

Social Mechanisms for the Collective Engineering of Ontologies¹

Daniele Porello, Nicolas Troquard, Rafael Peñaloza, Roberto Confalonieri, Pietro Galliani, Oliver Kutz

Abstract

Building ontologies collaboratively presents the advantage of allowing practitioners to share their expertise in the modelling of a domain. However, collaborative ontology engineering, seen as a form of knowledge integration, is prone to inconsistencies. We propose two techniques to deal with this situation. First, we study how to repair an inconsistent collective ontology that results from the views of heterogeneous experts, once they have been aggregated by means of voting. Second, we prevent the creation of any inconsistencies by letting the experts engage in a turn-based rational negotiation about the axioms to be added to the collective ontology.

1 Introduction

Ontology engineering is a hard and error-prone task, where even small changes may lead to unforeseen errors, in particular to inconsistency. Ontologies are not only growing in size, they are also increasingly being used in a variety of AI and NLP applications, e.g., [4, 13]. The need for repairing ontologies is all the more prominent when they are developed collaboratively because inconsistencies become inevitable. Nonetheless, collaborative ontologies present a certain advantage, in that the opinions of the experts can be favorably used to drive the repair. Building ontologies collaboratively requires some ways to compute compromises without sacrificing consistency. Relying on simple and well-known decision methods, say majority voting [5, 17], is not enough as they are prone to yield inconsistent ontologies [11]. We propose to obtain consistent compromises that reflect the opinions of voters (stakeholders, citizens, experts, etc.) by repairing the inconsistent collective ontologies.

In our setting, we have a set of experts who need to collaboratively build an ontology about a certain domain. We start with an agenda, which is a (typically inconsistent) set of statements about a domain, expressed as axioms in a description logic. Each expert i submits a (consistent) subset O_i of the agenda, and a preference $<_i$ over the agenda, which is a total ordering that reflects the agents' view of the importance of the statements for the description of the domain. We consider two social mechanisms: vote aggregation and turn-based. In the *vote aggregation mechanism*, we also fix an aggregation procedure F , e.g., majority. The subsets are aggregated into one unique subset $F((O_i)_i)$ of the agenda. As it is very likely that $F((O_i)_i)$ is an inconsistent set of formulas, or ontology, hence, we “repair” this collective ontology. Repairing an ontology can be done in various ways. One common way, a case of coarse repair, is to minimally remove axioms that are the cause of the inconsistency. Here, on the other hand, we propose a fine repair method based on axiom weakening. This has the putative advantage of retaining more information than the coarse repairs. One inconvenience is that axiom weakening must rely on a reference ontology to drive the weakening of the axioms. Moreover, to drive the weakening meaningfully, the reference ontology needs to be consistent and contain information that is sensible to some extent. For this purpose we take advantage of the preference profile ($<_i$) submitted by the experts. We propose a novel method for choosing a consistent subset O^{ref} of an agenda

¹A version of this paper will appear in the Proceedings of the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence. (IJCAI-ECAI 2018).

of formulas from a preference profile over the agenda. A reference ontology obtained this way enjoys a number of desirable formal properties as we will see in Section 3. Finally, the reference ontology O^{ref} is used to repair the aggregated ontology $F((O_i)_i)$. We present the vote aggregation mechanisms in Section 4.

In the *turn-based mechanism*, we do not aggregate the votes of the experts at once. Instead, the experts arbitrarily take turns adding their ‘favorite’ axiom to a set of previously selected axioms. When their favorite axiom cannot be added without causing an inconsistency, this axiom is weakened, using a reference ontology O^{ref} obtained from the orderings submitted by the experts. The procedure ends when all the axioms of the agenda that are supported by at least one expert have been considered (and so added as such or in a weakened form). We present the turn-based mechanism in Section 5.

Analogous mechanisms exist in social choice in the form of multiwinner rules, which among others find applications in parliamentary elections, portfolio/movie selection, or shortlisting [6]. Specific multiwinner elections do not always obey general requirements or principles. For instance, in portfolio selection, one should care mostly about diversity. This is quite the opposite in shortlisting, as one typically looks for a set of similar candidates. In parliamentary elections, we mostly value the proportional representation of the electorate. In the next section we introduce the required formal setting: Description Logic, knowledge refinements, and axiom weakening.

2 Preliminaries

An ontology is a set of formulas in an appropriate logical language with the purpose of describing a particular domain of interest. The precise logic used is not crucial for our approach as most techniques introduced apply to a variety of logics; however, for the sake of clarity we use description logics (DLs) as well-known ontology languages. We briefly introduce the basic DL \mathcal{ALC} ; for full details see [2]. Syntactically, \mathcal{ALC} is based on two disjoint sets N_C and N_R of *concept names* and *role names*, respectively. The set of \mathcal{ALC} *concepts* is generated by the grammar

$$C ::= A \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \forall R.C \mid \exists R.C ,$$

where $A \in N_C$ and $R \in N_R$. A *TBox* is a finite set of concept inclusions (GCIs) of the form $C \sqsubseteq D$ where C and D are concepts. It stores the terminological knowledge regarding the relationships between concepts.² An *ABox* is a finite set of assertions $C(a)$ and $R(a, b)$, which express knowledge about objects in the knowledge domain. An *ontology* is composed by a TBox and an ABox.

The semantics of \mathcal{ALC} is given by *interpretations* $I = (\Delta^I, \cdot^I)$, where Δ^I is a non-empty *domain*, and \cdot^I is a function mapping every individual name to an element of Δ^I , each concept name to a subset of the domain, and each role name to a binary relation on the domain. The interpretation \mathcal{I} is a *model* of the ontology \mathcal{T} if it satisfies all the GCIs and all the assertions in \mathcal{T} . \mathcal{T} is *consistent* if it has at least one model. Given two concepts C and D , C is *subsumed* by D w.r.t. the ontology \mathcal{T} ($C \sqsubseteq_{\mathcal{T}} D$) if $C^I \subseteq D^I$ for every model I of \mathcal{T} . We write $C \equiv_{\mathcal{T}} D$ when $C \sqsubseteq_{\mathcal{T}} D$ and $D \sqsubseteq_{\mathcal{T}} C$.

