

Music Emotion Identification from Lyrics

Dan Yang

Systems Science
University of Ottawa
Ottawa Canada K1N 6N5
dyang006@uottawa.ca

Won-Sook Lee

School of Information Technology and Engineering
University of Ottawa
Ottawa Canada K1N 6N5
wslee@uottawa.ca

ABSTRACT—Very large online music databases have recently been created by vendors, but they generally lack content-based retrieval methods. One exception is Allmusic.com which offers browsing by musical emotion, using human experts to classify several thousand songs into 183 moods. In this paper, machine learning techniques are used instead of human experts to extract emotions in Music. The classification is based on a psychological model of emotion that is extended to 23 specific emotion categories. Our results for mining the lyrical text of songs for specific emotion are promising, generate classification models that are human-comprehensible, and generate results that correspond to commonsense intuitions about specific emotions. Mining lyrics focused in this paper is one aspect of research which combines different classifiers of musical emotion such as acoustics and lyrical text.

Keywords: *Text Mining, Text Classification, Music Information Retrieval, Lyrical Text, Emotion.*

I. INTRODUCTION

Information retrieval (IR) capabilities have greatly improved in the past decade [29], particularly in extracting the denotative value of text. More needs to be done in identifying emotional value, style, presence and other qualities of text and multimedia content. Current music IR systems are generally unable to query on emotions, although media psychologists [35][30] have shown that emotionality is crucial to the entertainment experience of media users. Hansen and Hansen [7] describe the tremendous appeal of contemporary popular music in terms of the mood states induced in listeners, as well as the psychosocial benefits of being able to develop individual and group identities. Human expertise based on fine arts or psychology takes a long time to achieve, about one PhD per artist, and faster solutions need to be found. Advances in machine learning offer a faster, automated solution, given there is exists a large quantity of online music.

Our previous paper [34] gave a framework of our music emotion extraction system “EMO” in University of Ottawa for music data mining, ranging over several formats including acoustic waveform, song lyrics and artist reviews. In the function view of emotions, distinct positive and negative emotions can be separated on the basis of their different function (reward-approach and threat-avoidance) or dysfunction (psychotic). The framework is also based on a compositional theory of musical meaning [25], in which lyrics and other contextual information are understood to focus the listener’s attention on particular emotions. In this

paper we focus on mining lyrics as one part of research which combines different classifiers of musical emotion in “EMO”.

II. ANALYSIS OF EXISTING SYSTEMS

Text mining has largely focused on the classification of topics, independent of qualities such as emotion, style and opinion about the topic. Emotion identification may require a better understanding of the text as a whole, its structural workings and the context in which it is used. The generation of comprehensible results is particularly challenging, as common-sense understanding of emotion involves many distinct specific emotion categories, while machine learning models are often not intended to be comprehensible. Historically, psychological content analysis models have been significant. More recently, statistical natural language processing techniques have been created and Ogiwara’s work [15] is specifically related to emotion in popular music lyrics. Many of these techniques suffer the disadvantage of being ‘black box’ models and not intuitively comprehensible, which hinders the validation of these models with respect to commonsense intuitions about specific emotions.

A. Psychological feature identification in text

Lyrical text is distinct from ordinary text in the use of stylistic qualities such as rhyme, poetic form, and figurative language. Song lyrics help to focus the listener’s attention on specific emotions. Since the 1930s, psychologists have interpreted the affective value of words, based upon empirical surveys and expert judgments. Measurement scales were created to quantify the verbal reports of psychological state according to how many and which dimensions (e.g. intensity, valence, and dominance). A variety of ratings scales for affective words were developed, and documents were rated by summing the ratings of individual words. Whissell’s Dictionary of Affect in Language [32] has values for valence and intensity (both scaled from 1 to 3) for over 8,700 English words. The system has been used for comparing texts such as speeches and narratives. Claritech Corp’s Clairvoyance [26] has an extensive coverage of English words with ratings for specific affect (83 categories), degree of relatedness to the category, and degree of emotional intensity. The General Inquirer system [23] associates over 10,000 English words with one or more of 182 general psychological categories.

