Introduction to Judgement Aggregation

Ontology Aggregation

Daniele Porello

KRDB Research Center for Knowledge and Data
Free University of Bozen-Bolzano

daniele.porello@unibz.it

EASSS 2018

OVERVIEW

- Description Logics and Ontologies,
- Modelling ontology aggregation;
- Methods from JA;
- Voting procedures on ontologies;
- Balancing efficiency and fairness;
- Extensions: towards fine-grained aggregation.

This part is based on [PE14] and on [PTP$^+$18]

MOTIVATIONS

- We want to apply the analysis of the aggregation procedures provided by social choice theory and judgment aggregation to the problem of aggregating several ontologies.
- Different individual ontologies provide different and possibly contrasting information and we ask which ontology better represent the *group* information.
- Social choice theory shows that the notion of *group information* strongly depends on the aggregation procedure we use.
- Fairness conditions here can be interpreted as constraints of impartiality on agents or on propositions.
- E.g. anonymity assumes that we do not know which individual ontology is more reliable.

ONTOLOGIES

- Formal ontologies are increasingly used in a variety of domains in crucial applications of AI, as well as in Multiagent Systems, Conceptual Modelling, Database Design, NLP and Software Engineering.
- Ontologies are a way to express the information about a certain domain in a peculiar way: they intend to make the modelling choices and the assumptions of the modeller clear, justified, and sharable among the community of users.
- To achieve that, ontologies are ofter written in a logical language. A very well developed family of languages for ontologies is the family of *Description Logics* (DL).

- We use the Description Logic $\mathcal{ALC}$ as an important example, however this treatment can be adapted to other description logics.
- The language of $\mathcal{ALC}$ is based on an alphabet consisting of *atomic concepts*, *role names*, and *object names*.
- The set of *concept descriptions* is generated as follows (where A represents atomic concepts and R role names):

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

TBox and ABox

- A *TBox* is a finite set of formulas of the form $A \sqsubseteq C$ and $A \equiv C$ (where A is an atomic concept and C a concept description).
- An *ABox* is a finite set of formulas of the form $A(a)$ and $R(a, b)$.
- The semantics of $\mathcal{ALC}$ is given by *interpretations* that map each object name to an element of its domain, each atomic concept to a subset of the domain, and each role name to a binary relation on the domain.
- A set of (TBox and ABox) formulas is *satisfiable* if there exists an interpretation in which they are all true.

Example of an ontology

A very simple ontology that specifies (part of) the meaning of what a left policy is.

$\text{LeftPolicy} \sqsubseteq \text{RaiseWages}$
$\text{LeftPolicy} \sqsubseteq \text{RaiseWelfare}$
$\text{RaiseWages} \sqsubseteq \text{ReduceInequality}$
$\text{RaiseWelfare} \sqsubseteq \text{ReduceInequality}$
$\text{LeftPolicy} \sqsubseteq \text{ReduceInequality}$
$\text{ReduceInequality} \sqsubseteq \text{Policy}$
$\text{LeftPolicy} \sqsubseteq \text{Policy}$

- Let us fix a finite set $\mathcal{L}$ of $\mathcal{ALC}$ formulas over a given alphabet that includes all the possible ABox formulas.
- We call $\mathcal{L}$ the *agenda* and any set $O \subseteq \mathcal{L}$ an *ontology*.
- O can be divided into a TBox O^T and an ABox O^A.
- Let $\mathrm{On}(\mathcal{L})$ be the set of all those ontologies that are *satisfiable*.
- The *closure* of a set of formulas $\Phi \subseteq \mathcal{L}$ is given by $\mathrm{Cl}(\Phi) := \{\varphi \in \mathcal{L} \mid \Phi \models \varphi\}$

- Let $\mathcal{N} = \{1, \dots, n\}$ be a finite set of *agents*. Each agent $i \in \mathcal{N}$ provides a satisfiable ontology $O_i \in \mathrm{On}(\mathcal{L})$.
- An *ontology profile* $\boldsymbol{O} = (O_1, \dots, O_n) \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$ is a vector of such ontologies, one for each agent.
- We write $\mathcal{N}_{\varphi}^{\boldsymbol{O}} := \{i \in \mathcal{N} \mid \varphi \in O_i\}$ for the set of agents including φ in their ontology under profile $\boldsymbol{O}$.
- An *ontology aggregator* is a function $F : \mathrm{On}(\mathcal{L})^{\mathcal{N}} \to 2^{\mathcal{L}}$ mapping any profile of satisfiable ontologies to an ontology.

