J_k(\rho)$ for the special case when we are evaluating the cost of an alternative in the experimental set) the problem we are trying to solve is that of selecting the parameter ρ such that,

1) The pairwise comparisons are consistent with the costs, i.e. $J_i(\rho) < J_j(\rho),\ \forall (\mathcal{E}_i, \mathcal{E}_j) \in \mathcal{R}$.
2) ρ satisfies some feasibility constraint $\pi(\rho) = 0$.

We let $\Omega(\mathcal{E})$ denote the set of all such feasible ρ parameters and given that at least one feasible ρ exists, we want moreover to select $\rho \in \Omega(\mathcal{E})$ in such a way that it minimizes the smallest of all the alternative costs.

Summarizing these points, what we want to achieve is to solve the min-min problem

$$\min_{i \in \mathcal{P}} \left\{ \min_{\rho \in \mathbb{R}^N} J_i(\rho) \right\}$$

subject to the constraint

$$\rho \in \Omega(\mathcal{E}).$$

Before we can actually solve this problem, we first need to establish some necessary conditions for the existence of a solution associated with ensuring that the pairwise comparisons are rational in the sense that they induce a partial order on the alternatives. For this, we need to introduce the notion of a *directed alternative graph* $\mathcal{G}_{\mathcal{E}} = (\mathcal{E}, \mathcal{R})$, where the vertex set is equal the presented alternatives, and a directed edge between $\mathcal{E}_i$ and $\mathcal{E}_j$ exists if and only if there exists a $y^k \in \mathcal{R}$ such that $(\mathcal{E}_i, \mathcal{E}_j) = y^k$. In other words, each edge encodes a judgment about which of the vertices (alternatives) adjacent to it is "more orangelike."

Now, in order to ensure that we have indeed a partial order, i.e. that we can not end up with situations where

$$\mathcal{E}_i \succ \mathcal{E}_j,\ \mathcal{E}_j \succ \mathcal{E}_k,\ \mathcal{E}_k \succ \mathcal{E}_i,$$

we have to assume that $\mathcal{G}_{\mathcal{E}}$ is acyclic. Assuming that this is indeed the case, and that the feasible set $\Omega(\mathcal{E})$ is non-empty, then the min-min problem can in fact be solved by solving a total of at most $o(\mathcal{E}) \le P$ problems, where

$$o(\mathcal{E}) = \text{card}(\mathcal{O}(\mathcal{E})),$$

and $\mathcal{O}(\mathcal{E})$ is the set of all alternatives that won at least one comparison while not losing any comparison, i.e.

$$\mathcal{O}(\mathcal{E}) = \left\{ i \in \mathcal{P} \,\middle|\, \exists \mathcal{E}_j \text{ s.t. } \begin{array}{c} (\mathcal{E}_i, \mathcal{E}_j) \in \mathcal{R} \\ \text{and} \\ \nexists \mathcal{E}_k \text{ s.t. } (\mathcal{E}_k, \mathcal{E}_i) \in \mathcal{R} \end{array} \right\}.$$

Using the terminology from graph theory, what these nodes thus satisfy is that they have out-degree greater than zero and in-degree equal to zero.

B. The Transitive Reduction

In fact, for many graphs, the number of constraints when solving each of these subproblems can be reduced; i.e., we can remove edges from the graph, and thereby reduce the execution time of the optimization algorithm. For instance, consider the graph $\mathcal{G}_3$ given in Figure 2. For this graph, edge

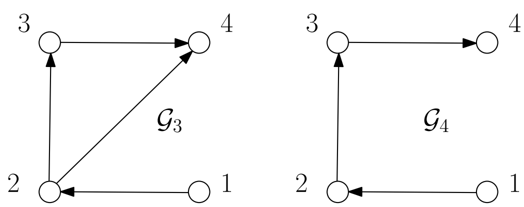

Fig. 2. Two equivalent alternative graphs.

