

Repairing Socially Aggregated Ontologies Using Axiom Weakening

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

Free University of Bozen-Bolzano, Italy
{Firstname.Lastname}@unibz.it

Abstract. Ontologies represent principled, formalised descriptions of agents' conceptualisations of a domain. For a community of agents, these descriptions may significantly differ. We propose an aggregative view of the integration of ontologies based on Judgement Aggregation (JA). Agents may vote on statements of the ontologies, and we aim at constructing a collective, integrated ontology, that reflects the individual conceptualisations as much as possible. As several results in JA show, many attractive and widely used aggregation procedures are prone to return inconsistent collective ontologies. We propose to solve the possible inconsistencies in the collective ontology by applying suitable *weakenings* of axioms that cause inconsistencies.

1 Introduction

Social choice theory is a branch of economic theory that deals with the design and analysis of mechanisms for aggregating opinions of individual agents to arrive at a basis for a collective decision [6]. An ubiquitous example of such a mechanism is voting, usually intended as voting on preferences in standard social choice. Recently, the model of aggregation has been applied to judgements, or more generally to propositional attitudes, expressed in some logical setting, in an area termed Judgement Aggregation (JA) [11, 13]. Ontologies are widely used in Knowledge Representation to provide principled descriptions of agents' knowledge, by presenting a clear formalisation of their conceptualisations. The meaning of the concepts is then represented by means of a number of axioms, which may be written in a variety of logical systems of varying expressivity [1, 9]. With the exception of [15], the typical approaches to JA are usually applied to propositional logics, modal logics, or even more general logics, but they do not touch the problem of the possibly heterogeneous definitions of concepts used by the agents to formalise their individual conceptualisation. Understanding what is the meaning of a concept for a community of agents and deciding how to elect a common conceptualisation out of possibly conflicting ones is an interesting open problem that has several applications, for instance, in the context of political applications of JA. In this setting, understanding what is the meaning of a concept for a community of agents is crucial for modelling electoral campaigning,

where parties try to maximise their electorate by appealing to widely shareable world views. In the context of ontology aggregation, we may think of each ontology as a vote submitted by a voter trying to ‘elect’ a collective ontology that adequately and fairly represents their conceptualisations. JA then provides the formal means to assess the suitable aggregation procedures for a given aggregation scenario, by defining a number of properties that aggregators may or may not satisfy. However, many results in JA show that a significant number of important aggregation procedures, e.g., the majority rule, fail to preserve the consistency of the individual inputs [13, 15]. This means that, although we assume that all ontologies that agents submit for aggregation are consistent, the outcome of the aggregation may not be. A number of strategies to circumvent inconsistency have been pursued in JA, for instance, abandoning well-known aggregators in favour of aggregators that indeed preserve consistency, or restricting the set of propositions about which the agents cast their vote to those for which consistency can be ensured.

In this paper, we propose a novel approach. We present a well-known justified aggregation procedure that is actually used in real collective decision problems, viz. absolute majority rule, and we propose a computational viable methodology based on *axiom weakening* to repair its possibly inconsistent outcomes. The idea of axiom weakening is to generalise or specialise possibly conflicting concepts with concepts that are, in some sense, as close as possible to the original ones, but do not yield an inconsistency. Notice that, in order to generalise or specialise a concept in an informative manner, we need to rely on a certain amount of background information about the concepts, which shall be encoded in what we term in this paper the *reference* ontology. In case we have no information about the concepts at issue, the only way of weaken an axiom to a less constraining one is to replace it with a trivial (i.e. tautological) axiom, that in fact imposes no constraint at all on the agent’s conceptualisation.

Preventing inconsistencies by appealing to ‘general’ concepts, which may then be prone to agreement although they have not been voted on by any individual, has been suggested and legitimated in the literature on social choice and deliberation [5, 12, 14]. This is an important issue, and it also relates to the distinction between fine vs. coarse integration of ontologies. In the case of a *coarse* integration, the ontology to be constructed will always contain some of the formulas included in the individual ontologies; in the *fine* integration, new formulas shall be constructed. The approach in [15] provides an example of coarse integration. In this paper, we are after a viable definition of fine integration.

