

Creatively Subverting Messages in Posters

Lorenzo Gatti¹, Marco Guerini¹, Charles Callaway¹, Oliviero Stock², Carlo Strapparava²

¹Trento-Rise, ²FBK-Irst

Via Sommarive 18, Povo, Trento - Italy

{l.gatti, marco.guerini, c.callaway}@trentorise.eu, {stock, strappa}@fbk.eu

Abstract

Creativity is widely used in advertisements, and is meant to be appreciated by people. However creativity can also be used as a defense. When we walk in the street we are overwhelmed by messages which try to get our attention with any persuasive device at hand. As messages get ever more aggressive, often our basic cognitive defense – trying not to perceive those messages – is not sufficient. One advanced defensive technique is based on transforming the perceived message into something different (for instance irony or hyperbole) from what was originally meant in the message. In this paper we describe an implemented application for smartphones that creatively modifies the linguistic expression in a virtual copy of the poster. The mobile system is inspired by the subvertising practice of counter-cultural art, and aims at experiencing aesthetic pleasure that relaxes the cognitive tension of the user.

Introduction

We are surrounded by linguistic expressions on the walls around us. Whenever we walk along a street, posters, signs and other similar advertisements are there trying to attract our attention and in most cases trying to influence our actions, beliefs and behavior. We may try to avoid those ads but it is not easy: even if the characteristics of our perceptive and cognitive system partially help us in being “banner blind” (Pagendam and Schaumburg 2001; Burke, Gorman, Nilsen and Hornof 2004), pervasive advertising often manages to overcome our barriers (Müller, Alt and Michelis 2011). One strategy to counter messages that forcefully grab our attention is to use our cognitive system to fight back and creatively alter the advertising message itself. This form of “reactive” creativity lies at the root of various phenomena, including some aspects of verbal humor, especially irony. The psychoanalytical approach to humor (Freud 1905), gives an attractive account of the release of energy that results from overcoming our inner censors through the appreciation of humorous expressions. A similarly liberating process can be attributed to other types of variations of linguistic expressions.

From an aesthetic point of view a given variation is more highly appreciated if the change is limited, as for instance suggested by the optimal innovation theory (Giora 2003).

When we humans entertain this creative, reactive modality to defend ourselves, we tend to intervene within our mind. Sometimes people even intervene on the physical object itself, the classic example being the poster, writing over it to correct an expression (or even add graphic symbols to images, such as moustaches added to a face). In countercultural art this is called *subvertising*.

As far as current technology is concerned, a lot of attention is being devoted to figuring out how to exploit smartphones for advertising. Here instead we propose a mainly defensive goal on behalf of the consumer. The aim is to exploit technology for producing linguistic expressions that slightly change the observed advertisement. The goal is to accommodate a message that is biased in a rather different direction. The system produces a new virtual poster so as to help the user relax the cognitive tension produced by the unduly attention-grabbing original message.

In particular we have developed a mobile application that allows users to take a picture of a poster, and then automatically produces a new virtual version with the same layout and visual aspect of the poster, but with a creative variation of the linguistic expression it originally expressed.

In our current prototype the user merely needs to point the camera of the smartphone at the poster, and the image, with the same appearance but the altered linguistic expression, is produced in a few short steps.

An image analysis and reconstruction component takes care of the graphic aspects, and an underlying program is called to obtain the actual variation of the given expression, which can have several different realizations. In this paper we utilize just one of the functionalities of VALENTINO, an affective valence shifting program (Guerini, Strapparava and Stock 2011); the creativity involved in the process is a necessary element for the successful impact of this defensive tool.

Background and relevant Work

The word *subvertising* is a portmanteau of the words “subvert” and “advertising”. Subvertising refers to the practice of making spoofs or parodies of corporate and political advertisements in order to make a statement. This can take the form of a new image, an alteration to an existing im-

age, or a modification/re-contextualization of an existing slogan (sometimes called a “meme hack”).

According to AdBusters, a Canadian magazine that is a leading proponent of counter-culture and subvertising, “A well produced 'subvert' mimics the look and feel of the targeted ad, promoting the classic 'double-take' as viewers suddenly realize they have been duped. Subverts create cognitive dissonance.”