$(2, 4)$ imposes the constraint that $J_2(\varphi) < J_4(\varphi)$, yet since $(2, 3)$ imposes $J_2(\varphi) < J_3(\varphi)$ and $(3, 4)$ imposes $J_3(\varphi) < J_4(\varphi)$, then by transitivity $(2, 3)$ and $(3, 4)$ collectively render $(2, 4)$ redundant, and hence the alternative graph $\mathcal{G}_3$ can be replaced by $\mathcal{G}_4$. That is to say, if we optimize the parametrized cost subject to all the constraints represented by $\mathcal{G}_4$, then all of the constraints represented by $\mathcal{G}_3$ will automatically be satisfied. From a graph-theoretic point of view, $\mathcal{G}_4$ is the *transitive reduction* of $\mathcal{G}_3$.

Formally, using Aho's definition [18], $\mathcal{G}^t$ is the transitive reduction of a graph $\mathcal{G}$ if,

1) there is a directed path from vertex u to vertex v in $\mathcal{G}^t$ if and only if there is a directed path from u to v in $\mathcal{G}$, and
2) there is no graph with fewer arcs than $\mathcal{G}^t$ satisfying condition 1.

In the case of a directed acyclic graph, the reduction $\mathcal{G}^t$ (which is unique) is a subgraph of $\mathcal{G}$. It was shown in [18] that computation of the transitive reduction is of the same complexity as transitive closure, and hence matrix multiplication; thus, the transitive reduction can be found in $O(n^{log_2 7})$ steps using Strassen's algorithm [19]. (See, e.g., [20], [21]).

III. Cost Models

In the following sections, we will present two different, related examples of cost functions, and investigate the implications of each choice.

A. Linear Cost Models

As an example, consider a situation in which the alternative costs are linear, i.e. $J_i = \rho^T \varphi(\mathcal{E}_i)$ and all feature vectors φ are non-negative. In that case, the min-min problem becomes

$$\min_{i \in \mathcal{P}} \left\{ \min_{\rho \in \mathbb{R}^N} \rho^T \varphi(\mathcal{E}_i) \right\},$$

subject to the constraints

$$\begin{cases} \rho^T \varphi(\mathcal{E}_i) \le \rho^T \varphi(\mathcal{E}_j),\ \forall (\mathcal{E}_i, \mathcal{E}_j) \in \mathcal{R} \\ \mathbf{1}^T \rho = 1 \\ \rho \ge 0, \end{cases}$$

where $\mathbf{1} = (1, \ldots, 1)^T$, and where we, for simplicity have assumed that $N = q$, i.e. the number of parameters (the dimension of ρ) is equal to the number of features (the dimension of φ). We moreover replaced the pairwise comparison constraints with non-strict inequalities.

We directly note that since $\rho \ge 0$, the notion of dominance allows us to reduce the number of constraints and possibly also $o(\mathcal{E})$ in the case when the problem is linear. In particular, an alternative $\mathcal{E}_i$ is said to linearly dominate alternative $\mathcal{E}_j$ if $(\mathcal{E}_i, \mathcal{E}_j) \in \mathcal{R}$ and $\varphi_i \le \varphi_j$, where the inequality is taken componentwise. And, $\rho \ge 0$ directly implies that if this is indeed the case then $\rho^T \varphi_i \le \rho^T \varphi_j$ and as such this constraint can be removed from the problem altogether. That is, the structure imposed by our choice of linear cost function allows for additional simplifications beyond those implied by transitivity alone.

B. Metric Cost Models

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, we would now like to consider *metric costs* of the form,

$$J_i = d(\varphi(\mathcal{E}_i), \varphi(\bar{a})) \tag{3}$$

where $\bar{a} \in \mathcal{A}$ is the "ideal" or "most orangelike" apple – which is unknown to us, the experimenters – and $d(\cdot\,,\cdot)$ is a metric in the inner product space $\mathcal{F}$ (We will assume the usual Euclidean metric and inner product, but what follows is readily generalizable to other inner products.) In this case, J is entirely parametrized by $\rho = \varphi(\bar{a})$, so the goal will be to determine this ideal feature vector from responses.

