

An Emotion Dimensional Model Based on Social Tags: Crossing Folksonomies and Enhancing Recommendations

Ignacio Fernández-Tobías, Iván Cantador, and Laura Plaza

Universidad Autónoma de Madrid, 28049 Madrid, Spain
{i.fernandez, ivan.cantador, laura.plaza}@uam.es

Abstract. In this paper we present an emotion computational model based on social tags. The model is built upon an automatically generated lexicon that describes emotions by means of synonym and antonym terms, and that is linked to multiple domain-specific emotion folksonomies extracted from entertainment social tagging systems. Using these cross-domain folksonomies, we develop a number of methods that automatically transform tag-based item profiles into emotion-oriented item profiles. To validate our model we report results from a user study that show a high precision of our methods to infer the emotions evoked by items in the movie and music domains, and results from an offline evaluation that show accuracy improvements on model-based recommender systems that incorporate the extracted item emotional information.

Keywords: emotions, folksonomies, cross domains, recommender systems.

1 Introduction

Emotions are intense feelings that are directed at someone or something. For instance, a person may be glad when she comes across an old friend, and may be excited when she receives a gift. Moods, in contrast, are feelings that tend to be less intense than emotions, and often – though not always – lack a contextual stimulus [8]. Moreover, emotions are more transitory than moods. Quoting the example given in [11], a person may feel angry when someone has been rude to her. This intense feeling of anger probably comes and goes quickly, maybe in a matter of seconds. In contrast, when a person is in a bad mood, she could feel bad for several hours.

Emotions and moods can be comprised in the generic concept of *affect* [11]. Emotions can turn into moods when there is a loss of focus on the contextual stimuli (people, objects or events) that started the feelings. In the opposite direction, moods can elicit more emotional responses to contextual stimuli. In this paper, since we aim to model the mostly ephemeral feelings caused by entertainment items – such as movies and music –, we use the term *emotions* to refer to both emotions and moods.

In adaptive and personalized systems, emotions are usually considered as contextual signals that can lead to enhanced approaches in a wide array of applications, such as constructing user behavior models [10], tailoring search results [13], and filtering and recommending items [18], to name a few. Hence, modeling, capturing and exploiting emotions present challenging problems that are addressed in distinct Computer Science

research areas that intersect with Psychology and Social Sciences, such as Human Computer Interaction [7], Artificial Intelligence and Robotics [3], Opinion Mining and Sentiment Analysis [6], and Information Access and Retrieval [16]. Here we focus on the **emotion modeling** task, and restrict our attention to situations where emotions are expressed in (and can be extracted from) **text contents** – such as reviews in blogs, and annotations in social tagging systems –, differently to e.g. situations where emotions are recognized in either the visual or auditory modalities [7][19].

Computational models of emotion can be categorized according to the emotional theories they adopt, namely the *categorical emotion theory* – which characterizes emotions as discrete units –, the *emotional dimension theory* – which conceive emotions as points in a continuous space –, and the *appraisal theory* – which represents emotions as outcomes of events and situations. Our model adopts the **emotional dimension theory** by representing an emotion as a vector, whose components correspond to terms in an emotion lexicon, and have assigned positive or negative weights depending on whether their terms are synonyms or antonyms of labels that describe the emotion. As we shall show in this paper, the projections of our emotion vectors into a two-dimension space are in accordance with the psychological **circumplex model of affect** [15].

The input data used for capturing and modeling emotions can also be used to categorize the existing computational models of emotion. Hence, we can distinguish linguistic approaches that extract emotions from *text*, image processing approaches that recognize emotions in facial expressions from *images* and *videos*, and speech recognition approaches that identify emotions on *audio* data. The linguistic approaches usually create or make use of text corpora and resources – such as lexicons, thesauri and ontologies – that provide specific vocabularies for describing emotions. In this paper we propose an approach that generates a **lexicon** and **folksonomies** to represent generic emotions and domain-specific emotional categories. These resources are automatically generated from a generic thesaurus and social tagging systems in entertainment domains, namely the movie, music and book domains.

Using the generated emotion lexicon and cross-domain folksonomies, we develop a number of methods that transform tag-based item profiles into **emotion-oriented item profiles**. We evaluated the quality of such profiles by conducting a user study, whose results show a high precision of our methods to infer the emotions evoked by items in the movie and music domains. Moreover, we performed an offline evaluation, whose results show that exploiting the extracted emotional information improves the accuracy of various **model-based recommender systems** on the above domains.

