Why do they vote that?

FLAIRS 2016 Paper: 217

Abstract

The mining of justifications to be recommended to visitors of deliberation fora used in decision making by constituents raises specific challenges. The problem lies at the confluence of Data Mining and Argumentation Frameworks. Graph-based representations can improve our understanding of the problem and enable reasoning with the available data. Our high level motivation is the societal problem concerning the need of a sustainable grassroots journalism as made available by the "wisdom of polls" to improve general population's as well as boards' decision making. The addressed technical problem consists in recommending sets of texts containing comprehensive arguments supporting or opposing *poll choices*, as mined from submissions of opinions in electronic deliberative polls. A graphical framework is proposed to enable the development of techniques for identification of relevant/encompassing arguments in debates.

Introduction

Decision making systems can be designed to provide incentives for each user to submit or support a candidate for the most comprehensive justification of the poll choice selection done by him rather than less comprehensive arguments. The justification text submitted by each user may contain the relevant worldview of the group of participants making the same selection, i.e., their view on all available poll choices, defining the balance that has tilted their decision in the direction selected by the group. The incentive for placing comprehensive arguments in one text submission can be provided, for example, by mechanisms where users can support only one justification at a time, in conjunction with ranking schemes for texts based on support.

We present a framework for ranking or extracting the k-most comprehensive justifications for each poll choice, where k is an estimation of the number of justifications that a user studies. Authors need to locate relevant justifications that others submitted for making the same choice as themselves, to identify and merge their arguments. Authors also need relevant justifications for remaining poll choices, to identify challenges from users making other selections, to which they can provide their group's answer. The argumentation framework we propose has classes/coalitions of arguments as a part of its definition rather than a result of some reasoning. As such these coalitions can be used as a robust support of further inferences.

Argumentation Frameworks for Debates

Mechanisms to extract data from forum statements into some argumentation framework can profit more from exploiting metadata, despite the existence of investigations into general level forum design and ranking schemes (Hsu, Khabiri, and Caverlee 2009; Delort, Arunasalam, and Paris 2011; Høgenhaven 2013; Momeni, Cardie, and Diakopoulos 2016). Data mining and intelligent recommendation is being used for tasks such as analysis of Twitter messages (Ediger et al. 2010) and recommendation of forum statements based on authors (Arnt and Zilberstein 2003), and significant effort is put by research in natural language and abstract argumentation frameworks to come up with formal models of argumentation that fit debates in various venues (Aakhus and Lewiński 2011), as well as for parsing statements to extract logical arguments (Palau and Moens 2009).

The research in abstract argumentation frameworks has been mainly dealing with dialectics (namely dialogues with arguments in a persuasion process to establish the truth). Much of the current research in abstract argumentation frameworks (AAFs) is based on Dung's seminal theory of argumentation. In his approach, a *(Dung) argument* is defined as *"an abstract entity whose role is solely determined by its relations to other arguments. No special attention is paid to the internal structure of the arguments."* (Dung 1995). The examples of arguments given by Dung are the natural language statements in the exchange: "My government cannot negotiate with your government because your government doesn't even recognize my government." and "Your government is a terrorist government."

Definition 1 (AAF) *An argumentation framework is a pair* $\langle \mathcal{A}, \mathcal{C} \rangle$ *where* $\mathcal{A}$ *is a set of arguments, and* $\mathcal{C}$ *is a binary relation representing attacks on* $\mathcal{A}$*, namely* $\mathcal{C} \subseteq \mathcal{A} \times \mathcal{A}$*.*

The relation $\mathcal{C}(\alpha, \beta)$ between two arguments, α and β, represents the fact that α attacks β. Based on such graphs, Dung defined the concepts of *conflict-free* and *admissible extension* (set of arguments which defends each argument in the set, i.e. *accepting* them, by attacking whatever other argument attacks them), as well as a characteristic function F that can filter admissible extensions of an input conflict-free set of arguments by retaining acceptable arguments with respect to the input. Procedures based on these concepts can infer maximal admissible extensions (aka preferred), complete extensions (fix points of F), stable extensions (attacking any other arguments), or the smallest complete extension (the least fix point of F). Each of these extensions defines a different semantic of acceptance.