In our work we focus on the creative textual modification task of subvertising. In particular, we want to implement a defense strategy for making the user aware of the subtle presuppositions implicit in advertising messages - by using *exaggeration* (or hyperbole) of the affective content of the message. The main resource used to implement such a defensive strategy is the VALENTINO prototype, a tool for affective modification of existing texts.

Affective variations of pre-existing texts have been studied and implemented in various domains, see for example (Mateas, Vanouse and Domike 2000; Guerini, Strapparava and Stock 2008b, 2011), or similarly funny variations (Stock and Strapparava 2003). The effectiveness of affective variations has also been assessed; in particular, Van Der Sluis and Mellish's (2010) evaluation shows that biased variations of a message work better than the neutral condition. With regard to output quality, Whitehead and Cavedon (2010) demonstrated that adding bigram frequencies for the insertion of valenced modifiers (chosen according to MAX function) significantly improve the perceived quality of the resulting texts.

Valentino

VALENTINO can modify existing textual expressions towards more positively or negatively valenced versions, given a numeric coefficient that represents the desired valence shifting for the final expression.

Since the system works in an open domain and without lexical restrictions, VALENTINO's linguistic resources are general purpose, and automatically built from large scale corpora and English lexical repositories.

For the task of modifying single words, we automatically built a resource that gathers these terms in vectors (Ordered Vectors of Valenced Terms - OVVTs). We used the WordNet *antonymy* relation as an indicator of terms that can be “graded”, and built four groups of terms that can be used (one group for each POS). Moreover, we populated the vectors using other specific WordNet semantic relations (the *similar_to* relation for adjectives, *hyponym* relation for verbs and nouns). Finally the valence of WordNet synsets, taken from SentiWordNet scores (Esuli and Sebastiani 2006), was added to the corresponding lemmata. An example OVVT for the antonymy pair (ugly ↔ beautiful) ordered from most negative to most positive is:

(hideous ... ugly ... unnatural) ↔ (pretty ... beautiful ... gorgeous)

For insertion or deletion of words that play the role of downtoners or intensifiers we created specific OVVTs (which we call Modifier-OVVTs). In this case the words were gathered according to a criterion of contextual, rather

than semantic, connection: we used the Google Web 1T 5-Grams Corpus (Brants and Franz 2006) to extract information about co-occurrences. In particular we created resources connecting terms with their modifiers (according to POS), thus obtaining adjective modifiers for nouns, and adverb modifiers both for adjectives and verbs. An example Modifier-OVVT for the term “dish”, ordered from most negative to most positive, is:

(disgusting ... mediocre ... tasty ... delicious ... exquisite).

Strategies

We undertook a preliminary qualitative study with human subjects, to understand how people modify the valence of existing texts. The insights gained showed that: (a) people usually modify single words, (b) sometimes add or subtract words that play the role of downtoners or intensifiers and (c) sometimes use paraphrases (Guerini, Strapparava and Stock 2008b).

As a first step VALENTINO performs POS tagging, named entity recognition, morphological analysis and chunking of the existing constituents (NPs, VPs, ADJPs, and so on). This task exploits the TextPro package (Pianta, Girardi and Zanolini 2008). Subsequently the strategies described in points a), b), and c) above are applied to the chunks, following some general guidelines.

Minimal variation: texts (chunks) are slanted as much as needed, but the target score should not be exceeded, limiting the variation as much as possible.

Modification of dependents: A constituent is modified considering first the dependents (from left to right) and then possibly the head. Consider the very positive and the slightly negative variations of the following sentence:

“ We ate [a very good dish]_{NP}”
“We ate [an incredibly delicious dish]_{NP}” (+)
“We ate [a good dish]_{NP}” (-)

The rationale is that in a constituent the element that bears the greatest part of the meaning is in the head, and it decreases the further we move into the constituent.

Candidates Selection: The selection of substitutes is a two-step process. Given a term to be modified (e.g. “good” in the example) there can be various candidates for the modification.

