

Representing and Reasoning with Multi-Stakeholder Qualitative Preference Queries

Samik Basu^{a,*}, Vasant Honavar^b, Ganesh Ram Santhanam^c and Jia Tao^d

^aDepartment of Computer Science, Iowa State University, Ames, IA, USA

^bCollege of Information Sciences and Technology, Pennsylvania State University, University Park, PA, USA

^cDepartment of Electrical and Computer Engineering, Iowa State University, Ames, IA, USA

^dDepartment of Computer Science, Lafayette College, Easton, PA, USA

ORCID ID: Vasant Honavar <https://orcid.org/0000-0001-5399-3489>, Jia Tao <https://orcid.org/0000-0002-2342-3271>

Abstract. Many decision-making scenarios, e.g., public policy, healthcare, business, and disaster response, require accommodating the preferences of multiple stakeholders. We offer the first formal treatment of reasoning with multi-stakeholder qualitative preferences in a setting where stakeholders express their preferences in a qualitative preference language, e.g., CP-net, CI-net, TCP-net, CP-Theory. We introduce a query language for expressing queries against such preferences over sets of outcomes that satisfy specified criteria, e.g., $\psi_1 \text{ } \mathcal{P}_A \text{ } \psi_2$ (read loosely as the set of outcomes satisfying ψ_1 that are preferred over outcomes satisfying ψ_2 by a set of stakeholders A). Motivated by practical application scenarios, we introduce and analyze several alternative semantics for such queries, and examine their interrelationships. We provide a provably correct algorithm for answering multi-stakeholder qualitative preference queries using model checking in alternation-free μ -calculus. We present experimental results that demonstrate the feasibility of our approach.

1 Introduction

The ability to express and reason about preferences over a set of alternatives is central to rational decision-making in a broad range of applications, including software design [40, 26, 35, 36, 18, 1], public policy, e.g., city planning [23, 37], healthcare [8], security [3, 21], privacy [28], among others. In general, the preferences can be quantitative [24, 19] or qualitative [7, 15]. But stakeholders often find it natural to express their preferences in qualitative terms [34], e.g., that a cheaper car is preferred to a more expensive car. Hence, there has been a growing interest in languages and tools for representing and reasoning with qualitative preferences [14, 34, 11]. For example, [32] leverage advances in model checking [10, 29, 9] to provide efficient and hence practically useful tools for reasoning with the qualitative preferences of single stakeholders [13, 33].

However, decision-making in real-world settings often needs to accommodate the preferences of *multiple* stakeholders. Consider, for example, the task of choosing a care plan for a critically ill patient. The stakeholders, in this case, may include the patient concerned with their health outcome and the cost of care, the physician committed to ensuring that the patient receives the best care available, the family members with an interest in the patient's well-being, the

hospital system seeking to maximize its profits, and the insurance provider seeking to minimize the reimbursements. A key challenge in extending the preference representation languages and reasoning tools from the single stakeholder setting to the multi-stakeholder setting has to do with maintaining, and reasoning with the (possibly conflicting) preferences of stakeholders. Furthermore, the preferences of some stakeholders in some settings may override those of others, e.g., due to their relative roles in an organization, or due to differences in their expertise as it relates to specific aspects of the application domain, etc. Ensuring transparency and accountability of decision-making requires that the system be able to explain *how* the stakeholders' preferences impact the outcomes.

Contributions. The key contributions of the paper are as follows: (i) We provide the first formal treatment of reasoning with multi-stakeholder qualitative preferences. We consider the setting where the stakeholders express their preferences in a qualitative preference language, e.g., CP-net, CI-net, TCP-net, CP-Theory. (ii) We introduce a query language for expressing queries with respect to the preferences of multiple stakeholders over outcomes that satisfy a set of specified criteria. (iii) We generalize the *induced preference graphs* that encode the qualitative preferences of a single stakeholder to *multi-stakeholder induced preference graphs* that encode the preferences of multiple stakeholders. (iv) We introduce and analyze several alternative semantics for such queries, motivated by the needs of different application scenarios, and examine their inter-relationships. (v) We provide a provably correct algorithm for answering multi-stakeholder preference queries using model checking in μ -calculus; and (vi) We present results of experiments that demonstrate the feasibility of our approach.

2 Qualitative Preference Languages

We consider settings in which stakeholders express preferences over a set of alternatives or outcomes, where each alternative is described by a set of attributes or (preference) variables. Stakeholders may directly express their preference between a pair of alternatives, by asserting that one valuation of the variables is preferred to another. In addition, preferences over sets of alternatives may be succinctly stated over (a) the possible valuations of each variable, i.e., *intra-variable preference*; or (b) the variables themselves indicating their *relative importance*. Several qualitative preference languages with

* Corresponding Author. Email: sbasu@iastate.edu.

Formally, let $X = \{X_i \mid 0 < i \leq n\}$ be a set of preference variables, D_i be the domain of X_i , and v_i be the assignment of X_i to a particular valuation in D_i . Let $O = \prod_{X_i \in X} D_i$ be the set of alternatives or outcomes, and $O^P = \prod_{X_i \in Y \subseteq X} D_i$ be the set of partial alternatives or outcomes. Each outcome $o \in O$ is represented as a tuple of valuations of each variable, i.e., $o = \langle v_1, v_2, \dots, v_n \rangle$. We use the following notation to represent a preference statement

where $c \in O^P$ is the condition under which this preference over X_i 's valuation holds, and $Y \subseteq X \setminus X_i$ is the set of variables less important than X_i . For brevity, we drop $[c]$ when $c = \text{true}$ and $[Y]$ when $Y = \emptyset$. A preference statement P specifies that when c holds, the valuation v_i is preferred to v_i' for variable X_i , regardless of the valuations and intra-variable preferences of the variables in Y .

$$\begin{array}{ll}
P_1 & E = \text{Code} \succ_E E = \text{No} - \text{Code} \\
P_2 & [E = \text{Code}] F = \text{No} - \text{Fix} \succ_F F = \text{Fix} \\
P_3 & [F = \text{Fix}] E = \text{Code} \succ_E E = \text{No} - \text{Code} \\
P_4 & [F = \text{Fix}] A = \text{Simple} \succ_A A = \text{Complex} \\
P_5 & \langle E = \text{No} - \text{Code}, A = \text{Simple}, F = \text{No} - \text{Fix} \rangle \succ \\
& \quad \langle E = \text{Code}, A = \text{Complex}, F = \text{No} - \text{Fix} \rangle \\
P_6 & A = \text{Simple} \succ_A A = \text{Complex} \\
P_7 & E = \text{Code} \succ_E E = \text{No} - \text{Code}[A, F]
\end{array}$$

Semantics of Preferences. The semantics of CP-nets, TCP-nets, and CP-theories is based on and extends the *ceteris-paribus* principle [22]. The preference statements induce a strict partial order over the alternatives. For instance, for $o, o' \in O$, a preference statement $P: [c] (X_i = v_i) \succ (X_i = v_i')$ induces a preference from o' to o (denoted $o' \prec o$) if both satisfy c ; their valuations for X_i are v_i and v_i' respectively; and their valuations for all other variables are identical.

