

Rating Comments on the Social Web using a Multi-Aspect Evaluation Framework

FORTH Technical Report TR 463, January 2016

Theodore Patkos
FORTH-ICS, Greece
patkos@ics.forth.gr

Antonis Bikakis
University College London, UK
a.bikakis@ucl.ac.uk

Giorgos Flouris
FORTH-ICS, Greece
fgeo@ics.forth.gr

Abstract

Users' reviews, comments and votes on the Social Web form the modern version of *word-of-mouth* communication, which has a huge impact on people's habits, businesses and the overall market. Nonetheless, there are only few attempts to formally model and analyse them using Computational Models of Argument, which achieved a first significant step in bringing these two fields closer. In this paper, we attempt their further integration by formalizing standard features of the Social Web, such as commentary and social voting, and by proposing methods for the evaluation of the comments' quality and acceptance. The proposed framework exhibits interesting formal properties, and is general enough to capture features of different review and debate websites.

Introduction

While a lot of research has been conducted on understanding and formalizing the interplay of arguments within the context of Computational Argumentation (Besnard and Hunter 2008), the types of arguments that populate the Social Web have not been formally studied and analyzed yet. These arguments usually have the form of comments, opinions or reviews, and are the main ingredients of online discussion forums, social networks, online rating and review sites, debate portals and other online communities - the electronic version of *word-of-mouth* communication. Since the emergence of the Social Web, their impact has been paramount to several aspects of human behavior, ranging from health-related (Chou et al. 2012; Shi, Messaris, and Cappella 2014), to buying (Cheung and Thadani 2012), travelling (Ye et al. 2011) or voting habits (Bond et al. 2012). They also significantly impact businesses and the overall market. Comments discussing products or brands have important consequences on the marketability of products themselves; as an example, a study on the impact of online reviews on the restaurant industry found that a one-star increase in Yelp.com rating led to a 5-9 percent increase in the revenue of the rated restaurants (Luca 2011).

In this paper, we are interested in investigating comments in online debates from the scope of two questions:

Q1. *How closely do the participants of an online debate share the opinion expressed by a given comment?*

Q2. *How helpful do the participants find a comment, and what exactly contributes to its helpfulness?*

Answering the first question is an attempt to value how universally acceptable the position reflected by the comment is, while the aim of the second question is to measure and analyze the comment's helpfulness or usefulness. Whereas traditional argumentation frameworks focus primarily on answering Q1, the majority of comment-enabled websites, such as eBay, Amazon, IMDb and others, partially answer Q2 by ranking comments based on the votes they have received. Regarding the second part of Q2, some empirical studies attempt to gain a deeper understanding of what makes online comments useful (Korfatis, García-Bariocanal, and Sánchez-Alonso 2012; Otterbacher, Hemphill, and Dekker 2011; Schindler and Bickart 2012; Willemsen et al. 2011), by identifying different content-related characteristics such as relevance, informativeness, clarity and conformity.

Recently, a number of frameworks have been suggested that aim to adapt the reasoning paradigm of Computational Argumentation to the needs of the Social Web. Initially conceived in the study of Leite and Martins (2011), these frameworks combine attacking and supporting relations among arguments with voting, which is a very popular practice in the Social Web. The result is a set of powerful frameworks and platforms, mostly aimed at enhancing question answering and decision making discussions (Evripidou and Toni 2014; Baroni et al. 2015). Nevertheless, these frameworks assign a single value when rating the strength of a comment, significantly restricting their capacity in handling common situations, such as the one illustrated below.

Example 1. In the imaginary SuggestYourWine forum, comment a_1 expresses the opinion that a particular Cabernet Sauvignon 1992 red Bordeaux wine is an excellent choice for stew dishes. The author of comment a_2 supports this position, by informing that she recently tasted it and noticed how nicely it paired with the Irish stew she cooked. Comment a_3 , written by what seems to be an expert in wines, further supports a_2 explaining that the full body of this type of grapes is a perfect match to dishes rich in fat and that the 1990s were golden years for Bordeaux wines. Another person attacks a_1 with opinion a_4 which states that consuming wine is a dangerous habit and should be taken with care. □

Although wine preference is largely a subjective matter,

one can expect that comments a_1, a_2, a_3 of Example 1 may eventually enjoy wide acceptance, as they refer to commonly held opinions about Bordeaux wines. Still, a_3 should stand in front of the other two in terms of quality or completeness, as it seems to express an expert and well-explained opinion. As for comment a_4 , although being true in principle, it does not seem relevant to the discussion and its attack should not significantly reduce the acceptance of a_1 .

Existing argumentation frameworks for the Social Web do not distinguish between the acceptance and the quality of arguments. They blend together the combined strength of attacking and supporting arguments with a fuzzy aggregation of votes, even though each of these features can carry different semantics that can lead to a more accurate valuation of arguments. Moreover, they have difficulty isolating irrelevant arguments (trolls) in an intuitive way.

In this paper, we suggest a generic framework that builds on and extends previous approaches, in order to accommodate comment evaluation for the Social Web. In particular:

- we formalize a framework that is flexible enough to model diverse features of comment-enabled sites, providing the machinery for extending them with new ones, if needed;
- we describe a simple mechanism that exploits users’ feedback, in order to distinguish between the score assigned to a comment for valuating the *position that it expresses* and that for valuating *how this position is presented*;
- we suggest a set of properties that guarantee an intuitive behavior for comment-enabled sites, and suggest a particular instantiation of our framework that satisfies these properties, yielding it suitable for a multitude of domains;
- we offer an intuitive means to isolate irrelevant or offensive comments.

In the rest of the paper, we first present the proposed framework, which consists of a model for online comments, social votes and their relationships, and a set of generic functions for calculating the comments’ quality and acceptance, and describe some desirable properties for these functions. We then describe an instantiation of the framework, determine its formal properties, and demonstrate its applicability using examples from the Social Web. Finally, we compare with related work from social argumentation, and we conclude. While this study focuses on comments appearing on the Web, we use the terms *comment* and *argument* interchangeably throughout the paper, due to the extensive use of theoretical models that have been developed within the context of Computational Argumentation.

Multi-Aspect Comment Evaluation

Our framework generalizes previous approaches in two ways. First, given an argument set \mathcal{A} , it assigns two different scores to characterize the strength of an argument $a \in \mathcal{A}$: the *quality score* ($QUA : \mathcal{A} \rightarrow \mathbb{I}$) and the *acceptance score* ($ACC : \mathcal{A} \rightarrow \mathbb{I}$). We assume $\mathbb{I} = [0, 1]$ for consistency with relevant literature, although different ranges can be used.

Second, it enables the definition of diverse criteria or *aspects* to calculate such scores, denoted as \mathcal{D}_{aspect} . Depending on the domain of interest, different aspects can be de-

finied. For instance, aspect \mathcal{D}_{rlv} may measure how relevant an argument is to the topic of a discussion; aspect \mathcal{D}_{rlb} may refer to how reliable and well-justified an argument is; aspect \mathcal{D}_{trt} may describe the degree to which an argument can be characterized as an “expert opinion”; and so on. Each of these aspects may influence the quality and acceptance score of a target argument in different ways. In order to calculate scores related to an aspect, one may decide to mix different features, such as positive votes, negative votes and/or supporting and attacking replies (i.e., other arguments). We next propose a generic scheme to formalize these notions.

Definition 1. An aspect \mathcal{D}_x corresponding to an argument set \mathcal{A} is a quadruple $\langle \mathcal{R}_x^{supp}, \mathcal{R}_x^{att}, V_x^+, V_x^- \rangle$, where $\mathcal{R}_x^{supp} \subseteq \mathcal{A} \times \mathcal{A}$ is a binary acyclic support relation on \mathcal{A} , $\mathcal{R}_x^{att} \subseteq \mathcal{A} \times \mathcal{A}$ is a binary acyclic attack relation on \mathcal{A} and $V_x^+ : \mathcal{A} \rightarrow \mathbb{N}^0$ and $V_x^- : \mathcal{A} \rightarrow \mathbb{N}^0$ are total functions mapping each argument to a number of positive and negative votes relative to this aspect, respectively.

The two relations are acyclic, as in online debates arguments are added in chronological order. The goal is to evaluate the strength of arguments considering one or more aspects. We propose a generic formal framework that can capture this intuition as follows:

Definition 2. An *mDiCE* (multi-Dimensional Comment Evaluation) framework is an $(N+1)$ -tuple $\langle \mathcal{A}, \mathcal{D}_{d1}, \dots, \mathcal{D}_{dN} \rangle$, where \mathcal{A} is a finite set of arguments and $\mathcal{D}_{d1}, \dots, \mathcal{D}_{dN}$ are aspects (dimensions), under which an argument is evaluated.

Using Definitions 1 and 2, we can formalize the forum of Example 1 as an mDiCE framework $\langle \mathcal{A}, \mathcal{D}_{crt}, \mathcal{D}_{inf}, \mathcal{D}_{rlv} \rangle$ where $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$ and \mathcal{D}_{crt} refers to correctness, \mathcal{D}_{inf} to informativeness, and \mathcal{D}_{rlv} to relevance. \mathcal{R}_{crt}^{supp} (the support relation with respect to correctness) contains $\{a_2, a_1\}$ and $\{a_3, a_2\}$, while \mathcal{R}_{crt}^{att} contains $\{a_4, a_1\}$. V_{inf}^+ is expected to assign a bigger value to a_3 compared to all other arguments assuming that participants will find it more informative, while $V_{rlv}^-(a_4)$ will probably be big assuming that many participants will find a_4 irrelevant. Note that the information related to attack/support with respect to each aspect should be provided by the users, via intuitive interfaces and interaction modes provided by the web site.

