Measuring Dissimilarity between Judgment Sets

Marija Slavkovik and Thomas Ågotnes

University of Bergen, PB. 7802, 5020 Bergen, Norway
{marija.slavkovik,thomas.agotnes}@infomedia.uib.no

Abstract. Distances and scores are widely used to measure (dis)similarity between objects of information such as preferences, belief sets, judgment sets, etc. Typically, measures are directly imported from information theory or topology, with little consideration for adequacy in the context of comparing logically related information. We propose a set of desirable properties for measures used to aggregate (logically related) *judgments*, and show which of the measures used for this purpose satisfy them.

1 Introduction

The aggregation of sets of logically related information is a problem that occurs in at least four disciplines with intersecting areas of interest with multiagent systems and artificial intelligence: judgement aggregation [3], belief revision [11], social choice [3] and abstract argumentation [2]. Many approaches to aggregating sets of information are based on comparing the information sets and measuring how similar they are. Furthermore, studies of complexity of various forms of manipulation, see for example [9,8,1], extensively rely on similarity comparisons. For an effective comparison, information sets cannot be viewed as inseparable units that are either entirely the same or entirely different from each other.

Simply counting the number of units on which collections of information differ, namely using the Hamming distance [10], is adequate only when these units are logically independent [12,7,2]. Although the Hamming distance is extensively used to aggregate them [12,8,11,16,17], in general, neither sets of beliefs, arguments labelings, votes, preferences, nor sets of judgments contain exclusively logically unrelated elements. How should logically related information sets be compared, namely, which properties should be satisfied by the (dis)similarity measures used?

We focus on sets of judgments and their comparison for the purpose of aggregation. Since judgment aggregation has known relations with belief merging [17], preference aggregation [15], voting [13], and aggregation of labelings within abstract argumentation [4,2], dissimilarity measures for sets of judgments can also be applied in these disciplines. In this paper we focus on three tasks: a) identifying a set of properties common to all dissimilarity measures used in the literature; b) defining desirable properties of dissimilarity measures that are apt for comparing sets containing logically related information; c) showing that there exist measures that satisfy both sets of properties.

In Section 2 we give the necessary preliminaries, while in Section 3 we first discuss related work and then we attend to tasks a) and b). In Section 4 we concern ourselves with task c). In Section 5 we draw conclusions and outline directions for future work.

E. Fermé and J. Leite (Eds.): JELIA 2014, LNAI 8761, pp. 609–617, 2014.

2 Preliminaries

Judgment aggregation problems are typically represented using a set $\mathcal{L}$ of well-formed propositional logic formulas, including $\top$ (tautology) and $\bot$ (contradiction). An *issue* is a pair of formulas $\{\varphi, \neg\varphi\} \subset \mathcal{L}$ where φ is neither a tautology nor a contradiction. For simplicity, we often abuse notation and write only the positive formula when we discuss issues. An *agenda* $\mathcal{A}$ is a finite set of issues, $\mathcal{A} = \{\varphi_1, \neg\varphi_1, \ldots, \varphi_m, \neg\varphi_m\}$. A sub-agenda $Y \subset \mathcal{A}$ is a subset of issues from $\mathcal{A}$, *e.g.*, $Y = \{p, \neg p\}$ is a sub-agenda for $\mathcal{A} = \{p, \neg p, q, \neg q\}$. A *judgment* on an issue $\{\neg\varphi, \varphi\} \in \mathcal{A}$ is one of φ or $\neg\varphi$.

A *judgment set* J is a subset of $\mathcal{A}$, *complete* iff for each $\{\neg\varphi, \varphi\} \in \mathcal{A}$ either $\varphi \in J$ or $\neg\varphi \in J$, and incomplete otherwise. A judgment set J is consistent iff it is a consistent set of formulas. For a given agenda $\mathcal{A}$, the set of all consistent nonempty judgment sets is $\mathcal{D}(\mathcal{A})$, while $\mathbb{D}(\mathcal{A}) \subset \mathcal{D}(\mathcal{A})$ is the set of all consistent and complete judgement sets. Judgment sets $J_1, J_2 \in \mathbb{D}(\mathcal{A})$ are *complementary* when for every $\varphi \in \mathcal{A}$, $\varphi \in J_1$ iff $\neg\varphi \in J_2$.