A multi-stakeholder induced preference graph is an induced preference graph where each edge in the graph is annotated by the set of stakeholders whose preferences induce that edge. That is, $I = (O \cup \{\perp\}, E, L, \mathcal{A})$ where the edge relation $E \subseteq (O \cup \{\perp\}) \times \mathcal{P}(\mathcal{A}) \times (O \cup \{\perp\})$. An edge $e = (o_1, A, o_2) \in E$ captures the fact that $o_1 \prec o_2$ for every agent in A . We note that $e = (\perp, \mathcal{A}, o)$ for every $o \in O$.

We will denote the edges in I as $o_1 \xrightarrow{A} o_2$, where A is the set of agents whose preferences have induced the edge from o_1 to o_2 .

Definition 3 ($\prec_A$ and $\prec_A^+$) We write $o \prec_A o'$ if there exists an edge $o \xrightarrow{A'} o'$ and $A \cap A' \neq \emptyset$. Similarly, $o \prec_A^+ o'$ if there exists a path $o = o_1 \xrightarrow{A_1} o_2 \xrightarrow{A_2} \dots \xrightarrow{A_k} o_{k+1} = o'$ where $\forall i \in [1, k]. (A \cap A_i) \neq \emptyset$.

When A is singleton ($A = \{a\}$), we will write $o \prec_a o'$.

3 Single Stakeholder Preference Queries

We first introduce a language for expressing queries with respect to single stakeholder qualitative preferences before proceeding to consider multi-stakeholder preferences queries. A key feature of this language is that it allows expressing queries against preferences over properties of outcomes, rather than the outcomes themselves. Thus, it can readily accommodate preferences expressed in existing qualitative preference languages such as CP-nets [4], TCP-nets [6], and CP-theories [41]. This allows us, for example, to query for outcomes with properties that are more preferred to all other outcomes. The resulting single stakeholder preference query language can express a range of preference queries (e.g., find the set of non-dominated outcomes) of common interest.

Syntax. The syntax of the query language is described over atomic propositions, propositional constants, boolean connectives and a (new) operator P : preference operator over properties. The language Ψ is defined by the grammar:

$$\psi \rightarrow \text{tt} \mid \text{ff} \mid \text{AP} \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \psi P_a \psi.$$

The answer to a query corresponds to the set of outcomes that belongs to the semantics of the query. For instance, all outcomes are returned for a query tt , while no outcome is returned for the query ff . A query involving an atomic proposition simply returns the outcomes that satisfy the proposition. Answers to queries involving Boolean connectives conform to the natural meaning of the connectives. The query $\psi_1 P_a \psi_2$ returns the outcomes satisfying ψ_1 that are more preferred than outcomes satisfying ψ_2 based on the preferences of the stakeholder a .

Semantics. The semantics of the query language is defined over the set of outcomes (states) in the preference graph I induced by the given preferences. Let $\mathcal{I}$ be the set of all preference graphs that can be induced by single stakeholder preferences with respect to which single stakeholder queries can be expressed given the syntax described above. We use the (semantic) function $\llbracket \cdot \rrbracket : \Psi \times \mathcal{I} \rightarrow \mathcal{P}(O)$, to define the semantics of $\psi \in \Psi$ in the context of an induced preference graph $I \in \mathcal{I}$. That is, $\llbracket \psi \rrbracket_I$ denotes the set of outcomes in I that satisfy the query expressed using the formula ψ . We will omit I from the definition unless it is explicitly necessary to distinguish between semantics in the context of two different induced preference graphs.

$$\begin{aligned} \llbracket \text{tt} \rrbracket &= O \\ \llbracket \text{ff} \rrbracket &= \emptyset \\ \llbracket p \rrbracket &= \{o \mid p \in L(o)\} \\ \llbracket \neg\psi \rrbracket &= O - \llbracket \psi \rrbracket \\ \llbracket \psi_1 \wedge \psi_2 \rrbracket &= \llbracket \psi_1 \rrbracket \cap \llbracket \psi_2 \rrbracket \\ \llbracket \psi_1 P_a \psi_2 \rrbracket &= \llbracket \psi_1 \rrbracket \cap \{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \succ_a^+ o'\} \\ &\quad \cap \{o \mid \forall o'. o' \in \llbracket \psi_2 \rrbracket \Rightarrow o \not\succ_a^+ o'\} \end{aligned}$$

Propositional constants tt and ff are satisfied by all and no outcomes, respectively. The proposition p is satisfied by any outcome that satisfies p . The formulas over Boolean connectives (negation,

conjunction, disjunction) conform to the standard set-based semantics (complement, intersection, union). The formula $\psi_1 P_a \psi_2$ is satisfied by outcomes that (i) satisfy ψ_1 , (ii) are preferred to at least one outcome that satisfies ψ_2 , and (iii) are not less preferred to any outcome that satisfies ψ_2 by the stakeholder a . In short, $\psi_1 P_a \psi_2$ is the set of outcomes satisfying ψ_1 that are more preferred to outcomes satisfying ψ_2 by the stakeholder a .

The resulting query language can be used to express queries such as:

- What is the set of outcomes that are preferred by the stakeholder a to outcomes that satisfy ψ ? The is expressed as $\text{tt} P_a \psi$.
- What is the non-dominated set of outcomes relative to stakeholder a 's preferences? The query can be expressed as $\text{tt} P_a \text{tt}$. What is the non-dominated set of outcomes for stakeholder a that satisfies ψ ? This can be expressed as $\psi P_a \text{tt}$.
- With respect to stakeholder a 's preferences, what are the best improvements to outcomes satisfying ψ ? The query can be expressed as $(\text{tt} P_a \text{tt}) \wedge (\text{tt} P_a \psi)$.

Example 3 If $\psi = \text{Code}$, then the semantics of $\text{tt} P_1 \psi$ (for stakeholder 1) is the set of outcomes $\{o_1, o_5\}$. This is because, while both o_1 and o_5 dominate some outcome satisfying Code with respect to stakeholder 1's preferences, they are not dominated by any outcome that satisfies Code . On the other hand, the query $\text{tt} P_2 \psi$ (for stakeholder 2) yields the set $\{o_2, o_6\}$.