Not every aspect is relevant to all domains. For example, how recent a comment is may not be important when discussing about a music band or a movie, but when it comes to rating a product, the effect of outdated comments may need to be neutralized. By making explicit which aspect is being evaluated by users when placing votes or supporting/attacking arguments, we enable a much more accurate evaluation, avoiding the correlation of unrelated aspects.

The Blank Argument Metaphor

Another way in which our framework extends previous ones is by introducing the following intuition: if votes on some argument a denote answers to an - explicit or implicit - aspect-related question, e.g., “is this a helpful argument?”, they themselves express an opinion that can be represented as a supporting argument to the target argument with a measurable strength. Since this new argument has no actual content,

rather it shares the same content with the target argument, we name it *blank argument* of a on aspect x and denote it as \hat{a}_x . Note that an attack to a is also an attack to \hat{a}_x , since they both share the same content and rationale.

Example 2. Consider the graph shown on Figure 1(a) presenting a debate involving three arguments, where argument a_2 supports a_1 and argument a_3 attacks it. Each argument is annotated with the number of positive and negative votes it received (shown next to the boxes) and with two values denoting its quality (left box, also depicted by the portion of the painted area of each circle) and acceptance scores (right box, also depicted by the size of the circle). For simplicity, we assume a single aspect in this example.

In order to calculate the quality and acceptance scores for a_1 , we introduce the blank argument \hat{a}_1 of a_1 (Figure 1(b)). The next section presents details about the computation of scores and explains how this model results in a more intuitive behavior compared to previous ones. Notice also how \hat{a}_1 is attacked by a_3 resulting in a weak support to a_1 . \square

In accordance with Leite and Martins’ notion of *social support* (2011), where votes denote the support of the audience, our approach relies on the intuition that negative votes act as a means to weaken the support towards a and not as a way to strengthen the attack¹. This is ascribed to the fact that positive votes denote a more self-explanatory response to the aspect-related question than negative ones. Positive votes have a very clear semantics, signifying congruence with the comment in terms of content, justification and stance towards the topic of the discussion. Arguably, the ideal comment has only positive votes and no supporting arguments, as the latter should ideally be asserted only in order to add material or explain better the opinion stated.

Negative votes, on the other hand, are more ambiguous. It is not clear if the person submitting a negative vote disagrees with the position stated; or if she finds it poorly explained to stand as an acceptable position; or if she just feels uncertain whether the comment qualifies for the aspect it is being asked for. Indeed, some rating sites try to interpret the meaning of negative responses using follow-up clarification questions. In Apple’s App Store, for example, when a user reports a concern, she is also asked to provide a justification (e.g., if it contains offensive material, is off-topic, etc.).

As a result, in our framework the strength of a blank argument is associated with the degree with which people have identified themselves with the target argument. And this strength, as already explained, is affected by the combined strength of arguments attacking a , as formalized next.

Definition 3. Let \mathcal{F} be an mDiCE framework and $\mathcal{D}_x = \langle \mathcal{R}_x^{supp}, \mathcal{R}_x^{att}, V_x^+, V_x^- \rangle$ be an aspect of \mathcal{F} . For each argument $a \in \mathcal{A}$, we define \hat{a}_x a new argument related to \mathcal{D}_x , called the blank argument of a on x , such that

¹In this paper, we assume a positive stance when expressing aspect-related questions, i.e., positive votes weight in favor of the target argument. One can configure the scheme given in this paper for questions of the form “Is this argument out-of-date?”, where positive votes characterize a negative stance towards the argument, by either interpreting negative votes and replies as positive ones (and vice-versa) or by placing blank arguments in the attacking set.




Figure 1: For the calculation of scores in a simple comment exchange graph (a), mDiCE requires additional features (b).

- $V_x^+(\hat{a}_x) = V_x^+(a)$, $V_x^-(\hat{a}_x) = V_x^-(a)$,
- $(\hat{a}_x, a) \in \mathcal{R}_x^{supp}$, and
- for all $(a_i, a) \in \mathcal{R}_x^{att}$ it also holds that $(a_i, \hat{a}_x) \in \mathcal{R}_x^{att}$.

For notational convenience, we use $\hat{\mathcal{A}}$ to refer to the set of blank arguments of an mDiCE framework \mathcal{F} and $\tilde{\mathcal{A}}$ for the set of user-generated arguments ($\mathcal{A} = \hat{\mathcal{A}} \cup \tilde{\mathcal{A}}$). Moreover, given an aspect $\mathcal{D}_x = \langle \mathcal{R}_x^{supp}, \mathcal{R}_x^{att}, V_x^+, V_x^- \rangle$, we define the set of direct supporters of an argument $a \in \mathcal{A}$ as $\mathcal{R}_x^+(a) = \{a_i : (a_i, a) \in \mathcal{R}_x^{supp}\}$. Similarly, the set of direct attackers of a is defined as $\mathcal{R}_x^-(a) = \{a_i : (a_i, a) \in \mathcal{R}_x^{att}\}$.

The Set of mDiCE Aggregation Functions

To calculate the different strength scores, we define a set of aggregation functions, which are summarized in Table 1: the left column presents those that drive the process of calculating intermediate scores, whereas the right column involves the ones that can be used to rank arguments from the scope of the motivating questions considered.

Table 1: Overview of the mDiCE Functions

Internal Functions	External Functions
$s_x^{cng} : \mathcal{A} \rightarrow \mathbb{I}$	$QUA : \mathcal{A} \rightarrow \mathbb{I}$
$s_x^{dlg} : \mathcal{A} \rightarrow \mathbb{I}$	$ACC : \mathcal{A} \rightarrow \mathbb{I}$
$g^{cng} : \mathbb{I} \times \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{I}$	$g^{QUA} : \mathbb{I}^N \rightarrow \mathbb{I}$
$g^{dlg} : \mathbb{I} \times \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{I}$	$g^{ACC} : \mathbb{I}^N \rightarrow \mathbb{I}$
$s^{vot} : \mathbb{N}^0 \times \mathbb{N}^0 \rightarrow \mathbb{I}$	
$s^{set} : (\mathbb{N}^0)^{\mathbb{I}} \rightarrow \mathbb{I}$	

The two core functions that we use to characterize the strength of an argument, either blank or user-generated, given a certain aspect x , are the *congruence strength* and the *dialogue strength*, denoted as $s_x^{cng}()$ and $s_x^{dlg}()$, respectively. The former aims to reflect the degree of people’s compliance with an argument along the given aspect. As explained in the previous section, the voting mechanism can be employed for this purpose, therefore this score will also characterize the supporting strength of the blank argument. The latter function aims to reflect the combined strength of supporting and attacking arguments that are attached to the target argument, i.e., the dialogue that it generated.

We first define the congruence strength as follows:

Definition 4. Let $\mathcal{F} = \langle \mathcal{A}, \mathcal{D}_{d1}, \dots, \mathcal{D}_{dN} \rangle$ be an mDiCE framework and $\mathcal{D}_x = \langle \mathcal{R}_x^{supp}, \mathcal{R}_x^{att}, V_x^+, V_x^- \rangle$ be an aspect

of \mathcal{F} . The congruence strength $s_x^{cng} : \mathcal{A} \rightarrow \mathbb{I}$ of an argument $a \in \mathcal{A}$ over aspect \mathcal{D}_x is given by

$$s_x^{cng}(a) = g^{cng}(s^{vot}(V_x^+(a), V_x^-(a)), s^{set}(\{s_x^{dlg}(a_i) : a_i \in \mathcal{R}_x^+(a) \cap \tilde{\mathcal{A}}\}), s^{set}(\{s_x^{dlg}(a_j) : a_j \in \mathcal{R}_x^-(a) \cap \tilde{\mathcal{A}}\})) \quad (1)$$

with

- the generic score function $s^{vot} : \mathbb{N}^0 \times \mathbb{N}^0 \rightarrow \mathbb{I}$ valuating the strength of an argument considering its positive and negative votes;
- the generic score function $s^{set} : (\mathbb{N}^0)^{\mathbb{I}} \rightarrow \mathbb{I}$ valuating the combined dialogue strength of a set of arguments;
- the generic score function $g^{cng} : \mathbb{I} \times \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{I}$ valuating the congruence score of an argument, considering the aggregation of the strength of the votes, the positive and the attacking arguments.

That is, the congruence strength can be determined by aggregating the strength of votes, the strength of supporting arguments and that of attacking arguments. Typically, $g^{cng}(x_v, x_s, x_a)$ should lay more emphasis on x_v and x_a , as already described, increasing on x_v and decreasing on x_a ; however, we keep the function generic to allow its instantiation to vary from system to system.