2 Related Work

The study and development of computational systems aimed to recognize, interpret and process human feelings is usually referred to as Affective Computing. This discipline involves a number of research fields and applications. In Artificial Intelligence, for instance, endowing robots with emotions for improving human-robot interaction has been largely studied [3].

Emotion recognition in natural language is becoming increasingly important as well. One of the most outstanding applications concerns discovering the affective relevance

of user online reviews of products and services [6]. Other works have focused on annotating texts (e.g. news items and tweets) with emotions [14].

The use of emotions in User Modeling and Recommender Systems is mainly concerned with detecting and modeling the user’s mood, and suggesting items according to such mood [18]. Kaminskas and Ricci [12] present an approach to recommend music compositions for places of interest by means of social tags that represent the user’s emotional state when listening to music and visiting places. To attach emotional tags to music, they use the Geneva Emotional Music Scale (GEMS) model [19]. Others have studied how to describe music in terms of the emotions it evokes. Feng et al. [9] map two dimensions of tempo and articulation into mood categories, such as *happiness*, *sadness* and *fear*. Shi et al. [16] propose a mood-specific movie similarity, which is exploited in a joint matrix factorization model for enhanced context-aware (mood-specific) recommendations.

As done by Baldoni et al. [2], we propose to extract emotional information from item annotations in social tagging systems. However, while they use ontologies and lexicons to assist the identification of emotions, we automatically derive emotions based on simple domain-specific emotional categories existing in specialized systems, such as the Jinni¹’s movie categories and GEMS’ music categories. Moreover, to make our approach generic and ensure cross-domain interoperability, the domain-specific emotional categories are mapped to the general and well accepted emotions of Russell’s circumplex model [15].

3 A Core Emotion Lexicon

Among the existing dimensional models of emotion, the circumplex model is a dominant one. It suggests that emotions are distributed in a two-dimensional circular space formed by two independent dimensions: *arousal* and *pleasure*. Figure 1a shows such distribution. Arousal represents the vertical axis and reflects the intensity of an emotion; and pleasure represents the horizontal axis and reflects if an emotion is positive or negative. The center of the circle represents medium levels of arousal and pleasure. Any emotion can be represented at any level of arousal and pleasure, including a neutral level of one or both of such factors. Hence, for instance, *happiness* and *sadness* can be considered as emotions with the highest and lowest levels of pleasure, respectively, but with neutral arousal levels, with respect to other emotions such as *tension* (with high arousal) and *calmness* (with low arousal). The figure shows the distribution of 16 core emotions. Our model also considers this set of emotions.

The dimensional model we propose is built upon an automatically generated lexicon $\mathcal{L} = \{t_1, \dots, t_K\}$ composed of synonym and antonym terms t_k of the core emotions’ names – which are adjectives (e.g. *happy*, *sad*), as shown in Figure 1a. The synonym and antonym terms of each emotion are obtained from the online thesaurus provided by Dictionary.com². Specifically, the lexicon is composed of the synonyms

¹ Jinni movie search and recommendation engine, <http://www.jinni.com>

² Dictionary.com thesaurus, <http://thesaurus.com>




Fig. 1. Two-dimensional distributions of core emotions established in the circumplex model (1a) and automatically obtained in our tag-based model (1b)

Table 1. Considered core emotions and some of their synonym terms

Emotion	Synonym terms	Emotion	Synonym terms
<i>alert</i>	alert, active, animated, lively, sprightly	<i>fatigued</i>	fatigued, tired, fatigued, drained, exhausted
<i>excited</i>	excited, stimulated, agitated, moved	<i>bored</i>	bored, apathetic, exasperated, indifferent
<i>elated</i>	elated, jubilant, overjoyed, exhilarated	<i>depressed</i>	depressed, dejected, despondent, disconsolate
<i>happy</i>	happy, merry, cheerful, joyful, bright	<i>sad</i>	sad, sorrowful, doleful, downcast, gloomy
<i>content</i>	content, satisfied, gratified, pleased, enjoyed	<i>upset</i>	upset, bother, disturbed, troubled, distressed
<i>serene</i>	serene, quiet, placid, tranquil, peaceful	<i>stressed</i>	stressed, tormented, harassed, vexed, irked
<i>relaxed</i>	relaxed, moderated, mitigated, loose, free	<i>nervous</i>	nervous, apprehensive, uneasy, disturbed
<i>calm</i>	calm, mild, appeased, smooth, soften	<i>tense</i>	tense, restless, uptight, jittery, restive

and antonyms of all noun, adjective and verb entries in the above thesaurus for the emotions' names. Table 1 shows some of the gathered synonyms for each emotion.

