

Two Approaches to Ontology Aggregation Based on Axiom Weakening

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

Free University of Bozen-Bolzano, Faculty of Computer Science, I-39100 Bozen-Bolzano, Italy
dporello@unibz.it, ntroquard@unibz.it, rpenaloza@unibz.it,
rconfalonieri@unibz.it, pgalliani@unibz.it, okutz@unibz.it

Abstract

Axiom weakening is a novel technique that allows for fine-grained repair of inconsistent ontologies. In a multi-agent setting, integrating ontologies corresponding to multiple agents may lead to inconsistencies. Such inconsistencies can be resolved after the integrated ontology has been built, or their generation can be prevented during ontology generation. We implement and compare these two approaches. First, we study how to repair an inconsistent ontology resulting from a voting-based aggregation of views of heterogeneous agents. Second, we prevent the generation of inconsistencies by letting the agents engage in a turn-based rational protocol about the axioms to be added to the integrated ontology. We instantiate the two approaches using real-world ontologies and compare them by measuring the levels of satisfaction of the agents w.r.t. the ontology obtained by the two procedures.

1 Introduction

Ontologies are increasingly being used in a variety of present-day applications. Nonetheless, ontology engineering is a hard and error-prone task, where even small changes may lead to unforeseen errors, in particular to inconsistency. The need for repairing ontologies is particularly crucial when they are developed by different agents. However, integrating ontologies provided by agents presents a certain advantage, in that their opinions can be used to drive the repair.

Ontology integration requires some ways to compute compromises without sacrificing consistency. Relying on simple and well-known decision methods like majority voting [Brandt *et al.*, 2016; Zwicker, 2016] is not enough as they are prone to yield inconsistent ontologies [Porello and Endriss, 2014]. We propose to obtain consistent compromises that reflect the opinions of agents (voters, citizens, experts, etc.) by repairing the inconsistent aggregated ontologies by weakening the axioms.

This work goes beyond existing methods for ontology aggregation based on ontology matching and alignment

(e.g., [Euzenat, 2015; 2017; Solimando *et al.*, 2017; Jiménez-Ruiz *et al.*, 2016]) and on social choice [Porello and Endriss, 2014]. On the one hand, we deal with global inconsistency, and we do not confine ourselves to ontology alignments, by actually integrating the agents' information; on the other hand, we use axiom weakening to solve in a *fine* way the inconsistencies that arise from the integration, namely, by building new formulas that balance the agents' opinions.

In our setting, we have a set of agents who need to build an ontology about a specific knowledge domain. We start with an agenda, which is a (typically inconsistent) set of statements about the domain, expressed as axioms in a description logic. Each agent submits a (consistent) subset of the agenda, and a preference profile, i.e., a total ordering over the agenda reflecting the importance they give to each axiom. We consider two social approaches to ontology aggregation: vote aggregation and turn-based.

In the *vote aggregation* approach, we fix an aggregation procedure, e.g. majority. The subsets are aggregated into one unique subset of the agenda, which is very likely inconsistent, and needs to be 'repaired.' 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. To enable axiom weakening one has to select a consistent reference ontology that contains enough information. For this purpose, we take advantage of the preference profile submitted by the agents. We propose a novel method for choosing a reference ontology from a preference profile over the agenda, and we study its formal properties. Finally, the reference ontology is used to repair the aggregated ontology.

In the *turn-based* approach, we let the agents take turns when adding their 'favorite' axiom to a set of previously selected axioms in a possibly arbitrary order. When their favorite axiom cannot be added without causing an inconsistency, this axiom is weakened, using the reference ontology obtained from the orderings submitted by the agents. The procedure ends when all the axioms of the agenda that are supported by at least one agent have been considered.

Analogous approaches exist in social choice in the form of

multiwinner rules, which among others find applications in parliamentary elections, portfolio/movie selection, or shortlisting [Elkind *et al.*, 2017]. 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. Different rules will perform differently depending on the application, and finding a more appropriate one is difficult. Experimentation can prove more useful than theory. In the paper, we lay out such an experimentation to compare our fine repair approaches against coarse repair approaches based on axiom removal and unanimity selection.

