

Aggregation operators to support collective reasoning

Juan A. Rodriguez-Aguilar¹ and Marc Serramia² and Maite Lopez-Sanchez²

¹ Artificial Intelligence Research Institute (IIIA-CSIC), Campus UAB. Bellaterra, Spain
jar@iiia.csic.es

² Mathematics and Computer Science Department, University of Barcelona (UB), Gran Via de les Corts Catalanes 585, Barcelona, Spain

Abstract. Moderation poses one of the main Internet challenges. Currently, many Internet platforms and virtual communities deal with it by intensive human labour, some big companies –such as YouTube or Facebook– hire people to do it, others –such as 4chan or fanscup– just ask volunteer users to get in charge of it. But in most cases the policies that they use to decide if some contents should be removed or if a user should be banned are not clear enough to users. And, in any case, typically users are not involved in their definition.

Nobel laureate Elinor Ostrom concluded that societies –such as institutions that had to share scarce resources– that involve individuals in the definition of their rules performed better –resources lasted more or did not deplete– than those organisations whose norms were imposed externally. Democracy also relies on this same idea of considering peoples’ opinions.

In this vein, we argue that participants in a virtual community will be more prone to behave correctly –and thus the community itself will be "healthier"– if they take part in the decisions about the norms of coexistence that rule the community. With this aim, we investigate a collective decision framework that: (1) structures (relate) arguments issued by different participants; (2) allows agents to express their opinions about arguments; and (3) aggregates opinions to synthesise a collective decision. More precisely, we investigate two aggregation operators that merge discrete and continuous opinions. Finally, we analyse the social choice properties that our discrete aggregator operator satisfies.

1 Introduction

With the advent of the Internet, a plethora of on-line communities, such as social networks, have emerged to articulate human interaction. Nonetheless, interactions are not frictionless. Thus, for instance, users may post inappropriate or offensive contents, or spam ads. Thus, typically the owners of on-line communities establish their own norms (terms and policies) to regulate interactions without the involvement of its participants. Moderators become then in charge of guaranteeing the enforcement of such norms disregarding what users may deem as fair or discomforting.

Here we take the stance that the participants in a social network must decide the norms that govern their interactions. Thus, we are in line with Nobel-prize winner E. Ostrom [4], who observed that involving a community’s participants in their decisions improves its long-term operation. Then, there is the matter of helping users agree on

their norms. As argued in [2, 3], argumentative debates are a powerful tool for reaching agreements in open environments such as on-line communities. On-line debates are usually organised as threads of arguments and counter-arguments that users issue to convince others. There are two main issues in the management of large-scale on-line debates. On the one hand, as highlighted by [2] and [3], there is simply too much noise when many individuals participate in a discussion, and hence there is the need for *structuring* it to keep the focus. On the other hand, the preferences on arguments issued by users must be aggregated to achieve a collective decision about the topic under discussion [1].

Against this background, here we consider that structured argumentative debates can also be employed to help users of a virtual community jointly agree on the norms that rule their interactions. With this aim, we present the following contributions:

- Based on the work in [3], we introduce an argumentative structure, the so-called *norm argument map*, to structure a debate focusing on the acceptance or rejection of a target norm. Figure 1 shows one example in an online sports community.
- A novel aggregation method to assess the *collective support for a single argument* by aggregating the preferences (expressed as ratings) issued by the participants in a discussion. Such method will consider that the impact of a single rating on the overall aggregated value will depend on the distance of that rating from neutrality. More precisely, our aggregation method abides by the following design principle: the farther a rating is from neutrality, the stronger its importance when computing the collective support for an argument.
- A novel aggregation method to compute the *collective support for a norm* based on the arguments issued by the participants in a discussion. This method is based on the following design principles: (1) the larger the support for an argument, the larger its importance on the computation of the collective support for a norm; and (2) only those arguments that are *relevant enough* (count on *sufficient support*) are worth aggregating. Technically, this method is conceived as a WOVA operator [7] because it allows to consider both the values and the information sources when performing the aggregation of argument supports.
- We compared our aggregation method with a more naive approach that simply averages participants' preferences on a collection of prototypical argumentation scenarios. We observe that our method obtains support values for norms that better capture the collective preference of the participants.