Examples

- An ontology aggregator is F with $F(\boldsymbol{O}) := O_1 \cup \cdots \cup O_n$, which returns the union of the individuals ontologies. The ontology obtained may not be satisfiable.
- An ontology aggregator is F with $F(\boldsymbol{O}) := O_1 \cap \cdots \cap O_n$, which returns the intersection of the individuals ontologies. The ontology obtained may be very poor.

An Example: the majority rule

- The *majority rule* (accept a formula if and only if a majority of the agents do) can lead to unsatisfiable outcomes, as we can easily simulate the *discursive dilemmas* from *Judgment Aggregation*.
- Suppose three agents share a common TBox with two formulas:

$$C_3 \equiv C_1 \sqcap C_2 \qquad C_4 \sqsubseteq \neg C_3$$

- Furthermore, suppose the three ABoxes are as follows:

	$C_1(a)$	$C_2(a)$	$C_3(a)$	$C_4(a)$
Agent 1	yes	yes	yes	no
Agent 2	yes	no	no	yes
Agent 3	no	yes	no	yes
Majority	yes	yes	no	yes

- Individual ontologies are satisfiable but the collective one is not.
- E.g. The original example of discursive dilemma can be viewed as a case in which O^T expresses the legal doctrine (Kornhauser and Sager, 1986).

The standard doctrinal paradox is slightly simpler than our example above.

The TBox only consists of the formula $C_3 \equiv C_1 \sqcap C_2$ and the individual ABoxes are as follows:

	$C_1(a)$	$C_2(a)$	$C_3(a)$
Agent 1	yes	yes	yes
Agent 2	yes	no	no
Agent 3	no	yes	no
Majority	yes	yes	no

- Observe that this is not a paradox in the same (strong) sense as the earlier example.

The TBox only consists of the formula $C_3 \equiv C_1 \sqcap C_2$ and the individual ABoxes are as follows:

	$C_1(a)$	$C_2(a)$	$C_3(a)$
Agent 1	yes	yes	yes
Agent 2	yes	no	no
Agent 3	no	yes	no
Majority	yes	yes	no

- As before, the group accepts that a belongs to both C_1 and C_2. Given $C_3 \equiv C_1 \sqcap C_2$, we can now *infer* that the group also accepts $C_3(a)$ to be true, even if this fact is not explicitly recorded in the collective ontology.
- That is, the only "paradox" we encounter here is that, even though the three individual ontologies are deductively closed (with respect to the set of four formulas under consideration here), this is not the case for the collective ontology.

Basic features of this modelling

- We restrict for now to “coarse” merging: the ontology to be constructed will be a list of some of the formulas included in the individual ontologies.
- We shall deal with “fine” merging later, where we might also want to construct entirely new formulas from those provided by the individuals.
- *Open vs Closed World Assumption*: we need to adapt the standard JA framework by dropping *completeness* (a agent accepts A or accepts $\neg A$): here it would entail that an agent cannot express her lack of knowledge concerning the application of both A and $\neg A$ to a certain object.

- *Syntactic vs Semantic axioms (explicit vs implicit knowledge)*. We define “syntactic” axioms, they relate to the formulas that occur *explicitly* in the ontologies of individual agents or in the collective ontology.
- We will contrast this with “semantic” axioms that make reference to the formulas that can be inferred *implicitly* from those ontologies.

