

A Cognitive-based Annotation System for Emotion Computing

Ying Chen, Sophia Y. M. Lee and Chu-Ren Huang

Department of Chinese & Bilingual Studies

The Hong Kong Polytechnic University

{chenying3176, sophiaym, churen.huang}@gmail.com

Abstract

Emotion computing is very important for expressive information extraction. In this paper, we provide a robust and versatile emotion annotation scheme based on cognitive emotion theories, which not only can annotate both explicit and implicit emotion expressions, but also can encode different levels of emotion information for the given emotion content. In addition, motivated by a cognitive framework, an automatic emotion annotation system is developed, and large and comparatively high-quality emotion corpora are created for emotion computing, one in Chinese and the other in English. Such an annotation system can be easily adapted for different kinds of emotion applications and be extended to other languages.

1 Introduction

Affective information is important for human language technology, and sentiment analysis, a coarse-grained affective computing (Shanahan et al., 2006), which is attitude assessment, has become the most salient trend. The polarity-driven approach in sentiment analysis is, however, often criticized as too general to satisfy some applications, such as advertisement design and robot design, and one way to capture more fine-grained affective information is to detect emotion expressions. Unlike sentiment, emotions are cognitive-based, which consistently occur across domains because of its human psychological activities. We believe that emotion computing, which is a fine-grained and cognitive-based framework of affective computing, will provide a more robust and versatile model for human language technology.

Since the concept of emotion is very complicated and subjective, comparing to some annotations such as POS annotation and Chinese word segmentation annotation, emotion annotation is highly labor intensive as it requires careful human judgment. Both explicit and implicit emotions must be recognized and tagged during emotion annotation, therefore, emotion annotation is not a simple assignment exercise as in POS annotation. Technically, emotion annotation can be divided into two subtasks: emotion detection (i.e. differentiate emotional content from neutral content), which is a very important task for affective information extraction, and emotion classification (i.e. assign emotion tags to emotional content.)

Emotion computing often requires a large and high-quality annotated data, however, there is a lack of this kind of corpus. This is not only because of the enormous human involvement, but also because of the unavailability of emotion annotation scheme, which is robust and versatile for both emotion annotation and emotion computing. Tokuhisa et al. (2008) is the only work that explores the issue of emotion detection while most of the previous studies concentrate on the emotion classification given a known emotion context (Mihalcea and Liu, 2006; Kozareva et al., 2007.) Even for emotion classification, some issues remain unresolved, such as the complicated relationships among different emotion types, emotion type selection, and so on. Thus, it is still far from solving the emotion problem if emotion annotation is just considered as emotion-tag assignment.

In this paper, we first explore the relationships among different emotion types with the support of a proposed emotion taxonomy, which combines some psychological theories and linguistic semantics. Based on the emotion taxonomy, a robust and versatile emotion annotation scheme is designed and used in both Chinese and English

emotion corpora. Our emotion annotation scheme is very flexible, which is only a layer added to a sentence, although it can easily be extended to a higher level of a text. Our annotation scheme not only can provide the emotion type information, but also can encode the information regarding the relationship between emotions. With this versatile annotated emotion information, different NLP users can extract different emotion information from a given annotated corpus according to their applications.

With such an emotion annotation scheme, a large and comparatively high-quality annotated emotion corpus is built for emotion computing through an unsupervised approach. Tokuhisa et al. (2008) pointed out that besides emotion corpus, neutral corpus (i.e. sentences containing no emotion) is also very important for emotion computing. Therefore, a high-quality neutral corpus is also automatically collected using contextual information. These two corpora are combined to form a complete emotion-driven corpus for emotion computing. Although the unsupervised method cannot provide a perfectly-annotated corpus, it can easily adapt for different emotion computing.

The remainder of this paper is organized as follows. In Section 2, we give an overview of the previous work on emotion annotation and some related psychological and linguistic theories. In Section 3, we describe our emotion taxonomy and emotion annotation scheme. Section 4 discusses how the unsupervised corpus is created. Section 5 presents the pilot experiments for emotion computing with our corpus, which suggests that the unsupervised approach of our corpus creation is effective. Finally, a conclusion is made in Section 5.

