

Multi-Attribute Preference Logic

Koen V. Hindriks
Man-Machine Interaction
Delft University of Technology
The Netherlands
k.v.hindriks@tudelft.nl

Wietske Visser
Man-Machine Interaction
Delft University of Technology
The Netherlands
wietske.visser@tudelft.nl

Catholijn M. Jonker
Man-Machine Interaction
Delft University of Technology
The Netherlands
c.m.jonker@tudelft.nl

ABSTRACT

Preferences for objects are commonly derived from ranked sets of properties or multiple attributes associated with these objects. There are several options or strategies to qualitatively derive a preference for one object over another from a property ranking. We introduce a modal logic, called multi-attribute preference logic, that provides a language for expressing such strategies. The logic provides the means to represent and reason about qualitative multi-attribute preferences and to derive object preferences from property rankings. The main result of the paper is a proof that various well-known preference orderings can be defined in multi-attribute preference logic.

1. INTRODUCTION

Preferences may be associated with various entities such as states of affairs, properties, objects and outcomes in e.g. games. Our main concern here are object preferences. A natural approach to obtain preferences about objects is to start with a set of properties of these objects and derive preferences from a ranking of these properties, where the ranking indicates the relative importance or priority of each of these properties. This approach to obtain preferences is typical in multi-attribute decision theory, see e.g. Keeney and Raiffa [10]. Multi-attribute decision theory provides a quantitative theory that derives object preferences from utility values assigned to outcomes which are derived from numeric weights associated with properties or attributes of objects. As it is difficult to obtain such quantitative utility values and weights, however, several qualitative approaches have been proposed instead, see e.g. [2, 4, 5, 6, 11]. There is also extensive literature on preference logic following the seminal work of Von Wright [12, 9], but such logics are not specifically suited for the multi-attribute case. To illustrate what we are after, we first present a motivating example that is used throughout the paper.

EXAMPLE 1. Suppose we want to buy a house. The properties that we find important are that we can afford the house, that it is close to our work, and that it is large, in that order. Consider three houses, $house_1$, $house_2$ and $house_3$, whose properties are listed in




Figure 1: Properties of three houses

Figure 1, which we have to order according to our preferences. It seems clear that we would prefer $house_1$ over the other two, because it has two of the most important properties, while both other houses only have one of these properties. But what about the relative preference of $house_2$ and $house_3$? $house_3$ has two out of three of the relevant properties where $house_2$ has only one. If the property that $house_2$ has is considered more important than both properties of $house_3$, $house_2$ would be preferred over $house_3$.

Key to a logic of multi-attribute preferences is the representation of property rankings. Encodings of property rankings have been explored in Coste-Marquis *et al.* [6] where they are called goal bases, and in Brewka [4] where they are called ranked knowledge bases. Such ranked goals are binary, and in this paper we also consider desired attributes that are binary (as opposed to numeric or ordinal ones). Coste-Marquis *et al.* and Brewka moreover discuss various options, or strategies, for deriving object preferences from a property ranking. The preference orderings thus obtained are not expressed in a logic, however. Brewka *et al.* [5] propose a non-monotonic logic called qualitative choice logic to reason about multi-attribute preferences. An alternative approach towards a logic of multi-attribute preferences is presented in Liu [11] where property rankings called priority sequences are encoded in first-order logic. Both approaches are based on one particular strategy, namely lexicographic ordering, and cannot be used to reason about preference orderings.

In this paper a generic logic of qualitative multi-attribute preferences is proposed in which property rankings and associated strategies for deriving object preferences from such rankings can be defined. In Section 2 the syntax and semantics of multi-attribute preference logic is introduced. Section 3 shows how various strategies to obtain object preferences from a property ranking can be defined

in the logic. Section 4 presents the main result of the paper and shows that property rankings encoded as ranked knowledge bases and a number of related strategies to obtain preference orderings can be equivalently translated into multi-attribute preference logic. Section 5 concludes the paper.

2. MULTI-ATTRIBUTE PREFERENCE LOGIC

2.1 Syntax and semantics

The logic of multi-attribute preferences that we introduce here is an extension of the modal binary preference logic presented in [7]. This logic is a propositional modal logic with a modal operator $\Box^{\leq}\varphi$, and its dual $\Diamond^{\leq}\varphi$. Here $\Box^{\leq}\varphi$ expresses that φ is true in all states that are at least as good as the current state. Binary preference relations over formulae are subsequently defined. One of the more natural binary preference statements is $\varphi <_{VV} \psi$ which expresses that any state where ψ is true is strictly better than any state where φ is true. That is, whenever φ is the case, ψ is preferred, and never vice versa. By adding a global modality U to the language, the binary preference operator $<_{VV}$ can be defined by $U(\psi \rightarrow \Box^{\leq}\neg\varphi)$, when it is assumed that the underlying order on worlds or states has been completely specified, i.e. is total.

Multi-attribute preference logic adds two operators to binary preference logic. First, multi-attribute preference logic, as in hybrid logic [1] adds names for objects to the language by adding nullary modal operators i, j to the language. The semantics of the operators introduced here, however, differs from the standard semantics of hybrid logic. Here i, j are used as names for objects which semantically are more complex entities than the usual worlds of modal semantics. In order to avoid confusion, we will refer to i, j as object names below. This language extension allows us to talk about objects and associated preferences explicitly.

Second, the logic introduces a new modal operator $\Box^{\#}$. The language of multi-attribute preference logic consists of four unary modal operators. Instead of the single operator \Box^{\leq} it is more convenient to introduce the two operators $\Box^{<}$ and $\Box^{\#}$: informally, $\Box^{<}\varphi$ expresses that at all worlds that are ranked higher than the current one φ is true, whereas $\Box^{\#}\varphi$ expresses that at all worlds that are equally ranked to the current one φ is true. The modal operator $\Box^{\#}$ is introduced to inspect worlds that are not ranked equally to the current one.

