Procesamiento del Lenguaje Natural, Revista nº 51, septiembre de 2013, pp 195-202 recibido 31-03-2013 revisado 16-06-2013 aceptado 21-06-2013

A social tag-based dimensional model of emotions: Building cross-domain folksonomies

Un modelo dimensional de emociones basado en etiquetas sociales: Construcción de folcsonomías en dominios cruzados

Ignacio Fernández-Tobías, Iván Cantador
Universidad Autónoma de Madrid
C/ Francisco Tomás y Valiente, 11
28049 Madrid, Spain
{i.fernandez,ivan.cantador}@uam.es

Laura Plaza
UNED NLP & IR Group
C/ Juan del Rosal, 16
28040 Madrid, Spain
lplaza@lsi.uned.es

Resumen: En este trabajo se presenta un modelo dimensional de emociones basado en etiquetas sociales. El modelo se construye sobre un léxico generado automáticamente que caracteriza emociones por medio de términos sinónimos y antónimos. Este léxico se enlaza con diversas folcsonomías emocionales específicas de dominio. Se propone una serie de métodos para transformar perfiles de objetos basados en etiquetas sociales en perfiles emocionales. El objetivo de estos perfiles es su uso por parte de sistemas adaptativos y de personalización que permitan recuperar o recomendar contenidos en función del estado de ánimo del usuario. Para validar el modelo, se muestra que la representación de un conjunto de emociones básicas se corresponde con la del aceptado modelo de Russell. También se reportan resultados de un estudio de usuario que demuestran una alta precisión de los métodos propuestos para inferir emociones evocadas por objetos en los dominios del cine y la música.
Palabras clave: emociones, léxico afectivo, etiquetado social, folcsonomías

Abstract: 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, which may be exploited by adaptation and personalization systems. To validate our model, we show that its representation of a number of core emotions is in accordance with the well known psychological circumplex model of affect. We also 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.
Keywords: emotions, affective lexicon, social tagging, folksonomies

1. *Introduction and background*

The study and development of computational systems aimed to recognize and interpret human feelings is usually referred to as *Affective Computing* (Picard, 1995). In Natural Language Processing, this discipline - often known as *Sentiment Analysis* - is becoming increasingly important with the development of the Social Web and the growing popularity of online forums, social networks, and collaborative tagging systems (Carrillo-De-Albornoz, Plaza, and Gervás, 2010; De Choudhury and Gamon, 2012).

Focusing on User Modeling and Recommender Systems, emotions (and moods) can be efficiently used in a wide range of applications; e.g. constructing user behaviour models (Hastings et al., 2011), tailoring the search results (Meyers, 2007) and filtering the recommending items (Winoto and Ya Tang, 2010). The user's mood has proven to have an important influence on the choice of items that the user is more likely to consume, and therefore the system should be able to suggest items according to that mood.

In this context, modeling and exploiting emotions present challenging problems. First, there is not agreement on the categorization of emotions to be used. Focusing on computational models of emotion, three main psy-

ISSN 1135-5948

chological theories have been adopted, namely the *categorical emotion theory* (James, 1984) - which characterizes emotions as discrete units with boundaries -, the *emotional dimension theory* (Russell, 1980) - which conceptualizes emotions as points in a continuous space -, and the *appraisal theory* (Scherer, Shorr, and Johnstone, 2001) - which represents emotions as outcomes of certain events and situations. Second, detecting emotions in text is extremely difficult. Most approaches to the problem use emotion lexicons that provide specific vocabularies for describing emotions. SentiWordNet (Baccianella, Esuli, and Sebastiani, 2010), for instance, associates to each WordNet synset three numerical scores *obj*, *pos* and *neg*, describing how objective, positive, and negative the terms in the synset are. More fine-grained is SentiSense (Carrillo-De-Albornoz, Plaza, and Gervás, 2012), which attaches specific emotions (e.g. *sad* or *like*) to WordNet concepts.

Existing emotion lexicons, usually developed for polarity classification of texts, present several drawbacks that limit their use in personalization and recommendation systems. First, they use a single set of generic emotions for categorizing all terms in the lexicon, but, as we will see, the emotions that are evoked by items usually are domain-specific - such as *scare* in the movie domain, and *peacefulness* in the music domain. Second, the vocabulary employed by users to characterize the different emotions varies with the application domain, and thus it is necessary to develop domain-specific emotional lexicons. Third, in order to perform cross-domain recommendation (i.e., to suggest items in a target domain using user feedback about items in a different source domain), an automatic method for translating emotional information from one domain to another is required.