A *profile* $P \subset \mathbb{D}(\mathcal{A})^n$ is a tuple $P = \langle J_1, \ldots, J_n \rangle$ of judgment sets for agents $1, \ldots, n$. An (irresolute) judgment aggregation rule is a correspondence $F : \mathbb{D}(\mathcal{A})^n \to 2^{\mathbb{D}(\mathcal{A})} \backslash \emptyset$. Namely, a judgment aggregation rule associates a set of complete and consistent judgment sets for an agenda $\mathcal{A}$, called *collective judgment sets*, to a profile of judgments for the same $\mathcal{A}$. Two very basic properties for judgment aggregation rules are *unanimity*, for every $P \in \mathbb{D}(\mathcal{A})^n$ s.t. $P = \langle J, J, \ldots, J \rangle$, $F(P) = \{J\}$, and *anonymity* for every permutation σ of P and every $P \in \mathbb{D}(\mathcal{A})^n$, $F(P) = F(\sigma(P))$.

Two existing classes of judgment aggregation rules make use of similarity or dissimilarity measures, respectively, selecting collective judgment sets: *scoring rules* [6], which here we refer to using F_s, and *distance-based rules* [7,8,16,17], here denoted with $F_{d,\odot}$. We give the respective definitions for these two classes of rules using our notation. For any $\mathcal{A}$, $P \in \mathbb{D}(\mathcal{A})^n$, $P = \langle J_1, \ldots, J_n \rangle$:

$$F_s(P) = \underset{J \in \mathbb{D}(\mathcal{A})}{\text{argmax}} \sum_{i \in [1,n]} \sum_{\varphi \in J \cap J_i} s(J_i, \varphi), \quad F_{d,\odot}(P) = \underset{J \in \mathbb{D}(\mathcal{A})}{\text{argmin}} \odot(d(J_1, J), \ldots, d(J_n, J)).$$

In the definition of F_s, s is a *scoring function* of type $s : \mathbb{D}(\mathcal{A}) \times \mathcal{A} \to \mathbb{R}$. Scoring functions assign judgment set-dependent scores for each possible judgment that can be cast. One set J is more similar to a given J_i than another j set J', if the sum of scores of the judgments in J, according to the J_i-respective scoring, is higher than that sum of the scores assigned to judgments in J'.

In the definition of $F_{d,\odot}$, two functions are used to determine the collective judgment sets. The function $\odot : \mathbb{R}^n \to \mathbb{R}$ is an n-ary aggregation function that assigns a unique value to an n-ary vector of values. *E.g.*, the function Σ in definition of F_s is an n-ary aggregation function; other examples include max, min, and Π. The second function used in the definition of $F_{d,\odot}$ is the dissimilarity function d, defined for every $\mathcal{A} \subset \mathcal{L}$, as $d : \mathbb{D}(\mathcal{A}) \times \mathbb{D}(\mathcal{A}) \to \mathbb{R}$, which assigns a higher number the more different J_i is from J.

Dissimilarity measures are typically defined as functions that take as arguments two sequences of equal length: for every $\mathcal{A} \in \mathcal{L}$, $d : \mathcal{A}^m \times \mathcal{A}^m \to \mathbb{R}$. Observe that d is being defined for all agendas; we stipulate that d compares judgment sets that are complete and consistent judgment sets for the *same* agenda.

Although scores and dissimilarity measures appear to be different functions, we show that for every scoring function s there exists a corresponding dissimilarity function d_s, and as a result, we show that the F_s rules are a special case of the $F_{d,\odot}$ rules.

Definition 1. *For every agenda $\mathcal{A} \subset \mathcal{L}$ and scoring function $s : \mathbb{D}(\mathcal{A}) \times \mathcal{A} \rightarrow \mathbb{R}$, for every $J, J' \in \mathbb{D}(\mathcal{A})$ we define $S(J, J') = \sum_{\varphi \in J \cap J'} s(J, \varphi)$. A dissimilarity measure $d_s : \mathbb{D}(\mathcal{A}) \times \mathbb{D}(\mathcal{A}) \rightarrow \mathbb{R}$ corresponds to a scoring function $s : \mathbb{D}(\mathcal{A}) \times \mathcal{A} \rightarrow \mathbb{R}$ iff, for every $J, J', J'' \in \mathbb{D}(\mathcal{A})$, $d_s(J, J') > d_s(J, J'')$ iff $S(J, J') < S(J, J'')$ and $d_s(J, J') = d_s(J, J'')$ iff $S(J, J') = S(J, J'')$.*

Proposition 1. *For every finite agenda $\mathcal{A} \subset \mathcal{L}$, for every scoring function s there exists a corresponding dissimilarity measure d_s.*

Proof. Observe that $S(J, J')$ gives a J-dependent score that measures *similarity* to a judgement set J' by summing the scores of all the judgments in $J \cap J'$. If instead we sum the scores for the judgments in $J \backslash J'$, we obtain a dissimilarity measure. For each s we can define a d_s as: $d_s(J, J') = \sum_{\varphi \in J \backslash J'}^{s} (J, \varphi)$. To show that d_s is corresponding to s, it is sufficient to observe that the judgment sets are finite, hence the maximal value that s can obtain for a given J is $S(J, J)$ and that $d_s(J, J') = S(J, J) - S(J, J')$.