The paper is organised as follows. Section 2 introduces some background on the aggregation operators employed, sections 3, 4, 5, and 6 introduce our formal notion of norm argument map and our functions to compute the support for an argument, a set of arguments and a norm. Section 7 details the analysis of our support functions on argumentation scenarios. Finally, section 8 draws conclusions and sets paths to future research.

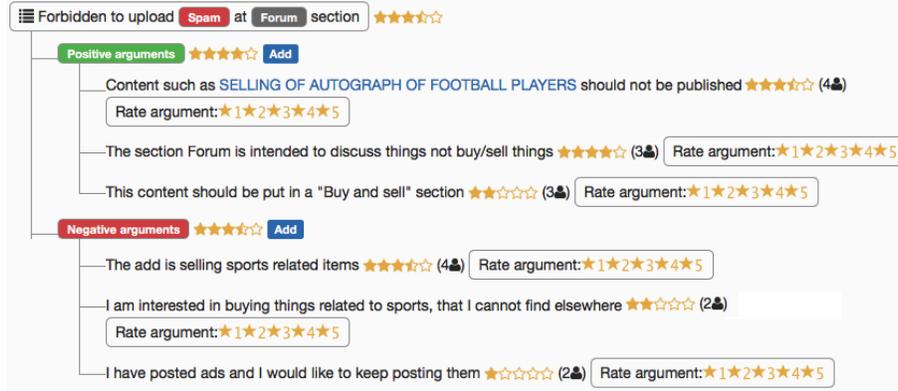


Fig. 1. Example of a norm argument map. Rated positive/negative arguments in favor/against a norm prohibiting to upload spam content at a social network forum.

2 Background

As previously stated, the main goal of this work is to compute an aggregated numerical score for a norm from its arguments and opinions³. Hence, aggregation operators become necessary to fuse all the numerical information participants provide. Next, we introduce the aggregation operators employed in this work, namely the *standard weighted mean (WM)* and the *weighted ordered weighted average (WOWA, an OWA's [8] variation)* from Torra [7], to compute the collective support for a norm.

Definition 1. A **weighting vector** w is a vector such that if $w = (w_1, \dots, w_n) \in \mathbb{R}^n$ then $w_i \in [0, 1]$ and $\sum_{i=1}^n w_i = 1$.

Definition 2. Let $w = (w_1, \dots, w_n) \in \mathbb{R}^n$ be a weighting vector and let $e = (e_1, \dots, e_n) \in \mathbb{R}^n$ be the vector of elements we want to aggregate. A **weighted mean** is a function $WM_w(e) : \mathbb{R}^n \rightarrow \mathbb{R}$, defined as $WM_w(e) = \sum_{i=1}^n w_i e_i$.

Notice that WM weighs the position of the elements, which amounts to concede different importance degrees to each particular (information) source. In order to weigh the values of aggregated elements in e we need an alternative operator.

Definition 3. Let $w = (w_1, \dots, w_n) \in \mathbb{R}^n$ and $q = (q_1, \dots, q_n) \in \mathbb{R}^n$ be two weighing vectors and let $e = (e_1, \dots, e_n) \in \mathbb{R}^n$ be the vector of elements we want to aggregate. A **weighted ordered weighted average, weighted OWA or WOWA** is a function $WOWA_{w,q}(e) : \mathbb{R}^n \rightarrow \mathbb{R}$ defined as:

$$WOWA_{w,q}(e) = \sum_{i=1}^n p_i e_{\sigma(i)}, \quad p_i = f^*\left(\sum_{j \leq i} w_{\sigma(j)}\right) - f^*\left(\sum_{j < i} w_{\sigma(j)}\right),$$

³ An argument's opinions are numerical values that, in the case of Figure 1, take the form of number of stars awarded to each argument.

where σ is a permutation of the elements in e so that $e_{\sigma(i)}$ is the i^{th} largest element in e and f^* is a non-decreasing interpolation function of the points: $\{(i/n, \sum_{j \leq i} q_j)\}_{i=1, \dots, n} \cup \{(0, 0)\}$ that has to be a straight line when the points can be interpolated that way.

