Automatizing Two Creative Functions for Advertising

Carlo Strapparava and **Alessandro Valitutti** and **Oliviero Stock**
FBK-irst, I-38050, Povo, Trento, ITALY
{strappa, alvalitu, stock}@itc.it

Abstract

The creation of advertising messages is a deep process of creative writing production. As far as the textual content is concerned, there are not many computational tools (besides the usual dictionaries, thesauri or program for performing of simple wordplays) that help the copywriter activity. In this work we explore the use of natural language processing and text animation techniques for proposing solutions to advertising professionals and improving the quality of advertising messages. In the proposed system, we consider two steps: (i) the creative variation of familiar expressions, taking into account the affective content of the produced text, (ii) the automatic animation (semantically consistent with the affective text content) of the resulting expression, using kinetic typography techniques.

Keywords: Natural Language Processing, Affective Text, Lexical Semantics, WORDNET, Text Animation.

1 Introduction

In modern advertising practice, it is common of "creatives" to be recruited and hired in pairs formed by a copywriter and an art director. They work in a creative partnership to conceive, develop and produce effective advertisements. While the copywriter is mostly responsible for the textual content of the creative product, the art director focalizes efforts on the graphical presentation of the message. Advertising messages tend to be quite short but, at the same time, rich of emotional meaning and persuasive power. While computational tools are an essential complement in many creative activities, e.g. graphical design, there are few tools for creation of textual messages (except for the usual dictionaries, thesauri or programs for performing simple wordplays).

In this paper we explore the development of computational tools to improve the quality of advertising messages, reducing the development time and possibly opening up the way to a full automatization of the whole process of creative writing production. We combine some computational functionalities for the creation of advertising messages. In particular, we implemented a strategy that is articulated in two steps. The first consists of the selection and creative variation of familiar common sense expressions (e.g. proverbs, idioms, cliches, movie titles, famous citations, etc.). The second step consists of the presentation of the expression through an automated text animation, and it is based on the use of kinetic typography. As we will see, the text animation can be built semantically consistent with the emotion we want to convey.

1.1 Advertising Messages and Optimal Innovation

An advertising message induces in the recipient a positive (or negative) attitude toward the subject to advertise (Petty and Wegener, 1998), for example through the evocation of an appropriate emotion. Another mandatory characteristic of an advertisement is its memorability. These two aspects of an ad increase the probability to induce some wanted behaviours, for example the purchase of some product, the choice of a specific brand, or the click on some specific web link. In the last case, it is crucial to make the recipient curious about the subject referred by the URL. The best way to realize in an ad both attitude induction and memorability is the generation of surprise, generally based on creative constraints.

In order to develop a strategy for surprise induction, we considered an interesting property of pleasurable creative communication that was named by Rachel Giora as the *optimal innovation hypothesis* (2003). According to this assumption, when the novelty is in a complementary relation to salience (familiarity), it is "optimal" in the sense that it has an aesthetics value and "induces the most pleasing effect".

Therefore the simultaneous presence of novelty and familiarity makes the message potentially surprising, because this combination allows the recipient's mind to oscillate between what is known and what is different from usual. For this reasons, an advertising message must be original but, at the same time, connected to what is familiar (Pricken, 2002). Familiarity causes expectations, while novelty violates them, and finally surprise arises.

1.2 Familiar Expression Variation

With "familiar expression variation" we indicate an expression (sentence or phrase) that is obtained as a linguistic change (e.g. substitution of a word, morphological or phonetic variation, etc.) of an expression recognized as familiar by recipients (e.g. selected by some collection of proverbs, famous movie titles, etc.). In this work we limited the variation to word substitution.

Moreover, an ad has to own a semantic connection with some concept of the target topic. At the same time, it has to be semantically related with some emotion of a prefixed valence (e.g. positive emotion as `joy` or negative emotion as `fear`) .

We combined all these constraints in a compatible way with the optimal innovation hypothesis. The "innovation" is provided by the semantic similarity with the target topic and with the emotion, and the "optimality" is guaranteed by the assonance (i.e. the old and the new word have to be assonant, e.g. rhymed).

We considered in this process some recent works in computational humor (e.g. (Stock and Strapparava, 2003)), in which incongruity theory is exploited to produce funny variations of given acronyms[1]. In this work we extend this approach, focussing on the affective load of lexicon for the variation production and generating automatically typographical animations that are coherent with the emotions we want to communicate.

