

# Globally Interconnecting Persuasive Arguments: The Vision of the Persuasive Web

Giorgos Flouris \*    Antonis Bikakis †    Theodore Patkos \*  
Dimitris Plexousakis \*

*FORTH Technical Report TR 438, November 2013*

## Abstract

The current Web evolves rapidly from a Web of information to a Web of Opinions. People are constantly uploading their viewpoints, ratings and comments on any conceivable topic, via blogs, social media websites, commercial websites, wikis and others. Despite some recent efforts to give a more “argumentative character” to the Web, the opinions are still expressed in purely textual form, without any metadata or any other type of annotation that would make them machine-understandable and processable. Moreover, their persuasive power, i.e., the characteristics that determine how an opinion affects the viewpoint of the audience, is largely neglected. Improving Web experience towards this direction would give immense added-value, as it would allow searching for opinions for or against a given claim, understanding the interlinking and interrelationships between opinions (e.g., support or opposition) and obtaining a better understanding of the various debates for a well-intended user. Ultimately, it can lead to increased participation of Web users in democratic, dialogical interchange of opinions, more informed decisions by professionals and decision-makers, as well as to an easier identification of biased, misleading or deceptive arguments. This paper presents the vision of the Persuasive Web, whose objective is to allow the above developments through the introduction and application of a formal, machine-readable representation of persuasive arguments that would allow the discovery, tracking, retrieval, combination, interrelation extraction and visualization of the vast variety of viewpoints that already exist on the Web, in a way that goes beyond simple keyword-based processing. In addition to explaining the vision and the added-value that we expect it to offer to its users, we also describe the associated research and technical challenges, as well as the technologies that could be potentially employed to address these challenges.

## 1 Introduction

The Web has changed. The way people use the Web and the software that is being developed to support their experience has changed. But Web itself is a kind of software as well; a huge, multi-participatory, continuously evolving software, which we like to version based on how we use it: Web 1.0<sup>1</sup> was primarily about publishing and consuming content; Web 2.0<sup>2</sup> is centered around a collaborative effort

---

\*Institute of Computer Science, FO.R.T.H.

†Department of Information Studies, University College London

<sup>1</sup>[http://en.wikipedia.org/wiki/Web\\_1.0](http://en.wikipedia.org/wiki/Web_1.0)

<sup>2</sup>[http://en.wikipedia.org/wiki/Web\\_2.0](http://en.wikipedia.org/wiki/Web_2.0)

to contribute knowledge through social media; Web 3.0<sup>3</sup>, although still not fully developed, aims at assigning meaning to data and enabling machines to provide better and more accurate services. Web 4.0, well ... is not fully crafted yet.

In 2009, Tim Berners-Lee<sup>4</sup> emphasized the role that linked data will play in the new Web, any type of data. Although still evolving, we argue here that a “beta version” of this new Web is already launched; it is the Web we are experiencing, full of colorful, significant, irrelevant, contradictory, life-changing *opinions*. And their linking is going to be an indispensable part.

People use the Web to rate a hotel or a restaurant; they make comments on the story and the writing style of a book; they use it to like or dislike a photograph, a video, or the whole lifework of a music band; they write opinions in blogs, discuss subjects of any matter in forums; they substantiate opinions in wikis citing sources of diverse reliability. The Web asphyxiates opinions and arguments touching topics related to just about everything important or insignificant that happens or has happened or will happen in our world. Its form is rapidly reshaping and what is most noteworthy about this change is that nobody designed it! It was the human users and the way they use the Web that unintentionally drove this change.

But the full new version is still long before being fully deployed. The opinions get lost; *sriptam manent*, yet opinions are currently not uploaded as machine-processable data, they are not interlinked, and it is extremely difficult for Web users to find arguments related to a particular subject, let alone to evaluate them, characterize them based on objective and subjective criteria and finally filter the ones that make a difference to each particular person.

To unlock its potential, people working on the future Web already make efforts to change that. Opinions and arguments need to be structured; they already start to take form as well-defined entities, and applications are in place now that visualize them as well as their attack/support relations using graphs. This is a first step towards allowing arguments to be searched by machines, to be mined, evaluated and linked. But these arguments are not just supposed to form a complex graph, a maze of interconnected nodes supporting or attacking other nodes without evident meaning. Arguments will obtain a *purpose* and their processing will be purposeful as well; the new Web will target at assisting humans in reaching conclusions using arguments that are not only formally structured, but also tailored to the particular characteristics of the person receiving the arguments in order to be better apprehensible. Moreover, they will be formulated considering the special characteristics of the context that constitute one argument more convincing for a particular situation than for another. In other words, arguments will not just be put forward to be heard; they will be adaptable in order to be *persuasive*. As the Web is becoming more and more personalized to accommodate each individual’s preferences and characteristics, so will the opinions, in order to refer to the logic, the sentiments and the background of each person issuing or receiving them.

Expressing opinions is a highly personalized process, and adopting them is no less. The new Web will depart from the formal processing of argumentation features and move closer to the human understanding. The ultimate outcome will be to offer the means for assisting humans in participating in debates and collective decision making processes in a persuasive manner.

We call this new Web the *Persuasive Web*. The aim of this paper is to describe this vision, and motivate work towards its realization. The not-so-distant-future example that follows illustrates how we envision the interaction with the new Web (Section 2). Subsequent sections give details on what the vision is (Section 3), show its relationships with currently existing or proposed technologies (Section 4),

---

<sup>3</sup>[http://en.wikipedia.org/wiki/Web\\_2.0#Web\\_3.0](http://en.wikipedia.org/wiki/Web_2.0#Web_3.0)

<sup>4</sup>[http://www.youtube.com/watch?v=OM6XIICm\\_qo](http://www.youtube.com/watch?v=OM6XIICm_qo)

describe the challenges that stand in the way of its realization and propose directions to overcome them (Section 5), and argue on its importance (Section 6).

