



HAL
open science

The complexity of unsupervised learning of lexicographic preferences

Hélène Fargier, Pierre-François Gimenez, Jérôme Mengin, Bao Ngoc Le
Nguyen

► **To cite this version:**

Hélène Fargier, Pierre-François Gimenez, Jérôme Mengin, Bao Ngoc Le Nguyen. The complexity of unsupervised learning of lexicographic preferences. MPREF 2022 - 13th Multidisciplinary Workshop on Advances in Preference Handling, Jul 2022, Vienne, Austria. pp.1-8. hal-03784693

HAL Id: hal-03784693

<https://inria.hal.science/hal-03784693v1>

Submitted on 23 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The complexity of unsupervised learning of lexicographic preferences

Hélène Fargier¹, Pierre-François Gimenez², Jérôme Mengin³ and Ngoc Bao Nguyen⁴

^{1,3}IRIT, Université de Toulouse, CNRS, Toulouse INP, UT3, Toulouse, France

²CentraleSupélec, Univ. Rennes, IRISA

⁴INSA Toulouse

{helene.fargier, jerome.mengin}@irit.fr, pierre-francois.gimenez@centralesupelec.fr,
nbnnguyen@etud.insa-toulouse.fr

Abstract

This paper considers the task of learning users' preferences on a combinatorial set of alternatives, as generally used by online configurators, for example. In many settings, only a set of selected alternatives during past interactions is available to the learner. Fargier et al. [2018] propose an approach to learn, in such a setting, a model of the users' preferences that ranks previously chosen alternatives as high as possible; and an algorithm to learn, in this setting, a particular model of preferences: lexicographic preferences trees (LP-trees). In this paper, we study complexity-theoretical problems related to this approach. We give an upper bound on the sample complexity of learning an LP-tree, which is logarithmic in the number of attributes. We also prove that computing the LP tree that minimises the empirical risk can be done in polynomial time when restricted to the class of linear LP-trees.

1 Introduction

Modern, interactive decision support systems like recommender systems or configurators often handle a very large set of possible decisions/alternatives. The task of finding the alternatives that best suit their preferences can be challenging for users, but the system can guide them towards their optimal decision if it has some knowledge of their preferences. In many settings, the users' preferences are not known in advance. This is especially the case with systems that enable anonymous users to browse the catalogues: such systems must be able to acquire users' preferences.

Preference learning has emerged as an important field; many interesting results are reported in, e.g., the book edited by Fürnkranz and Hüllermeier [2011], or the proceedings of recent Preference Learning or DA2PL (Decision Aid to Preference Learning) workshops. A general problem is: given a set of observed preferences, induce a model of preferences that best explains these observations within a certain class of models. As input, it is often assumed that the observed preferences are given as a set of pairwise comparisons or partial rankings of alternatives [Joachims, 2002]; or can be elicited online by asking the user to choose between two alternatives [Viappiani et al., 2006, Koriche and Zanuttini, 2009].

But in some circumstances, such input is not available. This is especially the case with some anonymous on-line configurators, where little information is stored about interactions. However, e-commerce companies generally keep a history of past sales. Users have chosen sold items, so they are probably ranked high in their preferences, but not necessarily at the very top. Indeed, a user may eventually choose an item that is not the optimal one in her preference order. For instance, because of the difficulty in grasping all possible options, a phenomenon called "mass confusion" [Huffman and Kahn, 1998], because of the influence of an advertisement, or because her preferred item is unavailable. Yet, this list of highly ranked items does provide information about the users' preferences.

Fargier et al. [2018] propose a model for learning preferences in such settings: the learning algorithm receives a multiset of alternatives that past users have chosen and induces a ranking of the alternatives to guide future users in their exploration of the set of possible alternatives. If, for instance, the colour red appears more often in the sales history than the colour yellow, then we want to induce a model that ranks alternatives with the colour red higher than alternatives with the colour yellow – maybe in association with some other criteria. This is similar to the usual, unsupervised setting when learning Bayesian networks [see e.g. Neapolitan, 2003]. This unsupervised setting is convenient from a machine learning point of view, as data is usually easy to obtain and does not necessitate any input labeling. The alternatives that users in past sessions have chosen are called "positive examples" in [Fargier et al., 2018], as opposed to possible "negative examples" that one could have in a setting where we would also have information about alternatives *rejected* by past users.

Research on the representation and learning of preferences has brought forward several types of models. Numerical models, like linear ranking functions or additive utilities [Joachims, 2002, Freund et al., 2003, Schiex et al., 1995, Gonzales and Perny, 2004, Braziunas, 2005], are rich families of models, especially if one allows high-dimensional feature spaces. Research in Artificial Intelligence has also brought forward ordinal models, like CP-nets [Boutilier et al., 2004a] and several extensions or variants. Lexicographic preferences are another family of ordinal models. This kind of preference is based on the importance of the attributes: when comparing two outcomes, their values for the most important attribute

are compared; if the two outcomes have different values for that attribute, then the one with the preferred value is deemed preferable to the other; otherwise one looks at the next most important attribute, and so on. This model can be extended by allowing the preferences on the values of an issue to depend on the values of more important ones. The relative importance of issues is no longer a linear order, but a “lexicographic preference” tree [Fraser, 1993, 1994, Brewka et al., 2006, Wallace and Wilson, 2009, Booth et al., 2010].

Lexicographic preference trees have several advantages over other preference representation models. First, this is an ordinal model, which is sufficient to represent a ranking of alternatives. Furthermore, they are generally an accurate representation of human behaviours [Gigerenzer and Goldstein, 1996]. Finally, one can quickly (in polytime) perform some interesting requests for recommendations, such as finding an optimal object or an optimal value for some attribute [Fargier and Mengin, 2018].

Learning lexicographic preference models with pairwise comparisons as inputs has been studied by e.g. Schmitt and Martignon [2006], Dombi et al. [2007], Yaman et al. [2008], while Booth et al. [2010], Bräuning and Hüllermeier [2012], Bräuning et al. [2017], Liu and Truszczynski [2015] studied learning of lexicographic preference trees. More recently, Fargier et al. [2018] have proposed a greedy algorithm for learning lexicographic preference trees from sales history. They reported some promising results on experimentation on both synthetic data and an industrial dataset from a car manufacturer (Renault), using clustering in a pre-processing step and a pruning pass in a post-processing step.

