Using Distances for Aggregation in Abstract Argumentation

Richard Booth [1] Mikołaj Podlaszewski [1]

CSC / SnT, University of Luxembourg
{richard.booth,mikolaj.podlaszewski}@uni.lu

Abstract

We continue recent work on the problem of aggregating labellings of an argumentation framework by adapting the *distance-based* framework of Miller and Osherson from binary judgment aggregation to the argumentation setting. To instantiate the framework we employ some notions of *labelling-distance* recently introduced by Booth et al., in the process generalising and extending some of the latters' results. We introduce some new postulates for distance methods based on the concept of *qualitative distance*, and check their validity.

1 Introduction

The aggregation of conflicting opinions into a collective one is fundamental in multi-agent systems. Individuals presented with the same set of conflicting arguments might take different rational positions. The problem of aggregation in abstract argumentation has been explored in a number of recent papers [1, 4, 10] which employ techniques from *judgment aggregation* (JA) [8] to the problem of aggregating 3-valued *argument labellings*. These works have shown that, as with classical JA, it is not possible to define general aggregation operators that satisfy a number of seemingly mild constraints while ensuring collective rationality of the outcome.

One way of getting around this problem in JA is to use one of the *distance-based solution methods* that were studied by Miller and Osherson (hereafter MO)[9] within the framework of binary judgment aggregation. As the name suggests these depend on a provided notion of *distance measure* between binary judgment sets. In this paper we show how this option can also be used in the argumentation setting. We first modify the MO framework for our purposes, before bringing in some notions of distance between argument labellings that have been defined in [2]. We thus illustrate the usefulness of the distance measures defined in that paper.

Along the way we generalise and extend some of the results in [2]. For example some of the MO aggregation methods require a distance to be defined between any two arbitrary labellings of an argumentation framework, whereas the most interesting distance measures of [2], such as the *issue-based* distance, are defined only between *complete* labellings. We thus extend the definition of these distances to apply to arbitrary labellings. We also look at two new postulates for distance measures between labellings, based on the *Normality* and *Increasing* properties introduced by MO in the JA setting, and confirm to what extent our distance methods conform with these postulates.

As well as a distance-measure, the MO methods also require up-front specification of an *initial* aggregation operator, which intuitively can be thought of as a *gold standard* operator that satisfies a number of basic postulates, without always yielding collectively rational results. In fact the MO methods can be viewed as offering recipes to *repair* the result of this operator in the cases when it does not give a collectively rational outcome. In the argumentation setting, Caminada and Pigozzi already suggested another way to carry out such a repair in these cases, using what they called the *down-admissible* and *up-complete* procedures [4]. Our work thus provides an alternative solution.

[1] Supported by the National Research Fund, Luxembourg (DYNGBaT and LAAMIcomp projects)

2 Preliminaries

We use the familiar setting of *abstract argumentation* [5]. We start by assuming a countably infinite set U of argument names, from which all possible argumentation frameworks are built. We restrict ourselves to finite argumentation frameworks.

Definition 1. *An* argumentation framework *(AF for short)* $\mathscr{A} = (Args, \rightharpoonup)$ *is a pair consisting of a finite set* $Args \subseteq U$ *of arguments and an attack relation* $\rightharpoonup \subseteq Args \times Args$*. We also use* $Args_{\mathscr{A}}$ *and* $\rightharpoonup_{\mathscr{A}}$ *to denote the arguments and attack relation of a given AF* $\mathscr{A}$*.*

Definition 2. *Let* $\mathscr{A} = (Args, \rightharpoonup)$ *be an AF. An* $\mathscr{A}$-labelling *is a function* $L : Args \to \{\texttt{in}, \texttt{out}, \texttt{undec}\}$*. The set of all* $\mathscr{A}$*-labellings is denoted by* $Labs(\mathscr{A})$*. Given* $A \subseteq Args$ *we denote by* $L[A]$ *the restriction of* L *to* A*.*

For notational purposes it is useful to define a unary "negation" operator on the set of labels by $\neg\texttt{in} = \texttt{out}$, $\neg\texttt{out} = \texttt{in}$ and $\neg\texttt{undec} = \texttt{undec}$.

Of course a *rational* labelling should somehow respect the attack relation. This is embodied in the following definition:

Definition 3. *Let* $\mathscr{A}$ *be an AF and* $L \in Labs(\mathscr{A})$*. For any argument* $a \in Args_{\mathscr{A}}$ *we say:*

- *a is legally* `in` *if* $L(a) = \texttt{in}$ *and* $L(b) = \texttt{out}$ *for all* $b \in Args_{\mathscr{A}}$ *s.t.* $b \rightharpoonup_{\mathscr{A}} a$*,*
- *a is legally* `out` *if* $L(a) = \texttt{out}$ *and* $L(b) = \texttt{in}$ *for some* $b \in Args_{\mathscr{A}}$ *s.t.* $b \rightharpoonup_{\mathscr{A}} a$*,*
- *a is legally* `undec` *if* $L(a) = \texttt{undec}$ *and there is no* $b \in Args_{\mathscr{A}}$ *s.t.* $b \rightharpoonup_{\mathscr{A}} a$ *and* $L(b) = \texttt{in}$*, and there exists* $c \in Args_{\mathscr{A}}$ *s.t.* $c \rightharpoonup_{\mathscr{A}} a$ *and* $L(c) = \texttt{undec}$*.*

L is complete *iff it has no illegally* `in` *and no illegally* `out` *and no illegally* `undec` *arguments. We denote the set of complete* $\mathscr{A}$*-labellings by* $Comp(\mathscr{A})$*.*

In the rest of this paper we identify rational $\mathscr{A}$-labellings with complete $\mathscr{A}$-labellings. This is because they form the basis for other semantics such as preferred, stable, semi-stable, etc (see [3]). This choice is also in line with other works on aggregation in argumentation [1, 4, 10]. However we will also make use later of the notion of *admissible* $\mathscr{A}$-labelling, which is an $\mathscr{A}$-labelling containing no illegally `in` or illegally `out` arguments (but possibly containing illegally `undec` arguments).