In this paper, we propose an automatic approach that generates a core lexicon and different folksonomies to represent both generic and domain-dependent emotion categories. These resources are generated from a generic thesaurus and social tagging systems in entertainment domains, namely the movie and music domains. More specifically, we propose a model in which emotions are represented as vectors of weighted synonym and antonym terms, and which enables computing (dis)similarities between emotions. In this way, it is possible to relate core emotions and domain-specific emotions, and thus to extrapolate emotional information from one domain to another. We think our model could be exploited by adaptive and personalized systems that are based on a keyword- or concept-based knowledge representations.

2. A core domain-independent emotion lexicon

Our model adopts the emotional dimension theory, and is based on the Russell's circumplex model of affect (Russell, 1980). This model understands emotions as a linear combination of two dimensions, *pleasure* and *arousal*, as shown in Figure 1. Arousal (in the vertical axis) reflects the intensity of an emotion; and pleasure (in the horizontal axis) reflects whether an emotion is positive or negative. With this representation, any emotion can be represented at any level of arousal and pleasure. Hence, for instance, *happiness* and *sadness* can be considered as emotions with the highest and lowest levels of pleasure, respectively, but with neutral levels of arousal, with respect to other emotions such as *tension* (with high arousal) and *calmness* (with low arousal). Russell proposed a set of 16 core (basic) emotions (see Figure 1). We will use this set of emotions in our model.

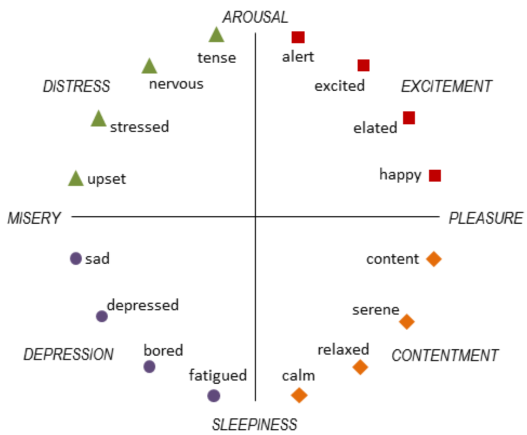

Figure 1: Distribution of core emotions in the circumplex model

Our dimensional model is built upon an automatically generated lexicon $L = \{t_1, ..., t_K\}$ composed of synonym and antonym terms t_k of the core emotions' names (e.g. *happy*, *sad*). The synonym and antonym terms of each emotion are obtained from the online thesaurus provided by Dictionary.com.[1] Table 1 shows some of the obtained synonyms for each emotion.

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

Table 1: Core emotions and their synonyms

Once the lexicon L is generated, a core emotion $e_i \in E$ is represented as a vector $e_i = (e_{i,1}, ..., e_{i,k}) \in R^K$, in which the component $e_{i,k}$ corresponds to the term $t_k \in L$ and is computed as shown in eq. (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 (Baeza-Yates and Ribeiro-Neto, 2011) of t_k computed by considering the lexicon L as the collection vocabulary, and the set E of emotions (described as sets of synonym and antonym terms) as the collection documents.

$$e_{i,k} = \begin{cases} tf \times idf(e_{i,k})\, si\, t_k \in synonyms(e_i) \\ -tf \times idf(e_{i,k})\, si\, t_k \in antonyms(e_i) \\ 0 \text{ otherwise} \end{cases} \quad (1)$$

With the proposed vector representation of emotions, we can measure (dis)similarities between emotions. Specifically, we can use the well known cosine similarity (Baeza-Yates and Ribeiro-Neto, 2011). Figure 2 shows the cosine similarity between each pair of core emotions. The cell colors indicate the sign of the similarity values - being black for positive values and white for negative ones -, and the cell intensities correspond to the similarity absolute values - being dark for values close to 1, and light for values close to 0. The emotions are sorted according to the quadrants of Russell's model (Figure 1). We can observe that emotions in the same quadrant have high similarities (e.g. *alert* and *excited*), while emotions in opposite quadrants have low similarities (e.g. *calm* and *excited*). These results show that our tag-based model is in accordance with the circumplex model.