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 [Baader *et al.*, 2003]. 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 general 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. 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 based on *interpretations* of the form $I = (\Delta^I, \cdot^I)$, 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. I satisfies $C \sqsubseteq D$ iff $C^I \subseteq D^I$ and I satisfies an assertion $C(a)$ ($R(a, b)$) iff $a^I \in C^I$ ($(a^I, b^I) \in R^I$). The interpretation 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 a 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 $D \sqsubseteq_{\mathcal{T}} C$.

We denote by \mathcal{DL} an arbitrary DL, and by $\mathcal{L}(\mathcal{DL}, N_C, N_R)$ the set of (complex) concepts built over N_C and N_R in \mathcal{DL} .

Refining knowledge and weakening axioms

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 construct specialisations of hypotheses, the latter construct generalisations [van der Laag and Nienhuys-Cheng, 1998].

Given the quasi-ordered set $\langle \mathcal{L}(\mathcal{DL}, N_C, N_R), \sqsubseteq \rangle$, a generalisation refinement operator satisfies

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

A specialisation refinement operator satisfies

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

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 more specific descriptions.

Our objective is not to propose new refinement operators. Instead, the proposal laid out in this paper can make use of any such operators. To ensure the termination of the procedures, we require that the following holds

$$\top \in \gamma_{\mathcal{T}}^*(C) \text{ and } \perp \in \rho_{\mathcal{T}}^*(C) \quad (1)$$

(where op^* denotes the unbounded finite iteration of the refinement operator op). That is, every concept can be generalised into \top and specialised into \perp in a finite number of steps. When specific refinement operators are needed, as in the examples and in the experiments, we use the refinement operators from [Troquard *et al.*, 2018].

We can now define the notion of axiom weakening. The set of all weakenings of an axiom w.r.t. 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 $D' = D$ or $C' = C$ 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 2. *For every axiom φ , if $\varphi' \in g_{\mathcal{T}}(\varphi)$, then $\varphi \models_{\mathcal{T}} \varphi'$.*

3 Computing the Reference Ontology

Our two axiom weakening-based approaches rely on a reference ontology. We define specifically how to choose this ontology and study its formal and computational aspects.

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. We assume 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 agents' preferences to determine the *best* repair. 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 3. *Let X be a set, $<$ a total ordering on X , and $W, W' \subseteq X$. W is lexicographically smaller than W' , $W \prec W'$, 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 particular in the context of finding maximally consistent sets of an ontology [Peñaloza and Sertkaya, 2017; Johnson *et al.*, 1988], 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 related to the problem of *ranking sets of objects* [Barberà *et al.*, 2004]. 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 3 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 [Barberà *et al.*, 2004]. We establish it in the following lemma.

Lemma 4. *The relation \prec is additive representable; that is, there exists a utility function u such that:*

$$W \prec W' \text{ iff } \sum_{x \in W} u(x) > \sum_{x \in W'} u(x) \quad (2)$$

One possible utility function that represents \prec is defined by $u(x) = 2^{u_0(x)}$, where for $x_m < x_{m-1} < \dots < x_1$, $u_0(x_j) = j$. To generalize this notion to the preferences of several agents (i.e., to consider several orderings simultaneously), 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 $\alpha = (\prec_1, \dots, \prec_k)$ over X , and $W \subseteq X$, we define for each n , $1 \leq n \leq m$ the value $[n]_\alpha^W = |\{i \mid [n]_{\prec_i} \in W\}|$. 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 5 (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 ℓ , $1 \leq \ell < n$, $\theta_\ell = \theta'_\ell$.*

Let X be a set with $|X| = m$, α a profile 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 3 is a special case of Definition 5 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 4, changing the utility function from $2^{u_0(x)}$ to $(k+1)^{u_0(x)}$.

Lemma 6. *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. $f : (\prec_1, \dots, \prec_k) \mapsto \prec_\alpha$. The collective ordering resulting from α by applying f is denoted by \prec_α . We focus in particular on the following properties of f .

- *weakly Pareto efficiency* for every $i \in \mathcal{A}$, $W \prec_i W'$, then $W \prec_\alpha W'$;
- *anonymity* for every permutation σ of the set of agents \mathcal{A} , $f(\prec_1, \dots, \prec_k) = f(\prec_{\sigma(1)}, \dots, \prec_{\sigma(k)})$; and
- *monotonicity* for every two profiles 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 6, 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 5, \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 4. Therefore, we have that for every agent the utility of W is strictly greater than the utility of W' . Since by Lemma 6, \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.¹ 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) set of axioms.