2 Related work

There is no clear consensus among many psychological and linguistic theories on the concept of emotions. Here, we limit our work by the classic definition of “emotions” (Cannon, 1927): Emotion is the felt awareness of bodily reactions to something perceived or thought.

Emotion is a complicated concept, and there are complicated relationships among different emotions. For example, the relationship between “discouraged” and “sad” is different with the one between “remorse” and “sad.” Hobbs & Gordon (2008) and Mathieu (2005) explore emotions mainly from a lexical semantics perspective, and Schröder et al. (2006) designed an annotation

scheme, EARL, mainly for speech processing. Because of the disagreements in emotion theories, EARL did not explore the relationships among emotion types. In this paper, we focus on emotions in written data, which is very different from that of in spoken data in terms of expressions. Here, we first adopt psychological theories (Plutchik, 1980; Turner, 2000) to create an emotion taxonomy, and then design an emotion annotation scheme based on the taxonomy.

Since most of the previous emotion corpora are either too small (Xu et al., 2008) or comparatively ineffective in terms of accuracy (Tokuhisa et al., 2008), they cannot satisfy the requirements of emotion computing. In this paper, based on Natural Semantic Metalanguage (NSM), a cognitive approach to human emotions (which will be discussed in the later section), we create an automatic emotion annotation system. While this annotation system needs only a little training data and does not require human supervision, the corpus still maintains a comparatively high quality. Another significant advantage of our automatic annotation system is that it can easily adapt to different emotion applications by simply supplying different training data.

Most of the existing emotion theories study emotions from the biological and psychological perspectives, hence they cannot easily apply to NLP. Fortunately, NSM, one of the prominent cognitive models exploring human emotions, offers a comprehensive and practical approach to emotions (Wierbicka 1996.) NSM describes complex and abstract concepts, such as emotions, in terms of simpler and concrete ones. In such a way, emotions are decomposed as complex events involving a cause and a mental state, which can be further described with a set of universal, irreducible cores called semantic primitives. This approach identifies the exact differences and connections between emotion concepts in terms of the causes, which provide an immediate cue for emotion detection and classification. We believe that the NSM model offers a plausible framework to be implemented for automatic emotion computing.

3 Emotion annotation scheme

3.1 The emotion taxonomy

Although there are many emotion theories developed in different fields, such as biology, psychology, and linguistics, most of them agree that emotion can be divided into primary emotions and complex emotions (i.e. the combinations of

primary emotions.) There is still controversy over the selection of primary emotions, nonetheless, “happiness”, “sadness”, “anger”, and “fear” are considered as primary emotions by most of emotion theories.

Plutchik’s emotion taxonomy (Plutchik 1980), one of the classic emotion taxonomies, also follows the division of primary emotions and complex emotions, and Turner’s taxonomy (Turner 2000), which is based on Plutchik’s work, allows more flexible combinations of primary emotions. In this paper, we adopt Turner’s taxonomy, with the two main points emphasized:

1) For each primary emotion, it is divided into three levels according to its intensity: high, moderate, and low. Besides “happiness,” “sadness,” “anger” and “fear,” Turner also suggests that “disgust” and “surprise” can be primary emotions (Turner 1996; Turner 2007). In Chinese, the character “惊” (“surprise”) has a strong ability to form many emotion words, such as 惊喜 (surprise and happiness), and 惊吓 (surprise and fear), which is consistent with the explanation of “surprise” emotion by Plutchik (1991): “when the stimulus has been evaluated, the surprise may quickly change to any other emotion.” Therefore, in our annotation scheme, we consider “happiness,” “sadness,” “anger,” “fear,” and “surprise” as primary emotions.

2) Complex emotion can be divided into first-order complex emotions (consisting of two primary emotions), second-order complex emotions (consisting of three primary emotions), and so on, according to the number of primary emotions that involves in the complex emotion. For example, “pride” (happiness + fear) is a first-order complex emotion, which contains a greater amount of “happiness” with a lesser amount of “fear.”