Once the lexicon \mathcal{L} is generated, a core emotion $e_i \in \mathcal{E}$ is represented as a vector $\mathbf{e}_i = (e_{i,1}, \dots, e_{i,K}) \in \mathbb{R}^K$, in which the component $e_{i,k}$ corresponds to the term $t_k \in \mathcal{L}$ (that can describe various emotions), and is a numeric value defined as:

$$e_{i,k} = \begin{cases} tfidf(t_k, e_i) & \text{if } t_k \in \text{synonyms}(e_i) \\ -tfidf(t_k, e_i) & \text{if } t_k \in \text{antonyms}(e_i) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The component $e_{i,k}$ is greater than 0 if the term t_k is a synonym of the emotion e_i , lower than 0 if t_k is an antonym of e_i , and 0 otherwise. Its absolute value corresponds to the TF-IDF weight of t_k computed by considering the lexicon \mathcal{L} as the *collection vocabulary*, and the set \mathcal{E} of emotions (described as sets of synonym and antonym terms) as the *collection documents*. Formally,

$$tfidf(t_k, e_i) = tf(t_k, e_i) \cdot idf(t_k, \mathcal{E})$$

where $tf(t_k, e_i)$ is the normalized term frequency of t_k for e_i , which measures how relevant the term is to describe the emotion, and which is defined as:

$$tf(t_k, e_i) = \frac{f(t_k, e_i)}{\max \{f(t, e_i) : t \in synonyms(e_i) \cup antonyms(e_i)\}}$$

being $f(t_k, e_i)$ the number of times that term t_k appears in the sets of synonyms and antonyms of e_i 's thesaurus entries; and where $idf(t_k, \mathcal{E})$ is the inverse document frequency of t_k in \mathcal{E} , which measures how rare (and thus informative) is the term across all the emotions' descriptions, and which is defined as:

$$idf(t_k, \mathcal{E}) = \log \frac{|\mathcal{E}|}{|\{e \in \mathcal{E} : t_k \in synonyms(e) \cup antonyms(e)\}|}$$

With the proposed vector representation, we can measure (dis)similarities between emotions. Specifically, we can use the cosine similarity $sim(e_i, e_j) = \cos(\mathbf{e}_i, \mathbf{e}_j)$.

To validate the correspondences between our computational model and the theoretic circumplex model, Figure 1b shows the projections of the emotion vectors into a two-dimensional space by applying Principal Component Analysis. We can see that our model locates all the 16 basic emotions in their corresponding quadrants. More interestingly, in our model the axes defined by the two most informative components are related to the *arousal* and *pleasure* factors of the circumplex model. Thus, positive emotions (e.g. *happy*, *calm*) are in the right quadrants, while negative emotions (e.g. *sad*, *upset*) are in the left ones, for the horizontal (pleasure) axis; and more intense emotions (e.g. *tense*, *alert*) are in the upper quadrants, while less intense emotions (e.g. *relaxed*, *bored*) are in the lower quadrants, for the vertical (arousal) axis.

We note that, as done in [4], we tested other term weighting methods – such as the BM25 probabilistic model – and emotion similarity functions – such as the Jaccard similarity. We finally used the TF-IDF weighting method and cosine similarity since they let generate the two-dimensional distribution of emotions closest to the circumplex model's. We also note that we did not perform any cleaning and filtering process on the original sets of synonyms and antonyms obtained from the online thesaurus. Such process may increase the quality of the representations (e.g. by discarding ambiguous terms), and thus may let generate a better emotion distribution.

4 A Cross-Domain Emotion Folksonomy

In a social tagging system users create or upload items, annotate them with freely chosen tags, and share them with other users. The whole set of tags constitutes an unstructured knowledge classification scheme that is known as *folksonomy*. This implicit classification is then used to search and recommend items. The purpose for tagging is manifold: describing the content of the items, providing contextual information about the items, expressing qualities and opinions about the items, or even stating self-references and personal tasks related to the items.

Within the set of tags that express qualities and opinions about the items, there are tags that refer to emotions caused by the annotated items. In most cases, however, such emotions are not the core emotions presented in Section 3, but domain-specific emotional categories – such as *suspense* in the movie domain, and *nostalgia* in the music domain –, which indeed may be related to one or more core emotions.