Example 1. *Consider the agenda Φ_{LP} on Table 1. 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 $\mathcal{O}^{\text{ref}} = \Phi_{LP} \setminus \{8\}$, that is, the agenda minus axiom 8 is a collectively best consistent subset of Φ_{LP} .*

¹Notice a significant difference w.r.t. Arrovian aggregation: the universal domain assumption fails due to Definition 3. For instance, the set of all axioms is always preferred by every agent to any of its subsets. Thus, we enable more choices of the collective ordering.

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

Table 1: Agenda Φ_{LP} .

First we study the complexity of finding a CBC ontology 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 $\mathcal{S}_\alpha \leq_{\text{lex}} \theta$. We call this the *optimal repair* problem.

Theorem 7. *If ontology consistency is in the class \mathcal{C} , then optimal repair is in the class $\text{NP}^{\mathcal{C}}$.*

Corollary 8. *Optimal repair problem is EXPTIME-complete in \mathcal{ALC} and PSPACE-complete in \mathcal{ALC} w.r.t. acyclic TBoxes.*

4 Vote Aggregation Mechanism

We present a model that extends Judgment Aggregation for the case of ontologies [Porello and Endriss, 2014]. 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 of the form $\mathbf{O} = (O_1, \dots, O_k) \in \text{On}(\Phi)^k$ of consistent ontologies, one for each agent. The set of agents including the axiom φ in their ontology under profile \mathbf{O} is $\mathcal{A}_\varphi^{\mathbf{O}} := \{i \in \mathcal{A} \mid \varphi \in O_i\}$.

Definition 9 (Ontology aggregators). *An ontology aggregator is a function $F : \text{On}(\Phi)^k \rightarrow 2^\Phi$ 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 e.g. the majority rule, which is nonetheless widely applied in any political scenarios. The majority rule is defined as follows.

Definition 10 (Absolute majority rule). *The absolute majority rule is the ontology aggregator F_m mapping each $\mathbf{O} \in \text{On}(\Phi)^k$ to the ontology $F_m(\mathbf{O}) = \{\varphi \in \Phi \mid |\mathcal{A}_\varphi^{\mathbf{O}}| > k/2\}$.*

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

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

```

 $O^{\text{ref}} \leftarrow \text{ReferenceOntology}(\Phi, (<_i)_i)$ 
 $R \leftarrow F((O_i)_i)$ 
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  $R$ 

```

It is easy to see that each voter's vote represents a consistent set of axioms. Nonetheless, the majority chooses axioms 1, 3, and 8 among others. The set $\{1, 3, 8\}$ is inconsistent.

We propose to repair inconsistent ontologies obtained from the aggregation of individual ontologies. When the aggregated 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 collective opinion of all the agents.

Once the reference ontology O^{ref} has been computed, and as long as $F(\mathbf{O})$ remains 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 φ .

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} .*

$$\begin{aligned}
<_1 &= 3 < 1 < 5 < 2 < 4 < 6 < 7 < 8 < 9 < 10 < 11 \\
<_2 &= 3 < 7 < 8 < 9 < 10 < 11 < 4 < 5 < 6 < 2 < 1 \\
<_3 &= 1 < 2 < 5 < 7 < 8 < 11 < 9 < 3 < 4 < 6 < 10
\end{aligned}$$

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 Equation 1 and Lemma 2 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

	Tolerant		Strict	
	Voting	Turn	Voting	Turn
dcmitype	0.91	0.94	0.78	0.73
elig	0.89	0.91	0.77	0.67
gco	0.91	0.93	0.80	0.73
icd11	0.88	0.91	0.76	0.69
pgxo	0.89	0.91	0.78	0.66
pseudo	0.90	0.93	0.77	0.64
ptrans	0.89	0.90	0.76	0.64

Table 2: Average agent happiness (tolerant and strict) for the ontologies considered and for voting and turn aggregation approaches.

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 turn-based rational negotiation about the axioms to be added to the collective ontology. As in the case of the voting procedure 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.