Note that w acts as the vector in the weighted mean, weighing the information source, while q weighs the value of the aggregated elements. For instance, $q_1 = q_n > q_2, \dots, q_{n-1}$, gives more importance to extreme (i.e., the highest and lowest) values.

3 Norm Argument Map

Next we formalise the notion of norm argument map as the argumentative structure that contains all arguments and opinions about a norm.

Definition 4. A *norm* is pair $n = (\phi, \theta(\alpha))$, where ϕ is the norm's precondition, θ is a deontic operator⁴ and α is an action that participants can perform.

Definition 5. An *argument* is a pair $a_i = (s, \mathbf{O}_{a_i})$ composed of a statement s , the argument itself, and a vector of opinions \mathbf{O}_{a_i} that contains all the opinion values participants issued.

Henceforth we will note the vector of opinions as $\mathbf{O}_{a_i} = (o_1^i, \dots, o_{n_i}^i)$, where o_j^i is the j^{th} opinion about argument a_i .

Definition 6. Given a norm n , the *argument set* for n is a non-empty collection of arguments A_n containing both arguments supporting and attacking the norm.

We will note the vector of all the opinions of the arguments in A_n as \mathbf{O}_{A_n} . For the sake of simplicity, we assume that all arguments in the argument set A_n of a norm are different. We divide the argument set of a norm into two subsets: the arguments in favor of the norm and the arguments against it.

We are now ready to define our argumentative structure as follows:

Definition 7. A *norm argument map* $M = (n, A_n, \kappa)$ is a triple composed of a norm n , a norm argument set A_n , and a function κ that classifies the arguments of A_n between the ones that are in favor of the norm and the ones that are against it.

Hereafter we will refer to the positive arguments of norm n as the set of arguments in favor of the norm and to its negative arguments as the set of arguments against the norm. These argument sets will be noted as A_n^+ and A_n^- respectively. A negative argument is distinguished from positive arguments by adding a bar over the argument (e.g. $\bar{a}_i \in A_n^-$).

Finally, we also define a framework wherein participants can simultaneously discuss over multiple norms.

Definition 8. A *norm argument map framework* $F = (P, N)$ is a pair of a set of participants P and a set of norm argument maps N , so that participants in P can deliberate about different norms by means of the norm argument maps in N .

⁴ A deontic operator stands for either prohibition, permission, or obligation.

4 Argument support

Having defined the norm argument map we aim now at aggregating arguments' opinions to calculate the support for each argument. In our case opinions will be numerical values defined in an opinion spectrum.

Definition 9. An *opinion spectrum* is a set of possible numerical values individual participants can assign to each argument meaning her opinion about the argument.

The spectrum will be considered a closed real number interval, and thus there exist a maximum, a minimum and a middle opinion values. Figure 2 shows an example of the opinion spectrum semantics considering $\lambda = [1, 5]$. Since opinions will have different values, we consider different semantics for them. The opinion spectrum will be divided into three subsets of opinions. Given an opinion spectrum $\lambda = [lb, ub]$ such that $lb, ub \in \mathbb{R}$ and $lb < ub$: $[lb, \frac{lb+ub}{2})$ contains the values for negative opinions, $(\frac{lb+ub}{2}, ub]$ contains the values for positive opinions, and $\{\frac{lb+ub}{2}\}$ contains the value for the neutral opinion. Note that an opinion $o_j^i = lb$ is the most extreme opinion against argument a_i , while another opinion $o_k^i = ub$ would represent the most extreme opinion in favor of the argument. Additionally, we consider the opinion laying in the middle of the spectrum $\frac{lb+ub}{2}$ as a neutral opinion.

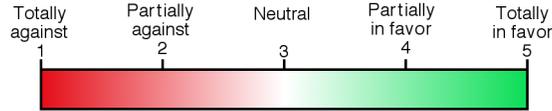


Fig. 2. Semantics of the opinion spectrum $\lambda = [1, 5]$.