Many other DLs exist. For instance, \mathcal{EL}_{\perp} is the restriction of \mathcal{ALC} allowing only conjunctions, existential restrictions, and the empty concept \perp [1]. It is widely used in biomedical ontologies for describing large terminologies and is the base of the OWL 2 EL profile. Another common restriction is to use only *acyclic* TBoxes [2]. It is known that consistency can be decided in polynomial time in \mathcal{EL}_{\perp} , in PSPACE in \mathcal{ALC} with acyclic TBoxes,

²Other TBox axioms like the OWL EquivalentClasses, DisjointClasses, DisjointUnion, ObjectPropertyRange, and ObjectPropertyDomain can be normalized into GCIs.

and in EXPTIME in (unrestricted) \mathcal{ALC} . In the following, \mathcal{DL} is an arbitrary DL, and $\mathcal{L}(\mathcal{DL}, N_C, N_R)$ denotes the set of (complex) concepts that can be built over N_C and N_R in \mathcal{DL} .

2.1 Refining Knowledge

Refinement operators are well-known in Inductive Logic Programming, where they are used to learn concepts from examples. In this setting, two types of refinement operators exist: specialisation refinement operators and generalisation refinement operators. While the former constructs specialisations of hypotheses, the latter constructs generalisations [16].

Given the quasi-ordered set $\langle \mathcal{L}(\mathcal{DL}, N_C, N_R), \sqsubseteq \rangle$, one possible generalisation refinement operator is defined as follows:

$$\gamma_{\mathcal{T}}(C) \subseteq \{C' \in \mathcal{L}(\mathcal{DL}, N_C, N_R) \mid C \sqsubseteq_{\mathcal{T}} C'\} .$$

Whereas a specialisation refinement operator is defined as follows:

$$\rho_{\mathcal{T}}(C) \subseteq \{C' \in \mathcal{L}(\mathcal{DL}, N_C, N_R) \mid C' \sqsubseteq_{\mathcal{T}} C\} .$$

Generalisation refinement operators take a concept C as input and return a set of descriptions that are more general than C by taking an ontology \mathcal{T} into account. A specialisation operator, instead, returns a set of descriptions that are more specific.

Our objective is not to propose new refinement operators. Instead, the proposal laid out in this paper can make use of any such operators. We only expect the refinement operators to satisfy a few formal properties (where op^* denotes the unbounded finite iteration of the refinement operator op): For every \mathcal{DL} ontology \mathcal{T} , and every concept C :

1. **trivial generalisability:** $\top \in \gamma_{\mathcal{T}}^*(C)$
falsehood specialisability: $\perp \in \rho_{\mathcal{T}}^*(C)$
2. **generalisation finiteness:** $\gamma_{\mathcal{T}}(C)$ is finite
specialisation finiteness: $\rho_{\mathcal{T}}(C)$ is finite.

When specific refinement operators are needed, as in the examples and in the experiments, we use the refinement operators from [14, 15], which satisfy these assumptions. We regard finiteness as important for computational purpose. Trivial generalisability and falsehood specialisability will be crucial for the termination of our mechanisms.

2.2 Axiom Weakening

We can now define the notion of axiom weakening. The set of all weakenings of an axiom with respect to a reference ontology \mathcal{T} is defined as follows.

Definition 1 (Axiom weakening). *Given a GCI $C \sqsubseteq D$ of \mathcal{T} , the set of (least) weakenings of $C \sqsubseteq D$ w.r.t. \mathcal{T} , denoted by $g_{\mathcal{T}}(C \sqsubseteq D)$ is the set of all axioms $C' \sqsubseteq D'$ such that $C' \in \rho_{\mathcal{T}}(C)$ and $D' \in \gamma_{\mathcal{T}}(D)$.*

Given an assertional axiom $C(a)$ of \mathcal{T} , the set of (least) weakenings of $C(a)$, denoted $g_{\mathcal{T}}(C(a))$ is the set of all axioms $C'(a)$ such that $C' \in \gamma_{\mathcal{T}}(C)$.

For every axiom φ , the axioms in the set $g_{\mathcal{T}}(\varphi)$ are weaker than φ .

Lemma 1. *For every axiom φ , if $\varphi' \in g_{\mathcal{T}}(\varphi)$, then $\varphi \models_{\mathcal{T}} \varphi'$.*

Table 1: Agenda Φ_{LP} .

- | |
|---|
| 1. RaiseWages(Switzerland) |
| 2. TaxHighIncomes(Sweden) |
| 3. RaiseWelfare(Switzerland) |
| 4. RaiseWages(Sweden) |
| 5. TaxHighIncomes(Switzerland) |
| 6. RaiseWelfare(France) |
| 7. RaiseWelfare \sqsubseteq LeftPolicy |
| 8. RaiseWelfare \sqsubseteq \neg RaiseWages |
| 9. TaxHighIncomes \sqsubseteq LeftPolicy |
| 10. LeftPolicy \sqsubseteq RaiseWages \sqcup RaiseWelfare \sqcup TaxHighIncomes |
| 11. RaiseWages \sqsubseteq LeftPolicy |

3 Reference ontologies

In Section 4 and Section 5, we present two methods for obtaining consistent collective ontologies. Both methods make critical use of the so-called reference ontology. Hence, we define specifically how we choose this reference ontology and study some of its formal and computational aspects. In the following section, we consider an arbitrary but fixed \mathcal{DL} ontology Φ , called the *agenda*, and a fixed integer number k , which refers to the number of voters building the reference ontology.