Example 4 For stakeholder 1, the non-dominated set of outcomes is $\{o_1, o_5\}$ (result of the query: $\text{tt} P_1 \text{tt}$), while for stakeholder 2, the non-dominated set is $\{o_1, o_2, o_6, o_8\}$. Note that the outcome o_8 neither dominates nor is dominated by any outcome. However, it dominates $\perp$ and hence is included as part of the non-dominated set.

Cycles in Induced Preference Graphs. Cycles in an induced preference graph are indicative of inconsistencies in the underlying preferences, the result being some outcome o both more and less preferred to an outcome o' . Does this pose any inconsistencies in the semantic interpretation of $\psi_1 P_a \psi_2$, when o satisfies ψ_1 and o' satisfies ψ_2 ? The answer is no. This is because semantics of $\psi_1 P_a \psi_2$ excludes all outcomes that are less preferred to outcomes satisfying ψ_2 . Hence, the outcome o will not be included in the set of outcomes returned by $\psi_1 P_a \psi_2$ as it is less preferred to o' .

4 Multi-Stakeholder Preference Queries

We proceed to extend the preceding language for expressing preference queries to allow preference queries with respect to the preferences of a set of stakeholders, as opposed to just a single stakeholder. Specifically, we add a new query construct $\psi_1 P_A \psi_2$ where $A \subseteq \mathcal{A}$, where $\mathcal{A}$ is the set of all stakeholders. When A is a singleton a , we use $\psi_1 P_a \psi_2$ to denote the query about the preferences of a single stakeholder a (as described in Section 3). In what follows, we describe the semantics of multi-stakeholder preference queries under several alternative interpretations of multi-stakeholder preferences.

Consensus Semantics. Consensus semantics, as the name suggests, is defined as the set of outcomes, whose preference over another set of outcomes, is decided by agreement among the set of stakeholders in question. Formally,

$$\llbracket \psi_1 P_A \psi_2 \rrbracket^{cs} = \bigcap_{a \in A} \llbracket \psi_1 P_a \psi_2 \rrbracket$$

Example 5 In Example 3, as per the consensus semantics the result of the query $(\text{tt } P_{\{1,2\}} \text{ Code})$ is the empty set as the stakeholders 1 and 2 do not agree on the outcomes that are more desirable than outcomes satisfying *Code*. On the other hand, stakeholders 1 and 2 agree on the non-dominated set $\{o_1\}$ computed as the semantics of $\text{tt } P_{\{1,2\}} \text{ tt}$ (see Example 4).

Collaborative Semantics. Unlike consensus semantics, which requires a complete agreement among the stakeholders, a collaborative semantics allows the stakeholders to arrive at a compromise that is not disagreeable to any stakeholder. There are several ways to realize such a compromise that correspond to different interpretations of the semantics of $\psi_1 \text{ P}_A \psi_2$. Recall that $\psi_1 \text{ P}_A \psi_2$ must return the set of outcomes that (i) satisfy ψ_1 , (ii) are preferred to at least one outcome that satisfies ψ_2 , and (iii) are not less preferred to any outcome that satisfies ψ_2 . We will refer to the last two conditions (ii and iii) as follows:

1. **Witness Condition (W)** for determining the set of outcomes that are preferred to at least one outcome satisfying ψ_2 .
2. **Agreement Condition (A)** for determining the set of outcomes that are not less preferred to any outcome satisfying ψ_2 .

Each of these conditions can be collaboratively decided in two ways:

1. **Collective Collaboration.** The set of outcomes that are preferred to at least one outcome satisfying ψ is chosen to be the **union** of outcomes preferred by each of the stakeholders to *at least one* outcome satisfying ψ .
2. **Constructive Collaboration.** An outcome o' is considered to be preferred to outcome o when there exists a path in the induced preference graph from o to o' where each edge along the path may be induced by the preferences of one or more stakeholders. Thus, there is no requirement that all of the edges along the path be induced by the preferences of the same stakeholder. Hence, the stakeholders collaboratively construct the path from o to o' by contributing one or more edges to the path based on their individual preferences. This can be viewed as **chaining** induced preference edges of different stakeholders to arrive at the result.

Constructive Collaboration is useful in situations where each stakeholder may not have complete information or expertise to determine a dominance relation between a pair of outcomes but they may be able to collaborate to arrive at a conclusion. For instance, healthcare providers (doctors, nurses) and hospital administrators may collaborate to develop an optimal placement strategy for hand sanitizers in the hospital. The healthcare providers present their preferences based on their knowledge of the usage of hand sanitizers at different times and locations, whereas the hospital administrators present their preferences based on the cost of procuring hand sanitizers.

Now we have two different choices for the witness (W) condition and agreement (A) condition:

- W₁. Collective collaboration for deciding witness condition for ψ_2 in $\psi_1 \text{ P}_A \psi_2$:

$$\bigcup_{a \in A} \{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \succ_a^+ o'\}$$

- W₂. Constructive collaboration for deciding witness condition for ψ_2 in $\psi_1 \text{ P}_A \psi_2$:

$$\{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \succ_A^+ o'\}$$

- A₁. Collective collaboration for deciding agreement condition for ψ_2 in $\psi_1 \text{ P}_A \psi_2$:

$$\begin{aligned} & O \setminus \bigcup_{a \in A} \{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \prec_a^+ o'\} \\ &= \bigcap_{a \in A} \{o \mid \forall o'. o' \in \llbracket \psi_2 \rrbracket \Rightarrow o \not\prec_a^+ o'\} \end{aligned}$$

- A₂. Constructive collaboration for deciding agreement condition for ψ_2 in $\psi_1 \text{ P}_A \psi_2$

$$\begin{aligned} & O \setminus \{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \prec_A^+ o'\} \\ &= \{o \mid \forall o'. o' \in \llbracket \psi_2 \rrbracket \Rightarrow o \not\prec_A^+ o'\} \end{aligned}$$

Example 6 Consider the induced preference graph in Figure 2. For stakeholder 1, the set of outcomes that dominate the outcomes satisfying *No-Code* is $\{o_1, o_2, o_4, o_5\}$. This is because $o_4 \succ_1 o_7$, $o_5 \succ_1 o_4$, $o_2 \succ_1 o_3$ and $o_1 \succ_1 o_2$. On the other hand, for stakeholder 2, the set of outcomes that dominate the outcomes satisfying *No-Code* is $\{o_2, o_3, o_4\}$.