Since different opinions in an opinion spectrum have different meanings and we aim at aggregating them in order to calculate the support for an argument, we need a function that weighs the importance of each opinion. Such *importance function* will consider neutral opinions less important than the extreme (strongly stated) ones. Formally,

Definition 10. Given an opinion spectrum $\lambda = [lb, ub]$, we say that a function $I : \lambda \rightarrow [0, 1]$ is an *importance function* iff it satisfies the following conditions:

- (C1) I continuous and piecewise differentiable
- (C2) $I(ub) = I(lb) = 1$
- (C3) $I(\frac{lb+ub}{2}) = 0$
- (C4)
$$\begin{cases} I'(x) < 0 & \text{if } x \in [lb, \frac{lb+ub}{2}) \text{ and } I \text{ is differentiable in } x \\ I'(x) = 0 & \text{if } x = \frac{lb+ub}{2} \text{ and } I \text{ is differentiable in } x \\ I'(x) > 0 & \text{if } x \in (\frac{lb+ub}{2}, ub] \text{ and } I \text{ is differentiable in } x \end{cases}$$

Given a opinion spectrum, we can construct an importance function either by interpolation or geometrically (parabola case). Here we follow the first approach. Below

we formally define the importance function that we propose in this paper, which is graphically depicted in Figure 3.

$$I(x) = \begin{cases} \frac{ub^2 - 1.8 ub lb - 0.2 ub x - 4lb^2 + 9.8 lb x - 4.8x^2}{(lb - ub)^2} & \text{if } x \in \left[lb, \frac{ub+3lb}{4} \right] \\ \frac{1.45ub + 1.75lb - 3.2x}{ub - lb} & \text{if } x \in \left(\frac{ub+3lb}{6}, \frac{3ub+5lb}{8} \right) \\ \frac{4ub^2 + 8 ub lb - 16 ub x + 4lb^2 - 16 lb x + 16x^2}{(lb - ub)^2} & \text{if } x \in \left[\frac{3ub+5lb}{8}, \frac{5ub+3lb}{8} \right] \\ \frac{1.75ub + 1.45lb - 3.2x}{lb - ub} & \text{if } x \in \left(\frac{5ub+3lb}{8}, \frac{3ub+lb}{4} \right) \\ \frac{-4ub^2 - 1.8 ub lb + 9.8 ub x + lb^2 - 0.2 lb x - 4.8x^2}{(lb - ub)^2} & \text{if } x \in \left[\frac{3ub+lb}{4}, ub \right] \end{cases}$$

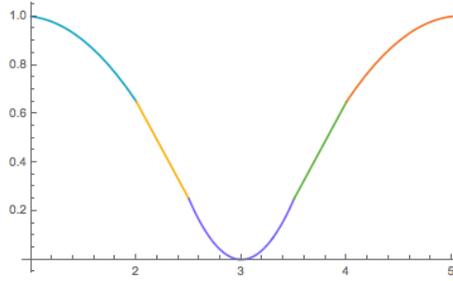


Fig. 3. Importance function (I) plot when $\lambda = [1, 5]$.

We can now weigh the importance of each opinion with our importance function to calculate the support of an argument as the weighted mean of its opinions.

Definition 11. Given an opinion spectrum $\lambda = [lb, ub]$, an **argument support function** $S_{arg} : A \rightarrow \lambda$ is a function that yields the collective support for each argument $a_i \in A$ as: $S_{arg}(a_i) = WM_w(\mathbf{O}_{a_i})$, where $w = \left(\frac{I(o_1^i)}{l_i}, \dots, \frac{I(o_{n_i}^i)}{l_i} \right)$ stands for a weighting vector for the opinions in \mathbf{O}_{a_i} , I is an importance function, and $l_i = \sum_{j=1}^{n_i} I(o_j^i)$.

Notice that o_j^i is the j^{th} opinion of argument a_i and l_i stands for the overall addition of all importance values associated to all opinions about argument a_i . Since the elements in w add up to one, w is a weighting vector.

5 Argument set support

So far we have learned how to aggregate an argument's opinions to compute its support. Next we face the problem of calculating the support for an argument set. To motivate the choice of our aggregation function, we start with an example.