To summarise, the contributions of this paper are as follows. We consider possible conceptualisations of agents’ opinions represented by means of ontologies written in Description Logic (DL). In particular, we focus on the basic DL \mathcal{ALC} [1], which is a popular language for ontology development. Secondly, we use the methodology of social choice theory and JA of [15] to define a framework for ontology aggregation. Thirdly, we use refinement operators for concept generalisations and specialisations [4], and we apply them to repair the collective ontology by selecting adequate refinements of the axioms that caused the inconsistency.

2 Ontologies and Description Logics

We take an ontology to be a set of formulas in an appropriate logical language, describing our domain of interest. A significant widely used basic description logic is \mathcal{ALC} , which is the logic we shall be working with here. For full details about this logic, we refer the interested reader to [1]. The language of \mathcal{ALC} is based on an alphabet consisting of *atomic concepts names* N_C , and *role names* N_R . The set of *concept descriptions* is generated by the following grammar, where A represents atomic concepts and R role names:

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

We collect all \mathcal{ALC} concepts over N_C and N_R in $\mathcal{L}(\mathcal{ALC}, N_C, N_R)$. We assume a linear order $\prec_{\mathcal{ALC}}$ over \mathcal{ALC} formulas. We do not need to attach any particular meaning to it, but it will be helpful for coping with non-determinism and for tie-breaking. A *TBox* is a finite set of concept inclusions of the form $C \sqsubseteq D$ (where C and D are concept descriptions). It is used to store terminological knowledge regarding the relationships between concepts. An *ABox* is a finite set of formulas of the form $A(a)$ (“object a is an instance of concept A ”) and $R(a, b)$ (“objects a and b stand to each other in the R -relation”). It is used to store assertional knowledge regarding specific objects. The semantics of \mathcal{ALC} is defined in terms of *interpretations* $I = (\Delta^I, \cdot^I)$ that map each object name to an element of its domain Δ^I , each atomic concept to a subset of the domain, and each role name to a binary relation on the domain. The truth of a formula in such an interpretation is defined in the usual manner [1].

3 Aggregating Ontologies

Consider an arbitrary but fixed finite set Φ of \mathcal{ALC} TBox statements over this alphabet.¹ We call Φ the *agenda* and any set $O \subseteq \Phi$ an *ontology*. We denote the set of all those ontologies that are *consistent* by $\text{On}(\Phi)$. Let $\mathcal{N} = \{1, \dots, n\}$ be a finite set of *agents*. Each agent $i \in \mathcal{N}$ provides a consistent ontology $O_i \in \text{On}(\Phi)$. An *ontology profile* is a vector $\mathbf{O} = (O_1, \dots, O_n) \in \text{On}(\Phi)^{\mathcal{N}}$ of consistent ontologies, one for each agent. We write $N_{\varphi}^{\mathbf{O}} := \{i \in \mathcal{N} \mid \varphi \in O_i\}$ for the set of agents that include φ in their ontology under profile \mathbf{O} . Our object of study are *ontology aggregators*; that is, functions $F : \text{On}(\Phi)^{\mathcal{N}} \rightarrow 2^{\Phi}$ mapping any profile of consistent ontologies to an ontology.

Observe that, according to this definition, the ontology we obtain as the outcome of an aggregation process might be inconsistent. Ontology aggregators that are *consistent* would be very desirable in general. Unfortunately, they also suffer certain drawbacks. The *unanimous aggregator*, that accepts a formula if every individual does, is one of these. It indeed preserves consistency: if every

¹ The finite set of TBox formulas in Φ might be all TBox formulas of a certain maximum length or the union of all TBox formulas that a given population of agents choose to include in their TBoxes.

Table 1. A voting scenario

	LeftPolicy \sqsubseteq RaiseWages	LeftPolicy \sqsubseteq RaiseWelfare	RaiseWages \sqcap RaiseWelfare $\sqsubseteq \perp$
1	yes	yes	no
2	yes	no	yes
3	no	yes	yes
Maj.	yes	yes	yes




Fig. 1. The TBox agenda of the agents

ontology O_j is consistent, so is $F_{un}(\mathbf{O})$. However, if the individual ontologies are heterogeneous enough, the unanimous aggregator is likely to provide a very poor collective ontology. At the opposite side of the spectrum, we can define the *union* aggregator, that accepts any piece of information provided by at least one agent. In this case, the collective ontology is very likely to be inconsistent.