Therefore, for $\psi_2 = \text{No-Code}$, we have:

$$W_1 : \bigcup_{a \in \{1,2\}} \{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \succ_a^+ o'\} = \{o_1, o_2, o_3, o_4, o_5\}$$

On the other hand,

$$W_2 : \{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \succ_{\{1,2\}}^+ o'\} = \{o_1, o_2, o_3, o_4, o_5, o_6\}.$$

Note that the set includes all outcomes whose inclusion is decided by stakeholders 1 and 2 on their own. Additionally, outcome o_6 is included because $o_4 \succ_1 o_7$, $o_5 \succ_1 o_4$ and $o_6 \succ_2 o_5$.

Example 7 For the induced preference graph in Figure 2, consider evaluating the agreement condition. The set of outcomes that are dominated by outcomes satisfying *No-Code* as per the stakeholder 1 is $\emptyset$. On the other hand, for stakeholder 2, the set is $\{o_5, o_7\}$ because $o_6 \succ_2 o_5$ and $o_3 \succ_2 o_7$.

Therefore, for $\psi_2 = \text{No-Code}$,

$$A_1 : O \setminus \bigcup_{a \in \{1,2\}} \{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \prec_a^+ o'\} = O \setminus \{o_5, o_7\}.$$

On the other hand,

$$A_2 : O \setminus \{o \mid \exists o'. o' \in \llbracket \psi_2 \rrbracket \wedge o \prec_{\{1,2\}}^+ o'\} = O \setminus \{o_4, o_5, o_7\}.$$

The membership of o_4 is decided from the relations: $o_6 \succ_2 o_5$ and $o_5 \succ_1 o_4$.

The combinations of W_1 and W_2 with A_1 and A_2 yield four different semantics for $\psi_1 \text{ P}_A \psi_2$. We will denote them by $\llbracket \psi_1 \text{ P}_A \psi_2 \rrbracket^{W_i A_j}$ where $i, j \in \{1, 2\}$.

Example 8 Using the Examples 6 and 7, we have the following when $\psi_1 = \text{tt}$ and $\psi_2 = \text{No-Code}$:

$$\begin{aligned} \llbracket \psi_1 \text{ P}_{\{1,2\}} \psi_2 \rrbracket^{W_1 A_2} &= \{o_1, o_2, o_3\}, \\ \llbracket \psi_1 \text{ P}_{\{1,2\}} \psi_2 \rrbracket^{W_1 A_1} &= \{o_1, o_2, o_3, o_4\}, \\ \llbracket \psi_1 \text{ P}_{\{1,2\}} \psi_2 \rrbracket^{W_2 A_2} &= \{o_1, o_2, o_3, o_6\}, \\ \llbracket \psi_1 \text{ P}_{\{1,2\}} \psi_2 \rrbracket^{W_2 A_1} &= \{o_1, o_2, o_3, o_4, o_6\}. \end{aligned}$$

Relationships Between Alternative Collaborative Semantics. The following Theorem shows the relationship between the two witness conditions and the relationship between the two agreement conditions.

Theorem 1 $W_1 \subseteq W_2$ and $A_2 \subseteq A_1$.

Proof. (i) $W_1 \subseteq W_2$.

Consider any $o_1 \in W_1$. Then, $o_1 \in \bigcup_{a \in A} \{o \mid \exists o'. o' \in \langle \psi_2 \rangle \wedge o \succ_a^+ o'\}$ by the definition of W_1 . Thus, there is an agent $a_1 \in A$ and an outcome $o_2 \in \langle \psi_2 \rangle$ such that $o_1 \succ_{a_1}^+ o_2$. Since $a \in A$, it then follows from the Definition 3 that $o_1 \succ_A^+ o_2$. Therefore, $o_1 \in W_2$, by the definition of W_2 .

(ii) $A_2 \subseteq A_1$. We first show that

$$\begin{aligned} & \bigcup_{a \in A} \{o \mid \exists o'. o' \in \langle \psi_2 \rangle \wedge o \prec_a^+ o'\} \\ & \subseteq \{o \mid \exists o'. o' \in \langle \psi_2 \rangle \wedge o \prec_A^+ o'\}. \end{aligned} \quad (1)$$

For any $o_1 \in \bigcup_{a \in A} \{o \mid \exists o'. o' \in \langle \psi_2 \rangle \wedge o \prec_a^+ o'\}$, there is an agent $a_1 \in A$ and an outcome $o_2 \in \langle \psi_2 \rangle$ such that $o_1 \prec_{a_1}^+ o_2$. Then, $o_1 \prec_A^+ o_2$ by Definition 3 because $a_1 \in A$. Hence, statement (1) is true. Thus, it follows from the definitions of A_1 and A_2 that

$$\begin{aligned} A_2 &= \{o \mid \forall o'. o' \in \langle \psi_2 \rangle \Rightarrow o \not\prec_a^+ o'\} \\ &= O \setminus \{o \mid \exists o'. o' \in \langle \psi_2 \rangle \wedge o \prec_a^+ o'\} \\ &\subseteq O \setminus \bigcup_{a \in A} \{o \mid \exists o'. o' \in \langle \psi_2 \rangle \wedge o \prec_a^+ o'\} \\ &= \bigcap_{a \in A} \{o \mid \forall o'. o' \in \langle \psi_2 \rangle \Rightarrow o \not\prec_a^+ o'\} = A_1. \quad \square \end{aligned}$$

The above theorem implies the relationships between different semantics shown in Figure 3.




Figure 3. Relative Ordering of Semantics of Preference Queries

5 Answering Preference Queries

We now proceed to show how to answer multi-stakeholder preference queries. Specifically, we show that multi-stakeholder preference queries can be reduced to evaluating a corresponding alternation-free modal μ -calculus expression. This allows us to take advantage of the state-of-the-art tools for μ -calculus model-checking to efficiently answer multi-stakeholder preference queries.

Modal μ -calculus. Modal μ -calculus [25, 16], $\mathcal{L}_\mu$, extends propositional modal logic by adding the least and the greatest fixed point operators. $\mathcal{L}_\mu$ uses explicit fixed point and modal operators to express temporal properties over events and states in a labeled transition system. *Labeled transition systems* consist of a set of states, a transition relation over state-pairs parameterized with events (transition annotations) and a labeling function that maps each state

$$\begin{aligned} [\text{tt}]_e &= O \cup \{\perp\} & [\text{ff}]_e &= \emptyset & [p]_e &= \{o \mid p \in L(o)\} \\ [\neg\varphi]_e &= O - [\varphi]_e & [\varphi_1 \wedge \varphi_2]_e &= [\varphi_1]_e \cap [\varphi_2]_e \\ [\langle A \rangle \varphi]_e &= \{s \mid \exists s'. ((A \cap A' \neq \emptyset) \wedge s \xrightarrow{A'} s') \wedge s' \in [\varphi]_e\} \\ [Z]_e &= e(Z) \\ [\mu Z. \varphi]_e &= f_{Z, \varphi, e}^{|\mathcal{O}|}(\emptyset) \text{ where } f_{Z, \varphi, e}(O') = [\varphi]_{e[Z \mapsto O']} \text{ and } O' \subseteq O \end{aligned}$$

Figure 4. μ -calculus Semantics

to a set of propositions that hold in that state. It is easy to see that an induced preference graph can be viewed as a labeled transition system over $O \cup \{\perp\}$, an annotated transition relation (edges being annotated with the set of stakeholders), and a labeling function mapping each outcome to the set of propositions satisfied by the outcome. The primary difference is that the edge-annotation is a set (in an induced preference graph) rather than a symbol (in a labeled transition systems). Note, however, that such a difference is purely syntactical as we can replace an edge annotated with a set by a set of edges, where each edge in the set is annotated by a distinct member of the set. We will use ‘states’ and ‘outcomes’ interchangeably in referring to an induced preference graph interpreted as a labeled transition system.

