

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Synthesis of collective tag-based opinions in the social web

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/131052> since 2016-06-30T18:19:00Z

Publisher:

SPRINGER

Published version:

DOI:10.1007/978-3-642-23954-0_27

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

This is the author's final version of the contribution published as:

F. Cena; S. Likavec; I. Lombardi; C. Picardi. Synthesis of collective tag-based opinions in the social web, in: LECTURE NOTES IN COMPUTER SCIENCE, SPRINGER, 2011, 9783642239533, pp: 286-298.

The publisher's version is available at:

http://link.springer.com/content/pdf/10.1007/978-3-642-23954-0_27

When citing, please refer to the published version.

Link to this full text:

<http://hdl.handle.net/2318/131052>

Synthesis of Collective Tag-based Opinions in the Social Web [★]

Federica Cena, Silvia Likavec, Ilaria Lombardi, Claudia Picardi

Dipartimento di Informatica, Università di Torino, Italy
{cena,likavec,lombardi,picardi}@di.unito.it

Abstract. This paper presents an approach to personalized synthesis of tag-based users' opinions in a social context. Our approach is based on an enhanced tagging framework, called $iT_{\mathcal{G}}$, where tags are enriched with structure and expressivity and can be addressed to different features of a resource and weighed by relevance. Our main contribution is a synthesis of the collective opinions that is *multi-faceted*: it shows different points of view on the same resource, rather than averaging the opposite opinions, or choosing the one with the most supporters. If the social tool provides user modeling and trust mechanisms, our synthesis can also be *personalized*, taking into account both the user's *social network* (considering only the opinions of trusted authors) and her *user model* (considering only the features the user likes). In addition, we propose an innovative visualization modality for $iT_{\mathcal{G}}$ s, which allows for an at-a-glance impression of all the opinions on a given resource, including significant differences in point of view. We evaluated the $iT_{\mathcal{G}}$ framework to test (i) its expressiveness for providing opinions, and (ii) the effectiveness of our synthesis with respect to traditional tag clouds.

Keywords: social web, tagging systems, personalized synthesis

1 Introduction

In the context of social applications, users often participate in the community life by providing their opinions on the resources the community life revolves around (e.g. books, music, pictures, etc.). To do so, they can *rate*, *tag* or write *free text* comments on items. Social applications could use such meta-data available on the resources for different purposes: to learn about users' preferences or to provide the other users with the synthesis of such a content. The possibility to do this depends on the typology of the user-generated content: ratings are the simplest one to be aggregated as average values, but they are not very informative on the qualities or shortcomings the ratings are based on. Free text comments are very informative but they are difficult to be effectively processed and synthesized. Tags lie in between ratings and free-text comments for richness of information and computability. Our work moves from the observation that (i) traditional tags are more suited to express facts (e.g. for content classification) rather than opinions, since they do not possess enough richness and structure to allow users

[★] This work has been supported by PIEMONTE Project - People Interaction with Enhanced Multimodal Objects for a New Territory Experience.

to express complex and multi-faceted opinions; (ii) the *tag clouds*, commonly used by most of the social applications to present tags in an aggregated form, are often difficult to browse and are not very informative.

Our goals were to overcome such limitations of tagging. First, our aim was to enable expressing of elaborate opinions using tags, i.e., giving users the possibility to use tags in order to express: *judgement* (liking or disliking a feature), *relevance* (saying that a feature is more important than another one) and *scope* (referring an opinion to only a part of a resource). Second, our goal was to merge opinions given by means of tags by different users into a synthesis representing an overview of the collective opinion. The synthesis should be (i) *multi-faceted*, i.e., present contrasting opinions, and (ii) *personalized*, i.e., take into account both the *social network* of the person it will be shown to (the synthesis will consider only the opinions of the users the person trusts)¹ and her *user model* (the synthesis will show only the features of a resource the user considers relevant). This brings about a need to find an innovative *visualization modality* which allows for an at-a-glance opinion of a large amount of users' tags on a resource, giving the possibility to discover different points of view on it.

The paper presents the $iT_{\lambda}G$ framework, an enhanced tagging framework, where tags are enriched with structure and expressivity, so that they can be addressed to different features of a resource and weighed by relevance, and where an approach to opinion synthesis is provided. We report the results of the evaluation of: (i) the expressiveness of the $iT_{\lambda}G$ framework for communicating opinions, (ii) the effectiveness of our synthesis with respect to traditional tag clouds, applied to a social environment for opinion-sharing on restaurants.