## 2 Motivating Example

The day begun with a feeling of unrest for Steffi. The new article she is about to prepare obtains added gravity in the prospect of the referendum that will be held next month. The topic is not unfamiliar to her; as a financial journalist she has written numerous articles in the past regarding the financial crisis and the impact of measures suggested by the International Monetary Fund (IMF) in other countries. This time though, her role needs to be more delicate: not only does she want to question and try to verify the diverse viewpoints that seem to dominate public opinion and flood the social media of her country these days, but also to present as objectively as possible well-justified and clearly-articulated persuasive arguments both in favor and against the controversial role of IMF.

She hits “IMF policies help countries recover from financial crises” in PerSE, the Persuasive Search Engine she mostly uses when seeking for arguments on the Web, and configures its settings in “debate mode”, in order to receive both supporting and refuting arguments. She has prepared a categorization of the different target groups she is interested in to drive the mining process, and has uploaded the corresponding profiles using the “Audience Characteristics” functionality of PerSE. For instance, she would like to know what arguments can be more meaningful for unemployed young people. She further added two more profile types to distinguish between those opinions being exchanged during the early days of the credit crisis in 2005, and those that seem more popular at the present day. Comparing the two, she expects to find those arguments that have a diachronic effect in influencing people’s opinions, as well as to verify which positive or negative expectations related to IMF have been met, along with the promises put forward in the past that have not been materialized. The expected audience profile that Steffi entered will determine the types of arguments that will be more convincing for the particular audience.

Furthermore, PerSE also considers Steffi’s profile, in order to proceed with a focused search, centered primarily around the sources of information she considers trustworthy, her origin, as well as other personal preferences. Her profile data guides PerSE to accurately decide on the level of detail to apply for the construction and presentation of arguments: her expertise in financial terms is sufficient to understand arguments on the connection between unemployment and inflation, but arguments regarding certain social aspects of unemployment require more detailed analysis in order to comprehend. To accomplish this, PerSE does not only mine arguments from different online sources, but has to appropriately combine them to achieve the level of persuasiveness needed before directing them to her.

Steffi navigates through the visual graph that represents the retrieved persuasive arguments. What PerSE has returned is a graphic showing in a visually appealing manner the arguments and counter-arguments, as well as their persuasive properties, including the sources (provenance) of each argument/premise, data that support the premises, the adequacy of each argument for a particular audience and the relationships between arguments (attack, support etc). It further identified categorizations that Steffi did not consider in the first place, classifying certain arguments to audience groups sharing similar characteristics.

A particular argument attracts her attention: IMF argues that “labor flexibility” helps a country attract foreign investments, which results in the creation of new jobs and an increase in the country’s productivity, and ultimately in economic growth. Steffi decides to look more deeply into it; in a dialogical style, she re-

quests PerSE to provide clarifications and further support for the premises that look doubtful to her, in particular the fact that the attraction of foreign investments increases the country's productivity. The system responds by returning persuasive arguments in favor (and against) this correlation.

The arguments returned by PerSE are characterized and annotated in various ways of both objective and subjective nature. The argument style (e.g., deductive or inductive), the consistency of the premises, the credibility of the argument and/or the source (as provided by Steffi's profile), the popularity of the argument itself or of the opinion it supports, and its appeal to the target groups (as presumed by analyzing data from social media or other sources using sentiment analysis methods) are all relevant for Steffi to assess the persuasiveness of each argument for her research. Using all this information, Steffi navigates more deeply in the complex graph of arguments, she filters, groups and organizes opinions, and eventually identifies and extracts the most persuasive ones.

A few hours later her article is ready. Her persuasion-enabled editor has assisted her in annotating the different parts of her text with a formal description of the persuasive arguments they refer to, so that search engines can identify and retrieve them, and in linking them with the respective online sources and arguments/facts they are based upon. Steffi's own conclusions, based on the correlation of facts she personally deduced during her research are also included (and annotated) in the text. This way, her annotated article and arguments can be posted in her company's argument repository for others to find and reuse. As she sends the article to her editor she feels confident that her audience will have the means to form a well-informed opinion before actively participating in the country's decision making process.

### 3 The Persuasive Web Vision

**Why: the need for Persuasive Web.** The current Web evolves rapidly from a Web of information to a Web of Opinions (also known as the Pragmatic Web [15]); and the more popular the social Web services grow, the more difficult it becomes to make meaning out of the plethora of opinions, i.e., to evaluate the credibility and coherency of information related to a subject of interest, understand why it is important, and ultimately decide whether to adopt or reject it.

Even today's Web contains the information necessary for Steffi to complete her article. However, this information, being in textual form, is not easily accessible, retrievable and processable; and therefore it is not appropriate for implementing the features presented in our example scenario. The Semantic Web [3] and the recent linked data hype<sup>5</sup> promise to overcome some of the limitations of natural-language Web pages by providing appropriate methodologies for posting interlinked semantical information (data) on the Web in a machine-readable way. However, the focus of the Semantic Web is on the representation of data, rather than arguments.

Similarly, the main tenets of computational argumentation [4, 12] and the extensive research conducted in this field have direct impact on the formulation of the new Web. The ideas being expressed around the concept of an Argument Web [13, 5] aim at combining the notions of argumentation theory with the current trends of the Web, such as for structuring and connecting arguments, thus offering a formal characterization of logical arguments to facilitate reasoning and retrieval. This is a decisive first step towards managing and organizing the opinions that exist on the Web in a machine-processable manner. Nonetheless, relying exclusively on the formal properties of factual information, they cannot satisfy the primal reason

---

<sup>5</sup><http://linkeddata.org/>

why opinions reach the Web in the first place, which is for *arguments to be persuasive*. This latter step is important, in order to depart from simple argument listings and create *debates with a purpose*, i.e., debates where each argument is formulated with a certain aim, namely to persuade a certain audience on some topic, as happens in real-world debates.

