From Preferences over Objects to Preferences over Concepts

Sergei Obiedkov[1]

Abstract. We present an approach to preference modeling and learning preferences from data based on formal concept analysis. We consider techniques to derive preferences over attribute subsets from preferences over objects, including ceteris paribus preferences.

1 INTRODUCTION

If John chooses a strawberry over an apple, would he choose a raspberry over a pear? If so, is it because, for John, red berries taste better than tree fruit or are there other factors involved? More generally, given a number of alternatives each described by a set of elementary features together with a preference relation over these alternatives, we would like to derive preferences among feature sets that would explain, at least partially, the observed preferences over individual alternatives.

A move from a strawberry and a raspberry to red berries is a move from individual objects to concepts; thus, our aim is to generalize from preferences over objects to preferences over concepts. A concept can be understood *extensionally*, through objects it covers, or *intensionally*, through attributes that define it. This duality reflects itself in preferences over object sets and preferences over attribute sets, so that the latter can be defined in terms of the former. The first step is then to get from preferences over objects to preferences over object sets, and for this there are various options extensively studied in, e.g., preference logics [21]. We pick up two such options and see what consequences they have for preferences over attribute sets.

In terms of preference logics, attributes can be regarded as atomic propositions and attribute sets as atomic conjunctions. Thus, what we present here is a simple version of propositional preference logic where only preferences over conjunctions of atomic formulae can be expressed. This limitation allows us to make a link to formal concept analysis (FCA) [12], through which we develop techniques for learning preferences from empirical data. FCA provides a wide range of computational tools, and, despite their often unattractive theoretical complexity, they are successfully used in practical data analysis [8].

The paper is organized as follows. We start by describing the types of preferences discussed throughout the paper including two types of global preferences and *ceteris paribus* preferences. We proceed with FCA definitions and then consider each type of preferences separately providing FCA-based semantics, discussing inference, and describing methods for learning preferences from data. Finally, we show a way to take into account the conceptual structure of the data and thus reduce the learning bias down to a well-defined point.

[1] National Research University Higher School of Economics, Moscow, Russia, email: sergei.obj@gmail.com

2 PREFERENCES IN PREFERENCE LOGICS

In modal preference logics [17, 21], preference relations are modeled by accessibility relations on possible worlds, which correspond to alternatives being compared. The preference relation is often assumed to be a preorder, that is, reflexive and transitive. We stick to this assumption and denote this relation by $\leq$. It can be extended to sets of possible worlds in several ways, of which we consider two.

In von Wright's version of preference logic [22], a set Y is preferred to a set X (notation: $X \trianglelefteq_\forall Y$) if

$$\forall x \in X \forall y \in Y (x \leq y), \tag{1}$$

that is, every alternative in Y is preferred to every alternative in X. The induced relation $\trianglelefteq_\forall$ is not necessarily reflexive or irreflexive: reflexivity is violated by a set containing two incomparable alternatives, while $\{x\} \trianglelefteq_\forall \{x\}$ for every single-element set $\{x\}$. Since $X \trianglelefteq_\forall \varnothing$ and $\varnothing \trianglelefteq_\forall Y$ for all X and Y, transitivity is not preserved, either. However, this is the only way transitivity may fail; by disallowing the empty set, we obtain a transitive relation. Besides, the $\trianglelefteq_\forall$ relation can be easily transformed into a strict partial order:

$$X \triangleleft_\forall Y \iff X \trianglelefteq_\forall Y \text{ and } Y \not\trianglelefteq_\forall X.$$

A different approach is to state that Y is preferred to X if, for each alternative from X, Y contains an alternative that is at least as good:

$$\forall x \in X \exists y \in Y (x \leq y).$$

We denote this by $X \trianglelefteq_\exists Y$. For some contexts, $\trianglelefteq_\exists$-preferences are more appropriate than $\trianglelefteq_\forall$-preferences. Consider the case of a two-person turn-based game. If X and Y are sets of positions reachable from the current position in one turn, preferences between X and Y are $\trianglelefteq_\exists$-like for the player whose turn it is, since this player has control over which position from X or Y gets chosen. For the other player, $\trianglelefteq_\forall$-preferences are more appropriate.

Preferences over propositions are defined as preferences over their sets of models. Thus, ϕ is preferred to ψ if every model of ϕ is preferred to every model of ψ (for $\trianglelefteq_\forall$-preferences) or, for every model of ψ, there is a "better" model of ϕ (for $\trianglelefteq_\exists$-preferences).

With a pinch of salt (ignoring empty sets and inconsistent propositions), $\trianglelefteq_\exists$-preferences can be viewed as a relaxation of $\trianglelefteq_\forall$-preferences, but both types are *global* in that propositions are compared w.r.t. all their models. *Ceteris paribus* preferences put restrictions on which models should be taken into account by assuming "other things being equal" when comparing ϕ and ψ. We might want to explicitly specify which other things must be equal. In the version of preference logic from [21], this is done by parameterizing the modal operator corresponding to the preference relation by a set of propositions Γ. The Γ-ceteris paribus version of $X \trianglelefteq_\forall Y$, which we denote by $X \trianglelefteq_\Gamma Y$, holds if

$$\forall x \in X \forall y \in Y (\forall \varphi \in \Gamma (x \models \varphi \iff y \models \varphi) \rightarrow x \leq y),$$

where $x \models \varphi$ means that φ is true in x. Thus, it is required that every alternative from Y is preferred to every alternative from X that satisfies exactly the same formulae from Γ. Clearly, this is a relaxation of the requirement specified by (1).

Interestingly, adding the ceteris paribus condition to the definition of $\trianglelefteq_\exists$-preferences results in stronger preferences. To say that $X \trianglelefteq_\exists Y$ holds ceteris paribus, we must find, for each alternative $x \in X$, an alternative in Y that is not only at least as good, but that is also sufficiently similar: it should satisfy exactly the same propositions from Γ that x does.

In this paper, we consider global $\trianglelefteq_\forall$- and $\trianglelefteq_\exists$-preferences and the ceteris paribus version of $\trianglelefteq_\forall$-preferences. Our discussion is restricted to rather simple propositions: we state preferences only over atomic conjunctions and allow only sets of atomic formulae as ceteris paribus conditions. Thus, in the ceteris paribus case, we consider preferences of the form $\phi \preccurlyeq_\Gamma \psi$, where ϕ and ψ are atomic conjunctions and Γ is a set of atomic formulae. We will work with ϕ, ψ, and Γ as with sets of attributes rather than as with logical formulae. The next section introduces formal concept analysis, which is the framework that we will use here.