Example. Consider a norm n with positive and negative arguments with opinions in the spectrum $\lambda = [1, 5]$. Say that there are three positive arguments a_1 , a_2 , a_3 , and a

single negative argument \bar{a}_4 . On the one hand, in the set of positive arguments a_1 has a support of 5, which comes from a single opinion while both a_2 and a_3 have a support of 1, which comes from aggregating 100 opinions. On the other hand, on the set of negative arguments, \bar{a}_4 's support is 5, which comes from aggregating 30 opinions:

$\frac{A_n^+}{a_1}$	$S_{arg}(a_i)$	$dim(\mathbf{O}_{a_i})$		$\frac{A_n^-}{\bar{a}_4}$	$S_{arg}(\bar{a})$	$dim(\mathbf{O}_{\bar{a}})$
5	1	100		5	30	
1	100					
1	100					

What should we consider to give the support for A_n^+ on this extreme case? We should discard a_2 or a_3 because they have bad (the minimum) support. People have decided these arguments are not appropriate or do not provide a valid reason to defend the norm under discussion. Since opinions' semantics can be applied to argument support: arguments with supports outside $(\frac{lb+ub}{2}, ub]$ are not accepted by participants and, therefore, should not be considered as valid arguments. We cannot consider a_1 either because, although it has the maximum possible support, it has only been validated by one person, hence it is negligible in front of the other arguments. Therefore, we propose to filter out arguments by just considering those having at least a number of opinions that corresponds to a significant fraction of the number of opinions of the argument with the largest number of opinions.

Thus, we tackle this argument relevance problem by creating a new subset of arguments containing only the arguments considered to be α -relevant (namely, relevant enough) and by defining the criteria needed to be considered as such:

Definition 12. Let A be a set of arguments and $\lambda = [lb, ub]$ an opinion spectrum, we say that an argument $a_i \in A$ is **relevant** iff $S_{arg}(a_i) > \frac{lb+ub}{2}$.

Definition 13. Let A be a set of arguments, $\lambda = [lb, ub]$ an opinion spectrum, $\alpha \in [0, 1]$ a relevance level, and $a_k \in A$ the argument with the largest number of opinions. We say that a relevant argument $a_i \in A$ is **α -relevant** iff $dim(\mathbf{O}_{a_i}) \geq \alpha dim(\mathbf{O}_{a_k})$.

Henceforth, $R_\alpha(A) = \{a_1^\alpha, \dots, a_r^\alpha\}$ will denote the set of α -relevant arguments in A . Notice that $r \leq |A|$ and that in general a_i^α is not equal to a_i .

We propose to aggregate the set of α -relevant arguments by weighting their supports with the importance function previously introduced in order to weight more those arguments that have received greater support than others. Moreover, since arguments count on different numbers of opinions, we consider the sum of importances of their opinions so that important opinions account for more weight than neutral opinions.

To aggregate the supports of the arguments weighting these two values we will use a WOWA operator. Hence, we define the argument set support function as follows:

Definition 14. Let λ be an opinion spectrum, an **argument set support function** S_{set} is a function that takes a non-empty argument set A , with $R_\alpha(A) \neq \emptyset$, and yields its support in λ as:

$$S_{set}(A) = S_{set}(R_\alpha(A)) = WOWA_{w,q}(S_{arg}(a_1^\alpha), \dots, S_{arg}(a_r^\alpha)),$$

where $R_\alpha(A) = \{a_1^\alpha, \dots, a_r^\alpha\}$, $w = \left(\frac{\sum_{j=1}^{\dim(\mathcal{O}_{a_1^\alpha})} I(o_j^1)}{\mathcal{I}_A^\alpha}, \dots, \frac{\sum_{j=1}^{\dim(\mathcal{O}_{a_r^\alpha})} I(o_j^r)}{\mathcal{I}_A^\alpha} \right)$,

$$\mathcal{I}_A^\alpha = \sum_{i=1}^r \left(\sum_{j=1}^{\dim(\mathcal{O}_{a_i^\alpha})} I(o_j^i) \right) \text{ with } o_j^i \in \mathcal{O}_{a_i^\alpha} = \{o_1^i, \dots, o_{n_i}^i\}$$

stands for the overall importance of all the opinions over arguments in A ,

$$q = \left(\frac{I(\text{Sarg}(a_{\sigma(1)}^\alpha))}{\mathcal{I}_A^{\text{arg}}}, \dots, \frac{I(\text{Sarg}(a_{\sigma(r)}^\alpha))}{\mathcal{I}_A^{\text{arg}}} \right)$$