DEFINITION 2. (Language)

Let At be a set of propositional atoms with typical element p and Nom be a set of names, with typical elements i, j . The language \mathcal{L}_{pref} is defined as follows:

$$\varphi \in \mathcal{L}_{pref} ::= p \mid i \mid \neg\varphi \mid \varphi \wedge \psi \mid \Box^{\leq}\varphi \mid \Box^{\#}\varphi \mid \Box^{<}\varphi \mid U\varphi$$

Disjunction \vee , implication \rightarrow , and bi-implication \leftrightarrow are defined as the usual abbreviations. $\Diamond^{<}\varphi, \Diamond^{\#}, \Diamond^{\leq}$ are abbreviations for $\neg\Box^{<}\neg\varphi, \neg\Box^{\#}\neg\varphi$, and $\neg\Box^{\leq}\neg\varphi$. $\Box^{\leq}\varphi$ is short for $\Box^{<}\varphi \vee \Box^{\#}\varphi$ and $\Diamond^{\leq}\varphi$ is its dual. The dual of the global modal operator, $E\varphi$, is defined as $\neg U\neg\varphi$. We also write $U_i\varphi$ for $U(i \rightarrow \varphi)$ and $E_i\varphi$ for $E(i \wedge \varphi)$ for $i \in Nom$. Finally, the set of purely propositional formulae is denoted by \mathcal{L}_0 and consists of all formulae without any occurrences of modal operators or names $i \in Nom$. $\varphi \in \mathcal{L}_0$ is also called an objective formula.

The basic concepts in the semantics for multi-attribute preference logic are objects and properties those objects may have. Properties are naturally represented by sets of worlds. As we want to use properties to classify the ranking of objects, properties are ordered in correspondence with their relative importance; such an order is called a property ranking here. To order properties, i.e. sets of worlds, it is required that properties are disjoint sets of worlds. Property rankings will be derived from an order on worlds below.

Objects are also identified with particular sets of worlds. The idea is that the properties (in the sense of the previous paragraph) of an object can be derived from the worlds which define the object. To ensure that objects are coherent, that is have a uniquely defined set of properties, the worlds that define the object need to be copies of each other, which means that these worlds need to assign the same truth values to propositional atoms. Objects are identified with equivalence classes of worlds with respect to a truth assignment.

DEFINITION 3. (Object)

Let W be a set of worlds and V be a mapping of W to truth assignments 2^{At} . An *object* is an equivalence class on W with respect to V . The set \mathcal{O}_V denotes the set of all objects defined by W and V and is formally defined by:

$$\mathcal{O}_V = \{[w]_V \mid w \in W\}$$

where $[w]_V = \{v \in W \mid V(w) = V(v)\}$. Whenever V is clear from the context, we drop the subscript V . As an object o is the equivalence class of a world w with respect to V , we also say that world w identifies object o .

DEFINITION 4. (Model)

A *multi-attribute preference model* \mathcal{M} is a tuple $\langle W, \lesssim, V, N \rangle$ where W is a set of *worlds* with typical elements u, v, w , \lesssim is a *total pre-order* (i.e. a reflexive, transitive and total relation) on W , V is a *valuation function* mapping worlds in W onto truth assignments in 2^{At} , and N is a *naming function*. The strict subrelation $<$ of \lesssim is defined by: $v < w := v \lesssim w \ \& \ w \not\lesssim v$. We write $v \sim w$ whenever $v \lesssim w$ and $w \lesssim v$.

Although the strict order $<$ derived from \lesssim indicates a ranking of worlds where $v < w$ means that w is ranked higher than v , we do not say that w is preferred over v , because we want to reserve this terminology for talking about objects. A preference between objects is derived from the ranking \lesssim over worlds. The naming function N maps names i to objects o .

The truth definition for propositional atoms and Boolean operators is standard. Given a model $\mathcal{M} = \langle W, \lesssim, V, N \rangle$, the semantics of names $i \in Nom$ is provided by the naming function N . The truth definitions for most modal operators are also standard definitions using the associated accessibility relations for these operators. The semantic clause for $\Box^{\#}$ is defined by means of the relation \sim , which is derived from the order \lesssim . Similarly, the semantic clause for $\Box^{<}$ is provided by means of the strict order $<$. The global operator U simply inspects all worlds in a model.

The truth definition for $\Box^{\#}$ is not directly defined in terms of a given relation on W . It inspects all worlds that (i) are not ranked equally as the current one, and (ii) are not copies of worlds that are ranked

equally as the current one. The motivation for this definition will become clear in Section 2.2 when clusters are introduced.

DEFINITION 5. (Truth Definition)

Let $\mathcal{M} = \langle W, \lesssim, V, N \rangle$ be an MPL model and $w \in W$ a world. The truth of a formula $\varphi \in \mathcal{L}_{pref}$ in \mathcal{M} at w is defined by:

$$\begin{aligned} \mathcal{M}, w \models p &\Leftrightarrow p \in V(w) \\ \mathcal{M}, w \models i &\Leftrightarrow w \in N(i) \\ \mathcal{M}, w \models \neg\varphi &\Leftrightarrow \mathcal{M}, w \not\models \varphi \\ \mathcal{M}, w \models \varphi \wedge \psi &\Leftrightarrow \mathcal{M}, w \models \varphi \ \& \ \mathcal{M}, w \models \psi \\ \mathcal{M}, w \models \Box^{\square} \varphi &\Leftrightarrow \forall v : w \sim v \Rightarrow \mathcal{M}, v \models \varphi \\ \mathcal{M}, w \models \Box^{\square^{\#}} \varphi &\Leftrightarrow \forall u \in \bigcup \{ [v]_V \mid w \sim v \} : \mathcal{M}, u \models \varphi \\ \mathcal{M}, w \models \Box^{\square^{\triangleleft}} \varphi &\Leftrightarrow \forall v : w < v \Rightarrow \mathcal{M}, v \models \varphi \\ \mathcal{M}, w \models U\varphi &\Leftrightarrow \forall v : \mathcal{M}, v \models \varphi \end{aligned}$$

A name $i \in Nom$ refers to an object o and, semantically, is true at a world w that identifies the object o , i.e. $w \in o$. A name thus is a special kind of operator that is true in all worlds that identify a certain object, and false in all other worlds. We can express that an object i has a property φ by $E_i\varphi = E(i \wedge \varphi)$. As we have $E(i)$ as a validity and the worlds that identify the corresponding object o are copies of each other, we have $E_i\varphi \leftrightarrow U_i\varphi$ for objective φ . This shows that an object is coherent in the sense that an object has a consistent set of objective properties and can be uniquely identified by this set.