What does an individual response $y = (\mathcal{E}_1, \mathcal{E}_2)$ tell us about the location of $\varphi(\bar{a})$?[1] Simply,

$$\varphi(\bar{a}) \in \left\{\varphi \mid n^T \varphi \geq b\right\} \Longleftrightarrow d(\varphi_2, \varphi(\bar{a})) \geq d(\varphi_1, \varphi(\bar{a})). \quad (4)$$

where $\varphi_1 = \varphi(\mathcal{E}_1)$, $\varphi_2 = \varphi(\mathcal{E}_2)$, $n = (\varphi_1 - \varphi_2)$, and $b = \frac{1}{2} n^T (\varphi_1 + \varphi_2)$. (This follows immediately from Lemma 3.1, which is given at the end of this section.) Hence, a sequence of n outputs $y^1, y^2, \ldots, y^n$ (with $y^k = (\mathcal{E}_1^k, \mathcal{E}_2^k), \varphi_j^k = \mathcal{E}_j^k \ \forall k = 1, 2, \ldots, n; \ j = 1, 2$) implies,

$$\bar{\varphi} \in \bigcap_{k=1}^{n} \left\{\varphi \mid (\varphi_1^k - \varphi_2^k)^T \varphi > \frac{1}{2}(\varphi_1^k - \varphi_2^k)^T (\varphi_1^k + \varphi_2^k)\right\} \quad (5)$$

or equivalently, $\bar{\varphi}$ is a solution to

$$\begin{bmatrix} (\varphi_1^1 - \varphi_2^1)^T \\ (\varphi_1^2 - \varphi_2^2)^T \\ \vdots \\ (\varphi_1^n - \varphi_2^n)^T \end{bmatrix} \bar{\varphi} > \frac{1}{2} \begin{bmatrix} (\varphi_1^1 - \varphi_2^1)^T (\varphi_1^1 + \varphi_2^1) \\ (\varphi_1^2 - \varphi_2^2)^T (\varphi_1^2 + \varphi_2^2) \\ \vdots \\ (\varphi_1^n - \varphi_2^n)^T (\varphi_1^n + \varphi_2^n) \end{bmatrix} \quad (6)$$

where "$>$" indicates componentwise inequality.

The geometric interpretation of (4) is that $\bar{\varphi}$ must lie within a half-plane in feature space. Likewise, (5) means that $\bar{\varphi}$ must lie within the intersection of the half-planes; this is a polytope in $\mathcal{F}$.

Before continuing, we now state the Lemma referred to earlier in this section; its statement is given more generally than (4). The geometric interpretation is that comparisons between distances relative to reference points can be interchanged with signed point-plane distance tests.

Lemma 3.1: Let φ_1, φ_2, $\bar{\varphi}$ be any vectors in the inner product space $\mathbb{R}^m$ for some $m \in \mathbb{N}$ (with the usual inner product), and let $\bigstar$ be a binary relation from the set, $\{=, <, >, \leq, \geq\}$. Then,

$$\bar{\varphi} \in \left\{\varphi \mid n^T \varphi \bigstar b\right\} \Longleftrightarrow d(\varphi_2, \bar{\varphi}) \bigstar d(\varphi_1, \bar{\varphi})$$

where $n = (\varphi_1 - \varphi_2)$, and $b = \frac{1}{2} n^T (\varphi_1 + \varphi_2)$.

The proof of this is based on the Polarization Identity and is straightforward.