Tables 1 and 2 show some keywords in Turner’s taxonomy, and the symbol “//” is to separate different emotion types. Table 1 lists the five most common English keywords and their corresponding primary emotions, and Table 2 lists the English keywords and their corresponding complex emotions. In Table 2, several emotion keywords, which express similar emotion meaning, are grouped into an emotion type. For example, the emotion keywords “awe, reverence, veneration” are grouped into emotion type “awe.” For a complex emotion, the order of primary emotions indicates the importance of those primary emotions for that complex emotion. For examples, “envy” is “fear + anger,” which con-

tains a greater amount of “fear” with a lesser amount of “anger” whereas “awe” is “fear + happiness,” which contains a greater amount of “fear” with a lesser amount of “happiness.”

For English emotion keywords, as Turner’s taxonomy missed some common emotion keywords, we add the emotion keywords from Plutchik’s taxonomy. Besides, unlike Chinese, English words have morphological variations, for example, the emotion keyword “pride” can occur in text with the various formats: “pride,” “prides,” “prided,” “proud,” “proudly.” As shown in Tables 1 and 2, there are 188 English lemmas in our taxonomy. In total, there are 720 emotion keywords if morphology is taken into account.

Since Turner’s emotion taxonomy is cognitive-based, it is versatile for different languages although there is no one-to-one mapping. We also explore Chinese emotion taxonomy in our previous work (Chen et al., 2009). We first select emotion keywords from the cognitive-based feeling words listed in Xu and Tao (2003), and then map those emotion keywords to Turner’s taxonomy with adaptation for some cases. Lastly, some polysemous emotion keywords are removed to reduce ambiguity, and 226 Chinese emotion keywords remain.

Moreover, Turner’s taxonomy is a comparatively flexible structure, and more extensions can be done for different applications. For example, for a complex emotion, not only its primary emotions are listed, but also the intensity of the primary emotions can be given. For instance, three emotion types, which belong to “anger + fear,” are extended as follows:

Jealousy: Anger (Moderate) + Fear (Moderate)

Suspicion: Anger (Low) + Fear (Low)

Abhorrence: Anger (High) + Fear (Low)

Finally, we should admit that the emotion taxonomy is still an on-going research topic and needs further exploration, such as the position of a given emotion keyword in the emotion taxonomy, whether and how to group similar emotion keywords, and how to decompose a complex emotion into primary emotions.

3.2 The emotion annotation scheme

Given Turner’s taxonomy, we design our annotation scheme to encode this kind of emotion information. Our emotion annotation scheme is XML scheme, and conforms with the Text Encoding Initiative (TEI) scheme with some modifications. The emotion scheme is a layer just

Primary Emotions	Keywords
Happiness	High: ecstatic, eager, joy, enthusiastic, happy//Moderate: cheerful, satisfy, pleased, enjoy, interest//Low: sanguine, serene, content, grateful
Fear	High: horror, terror//Moderate: misgivings, self-conscious, scare, panic, anxious//Low: bewilder, reluct, shy, puzzles, confuse
Anger	High: dislike, disgust, outrage, furious, hate//Moderate: contentious, offend, frustrate, hostile, angry//Low: contemptuous, agitate, irritate, annoy, impatient
Sadness	High: deject, despondent, sorrow, anguish, despair//Moderate: gloomy, dismay, sad, unhappy, disappointed//Low: dispirit, downcast, discourage
Surprise	High: astonish//Moderate: startled, amaze, surprise

Table1: Primary emotions and some corresponding keywords

Combinations	Keywords
Happiness + Fear	Wonder: wonder, wondering, hopeful//Pride: pride, boastful
Happiness + Anger	Vengeance: vengeance, vengeful//Calm: appeased, calmed, calm, soothed//Bemused: bemused
Happiness + Sadness	Yearning: nostalgia, yearning
Fear + Happiness	Awe: awe, reverence, veneration
Fear + Anger	Antagonism: antagonism, revulsed//Envy: envy
Fear + Sadness	Worried: dread, wariness, pensive, helpless, apprehension, worried
Anger +Happiness	Unfriendly: snubbing, mollified, rudeness, placated, apathetic, unsympathetic, unfriendly, unaffectionate//Sarcastic: sarcastic
Anger + Fear	Jealousy: jealous//Suspicion: suspicion, distrustful//Abhorrence: abhorrence
Anger + Sadness	Depressed: bitter, depression//Intolerant: intolerant
Sadness +Happiness	Acceptance: acceptance, tolerant//Solace: moroseness, solace, melancholy
Sadness+ Fear	Hopeless: forlorn, lonely, hopeless, miserable//Remorseful: remorseful, ashamed, humiliated
Sadness+ Anger	Discontent: aggrieved, discontent, dissatisfied, unfulfilled//Boredom: boredom//Grief: grief, sullenness
Surprise + Happiness	Delight: delight
Surprise + Sadness	Embarrassed: embarrassed