3 FORMAL CONCEPT ANALYSIS

We start with a few definitions from FCA [12]. Given a *(formal) context* $\mathbb{K} = (G, M, I)$, where G is called a set of *objects*, M is called a set of *attributes*, and the binary relation $I \subseteq G \times M$ specifies which objects have which attributes, the *derivation operators* $(\cdot)^I$ are defined for $A \subseteq G$ and $B \subseteq M$ as follows:

$$A^I = \{m \in M \mid \forall g \in A (gIm)\}$$
$$B^I = \{g \in G \mid \forall m \in B (gIm)\}$$

A^I is the set of attributes shared by objects of A, and B^I is the set of objects having all attributes of B. Often, $(\cdot)'$ is used instead of $(\cdot)^I$. The double application of $(\cdot)'$ is a closure operator: $(\cdot)''$ is extensive, idempotent, and increasing. Sets A'' and B'' are said to be *closed*.

The left-hand side of Fig. 1 shows a context where objects are lunch options and attributes are menu items.[2] For instance, l_3 corresponds to the choice of pumpkin soup, vegetables, and ice cream.

A *(formal) concept* of the context (G, M, I) is a pair (A, B), where $A \subseteq G$, $B \subseteq M$, $A = B'$, and $B = A'$. In this case, A and B are closed. The set A is called the *extent* and B is called the *intent* of the concept (A, B). A concept (A, B) is *less general* than (C, D) if $A \subseteq C$. The set of all concepts ordered by this generality relation forms a lattice, called the *concept lattice* of the context $\mathbb{K}$.

A line diagram of the concept lattice of the context from Fig. 1 is shown in Fig. 2. Nodes correspond to concepts, with more general concepts placed above less general ones. Two concepts are connected with a line if one is less general than the other and there is no concept between the two. The extent of a concept can be read off by looking at the labels immediately below the corresponding node and below all nodes reachable by downward arcs. The intent consists of attributes indicated just above the node and those above nodes reachable by upward arcs. For example, the top-right node corresponds to the concept of all lunch options with pumpkin soup as a starter (l_1, l_3, and l_5). The node just below corresponds to its subconcept $(\{l_1, l_3\}, \{\text{pumpkin soup, vegetables}\})$.

[2] The example is inspired by the menu of the Parisian restaurant *Derrière*: `http://derriere-resto.com/restaurant/paris/derriere/menus/`. We use only a small part of the menu, though.

	pumpkin soup	black carrots	pork	vegetables	ice cream	chocolate mousse
l_1	×			×		×
l_2		×		×		×
l_3	×			×	×	
l_4		×	×			×
l_5	×		×			×

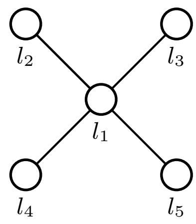

Figure 1. A preference context of lunch options

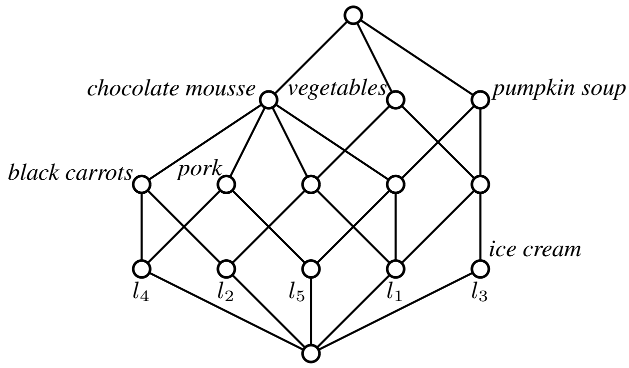

Figure 2. The concept lattice of the context in Fig. 1

It can be seen from this diagram that, if someone wants an ice cream for a dessert, she will have to have pumpkin soup as a starter and vegetables as the main dish. This is captured by the notion of an *implication*, which is, formally, an expression $A \rightarrow B$, where $A, B \subseteq M$ are attribute subsets. It *holds* or *is valid* in the context $\mathbb{K}$ (notation: $\mathbb{K} \models A \rightarrow B$) if $A' \subseteq B'$, i.e., every object of the context with all attributes from A also has all attributes from B.

If $A' = \varnothing$, then $(G, M, I) \models A \rightarrow M$. We use special notation for such zero-support implications: $A \rightarrow \bot$. Note that $A \rightarrow \bot$ is a headless Horn clause, whereas an implication $A \rightarrow B$ is a conjunction of definite Horn clauses with the same body.

The implications valid in a context are summarized by the Duquenne–Guigues basis [13], which has the minimal number of implications among equivalent implication sets. Nevertheless, it may be exponential in the size of the context, and determining its size is a #P-complete problem [16]. Other valid implications can be obtained from this basis using the Armstrong rules [4], which constitute a sound and complete inference system for implications.

The *preference context* $\mathbb{P} = (G, M, I, \leq)$ is defined in [18] as a context (G, M, I) supplied with a reflexive and transitive (as it is common in preference logics [21]) preference relation $\leq$ on G. We write $g < h$ if $g \leq h$ and $h \not\leq g$. The right-hand side of Fig. 1 shows preferences over lunch options: l_1 is better than l_4 and l_5, but worse than l_2 and l_3, while l_2 and l_3 are incomparable, as are l_4 and l_5. A preference context can be regarded as a combination of two formal contexts: (G, M, I) and $(G, G, \leq)$. We use $(\cdot)'$ for the derivation operators of (G, M, I) and $(\cdot)^{\leq}$ and $(\cdot)^{\geq}$ for the derivation operators of $(G, G, \leq)$: $X^{\leq}$ ($X^{\geq}$) is the set of all objects that are at least (at most) as good as all objects from $X \subseteq G$.

4 MODELING PREFERENCES IN FCA

In Sects. 4.1 and 4.2, we recall (without proofs) results from [18] concerning preferences based on the relations $\trianglelefteq_\forall$ and $\trianglelefteq_\exists$. In Sect. 4.3, we present a new approach to modeling ceteris paribus preferences.

