Emotion Analysis Using Latent Affective Folding and Embedding

Jerome R. Bellegarda

Speech & Language Technologies
Apple Inc.
Cupertino, California 95014, USA
jerome @ apple.com

Abstract

Though data-driven in nature, emotion analysis based on latent semantic analysis still relies on some measure of expert knowledge in order to isolate the emotional keywords or keysets necessary to the construction of affective categories. This makes it vulnerable to any discrepancy between the ensuing taxonomy of affective states and the underlying domain of discourse. This paper proposes a more general strategy which leverages two distincts semantic levels, one that encapsulates the foundations of the domain considered, and one that specifically accounts for the overall affective fabric of the language. Exposing the emergent relationship between these two levels advantageously informs the emotion classification process. Empirical evidence suggests that this is a promising solution for automatic emotion detection in text.

1 Introduction

The automatic detection of emotions in text is a necessary pre-processing step in many different fields touching on affective computing (Picard, 1997), such as natural language interfaces (Cosatto et al., 2003), e-learning environments (Ryan et al., 2000), educational or entertainment games (Pivec and Kearney, 2007), opinion mining and sentiment analysis (Pang and Lee, 2008), humor recognition (Mihalcea and Strapparava, 2006), and security informatics (Abbasi, 2007). In the latter case, for example, it can be used for monitoring levels of hateful or violent rhetoric (perhaps in multilingual settings). More generally, emotion detection is of great interest in human-computer interaction: if a system determines that a user is upset or annoyed, for instance, it could switch to a different mode of interaction (Liscombe et al., 2005). And of course, it plays a critical role in the generation of expressive synthetic speech (Schröder, 2006).

Emphasis has traditionally been placed on the set of six "universal" emotions (Ekman, 1993): ANGER, DISGUST, FEAR, JOY, SADNESS, and SURPRISE (Alm et al., 2005; Liu et al., 2003; Subasic and Huettner, 2001). Emotion analysis is typically carried out using a simplified description of emotional states in a low-dimensional space, which normally comprises dimensions such as valence (positive/negative evalution), activation (stimulation of activity), and/or control (dominant/submissive power) (Mehrabian, 1995; Russell, 1980; Strapparava and Mihalcea, 2008). Classification proceeds based on an underlying emotional knowledge base, which strives to provide adequate distinctions between different emotions. This affective information can either be built entirely upon manually selected vocabulary as in (Whissell, 1989), or derived automatically from data based on expert knowledge of the most relevant features that can be extracted from the input text (Alm et al., 2005). In both cases, the resulting system tends to rely, for the most part, on a few thousand annotated "emotional keywords," the presence of which triggers the associated emotional label(s).

The drawback of such confined lexical affinity is that the analysis tends to be hampered by the bias inherent in the underlying taxonomy of emotional states. Because this taxonomy only supports simplified relationships between affective words and emo-

Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, pages 1–9,

tional categories, it often fails to meaningfully generalize beyond the relatively few core terms explicitly considered in its construction. This has sparked interest in data-driven approaches based on latent semantic analysis (LSA), a paradigm originally developed for information retrieval (Deerwester et al., 1990). Upon suitable training using a large corpus of texts, LSA allows a similarity score to be computed between generic terms and affective categories (Strapparava et al., 2006). This way, every word can automatically be assigned some fractional affective influence. Still, the affective categories themselves are usually specified with the help of a reference lexical database like WordNet (Fellbaum, 1998).

The purpose of this paper is to more broadly leverage the principle of latent semantics in emotion analysis. We cast the problem as a general application of *latent semantic mapping* (LSM), an extrapolation of LSA for modeling global relationships implicit in large volumes of data (Bellegarda, 2005; Bellegarda, 2008). More specifically, we use the LSM framework to describe two distinct semantic levels: one that encapsulates the foundations of the domain considered (e.g., broadcast news, email messages, SMS conversations, etc.), and one that specifically accounts for the overall affective fabric of the language. Then, we leverage these two descriptions to appropriately relate domain and affective levels, and thereby inform the emotion classification process. This *de facto* bypasses the need for any explicit external knowledge.

The paper is organized as follows. The next section provides some motivation for, and gives an overview of, the proposed latent affective framework. In Sections 3 and 4, we describe the two main alternatives considered, latent folding and latent embedding. In Section 5, we discuss the mechanics of emotion detection based on such latent affective processing. Finally, Section 6 reports the outcome of experimental evaluations conducted on the "Affective Text" portion of the SemEval-2007 corpus (Strapparava and Mihalcea, 2007).