In principle, a reference ontology has to be informative about a certain domain of interest. In practice, a good choice is to select a maximally consistent subset of the agenda. We propose the following strategy to elect a maximally consistent subset of the agenda. Suppose that every agent i provides a *total ordering* $<_i$ over the axioms in the agenda, which represents the priority given to that axiom in the choice of the reference ontology; that is, axioms that are lower in the total ordering $<_i$ are more preferred by agent i . We want to select a maximally consistent subset (or *repair*) of the agenda on which the agents agree. We will use the preferences of agents to determine the *best* repair to consider. First, we introduce the following notion of lexicographic ordering, that extends the ordering on elements of a set X to elements of the power set of X .

Definition 2. *Let $<$ be a total ordering over the set X , and $W, W' \subseteq X$. We say that W is lexicographically smaller than W' wrt $<$, $W \prec W'$ iff there exists some $x \in X$ such that $x \in W \setminus W'$, and for all $y < x$ either $y \in W \cap W'$, or $y \notin W \cup W'$.*

This definition was introduced in the context of finding maximally consistent sets of an ontology [10, 8], motivated by its computational properties. In the literature of social choice, the problem of defining an ordering on sets from an ordering on objects is known as the problem of set extensions and to ranking sets of objects [3]. We assess our previous definition from a social choice theoretical perspective. Extending an ordering from objects to sets of objects requires deciding an interpretation of the preferences over sets. Suppose $x < y$. In principle, one could define an ordering on sets that satisfies either $\{x\} \prec \{x, y\}$ or $\{x, y\} \prec \{x\}$. In the former case, the intuition is that getting the set $\{x, y\}$ means receiving one between the mutually incompatible options x and y , without deciding which one. Therefore, since x is better than y , getting x is better than randomly getting one between x and y . In the latter case, the intuition is that getting $\{x, y\}$ means getting the mutually compatible x and y (or getting one between the two but we can choose which one). Hence $\{x, y\}$ is better than the sole $\{x\}$. Definition 2 embraces the second interpretation,

for which any super set of a set is better than the smaller set. In this context of mutually compatible objects, an important property is additive representability [3]. We establish it in the following lemma, by noticing that the lexicographical set extension can be represented by a function $u(x) = 2^{u_0(x)}$ where u_0 is the Borda score associated to $<$.

Lemma 2. *The relation \prec is additive representable; that is, there exists a utility function u such that:*

$$W \prec W' \text{ iff } \sum_{x \in W} u(x) > \sum_{x \in W'} u(x) \quad (1)$$

To generalise this notion to the preferences of several agents, we introduce some notation. Given a total ordering $<$ over X , $[n]_<$ denotes the n -th element of X according to $<$. Given a set X with $|X| = m$, a profile of total orders $\alpha = (<_1, \dots, <_k)$ over X , and $W \subseteq X$, we define for each $n, 1 \leq n \leq m$ the value $[n]_\alpha^W = |\{i \mid [n]_{<_i} \in W\}|$. That is, $[n]_\alpha^W$ expresses the number of orderings in α whose n -th element appears in W . We denote by W_α the m -tuple $([1]_\alpha^W, \dots, [m]_\alpha^W)$.

Definition 3 (Collective Ordering). *Let θ, θ' be two m -tuples. We say that θ is lexicographically smaller than θ' (denoted $\theta <_{\text{lex}} \theta'$) iff there is an $n, 1 \leq n \leq m$ such that $\theta_n > \theta'_n$ and for all $\ell, 1 \leq \ell < n$, $\theta_\ell = \theta'_\ell$.*

Let X be a set with $|X| = m$, α a set of lexicographic orderings over X , and $W, W' \subseteq X$. Then W is α -collectively better than W' (denoted by $W \prec_\alpha W'$) iff $W_\alpha <_{\text{lex}} W'_\alpha$.

Clearly, $W \prec_\alpha W'$ can be decided in linear time on the size of X and the number of agents: one can simply compute $[n]_\alpha^W$ and $[n]_\alpha^{W'}$ for all $n, 1 \leq n \leq m$ until these values differ. Note that Definition 2 is a special case of Definition 3 where α has only one ordering relation, and hence $[n]_\alpha^W$ is always either 1 (if the element belongs to the set) or 0 (if it does not). Contrary to standard lexicographic ordering, \prec_α may have several different minima, among a class of subsets of X ; thus, there may exist several collectively best repairs. However, all these minima are equally satisfying to the agents as a whole, according to their expressed priorities. Interestingly, \prec_α is also additively representable. We present a stronger result, whose proof is very similar to that of Lemma 2, changing the utility function from $2^{u_0(x)}$ to $(k+1)^{u_0(x)}$.

Lemma 3. *There exist utility functions $u_i, 1 \leq i \leq k$ such that $W \prec_\alpha W'$ iff $\sum_{i=1}^k \sum_{x \in W} u_i(x) > \sum_{i=1}^k \sum_{x \in W'} u_i(x)$.*

We study now the properties of \prec_α in terms of social choice. That is, we view the problem of deciding a collective ordering out of a profile α of orderings \prec_i provided by a set \mathcal{A} of k agents as a problem of studying the social welfare functions f that associate to a profile α a collective ordering \prec_α , i.e. $f : (\prec_1, \dots, \prec_n) \mapsto \prec_\alpha$. We denote the collective ordering resulting from α by applying f as simply \prec_α . We focus in particular on the following properties. The function f is *weakly Pareto efficient* iff for every $i \in \mathcal{A}$, if $W \prec_i W'$, then $W \prec_\alpha W'$. f is *anonymous* iff for every permutation σ of the set of agents \mathcal{A} , $f(\prec_1, \dots, \prec_k) = f(\prec_{\sigma(1)}, \dots, \prec_{\sigma(k)})$. f is *monotonic* iff for every profile α and α' such that $\{i \in \mathcal{A} \mid W \prec_i W'\}$ in α is included in $\{i \in \mathcal{A} \mid W \prec_i W'\}$ in α' , if $W \prec_\alpha W'$, then $W \prec_{\alpha'} W'$.