B. Linguistics research

1) Linguistics research on psychological issues

Linguistics disciplines provide expertise about the way the structured, rule-like nature of language provides clues that help to distinguish emotional and non-emotional sentences [11]. Mathieu [12] modeled French words in terms of about 100 psychological features where similar words used for the same emotion are understood in terms of how well these words fit constraints on word use in the same position in a sentence. Gordon et al. [6] encoded common psychological patterns of English as grammars. The patterns included planning, emotion, explanation, expectation management, reasoning, similarity, memory, etc. where the system is used for emotional detection of soldier speech in realistic, stressful training situations. Liddy [9] studied the way factual denotations in text are often accompanied by subjective colorations that are linguistically detectable, such as degree of certainty about a statement of fact. Certainty markers were studied in newspaper text. Another aspect of emotion identification studied was the effect of valence-reversing words such as 'however' used in the context of emotional polarity detection

2) Linguistics research about song lyrics

Scott and Matwin [21] applied text retrieval techniques to folk lyrics using Wordnet nouns and verbs [4]. Synonyms and hypernyms were used to transform words into a smaller number of features, which varies as few as 2, by generalizing. Their system produced discrimination rules for binary classification tasks, such as whether the folk song was about the topics 'politics' or 'religion', with about 70% accuracy. This work illustrated several problems with Wordnet, in that the concept hierarchy is very deep, and the appropriate level of generalization is difficult to select. Also, each part of speech has its own database, so identical concepts such as normalized-form verbs are duplicated in noun and verb databases, with different hierarchies.

Another approach is FrameNet [5] which lists about 600 schematic situations called frames. Within a frame, there are elements corresponding to participants, props, and other conceptual roles specific to the situation. For psychological subjects, these elements would correspond to typical emotions experienced by the subject. Taking the 'Revenge' frame as an example, an 'Offender' subject who causes 'Injury' may feel guilt, while the 'Injured_Party' may feel anger, and an 'Avenger' may feel the need to act to 'Punish'. FrameNet currently has coverage of over 4,000 English words, so it is a potentially useful method of organizing a large number of results for comprehensibility, by transforming texts into conceptual categories.

NLP systems have been developed that treat sentiment identification as a 2 class classification problem. Turney and Littman [28] trained a very large corpus to extract positive/negative polarity based on proximity of a topic to 7

positive words (nice, excellent, positive, fortunate, correct, superior) and 7 negative words (bad, nasty, poor, negative, unfortunate, wrong, and inferior). Rubin et al. [18] use machine learning to extract the positive/negative affective polarity of text based on up to 20 rhetorical features of sentences such as target verb, syntactic phrase type, voice, count of 'affect' words, and association with known set of words. Ang et al. [1] reported 65% accuracy with a language model differentiating frustration-annoyance and other, with a speech corpus of about 22,000 sentences. Baron and Hirst [2] applied a collocation technique to identify words with similar bipolar sentiment by their frequency of collocation with other words. Pang et al. [16] achieved 82.9% accuracy for the task of classifying 1400 film reviews from the web into positive and negative classes, using a feature vector of 16,165 words with no stemming or stop-lists. Using this movie review corpus, it was found that standard machine learning methods (Naive Bayes, Maximum Entropy classification, and Support Vector Machines) do not perform as well on sentiment classification as on traditional topic-based categorization. Lee et al [8] augmented speech acoustic emotion recognition with text classification in a help desk application.

3) Systems for sentiment identification in text

Few NLP systems have been developed for the multi-class emotion classification problem. Logan et al. [10] used latent semantic indexing of 15,589 pop lyrics by 399 artists for extracting genre. Polzin and Waibel [17] achieved 46.7% F-measure to classify 5,750 movie dialogue segments into 3 classes (neutral, angry, sad). Devillers et al. [3] reported 67.3% accuracy with 5 categories (anger, fear, satisfaction, excuse, neutral) using unigrams with stemming and compounding. Schuller et al. [20] used Bayesian belief networks to determine whether an automobile-task dialog was emotional, and if so categorized it by 6 primary emotions.

