Learning Intransitive Reciprocal Relations with Kernel Methods

Tapio Pahikkala [a,*] Willem Waegeman [b] Evgeni Tsivtsivadze [c]
Tapio Salakoski [a] Bernard De Baets [b]

[a] *Turku Centre for Computer Science (TUCS)*
University of Turku, Department of Information Technology
Joukahaisenkatu 3-5 B, FIN-20520 Turku, Finland
firstname.lastname@utu.fi

[b] *KERMIT, Department of Applied Mathematics, Biometrics and Process Control, Ghent University, Coupure links 653, B-9000 Ghent, Belgium*
firstname.lastname@ugent.be

[c] *Institute for Computing and Information Sciences, Radboud University, Heyendaalseweg 135, 6525 AJ Nijmegen, The Netherlands*
firstname.lastname@science.ru.nl

Abstract

In different fields like decision making, psychology, game theory and biology, it has been observed that paired-comparison data like preference relations defined by humans and animals can be intransitive. Intransitive relations cannot be modeled with existing machine learning methods like ranking models, because these models exhibit strong transitivity properties. More specifically, in a stochastic context, where often the reciprocity property characterizes probabilistic relations such as choice probabilities, it has been formally shown that ranking models always satisfy the well-known strong stochastic transitivity property. Given this limitation of ranking models, we present a new kernel function that together with the regularized least-squares algorithm is capable of inferring intransitive reciprocal relations in problems where transitivity violations cannot be considered as noise. In this approach it is the kernel function that defines the transition from learning transitive to learning intransitive relations, and the Kronecker-product is introduced for representing the latter type of relations. In addition, we empirically demonstrate on two benchmark problems, one in game theory and one in theoretical biology, that our algorithm outperforms methods not capable of learning intransitive reciprocal relations.

Key words: transitivity, reciprocal relations, utility functions, kernel methods, preference learning, decision theory, game theory

* Corresponding author. Tel: +358-2-2154204 Fax: +358-2-2410154 Email:

1 Introduction

We start with an introductory example in the field of sports games in order to describe the purpose of this paper. Let us assume that an online betting company for tennis games wants to build statistical models to predict the probability that a given tennis player will defeat his/her opponent in the next Grand Slam competition. The company could be interested in building such models to maximize its profit when defining the amount of money that a client gets if he/she is able to predict the outcome of the game correctly. To this end, different types of data could be collected in order to construct the model, such as previous game outcomes, strong and weak points of players, current physical and mental conditions of players, etc. Yet, which type of machinery is required to obtain accurate predictions in this type of data mining problems? Firstly, as we will discuss in more detail below, we are for this example looking for an algorithm capable of predicting reciprocal relations from data, i.e., a relation between couples of players leading to a probability estimate of the outcome of a game. Secondly, we are also looking for a model that can predict intransitive relations, since commonly in sports games it turns out that game outcomes manifest cycles such as player A defeating player B, B defeating a third player C, and simultaneously C winning from A.

So, this paper in general considers learning problems where intransitive reciprocal relations need to be learned. As mathematical and statistical properties of human preference judgments, reciprocity and transitivity have been a subject of study for researchers in different fields like mathematical psychology, decision theory, social choice theory, and fuzzy modeling. Historically, this kind of research has been motivated by the quest for a rational characterization of human judgments, and to this end, transitivity is often assumed as a crucial property [Diaz et al., 2008]. This property basically says that a preference of an object $\mathbf{x}_i$ over another object $\mathbf{x}_j$ and a similar preference of $\mathbf{x}_j$ over a third object $\mathbf{x}_k$ should always result in a preference of $\mathbf{x}_i$ over $\mathbf{x}_k$, if preference judgments are made in a rational way. Nevertheless, it has been observed in several psychological experiments that human preference judgments often violate this transitivity property (see e.g. [Azand, 1993, Tversky, 1998]), especially in a context where preference judgments are considered as uncertain, resulting in non-crisp [1] preference relations between objects.

Contrary to some approaches taken in fuzzy set theory and decision theory, we adopt a probabilistic view of expressing uncertainty in decision behavior, as it is for example the case in social choice theory and mathematical psychology, where preference relations are often called binary choice probabilities. In this probabilistic framework, it can be assumed that a preference relation defined on a space $\mathcal{X}$

Tapio.Pahikkala@utu.fi

[1] In this work a relation is called crisp, when it can take only three values, e.g. 0 if A wins from B, 1 if B wins from A and 0.5 in case of a tie.

satisfies the reciprocity property.

Definition 1.1 *A function $Q : \mathcal{X}^2 \to [0,1]$ is called a reciprocal relation if for any $(\mathbf{x}, \mathbf{x}') \in \mathcal{X}^2$ it holds that*

$$Q(\mathbf{x}, \mathbf{x}') + Q(\mathbf{x}', \mathbf{x}) = 1\,.$$

The reciprocity property was already taken into consideration in the early work of Luce and Suppes [1965] in mathematical psychology. In addition, the same authors also introduced several stochastic transitivity properties like weak, moderate and strong stochastic transitivity to characterize rational preference judgments in a probabilistic sense. Let us recall the definition of weak stochastic transitivity.

Definition 1.2 *A reciprocal relation $Q : \mathcal{X}^2 \to [0,1]$ is called weakly stochastically transitive if for any $(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k) \in \mathcal{X}^3$ it holds that*

$$\Big(Q(\mathbf{x}_i, \mathbf{x}_j) \geq 1/2 \wedge Q(\mathbf{x}_j, \mathbf{x}_k) \geq 1/2\Big) \Rightarrow Q(\mathbf{x}_i, \mathbf{x}_k) \geq 1/2\,. \tag{1}$$

This definition of transitivity for reciprocal relations naturally extends the basic definition of transitivity for crisp relations. Below, when we speak about intransitive reciprocal relations, we specifically allude to reciprocal relations violating weak stochastic transitivity. In addition, we will also utilize strong stochastic transitivity a few times in this paper. This stronger condition is defined as follows.

Definition 1.3 *A reciprocal relation $Q : \mathcal{X}^2 \to [0,1]$ is called strongly stochastically transitive if for any $(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k) \in \mathcal{X}^3$ it holds that*

$$\Big(Q(\mathbf{x}_i, \mathbf{x}_j) \geq 1/2 \wedge Q(\mathbf{x}_j, \mathbf{x}_k) \geq 1/2\Big) \Rightarrow Q(\mathbf{x}_i, \mathbf{x}_k) \geq \max(Q(\mathbf{x}_i, \mathbf{x}_j), Q(\mathbf{x}_j, \mathbf{x}_k))\,.$$

Many other transitivity properties for reciprocal relations have been put forward in recent years, but these properties will not be discussed here. Moreover, many of these properties can be elegantly expressed in the cycle-transitivity framework. We refer to De Baets et al. [2006] for an overview of this framework and the various transitivity properties it covers.

As for crisp relations, several authors observed that stochastic transitivity properties are often violated. This is definitely the case for strong stochastic transitivity [Garcia-Lapresta and Mesenes, 2005], but sometimes even weak stochastic transitivity can be violated [Switalski, 2000]. As a consequence, there has been a long debate on interpreting this absence of transitivity. If preference judgments are considered as rational human decisions, then one should neglect the transitivity violations and apply traditional transitive models to represent this type of data. Although

such a simplification makes sense in certain situations, different authors argued that often these transitivity violations describe potential truths of reasoned comparisons (see [Fishburn, 1991] for a review). As a result, several authors have constructed models for paired-comparison data that are able to represent intransitive judgments explicitly (e.g. [Carroll et al., 1990, Tsai and Böckenholt, 2006]).

The motivation for building intransitive reciprocal preference relations might be debatable in a traditional (decision-theoretic) context, but the existence of rational transitivity violations becomes more appealing when the notion of a reciprocal preference relation is defined in a broader sense, like in the introductory example, or generally as any binary relation satisfying the reciprocity property. For example, reciprocal relations in game theory violate weak stochastic transitivity, in situations where the best strategy of a player depends on the strategy of his/her opponent — see e.g. the well-known rock-scissors-paper game [Fisher, 2008], dice games [De Schuymer et al., 2003, 2006, 2009]) and quantum games in physics [Makowski and Piotrowski, 2006]. Furthermore, in biology many examples of intransitive reciprocal relations have been encountered, like in competition between bacteria [Kerr et al., 2002, Czárán et al., 2002, Nowak, 2002, Kirkup and Riley, 2004, Károlyi et al., 2005, Reichenbach et al., 2007] and fungi [Boddy, 2000], mating choice of lizards [Sinervo and Lively, 1996] and food choice of birds [Waite, 2001]. Other examples of intransitive reciprocal relations can be found in order theory, when considering mutual ranking probabilities of the elements of a partially ordered set [De Baets et al., De Loof et al., 2010].