We define semantics for preferences by describing conditions under which a preference π is said to be valid in a preference context $\mathbb{P}$, denoted by $\mathbb{P} \models \pi$. We say that a preference π *follows from* (or *is a semantic consequence of*) a set of preferences Π (notation: $\Pi \models \pi$) if, whenever all preferences from Π are valid in some preference context $\mathbb{P}$ (Π is *sound* for $\mathbb{P}$; $\mathbb{P} \models \Pi$), the preference π is also valid in $\mathbb{P}$ ($\mathbb{P} \models \pi$). A set Π of preferences (of a certain kind) is said to be *complete* for $\mathbb{P}$ if, for all preferences π (of this kind), $\mathbb{P} \models \pi$ if and only if $\Pi \models \pi$. If, in addition, none of the preferences in Π follows from the other preferences, we say that Π is a *preference basis* of $\mathbb{P}$.

4.1 Universal preferences

It is possible to summarize $\trianglelefteq_\forall$-preferences over subsets of G by the concept lattice of the formal context $(G, G, \leq)$. Indeed, $X \trianglelefteq_\forall Y$ holds for $X, Y \subseteq G$ if and only if $Y \subseteq X^{\leq}$. Sets X and Y are maximal with respect to this property if and only if (X, Y) is a formal concept of $(G, G, \leq)$. At the same time, if $X \trianglelefteq_\forall Y$, then $U \trianglelefteq_\forall V$ for every $U \subseteq X$ and $V \subseteq Y$. Thus, concepts of $(G, G, \leq)$ provide a complete representation of $\trianglelefteq_\forall$-preferences over object sets.

Having defined preferences over object sets, there is an easy way to translate the definition into preferences over attribute sets by associating each attribute set A with the set of objects that have all attributes from A or, to put it in terms of formal concept analysis, with A':

Definition 1. *A set of attributes $B \subseteq M$ is* universally preferred *to a set of attributes $A \subseteq M$ in a preference context $\mathbb{P} = (G, M, I, \leq)$ if $A' \trianglelefteq_\forall B'$, i.e.,*

$$\forall x \in A' \forall y \in B' (x \leq y).$$

Notation: $\mathbb{P} \models A \preccurlyeq_\forall B$.

That is, $A \preccurlyeq_\forall B$ holds (or is valid) in $(G, M, I, \leq)$ if every object with all attributes from B is preferred to every object with all attributes from A. This is precisely the approach used in preference logics as described in Sect. 2.

It is easy to obtain the following characterization of universal preferences in terms of the derivation operators of the preference context:

Proposition 1. $\mathbb{P} \models A \preccurlyeq_\forall B$ *if and only if* $B' \subseteq A'^{\leq}$.

To give an example, in the preference context from Fig. 1, we have $\{\text{pork}\} \preccurlyeq_\forall \{\text{vegetables}\}$, since every option with vegetables is preferred to every option with pork.

Proposition 2. *A sound and complete inference system for universal preferences consists of a single rule:*

$$\frac{A \preccurlyeq_\forall B}{A \cup C \preccurlyeq_\forall B \cup D},$$

which allows one to add arbitrary attributes to both sides of a valid preference.

A universal preference basis of $\mathbb{P}$ can be found by representing universal preferences of $\mathbb{P}$ as implications in another formal context.

Definition 2. *Let $\mathbb{P} = (G, M, I, \leq)$ be a preference context. The* universal translation *of $\mathbb{P}$ is a formal context $\mathbb{K}^{\mathbb{P}}_\forall = (G \times G, (M \times \{1, 2\}) \cup \{\leq\}, I_\forall)$, where*

$$\begin{aligned} (g_1, g_2) I_\forall m_1 &\iff g_1 I m, \\ (g_1, g_2) I_\forall m_2 &\iff g_2 I m, \\ (g_1, g_2) I_\forall \leq &\iff g_1 \leq g_2. \end{aligned}$$

Here, m_1 and m_2 stand for $(m, 1)$ and $(m, 2)$ respectively, $m \in M$. We denote the derivation operators of $\mathbb{K}^{\mathbb{P}}_\forall$ by $(\cdot)^\forall$.

$T_\forall(A \preccurlyeq_\forall B)$, the translation of a universal preference *$A \preccurlyeq_\forall B$, is the implication*

$$(A \times \{1\}) \cup (B \times \{2\}) \to \{\leq\}$$

of the formal context $\mathbb{K}^{\mathbb{P}}_\forall$.

Proposition 3. *A universal preference $A \preccurlyeq_\forall B$ is valid in a preference context $\mathbb{P} = (G, M, I, \leq)$ if and only if its translation is valid in $\mathbb{K}^{\mathbb{P}}_\forall$:*

$$\mathbb{P} \models A \preccurlyeq_\forall B \iff \mathbb{K}^{\mathbb{P}}_\forall \models T_\forall(A \preccurlyeq_\forall B).$$

The context resulting from the translation has as its objects pairs of objects of the preference context and contains two copies of each original attribute; (g_1, g_2) is associated with the first copy of m if g_1 has m and with the second copy if g_2 has m. For example, the preference $\{\text{pork}\} \preccurlyeq_\forall \{\text{vegetables}\}$ valid in the preference context $\mathbb{P}$ from Fig. 1 is translated into $\{(\text{pork}, 1), (\text{vegetables}, 2)\} \to \{\leq\}$, which is a valid implication of $\mathbb{K}^{\mathbb{P}}_\forall$.

The following proposition describes the basis of universal preferences:

Proposition 4. *Let $\mathbb{P}$ be a preference context. The set*

$$\Sigma = \{A \preccurlyeq_\forall B \mid (A \times \{1\}) \cup (B \times \{2\}) \text{ is minimal w.r.t. } \mathbb{K}^{\mathbb{P}}_\forall \models (A \times \{1\}) \cup (B \times \{2\}) \to \{\leq\}\}$$

is the minimal (in the number of preferences) basis of the universal preferences valid in $\mathbb{P}$.

In other words, to compute the basis of universal preferences of $\mathbb{P}$, we need to find minimal (by set-inclusion) attribute sets of $\mathbb{K}^{\mathbb{P}}_\forall$ that have $\leq$ in their closure. Note that this can be done without explicit construction of $\mathbb{K}^{\mathbb{P}}_\forall$.