Syntax of Modal μ -calculus. The syntax of μ -calculus involves propositional constants, atomic propositions, modalities, fixed point variables and expressions and Boolean connectives:

$$\phi \rightarrow \text{tt} \mid \text{ff} \mid \text{AP} \mid \neg\phi \mid \phi \wedge \phi \mid \phi \vee \phi \mid \langle A \rangle \phi \mid Z \mid \mu Z. \phi$$

In the above, the parameter A of the modal operator $\langle A \rangle$ is associated with the edge annotation of the labeled transition system. In an induced preference graph, each edge is annotated with a subset of all stakeholders. In our context, in the modal operators, A will represent a set of stakeholders. When A is singleton such as $A = \{a\}$, we will denote the modal condition as $\langle a \rangle$.

Semantics of Modal μ -calculus. The semantics of μ -calculus formula is given in terms of a set of states in a labeled transition system that satisfy the formula. The semantics is specified by the function $[\cdot] : \Phi \times \mathcal{E} \times \mathcal{I} \rightarrow \mathcal{P}(O)$ where $\mathcal{E}$ is the power set of mappings of fixed point variables to outcomes in O . This mapping is referred to as the environment: $e : \mathcal{Z} \rightarrow \mathcal{P}(O)$; $\mathcal{Z}$ being the set of fixed point variables in the formula whose semantics is being evaluated. We will use the notation $e[Z \mapsto O']$ to denote the environment where the mapping of fixed point variable Z in e is updated to $O' \subseteq O$. We omit I when it is not necessary to distinguish between different induced preference graphs.

Figure 4 shows the semantics of μ -calculus. The propositional constants tt and ff are satisfied by all states and no states, respectively. The atomic proposition p is satisfied in all states whose labeling includes p . The formula $\varphi_1 \wedge \varphi_2$ is satisfied by all states that satisfy both φ_1 and φ_2 . The formula $\langle A \rangle \varphi$ is satisfied by any state which has at least one next state (reachable via an edge annotated with a set that has a non-empty intersection with A) that satisfies φ .

The semantics of fixed point variable Z is given by the environment mapping e . The semantics of least fixed point formula $\mu Z. \varphi$ is computed by the $|\mathcal{O}|$ applications of function $f_{Z, \varphi, e}$ on $\emptyset$ (Tarski-Knaster fixed point theorem [39]). We omit the greatest fixed point construct as its semantics can be realized using the least fixed point and negation.

Model checking a labeled transition system against a given

μ -calculus formula amounts to identifying the set of states in the transition system that belong to the semantics of the μ -calculus formula.

Alternation-Free Modal μ -calculus. For our purposes, it turns out that we only need the alternation-free fragment $\mathcal{L}_\mu^{af}$ [17] of $\mathcal{L}_\mu$. An attractive property of $\mathcal{L}_\mu^{af}$ is that in it there is no real interaction between least and greatest fixpoint operators [27], which, at the expense of reduced expressive power relative to $\mathcal{L}_\mu$, yields more efficient reasoning [27, 16].

Translating Query Language to μ -calculus. We present a strategy to evaluate the proposed preference queries using model checking. We will augment the induced preference graph which encodes a labeled transition system with additional reverse edges; this will help in explaining the answers to multi-stakeholder preference queries in relation to the stakeholder preferences and the chosen semantics; however, in the implementation, such reverse edges can be handled implicitly. For every edge from o_j to o_i due to preference $o_i \succ_a o_j$ of stakeholder a , we will add a reverse edge from o_i to o_j .

Therefore, the set $\{o \mid \exists o'. o' \in \langle \psi \rangle \wedge o \succ_a^+ o'\}$ can be expressed in μ -calculus as: $\mu Z.(\langle a \rangle_r \psi \vee \langle a \rangle_r Z)$. The semantics captures the set of states which can reach some state satisfying ψ via one or more reverse edges; the modal requirement $\langle a \rangle_r$ is satisfied using reverse edges annotated with a .

Example 9 Consider the formula $\mu Z.(\langle 1 \rangle_r \text{Code} \vee \langle 1 \rangle_r Z)$ representing the set of all states that have a path to a state satisfying *Code* via one or more edges annotated with 1. We evaluate this expression using the induced preference graph shown in Figure 2. The semantic computation involves $|O|$ applications of the function $f_{Z,\varphi,e}(O') = [(\langle 1 \rangle_r \text{Code} \vee \langle 1 \rangle_r Z)]_{e[Z \mapsto O']}$. Since this is a least fixed point expression, Z is first mapped to $\emptyset$. Next, Z is mapped to $\{o_1, o_5\}$, the set of outcomes from which one can reach an outcome satisfying *Code* via one reverse edge annotated by 1. The next application of function f with this mapping of Z leaves it unchanged at $\{o_1, o_5\}$ (a fixed point) as there are no additional outcomes leading to some outcome in $\{o_1, o_5\}$ via edges annotated by 1.

Similarly, the set $\{o \mid \forall o'. o' \in \langle \psi \rangle \Rightarrow o \not\succ_a^+ o'\}$ is equal to $O \setminus \{o \mid \exists o'. o' \in \langle \psi \rangle \wedge o \prec_a^+ o'\}$, which can be expressed in μ -calculus as: $\neg \mu Z.(\langle a \rangle \psi \vee \langle a \rangle Z)$. This semantics yields the set of states which have no path to any state that satisfies ψ . Hence, a query of the form $p \text{ P}_a q$, where p and q are atomic propositions, can be expressed in μ -calculus as:

$$p \wedge [\mu Z.(\langle a \rangle_r q \vee \langle a \rangle_r Z)] \wedge [\neg \mu Z.(\langle a \rangle q \vee \langle a \rangle Z)]$$

Now, in *Collective Collaboration*, the set of outcomes that dominate at least one outcome satisfying ψ for a set A of stakeholders is given by $\bigcup_{a \in A} \{o \mid \exists o'. o' \in \psi \wedge o \succ_a^+ o'\}$ which in turn is reflected by the semantics of the μ -calculus formula: $\bigvee_{a \in A} (\mu Z.(\langle a \rangle_r \psi \vee \langle a \rangle_r Z))$. The preceding formula identifies the set of outcomes that have path(s) to some outcome satisfying ψ in the transpose-induced preference graph (i.e., using reversed edges) I . Along each path that decides reachability, each of the edges must be annotated by the same a .

