Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Computation and Complexity of Preference Inference Based on Hierarchical Models

Nic Wilson, Anne-Marie George and Barry O'Sullivan
Insight Centre for Data Analytics, School of Computer Science and IT
University College Cork, Ireland
{nic.wilson, annemarie.george, barry.osullivan}@insight-centre.org,

Abstract

Preference Inference involves inferring additional user preferences from elicited or observed preferences, based on assumptions regarding the form of the user's preference relation. In this paper we consider a situation in which alternatives have an associated vector of costs, each component corresponding to a different criterion, and are compared using a kind of lexicographic order, similar to the way alternatives are compared in a Hierarchical Constraint Logic Programming model. It is assumed that the user has some (unknown) importance ordering on criteria, and that to compare two alternatives, firstly, the combined cost of each alternative with respect to the most important criteria are compared; only if these combined costs are equal, are the next most important criteria considered. The preference inference problem then consists of determining whether a preference statement can be inferred from a set of input preferences. We show that this problem is coNP-complete, even if one restricts the cardinality of the equal-importance sets to have at most two elements, and one only considers non-strict preferences. However, it is polynomial if it is assumed that the user's ordering of criteria is a total ordering; it is also polynomial if the sets of equally important criteria are all equivalence classes of a given fixed equivalence relation. We give an efficient polynomial algorithm for these cases, which also throws light on the structure of the inference.

1 Introduction

There are increasing opportunities for decision making/support systems to take into account the preferences of individual users, with the user preferences being elicited or observed from the user's behaviour. However, users tend to have limited patience for preference elicitation, so such a system will tend to have a very incomplete picture of the user preferences. *Preference Inference* involves inferring additional user preferences from elicited or observed preferences, based on assumptions regarding the form of the user's preference relation. More specifically, given a set of input preferences Γ, and a set of preference models $\mathcal{M}$ (considered as candidates for the user's preference model), we infer a preference statement φ if every model in $\mathcal{M}$ that satisfies Γ also satisfies φ. Preference Inference can take many forms, depending on the choice of $\mathcal{M}$, and on the choices of language(s) for the input and inferred statements. For instance, if we just assume that the user model is a total order (or total pre-order), we can set $\mathcal{M}$ as the set of total [pre-]orders over a set of alternatives. This leads to a relatively cautious form of inference (based on transitive closure), including, for instance, the dominance relation for CP-nets and some related systems, e.g., [Boutilier *et al.*, 2004; Brafman *et al.*, 2006; Bouveret *et al.*, 2009; Bienvenu *et al.*, 2010].

Often it can be valuable to obtain a much less cautious form of inference, for example, in order to have some help in deciding which options to show to the user next in a recommender system [Bridge and Ricci, 2007; Trabelsi *et al.*, 2011]. This includes assuming that the user's preference relation in a multi-objective context is based on a simple weighted sum of objectives (as in a simple form of a Multi-Attribute Utility Theory model [Figueira *et al.*, 2005]); this kind of preference inference is considered in [Bridge and Ricci, 2007; Marinescu *et al.*, 2013]; or alternatively, assuming different lexicographic forms for the preference models as in [Wilson, 2009; Trabelsi *et al.*, 2011; Wilson, 2014]. Note that all these systems involve reasoning about what holds in a set of preference models. This contrasts with work in preference learning [Fürnkranz and (eds.), 2010; Dombi *et al.*, 2007; Flach and Matsubara, 2007; Bräuning and Hüllermeier, 2012; Booth *et al.*, 2010] that typically learns a single model.

In this paper we consider a situation in which alternatives have an associated vector of costs, each component corresponding to a different criterion, and are compared using a kind of lexicographic order, similar to the way alternatives are compared in a Hierarchical Constraint Logic Programming (HCLP) model [Wilson and Borning, 1993]. It is assumed that the user has some (unknown) importance ordering on criteria, and that to compare two alternatives, firstly, the combined cost of each alternative with respect to the most important criteria are compared; only if these combined costs are equal, are the next most important criteria considered.

We consider the case where the input preference statements are of a simple form that one alternative is preferred to another alternative, where we allow the expression of both strict and non-strict preferences (in contrast to most related preference logics, such as [Wilson, 2011; Boutilier *et al.*, 2004; Wilson, 2009; 2014]). This form of preference is natural in many contexts, including for conversational recommender systems [Bridge and Ricci, 2007]. The preference inference problem then consists of determining whether a preference statement can be inferred from a set of input preferences, i.e., if every model (of the assumed form) satisfying the inputs also satisfies the query. We show that this problem is coNP-complete, even if one restricts the cardinality of the equal-importance sets to have at most two elements, and one only considers non-strict preferences. However, it is polynomial if it is assumed that the user's ordering of criteria is a total ordering; it is also polynomial if the sets of equally important criteria are all equivalence classes of a given fixed equivalence relation. We give an efficient polynomial algorithm for these cases, which also throws light on the structure of the inference.

Section 2 defines our simple preference logic based on hierarchical models, along with some associated preference inference problems. Section 3 shows that in general the preference inference problem is coNP-complete. Section 4 considers the case where the importance ordering on criteria is a total order, and gives a polynomial algorithm for consistency. Section 5 concludes.

Proofs are included in a longer version of the paper available online [Wilson et al.*, 2015].*

2 A Preference Logic Based on Hierarchical Models

We consider preference models, based on an importance ordering of criteria, that is basically lexicographic, but involving a combination of criteria which are at the same level in the importance ordering. We call these "HCLP models", because models of a similar kind appear in the HCLP system [Wilson and Borning, 1993] (though we have abstracted away some details from the latter system).

HCLP structures: Define an HCLP structure to be a tuple $\mathcal{S} = \langle \mathcal{A}, \oplus, \mathcal{C} \rangle$, where $\mathcal{A}$ (the set of alternatives) is a finite set; $\oplus$ is an associative, commutative and monotonic operation ($x \oplus y \leq z \oplus y$ if $x \leq z$) on the non-negative rational numbers $\mathbb{Q}^+$, with identity element 0; and $\mathcal{C}$ (known as the set of ($\mathcal{A}$-)evaluations) is a set of functions from $\mathcal{A}$ to $\mathbb{Q}^+$. We also assume that operation $\oplus$ can be computed in linear time (which holds for natural definitions of $\oplus$, including addition and max). The evaluations in $\mathcal{C}$ may be considered as representing criteria or objectives under which the alternatives are evaluated. For $c \in \mathcal{C}$ and $\alpha \in \mathcal{A}$, if $c(\alpha) = 0$ then α fully satisfies the objective corresponding to c; more generally, the smaller the value of $c(\alpha)$, the better α satisfies the c-objective.