4.2 Existential preferences

In this section, we transfer the definition of $\trianglelefteq_\exists$-preferences to attribute sets similarly to how it was done for $\trianglelefteq_\forall$- preferences:

Definition 3. *A set of attributes $B \subseteq M$ is* existentially preferred *to a set of attributes $A \subseteq M$ in a preference context $\mathbb{P} = (G, M, I, \leq)$, denoted by $\mathbb{P} \models A \preccurlyeq_\exists B$, if $A' \trianglelefteq_\exists B'$, i.e.,*

$$\forall x \in A' \exists y \in B' (x \leq y).$$

Again, we can characterize existential preferences in terms of the derivation operators of the preference context:

Proposition 5. $\mathbb{P} \models A \preccurlyeq_\exists B$ *if and only if* $A' \subseteq \bigcup_{g \in B'} g^{\geq}$.

An example of an existential preference that does not hold universally in the context from Fig. 1 is $\varnothing \preccurlyeq_\exists$ {vegetables}: for every lunch option, there is one with vegetables that is at least as good. On the other hand, {pumpkin soup, vegetables} is preferred to {black carrots, pork} both universally and existentially.

Existential preferences generalize implications:

Proposition 6. *For a preference context* $\mathbb{P} = (G, M, I, \leq)$

1. *If* $(G, M, I) \models A \to B$*, then* $\mathbb{P} \models A \preccurlyeq_\exists B$.
2. *If* $\leq$ *is the identity relation and* $\mathbb{P} \models A \preccurlyeq_\exists B$*, then* $(G, M, I) \models A \to B$.

Proposition 7. *A system of three rules*

$$\frac{}{X \preccurlyeq_\exists X}, \qquad \frac{X \preccurlyeq_\exists Y \cup U}{X \cup V \preccurlyeq_\exists Y}, \qquad \frac{X \preccurlyeq_\exists Y, \quad Y \preccurlyeq_\exists Z}{X \preccurlyeq_\exists Z}$$

is sound and complete with respect to existential preferences.

As universal preferences, existential preferences can also be translated into implications of a formal context, although the translation is of exponential size compared to the size the preference context.

Definition 4. *The* existential translation *of a preference context* $\mathbb{P} = (G, M, I, \leq)$ *is a formal context* $\mathbb{K}_\exists^\mathbb{P} = (G, \mathfrak{P}(M), I_\exists)$*, where* $\mathfrak{P}(M)$ *is the power set of* M *and*

$$gI_\exists A \iff g^{\leq} \cap A' \neq \varnothing.$$

Definition 5. *The* translation of an existential preference $A \preccurlyeq_\exists B$*, denoted by* $T_\exists(A \preccurlyeq_\exists B)$*, is the implication*

$$\{A\} \to \{B\}$$

of the formal context $\mathbb{K}_\exists^\mathbb{P}$.

Thus, existential preferences are translated into implications with single-element premises and conclusions (both elements are sets of original attributes). Such translation preserves the validity:

Proposition 8. *An existential preference* $A \preccurlyeq_\exists B$ *is valid in a preference context* $\mathbb{P}$ *if and only if its translation is valid in* $\mathbb{K}_\exists^\mathbb{P}$*:*

$$\mathbb{P} \models A \preccurlyeq_\exists B \iff \mathbb{K}_\exists^\mathbb{P} \models T_\exists(A \preccurlyeq_\exists B).$$

The set $\{A \preccurlyeq_\exists B \mid A$ is minimal and B is maximal w.r.t. $\mathbb{K}_\exists^\mathbb{P} \models \{A\} \to \{B\}\}$ is sound and complete (but possibly redundant) for $\mathbb{P}$.

Clearly, the existential translation of $(G, M, I, \leq)$ is infeasible for all but very small M. However, the representation size can be reduced by making use of the dependencies in the data. We address this issue in Sect. 5.

4.3 Ceteris paribus preferences

We now turn to context-based semantics for the ceteris paribus version of universal ($\trianglelefteq_\forall$) preferences, as described in Sect. 2.

Definition 6. *A set of attributes* $B \subseteq M$ *is* preferred ceteris paribus *to a set of attributes* $A \subseteq M$ *with respect to a set of attributes* $C \subseteq M$ *in a preference context* $\mathbb{P} = (G, M, I, \leq)$ *if* $A' \trianglelefteq_C B'$*, i.e.,*

$$\forall g \in A' \forall h \in B'(\{g\}' \cap C = \{h\}' \cap C \to g \leq h).$$

In this case, we say that the ceteris paribus preference $A \preccurlyeq_C B$ *is* valid *in* $\mathbb{P}$.

The preference {chocolate mousse} $\preccurlyeq_{\{\text{pumpkin soup}\}}$ {ice cream} holds in the context from Fig. 1 even though ice cream is preferred to chocolate mousse neither universally nor existentially.

Definition 7. *The* ceteris paribus translation *of* $\mathbb{P} = (G, M, I, \leq)$ *is a formal context* $\mathbb{K}_\sim^\mathbb{P} = (G \times G, (M \times \{1, 2, 3\}) \cup \{\leq\}, I_\sim)$*, where*

$$\begin{aligned}
(g_1, g_2)I_\sim(m, 1) &\iff g_1 I m, \\
(g_1, g_2)I_\sim(m, 2) &\iff g_2 I m, \\
(g_1, g_2)I_\sim(m, 3) &\iff \{g_1\}' \cap \{m\} = \{g_2\}' \cap \{m\}, \\
(g_1, g_2)I_\sim \leq &\iff g_1 \leq g_2.
\end{aligned}$$

We denote the derivation operators of $\mathbb{K}_\sim^\mathbb{P}$ *by* $(\cdot)^\sim$.

$T_\sim(A \preccurlyeq_C B)$*, the* translation of a ceteris paribus preference $A \preccurlyeq_C B$*, is the implication*

$$(A \times \{1\}) \cup (B \times \{2\}) \cup (C \times \{3\}) \to \{\leq\}$$

of the formal context $\mathbb{K}_\sim^\mathbb{P}$.

This is similar to the universal translation, but here we have three copies of each original attribute. We associate (g_1, g_2) with the third copy of m if either both g_1 and g_2 have m or neither of them does.

Proposition 9. $A \preccurlyeq_C B$ *is valid in a preference context* $\mathbb{P} = (G, M, I, \leq)$ *if and only if its translation is valid in* $\mathbb{K}_\sim^\mathbb{P}$*:*

$$\mathbb{P} \models A \preccurlyeq_C B \iff \mathbb{K}_\sim^\mathbb{P} \models T_\sim(A \preccurlyeq_C B).$$

Proof. Suppose that $\mathbb{P} \models A \preccurlyeq_C B$ and $(A \times \{1\}) \cup (B \times \{2\}) \cup (C \times \{3\}) \subseteq (g_1, g_2)^\sim$ for some $g_1 \in G$ and $g_2 \in G$. Then, $A \subseteq \{g_1\}'$, $B \subseteq \{g_2\}'$, and $g_1 I c$ if and only if $g_2 I c$ for all $c \in C$. The latter means that $\{g_1\}' \cap C = \{g_2\}' \cap C$. Since $A \preccurlyeq_C B$ holds in $\mathbb{P}$, we have $g_1 \leq g_2$ and $(g_1, g_2)I_\sim \leq$ as required.