**How: the function and use of the Persuasive Web.** The Web today, as well as the Semantic Web, are mostly concerned with searching and retrieving documents and information. The Persuasive Web will have a different perspective: instead of information, the fundamental searchable component of the Persuasive Web is the *persuasive argument*. Persuasive arguments are not just logical formulae, or concrete arguments consisting of a conclusion and a justification for this conclusion (as in mainstream concrete argumentation frameworks). A long standing belief is that the acceptability and persuasiveness of an argument largely depends on the audience to which it is presented [11]. Therefore, a persuasive argument should also specify the audience that it is targeted at, but also its provenance, the context in which it was made, the values it promotes, the popularity of the claim that it supports, the conditions under which it is effective or valid, etc. It should also provide evidence for its believability, such as links to other documents, facts, or arguments that back it up. Moreover, persuasive arguments are interconnected, in the sense that they may attack or support each other, in various ways. Furthermore, unlike standard arguments, persuasive arguments place emphasis on emotions, trust and other extra-logical factors that may affect the *persuasiveness* of an argument. Therefore, a persuasive argument is a much more complex structure than a simple data item, or even an argument; despite its complexity, the persuasive argument is a standalone entity that forms the fundamental structure of the Persuasive Web.

The Persuasive Web will revolutionize the way the information that exists on the Web is organized and exploited. Persuasive arguments will be uploaded directly by content providers, but it will also be possible for persuasive arguments to be constructed on demand following formal methods, taking into account the intended audience. When a user requests persuasive arguments for or against a certain opinion or decision, Persuasive Web search engines will not only retrieve existing online arguments that exactly match the request, but will also attempt to construct new ones using other online arguments or data from the Semantic Web, which are not necessarily explicitly related to the request. In our motivating example, PerSE combines IMF's argument that labour flexibility ultimately results in an increase in the country's productivity with the fact that GDP is a metric for a country's productivity, and with another argument saying that the increase of GDP means economic growth, to construct a new argument linking labour flexibility and economic growth.

To allow content consumers to make the most out of the presented persuasive arguments, the Persuasive Web will exploit both objective and subjective information, such as the structure of an argument and any possible logical fallacies it may contain, relationships with other arguments (such as attack or support), its provenance, the credibility of the corresponding content provider (which could be provided by the content consumer and/or measured using some kind of objective credibility or reputation metric), and the consumer's personal preferences, background knowledge and interests (profile). This will allow a proper ranking of the presented arguments, so that the strongest, most relevant and most understandable arguments appear high in the list. In our motivating example, PerSE would return the official arguments of IMF, as well as counter-arguments put forward by leading economists and other people, provided that they are trustworthy enough (per Steffi's profile) and understandable (per Steffi's knowledge background). It would

also return arguments related to the original arguments and counter-arguments (e.g., supporting or attacking them), as well as counter-counter-arguments and so on. Ultimately, reading and understanding the natural-language arguments will not be entirely avoided, but the automated pre-processing and association of arguments would save lots of effort.

Persuasive arguments will be stored in and retrieved from “persuasive argument bases” (the persuasive analogous to knowledge bases and ontologies) and will be used to annotate online sources, such as a collection of sentences inside a document, information retrieved from a picture etc. In the context of our motivating example, people arguing for, or against, IMF’s positive role in mitigating the effects of the economic crisis, will have the ability to post and interrelate (support, attack, undermine) arguments (and counter-arguments) in a machine-readable way. Similarly, the IMF itself will be able to express its own arguments on the matter, stored in its own dedicated repository and uploaded on its website. Note that all types of digital artifacts (from financial reports to polls, simple text, images, videos, other arguments, datasets) can be used as evidence supporting a certain argument or position; thus, persuasive arguments and digital objects will be inter-related in two ways: persuasive arguments will be used to annotate digital objects, whereas digital objects will be used as parts of persuasive arguments (e.g., as supportive evidence).

The Persuasive Web will also enable certain forms of dialogical interaction with its users. As described in the motivating example, after receiving a set of persuasive arguments that best match her request, Steffi will be able to follow up by requesting more persuasive arguments, by asking for more clarifying information about a certain argument, or even by disputing the returned ones, e.g., using her own counter-arguments. PerSe will then be able to search again the Persuasive Web and respond back by presenting the next most persuasive arguments in the first case, data that back up or explain the persuasive argument in question in the second case, and data or arguments that invalidate (or support) the user’s counter-arguments in the third case. In addition to investigating the rationale that is leading to a particular conclusion in terms of content and sources, the user can also ask the system to explain the principles that characterize the rational process of generating the conclusions, such as the type of arguments used or the way they support their claims. This way the user can understand and potentially object a given inference before adopting the conclusions provided.

**Who: actors in the Persuasive Web.** The Persuasive Web will provide benefits for both the content provider and the content consumer. From the perspective of the content provider, the technologies behind the Persuasive Web will enable people to post persuasive arguments, rather than just documents or data. People, companies, governments and organizations will now have the ability to provide explicit persuasive arguments in their effort to persuade others into accepting a certain opinion or taking a certain action, thereby using the Web to argue in favor of the products, services or ideas that they promote.

From the perspective of the content consumer, the Persuasive Web will allow easy access to persuasive arguments and associated data. Most importantly, it will allow the content consumer to use automated services and tools to retrieve, process, visualize, understand and query the arguments uploaded by content providers, as well as their interrelationships. In our motivating scenario, Steffi is aided in her task by a graphic display summarizing the most persuasive arguments found in the Web, as well as their properties and interrelationships. In this way, she would be protected from malicious users and sloppy arguments. Moreover, she would be able to concentrate on the most important ones or those that are most relevant to the specific context or case that she is interested in, and she would be able to easily

identify poorly supported arguments.

**What: the goal of the Persuasive Web.** Note that the goal of Persuasive Web is not to impose any given opinion, but to provide the medium through which a user (content consumer) can “collect” different arguments in favor and/or against a certain claim, in order to form an opinion of her own by judging the individual arguments openly provided by various people and evaluating them against known facts, her own beliefs and sentiments, and her own personal priorities, aims, desires and objectives. The services offered by a search engine in a Persuasive Web environment are analogous to those of a journalist, whose role is to objectively reproduce the most prominent opinions expressed by other people or entities (e.g., political parties) on a subject. Even though the opinions and arguments themselves are subjective, the Persuasive Web services should just reproduce them in an objective and unbiased manner, adapting them to the profile of the content consumer and to the context of their use, letting the user decide on their quality and strength (based possibly on external information). This will eventually allow the user to reach informed conclusions or decisions on various matters of interest. In our example, PerSE only retrieves and presents arguments from sources that are trusted by Steffi. It also returns information associated to the persuasive strength of the arguments for audiences that match the profiles provided by Steffi. But it is up to Steffi to decide which of them would actually be the most influential for the readers of the newspaper she is working for.