By Lemma 3, we can represent \prec_α by means of a utility function u_α that sums the agents' utility levels. Therefore, when \prec_α is obtained by means of Definition 3, \prec_α satisfies weak Pareto efficiency, anonymity, and monotonicity. For instance, weak Pareto can be established as follows. Every \prec_i can be represented by a utility function according to Lemma 2. Therefore, we have that for every agent the utility of W is strictly greater than the utility of W' . Since by Lemma 3, \prec_α is represented by a utility function that sums

individual utilities, the collective value of W is strictly greater than the value of W' , which entails $W \prec_\alpha W'$.

These properties are appealing in our context: anonymity means that we do not have any information about the most reliable agents, monotonicity entails sensitivity to the consensus provided by the agents, and weak Pareto, as usual, provides a measure of the efficiency of the outcome. In particular, every minimal element w.r.t. the collective ordering is Pareto optimal: to increase the satisfaction of one agent requires decreasing that of another agent.³ Additionally, if $W \subset W'$, then W' is necessarily collectively *strictly* better than W . Hence, to find a collectively best repair, it suffices to find a collectively best *consistent* (CBC) subset of the ontology.

Example 1. Consider the agenda Φ_{LP} on Table 1. Observe that there are three maximally consistent sets in Φ_{LP} , which are $\Phi_{LP} \setminus \{1\}$, $\Phi_{LP} \setminus \{3\}$, and $\Phi_{LP} \setminus \{8\}$. Consider also two voters with orderings $11 <_1 10 <_1 9 <_1 8 <_1 7 <_1 6 <_1 5 <_1 4 <_1 3 <_1 2 <_1 1$ and $1 <_2 7 <_2 3 <_2 4 <_2 2 <_2 5 <_2 9 <_2 6 <_2 8 <_2 11 <_2 10$. The reference ontology $\mathcal{O}^{ref} = \Phi_{LP} \setminus \{8\}$, that is, the agenda minus axiom 8 is a collectively best consistent subset of Φ_{LP} .

First we study the complexity of this problem or, more precisely, its decision variant: given ontology \mathcal{T} with $|\mathcal{T}| = m$, a set α of lexicographic orderings over \mathcal{T} , and an m -tuple θ , decide whether there is a consistent subontology \mathcal{S} such that $\mathcal{S}_\alpha <_{lex} \theta$. We call this the *optimal repair* problem. Recall that we are considering an arbitrary ontology representation language \mathcal{DL} . Hence, our complexity results need to be parameterized with the complexity of deciding consistency in \mathcal{DL} .

Theorem 1. *If ontology consistency is in the class \mathfrak{C} , then optimal repair is in the class $\text{NP}^{\mathfrak{C}}$.*

Proof. To verify that such a subontology exists, we can guess $\mathcal{S} \subseteq \mathcal{T}$, verify in linear time that $\mathcal{S} <_{lex} \theta$, and check through a call to the consistency oracle that \mathcal{S} is consistent. \square

Corollary 1. *The optimal repair problem is EXPTIME-complete in \mathcal{ALC} and PSPACE-complete in \mathcal{ALC} w.r.t. acyclic TBoxes.*

We look now at other decision problems that are relevant in the context of the collective ordering. Consider the problem of *manipulation*: whether an agent can guarantee some properties of the reference ontology by proposing an adequate ordering. In its simplest form, *positive manipulation* asks whether one agent can provide an ordering such that a given axiom $\varphi \in \mathcal{T}$ appears in *all* CBCs; while *negative manipulation* asks whether there is an ordering such that φ appears in *no* CBC. For these problems there is also a *weak* variant, which ask for the existence of at least one CBC \mathcal{S} such that $\varphi \in \mathcal{S}$ or $\varphi \notin \mathcal{S}$, respectively.

Theorem 2. *If ontology consistency is in \mathfrak{C} , positive and negative manipulation are in $\Pi_p^{2\mathfrak{C}}$ and weak manipulation is in $\Sigma_p^{2\mathfrak{C}}$.*

Proof. We only prove weak positive manipulation; all other cases are similar. To show weak positive manipulation, we guess in polynomial time one ordering $<$, and one set \mathcal{S} with $\varphi \in \mathcal{S}$. Then we verify, in $\text{NP}^{\mathfrak{C}}$ (see Theorem 1) that \mathcal{S} is indeed a CBC. \square

³We leave the discussion of further properties of the collective ordering for a dedicated work. We only notice a significant difference with respect to the Arrovian setting. Here, the universal domain assumption fails. For instance, the set of all axioms is always preferred by every agent to any of its subsets. By giving up universal domain, we enable more choices of the collective ordering [7].

For weak positive manipulation, the best strategy for an agent wanting the existence of an CBC that contains an axiom φ is to provide φ as its most desired axiom; that is, as the first element in its ordering.

Theorem 3. *Let \mathcal{T} be an ontology, $\varphi \in \mathcal{T}$, and α a set of total orderings. Then, there exists an ordering $<$ such that φ is in some CBC w.r.t. $\alpha \cup \{<\}$ iff there exists an ordering \prec such that $[1]_{\prec} = \varphi$ and φ is in some CBC w.r.t. $\alpha \cup \{\prec\}$.*

Proof. The “if” direction is trivial, so we only show the converse. Suppose that there exists such an ordering $<$, and let \mathcal{S} be an CBC w.r.t. $\alpha \cup \{<\}$ such that $\varphi \in \mathcal{S}$. Define \prec to be the variant of $<$ with φ as its smallest element. Then, \mathcal{S} is also an CBC w.r.t. $\alpha \cup \{\prec\}$. \square

To conclude, we consider another problem associated to manipulation, called *axiom necessity*: decide whether an axiom $\varphi \in \mathcal{T}$ appears in *all* CBCs w.r.t. a given set of orderings α .