On the other hand, in the *Constructive Collaboration*, the domination of outcomes over at least one outcome satisfying ψ for a set A of stakeholders is decided by $\{o \mid \exists o'. o' \in \psi \wedge o \succ_A^+ o'\}$ which in turn corresponds to the semantics of the μ -calculus formula $\mu Z.(\langle A \rangle_r \psi \vee \langle A \rangle_r Z)$. This denotes the set of outcomes that have

path(s) to some outcome satisfying ψ in the transpose-induced preference graph I ; the reachability is determined by the edges annotated by at least one element from A .

Thus, the Witness and Agreement Conditions can be expressed in μ -calculus as follows:

$$\begin{aligned} W_1 : & \text{ semantics of } \bigvee_{a \in A} (\mu Z.(\langle a \rangle_r \psi \vee \langle a \rangle_r Z)) \\ W_2 : & \text{ semantics of } \mu Z.(\langle A \rangle_r \psi \vee \langle A \rangle_r Z) \\ A_1 : & \text{ semantics of } \bigwedge_{a \in A} (\neg \mu Z.(\langle a \rangle \psi \vee \langle a \rangle Z)) \\ A_2 : & \text{ semantics of } \neg \mu Z.(\langle A \rangle \psi \vee \langle A \rangle Z) \end{aligned}$$

Figure 5 shows the translation function that, given an expression in the multi-stakeholder preference query language and the chosen multi-stakeholder preference semantics as arguments, outputs the corresponding $\mathcal{L}_\mu^{af}$ expression. The run-time for translation is linear in the size (n) of the number of operators ($\wedge, \vee, \neg, P$) in the query. The size of the translation is of the order $O(|A|^k \times n^k)$, where $|A|$ is the number of stakeholders, k the nesting depth of the queries of the form $\psi_1 \text{ P}_A \psi_2$ and n the size of the query. For instance, for a query of the form $(p \text{ P}_A (q \text{ P}_B r))$, n and k are both equal to 2. The run-time for model checking $\mathcal{L}_\mu^{af}$ formula is linear in the size of the formula and the state space of the labeled transition system (induced preference graph). We expect the nesting depth of the query to be reasonably small and the run-time will be determined largely by the number of stakeholders in the query and the size of the number of outcomes (size of the induced preference graph). Note, however, that the number of outcomes is exponential in the number of attributes describing the outcomes, as in the case of reasoning with qualitative preferences [20].

The following theorem establishes the correctness of reduction of multi-stakeholder preference queries to $\mathcal{L}_\mu^{af}$ expressions.

Theorem 2 For a multi-stakeholder preference query ψ (as described in Section 4), $o \in \langle \psi \rangle_I^i$ if and only if $o \in [Tr^i(\psi)]_I$, where I is the preference graph induced by the stakeholder preferences and i denotes the type (consensus or variants of collaborative) of semantics used to answer ψ .

The proof of Theorem 2 proceeds by induction over the structure of the multi-stakeholder preference query.

6 Implementation and Experiments

We have implemented a multi-stakeholder preference reasoner in XSB tabled logic programming environment [38] to demonstrate the viability of our approach. The logical encoding of $\mathcal{L}_\mu^{af}$ used allows for *on-the-fly* evaluation of logical queries, circumventing the need for constructing the complete multi-stakeholder induced preference graph. In other words, only the portion of the induced preference graph relevant for answering the query is constructed, resulting in significant savings in computational and memory savings relative to a naive implementation.

Our experiments with synthetic stakeholder preferences with varying numbers of attributes, number of stakeholders, different choices of collaborative semantics and nesting depths of queries show that all of the queries returned results in less than 2 seconds; For instance, a query of the form

$$tt \text{ P}_{L_3} (tt \text{ P}_{L_2} (tt \text{ P}_{L_1} tt))$$

$$\begin{aligned}
Tr^t(X) &= X && \text{if } X \text{ is proposition or propositional constants} \\
Tr^t(\neg\psi) &= \neg Tr^t(\psi) \\
Tr^t(\psi_1 \text{ } b \text{ } \psi_2) &= Tr^t(\psi_1) \text{ } b \text{ } Tr^t(\psi_2) && \text{where } b \in \{\wedge, \vee\} \\
Tr^t(\psi_1 \text{ } P_A \text{ } \psi_2) &= \begin{cases} Tr^{cs}(\psi_1) \wedge \bigwedge_{a \in A} [(\mu Z.(\langle a \rangle_r Tr^{cs}(\psi_2) \vee \langle a \rangle_r Z)) \wedge (\neg \mu Z.(\langle a \rangle_r Tr^{cs}(\psi_2) \vee \langle a \rangle_r Z))] & \text{if } t = cs \\
Tr^{W_1 A_2}(\psi_1) \wedge \bigvee_{a \in A} (\mu Z.(\langle a \rangle_r Tr^{W_1 A_2}(\psi_2) \vee \langle a \rangle_r Z)) \wedge \neg \mu Z.(\langle A \rangle_r Tr^{W_1 A_2}(\psi_2) \vee \langle A \rangle_r Z) & \text{if } t = W_1 A_2 \\
Tr^{W_1 A_1}(\psi_1) \wedge \bigvee_{a \in A} (\mu Z.(\langle a \rangle_r Tr^{W_1 A_1}(\psi_2) \vee \langle a \rangle_r Z)) \wedge \bigwedge_{a \in A} (\neg \mu Z.(\langle a \rangle_r Tr^{W_1 A_1}(\psi_2) \vee \langle a \rangle_r Z)) & \text{if } t = W_1 A_1 \\
Tr^{W_2 A_2}(\psi_1) \wedge \mu Z.(\langle A \rangle_r Tr^{W_2 A_2}(\psi_2) \vee \langle A \rangle_r Z) \wedge \neg \mu Z.(\langle A \rangle_r Tr^{W_2 A_2}(\psi_2) \vee \langle A \rangle_r Z) & \text{if } t = W_2 A_2 \\
Tr^{W_2 A_1}(\psi_1) \wedge \mu Z.(\langle A \rangle_r Tr^{W_2 A_1}(\psi_2) \vee \langle A \rangle_r Z) \wedge \bigwedge_{a \in A} (\neg \mu Z.(\langle a \rangle_r Tr^{W_2 A_1}(\psi_2) \vee \langle a \rangle_r Z)) & \text{if } t = W_2 A_1 \end{cases}
\end{aligned}$$