This paper proposes a complexity theoretical analysis of the approach proposed by Fargier et al. [2018]. We derive an upper bound on the sample complexity of computing the optimal LP-tree w.r.t. a given sample of chosen alternatives. This complexity is, in particular, logarithmic in the number of attributes. We also prove that computing this optimal LP-tree can be done in polynomial time when restricted to the class of “linear” LP-trees (which correspond to usual lexicographic preferences). Finally, we propose an algorithm for computing it for the more expressive class of LP lists [Bräuning et al., 2017], where several attributes can be at the same importance level.

The paper is structured as follows. The next section recall some background on combinatorial domains, LP-trees and the learning model introduced by Fargier et al. [2018]. In section 3 we derive some results on the computation of the rank of alternatives w.r.t some LP-tree. The following three sections are devoted to three classes of LP-trees, in order of increasing generality.

2 Background and notations

2.1 Combinatorial Domain

We consider a combinatorial domain over a finite set \mathcal{X} discrete attributes that characterise the possible alternatives, each attribute $X \in \mathcal{X}$ having a finite set of possible values \underline{X} ; we assume that $|\underline{X}| \geq 2$ for every $X \in \mathcal{X}$; then $\underline{\mathcal{X}}$ denotes the Cartesian product of the domains of the attributes in \mathcal{X} , its elements are called alternatives; we often use the sym-

bols o, o', o_1, o_2, \dots to denote alternatives. In the sequel, n is the number of attributes in \mathcal{X} , and d is a bound on the size of the domains of the attributes: for every $X \in \mathcal{X}$, $|\underline{X}| \leq d$.

For a subset U of \mathcal{X} , we will denote by \underline{U} the cartesian product of the domains of the attributes in U , every $u \in \underline{U}$ is an instantiation of U , or partial instantiation (of \mathcal{X}). If v is an instantiation of some $V \subseteq \mathcal{X}$, $v[U]$ denotes the restriction of v to the attributes in $V \cap U$; we say that instantiation $u \in \underline{U}$ and v are compatible if $v[U \cap V] = u[U \cap V]$; if $U \subseteq V$ and $v[U] = u$, we say that v extends u .

Given a partial instantiation u , $\text{Var}(u)$ denotes the set of attributes, the values of which appear in u .

2.2 Preference relation

In this paper, a preference relation is a linear order over $\underline{\mathcal{X}}$, that is, a total, transitive, irreflexive binary relation over $\underline{\mathcal{X}}$, often denoted with curly symbol \succ . For alternatives $o, o' \in \underline{\mathcal{X}}$, $o \succ o'$ indicates that o is strictly more preferred to o' .

Because we consider linear orders over $\underline{\mathcal{X}}$, we can define the *rank* of $o \in \underline{\mathcal{X}}$ w.r.t. \succ : $\text{rank}(\succ, o) = 1 +$ the number of outcomes strictly preferred to o , so that the most preferred outcome has rank 1, the least preferred has rank $|\underline{\mathcal{X}}|$:

$$\text{rank}(\succ, o) = 1 + |\{o' \in \underline{\mathcal{X}} \mid o' \succ o\}|.$$

2.3 Learning model

We consider an unknown target linear order \succsim over $\underline{\mathcal{X}}$, representing the preferences of a decision maker, or of a group of decision makers: whenever these decision makers have to make a decision, they choose an alternative according to \succsim ; these decisions are not always the “top” alternative, maybe because of the difficulty of finding it, or because of some context making it not available for instance. However, we consider that there is some probability distribution p of drawing alternatives of $\underline{\mathcal{X}}$, unknown but supposed to be decreasing w.r.t. to $\text{rank}(\succsim, \cdot)$: if $o \succ o'$, then $p(o) \geq p(o')$ and $\text{rank}(o) < \text{rank}(o')$.

We want to learn a (representation of) a linear order \succ that is as close as possible to \succsim , that can be used to give good answers to queries about \succsim (for instance: “what is the optimal outcome?”, or “is o preferred to o' ?”)

In order to have a relevant measure of how close \succ is to \succsim , Fargier et al. [2018] introduce the notion of *ranking loss*, defined as the normalised difference between the expected ranks of the two relations according to the ground probability p :

$$\begin{aligned} \text{rloss}_p(\succ, \succsim) &= \frac{1}{|\underline{\mathcal{X}}|} (E_p[\text{rank}(\succ, \cdot)] - E_p[\text{rank}(\succsim, \cdot)]) \quad (1) \end{aligned}$$

$$= \frac{1}{|\underline{\mathcal{X}}|} \sum_{o \in \underline{\mathcal{X}}} p(o) (\text{rank}(\succ, o) - \text{rank}(\succsim, o)) \quad (2)$$

The aim of a learning process in this setting is to find, given an unknown target \succsim and an associated, unknown, probability distribution p , a linear order \succ that minimises this ranking loss. Eq. 1 indicates that this is equivalent to finding a linear order that has minimal expected rank.

Eq. 2 shows that this loss bears some similarity with the Spearman distance between linear orders; here however, the

contribution of the rank difference for each alternative o is weighted with its probability of being drawn. This is because we want to learn a model that orders more accurately preferred alternatives; it is less significant to make mistakes with the ordering of less preferred alternatives. Also, there will not be much information about the alternatives in the tail of the distribution, which are the less preferred one.

Proposition 1 (Fargier et al. 2018). *Let \succ and \succsim be two linear orders over \mathcal{X} and p a probability distribution decreasing w.r.t. $\text{rank}(\succ, \cdot)$. Then $0 \leq \text{rloss}_p(\succ, \succsim) < 1$. Furthermore, if p is strictly decreasing w.r.t. $\text{rank}(\succ, \cdot)$, then $\text{rloss}_p(\succ, \succsim) = 0$ iff $\succsim = \succ$.*

The target preference is unknown, and we try to learn it from some data. In our context, a sample \mathcal{S} is a multiset of alternatives, that may be, in the case of e-commerce for instance, a list of items that have been chosen by users of the system. Thus, \mathcal{S} is considered to be representative of the preference expressed with the target linear order: every alternative can have several occurrences in \mathcal{S} , and the higher the rank of an alternative o in \succ is, the more likely it is to find o in \mathcal{S} , and the higher $p(o)$ is, where p is the unknown, ground probability, supposed to be decreasing w.r.t. to $\text{rank}(\succ, \cdot)$. In other words, alternatives in \mathcal{S} are supposed to be drawn with replacement from \mathcal{X} according to p .