To balance the contributions of agents better than with the unanimous and the union aggregators, we can adopt the majority rule, which is widely applied in any political scenarios. In our setting, the majority rule is defined as follows: The absolute majority rule is the ontology aggregator F_m mapping any given profile $\mathbf{O} \in \text{On}(\Phi)^{\mathcal{N}}$ to the ontology

$$F_m(\mathbf{O}) := \{\varphi \in \Phi \mid \#N_{\varphi}^{\mathbf{O}} > \frac{n}{2}\} .$$

Under the absolute majority rule, a formula gets accepted if and only if more than half of the individual agents accept it. A simple generalisation of the majority rule provides the class of *quota* rules, where the threshold of $\frac{n}{2}$ is replaced by any threshold q . The majority rule, and more generally quota rules, return a consistent ontology only on very simple agendas [15].

4 Possibly Inconsistent Collective Ontologies

The following example shows that the absolute majority rule, which is widely used in practice, is not a consistent aggregator. Our example is a simple adaptation of the *doctrinal paradox* to the case of concept definitions [8, 13].

Consider three left-wing political leaders, i.e., three agents 1, 2, and 3, who must agree on what is a left policy in order to coordinate their campaigns. They vote on possible definitions of left-wing policy by casting their votes on the TBox agenda shown in Figure 1. Each individual ontology, in particular, formalises possible meanings that agents ascribe to what is a left-wing policy. Suppose that the agents vote as in Table 1.

Every individual set of axioms is consistent and the concept `LeftPolicy` is satisfiable in each of the individual ontologies. Agent 1, for instance, believes

LeftPolicy \sqsubseteq RaiseWages	LeftPolicy \sqsubseteq ReducelInequality
LeftPolicy \sqsubseteq RaiseWelfare	ReducelInequality \sqsubseteq Policy
RaiseWages \sqsubseteq ReducelInequality	LeftPolicy \sqsubseteq Policy
RaiseWelfare \sqsubseteq ReducelInequality	

Fig. 2. A reference ontology

that a left policy must raise both the wages and the levels of welfare, accordingly this agent believes that it is possible to promote the levels of both. Agent 2 believes that a left policy only has to raise wages, not the level of welfare, as they believe that it is not possible to do both. Agent 3 believes that what counts as a left policy is that it promotes the levels of welfare and that it is not possible to increase welfare and wages at the same time. Although all individual ontologies are consistent and the concept `LeftPolicy` is indeed satisfiable in each O_i , the ontology obtained by applying the absolute majority rule is not. The ontology $F_m(O_1, O_2, O_3)$ in this case coincides with the full agenda of Figure 1. By accepting both `LeftPolicy \sqsubseteq RaiseWages` and `LeftPolicy \sqsubseteq RaiseWelfare`, we infer `LeftPolicy \sqsubseteq RaiseWages \sqcap RaiseWelfare`, which together with `RaiseWages \sqcap RaiseWelfare \sqsubseteq \perp` makes the concept of `LeftPolicy` unsatisfiable. Moreover, as soon as we assume that there are indeed instances of left-wing policies, e.g., we add an ABox formula `LeftPolicy(a)`, for some constant `a`, to the ontology $F_m(O_1, O_2, O_3)$, then the collective ontology becomes inconsistent.

To repair the outcome of the majority rule, we assume that the agents share a certain amount of background information about the concepts at issue, that is they appeal a *reference* ontology (Figure 2). With respect to the reference ontology, there is more than one way of repairing the collective ontology. The concept `ReducelInequality` is a generalisation of `RaiseWelfare`, and of `RaiseWages`. So, one way of repairing is to weaken the axiom `LeftPolicy \sqsubseteq RaiseWages`, by replacing the concept `RaiseWages` with `ReducelInequality`. Symmetrically, one can weaken `LeftPolicy \sqsubseteq RaiseWelfare`, by generalising the concept `RaiseWelfare` also with `ReducelInequality`. In both cases, we obtain a consistent set of axioms. Another strategy is to weaken `RaiseWages \sqcap RaiseWelfare \sqsubseteq \perp` , for instance by specialising the concept `RaiseWages \sqcap RaiseWelfare` into `\perp` . However, the repaired ontology would contain the uninformative axiom `$\perp \sqsubseteq \perp$` . Although we effectively obtain a consistent ontology, a repair strategy would ideally avoid such an outcome when possible. Notice that if there is no viable reference ontology—that is, there is no information about the concepts that can be exploited for axiom weakening—replacing conflicting informative axioms with trivial or logical axioms is the only viable strategy.