2 Motivation and Overview

As alluded to above, lexical affinity alone fails to provide sufficient distinction between different emotions, in large part because only relatively few words have inherently clear, unambiguous emotional meaning. For example, *happy* and *sad* encapsulate JOY and SADNESS, respectively, in all conceivable scenarios. But is *thrilling* a marker of JOY or SURPRISE? Does *awful* capture SADNESS or DISGUST? It largely depends on contextual information: *thrilling* as a synonym for *uplifting* conveys JOY (as in *a thrilling speech*), while *thrilling* as a synonym for *amazing* may well mark SURPRISE (as in *a thrilling waterfall ride*); similarly, *awful* as a synonym for *grave* reflects SADNESS (as in *an awful car accident*), while *awful* as a synonym for *foul* is closer to DISGUST (as in *an awful smell*). The vast majority of words likewise carry multiple potential emotional connotations, with the degree of affective polysemy tightly linked to the granularity selected for the underlying taxonomy of emotions.

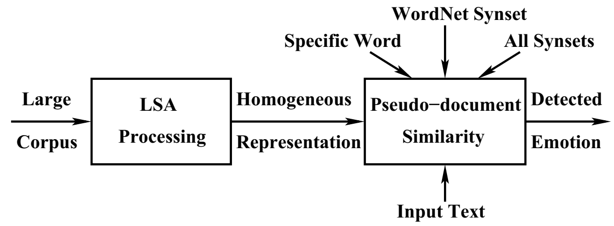

Figure 1: *Typical LSA-Based Emotion Analysis.*

Data-driven approaches based on LSA purport to "individuate" such indirect affective words via inference mechanisms automatically derived in an unsupervised way from a large corpus of texts, such as the British National Corpus (Strapparava et al., 2006). By looking at document-level cooccurrences, contextual information is exploited to encapsulate semantic information into a relatively low dimensional vector space. Suitable affective categories are then constructed in that space by "folding in" either the specific word denoting the emotion, or its associated synset (say, from WordNet), or even the entire set of words in all synsets that can be labelled with that emotion (Strapparava and Mihalcea, 2008). This is typically done by placing the relevant word(s) into a "pseudo-document," and map it into the space as if it were a real one (Deerwester et al., 1990). Finally, the global emotional affinity of a given input text is determined by computing similarities between all pseudo-documents. The resulting framework is depicted in Fig. 1.

This solution is attractive, if for no other reason than it allows every word to automatically be assigned some fractional affective influence. However, it suffers from two limitations which may well prove deleterious in practical situations. First, the inherent lack of supervision routinely leads to a latent semantic space which is not particularly representative of the underlying domain of discourse. And second, the construction of the affective categories still relies heavily on pre-defined lexical affinity, potentially resulting in an unwarranted bias in the taxonomy of affective states.

The first limitation impinges on the effectiveness of any LSA-based approach, which is known to vary substantially based on the size and quality of the training data (Bellegarda, 2008; Mohler and Mihalcea, 2009). In the present case, any discrepancy between latent semantic space and domain of discourse may distort the position of certain words in the space, which could in turn lead to subsequent sub-optimal affective weight assignment. For instance, in the examples above, the word *smell* is considerably more critical to the resolution of *awful* as a marker of DISGUST than the word *car*. But that fact may never be uncovered if the only pertinent documents in the training corpus happen to be about expensive fragrances and automobiles. Thus, it is highly desirable to derive the latent semantic space using data representative of the application considered. This points to a modicum of supervision.

The second limitation is tied to the difficulty of coming up with an *a priori* affective description that will work universally. Stipulating the affective categories using only the specific word denoting the emotion is likely to be less robust than using the set of words in all synsets labelled with that emotion. On the other hand, the latter may well expose some inherent ambiguities resulting from affective polysemy. This is compounded by the relatively small number of words for which an affective distribution is even available. For example, the well-known General Inquirer content analysis system (Stone, 1997) lists only about 2000 words with positive outlook and 2000 words with negative outlook. There are exactly 1281 words inventoried in the affective extension of WordNet (Strapparava and Mihalcea, 2008), and the affective word list from (Johnson–Laird and Oatley, 1989) comprises less than 1000 words. This considerably complicates the construction of reliable affective categories in the latent space.

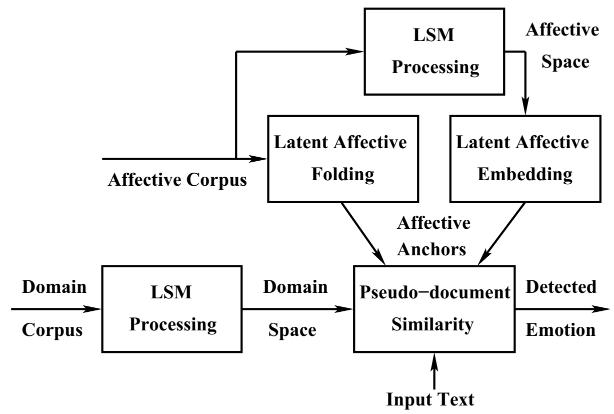

Figure 2: *Proposed Latent Affective Framework.*

To address the two limitations above, we propose to more broadly leverage the LSM paradigm (Bellegarda, 2005; Bellegarda, 2008), following the overall framework depicted in Fig. 2. Compared to Fig. 1, we inject some supervision at two separate levels: not only regarding the particular domain considered, but also how the affective categories themselves are defined. The first task is to exploit a suitable training collection to encapsulate into a (domain) latent semantic space the general foundations of the domain at hand. Next, we leverage a separate affective corpus, such as mood-annotated blog entries from LiveJournal.com (Strapparava and Mihalcea, 2008), to serve as a descriptive blueprint for the construction of affective categories.