With each subset C of $\mathcal{C}$ we define ordering $\preccurlyeq_C^{\oplus}$ on $\mathcal{A}$ by $\alpha \preccurlyeq_C^{\oplus} \beta$ if and only if $\bigoplus_{c \in C} c(\alpha) \leq \bigoplus_{c \in C} c(\beta)$. Relation $\preccurlyeq_C^{\oplus}$ represents how well the alternatives satisfy the multi-set of evaluations C if the latter are considered equally important. $\preccurlyeq_C^{\oplus}$ is a total pre-order (a weak order, i.e., a transitive and complete binary relation). If $\alpha \preccurlyeq_C^{\oplus} \beta$, we might also write $\beta \succcurlyeq_C^{\oplus} \alpha$. We write $\equiv_C^{\oplus}$ for the associated equivalence relation on $\mathcal{A}$, given by $\alpha \equiv_C^{\oplus} \beta \iff \alpha \preccurlyeq_C^{\oplus} \beta$ and $\beta \preccurlyeq_C^{\oplus} \alpha$. We write $\prec_C^{\oplus}$ for the associated strict weak ordering, defined by $\alpha \prec_C^{\oplus} \beta \iff \alpha \preccurlyeq_C^{\oplus} \beta$ and $\beta \not\preccurlyeq_C^{\oplus} \alpha$. Thus, $\alpha \equiv_C^{\oplus} \beta$ if and only if $\bigoplus_{c \in C} c(\alpha) = \bigoplus_{c \in C} c(\beta)$; and $\alpha \prec_C^{\oplus} \beta$ if and only if $\bigoplus_{c \in C} c(\alpha) < \bigoplus_{c \in C} c(\beta)$.

HCLP models: An HCLP model H based on $\langle \mathcal{A}, \oplus, \mathcal{C} \rangle$ is defined to be an ordered partition $(C_1, \ldots, C_k)$ of a (possibly empty) subset $\sigma(H)$ of $\mathcal{C}$. The sets C_i are called the *levels of* H, which are thus non-empty, disjoint and have union $\sigma(H)$. If $c \in C_i$ and $c' \in C_j$, and $i < j$, then we say that c *appears before* c' (and c' *appears after* c) in H. Associated with H is an ordering relation $\preccurlyeq_H^{\oplus}$ on $\mathcal{A}$ given by:

$\alpha \preccurlyeq_H^{\oplus} \beta$ if and only if either:

(I) for all $i = 1, \ldots, k$, $\alpha \equiv_{C_i}^{\oplus} \beta$; or

(II) there exists some $i \in \{1, \ldots, k\}$ such that (i) $\alpha \prec_{C_i}^{\oplus} \beta$ and (ii) for all j with $1 \leq j < i$, $\alpha \equiv_{C_j}^{\oplus} \beta$.

Relation $\preccurlyeq_H^{\oplus}$ is a kind of lexicographic order on $\mathcal{A}$, where the multi-set C_i of evaluations at the same level are first combined into a single evaluation. $\preccurlyeq_H^{\oplus}$ is a weak order on $\mathcal{A}$. We write $\equiv_H^{\oplus}$ for the associated equivalence relation (corresponding with condition (I)), and $\prec_H^{\oplus}$ for the associated strict weak order (corresponding with condition (II)), so that $\preccurlyeq_H^{\oplus}$ is the disjoint union of $\prec_H^{\oplus}$ and $\equiv_H^{\oplus}$. If $\sigma(H) = \emptyset$ then then the first condition for $\alpha \preccurlyeq_H^{\oplus} \beta$ holds vacuously (since $k = 0$), so we have $\alpha \preccurlyeq_H^{\oplus} \beta$ for all $\alpha, \beta \in \mathcal{A}$, and $\prec_H^{\oplus}$ is the empty relation.

Preference Language Inputs: Let $\mathcal{A}$ be a set of alternatives. We define $\mathcal{L}_{\leq}^{\mathcal{A}}$ to be the set of statements of the form $\alpha \leq \beta$, for $\alpha, \beta \in \mathcal{A}$ (the *non-strict* statements), we write $\mathcal{L}_{<}^{\mathcal{A}}$ for the set of statements of the form $\alpha < \beta$, for $\alpha, \beta \in \mathcal{A}$ (the *strict* statements), and let $\mathcal{L}^{\mathcal{A}} = \mathcal{L}_{\leq}^{\mathcal{A}} \cup \mathcal{L}_{<}^{\mathcal{A}}$. If φ is the preference statement $\alpha \leq \beta$ then $\neg\varphi$ is defined to be the preference statement $\beta < \alpha$. If φ is the preference statement $\alpha < \beta$ then $\neg\varphi$ is defined to be the preference statement $\beta \leq \alpha$.

Satisfaction of preference statements: For HCLP model H over HCLP structure $\langle \mathcal{A}, \oplus, \mathcal{C} \rangle$, we say that H satisfies $\alpha \leq \beta$ (written $H \models^{\oplus} \alpha \leq \beta$) if $\alpha \preccurlyeq_H^{\oplus} \beta$ holds. Similarly, we say that H satisfies $\alpha < \beta$ (written $H \models^{\oplus} \alpha < \beta$) if $\alpha \prec_H^{\oplus} \beta$. For $\Gamma \subseteq \mathcal{L}^{\mathcal{A}}$, we say that H satisfies Γ (written $H \models^{\oplus} \Gamma$) if H satisfies φ for all $\varphi \in \Gamma$.

Preference Inference/Deduction relation: We are interested in different restrictions on the set of models, and the corresponding inference relations. Let $\mathcal{M}$ be a set of HCLP models over HCLP structure $\langle \mathcal{A}, \oplus, \mathcal{C} \rangle$. For $\Gamma \subseteq \mathcal{L}^{\mathcal{A}}$, and $\varphi \in \mathcal{L}^{\mathcal{A}}$, we say that $\Gamma \models_{\mathcal{M}}^{\oplus} \varphi$, if H satisfies φ for every $H \in \mathcal{M}$ satisfying Γ. Thus, if we elicit some preference statements Γ of a user, and we assume that their preference relation is an HCLP model in $\mathcal{M}$ (based on the HCLP structure), then $\Gamma \models^{\oplus}_{\mathcal{M}} \varphi$ holds if and only if we can deduce that the user's HCLP model H satisfies φ.

Consistency: For set of HCLP models $\mathcal{M}$ over HCLP structure $\langle \mathcal{A}, \oplus, \mathcal{C} \rangle$, and set of preference statements $\Gamma \subseteq \mathcal{L}^{\mathcal{A}}$, we say that Γ is $(\mathcal{M}, \oplus)$-consistent if there exists $H \in \mathcal{M}$ such that $H \models^{\oplus} \Gamma$; otherwise, we say that Γ is $(\mathcal{M}, \oplus)$-inconsistent. In the usual way, because of the existence of a negation operator, deduction can be reduced to checking (in)consistency.