The concept of *preference* can be seen as a source of attacks on attacks, and has been introduced as an extension of AAFs:

Definition 2 (EAF (Modgil 2001)) *An Extended Argumentation Framework (EAF) is a tuple* $\langle \mathcal{A}, \mathcal{R}, \mathcal{D} \rangle$*, such that* $\mathcal{A}$ *is a set of arguments, and:*

- $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$*, showing attacks*
- $D \subseteq A \times R$*, showing arguments which defeat attacks*
- *If* $(X, (Y, Z)), (X', (Z, Y)) \in \mathcal{D}$ *then* $(X, X'), (X', X) \in \mathcal{R}$*.*

An alternative competing approach to preferences is quantitative, as defined by Value-based Argumentation, which also integrates the concept of audience:

Definition 3 (VAF (Bench-Capon 2003)) $\langle \mathcal{A}, \mathcal{R}, \mathcal{V}, val, \mathcal{P} \rangle$ *is a value-based argumentation framework (VAF), where val is a function from arguments* $\mathcal{A}$ *to a non-empty set of values* $\mathcal{V}$*, and* $\mathcal{P}$ *is a set* $\{a_1, ..., a_n\}$*, where each* a_i *names a total ordering (audience)* $>_{a_i}$ *on* $\mathcal{V} \times \mathcal{V}$*.*

An audience specific VAF (aVAF) is a tuple $\langle \mathcal{A}, \mathcal{R}, \mathcal{V}, val, a \rangle$ *where* $a \in \mathcal{P}$*.*

If $X, Y \in \mathcal{A}$*, then* X *defeats*$_a$ Y *iff* $(X, Y) \in \mathcal{R}$ *and it is not the case that* $val(Y) >_a val(X)$*.*

It has already been argued, as a practical desiderata, that *abstract argumentation frameworks (AAF) should also be studied under the assumption that they are motivated by requirements for modeling relations between locutions as used in common reasoning and debate* (Modgil 2013), as compared to their instantiation with logical theories, to make them more suitable for modeling dialogues appearing in practice. Abstract Locution Networks (ALNs) were proposed as an alternative where nodes are locutions rather than arguments. Previous research suggested that systems should be prompting users introducing locutions to clarify the relations intended (Brewka and Woltran 2010; Modgil 2013). An abstract framework was proposed for online debating systems, to enable reasoning about argument strength using votes (Eğilmez, Martins, and Leite 2013):

Definition 4 (ESAF) *An extended social argumentation framework is a 4-tuple* $F = \langle \mathcal{A}, \mathcal{R}, \mathcal{V}_{\mathcal{A}}, \mathcal{V}_{\mathcal{R}} \rangle$*, where* $\mathcal{A}$ *is a set of arguments,* $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ *is a binary attack relation between arguments,* $\mathcal{V}_{\mathcal{A}} : \mathcal{A} \rightarrow \mathbb{N} \times \mathbb{N}$ *stores the crowd's pro and con votes for each argument, and* $\mathcal{V}_{\mathcal{R}} : \mathcal{R} \rightarrow \mathbb{N} \times \mathbb{N}$ *stores the crowd's pro and con votes for each attack.*

An abstract bipolar argumentation framework (BAF) (Amgoud et al. 2008; Boella et al. 2010) introducing support relations has also been proposed:

Definition 5 (BAF) *An abstract bipolar argumentation framework* $\langle \mathcal{A}, \mathcal{R}_{def}, \mathcal{R}_{sup} \rangle$ *consists of a set* $\mathcal{A}$ *of arguments, a binary relation* $\mathcal{R}_{def}$ *on* $\mathcal{A}$ *called a defeat relation and another binary relation* $\mathcal{R}_{sup}$ *on* $\mathcal{A}$ *called a support relation: consider* A_i *and* $A_j \in \mathcal{A}$*,* $A_i \mathcal{R}_{def} A_j$ *(resp.* $A_i \mathcal{R}_{sup} A_j$*) means that* A_i *defeats* A_j *(resp.* A_i *supports* A_j*).*

Groups of arguments in a BAF that satisfy a coherence requirement define coalitions (Cayrol and Lagasquie-Schiex 2010):

Definition 6 (Coalition of BAF) $\mathcal{C} \subseteq \mathcal{A}$ *is a coalition of BAF iff:*
(i) *The subgraph of* $\mathcal{G}_{sup}$ *induced by* $\mathcal{C}$ *(the graph representing the AAF* $\langle \mathcal{A}, \mathcal{R}_{att} \rangle$*) is connected;*
(ii) $\mathcal{C}$ *is conflict-free for AF;*
(iii) $\mathcal{C}$ *is maximal (for* $\subseteq$*) among the sets satisfying* **(i)** *and* **(ii)***.*

The main AAF related problems addressed in literature are about finding a set of laws (rules) that are compatible and have support (Dokow and Holzman 2010), and finding the strongest chains of arguments (Amgoud and Devred 2011).