Syntactic Axioms (standard)

The usual social choice theory and judgment aggregation axioms can be restated as follows:

- **Unanimity:** F is called *unanimous* if $O_1 \cap \cdots \cap O_n \subseteq F(\boldsymbol{O})$ for every profile $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.
- **Anonymity:** F is called *anonymous* if for any profile $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$ and any permutation $\pi : \mathcal{N} \to \mathcal{N}$ we have that $F(O_1, \ldots, O_n) = F(O_{\pi(1)}, \ldots, O_{\pi(n)})$.
- **Independence:** F is called *independent* if for any $\varphi \in \mathcal{L}$ and profiles $\boldsymbol{O}, \boldsymbol{O}' \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$, we have that $\varphi \in O_i \Leftrightarrow \varphi \in O'_i$ for all $i \in \mathcal{N}$ implies $\varphi \in F(\boldsymbol{O}) \Leftrightarrow \varphi \in F(\boldsymbol{O}')$.
- **Monotonicity:** F is called *monotonic* if for any $i \in \mathcal{N}$, $\varphi \in \mathcal{L}$, and $\boldsymbol{O}, \boldsymbol{O}' \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$ with $O_j{=}O'_j$ for all $j \neq i$, we have that $\varphi \in O'_i \setminus O_i$ and $\varphi \in F(\boldsymbol{O})$ imply $\varphi \in F(\boldsymbol{O}')$.

SYNTACTIC AXIOMS (SPECIFIC)

We introduce the following specific axioms we might want for our ontology aggregators:

- **Groundedness:** F is called *grounded* if $F(\boldsymbol{O}) \subseteq O_1 \cup \cdots \cup O_n$ for every profile $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.
- **Exhaustiveness:** F is called *exhaustive* if there exists no satisfiable set $\Phi \subseteq O_1 \cup \cdots \cup O_n$ with $F(\boldsymbol{O}) \subset \Phi$ for any profile $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.
- **Group Closure:** F is called *group-closed* if there exists no set $\Phi \subseteq O_1 \cup \cdots \cup O_n$ with $F(\boldsymbol{O}) \models \Phi$ and $F(\boldsymbol{O}) \subset \Phi$ for any profile $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.

An important axiom is JA *neutrality*, which, intuitively, requires all formulas to be treated symmetrically. In fact, there are a number of possible interpretations of this notion, including these:

- **Neutrality:** F is called *neutral* if for any $\varphi, \psi \in \mathcal{L}$ and $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$ we have that $\varphi \in O_i \Leftrightarrow \psi \in O_i$ for all $i \in \mathcal{N}$ implies $\varphi \in F(\boldsymbol{O}) \Leftrightarrow \psi \in F(\boldsymbol{O})$.
- **Acceptance-Rejection Neutrality:** F is called *acceptance-rejection neutral* if for any $\varphi \in \mathcal{L}$ and $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$ we have that $\varphi \in O_i \Leftrightarrow \psi \notin O_i$ for all $i \in \mathcal{N}$ implies $\varphi \in F(\boldsymbol{O}) \Leftrightarrow \psi \notin F(\boldsymbol{O})$.

- In JA, when abandoning completeness, acceptance-rejection neutrality is usually assumed.
- Our objection to this axiom is stated as follows:

Proposition

Any ontology aggregator that satisfies acceptance-rejection neutrality violates exhaustiveness.

We propose the following definition of semantic axioms that model a form of implicit knowledge.

- **Semantic Unanimity:** F is called *semantically unanimous* if $\mathrm{Cl}(O_1) \cap \cdots \cap \mathrm{Cl}(O_n) \subseteq \mathrm{Cl}(F(\boldsymbol{O}))$ for every profile $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.
- **Semantic Groundedness:** F is called *semantically grounded* if $\mathrm{Cl}(F(\boldsymbol{O})) \subseteq \mathrm{Cl}(O_1) \cup \cdots \cup \mathrm{Cl}(O_n)$ for every $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.
- **Semantic Exhaustiveness:** F is called *semantically exhaustive* if there exists no satisfiable set $\Phi \subseteq \mathrm{Cl}(O_1) \cup \cdots \cup \mathrm{Cl}(O_n)$ with $\mathrm{Cl}(F(\boldsymbol{O})) \subset \Phi$ for any $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.