1) An asymptotic observer for metric cost models: Suppose we have access to a very long (infinite) sequence of comparisons $y^0, y^1, y^2, \ldots \in \mathcal{Y}$, perhaps as the result of passive monitoring over an extended period of time, and we would like to know the features $\bar{\varphi}$ of the ideal alternative. If alternatives are presented at random to the comparator, can we construct an asymptotic observer for $\bar{\varphi}$ which can avoid storing all of the very (infinitely) many constraints implied by this sequence? It turns out that the answer is yes, and exactly such an observer is given by,

$$\tilde{\varphi}^{k+1} = \begin{cases} P^k \tilde{\varphi}^k + \frac{\alpha^k b^k}{(n^k)^T (n^k)} n^k & \text{if } (n^k)^T \tilde{\varphi}^k < b^k \\ \tilde{\varphi}^k & \text{otherwise} \end{cases} \quad (7)$$

$$P^k = I - \alpha^k \frac{(n^k)(n^k)^T}{(n^k)^T (n^k)} \quad (8)$$

$$n^k = (\varphi_1^k - \varphi_2^k) \quad (9)$$

$$b^k = \frac{1}{2} n^T (\varphi_1^k + \varphi_2^k) \quad (10)$$

for any sequence of observer gains $\alpha^k \in (0, 2)$, regardless of $\tilde{\varphi}^0$. That is, $\tilde{\varphi}$ converges to $\bar{\varphi}$ in probability as $k \to \infty$, given a few assumptions; we will prove this shortly in Theorem 3.1. Moreover, note that, although (7-10) are broken down into separate expressions for clarity of presentation, they are in fact all functions of $\tilde{\varphi}^k$, so this observer can be implemented with only $\dim\{\mathcal{F}\}$ real memory elements.

Geometrically, the observer (7-10) operates through a series of projections (or under/over-projections, if $\alpha \neq 1$), as illustrated in Figure 3, with each projection bringing the estimate $\tilde{\varphi}^k$ of the ideal closer to the true ideal, $\bar{\varphi}$. A proof of convergence follows as Theorem 3.1.

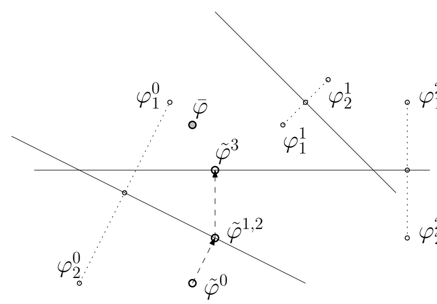

Fig. 3. A series of the observer's estimates, with $\alpha^k = 1 \ \forall k$. The initial estimate is $\tilde{\varphi}^0$, and the true ideal is given by $\bar{\varphi}$. In step 0, the observer projects $\tilde{\varphi}^0$ onto the plane (solid line) corresponding to the measured output $y^0 = (\varphi_1^0, \varphi_2^0)$ to produce $\tilde{\varphi}^1$. In step 1, the observer makes no changes to its estimate, because $\tilde{\varphi}^1$ is on the correct side of the plane corresponding to y^1; hence $\tilde{\varphi}^2 = \tilde{\varphi}^1$. In step 2, the observer projects $\tilde{\varphi}^2$ onto the plane corresponding to y^2 to create the estimate $\tilde{\varphi}^3$, which is yet closer to $\bar{\varphi}$.

Theorem 3.1: Let $u^k = (a_a^k, a_b^k)$ be a sequence of random alternatives issued as input to a comparator system with metric cost function as defined in (3), such that the features $\varphi_a^k, \varphi_b^k \in \mathcal{F}$ of these alternatives are i.i.d. random variables drawn according to any probability density function $p(\varphi)$ which is nonzero in an open ball $B(\bar{\varphi}, r) = B_r$ around the optimal alternative, $\bar{\varphi}$. Then, the asymptotic observer given by (7) converges to $\bar{\varphi}$ in probability.