Generally speaking, we believe that enough examples exist to justify the need for models that can represent intransitive reciprocal relations. In this article we will address the topic of constructing such models based on any type of paired-comparison data. Basically, one can interpret these models as a mathematical representation of a reciprocal preference relation, having parameters that need to be statistically inferred. The approach we take finds its origin in machine learning, as a generalization of existing utility or ranking models. These models have been popular in areas like information retrieval and marketing for predicting decisions of web users and clients of e-commerce applications (see e.g. [Kalish and Nelson, 1991, Joachims, 2002]). Utility or ranking models by construction possess weak (and often even strong) stochastic transitivity properties, rendering them unsuitable for representing intransitive preference judgments in an accurate way. As a solution, we will extend an existing kernel-based ranking algorithm that has been proposed recently by some of the present authors [Pahikkala et al., 2007, 2009b]. This algorithm has been called RankRLS, as it optimizes a regularized least-squares (RLS) objective function on paired-comparison data that is represented as a graph.

This article is organized as follows. In Section 2 we give a short review on the role of transitivity in decision making and its connection to ranking models. Using the notions weak and strong stochastic transitivity, we in particular claim that ranking methods always exhibit certain transitivity properties that makes them useless

for representing intransitive reciprocal relations. Then, in Section 3 we start with a brief introduction to kernel methods, followed by a discussion of a general kernel-based framework for learning reciprocal relations. We show that existing ranking models are included in this framework via a particular choice of kernel function. We prove that these models cannot learn intransitive reciprocal relations. Subsequently, we formally claim and prove how in our framework another kernel, based on the Kronecker-product, is able to represent intransitive reciprocal relations in a much more adequate way. Finally, we present in Section 4 experimental results for two benchmark problems, demonstrating the advantages of our approach over traditional (transitive) ranking algorithms.

2 From transitive to intransitive preference models

In order to model preference judgments one can distinguish two main types of models in decision making [Öztürk et al., 2005, Waegeman et al., 2009]:

(1) *Scoring methods:* these methods typically construct a continuous function of the form $f : \mathcal{X} \to \mathbb{R}$ such that:

$$\mathbf{x} \succeq \mathbf{x}' \Leftrightarrow f(\mathbf{x}) \geq f(\mathbf{x}')\,,$$

which means that alternative $\mathbf{x}$ is preferred to alternative $\mathbf{x}'$ if the highest value was assigned to $\mathbf{x}$. In decision making, f is usually referred to as a utility function, while it is called a ranking function in machine learning [2] .

(2) *Pairwise preference models:* here the preference judgments are modeled by one (or more) valued relations $Q : \mathcal{X}^2 \to [0, 1]$ that express whether $\mathbf{x}$ should be preferred over $\mathbf{x}'$. One can distinguish different kinds of relations such as crisp relations, fuzzy relations or reciprocal relations.

The former approach has been especially popular in machine learning for scalability reasons. The latter approach allows a flexible and interpretable description of preference judgments and has therefore been popular in decision theory and the fuzzy set community, see e.g. [Mousseau et al., 2001, Doumpos and Zopounidis, 2004, De Baets and De Meyer, 2005, Dias and Mousseau, 2006].

The semantics underlying reciprocal preference relations is often probabilistic: $Q(\mathbf{x}, \mathbf{x}')$ expresses the probability that object $\mathbf{x}$ is preferred to $\mathbf{x}'$. One can in general construct such a reciprocal or probabilistic preference relation from a utility model in the following way:

[2] It should be emphasized that the notion of a utility function is often ambiguously defined in the literature. We adopt in this paper the rather mild definition of Luce and Suppes [1965], while more recent papers in economics and decision theory sometimes require more restrictive properties for a utility function, such as monotonicity, concavity and continuity.

$$Q(\mathbf{x}, \mathbf{x}') = g(f(\mathbf{x}), f(\mathbf{x}'))\,, \tag{2}$$

with $g : \mathbb{R}^2 \to [0,1]$ usually increasing in its first argument and decreasing in its second argument [Switalski, 2003]. Examples of models based on reciprocal preference relations are Bradley-Terry models [Bradley and Terry, 1952, Agresti, 2002] and Thurstone-Case5 models [Thurstone, 1927]. They have been applied in a machine learning learning context by Chu and Ghahramani [2005], Herbrich et al. [2007], Radlinski and Joachims [2007] and Hüllermeier et al. [2008].

The representability of reciprocal and fuzzy preference relations in terms of a single ranking or utility function has been extensively studied in domains like utility theory [Fishburn, 1970], preference modeling [Öztürk et al., 2005], social choice theory [Dasgupta and Deb, 1996, Fono and Andjiga, 2007], fuzzy set theory [Billot, 1995] and mathematical psychology [Luce and Suppes, 1965, Doignon et al., 1986]. It has been shown that the notions of transitivity and ranking representability play a crucial role in this context.

Definition 2.1 *A reciprocal relation $Q : \mathcal{X}^2 \to [0,1]$ is called weakly ranking representable if there exists a ranking function $f : \mathcal{X} \to \mathbb{R}$ such that for any $(\mathbf{x}, \mathbf{x}') \in \mathcal{X}^2$ it holds that*

$$Q(\mathbf{x}, \mathbf{x}') \leq \frac{1}{2} \Leftrightarrow f(\mathbf{x}) \leq f(\mathbf{x}')\,.$$

Reciprocal preference relations for which this condition is satisfied have also been called weak utility models. Luce and Suppes [1965] proved that a reciprocal preference relation is a weak utility model if and only if it satisfies weak stochastic transitivity, as defined by (1). As pointed out by Switalski [2003], a weakly ranking representable reciprocal relation can be characterized in terms of (2) such that the function $g : \mathbb{R}^2 \to \mathbb{R}$ satisfies

$$g(a,b) > \frac{1}{2} \Leftrightarrow a > b\,, \quad g(a,b) = \frac{1}{2} \Leftrightarrow a = b\,.$$

Analogous to weak ranking representability or weak utility models, one can define other conditions on the relationship between Q and f, leading to (stronger) transitivity conditions like moderate and strong stochastic transitivity. These properties are satisfied respectively by moderately and strongly ranking representable reciprocal preference relations. For such reciprocal relations one imposes additional conditions on g, for example the following type of reciprocal relation satisfies strong stochastic transitivity [Luce and Suppes, 1965].

Definition 2.2 *A reciprocal relation $Q : \mathcal{X}^2 \to [0,1]$ is called strongly ranking representable if it can be written as in (2) with g given by*

$$g(f(\mathbf{x}), f(\mathbf{x}')) = G(f(\mathbf{x}) - f(\mathbf{x}'))\,, \tag{3}$$

where $G : \mathbb{R} \to [0, 1]$ *is a cumulative distribution function satisfying* $G(0) = \frac{1}{2}$.

In addition, other transitivity conditions and corresponding conditions on G have been defined, such as strict ranking representability. This last property of reciprocal preference relations is satisfied by the Bradley-Terry model, a classical model for paired-comparison data [Bradley and Terry, 1952]. A further discussion on ranking representability is however beyond the scope of this paper. More details and proofs can be found in Luce and Suppes [1965], Carroll et al. [1990], Ballinger and Wilcox [1997], Tversky [1998], Switalski [2003], Dhzafarov [2003], Zhang [2004] and Waegeman and De Baets, submitted.

3 Learning intransitive reciprocal relations

In this section we will show how intransitive reciprocal relations can be learned from data with kernel methods. During the last decade, a lot of interesting papers on preference learning have appeared in the machine learning community (see e.g. Herbrich et al. [2000], Freund et al. [2003], Crammer and Singer [2001], Chu and Keerthi [2007]). Many of these authors use kernel methods to design learning algorithms. The majority of them also considers utility approaches to represent the preferences. Only a few authors such as Hüllermeier and Fürnkranz [2003], Chu and Ghahramani [2005] talk about pairwise preference relations, assuming weak stochastic transitivity so that an underlying ranking function exists.

We first explain the basic ideas behind kernel methods, followed by a discussion of a general framework for learning intransitive reciprocal relations. In this framework ranking can be seen as a special case, with a particular choice of the kernel function. To learn intransitive reciprocal relations, we then define a new type of kernel over pairs of data objects. We will formally prove that using this kernel we always learn relations that are reciprocal, but do not necessarily fulfill weak stochastic transitivity. This new kernel can be seen as a general concept that can be plugged into other kernel-based ranking methods as well, but in this paper we will illustrate its usefulness with the RLS algorithm. As this method optimizes a least-squares loss function, it is very suitable for learning reciprocal relations if the mean squared error measures the performance of the algorithm.