It is now easy to show that Proposition 2 holds. The proof of the proposition consists in observing that, for each F_s, and $P \in \mathbb{D}(\mathcal{A})^n$, $F_s(P) = F_{d_s,\sum}(P)$.

Proposition 2. *For every scoring rule Fs there exists a rule $F_{d,\odot}$ such that for every $P \in \mathbb{D}(\mathcal{A})^n$, $F_s(P) = F_{d,\odot}(P)$.*

Due to Propositions 1 and 2 it is sufficient to consider dissimilarity measures when looking for desirable properties for both similarity and dissimilarity measures in judgment aggregation.

3 Measuring Dissimilarity between Judgment Sets

We first discuss related work on measuring dissimilarity between judgment sets and then what can be considered general requirements for such measures, before discussing how sensitivity to logic relations can be expressed as their property. The general requirements we outline are weak properties that most of the dissimilarity measures in use should satisfy. Although most are obvious, we do need to have them to show that they can be consistent with properties of sensitivity to the logic relations.

Given that it is obvious that a different (dis)similarity measure, even for the same aggregator $\odot$, yields a judgment aggregation rule with different properties, it is surprising that not even very general requirements on the d_s functions induced from scoring rules are required or discussed, meaning that anything goes as long as the scoring rule fits the signature $\mathbb{D}(\mathcal{A}) \times \mathbb{D}(\mathcal{A}) \rightarrow \mathbb{R}$. The situation appears to be better when the dissimilarity functions d are used directly in the distance-based rules, whereupon it is usually required that the function d is a pseudo-distance[1] or a distance [5]. It can be observed that the judgment aggregation rule $F_{d,\odot}$ works with any dissimilarity measure. The (pseudo-)distance requirements have been imported from the literature of belief merging [11], from where the $F_{d,\odot}$ rules originate, however the necessity of these requirements in judgment aggregation has never been justified.

The requirement of triangular inequality is easy to drop, since it is not required in belief merging either and no justification for it has been offered in judgment aggregation. The need for the symmetry property is not so clear. Five out of the six scoring functions presented in [6] give rise to d_s that is not symmetric, as it is simple to verify by looking at the examples in [6]. The exception is the simple scoring rule which corresponds to the Hamming distance. When one is measuring a distance, symmetry is necessary, but dissimilarity can be meaningful without symmetry as the scoring rules demonstrate. We therefore consider symmetry desirability to be context-dependent.

In [19] it has been identified that the non negativity and identity of indiscernibles properties of a pseudo-distance are necessary for the $F_{d,\odot}$ rule to satisfy unanimity. It was also established that the anonymity of $F_{d,\odot}$ does not depend on the function d, but on whether the aggregation function $\odot$ is commutative or not. As a consequence, all scoring rules will satisfy anonymity, however only those for which d_s is non-negative will satisfy unanimity.

Let us consider the sensitivity to logic relations for a similarity measure at this point before considering some more properties of measures from the literature. In [7] it is argued that if for an agenda $\mathcal{A}$ two agents cannot disagree on one issue without disagreeing on another issue, then these two disagreements in their judgment sets should not be counted as two disagreements, as the Hamming distance does, but only as one. In [7], this requirement is captured by Axiom 5. A $J_2 \in \mathbb{D}(\mathcal{A})$ is in-between $J_1, J_3 \in \mathbb{D}(\mathcal{A})$, $J_1 \neq J_2 \neq J_3$, if $J_1 \backslash J_2 \subset J_1 \backslash J_3$. Axiom 5 states that if $J_1, J_2 \in \mathbb{D}(\mathcal{A})$ are such that there exists no in-between $J \in \mathbb{D}(\mathcal{A})$, then $d(J_1, J_2) = 1$. Implicitly, in [7] it is advocated that similarity measures should consider the logical relations among issues and not only count disagreements. Here, we make explicit the logic relation sensitivity hinted on in [12,7,2] by defining it as a set of properties for measures.

An agenda $\mathcal{A}$ cannot contain tautologies or contradictions, but we may add arbitrarily many issues to it that are logically equivalent to existing agenda issues. Consider for example two hiring committee members that do not agree that "a candidate is good for the open position" (φ). Adding the issue "the candidate is not bad for the open position" ($\neg\neg\varphi$) to the agenda should not increase the quantity of disagreement between the positions of the agents. Regardless of how many times an issue is cloned in the agenda, the disagreement quantity should not increase, namely, a measure that is sensitive of the logic relations among issues should be *insensitive to agenda clones*.