Proposition 1 $\Gamma \models^{\oplus}_{\mathcal{M}} \varphi$ *if and only if* $\Gamma \cup \{\neg\varphi\}$ *is* $(\mathcal{M}, \oplus)$*-inconsistent.*

Let t be some number in $\{1, 2, \ldots, |\mathcal{C}|\}$. We define $\mathcal{C}(t)$ to be the set of all HCLP models $(C_1, \ldots, C_k)$ based on HCLP structure $\langle \mathcal{A}, \oplus, \mathcal{C} \rangle$ such that $|C_i| \leq t$, for all $i = 1, \ldots, k$. An element of $\mathcal{C}(1)$ thus corresponds to a sequence of singleton sets of evaluations; we identify it with a sequence of evaluations $(c_1, \ldots, c_k)$ in $\mathcal{C}$. Thus, $\Gamma \models^{\oplus}_{\mathcal{C}(t)} \varphi$ if and only if $H \models^{\oplus} \varphi$ for all $H \in \mathcal{C}(t)$ such that $H \models^{\oplus} \Gamma$. Note that for $t = 1$, these definitions do not depend on $\oplus$ (since there is no combination of evaluations involved), so we may drop any mention of $\oplus$.

Let $\equiv$ be an equivalence relation on $\mathcal{C}$. We define $\mathcal{C}(\equiv)$ to be the set of all HCLP models $(C_1, \ldots, C_k)$ such that each C_i is an equivalence class with respect to $\equiv$. It is easy to see that the relation $\models^{\oplus}_{\mathcal{C}(\equiv)}$ is the same as the relation $\models_{C'(1)}$ where C' is defined as follows. C' is in 1-1 correspondence with the set of $\equiv$-equivalence classes of $\mathcal{C}$. If E is the $\equiv$-equivalence class of $\mathcal{C}$ corresponding with $c' \in C'$ then, for $\alpha \in \mathcal{A}$, $c'(\alpha)$ is defined to be $\bigoplus_{c \in E} c(\alpha)$.

For $\models$ either being $\models^{\oplus}_{\mathcal{C}(t)}$ for some $t \in \{1, 2, \ldots, |\mathcal{C}|\}$, or being $\models^{\oplus}_{\mathcal{C}(\equiv)}$ for some equivalence relation $\equiv$ on $\mathcal{C}$, we consider the following decision problem.

HCLP-DEDUCTION FOR $\models$: Given $\mathcal{C}$, Γ and φ is it the case that $\Gamma \models \varphi$?

In Section 4, we will show that this problem is polynomial for $\models$ being $\models^{\oplus}_{\mathcal{C}(t)}$ when $t = 1$. Thus it is polynomial also for $\models^{\oplus}_{\mathcal{C}(\equiv)}$, for any equivalence relation $\equiv$. It is coNP-complete for $\models$ being $\models^{\oplus}_{\mathcal{C}(t)}$ when $t > 1$, as shown below in Section 3.

Theorem 1 HCLP-DEDUCTION FOR $\models^{\oplus}_{\mathcal{C}(t)}$ *is polynomial when* $t = 1$*, and is* coNP*-complete for any* $t > 1$*, even if we restrict the language to non-strict preference statements.* HCLP-DEDUCTION FOR $\models^{\oplus}_{\mathcal{C}(\equiv)}$ *is polynomial for any equivalence relation* $\equiv$*.*

Example

We consider an example with alternatives $\mathcal{A} = \{\alpha, \beta, \gamma\}$, using $\oplus$ as $+$ (ordinary addition), and with evaluations $\mathcal{C} = \{c_1, c_2, c_3\}$, defined as follows.

$c_1(\alpha) = 0;\ c_1(\beta) = 2;\ c_1(\gamma) = 1;$
$c_2(\alpha) = 2;\ c_2(\beta) = 0;\ c_2(\gamma) = 2;$
$c_3(\alpha) = 1;\ c_3(\beta) = 0;\ c_3(\gamma) = 0.$

Suppose that the user tells us that they consider α to be at least as good as β. We represent this as the non-strict preference $\alpha \leq \beta$. This eliminates some models; for instance, $\alpha \leq \beta$ is not satisfied by model H equalling $(\{c_1, c_2\}, \{c_3\})$, since $c_1(\alpha) \oplus c_2(\alpha) = 2 = c_1(\beta) \oplus c_2(\beta)$, and $c_3(\beta) < c_3(\alpha)$, which implies that $\beta \prec^{\oplus}_{H} \alpha$, and thus $H \models \beta < \alpha$ and $H \not\models \alpha \leq \beta$.

Now, $\alpha \leq \beta$ does not entail any preference between β and γ, as can be seen by considering the two models $(\{c_1\})$ and $(\{c_1, c_2\})$, each with just a single level, and both satisfying $\alpha \leq \beta$. Model $(\{c_1, c_2\})$ satisfies $\beta < \gamma$, since $c_1(\beta) \oplus c_2(\beta) = 2 < 3 = c_1(\gamma) \oplus c_2(\gamma)$, whereas model $(\{c_1\})$ satisfies $\gamma < \beta$, since $c_1(\gamma) < c_1(\beta)$. However, it can be shown that $\alpha \leq \beta \models^{\oplus}_{\mathcal{C}(3)} \alpha \leq \gamma$, so we infer that α is non-strictly preferred to γ.

A strict preference for α over β entails additional inferences, for instance, $\alpha < \beta \models^{\oplus}_{\mathcal{C}(3)} \gamma < \beta$. If we restrict the set of models to $\mathcal{C}(1)$ (thus assuming that the importance ordering on $\mathcal{C}$ is a total order) we get slightly stronger inferences still, obtaining in addition: $\alpha < \beta \models_{\mathcal{C}(1)} \alpha < \gamma$.

3 Proving coNP-completeness of HCLP-Deduction for $\models^{\oplus}_{\mathcal{C}(t)}$ for $t > 1$

Given an arbitrary 3-SAT instance we will show that we can construct a set Γ and a statement $\alpha \leq \beta$ such that the 3-SAT instance has a satisfying truth assignment if and only if $\Gamma \not\models^{\oplus}_{\mathcal{C}(t)} \alpha \leq \beta$ (see Proposition 2). This then implies that determining if $\Gamma \not\models^{\oplus}_{\mathcal{C}(t)} \alpha \leq \beta$ holds is NP-hard.

We have that $\Gamma \not\models^{\oplus}_{\mathcal{C}(t)} \alpha \leq \beta$ if and only if there exists an HCLP-model $H \in \mathcal{C}(t)$ such that $H \models^{\oplus} \Gamma$ and $H \not\models^{\oplus} \alpha \leq \beta$. For any given H, checking that $H \models^{\oplus} \Gamma$ and $H \not\models^{\oplus} \alpha \leq \beta$ can be performed in polynomial time. This implies that determining if $\Gamma \models^{\oplus}_{\mathcal{C}(t)} \alpha \leq \beta$ holds is coNP-complete.

For simplicity, we describe the construction for the $t = 2$ case. This construction also works for all $t \geq 2$ if $1 \oplus 1 > 1$. We indicate below how the construction is modified for the case when $t \geq 3$.

Consider an arbitrary 3-SAT instance based on propositional variables $p_1, \ldots, p_r$, consisting of clauses Λ_j, for $j = 1, \ldots, s$. For each propositional variable p_i we associate two evaluations q_i^+ and q_i^-, where q_i^- corresponds with literal $\neg p_i$ and q_i^+ corresponds with literal p_i.