Conversely, assume $\mathbb{K}_\sim^\mathbb{P} \models (A \times \{1\}) \cup (B \times \{2\}) \cup (C \times \{3\}) \to \{\leq\}$. We need to show that $g_1 \leq g_2$ whenever $A \subseteq \{g_1\}'$, $B \subseteq \{g_2\}'$, and $\{g_1\}' \cap C = \{g_2\}' \cap C$. Indeed, in this case, we have $(A \times \{1\}) \cup (B \times \{2\}) \cup (C \times \{3\}) \subseteq \{(g_1, g_2)\}^\sim$ and, consequently, $(g_1, g_2)I_\sim \leq$, i.e., $g_1 \leq g_2$. □

Definition 8. *We say that a ceteris paribus preference* $A \preccurlyeq_C B$ *is in* canonical form *if* $A \cap B = A \cap C = B \cap C$.

For every preference $A \preccurlyeq_C B$, there is a unique preference in canonical form equivalent to $A \preccurlyeq_C B$ in the sense that it holds precisely in the same preference contexts:

$$A \cup (B \cap C) \preccurlyeq_{C \cup (A \cap B)} B \cup (A \cap C).$$

Proposition 10. *Let* $\mathbb{P}$ *be a preference context. The set*

$$\Pi = \{A \preccurlyeq_C B \mid (A \times \{1\}) \cup (B \times \{2\}) \cup (C \times \{3\}) \text{ is minimal w.r.t. } \mathbb{K}_\sim^\mathbb{P} \models T_\forall(A \preccurlyeq_C B) \text{ and } A \cap B = A \cap C = B \cap C\}$$

is sound and complete for $\mathbb{P}$.

Proof. Due to Proposition 9, all ceteris paribus preferences from Π are valid in $\mathbb{P}$. To see that Π is complete, we consider, without loss of generality, an arbitrary preference $A \preccurlyeq_C B$ in the canonical form. If $\mathbb{P} \models A \preccurlyeq_C B$, then the implication $T_\forall(A \preccurlyeq_C B)$ holds and, therefore, either $A \preccurlyeq_C B \in \Pi$ or there are smaller sets $A_1 \subseteq A, B_1 \subseteq B$, and $C_1 \subseteq C$ such that $A_1 \preccurlyeq_{C_1} B_1 \in \Pi$. It is not hard to see that $\Pi \models A \preccurlyeq_C B$ holds then. □

We will not describe an inference system for ceteris paribus preferences similar to those provided by Propositions 2 and 7. Instead, we give an algorithm that decides whether a preference $A \preccurlyeq_C B$ follows from a set of preferences Π. By replacing A, B, and C in a valid preference $A \preccurlyeq_C B$ by their arbitrary supersets, we get valid preferences (cf. Proposition 2). The next definition captures preferences that can be obtained from other preferences in this way:

Definition 9. *Let Π be a set of ceteris paribus preferences. Then*

$$\Pi^\bullet = \{D \preccurlyeq_F E \mid \exists A \preccurlyeq_C B \in \Pi (A \subseteq D, B \subseteq E, C \subseteq F)\}.$$

Note that $\Pi \models \Pi^\bullet$. However, not all preferences that follow from Π are in $\Pi^\bullet$.

Proposition 11. *Let Π be a set of ceteris paribus preferences over M. For any preference $A \preccurlyeq_C B$ in canonical form, we have $\Pi \models A \preccurlyeq_C B$ if and only if $\Pi^\bullet$ contains all canonical-form preferences $D \preccurlyeq_F E$ such that $A \subseteq D, B \subseteq E, C \subseteq F$, and $M = D \cup E \cup F$.*

Proof. Let $D \preccurlyeq_F E \notin \Pi^\bullet$ be a preference satisfying the conditions above. Consider a preference context $\mathbb{P}$ with only two objects, $g_1 < g_2$, such that $\{g_1\}' = E$ and $\{g_2\}' = D$. The two objects have the same values for all attributes in F: each has all attributes in $E \cap F = D \cap F$ and none of the other attributes in F. The values of all attributes in $M \setminus F = (D \cup E) \setminus F$ are different for g_1 and g_2. Since $A \subseteq \{g_2\}', B \subseteq \{g_1\}'$, and $C \subseteq F$, we conclude that $\mathbb{P} \not\models A \preccurlyeq_C B$. Consider an arbitrary $P \preccurlyeq_R Q \in \Pi$. As $D \preccurlyeq_F E \notin \Pi^\bullet$, either $P \not\subseteq D$ or $Q \not\subseteq E$ or $R \not\subseteq F$. In all these cases, $\mathbb{P} \models P \preccurlyeq_R Q$. Thus, $\mathbb{P} \models \Pi$, but $\mathbb{P} \not\models A \preccurlyeq_C B$. It follows that $\Pi \not\models A \preccurlyeq_C B$.

For the other direction, suppose that $\Pi \not\models A \preccurlyeq_C B$. Then, there is a context $\mathbb{P}$ such that $\mathbb{P} \models \Pi$, but $\mathbb{P} \not\models A \preccurlyeq_C B$. This context must contain two objects, g_1 and g_2, for which $A \preccurlyeq_C B$ fails, i.e., $B \subseteq \{g_1\}', A \subseteq \{g_2\}', \{g_1\}' \cap C = \{g_2\}' \cap C$, but $g_2 \not\leq g_1$. Denote $D = \{g_2\}', E = \{g_1\}'$, and $F = (M \setminus (D \cup E)) \cup (D \cap E)$. Obviously, $D \preccurlyeq_F E$ is a canonical-form preference satisfying the conditions listed in the proposition, but $\mathbb{P} \not\models D \preccurlyeq_F E$ and, therefore, $\Pi^\bullet$ cannot contain $D \preccurlyeq_F E$, which concludes the proof. ☐