Towards this end, a number of tools are starting to be developed that facilitate the tasks of visualizing and linking semantically arguments. Debategraph<sup>6</sup>, Argunet<sup>7</sup>, Truthmapping<sup>8</sup> and Cohere<sup>9</sup> offer visual representations and connections of arguments and ideas, whereas Carneades<sup>10</sup> and Araucaria<sup>11</sup> formally analyze argumentation dialogues. The latter further provides a search engine over online argument repositories. The outcomes obtained by these tools make up a solid ground to base the generation of the next Web. However, the aspects of personalization and persuasion are missing completely, rendering infeasible the identifications of those characteristics that constitute an argument more convincing for a particular person than for another. As arguments are being uploaded on the Web with a purpose, Persuasive Web will focus primarily on the human aspect of argumentation, enabling the formulation and retrieval of persuasive arguments in a format tailored to the people involved, helping its users to form their individual opinions or take informed decisions, and freeing them to a large extent from time-consuming tasks of retrieving and combining relevant information from different sources.

**The big picture: the vision of the Persuasive Web.** Arguing on a matter of interest is much more informative than presenting facts about it without using them to reach a conclusion; asking an expert is more productive than going to the library or searching the Web for data on a subject of interest; and a documentary about IMF’s policies and their effect on countries under a financial crisis is much more informative for the average citizen than a large dataset containing all the facts that were used to create such a documentary. In all the above cases, opinions, being the outcome of a high-level aggregation of raw information, are much more concise and useful knowledge than the raw information required to form and support such opinions.

---

<sup>6</sup><http://debategraph.org/>

<sup>7</sup><http://www.argunet.org/>

<sup>8</sup><http://www.truthmapping.com/about.php>

<sup>9</sup><http://cohere.open.ac.uk/>

<sup>10</sup><http://carneades.github.io/>

<sup>11</sup><http://araucaria.computing.dundee.ac.uk/doku.php>

The Persuasive Web will enable Web users to publicly post explicit opinions and persuasive arguments (rather than just documents, or data) in a machine-readable and personalized manner. This way, the Web will become a large collection of searchable, machine-readable individual persuasive arguments, provided by experts (as well as non-experts) in various areas, essentially allowing the Web to behave like an expert itself, rather than as a large library of documents or data. This will help the user obtain easier access to the various opinions that are adequate for her background and profile, as well as the data that support them, essentially allowing more informed choices, better decision-making, increased involvement in democratic processes, better collective awareness on matters of interest, the mitigation of unjustified concerns or predispositions, as well as the easier identification of malicious or deceptive arguments.

As the current Web evolves into a Web of Opinions, where people express their viewpoints on any matter imaginable, the vision of the Persuasive Web is timely, as it aims at allowing the discovery, tracking, retrieval, combination, interrelation, extraction and visualization of the vast variety of these viewpoints in a way that goes beyond simple keyword-based searching. In the following sections, we will describe the challenges that stand in the way of this vision, as well as the currently existing technologies that could potentially be used to address these challenges; moreover, we elaborate on the importance of the vision.

## 4 Technologies

There are several state-of-the-art technologies today that can provide the substrate upon which the vision of Persuasive Web can be realized. Although the Persuasive Web has many similarities with such technologies, it also has some unique features that differentiate it. In this section, we disambiguate the connections with these technologies and show how the Persuasive Web is a distinct (although related) research field. In addition, we identify the elements of these technologies that are relevant to Persuasive Web and can be reused to achieve our vision, concluding that the Persuasive Web can only be made a reality by combining advances from technologies that are already there.

### 4.1 The (Semantic) Web

The Web is one of the most impressive technological success stories. The simple idea of interlinking documents and making them available from everywhere has spawned a great number of applications that most people use in their everyday lives. Building on the same principles, the Semantic Web [3] (and the closely related linked data initiative) provides languages and tools for representing and interlinking Web data based on their meaning, enabling the development of added-value services to its users.

The Persuasive Web is very similar to the above efforts in many aspects. First of all, all these technologies aim at providing the ability to post, interlink and interrelate some kind of information (documents, data, or persuasive arguments respectively) in a global scale, and at building tools and services that exploit this information to provide advanced services to its users. Further, both the Semantic and the Persuasive Web use this information to annotate Web pages in a manner that allows machines to interpret the corresponding data, or persuasive arguments. In this respect, Persuasive Web can be seen as an extension of the Semantic Web, in the sense that it will enrich the semantical information found in Semantic Web pages with persuasive arguments.



The main difference between the Persuasive Web and its predecessors, is in the type of information that they focus on: the Web focuses on documents, the Semantic Web deals with facts, data, knowledge and the associated reasoning, whereas the Persuasive Web concentrates on (persuasive) arguments and the reasoning associated with them; this can be seen as a more high-level type of information, since a persuasive argument is often the result of analysis, aggregation and summarization of raw data and/or commonsense knowledge.

Further, the interlinking between documents (via hyperlinks) or data (via URI references in the context of the Semantic Web and linked open data), will take much more complex forms in the Persuasive Web. In particular, arguments can be interlinked in many different ways, forming a complex Web of interrelationships (arguments may attack, defend, support, defeat, undercut or undermine other arguments, to give a few examples), which should be appropriately modelled and used by Persuasive Web applications and services. Moreover, arguments may consist of simpler arguments and may be related to counter-arguments that together form a coherent whole.

Despite the differences, the technological advances made in the context of the (Semantic) Web will contribute to the development of the Persuasive Web in a critical manner. In particular, the low-level infrastructure of the Persuasive Web is expected to be reusing the standard Web protocols (e.g., TCP/IP), whereas the Semantic Web technologies and representation languages, such as RDF/S and OWL, are expected to form the backbone of the representation languages of the Persuasive Web as well, by providing the low-level language upon which the logical part of persuasive arguments will be based.