Since the target preference is unknown, we cannot measure the ranking loss of an induced preference. However, since minimising the ranking loss of the induced preference amounts to minimising the sum of the ranks of the outcomes weighted by their probabilities of being drawn, we aim to minimise it by minimising the empirical mean rank of the training sample \mathcal{S} . For alternative $o \in \mathcal{X}$, $m(\mathcal{S}, o)$ denotes the multiplicity of alternative o in \mathcal{S} , that is, the number of occurrences of o in \mathcal{S} . We can define $p_{\mathcal{S}}$ to be the empirical probability distribution over \mathcal{X} such that $p_{\mathcal{S}}(o) = m(\mathcal{S}, o)/|\mathcal{S}|$ (in particular, $p_{\mathcal{S}}(o) = 0$ if $o \notin \mathcal{S}$). Then the empirical mean rank is defined as follows:

$$\overline{\text{rank}}(\succ, \mathcal{S}) = \sum_{o \in \mathcal{S}} p_{\mathcal{S}}(o) \times \text{rank}(\succ, o) = E_{p_{\mathcal{S}}}[\text{rank}(\succ, \cdot)]$$

In the sequel, given \mathcal{S} , \succ^* denotes a linear order that minimises, often within a given class \mathcal{C} of linear orders that will be clear from the context, the empirical mean rank: $\succ^* \in \mathcal{C}$ and for every $\succ \in \mathcal{C}$, $\text{rank}(\succ^*, \mathcal{S}) \leq \text{rank}(\succ, \mathcal{S})$. Or, equivalently:

$$\succ^* = \underset{\succ \in \mathcal{C}}{\text{argmin}} \overline{\text{rank}}(\succ, \mathcal{S})$$

From a complexity point of view, two questions arise :

1. How many examples in \mathcal{S} guarantee that \succ^* is probably a good approximation of \succ ?
2. What is the time complexity of computing a representation of \succ^* ?

The first question can be made more specific in the PAC setting: given a target linear order \succ ; a sample of outcomes \mathcal{S} drawn from \mathcal{X} according to some probability distribution p supposed to be decreasing with $\text{rank}(\succ, \cdot)$; a class \mathcal{C} of linear orders; and two real numbers $0 < \delta, \epsilon < 1$, what is the minimal function $S(\mathcal{C}, n, \delta, \epsilon)$ such that

$$\text{if } |\mathcal{S}| > S(\mathcal{C}, n, \delta, \epsilon) \text{ then } \Pr(\text{rloss}_p(\succ^*, \succ) \leq \epsilon) \geq 1 - \delta$$


The function $S(\mathcal{C}, \cdot, \cdot, \cdot)$ is called the *sample complexity* of learning a linear order in the class \mathcal{C} from “positive” examples (as opposed to learning from a set of pairwise comparisons). Parameter ϵ is the approximation that is wanted, here a bound on the ranking loss, and δ specifies the probability with which we want to attain this approximation.

2.4 Lexicographic preference trees

In this paper, we study the complexity of learning a specific class of preference relations: the preference relations that can be represented with lexicographic preference trees, or LP-trees. LP-trees generalise lexicographic orders, which have been widely studied in decision making – see e.g. Fishburn [1974]. As an inference mechanism, they are equivalent to search trees used by Boutilier et al. [2004b], and formalised by Wilson [2004, 2011]. As a preference representation, and elicitation, language, slightly different definitions for LP-trees have been proposed by Booth et al. [2010], Bräuning and Hüllermeier [2012], Fargier et al. [2018].

LP-trees provide a nice graphical representation of the corresponding preference relation. We illustrate that on an example before giving the formal definition that we use in this paper.

Example 1. (Example A in Wilson [2011], slightly extended) I am planning a holiday, with three choices / attributes: wait til next month ($W = w$) or leave now ($W = \bar{w}$), going to city 1, 2 or 3 ($C = c_1, C = c_2$ or $C = c_3$), travelling by plane ($P = p$) or by car ($P = \bar{p}$). The picture below shows an LP-tree φ_0 over $\mathcal{X} = \{W, C, P\}$ which defines a linear order \succ over \mathcal{X} as follows.



- The root of φ_0 is labelled with attribute W , meaning that this is the most important attribute in my decision. Associated to it is the *local preference* $\bar{w} > w$, that indicates that I would rather go now, irrespective of the other attributes: given two alternatives o and o' such that $o[W] = \bar{w}$ and $o'[W] = w$, $o \succ o'$.
- On the right branch, the next node is labelled with P ; it is connected to its parent with an edge labelled with w , and has the local preference $\bar{p} > p$: it indicates that if I go later (w) then the second most important attribute is the means of transport, and that I'd rather avoid flying; so whatever the cities c_i and c_j , $w\bar{p}c_i \succ wpc_j$.
- Still on the right branch, the leaf is labelled with C and has associated local preference $c_3 > c_1 > c_2$, indicating that, still when going later, if I must choose between alternatives that have the same means of transport, I prefer the alternative that has my most preferred destination c_3 , whereas my least preferred destination is c_2 ; so for instance $wpc_3 \succ wpc_1 \succ wpc_2$.

- The only node below the root on the left branch is labelled with the pair of attributes CP . It is connected to its parent with an edge labelled with \bar{w} , and has an associated local preference that linearly orders the cartesian product of the domains of C and P . It indicates that, if I go now (\bar{w}), I'd rather fly to c_1 than drive to c_3 , but drive to c_2 rather than fly to c_1 : $\bar{w}pc_1 \succ \bar{w}pc_3$ and $\bar{w}pc_2 \succ \bar{w}pc_1$.

Now, because the root indicates that W is the most important attribute and that $\bar{w} > w$, every alternative that has $W = \bar{w}$ will be preferred to every alternative that has $W = w$. So the most preferred alternative is $\bar{w}c_3p$, whereas the least preferred one is $w\bar{p}c_3$. And $w\bar{p}c_2$ is preferred to wpc_3 , because both alternatives have equal value for the attribute at the root ($W = \bar{w}$ for both), and the first node where they have differing values is the one labelled with P , and $\bar{p} > p$ at that node.