Proposition 11 paves the way for Algorithm 1, which checks whether a preference $A \preccurlyeq_C B$ is a consequence of the set Π of ceteris paribus preferences. The algorithm starts by computing the canonical form of $A \preccurlyeq_C B$ and putting the result, $A_1 \preccurlyeq_{C_1} B_1$, on a stack: $A \preccurlyeq_C B$ follows from Π if and only if $A_1 \preccurlyeq_{C_1} B_1$ does. It then tries to find a canonical-form preference $D \preccurlyeq_F E \notin \Pi^\bullet$ such that $A_1 \subseteq D, B_1 \subseteq E, C_1 \subseteq F$, and $M = D \cup E \cup F$. We know from Proposition 11 that $\Pi \not\models A_1 \preccurlyeq_{C_1} B_1$ and, consequently, $\Pi \not\models A \preccurlyeq_C B$, if and only if such a preference can be found. The algorithm searches for it in a depth-first manner, by replacing the first preference $D \preccurlyeq_F E$ on the stack with three extensions adding an arbitrary attribute from $M \setminus (D \cup E \cup F)$ to either of D, E, and F. Note that the resulting preferences are still in canonical form. On the other hand, if we add the same attribute to exactly two of D, E, and F, the resulting preference will not be in canonical form. By adding the same attribute to all the three sets, we obtain a weaker canonical-form preference, which we we can ignore, since it is not contained in $\Pi^\bullet$ only if neither of the three other extensions is. If, at some point, the algorithm comes across a preference $D \preccurlyeq_F E \in \Pi^\bullet$, it simply removes it from the stack, because all its extensions must also be in $\Pi^\bullet$. Thus, if the stack becomes empty, we know that all canonical-form preferences of the sort required by Proposition 11 are in $\Pi^\bullet$ and conclude that $\Pi \models A \preccurlyeq_C B$. If we find a preference that cannot be extended with additional attributes and is not in $\Pi^\bullet$, we conclude that $\Pi \not\models A \preccurlyeq_C B$.

Algorithm 1 CETERIS PARIBUS CONSEQUENCE$(A \preccurlyeq_C B, \Pi)$

Input: A ceteris paribus preference $A \preccurlyeq_C B$ and a set Π of ceteris paribus preferences (over a universal set M).
Output: **true**, if $\Pi \models A \preccurlyeq_C B$; **false**, otherwise.

```
S := [A ∪ (B ∩ C) ≼_{C∪(A∩B)} B ∪ (A ∩ C)]        {stack}
repeat
    D ≼_F E := pop(S)
    if D ≼_F E ∉ Π• then
        X := M \ (D ∪ E ∪ F)
        if X = ∅ then
            return false
        choose m ∈ X
        push(D ∪ {m} ≼_F E, S)
        push(D ≼_F E ∪ {m}, S)
        push(D ≼_{F∪{m}} E, S}
until empty(S)
return true
```

Algorithm 1 is exponential in $|M|$ in the worst case, but there is little hope to do better. The reason is that, although Proposition 10 makes it possible to represent ceteris paribus preferences as implications, or Horn formulae, these Horn formulae are not sufficient to generate the theory implied by the preferences: we must add the disjunctions $\neg m_i \vee \neg m_j \vee m_k$ for different $i, j, k \in \{1, 2, 3\}$ and, crucially, $m_1 \vee m_2 \vee m_3$ for each $m \in M$. The last disjunction is not a Horn clause, which makes inference hard. However, the algorithm is linear in $|\Pi|$, which makes it efficient in applications where the language for describing preferences (and, thus, the number of attributes) is fixed and small compared to the number of preferences that need to be taken into account.

5 REDUCING BIAS

The presented approach to deriving preferences assumes that the attribute combinations in the context are the only ones that matter. In practice, the data may cover only a small fraction of possible combinations. Derived preferences hold in the data, but may not hold in the entire domain, being *biased* towards the observed part of the data.

In our lunch context, $\{\text{pork}\} \preccurlyeq_\forall \{\text{vegetables}\}$, but every option with pork there comes with chocolate mousse. It may well be that the subject does not like the combination and the true preference is weaker: $\{\text{pork, chocolate mousse}\} \preccurlyeq_\forall \{\text{vegetables}\}$.

In this section, we outline a conservative approach to preference learning, which separates knowledge about preferences from knowledge about the structure of the underlying context and makes it possible to reduce bias down to a certain well-defined point. We start by extending the definition of semantic consequence to cover both implications and preferences under the same hood. If $\mathcal{H}$ is a set of implications over M and Π is a set of preferences (of a certain kind) over subsets of M, we say that $\pi \in \Pi$ *follows from* (or *is a semantic consequence of*) $\mathcal{H} \cup \Pi$ (notation: $\mathcal{H} \cup \Pi \models \pi$) if, whenever all preferences from Π are valid in some preference context $\mathbb{P}$ over M satisfying all implications from $\mathcal{H}$ (i.e., $\mathbb{P} \models \Pi$ and $\mathbb{P} \models \mathcal{H}$), the preference π is also valid in $\mathbb{P}$ (i.e., $\mathbb{P} \models \pi$).

Definition 10. *The* Horn bias *induced by a preference context $\mathbb{P} = (G, M, I, \leq)$ is the set of implications that hold in (G, M, I).*

The Horn bias induced by a preference context is simply the implicational (i.e., Horn) theory behind its "non-preferential" part.

Definition 11. *Let $\mathcal{H}$ be the Horn bias induced by $\mathbb{P} = (G, M, I, \leq)$ and Π be the set of all preferences (of a certain kind) that hold in $\mathbb{P}$. We say that a preference $\pi \in \Pi$ is* Horn-biased *in $\mathbb{P}$ if there is $\Pi_1 \subseteq \Pi \setminus \{\pi\}$ such that $\Pi_1 \not\models \pi$ and $\mathcal{H} \cup \Pi_1 \models \pi$.*

Intuitively, a Horn-biased preference is one that can be deduced from other—weaker—preferences given that we know the Horn theory behind the data, but not without this additional knowledge. In the example above, $\{\text{pork}\} \preccurlyeq_\forall \{\text{vegetables}\}$ is Horn-biased, since it is a consequence of $\mathcal{H} \cup \{\{\text{pork, chocolate mousse}\} \preccurlyeq_\forall \{\text{vegetables}\}\}$, where $\mathcal{H}$ is the set of all implications valid in the context including $\{\text{pork}\} \to \{\text{chocolate mousse}\}$.