In this section we extend our emotion model by linking the core emotions with domain-specific emotional categories described by tags in different folksonomies. Specifically, we focus on the movie and music entertainment domains by exploiting the MovieLens and Last.fm folksonomies provided in the HetRec’11 workshop [5] (Sections 4.1 and 4.2). With the extended model we propose to build emotion-oriented item profiles (Section 4.3) and cross-domain folksonomies (Section 4.4). We make all the generated data – lexicon, folksonomies, and item profiles – publicly available³.

4.1 An Emotion Folksonomy for Movies

To build an emotion folksonomy in the movie domain, we first select a total of 15 emotional categories listed under the *mood* topic in Jinni movie search and recommendation system. We describe each category by 4 to 6 associated feeling terms, and use them as seed terms (see Table 2). Next, we extend the seed terms with their synonyms and antonyms in Thesaurus.com, but restricted to those existing as social tags in the MovieLens dataset. Finally, we repeat the process explained in Section 3 to represent an emotional category as a vector of weighted terms. In this vector, positive components represent synonyms while negative components represent antonyms. In this way, each emotional category is represented as a set of tags that lets establish (dis)similarities with other categories.

Table 2. Considered movie emotional categories and seed terms

Category	Seed terms	Category	Seed terms
<i>clever</i>	clever, cerebral, reflective	<i>sexy</i>	sexy, erotic, sensual
<i>offbeat</i>	offbeat, quirky, surreal, witty	<i>sexual</i>	sexual, lascive, horny
<i>exciting</i>	exciting, energetic, frantic, forceful	<i>uplifting</i>	uplifting, inspirational, hope
<i>suspenseful</i>	suspenseful, tense	<i>bleak</i>	bleak, grim, depressing, hopeless
<i>captivating</i>	captivating, rousing, poignant	<i>gloomy</i>	gloomy, sad, melancholic, nostalgic
<i>emotional</i>	emotional, passionate, romantic	<i>rough</i>	rough, brutal, lurid, macabre, wry
<i>feel good</i>	cute, merry, happy	<i>scary</i>	scary, creepy, menacing, eerie
<i>humorous</i>	humorous, funny, comical		

Figure 2a depicts the cosine similarity values between each pair of emotional categories (green/red cells correspond to positive/negative values). It can be observed

³ Emotion lexicon, folksonomies, profiles, and online evaluation tool,
<http://ir.ii.uam.es/emotions>

that close emotional categories, such as *gloomy* and *bleak*, present high similarity, while very distinct categories, such as *gloomy* and *feel good*, present low similarity.

4.2 An Emotion Folksonomy for Music

To generate an emotion folksonomy in the music domain, we select as emotional categories the 9 emotions proposed in the GEMS model (see Table 3). As initial seed terms we use the category names and their associated feeling terms given in [19]. Next, we extend these terms with their synonyms and antonyms in Thesaurus.com, but restricted to those existing as social tags in the Last.fm dataset. The emotional category vectors are then created as for the movie domain. Table 3 shows some of the most informative tags for each emotional category.

Figure 2b shows the similarity values between each pair of emotional categories. Again, close categories, such as *tenderness* and *nostalgia*, present high similarity, while very distinct categories, such as *sadness* and *joy*, present low similarity.




Fig. 2. Cosine similarity values between movie and music emotional categories

Table 3. Considered music emotional categories and seed terms

Category	Main tags	Category	Main tags
<i>joy</i>	funny, happy, amusing, jolly	<i>tenderness</i>	tender, gentle, mellow, romantic
<i>nostalgia</i>	nostalgic, melancholic, sentimental	<i>tension</i>	tense, edgy, angry, fierce
<i>peacefulness</i>	peaceful, quiet, calm, gentle	<i>transcendence</i>	transcendent, fascinating, enchanting
<i>power</i>	powerful, strong, energetic, intense	<i>wonder</i>	wonderful, strange, fantastic
<i>sadness</i>	sad, sorrowful, unhappy, dismal, tearful		

4.3 Emotion-Oriented Tag-Based Profiles

The proposed representation of emotions lets transform tag-based item profiles (i.e., the items' annotation sets) into emotion-oriented profiles. In particular, we propose to perform such transformation in two stages. First, tag-based profiles are transformed into domain emotion-oriented profiles. Next, the obtained domain emotion-oriented profiles are transformed into core emotion-oriented profiles. Formally, let a core emotion $e_i^C \in \mathcal{E}$ and a domain-specific emotional category $e_i^D \in \mathcal{E}_D$ be defined as in formula (1). That is, they are vectors whose components represent lexicon terms and folksonomy tags that are synonyms and antonyms of the considered emotions. For an item (object) o_n , let $\mathbf{o}_n^T = (o_{n,1}, \dots, o_{n,|\mathcal{T}|}) \in \mathbb{R}^{|\mathcal{T}|}$ be the item's **tag-based profile**, where $o_{n,i}$ corresponds to the tag $t_i \in \mathcal{T}$ of the item's folksonomy. Then, from such profile, we define:

- the item's **domain emotion-oriented profile** as $\mathbf{p}_n^D = (p_{n,1}, \dots, p_{n,|\mathcal{E}_D|}) \in [-1,1]^{|\mathcal{E}_D|}$, where the i -th component corresponds to the domain emotional category $e_i^D \in \mathcal{E}_D$, and its weight is computed as $p_{n,i} = \cos(\mathbf{o}_n^T, \mathbf{e}_i^D)$, and
- the item's **core emotion-oriented profile** as $\mathbf{q}_n^C = (q_{n,1}, \dots, q_{n,|\mathcal{E}|}) \in [-1,1]^{|\mathcal{E}|}$, where the i -th component corresponds to the core emotion $e_i^C \in \mathcal{E}$, and its weight is computed as $q_{n,i} = \sum_{k=1}^{|\mathcal{E}_D|} p_{n,k} \cdot \cos(\mathbf{e}_i^C, \mathbf{e}_k^D)$.

Moreover, for each of these types of emotion-oriented profiles, we consider two alternatives for defining the (core and domain) emotion vectors \mathbf{e}_i^C and \mathbf{e}_i^D : **basic vectors**, whose components correspond to terms of the lexicon, as defined in formula (1), and **extended_N vectors**, whose components correspond to the N folksonomy tags that cooccur most frequently (in the tag-based item profiles) with the terms of the basic vectors. These tags are not necessarily synonyms/antonyms of the seed terms, and it is not clear whether they can be valuable to effectively assign emotions to items.

4.4 Crossing Emotion Folksonomies

In our model it is possible to relate core emotions and domain-specific emotional categories by computing the cosine similarity between their vector representations. Figure 3 shows the relation between some domain-specific emotional categories and the different core emotions for both the movie and music domains. It can be observed that, for instance, the emotional category *suspenseful* in the movie domain strongly overlaps with the *tense* and *nervous* core emotions, while the *peacefulness* category in the music domain intersects tightly with the *calm*, *relaxed* and *serene* core emotions.

Moreover, the intersection between cross domain-specific emotional categories could be computed to obtain a measure of similarity between them. Figure 2c shows the cosine similarity between pairs of cross-domain emotional categories. It can be seen that emotional categories such as *feel good-joy* and *gloomy-sadness*, which are very close in both valence and arousal, present very high similarity, while very distinct pairs of emotional categories, such as *joy-gloomy* and *sadness-uplifting*, present very low similarity. Other interesting pairs of very similar cross-domain emotional categories are *tension-suspenseful*, *power-exciting* and *nostalgia-emotional*.




Fig. 3. Relations between core emotions and domain-specific emotional categories

5 Experiments

5.1 User Study

To evaluate our methods that assign emotions to tagged items (see Section 4.3), we conducted a user study in which participants, recruited via social networking sites, were presented with sets of movies or musicians (no combinations of both), and were requested to freely select one or more core and domain-specific emotions for each item. A total of 72 users participated, evaluating 178 movies and 132 musicians. They generated 713 evaluation cases, assigning an average of 4.08 and 3.38 domain-specific emotional categories, and 3.30 and 4.18 core emotions, to items in the movie and music domains, respectively. To facilitate the evaluation, the users could select preferred movie and music genres and the language –English or Spanish– of the online evaluation tool (<http://ir.ii.uam.es/emotions>), and skip any item they did not want to evaluate.

We note that, as expressed by some of the participants, there are cases in which it is difficult to assign certain emotions to an item. Opposite emotions (e.g. *happiness* and *sadness*) can be evoked in different parts of a particular movie, and by different compositions of the same musician. This fact should be taken into account carefully in the future, and may have caused an underestimation of the precision of our methods; several participants decided not to assign certain emotions, which could have been retrieved by our methods, but were not considered as relevant.

Table 4 shows the top emotional categories assigned by the users to items belonging to some of the 26 genres considered from the Jinni and Last.fm systems, in the movie and music domains, respectively. Table 5 shows cooccurrence values of some core emotions in movie and music item profiles created by the users. These tables show coherent correspondences between domain emotions and genres (e.g. *exciting* for action movies, and *peacefulness* for ambient music), and between core emotions within the quadrants of Russell’s circumplex model (e.g. *happy* and *content*). It is interesting to note that there are emotions that barely relate with others (e.g. *bored* and *sad*).