Note that the domain of s^{set} is the set of multisets of numbers in \mathbb{I} . Moreover, although the congruence strength is defined for both blank and user-generated arguments, the valuation considers the dialogue strength of the underlying non-blank arguments only, avoiding redundancies. This dialogue strength is defined as follows:

Definition 5. Let $\mathcal{F} = \langle \mathcal{A}, \mathcal{D}_d, \dots, \mathcal{D}_{dN} \rangle$ be an mDiCE framework and $\mathcal{D}_x = \langle \mathcal{R}_x^{supp}, \mathcal{R}_x^{att}, V_x^+, V_x^- \rangle$ be an aspect of \mathcal{F} . The dialogue strength $s_x^{dlg} : \mathcal{A} \rightarrow \mathbb{I}$ of an argument $a \in \mathcal{A}$ over aspect \mathcal{D}_x is given by

$$s_x^{dlg}(a) = \begin{cases} g^{dlg}(s^{vot}(V_x^+(a), V_x^-(a)), s^{set}(\{s_x^{dlg}(a_i) : a_i \in \mathcal{R}_x^+(a)\}), s^{set}(\{s_x^{dlg}(a_j) : a_j \in \mathcal{R}_x^-(a)\})) \\ s_x^{cng}(a), \text{ if } a \in \hat{\mathcal{A}} \end{cases} \quad (2)$$

- with the generic score function $g^{dlg} : \mathbb{I} \times \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{I}$ valuating the dialogue strength of an argument for a given aspect, considering the aggregation of the strength of its votes, its supporting and its attacking arguments.

The idea is that the dialogue strength of blank arguments coincides with their congruence strength; for the rest, we can consider the aggregation of all supports and attacks that have been placed. In contrast with g^{cng} , $g^{dlg}(x_v, x_s, x_a)$ should lay more emphasis on x_s and x_a , increasing on x_s and decreasing on x_a . Different instantiations of these functions can be used; the ones that we present in the next section satisfy certain properties appropriate for many websites.

It is also clear that not all features are appropriate for every aspect. In a typical Facebook chat for instance, comments are characterized by their popularity based only on the

number of “likes” (i.e., positive votes) they attract. Most on-line rating websites, like Amazon or IMDb, measure “helpfulness” or “usefulness” of reviews based on both positive and negative answers. More sophisticated debate sites allow users to also attack or support a target argument, promoting a more lively and structured discussion. The developer may choose instantiations of g^{dlg} and g^{cng} considering only some of the features defined in an aspect.

Finally, by considering the strength of different aspects defined within a particular mDiCE framework, the external scores of an argument can be determined:

Definition 6. Let $\mathcal{F} = \langle \mathcal{A}, \mathcal{D}_{d1}, \dots, \mathcal{D}_{dN} \rangle$ be an mDiCE framework. The quality and acceptance scores of an argument $a \in \mathcal{A}$ is given by the functions $QUA : \mathcal{A} \rightarrow \mathbb{I}$ and $ACC : \mathcal{A} \rightarrow \mathbb{I}$, respectively, which aggregate the strength of its aspects, such that

$$QUA(a) = g^{QUA}(s_{d1}^{cng}(a), \dots, s_{dN}^{cng}(a)) \quad (3)$$

$$ACC(a) = g^{ACC}(s_{d1}^{dlg}(a), \dots, s_{dN}^{dlg}(a)) \quad (4)$$

with $g^{QUA}, g^{ACC} : \mathbb{I}^N \rightarrow \mathbb{I}$.

Desirable Properties

As is obvious from the above, our framework is generic enough to allow many different types of functions to be defined. Here, we discuss useful properties for such functions, properties that would guarantee a “reasonable” behaviour for the task at hand (e.g., one would expect that attacking arguments should reduce the acceptability score of the attacked arguments). Our properties are classified into *monotonicity* and *smoothness* requirements.

Monotonicity. This set of requirements is based on the idea that, regardless of how the effect of a new vote or argument is quantified, the relative effect should be fixed. For example, different instantiations may quantify the effect of a positive vote in a different way, but this effect should be non-negative regardless; similarly, the effect of a (set of) “strong” attacking argument(s) should be non-positive, and more powerful than the effect of a (set of) “weak” attacking argument(s). Formally, we require that:

- Function s^{vot} must be non-decreasing with respect to the number of positive votes, and non-increasing with respect to the number of negative votes.
- Function g^{dlg} must be non-decreasing with respect to the strength of votes and the strength of supporting arguments, and non-increasing with respect to the strength of attacking arguments. Similarly for g^{cng} , with the exception of supporting arguments, which is to the developer to decide whether they reduce the congruence strength.
- Functions g^{QUA}, g^{ACC} should be non-decreasing with respect to the strength of an argument over each aspect.

A similar condition should hold for function s^{set} , however this function applies over multisets, rather than numbers, therefore we require the following:

- If $A, B \in (\mathbb{N}^0)^{\mathbb{I}}$, $A \subseteq B$, then $s^{set}(A) \leq s^{set}(B)$.
- If $A, B, C \in (\mathbb{N}^0)^{\mathbb{I}}$, and $s^{set}(A) \leq s^{set}(B)$, then $s^{set}(C \uplus A) \leq s^{set}(C \uplus B)$.

The first condition guarantees that the addition of arguments cannot decrease the combined strength of a set of arguments. The second condition ensures that the addition of “stronger” arguments has a more powerful effect than the addition of “weaker” ones. Obviously, this behavior characterizes functions $s^{cng}(a)$, $s^{dlg}(a)$, $QUA(a)$, $ACC(a)$, as well. It also becomes clear from the above that when an argument’s acceptance score increases, this has a negative effect on the acceptance score of all the arguments it attacks, and a positive effect on the acceptance score of all the arguments it supports. This effect propagates along the tree of arguments using this pattern.

Smoothness. This set of requirements guarantees that the various functions will behave “smoothly”, i.e., small changes in some argument (e.g., a single new positive vote) cannot have large effects on the overall evaluation of arguments. This is an essential feature for any rating framework, in order to be adopted by the public, as the effect of an action on an argument’s acceptance and/or quality should be commensurate with the importance of the action itself; big (or very small) leaps that are not justified by the underlying changes may seem counterintuitive to users, causing them to lose their trust on the objectivity of the rating algorithms.

Before proceeding to the formal definitions, we will explain our intuition in a simple scenario: assume a function $f(x)$ with just one (real) variable. What we want is that when x changes by a certain amount, that the “distance” of the result $f(y)$ from $f(x)$ is somehow constrained by the “distance” between x and y . In other words we would like that, for some constant $\ell > 0$: $|f(y) - f(x)| \leq \ell \cdot |y - x|$.

The value of ℓ in the above inequality determines the “smoothness” of the function: a large ℓ implies that the function has at least some “abrupt” points, i.e., in our setting, that there are cases where simple (small) actions by the users (e.g., the addition of a few votes or a couple of attacking/supporting replies) would lead to major changes in the assessment result of the related arguments. On the other hand, small ℓ guarantees that a large number of changes are required to achieve a large effect on the assessment results, thus making the function more reluctant to change. To generalize the above notion to arbitrary functions and sets, we use *semi-metrics* to determine the “size” of a change:

Definition 7. Given a set S , a function $d_S : S \times S \mapsto \mathbb{R}$ is called a *semi-metric* for S iff for all $x, y \in S$: $d_S(x, y) \geq 0$, $d_S(x, y) = d_S(y, x)$ and $d_S(x, y) = 0$ iff $x = y$.

Definition 8. Consider two sets S, T equipped with semi-metrics d_S, d_T . A function $f : S \mapsto T$ is called ℓ -smooth (for d_S, d_T) iff $d_T(f(x), f(y)) \leq \ell \cdot d_S(x, y)$ for all $x, y \in S$.

The following theorem is trivial:

Theorem 1. If a function $f : S \mapsto T$ is ℓ -smooth (for d_S, d_T), then it is also ℓ' -smooth (for d_S, d_T), for all $\ell' \geq \ell$.²

Given a function f , when there is no ℓ such that f is ℓ -smooth, we will say that f is ∞ -smooth. Moreover, we will say that f is *exactly* ℓ -smooth when it is ℓ -smooth and there is no $\ell' < \ell$ such that f is ℓ' -smooth.

²All proofs appear at the Appendix of this Technical Report.

Instantiation of the mDiCE Score Functions

There are different ways to instantiate the aforementioned generic framework, given the objectives and scope of each system. We hereby suggest a set of instantiated functions that is broad enough for a multitude of domains, enjoying certain desirable properties.

Starting from the voting mechanism, commonly used by review sites to rank comments, many alternatives have been proposed. Average rating, though popular for its simplicity, is arguably a problematic solution, as for a example a comment with 1 positive and 0 negative votes has the same average rating with one with 100 positive and 0 negative votes.