- The first step requires filtering out all the terms that do not meet the target score. For example if the target score is higher than +0.5, all terms from -1 to +0.5 are discarded. Further possible constraints can be taken into account (e.g. if the reasoning is about “good dish” then only the *similar_to* “good” that co-occur with “dish”, and with score > 0.5, should be kept).
- Various strategies can then be used for choosing the best candidate: word persuasive impact (Guerini, Strapparava and Stock 2008a) word or n-gram frequency (Whitehead and Cavedon 2010), mutual-information, etc. Currently, the most used measure in VALENTINO is pointwise mutual information score, which yields modifiers specialized for the given term (e.g. “delicious” co-occurs less fre-

quently than “nice” with “dish”, but it is more specialized in this context).

As for metrics that help decide the best quality lexical choice, while we have converged so far on the *best mutual information measure*, we think in different situations different measures should be applied (although this is outside the scope of this paper). Furthermore specific n-gram patterns – see for example (Veale 2011) - for extracting *semantically* exaggerated variations are under development.

In the present scenario, the critical choice was deciding the suitable degree of the *affective* modification amongst those proposed by VALENTINO; i.e. which one is the best for obtaining a defensive effect. In fact, light modifications usually obtain the effect of strengthening the message, while stronger ones can weaken it (Guerini, Strapparava and Stock 2012). Obviously strengthening the message is not the aim of the present tool, which is why we chose maximum target scores for the *affective exaggeration* strategy in subvertising.



Figure 1



Figure 2

Interface

We have implemented SUBVERTISER, a mobile subvertising tool that allows a user to photograph an advertisement they see, select the text they wish to change. The system then replaces that text with a valenced version created by VALENTINO. SUBVERTISER tries to match the font face, size and color of the new slogan to the text in the original image in order to heighten the effect of presenting a message that subverts the original.

SUBVERTISER requires very little interaction on the part of the user: once a photo of a printed ad (e.g., poster, billboard or banner) is taken with the phone, only a few steps are required to swap the original advertising message with

a valenced version: selecting the region of the text in the photo to replace, correcting the text after scanning by in-built OCR, and selecting his or her preferred new version from of a list suggested by VALENTINO.

In a typical scenario, the user is walking with friends in a city, perhaps shopping or going to see a movie. When he is interrupted by a poster advertisement that bothers him, he uses his phone to take a snapshot of it, modifies it with SUBVERTISER, and can then show the new ad to his friends.

Algorithm

Behind the scenes, SUBVERTISER performs a number of steps to process both the language and the image of the advertisement. Given the photo taken by the user with the phone's camera (the image can also be chosen from a pre-existing library of images), the user selects the text region he wants to change containing the advertising message by moving and resizing a selection rectangle (Figure 1).

The image area is then passed to an OCR application on the smartphone itself, which scans for text within that rectangle¹. The OCR both detects the coordinates of the bounding boxes for every individual word as well as returns the recognized text string of the message.

From the bounding box information we obtain the rectangle containing the first line of text, which is then scaled down to 100 pixels in height, and uploaded to an online third-party (multi-step) font recognition service² using dedicated APIs. Meanwhile, the program applies an inpainting algorithm to each bounding box in the original text zone. This step reconstructs the background image that was underneath the original text, providing a blank background where new text can be written (Figure 2).

The user is then asked to correct OCR errors, which if left unchanged would lead to linguistic errors in the valenced text, via a text entry box on the smartphone. VALENTINO is queried with the corrected OCR text string, and four valenced sentences are returned, from the most positive to the most negative, and presented to the user (Figure 3) to choose from. Once we know the original text, we also send that information to the font recognition server, which needs to align known letters with the image in order to determine the font and then respond with that information.

Once the user selects one of the slanted messages, an algorithm decides how to divide the slanted text into lines, since VALENTINO typically changes the number of words in the sentence.

Then, we ask the online font service to generate a new image with the detected font and a transparent background,

¹ As much processing as possible is done directly on the mobile to avoid excessive bandwidth usage and associated costs, and to lower the needed time to complete the task. Image processing is done with the OpenCV library, while character recognition is provided by the Tesseract OCR engine, which are both open source.