This blueprint is then folded into the domain space in one of two ways. The easiest approach, called latent affective folding, is simply to superimpose *affective anchors* inferred in the space for every affective category. This is largely analogous to what happens in Fig. 1, with a crucial difference regarding the representation of affective categories: in latent affective folding, it is derived from a corpus of texts as opposed to a pre-specified keyword or keyset. This is likely to help making the categories more robust, but may not satisfactorily resolve subtle distinctions between emotional connotations. This technique is described in detail in the next section.

The second approach, called latent affective embedding, is to extract a distinct LSM representation

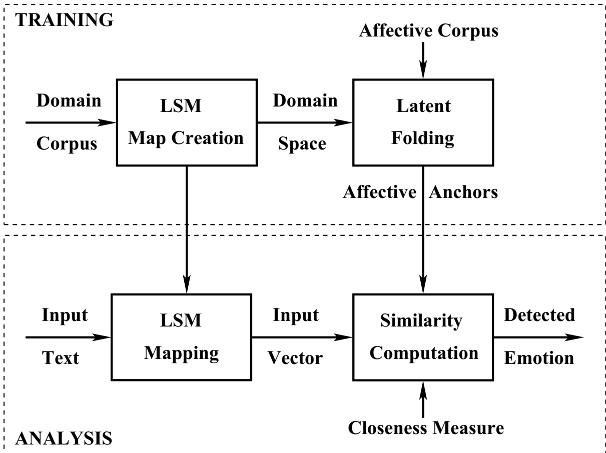

Figure 3: *Emotion Analysis Using Latent Folding.*

from the affective corpus, to encapsulate all prior affective information into a separate (affective) latent semantic space. In this space, affective anchors can be computed directly, instead of inferred after folding, presumably leading to a more accurate positioning. Domain and affective LSM spaces can then be related to each other via a mapping derived from words that are common to both. This way, the affective anchors can be precisely embedded into the domain space. This technique is described in detail in Section 4.

In both cases, the input text is mapped into the domain space as before. Emotion classification then follows from assessing how closely it aligns with each affective anchor.

3 Latent Affective Folding

Expanding the basic framework of Fig. 2 to take into account the two separate phases of training and analysis, latent affective folding proceeds as illustrated in Fig. 3.

Let $\mathcal{T}_1$, $|\mathcal{T}_1| = N_1$, be a collection of training texts (be they sentences, paragraphs, or documents) reflecting the domain of interest, and $\mathcal{V}_1$, $|\mathcal{V}_1| = M_1$, the associated set of all words (possibly augmented with some strategic word pairs, triplets, etc., as appropriate) observed in this collection. Generally, M_1 is on the order of several tens of thousands, while N_1 may be as high as a million.

We first construct a $(M_1 \times N_1)$ matrix W_1, whose elements w_{ij} suitably reflect the extent to which each word $w_i \in \mathcal{V}_1$ appeared in each text $t_j \in \mathcal{T}_1$. From (Bellegarda, 2008), a reasonable expression for w_{ij} is:

$$w_{i,j} = (1 - \varepsilon_i)\,\frac{c_{i,j}}{n_j}\,, \qquad (1)$$

where $c_{i,j}$ is the number of times w_i occurs in text t_j, n_j is the total number of words present in this text, and ε_i is the normalized entropy of w_i in $\mathcal{V}_1$. The global weighting implied by $1 - \varepsilon_i$ reflects the fact that two words appearing with the same count in a particular text do not necessarily convey the same amount of information; this is subordinated to the distribution of words in the entire set $\mathcal{V}_1$.

We then perform a singular value decomposition (SVD) of W_1 as (Bellegarda, 2008):

$$W_1 = U_1\, S_1\, V_1^T\,, \qquad (2)$$

where U_1 is the $(M_1 \times R_1)$ left singular matrix with row vectors $u_{1,i}$ $(1 \leq i \leq M_1)$, S_1 is the $(R_1 \times R_1)$ diagonal matrix of singular values $s_{1,1} \geq s_{1,2} \geq \ldots \geq s_{1,R_1} > 0$, V_1 is the $(N_1 \times R_1)$ right singular matrix with row vectors $v_{1,j}$ $(1 \leq j \leq N_1)$, $R_1 \ll M_1, N_1$ is the order of the decomposition, and ^{T} denotes matrix transposition.