Implicit vs explicit knowledge

- For most aggregators, syntactic and semantic axioms do not entail each other. E.g. The majority rule is syntactically unanimous but not semantically unanimous.
- Intuitively, semantic unanimity is the (much) stronger property. This intuition can be confirmed for the following aggregators:

Proposition

Any satisfiable and exhaustive ontology aggregator that is semantically unanimous is unanimous.

- Analogous results hold for the other axioms.
- In the next slides, we shall discuss some concrete voting procedures.

- Let $\mathrm{M}(\boldsymbol{O}) = \{\varphi \in \mathcal{L} \mid |\mathcal{N}_{\varphi}^{\boldsymbol{O}}| > \frac{n}{2}\}$ for all $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.
- In JA, the majority rule provides consistent outcomes on agendas that satisfy *the median property* (List and Puppe, 2009).
- We can refine this result for Description Logics as follows.

The majority Rule

- We say that $\mathcal{L}$ satisfies the $\mathcal{T}$*-median property* if and only if for every set of ABox formulas $X \subseteq \mathcal{L}^A$ such that $\mathcal{T} \cup X$ is unsatisfiable there exists a set $Y \subseteq X$ with cardinality at most 2 such such $\mathcal{T} \cup Y$ is also unsatisfiable.
- We obtain the following characterisation:

Proposition

The majority rule will return a satisfiable ontology for any profile with a common TBox $\mathcal{T}$ if and only if the agenda $\mathcal{L}$ satisfies the $\mathcal{T}$-median property.

- We can generalise the idea underlying the majority rule and accept a formula for the collective ontology whenever the number of agents who do so meet a certain quota.
- This gives rise to the family of quota rules:

Quota rules

Let $q \in [0, 1]$. The quota rule with quota q is the ontology aggregator F_q with $F_q(\boldsymbol{O}) = \{\varphi \in \mathcal{L} \mid |\mathcal{N}^{\boldsymbol{O}}_{\varphi}| \geqslant q \cdot n\}$ for all $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$.

Quota rules

- The majority rule violates semantic unanimity.
- In fact, any quota rule does, unless we lower the quota so far as to obtain the trivial union aggregator:

Proposition

A quota rule with quota q for n agents is semantically unanimous if and only if $q \leqslant \frac{1}{n}$.

Two procedures designed for ontologies

- We see now two procedures design for ontology aggregation.
- That is, they are designed to balance the preservation of consistency with the satisfaction of the other axioms:
- Support-based procedures
- Two-stages procedures

- We order the formulas in terms of the number of agents supporting them.
- We introduce a *priority rule* $\gg$ mapping each profile $\boldsymbol{O}$ to a strict linear order $\gg_{\boldsymbol{O}}$ on $\mathcal{L}$ such that $\varphi \gg_{\boldsymbol{O}} \psi$ implies $|\mathcal{N}_{\varphi}^{\boldsymbol{O}}| \geqslant |\mathcal{N}_{\psi}^{\boldsymbol{O}}|$ for all $\varphi, \psi \in \mathcal{L}$.

- Support-based procedures:

  Given a priority rule $\gg$, the support-based procedure with $\gg$ is the ontology aggregator $\mathrm{SBP}_{\gg}$ mapping any profile $\boldsymbol{O} \in \mathrm{On}(\mathcal{L})^{\mathcal{N}}$ to $\mathrm{SBP}_{\gg}(\boldsymbol{O}) := \Phi$ for the unique set $\Phi \subseteq \mathcal{L}$ for which $\varphi \in \Phi$ if and only if

  (*i*) $\mathcal{N}_{\varphi}^{\boldsymbol{O}} \neq \emptyset$ and

  (*ii*) $\{\psi \in \Phi \mid \psi \gg_{\boldsymbol{O}} \varphi\} \cup \{\varphi\}$ is satisfiable.

- The SBP clearly satisfies the axioms of anonymity, monotonicity, groundedness (due to condition (*i*)), and exhaustiveness (due to condition (*ii*)).
- Neutrality is violated by virtue of having to fix a priority rule $\gg$

Two-stages procedures

- We may give priority to the terminology or to the assertions.
- From JA, we have the *premise-based procedure*: individuals vote on the premises by majority and then draw the conclusions, and the *conclusion-based procedure*: individuals draw their own conclusions and then votes on them by majority.
- E.g. assertion-based procedures stress the information coming from the ABox.