5 Repairing Collective Ontologies

Our strategy for fixing the collective aggregated ontology relies on weakening the axioms present in a TBox w.r.t. an ontology. Weakening an axiom essentially amounts to refine its premise or its conclusion. In this setting, two types of

refinement operators exist: specialisation refinement operators and generalisation refinement operators [3, 10]. Given the quasi-ordered set $\langle \mathcal{L}(\mathcal{ALC}, N_c, N_R), \sqsubseteq \rangle$, a generalisation refinement operator is defined as follows:

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

Whereas a specialisation refinement operator is defined as follows:

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

A generalisation refinement operator takes a concept C as input and returns a set of descriptions that are more general than C , according to \mathcal{T} . A specialisation operator, instead, returns a set of descriptions that are more specific.

We note $\gamma_{\mathcal{T}}^*(C)$ and $\rho_{\mathcal{T}}^*(C)$ the iterated generalisation and the iterated specialisation of the concept C , respectively. As a minimal requirement, we assume that for every concept C , we have $\top \in \gamma_{\mathcal{T}}^*(C)$ and $\perp \in \rho_{\mathcal{T}}^*(C)$.

The following strategy is designed to use the novel generalisation and specialisation refinement operators of [4].

5.1 Axiom Weakening

Weakening an axiom $C \sqsubseteq D$ amounts to *enlarging* the set of interpretations that satisfy the axiom. This could be done in different ways: Either by substituting $C \sqsubseteq D$ with $C \sqsubseteq D'$, where D' is a more general concept than D (i.e., its interpretation is larger); or, by modifying the axiom $C \sqsubseteq D$ to $C' \sqsubseteq D$, where C' is a more specific concept than C ; or even by generalising and specialising simultaneously to obtain $C' \sqsubseteq D'$. Given an ontology O , we denote the set of concept names of O by N_O^C . We want to define a procedure to change axioms gradually by replacing them with less restrictive axioms. Recall that γ_O denotes the generalisation of a concept and ρ_O denotes its specialisation with respect to a given ontology O .

Definition 1 (Axiom weakening). *Given an axiom $C \sqsubseteq D$ of O , the set of weakenings of $C \sqsubseteq D$ in O , denoted by $g_O(C \sqsubseteq D)$ is the set of all axioms $C' \sqsubseteq D'$ such that*

$$C' \in \rho_O^*(C) \text{ and } D' \in \gamma_O^*(D) .$$

If the ontology O is consistent, the weakening of an axiom in O is always satisfied by a super set of the interpretations that satisfy the axiom. Let $I = (\Delta^I, \cdot^I)$ be an interpretation. By definition, the class of all entities that fulfill the axiom $C \sqsubseteq D$ is $(\Delta^I \setminus C^I) \cup D^I$. A weakening of $C \sqsubseteq D$ either specialises C , therefore restricting C^I , and accordingly extending $\Delta^I \setminus C^I$, or generalises D , therefore, extending D^I . Moreover, note that $\perp \sqsubseteq \top$ always belongs to $g_O(C \sqsubseteq D)$. We want to model how to repair any inconsistent set of axioms Y of \mathcal{ALC} , by appealing to a (consistent) reference ontology R . Notice that, even though it is not desirable, R can be dissociated from the axioms in the collective ontology.

Algorithm 1 Fixing ontologies through weakening.

Procedure FIX-ONTOLOGY(O, R) \triangleright O inconsistent ontology, R reference ontology

- 1: **while** O is inconsistent **do**
- 2: $\mathcal{Y} \leftarrow \text{MIS}(O)$ \triangleright find all minimally inconsistent subsets of O
- 3: **for** $Y \in \mathcal{Y}$ **do**
- 4: **choose** $\psi \in Y, \psi' \in g_R(\psi)$ with $Y \setminus \{\psi\} \cup \{\psi'\}$ consistent, $\lambda_O(\psi, \psi')$ minimal
- 5: $O \leftarrow (O \setminus \{\psi\}) \cup \{\psi'\}$
- 6: **return** O