Definition 1. An LP-tree φ over \mathcal{X} is a rooted tree with labelled nodes and edges, and a set of local preference tables; specifically

- every node N is labelled with a set of attributes, denoted $\text{Var}(N)$;
- if N is not a leaf, it can have one child, or $|\text{Var}(N)|$ children;
- in the latter case, the edges that connect N to its children are labelled with the instantiations in $\text{Var}(N)$;
- if N has one child only, the edge that connects N to its child is not labelled: all instantiations in $\text{Var}(N)$ lead to the same subtree;
- a local preference table $\text{CPT}(N)$ is associated with N , it specifies a linear order over $\text{Var}(N)$;
- moreover, every attribute must appear exactly once on every branch of φ .

We denote by $\text{Anc}(N)$ the set of *ancestors* of N : the attributes that appear in the nodes between the root and N (excluding those at N), and by $\text{Inst}(N)$ (resp. $\text{NonInst}(N)$) the set of attributes that appear in the nodes above N that have more than one children (resp. only one child). Finally, we denote by $\text{Desc}(N)$ the set of *descendants* of N : the attributes that appear below N ; thus $(\text{Anc}(N), \text{Var}(N), \text{Desc}(N))$ is a partition of \mathcal{X} .

Given an LP-tree φ and an alternative $o \in \mathcal{X}$, there is a unique way to traverse the tree, starting at the root, and along edges that are either not labelled, or labelled with instantiations that agree with o , until a leaf is reached.

LP-tree φ defines a linear order \succ_φ over \mathcal{X} as follows: given two distinct alternatives o, o' , it is possible to traverse the tree along edges that correspond to o and o' as long as o and o' agree, until a node N is reached which is labelled with some W such that $o[W] \neq o'[W]$: we say that N decides $\{o, o'\}$, and $o \succ_\varphi o'$ if and only if $o[W] > o'[W]$ in the linear order in $\text{CPT}(N)$. In the sequel, slightly abusing notations for ease of readability, we will often write φ where \succ_φ would be expected: for instance, we will denote by $\text{rank}(\varphi, o)$ the rank of alternative o in the linear order \succ_φ ; and write $\text{rloss}_p(\varphi, \check{\varphi})$

the ranking loss of \succ_φ with respect to target linear order $\check{\succ}$ if $\check{\varphi}$ is an LP-tree that represents $\check{\succ}$.

Definition 1 corresponds to the k -LP-trees of Fargier et al. [2018]. It is more general than the definition of LP-trees of Bräuning and Hüllermeier [2012] because it allows for nodes with one child only, and is more general than the definition of Booth et al. [2010] in that it allows for more than one attribute at every node. (However, Booth et al. [2010] allow *incomplete* LP-trees, where some attributes may be missing on some branch, leading to partial orders over \mathcal{X} .) Every linear order can be represented with an LP-tree as defined above, possibly with a single node that contains all attributes at the root.

Fargier et al. [2018] give a greedy, top-down unsupervised algorithm to learn such an LP-tree from a multiset of positive examples, and describe some experiments.

3 Computing the rank expectation

Many results in the sequel crucially rely on the possibility to decompose the rank of any alternative o in (the order represented by) any LP-tree. In this section, we slightly modify the decomposition given by Lang et al. [2018] in order to compute the rank expectation over all alternatives. We illustrate this on an example first, and then give the general formula.

Example 2. Consider the LP-tree φ_0 of Example 1, and suppose that we want to compute the rank of alternative wc_1p : this amounts to counting the number of alternatives that are “to its left” in φ_0 . wc_1p is less preferred than all alternative that have $W = \bar{w}$, there are $|\underline{CP}| = 3 \times 2 = 6$ of them. Alternatives that have $W = w$ and $P = \bar{p}$ are also preferred to wc_1p , there are $|\underline{C}| = 3$ of them. wc_1p is also less preferred than wc_3p . Finally $\text{rank}(\varphi_0, wc_1p) = 1 + (6 + 3 + 1) = 11$.

In general, the rank of any alternative can be decomposed as (one plus) a sum of contributions at every node: the root of φ_0 contributes 6 to the rank of wc_1p , the left-most leaf labelled CP contributes 0 because it is not on the branch that corresponds to wc_1p , the node labelled P contributes 3, and the node labelled C contributes 1.

More generally, given LP-tree φ and alternative o , it can be shown that:

$$\text{rank}(\varphi, o) = 1 + \sum_{N \in \text{nodes}(\varphi)} \llbracket o[\text{Inst}(N)] = \text{inst}(N) \rrbracket \quad (3)$$

$$\times (r(\text{Var}(N) >_N, o[\text{Var}(N)]) - 1) \times |\underline{\text{Desc}(N)}|$$

where :

- $\text{nodes}(\varphi)$ denotes the set of nodes of LP-tree φ ;
- $\llbracket o[\text{Inst}(N)] = \text{inst}(N) \rrbracket$ is an indicator function, that equals 1 when the condition $o[\text{Inst}(N)] = \text{inst}(N)$ is true; that is, when N is on the branch of φ that corresponds to o ; and equals 0 otherwise;
- $>_N$ is the linear order over $\text{Var}(N)$ specified in $\text{CPT}(N)$;
- $r(\text{Var}(N), >_N, o[\text{Var}(N)])$ denotes the rank in $\text{Var}(N)$ with respect to $>_N$ of the instantiation given by o to $\text{Var}(N)$; so that $r(\text{Var}(N), >_N, o[\text{Var}(N)]) - 1$ is the

number of subtrees rooted at children of N that are less preferred than o at N ;

- $\text{Desc}(N) = \mathcal{X} - (\text{Anc}(N) \cup \text{Var}(N))$ is the set of attributes that appear below N in that branch, so that $|\text{Desc}(N)|$ is the number of instantiations that are “contained” in every subtree of φ rooted at a child of N .

The difference with the decomposition of Lang et al. [2018] is that we sum over all nodes of φ , irrespective of the alternative. This is useful to express the expectation of the rank of φ with respect to some probability distribution p over \mathcal{X} . For set of attributes $V \subseteq \mathcal{X}$, linear order $>$ over V , and $v \in V$, let $r(V, >, v)$ be the rank of v in V with respect to $>$. Then the expectation of this rank, w.r.t. p , is

$$E_p[r(V, >, \cdot)] = \sum_{v \in V} p(v)r(V, >, v)$$

where $p(v)$ denotes the probability of drawing an alternative o such that $o[V] = v$.