The research reported in the emotion extraction in music lyrics is by Ogihara [15]. He used lyrics for identifying clusters of 45 artists and 55 albums (such as Carly Simon, James Taylor, Joni Mitchell, Suzanne Vega etc). Accuracy of lyrics was comparable with using sound (0.635 vs. 0.685), as was precision (0.572 vs. 0.654) and recall (0.622 vs. 0.714). The F-measure for lyrics (0.602) was comparable to that of sound (0.669). An important difference from the our work is the number of classes, as this previous research used 3 emotion classes while our emotion categories have 23 classes.

III. EXPERIMENTAL DESIGN

A. Analysis of song lyrics data

While it is common to use the bag-of-words representation of documents in text mining, the resulting

high-dimensionality of the feature vector causes problems for machine learning algorithms. This problem occurs in song lyrics, even though the problem is not related to a large domain-specific technical vocabulary, but each social group tends to create its own vocabulary to distinguish itself from other groups. With an “EMO” dataset of 65,000 song lyrics, the number of unique words increased more slowly than the number of documents.

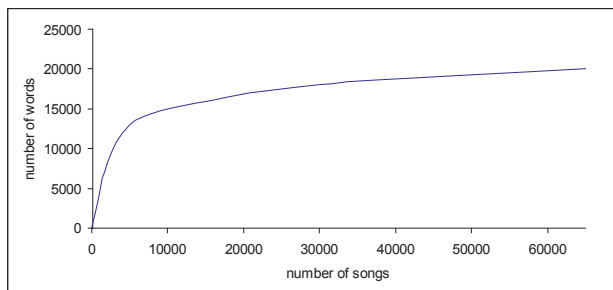


Figure 1. Relationship between numbers of songs and words

According to Zipf’s Law, many words occur with very low frequency. The information of these seldom-appearing words may have the most power to distinguish texts in applications such as emotion recognition. Furthermore, with statistical learning techniques, several training examples are needed to distinguish each new word. But because each additional song covers very few new words, it is hardly possible to collect a sufficient number of songs for enough examples to train the machine learning algorithm to distinguish large numbers of words, as shown in Figure 1.

To solve the problem of feature sparseness, a number of feature-reduction methods are commonly used in text mining. These reduce the high dimensionality of the data, but are more suitable for topic identification than emotion identification.

B. Modeling music emotion categories

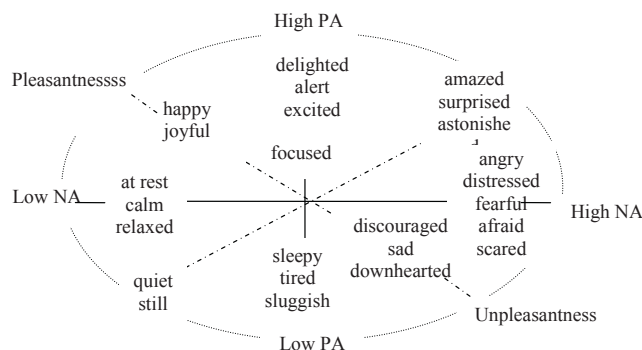


Figure 2. Psychological model of emotion (from [27])

Music emotion is generally seen as irreducible to simply one or two dimension ratings. For example, the online All Music Guide (allmusic.com) uses 168 different discrete emotion categories (e.g. trippy, quirky) to classify over 10,000 songs, albums and artists.

Table 1. Comparison of Allmusic.com mood categories (N=168) with the emotion descriptors in our model, EMO, (N=23)