The idea behind the construction is as follows: we generate a (polynomial size) set $\Gamma \subseteq \mathcal{L}^{\mathcal{A}}_{\leq}$ as the disjoint union of sets Γ_1, Γ_2 and Γ_3. Γ_1 is chosen so that if $H \models^{\oplus} \Gamma_1$ then, for each $i = 1, \ldots, r$, $\sigma(H)$ cannot contain both q_i^+ and q_i^-, i.e., q_i^+ and q_i^- do not both appear in H. (Recall H is an ordered partition of $\sigma(H)$, so that $\sigma(H)$ is the subset of $\mathcal{C}$ that appears in H.) If $H \models^{\oplus} \Gamma_2$ and $H \models^{\oplus} \beta < \alpha$ then $\sigma(H)$ contains either q_i^+ or q_i^-. Together, this implies that if $H \models^{\oplus} \Gamma$ and $H \not\models^{\oplus} \alpha \leq \beta$ then for each propositional variable p_i, model H involves either q_i^+ or q_i^-, but not both. Γ_3 is used to make the correspondence with the clauses. For instance, if one of the clauses is $p_2 \vee \neg p_5 \vee p_6$ then any HCLP model $H \in \mathcal{C}(t)$ of $\Gamma \cup \{\beta < \alpha\}$ will involve either q_2^+, q_5^-, or q_6^+.

Suppose that H satisfies Γ but not $\alpha \leq \beta$. We can generate a satisfying assignment of the 3-SAT instance, by assigning p_i to 1 (TRUE) if and only if q_i^+ appears in H.

We describe the construction more formally below.

Defining $\mathcal{A}$ and $\mathcal{C}$: The set of alternatives $\mathcal{A}$ is defined to be: $\{\alpha, \beta\} \cup \{\delta_i, \gamma_i, \alpha_i, \beta_i : i = 1, \ldots, r\} \cup \{\theta_j, \tau_j : j = 1, \ldots, s\}$. We define the set of evaluations $\mathcal{C}$ to be $\{c^*\} \cup \{a_i, q_i^+, q_i^- : i = 1, \ldots, r\}$. Both $\mathcal{A}$ and $\mathcal{C}$ are of polynomial size.

Satisfying $\beta < \alpha$: The evaluations on α and β are defined as follows:

- $c^*(\alpha) = 1$, and for all $c \in \mathcal{C} - \{c^*\}$, $c(\alpha) = 0$.
- For all $c \in \mathcal{C}$, $c(\beta) = 0$.

It immediately follows that: $H \models^{\oplus} \beta < \alpha \iff \sigma(H) \ni c^*$.

The construction of Γ_1: For each $i = 1, \ldots, r$, we define $\Gamma_1^i = \{\delta_i \leq \gamma_i, \gamma_i \leq \delta_i\}$, and we let $\Gamma_1 = \bigcup_{i=1}^r \Gamma_1^i$. We make use of auxiliary evaluation a_i. The values of the evaluations on γ_i and δ_i are defined as follows:

- $a_i(\gamma_i) = 1$, and for all $c \in \mathcal{C} - \{a_i\}$ we set $c(\gamma_i) = 0$.
- $q_i^+(\delta_i) = q_i^-(\delta_i) = 1$, and for other $c \in \mathcal{C}$, $c(\delta_i) = 0$.

Thus $(a_i \oplus q_i^+)(\delta_i) = a_i(\delta_i) \oplus q_i^+(\delta_i) = 0 \oplus 1 = 1$. Similarly, $(a_i \oplus q_i^-)(\delta_i) = 1$, and $(a_i \oplus q_i^+)(\gamma_i) = (a_i \oplus q_i^-)(\gamma_i) = 1 \oplus 0 = 1$.

Lemma 1 *$H \models^{\oplus} \Gamma_1^i$ if and only if either (i) $\sigma(H)$ does not contain a_i or q_i^+ or q_i^-, i.e., $\sigma(H) \cap \{a_i, q_i^+, q_i^-\} = \emptyset$; or (ii) $\{a_i, q_i^+\}$ is a level of H, and $\sigma(H) \not\ni q_i^-$; or (iii) $\{a_i, q_i^-\}$ is a level of H, and $\sigma(H) \not\ni q_i^+$. In particular, if $H \models^{\oplus} \Gamma_1^i$ then $\sigma(H)$ does not contain both q_i^+ and q_i^-.*

This holds because if a_i appears in H before q_i^+ and q_i^- then we will have that $H \not\models^{\oplus} \gamma_i \leq \delta_i$. If q_i^+ or q_i^- appear in H before a_i then we will have $H \not\models^{\oplus} \delta_i \leq \gamma_i$. Also, if e.g., $\{a_i, q_i^+\}$ is a level of H and q_i^- appears later in H then we will have $H \not\models^{\oplus} \delta_i \leq \gamma_i$.

(For the case when $t \geq 3$, we need a slightly more complex construction, involving evaluations a_i^k and alternatives δ_i^k and γ_i^k for $k = 1 \ldots, t-1$, with similar definitions as given above.)

The construction of Γ_2: For each $i = 1, \ldots, r$, define φ_i to be $\alpha_i \leq \beta_i$. We let $\Gamma_2 = \{\varphi_i : i = 1, \ldots, r\}$. The values of the evaluations on α_i and β_i are defined as follows. We define $c^*(\alpha_i) = 1$, and for all $c \in \mathcal{C} - \{c^*\}$, $c(\alpha_i) = 0$. Define $q_i^+(\beta_i) = q_i^-(\beta_i) = 1$, and for all $c \in \mathcal{C} - \{q_i^+, q_i^-\}$, $c(\beta_i) = 0$. The following result easily follows.

Lemma 2 *If q_i^+ or q_i^- appears before (or in the same level as) c^* in H then $H \models^{\oplus} \varphi_i$. If $\sigma(H) \ni c^*$ and $H \models^{\oplus} \varphi_i$ then $\sigma(H) \ni q_i^+$ or $\sigma(H) \ni q_i^-$.*

The construction of Γ_3: For each $i = 1, \ldots, r$, define $Q(p_i) = q_i^+$ and $Q(\neg p_i) = q_i^-$. This defines the function Q over all literals. Let us write the jth clause as $l_1 \vee l_2 \vee l_3$ for literals l_1, l_2 and l_3. Define $Q_j = \{Q(l_1), Q(l_2), Q(l_3)\}$. For example, if the jth clause were $p_2 \vee \neg p_5 \vee p_6$ then $Q_j = \{q_2^+, q_5^-, q_6^+\}$. We define ψ_j to be $\theta_j \leq \tau_j$, and $\Gamma_3 = \{\psi_j : j = 1, \ldots s\}$. Define $c^*(\theta_j) = 1$ and $c(\theta_j) = 0$ for all $c \in \mathcal{C} - \{c^*\}$. Define $q(\tau_j) = 1$ for $q \in Q_j$, and for all other c (i.e., $c \in \mathcal{C} - Q_j$), define $c(\tau_i) = 0$. This leads to the following properties.