3.1 A brief introduction to kernels

This section is primarily based on Schölkopf and Smola [2002], Shawe-Taylor and Cristianini [2004]. A better and much more detailed introduction to kernel methods

can be found in these works. Given a not further specified input space $\mathcal{E}$ that shows at this moment no correspondence with the space $\mathcal{X}$ defined in the previous section, let us consider mappings of the following form:

$$\begin{aligned}\Phi : \mathcal{E} &\rightarrow \mathcal{F} \\ \mathbf{e} &\mapsto \Phi(\mathbf{e}) \,.\end{aligned}$$

The function Φ represents a so-called feature mapping from $\mathcal{E}$ to $\mathcal{F}$ and $\mathcal{F}$ is called the associated feature space. Initially, kernels were introduced to compute the dot-product $\langle \cdot, \cdot \rangle$ in this feature space efficiently. Such a compact representation of the dot-products in a certain feature space $\mathcal{H}$ will in general be called a kernel with the notation

$$\langle \Phi(\mathbf{e}_1), \Phi(\mathbf{e}_2) \rangle = K(\mathbf{e}_1, \mathbf{e}_2) \,.$$

For a given sequence $\mathbf{e}_1, ..., \mathbf{e}_N$ of objects, let us define the Gram matrix $\mathbf{K}$ of a given kernel as $\mathbf{K}_{i,j} = K(\mathbf{e}_i, \mathbf{e}_j)$. Kernel functions resulting in positive semi-definite Gram matrices always yield a dot-product. As a consequence, data analysis methods based on dot-products can always be rewritten in terms of kernels. Kernel versions have been proposed for classification, regression, clustering, principal component analysis, independent component analysis and many other methods [Shawe-Taylor and Cristianini, 2004]. These algorithms are quite general, because the class of models considered is simply changed by replacing the kernel function.

Kernels can be interpreted as similarity measures, allowing to model similarity of complex data objects. The specific form of the kernel function is domain-dependent and usually constructed by the data analyst [Schölkopf and Smola, 2002]. Since the introduction of kernels, similarity measures have been proposed for a large number of complex data types, like trees, graphs, strings, text, sets, images, DNA-sequences, etc. In this paper we will restrict our discussion to kernels for vectorial data. The most basic kernel for vectors one can think of is the one for which Φ defines the identity mapping, i.e.

$$K(\mathbf{e}_1, \mathbf{e}_2) = \langle \mathbf{e}_1, \mathbf{e}_2 \rangle \,.$$

This similarity measure is called a linear kernel, since it defines linear models. Alternatively, if interactions up to d features are allowed, one can use a polynomial kernel of degree d, i.e.

$$K(\mathbf{e}_1, \mathbf{e}_2) = \langle \mathbf{e}_1, \mathbf{e}_2 \rangle^d \,.$$

Another popular kernel function for non-linear modeling is the Gaussian RBF kernel

$$K(\mathbf{e}_1, \mathbf{e}_2) = e^{-\gamma ||\mathbf{e}_1 - \mathbf{e}_2||^2} \,,$$

where γ is a parameter determining the width of the kernel, resulting in an infinite-dimensional feature map Φ. Many other kernels, like spline kernels and ANOVA kernels, exist, but are rarely employed in practice.

In the context of kernel methods, we also have the concept of so-called regularized bias (see e.g. Rifkin [2002]). With this, we refer to the approach in which an extra constant valued dimension is added to the feature mapping. Consequently, the kernel value then changes into $K(\mathbf{e}_1, \mathbf{e}_2)+\beta^2$, where β is the extra constant feature.

Following the standard notations for kernel methods, we formulate our learning problem as the selection of a suitable function $h \in \mathcal{H}$, with $\mathcal{H}$ a certain hypothesis space, in particular a kernel reproducing Hilbert space (RKHS). Hypotheses $h : \mathcal{E} \to \mathbb{R}$ are usually denoted as $h(\mathbf{e}) = \langle \mathbf{w}, \mathbf{e} \rangle$ with $\mathbf{w}$ a vector of parameters that needs to be estimated based on training data. Let us denote a training dataset as a sequence

$$E = (\mathbf{e}_i, y_i)_{i=1}^N \,, \tag{4}$$

of input-label pairs, then we formally consider the following variational problem in which we select an appropriate hypothesis h from $\mathcal{H}$ for training data E. Namely, we consider an algorithm

$$\mathcal{A}(E) = \operatorname*{argmin}_{h \in \mathcal{H}} \frac{1}{N} \sum_{i=1}^{N} L(h(\mathbf{e}_i), y_i) + \lambda \|h\|_{\mathcal{H}}^2 \tag{5}$$

with L a given loss function and $\lambda > 0$ a regularization parameter. The first term measures the performance of a candidate hypothesis on the training data and the second term, called the regularizer, measures the complexity of the hypothesis with the RKHS norm. In our framework below, a squared loss is optimized in (5):

$$L(h(\mathbf{e}), y) = (h(\mathbf{e}) - y)^2 \,. \tag{6}$$

Optimizing this loss function instead of the more conventional hinge loss has the advantage that the solution can be found by simply solving a system of linear equations. We do not describe in detail the mathematical properties and advantages of this approach compared to more traditional algorithms, since that is not in the scope of this paper. More details can be found for example in Rifkin [2002], Suykens et al. [2002].

According to the representer theorem [Schölkopf and Smola, 2002], any minimizer $h \in \mathcal{H}$ of (5) admits a dual representation of the following form:

$$h(\mathbf{e}) = \sum_{i=1}^{N} \alpha_i K(\mathbf{e}, \mathbf{e}_i) = \langle \Phi(\mathbf{e}), \mathbf{w} \rangle,$$

where $\alpha_i \in \mathbb{R}$, K is the kernel function associated with the RKHS mentioned above, Φ is the feature mapping corresponding to K, and

$$\mathbf{w} = \sum_{i=1}^{N} \alpha_i \Phi(\mathbf{e}_i).$$

We will alternate several times between the primal and dual representation for h in the remainder of this article.

3.2 Learning reciprocal relations

We will use the above framework in order to learn intransitive reciprocal relations. To this end, we associate in a preference learning setting with each input a couple of data objects, i.e. $\mathbf{e}_i = (\mathbf{x}_i, \mathbf{x}'_i)$, where $\mathbf{x}_i, \mathbf{x}'_i \in \mathcal{X}$ and $\mathcal{X}$ can be any set. Consequently, we have an i.i.d. dataset

$$E = (\mathbf{x}_i, \mathbf{x}'_i, y_i)_{i=1}^{N} \,,$$

so that for each couple in the training dataset a label is known. These labels will represent reciprocal relations observed on training data, but rescaled to the interval $[-1, 1]$. This means that the following correspondence holds

$$y = 2Q(\mathbf{x}, \mathbf{x}') - 1 \,, \quad \forall (\mathbf{x}, \mathbf{x}') \in \mathcal{X}^2 \,.$$

Such a conversion is primarily made for ease of implementation. This implies that we will minimize the regularized squared error so that a model of type $h : \mathcal{X}^2 \to \mathbb{R}$ is obtained. An additional mapping $G : \mathbb{R} \to [0, 1]$ is required to ensure that $[0, 1]$-valued relations are predicted:

$$Q(\mathbf{x}, \mathbf{x}') = G(h(\mathbf{x}, \mathbf{x}')). \tag{7}$$

In our framework, we will aim to find such a Q that minimizes the mean squared error between the true and predicted reciprocal relations. When using a least-squares loss function, we can equivalently search for a model h that minimizes the same mean squared error and choose G as follows:

$$G(a) = \begin{cases} 0, & \text{if } a < -1 \,, \\ (a+1)/2, & \text{if } -1 \leq a \leq 1 \,, \\ 1, & \text{if } a > 1 \,, \end{cases} \tag{8}$$

Furthermore, to guarantee that reciprocal relations are learned, let us suggest the following type of feature mapping:

$$\Phi(\mathbf{e}_i) = \Phi(\mathbf{x}_i, \mathbf{x}'_i) = \Psi(\mathbf{x}_i, \mathbf{x}'_i) - \Psi(\mathbf{x}'_i, \mathbf{x}_i),$$

where Φ is just the same feature mapping as before but now written in terms of couples and Ψ is a new (not further specified) feature mapping from $\mathcal{X}^2$ to a feature space. As shown below, this construction will result in a reciprocal representation of the corresponding $[0,1]$-valued relation. By means of the representer theorem, the above model can be rewritten in terms of kernels, such that two different kernels pop up, one for Φ and one for Ψ. Both kernels express a similarity measure between two couples of objects and the following relationship holds:

$$\begin{aligned}
K^{\Phi}(\mathbf{e}_i, \mathbf{e}_j) &= K^{\Phi}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}_j, \mathbf{x}'_j) \\
&= \langle \Psi(\mathbf{x}_i, \mathbf{x}'_i) - \Psi(\mathbf{x}'_i, \mathbf{x}_i), \Psi(\mathbf{x}_j, \mathbf{x}'_j) - \Psi(\mathbf{x}'_j, \mathbf{x}_j) \rangle \\
&= \langle \Psi(\mathbf{x}_i, \mathbf{x}'_i), \Psi(\mathbf{x}_j, \mathbf{x}'_j) \rangle + \langle \Psi(\mathbf{x}'_i, \mathbf{x}_i), \Psi(\mathbf{x}'_j, \mathbf{x}_j) \rangle \\
&\quad - \langle \Psi(\mathbf{x}_i, \mathbf{x}'_i), \Psi(\mathbf{x}'_j, \mathbf{x}_j) \rangle - \langle \Psi(\mathbf{x}'_i, \mathbf{x}_i), \Psi(\mathbf{x}_j, \mathbf{x}'_j) \rangle \\
&= K^{\Psi}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}_j, \mathbf{x}'_j) + K^{\Psi}(\mathbf{x}'_i, \mathbf{x}_i, \mathbf{x}'_j, \mathbf{x}_j) \\
&\quad - K^{\Psi}(\mathbf{x}'_i, \mathbf{x}_i, \mathbf{x}_j, \mathbf{x}'_j) - K^{\Psi}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}'_j, \mathbf{x}_j) \,.
\end{aligned}$$

Using this notation, the prediction function given by the representer theorem can be expressed as:

$$h(\mathbf{x}, \mathbf{x}') = \langle \mathbf{w}, \Psi(\mathbf{x}, \mathbf{x}') - \Psi(\mathbf{x}', \mathbf{x}) \rangle = \sum_{i=1}^{N} \alpha_i K^{\Phi}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}, \mathbf{x}') \,.$$

For this prediction function, we can easily show that it forms the basis of a reciprocal relation.