Figure 5. Translation of Multi-Stakeholder Preference Queries into μ -calculus

with 8 attributes where $L_1 = \{2, 3, 4\}$, $L_2 = \{5, 6, 7\}$ and $L_3 = \{8, 9, 10\}$ (where $\forall i, L_i \subseteq \mathcal{A}$ where $\mathcal{A} = \{1, \dots, N\}$ is the set of all stakeholders) can be evaluated in less than 2 seconds for all possible choices of collaborative semantics. Additional preliminary experiments show that queries on graphs with ≤ 200 vertices and with at most 400 edges per stakeholder (with nesting depth 3) can be answered within 20 seconds. Additional details of the experiments are given in [2]. These results are indicative of the practical feasibility of our approach.

7 Summary and Discussion

Summary. We provided the first formal treatment of reasoning with multi-stakeholder preferences in a setting where each stakeholder expresses their preferences in a qualitative preference language. We introduced a query language for expressing queries with respect to the preferences of a given set of stakeholders over sets of outcomes. Motivated by the needs of application scenarios, we introduced and analyzed several alternative semantics for such queries and examined their inter-relationships. We provided a provably correct algorithm for answering multi-stakeholder preference queries using model checking in alternation-free μ -calculus. Results of preliminary experiments demonstrate the feasibility of the approach.

Related Work. Existing approaches to reasoning about qualitative preferences of multi-stakeholders leverage voting-based social choice mechanisms [30, 31, 12], starting with the seminal work of Rossi et al. [30]. The applicability of such approaches is limited to settings where the stakeholder preferences are expressed over outcomes (rather than attributes of outcomes); or when they are expressed over attributes of an outcome, they are rather simple (e.g., expressible using CP-nets). A major focus of the social choice based approaches to multi-stakeholder preference reasoning is on voting strategies that are resistant to manipulation by some of the stakeholders and guarantee *fair* outcomes. The key aspects of our work that distinguish from social choice model such as mCP-net [30] are as follows: We seek to answer queries of the form $\psi_1 \text{ } P_A \text{ } \psi_2$, i.e., identify outcomes that satisfy ψ_1 and are more preferred to outcomes satisfying ψ_2 , and are not less preferred to any outcome that satisfies ψ_2 by the set A of stakeholders, whereas mCP-net queries are about whether one outcome is preferred to another by the given set of stakeholders. The precise conditions for deciding the answer to $\psi_1 \text{ } P_A \text{ } \psi_2$ depends on the type of semantics. In the special case where the set of outcomes satisfying ψ_2 is a singleton

set, then our semantics is similar to Pareto semantics defined in [30]. This raises the possibility of extending voting-based semantics where the set of outcomes satisfying ψ_2 is not a singleton set. In such as setting, one may use voting to identify an outcome (say o) that is preferred by a majority of the stakeholders, and include it in the solution set if it is preferred to one of the outcomes satisfying ψ_2 (similar to the witness condition in the paper), and all of them are not preferred to o (similar to the agreement condition in the paper), with the pair-wise outcome preferences decided using a voting mechanism.

Discussion. The framework introduced in this paper is especially useful in applications where it is necessary for multiple stakeholders to be able to express, explore and understand the implications of their preferences in settings where (i) the individual stakeholder preferences are naturally expressed over attributes of outcomes (as opposed to outcomes themselves), and are sufficiently nuanced to require more expressive preference languages e.g., TCP-nets [6] (which involve tradeoffs between conditional preferences), CI-nets [5] (which can express preferences between sets of objects), or their generalizations [34]; and (ii) there is a need for explanations of the role played by the preferences of different stakeholders in determining the outcomes of multi-stakeholder deliberations. One can envision extending this approach to allow individual stakeholders, once they understand the impact of their respective preferences, to minimally revise their preferences to arrive at a consensus that might otherwise have eluded them.

Work in progress. Work in Progress aims to (i) consider organizational structures that further constrain how preferences of multiple stakeholders influence outcomes (e.g., preferences of superiors overriding those of subordinates) (ii) generate targeted explanations of the answers to multi-stakeholder preference queries, (iii) support interactive revision of preferences by stakeholders in the search for consensus or compromise, (iv) further optimize the implementation of the multi-stakeholder preference reasoner, and rigorously assess its scalability as a function of the relevant factors, and (v) apply the resulting tools to support multi-stakeholder decision-making in public policy, healthcare, etc.

Acknowledgements

This work is supported in part by the National Science Foundation through grants IIS 2225823 and IIS 2225824.