Proof :

1. *If $(n^k)^T \tilde{\varphi}^k > b^k$, then $d(\tilde{\varphi}^{k+1}, \bar{\varphi}) < d(\tilde{\varphi}^k, \bar{\varphi})$.* The distances $d(\tilde{\varphi}^k, \bar{\varphi})$ and $d(\tilde{\varphi}^{k+1}, \bar{\varphi})$ are related through the Polarization Identity by (where $\Delta^k = \tilde{\varphi}^{k+1} - \tilde{\varphi}^k$),

$$||\tilde{\varphi}^{k+1} - \bar{\varphi}||^2 = ||\tilde{\varphi}^k + \Delta^k - \bar{\varphi}||^2 = \\ ||\tilde{\varphi}^k - \bar{\varphi}||^2 + ||\Delta^k||^2 + 2(\tilde{\varphi}^k - \bar{\varphi})^T \Delta^k$$

so, it order to show that $||\tilde{\varphi}^{k+1} - \bar{\varphi}|| < ||\tilde{\varphi}^k - \bar{\varphi}||$, it is sufficient to demonstrate

$$||\Delta^k||^2 + 2(\tilde{\varphi}^k - \bar{\varphi})^T \Delta^k < 0. \quad (11)$$

[1] In this section, the subscript 1 and 2 are used to denote the first and second elements of y (i.e., the more- and less- ideal alternatives), rather than the particular elements of $\mathcal{E}$ indexed by 1 and 2.

From (7, 8),

$$\begin{aligned}\Delta^k &= \left(I - a^k \frac{(n^k)(n^k)^T}{(n^k)^T(n^k)}\right)\tilde{\varphi}^k + \frac{\alpha^k b^k}{(n^k)^T(n^k)} n^k - \tilde{\varphi}^k \\ &= \frac{\alpha}{(n^k)^T(n^k)}\left(b^k - (n^k)^T\tilde{\varphi}^k\right) n^k \end{aligned} \tag{12}$$

so, substituting Δ into (11) (and dropping the superscript indices k),

$$\frac{\alpha^2}{n^T n}(b - n^T\tilde{\varphi})^2 + 2\frac{\alpha}{n^T n}(b - n^T\tilde{\varphi})n^T(\tilde{\varphi} - \bar{\varphi}) < 0 \tag{13}$$

or equivalently, so long as $\alpha > 0$ (as we require),

$$\left(b - n^T\tilde{\varphi}\right)\left[\alpha\left(b - n^T\tilde{\varphi}\right) + 2n^T\left(\tilde{\varphi} - \bar{\varphi}\right)\right] \quad < \quad 0. \tag{14}$$

Since by assumption $n^T\tilde{\varphi} < b$, this is satisfied iff the second factor is negative; that is,

$$\begin{aligned}\alpha\left(b - n^T\tilde{\varphi}\right) + 2n^T\left(\tilde{\varphi} - \bar{\varphi}\right) \quad &= \\ \alpha b + (2-\alpha)n^T\tilde{\varphi} - 2n^T\bar{\varphi} \quad &< \quad 0.\end{aligned} \tag{15}$$

or equivalently

$$\frac{1}{2}\alpha b + \left(1 - \frac{1}{2}\alpha\right) n^T\tilde{\varphi} < n^T\bar{\varphi}. \tag{16}$$

Since $n^T\tilde{\varphi} < b$, and by Lemma 3.1, $n^T\bar{\varphi} \geq b$, this is satisfied so long as $\alpha \in (0, 2)$, as we require.

2. *The sequence* $d^k = ||\tilde{\varphi}^k - \bar{\varphi}^k||,\ k = 0, 1, 2, \ldots$ *is nonincreasing.* In the second case of (7), $\tilde{\varphi}^{k+1} = \tilde{\varphi}^k$; this is nonincreasing. In the first case, $(n^k)^T\tilde{\varphi}^k > b^k$, so $d^{k+1} < d^k$ by point 1 above.