At a more general level, the lessons learnt regarding the modes of global-scale and almost instantaneous interaction between remote users, an unprecedented phenomenon in human history, will provide valuable feedback and help avoiding pitfalls in the design of the Persuasive Web.

In fact, the Persuasive Web is not going to replace the current (Semantic) Web. It will exploit existing technologies, languages and methodologies of the (Semantic) Web by building a new layer, that of persuasive arguments, on top of the knowledge layer of the Semantic Web. Arguments will be formulated on top of structured, semantically annotated data in order to be comprehended in the same way by all.

## 4.2 The Social Web

The set of Web-enabled services that form what is commonly known as the Social Web is experiencing remarkably rapid growth in popularity over the last years. Having as common objective to foster social interaction, a plethora of tools and platforms enable humans to communicate through blogging, tagging, Web content voting, social bookmarking and other means of social interaction. An open exchange of opinions, ideas and facts is dominating the Social Web, with blogs and wikis constituting prime examples of social and collaborative activity. Arguments are routinely created and shared among people in explicit and implicit ways, and the form that they obtain is not necessarily restricted by a “text-based” style.

The adoption of the Social Web by an increasing part of the population empowers people with access to a broader set of viewpoints regarding a topic of interest; nevertheless, in its current form, knowledge and opinion sharing is largely fragmented and seeking unbiased conclusions becomes all the more harder. The Persuasive Web will upgrade the role of the Social Web into a broader means of communicating opinions and carrying out debates. With a shift towards structured arguments, knowledge exchange will be carried out along the lines of logical consistency, factual accuracy and some degree of emotional appeal to the intended

audience. Even though the decision of adopting one conclusion over another will remain a subjective issue, Persuasive Web will facilitate the process of deliberation by filtering out irrational and logically incorrect expressions, while maintaining a significant degree of personalization in choosing the top-rated arguments for each user. In this respect, the Persuasive Web can be viewed as the “blog of tomorrow”, where people are able to not only express their viewpoints in a natural language, but also to annotate them in a machine-readable way.

### 4.3 The Pragmatic Web (Web of Opinions)

There is a growing feeling that the Semantic Web is evolving from a Web of Data to a Web of Opinions (also known as the Pragmatic Web [15]), where RDF statements do not actually represent factual data, but the subjective opinions of the people who upload them.

This is quite similar to our Persuasive Web vision, but still different; in particular, the Pragmatic Web looks into ways of handling incomplete data or resolving conflicting opinions about data, using non-traditional methodologies such as statistical and optimization techniques. Similarly with the Social Web and the Persuasive Web, it is concerned with the social dimension of Web data. The main problem that it deals with is the context-dependent representation and interpretation of data. Its aim is to enable users from different contexts to reach agreements on the meaning of data using tools for ontology negotiation, collaborative ontology building and others.

The Persuasive Web, on the other hand, does not focus on data, but rather on the persuasive arguments that one can build on top of it. Instead of handling incomplete data, it deals with incomplete and approximate arguments (also known as enthymemes [18]), i.e., arguments that explicitly represent only some of the premises for entailing their claims; and instead of resolving data conflicts it selects the most relevant arguments for and against a given claim or decision and lets the user decide.

Some of the techniques used in the Pragmatic Web could, however, be relevant and useful for the Persuasive Web as well. For example, crowdsourcing is becoming a popular methodology for several tasks such as annotating a Web document, populating an ontology, adding relationships between Web resources, etc. The same methodology can be used in the context of the Persuasive Web to fill incomplete arguments, find relationships among arguments or link arguments to specific data. The Games With A Purpose (GWAP) hype [16] could also be used for such a purpose, i.e., as an approach to encourage the annotation of Web pages with persuasive arguments. And more generally, the systems for ontology negotiations and pragmatic ontology building that are being developed in the context of the Pragmatic Web [15] will contribute to the development of consensual ontologies, which are essential for the successful deployment of the Persuasive Web.

### 4.4 Argumentation Systems and the Argument Web

Argumentation theory studies how conclusions can be reached through logical reasoning; and argumentation systems are logic-based computational systems that aim to reflect this process (a recent overview of research in argumentation systems is available in [12]). They are valuable tools for the realization of our vision, as they can provide the models, the methods and the tools for constructing arguments, specifying relations among arguments, and drawing logical conclusions from arguments. However, there are significant differences between argumentation and persuasion, and Persuasive Web is a lot more than an argumentation system deployed in a global scale.

The difference between persuasion and argumentation is analogous to the difference between a political speech and a mathematical proof. In contrast to argumentation, persuasion does not only appeal to the logic of the audience, but also to its emotions. It is only partly based on facts and data; persuasive arguments also employ other techniques including the clever use of verbal cues and the semantic structure of text/speech (politeness, aggressiveness etc), appeal to authority, and others. Persuasion also involves a process of selecting the *best* or *strongest* arguments to present first taking into account their provenance, the context in which they are made, but also the profile, the preferences, the background knowledge and the cognitive state of the audience. In this sense, persuasion is more context-aware and more personalized.

Moreover, there are differences with respect to the forms and aims of the two fields. Argumentation can be monological, i.e. an internal process where an agent constructs and evaluates arguments given some knowledge and draws conclusions from the assembled arguments; or dialogical, where a set of agents interact to construct arguments for and against a particular claim [4]. Persuasion on the other hand always involves at least two parties with different roles: the persuader and the audience. Note that, given its dialogical nature, the persuader and the audience may exchange roles during the interaction; it could also be the case that the audience is (a group of) person(s) (e.g., a judge, or a committee) that try to take a decision and there are several persuaders that try to influence the group towards a certain decision. The goal of argumentation is to reach a consensus on a set of conclusions that can be logically drawn; the aim of persuasion is for the audience to accept the claims of the persuader.

The Argument Web [13, 5] is an effort to deploy argumentation theory on the Web. The effort consists of the design and implementation of an ontology for arguments called the Argument Interchange Format (AIF, see [7, 13]), which captures different forms of arguments, argumentation schemes and acceptability semantics; and the development of Web-based tools for argument visualisation and analysis (e.g., Argument Analysis Wall<sup>12</sup>), argument-based debates (e.g., Arvina<sup>13</sup>), and argument blogging (e.g., ArguBlogging<sup>14</sup>).