As is well known, both left and right singular matrices U_1 and V_1 are column-orthonormal, i.e., $U_1^T U_1 = V_1^T V_1 = I_{R_1}$ (the identity matrix of order R_1). Thus, the column vectors of U_1 and V_1 each define an orthornormal basis for the space of dimension R_1 spanned by the $u_{1,i}$'s and $v_{1,j}$'s. We refer to this space as the *latent semantic space* $\mathcal{L}_1$. The (rank-R_1) decomposition (2) encapsulates a mapping between the set of words w_i and texts t_j and (after apropriate scaling by the singular values) the set of R_1-dimensional vectors $y_{1,i} = u_{1,i}S_1$ and $z_{1,j} = v_{1,j}S_1$.

The basic idea behind (2) is that the rank-R_1 decomposition captures the major structural associations in W_1 and ignores higher order effects. Hence, the relative positions of the input words in the space $\mathcal{L}_1$ reflect a parsimonious encoding of the semantic concepts used in the domain considered. This means that any new text mapped onto a vector "close" (in some suitable metric) to a particular set of words can be expected to be closely related to the concept encapsulated by this set. If each of these words is then scored in terms of their affective affinity, this offers a way to automatically predict the overall emotional affinity of the text.

In order to do so, we need to isolate regions in that space which are representative of the underlying taxonomy of emotions considered. The centroid of each such region is the *affective anchor* associated with that basic emotion. Affective anchors are superimposed onto the space $\mathcal{L}_1$ on the basis of the affective corpus available.

Let $\mathcal{T}_2$, $|\mathcal{T}_2| = N_2$, represent a separate collection of mood-annotated texts (again they could be sentences, paragraphs, or documents), representative of the desired categories of emotions (such as JOY and SADNESS), and $\mathcal{V}_2$, $|\mathcal{V}_2| = M_2$, the associated set of words or expressions observed in this collection. As such affective data may be more difficult to gather than regular texts (especially in annotated form), in practice $N_2 < N_1$.

Further let $\mathcal{V}_{12}$, $|\mathcal{V}_{12}| = M_{12}$, represent the intersection between $\mathcal{V}_1$ and $\mathcal{V}_2$. We will denote the representations of these words in $\mathcal{L}_1$ by $\lambda_{1,k}$ ($1 \leq k \leq M_{12}$).

Clearly, it is possible to form, for each $1 \leq \ell \leq L$, where L is the number of distinct emotions considered, each subset $\mathcal{V}_{12}^{(\ell)}$ of all entries from $\mathcal{V}_{12}$ which is aligned with a particular emotion.[1] We can then compute:

$$\hat{z}_{1,\ell} = \frac{1}{|\mathcal{V}_{12}^{(\ell)}|} \sum_{\mathcal{V}_{12}^{(\ell)}} \lambda_{1,k} \, , \tag{3}$$

as the affective anchor of emotion ℓ ($1 \leq \ell \leq L$) in the domain space. The notation $\hat{z}_{1,\ell}$ is chosen to underscore the connection with $z_{1,j}$: in essence, $\hat{z}_{1,\ell}$ represents the (fictitious) text in the domain space that would be perfectly aligned with emotion ℓ, had it been seen the training collection $\mathcal{T}_1$. Comparing the representation of an input text to each of these anchors therefore leads to a quantitative assessment for the overall emotional affinity of the text.

A potential drawback of this approach is that (3) is patently sensitive to the distribution of words within $\mathcal{T}_2$, which may be quite different from the distribution of words within $\mathcal{T}_1$. In such a case, "folding in" the affective anchors as described above may well introduce a bias in the position of the anchors in the domain space. This could in turn lead to an inability to satisfactorily resolve subtle distinctions between emotional connotations.

[1]Note that one entry could conceivably contribute to several such subsets.

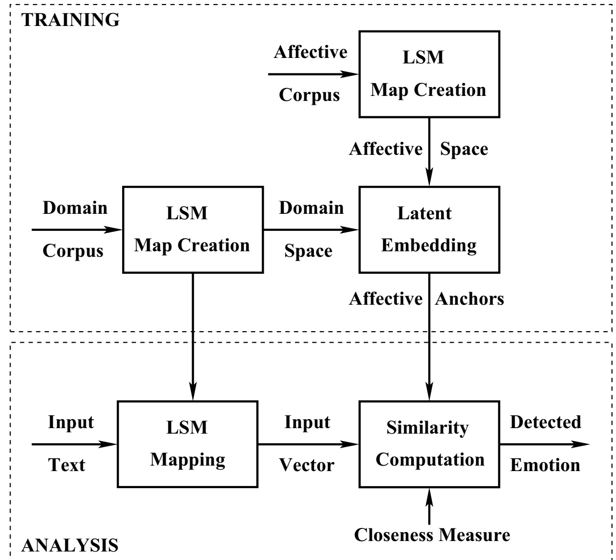

Figure 4: *Emotion Analysis Using Latent Embedding.*

4 Latent Affective Embedding

To remedy this situation, a natural solution is to build a separate LSM space from the affective training data. Referring back to the basic framework of Fig. 2 and taking into account the two separate phases of training and analysis as in Fig. 3, latent affective embedding proceeds as illustrated in Fig. 4.