Positive emotions		Negative emotions	
Allmusic	EMO	Allmusic	EMO
Brash, Bravado, Swaggering	Proud	Acerbic, Aggressive, Bitter, Fiery, Nasty, Outraged, Rebellious, Snide	Anger
Bright, Effervescent, Gleeful, Humorous, Irreverent, Party/Celebratory, Rambunctious, Raucous, Rollicking, Silly, Sweet, Whimsical, Witty	Joy	Aggressive	Aggressive
Yearning, Sexy, Provocative, Sleazy, Sensual	Arousal	Bleak, Distraught, Somber, Poignant, Melancholy, Plaintive	Sad
Exuberant	Exuberant	Gloomy	Gloomy
Passionate	Passionate	Cynical, Ironic	Alienation
Happy	Happy	Manic, Paranoid, Spooky, Unsettling	Fear
Wistful	Reflective	Ominous	Ominous
Romantic, Precious, Reverent, Sentimental, Warm	Love	Eerie	Eerie
Soothing	Soothing	Tense	Tense
Laidback/Mellow, Gentle, Refined/Mannered, Reserved, Restrained	Calm	Greasy	Disgust
Confident	Confident	Eccentric, Fractured, Trippy, Freakish, Spacey	Psychotic
Detached	Reflective	Aggressive, Boisterous, Brittle, Brooding, Cold, Confrontational, Harsh, Wry, Dramatic, Theatrical, Volatile, Thuggish, Trashy, Visceral, Earthy, Rustic	
Complex, Enigmatic, Eccentric, Street-smart, Stylish, Earnest, Ramshackle, Smooth, Slick, Sophisticated, Gritty, Spiritual, Searching, Playful, Rousing, Free-wheeling			

Psychological research has explored ways of unifying the dimensional and discrete approaches to emotion ratings. Sloboda and Juslin [22] note that dimensional and discrete models of music emotion can be complementary. One accessible approach is the PANAS-X [31] test scale which

has two dimensional ratings called Positive Affect (PA) and Negative Affect (NA). The dimensional ratings function as entry points to more detailed ratings of discrete emotions under each axis (e.g. Fear under NA). The two PANAS-X dimensions can be mathematically related to Russell’s

circumplex model [19]. Russell's Arousal is the sum of PA and NA, while Russell's Valence is the difference (PA - NA). Tellegen, Watson and Clark [27] use the Valence dimension (pleasant-unpleasant) as the top-level entry point of a 3-layer model. This unified model offers the benefits of dimensional ratings, plus a theoretical basis that navigates from the dimensional entry-point of the hierarchy to 12 discrete emotion categories.

The 168 Allmusic categories with general psychological emotions started from the PANAS-X model is refined to 23 emotions by distinguishing between alienation, sympathy and sadness; dread, pride and fear; and frustration, annoyance and anger as shown in Table 1. There are some Allmusic categories that could not be classified as specific emotions in "EMO" category, depending on their PA/NA dimensional value. In refining the list of emotions, we excluded descriptors related more to the surface sound perception than to emotion.

C. General Inquirer psychological features

From the literature review, it appeared that using psychological features of text, rather than a bag-of-words representation, would be an effective transformation to reduce the high-dimensionality of the feature vector. Initial tests with uni-gram and bi-grams using Rainbow [13] showed there was a need to generalize beyond individual words as features in order to obtain a comprehensible model. A different technique was needed in specific emotion detection than in emotional polarity detection, where a small number of adjectives are often used such as 'good' or 'bad'. Whissell's Dictionary of Affect in Language [32] is also found that each song's numerical value is too close to the average to distinguish emotional valence from others. Finally the General Inquirer is chosen because of its general scope of psychological features, its coverage of over 11,700 English words, its long track record in previous content analysis, its basis in psychological theory and its comprehensible results. The General Inquirer was often used for emotional polarity detection as well, but it also had extensive features for specific emotion identification.

There are 182 psychological features in General Inquirer Harvard IV.4 model [24], grouped in the following categories: power/strength/dominance vs. weak, active vs. passive, feeling, understatement vs. exaggeration of expression, types of jargon (legal, military, religious etc), words related to roles, demographic/kinship words, objects, places, communication, motivation, process/change, cognitive orientation, pronouns, state verbs, action verbs, social adjectives, asocial adjectives, moral words, respect status, affection status, wealth, well-being, enlightenment, skill, and others.