The paper is structured as follows. In Section 2 we introduce the resources used in the system, in particular (i) WORDNET-AFFECT, an extension of the WordNet database in which some affective labels are assigned to a number of synsets; (ii) an affective semantic similarity, based on a Latent Semantic Analysis, which gives us an indication of the affective weight of generic terms; (iii) databases of familiar expressions and assonance tools; and (iv) a kinetic typography scripting language used for the final sentence animation. Section 3 describes the algorithm to variate familiar expressions and Section 4 displays some examples. Conclusions and future works are reported in Section 5.

2 Resources

2.1 Affective Semantic Similarity

All words can potentially convey affective meaning. Each of them, even those more apparently neutral, can evoke pleasant or painful experiences, because of their semantic relation with emotional concepts. While some words have emotional meaning with respect to the individual story, for many others the affective power is part of the collective imagination (e.g. words "mum", "ghost", "war" etc.).

We are interested in this second group, because their affective meaning is part of common sense knowledge and can be detected in the linguistic usage. For this reason, we studied the use of words in textual productions, and in particular their co-occurrences with the words in which the affective meaning is explicit. As claimed by Ortony et al. (Ortony et al., 1987), we have to distinguish between words directly referring to emotional states (e.g. "fear", "cheerful") and those having only an indirect reference that depends on the context (e.g. words that indicate possible emotional causes as "killer" or emotional responses as "cry"). We call the former *direct affective words* and the latter *indirect affective words* (Strapparava et al., 2006).

In order to manage affective lexical meaning, we (i) organized the direct affective words and synsets inside WORDNET-AFFECT, an affective lexical resource based on an extension of WORDNET, and (ii) implemented a selection function (named *affective weight*) based on a semantic similarity mechanism automatically acquired in an unsupervised way from a large corpus of texts (100 millions of words), in order to individuate the indirect affective lexicon.

Applied to a concept (e.g. a WORDNET synset) and an emotional category, this function returns a value representing the semantic affinity with that emotion. In this way it is possible to assign a value to the concept with respect to each emotional category, and eventually select the emotion with the highest value. Applied to a set of concepts that are semantically similar, this function selects subsets characterized by some given affective constraints (e.g. referring to a particular emotional category or valence).

As we will see, we are able to focus selectively on positive, negative, ambiguous or neutral types of emotions. For example, given "difficulty" as input term, the system suggests as related emotions: IDENTIFICATION, NEGATIVE-CONCERN, AMBIGUOUS-EXPECTATION, APATHY. Moreover, given an input word (e.g. "university") and the indication of an emotional valence (e.g. positive), the system suggests a set of related words through some positive emotional category (e.g. "professor" "scholarship" "achievement") found through the emotions ENTHUSIASM, SYMPATHY, DEVOTION, ENCOURAGEMENT.

This fine-grained affective lexicon selection can open up new possibilities in many applications that exploit verbal communication of emotions. For example, (Valitutti et al., 2005) exploited the semantic connection between a generic word and an emotion for the generation of affective evaluative predicates and sentences.

WORDNET-AFFECT and the Emotional Categories. WORDNET-AFFECT is an extension of the WordNet database (Fellbaum, 1998), including a subset of synsets suitable to represent affective concepts. Similarly to what was done for domain labels (Magnini and Cavaglià, 2000), one or more affective labels (*a-labels*) are assigned to a number of WordNet synsets. In particular, the affective concepts representing an emotional state are individuated by synsets marked with the a-label EMOTION. There are also other a-labels for those concepts representing moods, situations eliciting emotions, or emotional responses. WORDNET-AFFECT is freely available for research purpose at `http://wndomains.itc.it`. See (Strapparava and Valitutti, 2004) for a complete description of the resource.

We extended WORDNET-AFFECT with a set of additional a-labels (i.e. the *emotional categories*), hierar-

[1] As far as computational humor is concerned and in particular funny variations of existing expressions, it is worth mentioning the work on pun creation of Binsted and Ritchie (1997).