A property called *disagreement monotonicity*, is considered in [2], for aggregation of labelings in argumentation, but applicable to judgment sets as well. A dissimilarity

[1] A pseudo-distance f is a function that (for every x, y, z in its codomain) satisfies $f(x, y) \geqslant 0$ (nonnegativity), $f(x, y) = 0$ iff $x = y$ (identity of indiscernibles) and $f(x, y) = f(y, x)$(symmetry). A distance additionally satisfies $f(x, y) + f(y, z) \geqslant f(x, z)$ (triangular inequality).

measure is *disagreement monotonic* if for all $J_1, J_2, J_3 \in \mathbb{D}(\mathcal{A})$ s.t. J_2 is in-between J_1 and J_3, $d(J_1, J_2) < d(J_1, J_3)$. Requiring that the amount of disagreement is strictly increasing with the number of judgments on which two judgment sets differ is not compatible with the insensitivity to agenda clones requirement. Indeed when a clone is added to the agenda, the number of issues on which two judgment sets disagree will increase, but the amount of disagreement will not. Therefore we propose weak disagreement monotonicity, requiring that the amount of disagreement between judgment sets does not decrease with an increase in the number of disagreeing issues in the sets.

In addition to Axiom 5, another requirement was considered in [7], the Axiom 4. Axiom 4 states that, if J_2 is in-between J_1 and J_3, then $d(J_1, J_3) = d(J_1, J_2) + d(J_2, J_3)$. Axiom 4 is strictly stronger than the disagreement monotonicity requirement of [2], namely, in-between together with non-negativity implies disagreement monotonicity, but the implication in the other direction does not hold. In judgment aggregation, disagreement monotonicity is easy to justify, however requiring Axiom 5, that for every judgment set there exists a judgment set at distance 1, is arbitrary outside of the scope of the [7], where this property is needed to characterise the introduced distance. In addition, forcing Axiom 5 limits the domain of the dissimilarity measure to natural numbers. We therefore consider only disagreement monotonicity to be a basic requirement.

The insensitivity to clones requirements can be made stronger. Assume that two committee members agree "not to hire any more academic staff until 2015"($\neg\varphi'$) but do not agree on "increasing the number of administrative staff" (ψ). The committee members have no need to vote regarding the issue of "increase the number of academic staff and hire John for a lecturer position" ($\varphi' \wedge \varphi$") as agreeing on $\neg\varphi'$ also means an agreement on $\neg(\varphi' \wedge \varphi'')$. Removing an implied judgment from the judgment sets should not change the amount of disagreement between them. This property we call *insensitivity to consequents*.

Definition 2. *Let $d : \mathbb{D}(\mathcal{A}) \times \mathbb{D}(\mathcal{A}) \to \mathbb{R}$ be a dissimilarity measure defined for every $\mathcal{A} \subset \mathcal{L}$. The function d is an adequate dissimilarity measure for judgment aggregation if, for every $J_1, J_2, J_3 \in \mathbb{D}(\mathcal{A})$, properties (p1)-(p3) hold, and an adequate dissimilarity measure for logically related sets of formulas if properties (p1)-(p5) hold. Desirability of (p6) is context-dependent.*

Nonnegativity:	$d(J_1, J_2) \geqslant 0$.	**(p1)**
Identity of indiscernibles:	$d(J_1, J_2) = 0$ *iff* $J_1 = J_2$.	**(p2)**
Agreement monotonicity:	*For every $J_1 \neq J_2 \neq J_3$, if $J_3 \backslash J_2 \subset J_3 \backslash J_1$, then $d(J_2, J_3) \leqslant d(J_1, J_3)$.*	**(p3)**
Insensitivity to clones:	*If there are $\psi, \varphi \in \mathcal{A}$ s.t. $\psi \equiv \varphi$, then for every $J_1, J_2 \in \mathbb{D}(\mathcal{A})$ it holds that $d(J_1, J_2) = d(J_1 \backslash \{\varphi, \neg\varphi\}, J_2 \backslash \{\varphi, \neg\varphi\})$.*	**(p4)**
Insensitivity to consequents:	*If there exist $S \subset J_1$ and $\varphi \in J_1$ s.t. $J_1 \in \mathbb{D}(\mathcal{A})$, $S \vdash \varphi$, and there is no $S' \subset S$ s.t. $S' \vdash \varphi$, then for every $J_2 \in \mathbb{D}(\mathcal{A})$, it holds that $d(J_1, J_2) = d(J \backslash \{\varphi, \neg\varphi\}, J' \backslash \{\varphi, \neg\varphi\})$.*	**(p5)**
Symmetry:	$d(J_1, J_2) = d(J_2, J_1)$.	**(p6)**