IV. EMOTIONAL IDENTIFICATION IN LYRICS

For the dataset of 1032 songs that randomly chosen from the allmusic.com site which is an online service for browsing music by mood, and the lyrics for these songs were located from popular music lyric websites, the number of unique words was over 5,000 so these were transformed into a much smaller set of 182 features of General Inquirer. For each of the 182 features, the feature vector recorded the incidence of the feature in the song divided by the total number of words in the song using a Perl script. Zero values were replaced with small random values using WEKA[33]'s built-in Laplacian facility. We also chose decision tree or rule methods for deriving a human-comprehensible model. The best accuracy was achieved with the DECORATE (Diverse Ensemble Creation by Oppositional Relabeling of Artificial Training Examples) algorithm [14] for generating classifier ensembles. Accuracy of 67% is estimated by 10-fold cross-validation [33], a holdout method which reserves a certain amount for testing and uses the remainder for training. This result of 23-class classification on 1032 songs with roughly 5,000 unique words is similar to that reported in the literature such as Ang et.al [1], which does 2-class classification on a speech corpus of about 22000 sentences. The biggest benefit of our approach is human comprehensibility of the classification model generated, which is useful for further refinement in a particular application such as combining classifiers. The decision tree derived from this model was inspected to check that the relationships between General Inquirer features and the class variable made sense from a common-sense perspective. The rules shown in Table 2 and table 3 do look valid, and relate to everyday ideas about emotion.

Like-valenced emotions are difficult to distinguish into specific discrete emotions, especially for negative emotions [19][27]. Specific discrete emotions are difficult to distinguish using acoustic data alone [25], even though acoustic features alone are adequate for broadly distinguishing between positive and negative emotions. Our results for positive and negative emotion identification in lyrics reflect show the similar success rate, differently from acoustic feature extraction [34] where the negative emotions are difficult to distinguish on the basis of acoustic features alone. This confirms our approach of considering emotion in music in terms of the composition of both acoustic and textual effects, corresponding to different levels of emotion processing in the brain.

A. Identifying Specific Positive/Negative Emotions

Table 2 shows the psychological features that WEKA mined from song lyrics associated with positive emotions. These results are taken from shallower parts of the classification decision tree. Informally, these results are recognizable in terms our common-sense understanding of positive emotions that are often directed to reaching some goal such as love from others, or respect from others.

Table 2. General Inquirer [23] features that six most distinguished positive emotion in lyrics

Emotion	WEKA Rule	Explanation of WEKA Rule and General Inquirer tag
Love	~KNOW	Not knowing-type words
	~Politics	Not political
	~LossWB	Not loss of well-being
	~Negativ	Not negative
	~FAIL	Not failing
	GainLF	Gains from love and friendship
	Passive	Passive-type word
Excitement	~SAY	Not saying-type word
	GainWB	Gains of well-being from relationship
Pride	ANI	Animal-type words
	Politics	Political words
	RESP	Respect
	BEGIN	Initiate change
Attentive	KNOW	Knowing-type words
	KNOW	Knowing-type words
Reflective	COLOR	Color words
	Passive	Passive type words
Calm	Comp	Completion of a goal-type words

Table 3. General Inquirer [23] features of lyrics text that three most distinguish emotion in lyrics by negative emotion

Emotion	WEKA Rule	Explanation of WEKA Rule and General Inquirer tag
Hostility	~GainLF	Not Words about gain from love or friendship
	~PartLF	Not Being a participant in love or friendship
	~Understand	Not Words about understanding
Sadness	NEED	Words expressing a need or intent
	LossWB	Loss of well-being
	GainWB	Words about a gain of well-being
	~COLOR	Not About colors
Guilt	~Relation	Not About relationship
	SAY	Saying-type words
	SAY	About Talking
	GainLF	About gains from love or friendship
	Passive	Passive-type words

Table 3 shows the psychological features that WEKA

mined from song lyrics with a specific negative emotion such as hostility, sadness or guilt. The types of features WEKA found appear reasonable, in commonsense terms, for instance sadness does focus on a loss of well-being. In data mining, the most discriminating features may not be the most descriptive from a commonsense understanding, however.

V. CONCLUSION AND FUTURE DIRECTION

We proposed the novel approach of identifying discrete music emotions from lyrics. Much of the previous literature in verbal emotion identification has been based on everyday

conversations such as call centers and customer service, not song lyrics; has been about emotional polarity of text, not specific emotion categories; has used bag-of-words representation not conceptual transformation; or has resulted in black-box models rather than human-comprehensible models. The study applies statistical text processing tools to the novel application area of identifying emotion in music using lyrics.