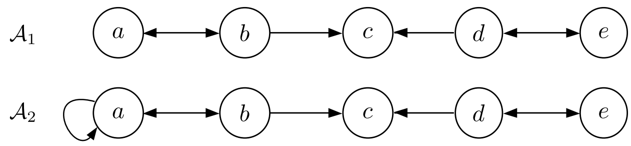

	l_1	l_2	l_3	l_4	l_5	l_6	l_7	l_8	l_9	n_1	n_2	n_3
a	o	o	o	u	u	u	i	i	i	u	i	u
b	i	i	i	u	u	u	o	o	o	u	o	u
c	o	o	o	o	u	u	o	i	u	o	u	i
d	i	o	u	i	o	u	i	o	u	u	o	o
e	o	i	u	o	i	u	o	i	u	u	i	i

Figure 1: Two aggregation frameworks with all their complete labellings ($\mathscr{A}_1$: l_1-l_9, $\mathscr{A}_2$: l_1-l_6), and three other labellings (n_1-n_3).

Example 1. *AFs can be visualised as directed graphs with nodes and edges representing arguments and attacks respectively. Throughout the paper we will use the following two running example AFs* $\mathscr{A}_1$ *and* $\mathscr{A}_2$*, both consisting of the five arguments* $\{a, b, c, d, e\}$*. We represent labellings by a string of letters* `i`*,* `u` *and* `o` *corresponding to* `in`*,* `undec` *and* `out` *respectively. AF* $\mathscr{A}_1$ *has 9 complete labellings* l_1*-*l_9*. In* $\mathscr{A}_2$ *labellings* $l_7 - l_9$ *are no longer complete because argument a is illegally-*`in` *due to the additional attack. For both AFs the labellings* n_1 *and* n_3 *are not admissible (and hence not complete). For* $\mathscr{A}_1$ *the labelling* n_2 *is admissible but not complete. For* $\mathscr{A}_2$ *it is even not admissible because argument a is illegally-*`in`*.*

Now we come to aggregation. We assume a set of agents $Ag = \{1, \ldots, n\}$ (with $n \geq 2$) is fixed.

Definition 4. *Let* $\mathscr{A}$ *be an AF. An* $\mathscr{A}$-profile *is any n-tuple of* $\mathscr{A}$*-labellings* $\mathbf{L} = (L_1, \ldots, L_n)$*. If every* L_i *is a complete* $\mathscr{A}$*-labelling then we call* $\mathbf{L}$ *a* complete $\mathscr{A}$*-profile.*

What we seek is a way to construct aggregation operators that, given any $\mathscr{A}$-profile $\mathbf{L}$ as input, return a set of $\mathscr{A}$-labellings $F_{\mathscr{A}}(\mathbf{L})$. Note that an aggregation operator is always defined in a context of some specific AF $\mathscr{A}$. More generally we are interested in an aggregation *method* that given any context AF $\mathscr{A}$ will return in a systematic way an aggregation operator for $\mathscr{A}$.

Definition 5. *Let $\mathscr{A}$ be an AF. An* (irresolute) aggregation operator *(for $\mathscr{A}$) is a function $F_{\mathscr{A}}$ that assigns, to each $\mathscr{A}$-profile $\mathbf{L}$ a set $F_{\mathscr{A}}(\mathbf{L}) \subseteq Labs(\mathscr{A})$. An* aggregation method *is a mapping that, given any context AF $\mathscr{A}$, returns an aggregation operator $F_{\mathscr{A}}$ for $\mathscr{A}$.*

From now on we will drop the *irresolute* and just say *aggregation operator*. In previous works on aggregation in argumentation the output is usually taken to be a single labelling, but we relax that here. When important, we say a *resolute* aggregation operator for an operator which returns always a singleton set. Also note for each $\mathscr{A}$, $F_{\mathscr{A}}$ is defined for *all* $\mathscr{A}$-profiles (not necessarily just the complete ones), and that the output of $F_{\mathscr{A}}(\mathbf{L})$ is allowed to be *any* subset of $\mathscr{A}$-labellings. Ideally, of course, we would like the output to consist only of *complete* $\mathscr{A}$-labellings, i.e., we want the following to hold:

Collective Rationality For all $\mathscr{A}$ and $\mathscr{A}$-profiles $\mathbf{L}$, $F_{\mathscr{A}}(\mathbf{L}) \subseteq Comp(\mathscr{A})$

3 Miller and Osherson aggregation methods.

Miller and Osherson (hereafter MO) [9] described a framework for using *distance measures* to define aggregation methods in binary judgment aggregation. In that setting agents evaluate a set of logical propositions called an *agenda* by providing a *judgement set* which is an assignment of either True or False to each proposition in the agenda. Their framework requires specification of two things:

1. An *initial* ***resolute*** *aggregation method M* that is able to give collectively rational answers in simple cases. In their case they only considered the proposition-wise majority method, but in principle any *M* can be used. The intuition is that if the outcome produced by *M* happens to be collectively rational then there is no need to choose a different outcome.