For universal and existential preferences, the Horn bias can be avoided by considering only preferences over closed attribute sets. Any preference of the form $A'' \preccurlyeq_\forall B''$ or $A'' \preccurlyeq_\exists B''$ is guaranteed not to be Horn-biased, but all other universal and existential preferences are Horn-biased. A'' and B'' are concept intents of (G, M, I); thus, unbiased preferences are preferences over formal concepts.

Technically, there are at least two ways to achieve an unbiased representation of universal preferences without constructing the basis from Proposition 4. One is to build the Duquenne–Guigues basis of (G, M, I), transform its implications into *background knowledge*, and then build the basis of $\mathbb{K}_\forall^\mathbb{P}$ relative to this background knowledge (see [18] for more details). The other is to build the so-called *minimal hypotheses* for $\leq$ in $\mathbb{K}_\forall^\mathbb{P}$ [11]. The results of the two approaches are identical: it is the minimal basis of unbiased universal preferences.

If we want to keep preferences unbiased, but be able to derive biased preferences, too, we can do this using implications and a hybrid inference system that combines the Armstrong rules [4] for implications, the rule from Proposition 2, and three additional rules:

$$\frac{X \to \bot}{\varnothing \preccurlyeq_\forall X, \quad X \preccurlyeq_\forall \varnothing}, \quad \frac{X \to Y, \quad X \cup Y \preccurlyeq_\forall Z}{X \preccurlyeq_\forall Z},$$

$$\frac{X \to Y, \quad Z \preccurlyeq_\forall X \cup Y}{Z \preccurlyeq_\forall X}.$$

For existential preferences, an unbiased representation is actually easier to compute than a biased one: in this case, not all attribute sets are needed for the translation, but only concept intents of (G, M, I).

Definition 12. *The* conceptual existential translation *of a preference context $\mathbb{P}$ is a formal context $\mathbb{C}_\exists^\mathbb{P} = (G, \mathfrak{B}(G, M, I), I_\exists)$, where $\mathfrak{B}(G, M, I)$ is the concept set of (G, M, I) and*

$$gI_\exists(A, B) \iff g^\leq \cap A \neq \varnothing.$$

For the lunch example, this translation produces a context with fifteen attributes corresponding to the concepts shown in Fig. 2 compared to 64 attributes produced by the existential translation.

Definition 13. *The* conceptual translation of an existential preference *$A \preccurlyeq_\exists B$, denoted by $T_\exists^\mathbb{C}(A \preccurlyeq_\exists B)$, is the implication*

$$\{(A', A'')\} \to \{(B', B'')\}$$

of the formal context $\mathbb{C}_\exists^\mathbb{P}$.

The conceptual translation preserves the validity of existential preferences and provides another way to summarize them:

$$\{A \preccurlyeq_\exists B \mid \mathbb{C}_\exists^\mathbb{P} \models \{(A', A)\} \to \{(B', B)\} \text{ and } B \not\subseteq A\}$$

is a complete set of existential preferences *relative* to the implications of (G, M, I). A hybrid inference system for implications and existential preferences includes Armstrong rules [4], the rules for existential preferences from Proposition 7, and the rule

$$\frac{A \to B}{A \preccurlyeq_\exists B}.$$

For ceteris paribus preferences, bias can be reduced even further.

Definition 14. *We call the expression $[A, B]C \Rightarrow D[E, F]$ a* doubly conditional functional dependency *and say that it* holds *in (G, M, I) if, for every $g, h \in G$ such that $g \in A', h \in B'$, and $\{g\}' \cap C = \{h\}' \cap C$, we have $g \in E', h \in F'$, and $\{g\}' \cap D = \{h\}' \cap D$.*

This generalizes both implications and conditional functional dependencies from [9]. Thus, the induced bias, which we call the 2CFD bias, includes the Horn bias.

Definition 15. *The* 2CFD bias *induced by a preference context $\mathbb{P} = (G, M, I, \leq)$ is the set of doubly conditional functional dependencies that hold in the formal context (G, M, I).*

Doubly conditional functional dependencies are in one-to-one correspondence with implications of the context obtained from $\mathbb{K}_\sim^\mathbb{P}$ by removing the $\leq$ attribute. To avoid the 2CFD bias, we should consider only preferences translated into implications of $\mathbb{K}_\sim^\mathbb{P}$ whose left-hand side X is minimal w.r.t. $X \cup \{\leq\}$ being a concept intent of $\mathbb{K}_\sim^\mathbb{P}$. These correspond to minimal hypotheses for $\leq$ [11].

6 CONCLUSION

We have proposed a formalism based on concept lattices for modeling several types of preferences, including preferences that hold only *ceteris paribus* and showed how such preference models can be learned from data. Our approach may seem limited for, taken literally, it is only concerned with preferences over conjunctions of boolean variables; even negations of variables are not covered. Compare this to other approaches, such as *cp-theories* as defined in [23]. In this framework, one works with a set of variables V, each of which has an associated set of values. A *conditional preference* is a statement of the form $u : x_1 > x_2[W]$, where u is an assignment to $U \subseteq V$, x_1 and x_2 are different assignments to some $X \in V$, and W is a subset of $V \setminus (U \cup \{X\})$. Such preference is interpreted as follows: between two alternatives satisfying u, the one with $X = x_1$ is preferred to the one with $X = x_2$ provided that they agree on all other variables with a possible exception of those in W. This may be regarded as a generalization of CP-nets [6] and TCP-nets [7].

To model such preferences in our framework, we can build a preference context whose attribute set M consists of expressions of the form $X = x$, where $X \in U$ and x ranges over possible values of X. For the boolean case, this would mean adding a negated copy for each attribute. Then, a strict conditional preference $u : x_1 > x_2[W]$ would have the following weak counterpart in our framework:

$$u \cup \{X = x_2\} \preccurlyeq_{\{M \setminus W\}} u \cup \{X = x_1\}.$$

To express strict conditional preferences, we can start with a strict preference relation over objects. On the other hand, the language of conditional preferences only allows preferences of a single variable, whereas, with our approach, we can express (and learn from data) more general preferences such as

$$u \cup \{X = x_2, Y = y_2\} \preccurlyeq_{\{M \setminus W\}} u \cup \{X = x_1, Y = y_1\}.$$

Furthermore, for variables with ordinal values, we could use attributes of the form, e.g., $x_1 \leq X \leq x_2$ instead of just $X = x$. In FCA, this is done by *scaling* so-called *many-valued contexts*, in which attributes are not necessarily boolean [12]. Also, the ceteris paribus conditions in the translated context from Definition 7, which are specified through attributes from $M \times \{3\}$, could be customized to specify relations other than equality. This would make it possible to express preferences like the following: "Between two ways of travel, I prefer a cheap one provided that it is at least as fast as the other." We leave a thorough treatment of these issues and a proper comparison to other approaches to preference modeling for further research.