² URL: www.myfonts.com/WhatTheFont/

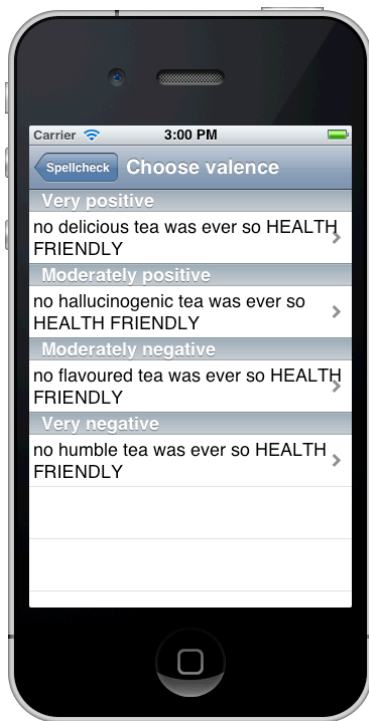


Figure 3

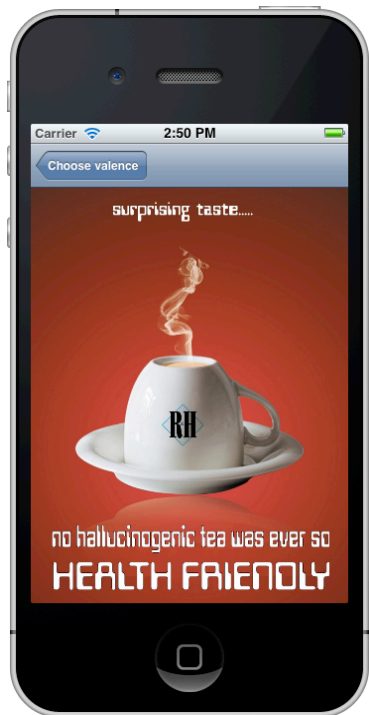


Figure 4

since we often may not have access to the original font due to limitations of the smartphone, and the image containing the new text is downloaded to the phone.

We next identify the original text color by looking at the first line of text in the original image, and treat that as the color for the whole text (even if the ad is written in multiple colors), to save processing time. This identification is performed by clustering the colors of the pixels in two groups with k-means, then considering the two means as colors of text and background.

Finally the new message is copied inside the bounding boxes of the original text. The image is shown onscreen (Figure 4) and the user can save it to the image library or share it by mail or MMS.

SUBVERTISER is currently implemented on the iPhone. An Android version is under development. The only external resources needed are the font recognition service and VALENTINO's server.

Further work

From the technical, NLP processing point of view, the basic capabilities of VALENTINO are currently being expanded. First, we are now starting to take into account the sentence structure, so that it can be better focused. An analysis of rhetorical aspects of a given sentence will also benefit the quality of the intervention. As for metrics that help decide the best quality for lexical choice, we have converged so far on the best mutual information measure (see above), although we think in different situations different measures should be applied.

More generally we can note that an extended VALENTINO could be parameterized to achieve different goals, including: a) generic valence shifting; b) focused biased language to influence the audience's view on one element (e.g., a human or a thing) in the sentence; c) "cleansing" of biases present in the original expression; or d) special effects, such as ironic or hyperbolic reconstructions.

As for point b, specifically for evaluative expressions, subjective evaluations can be along different dimensions: the ethical aspect (related to moral values), the epistemological aspect (related to truth), aesthetics (related to beauty or pleasure), and the utilitarian (related to utility, resources, results), and can be especially reflected in the lexical choice.

The aim is also to link a set of preferences and information about the context to the system. For instance audience preferences can shift its behavior in line with the user's attitude; also independent preferences (e.g. originating from a social institution) might produce expressions that could influence the audience towards a specific direction.

As for the overall mobile application we have described, some technical improvements, like automatic spell checking of the OCR output, would enhance the app's usability.

We would like to mention that additional uses can be envisaged, apart from defense against unwanted advertising expressions. In a sophisticated but not unusual twist of fate, the same technology can be used by the advertisers themselves. A new form of promotion could be based on an active role on the part of their target, which, by adding a

creative variation, contributes to the reinforcement of the basic advertising goal.