2. A measure of *distance* between any two judgment sets of any given agenda. This measure is assumed to be a *metric*.

Definition 6. *A* metric *on a set X is a function $d : X \times X \to \mathbb{R}$ such that for all $x,y,z \in X$,* (i) *if $d(x,y) = 0$ then $x = y$,* (ii) $d(x,x) = 0$, (iii) $d(x,y) = d(y,x)$, (iv) $d(x,z) \leq d(x,y) + d(y,z)$. *A function d satisfying* (ii)-(iv) *but not* (i) *is called a* pseudometric.

In our argumentation setting, the roles of the agenda and the judgment set are filled by the AF $\mathscr{A}$ and $\mathscr{A}$-labelling respectively. We will make use of a *distance method*, which for each AF $\mathscr{A}$ returns a distance measure over $\mathscr{A}$-labellings.

Definition 7. *Let $\mathscr{A}$ be an AF. A* distance measure $d_{\mathscr{A}}$ for $\mathscr{A}$ *is a metric on the set $Labs(\mathscr{A})$. A* distance method *d is a function that associates a distance measure $d_{\mathscr{A}}$ to each AF $\mathscr{A}$.*

We can extend a distance measure $d_{\mathscr{A}}$ so that it also returns distance from an $\mathscr{A}$-profile to an $\mathscr{A}$-labelling, as well as distance between 2 $\mathscr{A}$-profiles. For $L \in Labs(\mathscr{A})$ and profiles $\mathbf{L} = (L_1, \ldots, L_n)$, $\mathbf{L}' = (L'_1, \ldots, L'_n) \in Labs(\mathscr{A})^n$ we define

$$d_{\mathscr{A}}(\mathbf{L}, L) = \sum_{i=1}^{n} d_{\mathscr{A}}(L_i, L), \qquad d_{\mathscr{A}}(\mathbf{L}, \mathbf{L}') = \sum_{i=1}^{n} d_{\mathscr{A}}(L_i, L'_i).$$

Also required for MO is the notion of *M-consistent* profile. These are the profiles that, when passed to *M*, result in a collectively rational outcome.

Definition 8. *Let $\mathscr{A}$ be an AF, $\mathbf{L} \in Labs(\mathscr{A})^n$ and M a resolute aggregation method. Then $\mathbf{L}$ is M-*consistent (for $\mathscr{A}$) *iff $M_{\mathscr{A}}(\mathbf{L}) \in Comp(\mathscr{A})$. We denote by $Cons_{\mathscr{A}}(M)$ the set of M-consistent $\mathscr{A}$-profiles, and by $Cons_{\mathscr{A}}(M, Comp)$ the set $Cons_{\mathscr{A}}(M) \cap Comp(\mathscr{A})^n$.*

MO describe four different ways in which all the above ingredients can be combined, resulting in four classes of aggregation methods which we now describe. First some notation: For any function $f : X \to Y$ and sets $D \subseteq X$, $C \subseteq Y$ we denote the image of D by $f(D) = \{f(x) \mid x \in D\}$ and the inverse image of C by $f^{-1}(C) = \{x \in X \mid f(x) \in C\}$. The subset of D for which f obtains its minimal value is returned by the operator $\arg\min_{x \in D} f(x) = \{x \in D \mid f(x) \leq f(x') \text{ for all } x' \in D\}$.

Definition 9 ([9]). *Let d be distance method and M a resolute aggregation method. The four aggregation methods* $\mathtt{P}^d$, $\mathtt{E}^{M,d}$, $\mathtt{F}^{M,d}$ *and* $\mathtt{O}^{M,d}$ *are defined by setting, for each AF* $\mathscr{A}$ *and* $\mathscr{A}$*-profile* $\mathbf{L}$*:*

$$\mathtt{P}^d_{\mathscr{A}}(\mathbf{L}) = \underset{L \in Comp(\mathscr{A})}{\arg\min}\ d_{\mathscr{A}}(\mathbf{L}, L) \qquad \mathtt{F}^{M,d}_{\mathscr{A}}(\mathbf{L}) = M_{\mathscr{A}}\left(\underset{\mathbf{L}' \in Cons_{\mathscr{A}}(M,Comp)}{\arg\min}\ d_{\mathscr{A}}(\mathbf{L}, \mathbf{L}')\right)$$

$$\mathtt{E}^{M,d}_{\mathscr{A}}(\mathbf{L}) = \underset{L \in Comp(\mathscr{A})}{\arg\min}\ d_{\mathscr{A}}(M_{\mathscr{A}}(\mathbf{L}), L) \qquad \mathtt{O}^{M,d}_{\mathscr{A}}(\mathbf{L}) = M_{\mathscr{A}}\left(\underset{\mathbf{L}' \in Cons_{\mathscr{A}}(M)}{\arg\min}\ d_{\mathscr{A}}(\mathbf{L}, \mathbf{L}')\right)$$

All four MO aggregation methods minimise distance to ensure the collective outcome is rational. The P and E methods minimise the distance over all complete labellings. P (or Prototype) returns the complete labellings closest to the profile $\mathbf{L}$. E (or Endpoint) returns the complete labellings closest to the labelling returned by initial aggregator $M_{\mathscr{A}}(\mathbf{L})$, which possibly is not complete. The F and O (for Full and Output resp.) methods select the M-consistent profiles closest to $\mathbf{L}$ and then applies $M_{\mathscr{A}}$ to them. The difference between these two is that O performs its selection from among *all* M-consistent profiles, while F selects only from those that are, in addition, themselves complete.