Table 2: First-order complex emotions and some corresponding keywords

beyond a sentence, and encodes emotion information for a sentence. This annotation scheme can be compatible for any TEI-based annotated corpora as long as sentences are clearly marked.

The emotion-related elements (tags) in our annotation scheme are described as follows. For easy demonstration, our elements are defined with the format of British National Corpus (BNC) annotation scheme¹, and our examples are also based on BNC annotated text. Figure 1 gives the definition of each element, and Figure 2 shows several examples using our annotation scheme. Note that <s> element is a tag for a sentence-like division of a text, and its attribute “n” gives the sentence index. In Figure 2, Sentence 1, which expresses emotions by emotion keywords, contains two types of emotions: “surprise” (primary emotion) and “jealousy” (complex emotion); Sentence 2 is a neutral sentence.

<emotion> element

It is used only when the sentence expresses emotions. It contains a list of <emotionType> elements and a <s> element. As a sentence may

express several emotions, an <emotion> element can contain several <emotionType> elements, and each <emotionType> element describes an emotion occurring in that sentence separately.

<neutral> element

It is used only when the sentence does not contain any emotion expression. It contains only a <s> element.

<emotionType> element

It describes a type of emotion in that sentence. It contains an ordered sequence of <primaryEmotion> elements. Attribute “name” provides the name of the emotion type, such as “surprise”, “jealousy,” and so on, and it is optional. If the emotion type is a primary emotion, the <emotionType> element will have only one <primaryEmotion> element, which encodes the information of this primary emotion. If the emotion is a complex emotion, the <emotionType> element will have several <primaryEmotion> elements (each of them describes the primary emotion involved in that complex emotion.) Attribute “keyword” is an optional attribution if annotators want to provide the indicator of a text for that emotion.

¹ <http://www.natcorp.ox.ac.uk/XMLedition/URG/>

```

<emotion>
  <emotionType name = "surprise" keyword = "surprised">
    <primaryEmotion order = "1" name = "surprise" intensity = "moderate"></primaryEmotion>
  </emotionType>
  <emotionType name = "jealousy" keyword = "jealousy">
    <primaryEmotion order = "1" name = "anger" intensity = "moderate"></primaryEmotion>
    <primaryEmotion order = "2" name = "fear" intensity = "moderate"></primaryEmotion>
  </emotionType>
  <s n = "1"> Hari was surprised at the rush of pure jealousy that swept over her at the mention of Emily Grenfell .</s>
</emotion>
<neutral>
  <s n = "2"> By law no attempts may be made to hasten death or prolong the life of the sufferer . </s>
</neutral>
<emotion>
  <emotionType>
    <primaryEmotion name = "sadness"></primaryEmotion>
  </emotionType>
  <s n = "3">He looked hurt when she did n't join him , his emotions transparent as a child 's . </s>
</emotion>

```

Figure 2: The example of sentence annotation

```

element emotion
{
  (emotionType)+,
  <s>
}
element emotionType
{
  attribute name (optional),
  attribute keyword (optional),
  (primaryEmotion)+
}
element primaryEmotion
{
  attribute order (optional),
  attribute name (necessary),
  attribute intensity (optional)
}
element neutral
{
  <s>
}

```

Figure 1: The definition of emotion-related elements

<primaryEmotion> element

It describes the property of a primary emotion involved in the emotion type. There are three attributes: “order,” “name,” and “intensity.” “Order” gives the weight of this primary emotion in the emotion type, and the weight value decreases with the ascending “order” value. “Name” and “intensity” provide the name and intensity of a primary emotion. To encode the information in our emotion taxonomy, the value of “order” is {1,2,3,4,5}, the value of “name” is {“happiness,” “sadness,” “anger,” “fear,” “surprise”}, and the value of “intensity” is {“high,” “moderate,” “low”}.