A psychological model of emotion is developed that covers 23 specific emotions and this is used to classify 168 different music moods used to describe thousands of songs online at allmusic.com. A training set of 1032 songs was randomly chosen from the allmusic.com site which provides online service for browsing music by mood and the lyrics for these songs were located from popular music lyric websites. Each song lyric was transformed into a feature vector of 182 psychological features using a content analysis package. Text features most related to specific emotions are chosen, and classification decision trees were used to understand the classification models produced. Our results for mining the lyrical text for specific emotion are promising, generate classification models that are humanly-comprehensible, and generate results that correspond to commonsense intuitions about specific emotions. This study is one component of a data fusion experiment that combined classifiers, such as audio and text, with different errors.

In future work, we will experiment with transforming text with different feature typologies and compare their performance to acoustic features in identifying specific emotions.

VI. ACKNOWLEDGMENTS

We would like to thank Professor Philip J. Stone of Harvard University, for letting us use the Java version 1.0 of the General Inquirer.

REFERENCES

- [1] ANG, J., DHILLON, R., KRUPSKI, A., SHRIBERG, E., AND STOLCKE, A. 2002. Prosody-based automatic detection of annoyance and frustration in human-computer dialog. In Proceedings of ICSLP2002. vol. 3, pp. 2037–2040.
- [2] BARON, F. AND HIRST, G. 2004. Collocations as Cues to Semantic Orientation. In AAAI Exploring Emotion and Affect in Text workshop.
- [3] DEVILLERS, L., VASILESCU, I., AND LAMEL, L. 2002. Annotation and detection of emotion in a task-oriented human-human dialog corpus. In Proceedings of ISLE Workshop.
- [4] FELLBAUM, C. 1998. Wordnet: an electronic lexical database. MIT Press, Cambridge, MA, USA.
- [5] FILLMORE, C. 2004. Berkeley Framenet Project. Online www.icsi.berkeley.edu/~framenet.
- [6] GORDON, A., KAZEMZADEH, A., NAIR, A., AND PETROVA, M. 2003. Recognizing Expressions of Commonsense Psychology in English Text. In Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics ACL-2003. Sapporo, Japan.