Theorem 4. *If consistency is in \mathfrak{C} , axiom necessity is in $\Pi_p^{2^{\mathfrak{C}}}$.*

Proof. To negate axiom necessity, we guess a set \mathcal{S} , such that $\varphi \notin \mathcal{S}$, and verify in $\text{NP}^{\mathfrak{C}}$ that \mathcal{S} is indeed a CBC. \square

Recall that consistency in \mathcal{ALC} is PSPACE-complete for acyclic TBoxes, and EXPTIME-complete in general. Using the upper bounds from the previous theorems, we obtain tight complexity results for all these decision problems in these logics, too.

Corollary 2. *Manipulation and axiom necessity are PSPACE-complete for \mathcal{ALC} w.r.t. acyclic TBoxes, and EXPTIME-complete for (unrestricted) \mathcal{ALC} .*

3.1 Hardness Results

Notice that the tight complexity results showcased in Corollaries 1 and 2 were derived mainly because reasoning in \mathcal{ALC} is computationally hard. One interesting question is whether the upper bounds shown before are also tight for logics in a lower complexity class. For that reason, we now provide lower bounds for the DL \mathcal{EL}_{\perp} , and the even smaller Horn logic, for which consistency can be decided in polynomial time. We start by showing that optimal repair is NP-hard.

Theorem 5. *The optimal repair problem in Horn logic is NP-complete.*

Proof. The proof is a reduction from the following NP-hard problem: given an inconsistent set of Horn clauses \mathcal{T} , and $h \geq 0$, decide whether there is a consistent subset $\mathcal{S} \subseteq \mathcal{T}$ with $|\mathcal{S}| > h$ [9]. Let \mathcal{T}, h be an instance of this problem, and define $\mathcal{T} = \{\varphi_1, \dots, \varphi_m\}$. We consider m agents, and any arbitrary, but fixed set of total orderings $\alpha = \{<_1, \dots, <_m\}$ such that $[1]_{<_i} = \varphi_i$ for all $i, 1 \leq i \leq m$. Define the m -tuple $\theta = \{h\} \times \{m\}^{m-1}$. Then, for every subontology $\mathcal{S} \subseteq \mathcal{T}$ it holds that $|\mathcal{S}| > h$ iff $\mathcal{S}_{\alpha} <_{\text{lex}} \theta$. \square

Theorem 6. *Negative manipulation and weak negative manipulation are NP-hard; positive manipulation and weak positive manipulation are coNP-hard in Horn logic, even if $k = 1$.*

Proof. Given an inconsistent set of Horn clauses \mathcal{T} and $\varphi \in \mathcal{T}$, it is NP-hard to decide whether there exists a maximal consistent subset $\mathcal{S} \subseteq \mathcal{T}$ not containing φ [9]. For every maximal consistent subset $\mathcal{S} \subseteq \mathcal{T}$, there exists a total ordering $<$ over \mathcal{T} such that \mathcal{S} is the (only) LMC: simply put the axioms in \mathcal{S} as the first elements in $<$. Hence, there exists a maximal consistent subset \mathcal{S} such that $\varphi \notin \mathcal{S}$ iff there exists an ordering $<$ such that φ does not appear in the only LMC, which means that φ appears in none of the LMCs. \square

Table 2: Complexity results for different DLs.

Problem	\mathcal{EL}_\perp	$\mathcal{ALC}(a)$	\mathcal{ALC}
Optimal repair	NP-c	PSP-c	EXPT-c
Weak pos manipulation	coNP-h ¹ - Σ_p^2	PSP-c	EXPT-c
Weak neg manipulation	NP-h ¹ - Σ_p^2	PSP-c	EXPT-c
Positive manipulation	coNP-h ¹ - Π_p^2	PSP-c	EXPT-c
Negative manipulation	NP-h ¹ - Π_p^2	PSP-c	EXPT-c
Axiom necessity	coNP-h- Π_p^2	PSP-c	EXPT-c

Consider now the following problem: given an inconsistent ontology \mathcal{T} and an axiom $\varphi \in \mathcal{T}$, decide whether there is a *maximum* (i.e., of maximal cardinality) consistent subset $\mathcal{S} \subseteq \mathcal{T}$ such that $\varphi \notin \mathcal{S}$. We call this the *repair irrelevance* problem.

Lemma 4. *Repair irrelevance is NP-complete in Horn logic.*

Proof. We provide a reduction from the following NP-complete problem: given an inconsistent set \mathcal{T} of clauses of the form $p \rightarrow q$, where p, q are propositional variables, \top , or \perp , and a clause $\varphi = p_0 \rightarrow q_0 \in \mathcal{T}$, decide whether there is a maximal (w.r.t. set inclusion) consistent subset $\mathcal{S} \subseteq \mathcal{T}$ such that $\varphi \notin \mathcal{S}$ [10]. Given an instance of this problem, with $|\mathcal{T}| = m$, we introduce m new propositional variables p_1, \dots, p_m , and construct the new set of Horn clauses

$$\mathcal{T}' := \{p \rightarrow q \in \mathcal{T} \setminus \{p_0 \rightarrow q_0\} \mid q \neq p_0\} \cup \{p \rightarrow p_i \mid p \rightarrow p_0 \in \mathcal{T}, 1 \leq i \leq m\} \cup \left\{ \bigwedge_{i=1}^m p_i \rightarrow q \right\}$$

It is easy to see that there is a maximal $\mathcal{S} \subseteq \mathcal{T}$ with $\varphi \notin \mathcal{S}$ iff there is a maximum $\mathcal{S}' \subseteq \mathcal{T}'$ such that $\bigwedge_{i=1}^m p_i \rightarrow q \notin \mathcal{S}'$. \square

Using this result, we can prove that axiom necessity is coNP-hard in this same logic.