$\mathcal{I}_A^{\text{arg}} = \sum_{i=1}^r I(\text{Sarg}(a_{\sigma(i)}^\alpha))$ stands for the overall importance of the collective supports received by the arguments in A , $a_{\sigma(i)}^\alpha \in R_\alpha(A) = \{a_1^\alpha, \dots, a_r^\alpha\}$, and $a_{\sigma(i)}^\alpha$ is the α -relevant argument with the i^{th} largest support.

Notice that, if there are no α -relevant arguments then we cannot assess the support for the set, hence we consider $S_{\text{set}}(\emptyset)$ to be not defined.

Also note that the w vector is used to weigh the importance of the arguments as the sum of the importances of its opinions. After that we have to divide by \mathcal{I}_A^α so we get a weighting vector. The q vector uses the importance of the supports for the arguments. We have to order the arguments with the σ permutation because the WOWA orders the values being aggregated. This way each weight in the q vector weighs its corresponding element. With this modification, we get the WOWA to aggregate the elements using two weighting vectors. Note that the weighting vector w does not have to be ordered because the WOWA itself orders it.

6 Computing the collective support for a norm

To compute the collective support for a norm, we will use the support for its positive and negative argument sets, namely $S(A_n^+)$ and $S(A_n^-)$. In general, a large support for the negative arguments of a norm is expected to negatively impact the norm's support. Thus, instead of directly aggregating $S(A_n^-)$, we will aggregate the symmetric value of the support in the spectrum with respect to the center of the spectrum, namely $ub + lb - S(A_n^-)$.

Analogously to the computation of the support for an argument set, here we have to weigh the importance of the values aggregated as well as the importance of each argument set as information source. Thus, we will also employ a WOWA operator to compute the collective support for a norm, which we define as follows:

Definition 15. A *norm support function* is a function S_{norm} that takes a norm n , and uses the supports of its positive and negative arguments to obtain the support for the norm in $\lambda = [lb, ub]$. If $R_\alpha(A_n^+) \neq \emptyset$ and $R_\alpha(A_n^-) \neq \emptyset$, the function is defined as follows:

$$S_{\text{norm}}(n) = \text{WOWA}_{w,q}(S_{\text{set}}(A_n^+), ub + lb - S_{\text{set}}(A_n^-))$$

such that the information source is weighed by

$$w = \left(\frac{\sum_{i=1}^{|R_\alpha(A_n^+)|} (\sum_{j=1}^{n_i} I(o_j^i))}{\mathcal{I}_n^o}, \frac{\sum_{i=1}^{|R_\alpha(A_n^-)|} (\sum_{j=1}^{\bar{n}_i} I(\bar{o}_j^i))}{\mathcal{I}_n^o} \right)$$

and the aggregated values are weighed by

$$q = \left(\frac{I(S_{set}(A_n^+))}{\mathcal{I}_n^{set}}, \frac{I(ub + lb - S_{set}(A_n^-))}{\mathcal{I}_n^{set}} \right)$$

where $\mathcal{I}_n^o = \sum_{i=1}^{|R_\alpha(A_n^+)|} (\sum_{j=1}^{n_i} I(o_j^i)) + \sum_{i=1}^{|R_\alpha(A_n^-)|} (\sum_{j=1}^{\bar{n}_i} I(\bar{o}_j^i))$, o_j^i is the j^{th} opinion in $\mathbf{O}_{\alpha_i^+} = \{o_1^i, \dots, o_{n_i}^i\}$, $a_i^\alpha \in R_\alpha(A_n^+) = \{a_1^\alpha, \dots, a_{k_1}^\alpha\}$, \bar{o}_j^i is the j^{th} opinion in $\mathbf{O}_{\alpha_i^-} = \{\bar{o}_1^i, \dots, \bar{o}_{\bar{n}_i}^i\}$, $\bar{a}_i^\alpha \in R_\alpha(A_n^-) = \{\bar{a}_1^\alpha, \dots, \bar{a}_{k_2}^\alpha\}$, and $\mathcal{I}_n^{set} = I(S_{set}(A_n^+)) + I(ub + lb - S_{set}(A_n^-))$.