The language also allows us to express properties that concern comparison of objects. For example, $U(i \rightarrow \Box^{\square} j)$ expresses that for every property of object i object j has a property that is strictly better. The formula $E(j \wedge \neg \Box^{\square} i)$ expresses that object j has a property that object i cannot match, i.e. i has no property that is strictly better than this property of j . We have $E(j \wedge \neg \Box^{\square} i) \rightarrow U(i \rightarrow \Box^{\square} j)$ in multi-attribute preference logic. This validity is based on the assumption that the pre-order in models for \mathcal{L}_{pref} is total.

Recall that the binary preference operator $\varphi <_{\forall V} \psi$ can be defined as $U(\psi \rightarrow \Box^{\square} \neg\varphi)$. Using $<_{\forall V}$ it is possible to define property rankings and express that a property ψ is ranked higher than property φ . Using the truth definitions for $U\varphi$, $\Box^{\square}\varphi$ and $\Box^{\square^{\triangleleft}}\varphi$ and the definition of $\Box^{\square}\varphi$ as $\Box^{\square}\varphi \wedge \Box^{\square^{\triangleleft}}\varphi$, it can be shown that $\varphi <_{\forall V} \psi$ has the following truth definition:

$$\mathcal{M}, w \models \varphi <_{\forall V} \psi \Leftrightarrow \forall u, v : \mathcal{M}, u \models \varphi \ \& \ \mathcal{M}, v \models \psi \Rightarrow u < v$$

The intuitive reading of $\varphi <_{\forall V} \psi$ is that every ψ -state is ranked higher than every φ -state (cf. [7]). Returning to the comparison of objects again, $i <_{\forall V} j$ expresses that object j is preferred over i . The preference expressed in this way is a very strong kind of preference, however. It requires that all of object j 's relevant properties are considered more important than objects i 's properties, which corresponds with the definition of $i <_{\forall V} j$ by $U(j \rightarrow \Box^{\square} \neg i)$. In contrast, multi-attribute preference logic is able to specify principles that allow to derive preferences over objects from their properties in a weaker sense. It enables, for example, to specify orderings where object j is preferred over object i even when object i has at least one property that is considered more important than a property that object j has (compare e.g. object c and f in Figure 2). The logic thus facilitates the specification of different ordering strategies, and, given such a specification, provides the means to derive a preference of one object over another from a property ranking and an additional specification of the objects' properties.

Proposition 6 supports our claim that multi-attribute preference lo-

gic extends binary preference logic as all listed axioms of this logic are valid in multi-attribute preference logic as well (cf. [7], p. 66). We have listed only those axioms that can straightforwardly be expressed without the need to introduce additional definitions of other binary preference operators; all of the remaining axioms are valid as well in multi-attribute preference logic when such definitions are added. Below we use that \wedge and \vee bind their arguments stronger than \rightarrow to be able to remove some brackets.

PROPOSITION 6. We have the following validities:

1. $\models E_i\varphi \leftrightarrow U_i\varphi$ for $\varphi \in \mathcal{L}_0$.
2. $\models \varphi <_{\forall V} \psi \wedge U(\xi \rightarrow \psi) \rightarrow \varphi <_{\forall V} \xi$
3. $\models \varphi <_{\forall V} \psi \wedge U(\xi \rightarrow \varphi) \rightarrow \xi <_{\forall V} \psi$
4. $\models \varphi <_{\forall V} \psi \wedge \psi <_{\forall V} \xi \wedge E\xi \rightarrow \varphi <_{\forall V} \xi$
5. $\models U\neg\varphi \vee U\neg\psi \rightarrow \varphi <_{\forall V} \psi$
6. $\models \varphi <_{\forall V} \psi \rightarrow U(\varphi <_{\forall V} \psi)$

What multi-attribute preference logic adds to binary preference logic are names for objects, and most importantly, the $\Box^{\square^{\#}}$ operator that allows us to define clusters (see Section 2.2) that represent desirable attributes. All of the modal operators \Box^{\square} , $\Box^{\square^{\triangleleft}}$, $\Box^{\square^{\#}}$ and U are normal modal operators and satisfy the K axiom. In addition, we prove some properties of the \Box^{\square} and $\Box^{\square^{\#}}$ operators (some of the more obvious axioms have not been listed below). Proposition 7.3 shows that multi-attribute preference logic is related to the logic of only knowing, see [8].