	# Synsets	# Words	# Senses
Nouns	280	539	564
Adjectives	342	601	951
Verbs	142	294	430
Adverbs	154	203	270
Total	918	1637	2215

Table 1: Number of elements in the emotional hierarchy.

chically organized, in order to specialize synsets with a-label EMOTION. In a second stage, we introduced some modifications, in order to distinguish synsets according to emotional valence. We defined four additional a-labels: POSITIVE, NEGATIVE, AMBIGUOUS, NEUTRAL. The first one corresponds to "positive emotions", defined as emotional states characterized by the presence of positive edonic signals (or pleasure). It includes synsets such as `joy#1` or `enthusiasm#1`. Similarly the NEGATIVE a-label identifies "negative emotions" characterized by negative edonic signals (or pain), for example `anger#1` or `sadness#1`. Synsets representing affective states whose valence depends on semantic context (e.g. `surprise#1`) were marked with the tag AMBIGUOUS. Finally, synsets referring to mental states that are generally considered affective but are not characterized by valence, were marked with the tag NEUTRAL.

Computing Lexical Affective Semantic Similarity. There is an active research direction in the NLP field about sentiment analysis and recognition of semantic orientation from texts (e.g. (Turney and Littman, 2003; Liu et al., 2003; Mihalcea and Liu, 2006)). In our opinion, a crucial issue is to have a mechanism for evaluating the semantic similarity among generic terms and affective lexical concepts. To this aim we estimated term similarity from a large scale corpus. In particular we implemented a variation of Latent Semantic Analysis (LSA) in order to obtain a vector representation for words, texts and synsets.

In LSA (Deerwester et al., 1990), second order relations among terms and documents of the corpus are captured by means of a dimensionality reduction operated by a Singular Value Decomposition (SVD) on the term-by-document matrix. For the experiments reported in this paper, we run the SVD operation on the full British National Corpus[2].

SVD is a well-known operation in linear algebra, which can be applied to any rectangular matrix in order to find correlations among its rows and columns. SVD decomposes the term-by-document matrix $\mathbf{T}$ into three matrices $\mathbf{T} = \mathbf{U}\mathbf{\Sigma_k}\mathbf{V}^T$ where $\mathbf{\Sigma_k}$ is the diagonal $k \times k$ matrix containing the k singular values of $\mathbf{T}$, $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_k$, and $\mathbf{U}$ and $\mathbf{V}$ are column-orthogonal matrices. When the three matrices are multiplied together the original term-by-document matrix is re-composed. Typically we can choose $k' \ll k$ obtaining the approximation $\mathbf{T} \simeq \mathbf{U}\mathbf{\Sigma_{k'}}\mathbf{V}^T$. More specifically, in the experiments for this paper we use the matrix $\mathbf{T}' = \mathbf{U}\mathbf{\Sigma_{k'}}$, whose rows represent the term vectors in the reduced space, taking into account the first 100 dimensions (i.e. $k' = 100$).

LSA can be viewed as a way to overcome some of the drawbacks of the standard vector space model (sparseness and high dimensionality). In fact, the LSA similarity is computed in a lower dimensional space, in which second-order relations among terms and texts are exploited. The similarity in the resulting vector space can be measured with the standard cosine similarity. Note also that LSA yields a vector space model that allows for a *homogeneous* representation (and hence comparison) of words, word sets, sentences and texts.

For representing word sets and texts by means of a LSA vector, we used a variation of the *pseudo-document* methodology described in (Berry, 1992). This variation takes into account also a *tf-idf* weighting schema (see (Gliozzo and Strapparava, 2005) for more details). In practice, each document can be represented in the LSA space by summing up the normalized LSA vectors of all the terms contained in it. Also a synset in WORDNET (and then an emotional category) can be represented in the LSA space, performing the pseudo-document technique on all the words contained in the synset. Thus it is possible to have a vectorial representation of each emotional category in the LSA space (i.e. the *emotional vectors*), and consequently we can compute a similarity measure among terms and affective categories. We defined the *affective weight* as the similarity value between an emotional vector and an input term vector (e.g. we can check how a generic term is similar to a given emotion).

For example, the noun "gift" is highly related to the emotional categories: LOVE (with positive valence), COMPASSION (with negative valence), SURPRISE (with ambiguous valence), and INDIFFERENCE (with neutral valence).

In summary, the vectorial representation in the Latent Semantic Space allows us to represent, in a *uniform* way, emotional categories, generic terms and concepts (synsets), and eventually full sentences.

2.2 Database of Familiar Expressions

The base for the strategy of "familiar expression variation" is the availability of a set of expressions that are recognized as familiar by English speakers.