The Argument Web and the Persuasive Web share similar features but have different aims: the aim of the Argument Web is to enable argumentation on the Web, while Persuasive Web further aims at implementing the aspects of persuasion as discussed above. Consequently, they also differ in their main design characteristics. The core element of the Persuasive Web is the *persuasive argument*, which compared to the *logical argument* used in the Argument Web, carries additional information associated to its persuasive strength, intended audience etc. Furthermore, additionally to the logic-based reasoning processes used in the Argument Web, the Persuasive Web also involves extra-logical processes such as profile and context analysis, audience analysis, sentiment analysis, trust analysis, reputation analysis and others. Certainly, the Persuasive Web will exploit the developments made in the Argument Web with respect to argument modeling, annotation and visualization, but will further extend them to integrate the distinct characteristics of persuasive arguments and persuasion.

## 5 Realizing the Vision

To achieve the vision of Persuasive Web, several technological advances in different areas will be necessary. In this section, we explore some of the research chal-

---

<sup>12</sup>[http://www.arg.dundee.ac.uk/?page\\_id=645](http://www.arg.dundee.ac.uk/?page_id=645)

<sup>13</sup>[http://www.arg.dundee.ac.uk/?page\\_id=492](http://www.arg.dundee.ac.uk/?page_id=492)

<sup>14</sup><http://www.argublogging.com>

lenges that stand in the way of creating this new generation of Web content, and highlight opportunities for overcoming these obstacles through extensions, adaptations or improvements of existing research efforts. Note that most such opportunities will come through research efforts performed in the context of the technologies described in Section 4.

**Understanding persuasion.** A first challenge is to understand the process of persuasion, its different aspects, and the characteristics, the form and the structure of its core element, the persuasive argument. At the moment, it is clear that advances in the field of argumentation and the understanding of how logical arguments are structured will form the basis for understanding and defining persuasive arguments. However, logical arguments are not the same as persuasive arguments, and the increased emphasis on extra-logical factors that persuasive arguments employ should also be considered. What are exactly the types of information that define the persuasiveness of an argument? How are such types of information modeled and attached to an argument? How can we use this information to evaluate the persuasive strength of an argument (for a given context and audience)? How do we characterize the audience? How do we treat the combined strength of argument collections, as well as degrees of preference among them, to capture their persuasive power? What are the relations between persuasive arguments (e.g., do we just have to consider the attack/support relations used in current argumentation frameworks)? And how do these relations affect an argument's persuasive strength? Most of these issues are open research topics in computational argumentation, with numerous interesting frameworks being proposed the last few years (e.g., see [1]).

To address these challenges we need to exploit relevant research from the fields of Psychology and Philosophy that will help us better understand the internal mechanics of persuasion for humans, and address questions like the ones presented above. Related research on the argumentation field is also quite relevant in this respect, even though it only focuses on the logical part of arguments. The aim of argumentation systems, such as Dung's abstract argumentation system [8] (probably the most influential work in this area), is to compute a set of acceptable arguments and conclusions given an initial set of arguments and their relations. However, unlike what is assumed in [8] the strength of a persuasive argument or its strength relative to its counter-arguments do not only depend on their content and logical structure. As Walton explains in [17], other factors are also considered: appeal to expert opinion, popularity of arguments or of the opinions that they support, analogous arguments, non-logical (e.g., statistical) correlations between different arguments, expected outcome of accepting/rejecting an argument, arguments from sign, arguments from commitment, ad hominem arguments, and arguments from verbal classification.

Another challenge is the definition of the "audience" of a persuasive argument, and its most important characteristics with respect to persuasion [11]. In Persuasive Web, as in real life, when someone makes or uses (or uploads) an argument, she will have to consider its impact for the particular audience that the argument is addressed. Types of information that are relevant in this respect are: the values or aspirations that the audience considers significant; the background knowledge and beliefs of the audience; its general preferences and goals. There are some recent works that study and model the different characteristics of the audience in the context of argumentation dialogues [2, 9, 14, 6]. These works can provide the basis for the development of a holistic audience model that combines all different aspects that are relevant to persuasion, and algorithms for matching the persuasive arguments with the audience's profile.

Finally, Bench-Capon [1] signifies the role of other issues that have been put

forward in the past, such as the accrual of arguments, i.e., the inclination of humans to consider arguments accumulatively when deciding which conclusion to choose. All the aforementioned factors are relevant, and should be taken into account in the design of the Persuasive Web.

**Representation and Storage.** As also explained in Section 4, the Persuasive Web will use the Semantic Web infrastructure to enable users to upload, combine, interlink and consume persuasive arguments from different sites of the Web using different applications. This requires the development of a semantically explicit representation model for persuasive arguments -like an ontology for persuasive arguments- so that different independently developed applications will be able to comprehend persuasive arguments in a common manner and interoperate within an integrated environment. The Argument Interchange Format [7, 13] is a first significant step towards this direction. For the needs of the Persuasive Web, it will have to be extended to capture information associated to the persuasive nature of arguments, and combined with profile and context ontologies. In this way it will enable, for example, retrieving or selecting persuasive arguments that are most appropriate for a specific context and audience.

Persuasive arguments will be stored in what we call “persuasive argument bases”, i.e., the persuasive analogous of knowledge bases. Their structure must enable storing persuasive arguments, as well as any other information which is relevant to their persuasiveness, and that is represented in the ontology of persuasive arguments. Apart from storing relevant data, a persuasive argument base must also provide: (a) inference support; (b) query support; (c) support for data management tasks such as updating, repairing and change monitoring; (d) alignment with related ontologies; and (e) propagation of relevant persuasive information among different systems. For the development of such systems, the experience gained from the deployment of triple stores and other semantic data management systems in the Semantic Web will be exploited. The AIFdb database system [10], which was developed for storing and managing arguments described in the AIF ontology, supports some of the desired functionalities: it enables semantic processing and visualisation of arguments, query management and dialogue control. Its main limitation with respect to the needs of the Persuasive Web, is that, being developed on top of AIF, it does not capture the persuasive aspects of arguments.