Some observations in these definitions: *(i)* All four aggregation methods are potentially irresolute. *(ii)* P doesn't require an initial aggregator M, only a distance method d. The other three all rely on M. *(iii)* A distance method d used in E and O needs to return the distance between *all* labellings. In contrast, for P and F it is enough that d is defined only between complete labellings.

If we want to apply MO to our problem of aggregating labellings we need to instantiate the two parameters M and d. Let's look at each in turn.

3.1 Initial aggregation methods

An interesting family of resolute aggregation methods has been defined in [1], namely the *interval aggregation methods*. Formally, let Int_n be the set of *intervals* of non-zero length in $\{0,1,\ldots,n\}$, i.e., $Int_n = \{(k,l) \mid k < l,\ k,l \in \{0,1,\ldots,n\}\}$. Let $Y \subseteq Int_n$ be some subset of distinguished intervals in Int_n. Then we define aggregation method F^Y by setting, for each $\mathscr{A}$, $\mathscr{A}$-profile $\mathbf{L}$ and $a \in Args_{\mathscr{A}}$:

$$[F^Y_{\mathscr{A}}(\mathbf{L})](a) = \begin{cases} \mathtt{x} & \text{if } \mathtt{x} \in \{\mathtt{in}, \mathtt{out}\} \text{ and } (|V^{\mathbf{L}}_{a:\neg\mathtt{x}}|, |V^{\mathbf{L}}_{a:\mathtt{x}}|) \in Y \\ \mathtt{undec} & \text{otherwise,} \end{cases}$$

where, for any $\mathtt{x} \in \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$, $V^{\mathbf{L}}_{a:\mathtt{x}}$ denotes $\{i \in Ag \mid L_i(a) = \mathtt{x}\}$. A particular member of this family, which we will use in our examples, is the *credulous* aggregation method [4], denoted by *cio*. This is obtained by taking $Y_{cio} = \{(0,l) \in Int_n \mid l > 0\}$. The credulous method *cio* returns a collective label of in (resp. out) to an argument if at least one agent votes for in (resp. out) while none vote for the opposite label out (resp. in). Otherwise it returns undec.

Interval methods may be characterised by a number of postulates such as *Anonymity*, *Unanimity* and *AF-Independence* (the collective label of a is calculated independently of which other arguments might be present or absent from $\mathscr{A}$). However, despite their simplicity, there is *no* interval method that satisfies *Collective Rationality* (we refer the reader to [1] for the details).

3.2 Labelling distance methods

In [2] a few distance methods were proposed, although that paper was concerned only in defining distances between *complete* $\mathscr{A}$*-labellings* rather than between any two arbitrary $\mathscr{A}$-labellings. (This is a point we will return to later.) The idea is first to define a distance *diff* over the set $\{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$ of *labels* and then define the distance between two $\mathscr{A}$-labellings as a sum over some set of arguments of the distances between labels assigned by those $\mathscr{A}$-labellings to arguments from the set. Formally, all of the distance methods of [2] shared the following form, given $\mathscr{A}$ and $L_1, L_2 \in Comp(\mathscr{A})$:

$$d^{diff,\mathfrak{S}}_{\mathscr{A}}(L_1, L_2) = \sum_{a \in \mathfrak{S}(\mathscr{A})} diff(L_1(a), L_2(a)) \tag{1}$$

where *(i) diff* is a metric over the set of labels and *(ii)* $\mathfrak{S}$ is a function that, for each AF $\mathscr{A}$ selects a subset $A \subseteq Args_{\mathscr{A}}$ of "important" arguments in $\mathscr{A}$. A number of different combinations of *diff* and $\mathfrak{S}$ were considered in [2].

In particular two distance metrics between labels were considered - *discrete* metrics, i.e. $diff_D(\mathtt{x},\mathtt{y}) = 1$ if $\mathtt{x} \neq \mathtt{y}$, $diff_D(\mathtt{x},\mathtt{y}) = 0$ if $\mathtt{x} = \mathtt{y}$ (which, when plugged into the method sd^{diff} below results in taking the Hamming distance between labellings), and a distance metric which assigns 2 to the hard conflict, i.e. $diff_{rh}(\mathtt{in},\mathtt{out}) = 2$, and 1 to soft conflicts, i.e. $diff_{rh}(\mathtt{in},\mathtt{undec}) = diff_{rh}(\mathtt{out},\mathtt{undec}) = 1$ (which when plugged into sd^{diff} results in the *refined Hamming* distance - see [2]).

Regarding $\mathfrak{S}$, three options were considered. The most obvious idea is to take $\mathfrak{S}(\mathscr{A}) = Args_{\mathscr{A}}$. We denote by sd^{diff} the resulting distance method. As was shown in [2][Sect. 6], the drawback with this method is that it somehow leads to "double-counting" of label differences that are already in some sense implied by others. This leads us to focus on critical sets (originally due to [6] but here generalised to make it relative to $\mathscr{X}$).