Proposition 2. For LP-tree φ and probability distribution p :

$$E_p[\text{rank}(\varphi, \cdot)] = 1 + \sum_{N \in \text{nodes}(\varphi)} \frac{|\text{Desc}(N)|}{|\underline{\mathcal{X}}|} \times p(\text{inst}(N)) \times E_{p|_{\text{inst}(N)}}[r(\text{Var}(N), >_N, \cdot) - 1] \quad (4)$$

where $p|_{\text{inst}(N)}$ denotes the probability distribution marginalized to $\text{inst}(N)$, that is: $p|_{\text{inst}(N)}(v) = p(v|_{\text{inst}(N)})$.

Proof sketch. By definition, $E_p[\text{rank}(\varphi, \cdot)] = \sum_{o \in \underline{\mathcal{X}}} (p(o) \times \text{rank}(\varphi, o))$. Equation 3 shows that $\text{rank}(\varphi, o)$ can be decomposed as a sum of contributions over all nodes of φ . It is not difficult to see that it is possible to invert the sum over $\underline{\mathcal{X}}$ and the sum over nodes(φ), which yields the result. \square

Locally optimal LP-trees Proposition 2 above suggests that, given some LP-tree φ , and a probability distribution p , one can reduce the rank expectation of φ by re-ordering the values at every node N , so as to minimize $E_{p|_{\text{inst}(N)}}[r(\text{Var}(N), >_N, \cdot)]$: essentially, one only has to order the instantiations of $\text{Var}(N)$ in order of non-increasing probability. Thus we have here a desirable property shared by many graphical models of preferences, namely that learning can often be decomposed in two parts: 1) learning the structure, often a hard problem; and 2) learning the *local* preferences, often an easy task.

Definition 2. Given a probability distribution p over $\underline{\mathcal{X}}$, we say that LP-tree φ is *p-locally optimal* if at every node N of φ , the linear order $>_N$ orders the instantiations of $\text{Var}(N)$ in order of non-increasing probability.

In the sequel, given set of attributes $V \subseteq \mathcal{X}$, and probability distribution p over $\underline{\mathcal{X}}$, we denote by $E_p^*(V)$ the minimum expectation that is possible for the ranks of the instantiations of V given p , that is, when ordering these values in order of non-increasing probability p :

$$E_p^*(V) = \underset{>}{\text{argmin}} E_p[r(V, >, \cdot)]$$

In the next sections, we study, for different classes of LP-trees, how difficult it is to compute an LP-tree that has minimum mean rank with respect to sample \mathcal{S} , and how many alternatives in \mathcal{S} guarantee that this \mathcal{S} -optimal LP-tree is close to the target one.

4 Linear lexicographic preference trees with univariate nodes

We consider in this section the class of *linear* LP-trees, that is, LP trees that have a single branch; and also such that every node is labelled with a single attribute. Let us denote by $\text{LPT}_{\text{lin}}^1$ this class. It is in fact the class of the usual lexicographic preference relations Fishburn [1974].

4.1 Greedy algorithm for the optimal LP-tree

Note first that at every node N of some $\varphi \in \text{LPT}_{\text{lin}}^1$, $\text{Inst}(N)$ is empty, so $\text{inst}(N)$ is the empty instantiation, and $p|_{\text{inst}(N)} = p$. Thus :

$$E_p[\text{rank}(\varphi, \cdot)] = 1 + \sum_{N \in \text{nodes}(\varphi)} \frac{|\text{Desc}(N)|}{|\underline{\mathcal{X}}|} \times E_p[r(\text{Var}(N), >_N, \cdot) - 1]$$

We define, for every attribute $X \in \mathcal{X}$ and probability distribution p over $\underline{\mathcal{X}}$:

$$\text{Score}(p, X) = E_p^*(X) / (|\underline{\mathcal{X}}| - 1)$$

Proposition 3. Let p be a probability distribution over $\underline{\mathcal{X}}$, the LP-tree that has minimal expected rank in $\text{LPT}_{\text{lin}}^1$ has the attributes in order of non-decreasing values for $\text{Score}(p, \cdot)$ along its single branch from the root down to its single leaf, and is locally optimal.

Proof. Let φ be a locally optimal tree of $\text{LPT}_{\text{lin}}^1$. Suppose that there are two nodes of φ , N and N' , labelled with X and Y respectively, such that N' is the child of N but $\text{Score}(p, X) > \text{Score}(p, Y)$. Let φ' be identical to φ except that X and Y have been inverted: X is at N' in φ' , and Y at N . Then $\text{Desc}(N) = \text{Desc}(N') \cup \text{Var}(N')$, where $\text{Var}(N') = \{Y\}$ in φ , and $\text{Var}(N) = \{X\}$ in φ' . Thus:

$$E_p[\text{rank}(\varphi, \cdot)] - E_p[\text{rank}(\varphi', \cdot)] = \frac{|\text{Desc}(N')|}{|\underline{\mathcal{X}}|} \times (E_p^*(X)(|Y| - 1) - E_p^*(Y)(|X| - 1))$$

On the other hand:

$$\begin{aligned} \text{Score}(p, X) &> \text{Score}(p, Y) \\ \Leftrightarrow \frac{E_p^*(X)}{|\underline{\mathcal{X}}| - 1} - \frac{E_p^*(Y)}{|\underline{\mathcal{Y}}| - 1} &> 0 \\ \Leftrightarrow E_p^*(X) \times (|Y| - 1) - E_p^*(Y) \times (|X| - 1) &> 0 \\ \Leftrightarrow E_p[\text{rank}(\varphi, \cdot)] - E_p[\text{rank}(\varphi', \cdot)] &> 0 \end{aligned}$$

This shows that, by applying a kind of *bubble sort* to the nodes of φ so as to order them in order of non-decreasing score, we obtain a new tree that has better expected rank. \square

As a consequence:

Corollary 4. Let $\check{\varphi}$ be some target LP-tree in LPT_{lin}^1 and p some probability distribution over \mathcal{X} non-decreasing w.r.t. $\text{rank}(\check{\varphi}, \cdot)$, then the attributes appear in $\check{\varphi}$ in order of non-decreasing values for $\text{Score}(p, \cdot)$ from the root down to the leaf.

Corollary 5. Let $\mathcal{S} \subseteq \mathcal{X}$, recall that $p_{\mathcal{S}}$ denotes the probability distribution over \mathcal{X} such that $p_{\mathcal{S}}(o)$ is the frequency of o in \mathcal{S} . Let φ^* be the LP-tree in LPT_{lin}^1 that has minimal mean rank w.r.t. \mathcal{S} , then the attributes appear in φ^* in order of non-decreasing values for $\text{Score}(p_{\mathcal{S}}, \cdot)$ from the root down to the leaf.