We considered three types of familiar expressions: proverbs, movie titles, clichés. We collected 1836 familiar expressions from the Web, organized in three types: common use proverbs (628), famous movie titles (290), and clichés (918). Proverbs were retrieved in some of many web sites in which they are grouped (e.g. `http://www.francesfarmersrevenge.com/stuff/proverbs.htm` or `www.manythings.org/proverbs`). We considered only proverbs of common use. In a similar way we collected clichés, that are sentences whose overuse often makes them humorous (e.g. home sweet home, I am playing my own game). Finally, movie titles were selected from the Internet Movie Database (`www.imdb.com`). In particular, we considered the list of the best movies in all sorts of categories based on votes from users.

[2]The British National Corpus is a very large (over 100 million words) corpus of modern English, both spoken and written (see `http://www.hcu.ox.ac.uk/bnc/`).

The list of familiar expressions is composed mostly of sentences (in particular, proverbs and clichés), but part of them are phrases (in particular, movie title list includes a significant number of noun phrases)

2.3 Assonance Tool

To cope with this aspect we got and re-organized the CMU pronouncing dictionary (`http://www.speech.cs.cmu.edu/cgi-bin/cmudict`) with a suitable indexing. The CMU Pronouncing Dictionary is a machine-readable pronunciation dictionary for North American English that contains over 125,000 words and their transcriptions.

Its format is particularly useful for speech recognition and synthesis, as it has mappings from words to their pronunciations in the given phoneme set. The current phoneme set contains 39 phonemes; vowels may carry lexical stress.

2.4 Kinetic Typography Scripting Language

Kinetic typography is the technology of text animation, i.e. text that uses movement or other changes over time. The advantage of kinetic typography consists in a further communicative dimension, combining verbal and visual communication, and providing opportunities to enrich the expressiveness of static texts. According to (Lee et al., 2002), kinetic typography can be used for three different communicative goals: capturing and directing attention of recipients, creating characters, and expressing emotions. A possible way of animating a text is mimicking the typical movement of humans when they express the content of the text (e.g. “Hi” with a jumping motion mimics exaggerated body motion of humans when they are really glad).

We explore the idea to have a link between lexical semantics of texts (automatically discerned through NLP techniques) and some kinetic properties exploited for animating the words. In this paper, we consider affective connotation of texts by exploiting the affective semantic similarity introduced above. This holds particularly for “indirect affective words” (Strapparava et al., 2006). For example, these words may indicate possible emotional causes (e.g. “monster”) or emotional responses (e.g. “cry”). Thus kinetic typography allows us to make the indirect affective meaning explicit in order to automatically augment the affective expressiveness of texts.

A first step was the individuation of an appropriate tool for the authoring and visualization of text animations. In particular, we wanted to act in an environment that allows us to realize animations in a very simple manner and to represent them in an easily exportable format. Functionalities for the automated composition of animations were our specific concern. To this aim we considered the Kinetic Typography Engine (KTE), a Java package developed at the Design School of Carnegie Mellon University (Lee et al., 2002). It allows us to create a potentially wide range of animations. Taking this engine as a starting point, we first realized a development environment for the creation and the visualization of text animations. Our model for the animation representation is a bit simpler than the KTE model. The central assumption consists of the representation of the animation as a composition of elementary animations (e.g. linear, sinusoidal or exponential variation). In particular, we consider only one operator for the identification of elementary animations (K-BASE) and three composition operators: kinetic addition (K-ADD), kinetic concatenation (K-JOIN), and kinetic loop (K-LOOP).

The K-BASE operator selects an elementary animation (named *elementary kinetic behavior*) as a temporal variation of some kinetic property. Elementary kinetic behaviors correspond to a subset of dynamic variations implemented in KTE, for example linear variation (*linear*), sinusoidal variation (*oscillate*), and exponential variation (*exponential*).

`linear`	linear variation
`oscillate`	sinusoidal variation
`pulse`	impulse
`jitter`	sort of “chaotic” vibration
`curve`	parabolic variation
`hop`	parabolic variation with small impulses at the endpoints
`hop-secondary`	derivative of hop, used as secondary effect to simulate elastic movements

Table 2: Some elementary kinetic behaviors

The kinetic addition (K-ADD) of two animations with the same start time is obtained by adding, for each kinetic property of text, the corresponding dynamical variation of each single animation. The kinetic concatenation (K-JOIN) consists in the temporal shifting of the second animation, so that the ending time of the first is the starting time of the second. The kinetic loop (K-LOOP) concatenates an animation with itself a fixed number of times. In the development environment it is possible to freely apply these operators for the real time building of new animations. Compositional structure of animations can be represented in XML format and then easily exported. Finally, an interpreter allows us to generate in real time the animation starting from its structural representation.