Definition 10. *Given $\mathscr{A}$ and $\mathscr{X} \in \{Labs, Comp\}$, a set of arguments $A \subseteq Args_{\mathscr{A}}$ is $\mathscr{X}$-*critical (for $\mathscr{A}$) *iff for any $L_1, L_2 \in \mathscr{X}(\mathscr{A})$, if $L_1[A] = L_2[A]$ then $L_1 = L_2$.*

A $\mathscr{X}$-critical set is a set of arguments such that any two $\mathscr{X}$-labellings are different iff they label at least one argument in the set differently. Note there is only one *Labs*-critical set for $\mathscr{A}$, namely $Args_{\mathscr{A}}$. An idea explored in [2] was for $\mathfrak{S}(\mathscr{A})$ to always select some $\subseteq$-*minimal Comp*-critical set for $\mathscr{A}$. This leads to the definition of the *critical sets distance method*, denoted by $cd^{diff,\mathfrak{S}}$. A problem with this is that different choices of minimal critical set can yield quite different results. The distance method finally arrived at in [2] uses the notion of *issue* (again generalised below to make it relative to $\mathscr{X}$).

Definition 11. *Given $\mathscr{X} \in \{Labs, Comp\}$ we define an equivalence relation $\equiv_{\mathscr{X}}$ over $Args_{\mathscr{A}}$ by setting $a \equiv_{\mathscr{X}} b$ iff* either *[$L(a) = L(b)$ for all $L \in \mathscr{X}(\mathscr{A})$]* or *[$L(a) = \neg L(b)$ for all $L \in \mathscr{X}(\mathscr{A})$]. Each $\equiv_{\mathscr{X}}$-equivalence class is called an $\mathscr{X}$-*issue*, and we denote the $\equiv_{\mathscr{X}}$-equivalence class to which a belongs by $[a]_{\mathscr{X}}$.*

If $a \equiv_{\mathscr{X}} b$ then either a, b are always labelled the same by every labelling in $\mathscr{X}$, or they are always labelled "opposite". Thus changing the label of one of them always produces a change of equal magnitude in the label of the other. The idea behind the *issue-based distance method* $id^{diff,\mathfrak{S}}$ is to form $\mathfrak{S}(\mathscr{A})$ by selecting one representative from each *Comp*-issue. Note $\mathfrak{S}(\mathscr{A})$ so defined will return a *Comp*-critical set for $\mathscr{A}$, though not necessarily a minimal one. It was also shown in [2] that the resulting distance is independent of the choice of representative from the *Comp*-issue. In [2] it was shown that $id^{diff,\mathfrak{S}}$ forms a metric over $Comp(\mathscr{A})$. Below we give a more general result.

Proposition 1. *Let $\mathscr{X} \in \{Labs, Comp\}$. For any label metric diff and any function $\mathfrak{S}$, the function $d^{diff,\mathfrak{S}}_{\mathscr{A}}$ defined in (1) defines a pseudometric over $\mathscr{X}(\mathscr{A})$. Moreover, it defines a metric over $\mathscr{X}(\mathscr{A})$ iff $\mathfrak{S}(\mathscr{A})$ is $\mathscr{X}$-critical.*

In [2] it was assumed the above distances were defined only between complete labellings. In this case $id^{diff,\mathfrak{S}}$ seems to be a good candidate to use in MO. But two of the MO aggregation methods, namely E and O, require distance to be defined between *all* labellings, and in this case Prop. 1 gives us a problem, for it tells us that the **only** way for any distance method of the form (1) to yield a metric over the whole set $Labs(\mathscr{A})$ is if $\mathfrak{S}(\mathscr{A}) = Args_{\mathscr{A}}$. The question is, is there any alternative way to define a distance method such that $d_{\mathscr{A}}$ is a metric over $Labs(\mathscr{A})$, but which agrees with $id^{diff,\mathfrak{S}}$ on $Comp(\mathscr{A})$? Here we give one possibility. The idea is to take a sum over all arguments, but to weight the contribution of a in the sum by the inverse of the size of the *Comp*-issue to which a belongs. This gives rise to the *extended issue-based distance method* eid^{diff}.

$$eid^{diff}_{\mathscr{A}}(L_1, L_2) = \sum_{a \in Args_{\mathscr{A}}} \frac{diff(L_1(a), L_2(a))}{|[a]_{Comp}|}$$

Proposition 2. (i). *$eid^{diff}_{\mathscr{A}}$ is a metric over $Labs(\mathscr{A})$* (ii). *$eid^{diff}_{\mathscr{A}}(L_1, L_2) = id^{diff,\mathfrak{S}}_{\mathscr{A}}(L_1, L_2)$ for all $L_1, L_2 \in Comp(\mathscr{A})$.*

Therefore we have two distance methods that can be used freely on $Labs(\mathscr{A})$: sd^{diff} and eid^{diff}.