The <primaryEmotion> element seems to be redundant because its encoded information can be obtained from the given emotion taxonomy if the name of the emotion type is available, but the presence of this element can make our annotation scheme more robust. Sometimes emotion

is so complicated (especially for those emotion expressions without any explicit emotion keyword) that an annotator may not be able to find an exact emotion type to match this emotion, or to list all involved primary emotions. For those subtle cases, emotion annotation can be simplified to list the involved primary emotions as many as possible through <primaryEmotion> elements. For example, in Sentence 3 in Figure 2, although there is no emotion keyword occurring, the word “hurt” indicates the presence of an emotion, which at least involves “sadness.” However, because it is hard to explicitly list other primary emotions, therefore, we give only the annotation of “sadness.”

Our annotation scheme has the versatility to provide emotion data for different applications. For example, if textual information input annotated with our scheme is provided for the Japanese robot Saya (Hashimoto et al, 2006) to control her facial emotion expression, a simple mapping from our 24 emotion types can be done automatically to Saya’s six emotion types, i.e. surprise, fear, disgust, anger, happiness, and sadness. As four of her emotion types are also unique primary emotions, using information encoded in <emotionType> element and <primaryEmotion> element will ensure unique many-to-one mapping and the correct robotic expressions. A trickier case involves her “anger” and “disgust” emotions. The emotion type “anger” in our taxonomy includes emotion words “anger” and “disgust”. However, with the “keyword” information provided in <emotionType> element, a small subset of “anger” emotion in our taxonomy can be mapped to “disgust” in Saya’s system. For example, we could map keywords “dislike, disgust, hate” to “disgust”,

and all the remaining ones, such as “outrage, furious,” to “anger.”

4 Emotion-driven corpus creation

Similar to most corpora, our corpus creation is designed to satisfy the requirements of real emotion computing. Emotions can be expressed with or without emotion vocabulary in the text. It seems to be intuitive that emotion computing for a context with emotion keywords can be satisfactory when the collection of emotion vocabulary is comprehensive, such as “joyful” indicates the presence of “happiness” emotion. However, this intuitive approach cannot work well because of the ambiguity of some emotion keywords and the emotion context shift as the sentiment shift (Polanyi and Zaenen, 2004). Moreover, the detection of emotions in a context without emotion keywords is very challenging. To deal with these problems, we build the emotion corpus, which is motivated by the NSM theory.

According to the NSM theory, an emotion is provoked by a stimulus. This indicates one possible way to detect emotions in text, i.e. the detection of emotional stimulus, which is often provided in the text. In other words, emotion corpus is a collection of emotion stimuli. Since emotion is subjective, the stimulus-based approach works only when its context is provided. For example, the stimulus – “build a gym for this community” – may cause different emotions, such as “surprise”, “happy” and so on, depending on its context. We also notice that the text containing an emotion keyword may contain emotional stimulus and its context. Thus, a natural corpus creation approach comes out.

In our system, a pattern-based approach is used to collect the emotion corpus, which is similar to the one used in Tokuhisa et al. (2008), but we do not limit to event-driven emotions (Kozareva et al., 2008), and adjust our rules to improve the quality of emotion annotation. There are five steps in our emotion sentence annotation as given below, and Steps (2) and (3) are to improve the annotation quality.

- 1) Extract emotion sentences: sentences containing emotion keywords are extracted by keyword matching.
- 2) Delete ambiguous structures: some ambiguous sentences, which contain structures such as negation and modal, are filtered out.
- 3) Delete ambiguous emotion keywords: if an emotion keyword is very ambiguous, all sen-

tences containing this ambiguous emotion keyword are filtered out.