The paper is organized as follows. Section 2 introduces the notion of $iT_{\lambda}G$ s. Section 3 describes our approach to $iT_{\lambda}G$ s semantic interpretation, used further in $iT_{\lambda}G$ s synthesis in Section 4. Section 5 discusses the results of the $iT_{\lambda}G$ framework evaluation, followed by the discussion of related work in Section 6 and conclusions in Section 7.

2 Introducing the $iTag$ concept

Let us consider an object $O \in \mathbb{O}$ the user wishes to comment upon, where \mathbb{O} is the set of all the objects in the domain. We assume that: (i) the object O is of a distinguished type, $type(O)$; (ii) a hierarchy of *facets* \mathbb{F} is associated with $type(O)$, where each facet $F \in \mathbb{F}(type(O))$ denotes something about the objects of this type the user may want to comment on. \mathbb{F} contains $type(O)$ itself, as the root of the facet hierarchy. Figure 1 shows a possible hierarchy of facets associated with the *restaurant* type.

An $iT_{\lambda}G$ can be assigned to a specific facet of a given object. We represent an $iT_{\lambda}G$ as a labelled circle of a given size, placed above (positive impression) or below (negative impression) the given facet (see Figure 2). In other words, an $iT_{\lambda}G$ can express an *opinion* ($iT_{\lambda}G$ label, typically an adjective), *scope* (choice of a given facet), *judgement* ($iT_{\lambda}G$ placement) and *relevance* ($iT_{\lambda}G$ size). Formally, an $iT_{\lambda}G$ is defined as follows:

Definition 1 ($iT_{\lambda}G$) An $iT_{\lambda}G$ I is a tuple $I = \langle a, O, F, L, p, S \rangle$ where

¹ Users prefer the recommendations generated by the users they trust rather than the suggestions generated by computer programs, see [1, 3].

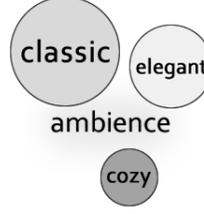
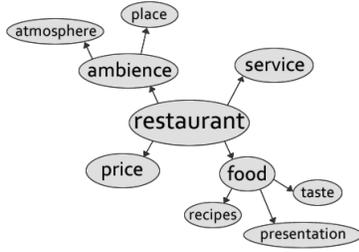


Fig. 1. A hierarchy of facets for type *restaurant*. **Fig. 2.** Representations of iTags as spheres.

- a is the $iT_{\mathbb{A}}^G$ author;
- $O \in \mathbb{O}$ is the object the $iT_{\mathbb{A}}^G$ refers to;
- $F \in \mathbb{F}(\text{type}(O))$ is the facet of $\text{type}(O)$ the $iT_{\mathbb{A}}^G$ is assigned to;
- L is the label (typically an adjective) that refers to facet F of object O ;
- p is the polarity (positive or negative) which describes the user’s judgement of a given facet: $iT_{\mathbb{A}}^G$ placed above (below) the given facet denotes a positive (negative) opinion on a given facet;
- $S \in [\theta, 1]$ is the size that expresses the relevance the $iT_{\mathbb{A}}^G$ author gives to L w.r.t. her global impression of O ($\theta > 0$ is the minimum threshold provided by the system).²

For example, in the following $iT_{\mathbb{A}}^G$ I , Jo comments positively on Alfredo’s food, which she finds *tasty* in a relevant way (size is 0.8):

$$I = \langle \text{Jo}, \text{Alfredo's}, \text{food}, \text{tasty}, +, 0.8 \rangle.$$

3 iTags interpretation

We assume that each facet F is associated with a set of properties $\mathbb{P}(F)$. Properties represent the attributes which are relevant when commenting on the facet; they are known to the system but invisible to the user. For example, the *ambience* facet of type *restaurant* could be associated with the following property set: $\{\text{classicism}, \text{elegance}, \text{comfort}, \text{spaciousness}, \text{lighting}, \text{quietness}, \text{cleanliness}\}$.

Interpreting an $iT_{\mathbb{A}}^G$ $I = \langle a, O, F, L, p, S \rangle$ means finding out which property P of the facet F the label L is addressing. The label L in I can relate to the property P in two ways: it can confirm the property (e.g. *calm* w.r.t. *quietness*) or oppose it (e.g. *noisy* w.r.t. *quietness*). Therefore, we define an *interpreted* $iT_{\mathbb{A}}^G$ in the following way:

Definition 2 (iTags interpretation) For a given $iT_{\mathbb{A}}^G$ $I = \langle a, O, F, L, p, S \rangle$, an interpreted $iT_{\mathbb{A}}^G$ $\hat{I} = \langle a, O, F, L, p, (P, r), S \rangle$ where

- a, O, F, L, p, S are the same as in Definition 1;
- $P \in \mathbb{P}(F)$ is the property the label L refers to;

² The user may give a high relevance value to a given impression for several reasons. She may consider it relevant because the corresponding facet F is very important to her, or because the feature she is considering is very prominent. For example, the label *expensive* may be large because the restaurant is very expensive or because the tagger thinks that it being expensive is important in her judgment. We do not distinguish the reasons behind a given relevance value.