Figure 2: Jittering *anger*

After building the development tool, we selected a set of emotional categories and, for each of them, we created the corresponding text animations.

In particular, we focused on five emotional categories: joy, fear, surprise, anger, sadness (i.e. a subset of Ekman emotions (Ekman, 1977)).

The kinetic animation to associate to a fixed emotion

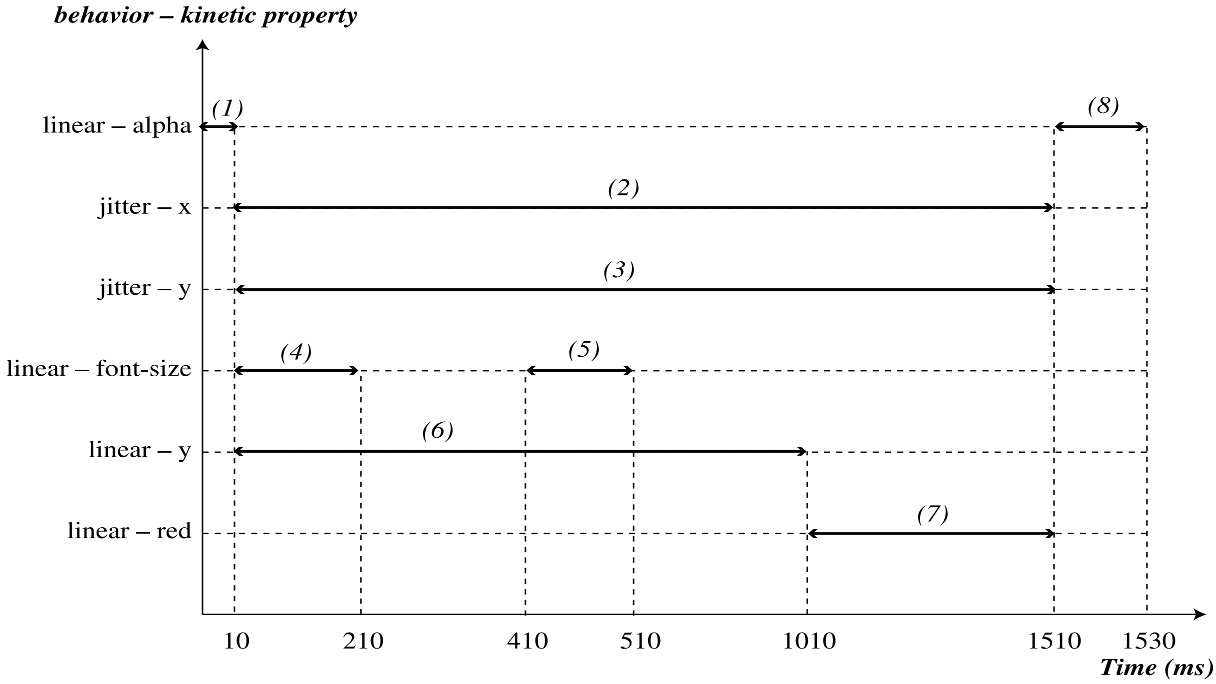

Figure 1: Kinetic behavior description for “anger” emotion

can be realized imitating either emotional and physiological responses (*analogous motion* technique), or tone of voice. We consider only animations of the first type, i.e. we represent each emotion with an animation that simulates a particular emotional behavior. In particular, JOY is represented with a sequence of hops, FEAR with palpitations, ANGER with a strong tremble and blush, SURPRISE with a sudden swelling of text, and finally SADNESS with text deflation and getting squashed. Thus we annotated the corresponding emotional categories in WORDNET-AFFECT with these kinematic properties.

Figure 1 displays in detail the behavior of the anger emotion, showing the time-dependent composition graph of the basic animations. The string appears (1) and disappears (8) with a linear variation of the alpha property (that defines the transparency of a color and can be represented by a float value). The animation is contained between these two intervals and its duration is 1500 ms. The first component is a tiny random variation of the position (2) (3), represented by x and y kinetic properties, with jitter behavior. The second component consists of an expansion of the string (4) and a subsequent compression (5). The third component is given by a slow rise up (6). The last component, before disappearing, is a color change to red (7). The whole behavior is then described and implemented using the scripting language introduced above.