- 4) Give emotion tags: each remaining sentence is marked with its emotion tag according to the emotion type which the focus emotion word belongs to (refer to Tables 1 and 2.)
- 5) Ignore the focus emotion keywords: for emotion computing, the emotion word is removed from each sentence.

Polanyi and Zaenen (2004) addressed the issue of polarity-based sentiment context shift, and the similar phenomenon also exists in emotion expressions. In our corpus creation, two kinds of contextual structures are handled with: the negation structure and the modal structure. In both English and Chinese, a negated emotion expression can be interpreted as one of the three possible meanings (as shown in Figure 3): opposite to the target emotion (S1), deny the existence of the target emotion (S2), or confirm the existence of the target emotion (S3). The modal structure often indicates that the emotion expression is based on the counter-factual assumption, hence the emotion does not exist at all (S4 and S5 in Figure 3). Although Chinese and English have different interpretations about the modal structure, for emotion analysis, those sentences often do not express an emotion. Therefore, to ensure the quality of the emotion corpus, all sentences containing a negation structure or a modal structure, which are detected by some rules plus a list of keywords (negation polarity words for the negation structure, and modal words for the modal structure), are removed.

S1 (Neg_Happiness): I am not happy about that.
S2 (Netural): Though the palazzo is our family home, my father had never been very happy there.
S3 (Pos_Happiness): I've never been so happy.
S4 (Netural): I can die happy if you will look after them when I have gone.
S5 (Netural): Then you could move over there and we'd all be happy.

Figure 3: Structures for emotion shift

To overcome the high ambiguity of some emotion keywords, after Step (2), for each emotion keyword, five sentences are randomly selected and annotated by two annotators. If the accuracy of five sentences is lower than 40%, this emotion keyword is removed from our emotion taxonomy. Finally, 191 Chinese keywords and 645 English keywords are remained.

Tokuhisa et al. found that a big challenge for emotion computing, especially for emotion detection, is to collect neutral sentences. Since neutral sentences are unmarked and hard to detect, we develop a naïve yet effective algorithm

to create a neutral corpus. A sentence is considered as neutral only when the sentence itself and its context (i.e. the previous sentence and the following sentence) do not contain any of the given emotion keywords.

We run our emotion sentence extraction and neutral sentence extraction on three corpora: the Sinica Corpus (Chinese), the Chinese Gigaword Corpus, and the British National Corpus (BNC, English), and create three emotion corpora and three neutral corpora separately. The Sinica Corpus is a balanced Chinese corpus, which includes documents in 15 kinds of genres; The Chinese Gigaword Corpus is a huge collection of news reports; The BNC is also a balanced corpus, which collects documents from different domains.

To estimate the accuracy of our emotion sentence extraction, we randomly select about 1000 sentences from the three emotion corpora, and have two annotators to check it. Table 3 lists the accuracy of those emotions sentences (emotion corpus.) To test how good this straightforward neutral sentence extraction strategy is, about 1000 sentences are randomly selected from each of the three neutral corpora and are checked by two annotators. Table 3 lists the accuracy of those neutral sentences (neutral corpus.)

	Emotion corpus	Neutral corpus
Gigaword	82.17	98.61
Sinica	77.56	98.39
BNC	69.36	99.50

Table 3: The accuracy of the emotion-driven corpora

From Table 3, the high accuracy of neutral corpus proves that our approach is effective in extracting neutral sentences from the document-based corpus which contains contextual information. Although the accuracy of emotion corpus is lower, it is still much higher than the one reported by Kozareva et al. (2008), i.e. 49.4. The accuracy is significantly increased by deleting ambiguous emotion keywords in Step (3). For the 2,474 randomly selected Chinese sentences, the overall accuracy of the remaining 1,751 sentence is increased by about 14% after Step (3). For the 803 randomly selected English sentences, the accuracy of the remaining 473 sentence is increased about 21% after Step (3). Whether or how the ambiguous emotion keywords in Step 3 are removed is a tradeoff between the coverage and the accuracy of the emotion corpus.