PROPOSITION 7. We have:

1. $\models \Box^{\square} \Box^{\square^{\#}} \varphi \leftrightarrow \Box^{\square^{\#}} \varphi$
2. $\models \Box^{\square} \Box^{\square^{\triangleleft}} \varphi \leftrightarrow \Box^{\square^{\triangleleft}} \varphi$
3. $\models \Box^{\square} \varphi \rightarrow \neg \Box^{\square^{\#}} \varphi$ where $\neg\varphi \in \mathcal{L}_0$ is consistent

PROOF. We prove item 3. Suppose $\Box^{\square} \varphi$ is true at world w . Then φ is true in all worlds $v \sim w$. Since the truth of objective formulae is the same within an object, φ is also true in every world $u \in \{ [v]_V \mid w \sim v \}$. Since $\neg\varphi$ is a consistent objective formula and all valuations are present in the model, $\neg\varphi$ must be true in some world in the model. So there must be some world in $\{ [v]_V \mid w \sim v \}$ that satisfies $\neg\varphi$, so we have $\neg \Box^{\square^{\#}} \varphi$ at world w . \square

2.2 Clusters

The total pre-order \lesssim in a multi-attribute preference model induces a strict linear order on sets of worlds, which we call clusters. Formally, a cluster is an equivalence class induced by \lesssim . Intuitively, such clusters represent the properties or attributes considered relevant for deriving object preferences. The order on clusters induced by \lesssim represents a property ranking, i.e. the relative importance of one property compared to another. The relation between objects and properties may now be clarified as follows. The idea is that if an object has a particular property it should be represented within the cluster of worlds that represents the property. Technically, this is realized by making sure that (at least) one of the copies of a world that identifies the object is an element of the cluster that represents the property. The worlds that identify an object act as representatives for the object within a certain cluster and thus indicate that the object has that property. As clusters are disjoint and objects may have multiple properties, this also explains the need for introducing copies of worlds.

DEFINITION 8. (Cluster)

Let \lesssim be a total pre-order on W . A *cluster* c is an equivalence class induced by \lesssim , i.e. $c = [w]_{\lesssim} = \{ v \mid w \sim v \}$ for some $w \in W$.

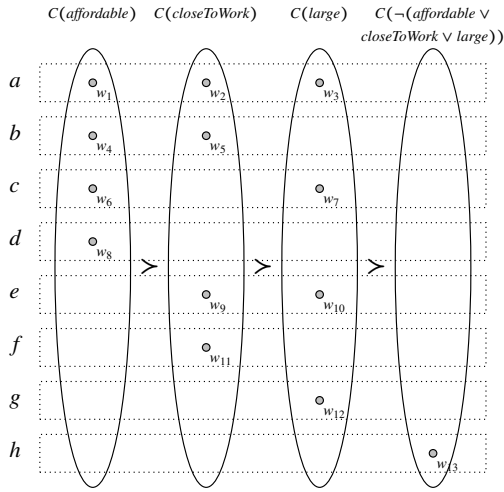


Figure 2: Visualization of an MPL model

EXAMPLE 9. The relation between clusters (properties) and sets of copies (objects) is visualized in Figure 2 (this is a model of the theory in Example 18). The ellipses (columns) represent the clusters or properties and the boxes (rows) represent objects. Objects in this case are supposed to be houses. For example, the house labelled b consists of two worlds, w_4 and w_5 . As these worlds are part of the same object, they must be copies of each other. One of these worlds, w_4 , is also part of the cluster representing the property of being affordable. This means that house b is affordable, as *affordable* is true at w_4 (and thus also at w_5). Similarly, it follows that house b is close to work, a property that is true at w_5 (and thus at w_4). As there is no world that is part of object b as well as in the cluster representing the property *large*, house b is not large. The ranking of the properties is indicated by the $<$ symbol: property *affordable* is more important than *close to work* which in turn is more important than *large*. As a result, in any natural preference ordering based on this ranking one would expect house b to be preferred over house c .

The modality $\Box^{\#}$ can be used to express a property of a cluster. For example, $E\Box^{\#}\varphi$ expresses that there is a cluster where φ is true everywhere. $\Box^{\#}\varphi$ expresses that at least φ is true in the cluster. In Figure 2, for example, in the third cluster we have that $\Box^{\#}large$ is true. This means that every object that is represented by a world in this cluster is *large*. But we also want every object that is *large* to be represented in the cluster. To specify this, we use the modality \Box^{\neq} . We can now explain why simply defining the truth of $\Box^{\neq}\varphi$ in terms of truth of φ in all worlds that are not equally ranked to the current one does not work. The point is that there may be copies v of worlds w that have a different ranking than world w . As copies have the same truth assignment, at such copies a propositional formula φ would be assigned the same truth value. This is illustrated in Figure 3, where *large* is true in all worlds in the shaded area. The key observation here is that worlds of a particular ranking identify a set of objects, i.e. copies of these worlds which must be part of these objects (by Definition 3 of an object). This is why $\Box^{\neq}\varphi$ evaluates φ at all objects, or, more precisely, the worlds that define these objects, that are not identified by any of the worlds that have the same ranking as the current one.




Figure 3: Visualization of an MPL model. All worlds where *large* is true are in the shaded section.

By combining both operators we are able to characterize a cluster. For the third cluster in Figure 2, we have that $\Box^{\#}large \wedge \Box^{\neq}\neg large$ where *large* exactly characterizes the cluster. The characterization of a cluster by φ is abbreviated as $C\varphi$, and defined by:

$$C\varphi ::= \Box^{\#}\varphi \wedge \Box^{\neq}\neg\varphi$$

φ is true for all objects identified by (worlds in) the cluster and not true in all worlds that identify other objects. As an object may consist of several copies to represent that it has various properties represented by different clusters, copies of such worlds outside the cluster need to be excluded in the evaluation of $\neg\varphi$ which explains the truth condition for \Box^{\neq} .

Proposition 10 shows that properties and objects are related in such a way that object preferences can be derived. The first item of the proposition states that if there is an object that has property φ and the current world identifies a cluster characterized by φ , then within the cluster there is a world that is named i , i.e. identifies the object i . The second item states that the converse is true for an object that does not satisfy a property φ that characterizes a cluster. That is, if object i does not satisfy φ and the current world identifies a cluster characterized by φ , then no world that identifies the object labelled i is part of that cluster. The third item generalizes the first item. It states that if there is a cluster characterized by φ , and there is an object named i that satisfies φ , then there is an i -world in that cluster. The last item states that when a world satisfies $C(\varphi)$, then all worlds within the same cluster satisfy $C(\varphi)$.