Theorem 7. *Axiom necessity is coNP-hard in Horn logic.*

Proof. We reduce the repair irrelevance problem to this case. As in the proof of Theorem 5, we consider m agents, where m is the cardinality of the agenda $\mathcal{T} = \{\varphi_1, \dots, \varphi_m\}$, and any fixed set of orderings $\alpha = \{<_1, \dots, <_m\}$ such that $[1]_{<_i} = \varphi_i$ for all $i, 1 \leq i \leq m$. Then, given $\varphi \in \mathcal{T}$, there is a maximum consistent subset $\mathcal{S} \subseteq \mathcal{T}$ such that $\varphi \notin \mathcal{S}$ iff there is a CBC that does not contain φ . \square

All these results are summarized in Table 2. In the table, hardness results marked with the superscript 1 refer to the problem being hard already for one agent.

4 Vote aggregation mechanism

We present a model that extends Judgment Aggregation for the case of ontologies [11, 12]. Recall that the agenda Φ may be inconsistent. We denote the set of all the *consistent* subontologies of Φ as $\text{On}(\Phi)$.

Given a set $\mathcal{A} = \{1, \dots, k\}$ of agents, the voting mechanism asks each agent $i \in \mathcal{A}$ to provide a consistent ontology $O_i \in \text{On}(\Phi)$. An *ontology profile* is a vector $\mathbf{O} = (O_1, \dots, O_k) \in \text{On}(\Phi)^{\mathcal{A}}$ of consistent ontologies, one for each agent. We denote the set of agents that include the axiom φ in their ontology under profile \mathbf{O} by $\mathcal{A}_\varphi^{\mathbf{O}} := \{i \in \mathcal{A} \mid \varphi \in O_i\}$. We consider *ontology aggregators*.

Algorithm 1 VoteBasedCollectiveOntology($\Phi, (\prec_i)_i, (O_i)_i$)

$O^{\text{ref}} \leftarrow \text{ReferenceOntology}(\Phi, (\prec_i)_i)$
 $R \leftarrow F((O_i)_i)$
while R is inconsistent **do**
 $\text{BadAx} \leftarrow \text{FindBadAxiom}(R)$
 $\text{WeakerAx} \leftarrow \text{WeakenAxiom}(\text{BadAx}, O^{\text{ref}})$
 $R \leftarrow R \setminus \{\text{BadAx}\} \cup \{\text{WeakerAx}\}$
Return R

Definition 4 (Ontology aggregators). *An ontology aggregator is a function $F : \text{On}(\Phi)^A \rightarrow 2^\Phi$ mapping any profile of consistent ontologies to an ontology.*

According to this definition, the ontology we obtain as the outcome of an aggregation process may be inconsistent. This is the case of the majority rule, which is nonetheless widely applied in any political scenarios. The majority rule is defined as follows.

Definition 5 (Absolute majority rule). *The absolute majority rule is the ontology aggregator F_m mapping each profile $\mathbf{O} \in \text{On}(\Phi)^A$ to the ontology $F_m(\mathbf{O}) := \{\varphi \in \Phi \mid |\mathcal{A}_\varphi^{\mathbf{O}}| > n/2\}$.*

We illustrate it with an example.

Example 2. *Consider three voters, voting on the agenda Φ_{LP} of Figure 1 as follows:*

	1	2	3	4	5	6	7	8	9	10	11
Voter 1	1	0	1	0	1	0	0	0	0	0	0
Voter 2	0	0	1	0	0	0	1	1	1	1	1
Voter 3	1	1	0	0	1	0	1	1	0	0	1
Majority	1	0	1	0	1	0	1	1	0	0	1

Each voter's vote represents a consistent set of axioms. Nonetheless, the majority chooses axioms 1, 3, and 8 among others, and the set of axioms $\{1, 3, 8\}$ is inconsistent.

We propose to repair inconsistent collective ontologies obtained from the aggregation of individual ontologies. When the collective ontology $F(\mathbf{O})$ is inconsistent, we can adopt a general strategy based on axiom weakening to repair it. The first step is to compute a reference ontology that the agents agree with. For this step, we require all agents to express their preferences in the form of a total ordering between the axioms in the agenda Φ , and compute one of the CBCs, as described in Section 3. This reference ontology, denoted as O^{ref} , will be used as the basis for the definition of the refinement operators. In this way, the generalisations of the axioms will take into account the views of all the agents.

Once the collective reference ontology O^{ref} has been computed, and as long as $F(\mathbf{O})$ is inconsistent, we select a “bad axiom” and replace it with a random weakening of it with respect to O^{ref} ; see Algorithm 1. The subprocedure $\text{FindBadAxiom}(O)$ samples a number of minimally inconsistent subsets $I_1, I_2, \dots, I_k \subseteq O$ and returns one axiom from the ones occurring the most often, i.e., an axiom from the set $\text{argmax}_{\varphi \in O} (|\{j \mid \varphi \in I_j \text{ and } 1 \leq j \leq k\}|)$. The subprocedure $\text{WeakenAxiom}(\varphi, O^{\text{ref}})$ randomly returns one axiom in $g_{O^{\text{ref}}}(\varphi)$ which is weaker than axiom t .