As it is difficult to enjoy the animations on *static* paper, please visit the web page `http://tcc.itc.it/people/strapparava/affective-KT` where some downloadable short movies are available.

3 Algorithm

In this section, we describe the algorithm developed to perform a creative variation of an existing expression.

1. **Insertion of an input concept.** The first step of the procedure consists of the insertion of an input concept. This is represented by one or more words, a set of synonyms, or a WordNet synset. In the latter case, it is individuated through a word, the part of speech (noun, adjective, verb, or adverb), and the sense number, and it corresponds to a set of synonyms. Using the pseudo-document representation technique described above, the input word list is represented as a vector (named input-vector) in the LSA vectorial space.

2. **Generation of the target-list.** A list of terms (named target list) that are semantically connected (in the LSA space) with the input concept is generated. This target list represents a semantic domain that includes the input concept.

3. **Association of assonant words.** For each word of the target-list one or more possible *assonant words* are found. Then a list of word pairs (named *variation-pairs*) is created. Each pair has a *target word* as first element and an *assonant-word* as second element. At this point, the list of variation-pairs is filtered according to some constraints. The first one is syntactic (target-word and assonant-word must have the same part of speech). The second one is semantic (assonant word must not be included in the target-list), and its function is to maximize the probability to realize a semantic opposition between the elements of a variation pair. Finally, to each variation pair is associated an *emotion-label* (representing

the emotional category most semantically similar to the assonant word) with the corresponding value of affective weight. If a target word has more possible assonant words, we selected only that one having higher value of affective weight.

4. **Creative variation of familiar expressions.** In this step, the procedure gets as input a set of familiar expressions (in particular, proverbs and movie titles) and, for each of them, generate all possible variations. If an expression includes a word that is an element of at least one of the variation-pairs, then that word is substituted by the other element of the same pair.

5. **Ordering of familiar expressions.** The list of varied expressions is ordered according to the value of affective weight associated to the assonant word. The dimension of the input set of familiar expressions is crucial because it is related to the probability of generating a satisfactory creative variation.

6. **Text animation.** Finally, the varied expression is animated with kinetic typography technique. In particular, the assonant-word is animated according to the underlying emotion to emphasize the affective connotation.

4 Examples

In this section we want to show some examples of the creative function developed in our work and how it is useful for creating advertisements.

Simple creative variations. Using the affective weight function, it is possible to select a variation according to the valence (e.g. the substitution of the word *bad*, detected as negative, with *glad*, recognized as positive) or to some wanted affective direction. In Table 3 there are three variation of a movie title, according to three different emotions, to show that we can constrain the word substitution toward a word semantically similar to the desired emotional category.

Original	Variation	Category
Notting Hill	Notting *Thrill*	`Exhilaration`
	Notting *Still*	`Calmness`
	Notting *Chill*	`Gladness`

Table 3: Simple variation of a movie title

Humorous effects. Table 4 shows how word substitution may propagate the change of connotation at the level of the entire expression, and may also produce humorous effects. In particular, we observed that the semantic opposition, determined by switching affective polarity, generates another more complex semantic opposition at phrase (or sentence) level. In a possible scenario in which the creative user interacts with the system to generate creative expressions, the human recognition of high level humorous effects may be part of the creative interaction. The system proposes a list of possible candidates and the user makes the ultimate decision, selecting the creative variations that seem more meaningful[3].

Original	Variation	Category
when all else fails, read the instructions	when all else fails, *dread* the instructions	`Fear`
children and fools tell the truth	children and fools *repel* the truth	`Repugnance`
divide and rule	divide and *cool*	`Coolness`
a guilty conscience feels continual fear	a guilty conscience feels continual *cheer*	`Cheerfulness`

Table 4: Humorous variations

Advertising. In Table 5 there are some examples of automatically generated advertising messages. The creative variation has a semantic connection with a target topic and it is suitable for advertising purposes. In the first example, the original word *park* is substituted by the work *dark*, that have high semantic similarity with a target topic (*clothes*) and has a negative affective weight. The global expression communicates the idea that the colours for the new fashion must be clear and the dark clothes are old fashioned. The second example shows the substitution of the original word *night* with the word *fright* , that is semantically similar to the target topic *crash* and has a negative affective weight. The entire phrase can be used to warn young drivers about alcohol related driving accidents.