PROPOSITION 10. We have:

1. $\models C(\varphi) \wedge E_i\varphi \rightarrow \Diamond^{\#}i$
2. $\models C(\varphi) \wedge \neg E_i\varphi \rightarrow \neg\Diamond^{\#}i$
3. $\models EC(\varphi) \wedge E_i\varphi \rightarrow E_iC(\varphi)$
4. $\models C(\varphi) \rightarrow \Box^{\#}C(\varphi)$

PROOF. We prove item 1. Suppose $\mathcal{M}, w \models C(\varphi) \wedge E_i\varphi$. This means that $\mathcal{M}, w \models \Box^{\#}\varphi$. By the truth definition for $\Box^{\#}$, this is equivalent to $\forall u \in \bigcup\{[v]_{\sim} \mid w \sim v\} : \mathcal{M}, u \models \varphi$. By the definition of $E_i\varphi$ we must also have a world u' such that $\mathcal{M}, u' \models i \wedge \varphi$. This means that we cannot have $u' \in \bigcup\{[v]_{\sim} \mid w \sim v\}$ and we have that

$u' \in \cup\{[v]_V \mid w \sim v\}$. It follows that $u' \in [v]_V$ for some $v \sim w$; as u' must be a copy of v this means that we have $\mathcal{M}, v \models i$ and, by the truth definition for $\diamond^=$, we have $\mathcal{M}, w \models \diamond^= i$. \square

The operator C provides exactly what we need to define property rankings. Semantically, we have already seen that the pre-order \lesssim induces a strict linear order on clusters. The formula $C\varphi$ allows us to express that a cluster is characterized by a formula φ . Using this operator and the binary preference operator $<_{\forall\forall}$ we can express that property ψ (represented by a cluster) is ranked higher than another property φ (represented by another cluster) by $C\varphi <_{\forall\forall} C\psi$. For example, in Figure 2, we have $C(\text{large}) <_{\forall\forall} C(\text{closeToWork}) <_{\forall\forall} C(\text{affordable})$. By combining this with specifications of particular preferences orderings and statements that an object has a particular property (cf. Proposition 10), this will allow the derivation of object preferences from a property ranking.

3. PREFERENCE ORDERINGS

In this Section, we show how to use multi-attribute preference logic to define multi-attribute preference orderings derived from property rankings. Coste-Marquis *et al.* [6] describe three frequent orderings based on prioritized goals: best-out, discrimin and leximin ordering. Brewka [4] defines a preference language in which different basic preference orderings can be combined and identifies four ‘fundamental strategies’ for deriving preferences from what he calls a ranked knowledge base: \top , κ , \subseteq and $\#$. As best-out is the same as κ , discrimin is \subseteq , and leximin is $\#$, we will base the remainder of our discussion on Brewka [4].

We first informally introduce these orderings and then present definitions for each of them in the logic. Section 4 presents the definitions of [4] and a proof that the definitions in multi-attribute preference logic match those provided in [4]. The advantage of defining preference orderings in a logic instead of providing set-theoretical definitions is that it formalizes the reasoning about object preferences. From a practical point of view, the logic allows us to provide rigorous formal proofs for object preferences derived from property rankings. From a theoretical point of view, it provides the tools to reason *about* preference orderings and allows, for example, to prove that whenever an object is preferred over another by the \top strategy it also is preferred by the $\#$ strategy (see Proposition 12 below).

The two orderings \subseteq and $\#$ first consider the most important property. If some object has that property and another does not, then the first is preferred over the second. So in the example, both $house_1$ and $house_2$ would be preferred over $house_3$. If two houses both have the property or if neither of them has it, the next property is considered. $house_1$ and $house_2$ are both affordable, but $house_1$ is close to work and $house_2$ is not, so $house_1$ would be preferred over $house_2$. Note that although $house_3$ satisfies two properties and house $house_2$ only satisfies one property, $house_2$ is still preferred over $house_3$ because the single property of $house_2$ is considered more important than both properties of $house_3$. The \subseteq and $\#$ orderings only differ if multiple properties are equally important. As we will make the assumption that no two properties can have the same importance, we will not discuss the difference and only refer to the $\#$ ordering in the following.

The \top ordering looks at the highest ranked or most important property that *is* satisfied. If that property of one object is ranked higher than that of another object, then the first object is preferred over the

second. If those properties are equally ranked, then both objects are equally preferred. In our running example, $house_1$ and $house_2$ are both preferred over $house_3$, since the property ranked highest that is satisfied by both $house_1$ and $house_2$ is *affordable*, and this property is ranked higher than the highest ranked property satisfied by $house_3$, i.e. *closeToWork*. Since the most important property satisfied by $house_1$ is the same as the most important property satisfied by $house_2$, $house_1$ and $house_2$ are equally preferred.

The κ ordering looks at the most important property that *is not* satisfied. If that property of one object is less important than the property of another object, then the first object is preferred over the second. If those properties are equally important, then both objects are equally preferred. In our running example, the highest ranked property that is not satisfied by $house_1$ is *large*, that of $house_2$ is *closeToWork* and that of $house_3$ is *affordable*. Since *large* is the least important property of these properties, $house_1$ is preferred over both other houses. As *closeToWork* is less important than *affordable*, $house_2$ is preferred over $house_3$.

All preference orderings introduced can be defined in multi-attribute preference logic. We use $pref^s(i, j)$ to stand for: object i is weakly preferred over object j according to strategy s , where s is one of \top , κ and $\#$; $pref^s(i, j)$ is used to express strict preference.

