AVAPla: A Tool for Supporting Public Interest Communication via Computational Argumentation

Irene Russo[1], Carlo Taticchi[2] and Paola Vernillo[1,3,*]

[1]Istituto di Linguistica Computazionale "Antonio Zampolli", Consiglio Nazionale delle Ricerche
[2]Dipartimento di Matematica e Informatica, Università degli Studi di Perugia
[3]Alma Mater Studiorum - Università di Bologna

Abstract

Public Interest Communication (PIC) seeks to deliver messages to the public through communication strategies that align with ethical and social values. Despite its social relevance, the field currently lacks formal tool-supported methods for designing value-sensitive campaigns. This paper presents AVAPla, a web-based tool developed within the EPICA project to assist in planning PIC campaigns through computational argumentation. AVAPla enables users to construct arguments that consider both audience characteristics and underlying values, offering visual and analytical features to assess their potential impact and effectiveness. A use case is provided to demonstrate how the tool can support more targeted and persuasive public interest communication.

Keywords

public interest communication, computational argumentation, value-based argumentation

1. Introduction

Public Interest Communication (PIC) is an emerging field focused on how to effectively communicate about social issues in ways that foster public engagement and debate. It draws on disciplines such as public relations, political communication, advocacy, and activism, all of which are concerned with the ethical design and governance of public communication campaigns [1]. Although various best practices and guidelines have been developed, the field still lacks practical tools to implement key strategic choices that are increasingly recognised as essential in communication management, such as the selection of effective arguments' based on audience's values. Individuals, groups, or societies can be guided by different value priorities. Literature has shown [2, 3] that basic values (e.g., benevolence, tradition, etc.) are universally recognised by people of all cultures. This consideration is particularly important when addressing composite audiences, i.e., those whom the speaker wishes to influence through their argument [4], since values, opinions, and beliefs are inextricably connected. The EPICA project (Empowering Public Interest Communication with Argumentation) started from the observation that the development of formal models and the use of information technology to support PIC activities are quite limited. For these reasons, we introduce AVAPla (*Audiences and Values Arguments Planner*), a tool designed to support the planning phase of public interest campaigns by enabling a preliminary analysis of foundational elements through the lens of computational argumentation.

AVAPla is a computational tool designed to support the structured design, selection, and organization of arguments to achieve a communicative goal. It is specifically created to help design public interest campaigns that are tailored to value-sensitive audiences. While web-based interfaces for creating argument maps as graphs are available [5], they generally do not integrate formal principles from computational argumentation into their design. This leaves users responsible for manually managing the higher-level qualitative aspects of the argument set. In contrast, AVAPla provides an independent evaluation informed by computational argumentation principles, improving user understanding and supporting more focused decision making based on the relationship between arguments. Formal

AI³ 2025: 9th Workshop on Advances in Argumentation in Artificial Intelligence, September 13, 2025, Rende, Italy
*Corresponding author.
irene.russo@ilc.cnr.it (I. Russo); carlo.taticchi@unipg.it (C. Taticchi); paola.vernillo@unibo.it (P. Vernillo)
0000-0003-1260-4672 (C. Taticchi)

approaches to argumentation emphasize normative structures, focusing on how arguments should be constructed and exploring theoretical scenarios, such as when a claim is subject to multiple opposing counterarguments from different sources. Given that public interest communication naturally involves dialogic and argumentative exchange, computational argumentation has the potential to enrich the field with new analytical tools and innovative methodologies.

This paper introduces AVAPla, presents its features in Section 2 and reports a case study in Section 3. Future developments are reported in the concluding section.

2. System Description

The tool is provided as a web interface[1], designed to support the modelling and analysis of public interest communication through computational argumentation techniques. Each argument is associated with a vector that captures its affinity for different values that influence public perception of the communication. By exploring these value arrays, the tool offers insights into how different argumentative strategies may resonate with specific audience segments. User interaction unfolds in three main steps: campaign definition, visualisation of various computed metrics, and analysis through predetermined goals.