- $r \in \{0, 1\}$ is the relationship between L and P ; $r = 1$ means that L confirms P and $r = 0$ means that L opposes P .

As an example, consider the following iT_AGs:

$$I_1 = \langle Jo, Alfredo's, ambience, refined, +, 0.7 \rangle \quad I_2 = \langle Meg, Alfredo's, ambience, simple, +, 0.4 \rangle.$$

Both iT_AGs can be related to the property *elegance* of the facet *ambience*, but the labels *refined* and *simple* express opposite meanings. Therefore, the interpretation of these two iT_AGs would result in:

$$\hat{I}_1 = \langle Jo, Alfredo's, ambience, refined, (elegance, 1), +, 0.7 \rangle$$

$$\hat{I}_2 = \langle Meg, Alfredo's, ambience, simple, (elegance, 0), +, 0.4 \rangle.$$

We propose an automated interpretation method based on WordNet [10], which works on iT_AGs whose labels are descriptive adjectives, possibly combined with the negation *not* or with an adverb of degree, such as *very*, *scarcely*, etc. Other iT_AGs are left uninterpreted; that they will not be used in the iT_AG synthesis but will be individually visible to users.

We briefly recall that WordNet organizes adjectives in synset clusters. Each cluster C is characterized by a focal synset $foc(C)$, expressing the “main” adjective, while the other “satellite” synsets express similar, more specialized notions (e.g., if the focal synset is represented by *fast*, its satellites are *prompt*, *alacritous*, etc.). The most relevant semantic relation between adjectives is that of *antonymy*. Synset clusters come in pairs (C, \bar{C}) , where the two focal synsets are direct antonyms. Given C , we can determine its opposite \bar{C} , such that $foc(C)$ is the antonym of $foc(\bar{C})$. Satellite synsets are not considered direct antonyms, rather conceptual opposites or, as WordNet puts it, *indirect* antonyms. For example, *slow* is the direct antonym of *fast*, while *sluggish* is conceptually opposite to *alacritous*, but they are not antonyms. Hence, WordNet uses a bipolar adjective structure, the two poles being direct antonyms, each surrounded by satellites representing similar adjectives.

For our purposes, each bipolar structure in the WordNet adjective organization corresponds to a property. One of the two poles is selected as representative; any word in that pole synset or in one of its satellites *confirms* the property, while any word in the opposite pole synset or in one of its satellites *opposes* the property.

Since we also consider adverbs of degree as adjective modifiers, and WordNet does not offer any means to derive the “direction” of the modification, we pre-partition the set of adverbs of degree in two: *positive* adverbs enhance or intensify the meaning of the adjective, while *negative* ones diminish or negate it. The negative set obviously contains *not*. Therefore, our approach can be summarized as follows:

- Given a label L used for a facet F , we search for the words contained in it in WordNet, to find out whether L is indeed an adjective, possibly accompanied by an adverb of degree ad . Any other combination of words is discarded.
- If L is an adjective, we consider the WordNet cluster C it belongs to. If the noun obtained from $foc(C)$ or $foc(\bar{C})$ belongs to $\mathbb{P}(F)$, then the property we seek is represented by the pair (C, \bar{C}) . If $foc(C) \in \mathbb{P}(F)$, $r = 1$. If $foc(\bar{C}) \in \mathbb{P}(F)$, $r = 0$.
- If neither $foc(C)$ nor $foc(\bar{C})$ belongs to $\mathbb{P}(F)$, then $foc(C)$ is added in as a new property representative, and r is set to 1.

- In case L is accompanied by an adverb of degree ad , if ad is a negative adverb according to our partition, r is reversed (it becomes 1 if it was 0, and vice versa).

Interpretation allows us to understand which property the tag author is addressing, and whether she thinks the property is present or not, but it does not say whether the tag author *likes* the presence or absence of that property. In the above example, I_1 and I_2 express opposite opinions on the *elegance* property. However, two $iT_{\mathcal{A}}^G$ s may express the same opinion with opposite judgements: two $iT_{\mathcal{A}}^G$ authors may both think that the resource or the facet has a given property, but one of them likes it, the other one does not. This is a difference in *polarity*. If we consider the *relationship* r between label and property, and the *polarity* p of the $iT_{\mathcal{A}}^G$ author's impression, we have four different possibilities. Each of these possible combinations is called an *aspect* of the property. For example (see Figure 3), the four aspects of the *elegance* property could be represented by the labels *chic* (“it’s elegant, I like this”, $p = +$, $r = 1$), *sophisticated* (“it’s too elegant, I don’t like this”, $p = -$, $r = 1$), *simple* (“it’s not elegant, but I like this”, $p = +$, $r = 0$), *shabby* (“it’s not elegant at all, I don’t like this”, $p = -$, $r = 0$).