Clearly when comparing $d(J_1, J_2)$ and $d(J_1 \backslash \{\varphi, \neg\varphi\}, J_2 \backslash \{\varphi, \neg\varphi\})$, the sets J_1 and $J_1 \backslash \{\varphi, \neg\varphi\}$ are not complete and consistent judgment sets for the same agenda: $J_1 \in \mathbb{D}(\mathcal{A})$ and $J_1 \backslash \{\varphi, \neg\varphi\} \in \mathbb{D}(\mathcal{A} \backslash \{\varphi, \neg\varphi\})$. However, these two sets do not need to be from the same agenda, we are only comparing two rational numbers. Also observe that while (p4) is a property that refers to issues (that are pairs of judgments), the (p5) property refers to judgments. Lastly, we mention two obvious relationships.

Proposition 3. *If d satisfies insensitivity to consequents, then it also satisfies insensitivity to agenda clones. The reverse does not hold. If d satisfies Axiom 4 as defined in [7], then it also satisfies agreement monotonicity.*

4 Compliance

In this section we analyse existing (dis)similarity measures from the literature and identify which satisfy the desirable properties we outlined. We demonstrate that there does exist a measure that satisfies all (p1)-(p6). We found the following measures: the Hamming and drastic distances d_H and d_D, see *e.g.*, [16], defined respectively as $d_H(J_1, J_2) = |J_2 \backslash J_1| (= |J_1 \backslash J_2|)$ and $d_D(J_1, J_2)$ is 0 iff $J_2 \backslash J_1 = \varnothing$ and 1 otherwise; the five scoring rules from [6]: reversal scoring, entailment scoring, disjoint entailment scoring, minimal entailment scoring and irreducible entailment scoring, giving rise to d_{rv}, d_{et}, d_{ds}, d_{md}, and d_{ir} respectively; the critical subsets distance d_{CS} from [2] and the minimal prime implicant measures introduced in [18], definitions follow. We omit here the definitions of the five scoring functions from [6], and resulting distances, due to space restrictions. These are fairly simple to retrieve from the original work [6], and the proofs involving them are straightforward. We give the definitions of the rest.

We repeat the concept of prime implicants of judgment sets introduced in [18]. Consider an agenda $\mathcal{A}$ and $J \in \mathcal{D}(\mathcal{A})$ with a subset $I \subseteq J$. The set I is an implicant of J if for every $\varphi \in J$ it holds that $I \vdash \varphi$. I is a *prime J-implicant* if I is an implicant of J and there is no $I' \subset I$ s.t. $I' \vdash \varphi$ for every $\varphi \in J$. Intuitively, the prime J-implicant is a set of judgments which when known, all the judgments in J can be known as well. We assume that $PI(\varnothing) = \varnothing$. The minimal prime J-implicant is defined as that prime J-implicant that, among all of the prime J-implicants, has the minimal cardinality. We denote the set of prime implicants for J with $PI(J)$ and the minimal prime J-implicant with $MPI(J)$: $MPI(J) = \underset{I \in PI(J)}{\text{argmin}} |I|$. The Minimal sum prime implicant measure is defined as $d_{msp}(J_1, J_2) = |MPI(J_2 \backslash J_1)| + |MPI(J_1 \backslash J_2)|$.

A *critical set* for an agenda $\mathcal{A}$ is a sub-agenda $Y \subseteq \mathcal{A}$, s.t. for every $J \in \mathbb{D}(\mathcal{A})$, $J \cap Y \vdash \bigwedge J$ and there exists no sub-agenda $Y' \subset Y$ s.t. $J \cap Y' \vdash \bigwedge J$. *E.g.*, for the agenda $\mathcal{A} = \{p, \neg p, q, \neg q, p \wedge q), \neg(p \wedge q)\}$, there is one critical set $Y = \{p, \neg p, q, \neg q, \}$. For a critical subset $Y \subset \mathcal{A}$, $d_{CS}(J_1, J_2) = d_H(J_1 \cap Y, J_2 \cap Y)$.

We also consider the Duddy-Piggins distance d_{DP} from [7]. Let $G = \langle V, E \rangle$ be a graph in which the set of vertices is $V = \mathbb{D}(\mathcal{A})$ and there exists an edge, in the set of edges $E \subseteq \mathbb{D}(\mathcal{A}) \times \mathbb{D}(\mathcal{A})$, between two judgement sets J_1 and J_2 iff there exists no $J_3 \in \mathbb{D}(\mathcal{A})$ in-between J_1 and J_2. For any $J_1, J_2 \in \mathbb{D}(\mathcal{A})$, $d_{DP}(J_1, J_2)$ is the minimal number of edges between J_1 and J_2.