Any inconsistent set of axioms Y can in principle be repaired by means of a sequence of weakenings of the axioms in Y with respect to R : in the worst case these axioms are weakened to become a tautology (e.g., $\perp \sqsubseteq \top$). However, we are interested in weakening axioms as little as possible to remain close to the original axioms. Since every axiom in $g_O(C \sqsubseteq D)$ is obtained by applying γ and ρ a finite number of times, we can define λ_O to be a *refinement distance* in an ontology O . Repair strategies can exploit this distance to guide the weakening of axioms that are the least stringent. Moreover, by trying to minimise the distance, we are trying to prevent non-informative (i.e., tautological) axioms to be selected as weakenings. In principle, we can also provide refined constraints on the generalisation and specialisation paths, e.g., by fixing an ordering of the concepts of the ontology O that determines which concepts are to be generalised or specialised first.

5.2 Fixing Collective Ontologies via Axiom Weakenings

When $F(\mathbf{O})$ is inconsistent, we can adopt the general strategy described in Algorithm 1 to repair it w.r.t. a given (fixed) reference ontology R .

The algorithm finds all the minimally inconsistent subsets Y_1, \dots, Y_n of $F(\mathbf{O})$ (e.g., using the methods from [2, 16]) and repairs each of them by weakening one of its axioms to regain consistency. From all the possible choices made to achieve this goal, the algorithm selects one that minimizes the distance λ_O (line 4). This process corrects all original causes for inconsistency, but may still produce an inconsistent ontology [7]. Hence, the process is repeated until a consistent ontology is found. Notice that the algorithm is non-deterministic, since it depends on the choice of the axiom to weaken, and the weakening selected. As such, it can also be seen as a strategy returning a non-singleton *set* of ontologies (i.e. the procedure is termed non-resolute in social choice [15]). To make it resolute, two policies for breaking ties are required. For both, we can capitalize on the linear order over formulas $\prec_{\mathcal{ALC}}$ introduced earlier. We can define a linear order $\prec_{\mathcal{ALC}}^x$ over axioms as follows: $C \sqsubseteq D \prec_{\mathcal{ALC}}^x E \sqsubseteq F$ iff $C \prec_{\mathcal{ALC}} E$, or $C = D$ and $D \prec_{\mathcal{ALC}} F$.

Now, with a reference ontology R and the linear order $\prec_{\mathcal{ALC}}$ fixed, the strategy returns an aggregation procedure $g_{R, \prec_{\mathcal{ALC}}}(F(\mathbf{O}))$: firstly, aggregate the individual ontologies in \mathbf{O} , then generalise the axioms in any possible inconsistent

set of $F(\mathbf{O})$ with respect to the reference ontology R , and obtain $g_{R, \prec_{ACC}}(F(\mathbf{O}))$. We leave a detailed presentation for future work.

5.3 An Application

We illustrate our strategy by discussing the example in Section 4. We have seen that the absolute majority rule returns an inconsistent collective ontologies. The inconsistent ontology $F_m(\mathbf{O})$ coincides, in this example, with the agenda of the agents (Figure 1).

To apply our strategy, we have firstly to select a reference ontology R . Suppose we choose the ontology in Figure 2. We exemplify how $g_R(F_m(\mathbf{O}))$ works by assuming in this case that it is non-resolute. We start by choosing an axiom in a minimally inconsistent subset of $F_m(\mathbf{O})$ that needs to be weakened. The whole collective ontology $F_m(\mathbf{O})$ is a minimally inconsistent set. Assume we start by $\text{LeftPolicy} \sqsubseteq \text{RaiseWages}$. Then, we have to select a concept to generalise or specialise. Suppose we select RaiseWages . Thus, to generalise the axiom $\text{LeftPolicy} \sqsubseteq \text{RaiseWages}$ we can replace it by $\text{LeftPolicy} \sqsubseteq \text{Reducelnequality}$, since Reducelnequality is the closest generalisation to RaiseWages in the reference ontology R . We obtain the new ontology, where the axiom $\text{LeftPolicy} \sqsubseteq \text{RaiseWages}$ has been replaced by the weaker $\text{LeftPolicy} \sqsubseteq \text{Reducelnequality}$.