DEFINITION 11. (Preference Orderings)

$pref^{\kappa}(i, j)$, $pref^{\subseteq}(i, j)$, $pref^{\#}(i, j)$, $pref^{\sim}(i, j)$, $pref^{\top}(i, j)$ and $pref^{\top}(i, j)$ are defined by:

$$\begin{aligned} pref^{\top}(i, j) &::= E(i \wedge \neg \diamond^= j \wedge \square^< (\neg i \wedge \neg j)) \\ pref^{\sim}(i, j) &::= pref^{\top}(i, j) \vee \\ &\quad U((\diamond^= i \wedge \square^< \neg i) \leftrightarrow (\diamond^= j \wedge \square^< \neg j)) \\ pref^{\kappa}(i, j) &::= E(i \wedge \neg \diamond^= j \wedge \square^< (\diamond^= i \wedge \diamond^= j)) \\ pref^{\subseteq}(i, j) &::= pref^{\kappa}(i, j) \vee \\ &\quad U((\neg \diamond^= i \wedge \square^< \diamond^= i) \leftrightarrow (\neg \diamond^= j \wedge \square^< \diamond^= j)) \\ pref^{\#}(i, j) &::= E(i \wedge \neg \diamond^= j \wedge \square^< (\diamond^= i \leftrightarrow \diamond^= j)) \\ pref^{\sim}(i, j) &::= pref^{\#}(i, j) \vee U(\diamond^= i \leftrightarrow \diamond^= j) \end{aligned}$$

To understand these definitions, recall that we say that a world identifies an object when it is part of that object and the object consists of copies of one and the same world. These copies are used to represent that an object has a property present in a property ranking. In Figure 2, for example, world w_7 is a representative of object c for the property *large*. Thus, the formula $E_{i \neg \diamond^= j}$ may be read as ‘object i has a property that object j does not have’. Similarly, $\diamond^< i$ can be read as ‘there is a more important property (than the current one) that object i has’. These readings may help explain the definitions. $pref^{\top}(i, j)$ may be read as ‘there is a property such that i has it and j does not, and for all more important properties, neither i nor j has any of them’. The second disjunct in the definition of $pref^{\sim}(i, j)$ defines when two objects are equally preferred with respect to \top , and may be read as ‘if there is a property that i has, but i does not have any more important properties, then j has that property too and does not have any more important properties either, and vice versa’. Similar readings can be provided for the other preference operators.

Proposition 12 shows that the relation between weak and strict preference is as usual, and, moreover, a strict preference according to \top or κ implies a strict preference according to $\#$.

PROPOSITION 12. We have:

1. $\models \text{pref}^s(i, j) \leftrightarrow \text{pref}^s(i, j) \wedge \neg \text{pref}^s(j, i)$ for $s \in \{\top, \kappa, \#\}$.
2. $\models \text{pref}^\top(i, j) \rightarrow \text{pref}^\#(i, j)$
3. $\models \text{pref}^\kappa(i, j) \rightarrow \text{pref}^\#(i, j)$

EXAMPLE 13. Given the model of Figure 2, we can derive that $\text{pref}^\#(b, d)$. By definition, this is the case when $E(b \wedge \neg \diamond^\# d \wedge \square^\# (\diamond^\# b \leftrightarrow \diamond^\# d))$ is true. This means that there must be a world w that is named b that has no equally ranked world named d , and, moreover, for every higher ranked world v there is an equally ranked world named b if and only if there is an equally ranked world with name d . By inspection of Figure 2, world w_5 fits the description.

4. MPL DEFINES RANKED KNOWLEDGE BASES

Here we prove that the preference orderings of Definition 11 define those of Brewka [4]. Brewka [4] calls property rankings *ranked knowledge bases*, defined as follows:

DEFINITION 14. (**Ranked Knowledge Base**)

A *ranked knowledge base* (RKB) is a set $F \subseteq \mathcal{L}_0$ of objective formulae together with a total pre-order \geq on F . Ranked knowledge bases are represented as a set of ranked formulae (f, k) , where f is an objective formula and k , the rank of f , is a non-negative integer such that $f_1 \geq f_2$ iff $\text{rank}(f_1) \geq \text{rank}(f_2)$. That is, higher rank is expressed by higher indices.

In the setting of [4], comparing objects given a ranked knowledge base means comparing *truth assignments* which represent these objects, analogously to the representation of the three houses used in Figure 1. It is easy to see that this example is represented by the following ranked knowledge base: $\{(affordable, 3), (closeToWork, 2), (large, 1)\}$.

Object preferences can be derived in multiple ways from a ranked knowledge base. In order to define these strategies, some auxiliary definitions are introduced next. Below, $K^n(m)$ denotes the set of properties of a certain rank n that are satisfied with respect to truth assignment m ; $\text{maxsat}^K(m)$ denotes the highest rank associated with the properties that are satisfied by assignment m , and $\text{maxunsat}^K(m)$ denotes the highest rank associated with the properties that are not satisfied by m .

DEFINITION 15.

Let K be a ranked knowledge base and $m \in 2^{A^r}$.

$$\begin{aligned} K^n(m) &::= \{f \mid (f, n) \in K, m \models f\} \\ \text{maxsat}^K(m) &::= -\infty \text{ if } m \not\models f_i \text{ for all } (f_i, v_i) \in K, \\ &\quad \max\{i \mid (f, i) \in K, m \models f\} \text{ otherwise} \\ \text{maxunsat}^K(m) &::= -\infty \text{ if } m \models f_i \text{ for all } (f_i, v_i) \in K, \\ &\quad \max\{i \mid (f, i) \in K, m \not\models f\} \text{ otherwise} \end{aligned}$$

Using these auxiliary definitions, preference orderings $m_1 \geq_s^K m_2$ are defined which mean that object (truth assignment) m_1 is (weakly) preferred over object m_2 according to strategy s .