Example 3. *We continue Example 2, where the majority elected an inconsistent subset of Φ_{LP} , viz., $\{1, 3, 5, 7, 8, 11\}$. Consider the following preference orderings over Φ_{LP} .*

$$\begin{aligned} \prec_1 &= 3 < 1 < 5 < 2 < 4 < 6 < 7 < 8 < 9 < 10 < 11 \\ \prec_2 &= 3 < 7 < 8 < 9 < 10 < 11 < 4 < 5 < 6 < 2 < 1 \\ \prec_3 &= 1 < 2 < 5 < 7 < 8 < 11 < 9 < 3 < 4 < 6 < 10 \end{aligned}$$

Algorithm 2 TurnBasedCollectiveOntology($\Phi, (<_i)_i, (O_i)_i$)

```
 $O^{\text{ref}} \leftarrow \text{ReferenceOntology}(\Phi, (<_i)_i)$ 
 $R \leftarrow \emptyset$ 
TreatedAxioms  $\leftarrow \emptyset$ 
Agent  $\leftarrow 1$ 
while not all agents have finished do
  if every axiom in  $O_{\text{Agent}}$  is treated then
    Agent has finished
  else
    Ax  $\leftarrow \text{FavoriteUntreatedAxiom}(<_{\text{Agent}}, O_{\text{Agent}})$ 
    SetToTreated(Ax)
    while  $R \cup \{Ax\}$  is inconsistent do
      Ax  $\leftarrow \text{WeakenAxiom}(Ax, O^{\text{ref}})$ 
     $R \leftarrow R \cup \{Ax\}$ 
  Agent  $\leftarrow (\text{Agent} \bmod |\mathcal{A}|) + 1$ 
Return  $R$ 
```

The subset $\Phi_{LP} \setminus \{8\}$ is again a CBC, and is chosen as reference ontology O^{ref} .

Out of $\{1, 3, 5, 7, 8, 11\}$, the algorithm then randomly chooses between the axioms 1, 3, and 8, which are the “bad” axioms responsible for the inconsistency. Say it picks axiom 1, RaiseWages(Switzerland). Among the weakenings of axiom 1, there is LeftPolicy(Switzerland) which is used to replace axiom 1. The set of axioms $\{\text{LeftPolicy}(\text{Switzerland}), 3, 5, 7, 8, 11\}$ is consistent, and the vote aggregation mechanism is over.

Clearly, substituting an axiom φ with an element from $g_{O^{\text{ref}}}(\varphi)$ cannot diminish the set of models of an ontology. By our assumption 1 and Lemma 1 any GCI is a finite number of refinement steps away from the trivial axiom $\perp \sqsubseteq \top$. Any assertion $C(a)$ is also a finite number of generalisations away from the trivial assertion $\top(a)$. It follows that by repeatedly replacing an axiom with one of its weakenings, the weakening procedure will eventually obtain an ontology with some interpretations. Hence, the algorithm terminates.

5 Turn-based Mechanism

In the turn-based procedure, the agents engage in a rational negotiation about the axioms to be added to the collective ontology. Again, the agents share an agenda and furnish a total order over the axioms in the agenda. They also choose which axioms, among the most preferred ones, they want to propose during their turns. This procedure, described in Algorithm 2, works as follows:

1. Compute a reference ontology O^{ref} of the agenda using the orders provided by the agents. Set the collective ontology R to empty.
2. By turn, each agent i considers their next preferred axiom in their set O_i of chosen axioms.
3. If agent i does not have any more axioms to propose (when $\text{TreatedAxiom} \cap O_i = \emptyset$), then they skip. Agent i has finished.
4. Otherwise, agent i picks $\text{FavoriteUntreatedAxiom}(<_i, O_i)$, which is his favorite axiom Ax in the set $(\Phi \setminus \text{TreatedAxioms}) \cap O_i$. Then, as long as $R \cup \{Ax\}$ is inconsistent, weaken it w.r.t. the collective reference ontology: Ax is set to one of its weakenings. R is then set to $R \cup \{Ax\}$.
5. Agents repeat steps 2–4 until they have processed all their chosen axioms.

Example 4. Consider the agenda Φ_{LP} from Figure 1. Suppose that three experts submit their opinions on this agenda, which is represented below:

Voter 1	$\langle_1 = 1 < 2 < 3 < 4 < 5 < 6 < 7 < 8 < 9 < 10 < 11$ $O_1 = \{1, 2, 3, 4, 5, 6\}$
Voter 2	$\langle_2 = 11 < 10 < 9 < 8 < 7 < 6 < 5 < 4 < 3 < 2 < 1$ $O_2 = \{7, 8, 9, 10, 11\}$
Voter 3	$\langle_3 = 1 < 7 < 11 < 8 < 2 < 3 < 9 < 4 < 6 < 5 < 10$ $O_3 = \{1, 2, 7, 8, 11\}$

The turn-based mechanism proceeds as follows. Considering the experts' profiles, the computed collective reference ontology \mathbf{O}^{ref} is $\Phi_{LP} \setminus \{8\}$, that is, the agenda minus axiom 8 saying that RaiseWelfare and RaiseWages are incompatible. We initialize $R = \emptyset$.

1. Voter 1 chooses axiom 1. $R = \{1\}$.
2. Voter 2 chooses axiom 11. $R = \{1, 11\}$.
3. Voter 3 chooses axiom 7. His first choice, axiom 1, is already treated. $R = \{1, 7, 11\}$.
4. Voter 1 chooses axiom 2. $R = \{1, 2, 7, 11\}$.
5. Voter 2 chooses axiom 10. $R = \{1, 2, 7, 10, 11\}$.
6. Voter 3 chooses axiom 8 (RaiseWelfare \sqsubseteq \neg RaiseWages). This is the axiom he ranked fourth. $R = \{1, 2, 7, 8, 10, 11\}$.
7. Voter 1 chooses axiom 3 (RaiseWelfare(Switzerland)). However, this is inconsistent with axiom 1 and axiom 8 that are already present in R . It is then weakened into

$$(\text{RaiseWages} \sqcup \text{RaiseWelfare} \sqcup \text{TaxHighIncomes})(\text{Switzerland}),$$

that we denote axiom 3w. $R = \{1, 2, 3w, 7, 8, 10, 11\}$.