Lemma 3 *If some element of Q_j appears in H before c^* then $H \models^{\oplus} \psi_j$. If $\sigma(H) \ni c^*$ and $H \models^{\oplus} \psi_j$ then $\sigma(H)$ contains some element of Q_j.*

We set $\Gamma = \Gamma_1 \cup \Gamma_2 \cup \Gamma_3$. The following result implies that the HCLP deduction problem is coNP-hard (even if we restrict to the case when $\Gamma \cup \{\varphi\} \subseteq \mathcal{L}_{\leq}^{\mathcal{A}}$).

Proposition 2 *Using the notation defined above, the 3-SAT instance is satisfiable if and only if $\Gamma \not\models_{\mathcal{C}(t)}^{\oplus} \alpha \leq \beta$.*

Proof: First let us assume that $\Gamma \not\models_{\mathcal{C}(t)}^{\oplus} \alpha \leq \beta$. Then by definition, there exists an HCLP model $H \in \mathcal{C}(t)$ with $H \models^{\oplus} \Gamma$ and $H \not\models^{\oplus} \alpha \leq \beta$. Since $H \not\models^{\oplus} \alpha \leq \beta \iff H \models^{\oplus} \beta < \alpha$, we have $H \models^{\oplus} \Gamma \cup \{\beta < \alpha\}$. Because $H \models^{\oplus} \beta < \alpha$, $\sigma(H) \ni c^*$.

Because also $H \models^{\oplus} \Gamma_2^i$, either $\sigma(H) \ni q_i^+$ or $\sigma(H) \ni q_i^-$, by Lemma 2. Since $H \models^{\oplus} \Gamma_1^i$, the set $\sigma(H)$ does not contain both q_i^+ and q_i^-, by Lemma 1.

Let us define a truth function $f : \mathcal{P} \to \{0, 1\}$ as follows: $f(p_i) = 1 \iff \sigma(H) \ni q_i^+$. Since $\sigma(H)$ contains exactly one of q_i^+ and q_i^-, we have $f(p_i) = 0 \iff \sigma(H) \ni q_i^-$. We extend f to negative literals in the obvious way: $f(\neg p_i) = 1 - f(p_i)$.

Since $H \models^{\oplus} \Gamma_3 \cup \{\beta < \alpha\}$, $\sigma(H)$ contains at least one element of each Q_j, by Lemma 3. Thus for each j, $f(l) = 1$ for at least one literal l in the jth clause, and hence f satisfies clause Λ_j. We have shown that f satisfies each clause of the 3-SAT instance, proving that the instance is satisfiable.

Conversely, suppose that the 3-SAT instance is satisfiable, so there exists a truth function f satisfying it.

We will construct an HCLP model $H \in \mathcal{C}(t)$ such that $H \models^{\oplus} \Gamma \cup \{\beta < \alpha\}$, and thus $H \not\models^{\oplus} \alpha \leq \beta$, proving that $\Gamma \not\models_{\mathcal{C}(t)}^{\oplus} \alpha \leq \beta$.

For $i = 1, \ldots, r$, let $S_i = \{a_i, q_i^+\}$ if $f(p_i) = 1$, and otherwise, let $S_i = \{a_i, q_i^-\}$. We then define H to be the sequence $S_1, S_2, \ldots, S_r, \{c^*\}$. Since $\sigma(H) \ni c^*$, we have that $H \models^{\oplus} \beta < \alpha$.

By Lemma 1, for all $i = 1, \ldots, r$, $H \models^{\oplus} \Gamma_1^i$ and so $H \models^{\oplus} \Gamma_1$. By Lemma 2, for all $i = 1, \ldots, r$, $H \models^{\oplus} \varphi_i$, so $H \models^{\oplus} \Gamma_2$.

Consider any $j \in \{1, \ldots, s\}$, and, as above, write the jth clause as $l_1 \vee l_2 \vee l_3$. Truth assignment f satisfies this clause, so there exists $k \in \{1, 2, 3\}$ such that $f(l_k) = 1$, where f is extended to literals in the usual way. Then $Q(l_k)$ appears in H before c^*, so, by Lemma 3, $H \models^{\oplus} \psi_j$. Thus $H \models^{\oplus} \Gamma_3$.

Since $\Gamma = \Gamma_1 \cup \Gamma_2 \cup \Gamma_3$, we have shown that $H \models^{\oplus} \Gamma \cup \{\beta < \alpha\}$, completing the proof. $\square$

4 Sequence-of-Evaluations Models

In this section, we consider the case where we restrict to HCLP models which consist of a sequence of singletons; thus each model corresponds to a sequence of evaluations, and generates a lexicographic order based on these.

Let $\mathcal{C}$ be a set of evaluations on $\mathcal{A}$. A $\mathcal{C}(1)$-model is a sequence of different elements of $\mathcal{C}$. The operation $\oplus$ plays no part, so we can harmlessly abbreviate ordering $\preccurlyeq_H^{\oplus}$ to just $\preccurlyeq_H$, for any $\mathcal{C}(1)$-model H, and similarly for $\prec_H$ and $\equiv_H$.

4.1 Some Basic Definitions and Results

We write $\varphi \in \mathcal{L}^{\mathcal{A}}$ as $\alpha_\varphi < \beta_\varphi$, if φ is strict, or as $\alpha_\varphi \leq \beta_\varphi$, if φ is non-strict. We consider a set $\Gamma \subseteq \mathcal{L}^{\mathcal{A}}$, and a set $\mathcal{C}$ of evaluations on $\mathcal{A}$. For $\varphi \in \Gamma$, define $Supp_{\mathcal{C}}^{\varphi}$ to be $\{c \in \mathcal{C} : c(\alpha_\varphi) < c(\beta_\varphi)\}$; define $Opp_{\mathcal{C}}^{\varphi}$ to be $\{c \in \mathcal{C} : c(\alpha_\varphi) > c(\beta_\varphi)\}$; and define $Ind_{\mathcal{C}}^{\varphi}$ to be $\{c \in \mathcal{C} : c(\alpha_\varphi) = c(\beta_\varphi)\}$. We may sometimes abbreviate $Supp_{\mathcal{C}}^{\varphi}$ to $Supp^{\varphi}$, and similarly for $Opp_{\mathcal{C}}^{\varphi}$ and $Ind_{\mathcal{C}}^{\varphi}$. $Supp^{\varphi}$ are the evaluations that support φ; Opp^{φ} are the evaluations that oppose φ. Ind^{φ} are the other evaluations, that are indifferent regarding φ. For a model H to satisfy φ it is necessary that no evaluation that opposes φ appears before all evaluations that support φ.

The following defines *inconsistency bases*, which are concerned with evaluations that cannot appear in any model satisfying the set of preference statements Γ (see Proposition 3 below).