4 iTags Synthesis

The aim of $iT_{\mathcal{A}}^G$ synthesis is to provide users with a comprehensive and immediate aggregation of what people think about a given object, and to offer an effective representation of the overall opinion, which is the most meaningful for the user. In doing this, we take into account:

- the existence of niches of people whose opinions differ from the majority³;
- the *social network* of the target user, since other people’s opinions weigh differently depending on how much the user trusts them on the topic⁴;
- the *user model* of the target user, considering only the facets relevant for her.

In order to produce a meaningful synthesis that takes into account the above issues, we partition the set of $iT_{\mathcal{A}}^G$ s associated with a given object O first according to the facet F and property P , and then according to the *aspect*, i.e. the (*relationship, polarity*) pair. The rationale behind this lies in our approach to synthesis, which is the following:

1. We merge all $iT_{\mathcal{A}}^G$ s that refer to the same facet, the same property, and have the same relationship and polarity. These $iT_{\mathcal{A}}^G$ s are essentially stating the same concept, only in different words (labels) and with different relevance (sizes). In order to merge them, we need to select a representative label and find an average size. As we will see, in doing this we will take into account the social network and the trust level.
2. We decide which facets we want to show to a given user, by considering her interest in them as expressed in the user model.

³ As an example, suppose that 30 people out of 50 think that a restaurant is *cheap* while the remaining 20 think it is *expensive*. Going with the prevailing opinion would mean showing only the *cheap* fraction. On the other hand, computing an average of the users’ impressions, imagining that *cheap* and *expensive* lie on the same scale of *cheapness* with opposite signs, would lead to showing something like *moderately cheap*. We think that none of these solutions correctly portrays the collective impression on the restaurant.

⁴ In the case of a restaurant, one would probably trust more the impression of a well-known enogastronomic journalist, than those of her gym friend who usually goes for the cheapest meal around.

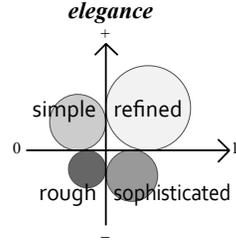


Fig. 3. The four synthesis aspects of the *elegance* property

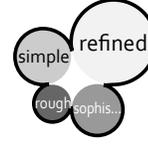


Fig. 4. The synthesized itag for the *elegance* property

3. For such facets, we show the merged iTAGs for all four aspects, provided they are relevant enough (i.e. the resulting merged size is above a given threshold). We wish to show a minority's opinion only if it is a significant.

As an example of a possible result, Figure 3 shows the merged iTAGs for the four aspects of the *elegance* property in the *ambience* facet, while Figure 4 shows the visualization in the iTAG system of the synthesized iTAG. Where a single iTAG appears as a circle around its facet, a synthesized iTAG resembles a flower with at most four petals.

Let us now formalize how we compute a label L and a size S for a merged iTAG. Recall that we merge a set of iTAGs that refer to the same object \bar{O} , to the same facet \bar{F} and to the same property \bar{P} , and that belong to the same aspect (\bar{r}, \bar{p}) . We will denote the set of such iTAGs as $iTAGs(\bar{O}, \bar{F}, \bar{P}, \bar{r}, \bar{p}) = \{\hat{I}_1, \dots, \hat{I}_n\}$.

Both label and size depend on the user u asking for the synthesis (*target user*), and on her social relationship with the iTAGs *authors*, since the u may value more the opinions of specific people (e.g. official experts or trusted people in her social network).

In what follows we clarify how people associated with the target user u influence the iTAG synthesis. The people belonging to the model of the target user u are divided into groups g_1^u, \dots, g_m^u called *trust groups*. Each group g_j^u , $j = 1, \dots, m$, has a *weight* $\omega_j^u \in [0, 1]$ associated with it. The user u trusts differently people in each group, e.g. she could have a group of “friends” weighted 0.6 and a group of “experts” weighted 1. The weight assigned by the target user to each group can be initially set to default values and then tuned according to her preferences or her behavior in the social application.⁵ We denote by $\{a_1, \dots, a_k\}$ the tagging authors who have tagged the object O .