2.1. Campaign Definition

The campaign is modelled starting from a Value-based Argumentation Framework $\langle A, \to, A^{\text{pos}} \rangle$, defined in [6] as an extension of [7]. In this framework, A denotes a set of arguments, $\to$ is a binary attack relation between arguments, and $A^{\text{pos}} \subseteq A$ is the subset representing the core arguments that express the campaign's objectives. To enrich the expressiveness of the framework, [6] also incorporates values and audiences in the model. Arguments can support several values at the same time, each to a different extent. This is modelled using a value space $V = [0, 1]^n$, where each dimension represents a specific value. A value function $\text{val} : A \to V$ assigns to each argument $a \in A$ a vector that shows how strongly it relates to each value. The set of audiences is denoted by $I = \{1, 2, \ldots, k\}$, where k is the number of distinct audiences considered. An audience $i \in I$ is associated with a weight p_i to represent the proportion of the total population that shares similar value preferences. These weights satisfy normalisation and non-negativity conditions:

$$\sum_{i=1}^{k} p_i = 1 \quad \text{and} \quad \forall i \leq k,\ p_i \geq 0.$$

Each audience has its own preferences over the values. These are represented by a function $\text{asv} : I \to V$, which assigns to each audience i a vector. The j-th entry of $\text{asv}(i)$ shows how important value j is to audience i.

In the tool, the campaign structure is defined through a JSON schema[2] that formalises the components of a framework $\langle A, \to, A^{\text{pos}} \rangle$ enriched with values and audiences. The schema requires the specification of seven main elements: `positiveArguments`, `arguments`, `attacks`, `values`, `audiences`, `argumentValues`, and `audienceValues`. Arguments are divided into two sets: `positiveArguments`, which contain the core arguments A^{pos}, and a general `arguments` set, which includes all other relevant arguments from $A \setminus A^{\text{pos}}$. Each argument includes a unique identifier and an array of sentences representing its textual content. The `attacks` component defines the binary relation $\to$, where each attack is specified by an identifier of the source and target argument. The `values` array lists value identifiers used in the campaign, then associated with arguments via `argumentValues`, which, like the val function, links each argument to one or more values, using a weight between 0 and 1 to represent the strength of the connection. Similarly, the `audiences` array defines each audience $i \in I$ by a name and a weight p_i, while `audienceValues` shows how much each audience cares about

[1]Tool webpage: https://epica.dmi.unipg.it/tool.
[2]Schema available at: https://epica.dmi.unipg.it/tool/script/schema.js.

each value, using weights to express these preferences, effectively implementing the function asv. To define a public interest campaign, users may choose from one of the available input methods: uploading or pasting a JSON file, or using a guided input form.

2.2. Visualisation

The second step allows users to explore the campaign by computing and visualising in tables some of the measures proposed by [6] to evaluate the effectiveness of the arguments with respect to different audiences. After the computation, the tool displays the following elements.

Impact measure table For each audience $i \in I$ and argument $a \in A$, AVAPla computes the impact function

$$\|a\|_i = \frac{1}{\sqrt{n}} \|\text{asv}(i) \odot \text{val}(a)\|$$

that measures the influence of a on i, based on how well the values promoted by the argument align with those prioritised by the audience.

Defeat relation table To include values and audience preferences in the framework, AVAPla computes a defeat relation $\twoheadrightarrow_i$ for each audience $i \in I$. Given two arguments $a, b \in A$, we use this definition of defeat

$$a \twoheadrightarrow_i b \iff (a \rightarrow b \wedge \|a\|_i \geq \|b\|_i).$$

This ensures that an argument can only defeat another if it attacks it and has at least as much impact on the audience.

Acceptability table AVAPla checks whether an argument $a \in A$ convinces a specific audience i by computing the grounded semantics [8] on $\langle A, \twoheadrightarrow_i \rangle$. The result is $\text{con}_i(a)$, which is true if and only if the argument is accepted in the grounded extension. The tool computes and shows the value of $\text{con}_i(a)$ for every argument and audience.