Definition 1 *Let $\Gamma \subseteq \mathcal{L}^{\mathcal{A}}$, and let $\mathcal{C}$ be a set of $\mathcal{A}$-evaluations. We say that (Γ', C') is an inconsistency base for $(\Gamma, \mathcal{C})$ if $\Gamma' \subseteq \Gamma$, and $C' \subseteq \mathcal{C}$, and*

(i) for all $\varphi \in \Gamma'$, $Supp_{\mathcal{C}}^{\varphi} \cup Opp_{\mathcal{C}}^{\varphi} \subseteq C'$ (and thus $\mathcal{C} - C' \subseteq Ind_{\mathcal{C}}^{\varphi}$); and

(ii) for all $c \in C'$, there exists $\varphi \in \Gamma'$ such that $Opp_{\mathcal{C}}^{\varphi} \ni c$.

Thus, for all $\varphi \in \Gamma'$, C' contains all evaluations that are not indifferent regarding φ, and for all $c \in C'$ there is some element of Γ' that is opposed by c. The following result motivates the definition:

Proposition 3 *Let H be an element of $\mathcal{C}(1)$. Suppose that $H \models \Gamma$, and let (Γ', C') be an inconsistency base for $(\Gamma, \mathcal{C})$. Then $C' \cap \sigma(H) = \emptyset$. Thus no $\mathcal{C}(1)$ model of Γ can involve any element of C'. Also, we have for any $\varphi \in \Gamma'$, $\alpha_\varphi \equiv_H \beta_\varphi$, so $H \not\models \alpha_\varphi < \beta_\varphi$.*

Thus H does not strictly satisfy any element of Γ'.

For inconsistency bases (Γ_1, C_1) and (Γ_2, C_2) for $(\Gamma, \mathcal{C})$, define $(\Gamma_1, C_1) \cup (\Gamma_2, C_2)$ to be $(\Gamma_1 \cup \Gamma_2, C_1 \cup C_2)$. It is easy to show that if (Γ_1, C_1) and (Γ_2, C_2) are both inconsistency bases for $(\Gamma, \mathcal{C})$ then $(\Gamma_1, C_1) \cup (\Gamma_2, C_2)$ is one also. Define $MIB(\Gamma, \mathcal{C})$, the Maximal Inconsistency Base for $(\Gamma, \mathcal{C})$, to be the union of all inconsistency bases for $(\Gamma, \mathcal{C})$, i.e., $\bigcup\{(\Gamma_1, C_1) \in \mathcal{I}\}$, where $\mathcal{I}$ is the set of inconsistency bases for $(\Gamma, \mathcal{C})$. The next result follows.

Proposition 4 *$MIB(\Gamma, \mathcal{C})$ always exists, and is an inconsistency base for $(\Gamma, \mathcal{C})$, which is maximal in the following sense: if (Γ_1, C_1) is an inconsistency base for $(\Gamma, \mathcal{C})$ then $\Gamma_1 \subseteq \Gamma^{\perp}$ and $C_1 \subseteq C^{\perp}$, where $MIB(\Gamma, \mathcal{C}) = (\Gamma^{\perp}, C^{\perp})$.*

It can also be easily shown that if Γ is $\mathcal{C}(1)$-consistent then $\Gamma^{\perp}$ contains no strict elements. Theorem 2 below implies the converse of this result.

4.2 A Polynomial Algorithm for Consistency and Deduction

Throughout this section we consider a set $\Gamma \subseteq \mathcal{L}^{\mathcal{A}}$ of input preference statements, and a set $\mathcal{C}$ of $\mathcal{A}$-evaluations.

Define $Opp_{\Gamma}(c)$ (abbreviated to $Opp(c)$) to be the set of elements opposed by c, i.e., $\varphi \in \Gamma$ such that $c(\alpha_\varphi) > c(\beta_\varphi)$, and define $Supp_{\Gamma}(c)$ (usually abbreviated to $Supp(c)$) to be the set of elements φ of Γ supported by c, (i.e., $c(\alpha_\varphi) < c(\beta_\varphi)$). Also, for sequence of evaluations $(c_1, \ldots, c_k)$, we define $Supp(c_1, \ldots, c_k)$ to be $\bigcup_{i=1}^{k} Supp(c_i)$.

The idea behind the algorithm is as follows: suppose that we have picked a sequence of elements of $\mathcal{C}$, C' being the set picked so far. We need to choose next an evaluation c such that, if c opposes some φ in Γ, then φ is supported by some evaluation in C' (or else the generated sequence will not satisfy φ).

The Algorithm

H is initialised as the empty list $()$ of evaluations. $H \leftarrow H + c$ means that evaluation c is added to the end of H.

Function *Cons-check*$(\Gamma, \mathcal{C})$

$H \leftarrow ()$
for $k = 1, \ldots, |\mathcal{C}|$ do
  if $\exists\, c \in \mathcal{C} - \sigma(H)$ such that $Opp(c) \subseteq Supp(H)$
    then choose some such c; $H \leftarrow H + c$
    else stop
  end for
return H

The algorithm involves often non-unique choices. However, if we wish, the choosing of c can be done based on an ordering $c_1, \ldots, c_m$ of $\mathcal{C}$, where, if there exists more than one $c \in \mathcal{C} - \sigma(H)$ such that $Opp(c) \subseteq Supp(H)$, we choose the element c_i with smallest index i. The algorithm then becomes deterministic, with a unique result following from the given inputs.

A straight-forward implementation runs in $O(|\Gamma||\mathcal{C}|^2)$ time; however, a more careful implementation runs in $O(|\Gamma||\mathcal{C}|)$ time.

Properties of the Algorithm

Any $\varphi \in \Gamma$ which is opposed by some evaluation c in H is supported by some earlier evaluation in H. Consider any φ in $Supp(H)$. Let c_j be the earliest evaluation in H that supports φ. None of the earlier evaluations than c_j oppose φ, and thus H strictly satisfies φ. A similar argument shows that H satisfies $\Gamma^{(\leq)}$, defined to be $\{\alpha_\varphi \leq \beta_\varphi : \varphi \in \Gamma\}$, i.e., Γ where the strict statements are replaced by corresponding non-strict statements.

The algorithm will always generate an HCLP model satisfying Γ if Γ is consistent. It can also be used for computing the Maximal Inconsistency Base. The following result sums up some properties related to the algorithm.