If one or both relevant argument sets are empty the function is defined as follows:

$$S_{norm}(n) = \begin{cases} ub + lb - S_{set}(A_n^-) & \text{if } R_\alpha(A_n^+) = \emptyset \text{ and } R_\alpha(A_n^-) \neq \emptyset \\ S_{set}(A_n^+) & \text{if } R_\alpha(A_n^+) \neq \emptyset \text{ and } R_\alpha(A_n^-) = \emptyset \\ \text{not defined} & \text{if } R_\alpha(A_n^+) = \emptyset \text{ and } R_\alpha(A_n^-) = \emptyset \end{cases}$$

At this point, once we compute the collective support for a norm, we can decide whether the norm should be enacted or not. Given a predefined norm acceptance level μ , a norm will be enacted if $S_{norm}(n) > \mu$. For the norm to be enacted, its support should be laying on the positive side of the spectrum, hence μ should be picked so that $\mu \in (\frac{lb+ub}{2}, ub]$.

7 Case study: A virtual community

In this section we qualitatively compare the outcome of our norm support function with that of a naive average support function. This naive average support function obtains the support for a norm n as $S_{avg}(n) = \frac{1}{\dim(\mathbf{O}_{A_n})} \left(\sum_{i=1}^{|A_n^+|} \sum_{j=1}^{n_i} o_j^i + \sum_{i=1}^{|A_n^-|} \sum_{j=1}^{\bar{n}_i} ub + lb - \bar{o}_j^i \right)$. Our comparison encompasses a collection of Norm Argument Maps (NAM) that we characterise based on the opinions about their positive and negative arguments. Table 1 summarises the results of our comparison, which we detail next through some examples which invoke our norm support function⁵ with an opinion spectrum $\lambda = [1, 5]$, the importance function I , and a relevance level $\alpha = 0.3$.

NAM 1. Consider a norm n with one positive argument which is highly supported by opinions (e.g. with values 5, 5, 5), and three negative arguments that count on neutral supports (e.g. one with opinions 3.15, 3.2, 2.8; another one with opinions 3, 3.5, 2.6; and a third one with opinions 2.5, 3.5, 3.2). Thus, while the average support function would yield a rather neutral norm support ($S_{avg}(n) = 3.5375$), our norm support function would compute a strong support ($S_{norm}(n) = 4.9842$). Note that, since participants

⁵ This implementation is based on [6] and we have made it publicly available in [5].

Norm argument map	Argument sets		Norm support	
	Positive arguments	Negative arguments	S_{avg}	S_{norm}
NAM 1	one strong argument	several neutral arguments	neutral	strong
NAM 2	one neutral argument	one weak argument	strong	neutral
NAM 3	weak arguments	none	weak	undefined
NAM 4	none	weak arguments	strong	undefined
NAM 5	strong with few opinions	weak with lots of opinions	strong	undefined

Table 1. Norm supports computed by the average approach S_{avg} and our approach S_{norm} .

have not issued negative arguments that are strong enough to attack the norm, whereas they have found a strong argument to support it, the norm support should be favorable to the enacting of the norm. This is captured by our norm support function, while the average support function remains neutral. This happens because it is fundamental to weigh the importance of the arguments as well as the importance of the argument sets. In this way neutral arguments do not weigh much in the overall norm support. **NAM 2.** Consider the case of a norm with one positive argument with neutral opinions (e.g. 3.5, 3.25, 3.5, 3, 2.5) and one negative argument with a similar number of opinions but with weak support (e.g. opinions with values 1, 1, 1.2, 1.3, 1.25). The average support function would yield a strong support for the norm ($S_{avg}(n) = 4$) because of the weak support received by the negative argument. Unlike the average, our support function would obtain a neutral support for the norm ($S_{norm}(n) = 3.1731$) because the negative arguments are weakly supported and the positive one counts on neutral support.