- Assertion-based procedure:

An (irresolute) assertion-based procedure maps each profile $\boldsymbol{O}$ to the set of ontologies obtained as follows:

1. Choose an aggregator F_A restricted to ABox formulas, and let $F_A(\boldsymbol{O})$ be the outcome.
2. Then the TBox is defined as follows:

$$F_T(\boldsymbol{O}) = \mathrm{argmin}_{O \in \mathrm{On}(\mathcal{L})} \sum_{i \in \mathcal{N}} d(F_A(\boldsymbol{O}) \cup O_i^T, O)$$

(where d is a distance)

SUMMING UP

- We presented a model inspired by judgment aggregation to define aggregation of individual ontologies for the case of the coarse merging.
- Different aggregation procedures define different notion of *group information* and the axiomatic analysis spells out the properties of such notions.
- We presented our analysis distinguishing between implicit and explicit knowledge by the distinction between semantic and syntactic axioms.
- We introduced and discussed voting rules and properties of aggregators with the aim of balancing the satisfiability and fairness.

TOWARDS FINE MERGING

- We approach the fine merging of ontologies as follows.
- We fix an aggregation procedure such as the majority rule.
- We know that it may return inconsistent ontologies.
- We propose a strategy to repair the collective ontology by building up formulas that are close enough to the individuals' original formulas. That is,
- We let agents votes on the formulas of the ontology by relying on *Judgment Aggregation*.
- We introduce a strategy for repairing the possibly inconsistent ontology by means of *axiom weakening*.

Suppose three agents submit their views about what a left policy is and agree to elect a collective opinion by means of the majority rule.

TABLE: 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

In this case, the concept of left policy is not satisfiable, although it is for each agent.

We propose to repair possibly inconsistent collective ontology by means of *Axiom Weakening*.

The weakening of an axiom $C \sqsubseteq D$ is given by *specialising* the concept C or by *generalising* the concept D.

Specialisations and generalisations make sense only with respect to a *reference* ontology, e.g.

LeftPolicy $\sqsubseteq$ RaiseWages	LeftPolicy $\sqsubseteq$ ReduceInequality
LeftPolicy $\sqsubseteq$ RaiseWelfare	ReduceInequality $\sqsubseteq$ Policy
RaiseWages $\sqsubseteq$ ReduceInequality	LeftPolicy $\sqsubseteq$ Policy
RaiseWelfare $\sqsubseteq$ ReduceInequality	

FIGURE: A reference ontology

Algorithm for repairing collective ontologies

- Given a profile of voting on agents' ontologies, an aggregation procedure (the majority rule), and a reference ontology,
- the following algorithm describes how to fix a possibly inconsistent collective ontology:

Algorithm 1 Fixing ontologies through weakening.

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

while O is inconsistent **do**
  $\mathcal{Y} \leftarrow \text{MIS}(O)$ $\triangleright$ find all minimally inconsistent subsets of O
  for $Y \in \mathcal{Y}$ **do**
    choose $\psi \in Y, \psi' \in g_{\mathrm{R}}(\psi)$ **with** $Y \setminus \{\psi\} \cup \{\psi'\}$ consistent, $\lambda_O(\psi, \psi')$ minimal
    $O \leftarrow (O \setminus \{\psi\}) \cup \{\psi'\}$
return O

- The procedure ensures that the output is a consistent ontology.
- Moreover, the algorithm selects candidates for axiom weakening that are as close as possible to the original axiom.
- The complexity of the algorithm depends on the complexity of logic used to formalise the conceptualisation.

References I

Daniele Porello and Ulle Endriss.
Ontology merging as social choice: judgment aggregation under the open world assumption.
J. Log. Comput., 24(6):1229–1249, 2014.

Daniele Porello, Nicolas Troquard, Rafael Penaloza, Roberto Confalonieri, Pietro Galliani, and Oliver Kutz.
Two approaches to ontology aggregation based on axiomweakening.
In *IJCAI, to appear*, 2018.