**Presentation and Visualization.** Given the sheer size of the Web, one is expected to find a large number of persuasive arguments in favor (or against) a certain opinion, so presenting everything to the user is certainly not productive. Some kind of aggregation or summarization is necessary, along with a ranking process that will present the most persuasive ones. It should be emphasized here that ranking should not filter any arguments, but only aims at the practical necessity to present some arguments before some others; the user should be potentially capable of viewing all arguments, and no filtering or censorship should take place as part of the ranking process.

Along similar lines, the automated generation of persuasive arguments on the basis of data or other arguments found in the Web would be an interesting feature for many related applications. This will create additional value from existing arguments, via aggregation, summarization and generation of new knowledge in the form of persuasive arguments.

Another challenge, which is also related to improving the experience of the content consumer while accessing persuasive arguments found in the Persuasive Web, is visualization. Visualization of persuasive arguments and their relationships is important for the content consumer in order to understand the structure

of a complex set of persuasive arguments. Recently, there have been some efforts to visualize arguments as trees or graphs (e.g., DebateGraph, Carneades and others), which are based on the classical visualization of abstract argumentation frameworks [8]. However, such visualizations focus on the logical part of arguments, as well as on their relationships; even though this is useful for persuasive arguments as well, the evaluation of their usability when taking into account the idiosyncrasies of persuasive arguments (where, e.g., the source of the argument, its expected audience, its persuasiveness etc play a prominent role), should be a subject of future work.

**Populating the Persuasive Web.** As with the Semantic Web, encouraging simple content providers to upload their data/arguments in the Web on the proper format has always been an issue. The size of the Persuasive Web must reach a critical mass to make itself useful (as with all added-value technologies), so technologies like automated mining of arguments from blogs, forums or other social media, Natural Language Processing (NLP) and others should be employed once the necessary infrastructure is in place. In addition, human help could be enlisted for this task, by adapting existing technologies such as the GWAP hype [16].

Along the same lines, the fact that textual persuasive arguments may be expressed in various different languages constitutes a problem. NLP technologies are mostly focusing on the English language, so, unless this focus shifts to other languages in the future, it will generally be difficult to use NLP for annotating non-English arguments. However, the improving quality of automated translation tools (even online ones) will probably allow us to overcome this obstacle eventually. To this end, other techniques, like GWAP, could be used to allow the annotation of such arguments and mitigate the effects of multi-linguality. Another relevant note is that that, in addition with text, other media, which are immune to multi-linguality issues, can play the role of evidence or arguments to support or oppose opinions. The annotation of images, sounds or complete documents with the arguments that characterize them is equally critical for a Web where knowledge can take various forms.

## 6 Importance

### 6.1 Importance of the Persuasive Web

The expression of arguments in formal, machine-readable terms, as well as their interlinking (via support, attack, or other relationships) will create significant added-value benefits. In the same way that linked data in the current Semantic Web has led to the discovery of new, previously unseen connections, correlations and knowledge (e.g., business analytics), we expect the interlinking of arguments to lead to a better understanding of the various debates and the generation of new, aggregated or previously unknown persuasive arguments and insights.

The abundance of Web data, combined with machine-readable and processable persuasive arguments, will allow the envisioned version of the Web to not only provide information (as when reading a book), but to also combine available data in order to provide persuasive arguments in favour of (or against) different alternative options (as done by a knowledgeable expert). This way, people will be better informed on matters of interest or debates, thus promoting collective awareness on community problems and enabling better decision making for professionals or companies.

At the community level, automated persuasion can enable public authorities to reach a broader audience in a more personalized way, in order to foster policies

of societal value (e.g., healthy lifestyle, sound environmental behavior), to target unjustified concerns, or to promote participation in community matters and democratic processes (e-democracy).

At the individual level, persuasion systems will form a critical component of future autonomous entities endowed with socio-cognitive intelligence, which are used in the emerging market of smart spaces. This can find applications ranging from service robots for domestic use, to smart environments related to domestic care and work, education, healthcare, communication and entertainment.

In addition, there is a wide range of potential applications suitable for the private sector; these generally fall under marketing, e.g., persuading customers to buy products/services, convincing people to donate to a charity, etc. Similarly, the Persuasive Web can also be used as an assistive tool for individuals that practice persuasion as part of their professional life, such as lawyers, business executives etc, or for decision-makers in general, as it would allow better and more informed choices by combining information found in the Web, and also possibly in local databases, to build persuasive arguments and suggestions.

To achieve all these, various tools and applications can be envisioned, which will allow the retrieval, enhancement, connection, adaptation, personalization and visualization of persuasive arguments in blogs, social media, Web forums or other contexts where people express opinions. In addition, one can envision various personalized applications for certain contexts, such as: expert advice in specific domains, advertising and marketing, expert or personalized commentary on products or services, persuasive reasoning in social robotics, novel methods of human-computer interaction and others.

## 6.2 Ethical Issues

We acknowledge the fact that persuasion (that underlies the Persuasive Web), as well as the development of automated persuasion systems, would, by their very nature, be open for misuse by governments, businesses, individuals or organizations (e.g., for coercion, control or opinion enforcement). For example, one potential issue would be the usage of the Persuasive Web as a means to promote extremist or other dangerous opinions, or the incorporation of false, deceptive or misleading arguments by malicious content providers in the Persuasive Web. In both cases, naive content consumers could be deceived, thus causing disillusionment to well-intentioned users and jeopardizing the usefulness of Persuasive Web.

Despite the fact that such opportunities for abuse are admittedly present, this is the case for most useful technologies, so we argue that this should not be a deterring factor towards realizing this technology. As a most striking example, one could mention today's Web, where all such features exist (extremist opinions, inaccurate or false facts etc). However, we argue that the Persuasive Web will in fact improve the situation, and will be helpful towards mitigating this problem.

In particular, it should be noted that it is not the aim of the Persuasive Web to provide any kind of censorship or checking on different opinions. On the contrary, the Persuasive Web will allow all opinions, including extreme or deceptive ones, to be more easily publishable and accessible. We argue that this feature will in fact reduce the opportunities for censorship, coercion, or deception, in the sense that access to different opinions, as well as the verification of the validity of arguments associated with these opinions, will be easier for open-minded content consumers, so the power of extreme opinions will be mitigated.