4.2 Time and sample complexity

According to Corollary 4 above, in order to compute φ^* , we can, for every $X \in \mathcal{X}$, compute the number of occurrences in \mathcal{S} of every $x \in \underline{X}$, and compute $>_X^{p_{\mathcal{S}}}$ and $\text{Score}(p_{\mathcal{S}}, X)$, and then order the attributes w.r.t. $\text{Score}(p_{\mathcal{S}}, \cdot)$. Thus:

Proposition 6. Given sample $\mathcal{S} \subseteq \mathcal{X}$, computing φ^* in LPT_{lin}^1 can be done in time in $O(n \log n |\mathcal{S}| d \log d)$, where d is an upper bound on the size of domains of each attribute $X \in \mathcal{X}$.

Proposition 7. $S(LPT_{\text{lin}}^1, n, \delta, \epsilon) = O(\frac{1}{\epsilon^2} (\ln \frac{1}{\delta} + \ln d + \ln(n+1)))d^4$.

Proof. This is simple consequence of the more general Proposition 10 given in section 6 for the sample complexity of learning LP-trees with bounded number of leaves and bounded number of attributes at every node: here both bounds equal 1. \square

5 Linear lexicographic preference trees with multivariate nodes

We consider now the more general settings of the class LPT_{lin}^k , for some fixed $k \leq n$: linear LP-trees (with a single leaf), where each node can be labelled with a non-empty set of no more than k attributes.

The score defined in the case of univariate nodes in the previous section can easily be extended to multivariate nodes: for $V \subseteq \mathcal{X}$ and probability distribution p over \mathcal{X} :

$$\text{Score}(p, V) = E_p^*[V] / (|V| - 1)$$

Proposition 3 can be generalised in this settings as follows:

Proposition 8. Let p be a probability distribution over \mathcal{X} , let \mathcal{P} be a partition of \mathcal{X} . The LP-tree that has minimal expected rank, within the class of LP-trees with a single leaf, and where each node is labelled with an element of \mathcal{P} , is locally optimal and has the set of attributes that label its nodes in order of non-increasing values for $\text{Score}(p, \cdot)$ along its single branch from the root down to its single leaf.

Proof. Given partition \mathcal{P} of \mathcal{X} , consider each part $P \in \mathcal{P}$ as a new attribute with domain \underline{P} , Prop. 3 gives the result. \square

Based on this, algorithm 1 computes the tree in LPT_{lin}^k that minimises the empirical mean rank with respect to some sample \mathcal{S} : it enumerates the k -partitions of \mathcal{X} , i.e. the partitions

such that no part has more than k attributes; and, for every partition, orders the parts in order of non-increasing score and compute the empirical mean rank of the tree thus obtained.

Algorithm 1 Optimal linear multivariate LP-tree

Input : sample \mathcal{S} ;

1. $r^* \leftarrow +\infty$;
 2. for every k -partition \mathcal{P} of \mathcal{X} do:
 - a. order the parts in \mathcal{P} in order of non-decreasing $\text{Score}(p_{\mathcal{S}}, \cdot)$;
 - b. $\varphi \leftarrow$ the optimal linear univariate LP-tree over set of attributes \mathcal{P} ;
 - c. if $\overline{\text{rank}}(\varphi, \mathcal{S}) < r^*$: $r^* \leftarrow \overline{\text{rank}}(\varphi, \mathcal{S})$ and $\varphi^* \leftarrow \varphi$;
 3. return φ^* .
-

Proposition 9. For fixed k , the time complexity of Algorithm 1 is in $\Omega(k^n n (\log n + |\mathcal{S}| d^k \log d))$ and in $O(k^{n+1} n (\log n + |\mathcal{S}| d^k \log d))$.

Proof. The number of k -partitions of \mathcal{X} is also the number of partitions of \mathcal{X} into at most k non empty parts (in the same way as the number of partitions of integer n into positive integers no greater than k is equal to the number of partitions of n into k non-negative integers). Let $S(n, k)$ denote the number of partitions of a set of n elements into exactly k non-empty parts, it is known as the *Stirling number of the second kind*, an asymptotic approximation is $S(n, k) \sim k^n/k!$ [e.g. Knuth, 1997]. Therefore, for fixed k , the number of k -partitions of \mathcal{X} is asymptotically equivalent to $\sum_{i=1}^k i^n/i!$, which is $\geq k^n/k!$. We also have that $i^n/i! < j^n/j!$ for $i < j$, thus $\sum_{i=1}^k i^n/i! < k \times k^n/k!$. So the number of k -partitions is in $\Theta(k^n)$.

For a given partition, at step 2.a., one has to compute the scores of the parts, and order them. There are no more than n parts; each part has no more than d^k possible instantiations, so computing and ordering the numbers of occurrences of for all instantiations takes worse time in $\Theta(n |\mathcal{S}| d^k \log(d^k))$. Ordering the parts then takes time in $\Theta(n \log n)$. \square

6 Lexicographic preference trees with bounded number of leaves

We now turn to the class of LP-trees with a bounded number of attributes at each node, and a bounded number of leaves. Specifically, LPT_l^k denotes the set of LP-trees that have no more than k attributes at each node, and no more than l leaves. (In particular, when $l = 1$, $LPT_1^k = LPT_{\text{lin}}^k$)

Proposition 10. $S(LPT_l^k, n, \delta, \epsilon) = O(\frac{1}{\epsilon^2} (\ln \frac{1}{\delta} + k(\ln d + \ln(n+1))))l^2 d^{4k}$.