Example 2. *Table 1 presents some distances returned by $d = eid^{diff}_{\mathscr{A}_2}$, where $diff = diff_{rh}$ and $\mathscr{A}_2$ is from Fig. 1. $Args_{\mathscr{A}_2}$ partitions into three Comp-issues $\{a, b\}, \{c\}$ and $\{d, e\}$. Let us calculate a few entries as an example: $d(l_1, l_2) = 2$ because there is no conflict over first two issues and there is a hard conflict over the last one $(0+0+2)$; $d(l_3, l_4) = 2$ because there are two soft conflicts over the first and the last issue and no conflict on the middle one $(1+0+1)$; $d(l_5, n_4) = 1.5$ because there is half of a soft conflict over the first issue and a soft conflict over the second one $(0.5+1+0)$; etc.*

	l_1	l_2	l_3	l_4	l_5	l_6	n_1	n_2	n_3	n_4
l_1: oioio	0.0	2.0	1.0	1.0	4.0	3.0	2.0	5.0	5.0	4.5
l_2: oiooi	2.0	0.0	1.0	3.0	2.0	3.0	2.0	3.0	3.0	2.5
l_3: oiouu	1.0	1.0	0.0	2.0	3.0	2.0	1.0	4.0	4.0	3.5
l_4: uuoio	1.0	3.0	2.0	0.0	3.0	2.0	1.0	4.0	4.0	4.5
l_5: uuuoi	4.0	2.0	3.0	3.0	0.0	1.0	2.0	1.0	1.0	1.5
l_6: uuuuu	3.0	3.0	2.0	2.0	1.0	0.0	1.0	2.0	2.0	2.5
n_1: uuouu	2.0	2.0	1.0	1.0	2.0	1.0	0.0	3.0	3.0	3.5
n_2: iouoi	5.0	3.0	4.0	4.0	1.0	2.0	3.0	0.0	2.0	2.5
n_3: uuioi	5.0	3.0	4.0	4.0	1.0	2.0	3.0	2.0	0.0	0.5
n_4: uiioi	4.5	2.5	3.5	4.5	1.5	2.5	3.5	2.5	0.5	0.0

Table 1: Distance between labellings by eid^{diff} with refined Hamming distance $diff_{rh}$ over the labels.

	$cio(\mathbf{L})$	P($\mathbf{L}$)	E($\mathbf{L}$)			O($\mathbf{L}$)		F($\mathbf{L}$)
	n_1	l_5	l_3	l_4	l_6	l_3	l_6	l_3
a	u	u	o	u	u	o	u	o
b	u	u	i	u	u	i	u	i
c	o	u	o	o	u	o	u	o
d	u	o	u	i	u	u	u	u
e	u	i	u	o	u	u	u	u

	$cio(\mathbf{L})$	P($\mathbf{L}$)	E($\mathbf{L}$)	O($\mathbf{L}$)	F($\mathbf{L}$)
	n_1	l_5	l_6	l_6	l_3
a	u	u	u	u	o
b	u	u	u	u	i
c	o	u	u	u	o
d	u	o	u	u	u
e	u	i	u	u	u

Table 2: Aggregation of the profile $\mathbf{L} = (l_4\colon \texttt{uuoio}, l_5\colon \texttt{uuuoi}, l_5\colon \texttt{uuuoi})$ - outcomes for different MO aggregation methods used with: extended issue-based labelling distance (left table), and sum over all arguments (right table). In both columns refined hamming distance over the labels and credulous initial operator were used.

3.3 Example of the MO methods

We illustrate the MO methods in our setting by continuing with the AF $\mathscr{A}_2$ from Fig. 1. The $\mathscr{A}_2$-profile $\mathbf{L} = (l_4\colon \texttt{uuoio}, l_5\colon \texttt{uuuoi}, l_5\colon \texttt{uuuoi})$ aggregated with *cio* results in the non-complete labelling $n_1\colon \texttt{uuouu}$. The result of repairing it with MO methods is listed in Table 2.

In the left column the MO methods were instantiated with eid^{diff} with $diff = diff_{rh}$. Method P returns the closest complete labelling to the profile $\mathbf{L}$. We calculate the distance between $\mathbf{L}$ and labellings $l_1 - l_6$ by adding distances from the row l_4 (Table 1) to the doubled distances from the row l_5 and receive $9, 7, 8, 6, 3, 4$ respectively. The minimum is obtained for labelling $l_5\colon \texttt{uuuoi}$.

The E procedure returns the closest complete labelling to $cio(\mathbf{L})$. We inspect the row n_1 in Table 1 to find that the minimum distance 1 is obtained for l_3, l_4 and l_6.

The F and O procedures search for closest *cio*-consistent profiles $\mathbf{L}'$ to profile $\mathbf{L}$. The F procedure is restricted to the complete *cio*-consistent profiles. Consider $\mathbf{L}' = (l_1\colon \texttt{oioio}, l_5\colon \texttt{uuoio}, l_5\colon \texttt{uuoio})$. It is a *cio*-consistent profile with $cio(\mathbf{L}') = l_3\colon \texttt{oiouu}$. It is also minimal. The profiles $\mathbf{L}$ and $\mathbf{L}'$ differ just on l_1 and l_4 with distance 1 (soft conflict on issue $\{a, b\}$). There are no other complete *cio*-consistent profiles with different *cio* outcome and same distance, because any change of labelling l_4 to another complete labelling costs more than 1. The only other candidate for change is labelling l_5. It can be changed to $l_6\colon \texttt{uuuuu}$ for a cost of 1 but to affect the outcome of *cio* both occurrences of l_5 in $\mathbf{L}$ would have to be changed with a total cost of 2. The profile $\mathbf{L}'$ works for the O procedure as well but O is not restricted to complete *cio*-consistent profiles. Changing one of the l_5 labellings to a non-complete labelling $n_3\colon \texttt{uuioi}$ with distance 1 creates additional conflict on argument c. As a result the labelling $l_6\colon \texttt{uuuuu}$ is produced.