Proposition 3.1 *Let $G : \mathbb{R} \to [0,1]$ be a cumulative distribution function satisfying $G(0) = 0.5$ and $G(-a) = 1 - G(a)$, then the function $Q : \mathcal{X}^2 \to [0,1]$ defined by (7) with $h : \mathcal{X}^2 \to \mathbb{R}$ given by (9), is a reciprocal relation.*

PROOF. One can easily see that $h(\mathbf{x}, \mathbf{x}') = -h(\mathbf{x}', \mathbf{x})$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$. The proof then immediately follows from:

$$\begin{aligned}
Q(\mathbf{x}, \mathbf{x}') + Q(\mathbf{x}', \mathbf{x}) &= G(h(\mathbf{x}, \mathbf{x}')) + G(h(\mathbf{x}', \mathbf{x})) \\
&= G(h(\mathbf{x}, \mathbf{x}')) - G(h(\mathbf{x}, \mathbf{x}')) + 1 \\
&= 1 \,.
\end{aligned}$$

□

Using the above notation, utility or ranking functions are usually written as

$$f(\mathbf{x}) = \langle \mathbf{w}, \phi(\mathbf{x}) \rangle \,. \tag{9}$$

They can be elegantly expressed in our framework by defining a specific feature mapping and corresponding kernel function.

Proposition 3.2 *If K^{Ψ} corresponds to the transitive kernel K_T^{Ψ} defined by*

$$K_T^{\Psi}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}_j, \mathbf{x}'_j) = K^{\phi}(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle \,,$$

with K^{ϕ} any two-dimensional kernel function on $\mathcal{X}^2$, whose value depends only on the arguments $\mathbf{x}_i$ and $\mathbf{x}_j$ and their feature representations $\phi(\mathbf{x}_i)$ and $\phi(\mathbf{x}_j)$, then the reciprocal relation $Q : \mathcal{X}^2 \to [0,1]$ given by (7) is strongly stochastically transitive.

PROOF. With the above representation for K^{Ψ} and $K(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$, the feature mapping Ψ, further denoted as Ψ_T, is for this model defined by

$$\Psi_T(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x}) \,.$$

So, only the first element of the couple is taken and the second element is simply ignored. Because of that, the model can be written as

$$\begin{aligned} h(\mathbf{x}, \mathbf{x}') &= \langle \mathbf{w}, \phi(\mathbf{x}) \rangle - \langle \mathbf{w}, \phi(\mathbf{x}') \rangle \\ &= f(\mathbf{x}) - f(\mathbf{x}') \,, \end{aligned}$$

with f defined by (9) such that Q takes the form of (3) without any further specified G. So, Q is strongly ranking representable. As a consequence, Q is also strongly stochastically transitive. We refer to [Luce and Suppes, 1965] for this proof. □

For this choice of K^{Ψ}, our framework is reduced to a popular type of kernel function that has been introduced by Herbrich et al. [2000]. The insight of the proposition is that the use of this kernel is equivalent to constructing a ranking for the individual inputs. This ranking function is in the dual representation given by:

$$f(\mathbf{x}) = \langle \mathbf{w}, \phi(\mathbf{x}) \rangle = \sum_{i=1}^{N} \alpha_i \Big(K^{\phi}(\mathbf{x}_i, \mathbf{x}) - K^{\phi}(\mathbf{x}'_i, \mathbf{x}) \Big) \,.$$

As explained in Section 2, ranking results in a reciprocal relations that satisfies the weak stochastic transitivity property. Due to the above proposition, we can even claim that the resulting reciprocal relation satisfies strong stochastic transitivity. Different ranking methods are obtained with different loss functions, such as RankSVM [Joachims, 2002] for the hinge loss and RankRLS [Pahikkala et al., 2007, 2009b] for the least-squares loss.

3.4 Learning intransitive reciprocal relations

Since the above choice for Ψ forms the core of all kernel-based ranking methods, these methods cannot generate intransitive relations, i.e. relations violating weak stochastic transitivity. In order to derive a model capable of violating weak stochastic transitivity, we introduce the following feature mapping Ψ_I for couples of objects:

$$\Psi_I(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x}) \otimes \phi(\mathbf{x}')\,,$$

where $\phi(\mathbf{x})$ is again the feature representation of the individual object $\mathbf{x}$ and $\otimes$ denotes the Kronecker-product, which is defined as follows:

$$A \otimes B = \begin{pmatrix} A_{1,1}B & \cdots & A_{1,n}B \\ \vdots & \ddots & \vdots \\ A_{m,1}B & \cdots & A_{m,n}B \end{pmatrix},$$

where A and B are matrices and $A_{i,j}$ is the i,jth element of A. Kernel functions induced by this type of feature maps have also been considered under the name tensor product kernels (see e.g.Weston et al. [2005]) or Kronecker kernels (see e.g. Kashima et al. [2009]) and the Kronecker product has also been used to construct kernels based of linear feature transformation (see e.g. Pahikkala et al. [2009a]).

We use the following property of the Kronecker product:

$$(A \otimes B)(C \otimes D) = (AC) \otimes (BD),$$

where $A \in \mathbb{R}^{a \times b}$, $B \in \mathbb{R}^{c \times d}$, $C \in \mathbb{R}^{b \times e}$, and $D \in \mathbb{R}^{b \times f}$. The Kronecker-product establishes joint feature representations Φ_I and Ψ_I that depend on both arguments of Φ and Ψ. Instead of ignoring the second argument of Φ and Ψ, we now represent all pairwise interactions between individual features of the two data objects in the joint feature representation. Using the notation K_I^{Ψ}, this leads to the following expression:

$$\begin{aligned} K_I^{\Psi}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}_j, \mathbf{x}'_j) &= \langle \phi(\mathbf{x}_i) \otimes \phi(\mathbf{x}'_i), \phi(\mathbf{x}_j) \otimes \phi(\mathbf{x}'_j) \rangle \\ &= \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle \otimes \langle \phi(\mathbf{x}'_i), \phi(\mathbf{x}'_j) \rangle \\ &= K^{\phi}(\mathbf{x}_i, \mathbf{x}_j) K^{\phi}(\mathbf{x}'_i, \mathbf{x}'_j), \end{aligned}$$

with again K^{ϕ} any kernel function defined over $\mathcal{X}^2$. As a result, using the Kronecker-product as feature mapping basically leads to a very simple kernel in the dual representation, consisting of just a regular product between two traditional kernels K^{ϕ}. Remark that K^{ϕ} can be any existing kernel, such as the linear kernel, the RBF-kernel, etc. As a result of the above construction, the kernel function K^{Φ} becomes:

$$K_I^{\Phi}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}_j, \mathbf{x}'_j) = 2K^{\phi}(\mathbf{x}_i, \mathbf{x}_j) K^{\phi}(\mathbf{x}'_i, \mathbf{x}'_j) - 2K^{\phi}(\mathbf{x}'_i, \mathbf{x}_j) K^{\phi}(\mathbf{x}_i, \mathbf{x}'_j).$$

We further refer to K_I^{Φ} as the intransitive kernel.

Indeed, in the above extension of the ranking framework, two different kernels K^{Ψ} and K^{ϕ} must be specified by the data analyst, while the third kernel K^{Φ} is defined by the choice for K^{Ψ}. On the one hand, the choice for K^{Ψ} (and hence K^{Φ}) determines whether the model is allowed to violate weak stochastic transitivity. On the other hand, the kernel function K^{ϕ} acts as the traditional similarity measure on $\mathcal{X}$, resulting in a linear, polynomial, radial basis function or any other representation of the data.

We now present a result indicating that the intransitive kernel K_I^{Φ} can be used to learn arbitrary reciprocal preference relations, provided that the feature representation ϕ of the individual objects is powerful enough. It is important to emphasize that the proposition does not impose any assumption on the loss function.