Original	Variation	Category
Jurassic Park	Jurassic *Dark*	`Gloom`
Saturday Night Fever	Saturday *Fright* Fever	`Fear`

Table 5: Variations for advertising messages

5 Conclusions

Exploiting some state-of-the-art natural language processing techniques, we described a system that produces creative variations of familiar expressions and animates them accordingly to the affective content. The creative textual variations are based on lexical semantics techniques such as affective similarity, while the animation makes use of a kinetic typography dynamic scripting language.

From an applied point of view, we believe that a thorough environment for proposing solutions to advertising

[3]In future work we are interested to refine a computational model that suggests the best semantic opposition or *incongruity* for humorous effect generation. Some useful considerations about the issue of incongruity can be found in (Ritchie, 1999; Veale, 2004)

professionals can be a practical development of this work, for the moment leaving the last word to the human professional. In the future, the potential of fully automatic production will find a big opportunity if advertisements are to be linked to an evolving context, such as incoming news, or changing of location of the audience, until a full user personalization of advertisements.

References

Berry, M. (1992). Large-scale sparse singular value computations. *International Journal of Supercomputer Applications*, 6(1):13–49.

Binsted, K. and Ritchie, G. (1997). Computational rules for punning riddles. *Humor*, 10(1).

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T., and Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407.

Ekman, P. (1977). Biological and cultural contributions to body and facial movement. In Blacking, J., editor, *Anthropology of the Body*, pages 34–84. Academic Press, London.

Fellbaum, C. (1998). *WordNet. An Electronic Lexical Database*. The MIT Press.

Giora, R. (2003). *On Our Mind: Salience, Context and Figurative Language*. Oxford University Press, New York.

Gliozzo, A. and Strapparava, C. (2005). Domains kernels for text categorization. In *Proc. of the Ninth Conference on Computational Natural Language Learning (CoNLL-2005)*, Ann Arbor.

Lee, J., Forlizzi, J., and Hudson, S. (2002). The kinetic typography engine: An extensible system for animating expressive text. In *Proc. of ACM UIST 2002 Conference*.

Liu, H., Lieberman, H., and Selker, T. (2003). A model of textual affect sensing using real-world knowledge. In *Proc. of the Seventh International Conference on Intelligent User Interfaces (IUI 2003)*, Miami.

Magnini, B. and Cavaglià, G. (2000). Integrating subject field codes into wordnet. In *Proc. of the 2^{nd} International Conference on Language Resources and Evaluation (LREC2000)*, Athens, Greece.

Mihalcea, R. and Liu, H. (2006). A corpus-based approach to finding happiness. In *Proc. of Computational approaches for analysis of weblogs, AAAI Spring Symposium 2006*, Stanford.

Ortony, A., Clore, G. L., and Foss, M. A. (1987). The psychological foundations of the affective lexicon. *Journal of Personality and Social Psychology*, 53:751–766.

Petty, R. and Wegener, D. (1998). Attitude change: Multiple roles for persuasion variables. In Gilbert, D., Fiske, S., and Lindzey, G., editors, *The handbook of social psychology*, pages 323–390. McGraw-Hill, New York, 4^{th} edition.

Pricken, M. (2002). *Creative Advertising*. Thames & Hudson.

Ritchie, G. (1999). Developing the incongruity-resolution theory. In *Proceedings of the AISB Symposium on Creative Language: Stories and Humour*, Edinburgh.

Stock, O. and Strapparava, C. (2003). Getting serious about the development of computational humour. In *Proceedings of the 8^{th} International Joint Conference on Artificial Intelligence (IJCAI-03)*, Acapulco, Mexico.

Strapparava, C. and Valitutti, A. (2004). WordNet-Affect: an affective extension of WordNet. In *Proc. of 4^{th} International Conference on Language Resources and Evaluation (LREC 2004)*, Lisbon.

Strapparava, C., Valitutti, A., and Stock, O. (2006). The affective weight of lexicon. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC 2006)*, Genoa, Italy.

Turney, P. and Littman, M. (2003). Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems (TOIS)*, 21(4):315–346.

Valitutti, A., Strapparava, C., and Stock, O. (2005). Lexical resources and semantic similarity for affective evaluative expressions generation. In *Proc. of the First International Conference on Affective Computing & Intelligent Interaction (ACII 2005)*, Beijing, China.

Veale, T. (2004). Incongruity in humor: Root-cause or epiphenomenon? *The International Journal of Humor*, 17(4).