A more intuitive behavior can be obtained by considering the Wilson Score Interval (Wilson 1927), which estimates the actual fraction of positive votes with a certain confidence degree, given the votes obtained up to the current point. This is a popular methodology among sites that wish to exploit user ratings. In accordance with similar frameworks, e.g., (Evrpidou and Toni 2014), we use the lower bound of this interval, given by the following formula:

$$ws(x, y) = \frac{n}{n + z^2} \left[\hat{p} + \frac{z^2}{2n} - z \sqrt{\frac{\hat{p}(1 - \hat{p})}{n} + \frac{z^2}{4n^2}} \right] \quad (5)$$

where $n = x + y$, $\hat{p} = x/n$ and $z = 1.96$ for a confidence level of 95% (see Fig. 2(a)). By means of this function, we calculate the voting strength of a comment as follows:

$$s^{vot}(v^+, v^-) = \begin{cases} 0, & \text{if } v^+ = v^- = 0 \\ ws(v^+, v^-), & \text{otherwise.} \end{cases} \quad (6)$$

In addition to aggregating votes, we also need to combine the strength of arguments into a single value. The T-CoNorm function $\perp_{sum} : \mathbb{I} \times \mathbb{I} \rightarrow \mathbb{I}$, also known as the probabilistic sum, satisfies a number of useful properties (Klement, Mesiar, and Pap 2000) and is defined as follows:

$$\perp_{sum}(x_1, x_2) = x_1 + x_2 - x_1 \cdot x_2 \quad (7)$$

For a multiset S of numbers, we define:


$$\perp_{sum}^*(S) = \begin{cases} 0, & \text{if } S = \emptyset \\ \perp_{sum}(x_1, \perp_{sum}^*({x_2, \dots, x_n})), & \text{if } S = \{x_1, x_2, \dots, x_n\} \text{ with } n > 0 \end{cases} \quad (8)$$

Consequently, we can instantiate the s^{set} function for the multiset of argument strengths as follows:

$$s^{set}(S) = \perp_{sum}^*({x_i : x_i \in S \text{ and } x_i \geq \vartheta}) \quad (9)$$

Constant ϑ can be used to discard arguments that fall below a given threshold, rendering them ineffective in changing the strength scores of other arguments. This way, irrelevant (troll) arguments can easily be neutralized. For $\vartheta = 0$, the identity element of \perp_{sum}^* is zero. As a result an argument with zero strength (e.g. according to Eq. (6) one that has received no votes) does not affect the overall strength of the set.

For aggregating vote and argument strengths, we use two simple functions that have been shown to work intuitively in practice. In particular, we ignore supporting arguments for



(a) Eq. (6) with $0 < x, y \leq 10^3$ and $0 \leq z \leq 1$. (b) Eq. (11) with $0 \leq x, y, z \leq 1$.

Figure 2: Plots for (a) $s^{vot}(x, y)$ and (b) $g^{dlg}(v, x, y)$.

estimating the congruence strength of an argument, permitting the strength of the attacking arguments to reduce appropriately the strength assigned by the votes:

$$g^{cng}(x_v, x_s, x_a) = x_v \cdot (1 - x_a) \quad (10)$$

Similarly, dialogue strength ignores votes, since their strength is reflected by the blank argument. An intuitive behavior for g^{dlg} would be to average the strength of supporting and attacking arguments, while maintaining smooth changes at the extremes. We suggest the following function:

$$g^{dlg}(x_v, x_s, x_a) = \frac{x_s^3 - x_a^3 - x_a \cdot x_s^3 + x_s \cdot x_a^3 + 1}{2} \quad (11)$$

This formulation infuses certain desirable properties to the behavior of comment-enabled systems. Notice that when the value of one parameter is small (e.g., supporting arguments have weak strength), a change of the other parameter causes a polynomial change in the function (Fig. 2(b)). This practically enforces the system to start trusting the opinion of users once some clear inclination towards an opinion starts to appear, either in favor or against it. In fact, the 3rd-degree (rather than exponential) polynomial achieves a steeper increase in the confidence as matters become clearer.

On the other hand, when the value of one parameter is big, changes in the other parameter have a linear effect. This way, we avoid being overly biased by strong opinions. Keep in mind that the debate forums we aim for are dynamic systems, where new information may arrive at different times. We want to allow new arguments that have not attracted a lot of support yet, to be able to challenge previously accepted positions, thus maintaining the liveliness of the discussion.

Ultimately, the goal is to allow each aspect to contribute to the quality and acceptance scores of an argument with a different degree of influence. We suggest a simple, weight-based solution that balances the influence among aspects, where w_i denote weights assigned by the system moderator based on other metrics or by experience:

$$g^{QUA}(x_1, \dots, x_n) = \sum_{i=1}^n w_i^{QUA} \cdot x_i, \text{ with } \sum_{i=1}^n w_i^{QUA} = 1 \quad (12)$$

$$g^{ACC}(x_1, \dots, x_n) = \sum_{i=1}^n w_i^{ACC} \cdot x_i, \text{ with } \sum_{i=1}^n w_i^{ACC} = 1 \quad (13)$$

Example 3 Consider argument a_1 in Figure 1. In order to determine its scores, we start by calculating the strength of the blank argument \hat{a}_1 of a_1 shown in Figure 1(b). From Definitions 4, 5, along with Eq. (6-11), we obtain

$$\begin{aligned} s^{cng}(\hat{a}_1) &= g^{cng}(s^{vot}(18, 16), s^{set}(\{\}), s^{set}(\{0.2\})) \\ &= g^{cng}(0.37, 0, 0.2) = 0.37 \cdot (1 - 0.2) = 0.3 \\ s^{dlg}(\hat{a}_1) &= s^{cng}(\hat{a}_1) = 0.3 \end{aligned}$$

Similarly, for argument a_1 we have

$$\begin{aligned} s^{cng}(a_1) &= g^{cng}(s^{vot}(18, 16), s^{set}(\{0.8, 0.3\}), s^{set}(\{0.2\})) \\ &= g^{cng}(0.37, 0.86, 0.2) = 0.37 \cdot (1 - 0.2) = 0.3 \\ s^{dlg}(a_1) &= g^{dlg}(0.37, 0.86, 0.2) \\ &= \frac{0.86^3 - 0.2^3 - 0.2 \cdot 0.86^3 + 0.86 \cdot 0.2^3 + 1}{2} \\ &= 0.75 \end{aligned}$$

Considering that for this example a single aspect has been used, we can assume, as shown in Figure 1(b), that

$$\begin{aligned} ACC(a_1) &= s^{dlg}(a_1) = 0.75 \\ QUA(a_1) &= s^{cng}(a_1) = 0.3 \end{aligned}$$

Apart from the simplicity of calculations, this example reveals also a notable behavior of the framework that differentiates it from previous approaches. While the opinion of users for a_1 , as expressed by the votes, is balanced, this argument receives a solid support by a_2 . As a result, the opinion expressed by a_1 enjoys a high acceptance score, given that the attack against it is considerably weak. The quality of a_1 though is mediocre; one should lean towards a_2 rather than a_1 in order to better understand the opinion they express.

State-of-the-art frameworks combine vote strength with argument strength resulting in a single score; in the case of a_1 above, this could cause misinterpretations. For instance, in (Evripidou and Toni 2014), as well as in (Baroni et al. 2015), the strength would be equal to 0.58, giving a significant reduction in the argument's acceptance.

Properties of the Proposed Functions

The aim of any Social Web platform should be to identify frameworks that yield intuitive results, an aspect that directly affects their understandability by the user, and therefore the trust that the user places on them. We claim that no given set of functions can be universally applicable, and that an appropriate fine-tuning is necessary to generate systems that are tailored to the particular needs of diverse Social Web platforms, exhibiting an intuitive behavior for the domain at hand. Here, we provide this understanding for the score functions proposed in the previous section. In the rest of this section, we assume a fixed aspect $\mathcal{D}_x = \langle \mathcal{R}_x^{supp}, \mathcal{R}_x^{att}, V_x^+, V_x^- \rangle$ of an mDiCE framework, and we will omit the x subscript in the various functions for simplicity.

Initially, it is trivial to show that all proposed functions have the monotonicity properties required. Moreover, Table 2 shows some characteristic results of the proposed framework for some special cases. For example, congruence does

not make sense if no positive votes have been placed or if the attacks on the target argument are very strong. Moreover, supporting and attacking arguments have a dual behavior.

Table 2: Special cases of the proposed functions

Conditions
$s^{set}(\emptyset) = 0$
$g^{cng}(0, x, y) = g^{cng}(x, y, 1) = 0$
$g^{cng}(1, 0, 0) = 1$ (the ideal argument in terms of quality)
$g^{dlg}(x, y, y) = 0.5$
$g^{dlg}(x, 1, 0.5) = 0.75, g^{dlg}(x, 0.5, 1) = 0.25$
$g^{QUA}(0.5, \dots, 0.5) = g^{ACC}(0.5, \dots, 0.5) = 0.5$

Another important aspect of the functions’ behavior is related to how “sensitive” the acceptance and quality scores of an argument are to different user actions; the notion of smoothness is appropriate for such quantification.

Note that different sites may have different requirements in this respect. As an example, increasing the smoothness of the used functions will allow sites that attract few users to lay more emphasis on maintaining the liveness of discussions by having users’ feedback cause reasonable, yet evident, effects in the course of the discussions; on the contrary, portals that lay emphasis on the reliability of the outcome of a debate on a particular topic, such as product/services rating sites, probably want to disallow small changes to significantly impact the outcome, in order to secure reliable results, therefore requiring functions that are less sensitive to user input. Our aim is to determine how different parameters affect smoothness in the proposed functions, enabling the fine-tuning of our framework to guarantee appropriate results.

To study smoothness, we will first define the semi-metrics required for the various sets to quantify the distance between two quantities, in a manner consistent with our intentions. First, we study \mathbb{I} , which is the range of all considered functions. In \mathbb{I} , we define $d_{\mathbb{I}}(x, y) = |x - y|$, i.e., the distance between two numbers is their difference.