- [7] HANSEN, C.H. AND HANSEN, R.D. 2000. Music and Music Videos, In ZILLMAN, D. AND VORDERER, P. (Eds.) 2000 Media Entertainment: the psychology of its appeal, Lawrence Erlbaum Associates Publishers, Mahwah NJ, USA.
- [8] LEE, C.M., NARAYANAN, S.S. AND PIERACCINI, R. 2002. Combining acoustic and language information for emotion recognition. In Proceedings of ICSLP 2002. Denver CO USA
- [9] LIDDY, E.D. 2004. Advances in Information Extraction of Elusive Information from Text; Presentation to the Armed Forces Communications and Electronics Association. Clinton, NY USA. Online www.cnlp.org/presentations/slides/AFCEA.2004.Liddy.ppt
- [10] LOGAN, B., KOSITSKY, A. AND MORENO, A. 2004. Semantic analysis of song lyrics. HP Labs Tech Report HPL-2004-66. Cambridge, Mass. USA.
- [11] MANNING, C. AND SCHUTZE, H. 1999. Foundations of statistical natural language processing. MIT Press Cambridge MA, USA.
- [12] MATHIEU, Y.Y. 2004. FEELING : A Semantic Lexicon for Emotions and Feelings. In AAAI Exploring Emotion and Affect in Text workshop 2004.
- [13] MCCALLUM, A. 2004. Bow: A Toolkit for Statistical Language Modeling, Text Retrieval, Classification and Clustering. Online www-2.cs.cmu.edu/~McCallum/bow/.
- [14] MELVILLE, P. AND MOONEY, R. 2003. Constructing diverse classifier ensembles using artificial training examples. In Proc. of 18th Intl. Joint Conf. on Artificial Intelligence, Acapulco, Mexico, August 2003. 505–510
- [15] OGIHARA, M. 2004. Data mining for studying large databases: Course presentation Online location: www.cs.rochester.edu/u/scott/200/2004-04-08_ogihara_slides.pdf.
- [16] PANG, B., LEE, L. AND VAITHYANATHAN, S. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. In Proceedings of EMNLP2002.
- [17] POLZIN, T. AND WAIBEL, A. 2000. Emotion-sensitive human-computer interfaces. In Proceedings of the ISCA-Workshop on Speech and Emotion.
- [18] RUBIN, V.L., STANTON, J. M. AND LIDDY, E.D. 2004. Discerning Emotions In Texts. In AAAI Exploring Emotion and Affect in Text Workshop 2004.
- [19] RUSSELL, J.A. 2003. Core affect and the psychological construction of emotion. *Psychological Review* Vol. 110, No. 1, 145-172, Jan 2003.
- [20] SCHULLER, B., RIGOLL, G., AND LANG, M. 2004. Speech emotion recognition combining acoustic features and linguistic information in a hybrid support vector machine – belief network architecture. IEEE International Conference on Speech, Acoustics and Signal Processing 2004.
- [21] SCOTT, S. AND MATWIN, S. 1998. Text classification using WordNet hypernyms. In Usage of WordNet in Natural Language Systems. Online acl.ldc.upenn.edu/W/W98/W98-0706.pdf.
- [22] SLOBODA, J.A. AND JUSLIN, P.N. 2001. Psychological perspective on emotion. In *Music and Emotion*. Oxford University Press, New York, NY, USA.
- [23] STONE, P.J., DUNPHY, D.C., SMITH, M.S., AND OGILVIE, D.M. (EDS.) 1966. *The General Inquirer: A computer approach to content analysis*. MIT Press, Cambridge, MA, USA.
- [24] STONE, P.J. 2000. *General Inquirer Categories*, Harvard University, Cambridge MA, USA. Online www.wjh.harvard.edu/~inquirer/homecat.htm.
- [25] STRATTON, V.N. AND ZALANOWSKI, A.H. 1994. Affective Impact of Music vs. Lyrics. *Empirical Studies of the Arts* 12:2, 1994, 173-184.
- [26] SUBASIC, P. AND HUETTNER, A. 2004. Method and apparatus for information management using fuzzy typing. US Patent 6,721,734.
- [27] TELLEGEN, A., WATSON, D. AND CLARK, L.A. 1999. On the dimensional and hierarchical structure of affect. *Psychological Science*, Vol. 10, No. 4, July 1999.
- [28] TURNEY, P.D. AND LITTMAN, M.L. 2002. Unsupervised learning of semantic orientation from a hundred-billion word corpus. National Research Council, Ottawa, Canada.
- [29] VOORHEES, E. (ED.) 2003. Proceedings of the Twelfth Text Retrieval Conference TREC2003. NIST Special Publication 500-255. US Department of Commerce, Gaithersburg, MD, USA.
- [30] VORDERER, P. 2001. It's all entertainment – Sure. But what exactly is entertainment? Communication research, media psychology and the explanation of entertainment experiences. *Poetics* 29 2001. 247-261.
- [31] WATSON, D. AND CLARK, L.A. 1999. The PANAS-X Manual for the Positive and Negative Affect Schedule – Expanded Form. Online <http://www.psychology.uiowa.edu/Faculty/Clark/PANAS-X.pdf>.
- [32] WHISSELL, C. 1989. *Dictionary of Affect in Language* In PLUTCHIK AND KELLERMAN (Eds.) *Emotion: Theory, Research and Experience*. Vol 4 Academic Press, NY.
- [33] WITTEN, I.H., AND FRANK, E., 2000. *Data Mining: Practical machine learning tools and techniques with Java implementations*. Morgan Kaufmann, San Francisco, CA, USA.
- [34] YANG, D. AND LEE, W.S. 2004. Disambiguating Music Emotion Using Software Agents. 5TH International Conference on Music Information Retrieval ISMIR2004. Barcelona, Spain,
- [35] ZILLMAN, D., AND VORDERER, P. 2000. *Media Entertainment: the psychology of its appeal*. Lawrence Erlbaum Associates Publishers, Mahwah NJ, USA.