The first task is to group all N_2 documents present in $\mathcal{T}_2$ into L bins, one for each of the emotions considered. Then we can construct a $(M_2 \times L)$ matrix W_2, whose elements $w'_{k,\ell}$ suitably reflect the extent to which each word or expression $w'_k \in \mathcal{V}_2$ appeared in each affective category c_ℓ, $1 \leq \ell \leq L$. This leads to:

$$w'_{k,\ell} = (1 - \varepsilon'_k) \, \frac{c'_{k,\ell}}{n'_\ell} \, , \tag{4}$$

with $c'_{k,\ell}$, n'_ℓ, and ε'_k following definitions analogous to (1), albeit with domain texts replaced by affective categories.

We then perform the SVD of W_2 in a similar vein as (2):

$$W_2 = U_2 \, S_2 \, V_2^T \, , \tag{5}$$

where all definitions are analogous. As before, both left and right singular matrices U_2 and V_2 are column-orthonormal, and their column vectors each define an orthornormal basis for the space of dimension R_2 spanned by the $u_{2,k}$'s and $v_{2,\ell}$'s. We refer to this space as the *latent affective space* $\mathcal{L}_2$. The (rank-R_2) decomposition (5) encapsulates a mapping between the set of words w'_k and categories c_ℓ and (after apropriate scaling by the singular values) the set of R_2-dimensional vectors $y_{2,k} = u_{2,k} S_2$ and $z_{2,\ell} = v_{2,\ell} S_2$.

Thus, each vector $z_{2,\ell}$ can be viewed as the centroid of an emotion in $\mathcal{L}_2$, or, said another way, an affective anchor in the affective space. Since their relative positions reflect a parsimonious encoding of the affective annotations observed in the emotion corpus, these affective anchors now properly take into account any accidental skew in the distribution of words which contribute to them. All that remains to do is map them back to the domain space.

This is done on the basis of words that are common to both the affective space and the domain space, i.e., the words in $\mathcal{V}_{12}$. Since these words were denoted by $\lambda_{1,k}$ in $\mathcal{L}_1$, we similarly denote them by $\lambda_{2,k}$ ($1 \leq k \leq M_{12}$) in $\mathcal{L}_2$.

Now let μ_1, μ_2 and Σ_1, Σ_2 denote the mean vector and covariance matrix for all observations $\lambda_{1,k}$ and $\lambda_{2,k}$ in the two spaces, respectively. We first transform each feature vector as:

$$\bar{\lambda}_{1,k} = \Sigma_1^{-1/2} \left(\lambda_{1,k} - \mu_1 \right), \tag{6}$$

$$\bar{\lambda}_{2,k} = \Sigma_2^{-1/2} \left(\lambda_{2,k} - \mu_2 \right), \tag{7}$$

so that the resulting sets $\{\bar{\lambda}_{1,k}\}$ and $\{\bar{\lambda}_{2,k}\}$ each have zero mean and identity covariance matrix.

For this purpose, the inverse square root of each covariance matrix can be obtained as:

$$\Sigma^{-1/2} = Q \Delta^{-1/2} Q^T, \tag{8}$$

where Q is the eigenvector matrix of the covariance matrix Σ, and Δ is the diagonal matrix of corresponding eigenvalues. This applies to both domain and affective data.

We next relate each vector $\bar{\lambda}_{2,k}$ in the affective space to the corresponding vector $\bar{\lambda}_{1,k}$ in the domain space. For a relative measure of how the two spaces are correlated with each other, as accumulated on a common word basis, we first project $\bar{\lambda}_{1,k}$ into the unit sphere of same dimension as $\bar{\lambda}_{2,k}$, i.e., $R_2 = \min(R_1, R_2)$. We then compute the (normalized) cross-covariance matrix between the two unit sphere representations, specified as:

$$K_{12} = \sum_{k=1}^{M_{12}} P \bar{\lambda}_{1,k} P^T \bar{\lambda}_{2,k}^T, \tag{9}$$

where P is the R_1 to R_2 projection matrix. Note that K_{12} is typically full rank as long as $M_{12} > R_2^2$. Performing the SVD of K_{12} yields the expression:

$$K_{12} = \Phi \, \Omega \, \Psi^T, \tag{10}$$

where as before Ω is the diagonal matrix of singular values, and Φ and Ψ are both unitary in the unit sphere of dimension R_2. This in turn leads to the definition:

$$\Gamma = \Phi \Psi^T, \tag{11}$$

which can be shown (cf. (Bellegarda et al., 1994)) to represent the least squares rotation that must be applied (in that unit sphere) to $\bar{\lambda}_{2,k}$ to obtain an estimate of $P \bar{\lambda}_{1,k} P^T$.