3. $\text{g.l.b.}(d^k) = 0$ *with unit probability.* By positivity of $d(\cdot, \cdot)$, zero is a lower bound. To show that this is the greatest such bound, consider some $\epsilon > 0$ and suppose that, at iteration m, $d(\tilde{\varphi}^m, \bar{\varphi}) = \epsilon$. Now, let $z = \min(r, \epsilon/2)$, and consider the open balls $B_1 = B(c_1, z/4), B_2 = B(c_2, z/4)$, where the center points c_1, c_2 are defined,

$$c_j = \bar{\varphi} + \frac{\tilde{\varphi} - \bar{\varphi}}{||\tilde{\varphi} - \bar{\varphi}||}\frac{(2j-1)}{4}z;$$

additionally, let $\varphi_1 \in B_1, \varphi_2 \in B_2$. Then by Lemma 3.1, we can confirm that $\bar{\varphi}$ and $\tilde{\varphi}$ are on opposite sides of the plane (and hence, that a projection will occur) by verifying that,

$$||\varphi_2 - \tilde{\varphi}|| < ||\varphi_1 - \tilde{\varphi}|| \tag{17}$$

$$||\varphi_2 - \bar{\varphi}|| > ||\varphi_1 - \bar{\varphi}||. \tag{18}$$

Considering the first of these, we note by the triangle inequality,

$$||\varphi_2 - \tilde{\varphi}|| \leq ||\varphi_2 - c_2|| + ||c_2 - \tilde{\varphi}|| < \tfrac{1}{4}z + ||c_2 - \tilde{\varphi}||$$

whereas, by the inverse triangle inequality,

$$\begin{aligned}||\varphi_1 - \tilde{\varphi}|| \quad &\geq \quad |\ ||\varphi_1 - c_1|| + ||c_1 - \tilde{\varphi}||\ | \\ &\geq \quad ||c_1 - \tilde{\varphi}|| = \tfrac{1}{2}z + ||\varphi_2 - c_2||\end{aligned}$$

so this is indeed the case. Considering the second inequality (18), we have likewise,

$$||\varphi_1 - \bar{\varphi}|| \leq ||\varphi_1 - c_1|| + ||c_1 - \bar{\varphi}|| < \tfrac{1}{4}z + \tfrac{1}{4}z = \tfrac{1}{2}z$$

and

$$||\varphi_2 - \bar{\varphi}|| \geq |\ ||\varphi_2 - c_2|| - ||c_2 - \bar{\varphi}||\ | \geq \tfrac{3}{4}z$$

so this inequality holds as well. Therefore, *any* φ_1, φ_2 from B_1, B_2 are associated with a plane which separates $\tilde{\varphi}$ from $\bar{\varphi}$ and hence triggers a projection. Since B_1 and B_2 have nonzero measure, and are subsets of B_r in which $p(\cdot)$ is nonzero, then the probabilities for this iteration $P_1 = \Pr($"a point is selected in B_1") and $P_2 = \Pr($"a point is selected in B_2") are both nonzero, and therefore, since the u^k are independent, $P_{\text{both}} = \Pr($"one point is selected in B_1 and the other is selected in B_2"$) = P_1 P_2$ is nonzero, and the probability that this occurs for *at least* one iteration $k > m$ is given by $1 - \prod_{k=m}^{\infty}\left(1 - P_{\text{both}}^k\right) = 1$ or in other words, with probability one, there exists a $q > m$ such that $P((n^q)^T\tilde{\varphi}^q > b^q)$. Then, by point 1, $d(\tilde{\varphi}^q, \bar{\varphi}) < d(\tilde{\varphi}^m, \bar{\varphi}) = \epsilon$, and so ϵ, with unit probability, cannot be a lower bound. Since d^k is a nonincreasing sequence in $\mathbb{R}$ and $\text{g.l.b.}(d^k) = 0$, d^k converges to 0 and thus $\tilde{\varphi}$ converges to $\bar{\varphi}$ in probability. ∎

An example of the estimate trajectory in feature space generated by such an observer is given in Figure 4.For this example, $\mathcal{F} = \mathbb{R}^2$, and features were drawn from a uniform distribution in the square $[-20, 20] \times [-20, 20]$. The estimate evolves from its initial condition, $\tilde{\varphi}^0 = (-15, 15)^T$ to near the ideal $\bar{\varphi} = (17, 0)^T$.