We also plan to develop algorithms for learning preferences from queries [2]. Such algorithms exist, e.g., for CP-nets [15]. Since, in our framework, preferences can be translated into Horn clauses, it might be possible to adapt the output-polynomial algorithm for learning Horn theories from [3] (adaptation is needed, because the translation is not surjective, i.e., not all Horn clauses over a given set of variables correspond to preferences). However, this algorithm uses equivalence queries, which are hard to answer. An alternative approach is a similar technique from FCA, called *attribute exploration* [10, 19, 20], which only uses queries on the validity of implications (even though, in theory, the number of such questions may be exponentially large). Note that, with this approach, the user is not asked to specify preferences between two given examples, but rather to confirm or reject a stated preference. When rejecting a preference, the user must point out two objects contradicting this preference. A precise specification of such query learning algorithm and its computational complexity are a matter of further research.

In application to real-life data analysis, it may be useful to introduce some statistical considerations into the theory presented here. One obvious approach is to replace the semantics based on implications by one based on association rules [1], thus, allowing exceptions in derived preferences, but making sure that these preferences are supported by a sufficiently large volume of data. On the other hand, methods for pruning concept lattices by selecting only the most interesting (in some sense) concepts [14, 5] may be of value in deriving "unbiased" preferences from Sect. 5, which are interpreted as preferences over concepts.

REFERENCES

[1] Rakesh Agrawal, Tomasz Imieliński, and Arun Swami, 'Mining association rules between sets of items in large databases', *SIGMOD Rec.*, **22**(2), 207–216, (June 1993).

[2] Dana Angluin, 'Queries and concept learning', *Machine Learning*, **2**, 319–342, (1988).

[3] Dana Angluin, Michael Frazier, and Leonard Pitt, 'Learning conjunctions of Horn clauses', *Machine Learning*, **9**, 147–164, (1992).

[4] William Ward Armstrong, 'Dependency structures of data base relationships', *Proc. IFIP Congress*, 580–583, (1974).

[5] Mikhail Babin and Sergei O. Kuznetsov, 'Approximating concept stability', in *Formal Concept Analysis*, eds., Florent Domenach, Dmitry Ignatov, and Jonas Poelmans, volume 7278 of *Lecture Notes in Computer Science*, 7–15, Springer Berlin / Heidelberg, (2012).

[6] Craig Boutilier, Ronen I. Brafman, Holger H. Hoos, and David Poole, 'Reasoning with conditional ceteris paribus preference statements', in *UAI*, eds., Kathryn B. Laskey and Henri Prade, pp. 71–80. Morgan Kaufmann, (1999).

[7] Ronen I. Brafman and Carmel Domshlak, 'Introducing variable importance tradeoffs into cp-nets', in *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*, pp. 69–76. Morgan Kaufmann Publishers Inc., (2002).

[8] Claudio Carpineto and Giovanni Romano, *Concept Data Analysis: Theory and Applications*, Wiley, 2004.

[9] Wenfei Fan, Floris Geerts, Laks V.S. Lakshmanan, and Ming Xiong, 'Discovering conditional functional dependencies', in *Data Engineering, 2009. ICDE'09. IEEE 25th International Conference on*, pp. 1231–1234, (2009).

[10] Bernhard Ganter, 'Attribute exploration with background knowledge', *Theoretical Computer Science*, **217**(2), 215–233, (1999).

[11] Bernhard Ganter and Sergei O. Kuznetsov, 'Formalizing hypotheses with concepts', in *ICCS*, eds., Bernhard Ganter and Guy W. Mineau, volume 1867 of *Lecture Notes in Computer Science*, pp. 342–356. Springer, (2000).

[12] Bernhard Ganter and Rudolf Wille, *Formal Concept Analysis: Mathematical Foundations*, Springer, Berlin/Heidelberg, 1999.

[13] Jean-Luc Guigues and Vincent Duquenne, 'Famille minimale d'implications informatives résultant d'un tableau de données binaires', *Mathématiques et Sciences Humaines*, **24**(95), 5–18, (1986).

[14] Mikhail Klimushkin, Sergei Obiedkov, and Camille Roth, 'Approaches to the selection of relevant concepts in the case of noisy data', in *Formal Concept Analysis*, eds., Léonard Kwuida and Baris Sertkaya, volume 5986 of *Lecture Notes in Computer Science*, 255–266, Springer Berlin / Heidelberg, (2010).

[15] Frédéric Koriche and Bruno Zanuttini, 'Learning conditional preference networks', *Artificial Intelligence*, **174**(11), 685 – 703, (2010).

[16] Sergei O. Kuznetsov and Sergei Obiedkov, 'Some decision and counting problems of the Duquenne–Guigues basis of implications', *Discrete Applied Mathematics*, **156**(11), 1994–2003, (2008).

[17] Fenrong Liu, *Changing for the Better. Preference Dynamics and Agent Diversity*, Ph.D. dissertation, Universiteit van Amsterdam, 2008.

[18] Sergei Obiedkov, 'Modeling preferences over attribute sets in formal concept analysis', in *Formal Concept Analysis*, volume 7278 of *Lecture Notes in Artificial Intelligence*, 227–243, Springer Berlin / Heidelberg, (2012).

[19] Sergei Obiedkov, Derrick G. Kourie, and J.H.P. Eloff, 'Building access control models with attribute exploration', *Computers & Security*, **28**(1-2), 2–7, (2009).

[20] Artem Revenko and Sergei O. Kuznetsov, 'Attribute exploration of properties of functions on sets', *Fundamenta Informaticae*, **115**(4), 377–394, (2012).

[21] Johan van Benthem, Patrick Girard, and Olivier Roy, 'Everything else being equal: A modal logic for *ceteris paribus* preferences', *J. Philosophical Logic*, **38**(1), 83–125, (2009).

[22] Georg H. von Wright, *The Logic of Preference*, Edinburgh University Press, 1963.

[23] Nic Wilson, 'Computational techniques for a simple theory of conditional preferences', *Artif. Intell.*, **175**(7-8), 1053–1091, (2011).