Social Media Given a graph, a number of researchers in different areas have studied how to find the set of most "important" nodes in the graph. In social networks, important nodes are influential people (Kempe, Kleinberg, and Tardos 2003). Various algorithms have been proposed for this problem, such as the Flow Authority model (Aggarwal, Khan, and Yan 2011), and IMM (Tang, Shi, and Xiao 2015).

Diversifying results from search engines to help users access opposing opinions, was seen as a contribution to the reduction of public polarization (Yom-Tov, Dumais, and Guo 2013; LaCour 2015).

Proposed Framework

We start by introducing the problem setting and usage of the involved terms: justification, reason, group and worldview.

Problem Setting Specification: Poll Dialogues

We address decision supporting processes via *fora associated with electronic polls*, where a discussion is focused on a single question Q (i.e. claim). In general polls, questions can be either open-ended or restricted to a finite number of allowed *choices*: $c_1, ..., c_k$. In a common case, the answer choices are just *Yes* and *No*.

A statement submitted by a user can contain multiple explanations for different choices and for various audiences. Unlike for Dung's abstract arguments, we refine the term argument in the sense of minimality. To avoid confusion, we refer to a minimal argument as a *"reason"*:

Definition 7 (Reason) *A* reason for a choice *is a minimal statement in support of the given choice.*

For example, we will say that the statement: *"We should keep buying US bonds because US economy is flourishing and EU is in an existential crisis"* is a *justification*, but not a *reason* (minimal argument) for the choice of *"buying US bonds"*. Examples od reasons are: (1) *"US economy is flourishing"*, and (2) *"EU is in an existential crisis"*. While Dung's definition could account for the above statement, α, as containing one argument (given Dung's definition of argument as an opaque entity), we will count it as containing two reasons. This enables us to model and evaluate the comprehensiveness of a *portfolio of statements.*

When we talk about a *group* of users we generally refer to the set of participants selecting the same choice of the given poll. By the *worldview* of a given group with respect to a specific poll question we mean the union of the offered sets of reasons due to which members select the corresponding poll choice.

Frequently such groups are not monolithic and the literature has looked into considering them as separate audiences or groups (Lewinski and Blair 2011). Various levels of resolution in the analysis of these people could yield different sets of sub-groups. This is partly accounted in our work by the fact that we do not look to reconstruct the worldview of the group (defined by a poll choice) as a single justification, but as a portfolio of justifications. The actions a user is expected to perform are: (1) submit an argument, (2) reply/improve/respond to an argument (with comments), (3) like an argument (without comments).

Framework modeling justification content in poll debates We introduce a new type of relation between arguments, namely *enhance*. A definition of this relation is:

Definition 8 (Enhance Relation) *Argument α enhances another argument β, if α contains all the reasons of β, but beta does not contain all the reasons of α.*

While in practice enhancement can be just an improvement on the clarity of the description or organization of some reasons, that can still be accounted as an increase of the number of arguments (with a numerical fraction).

Definition 9 (Justification) *A justification is an argument for voting a choice c_i of a poll question s with possible choices $c_1, ..., c_k$. It consists of a statement $\langle \Sigma, \Pi \rangle$, with the semantic $\langle \Sigma, \Pi \rangle \rightarrow (s = c_i)$, where $\Sigma = \{r_1, ..., r_k\}$ such that each r_j is a set of reasons for voting choice c_j of s, and $\Pi : r_1 \cup ... \cup r_k \rightarrow \mathcal{O}$ is a function mapping each reason to an element of a partially ordered set $\mathcal{O}$ specifying a qualitative or quantitative order of significance between reasons.*

The significance component Π can be factored into the computation of relevance for reasons, and can be provided by meta-data or by NLP.

We have to account for the fact that different shareholders may own different amounts of shares in the corresponding company whose business decisions are debated, and therefore that the votes have different *weights* (as given by the shares owned). We will consider that the voting weights are non-negative real numbers. This does not preclude non-shareholders from participating in decision-making. Other stakeholders and volunteer advisers can be enabled to contribute statements with zero voting weight. We start by addressing the case of Boolean poll questions, namely with two possible answers, such as *Yes* and *No*.