Proof. Recall that $\check{\varphi}$ is an LP-tree, supposed to be in LPT_l^k , that represents the target, unknown preference relation \succ ; p is a probability distribution that is non-increasing w.r.t. $\text{rank}(\check{\varphi}, \cdot)$; and φ^* denotes the LP-tree in LPT_l^k that has minimal normalised empirical mean rank with respect to some

sample \mathcal{S} . Then:

$$\begin{aligned} \text{rloss}(\varphi^*, \check{\varphi}) &= \frac{1}{|\underline{\mathcal{X}}|} (E_p[\text{rank}(\varphi^*, \cdot)] - E_p[\text{rank}(\check{\varphi}, \cdot)]) \\ &\leq \frac{1}{|\underline{\mathcal{X}}|} (|E_p[\text{rank}(\varphi^*, \cdot)] - E_{p_{\mathcal{S}}}[\text{rank}(\varphi^*, \cdot)]| \\ &\quad + E_{p_{\mathcal{S}}}[\text{rank}(\varphi^*, \cdot)] - E_{p_{\mathcal{S}}}[\text{rank}(\check{\varphi}, \cdot)] \\ &\quad + |E_{p_{\mathcal{S}}}[\text{rank}(\check{\varphi}, \cdot)] - E_p[\text{rank}(\check{\varphi}, \cdot)]|) \\ &\leq \frac{2}{|\underline{\mathcal{X}}|} \max_{\varphi \in \text{LPT}_l^k} |E_p[\text{rank}(\varphi, \cdot)] - E_{p_{\mathcal{S}}}[\text{rank}(\varphi, \cdot)]| \end{aligned}$$

because, by definition of φ^* :

$$E_{p_{\mathcal{S}}}[\text{rank}(\varphi^*, \cdot)] - E_{p_{\mathcal{S}}}[\text{rank}(\check{\varphi}, \cdot)] \leq 0.$$

Now, for $\varphi \in \text{LPT}_l^k$:

$$\begin{aligned} &|E_p[\text{rank}(\varphi, \cdot)] - E_{p_{\mathcal{S}}}[\text{rank}(\varphi, \cdot)]| \\ &\leq \sum_{N \in \text{nodes}(\varphi)} \left(|\text{Desc}(N)| \times p(\text{inst}(N)) \right. \\ &\quad \times |E_p[\text{inst}(N)[r(\text{Var}(N), >_N, \cdot)]] \\ &\quad \left. - E_{p_{\mathcal{S}}}[\text{inst}(N)[r(\text{Var}(N), >_N, \cdot)]] \right) \\ &\leq \sum_{N \in \text{nodes}(\varphi)} \left(|\text{Desc}(N)| \times \sum_{v \in \text{Var}(N)} (r(v) \right. \\ &\quad \left. \times |p(v \wedge \text{inst}(N)) - p_{\mathcal{S}}(v \wedge \text{inst}(N))| \right) \\ &\leq M \times \frac{d^k(d^k + 1)}{2} \times \sum_{N \in \text{nodes}(\varphi)} |\text{Desc}(N)| \end{aligned}$$

where d is a bound on the domain size of the attributes, and M is an upper bound on $|p(v \wedge \text{inst}(N)) - p_{\mathcal{S}}(v \wedge \text{inst}(N))|$ for every $N \in \text{nodes}(\varphi)$, every $v \in \text{Var}(N)$.

Now consider a single branch ψ of φ , let V_1, V_2, \dots, V_t be the set of attributes that label its t nodes from the root down to its leaf (at unknown depth t), then:

$$\begin{aligned} &\sum_{N \in \text{nodes}(\psi)} |\text{Desc}(N)| \\ &= (|V_2 \times V_3 \times \dots \times V_t| + \dots + |V_{t-1} \times V_t| + |V_t| + 0) \\ &= |\underline{\mathcal{X}}| \times \left(\frac{1}{|V_1|} + \frac{1}{|V_1 \times V_2|} + \dots + \frac{1}{|V_1 \times V_2 \times \dots \times V_{t-1}|} \right) \\ &\leq |\underline{\mathcal{X}}| \end{aligned}$$

because for every node, $|\text{Var}(N)| \geq 2$. Therefore, since we consider LP-trees with a bounded number of leaves (and branches), for every $\varphi \in \text{LPT}_l^k$:

$$|E_p[\text{rank}(\varphi, \cdot)] - E_{p_{\mathcal{S}}}[\text{rank}(\varphi, \cdot)]| \leq M \times \frac{d^k(d^k + 1)}{2} \times l \times |\underline{\mathcal{X}}|$$

and $\text{rloss}_p(\varphi^*, \check{\varphi}) \leq l \times d^k(d^k + 1) \times \max_{V \subseteq \underline{\mathcal{X}}, v \in V} |p(v) - p_{\mathcal{S}}(v)|$.

Thus, if $\text{rloss}_p(\varphi^*, \check{\varphi}) \geq \epsilon$, there must be some $V \subseteq \underline{\mathcal{X}}$ and $v \in V$ such that $|p(v) - p_{\mathcal{S}}(v)| \geq \epsilon / (l \times d^k(d^k + 1))$, which implies that:

$$\begin{aligned} &Pr(\text{rloss}(\varphi^*, \check{\varphi}) \geq \epsilon) \\ &\leq Pr\left(\bigcup_{\substack{V \subseteq \underline{\mathcal{X}} \\ v \in V}} |p(v) - p_{\mathcal{S}}(v)| \geq \epsilon / (l \times d^k(d^k + 1))\right) \\ &\leq \sum_{\substack{V \subseteq \underline{\mathcal{X}} \\ v \in V}} Pr(|p(v) - p_{\mathcal{S}}(v)| \geq \epsilon / (l \times d^k(d^k + 1))) \end{aligned}$$

For every $V \subseteq \underline{\mathcal{X}}$ and every $v \in V$, $p_{\mathcal{S}}(v)$ is an estimate, from sample \mathcal{S} , of the ground probability $p(v)$ of drawing an alternative o such that $o[V] = v$. Hoeffding's inequality states that for every $\alpha > 0$:

$$Pr(|p(v) - p_{\mathcal{S}}(v)| \geq \alpha) \leq e^{-2|\mathcal{S}|\alpha^2}$$

For every $i \in \{1, \dots, k\}$, there are $\binom{n}{i}$ ways of choosing a subset V of $\underline{\mathcal{X}}$ of cardinality i , then $|V| \leq d^i$; therefore:

$$\begin{aligned} &Pr(\text{rloss}(\varphi^*, \check{\varphi}) \geq \epsilon) \\ &\leq \left(\sum_{i=1}^k \binom{n}{i} d^i \right) \exp(-2|\mathcal{S}|(\epsilon / (l \times d^k(d^k + 1)))^2) \\ &\leq d^k(1+n)^k \exp(-2|\mathcal{S}|(\epsilon / (l \times d^k(d^k + 1)))^2) \end{aligned}$$

Therefore, in order to have $Pr(\text{rloss}(\varphi^*, \check{\varphi}) \leq \epsilon) \geq 1 - \delta$, it is sufficient to have $1 - d^k(1+n)^k \exp(-2|\mathcal{S}|(\epsilon / (l \times d^k(d^k + 1)))^2) \geq 1 - \delta$, which is equivalent to:

$$|\mathcal{S}| \geq (k(\ln d + \ln(n+1))) + \ln \frac{1}{\delta} \frac{(ld^k(d^k + 1))^2}{2\epsilon^2} \quad \square$$

7 Conclusion

The bound on the sample complexity given in the last section shows several interesting properties of the problem of unsupervised learning LP-trees from sales history. First, it is logarithmic in the number of attributes. Also the factor d^{4k} may, in general, be overly pessimistic, as it assumes that all attributes have the same domain size. A finer analysis would show that what counts is the largest domain size of any combination of k attributes. Note that k should be kept small: Bräuning et al. [2017], Fargier et al. [2018] report promising results on real datasets with $k < 4$; larger values of k would improve the expressiveness, but also greatly increase the model size and may lead to overfitting.