Now what is needed is to apply this transformation to the centroids $z_{2,\ell}$ ($1 \leq \ell \leq L$) of the affective categories in the affective space, so as to map them to the domain space. We first project each vector into the unit sphere, resulting in:

$$\bar{z}_{2,\ell} = \Sigma_2^{-1/2} \left(z_{2,\ell} - \mu_2 \right), \tag{12}$$

as prescribed in (7). We then synthesize from $\bar{z}_{2,\ell}$ a unit sphere vector corresponding to the estimate in the projected domain space. From the foregoing, this estimate is given by:

$$\hat{\bar{z}}_{1,\ell} = \Gamma \, \bar{z}_{2,\ell} \,. \tag{13}$$

Finally, we restore the resulting contribution at the appropriate place in the domain space, by reversing the transformation (6):

$$\hat{z}_{1,\ell} = \Sigma_1^{1/2} \, \hat{\bar{z}}_{1,\ell} \, + \, \mu_1 \,. \tag{14}$$

Combining the three steps (12)–(14) together, the overall mapping can be written as:

$$\hat{z}_{1,\ell} \; = \; (\Sigma_1^{1/2} \Gamma \Sigma_2^{-1/2}) \, z_{2,\ell} + (\mu_1 - \Sigma_1^{1/2} \Gamma \Sigma_2^{-1/2} \mu_2) \,. \tag{15}$$

This expression stipulates how to leverage the *observed* affective anchors $z_{2,\ell}$ in the affective space to obtain an estimate of the *unobserved* affective anchors $\hat{z}_{1,\ell}$ in the domain space, for $1 \leq \ell \leq L$. The overall procedure is illustrated in Fig. 5 (in the simple case of two dimensions).

Once the affective anchors are suitably embedded into the domain space, we proceed as before to compare the representation of a given input text to each of these anchors, which leads to the desired quantitative assessment for the overall emotional affinity of the text.

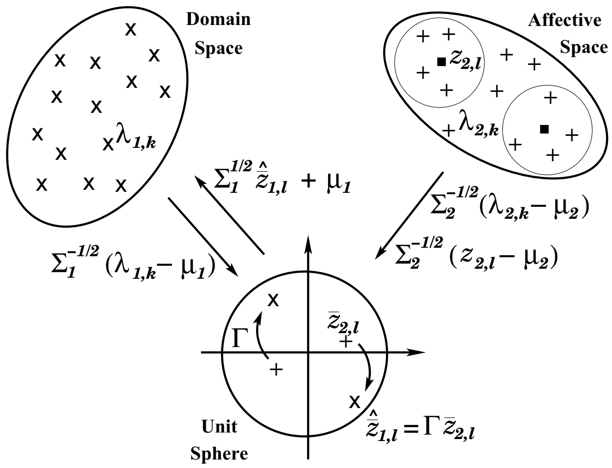

Figure 5: *Affective Anchor Embedding (2-D Case).*

5 Emotion Classification

To summarize, using either latent affective folding or latent affective embedding, we end up with an estimate $\hat{z}_{1,\ell}$ of the affective anchor for each emotion ℓ in the domain space $\mathcal{L}_1$. What remains to be described is how to perform emotion classification in that space.

To proceed, we first need to specify how to represent in that space an input text not seen in the training corpus, say t_p (where $p > N_1$). For each entry in $\mathcal{T}_1$, we compute for the new text the weighted counts (1) with $j = p$. The resulting feature vector, a column vector of dimension N_1, can be thought of as an additional column of the matrix W_1. Assuming the matrices U_1 and S_1 do not change appreciably, the SVD expansion (2) therefore implies:

$$t_p = U_1 \, S_1 \, v_{1,p}^T \,, \tag{16}$$

where the R_1-dimensional vector $v_{1,p}^T$ acts as an additional column of the matrix V_1^T. Thus, the represention of the new text in the domain space can be obtained from $z_{1,p} = v_{1,p} S_1$.

All is needed now is a suitable closeness measure to compare this representation to each affective anchor $\hat{z}_{1,\ell}$ ($1 \leq \ell \leq L$). From (Bellegarda, 2008), a natural metric to consider is the cosine of the angle between them. This yields:

$$\mathcal{C}(z_{1,p}, \hat{z}_{1,\ell}) = \frac{z_{1,p} \, \hat{z}_{1,\ell}^T}{\|z_{1,p}\| \, \|\hat{z}_{1,\ell}\|} \,, \tag{17}$$

for any $1 \leq \ell \leq L$. Using (17), it is a simple matter to directly compute the relevance of the input text to each emotional category. It is important to note that word weighting is now implicitly taken into account by the LSM formalism.

6 Experimental Evaluation

In order to evaluate the latent affective framework described above, we used the data set that was developed for the SemEval 2007 task on "Affective Text" (Strapparava and Mihalcea, 2007). This task was focused on the emotion classification of news headlines. Headlines typically consist of a few words and are often written by creative people with the intention to "provoke" emotions, and consequently attract the readers' attention. These characteristics make this kind of data particularly suitable for use in an automatic emotion recognition setting, as the affective/emotional features (if present) are guaranteed to appear in these short sentences. The test data accordingly consisted of 1,250 short news headlines[2] extracted from news web sites (such as Google news, CNN) and/or newspapers, and annotated along $L = 6$ emotions (ANGER, DISGUST, FEAR, JOY, SADNESS, and SURPRISE) by different evaluators.