Now we can define a family of argumentation frameworks, that we refer to as: *Bipolar Abstract Poll Debate Frameworks (BAPDFs)*. Unlike previous extensions of Dung AAFs, BAPDFs enable us to specify as input the labels provided by decision-makers to their statements/justifications. Namely, we assume that authors label all their statements with meta-data specifying at least the coalition (choice of the claim s) that the statement supports. Further the BAPDFs model the new relations (e.g., enhance), besides attack and support relations available in BAF and besides the corresponding votes defined with ESAF. Here, rather than specifying formally instances of BAPDFs, we limiting ourselves to state that BAPDFs may define a binary relation representing attacks between statements of different types, $\mathcal{R}_{att}$ and a binary relation representing enhancement between statements of the same type $\mathcal{R}_{enh}$. It may also specify functions $\mathcal{V}_{\mathcal{A}}$ and $\mathcal{V}_{\mathcal{R}}$ assigning weights/votes to arguments and relations.

When a justification for voting choice c of poll question s is studied as an opaque entity j, its representation from Definition 9 becomes $j \rightarrow (s = c)$. The other types of relations we plan to exploit are *attack* relations between justifications of different types, denoted $j_1 \not\longrightarrow j_2$, and *enhance* relations between justifications of the same type, denoted $j_1 \rightarrowtail j_2$. The attack relation $j_1 \not\longrightarrow j_2$ suggests that j_1 *contradicts* j_2. The enhance relation $j_1 \rightarrowtail j_2$ suggests that j_1 *includes* j_2. For both cases we say that the relations *start at* the justification j_1, the one on their left-hand side, and *point to* j_2, being *pointed by* j_1.

The relation between BAPDF and existing argumentation frameworks has to be explored from the perspective of both generality and semantics. We note that simplifications of the BAPDF (e.g., without enhancement relations) can be modeled as instances of ESAF, namely for a special structure on arguments. Further, BAPDF without votes can also be seen as a modification of BAF with different semantics for relations and coalitions. As such, the concepts of *admissible extensions* from AAF can be directly used with the above simplified view of BAPDFs. However, additional attacking rules may be inferred by observing that from relations like j_1 *enhances* j_2, j_2 *attacks* j_3, and j_3 *enhances* j_4 ($j_1 \rightarrowtail j_2$, $j_2 \not\longrightarrow j_3$, $j_3 \rightarrowtail j_4$) it can be extracted that likely j_1 *attacks* j_3 and that j_2 *attacks* j_4 ($j_1 \not\longrightarrow j_3$, $j_2 \not\longrightarrow j_4$).

The conclusion is that more general frameworks can be obtained by merging the enhancement and group concepts from BAPDF with the negative votes of ESAF, with the support relations and inferred coalitions in general graph structures of BAF, and with the preferences from EAF or values in VAF.

Problems Enabled

A central problem that occurs in various flavors in multiple applications that we plan to address, is the identification of *the most encompassing K justifications* for each choice. We define *the most encompassing K arguments* for a poll choice from two different perspectives, obtaining results with different semantics:

(a) *Relevance or Representativity:* A set of K arguments that together describe the reasons of a set of voting shareholders with the highest total number of shares/stock (among all sets of K arguments).

(b) *Comprehensiveness:* A set of K arguments that together contain the largest number of reasons for a given choice of a poll (among all sets of K arguments).

A version of this problem consists in finding *the most encompassing any-time $\mathcal{K}$ ranking of justifications*:

Definition 10 (Any-Time $\mathcal{K}$ Ranking) *The most encompassing any-time $\mathcal{K}$ ranking of a set of justifications for a set of integers $\mathcal{K}$ is an ordering of the justifications such that for any number K, $K \in \mathcal{K}$, the first K justifications for the most encompassing set among all possible sets of K justifications.*

Theorem 1 (Impossibility) *In general it is impossible to have a most encompassing any-time ranking of justifications for all possible sets $\mathcal{K}$.*

Proof: A counterexample can be built as follows. Let us consider the problem defined by the set of reasons $\{a, b, c, d, e\}$ and the set of justifications: $j_1 = \{a, b, c, d\}, j_2 = \{a, b, e\}, j_3 = \{c, d, f\}$. The single most encompassing $K = 1$ justification is j_1, and the single most encompassing set of $K = 2$ justifications is $\{j_2, j_3\}$. Therefore no extension of the solution for $K = 1$ is a solution for $K = 2$, and finding a most encompassing any-time ranking for $\mathcal{K} = \{1, 2\}$ is impossible. □