Theorem 2 *Let H be a sequence returned by the algorithm with inputs Γ and $\mathcal{C}$, and write MIB$(\Gamma, \mathcal{C})$ as $(\Gamma^\perp, C^\perp)$. Then $C^\perp = \mathcal{C} - \sigma(H)$ (i.e., the evaluations that don't appear in H), and $\Gamma^\perp = \Gamma - Supp(H)$. We have that $H \models \Gamma^{(\leq)}$. Also, Γ is $\mathcal{C}(1)$-consistent if and only if Supp(H) contains all the strict elements of Γ, which is if and only if $\Gamma^\perp \cap \mathcal{L}^{\mathcal{A}}_{<} = \emptyset$. If Γ is $\mathcal{C}(1)$-consistent then $H \models \Gamma$.*

The algorithm therefore determines $\mathcal{C}(1)$-consistency, and hence $\mathcal{C}(1)$-deduction (because of Proposition 1), in polynomial time, and also generates the Maximal Inconsistency Base. For the case when Γ is not $\mathcal{C}(1)$-consistent, the output H of the algorithm is a model which, in a sense, comes closest to satisfying Γ: it always satisfies $\Gamma^{(\leq)}$, the non-strict version of Γ, and if any model $H' \in \mathcal{C}(1)$ satisfies $\Gamma^{(\leq)}$ and any element φ of Γ, then H also satisfies φ. These properties suggest the following way of reasoning with inconsistent Γ. Let us define Γ' to be equal to $(\Gamma - \Gamma^\perp) \cup \Gamma^{(\leq)}$. By Theorem 2, this is equal to *Supp*$(H) \cup \Gamma^{(\leq)}$, where H is a model generated by the algorithm, enabling easy computation of Γ'. Γ' is consistent, since it is satisfied by H. We might then (re-)define the (non-monotonic) deductions from inconsistent Γ to be the deductions from Γ'.

4.3 Strong Consistency

In the set of models $\mathcal{C}(1)$, we allow models involving any subset of $\mathcal{C}$, the set of evaluations. We could alternatively consider a semantics where we only allow models H that involve all elements of $\mathcal{C}$, i.e., with $\sigma(H) = \mathcal{C}$.

Let $\mathcal{C}(1^*)$ be the set of elements H of $\mathcal{C}(1)$ with $\sigma(H) = \mathcal{C}$. Γ is defined to be *strongly $\mathcal{C}(1)$-consistent* if and only if there exists a model $H \in \mathcal{C}(1^*)$ such that $H \models \Gamma$. Let *MIB*$(\Gamma, \mathcal{C}) = (\Gamma^\perp, C^\perp)$. Proposition 3 implies that, if Γ is strongly $\mathcal{C}(1)$-consistent then $C^\perp$ is empty, and $\Gamma^\perp$ consists of all the elements of Γ that are indifferent to all of $\mathcal{C}$, i.e., the set of $\varphi \in \Gamma$ such that $c(\alpha_\varphi) = c(\beta_\varphi)$ for all $c \in \mathcal{C}$.

We write $\Gamma \models_{\mathcal{C}(1^*)} \varphi$ if $H \models \varphi$ holds for every $H \in \mathcal{C}(1^*)$ such that $H \models \Gamma$. The next result shows that the non-strict $\models_{\mathcal{C}(1^*)}$ inferences are the same as the non-strict $\models_{\mathcal{C}(1)}$ inferences, and that (in contrast to the case of $\models_{\mathcal{C}(1)}$), the strict $\models_{\mathcal{C}(1^*)}$ inferences almost correspond with the non-strict ones. The result also implies that the algorithm in Section 4.2 can be used to efficiently determine the $\models_{\mathcal{C}(1^*)}$ inferences.

Proposition 5 *Let MIB$(\Gamma, \mathcal{C}) = (\Gamma^\perp, C^\perp)$. Γ is strongly $\mathcal{C}(1)$-consistent if and only if $C^\perp = \emptyset$. Suppose that Γ is strongly $\mathcal{C}(1)$-consistent. Then,*

(i) $\Gamma \models_{\mathcal{C}(1)} \alpha \leq \beta \iff \Gamma \models_{\mathcal{C}(1^*)} \alpha \leq \beta$;

(ii) $\Gamma \models_{\mathcal{C}(1^*)} \alpha \equiv \beta$ *if and only if α and β agree on all of $\mathcal{C}$, i.e., for all $c \in \mathcal{C}$, $c(\alpha) = c(\beta)$;*

(iii) $\Gamma \models_{\mathcal{C}(1^*)} \alpha < \beta$ *if and only if* $\Gamma \models_{\mathcal{C}(1)} \alpha \leq \beta$ *and α and β differ on some element of $\mathcal{C}$, i.e., there exists $c \in \mathcal{C}$ such that $c(\alpha) \neq c(\beta)$.*

We obtain a similar (and generalisation of this) result, if we consider only the maximal models in $\mathcal{C}(1)$ that satisfy Γ (these all have the same cardinality: $|\mathcal{C}| - |C^\perp|$).

The next result shows that $\models_{\mathcal{C}(1)}$ inference is not affected if one removes the evaluations in the MIB.

Proposition 6 *Suppose that Γ is $\mathcal{C}(1)$-consistent, let MIB$(\Gamma, \mathcal{C}) = (\Gamma^\perp, C^\perp)$, and let $\mathcal{C}' = \mathcal{C} - C^\perp$. Then Γ is strongly $\mathcal{C}'(1)$-consistent, and $\Gamma \models_{\mathcal{C}(1)} \varphi$ if and only if $\Gamma \models_{\mathcal{C}'(1)} \varphi$.*

4.4 Orderings on evaluations

The preference logic defined here is closely related to a logic based on disjunctive ordering statements. Given set of evaluations $\mathcal{C}$, we consider the set of statements $\mathcal{O}_\mathcal{C}$ of the form $C_1 < C_2$, and of $C_1 \leq C_2$, where C_1 and C_2 are disjoint subsets of $\mathcal{C}$.

We say that $H \models C_1 < C_2$ if some evaluation in C_1 appears in H before every element of C_2, that is, there exists some element of C_1 in H (i.e., $C_1 \cap \sigma(H) \neq \emptyset$) and the earliest element of $C_1 \cup C_2$ to appear in H is in C_1.

We say that $H \models C_1 \leq C_2$ if either $H \models C_1 < C_2$ or no element of C_1 or C_2 appears in H: $(C_1 \cup C_2) \cap \sigma(H) = \emptyset$. Then we have that

$H \models \alpha_\varphi < \beta_\varphi \iff H \models Supp^\varphi_\mathcal{C} < Opp^\varphi_\mathcal{C}$,
and $H \models \alpha_\varphi \leq \beta_\varphi \iff H \models Supp^\varphi_\mathcal{C} \leq Opp^\varphi_\mathcal{C}$.

This shows that the language $\mathcal{O}_\mathcal{C}$ can express anything that can be expressed in $\mathcal{L}^\mathcal{A}$. It can be shown, conversely, that for any statement τ in $\mathcal{O}_\mathcal{C}$, one can define α_φ and β_φ, and the values of elements of $\mathcal{C}$ on these, such that for all $H \in \mathcal{C}(1)$, $H \models \tau$ if and only if $H \models \varphi$. For instance, if τ is the statement $C_1 < C_2$, we can define $c(\alpha_\varphi) = 1$ for all $c \in C_2$, and $c(\alpha_\varphi) = 0$ for $c \in \mathcal{C} - C_2$; and define $c(\beta_\varphi) = 1$ for all $c \in C_1$, and $c(\beta_\varphi) = 0$ for $c \in \mathcal{C} - C_1$.