Campaign graph The tool also provides an interactive graph to explore the campaign and facilitate the facilitate interpretation of the results. The graph shows the campaign from the perspective of a selected audience, which can be chosen and toggled by the user via a dropdown menu. Visual styles in the graph help distinguish between elements from `positiveArguments` and `arguments` (the nodes), `attacks` and the computed defeat relation $\twoheadrightarrow_i$ (the edges), and the arguments accepted as convincing by $\text{con}_i(a)$.

Users can interact with the graph in various ways to add arguments, create attacks, delete elements, and move arguments around. Each interaction triggers a recomputation of all the measures from step 2, ensuring that the graph remains consistent with the campaign representation. For example, when an attack is added, the system checks whether it qualifies as a defeat and updates its visual style accordingly. When adding a new argument node to the campaign graph, the user is asked to provide some information to define the argument. These include: an ID, the type (either `positiveArguments` or `arguments`), a natural language sentence expressing the argument's core idea, and a set of weights associated with the different values. This ensures that the argument is fully integrated into the framework and enables a meaningful computation of its impact and interactions with other arguments. In addition to these interactions, a tooltip appears when hovering over an argument in the graph, displaying information including the argument's computed impact for the selected audience.

2.3. Analysis

The final step allows users to analyse the campaign through two distinct goal functions. *Overall Effectiveness* selects the argument $a \in A^{\text{pos}}$ that achieves the highest impact across all audiences, weighted by p_i. For each positive argument, the tool computes the quantity

$$\sum_{i=1}^{k} p_i \cdot \|a\|_i.$$

On the other hand, the *Convinced People* goal selects $a \in A^{\text{pos}}$ that convinces the largest share of the audience. In this case, the tool computes the quantity

$$\sum_{i=1}^{k} p_i \cdot [\text{con}_i(a)] \quad \text{(with } [\varphi] = 1 \text{ if } \varphi \text{ is true, 0 otherwise).}$$

Once the analysis is computed, a table is shown with the weighted scores of all positive arguments according to each goal function.

3. Use Case

In this section, we present a simple illustrative use case to show the functionalities and analytical capabilities of the AVAPla tool in the context of planning a Public Interest Communication (PIC) campaign. A government agency commissions a team of communication experts to design a campaign — referred to as *MoreGreens* — aimed at encouraging the consumption of fruits and vegetables. The target audience includes two particularly resistant demographic groups: individuals under 25 years old (hereafter, *young people*) and individuals over 60 (hereafter, *older adults*). Survey data previously collected on these two groups reveals distinct value orientations. *Young people* tend to pursue lifestyles centred on pleasure, and are thus motivated by the value of *hedonism*. In contrast, *older adults* focus on maintaining their health and physical well-being, making *security* their guiding value.

Although these two values are not inherently incompatible and can, in theory, be addressed by the same argument, in this scenario, the communication experts have identified seven arguments for the *MoreGreens* campaign, each of which exclusively appeals to a single value. Of these, two are strongly aligned with hedonism and three with security. The remaining two arguments act as counterarguments: one challenges two pro-*hedonism* arguments, and the other a pro-*security* argument. Both counterarguments strongly promote hedonism or security, respectively.

In the first case, focusing on the *young people* audience, the counterargument is able - at least in one instance - to defeat the corresponding positive argument, thereby reducing its persuasive impact on the target group (see Figure 1).

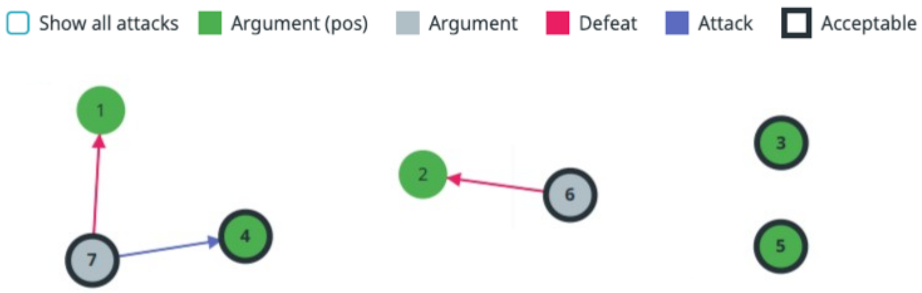

Figure 1: Visualisation of the simple case study.