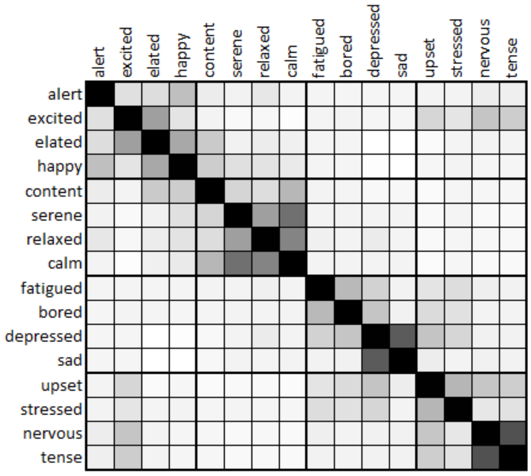

Figure 2: Similarity between core emotions

To better show the correspondences between our computational model and the theoretical circumplex model, Figure 3 shows the projections of our emotion vectors into a two-dimension space by applying *Principal Component Analysis*. We observe that our model locates all 16 core emotions in their corresponding quadrants.

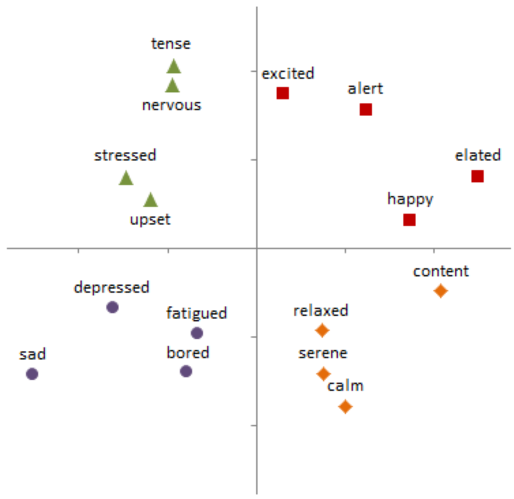

Figure 3: Distribution of core emotions in our tag-based model

More interestingly, we observe that in our model the axes defined by the two most informative principal components are related to the arousal and pleasure factors of the circumplex model. Hence, positive emotions

[1] Dictionary.com thesaurus, http://thesaurus.com

(e.g. *happy*, *content*) are in right quadrants, while negative emotions (e.g. *sad*, *upset*) are in left quadrants; and more intense emotions (e.g. *tense*, *excited*) are in the upper quadrants, while less intense emotions (e.g. *relaxed*, *fatigated*) are in the lower quadrants.

3. Cross-domain emotion folksonomies

In a social tagging system users create items and annotate them with freely chosen tags. 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 (Cantador, Brusilovsky, and Kuflik, 2011). Within the set of tags that express qualities and opinions about the items, there are tags that refer to emotions caused by the items. In most cases, however, such emotions are not the core emotions presented in the previous section, 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 (Cantador, Konstas, and Jose, 2011). With the extended model we propose to build emotion-oriented item profiles and cross-domain folksonomies. This process is illustrated in Figure 4. We make all the data publicly available.[2]

3.1. An emotion folksonomy for the movie domain

To build the emotion folksonomy for movies, we first select a total of 15 emotional categories listed under the mood topic in Jinni [3] 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 from Thesaurus.com, but restricted to those existing as social tags in the MovieLens dataset. Finally, we repeat the process in Section 2 to represent an emotional category as a vector of weighted terms. Table 2 shows the number of terms per category that we collected.

Category	Seed terms	#
clever	clever, cerebral, reflective	71
offbeat	offbeat, quirky, surreal	83
exciting	exciting, energetic, frantic	104
suspenseful	suspenseful, tense	34
captivating	captivating, rousing, poignant	83
emotional	emotional, passionate, romantic	185
feel good	cute, merry, happy	41
humorous	humorous, funny, comical	101
sexy	sexy, erotic, sensual	39
sexual	sexual, lascive, horny	16
uplifting	uplifting, inspirational, hope	32
bleak	bleak, grim, depressing	84
gloomy	gloomy, sad, melancholic	85
rough	rough, brutal, lurid, macabre	126
scary	scary, creepy, menacing	57

Table 2: Movie emotional categories, seed terms and number of terms per category

Figure 5 shows the cosine similarity between each pair of emotional categories. It can be observed 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.