For the set $\mathbb{N}^0 \times \mathbb{N}^0$ (used in s^{vot}), we define $d_{\mathbb{N}^0 \times \mathbb{N}^0}((x_1, y_1), (x_2, y_2)) = |x_1 - x_2| + |y_1 - y_2|$. In practice, this semi-metric specifies that the distance between two pairs of positive/negative votes is equal to the number of votes that one has to place (or retract) to reach from one to the other. This is reasonable, because in our case the “distance” is equal to the amount of operations performed by a user (e.g., placing a vote or adding a new argument).

Theorem 2. *Function s^{vot} is exactly $\frac{1}{1+z^2}$ -smooth.*

Theorem 2 implies that a new vote (positive or negative) cannot modify the value of s^{vot} by more than $\frac{1}{1+z^2}$. In fact, it can be shown that this extreme is reached only for the first positive vote placed; all subsequent votes have strictly smaller effects (Fig. 2(a)). Note how different z cause different smoothness properties on s^{vot} .

For s^{set} we need to define a semi-metric for $(\mathbb{N}^0)^{\mathbb{I}}$; our notion of distance will be based on the strength (based on s^{set}) of the symmetric difference between the sets compared, in particular: $d_{(\mathbb{N}^0)^{\mathbb{I}}}(X, Y) = s^{set}(X \setminus Y \uplus Y \setminus X)$.

Theorem 3. *For all $S_1, S_2, S_3 \in (\mathbb{N}^0)^{\mathbb{I}}$, it holds that: $|s^{set}(S_1 \uplus S_2) - s^{set}(S_1 \uplus S_3)| = (1 - s^{set}(S_1)) \cdot |s^{set}(S_2) - s^{set}(S_3)|$.*

Theorem 4. *Function s^{set} is exactly 1-smooth.*

Note that Theorem 4 is a direct corollary of Theorem 3. A closer look on those theorems implies that, when no arguments are present, the addition of a new one would cause s^{set} to change linearly (i.e., by a value equal to the strength of the added argument); on the contrary, subsequent additions have sublinear effects, namely analogous to $1 - s^{set}(S)$, where S is the current set of argument strengths.

To study g^{cng}, g^{dlg} , we will define the following semi-metric over \mathbb{I}^3 : $d_{\mathbb{I}^3}((x_1, y_1, z_1), (x_2, y_2, z_2)) = |x_1 - x_2| + |y_1 - y_2| + |z_1 - z_2|$. Now the following can be shown:

Theorem 5. *Function g^{cng} is exactly 1-smooth, whereas g^{dlg} is exactly $\frac{3}{2}$ -smooth.*

Thus, g^{cng} is (at most) linearly affected by changes in its input (votes and attacking arguments). The maximum effects on $g^{cng}(x_v, x_s, x_a)$ appear when either x_v is close to 1 and x_a changes; or when x_a is close to 0 and x_v changes. This observation also implies that the effect of an action (addition of a vote or attacking argument) on g^{cng} increases when this action is contrary to the current sentiment.

Similarly, the maximum effect on g^{dlg} is reached when the current sentiment is very positive (close to 1) or very negative (close to 0) and someone casts an opposing argument (Fig. 2(b)). This means that it is easier to cast doubts on the strength of a strong argument, than to quickly trust a doubtful one. As explained in the previous section, this aims at promoting the liveness of the discussion without damaging the credibility of conclusions; yet, one can easily adapt this behavior by changing the degree of the polynomial in Eq. (11) or by applying a different function altogether.

Finally, in order to study g^{QUA}, g^{ACC} we define semi-metrics over \mathbb{I}^n as follows:

$$d_{\mathbb{I}^n}(\langle x_1, \dots, x_n \rangle, \langle y_1, \dots, y_n \rangle) = \sum_{i=1, \dots, n} |x_i - y_i|.$$

Theorem 6. *Functions g^{QUA}, g^{ACC} are exactly w_M -smooth, where $w_M = \max\{w_i\}$ (for said function).*

An interesting conclusion from Theorem 6 is that “balanced” functions (where w_i are close or equal) are smoother, i.e., less sensitive to input changes. Moreover, as expected, the largest effects appear when the changed aspects are those that have the largest weights. Therefore, one needs to carefully balance the different aspects; in the attempt to achieve accurate results by laying more trust on certain aspects than others, intuitiveness of the framework may be lost.

Next, we study how we can combine the above results to “predict” the maximum possible effect of an action (addition of votes and/or arguments) on the acceptability and quality score of a given argument. This will enable site owners to properly adapt the behavior of their system based on their estimate of the number of users they attract; when the audience is small, high smoothness values can be tolerated, but as more users participate in a debate high smoothness will cause the behavior to seem unreasonable.

In the following theorems, we assume an mDiCE framework, instantiated using the functions presented earlier, and

an argument a such that $ACC(a) = u_a$, $QUA(a) = u_q$. We set w_a, w_q the maximum weights appearing in ACC , QUA respectively and z the constant appearing in s^{vot} .

Theorem 7. *If n_v new votes are added on a , and the new scores are $ACC(a) = u'_a$ and $QUA(a) = u'_q$, then: $|u'_a - u_a| \leq \frac{3 \cdot w_a \cdot n_v}{2 \cdot (1+z^2)}$ and $|u'_q - u_q| \leq \frac{w_q \cdot n_v}{1+z^2}$.*

Theorem 8. *If a set S of supporting arguments are added on a , such that $s^{set}(A) = s$, and the new scores for a are $ACC(a) = u'_a$ and $QUA(a) = u'_q$, then: $|u'_a - u_a| \leq \frac{3w_a \cdot s}{2}$ and $u'_q = u_q$. If, on the other hand, S is a set of attacking arguments, then: $|u'_a - u_a| \leq 3 \cdot w_a \cdot s$ and $|u'_q - u_q| \leq s \cdot w_q$.*

When a user performs multiple actions (e.g., adds both an attacking argument and a vote), the total deviation of each function (ACC , QUA) is at most equal to the sum of the bounds described in Theorems 7, 8:

Theorem 9. *If n_v new votes, and two sets S_1, S_2 of supporting/attacking (respectively) arguments are added on a , such that $s^{set}(S_1) = s_1$, $s^{set}(S_2) = s_2$, and that the new scores for a are $ACC(a) = u'_a$ and $QUA(a) = u'_q$, then: $|u'_a - u_a| \leq \frac{3 \cdot w_a \cdot n_v}{2 \cdot (1+z^2)} + \frac{3 \cdot w_a \cdot s_1}{2} + 3 \cdot w_a \cdot s_2$ and $|u'_q - u_q| \leq \frac{w_q \cdot n_v}{(1+z^2)} + w_q \cdot s_2$.*

Obviously, these are the maximum possible effects, which are reached for the special cases described above. It is also interesting to study how an action's effects propagate along the tree of arguments, as the acceptance of an argument affects other arguments that it attacks/supports:

Theorem 10. *Consider a sequence of arguments a_1, \dots, a_n , a such that for each $i \in [1, n-1]$, a_i either supports or attacks a_{i+1} , a_n supports or attacks a , and the total numbers of supports/attacks in the chain are m_S, m_A respectively. If the acceptance score of a_1 changes by s , then the acceptance and quality scores of a will change respectively from u_a to u'_a , and from u_q to u'_q as follows: $|u'_a - u_a| \leq s \cdot (\frac{3 \cdot w_a}{2})^{m_S} \cdot (3 \cdot w_a)^{m_A}$ and $|u'_q - u_q| \leq s \cdot w_q^{m_A}$.*

Attacking arguments provide a whole new view on the discussion, which explains why their effect is (generally) larger than supporting ones. Similarly, the effect of a vote is even smaller, as a vote has no argumentative character.

Another interesting observation is that the effects are generally maximized in early interactions (e.g., adding votes/arguments when no votes/arguments are present). This is expected, because, for example, the effect of a new vote over the strength of an argument that has already 1000 votes should be smaller than the effect of a new vote over an argument that has a small number of votes. To formalize this effect, one could directly use the smoothness concept, but with a slightly different set of semi-metrics, which would quantify change in relative (rather than absolute) terms. For example, the "relativized" version of the semi-metric over $\mathbb{N}^0 \times \mathbb{N}^0$ would be: $d'_{\mathbb{N}^0 \times \mathbb{N}^0}((x_1, y_1), (x_2, y_2)) = 2 \cdot \frac{|x_1 - x_2| + |y_1 - y_2|}{x_1 + x_2 + y_1 + y_2}$, which divides "absolute" distance with the average number of votes (before and after the change) to get a "relative" distance³. Studying this type of smoothness

³Set $d'_{\mathbb{N}^0 \times \mathbb{N}^0}((0, 0), (0, 0)) = 0$, when all variables are 0.

is reserved for future work.

Application of mDiCE on the Social Web

To demonstrate the applicability of the mDiCE framework on the Social Web, we show how existing review and debate web sites can be modeled as mDiCE frameworks. Based on the features they support, we group them in three categories.