If the campaign's objective is to maximise the persuasiveness of selected arguments across both audiences —each guided by different value systems— it becomes essential to identify arguments that resonate with the largest group of people. These are determined through AVAPla's *Convinced People* function. Naturally, arguments that are defeated (i.e., successfully countered) are excluded from this selection process. AVAPla thus supports the identification of the most promising arguments on which an effective campaign can be built.

4. Conclusion and Future Work

In this paper, we introduced AVAPla, a tool designed to support the planning phase of public interest campaigns, providing insights into argument effectiveness based on measures from computational argumentation. We believe that PIC could benefit from dynamic tools that support the effective planning of value-sensitive campaigns. At the same time, we believe that computational argumentation should investigate interesting aspects — such as the dynamic construction of the audience — that emerge as relevant in the actual practice of PIC but would be considered marginal in a purely formal approach.

As a future development, we plan to integrate AVAPla with a natural language processing pipeline capable of automatically extracting core arguments from a set of documents, identifying attack relations, and assigning each argument a graded value based on the finite set of universal values. Within this broader framework, we also aim to model the perceived credibility of the source of the arguments — as judged by the target audiences of the PIC campaign — as an additional parameter that can influence the overall effectiveness of the positive arguments.

Acknowledgments

Carlo Taticchi is a member of the INdAM Research group GNCS and of Consorzio CINI. This work has been partially supported by INdAM - GNCS Project, codice CUP_E53C24001950001 and MUR project PRIN 2022TXPK39 - PNRR M4.C2.1.1. "Empowering Public Interest Communication with Argumentation (EPICA)" CUP H53D23003660006, funded by the European Union - Next Generation EU, Missione 4 Componente 1.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] J. Johnston, M. Pieczka (Eds.), Public Interest Communication: Critical Debates and Global Contexts, Routledge, London; New York, 2024.

[2] S. H. Schwartz, Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries, Advances in Experimental Social Psychology 25 (1992) 1–65.

[3] J. Graham, J. Haidt, S. Koleva, M. Motyl, R. Iyer, S. P. Wojcik, P. H. Ditto, Moral foundations theory: The pragmatic validity of moral pluralism, 2012.

[4] C. Perelman, L. Olbrechts-Tyteca, The New Rhetoric: A Treatise on Argumentation, University of Notre Dame Press, Notre Dame, IN, 1969.

[5] M. Janier, J. Lawrence, C. Reed, OVA+: An argument analysis interface, in: S. Parsons, N. Oren, C. Reed, F. Cerutti (Eds.), Proceedings of the 5th International Conference on Computational Models of Argument (COMMA 2014), volume 266 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2014, pp. 463–464. doi:10.3233/978-1-61499-436-7-463.

[6] P. Baroni, G. Fellin, M. Giacomin, C. Proietti, A vector-based extension of value-based argumentation for public interest communication, in: C. Proietti, C. Taticchi (Eds.), Proceedings of the 8th Workshop on Advances in Argumentation in Artificial Intelligence 2024 co-located with the 23rd International Conference of the Italian Association for Artificial Intelligence (AIxIA 2024), Bozen, Italy, November 28, 2024, volume 3871 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024. URL: https://ceur-ws.org/Vol-3871/paper3.pdf.

[7] T. J. M. Bench-Capon, Persuasion in practical argument using value-based argumentation frameworks, J. Log. Comput. 13 (2003) 429–448. URL: https://doi.org/10.1093/logcom/13.3.429. doi:10.1093/LOGCOM/13.3.429.

[8] P. M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning and logic programming, in: R. Bajcsy (Ed.), Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993, Morgan Kaufmann, 1993, pp. 852–859. URL: http://ijcai.org/Proceedings/93-2/Papers/003.pdf.