For baseline purposes, we considered the following approaches: (i) a simple word accumulation system, which annotates the emotions in a text based on the presence of words from the WordNet-Affect lexicon; and (ii) three LSA-based systems implemented as in Fig. 1, which only differ in the way each emotion is represented in the LSA space: either based on a specific word only (e.g., JOY), or the word plus its WordNet synset, or the word plus all WordNet synsets labelled with that emotion in WordNet-Affect (cf. (Strapparava and Mihalcea, 2007)). In all three cases, the large corpus used for LSA processing was the Wall Street Journal text collection (Graff et al., 1995), comprising about 86,000 articles.

For the latent affective framework, we needed to select two separate training corpora. For the "domain" corpus, we selected a collection of about $N_1 = 8,500$ relatively short English sentences (with a vocabulary of roughly $M_1 = 12,000$ words) originally compiled for the purpose of a building a concatenative text-to-speech voice. Though not

[2]Development data was merged into the original SemEval 2007 test set to produce a larger test set.

Table I: Results on SemEval-2007 Test Corpus.

Approach Considered	Precision	Recall	F-Measure
Baseline Word Accumulation	44.7	2.4	4.6
LSA (Specific Word Only)	11.5	65.8	19.6
LSA (With WordNet Synset)	12.2	77.5	21.1
LSA (With All WordNet Synsets)	11.4	89.6	20.3
Latent Affective Folding	18.8	90.1	31.1
Latent Affective Embedding	20.9	91.7	34.0

completely congruent with news headlines, we felt that the type and range of topics covered was close enough to serve as a good proxy for the domain. For the "affective" corpus, we relied on about $N_2 = 5,000$ mood-annotated blog entries from LiveJournal.com, with a filtered[3] vocabulary of about $M_2 = 20,000$ words. The indication of mood being explicitly specified when posting on LiveJournal, without particular coercion from the interface, mood-annotated posts are likely to reflect the true mood of the blog authors (Strapparava and Mihalcea, 2008). The moods were then mapped to the $L = 6$ emotions considered in the classification.

Next, we formed the domain and affective matrices W_1 and W_2 and processed them as in (2) and (5). We used $R_1 = 100$ for the dimension of the domain space $\mathcal{L}_1$ and $R_2 = L = 6$ for the dimension of the affective space $\mathcal{L}_2$. We then compared latent affective folding and embedding to the above systems. The results are summarized in Table I.

Consistent with the observations in (Strapparava and Mihalcea, 2008), word accumulation secures the highest precision at the cost of the lowest recall, while LSA-based systems achieve high recall but significantly lower precision. Encouragingly, the F-measure obtained with both latent affective mapping techniques is substantially higher than with all four baseline approaches. Of the two techniques, latent embedding performs better, presumably because the embedded affective anchors are less sensitive than the folded affective anchors to the distribution of words within the affective corpus. Both techniques seem to exhibit an improved ability to resolve distinctions between emotional connotations.

[3]Extensive text pre-processing is usually required on blog entries, to address typos and assorted creative license.

7 Conclusion

We have proposed a data-driven strategy for emotion analysis which focuses on two coupled phases: (i) separately encapsulate both the foundations of the domain considered and the overall affective fabric of the language, and (ii) exploit the emergent relationship between these two semantic levels of description in order to inform the emotion classification process. We address (i) by leveraging the latent topicality of two distinct corpora, as uncovered by a global LSM analysis of domain-oriented and emotion-oriented training documents. The two descriptions are then superimposed to produce the desired connection between all terms and emotional categories. Because this connection automatically takes into account the influence of the entire training corpora, it is more encompassing than that based on the relatively few affective terms typically considered in conventional processing.

Empirical evidence gathered on the "Affective Text" portion of the SemEval-2007 corpus (Strapparava and Mihalcea, 2007) shows the effectiveness of the proposed strategy. Classification performance with latent affective embedding is slightly better than with latent affective folding, presumably because of its ability to more richly describe the affective space. Both techniques outperform standard LSA-based approaches, as well as affectively weighted word accumulation. This bodes well for the general deployability of latent affective processing across a wide range of applications.

Future efforts will concentrate on characterizing the influence of the parameters R_1 and R_2 on the vector spaces $\mathcal{L}_1$ and $\mathcal{L}_2$, and the corresponding trade-off between modeling power and generalization properties. It is also of interest to investigate how incorporating higher level units (such as common lexical compounds) into the LSM procedure might further increase performance.

References

A. Abbasi (2007), "Affect Intensity Analysis of Dark Web Forums," in *Proc. IEEE Int. Conf. Intelligence and Security Informatics (ISI)*, New Brunswick, NJ, 282–288.

C. Ovesdotter Alm, D. Roth, and R. Sproat (2005), "Emotions from Text: Machine Learning for Text–Based Emotion Prediction," in *Proc. Conf. Human Language Technology and Empirical Methods in NLP*, Vancouver, BC, 579–586.