Definition 3 (Trust value) For the tagging author a_j , $j \in \mathbb{N}$, its trust value $\tau_{a_j}^u \in [0, 1]$ w.r.t. the target user u is computed as:

1. If $\forall i \in \mathbb{N}$, $a_j \notin g_i^u$, $\tau_{a_j}^u = 0$;
2. If $\exists ! i \in \mathbb{N}$, $a_j \in g_i^u$, $\tau_{a_j}^u = \omega_i^u$;

⁵ Notice that the way the trust level is computed is out of the scope of this paper. It is understood that the trust values are topic dependent. In case there is no trust value, the approach works by considering a unique group.

3. If $\exists k \in \mathbb{N}$, $a_j \in g_1^u \cap \dots \cap g_k^u$ and $\omega_1^u, \dots, \omega_k^u$ are the weights for g_1^u, \dots, g_k^u ,

$$\tau_{a_j}^u = \omega_1^u + \sum_{j=2}^k \omega_j^u \prod_{h=1}^{j-1} (1 - \omega_h^u). \quad (1)$$

This means that: (i) the opinion of an author not belonging to any group of interest for the target user is discarded; (ii) if the author belongs just to one interest group her trust value is equal to the weight of that group; and (iii) if the author belongs to several interest groups for the target user, her trust value is computed according to Formula 1.

The following definition computes a weight for each interpreted $iTAG$, taking into account the sizes of the interpreted $iTAG$ s, and the trust values of the tagging authors.⁶ The weight represents the contribution that each $iTAG$ gives to the synthesis.

Definition 4 (Weight of interpreted $iTAG$) For a given target user u , the weight of an interpreted $iTAG$ $\hat{I}_j \in iTAGS(\bar{O}, \bar{F}, \bar{P}, \bar{r}, \bar{p}) = \{\hat{I}_1, \dots, \hat{I}_n\}$ is computed as follows:

$$W(\hat{I}_j) = S_j \tau_{a_j}^u / \sum_{a \in AUTH(\bar{O}, \bar{F})} \tau_a^u \quad (2)$$

- $\tau_{a_j}^u$ is the trust value of the author a_j w.r.t. the target user u (see Eq. 1);
- S_j is the size of the $iTAG$ provided by the author a_j (we assume $S_j \in [\theta, 1], \theta > 0$).
- $AUTH(\bar{O}, \bar{F})$ is the set of the authors that tagged the facet $\bar{F} \in \mathbb{F}(type(\bar{O}))$.

The weight of each $iTAG$ can be computed w.r.t. all tagging authors, w.r.t. all tagging authors who have tagged F , or w.r.t. the tagging authors who have tagged P . In our opinion the best option to consider is the second one: a tagging user neglecting a whole facet probably means she does not find it relevant, while a tagging user not mentioning a property in a facet she is tagging probably means she has a neutral opinion with respect to that property.

We use the computed weight to select a label L for the merged $iTAG$ in the set $\{L_1, \dots, L_h\}$ of all labels used in $iTAGS(\bar{O}, \bar{F}, \bar{P}, \bar{r}, \bar{p}) = \{\hat{I}_1, \dots, \hat{I}_n\}$, then for a given label L_i we can assume without losing generality that $\{\hat{I}_1, \dots, \hat{I}_{h-1}\}$ are the ones using L_i while $\{\hat{I}_i, \dots, \hat{I}_n\}$ are the ones using some other label.

Definition 5 (Resulting label) The weight associated with L_i is given by:

$$W(L_i) = \sum_{j=1}^{h-1} W(\hat{I}_j) \quad (3)$$

Then we select as label L for the merged $iTAG$ the L_i with the highest value for $W(L_i)$.

The size S of a merged $iTAG$ is calculated by adding the weights of the corresponding interpreted $iTAG$ s. To avoid $iTAG$ s of extremely small size, we introduce a minimum threshold $\theta \in (0, 1]$.

⁶ The sum of the trust values is not 1, therefore we use it to normalize the weighted size.

Definition 6 (Resulting size) The size S of the merged iTAG for $iTAGS(\bar{O}, \bar{F}, \bar{P}, \bar{r}, \bar{p}) = \{\hat{I}_1, \dots, \hat{I}_n\}$ is:

$$S = \begin{cases} \sum_{i=1}^n W(\hat{I}_i) & \text{if } \sum_{i=1}^n W(\hat{I}_i) \geq \theta \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

5 A preliminary evaluation of the approach

In order to evaluate whether our iTAG framework effectively achieves the goals of (i) allowing the communication of opinion and (ii) providing effective collective opinion’s synthesis, we carried out a preliminary evaluations with users, targeted at answering the following research questions:

Research Question 1 (RQ1). Does our tagging framework allow people to express complex opinions? Is it intuitive to use and easily understandable?

Research Question 2 (RQ2). Is the personalized multi-facet synthesis more informative than traditional tag clouds?