References

- [1] Abdulaziz Alashaikh and Eisa Alanazi, 'Conditional preference networks for cloud service selection and ranking with many irrelevant attributes', *IEEE Access*, **9**, (2021).
- [2] Samik Basu, Vasant Honavar, Ganesh Ram Santhanam, and Jia Tao. Representing and reasoning with multi-stakeholder qualitative preference queries, 2023. <https://arxiv.org/abs/2307.16307>.
- [3] Stefano Bistarelli, Pamela Peretti, and Irina Trubitsyna, 'Answer set optimization for and/or composition of cp-nets: A security scenario', in *Principles and Practice of Constraint Programming – CP 2007*, ed., Christian Bessière, pp. 773–781, Berlin, Heidelberg, (2007). Springer Berlin Heidelberg.
- [4] 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).
- [5] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang, 'Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods', in *International Joint Conference on Artificial Intelligence*, pp. 67–72, (2009).
- [6] R. I. Brafman, C. Domshlak, and S. E. Shimony, 'On graphical modeling of preference and importance', *Journal of Artificial Intelligence Research*, **25**, 389–424, (2006).
- [7] Ronen Brafman and Carmel Domshlak, 'Preference handling - an introductory tutorial', *AI magazine*, **30**(1), (2009).
- [8] Patricia Flatley Brennan and Indiana Strombom, 'Improving health care by understanding patient preferences: the role of computer technology', *Journal of the American Medical Informatics Association*, **5**(3), 257–262, (1998).
- [9] A. Cimatti, E. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, and A. Tacchella, 'NuSMV Version 2: An OpenSource Tool for Symbolic Model Checking', in *Proc. Intl. Conf. on Computer-Aided Verification*, Copenhagen, Denmark, (July 2002). Springer.
- [10] Edmund Clarke, Orna Grumberg, and Doron Peled, *Model Checking*, MIT Press, January 2000.
- [11] Cristina Cornelio, Judy Goldsmith, Umberto Grandi, Nicholas Mattei, Francesca Rossi, and K. Brent Venable, 'Reasoning with PCP-nets', *Journal of Artificial Intelligence Research*, **72**, 1103–1161, (2021).
- [12] Cristina Cornelio, Umberto Grandi, Judy Goldsmith, Nicholas Mattei, Francesca Rossi, and K. Brent Venable, 'Reasoning with pcp-nets in a multi-agent context', in *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pp. 969–977. International Foundation for Autonomous Agents and Multi-agent Systems, (2015).
- [13] CRISNER: conditional and relative importance statement network preference reasoner. <http://www.ece.iastate.edu/~gsanathan/crisner.html>.
- [14] Carmel Domshlak, Eyke Hüllermeier, Souhila Kaci, and Henri Prade, 'Preferences in ai: An overview', *Artif. Intell.*, **175**(7-8), 1037–1052, (2011).
- [15] Jon Doyle and Richmond H. Thomason, 'Background to qualitative decision theory', *AI magazine*, **20**, 55–68, (1999).
- [16] E. A. Emerson, C. S. Jutla, and A. P. Sistla, 'On model checking for the μ -calculus and its fragments', *Theoretical Computer Science*, **258**(1-2), 491–522, (2001).
- [17] EA Emerson and CL Lei, 'Efficient model checking in fragments of the propositional mu-calculus', in *IEEE Symposium on Logic in Computer Science*, pp. 267–278. IEEE Computer Society Press, (1986).
- [18] Sheik Mohammad Mostakim Fattah, Athman Bouguettaya, and Sajib Mistry, 'A cp-net based qualitative composition approach for an iaas provider', in *International Conference on Web Information Systems Engineering*, (2021).
- [19] Simon French, *Decision Theory: An Introduction to the Mathematics of Rationality*, Ellis Horwood Limited, 1986.
- [20] Judy Goldsmith, Jérôme Lang, Mirosław Truszczyński, and Nic Wilson, 'The computational complexity of dominance and consistency in CP-nets', *JAIR*, **33**, 403–432, (2008).
- [21] Maheedhar Gunasekharan, Samik Basu, and Ganesh Ram Santhanam, 'Selecting the minimal set of preferred responses to counter detected intrusions', in *Proceedings of the 12th Annual Conference on Cyber and Information Security Research*, CISRC '17, pp. 5:1–5:8, New York, NY, USA, (2017). ACM.
- [22] Sven Ove Hansson, 'Changes in preference', *Theory and Decision*, **38**, 1–28, (1995).
- [23] Piotr Jankowski, Michał Czepkiewicz, Marek Młodkowski, and Zbigniew Zwoliński, 'Geo-questionnaire: A method and tool for public preference elicitation in land use planning', *Transactions in GIS*, **20**(6), 903–924, (2016).
- [24] R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*, Cambridge University Press, 1993.
- [25] D. Kozen, 'Results on the propositional μ -calculus', *Theoretical Computer Science*, (1983).
- [26] Sotirios Liaskos, Sheila A. McIlraith, Shirin Sohrabi, and John Mylopoulos, 'Integrating preferences into goal models for requirements engineering', in *RE*, pp. 135–144. IEEE Computer Society, (2010).
- [27] Johannes Marti and Yde Venema, 'A focus system for the alternation-free μ -calculus', in *International Conference on Automated Reasoning with Analytic Tableaux and Related Methods*, pp. 371–388. Springer, (2021).
- [28] Zachary J. Oster, Ganesh Ram Santhanam, Samik Basu, and Vasant Honavar, 'Model checking of qualitative sensitivity preferences to minimize credential disclosure', in *FACS*, pp. 205–223, (2012).
- [29] Jean-Pierre Queille and Joseph Sifakis, 'Specification and verification of concurrent systems in cesar', in *Proceedings of the 5th Colloquium on International Symposium on Programming*, pp. 337–351, London, UK, UK, (1982). Springer-Verlag.
- [30] F. Rossi, K. B. Venable, and T. Walsh, 'mcp nets: Representing and reasoning with preferences of multiple agents', in *Proceedings of the 19th National Conference on Artificial Intelligence*, AAAI'04, pp. 729–734. AAAI Press, (2004).
- [31] Francesca Rossi, Kristen Brent Venable, and Toby Walsh, 'A short introduction to preferences between artificial intelligence and social choice', *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 102, (2011).
- [32] Ganesh Ram Santhanam, Samik Basu, and Vasant Honavar, 'Dominance testing via model checking', in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, pp. 357–362. AAAI Press, (2010).
- [33] Ganesh Ram Santhanam, Samik Basu, and Vasant Honavar, 'CRISNER: A practically efficient reasoner for qualitative preferences', *CoRR*, [abs/1507.08559](https://arxiv.org/abs/1507.08559), (2015).
- [34] Ganesh Ram Santhanam, Samik Basu, and Vasant Honavar, 'Representing and reasoning with qualitative preferences: Tools and applications', *Synthesis Lectures on Artificial Intelligence and Machine Learning*, **10**(1), 1–154, (2016).
- [35] Shirin Sohrabi, Jorge A. Baier, and Sheila A. McIlraith, 'Preferred explanations: Theory and generation via planning', in *Proceedings of the 25th Conference on Artificial Intelligence (AAAI-11)*, pp. 261–267, San Francisco, USA, (August 2011).
- [36] Shirin Sohrabi and Sheila A. McIlraith, 'Preference-based web service composition: A middle ground between execution and search', in *Proceedings of the 9th International Semantic Web Conference (ISWC10)*, Shanghai, China, (November 2010).
- [37] Tran Cao Son, Enrico Pontelli, and Chitta Baral, 'A non-monotonic goal specification language for planning with preferences', in *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, eds., Thomas Eiter, Hannes Strass, Mirosław Truszczyński, and Stefan Woltran, volume 9060 of *Lecture Notes in Computer Science*, pp. 202–217. Springer, (2015).
- [38] Terrance Swift and David S Warren, 'Xsb: Extending prolog with tabled logic programming', *Theory and Practice of Logic Programming*, **12**(1-2), 157–187, (2012).
- [39] A. Tarski, 'A lattice-theoretical fixpoint theorem and its applications', *Pacific Journal of Mathematics*, (1955).
- [40] Axel van Lamsweerde, 'Goal-oriented requirements engineering: A guided tour', in *RE*, pp. 249–263, (2001).
- [41] Nic Wilson, 'Extending CP-nets with stronger conditional preference statements', in *AAAI*, pp. 735–741, (2004).