Table 4. Top emotional categories implicitly assigned to some movie and music genres

Movie genre	Top emotional categories	Music genre	Top emotional categories
<i>action</i>	exciting, suspenseful, offbeat	<i>ambient</i>	peacefulness, nostalgia, transcendence
<i>comedy</i>	humorous, feel good, offbeat	<i>classical</i>	nostalgia, peacefulness, joy
<i>crime</i>	suspenseful, clever, bleak	<i>jazz</i>	nostalgia, peacefulness, power
<i>drama</i>	emotional, captivating, gloomy	<i>rock</i>	power, tension, joy
<i>horror</i>	scary, rough, exciting	<i>pop</i>	joy, power, tenderness
<i>war</i>	emotional, captivating, rough	<i>world</i>	wonder, transcendence, power

Table 5. Cooccurrence values of some core emotions in movie and music profiles

	movies								music							
	<i>excited</i>	<i>happy</i>	<i>content</i>	<i>serene</i>	<i>bored</i>	<i>sad</i>	<i>nervous</i>	<i>tense</i>	<i>excited</i>	<i>happy</i>	<i>content</i>	<i>serene</i>	<i>bored</i>	<i>sad</i>	<i>nervous</i>	<i>tense</i>
<i>excited</i>	55	11	9			3	12	25	38	15	10	2			5	8
<i>happy</i>	11	108	89	12			2	3	15	54	26	8			1	3
<i>content</i>	9	89	113	12			1	1	10	26	53	12				1
<i>serene</i>		12	12	17					2	8	12	37		2		
<i>bored</i>					7		1						1			
<i>sad</i>	3					11	3	6				2		7		
<i>nervous</i>	12	2	1	1	4	24	18		5	1					7	6
<i>tense</i>	25	3	1			3	18	44	8	3	1				6	10

5.2 Evaluating Emotion-Oriented Tag-Based Profiles

In the user study, participants stated which core and domain-specific emotions they consider as relevant for each item (movie or musician), thus manually (and collectively) creating emotion-oriented item profiles, which we consider as ground truth.

To evaluate the quality of the emotion-oriented profiles generated by our methods (Section 4.3) with respect to the ground truth profiles, we compared them with precision metrics. Specifically, we computed Precision at position k , $P@k$, which, for a particular item, is defined as the percentage of the top k emotions returned by a method that are relevant for the item, as stated by the users of our study. We also computed R -precision, which is defined as the precision of the top R emotions returned by a method for an item, being R the number of emotions that are relevant, as stated by the users of our study. Under a reasonable set of assumptions, R -precision approximates the area under the precision-recall curve [1]. Table 6 shows average precision values of the different methods (and a random emotion ranking method) on the movie and music domains.

The **basic method** was the best performing approach in both domains (with highest $P@1$ values around 70%), only outperformed by the `extended_10` method in the movie domain for the core emotion-oriented profiles. In general, the methods performed in the **music domain** better than in the movie domain, and were able to identify **domain emotional categories** more effectively than core emotions in both domains.

Table 6. P@k and R-precision values of the considered emotion-oriented profiles

Profile type	Emotion vector model	movies					music				
		#evals	P@1	P@2	P@3	R-Prec	#evals	P@1	P@2	P@3	R-Prec
<i>core emotion-oriented</i>	<i>random</i>	165	0.297	0.305	0.302	0.300	129	0.327	0.339	0.345	0.348
	<i>basic</i>	107	0.598	0.528	0.514	0.481	109	0.606	0.670	0.636	0.547
	<i>extended_10</i>	77	0.675	0.643	0.589	0.519	11	0.636	0.636	0.546	0.497
	<i>extended_50</i>	142	0.373	0.324	0.406	0.365	44	0.546	0.625	0.568	0.502
<i>domain emotion-oriented</i>	<i>random</i>	165	0.379	0.382	0.377	0.380	129	0.418	0.416	0.414	0.414
	<i>basic</i>	108	0.722	0.625	0.571	0.579	109	0.743	0.587	0.532	0.546
	<i>extended_10</i>	77	0.675	0.656	0.554	0.399	11	0.727	0.546	0.455	0.503
	<i>extended_50</i>	144	0.507	0.490	0.463	0.412	44	0.682	0.443	0.394	0.428

5.3 Evaluating Emotion-Oriented Recommendations

In the user study, participants initially stated which movie and music genres they were interested in, and, in addition to emotions, they assigned to movies and musicians numeric ratings in the range [1, 10] according to their tastes.