Solving algorithms can be designed to exploit these relations as arcs in a bipartite graph. The occurrence of bipartite-graphs with reply links was previously observed from the perspective of coalition detection (Agrawal et al. 2003). These algorithms search for the best (i.e., most relevant) subsuming candidates for justifications of the opposing conclusions. The search can look directly for K justifications at a time, or for a ranking, e.g., by interleaving as appropriate extractions of one justification at a time, with extractions of K justifications at a time.

We formalize now the set of situations when it can be argued that a justification discusses the reasons or issues raised by a given voter, under the aforementioned interaction assumptions of this work:

(i) that a voter can sign support for a single justification which comprehensively describes his reasons,

(ii) and that a voter can also sign attack (called **decl_attacks**) relations and/or enhances relations (called **decl_enhances**) between the supported justification and other justifications.

Definition 11 (Answer) *A justification is said to* **answer** *a given* voter *if either it is signed by that voter, or if it points, by a chain of* **decl_attacks** *or* **decl_enhances** *relations, to some justification signed by that voter.*

A simple version of the problem is where there exist no **decl_enhances** relation and the **decl_attacks** attack relation is functional, being specified only by the author of the justification on its left-hand side.

The Subsuming Justification Problem (SJP) is specified by a pair $\langle \mathcal{B}, K \rangle$ where $\mathcal{B}$ is a BAPDF for a Boolean poll question s with a functional **decl_attacks** relation $\mathcal{R}_{att}$ and argument weight function $\mathcal{V}_{\mathcal{A}}$. The SJP problem is to find a set of at most K supporting justifications that **answer** to a maximum weighed number of voters supporting s, and a set of at most K opposing justifications that **answer** to a maximum weighted number of voters opposing s.

Note that, in SJP, for ranking endorsing justifications we are only interested in how they represent the participants approving the Boolean claim. The symmetric approach is used for ranking opposing justifications. This design is based on the principle of limiting the incentives of opponents for voting strategically on relations and justifications.

An example of a SJP's bipartite graph for 6 justifications is shown in Figure 1a. Nodes on the left represent endorsing justifications $\{j_1, j_2, j_3\}$ and those on the right represent opposing justifications $\{j_4, j_5, j_6\}$. Each node j_i is shown with a weight w_i representing the shares of shareholders signing justification j_i, $w_i = \mathcal{V}_{\mathcal{A}}(j_i)$. An alternative representation is shown in Figure 1b, based on a heterogeneous graph adding nodes for votes and nodes for voters (i.e., shareholders). It is possible to use one node per shareholder, labeled with its voting share, or one node for all shareholders with the same vote and labeled with the sum of their voting shares. The shareholders represented (answered to) by a justification are those that can be reached by traversing the obtained directed graph starting with the justification. The approximations in the assumption that each justification subsumes the justification that it attacks can be quantified by a discount factor/ For example, in the given graph, if weights $w_i = i$,

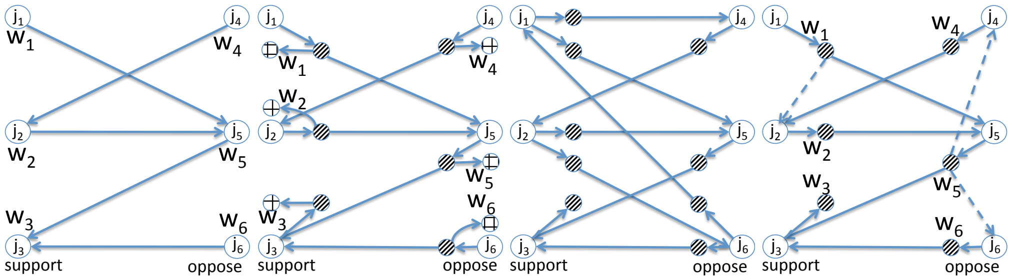

a) SJP b) SJP (alternative) c) WSJP d) CSJP

Figure 1: Large circles show "justification" nodes, diagonally hatched circles show "vote" nodes, and cross hatched circles show "voter" nodes (aka "shareholder" nodes). Continuous arrows between "justification" nodes, or from "votes" to "justifications" show attack relations. Dashed arrows show enhance relations. The arrows from "justification" nodes to "vote" nodes are called "vote" arcs.

then the most encompassing supporting justification is j_2 which answers all voters represented by shares in w_2, w_3. The *most encompassing 2 supporting justifications* are $\{j_1, j_2\}$, answering all supporting voters. The most encompassing opposing justification is j_4.