3.2. An emotion folksonomy for the music domain

To generate an emotion folksonomy in the music domain, we select as emotional categories the 9 emotions proposed in the GEMS (Geneva Emotional Music Scales) model (see Table 3). As initial seed terms we use the category names and their associated feeling terms given in (Zentner, Grandjean, and Scherer, 2008). 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, along with the total number of tags we collected for each category.

Figure 5 shows the similarity 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.

[2]http://ir.ii.uam.es/emotions/
[3]http://www.jinni.com

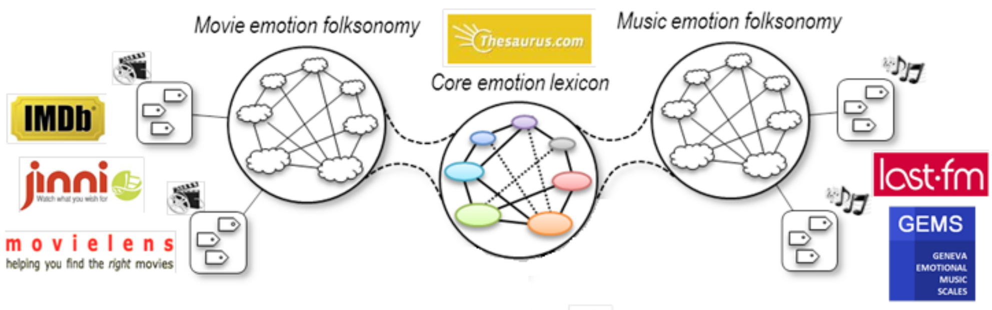

Figure 4: Crossing domain-dependent folksonomies

Category	Seed terms	#
joy	funny, happy, amusing, jolly	84
nostalgia	nostalgic, melancholic, sentimental	49
peacefulness	peaceful, quiet, calm, gentle	71
power	powerful, strong, energetic, intense	97
sadness	sad, sorrowful, unhappy, dismal	51
tenderness	tender, gentle, mellow, romantic	41
tension	tense, edgy, angry, fierce	58
transcendence	fascinating, enchanting	45
wonder	wonderful, strange, fantastic	24

Table 3: Music emotional categories, seed terms and number of terms per category

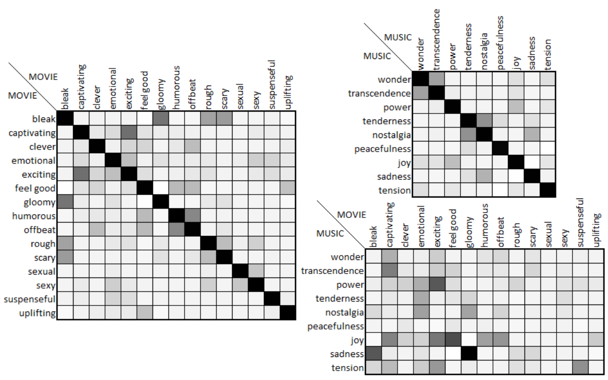

Figure 5: Similarity values between movie and music emotional categories

3.3. 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. To this end, we first transform the tag-based profiles into domain emotion-oriented profiles. Next, the domain emotion-oriented profiles are transformed into core emotion-oriented profiles. Formally, let a core emotion $e_i^C \in E$ and a domain-specific emotional category $e_j^D \in E_D$ be defined as in eq. (1). For an item (o_n), let $o_n^T = (o_{n,1}, ..., o_{n,|\tau|}) \in \Re^{|\tau|}$ be the item's tag-based profile. Then, from such profile, we define:

- the item's **domain emotion-oriented profile** as $p_n^D = (p_{n,1}, ..., p_{n,|E_D|}) \in [-1, 1]$, where each component represents a domain emotion, and its weight is computed as $p_{n,i} = cos(o_n^\tau, e_i^D)$, and
- the item's **core emotion-oriented profile** as $q_n^C = (q_{n,1}, ..., q_{n,|E|}) \in [-1, 1]$, where each component corresponds a core emotion, and its weight is computed as $q_{n,i} = \sum_{k=1}^{|E_D|} p_{n,k} \times cos(e_i^C, e_k^D)$.