Proposition 3.3 *Let E be a training dataset of type (4), let $L : \mathbb{R}^2 \to \mathbb{R}^+$ be a loss function, and let $\mathcal{H}_R : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ be the set of all hypotheses inducing a reciprocal relation on $\mathcal{X}$. Moreover, let*

$$h^*(\mathbf{x}, \mathbf{x}') = \operatorname*{argmin}_{h \in \mathcal{H}_R} \sum_{i=1}^{N} L(y_i, h(\mathbf{x}_i, \mathbf{x}'_i)) \tag{10}$$

be the set of hypotheses inducing a reciprocal relation on $\mathcal{X}$ that have a minimal empirical loss on E. Further, let

$$\begin{aligned} h(\mathbf{x}, \mathbf{x}') &= \sum_{i=1}^{N} \alpha_i K_I^{\Phi}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}, \mathbf{x}') \\ &= \sum_{i=1}^{N} \alpha_i 2 \left(K^{\phi}(\mathbf{x}_i, \mathbf{x}) K^{\phi}(\mathbf{x}'_i, \mathbf{x}') - K^{\phi}(\mathbf{x}'_i, \mathbf{x}) K^{\phi}(\mathbf{x}_i, \mathbf{x}') \right), \end{aligned} \tag{11}$$

where $\alpha_i \in \mathbb{R}$, be the set of hypotheses we can construct using the intransitive kernel K_I^Φ and a given feature representation ϕ of a base kernel K^ϕ.

There exists a feature representation ϕ and coefficients α_i for which the corresponding hypothesis (11) is one of the minimizers of (10).

PROOF. First of all, remark that the above proposition does not make any assumption on the loss function. As mentioned in Section 3.1, in experiments we will consider the squared loss (6), but the proof we give here holds for other loss functions too. We start by defining the reciprocal relation which is the solution to (10).

The training set E may contain several couples that have the same two data objects either in the same or in the opposite order, while their labels may be noisy in such a way that there would be no reciprocal relation that would have a zero loss on the whole training set. Therefore, we define

$$Z_i^+ = \{j \mid j \in \{1, \ldots, N\}, \mathbf{x}_j = \mathbf{x}_i, \mathbf{x}'_j = \mathbf{x}'_i\}\,,$$
$$Z_i^- = \{j \mid j \in \{1, \ldots, N\}, \mathbf{x}_j = \mathbf{x}'_i, \mathbf{x}'_j = \mathbf{x}_i\}\,,$$

that is, Z_i^+ is the set of indices of the couples in the training set having $\mathbf{x}_i$ as the first and $\mathbf{x}'_i$ as the second data object and Z_i^- is the corresponding index set of the couples having $\mathbf{x}_i$ and $\mathbf{x}'_i$ in the opposite order. Moreover, for all $(\mathbf{x}_i, \mathbf{x}'_i, y_i) \in E$, let

$$\overline{y}_i = \operatorname*{argmin}_{y \in \mathbb{R}} \left(\sum_{j \in Z_i^+} L(y_j, y) + \sum_{j \in Z_i^-} L(y_j, -y) \right),$$

that is, $\overline{y}_i$ minimizes the sum of losses for the couples in E having the same two data objects as the i-th couple. Now, the function

$$h(\mathbf{x}_i, \mathbf{x}'_i) = \overline{y_i}\,, \quad \forall i \in \{1, ..., N\}\,,$$

obviously determines a solution to (10).

Next, let us define K^ϕ as follows:

$$K_S^\phi(\mathbf{x}, \mathbf{x}') = \begin{cases} 1, & \text{if } \mathbf{x} = \mathbf{x}'\,, \\ 0, & \text{if } \mathbf{x} \neq \mathbf{x}\,. \end{cases}$$

This kernel can be interpreted as a limit case of the Gaussian RBF kernel with $\gamma \to +\infty$. So, we define ϕ so that $\langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle = 1$ if $\mathbf{x} = \mathbf{x}'$ and $\langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle =$

0 otherwise. Then, $\langle \Phi(\mathbf{x}, \mathbf{x}'), \Phi(\mathbf{x}, \mathbf{x}') \rangle = 1$, $\langle \Phi(\mathbf{x}, \mathbf{x}'), \Phi(\mathbf{x}', \mathbf{x}) \rangle = -1$, and the value of the inner product is zero in all other cases. K^{Φ} is in this case given by

$$K^{\Phi}_{I,S}(\mathbf{x}_i, \mathbf{x}'_i, \mathbf{x}_j, \mathbf{x}'_j) = \begin{cases} 1, & \text{if } \mathbf{x}_i = \mathbf{x}_j \wedge \mathbf{x}'_i = \mathbf{x}'_j \,, \\ -1, & \text{if } \mathbf{x}_i = \mathbf{x}'_j \wedge \mathbf{x}'_i = \mathbf{x}_j \,, \\ 0, & \text{otherwise} \,. \end{cases}$$

In this construction, choosing

$$\alpha_i = \frac{\overline{y}_i}{|Z_i^+| + |Z_i^-|} \,,$$

satisfies $h(\mathbf{x}_i, \mathbf{x}'_i) = \overline{y}_i$ for all couples in the training set. $\square$

The above result indicates that this type of model is flexible enough to obtain an as low as possible empirical error on training data, while maintaining the reciprocity property. Hence, the algorithm can learn intransitive reciprocal relations, because intransitive reciprocal relations will minimize the empirical loss in intransitive problem settings.

4 Experiments

4.1 Rock-paper-scissors

In order to test our approach, we consider a semi-synthetic benchmark problem in game theory, a domain in which intransitive reciprocal relations between players is often observed. In such a context, a pure strategy provides a complete description of how a player will play a game. In particular, it determines the move a player will make for any situation (s)he could face. A player's strategy set is the set of pure strategies available to that player. A mixed strategy is an assignment of a probability to each pure strategy. This allows for a player to randomly select a pure strategy. Since probabilities are continuous, there are infinite mixed strategies available to a player, even if the strategy set is finite.

We consider learning the reciprocal relation of the probability that one player wins from another in the well-known rock-paper-scissors game. To test the performance of the learning algorithm in such a non-linear task, we generated the following synthetic data. First, we generate 100 individual objects for training and 100 for testing. The data objects are three-dimensional vectors representing players of the

rock-paper-scissors game. The three attributes of the players are the probabilities that the player will choose 'rock', 'paper', or 'scissors', respectively. The probability $P(r \mid \mathbf{x})$ of player $\mathbf{x}$ choosing rock is determined by $P(r \mid \mathbf{x}) = \exp(wu)/z$, where u is a random number between 0 and 1, w is a steepness parameter, and z is a normalization constant ensuring that the three probabilities sum up to one. By varying the width w of the exponential function, we can generate players tending to favor one of the three choices over the others or to play each choice almost equally likely.

We generate 1000 player couples for training by randomly selecting the first and the second player from the set of training players. Each couple represents a game of rock-paper-scissors and the outcome of this game can be considered as stochastic in nature, because the strategy of a player is chosen in accordance with the probabilities of picking a particular fixed strategy from that player's set of mixed strategies. For example, when a fixed rock player plays against a mixed strategy player that plays scissors with probability 0.8 and paper with probability 0.2, then we have a higher chance of observing a game outcome for which the fixed rock player wins from the second player. Yet, the same couple of players with different outcomes can simultaneously occur in the training data. During training and testing, the outcome of a game is -1, 0, or 1 depending on whether the first player loses the game, the game ends in a tie, or the first player wins the game, respectively. We use the game outcomes as the labels of the training couples.

For testing purposes, we use each possible couple of test players once, that is, we have a test set of 10000 games. However, instead of using the outcome of a single simulated game as label, we assign for each test couple the element of the reciprocal relation that corresponds to the probability that the first player wins:

$$\begin{aligned} Q(\mathbf{x}, \mathbf{x}') = {} & P(p \mid \mathbf{x})P(r \mid \mathbf{x}') + \frac{1}{2}P(p \mid \mathbf{x})P(p \mid \mathbf{x}') + P(r \mid \mathbf{x})P(s \mid \mathbf{x}') \\ & + \frac{1}{2}P(r \mid \mathbf{x})P(r \mid \mathbf{x}') + P(s \mid \mathbf{x})P(p \mid \mathbf{x}') + \frac{1}{2}P(s \mid \mathbf{x})P(s \mid \mathbf{x}') \,. \end{aligned}$$

The task is to learn to predict this reciprocal relation. The algorithm estimates the relation by rescaling the predicted outputs that lie in the interval $[-1, 1]$, as discussed above.

As mentioned above, we optimize a squared loss function on training data [3]. The values of the regularization and bias parameters are selected with a grid search and cross-validation performed on the training set. We conduct experiments with three data sets generated using the values 1, 10, and 100 for the parameter w. These parameterizations are illustrated in Figure 1. The value $w = 1$ corresponds to the situation where each player tends to play 'rock', 'paper', or 'scissors' almost equally

[3] For running the experiments, we use our RLScore software package, available at `http://www.tucs.fi/rlscore`.