DEFINITION 16. (**Preference Orderings**)

Let K be a ranked knowledge base. Then the following preference

orderings over truth assignments are defined:

- $m_1 \geq_\top^K m_2$ iff $\text{maxsat}^K(m_1) \geq \text{maxsat}^K(m_2)$.
- $m_1 \geq_\kappa^K m_2$ iff $\text{maxunsat}^K(m_1) \leq \text{maxunsat}^K(m_2)$.
- $m_1 \geq_\#^K m_2$ iff $|K^n(m_1)| = |K^n(m_2)|$ for all n , or there is n s.t. $|K^n(m_1)| > |K^n(m_2)|$, and for all $j > n$: $|K^j(m_1)| = |K^j(m_2)|$.

To simplify, we make the assumption here that different properties cannot have the same ranking. In that case, the set of all satisfied properties of a given rank is a singleton set or the empty set, we have that \geq is a strict linear order on F - also denoted by $>$, and, as a result, the \subseteq and $\#$ orderings coincide. We also assume that properties in a ranked knowledge base are consistent. Finally, we may assume that a ranked knowledge base does not contain logically equivalent properties with different ranks since such occurrences except for the one ranked highest can be discarded as it has no influence on any of the preference orderings.

DEFINITION 17. (**Translation Function**)

The function τ translates ranked knowledge bases $K = \langle F, \geq \rangle$ and truth assignments m to formulae and is defined by:

- $\tau(K) ::= \bigwedge \{EC(\varphi) \mid \varphi \in F\} \wedge$
 $U(\bigvee \{C(\varphi) \mid \varphi \in F \text{ or } \varphi = \neg \bigvee \{\chi \mid \chi \in F\}\})$
 $\bigwedge \{C(\varphi) <_{\forall \vee} C(\psi) \mid \varphi, \psi \in F \ \& \ \psi > \varphi\} \wedge$
 $\bigwedge \{C(\neg \bigvee \{\varphi \mid \varphi \in F\}) <_{\forall \vee} \psi \mid \psi \in F\} \wedge$
- $\tau_{name}(m) \in Nom$
- $\tau(m) ::= \bigwedge \{E_i \varphi \mid m \models \varphi\} \cup \{\neg E_i \varphi \mid m \not\models \varphi\}$ with $i = \tau_{name}(m)$

The translation of a ranked knowledge base K expresses that for each property φ in K , there exists a corresponding cluster by $C\varphi$, that there are no other clusters than those specified by the properties, and one extra cluster for the case in which none of the properties is satisfied. It forces the ranking of these clusters to be the same as the property ranking induced by K , with the added extra cluster as least important one. The translation also associates an object name with a truth assignment and states for each property whether the object (truth assignment) has the property or not.

EXAMPLE 18. Using the translation function, and assuming that

$\tau_{name}(house_1) = b$, $\tau_{name}(house_2) = d$ and $\tau_{name}(house_3) = e$, the RKB $\{(affordable, 3), (closeToWork, 2), (large, 1)\}$ translates into:

1. $E(C(affordable)) \wedge E(C(closeToWork)) \wedge E(C(large))$
2. $U(C(affordable) \vee C(closeToWork) \vee C(large) \vee C(\neg(affordable \vee closeToWork \vee large)))$
3. $C(\neg(affordable \vee closeToWork \vee large)) <_{\forall \vee} C(large) <_{\forall \vee} C(closeToWork) <_{\forall \vee} C(affordable)$
4. $E_b(affordable) \wedge E_b(closeToWork) \wedge \neg E_b(large)$
5. $E_d(affordable) \wedge \neg E_d(closeToWork) \wedge \neg E_d(large)$
6. $\neg E_e(affordable) \wedge E_e(closeToWork) \wedge E_e(large)$

A model of this theory is shown in Figure 2. Although only objects b , d and e are specified in the theory, for illustrative reasons this model contains all possible objects (there is a world, and hence an object, for every possible valuation of the three propositional atoms). Every property has its own cluster, which means that every object satisfying that property has a world in that cluster, and that every world in that cluster satisfies that property. No worlds exist outside the four specified clusters, and the order among clusters is fixed. The only ways a model of this theory can be structurally different from the one shown are by removing objects that are not b , d or e (but then all worlds belonging to that object have to be

removed at once), or by adding more worlds, but only at the same 'places' as the worlds shown.

Theorem 19 shows that every multi-attribute preference model that is a model of the translation of a particular RKB yields the same preference ordering as the original RKB.

THEOREM 19. $m_1 \succeq_s^K m_2$ iff
 $\models \tau(K) \wedge \tau(m_1) \wedge \tau(m_2) \rightarrow pref^s(\tau_{name}(m_1), \tau_{name}(m_2))$
 where $s \in \{\top, \kappa, \#\}$.

PROOF. Assume that $\tau_{name}(m_1) = i$ and $\tau_{name}(m_2) = j$, and observe that the translation of $K = \langle F, \succeq \rangle$ is equivalent to:

- (1) $C(\neg(f_1 \vee \dots \vee f_n)) <_{\forall\forall} C(f_1) <_{\forall\forall} \dots <_{\forall\forall} C(f_n)$,
- (2) $\forall f \in F : E(C(f))$ and
- (3) $U(C(f_1) \vee \dots \vee C(f_n) \vee C(\neg(f_1 \vee \dots \vee f_n)))$.

For brevity, we only prove the left to right direction for the case $m_1 \succ_{\kappa}^K m_2$. Then we have $maxunsat^K(m_1) < maxunsat^K(m_2)$ and $maxunsat^K(m_2) > -\infty$, so there is a formula f_k in F such that

- (4) $m_2 \not\models f_k$,
- (5) $m_1 \models f_k$ and
- (6) $\forall f' > f_k : m_1 \models f' \ \& \ m_2 \models f'$.

Applying the translation function τ , we then get:

- (4) $\neg E_j f_k$,
- (5) $E_i f_k$ and
- (6) $\forall f' > f_k : E_i f' \wedge E_j f'$.