Moreover, for each of these types of emotion-oriented profiles, we consider two alternatives for defining the emotion vectors: **basic vectors**, whose components correspond to terms of the lexicon, 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.

3.4. Crossing folksonomies in different domains

The proposed model let us to relate core emotions and domain-specific emotional categories by computing the cosine similarity between their vector representations. Figure 6 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 movies 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

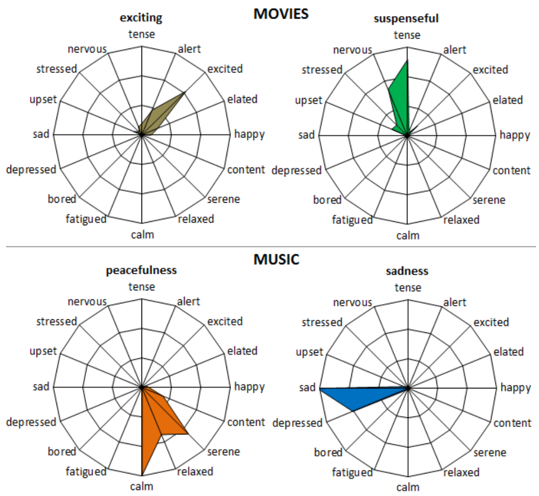

Figure 6: Relation between core emotions and domain-specific emotional categories

be computed to obtain a measure of similarity between them. Figure 5 shows the cosine similarity between pairs of cross-domain emotional categories. It can be seen that pairs of categories such as *feel good-joy* and *gloomy-sadness*, which are very close in pleasure and arousal, present very high similarity, while very distinct categories, such *joy-gloomy* and *sadness-uplifting*, present very low similarity.

4. Experiments and results

To evaluate our emotional model, 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 domain-specific and core emotions for each item. A total of 71 users participated, evaluating 165 movies and 129 musicians. They generated 703 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 movies 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[4] (see Figure 7), 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.

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 3.3) with respect to the ground truth profiles, we compared them by means of IR precision metrics. Specifically, we computed Precision at position k, $P@k$ (Baeza-Yates and Ribeiro-Neto, 2011), 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 $R - precision$ (Buckley and Voorhees, 2005), 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 for the item, as stated by the users of our study. That is, $R - precision$ is $P@R$, i.e. it is the break-even point in the precision-recall curve where precision is equal to recall.

Table 4 shows average precision for the different methods (and a random method) on the movie and music domains. The basic method was the best performing one 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 5 shows the top two emotional categories assigned by the users to items belonging to some of the 26 genres considered from the Jinni and Last.fm systems. This table also shows the two predominant emotional categories for each genre, according to our emotion-tag based profiles. It can be seen that the emotions assigned by our model are very similar to that assigned by the users.

Finally, and concerning the frequency with which the different emotions are associated with the movies and musicians in our experiment, Figure 8 shows the percentage of items that have been assigned a given emotion (both core and domain-specific). As it