Single-aspect voting-based sites. The first category is characterized by the following features: users are able to vote on the overall quality/helpfulness of a review or comment; they may also reply to other reviews, but not to explicitly support or dispute them. It includes most of the available online review sites, such as Amazon, IMDb, TripAdvisor, App Store and Google Play. These sites may differ in the description of the aspect that the users are asked to vote on (e.g. helpfulness, usefulness, etc.). However they can all be modeled in a similar way: as mDiCE frameworks with a single aspect, with two functions returning for each comment the number of positive and negative votes respectively, and empty support and attack relations: $\langle \mathcal{A}, \mathcal{D}_{hlp} \rangle$, where

$$\mathcal{D}_{hlp} = \langle \emptyset, \emptyset, V_{hlp}^+, V_{hlp}^- \rangle$$

Multiple-aspect voting-based sites. Sites in the second category enable users to vote on multiple aspects of a review or comment. Similarly with the first category, users may respond to other people's comments, but there is no distinction between supporting and attacking comments. Two sites in this category are Yelp⁴, which publishes customer reviews about local businesses; and Slashdot⁵, which features news stories on science and technology that are submitted and evaluated by its users. In Yelp, users can state whether they find a review "useful", "funny" or "cool". Slashdot users can annotate their votes with one or more tags, such as "offtopic", "flamebait", "troll", "redundant", "insightful", "interesting", "informative", "funny", "overrated" etc.

We can model sites in this category as mDiCE frameworks with multiple aspects. For example, Slashdot can be modeled as a mDiCE framework with three aspects: relevance, informativeness and underestimate:

$$\langle \mathcal{A}, \mathcal{D}_{rlv}, \mathcal{D}_{inf}, \mathcal{D}_{est} \rangle$$

The three aspects of the framework are defined as follows:

$$\begin{aligned} \mathcal{D}_{rlv} &= \langle \emptyset, \emptyset, V_{rlv}^+, V_{rlv}^- \rangle, \\ \mathcal{D}_{inf} &= \langle \emptyset, \emptyset, V_{inf}^+, V_{inf}^- \rangle, \\ \mathcal{D}_{est} &= \langle \emptyset, \emptyset, V_{est}^+, V_{est}^- \rangle \end{aligned}$$

where for each comment V_{rlv}^+ refers to votes tagged as interesting, V_{rlv}^- those tagged as offtopic, troll or flambait, V_{inf}^+ the insightful or informative, V_{inf}^- the redundant, V_{est}^+ the underrated, and V_{est}^- the overrated.

⁴<http://www.yelp.com>

⁵<http://slashdot.org>

Debate-based sites The distinctive feature of sites in the third category is that they enable users to explicitly state whether their comments support or attack other users' comments. One such example is CreateDebate⁶, a portal for online debates. Users can participate in the debates by posting arguments that support or dispute other users' arguments, by voting up or down arguments or by asking other users to clarify their arguments.

We can model CreateDebate as a mDiCE framework with two aspects, \mathcal{D}_{agr} and \mathcal{D}_{clr} , the first one referring to the level to which the participants agree with the content of an argument, and the second referring to the clarity of an argument:

$$\mathcal{D}_{agr} = \langle \mathcal{R}_{agr}^{supp}, \mathcal{R}_{agr}^{att}, V_{agr}^+, V_{agr}^- \rangle$$

$$\mathcal{D}_{clr} = \langle \emptyset, \mathcal{R}_{clr}^{att}, V_{clr}^+, \emptyset \rangle$$

where \mathcal{R}_{agr}^{supp} contains all pairs of arguments (b, a) such that b supports a , \mathcal{R}_{agr}^{att} contains all pairs (c, a) such that c disputes a , V_{agr}^+ returns for each argument the number of positive votes, V_{agr}^- the number of negative votes, \mathcal{R}_{clr}^{att} contains all pairs of arguments (d, a) such that d requests clarification of a , and $V_{clr}^+ = V_{agr}^+$ (based on the intuition that participants who vote positively for an argument also find it clear).

Exploiting the mDiCE features in the Social Web

Given the range of features supported by the mDiCE framework, it is obvious that the above sites can be extended by enabling users to vote or argue on multiple aspects of their comments. Such extensions are consistent with the findings of several empirical studies on the use of social comments and votes. For example, in the user study presented in (Otterbacher, Hemphill, and Dekker 2011), where participants were asked to rate reviews and comments posted on Amazon.com and IMDb.com, most of them agreed that "helpful reviews were clearly written, were relevant to their information needs, contained an appropriate amount of information, and conformed to the general opinion of the item reviewed". Based on this, one obvious way to extend these two sites is by allowing users to vote on different aspects of a comment: clarity, relevance, informativeness and conformance to the general opinion. The latter is similar to the notion of acceptance used in our framework. It can therefore be evaluated by allowing users to explicitly indicate whether their comments support or attack those other comments that they reply to.

Related Work

Computational Argumentation has found fruitful application in Artificial Intelligence (Rahwan and Simari 2009) and decision support systems (Lucas Carstens 2015). Recently, it was also applied to the Social Web for better filtering and ranking online arguments, using bipolar frameworks (Amgoud et al. 2008; Cayrol and Lagasque-Schiex 2005) which consider both supporting and attacking arguments.

The first framework that adapted the basic tenets of abstract argumentation to the features and needs of the Social Web was the Social Abstract Argumentation Framework (SAF) proposed by Leite and Martins (2011). SAF

⁶<http://www.createdebate.com>

assesses quantitatively the strength of arguments by combining the votes with the attacks that an argument has received. Although it does not accommodate supporting arguments, the framework is generic enough to model dialogues represented as directed graphs, whereas most other frameworks, including mDiCE, concentrate on trees. Different functions can be used to instantiate the framework, motivating much of the research that followed this study. The framework was later extended to accommodate the valuation of strength on attacks (Eilmez, Martins, and Leite 2014), in an attempt to isolate irrelevant arguments. Although this practice may seem confusing to users of the Social Web, it spots a crucial issue that most similar frameworks fail to handle.

An initial attempt to incorporate supporting arguments in social argumentation was presented in (Evripidou and Toni 2012). Their framework relies on the notion of treating supporting arguments as negative attacks. The paper describes the main ideas and an example application, but doesn't provide any formal definitions or properties of the framework.

Very close to our study are the frameworks proposed in (Evripidou and Toni 2014) and (Baroni et al. 2015), which were applied for the social debating platform *Quaestio-it*⁷, and the issue-based information system *designVUE*⁸, respectively. They both extend bipolar argumentation with the notion of social support and employ the idea that supporting/attacking arguments increase/decrease social support (i.e., the strength of votes) in a symmetric manner. These frameworks are applied on debates with a tree structure, and satisfy many intuitive properties. However, the aggregation of supporting and attacking arguments is defined in a way that induces discontinuity in certain cases.

All these frameworks aim to assign a single score to characterize the strength of an argument. Our framework generalizes them, combining their strong points with a novel, multi-aspect evaluation of comments, which models more accurately debates such as the one described in Example 1.

Conclusion

Despite the profound impact of online comments and reviews in the market as well as in several of our daily habits, there have been only few recent attempts to formally model and evaluate them. Building on previous work on social abstract argumentation, we proposed a formal multi-dimensional framework for evaluating online comments taking into account the responses and votes that they receive from users of the Social Web. Compared to previous efforts, we distinguish between the quality and the acceptance of a comment, and we consider different ways with which web users may assess a comment in a certain discussion.

Our plans for future work include a further generalization of the proposed framework to support different kinds of user ratings, e.g., star-ratings instead of boolean votes. This, for example, would enable someone to state how useful a comment is, instead of just saying whether a comment is useful or not. We also plan to study alternative instantiations of the

⁷<http://www.quaestio-it.com/>

⁸<http://www3.imperial.ac.uk/designengineering/tools/designvue>

framework and evaluate them with real data such as reviews, comments and votes from the Yelp dataset⁹.

References

- Amgoud, L.; Cayrol, C.; Lagasquie-Schiex, M.-C.; and Livet, P. 2008. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems, Bipolar Representations of Information and Preference (Part 2: reasoning and learning)* 23(10):1062–1093.
- Baroni, P.; Romano, M.; Toni, F.; Aurisicchio, M.; and Bertanza, G. 2015. Automatic evaluation of design alternatives with quantitative argumentation. *Argument & Computation* 6(1):24–49.
- Besnard, P., and Hunter, A. 2008. *Elements of Argumentation*. The MIT Press.
- Bond, R. M.; Fariss, C. J.; Jones, J. J.; Kramer, A. D. I.; Marlow, C.; Settle, J. E.; and Fowler, J. H. 2012. A 61-million-person experiment in social influence and political mobilization. *Nature* 489(7415):295–298.
- Cayrol, C., and Lagasquie-Schiex, M. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, volume 3571 of *Lecture Notes in Computer Science*. 378–389.
- Cheung, C. M., and Thadani, D. R. 2012. The impact of electronic word-of-mouth communication: A literature analysis and integrative model. *Decision Support Systems* 54(1):461 – 470.
- Chou, W.-y. S.; Prestin, A.; Lyons, C.; and Wen, K.-y. 2012. Web 2.0 for health promotion: Reviewing the current evidence. *American Journal of Public Health* 103(1):e9–e18.
- Eilmez, S.; Martins, J. a.; and Leite, J. a. 2014. Extending social abstract argumentation with votes on attacks. In Black, E.; Modgil, S.; and Oren, N., eds., *Theory and Applications of Formal Argumentation*, volume 8306 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 16–31.
- Evripidou, V., and Toni, F. 2012. Argumentation and voting for an intelligent user empowering business directory on the web. In Krtzsch, M., and Straccia, U., eds., *Web Reasoning and Rule Systems*, volume 7497 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 209–212.
- Evripidou, V., and Toni, F. 2014. Quaestio-it.com: a social intelligent debating platform. *Journal of Decision Systems* 23(3):333–349.
- Klement, E. P.; Mesiar, R.; and Pap, E. 2000. *Triangular Norms*. Springer, 1 edition.
- Korfiatis, N.; García-Bariocanal, E.; and Sánchez-Alonso, S. 2012. Evaluating content quality and helpfulness of online product reviews: The interplay of review helpfulness vs. review content. *Electronic Commerce Research and Applications* 11(3):205 – 217.
- Leite, J. a., and Martins, J. a. 2011. Social abstract argumentation. In *Proceedings of the Twenty-Second International*

Joint Conference on Artificial Intelligence - Volume Volume Three, IJCAI’11, 2287–2292. AAAI Press.