A practical problem with the simplifying assumptions in SJPs is that an older justification cannot be marked as a valid **decl_attacks** attack of newer justifications without its author's help. In the next extension with practical relevance, each voter can specify explicitly the justification that his selected justification **decl_attacks** (rather than inheriting the one specified at the creation of his justification).

A Weighted Subsuming Justification Problem (WSJP) is specified by a pair $\langle \mathcal{B}, K \rangle$ where $\mathcal{B}$ is a BAPDF for a Boolean poll question s, with a **decl_attacks** attack relation $\mathcal{R}_{att}$, where functions $\mathcal{V}_{\mathcal{A}}$ and $\mathcal{V}_{\mathcal{R}}$ assign weights for arguments and relations. The WSJP problem is to find a set of at most K supporting justifications that **answer** to a maximum weighted number of voters supporting s, and a set of at most K opposing justifications that **answer** to a maximum weighted number of voters opposing s.

A sample graph depicting a WSJP is shown in Figure 1c. In this version of the WSJP graph we do not show "shareholder"/"voter" nodes but only "vote" nodes. A "vote" node can be labeled with the number of shares of the shareholders submitting the corresponding vote, or alternatively it can be linked with directed weighted arcs to "shareholder" nodes. Each "vote" node is pointed by a "voter" arrow from the justification j_i signed by the corresponding vote, and may be linked by an outgoing "attack" arrow to another justification that is marked as attacked by j_i in $\mathcal{R}_{att}$. The weight of a justification returned by $\mathcal{V}_{\mathcal{A}}$ is the sum of weights for votes on attacks from that justification, as returned by $\mathcal{V}_{\mathcal{R}}$, potentially summed with the weight of a "vote" node that does not lead to any attack. Unlike for SJPs, with WSJPs there may be multiple arrows exiting a justification. Here the contribution of a justification in terms of representativity is given by the sum of weights of "vote" nodes (or "shareholder" nodes, if shown) that can be reached from it along directed arcs. The *most encompassing supporting justifications* in this example are either j_1 or j_2, function of discount factors and operators used to factor for the votes on attacks.

The last problem with practical relevance introduced here is the extension of the simple SJP problem where **decl_enhances** relations are added.

The Components Subsuming Justification Problem (CSJP) is specified by a pair $\langle \mathcal{B}, K \rangle$ where $\mathcal{B}$ is a BAPDF for a Boolean poll question s with a functional **decl_attacks** relation $\mathcal{R}_{att}$, a **decl_enhances** enhancement relation $\mathcal{R}_{att}$, and a function $\mathcal{V}_{\mathcal{A}}$ assigning weights to arguments. The CSJP problem is to find a set of at most K supporting justifications that **answer** to a maximum weighted number of voters supporting s, and a set of at most K opposing justifications that **answer** to a maximum weighted number of voters opposing s.

The example of graph showing a CSJP in Figure 1d also displays "shareholder" nodes. The main difference with the SJP in Figure 1b is that it also contains "enhance" arrows representing relations in $\mathcal{R}_{enh}$. These arrows lead from "vote" nodes to the enhanced "justification" nodes. Again, the shareholders *answered* by a justification are those reachable along directed arcs from the node of that justification.

Theorem 2 *The SJP problem is NP-hard.*

The proof follows from the reduction of the NP-hard influence-maximization problem (Kempe, Kleinberg, and Tardos 2003). A SJP can be mapped from an influence maximization problem by interpreting the phenomena in the opposite direction. Namely we interpret that the reasons have propagated along attack relations in the opposite directions of arrows.

Conclusion

We propose a graphical framework which enables techniques that will directly benefit the next generation of debate forums, governmental sites asking for suggestions, as well as commercial and public-domain decision support systems.

Existing argumentation frameworks focus on problems related to the automatic identification of admissible extensions, i.e. sets of arguments/statements satisfying certain consistency criteria. An argumentation framework is introduced here to help define problems related to retrieving the most encompassing sets of K arguments, or to help generating any-time rankings of arguments.