In a second part of our experiment, we evaluated **whether emotional information of items can be used to increase the accuracy of recommendation based on the users' past ratings**. For such purpose, and due to the limited number of ratings in the study, we addressed the recommendation problem as a (binary) classification task instead of as a rating prediction task, in which collaborative filtering strategies could be applied. For each user, we considered as *relevant* those items to which she assigned a rating over her average rating value, and as *non-relevant* to the reminder items she rated. We then built patterns datasets in which each pattern was associated to an evaluation case [*user* *u*, *item* *i*, *rating* *r*, *core emotions* $\{e_1^C, \dots, e_{|\mathcal{E}|}^C\}$, *domain-specific emotions* $\{e_1^D, \dots, e_{|\mathcal{E}_D|}^D\}$]. A pattern's class was 0 or 1, *relevant* or *non-relevant*, based on *r* and what was explained above. Its attributes were binary values associated to *u*'s genres of interest, binary values associated to *i*'s core emotions, and binary values associated to *i*'s domain-specific emotions. To assess the impact of emotions in recommendation, we separately evaluated classification cases with/without the consideration of the emotion attributes. We used such cases (patterns) to build and evaluate several well known classifiers, namely Naïve Bayes, Random forest, Multilayer Perceptron (MLP), and Support Vector Machine (SVM), which we used as model-based recommender systems.

Table 7 shows the best average (10-fold cross validation) performance values of the classifiers for the distinct pattern attribute configurations and domains. In addition to accuracy values, we also report $g = \sqrt{acc^+ \cdot acc^-}$ and AUC values to take the class balance levels into account. Classifiers incorporating emotion attributes outperformed those built with only the users' genres of interest. For movies, **core emotions** were more valuable, whereas for musicians, **domain-specific emotions** were better. **Random forest** (for movies) and **SVM** (for music) were the best performing classifiers.

Table 7. Performance values obtained by the model-based recommender systems built with the different profile types (attribute configurations). Global top values are in bold, and best values for each profile type are underlined.

Profile type	Classifier	movies					music				
		acc	acc+	acc-	g	AUC	acc	acc+	acc-	g	AUC
-	Majority class	56.009	100.000	0.000	0.000	0.402	57.273	100.000	0.000	0.000	0.417
emotion-unaware	Naïve Bayes	53.648	77.011	23.902	42.904	0.502	49.545	53.175	44.681	48.743	0.497
	Random forest	<u>59.442</u>	<u>56.322</u>	<u>63.415</u>	<u>59.763</u>	<u>0.596</u>	<u>55.909</u>	<u>65.079</u>	<u>43.617</u>	<u>53.278</u>	<u>0.556</u>
	MLP	59.227	64.751	52.195	58.135	0.592	50.909	53.968	46.809	50.261	0.511
	SVM	57.296	61.303	52.195	56.566	0.573	50.455	56.349	42.553	48.968	0.505
core emotion-aware	Naïve Bayes	58.798	73.946	39.512	54.054	0.575	52.727	61.111	41.489	50.353	0.525
	Random forest	61.588	55.556	69.268	62.034	0.616	<u>54.091</u>	61.111	44.681	<u>52.254</u>	<u>0.540</u>
	MLP	62.876	69.732	54.146	61.447	0.627	48.636	55.556	39.362	46.763	0.486
	SVM	59.871	63.602	55.122	59.210	0.599	50.000	50.794	48.936	49.856	0.503
domain emotion-aware	Naïve Bayes	57.940	80.077	29.756	48.814	0.550	52.727	62.698	39.362	49.678	0.523
	Random forest	<u>60.515</u>	55.939	66.341	<u>60.918</u>	<u>0.606</u>	58.182	65.079	48.936	56.433	0.581
	MLP	59.657	66.667	50.732	58.156	0.595	55.455	60.317	48.936	54.330	0.555
	SVM	57.511	63.218	50.244	56.359	0.574	<u>59.091</u>	62.698	54.255	58.324	0.592

6 Future Work

The next step in our research is to exploit the generated emotion-oriented profiles for developing mood-based and cross-domain recommendation strategies. We are interested in determining which items (according to the emotions they evoke) should be suggested to a user based on her current mood, and which items in a (target) domain should be suggested to a user whose preferences in a distinct (source) domain are available. We believe our emotion model and its cross-domain folksonomies could help address such problems independently or in combination with existing approaches [2][17].