Another prospect is in an artistic direction. For instance the mobile application could be monitored on a large display by a crowd at an exhibition, where the audience could see posters in the city being continuously and automatically changed by different individuals walking around with their smartphones, so as to counter the messages on the walls and introducing a collective sense of liberation.

Another setting is with mobile games, where the user may interact with existing linguistic expressions to produce anagrams, wordplay and so on.

Conclusions

Creativity is widely used in advertising, which must appeal to people of all walks of life in every imaginable situation. But advertising also tries to change people's actions, beliefs and behavior, which they rightfully resist as an invasion of their time and attention. The resulting conflict leads to increasingly pervasive, aggressive and frequent advertisements on one hand, while on the other to a conscious refusal to pay attention to those ads or a profaning transformation of the message. Inspired by the latter, a system, even if just based on some degree of combinational creativity (Boden 2009) can aid people in defending themselves against elements in their environment.

We thus built SUBVERTISER, a creative subvertising system, which assists consumers in proactively “taking back” their daily outdoor routine. SUBVERTISER allows consumers to use the power of satire and virtual profaning to push back at advertisers. By combining the utility and pervasiveness of smartphones with the capability of the VALENTINO affective valencing system, consumers can take a picture of an advertisement in public that offends them, select the wording they want to change, use VALENTINO to supply them with language variations that subverts the intended message, and then modify the advertisement with their chosen variations to look just like the original. By sending the new version to their friends, they can join in a collective release of tension from the perpetual barrage of advertisements.

References

Boden, M. 2009. Computer Models of Creativity. *AI Magazine*, 30(3):23-34.

Brants, T., and Franz, A. 2006. Web 1T 5-Gram Version 1. *Linguistic Data Consortium*, Philadelphia.

Burke, M., Gorman, N., Nilsen, E., and Hornof, A. 2004. Banner Ads Hinder Visual Search and Are Forgotten. In *Proceedings of CHI 2004*, 1139-1142.

Esuli, A., and Sebastiani, F. 2006. SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining. In *Proceedings of LREC 2006*, 417-422.

Freud S. 1905. *Der Witz und Seine Beziehung zum Unbewussten*. Deutike, Vienna.

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

Guerini, M., Strapparava, C., and Stock, O. 2012. Ecological Evaluation of Persuasive Messages Using Google AdWords. In *Proceedings of ACL 2012*.

Guerini, M., Strapparava, C., and Stock, O. 2011. Slanting Existing Text with Valentino. In *Proceedings of IUI 2011*, 439-440.

Guerini, M., Strapparava, C., and Stock, O. 2008a. CORPS: A Corpus of Tagged Political Speeches for Persuasive Communication Processing. *Journal of Information Technology & Politics*, 5(1):19-32.

Guerini M., Strapparava C., and Stock O. 2008b. Valentino: A tool for Valence Shifting of Natural Language Texts, In *Proceedings of LREC 2008*, 243-246.

Mateas, M., Vanouse, P., and Domike, S. 2000. Generation of Ideologically-Biased Historical Documentaries. In *Proceedings of AAAI 2000*, 36-42.

Müller, J., Alt, F., and Michelis, D. (eds.). 2011. *Pervasive Advertising*. HCI Series - Springer.

Pagendarm, M., and Schaumburg, H. 2001. Why are users banner-blind? The impact of navigation style on the perception of web banners. *Journal of Digital Information*, 2(1).

Pianta, E., Girardi, C., and Zanoli, R. 2008. The TextPro tool suite, In *Proceedings of LREC 2008*, 2603-2607.

Stock, O., and Strapparava, C. 2003. Getting Serious about the Development of Computational Humour. In *Proceedings of IJCAI 2003*, 59-64.

Veale T. 2011. Creative Language Retrieval: A Robust Hybrid of Information Retrieval and Linguistic Creativity. In *Proceedings of ACL 2011*, 278-287.

Van der Sluis, I., and Mellish, C. 2009. Towards empirical evaluation of affective tactical NLG, in *Proceedings of ENLG 2009*, 146-153.

Whitehead, S., and Cavedon, L. 2010. Generating Shifting Sentiment for a Conversational Agent, in *Proceedings of NAACL HLT 2010*, 89-97.