Proposition 4. *The compliance of the dissimilarity measures d_D, d_{msp}, d_{DP}, d_{CS}, d_H, d_{rv}, d_{et}, d_{ds}, d_{md}, and d_{ir} with properties (p1)-(p6) is as in Table 1.*

Proof. We prove only the non-obvious entries.

Table 1. Compliance of existing distances with proposed properties

	(p1)	(p2)	(p3)	(p4)	(p5)	(p6)
d_D, d_{msp}	✓	✓	✓	✓	✓	✓
d_{DP}, d_{CS}	✓	✓	✓	✓	×	✓
d_H	✓	✓	✓	×	×	✓
$d_{rv}, d_{et}, d_{ds}, d_{md}, d_{ir}$	✓	✓	✓	×	×	×

Measures d_{CS} and d_{msp} are agreement monotonic.
Consider $J_1, J_2, J_3 \in \mathbb{D}(\mathcal{A})$, $J_1 \neq J_2 \neq J_3$ s.t. J_2 is in-between J_1 and J_3. For d_{CS}, observe that if J_2 is in-between J_1 and J_3, then also $J_2 \cap Y$ is in-between $J_1 \cap Y$ and $J_3 \cap Y$. Consequently d_{CS} behaves as Hamming distance and is as such agreement monotonic.

For d_{msp} we show that when J_2 is in-between J_1 and J_3 (and $J_1 \neq J_2 \neq J_3$) it holds that $MPI(J_3 \backslash J_2) \subset MPI(J_3 \backslash J_1)$. Observe that when J_2 is in-between J_1 and J_3 we can represent the sets J_1, J_2, J_3 as a union of mutually exclusive sets $S_1, S_2, S_3, \overline{S_2}, \overline{S_3}$: $J_1 = S_1 \cup \overline{S_2} \cup \overline{S_3}$, $J_2 = S_1 \cup S_2 \cup \overline{S_3}$, $J_3 = S_1 \cup S_2 \cup S_3$, where $S_1 = J_1 \cap J_2 \cap J_3$, $S_2 = J_3 \cap J_2$, $\overline{S_3} = \{\neg\varphi \mid \varphi \in S_3\}$, and $\overline{S_3} = J_2 \cap J_1$, $S_2 = \{\neg\varphi \mid \varphi \in S_2\}$. We have that $J_3 \backslash J_1 = S_2 \cup S_3$ and $J_3 \backslash J_2 = S_3$. Since S_2 is consistent with both S_3 and $\overline{S_3}$, clearly neither of these subsets implies the other. Therefore the $MPI(J_3 \backslash J_1) = MPI(J_3 \backslash J_2) \cup MPI(S_2)$ and $MPI(J_3 \backslash J_2) \subset MPI(J_3 \backslash J_1)$. Observe further that $MPI(J_1 \backslash J_3) = MPI(J_2 \backslash J_3) \cup MPI(\overline{S_2})$. Consequently $MPI(J_2 \backslash J_3) \subset MPI(J_1 \backslash J_3)$.

We now have that $d_{msp}(J_2, J_3) = |MPI(S_3)| + |MPI(\overline{S_3})|$ and $d_{msp}(J_1, J_3) = |MPI(S_3)| + |MPI(S_2)| + |MPI(\overline{S_3})| + |MPI(\overline{S_2})|$.

From this proof it also follows that d_{msp} satisfies the Axiom 4 of [7]: observe that $MPI(J_3 \backslash J_2) = MPI(J_3 \backslash J_2) \cup MPI(S_2) = MPI(J_3 \backslash J_2) \cup MPI(J_2 \backslash J_1)$.

The measures d_{DP}, d_{CS} and d_{msp} are insensitive to clones.
Let $\mathcal{A}' = \mathcal{A} \cup \{\varphi, \neg\varphi\}$ where $\varphi \equiv \varphi_i$ for some $\varphi_i \in \mathcal{A}$. Observe that there exists an isomorphism between $\mathbb{D}(\mathcal{A})$ and $\mathbb{D}(\mathcal{A}')$. For every $J \in \mathbb{D}(\mathcal{A})$ there exists exactly one $J' \in \mathbb{D}(\mathcal{A}')$, furthermore $J' = J \cup \{\varphi\}$ iff $\varphi \in J$ and $J' = J \cup \{\neg\varphi\}$ iff $\neg\varphi \in J$. Consequently, there is an edge between J_1 and J_2 in the graph G built for $\mathcal{A}$ iff there is an edge between the corresponding J'_1 and J'_2 in the graph G' built for $\mathcal{A}'$. Hence for every $J_1, J_2 \in \mathbb{D}(\mathcal{A})$ and corresponding J'_1, J'_2, $d_{DP}(J_1, J_2) = d_{DP}(J'_1, J'_2)$.