From Table 3, we also find that the accuracy of English emotion corpus is much lower than Chinese emotion corpus, which indicates Eng-

lish emotion sentences expressed by emotion keywords are more ambiguous than that of Chinese. Moreover, during our emotion corpus building, 20.2% of Sinica sentences and 22.4% of Gigaword sentences are removed in Step (2) and (3), on the contrary, 41.2% of BNC sentences are deleted. Although it is more difficult to develop the rules in Step (2) and (3) for Chinese than for English, it also confirms the higher ambiguity of emotion expressions in English due to the ambiguity of emotion keyword. Finally, because of the comparatively-high percentage of the sentences removed in Step (2) and (3), more exploration about those sentences is needed, such as the emotion distribution, the expression patterns and so on, and how to re-incorporate them into the emotion corpus without hurting the whole quality is also our future work.

We also explore emotions through the sentences (no-emotion-keyword sentences) that do not contain any given emotion keyword, because our approach extracts only partial neutral sentences and partial emotion sentences in reality. For each corpus, about 1000 no-emotion-keyword sentences are randomly selected and checked by two annotators. It is surprising that only about 1% of those sentences express emotions. This indicates that it is important for real emotion computing, which mainly works on formal written text, to deal with the emotion expressions which contain emotion keywords and however are ambiguous, such as the sentences deleted in Steps (2) and (3). More exploration is needed for the emotion and neutral sentence distribution on other kinds of written text, such as blogs, and on spoken text.

The unsupervised corpus creation approach can easily be adapted for different languages and different emotion applications, provided that the keyword collection and patterns in Step (2) and (3) need some changes. Moreover, another big advantage of our approach is that it can avoid the controversy during emotion annotation. Emotion is subjective, and therefore disagreement for emotion types often arises if the emotion is not expressed through an explicit emotion keyword.

Overall, the annotated corpus created by the unsupervised approach has a comparatively high quality, and is suitable for the emotion computing. As the size of the neutral corpus is much bigger than its corresponding emotion corpus, to avoid model bias, we randomly select some neutral sentences from the neutral corpus, combin-

ing with its corresponding emotion sentences to form a complete emotion-driven corpus.

5 Emotion computing system

In this paper, we present some pilot work to prove that our emotion-driven corpus is useful for emotion computing. With the inclusion of neutral sentences, emotion detection and classification is simplified into a general classification problem, and a supervised machine learning method can be directly applied if enough annotated data are obtained. Here, we choose the MaxEnt learning in Mallet as a classifier.

Both the Sinica Corpus and the Chinese Gigaword Corpus are segmented, and POS-tagged. This allows us to implement the bag-of-words approach in the focus sentences in both Chinese and English. However, emotions are mostly human attitudes or expectations arising from situations, where situations are often expressed in more than a single word. Such kind of situations tends to be more easily extracted by word bi-grams (2-gram word) than by word unigram (1-gram word.) To take this into account, besides 1-gram words, we also extract word bi-grams from the focus sentences.

There are too many emotion types in our corpus, which can cause data sparse; therefore, we choose the most frequent emotions to do exploration. Besides the five primary emotions, for Chinese, we select another nine complex emotions, and for English, we select another four complex emotions. Other emotion types are re-named as “Other Emotions.”

Since Chinese emotion-driven corpus is much larger than the English one, to fairly compare the performance, we reduce the size of Chinese corpus in our experiments. Then, for each corpus, we reserve 80% as the training data, 10% as the development data, and 10% as the test data (there are two sets of test data as follows.) In the evaluation, for each emotion sentence, if our system detects one of its emotion tags, we consider this sentence is correctly tagged.

Test data set 1 (TDS 1): contains about 10% of the sentences from the complete emotion-driven corpus, and emotion tags are automatically given during the corpus creation.

Test data set 2 (TDS 2): contains the sentences used in Table 3, which is checked by two annotators. If more than one emotion tags co-exist in a sentence, all of them are chosen to label the sentence. If there exists an emotion that

does not belong to any of the emotion types, it is labeled as “Other Emotions.”