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

```

Algorithm 2 works as follows. We compute a reference ontology O^{ref} of the agenda using the orders provided by the agents, and set the collective ontology R to empty.

By turn, the agents consider their next preferred axiom in their set O_i of chosen axioms. If an agent i does not have any more axioms to propose (i.e., $\text{TreatedAxioms} \subseteq O_i$), then it has finished. Otherwise, agent i picks its favorite axiom Ax in the set $(\Phi \setminus \text{TreatedAxioms}) \cap O_i$ through $\text{FavoriteUntreatedAxiom}(<_i, O_i)$. As long as $R \cup \{Ax\}$ is inconsistent, the agent weakens it w.r.t. the collective reference ontology. Then a weaker version of Ax is added to R . This is repeated until all agents have processed all their chosen axioms. It is easy to see that this algorithm terminates.

6 Experiments

With the experiments we aim at understanding which one of the proposed collective repairing procedures is preferable under which conditions. To this end, we compare them by defining the concept of *agent happiness* from two perspectives:

- **Tolerant Happiness:** an agent is *tolerantly happy* when their chosen axioms are in (or follow from) the collective repaired ontology, regardless of whether this ontology also entails further statements.
- **Strict Happiness:** an agent is *strictly happy* when their chosen axioms are in (or follow from) the collective ontology *and* everything that follows from the collective ontology also follows from their own chosen axioms.

To formally define these concepts, we take into account both the logical axioms in a given ontology and the set of *inferred class hierarchies* that can be obtained from it. To this end, given a consistent ontology O , we define

$$\text{Inf}(O) = \{A \sqsubseteq B : A, B \in N_C, O \models A \sqsubseteq B\}.$$

The agent happiness can then be defined by checking how many logical axioms in the ontology of the agent and how many inferred class hierarchies are provable w.r.t. the collective ontology.

Definition 11. Let O_i be the set of axioms chosen by an agent i , and O be a consistent collective ontology, we define:

- the tolerant agent happiness $\text{TotH}(O|O_i)$ as

$$\frac{|\{\varphi \in O_i \cup \text{Inf}(O_i) \text{ s.t. } O \models \varphi\}|}{|O_i \cup \text{Inf}(O_i)|};$$

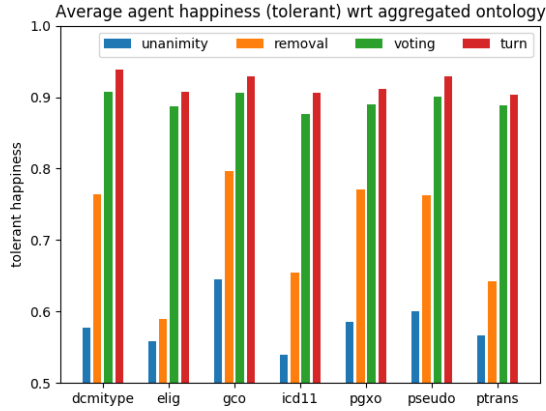
- the strict agent happiness $\text{StrH}(O|O_i)$ as

$$\frac{|\{\varphi \in O_i \cup \text{Inf}(O_i) \cup O \cup \text{Inf}(O) \text{ s.t. } O \models \varphi \text{ and } O_i \models \varphi\}|}{|O_i \cup \text{Inf}(O_i) \cup O \cup \text{Inf}(O)|}.$$

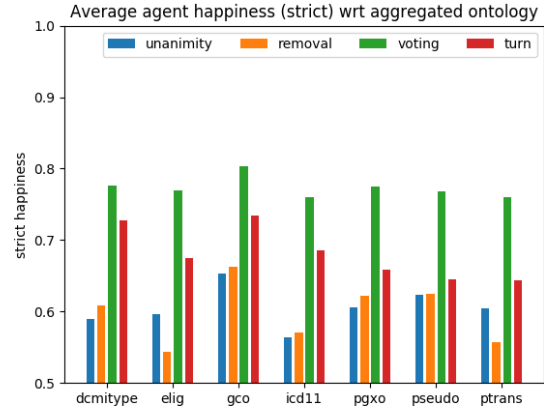
We evaluated our aggregation approaches over 7 ontologies from BioPortal [Matentzoglou and Parsia, 2017]. We compared the agents' happiness of the collective repairing by voting and weakening (Algorithm 1) with the ones of the turn-based procedure (Algorithm 2) by first making the ontologies inconsistent through the addition of random axioms. We also considered, as baselines, the approach *unanimity* which takes the intersection of the votes of all agents (i.e., the ontology aggregator $F_u(O) := \{\varphi \in \Phi \mid |\mathcal{A}_\varphi^O| = k\}$), and the approach *removal* which removes random axioms from the (inconsistent) global agenda until consistency is achieved.

We made each ontology inconsistent 250 times, by selecting 10% of the axioms in an ontology to generate inconsistent subclass axiom chains. Each time and for each agent i , we randomly generated a preference order $<_i$, and chose the individual ontology to be a consistent set of the best axioms of the agenda w.r.t. $<_i$. We then built up two collective ontologies running Algorithm 1 and Algorithm 2. Finally, we measured the average agent happiness (strict and tolerant) for our aggregation procedures.

Our hypothesis was that our procedure would yield better agents' happiness than the removal- and unanimity-based



(a) Average agent tolerant happiness.



(b) Average agent strict happiness.