The results change when we switch to using sd^{diff} (with $diff = diff_{rh}$) rather than eid^{diff} (right column). In this case size of the issues does matter. The E procedure only selects l_6 because it differs with n_1 over issue $\{c\}$ with one argument, while the other two labellings l_3, l_4 differ over the two-argument issues $\{a, b\}$ and $\{d, e\}$ respectively. Similarly in the case of O changing l_5 into n_3 over a single-argument issue is closer than the change of l_4 into l_1 over an issue with two arguments.

4 Qualitative distance postulates

In this section we consider two new postulates for distance measures in argumentation, which are inspired by similar postulates considered in the binary aggregation of MO.

One intuition that MO have about distance between judgment sets in their setting is that the *quantitative* distance between two judgment sets X,Y, expressed by $d(X,Y)$, should depend only on the *qualitative distance* between them. In their binary setting in which every judgment set either accepts or rejects each agenda item, the natural way to measure such qualitative distance is simply to take the set of agenda items that X and Y differ on. In our 3-valued setting we can have two labellings that disagree on an argument to differing *degrees*. We thus adapt the notion of qualitative distance as follows:

Definition 12. *Let $\mathscr{A}$ be an AF and $L_1, L_2 \in Labs(\mathscr{A})$. The* qualitative distance between L_1 and L_2, *denoted $L_1 \ominus L_2$, is defined as $L_1 \ominus L_2 = \langle C(L_1,L_2), H(L_1,L_2)\rangle$, where $C(L_1,L_2)$ and $H(L_1,L_2)$ are resp. the sets of* conflicts *and of* hard conflicts *between L_1 and L_2, i.e., $C(L_1,L_2) \stackrel{\text{def}}{=} \{a \in Args_{\mathscr{A}} \mid L_1(a) \neq L_2(a)\}$ and $H(L_1,L_2) \stackrel{\text{def}}{=} \{a \in Args_{\mathscr{A}} \mid L_1(a) = \neg L_2(a) \neq \texttt{undec}\}$.*

Note that $H(L_1,L_2) \subseteq C(L_1,L_2)$. We can define a natural ordering between different qualitative distances as follows:

$$L_1 \ominus L_2 \preceq N_1 \ominus N_2 \text{ iff } [C(L_1,L_2) \subseteq C(N_1,N_2) \text{ and } H(L_1,L_2) \subseteq H(N_1,N_2)]$$

So if every conflict, resp. hard conflict, between L_1, L_2 remains a conflict, resp. hard conflict, between N_1, N_2 then the disagreement between N_1, N_2 is qualitatively at least as large as that between L_1, L_2. One can check that $\preceq$ forms a preorder on $Labs(\mathscr{A})^2$. We denote the strict part of $\preceq$ by $\prec$. We can now adapt the distance property of *Normality* [9] to our setting. It expresses that a (weak) increase in qualitative distance should yield a (weak) increase in quantitative distance.

Normality If $L_1 \ominus L_2 \preceq N_1 \ominus N_2$ then $d_{\mathscr{A}}(L_1,L_2) \leq d_{\mathscr{A}}(N_1,N_2)$

If *Normality* holds then clearly we also have that if $L_1 \ominus L_2 = N_1 \ominus N_2$ then $d_{\mathscr{A}}(L_1,L_2) = d_{\mathscr{A}}(N_1,N_2)$. Thus if *Normality* is required then the problem of defining a distance measure $d_{\mathscr{A}}$ for $\mathscr{A}$ essentially reduces to the problem of assigning a number to $L_1 \ominus L_2$ for each $L_1, L_2 \in Labs(\mathscr{A})^2$.

Example 3. *For $\mathscr{A}_2$ from Fig. 1 we have $l_2 \ominus l_3 = l_5 \ominus l_6 = \langle\{d,e\},\emptyset\rangle \preceq \langle\{d,e\},\{d,e\}\rangle = l_1 \ominus l_2$. Hence if d satisfies* Normality *we must have $d_{\mathscr{A}_2}(l_2,l_3) = d_{\mathscr{A}_2}(l_5,l_6) \leq d_{\mathscr{A}_2}(l_1,l_2)$.*

Similarly we can postulate that a *strict* increase in qualitative distance should lead to a *strict* increase in quantitative distance, leading to the following adaptation of the *Increasing* property of MO:

Increasing If $L_1 \ominus L_2 \prec N_1 \ominus N_2$ then $d_{\mathscr{A}}(L_1,L_2) < d_{\mathscr{A}}(N_1,N_2)$

How do the distance methods described in the previous section fare with respect to *Normality* and *Increasing*? As we see, provided *diff* satisfies certain conditions then at least some of them validate these postulates.