Similarly, understanding persuasion (in general) can reduce the opportunities of coercion, control or manipulation that may potentially be exercised by businesses, individuals or organizations to unaware citizens. Research on persuasion can help in identifying how and when this happens, as well as in preventing it, by

allowing humans and intelligent systems to argue together.

At a more technical level, advances in the fields of trust and automated fact-checking<sup>15</sup>, as well as the incorporation of provenance information in persuasive arguments could help users in the task of identifying deceptive or misleading arguments. This is similar to how the current Web has allowed recent advances in technology where facts and statements can be more easily checked for validity against the vast amount of the information available on the Web, using fact-checkers of a manual nature<sup>16</sup>, or even automated ones<sup>17</sup>.

## 7 Conclusion

Not long ago, the problem of information overload attracted the attention of different scientific communities, fueled by the increasing number of people posting and accessing information existing on the Web; nowadays, the increasing participation of users in the generation of new information in the Web, in the form of opinions, may lead to a similar problem, that of opinion overload. In this paper, we looked ahead to a future version of the Web, where this problem can be overcome by exploiting the structure and intension of online opinions. We presented the Persuasive Web, a vision for the future Web, where machine-readable arguments (and data) can be created, searched and exploited to provide a personalized and persuasive browsing experience, useful for decision support and consultation. We outlined the main tenets of a new means of communicating opinions online based on the notion of persuasive arguments, along with the critical challenges that lie ahead and the prominent existing technologies that can offer solutions to make this vision a reality.

We strongly believe that the realization of this vision will stipulate research in a wide range of domains, scientific, academic and commercial, and can lead to the development of innovative applications that will revolutionize Web experience. Apart from its evident impact on the organization of opinions and knowledge exchange on the Web, this effort opens up a way to serve a higher-level purpose: by enabling people to locate the valid rational arguments in the sea of opinions of questionable credibility, as well as those arguments that better support them, it will empower critical thinking and facilitate the active participation of humans in collective governance processes. Ultimately, we see the Persuasive Web as the platform of ideas that holds the promise for promoting the role of humans in collective decision-making and e-democracy, able to have significant impact at both the individual and the societal level.

## Acknowledgements

We would like to thank the following people for inspiring discussions in various stages of this work: Thomas Eiter, Michael Fink, Stefan Woltran, Irimi Fundulaki, Panagiotis Papadakos, Anthony Hunter, Grigoris Antoniou, Sarit Kraus.

---

<sup>15</sup>[http://en.wikipedia.org/wiki/Fact\\_checking](http://en.wikipedia.org/wiki/Fact_checking)

<sup>16</sup><http://factcheck.org/>, <http://www.politifact.com/>, <http://decodeurs.blog.lemonde.fr/>

<sup>17</sup><http://truthteller.washingtonpost.com/>



## References

- [1] Trevor Bench-Capon. Logic programs, norms and action. chapter Open Texture and Argumentation: What Makes an Argument Persuasive?, pages 220–233. Springer-Verlag, Berlin, Heidelberg, 2012.
- [2] Trevor J. M. Bench-Capon, Sylvie Doutre, and Paul E. Dunne. Audiences in argumentation frameworks. *Artif. Intell.*, 171(1):42–71, 2007.
- [3] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American-American Edition*, 2001.
- [4] Philippe Besnard and Anthony Hunter. *Elements of Argumentation*. The MIT Press, 2008.
- [5] Floris Bex, John Lawrence, Mark Snaith, and Chris Reed. Implementing the Argument Web. *Communications of the ACM*, 56(10):66–73, 2013.
- [6] Elizabeth Black and Katie Atkinson. Choosing persuasive arguments for action. In *The 10th International Conference on Autonomous Agents and Multi-agent Systems - Volume 3*, AAMAS '11, pages 905–912, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.
- [7] Carlos Chesñevar, Jarred McGinnis, Sanjay Modgil, Iyad Rahwan, Chris Reed, Guillermo Simari, Matthew South, Gerard Vreeswijk, and Steven Willmott. Towards an argument interchange format. *Knowl. Eng. Rev.*, 21(4):293–316, December 2006.
- [8] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, September 1995.
- [9] Anthony Hunter. Making argumentation more believable. In *AAAI*, pages 269–274, 2004.
- [10] J. Lawrence, M. Snaith, F.J. Bex, and C. Reed. ArguBlogging, Arvina, and TOAST. *Frontiers in Artificial Intelligence and Applications*, 245:511–516, 2012.
- [11] Chaim Perelman and L. Olbrechts-Tyteca. *The New Rhetoric: A Treatise on Argumentation*. University of Notre Dame Press, June 1969.
- [12] Iyad Rahwan and Guillermo R. Simari. *Argumentation in Artificial Intelligence*. Springer Publishing Company, Incorporated, 1st edition, 2009.
- [13] Iyad Rahwan, Fouad Zablith, and Chris Reed. Laying the foundations for a world wide argument web. *Artif. Intell.*, 171(10-15):897–921, July 2007.
- [14] Régis Riveret, Henry Prakken, Antonino Rotolo, and Giovanni Sartor. Heuristics in argumentation: A game-theoretical investigation. In *Proceedings of the 2008 Conference on Computational Models of Argument: Proceedings of COMMA 2008*, pages 324–335, Amsterdam, The Netherlands, The Netherlands, 2008. IOS Press.
- [15] Mareike Schoop, Aldo de Moor, and Jan L.G. Dietz. The pragmatic web: A manifesto. *Communications of the ACM*, 49(5):75–76, 2006.
- [16] Luis von Ahn and Laura Dabbish. Designing games with a purpose. *Communications of the ACM*, 51(8):58–67, 2008.
- [17] Douglas Walton. *Fundamentals of Critical Argumentation*. Cambridge University Press, Cambridge, UK, 2006.
- [18] Douglas Walton. *Informal Logic: A Pragmatic Approach*. Cambridge University Press, 2008.