Alternatively, we could have started by generalising $\text{RaiseWages} \sqcap \text{RaiseWelfare} \sqsubseteq \perp$. In this case, we have two choices, either we generalise \perp , or we specialise $\text{RaiseWages} \sqcap \text{RaiseWelfare}$. \perp can be generalised by any concept in the reference ontology. $\text{RaiseWages} \sqcap \text{RaiseWelfare}$ can here be specialised only by replacing it with \perp , obtaining therefore $\perp \sqsubseteq \perp$, which is a (non-informative) logical axiom. By replacing an axiom with a logical one, the effect on the final ontology is the same as removing the original axiom (a logical axiom does not restrict the models of the ontology). Thus, in this case, the repaired ontology contains $\text{LeftPolicy} \sqsubseteq \text{Reducelnequality}$ and $\text{LeftPolicy} \sqsubseteq \text{RaiseWelfare}$.

6 Conclusion and Future Work

We proposed a novel approach to repair an inconsistent ontology, obtained by aggregating the individual ontologies of a community of agents. Our approach is based on the notion of axiom weakening, which amounts to the generalisation or specialisation of concepts found in axioms that belong to minimally inconsistent subsets. Whilst we presented a viable solution, a more extensive evaluation is needed. Firstly, discussing good strategies for deciding a reference ontology is crucial for the present approach. Secondly, the study of the formal properties of the proposed algorithm and its computational complexity is required. Finally, it is important to extend the proposed approach to a large class of description logics and to a variety of important aggregation procedures. We leave these points for future work.

References

1. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, New York, NY, USA (2003)
2. Baader, F., Peñaloza, R.: Axiom pinpointing in general tableaux. *Journal of Logic and Computation* 20(1), 5–34 (February 2010), special Issue: Tableaux and Analytic Proof Methods
3. Confalonieri, R., Eppe, M., Schorlemmer, M., Kutz, O., Peñaloza, R., Plaza, E.: Upward Refinement Operators for Conceptual Blending in \mathcal{EL}^{++} . *Annals of Mathematics and Artificial Intelligence* (2016), doi:10.1007/s10472-016-9524-8
4. Confalonieri, R., Kutz, O., Troquard, N., Galliani, P., Porello, D., Peñaloza, R., Schorlemmer, M.: Coherence, Similarity, and Concept Generalisation. In: *Proc. of DL 2017*. CEUR (2017)
5. Dietrich, F., List, C.: Judgment aggregation by quota rules: Majority voting generalized. *Journal of Theoretical Politics* 19(4), 391–424 (2007)
6. Gaertner, W.: *A Primer in Social Choice Theory*. Oxford University Press (2006)
7. Horridge, M., Parsia, B., Sattler, U.: Justification masking in ontologies. In: *KR 2012*. AAAI Press (2012)
8. Kornhauser, L.A., Sager, L.G.: The one and the many: Adjudication in collegial courts. *California Law Review* 81(1), 1–59 (1993)
9. Kutz, O., Mossakowski, T., Lücke, D.: Carnap, Goguen, and the Hyperontologies: Logical Pluralism and Heterogeneous Structuring in Ontology Design. *Logica Universalis* 4(2), 255–333 (2010), special Issue on ‘Is Logic Universal?’
10. van der Laag, P.R., Nienhuys-Cheng, S.H.: Completeness and properness of refinement operators in inductive logic programming. *The Journal of Logic Programming* 34(3), 201 – 225 (1998)
11. List, C., Pettit, P.: Aggregating sets of judgments: An impossibility result. *Economics and Philosophy* 18(1), 89–110 (2002)
12. List, C.: Deliberation and agreement. In: *Deliberation, Participation and Democracy*, pp. 64–81. Springer (2007)
13. List, C., Puppe, C.: Judgment aggregation: A survey. In: *Handbook of Rational and Social Choice*. Oxford University Press (2009)
14. Ottonelli, V., Porello, D.: On the elusive notion of meta-agreement. *Politics, Philosophy and Economics* 12(1), 68–92 (2013)
15. Porello, D., Endriss, U.: Ontology merging as social choice: Judgment aggregation under the open world assumption. *Journal of Logic and Computation* 24(6), 1229–1249 (2014)
16. Schlobach, S., Cornet, R.: Non-standard reasoning services for the debugging of description logic terminologies. In: *Proc. IJCAI-03*. pp. 355–362. Morgan Kaufmann (2003)