For d_{CS}, observe that a logically equivalent issue (to some agenda issue) would never be part of the critical set. The insensitivity to clones of d_{msp} is obtained as a consequence of its insensitivity to consequents; proof follows.

The d_{DP} and the d_{CS} are not insensitive to consequents.
As a counter example for d_{DP}, consider an $\mathcal{A} = \{p, \neg p, p \wedge q, \neg(p \wedge q)\}$ and $J_1, J_2 \in \mathbb{D}(\mathcal{A})$ s.t. $J_1 = \{p, p \wedge q\}$ and $J_2 = \{\neg p, \neg(p \wedge q)\}$. Observe that $\{\neg p\} \vdash \neg(p \wedge q)$. We have that $d_{DP}(J_1, J_2) = 2$, because $J_3 = \{p, \neg(p \wedge q)\}$ is in between J_1 and J_2 ($\mathbb{D}(\mathcal{A}) = \{J_1, J_2, J_3\}$). However $d_{DP}(J_1 \backslash \{p \wedge q\}, J_2 \backslash \{\neg(p \wedge q)\}) = 1$, since $J_3 \backslash \{p \wedge q\} = J_2 \backslash \{\neg(p \wedge q)\}$ and these two points in the graph for $\mathcal{A}$ collapse into one point in the graph for $\mathcal{A} \backslash \{p \wedge q, \neg(p \wedge q)\}$.

As a counter example for d_{CS} consider $\mathcal{A} = \{p, \neg p, p \vee q, \neg(p \vee q)\}$. Observe that the critical set $Y = \mathcal{A}$. Consider $J_1 = \{\neg p, \neg(p \vee q)\}$ and $J_2 = \{p, p \vee q\}$. We have $\{p\} \vdash p \vee q$. We have that $d_{CS}(J_1, J_2) = 2$ and $d_{CS}(J_1 \backslash \{p \vee q, \neg(p \vee q)\}, J_2 \backslash \{p \vee q, \neg(p \vee q)\}) = 1$.

The d_{msp} is insensitive to consequents.
Let $\mathcal{A}$ be s.t. $S \subset \mathcal{A}$, $\varphi \in \mathcal{A}$ and $S \vdash \varphi$. Let $\mathcal{A}' = \mathcal{A} \backslash \{\varphi, \neg\varphi\}$ be a sub-agenda of $\mathcal{A}$. Consider a $J_1, J_2 \in \mathbb{D}(\mathcal{A})$ with $S \subset J_2$ and corresponding $J_1' = J_1 \cap \mathcal{A}'$ and $J_2' = J_2 \cap \mathcal{A}'$. There are two possible cases: (a)$\varphi \in J_1$ and (b) $\varphi \notin J_1$. If (a) is the case, then $J_2 \backslash J_1 = J_2' \backslash J_1'$ (also $J_1 \backslash J_2 = J_1' \backslash J_2'$) and thus $d_{msp}(J_1, J_2) = d_{msp}(J_1', J_2')$. If (b) is the case, then $S \not\subset J_1$. We have that $\varphi \notin MPI(J_2 \backslash J_1)$, thus $MPI(J_2 \backslash J_1) = MPI(J_2' \backslash J_1')$. If there exists an $MPI(J_1 \backslash J_2)$ s.t. $\neg\varphi \notin MPI(J_1 \backslash J_2)$, then $|MPI(J_1 \backslash J_2)| = |MPI(J_1' \backslash J_2')|$. If for all $MPI(J_1 \backslash J_2)$, $\neg\varphi \in MPI(J_1 \backslash J_2)$, then there will be exactly one element of S not in $MPI(J_1 \backslash J_2)$ because S minimally entails φ, thus $|MPI(J_1 \backslash J_2)| = |MPI(J_1' \backslash J_2')|$.

5 Summary

Functions are used to quantify the (dis)similarity between different types of information collections, in *e.g.*, belief merging, judgment aggregation, preference aggregation and abstract argumentation. This is the first work to consider the assembly of desirable properties for dissimilarity measures in judgment aggregation, as well as defining properties that identify measures sensitive to the logic relations among judgments.