We selected 38 users⁷ which were divided into two groups. The first one, consisting of 18 participants, used the framework to provide and examine single iTAGs; their experience allowed us to investigate **RQ1**. The second group, with the remaining 20 participants, had to evaluate synthesized iTAGs and thus provide information concerning **RQ2**.

We asked the users in the first group to comment on 3 restaurants (tagging at least 2 facets for each restaurant), using iTAGs. Afterwards, all the users were asked to look at 3 iTAGs provided by other people, and describe their understanding of the other people’s opinion on those restaurants.

The 20 users of the second group were then asked to chose 2 restaurants from a list of 20 restaurants that were tagged in the previous phase. Each restaurant had a detailed description of different facets (price, cuisine, ambient, etc). The tags for the restaurants of their choice were presented both as a traditional tag cloud and as synthesized iTAGs; the users were asked to compare the two presentations.

Finally, all users answered a questionnaire about their experience with the iTAG system.

The goal of the evaluation was to answer address iTAGs’ expressiveness and immediacy in conveying opinions, the overall usability of the interface, and the correct interpretation of iTAGs. More specifically we examined:

1. *Expressiveness of iTAG editing interface.* This implies two further questions:
 - Does the iTAG framework allow users to freely express opinions?
 - Is the specific interface we developed usable?

Regarding both questions, users had to provide a number on a 4 point scale (1 being “absolutely no” and 4 being “definitely yes”). For the first experiment, we obtained a mean of 3.3 and a mode of 4. For the second experiment we obtained a mean of 3.2 with a mode of 3. We can conclude that the system expressiveness is good.

⁷ Users were recruited among the contacts and colleagues of the authors, according to an availability sampling strategy. Even if non-random samples are not statistically representative, they are often used in psychology research and usability testing, during early evaluation phases.

2. *Expressiveness of iTAG viewing interface.* This implies two further questions:
 - Does the iTAG framework allow users to correctly understand the opinion the iTAG authors wanted to communicate?
 - Are iTAGs immediate and do they communicate opinions at a glance?

Users had the chance to express a first impression looking only briefly at the iTAG, and then to examine in more depth the iTAG structure possibly zooming over the smaller labels. All users correctly understood the taggers' opinions after the first brief examination. We also asked the users what they thought the circle size expressed. Most of the users (87% in the first experiment, 83% in the second one) answered that they interpreted the circle size as the relevance of the iTAG with respect to the overall comment. The remaining people either saw no particular meaning associated to size, or thought it was a quantification of the label. Finally, we asked the users which was in their opinion the major advantage of iTAG system (if any) with respect to traditional tagging systems. For 61% the advantage of iTAG system is immediacy: opinions can be understood at a glance; for 17% the advantage is the possibility to refer the words to different facets of the resource. Hence, we can conclude that iTAG framework is expressive and communicatively rich.
3. *Expressiveness of the synthesis.* 89% of the users preferred our synthesis to the traditional tag cloud, for the following reasons: (i) the overall opinion is clear at a glance (78%); (ii) it is quicker to read (60%); (iii) it is more informative (71%); (iv) it presents also the niche opinions (82%).

6 Related Work

Our work aims at enhancing tagging with capability of express complex opinions, in order to provide a personalized synthesis of tags in social applications. This implies (i) interpreting and synthesizing the tags, (ii) adapting the synthesis to target users and (iii) visualizing the synthesized tags.

Interpretation and synthesis. The interpretation techniques for tags depend on the tags typology (free, structured or facet-based tags). *Free tags* give the highest freedom to users, but they are difficult to process and interpret. Some work uses techniques from machine learning and artificial intelligence [13]; others use clustering methods or similar mathematical approaches [4]. Yet another approach is to map the tags to an existing domain ontology or semantic knowledge base as DBpedia [15], using some similarity measures to compute the distances between words from a syntactic [6] and semantic point of view [5]. *Structured tagging* provides more information, since it forces users to focus on a specific subject and to assign values to a set of predefined metadata fields (see BibSonomy [13] for documents and VRA Core Vr4⁸ for multimedia). Although tag interpretation is easier for structured tagging, too much complexity discourages users from providing tags. To solve the processing problem, we adopted a compromise between freedom and structure using *facets* [17], which involve creating a bottom-up classification. There are several proposals for applying facets to social tagging applications mostly with the aim to classify a tag associating it to one or more facets [18]. In this case, facets are tag categories (people, time, place, etc.), possibly organized in a

⁸ www.vraweb.org/projects/vracore4/VRA.Core4_Intro.pdf

hierarchy, that can help to clarify the meaning of the tag [16]. Even though our facets are organized in a hierarchical structure, they do not serve as tag classifiers. They rather represent different features the user can comment on by expressing her opinion with the iT_AGs.