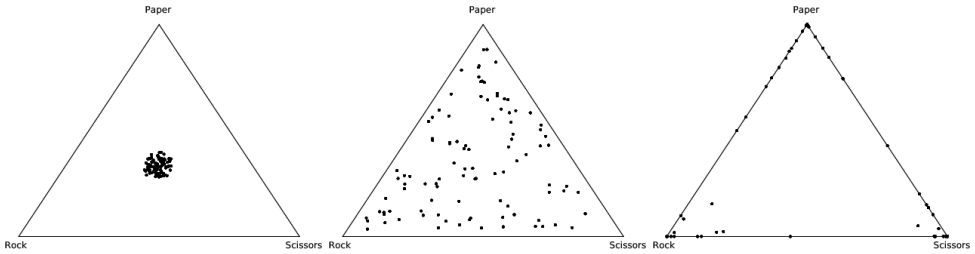

Fig. 1. Illustration of the players in the three data sets generated using the values 1 (left), 10 (middle), and 100 (right) for the parameter w.

	$w = 1$	$w = 10$	$w = 100$
I	0.000209	0.000445	0.000076
II	0.000162	0.006804	0.131972
III	0.000001	0.006454	0.125460

Table 1
Mean-squared error obtained with three different approaches: regularized least-squares with the kernel K_I^Φ (I), regularized least-squares with the kernel K_T^Φ (II) and a naive approach consisting of always predicting $1/2$ (III).

	$w = 1$	$w = 10$	$w = 100$
I	0.538200	0.957800	0.995000
II	0.5	0.5	0.592950
III	0.5	0.5	0.5

Table 2
Classification accuracy obtained with three different approaches: regularized least-squares with the kernel K_I^Φ (I), regularized least-squares with the kernel K_T^Φ (II) and a naive approach consisting of always predicting a tie (III).

likely, that is, the players are concentrated in the center of the triangle in the figure. For $w = 100$ the players always tend to play only their favorite item, that is, the players' strategies are concentrated near the three corners of the triangle. Finally, $w = 10$ corresponds to a setting between these two extremes.

The regression results are presented in Table 1. We report the mean squared-error obtained by regularized least-squares in a transitive and intransitive setting, respectively by specifying the kernels K_T^Φ and K_I^Φ. For K^ϕ a simple linear kernel is chosen in both cases. In addition, in Table 2 we present the results of binary classification experiments, in which the aim is to correctly predict the direction of preference, that is, whether the first player is more likely to win the second player or vice versa. In both regression and classification experiments, we also compare these two approaches with a naive heuristic consisting of always predicting $1/2$ (a tie).

In the regression experiments, the heuristic of always predicting a tie can be inter-

preted as quite optimal for $w = 1$, because in that case all players are located in the center of the triangle. This explains why neither the transitive nor the intransitive regularized least-squares algorithm can outperform this naive approach when $w = 1$. We conclude that there is not much to learn in this case. For the other two values of w, the situation is different, with the regularized least-squares algorithm with the intransitive kernel performing substantially better than the naive approach, while the performance with the transitive kernel being close to that of the naive one. Unsurprisingly, learning the intransitive reciprocal relations is more difficult when the probabilities of the players are close to the uniform distribution ($w = 10$) than in case the players tend to always play their favorite strategy ($w = 100$). Especially in this last case, regularized least-squares with an intransitive kernel performs substantially better than its transitive counterpart. This supports the claim that our approach works well in practice, when the reciprocal relation to be learned indeed violates weak stochastic transitivity. The stronger this violation, the more the advantage of an intransitive kernel will become visible.

In the binary classification experiments, the heuristic of always predicting a tie or the same class has the classification accuracy equal to 0.5 in all cases, because each pair of players is twice in the test data. For the trained predictors, the most difficult case is now with the parameter $w = 1$, because the game outcomes are almost random when both players are in the center of the triangle. For the values $w = 10$ and $w = 100$, RLS with the kernel K_I^Φ is able to perform the classification almost perfectly, while the classification accuracy with the kernel K_T^Φ is of random level.

4.2 Theoretical biology

Non-transitive competition between species has recently received attention in theoretical biology. This phenomenon has been observed in many natural systems (see e.g. Sinervo and Lively [1996], Boddy [2000], Kerr et al. [2002], Czárán et al. [2002], Nowak [2002], Kirkup and Riley [2004], Károlyi et al. [2005], Reichenbach et al. [2007]) and it has been studied and analyzed with computer simulations (see e.g. Frean and Abraham [2001], Frean [2006]). The simulations usually consist of an initial population of individuals or species and some limited resource, such as space, for which they compete. Most of the studies of non-transitive systems have considered rock-paper-scissors type of relationships between the competing species. Some of the studies and simulations also address competition of mutated individuals of a single species having a similar type of non-transitive fashion as that of interspecific competition. Below, we use the term species when referring to the individuals in a cyclic competitive structure.

Inspired by the simulations made by Frean [2006], we consider the following setting. Suppose we have a number of competing species, each of them having two features. Namely, a species $\mathbf{x}$ has a strong point denoted by $s(\mathbf{x})$ and a weak point

denoted by $w(\mathbf{x})$, and the values of both features are between 0 and 1. Then, for a couple of individuals, say $(\mathbf{x}, \mathbf{x}')$, we define a label y, whose value equals 1 if $\mathbf{x}$ dominates $\mathbf{x}'$ and -1 in the opposite case. The dominance is determined by the following formula:

$$y = \text{sign}(u(s(\mathbf{x}'), w(\mathbf{x})) - u(s(\mathbf{x}), w(\mathbf{x}'))) \tag{12}$$

where sign is the signum function and

$$u(a, b) = \min(|a - b|, 1 - |a - b|)\,. \tag{13}$$

We observe that the species $\mathbf{x}$ dominates $\mathbf{x}'$ if and only if the strong point $s(\mathbf{x})$ of $\mathbf{x}$ is closer to the weak point $w(\mathbf{x}')$ of $\mathbf{x}'$ than $s(\mathbf{x}')$ is to $w(\mathbf{x})$, the closeness being defined by (13), expressing distance considering the unit interval $[0, 1]$ as a circular domain.

We elucidate the strong and weak points of a species with the following toy example about animals. The strong point of an animal can be considered, for example, as the colour that the animal is best able to see and the weak point the colour of the animal. Further, the colour can be considered as a continuous variable so that the smaller the distance between the colours $s(\mathbf{x})$ and $w(\mathbf{x}')$ is, the better $\mathbf{x}$ is able to see $\mathbf{x}'$. Then, an animal $\mathbf{x}$ can dominate animal $\mathbf{x}'$ if the distance is small enough.

We set up an experiment in which we randomly generate an initial population of $2,500$ species so that their strong and weak points have been drawn from a uniform distribution between 0 and 1. Then, we select randomly two species $\mathbf{x}$ and $\mathbf{x}'$ from the population for which we compute a label y with (12). In the confrontation of these two species, we say that $\mathbf{x}$ is the winner and $\mathbf{x}'$ is the loser if $y = 1$ and vice versa if $y = -1$. After the confrontation, the loser is replaced with a mutation $\hat{\mathbf{x}}$ of the winner. The strong and weak points of the mutant are obtained from the strong and weak points of the winner by shifting them by small amounts whose sizes are drawn from a normal distribution having zero mean and standard deviation 0.005.

Unlike in the experiments done by Frean [2006], we adopt an approach in which we do not consider any local neighborhood of the species, that is, the two confronting species are randomly selected from the current population of $2,500$ species. This is done in order to simplify the experimental setting. In addition, for each confrontation of two species, there is always a winner and a loser, while this was the case in the experiments of Frean [2006] only if the value of (13) for $s(\mathbf{x})$ and $w(\mathbf{x}')$ was smaller than a certain threshold. A next couple was randomly selected in case of the value being larger than the threshold. Finally, our closeness function (13) differs from the one used by Frean [2006] so that the strong and the weak points are cyclic in the sense that values 0 and 1 can be considered to be equal. We adopted the cyclic property of the weak and strong points in order to eliminate the special case of the values being close to 0 and 1.

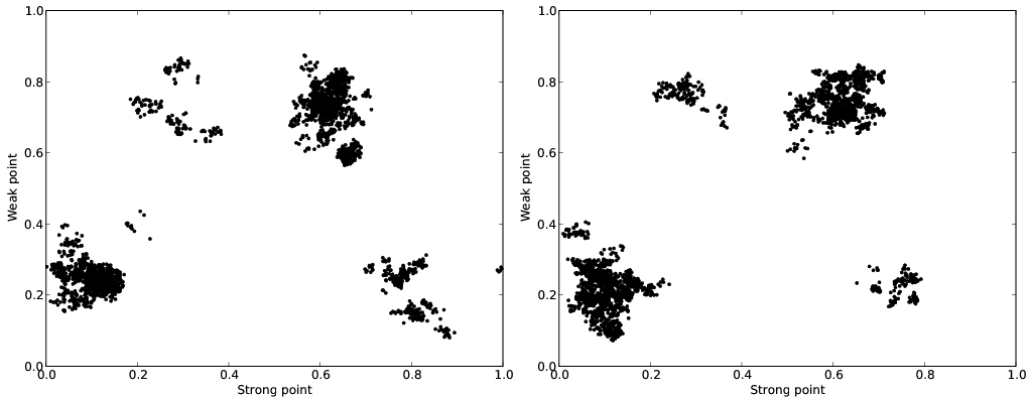

Fig. 2. The set of $2,500$ species after $900,000$ (left) and $1,000,000$ (right) confrontations.