Acknowledgements. This work was supported by the Spanish Government (TIN2011-28538-C02) and the Regional Government of Madrid (S2009TIC-1542).

References

1. Aslam, J.A., Yilmaz, E.: A Geometric Interpretation and Analysis of R-precision. In: 14th ACM Conference on Information and Knowledge Management, pp. 664–671 (2005)
2. Baldoni, M., Baroglio, C., Patti, V., Rena, P.: From Tags to Emotions: Ontology-driven Sentiment Analysis in the Social Semantic Web. *Intell. Artificiale* 6(1), 41–54 (2012)
3. Breazeal, C.: Emotion and Sociable Humanoid Robots. *International Journal of Human-Computer Studies* 59, 119–155 (2003)
4. Cantador, I., Bellogín, A., Fernández-Tobías, I., López-Hernández, S.: Semantic Contextualisation of Social Tag-Based Profiles and Item Recommendations. In: Huemer, C., Setzer, T. (eds.) *EC-Web 2011*. LNBP, vol. 85, pp. 101–113. Springer, Heidelberg (2011)

5. Cantador, I., Brusilovsky, P., Kuflik, T.: Second Workshop on Information Heterogeneity and Fusion in Recommender Systems. In: 5th ACM Conference on Recommender Systems, pp. 387–388 (2011)
6. Carrillo de Albornoz, J., Plaza, L., Gervás, P.: A Hybrid Approach to Emotional Sentence Polarity and Intensity Classification. In: 14th Intl. Conference on Computational Natural Language Learning, pp. 153–161 (2010)
7. Cassell, J.: Nudge Nudge Wink Wink: Elements of Face-to-Face Conversation for Embodied Conversational Agents. In: Embodied Conversational Agents, pp. 1–27 (2003)
8. Ekman, P., Davidson, R.J. (eds.): The Nature of Emotions: Fundamental Questions. Oxford University Press (1994)
9. Feng, Y., Zhuang, Y., Pan, Y.: Popular Music Retrieval by Detecting Mood. In: 6th ACM SIGIR Conf. on Research and Development in Information Retrieval, pp. 375–376 (2003)
10. Hastings, J., Ceusters, W., Smith, B., Mulligan, K.: The Emotion Ontology: Enabling Interdisciplinary Research in the Affective Sciences. In: Beigl, M., Christiansen, H., Roth-Berghofer, T.R., Kofod-Petersen, A., Coventry, K.R., Schmidtke, H.R. (eds.) CONTEXT 2011. LNCS, vol. 6967, pp. 119–123. Springer, Heidelberg (2011)
11. Hume, D.: Emotions and Moods. In: Robbins, S.P., Judge, T.A. (eds.) Organizational Behavior, pp. 258–297
12. Kaminskis, M., Ricci, F.: Location-Adapted Music Recommendation Using Tags. In: Konstan, J.A., Conejo, R., Marzo, J.L., Oliver, N. (eds.) UMAP 2011. LNCS, vol. 6787, pp. 183–194. Springer, Heidelberg (2011)
13. Meyers, O.C.: A Mood-based Music Classification and Exploration System. MSc Thesis. School of Architecture and Planning, MIT (2007)
14. Roberts, K., Roach, M.A., Johnson, J., Guthrie, J., Harabagiu, S.M.: EmpaTweet: Annotating and Detecting Emotions on Twitter. In: 8th International Conference on Language Resources and Evaluation, pp. 3806–3813 (2012)
15. Russell, J.A.: A Circumplex Model of Affect. *Journal of Personality and Social Psychology* 39(6), 1161–1178 (1980)
16. Shi, Y., Larson, M., Hanjalic, A.: Mining Mood-specific Movie Similarity with Matrix Factorization for Context-aware Recommendation. In: Workshop on Context-Aware Movie Recommendation, pp. 34–40 (2010)
17. Shi, Y., Larson, M., Hanjalic, A.: Tags as Bridges between Domains: Improving Recommendation with Tag-Induced Cross-Domain Collaborative Filtering. In: Konstan, J.A., Conejo, R., Marzo, J.L., Oliver, N. (eds.) UMAP 2011. LNCS, vol. 6787, pp. 305–316. Springer, Heidelberg (2011)
18. Winoto, P., Ya Tang, T.: The Role of User Mood in Movie Recommendations. *Expert Systems with Applications* 37(8), 6086–6092 (2010)
19. Zentner, M., Grandjean, D., Scherer, K.: Emotions Evoked by the Sound of Music: Characterization, Classification, and Measurement. *Emotion* 8, 494–521 (2008)