Concerning the time complexity, we conjecture that the problem of computing the LP-tree that has minimal empirical risk may be NP-hard in general, but that is an important question for future work. Another avenue for future investigation is a finer analysis of the stochastic process that leads to the probability distribution p .

Acknowledgements

We thank the reviewers for their interesting and helpful comments. The authors gratefully acknowledge the support of the Artificial and Natural Intelligence Toulouse Institute – ANITI. ANITI is funded by the French "Investing for the Future – PIA3" program under grant agreement ANR-19-PI3A-0004.

References

- Richard Booth, Yann Chevaleyre, Jérôme Lang, Jérôme Mengin, and Chattrakul Sombatheera. Learning conditionally lexicographic preference relations. In *Proceedings of ECAI'10*, pages 269–274, 2010.
- Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004a.
- Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. Preference-based constrained optimization with cp-nets. *Computational Intelligence*, 20(2):137–157, 2004b.
- Michael Bräuning and Eyke Hüllermeier. Learning conditional lexicographic preference trees. In Johannes Fürnkranz and Eyke Hüllermeier, editors, *Proceedings of ECAI'12 Workshop*, 2012.
- Michael Bräuning, Eyke Hüllermeier, Tobias Keller, and Martin Glaum. Lexicographic preferences for predictive modeling of human decision making: A new machine learning method with an application in accounting. *European Journal of Operational Research*, 258(1):295–306, 2017.
- Darius Braziunas. Local utility elicitation in GAI models. In *Proceedings of UAI'05*, pages 42–49, 2005.
- G Brewka et al. An efficient upper approximation for conditional preference. In *Proceedings of ECAI'06*, volume 141, page 472, 2006.
- József Dombi, Csanád Imreh, and Nándor Vincze. Learning lexicographic orders. *European Journal of Operational Research*, 183:748–756, 2007.
- Hélène Fargier and Jérôme Mengin. A knowledge compilation map for conditional preference statements-based languages. In *Proceedings AAMAS 2021*, pages 492–500, 2018.
- Hélène Fargier, Pierre Francois Gimenez, and Jérôme Mengin. Learning lexicographic preference trees from positive examples. In *Proceedings AAAI 2018*, pages 2959–2966. ACM, 2018.
- Peter C. Fishburn. Lexicographic orders, utilities and decision rules: A survey. *Management Science*, 20(11):1442–1471, 1974.
- Niall M Fraser. Applications of preference trees. In *Proceedings of SMC'93*, pages 132–136, 1993.
- Niall M. Fraser. Ordinal preference representations. *Theory and Decision*, 36(1):45–67, 1994.
- Yoav Freund, Raj D. Iyer, Robert E. Schapire, and Yoram Singer. An efficient boosting algorithm for combining preferences. *Journal of Machine Learning Research*, 4:933–969, 2003.
- Johannes Fürnkranz and Heike Hüllermeier, editors. *Preference learning*. Springer, 2011.
- Gerd Gigerenzer and Daniel G. Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669, 1996.
- Christophe Gonzales and Patrice Perny. GAI networks for utility elicitation. In *Proceedings of KR'04*, pages 224–234, 2004.
- Cynthia Huffman and Barbara E Kahn. Variety for sale: Mass customization or mass confusion? *Journal of retailing*, 74(4):491–513, 1998.
- Thorsten Joachims. Optimizing search engines using click-through data. In *Proceedings of SIGKDD'02*, pages 133–142, 2002.
- Donald E. Knuth. *The Art of Computer Programming, Volume I: Fundamental Algorithms*. Addison-Wesley, 3rd edition, 1997. ISBN 0201896834.
- Frédéric Koriche and Bruno Zanuttini. Learning conditional preference networks with queries. In *Proceedings of IJCAI'09*, pages 1930–1935, 2009.
- Jérôme Lang, Jérôme Mengin, and Lirong Xia. Voting on multi-issue domains with conditionally lexicographic preferences. *Artificial Intelligence*, 265:18–44, 2018.
- Xudong Liu and Miroslaw Truszczyński. Learning partial lexicographic preference trees over combinatorial domains. In *Proceedings of AAAI'15*, volume 15, pages 1539–1545, 2015.
- Richard E. Neapolitan. *Learning bayesian networks*. Prentice Hall, 2003.
- Thomas Schiex, Hélène Fargier, and Gérard Verfaillie. Valued constraint satisfaction problems: Hard and easy problems. In *Proceedings of IJCAI'95*, pages 631–637, 1995.
- Michael Schmitt and Laura Martignon. On the complexity of learning lexicographic strategies. *Journal of Machine Learning Research*, 7:55–83, 2006.
- Paolo Viappiani, Boi Faltings, and Pearl Pu. Preference-based search using example-critiquing with suggestions. *Journal of Artificial Intelligence Research*, 27:465–503, 2006.
- Richard J. Wallace and Nic Wilson. Conditional lexicographic orders in constraint satisfaction problems. *Annals of Operations Research*, 171(1):3–25, 2009.
- Nic Wilson. Consistency and constrained optimisation for conditional preferences. In *Proceedings ECAI 2004*, pages 888–892, 2004.
- Nic Wilson. Computational techniques for a simple theory of conditional preferences. *Artificial Intelligence*, 175:1053–1091, 2011.
- Fusun Yaman, Thomas J Walsh, Michael L Littman, and Marie Desjardins. Democratic approximation of lexicographic preference models. In *Proceedings of ICML'08*, pages 1200–1207, 2008.