Method	Accuracy	MSE
I	0.960600	0.000146
II	0.680000	0.007835
III	0.500000	0.007757

Table 3
Classification accuracy (left) and mean squared regression error (right) obtained with three different approaches: regularized least-squares with the kernel K_I^Φ (I), regularized least-squares with the kernel K_T^Φ (II) and a naive approach consisting of always predicting a tie in classification and 0.5 in regression.

We perform altogether $1,000,000$ subsequent confrontations of two species. In the beginning, there are no clusters, since the strong and weak points of the species are uniformly distributed. However, the species start to form small clusters after a couple of tens of thousands of confrontations and large clusters when a couple of hundreds of thousands of confrontations has passed. We sample our training and test sets from the $100,000$ last confrontations, since at this point the simulation has already formed quite stable clusters. Namely, we randomly sample without replacement $1,000$ couples for a training set and $10,000$ for a test set. The clusters formed after $900,000$ and $1,000,000$ confrontations are depicted in Figure 2. The figures are two consecutive snapshots of a movie which is available online at `http://staff.cs.utu.fi/~aatapa/tbmovie.avi`.

We train two RLS classifiers with the training set of 1000 confrontations and use them for predicting the outcomes of the unseen 10000 confrontations in the test set. The first classifier uses a transitive kernel K_T^Φ and the second one an intransitive kernel K_I^Φ. The base kernel K^ϕ is chosen to be the Gaussian radial basis function kernel for both the cases, that is,

$$K^\phi(\mathbf{x}, \mathbf{x}') = e^{-\gamma((s(\mathbf{x})-s(\mathbf{x}'))^2+(w(\mathbf{x})-w(\mathbf{x}'))^2)}\,. \tag{14}$$

The value of the regularization and bias parameters, and the width γ of the Gaussian kernel are selected with a grid search and cross-validation performed on the training

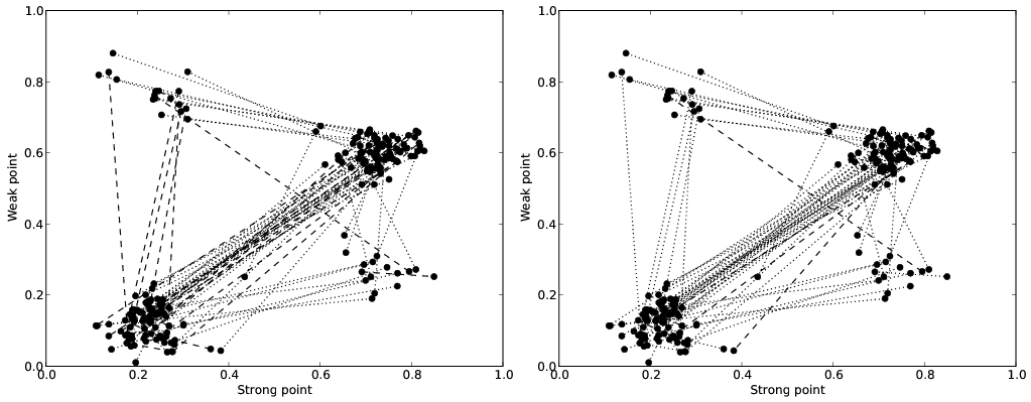

Fig. 3. Illustration of 100 randomly selected test couples. Left: the dotted lines denote the 69 couples classified correctly and the dashed lines denote the 31 incorrectly classified ones using RLS with the transitive kernel. Right: the dotted lines denote the 89 couples classified correctly and the dashed lines denote the 11 incorrectly classified ones using RLS with the intransitive kernel.

set.

The experimental results are listed in Table 3. Moreover, a random sample of 100 test couples and their classifications by the transitive and intransitive RLS classifier are illustrated in Figure 3. From the results, we observe that the classifier using the transitive kernel can learn the relation to some extent, but the intransitive kernel is clearly better for this purpose. In addition to the classification experiments, we also conduct experiments in which we aim to correctly regress the value of the reciprocal relation between two species. The value of the relation for two species is obtained via using the scaling function (8) in place of the signum function in (12). One can clearly observe that here as well the intransitive kernel performs clearly better than its transitive counterpart.

5 Conclusion

In this paper the problem of learning intransitive reciprocal relations was tackled. To this end, we showed that existing approaches for preference learning typically exhibit strong stochastic transitivity as property, and we introduced an extension of the existing RankRLS framework to predict reciprocal relations that can violate weak stochastic transitivity. In this framework, the choice of kernel function defines the transition from transitive to intransitive models. By choosing a feature mapping based on the Kronecker-product, we are able to predict intransitive reciprocal relations. Experiments on benchmark problems in game theory and theoretical biology confirmed that our approach substantially outperforms the ranking approach when intransitive reciprocal relations are present in the data. Given the absence of publicly available datasets on learning intransitive reciprocal relations, we are willing to share our data with other researchers, and in the future we hope to apply our

algorithm in other domains as well.

From a decision making point of view, one might argue that the models proposed in this paper suffer from some lack of interpretability. However, for most real-world data modeling problems, a clear trade-off between interpretability and performance can be expected. Our methods rather incline to the latter side of the balance: state-of-the-art predictive performance but a limited interpretability. Nevertheless, interpretability can simply be preserved by choosing a linear kernel for K^{ϕ}, while still not necessarily imposing transitivity by using the Kronecker product kernel on top of this linear kernel. Furthermore, as another important property of commonly used decision models, monotonicity can also be guaranteed in the same way, since a linear model always satisfies monotonicity, but unfortunately not vice versa, because monotone models do not necessarily have to be linear models. For kernel methods in particular, it is widely accepted that monotonicity cannot be easily guaranteed in the dual formulation. As a specific form of incorporating domain knowledge into kernel methods, Le et al. [2006] recently proposed an algorithm capable of enforcing monotonicity in regression problems, but the topic definitely remains an open challenge that deserves further attention in future work.

References

A. Agresti. *Categorical Data Analysis, 2nd version*. John Wiley and Sons, 2002.

P. Azand. The philosophy of intransitive preferences. *The Economic Journal*, 103: 337–346, 1993.

T. Ballinger and P. Wilcox. Decisions, error and heterogeneity. *The Economic Journal*, 107:1090–1105, 1997.

A. Billot. An existence theorem for fuzzy utility functions: A new elementary proof. *Fuzzy Sets and Systems*, 74:271–276, 1995.

L. Boddy. Interspecific combative interactions between wood-decaying basidiomycetes. *FEMS Microbiology Ecology*, 31:185–194, 2000.

R. Bradley and M. Terry. Rank analysis of incomplete block designs. I: The method of paired comparisons. *Biometrika*, 39:324–345, 1952.

J. Carroll, G. De Soete, and W. De Sarbo. Two stochastic multidimensional choice models for marketing research. *Decision Sciences*, 21:337–356, 1990.

W. Chu and Z. Ghahramani. Preference learning with gaussian processes. In *Proceedings of the International Conference on Machine Learning, Bonn, Germany*, pages 137–144, 2005.

W. Chu and S. Keerthi. Support vector ordinal regression. *Neural Computation*, 19 (3):792–815, 2007.

K. Crammer and Y. Singer. Pranking with ranking. In *Proceedings of the Conference on Neural Information Processing Systems, Vancouver, Canada*, pages 641–647, 2001.

T. Czárán, R. Hoekstra, and L. Pagie. Chemical warfare between microbes pro-

motes biodiversity. *Proceedings of the National Academy of Sciences*, 99(2): 786–790, 2002.

M. Dasgupta and R. Deb. Transitivity and fuzzy preferences. *Social Choice and Welfare*, 13:305–318, 1996.

B. De Baets and H. De Meyer. Transitivity frameworks for reciprocal relations: cycle-transitivity versus FG-transitivity. *Fuzzy Sets and Systems*, 152:249–270, 2005.

B. De Baets, H. De Meyer, and K. De Loof. On the cycle-transitivity of the mutual rank probability relation of a poset. *Fuzzy Sets and Systems*. submitted.

B. De Baets, H. De Meyer, B. De Schuymer, and S. Jenei. Cyclic evaluation of transitivity of reciprocal relations. *Social Choice and Welfare*, 26:217–238, 2006.

K. De Loof, B. De Baets, and H. De Meyer. Counting linear extension majority cycles in posets on up to 13 points. *Computers and Mathematics with Applications*, 59:1541–1547, 2010.

B. De Schuymer, H. De Meyer, B. De Baets, and S. Jenei. On the cycle-transitivity of the dice model. *Theory and Decision*, 54:164–185, 2003.

B. De Schuymer, H. De Meyer, and B. De Baets. Optimal strategies for equal-sum dice games. *Discrete Applied Mathematics*, 154:2565–2576, 2006.