The two cases above show that the norm support of our method is in line with positive arguments because negative arguments are not strong enough. The next three cases show the importance of counting on relevant arguments. **NAM 3.** Consider now a norm with weak positive arguments and no negative arguments. The average approach would yield weak support for the norm. However, notice that the lack of strongly-supported positive arguments does not imply that the norm is not good. If the norm was not good, we should expect that participants eventually issue strongly-supported negative arguments. Since there is not enough relevant information to decide whether the norm is good or not, the norm support would be undefined for our norm support function. This seems more reasonable than the weak support computed by the naive approach. **NAM 4.** Consider now the dual of our last NAM: a norm counts on weak negative arguments and no positive arguments. Here the average support function would obtain a strong norm support. Again, like in the previous case, our norm support function would be undefined, which seems more adequate due to the lack of relevance of the arguments issued so far. **NAM 5.** Consider the case of a norm with positive arguments, each one counting with a few high-valued opinions, and negative arguments, each one counting on a much larger number of low-valued opinions. The average support function $S_{avg}(n)$ would produce a strong norm support. However, notice that weak negative arguments should not favorably support a norm. Moreover, the positive arguments count on few opinions. If the norm was good enough, we should have expected to receive more supporting opinions, which is not the case. This is why our norm support function $S_{norm}(n)$ yields an undefined support. Overall, the three last examples show that the

lack of enough relevant information leads our norm support function to an undefined norm support, which seems more reasonable (and cautious) than that of a naive average support function.

7.1 A test with human users

We conducted a test to evaluate the functionality of the norm argument map. Our test encompassed eleven people debating on norms similar to the one in Figure 1 within a prototyped football social network. Users debated normally for several rounds and, afterwards, a satisfaction survey asked them if resulting aggregated ratings were reasonable. In a scale from 1 to 5, the answers' mean was 3.36, which we can consider as a positive preliminary result if we take into account the usability deficiencies of our prototype.

8 Conclusions and future work

To provide a more democratic way of moderating virtual communities, we propose a new argumentative structure, the so-called norm argument map. We also faced the problem of computing the collective support for a norm from the opinions of an argument's participants. We have identified two core concepts when computing a norm's support: the relevance of arguments and their importance. Thus, we argue that we must only consider relevant enough arguments and weigh opinions based on their importance (strength).

As to future work, we are currently working on identifying similar arguments that should be collapsed, but some other issues, such as when to close the argumentation process or how to define the norm acceptance level μ , still need to be studied. Moreover, we also plan to apply it to other social participation situations such as direct democracy.

Acknowledgments

Work funded by Spanish National project CollectiveWare code TIN2015-66863-C2-1-R (MINECO/FEDER).

References

1. Edmond Awad, Richard Booth, Fernando Tohmé, and Iyad Rahwan. Judgment aggregation in multi-agent argumentation. *CoRR*, abs/1405.6509, 2014.
2. Simone Gabbriellini and Paolo Torrioni. Microdebates: Structuring debates without a structuring tool. *AI Commun.*, 29(1):31–51, 2015.
3. Mark Klein. Enabling large-scale deliberation using attention-mediation metrics. *Computer Supported Cooperative Work*, 21(4-5):449–473, 2012.
4. Elinor Ostrom. *Governing the commons : the evolution of institutions for collective action*. Cambridge University Press, 1990.
5. Marc Serramia. Java implementation for the norm argument map. <https://bitbucket.org/msamsa/norm-argument-map.git>, 2016.

6. Vicenç Torra. Java implementation of the WOWA, OWA and WM aggregation operators. <http://www.mdai.cat/ifao/wowa.php>, 2000–2004.
7. Vicenç Torra and Yasuo Narukawa. *Modeling decisions - information fusion and aggregation operators*. Springer, 2007.
8. R. R. Yager. On ordered weighted averaging aggregation operators in multicriteria decision-making. *IEEE Transactions on Systems, Man, and Cybernetics*, 18(1):183–190, Jan 1988.