[4]Evaluation tool, url omitted to ensure anonymity

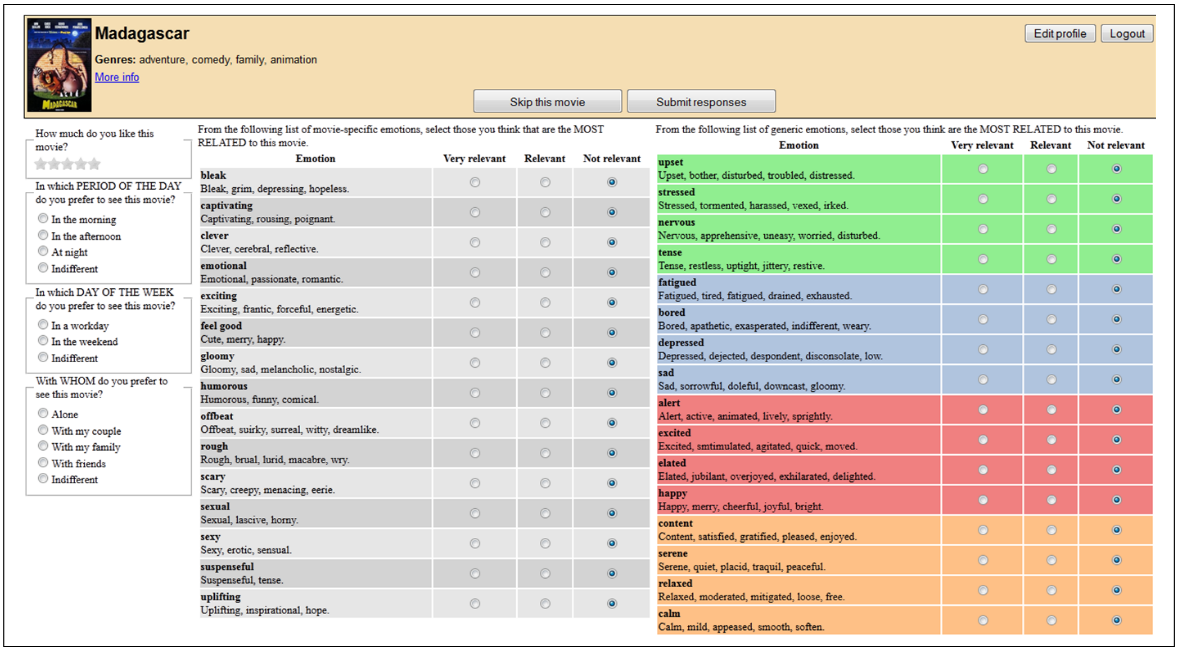

Figure 7: User study - online evaluation tool

Profile type	Vector model	Movie domain #evals	$P@1$	P@2	P@3	R-Pr	Music domain #evals	$P@1$	$P@2$	$P@3$	R-Pr
core emotion-oriented	random	165	0.297	0.305	0.302	0.300	129	0.327	0.339	0.345	0.348
	basic	107	0.598	0.528	0.514	0.481	109	0.606	**0.670**	**0.636**	**0.547**
	extended_10	77	**0.675**	**0.643**	**0.589**	**0.519**	11	**0.636**	0.636	0.546	0.497
	extended_50	142	0.373	0.324	0.406	0.365	44	0.546	0.625	0.568	0.502
	extended_100	155	0.419	0.390	0.411	0.399	79	0.557	0.620	0.582	0.546
domain emotion-oriented	random	165	0.379	0.382	0.377	0.380	129	0.418	0.416	0.414	0.414
	basic	108	**0.722**	0.625	**0.571**	**0.579**	109	**0.743**	**0.587**	**0.532**	**0.546**
	extended_10	77	0.675	**0.656**	0.554	0.399	11	0.727	0.546	0.455	0.503
	extended_50	144	0.507	0.490	0.463	0.412	44	0.682	0.443	0.394	0.428
	extended_100	158	0.551	0.532	0.513	0.449	79	0.696	0.494	0.426	0.463

Table 4: Avg. $P@k$ and $R - precision$ values of the considered emotion-oriented profiles

Movie domain		
action	Users	exciting, suspenseful
	Model	suspenseful, captivating
comedy	Users	humorous, feel good
	Model	humorous, feel good
horror	Users	scary, rough
	Model	scary, exciting
Music domain		
classical	Users	nostalgia, peacefulness
	Model	nostalgia, peacefulness
rock	Users	power, tension
	Model	power, joy
jazz	Users	nostalgia, peacefulness
	Model	tension, peacefulness

Table 5: Top emotional categories assigned to some movie and music genres by (a) the users (b) inferred using the proposed model

can be observed, the most frequent core emotion is *content*, followed by *happy*. In the music domain, the predominant emotional category is *power*, while the most frequent one in the movie domain is *humorous*.

5. *Conclusions and future work*

We have presented a computational model that represents emotions as vectors of weighted synonym and antonym terms, which are automatically obtained from an online thesaurus and social tagging systems in different entertainment domains. Our model distinguishes and relates generic core emotions (e.g. *happiness*, *sadness*) with domain-specific emotional categories (e.g. *suspense* in the movie domain, and *nostalgia* in the music domain). This lets transform tag-based profiles into emotion-oriented profiles, and build cross-domain emotion folksonomies.