Figure 1: Average agent happiness using our two methods (voting, turn) and two baselines (unanimity, removal).

procedures. Furthermore, we expected the turn-based approach to prove itself better than the voting-based approach insofar as *tolerant*-based agent happiness is concerned, but not with respect to *strict* agent happiness, due to the fact that voting allows a majority of agents to reject axioms proposed by others while this is not the case in the turn-based approach. On the other hand, we expected the voting-based procedure to be preferable w.r.t. strict happiness: under this regime a majority of agents can force any axiom they disagree with not to appear in the collective ontology, but in the turn-based approach there is no such guarantee. As shown in Figure 1 and Table 2, the results of our experiments strongly confirm our hypothesis² and illustrate the benefits of our weakening-based approach to ontology aggregation.

7 Related Work

Axiom weakening was introduced as a technique to repair inconsistent ontologies from the perspective of a single agent [Troquard *et al.*, 2018]. Here, we repair socially aggregated ontologies. Other works relate to axiom weakening from different perspectives. E.g., a similarly general approach to axiom weakening (and generalization) has been considered in ontology learning [Lehmann and Hitzler, 2010] and in conceptual blending [Confalonieri *et al.*, 2018], whereas finding explanations in ontologies to identify parts of axioms that caused inconsistency—to then modify those axioms instead of removing them—was explored in [Horridge *et al.*, 2008; 2012].

The problem of integrating ontologies has been addressed from the perspective of ontology alignment and matching (e.g., [Euzenat, 2015; 2017; Solimando *et al.*, 2017; Jiménez-Ruiz *et al.*, 2016]). Although we share the main motivation—viz., aggregating ontologies—our approach is intrinsically

different from ontology alignment approaches in the literature. Using the terminology of [Euzenat, 2015], our ontologies are always “locally” consistent, and typically “globally” inconsistent. This is the assumption in the setting of judgment aggregation and computational social choice. Since we deal with global inconsistency, we repair the integrated ontology straightaway by means of axiom weakening. Furthermore, unlike [Euzenat, 2017; Jiménez-Ruiz *et al.*, 2016; Solimando *et al.*, 2017], our framework is not looking for ontology alignments: all individual ontologies are indeed subsets of the same agenda, therefore, we do not focus on identifying correspondences between entities of agents’ ontologies.

Another approach worth mentioning is [Trojahn *et al.*, 2010]. The authors compare different argumentation frameworks where agents use argumentation to decide which ontology matching technique to adopt. In our turn-based procedure, all the agents share the same repairing techniques, and they do not negotiate over the axioms. It suffices to guarantee that added axioms (or their weakened version) keep the integrated ontology consistent. On the other hand, arguing about the quality and informativeness of the axioms to be included in the repaired ontology can be an interesting aspect to be investigated as an extension of our framework.

8 Conclusions and Future Work

We presented two approaches to ontology aggregation. In the voting-based approach, the vote of agents (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 approach, the agents are involved in a turn-based dialogue, where at each step they add axioms to the aggregated ontology, and weaken them if needed. Both approaches have the advantage of generating an aggregated ontology that is consistent. Both methods make critical use of a reference ontology to weaken axioms, and we formally defined a principled method to select it.

As shown in this work, axiom weakening can play an important role in both of these scenarios. In voting-based ag-

²Via Wilcoxon signed-rank significance testing with Holm-Bonferroni correction for multiple tests. For *pseudo*, we did not observe a significant difference between removal and turn-based approaches; otherwise, our hypotheses were all confirmed.