Proposition 3. (i). *Assume the label metric diff satisfies $diff(\texttt{in},\texttt{out}) \geq diff(\texttt{in},\texttt{undec}) = diff(\texttt{out},\texttt{undec})$. Then all of $sd^{diff}_{\mathscr{A}}$, $cd^{diff,\mathfrak{S}}_{\mathscr{A}}$, $id^{diff,\mathfrak{S}}_{\mathscr{A}}$ and $eid^{diff}_{\mathscr{A}}$ satisfy* Normality *over $Labs(\mathscr{A})$ (and hence also over $Comp(\mathscr{A})$).*
(ii). *Assume the label metric diff satisfies $diff(\texttt{in},\texttt{out}) > diff(\texttt{in},\texttt{undec}) = diff(\texttt{out},\texttt{undec})$. Then* (1). *$sd^{diff}_{\mathscr{A}}$ and $eid^{diff}_{\mathscr{A}}$ satisfy* Increasing *over $Labs(\mathscr{A})$ (and hence also $Comp(\mathscr{A})$).* (2). *$id^{diff,\mathfrak{S}}_{\mathscr{A}}$ satisfies* Increasing *over $Comp(\mathscr{A})$, but not over $Labs(\mathscr{A})$.* (3) *$cd^{diff,\mathfrak{S}}_{\mathscr{A}}$ does not satisfy* Increasing *over $Comp(\mathscr{A})$ (and hence also not over $Labs(\mathscr{A})$).*

Observe that $diff_{rh}$ satisfies both conditions mentioned in *(i)*, *(ii)* above, and so can be plugged into sd^{diff} or eid^{diff} to obtain a distance method satisfying both *Normality* and *Increasing*.

Surprisingly, $cd^{diff,\mathfrak{S}}$ fails *Increasing* even when restricted to $Comp(\mathscr{A})$, as the next example demonstrates.

Example 4. *Consider $\mathscr{A}_1$ from Fig. 1. There are three Comp-issues: $\{a,b\},\{c\},\{d,e\}$. But the label of e is determined by the labels of the other arguments, so there are 4 possible minimal Comp-critical sets $\{a,b\} \times \{d,e\}$. Let's take $\{a,d\}$ as our selected minimal set (though the counterexample will also work for any of the other three possible choices). Consider labellings l_1:* `oioio`*, l_8:* `ioioi` *and l_2:* `oiooi`*, l_7:* `iooio`*. We have $l_2 \ominus l_7 = \langle\{a,b,d,e\},\{a,b,d,e\}\rangle \prec \langle\{a,b,c,d,e\},\{a,b,c,d,e\}\rangle = l_1 \ominus l_8$ but $d(l_2,l_7) = d(l_1,l_8) = 2 \times diff(\texttt{in},\texttt{out})$.*

5 Conclusions

We have continued work on distance methods for argumentation that was initiated in [2], illustrating how they can be employed to address problems of aggregation in argumentation. To do this we adapted the framework of Miller and Osherson from binary judgment aggregation to our setting, defining several operators for aggregating argument labellings. In the process we extended and generalised some results on the distance measures from [2]. We also examined an adaptation of the MO postulates *Normality* and *Increasing* to our setting and were able to identify the circumstances under which they hold. Under some mild conditions on the label metric *diff*, it turned out that only two of our distance methods, namely sd^{diff} and the *extended issue-based* method eid^{diff} are able to satisfy both without restriction on their domain $Labs(\mathscr{A})$.

There are several avenues open for future work. Firstly, a feature of our examples (see Sect. 3.3) is that the different aggregation methods can all yield quite different results. This raises the question of which method to prefer. We plan to classify the different methods in terms of the postulates they satisfy. Previous works on labelling aggregation [1, 10] have examined postulates for such operators (inspired by postulates from JA), but have done so only for *resolute* aggregation methods. We will generalise these to irresolute methods, perhaps taking a lead from similar generalisations from JA [7].

Secondly, it can be shown that aggregation methods which use the *down-admissible* and *up-complete* procedures of [4] can be represented as an instance of Endpoint method. We would like to investigate the classes of aggregation operators given by different distance methods and relate them to the ones which use *down-admissible* and *up-complete* procedures.

Another interesting question is to consider what happens if you aggregate *all* complete $\mathscr{A}$-labellings of an AF. This question was considered in [4], but again only for resolute operators. In this way the operators of [4] were able to characterise certain *single-status* argumentation semantics (i.e., grounded and ideal). Our move to irresolute aggregation opens the possibility that we might be able to capture also some *multiple-status* semantics such as *preferred* [5]. That is, does aggregating all complete $\mathscr{A}$-labellings yield precisely the set of preferred $\mathscr{A}$-labellings?

References

[1] R. Booth, E. Awad, and I. Rahwan. Interval methods for judgment aggregation in abstract argumentation. In *Proc. KR 2014*, pages 594–597, 2014.

[2] R. Booth, M. Caminada, M. Podlaszewski, and I. Rahwan. Quantifying disagreement in argument-based reasoning. In *Proc. AAMAS 2012*, pages 493–500, 2012.

[3] M. Caminada and D. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2-3):109–145, 2009.

[4] M. Caminada and G. Pigozzi. On judgment aggregation in abstract argumentation. *Autonomous Agents and Multi-Agent Systems*, 22(1):64–102, 2011.

[5] P.M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.

[6] D. Gabbay. Fibring argumentation frames. *Studia Logica*, 93(2-3):231–295, 2009.

[7] U. Grandi and G. Pigozzi. On compatible multi-issue group decisions. *Proc. LOFT-2012*, 2012.

[8] C. List and C. Puppe. Judgment aggregation: A survey. In *Handbook of Rational and Social Choice*. Oxford University Press, 2009.

[9] M. Miller and D. Osherson. Methods for distance-based judgment aggregation. *Social Choice and Welfare*, 32(4):575–601, 2009.

[10] I. Rahwan and F. Tohmé. Collective argument evaluation as judgement aggregation. In *Proc. AAMAS 2010*, pages 417–424, 2010.