From (5), (2) and Prop. 10.3 it then follows that

- (8) $E_i C(f_k)$.

From (8), (4) and Prop. 10.2 it follows that

- (9) $E_i \neg \diamond^= j \wedge C(f_k)$.

And from (6) and Prop. 10.1 it follows that

- (10) $\forall f' > f_k : \diamond^= i \wedge \diamond^= j$.

Using (1) and (3) we obtain

- (11) $C(f_k) \rightarrow \square^< (C(f_{k+1}) \vee \dots \vee C(f_n))$.

From (10) and (11) we obtain

- (12) $C(f_k) \rightarrow \square^< \diamond^= i \wedge \square^< \diamond^= j$.

Then (9) and (12) can be combined into $E(i \wedge \neg \diamond^= j \wedge \square^< (\diamond^= i \wedge \diamond^= j))$, which is the definition of $pref^K(i, j)$. \square

EXAMPLE 20. We now show how to formally derive a preference statement from the formulae obtained by translating a ranked knowledge base in Example 18. As an illustration, we show that $pref^K(b, d)$ can be derived.

From (18.4) $E_b(closeToWork)$, (18.1) $E(C(closeToWork))$ and Proposition 10.3 we obtain

- (1) $E_b C(closeToWork)$.

From (18.5) $\neg E_d(closeToWork)$ and Proposition 10.2 it follows that

- (2a) $C(closeToWork) \rightarrow \neg \diamond^= d$.

From 18.3 and 18.2 we can derive that

- (2b) $C(closeToWork) \rightarrow \square^< C(affordable)$.

By combining (1), (2a) and (2b) we derive

- (3) $E_b(\neg \diamond^= d \wedge \square^< C(affordable))$.

Now, from Proposition 10.1, (18.4) $E_b(affordable)$ and (18.5) $E_d(affordable)$, we derive

- (4a) $C(affordable) \rightarrow \diamond^= b$ and
- (4b) $C(affordable) \rightarrow \diamond^= b$.

Using (3), (4a), and (4b), we obtain $E_b(\neg \diamond^= d \wedge \square^< (\diamond^= b \wedge \diamond^= d))$, which is the definition of $pref^K(b, d)$.

5. CONCLUSION

In this paper we introduced a modal logic for qualitative multi-attribute preferences. The logic is based on Girard's binary preference logic [7], but extends this logic with objects and clusters that introduce the possibility to reason explicitly about multiple attributes. We showed that multi-attribute preference logic is expressive enough to define various natural preference orderings based on property rankings [4, 6]. The additional value of the logic is that it is possible to reason about these different preference orderings within the logic. This means we cannot only reason about which objects are preferred according to a certain ordering, but also about the relation between different orderings as is shown in Proposition 12.

One possible extension to multi-attribute preference logic is the introduction of indices for different agents. In this way, distinct preference orderings for several agents can be expressed. This introduces the possibility to reason about properties such as pareto-optimality of objects (an object is pareto-optimal if there is no other object that is better for at least one agent and not worse for the other agents), which is useful in the context of e.g. joint decision making or negotiation.

We have made the assumptions that attributes are binary, and that priority orderings are total linear orders. In future work we plan to investigate how we can loosen these assumptions. For example, if multiple attributes can have the same importance, the $\#$ and \subseteq orderings will differ and we will be able to encode trade-offs between attributes.

Our main concern in this paper has been the expressiveness of multi-attribute preference logic. Other questions such as a complete axiomatization of the logic, succinctness and complexity remain future work. We plan to develop a reasoning system in which agents can reason about qualitative multi-attribute preferences in various settings. In our future work we will focus more on the reasoning mechanism and how different domains can be modelled accurately in our approach.

A more detailed comparison of multi-attribute preference logic with other preference logics such as Qualitative Choice Logic [5] is planned. Other areas for future work concern the representation of dependent properties and the relation of multi-attribute preference logic to e.g. CP-nets [3].

6. ACKNOWLEDGEMENTS

This research is supported by the Dutch Technology Foundation STW, applied science division of NWO and the Technology Program of the Ministry of Economic Affairs. It is part of the Pocket Negotiator project with grant number VICI-project 08075.

7. REFERENCES

- [1] P. Blackburn and J. Seligman. Hybrid languages. *Journal of Logic, Language and Information*, 4(3):251–272, 1995.
- [2] C. Boutilier. Toward a logic for qualitative decision theory. In *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR 1994)*, pages 75–86, 1994.
- [3] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

- [4] G. Brewka. A rank based description language for qualitative preferences. In *Proceedings of the Sixteenth European Conference on Artificial Intelligence (ECAI 2004)*, pages 303–307, 2004.
- [5] G. Brewka, S. Benferhat, and D. Le Berre. Qualitative choice logic. *Artificial Intelligence*, 157(1-2):203–237, 2004.
- [6] S. Coste-Marquis, J. Lang, P. Liberatore, and P. Marquis. Expressive power and succinctness of propositional languages for preference representation. In *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR 2004)*, pages 203–212, 2004.
- [7] P. Girard. *Modal Logic for Belief and Preference Change*. PhD thesis, Universiteit van Amsterdam, 2008.
- [8] J. Y. Halpern and G. Lakemeyer. Multi-agent only knowing. *Journal of Logic and Computation*, 11(1):41–70, 2001.
- [9] S. O. Hansson. Preference logic. In D. Gabbay and F. Günthner, editors, *Handbook of Philosophical Logic*, volume 4, pages 319–393. Kluwer Academic Publishers, 2nd edition, 2001.
- [10] R. L. Keeney and H. Raiffa. *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge University Press, 1993.
- [11] F. Liu. *Changing for the Better: Preference Dynamics and Agent Diversity*. PhD thesis, Universiteit van Amsterdam, 2008.
- [12] G. H. von Wright. *The Logic of Preference: An Essay*. Edinburgh University Press, 1963.