gregation 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 aggregated ontology, if the stronger versions would cause inconsistencies. Our results are strongly indicative of the relevance of axiom weakening to ontology aggregation.

As future work, we aim at exploring different weakening operators (which, in practical applications, would also require expert input in order to select among different weakenings of a given axiom) and studying the effect of various parameters (e.g., the sparsity of the agent ontologies) on the relative effectiveness of the approaches described here.

References

- [Baader *et al.*, 2003] 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.
- [Barberà *et al.*, 2004] Salvador Barberà, Walter Bossert, and Prasanta K Pattanaik. Ranking sets of objects. In *Handbook of utility theory*, pages 893–977. Springer, 2004.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Confalonieri *et al.*, 2018] Roberto Confalonieri, Manfred Eppe, Marco Schorlemmer, Oliver Kutz, Rafael Peñaloza, and Enric Plaza. Upward Refinement Operators for Conceptual Blending in \mathcal{EL}^{++} . *Annals of Mathematics and Artificial Intelligence*, 82(1):69–99, 2018.
- [Elkind *et al.*, 2017] 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*, pages 494–501, 2017.
- [Euzenat, 2015] Jérôme Euzenat. Revision in networks of ontologies. *Artificial Intelligence*, 228:195–216, 2015.
- [Euzenat, 2017] Jérôme Euzenat. Interaction-based ontology alignment repair with expansion and relaxation. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 185–191, 2017.
- [Horridge *et al.*, 2008] Matthew Horridge, Bijan Parsia, and Ulrike Sattler. Laconic and precise justifications in owl. In Amit Sheth, Steffen Staab, Mike Dean, Massimo Paolucci, Diana Maynard, Timothy Finin, and Krishnaprasad Thirunarayan, editors, *The Semantic Web - ISWC 2008*, pages 323–338, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [Horridge *et al.*, 2012] Matthew Horridge, Bijan Parsia, and Ulrike Sattler. Justification masking in ontologies. In *KR*. AAAI Press, 2012.
- [Jiménez-Ruiz *et al.*, 2016] Ernesto Jiménez-Ruiz, Terry R. Payne, Alessandro Solimando, and Valentina Tamma. Limiting logical violations in ontology alignment through negotiation. In *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning, KR’16*, pages 217–226. AAAI Press, 2016.
- [Johnson *et al.*, 1988] David S Johnson, Mihalis Yannakakis, and Christos H Papadimitriou. On generating all maximal independent sets. *Information Processing Letters*, 27(3):119–123, 1988.
- [Lehmann and Hitzler, 2010] Jens Lehmann and Pascal Hitzler. Concept learning in description logics using refinement operators. *Machine Learning*, 78(1-2):203–250, 2010.
- [Matentzoglou and Parsia, 2017] Nicolas Matentzoglou and Bijan Parsia. BioPortal snapshot 30.03.2017, 2017. last accessed, 2017/08/04.
- [Peñaloza and Sertkaya, 2017] Rafael Peñaloza and Barış Sertkaya. Understanding the complexity of axiom pinpointing in lightweight description logics. *Artificial Intelligence*, 250:80–104, 2017.
- [Porello and Endriss, 2014] 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.
- [Solimando *et al.*, 2017] Alessandro Solimando, Ernesto Jiménez-Ruiz, and Giovanna Guerrini. Minimizing conservativity violations in ontology alignments: algorithms and evaluation. *Knowledge and Information Systems*, 51(3):775–819, Jun 2017.
- [Trojahn *et al.*, 2010] Cássia Trojahn, Paulo Quaresma, Renata Vieira, and Antoine Isaac. Comparing argumentation frameworks for composite ontology matching. In Peter McBurney, Iyad Rahwan, Simon Parsons, and Nicolas Maudet, editors, *Argumentation in Multi-Agent Systems*, pages 305–320, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [Troquard *et al.*, 2018] 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*, pages 1981–1988, 2018.
- [van der Laag and Nienhuys-Cheng, 1998] 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.
- [Zwicker, 2016] William S. Zwicker. Introduction to the theory of voting. In Brandt *et al.* [2016], pages 23–56.