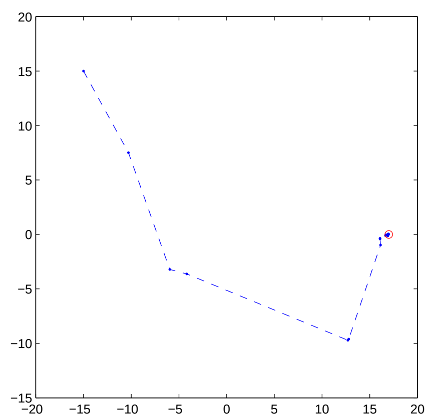

Fig. 4. Example estimate trajectory for observer (7-10) for $\alpha^k = \alpha = 1$, with $\mathcal{F} = \mathbb{R}^2$. The estimate begins at $\tilde{\varphi}^0 = (-15, 15)^T$, and approaches the ideal $\bar{\varphi} = (17, 0)^T$.

IV. Apples and Oranges

To demonstrate the application of these ideas, photos of nine apples were shown to an audience of thirteen people in a number of pairwise experiments.

Each apple was described by a 15-dimensional feature vector, containing (1-3) the average color in HSB (hue, saturation, brightness) color space, (4-6) the average color in RGB color space, (7) the color variance, (8-10) width, height, and the ratio of the two, (11-12) stem length, and

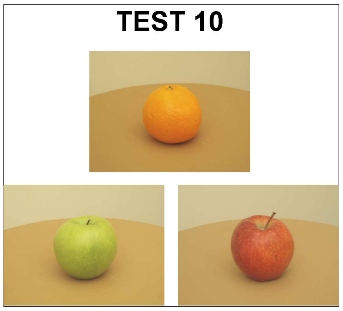

Fig. 5. An example of a pairwise comparison between two apples, relative to the orange.

angle relative to apple, (13-14) dimple angle and depth, and (15) roundness.

The partial order over the apples was thus generated by having a group of people make a number of randomly selected, pairwise comparisons (as the one depicted in Figure 5). Represented as a directed alternative graph, the results of these experiments are given as Figure 6.

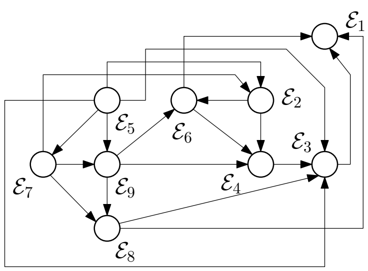

Fig. 6. The DAG corresponding to the apple experiments.

This results in the following optimal cost parameter ρ (all components of ρ not listed below are 0.0000):

$$\begin{aligned} \rho_1 &= 0.0505 && \text{(Hue)} \\ \rho_3 &= 0.1861 && \text{(Brightness)} \\ \rho_5 &= 0.2846 && \text{(Green)} \\ \rho_8 &= 0.2834 && \text{(Width)} \\ \rho_{11} &= 0.1953 && \text{(Stem Length)} \end{aligned}$$

which tells us that the single most important attribute that distinguishes apples from each other relative to oranges is the fifth dimension of the parameter space, namely, the amount of green in RGB colorspace; this is closely followed, perhaps surprisingly, by the width of the apple.

V. Conclusions

In this paper, we present a method for inferring the underlying cost structure that we assume is implicitly computed when people make comparisons between alternatives. In particular, given a collection of such comparisons, we produce a partial order over the set of alternatives, which, in turn, allows to infer the corresponding cost function (given a parametrized cost model and certain regularity assumptions on how people act.)