The algorithm adapts in the obvious way to the case where we have Γ consisting of (or including) elements in $\mathcal{O}_\mathcal{C}$. When viewed in this way, the algorithm can be seen as a simple extension of a topological sort algorithm; the standard case corresponds to when the ordering statements only consist of singleton sets.

5 Discussion and Conclusions

We defined a class of relatively simple preference logics based on hierarchical models. These generate an adventurous form of inference, which can be helpful if there is only relatively sparse input preference information. We showed that the complexity of preference deduction is coNP-complete in general, and polynomial for the case where the criteria are assumed to be totally ordered (the sequence-of-evaluations case, Section 4).

The latter logic has strong connections with the preference inference formalism described in [Wilson, 2014]. To clarify the connection, for each evaluation $c \in \mathcal{C}$ we can generate a variable X_c, and let V be the set of these variables. For each alternative $\alpha \in \mathcal{A}$ we generate a complete assignment α^* on the variables V (i.e., an outcome as defined in [Wilson, 2014]) by $\alpha^*(X_c) = c(\alpha)$ for each $X_c \in V$. Note that values of $\alpha^*(X_c)$ are non-negative numbers, and thus have a fixed ordering, with zero being the best value. A preference statement $\alpha \leq \beta$ in $\mathcal{L}^\mathcal{A}_\leq$ then corresponds with a basic preference formula $\alpha^* \geq \beta^*$ in [Wilson, 2014]. Each model $H \in \mathcal{C}(1)$ corresponds to a sequence of evaluations, and thus has an associated sequence of variables; this sequence together with the fixed value orderings, generates a lexicographic model as defined in [Wilson, 2014].

In contrast with the lexicographic inference system in [Wilson, 2014], the logic developed in this paper allows strict (as well as non-strict) preference statements, and allows more than one variable at the same level. However, the lexicographic inference logic from [Wilson, 2014] does not assume a fixed value ordering (which, translated into the current formalism, corresponds to not assuming that the values of the evaluation function are known); it also allows a richer language of preference statements, where a statement can be a compact representation for a (possibly exponentially large) set of basic preference statements of the form $\alpha \leq \beta$. Many of the results of Section 4 immediately extend to richer preference languages (by replacing a preference statement by a corresponding set of basic preference statements). In future work we will determine under what circumstances deduction remains polynomial when extending the language, and when removing the assumption that the evaluation functions are known.

The coNP-hardness result for the general case (and for the $\models^{\oplus}_{\mathcal{C}(t)}$ systems with $t \geq 2$) is notable and perhaps surprising, since these preference logics are relatively simple ones. The result obviously extends to more general systems. The preference inference system described in [Wilson, 2009] is based on much more complex forms of lexicographic models, allowing conditional dependencies, as well as having local orderings on sets of variables (with bounded cardinality). Theorem 1 implies that the (polynomial) deduction system in [Wilson, 2009] is not more general than the system described here (assuming $\mathsf{P} \neq \mathsf{NP}$). It also implies that if one were to extend the system from [Wilson, 2009] to allow a richer form of equivalence, generalising e.g., the $\models^{\oplus}_{\mathcal{C}(2)}$ system, then the preference inference will no longer be polynomial.

Acknowledgments

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289. Nic Wilson is also supported by the School of EEE&CS, Queen's University Belfast. We are grateful to the reviewers for their helpful comments.

References

[Bienvenu *et al.*, 2010] M. Bienvenu, J. Lang, and N. Wilson. From preference logics to preference languages, and back. In *Proc. KR 2010*, 2010.

[Booth *et al.*, 2010] R. Booth, Y. Chevaleyre, J. Lang, J. Mengin, and C. Sombattheera. Learning conditionally lexicographic preference relations. In *Proc. ECAI-2010*, pages 269–274, 2010.

[Boutilier *et al.*, 2004] C. Boutilier, R. I. Brafman, C. Domshlak, H. Hoos, and D. Poole. CP-nets: A tool for reasoning with conditional *ceteris paribus* preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[Bouveret *et al.*, 2009] S. Bouveret, U. Endriss, and J.Lang. Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods. In *Proc. IJCAI-09*, pages 67–72, 2009.

[Brafman *et al.*, 2006] R. Brafman, C. Domshlak, and E. Shimony. On graphical modeling of preference and importance. *Journal of Artificial Intelligence Research*, 25:389–424, 2006.

[Bräuning and Hüllermeier, 2012] M. Bräuning and E. Hüllermeier. Learning conditional lexicographic preference trees. In *Preference Learning (PL-12), ECAI-12 workshop*, 2012.

[Bridge and Ricci, 2007] D. Bridge and F. Ricci. Supporting product selection with query editing recommendations. In *RecSys '07*, pages 65–72, New York, NY, USA, 2007. ACM.

[Dombi *et al.*, 2007] J. Dombi, C. Imreh, and N. Vincze. Learning lexicographic orders. *European Journal of Operational Research*, 183(2):748–756, 2007.

[Figueira *et al.*, 2005] J. Figueira, S. Greco, and M. Ehrgott. *Multiple Criteria Decision Analysis—State of the Art Surveys*. Springer International Series in Operations Research and Management Science Volume 76, 2005.

[Flach and Matsubara, 2007] P. A. Flach and E.T. Matsubara. A simple lexicographic ranker and probability estimator. In *Proc. ECML-2007*, pages 575–582, 2007.

[Fürnkranz and (eds.), 2010] J. Fürnkranz and E. Hüllermeier (eds.). *Preference Learning*. Springer-Verlag, 2010.

[Marinescu *et al.*, 2013] R. Marinescu, A. Razak, and N. Wilson. Multi-objective constraint optimization with tradeoffs. In *Proc. CP-2013*, pages 497–512, 2013.

[Trabelsi *et al.*, 2011] W. Trabelsi, N. Wilson, D. Bridge, and F. Ricci. Preference dominance reasoning for conversational recommender systems: a comparison between a comparative preferences and a sum of weights approach. *International Journal on Artificial Intelligence Tools*, 20(4):591–616, 2011.

[Wilson and Borning, 1993] Molly Wilson and Alan Borning. Hierarchical constraint logic programming. *The Journal of Logic Programming*, 16(34):277–318, 1993.

[Wilson *et al.*, 2015] N. Wilson, A.-M. George, and B. O'Sullivan. *Computation and Complexity of Preference Inference Based on Hierarchical Models (extended version of current paper including proofs)*. http://ucc.insight-centre.org/nwilson/PrefInfHCLPproofs.pdf), 2015.

[Wilson, 2009] N. Wilson. Efficient inference for expressive comparative preference languages. In *Proc. IJCAI-09*, pages 961–966, 2009.

[Wilson, 2011] N. Wilson. Computational techniques for a simple theory of conditional preferences. *Artificial Intelligence*, 175(7-8):1053–1091, 2011.

[Wilson, 2014] N. Wilson. Preference inference based on lexicographic models. In *Proc. ECAI-2014*, pages 921–926, 2014.