B. De Schuymer, H. De Meyer, and B. De Baets. Optimal strategies for symmetric matrix games with partitions. *Bulletin of the Belgian Mathematical Society - Simon Stevin*, 16:67–89, 2009.

E. Dhzafarov. Thurstonian-type representations for same-different discriminations: Deterministic decisions and independent images. *Journal of Mathematical Psychology*, 47:184–204, 2003.

L. Dias and V. Mousseau. Inferring ELECTRE's veto-related parameters from outranking examples. *European Journal of Operational Research*, 170:172–191, 2006.

S. Diaz, J.-L. Garcia-Lapresta, and S. Montes. Consistent models of transitivity for reciprocal preferences on a finite ordinal scale. *Information Sciences*, 178(13): 2832–2848, 2008.

J.-P. Doignon, B. Monjardet, M. Roubens, and Ph. Vincke. Biorder families, valued relations and preference modelling. *Journal of Mathematical Psychology*, 30: 435–480, 1986.

M. Doumpos and C. Zopounidis. A multicriteria classification approach based on pairwise comparisons. *European Journal of Operational Research*, 158:378–389, 2004.

P. Fishburn. *Utility Theory for Decision Making*. Wiley, 1970.

P. Fishburn. Nontransitive preferences in decision theory. *Journal of Risk and Uncertainty*, 4:113–134, 1991.

L. Fisher. *Rock, Paper, Scissors: Game Theory in Everyday Life*. Basic Books, 2008.

L. Fono and N. Andjiga. Utility function of fuzzy preferences on a countable set under max-*-transitivity. *Social Choice and Welfare*, 28:667–683, 2007.

M. Frean. Emergence of cyclic competitions in spatial ecosystems. In P.A. Whigham, editor, *SIRC 2006: Interactions and Spatial Process (Eighteenth Col-*

loquium hosted by the Spatial Information Research Centre), pages 1–9, 2006.

M. Frean and E. Abraham. Rock-scissors-paper and the survival of the weakest. *Proceedings of the Royal Society (London), Series B*, 268:1323–1328, 2001.

Y. Freund, R. Yier, R. Schapire, and Y. Singer. An efficient boosting algorithm for combining preferences. *Journal of Machine Learning Research*, 4:933–969, 2003.

J. Garcia-Lapresta and C. Mesenes. Individual-valued preferences and their aggregation: Consistensy analysis in a real case. *Fuzzy Sets and Systems*, 151: 269–284, 2005.

R. Herbrich, T. Graepel, and K. Obermayer. Large margin rank boundaries for ordinal regression. In A. Smola, P. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 115–132. MIT Press, 2000.

R. Herbrich, T. Minka, and T. Graepel. Trueskill: a bayesian skill rating system. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 569–576. The MIT Press, Cambridge, MA, 2007.

E. Hüllermeier and J. Fürnkranz. Pairwise preference learning and ranking. In *Proceedings of the European Conference on Machine Learning, Dubrovnik, Croatia*, pages 145–156, 2003.

E. Hüllermeier, J. Fürnkranz, W. Cheng, and K. Brinker. Label ranking by learning pairwise preferences. *Artificial Intelligence*, 172:1897–1916, 2008.

T. Joachims. Optimizing search engines using clickthrough data. In D. Hand, D. Keim, and R. Ng, editors, *Proceedings of the 8th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'02)*, pages 133–142. ACM Press, 2002.

S. Kalish and P. Nelson. A comparison of ranking, rating and reservation price measurement in conjoint analysis. *Marketing Letters*, 2(4):327–335, 1991.

G. Károlyi, Z. Neufeld, and I. Scheuring. Rock-scissors-paper game in a chaotic flow: The effect of dispersion on the cyclic competition of microorganisms. *Journal of Theoretical Biology*, 236(1):12 – 20, 2005.

H. Kashima, S. Oyama, Y. Yamanishi, and K. Tsuda. On pairwise kernels: An efficient alternative and generalization analysis. In Thanaruk Theeramunkong, Boonserm Kijsirikul, Nick Cercone, and Tu Bao Ho, editors, *PAKDD*, volume 5476 of *Lecture Notes in Computer Science*, pages 1030–1037. Springer, 2009.

B. Kerr, M. Riley, M.Feldman, and B. Bohannan. Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature*, 418:171–174, 2002.

B. Kirkup and M. Riley. Antibiotic-mediated antagonism leads to a bacterial game of rock-paper-scissors in vivo. *Nature*, 428:412–414, 2004.

Q. Le, A. Smola, and T. Grtner. Simpler knowledge-based support vector machines. In *Proceedings of the 23rd international conference on Machine learning, Pittsburgh, PA, USA*, pages 521–528, 2006.

R. Luce and P. Suppes. Preference, utility and subjective probability. In R. Luce, R. Bush, and E. Galanter, editors, *Handbook of Mathematical Psychology*, pages 249–410. Wiley, 1965.

M. Makowski and E. Piotrowski. Quantum cat's dilemma: an example of intransitivity in quantum games. *Physics Letters A*, 355:250–254, 2006.

V. Mousseau, J. Figueira, and J. Naux. Using assignment examples to infer weights for ELECTRE TRI method: Some experimental results. *European Journal of Operational Research*, 130(2):263–275, 2001.

M. Nowak. Biodiversity: Bacterial game dynamics. *Nature*, 418:138–139, 2002.

M. Öztürk, A. Tsoukiàs, and Ph. Vincke. Preference modelling. In J. Figueira, S. Greco, and M. Ehrgott, editors, *Multiple Criteria Decision Analysis. State of the Art Surveys*, pages 27–71. Springer-Verlag, 2005.

T. Pahikkala, E. Tsivtsivadze, A. Airola, J. Boberg, and T. Salakoski. Learning to rank with pairwise regularized least-squares. In Thorsten Joachims, Hang Li, Tie-Yan Liu, and ChengXiang Zhai, editors, *SIGIR 2007 Workshop on Learning to Rank for Information Retrieval*, pages 27–33, 2007.

T. Pahikkala, S. Pyysalo, J. Boberg, J. Järvinen, and T. Salakoski. Matrix representations, linear transformations, and kernels for disambiguation in natural language. *Machine Learning*, 74(2):133–158, 2009a.

T. Pahikkala, E. Tsivtsivadze, A. Airola, J. Järvinen, and J. Boberg. An efficient algorithm for learning to rank from preference graphs. *Machine Learning*, 75 (1):129–165, 2009b.

F. Radlinski and T. Joachims. Active exploration for learning rankings from clickthrough data. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining, San Jose, CA, U.S.A.*, pages 570–579, 2007.

T. Reichenbach, M. Mobilia, and E. Frey. Mobility promotes and jeopardizes biodiversity in rock-paper-scissors games. *Nature*, 448:1046–1049, 2007.

R. Rifkin. *Everything Old Is New Again: A Fresh Look at Historical Approaches in Machine Learning*. PhD thesis, Massachusetts Institute of Technology, 2002.

B. Schölkopf and A. Smola. *Learning with Kernels, Support Vector Machines, Regularisation, Optimization and Beyond*. The MIT Press, 2002.

J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.

S. Sinervo and C. Lively. The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 340:240–246, 1996.

J. Suykens, T. Van Gestel, J. De Brabanter, B. De Moor, and J. Vandewalle. *Least Squares Support Vector Machines*. World Scientific Pub. Co., Singapore, 2002.

Z. Switalski. Transitivity of fuzzy preference relations - an empirical study. *Fuzzy Sets and Systems*, 118:503–508, 2000.

Z. Switalski. General transitivity conditions for fuzzy reciprocal preference matrices. *Fuzzy Sets and Systems*, 137:85–100, 2003.

L. Thurstone. A law of comparative judgment. *Psychological Review*, 79:281–299, 1927.

R.-C. Tsai and U. Böckenholt. Modelling intransitive preferences: A random-effects approach. *Journal of Mathematical Psychology*, 50:1–14, 2006.

A. Tversky. *Preference, Belief and Similarity*. MIT Press, 1998.

W. Waegeman and B. De Baets. On the ERA ranking representability of multi-class classifiers. *Artificial Intelligence*. submitted.

W. Waegeman, B. De Baets, and L. Boullart. Kernel-based learning methods for preference aggregation. *4OR*, 7:169–189, 2009.

T. Waite. Intransitive preferences in hoarding gray jays (*Perisoreus canadensis*). *Journal of Behavioural Ecology and Sociabiology*, 50:116–121, 2001.

J. Weston, B. Schölkopf, and O. Bousquet. Joint kernel maps. In Joan Cabestany, Alberto Prieto, and Francisco Sandoval Hernández, editors, *Computational Intelligence and Bioinspired Systems, 8th International Work-Conference on Artificial Neural Networks (IWANN 2005)*, volume 3512 of *Lecture Notes in Computer Science*, pages 176–191, Berlin Heidelberg, Germany, 2005. Springer-Verlag.

J. Zhang. Binary choice, subset choice, random utility, and ranking: A unified perspective using the permutahedron. *Journal of Mathematical Psychology*, 48: 107–134, 2004.