Three different types of problems relevant for the new argumentation framework have been proposed, for identifying *the most subsuming justifications* among those submitted with votes in a threaded debate focusing on a single question. These BP-hard problems differ from the perspective of the flexibility allowed for argument meta-data.

We have also proven that it is not always possible to generate any-time rankings, namely rankings where reading the justifications in the provided order is a strategy maximizing the learning process. In this cases, algorithms can generate approximate solutions.

References

Aakhus, M., and Lewiński, M. 2011. Argument analysis in large-scale deliberation. *Keeping in touch with pragma-dialectics* 165–183.

Aggarwal, C.; Khan, A.; and Yan, X. 2011. On flow authority discovery in social networks. In *Proc. SDM*, 522–533.

Agrawal, R.; Rajagopalan, S.; Srikant, R.; and Xu, Y. 2003. Mining newsgroups using networks arising from social behavior. In *World Wide Web*, 529–535.

Amgoud, L., and Devred, C. 2011. Argumentation frameworks as constraint satisfaction problems. In Benferhat, S., and Grant, J., eds., *SUM*, volume 6929 of *LNAI*. Springer. 110–122.

Amgoud, L.; Cayrol, C.; Lagasquie-Schiex, M.-C.; and Livet, P. 2008. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems* 23(10):1062–1093.

Arnt, A., and Zilberstein, S. 2003. Learning to perform moderation in online forums. In *Web Intelligence*, 637–641.

Bench-Capon, T. 2003. Persuasion in practical argumentation using value-based argumentation frameworks. 3(13):429–448.

Boella, G.; Gabbay, D. M.; van der Torre, L. W.; and Villata, S. 2010. Support in abstract argumentation. *COMMA* 216:111–122.

Brewka, G., and Woltran, S. 2010. Abstract dialectical frameworks. In *Principles of Knowledge Representation and Reasoning*.

Cayrol, C., and Lagasquie-Schiex, M.-C. 2010. Coalitions of arguments: A tool for handling bipolar argumentation frameworks. *International Journal of Intelligent Systems* 25(1):83–109.

Delort, J.-Y.; Arunasalam, B.; and Paris, C. 2011. Automatic moderation of online discussion sites. *International Journal of Electronic Commerce* 15(3):9–30.

Dokow, E., and Holzman, R. 2010. Aggregation of binary evaluations. *Journal of Economic Theory* 145:495–511.

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

Ediger, D.; Jiang, K.; Riedy, J.; Bader, D. A.; and Corley, C. 2010. Massive social network analysis: Mining twitter for social good. In *2010 39th International Conference on Parallel Processing*, 583–593. IEEE.

Eğilmez, S.; Martins, J.; and Leite, J. 2013. Extending social abstract argumentation with votes on attacks. In *International Workshop on Theorie and Applications of Formal Argumentation*, 16–31.

Høgenhaven, T. 2013. *Open Government Communities: Does Design Affect Participation?* Copenhagen Business School.

Hsu, C.-F.; Khabiri, E.; and Caverlee, J. 2009. Ranking comments on the social web. In *Computational Science and Engineering*, volume 4, 90–97.

Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *SIGKDD international conference on Knowledge discovery and data mining*, 137–146.

LaCour, M. J. 2015. A balanced news diet, not selective exposure: Evidence from a direct measure of media exposure. In *APSA 2012 Annual Meeting Paper*.

Lewinski, M., and Blair, J. A. 2011. Monologue, dilogue or polylogue: Which model for public deliberation?

Modgil, S. 2001. Reasoning about preferences in argumentation frameworks.

Modgil, S. 2013. Revisiting abstract argumentation frameworks. In *Theory and Applications of Formal Argumentation*. 1–15.

Momeni, E.; Cardie, C.; and Diakopoulos, N. 2016. A survey on assessment and ranking methodologies for user-generated content on the web. *ACM Computing Surveys (CSUR)* 48(3):41.

Palau, R. M., and Moens, M.-F. 2009. Argumentation mining: the detection, classification and structure of arguments in text. In *International Conference on Artificial Intelligence and Law*, 98–107.

Tang, Y.; Shi, Y.; and Xiao, X. 2015. Influence maximization in near-linear time: A martingale approach. In *Proc. SIGMOD*, 1539–1554.

Yom-Tov, E.; Dumais, S.; and Guo, Q. 2013. Promoting civil discourse through search engine diversity. *Social Science Computer Review*.