It is straightforward to show that neither of the scoring distances d_H, d_{rv}, d_{et}, d_{ds}, d_{md}, and d_{ir} are insensitive to clones. Any scoring function s can be transformed into a clone insensitive version s_{ci} using $s_{ci}(J_i, \varphi) = \frac{s(J,\varphi)}{|S_\varphi|}$, where $S_\varphi = \{\psi \mid \psi \in \mathcal{A}, \psi \equiv \varphi\}$. It remains to be explored whether the clone insensitive scores still generalise known voting rules, as studied in [6,13].

Interesting future work arises from looking into how the "logic relation sensitivity" properties of a measure interact with the properties of a judgment aggregation operator that uses them. The first obvious property to investigate is the property of majority-preservation. A profile is majority-consistent when the majoritarian set is consistent. The majoritarian set is the judgment set in which each judgment is supported by a majority in the profile. A judgment aggregation rule is majority-preserving when it selects as a unique collective judgment set the majoritarian set whenever the profile is majority-consistent. Can a distance-based rule using a logic relation sensitive dissimilarity measure be majority-preserving? It can be shown that neither the Duddy-Piggins distance, nor the d_{msp} combined with the sum yield a majority-preserving rule. We conjecture that this result scales to all sensitive measures and aggregators. Other possible dependencies between the judgment aggregation rule and the constituting d, such as agenda separability [14] are likely to exist.

Acknowledgment. We are grateful to Jérôme Lang for his valuable comments on various versions of this paper.

References

1. Baumeister, D., Erdélyi, G., Erdélyi, O.J., Rothe, J.: Computational aspects of manipulation and control in judgment aggregation. In: Perny, P., Pirlot, M., Tsoukiàs, A. (eds.) ADT 2013. LNCS, vol. 8176, pp. 71–85. Springer, Heidelberg (2013)
2. Booth, R., Caminada, M., Podlaszewski, M., Rahwan, I.: Quantifying disagreement in argument-based reasoning. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Richland, SC, pp. 493–500. International Foundation for Autonomous Agents and Multiagent Systems (2012)
3. Brandt, F., Conitzer, V., Endriss, U.: Computational social choice. In: Weiss, G. (ed.) Multiagent Systems, pp. 213–283. MIT Press (2013)
4. Caminada, M., Pigozzi, G.: On judgment aggregation in abstract argumentation. Autonomous Agents and Multi-Agent Systems 22(1), 64–102 (2011)
5. Deza, M.M., Deza, E.: Encyclopedia of Distances. Springer (2009)
6. Dietrich, F.: Scoring rules for judgment aggregation. Social Choice and Welfare, 1–39 (2013)
7. Duddy, C., Piggins, A.: A measure of distance between judgment sets. Social Choice and Welfare 39(4), 855–867 (2012)
8. Endriss, U., Grandi, U., Porello, D.: Complexity of judgment aggregation. Journal Artificial Intelligence Research (JAIR) 45, 481–514 (2012)
9. Everaere, P., Konieczny, S., Marquis, P.: The strategy-proofness landscape of merging. Journal of Artificial Intelligence Research (JAIR) 28, 49–105 (2007)
10. Hamming, R.W.: Error detecting and error correcting codes. Bell System Technical Journal 29(2), 147–160 (1950)
11. Konieczny, S., Pino-Pérez, R.: Logic based merging. Journal of Philosophical Logic 40(2), 239–270 (2011)
12. Lafage, C., Lang, J.: Propositional distances and compact preference representation. European Journal of Operational Research 160(3), 741–761 (2005)
13. Lang, J., Slavkovik, M.: Judgment aggregation rules and voting rules. In: Perny, P., Pirlot, M., Tsoukiàs, A. (eds.) ADT 2013. LNCS, vol. 8176, pp. 230–243. Springer, Heidelberg (2013)
14. Lang, J., Slavkovik, M., Vesic, S.: A weakening of independence in judgment aggregation: agenda separability (extended abstract). In: Schaub, T. (ed.) Proceedings of the 21st European Conference on Artificial Intelligence, page forthcoming (2014)
15. List, C., Puppe, C.: Judgment aggregation: A survey. In: Anand, P., Puppe, C., Pattanaik, P. (eds.) Oxford Handbook of Rational and Social Choice. Oxford (2009)
16. Miller, M.K., Osherson, D.: Methods for distance-based judgment aggregation. Social Choice and Welfare 32(4), 575–601 (2009)
17. Pigozzi, G.: Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation. Synthese 152(2), 285–298 (2006)
18. Slavkovik, M., Agotnes, T.: A judgment set similarity measure based on prime implicants. In: Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems, AAMAS, page forthcoming (2013)
19. Slavkovik, M., Jamroga, W.: Distance-based rules for weighted judgment aggregation (extended abstract). In: Proceedings of AAMAS, pp. 1405–1406 (2012)