Luca, M. 2011. Reviews, reputation, and revenue: The case of yelp.com. Technical Report 12-016, Harvard Business School.

Lucas Carstens, Xiuyi Fan, Y. G. F. T. 2015. An overview of argumentation frameworks for decision support. In *The Fourth IJCAI International Workshop on Graph Structures for Knowledge Representation and Reasoning. To appear*, GKR 2015.

Otterbacher, J.; Hemphill, L.; and Dekker, E. 2011. Helpful to you is useful to me: The use and interpretation of social voting. *Proceedings of the American Society for Information Science and Technology* 6(1):1–10.

Rahwan, I., and Simari, G. R. 2009. *Argumentation in Artificial Intelligence*. Springer Publishing Company, Incorporated, 1st edition.

Schindler, R. M., and Bickart, B. 2012. Perceived helpfulness of online consumer reviews : the role of message content and style. *Journal of Consumer Behaviour* 11(3, (5/6)):234–243.

Shi, R.; Messaris, P.; and Cappella, J. N. 2014. Effects of online comments on smokers’ perception of antismoking public service announcements. *Journal of Computer-Mediated Communication* 19(4):975–990.

Willemsen, L. M.; Neijens, P. C.; Bronner, F.; and de Ridder, J. A. 2011. “highly recommended!” the content characteristics and perceived usefulness of online consumer reviews. *Journal of Computer-Mediated Communication* 17(1):19–38.

Wilson, E. B. 1927. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association* 22(158):209–212.

Ye, Q.; Law, R.; Gu, B.; and Chen, W. 2011. The influence of user-generated content on traveler behavior: An empirical investigation on the effects of e-word-of-mouth to hotel online bookings. *Computers in Human Behavior* 27(2):634–639.

Proof Sketches for Theorems

Proof for Theorem 1.

Obvious from Definition 8. □

Proof for Theorem 2.

It suffices to show that: $\frac{|s^{vot}(v_1^+, v_1^-) - s^{vot}(v_2^+, v_2^-)|}{|v_1^+ - v_2^+| + |v_1^- - v_2^-|} \leq \frac{1}{1+z^2}$ for all $v_1^+, v_2^+, v_1^-, v_2^- \in \mathbb{N}^0$. To do that, we will identify the maximum value that the quantity $\frac{|s^{vot}(v_1^+, v_1^-) - s^{vot}(v_2^+, v_2^-)|}{|v_1^+ - v_2^+| + |v_1^- - v_2^-|}$ can take. Obviously, the relation holds when $v_1^+ = v_2^+, v_1^- = v_2^-$. We will show it for the other cases.

Case 1: We will initially assume that $v_1^+ = v_2^+ = x, v_1^- = y, v_2^- = y + \epsilon$ for some $\epsilon > 0$. We set $g_1(x, y, \epsilon) = \frac{|s^{vot}(x, y) - s^{vot}(x, y + \epsilon)|}{\epsilon}$. We will show that $g_1(x, y, \epsilon) \leq \frac{1}{1+z^2}$ for all x, y, ϵ . For $x = 0$ it is trivial to show that the relation holds. For $x > 0$, we compute $\frac{\partial g_1}{\partial y}$ and find that

⁹http://www.yelp.com/dataset_challenge

$\frac{\partial g_1}{\partial y} \leq 0$ (using the monotonicity properties of s^{vot}). Therefore $g_1(x, 0, \epsilon) \geq g_1(x, y, \epsilon)$ for all x, y, ϵ . Then, we set $g_2(x, \epsilon) = g_1(x, 0, \epsilon)$, and determine that it is decreasing with respect to ϵ , so $g_2(x, 1) \geq g_2(x, \epsilon)$ for all x, ϵ . Finally, we set $g_3(x) = g_2(x, 1)$ and determine that $g_3(x) \leq \frac{1}{2 \cdot z \cdot \sqrt{1+z^2} + 2 \cdot z^2 + 1} \leq \frac{1}{1+z^2}$. So, the theorem holds for this case (the case where $v_1^+ = v_2^+ = x, v_1^- = y + \epsilon, v_2^- = y$ for some $\epsilon > 0$ is symmetric).

Case 2: We will now assume that $v_1^+ = x + \epsilon, v_2^+ = x, v_1^- = v_2^- = y$ for some $\epsilon > 0$. To address this case, we first establish that $s^{vot}(v^+, v^-) = s^{vot}(v^-, v^+) + \frac{v^+ - v^-}{v^+ + v^- + z^2}$. Thus our initial formula becomes for this case: $\frac{|s^{vot}(x+\epsilon, y) - s^{vot}(x, y)|}{\epsilon}$. Using the monotonicity of s^{vot} and the previous result, this quantity can be written as: $\frac{s^{vot}(y, x+\epsilon) - s^{vot}(y, x)}{\epsilon} + \frac{x-y+\epsilon}{\epsilon \cdot (x+y+z^2+\epsilon)} + \frac{x-y}{\epsilon \cdot (x+y+z^2)}$. Using partial differentials with respect to x, ϵ we establish that the quantity $\frac{x-y+\epsilon}{\epsilon \cdot (x+y+z^2+\epsilon)} + \frac{x-y}{\epsilon \cdot (x+y+z^2)} \leq \frac{1}{1+z^2}$. Moreover, $\frac{x-y+\epsilon}{\epsilon \cdot (x+y+z^2+\epsilon)} + \frac{x-y}{\epsilon \cdot (x+y+z^2)} \geq 0$. Similarly, using the previous result (case 1) and the monotonicity properties of s^{vot} , we can easily show that $-\frac{1}{1+z^2} \leq \frac{s^{vot}(y, x+\epsilon) - s^{vot}(y, x)}{\epsilon} \leq 0$. Thus, the initial quantity is necessarily confined in $[-\frac{1}{1+z^2}, \frac{1}{1+z^2}]$, which completes the proof for this case also.

Case 3: As a final case, we will consider the most general case where both variables of s^{vot} fluctuate ($v_1^+ \neq v_2^+$ and $v_1^- \neq v_2^-$). This case is split into 4 sub-cases:

Case 3.1: Assume initially that $v_1^+ > v_2^+$ and $v_1^- < v_2^-$. Then, using the monotonicity properties of s^{vot} : $|s^{vot}(v_1^+, v_1^-) - s^{vot}(v_2^+, v_2^-)| = s^{vot}(v_1^+, v_1^-) - s^{vot}(v_1^+, v_2^-) + s^{vot}(v_1^+, v_2^-) - s^{vot}(v_2^+, v_2^-) =$, so, the case is proved using the previous results.

Case 3.2: Assume now that $v_1^+ > v_2^+$ and $v_1^- > v_2^-$. Then, using the monotonicity properties of s^{vot} : $|s^{vot}(v_1^+, v_1^-) - s^{vot}(v_2^+, v_2^-)| = s^{vot}(v_1^+, v_1^-) - s^{vot}(v_2^+, v_2^-) \leq s^{vot}(v_1^+, v_2^-) - s^{vot}(v_2^+, v_2^-)$, and the result follows from case 1.

Case 3.3: The case where $v_1^+ < v_2^+$ and $v_1^- < v_2^-$, is similar to case 3.2 (we replace $s^{vot}(v_1^+, v_1^-)$ with $s^{vot}(v_2^+, v_1^-)$).

Case 3.4: Finally, the case where $v_1^+ < v_2^+$ and $v_1^- > v_2^-$, is handled similarly to case 3.1, by adding and subtracting $s^{vot}(v_2^+, v_1^-)$.

The above steps have shown that s^{vot} is $\frac{1}{1+z^2}$ -smooth; to show that it is exactly $\frac{1}{1+z^2}$ -smooth, set $v_1^+ = v_2^+ = v_1^- = 0$ and $v_2^- = 1$. \square

Lemma 1. For all finite $S_1, S_2 \in (\mathbb{N}^0)^\mathbb{I}$, it holds that $s^{set}(S_1 \uplus S_2) = s^{set}(S_1) + s^{set}(S_2) - s^{set}(S_1) \cdot s^{set}(S_2)$.

Proof for Lemma 1.

We will use induction on the size of S_2 . It is easy to show the cases where $|S_2| = 0$ or $|S_2| = 1$. Suppose that it holds for $|S_2| < n$ ($n > 1$), we will show it for $|S_2| = n$. Take any $x \in S_2$, and set $S_2' = S_2 \setminus \{x\}$. Then: $s^{set}(S_1 \uplus S_2) = s^{set}((S_1 \uplus S_2') \uplus \{x\})$. Using the definition of s^{set} and the induction hypothesis on $S_1 \uplus S_2'$, it is easy to

show the result. \square

Lemma 2. For all finite $S_1, S_2 \in (\mathbb{N}^0)^\mathbb{I}$, $S_1 \subseteq S_2$ implies that $s^{set}(S_1) \leq s^{set}(S_2)$.