An example application of this is given in terms of comparing apples and oranges, and we recognize that this may not be the world's most compelling application in itself. Instead, we view this as a first step towards understanding and solving the very important question of Programming by Demonstration in robotics, where a robot is asked to act "similarly" to a human operator.

Acknowledgments

This work was sponsored by the US 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."

References

[1] D. Martinez and D. Kragic, "Modeling and recognition of actions through motor primitives," in *IEEE International Conference on Robotics and Automation (ICRA'08)*, May 2008, pp. 1704–1709.

[2] J. Tegin, J. Wikander, S. Ekvall, D. Kragic, and B. Illev, "Demonstration based learning and control for automatic grasping," in *The IEEE International Conference on Advanced Robotics (ICAR'07)*.

[3] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, *Robot Programming by Demonstration*, chapter 59, MIT Press, 2007.

[4] Cynthia Breazeal and Brian Scassellati, "Robots that imitate humans," *Trends in Cognitive Sciences*, vol. 6, no. 1, pp. 481–487, 2002.

[5] V. Kruger, D. Kragic, A. Ude, and C. Geib, "The meaning of action: A review on action recognition and mapping," *Advanced Robotics*, vol. 21, no. 13, pp. 1437–1501, 2007.

[6] *IEEE Transaction on Systems, Man, and Cybernetics, Special Issue on Robot Learning by Observation, Demonstration, and Imitation, Part B*, vol. 37, Apr 2007.

[7] S. Schaal, A. Ijspeert, and A. Billard, "Computational approaches to motor learning by imitation," *Philosophical Transactions of the Royal Society*, vol. 358, no. 1431, Mar. 2003.

[8] S. Schaal, "Is imitation learning the route to humanoid robots?," *Trends in Cognitive Sciences*, vol. 3, no. 6, pp. 233–242, Jun. 1999.

[9] T. Mehta and M. Egerstedt, "Learning multi-modal control programs," *Hybrid Systems: Computation and Control*, Mar 2005.

[10] A. Coates, P. Abbeel, and A. Y. Ng, "Learning for control from muliple demonstrations," in *Proceedings of the Twenty-fifth International Conference on Machine Learning*, 2008.

[11] A. Y. Ng P. Abbeel, M. Quigley, "Using inaccurate models in reinforcement learning," in *Proceedings of the Twenty-third International Conference on Machine Learning*, 2006.

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

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

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

[15] J. Diez, J. del Coz, O. Luances, and A. Bahamonde, "Clustering people according to their preference criteria," *Expert Systems with Applications*, vol. 34, pp. 1274–1284, 2008.

[16] Mian Qin, Scott Buffett, and Michael W. Fleming, "Predicting user preferences via similarity-based clustering," vol. 5032, pp. 222–233, 2008.

[17] A. Bahamonde, J. Diez, J. Quevedo, O. Luances, and J. del Coz, "How to learn consumer preferences from the analysis of sensory data by means of support vector machines (svm)," *Trends in Food Science and Technology*, vol. 18, pp. 20–28, 2007.

[18] A. V. Aho, M. R. Garey, and J. D. Ullman, "The transitive reduction of a directed graph," *SIAM Journal on Computing*, vol. 1, no. 2, pp. 131–137, Jun 1972.

[19] V. Strassen, "Gaussian elimination is not optimal," *Numer. Math*, vol. 13, no. 3, pp. 354–356, 1969.

[20] M. J. Fischer and A. R. Meyer, "Boolean matrix multiplication and transitive closure," in *Twelfth Annual Symposium on Switching and Automata Theory*, 1971, pp. 129–131.

[21] I. Munro, "Efficient determination of the strongly connected components and the transitive closure of a graph," Unpublished manuscript, 1971.