Personalization. To our knowledge, there are no other proposals to personalize tag clouds according to trust measures. Some works in information retrieval proposed personalized search tools that return only the tags that agree with the user main interests given in her user profile [7]. In social bookmarking systems several authors propose to recommend tags to users, i.e., to propose the tags that better fit a user's needs. Our aim is different: we do not suggest resources to users, we rather provide a synthesis of the opinions of the people they trust. In this sense, our work is similar to trust-based recommender systems, which generate personalized recommendations by aggregating the opinions of users in the trust network. Even if it is not our main goal, another application of our work is to use iT_AGs for recommending resources to users, as it has been recently proposed for example in collaborative filtering approaches, i.e., to recommend an item given the similarity of its tags with the tags used for another item she liked [9], or to compute users' similarity starting from the tags they used [14].

Visualization. In social web, the method used the most for visually representing information are fuzzy aggregations of *tag clouds* [12], where terms are organized in alphabetical order and presented in a compact space. Tag clouds enable visual browsing by showing the list of the most popular tags, alphabetically ordered and weighted by font size and thickness. The selection of tags to be shown in clouds is based on frequency, which results in high semantic density and in a limited number of different topics dominating the whole cloud [4]. Moreover, alphabetical order is convenient only when the user already knows what she is looking for. In fact, tag clouds facilitate neither visual scanning nor representation of semantic relationships among tags.

In [19] the authors present *tag expression*, a tag cloud-based interface that allows users to rate a movie with a tag and an associated feeling (like, dislike, or neutral) which measures the user's opinion about the movie. Similarly to us, the user can express an opinion on one of item features (tags in *tag expression* and facets in iT_AG), but the synthesis approaches are different. In fact, they simply flatten the opinions in a unique average value, not considering different points of view.

In order to overcome the limitations of traditional tag clouds, several methods have been proposed. [11] presents a new tag cloud layout that shows tag similarity⁹ at a glance. Based on co-occurrence similarity, data clustering techniques are used to aggregate tags into clusters whose members are similar to each other and dissimilar to members of other clusters. The result is a tag cloud where semantically similar tags are grouped horizontally whereas similar clusters are vertical neighbors. Instead of clustering, we propose the four aspects layout as a form of iT_AG aggregation, since the aim of our synthesis is to summarize the impressions of different users about a certain property of a facet, by condensing agreeing options and by relating the opposite ones. An alternative approach is proposed by *tagFlakes* [8], a system that helps the user navigate tags in a hierarchical structure, where descendant terms occur within the context defined by

⁹ One easy and commonly used technique to evaluate the similarity of two tags is to count their co-occurrences, i.e., how many times they are used to annotate the same resource

the ancestor terms. Similarly, we adopt a form of aggregation, in which we group tags that refer to the same property. However, our aggregation is *multi-aspect*, in order to show and highlight disagreements, as well as similarities.

7 Conclusions and future work

We presented an enhanced tagging framework, called iT_A^G , which allows the users of a social application to share their opinions on resources, and that allows for personalized synthesis of users' opinions, enabling easier understanding of a huge amount of users' tags on a resource. We introduce a method for interpreting tags, which partitions the tags according to the *property* of the resource they describe and their *relationship* with the given property. Next, we propose a personalized synthesis method which takes into account how much the user trusts the tag authors. We assume that the trust measure is provided by the social application using some existing state of the art approach (such as [2]). This is out of the scope of the paper. The main contributions of our work are the following:

- (i) A novel tagging modality, which enables expressing complex opinions on a resource, and at the same time enables their interpretation by the system.
- (ii) A method for iT_A^G s interpretation, which assigns iT_A^G s to specific properties, without using Natural Language Processing techniques.
- (iii) A personalized synthesis of users' opinions using the above interpretation. The iT_A^G s synthesis is (i) *multi-faceted*, it maintains the differences in opinions of different users, and (ii) *personalized*, it is based on the user model and her social network.
- (iv) A visualization of the synthesis that allows for an immediate understanding of the collective opinion.

Preliminary evaluation showed the users appreciated the iT_A^G framework. We intend to perform a more rigorous and comprehensive evaluation in the future.

Another open problem is how to resolve the problem of labels polysemy in the process of mapping labels to properties.

Currently we are working on the use of iT_A^G s for recommendation purposes, by defining a notion of *distance* between iT_A^G s. The goal is twofold: being able to recommend a resource to a user due to the similarity between the resource reputation and the user's tastes, and being able to find similar users for social recommendation purposes. At the same time, we are working on exploiting such tag-based opinions to enrich the user model, also with dislike values.