Proof for Lemma 2.

We set $S_2 = S_1 \uplus (S_2 \setminus S_1)$ and apply Lemma 1. The result is now obvious by the fact that $s^{set}(S_2 \setminus S_1) \leq 1$. \square

Proof for Theorem 3.

Trivial using Lemma 1. \square

Proof for Theorem 4.

It suffices to show that, for any given $S_1, S_2 \in (\mathbb{N}^0)^\mathbb{I}$, it holds that $|s^{set}(S_1) - s^{set}(S_2)| \leq 1 \cdot s^{set}((S_1 \setminus S_2) \uplus (S_2 \setminus S_1))$. Take $S = S_1 \cap S_2$. Now the left-hand side formula becomes:

$$|s^{set}(S_1) - s^{set}(S_2)| = (\text{by definition}) \\ |s^{set}(S \cup (S_1 \setminus S_2)) - s^{set}(S \cup (S_2 \setminus S_1))| = (\text{Theorem 3}) \\ |(1 - s^{set}(S)) \cdot (s^{set}(S_1 \setminus S_2) - s^{set}(S_2 \setminus S_1))|.$$

Applying Theorem 3 and the fact that $0 \leq s^{set}(S_1 \setminus S_2) \leq 1$, we get: $|s^{set}(S_1 \setminus S_2) - s^{set}(S_2 \setminus S_1)| = |s^{set}((S_1 \setminus S_2) \uplus (S_2 \setminus S_1)) + s^{set}(S_2 \setminus S_1) \cdot (s^{set}(S_1 \setminus S_2) - 2)| \leq |s^{set}((S_1 \setminus S_2) \uplus (S_2 \setminus S_1))|$. Finally, the result follows by replacing the above inequality in the original formula and using the fact that $0 \leq s^{set}(S) \leq 1$. \square

Proof for Theorem 5.

We break the proof in two parts, one per function.

For g^{cng} .

We need to show that, for all $x_v, x_s, x_a, x'_v, x'_s, x'_a$, $|g^{cng}(x_v, x_s, x_a) - g^{cng}(x'_v, x'_s, x'_a)| \leq 1 \cdot (|x_v - x'_v| + |x_s - x'_s| + |x_a - x'_a|)$. By the definition of g^{cng} it is obvious that if the above holds for $x_s = x'_s = 0$ it will hold for all x_s, x'_s . So we assume that $x_s = x'_s = 0$.

If $x_a = x'_a$, the result follows easily by observing that $|g^{cng}(x_v, 0, z) - g^{cng}(x'_v, 0, z)| = g^{cng}(|x_v - x'_v|, 0, z)$.

If $x_v = x'_v$, the result follows easily using the above result and the observation that $g^{cng}(x_v, x_s, x_a) = g^{cng}(1 - x_a, x_s, 1 - x_v)$.

Finally, if $x_v \neq x'_v$ and $x_a \neq x'_a$, then the result follows from the fact that: $|g^{cng}(x_v, 0, x_a) - g^{cng}(x'_v, 0, x'_a)| = |g^{cng}(x_v, 0, x_a) - g^{cng}(x_v, 0, x'_a) + g^{cng}(x_v, 0, x'_a) - g^{cng}(x'_v, 0, x'_a)| \leq |g^{cng}(x_v, 0, x_a) - g^{cng}(x_v, 0, x'_a)| + |g^{cng}(x_v, 0, x'_a) - g^{cng}(x'_v, 0, x'_a)|$ and the application of the above cases.

To show that the bound $\ell = 1$ is tight, take $x_v = 1, x_s = x_a = x'_v = x'_s = x'_a = 0$.

For g^{dlg} .

We need to show that, for all $x_v, x_s, x_a, x'_v, x'_s, x'_a$, $|g^{dlg}(x_v, x_s, x_a) - g^{dlg}(x'_v, x'_s, x'_a)| \leq \frac{3}{2} \cdot (|x_v - x'_v| + |x_s - x'_s| + |x_a - x'_a|)$. By the definition of g^{dlg} it is obvious that if the above holds for $x_v = x'_v = 0$ it will hold for all x_v, x'_v . So we assume that $x_v = x'_v = 0$.

The relation also holds when $x_s = x'_s$ and $x_a = x'_a$.

Let's assume that $x_s = x + \epsilon, x'_s = x$ and $x_a = x'_a = y$ for $\epsilon \neq 0$. Replacing these values in the initial formula, we get that it suffices to show that:

$$(x + \epsilon)^3 - x^3 + \epsilon \cdot y^3 - y \cdot ((x + \epsilon)^3 - x^3) - 3 \cdot \epsilon \leq 0.$$

Thus, we seek the maximum value of $g_1(x, y, \epsilon) = (x+\epsilon)^3 - x^3 + \epsilon \cdot y^3 - y \cdot ((x+\epsilon)^3 - x^3) - 3 \cdot \epsilon$ for $x, y, \epsilon \in [0, 1]$. We observe that $\frac{\partial g_1}{\partial y} = 3 \cdot \epsilon \cdot y^2 - ((x+\epsilon)^3 - x^3)$, i.e., a second-order polynomial, with two roots (say y_1, y_2) such that $y_1 = -y_2 < 0$. Thus, the original function is increasing (with respect to y) when $y \in (-\infty, y_1) \cup (y_2, \infty)$ and decreasing when $y \in (y_1, y_2)$. Regardless of whether $y_2 \in [0, 1]$ or not, the maximum value of $g_1(x, y, \epsilon)$ with respect to y , for $y \in [0, 1]$, is reached for $y = 0$ or for $y = 1$, i.e., for all x, y, ϵ it holds that, either $g_1(x, y, \epsilon) \leq g_1(x, 0, \epsilon)$ or $g_1(x, y, \epsilon) \leq g_1(x, 1, \epsilon)$.

For $y = 1$ we get that $g_1(x, 1, \epsilon) = -2 \cdot \epsilon < 0$.

For $y = 0$ we get that:

$$\begin{aligned} g_1(x, 0, \epsilon) &= \epsilon \cdot (\epsilon^2 + 3 \cdot x \cdot \epsilon + 3 \cdot x^2 - 3) \leq \\ &\epsilon \cdot (\epsilon^2 - 3 + 3 \cdot x(x + \epsilon)) \leq (\text{because } x + \epsilon \leq 1) \\ &\epsilon \cdot (\epsilon^2 - 3 + 3 \cdot x) \leq (\text{because } \epsilon^2 \leq \epsilon) \\ &\epsilon \cdot (\epsilon - 3 + 3 \cdot x) \leq (\text{because } x + \epsilon \leq 1) \\ &\epsilon \cdot (-2 + 2 \cdot x) \leq 0 (\text{because } x \leq 1). \end{aligned}$$

Thus, in any case, $g_1(x, y, \epsilon) \leq 0$, so the result holds.

The case where $x_s = x'_s$ and $x_a \neq x'_a$ follows from the above result and the observation that $g^{dlg}(x_v, x_s, x_a) = 1 - g^{dlg}(x_v, x_a, x_s)$.

Finally, if $x_s \neq x'_s$ and $x_a \neq x'_a$, then: $|g^{dlg}(x_v, x_s, x_a) - g^{dlg}(x'_v, x'_s, x'_a)| = |g^{dlg}(x_v, x_s, x_a) - g^{dlg}(x_v, x'_s, x_a) + g^{dlg}(x_v, x'_s, x_a) - g^{dlg}(x'_v, x'_s, x'_a)| \leq |g^{dlg}(x_v, x_s, x_a) - g^{dlg}(x_v, x'_s, x_a)| + |g^{dlg}(x_v, x'_s, x_a) - g^{dlg}(x'_v, x'_s, x'_a)|$, and the result follows from the above cases.

To show that the bound $\ell = \frac{3}{2}$ is tight, suppose that g^{dlg} is ℓ' -smooth, for some $\ell' < \frac{3}{2}$. Take $x_v = 1, x'_v = \frac{2}{3} \cdot \ell', x_s = x_v = x'_s = x'_v = 0$. Then $0 \leq x'_v < 1$, and, using the definition of smoothness, we get to an absurdity. \square

Proof for Theorem 6.

The result follows easily by observing that $w_i \leq w_M$ and replacing w_i by w_M in the formulas. To show that the functions are exactly w_M -smooth, suppose, without loss of generality, that $w_1 = w_M$, and consider the case: $|g^{QUA}(1, 0, \dots, 0) - g^{QUA}(0, \dots, 0)|$ (same for g^{ACC}). \square

Proof for Theorem 7.

The proof follows using Theorems 2, 3, 5, 6, and considering the effects of adding a vote on the different structures (votes, supporting arguments). \square

Proof for Theorem 8.

The proof follows using Theorems 3, 5, 6, and considering the effects of adding a supportive argument. \square

Proof for Theorem 9.

The proof follows using Theorems 3, 5, 6, and considering the effects of the addition of an attacking argument on the different structures (votes, supporting arguments, attacking arguments). \square

Proof for Theorem 10.

The proof follows using Theorems 8, 9 and considering the

effects of changing the strength of an argument. \square