8. Voter 2 chooses axiom 9. $R = \{1, 2, 3w, 7, 8, 9, 10, 11\}$.
9. Voter 3 is happy with what is already in R and does not care to add anything more. He skips. $R = \{1, 2, 3w, 7, 8, 9, 10, 11\}$.
10. Voter 1 chooses axiom 4. $R = \{1, 2, 3w, 4, 7, 8, 9, 10, 11\}$.
11. Voter 2 is happy with what is already in R and does not care to add anything more. He skips. $R = \{1, 2, 3w, 4, 7, 8, 9, 10, 11\}$.
12. Voter 3 skips.
13. Voter 1 chooses axiom 5. $R = \{1, 2, 3w, 4, 5, 7, 8, 9, 10, 11\}$.
14. Voter 2 skips.
15. Voter 3 skips.
16. Voter 1 chooses axiom 6. $R = \{1, 2, 3w, 4, 5, 6, 7, 8, 9, 10, 11\}$.

The ontology R resulting from this mechanism is very much like the agenda, but axiom 3 (RaiseWelfare(Switzerland)) was weakened to $(\text{RaiseWages} \sqcup \text{RaiseWelfare} \sqcup \text{TaxHighIncomes})(\text{Switzerland})$. In this ontology, it is acknowledged that Switzerland is in the class $\text{RaiseWages} \sqcup \text{RaiseWelfare} \sqcup \text{TaxHighIncomes}$.

The termination of TurnBasedCollectiveOntology is easy to see: at each step, an agent finishes, or an axiom from Φ is set as treated. The justification of the termination of the inner while-loop is the same as we provided for the algorithm VoteBasedCollectiveOntology: for every axiom a weakening will eventually be found that can be added to R without causing an inconsistency. Eventually, all agents will skip and the procedure terminates.

6 Conclusions and Future Work

We presented two procedures for the collaborative and social engineering of ontologies. In the voting-based procedure, the vote of experts (over the axioms of an agenda) is aggregated to build a collective ontology, and, then, the ontology is repaired by weakening its axioms. In the turn-based procedure, the experts are involved in a turn-based negotiation, where at each step they propose to add axioms to the collective ontology, and weaken them if needed.

Both procedures have the advantage of generating a collective ontology that is consistent. Since both methods make critical use of a reference ontology to weaken axioms, we formally defined how to choose it and studied its formal and computational aspects.

As shown in this work, axiom weakening can play an important role in both of these scenarios. In voting-based aggregation mechanisms, weakening allows us to make use of a much wider range of aggregation procedures, regardless of whether they can preserve consistency. In turn-based aggregation mechanisms, weakening allows agents to “settle” on weaker versions of their preferred axioms for addition to the collective ontology, in case the stronger versions would yield inconsistency.

References

- [1] Franz Baader, Sebastian Brandt, and Carsten Lutz. Pushing the \mathcal{EL} envelope. In *IJCAI*, pages 364–369, 2005.
- [2] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, New York, NY, USA, 2003.
- [3] Salvador Barberà, Walter Bossert, and Prasanta K Pattanaik. Ranking sets of objects. In *Handbook of utility theory*, pages 893–977. Springer, 2004.
- [4] J. Bateman, J. Hois, R. Ross, and T. Tenbrink. A Linguistic Ontology of Space for Natural Language Processing. *Artificial Intelligence*, 174(14):1027–1071, 2010.
- [5] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [6] Edith Elkind, Piotr Faliszewski, Jean-François Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. What do multiwinner voting rules do? an experiment over the two-dimensional euclidean domain. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 494–501, 2017.
- [7] Wulf Gaertner. *A Primer in Social Choice Theory*. Oxford University Press, 2006.
- [8] David S Johnson, Mihalis Yannakakis, and Christos H Papadimitriou. On generating all maximal independent sets. *Information Processing Letters*, 27(3):119–123, 1988.
- [9] Rafael Peñaloza. *Axiom pinpointing in description logics and beyond*. PhD thesis, Dresden University of Technology, 2009.
- [10] Rafael Peñaloza and Barış Sertkaya. Understanding the complexity of axiom pinpointing in lightweight description logics. *Artificial Intelligence*, 250:80–104, 2017.
- [11] Daniele Porello and Ulle Endriss. Ontology merging as social choice: Judgment aggregation under the open world assumption. *Journal of Logic and Computation*, 24(6):1229–1249, 2014.
- [12] Daniele Porello, Nicolas Troquard, Roberto Confalonieri, Pietro Galliani, Oliver Kutz, and Rafael Peñaloza. Repairing Socially Aggregated Ontologies Using Axiom Weakening. In *20th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA 2017)*, LNCS. Springer, 2017.

- [13] Edson Prestes et al. Towards a core ontology for robotics and automation. *Robotics and Autonomous Systems*, 61(11):1193 – 1204, 2013. Ubiquitous Robotics.
- [14] Nicolas Troquard, Roberto Confalonieri, Pietro Galliani, Rafael Peñaloza, Daniele Porello, and Oliver Kutz. Repairing Ontologies via Axiom Weakening. *ArXiv e-prints*, November 2017.
- [15] Nicolas Troquard, Roberto Confalonieri, Pietro Galliani, Rafael Peñaloza, Daniele Porello, and Oliver Kutz. Repairing ontologies via axiom weakening. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*, 2018.
- [16] Patrick R.J. van der Laag and Shan-Hwei Nienhuys-Cheng. Completeness and properness of refinement operators in inductive logic programming. *The Journal of Logic Programming*, 34(3):201 – 225, 1998.
- [17] William S. Zwicker. Introduction to the theory of voting. In Brandt et al. [5], pages 23–56.