J.R. Bellegarda (2005), "Latent Semantic Mapping: A Data–Driven Framework for Modeling Global Relationships Implicit in Large Volumes of Data," *IEEE Signal Processing Magazine*, 22(5):70–80.

J.R. Bellegarda (2008), *Latent Semantic Mapping: Principles & Applications*, Synthesis Lectures on Speech and Audio Processing Series, Fort Collins, CO: Morgan & Claypool.

J.R. Bellegarda, P.V. de Souza, A. Nadas, D. Nahamoo, M.A. Picheny and L.R. Bahl (1994), "The Metamorphic Algorithm: A Speaker Mapping Approach to Data Augmentation," *IEEE Trans. Speech and Audio Processing*, 2(3):413–420.

E. Cosatto, J. Ostermann, H.P. Graf, and J. Schroeter (2003), "Lifelike talking faces for interactive services," in *Proc. IEEE*, 91(9), 1406–1429.

S. Deerwester, S.T. Dumais, G.W. Furnas, T.K. Landauer, and R. Harshman (1990), "Indexing by Latent Semantic Analysis," *J. Amer. Soc. Information Science*, 41:391–407.

P. Ekman (1993), "Facial Expression and Emotion", *American Psychologist*, 48(4), 384–392.

C. Fellbaum, Ed., (1998), *WordNet: An Electronic Lexical Database*, Cambridge, MA: MIT Press.

D. Graff, R. Rosenfeld, and D. Paul (1995), "CSR-III Text," Linguistic Data Consortium, #LDC95T6.

P. Johnson–Laird and K. Oatley (1989), "The Language of Emotions: An Analysis of a Semantic Field," *Cognition and Emotion*, 3:81–123.

J. Liscombe, G. Riccardi, and D. Hakkani-Tür (2005), "Using Context to Improve Emotion Detection in Spoken Dialog Systems," *Proc. Interspeech*, Lisbon, Portugal, 1845–1848.

H. Liu, H. Lieberman, and T. Selker (2003), "A Model of Textual Affect Sensing Using Real-World Knowledge," in *Proc. Intelligent User Interfaces (IUI)*, Miami, FL, 125–132.

A. Mehrabian (1995), "Framework for a Comprehensive Description and Measurement of Emotional States," *Genetic, Social, and General Psychology Monographs*, 121(3):339–361.

R. Mihalcea and C. Strapparava (2006), "Learning to Laugh (Automatically): Computational Models for Humor Recognition," *J. Computational Intelligence*, 22(2):126–142.

M. Mohler and R. Mihalcea (2009), "Text-to-text Semantic Similarity for Automatic Short Answer Grading," in *Proc. 12th Conf. European Chap. ACL*, Athens, Greece, 567–575.

B. Pang and L. Lee (2008), "Opinion Mining and Sentiment Analysis," in *Foundations and Trends in Information Retrieval*, 2(1-2):1–135.

R.W. Picard (1997), *Affective Computing*, Cambridge, MA: MIT Press.

M. Pivec and P. Kearney (2007), "Games for Learning and Learning from Games," *Informatica*, 31:419–423.

J.A. Russell (1980), "A Circumplex Model of Affect," *J. Personality and Social Psychology*, 39:1161–1178.

S. Ryan, B. Scott, H. Freeman, and D. Patel (2000), *The Virtual University: The Internet and Resource-based Learning*, London, UK: Kogan Page.

M. Schröder (2006), "Expressing Degree of Activation in Synthetic Speech," *IEEE Trans. Audio, Speech, and Language Processing*, 14(4):1128–1136.

P.J. Stone (1997), "Thematic Text Analysis: New agendas for Analyzing Text Content," in *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences from Texts and Transcripts*, C.W. Roberts, Ed., Mahwah, NJ: Lawrence Erlbaum Assoc. Publishers, 35–54.

C. Strapparava and R. Mihalcea (2007), "SemEval-2007 Task 14: Affective Text," in *Proc. 4th Int. Workshop on Semantic Evaluations (SemEval 2007)*, Prague, Czech Republic.

C. Strapparava and R. Mihalcea (2008), "Learning to Identify Emotions in Text," in *Proc. 2008 ACM Symposium on Applied Computing*, New York, NY, 1556–1560.

C. Strapparava, A. Valitutti, and O. Stock (2006), "The Affective Weight of Lexicon," in *Proc. 5th Int. Conf. Language Resources and Evaluation (LREC)*, Lisbon, Portugal.

P. Subasic and A. Huettner (2001), "Affect Analysis of Text Using Fuzzy Semantic Typing," *IEEE Trans. Fuzzy Systems*, 9(4):483–496.

C.M. Whissell (1989), "The Dictionary of Affect in Language," in *Emotion: Theory, Research, and Experience*, R. Plutchik and H. Kellerman, Eds., New York, NY: Academic Press, 13–131.