The next step in our research is to exploit the generated emotion-oriented profiles in adaptation and personalization systems. In

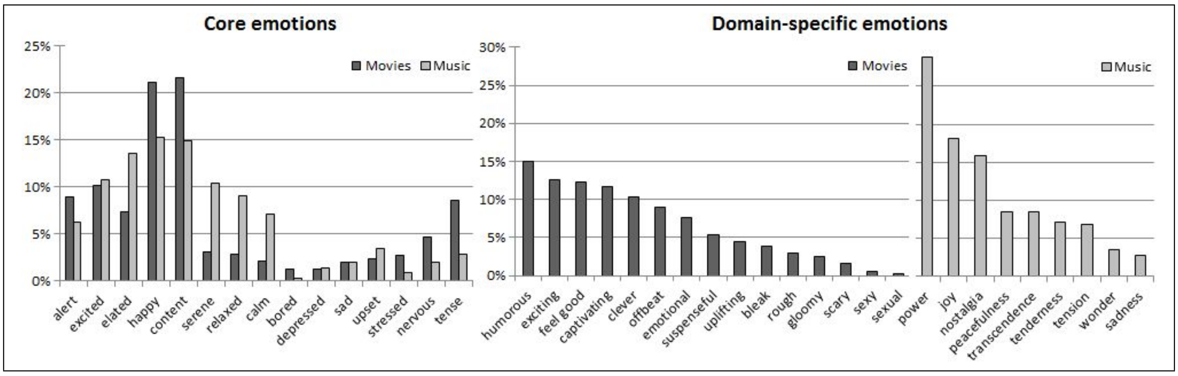

Figure 8: Distribution of core and domain-dependent emotions in the evaluation collection

particular, we plan to use them for developing mood-based and cross-domain recommendation strategies.

References

Baccianella, S., A. Esuli, and F. Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the 7th Conference on Language Resources and Evaluation (LREC'10).*

Baeza-Yates, R. A. and B.A. Ribeiro-Neto. 2011. *Modern Information Retrieval - The Concepts and Technology behind Search.* Pearson Education.

Buckley, C. and E. Voorhees, 2005. *TREC: Experiment and Evaluation in Information Retrieval*, chapter Retrieval System Evaluation. MIT Press.

Cantador, I., P. Brusilovsky, and T. Kuflik. 2011. Second workshop on information heterogeneity and fusion in recommender systems. In *Proceedings of the 5th ACM Conference on Recommender Systems (RecSys'11)*, pages 387–388.

Cantador, I., I. Konstas, and J. Jose. 2011. Categorising social tags to improve folksonomy-based recommendations. *Journal of Web Semantics*, 9(1):1–15.

Carrillo-De-Albornoz, J., L. Plaza, and P. Gervás. 2010. A hybrid approach to emotional sentence polarity and intensity classification. pages 153–161.

Carrillo-De-Albornoz, J., L. Plaza, and P. Gervás. 2012. Sentisense: An easily scalable concept-based affective lexicon for sentiment analysis. In *Proceedings of the 8th Conference on Language Resources and Evaluation (LREC'12).*

De Choudhury, M., Counts S. and M. Gamon. 2012. Not all moods are created equal! exploring human emotional states in social media. In *Proceedings of the 6th International AAAI Conference on Weblogs and Social Media.*

Hastings, J., W. Ceusters, B. Smith, and K. Mulligan. 2011. The emotion ontology: Enabling interdisciplinary research in the affective sciences. In *Proceedings of the 7th International and Interdisciplinary Conference on Modeling and Using Context (Context'11).*

James, W. 1984. What is emotion? *Mind*, 9:188–205.

Meyers, O. C. 2007. A mood-based music classification and exploration system. Master's thesis, School of Architecture and Planning, MIT.

Picard, R. W. 1995. Affective computing. Technical Report 321, MIT Media Laboratory, Perceptual Computing Section.

Russell, J.A. 1980. A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6):1161–1178.

Scherer, K. R., A. Shorr, and T. (Eds.) Johnstone. 2001. *Appraisal Processes in Emotion: Theory, Methods, Research.*

Winoto, P. and T. Ya Tang. 2010. The role of user mood in movie recommendations. *Expert Systems with Applications*, 37(8):6086–6092.

Zentner, M., D. Grandjean, and K. Scherer. 2008. Emotions evoked by the sound of music: Characterization, classification, and measurement. *Emotion*, 8:494–521.