As future work, we aim at investigating the possibility to combine our research with the results of the Sentiment Analysis field, which aims to determine the attitude of a person with respect to some topic. In particular, we plan to exploit in our framework the SentiWordnet system,¹⁰ a lexical resource for opinion mining.

¹⁰ <http://sentiwordnet.isti.cnr.it/>

References

1. R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tenenholzt. Trust-based recommendation systems: an axiomatic approach. In *17th Int. Conference on World Wide Web, WWW 2008*, pages 199–208, New York, USA, 2008. ACM.
2. L. Aroyo, P. D. Meo, and D. Ursino. Trust and reputation in social internetworking systems. In *Proc. of the Workshop on Adaptation in Social and Semantic Web (at UMAP10)*, 2010.
3. J. Bar-Ilan, S. Shoham, A. Idan, Y. Miller, and A. Shachak. Structured versus unstructured tagging: a case study. *Online Information Review*, 32(5):635–647, 2008.
4. G. Begelman, P. Keller, and F. Smadja. Automated Tag Clustering: Improving search and exploration in the tag space. In *Collaborative Web Tagging Workshop at WWW 2006*, 2006.
5. C. Cattuto, D. Benz, A. Hotho, and G. Stumme. Semantic grounding of tag relatedness in social bookmarking systems. In *7th Int. Semantic Web Conference, ISWC '08*, pages 615–631. Springer-Verlag, 2008.
6. F. J. Damerau. A technique for computer detection and correction of spelling errors. *Communications of the ACM*, 7(3):171–176, 1964.
7. M. Daoud, L. Tamine, and M. Boughanem. A personalized graph-based document ranking model using a semantic user profile. In P. D. Bra, A. Kobsa, and D. N. Chin, editors, *18th Int. Conference on User Modeling, Adaptation, and Personalization, UMAP '10*, volume 6075 of *Lecture Notes in Computer Science*, pages 171–182. Springer, 2010.
8. L. Di Caro, K. S. Candan, and M. L. Sapino. Using tagFlake for Condensing Navigable Tag Hierarchies from Tag Clouds. In *14th ACM SIGKDD Int. Conference on Knowledge Discovery and Data Mining KDD '08*, pages 1069–1072. ACM, 2008.
9. F. Durao and P. Dolog. A personalized tag-based recommendation in social web systems. In *Int. Workshop on Adaptation and Personalization for Web 2.0 at UMAP 2009*, 2009.
10. C. Fellbaum, editor. *WordNet: An Electronic Lexical Database*. MIT Press, 1998.
11. Y. Hassan-Montero and V. Herrero-Solana. Improving Tag-Clouds as Visual Information Retrieval Interfaces. In *Int. Conference on Multidisciplinary Information Sciences and Technologies InScit2006*, 2006.
12. O. Hoeber and H. Liu. Comparing tag clouds, term histograms, and term lists for enhancing personalized web search. In *Int. Workshop on Web Information Retrieval Support Systems at IEEE/WIC/ACM Int. Conference on Web Intelligence*, pages 309–313, 2010.
13. A. Hotho, R. Jäschke, C. Schmitz, and G. Stumme. Emergent semantics in bibsonomy. In *Informatik 2006 – Informatik für Menschen. Band 2*, volume P-94 of *Lecture Notes in Informatics*. Gesellschaft für Informatik, 2006.
14. H. Liang, Y. Xu, Y. Li, and R. Nayak. Tag based collaborative filtering for recommender systems. In *Proc. 4th Int. Conference on Rough Sets and Knowledge Technology, RSKT '09*, *Lecture Notes in Computer Science*, pages 666–673. Springer-Verlag, 2009.
15. R. Mirizzi, A. Ragone, T. Di Noia, and E. Di Sciascio. Semantic tag cloud generation via dbpedia. In *11th Int. Conference on Electronic Commerce and Web Technologies, EC-Web '10*, *Lecture Notes in Business Information Processing*. Springer Verlag, 2010.
16. E. Quintarelli, A. Resmini, and L. Rosati. FaceTag: Integrating Bottom-up and Top-down Classification in a Social Tagging System. In *IA Summit 2007: Enriching Information Architecture*, 2007.
17. S. R. Ranganathan. *Prolegomena to Library Classification*. New York: Asia Publishing House, 1967.
18. L. Spiteri. Incorporating facets into social tagging applications: An analysis of current trends. *Cataloguing and Classification Quarterly*, 48(1):94–109, 2010.
19. J. Vig, M. Soukup, S. Sen, and J. Riedl. Tag expression: tagging with feeling. In *23rd ACM Symp. on User interface software and technology*, pages 323–332, 2010.