Table 4 shows the performance (accuracy) of our system for Test data set 1 and 2 for both Chinese and English. We notice that our corpus creation approach is effective for emotion computing. As we expect, the 2-gram words can partially catch the emotion stimulus, and improves the performances. However, the overall performance is still very low, which indicates that emotion computing is a difficult task. From the error analysis, it is surprised that for Chinese, the mislabeling of emotion sentences as neutral sentences (“emotion” vs. “neutral”) is a common error, and whereas, for English, two kinds of errors: “emotion” vs. “neutral” and “focus emotions” vs. “Other emotions” (the mislabeling of a sentence with a focus emotion as “Other emotions,”) occupy at least 50%. The error distribution confirms the importance of emotion detection during emotion computing. The high frequency of the error of “focus emotions” vs. “Other Emotions” in English may be because there are fewer focus emotion types for English.

	1-gram words	{1,2}-gram words
Chinese TDS 1	53.92	58.75
English TDS 1	44.02	48.20
Chinese TDS 2	37.18	39.95
English TDS 2	33.24	36.31

Table 4: The performances of our system for the test data

6 Conclusion

Emotion, no matter its annotation or computing, is still a new and difficult topic. In this paper, we apply emotion theories to design a cognitive-based emotion annotation scheme, which are robust and versatile so that it can encode different levels of emotion information for different emotion computing. Moreover, motivated from NSM, we develop an unsupervised approach to create a large and comparatively high-quality corpus for emotion computing, which is proven in our pilot experiments to be useful. Moreover, this approach makes emotion computing for different applications possible through a little modification.

Certainly, there are some issues remaining unsolved. For corpus construction, we will explore emotion distribution in other kinds of corpora, such as blog and dialog, and make analysis of ambiguous emotion sentences, such as negation structure and modal structure. For emotion computing, we did only pilot experiments and more work needs to be done, such as feature extraction.

References

- W. B. Cannon. 1927. The James-Lange theory of emotions: A Critical Examination and an Alternative Theory. *American Journal of Psychology*, 39, 106-124.
- Y. Chen, S. Y. M. Lee and C. R. Huang, 2009. Construction of Chinese Emotion Corpus with an Unsupervised Approach. In CNCCL-2009, 2009. (in Chinese)
- T. Hashimoto, S. Hiramatsu, T. Tsuji and H. Kobayashi. 2006. Development of the Face Robot SAYA for Rich Facial Expressions. SICE-ICASE International Joint Conference, Busan, Korea.
- J. Hobbs and A. Gordon. 2008. The Deep Lexical Semantics of Emotions. Workshop on Sentiment Analysis: Emotion, Metaphor, Ontology and Terminology (EMOT-08), 6th International conference on Language Resources and Evaluation (LREC-08), Marrakech, Morocco, May 27, 2008.
- P. Livia, A. Zaenen. 2004. Contextual Valence Shifters. In Shanahan, J. G., Y. Qu, and J. Wiebe (Eds.), *Computing Attitude and Affect in Text: Theory and Applications*, pp. 1-10.
- Z. Kozareva, Borja Navarro, Sonia Vazquez, and Andres Nibtoyo. 2007. UA-ZBSA: A Headline Emotion Classification through Web Information. In *Proceedings of the 4th International Workshop on Semantic Evaluations*.
- Y. Y. Mathieu. 2005. Annotations of Emotions and Feelings in Texts. In Conference on Affective Computing and intelligent Interaction (ACII2005), Beijing, Springer Lecture Notes in Computer Science, pp. 350-357.
- R. Mihalcefa, and Hugo Liu. 2006. A Corpus-based Approach to Finding Happiness. In *Proceedings of AAAI*.
- R. Plutchik. 1991. *The Emotions*. University Press of America, Inc.
- R. Plutchik. 1980. *Emotions: A psychoevolutionary synthesis*. New York: Harper & Row.
- M. Schröder, H. Pirker and M. Lamolle. 2006. First suggestions for an emotion annotation and representation language. In L. Deviller et al. (Ed.), *Proceedings of LREC'06 Workshop on Corpora for Research on Emotion and Affect* (pp. 88-92). Genoa, Italy.
- J. G. Shanahan, Y. Qu and J. Wiebe. 2006. *Computing attitude and affect in text: theory and applications*, Springer.
- R. Tokuhisa, K. Inui, and Y. Matsumoto (Eds.) 2008. *Emotion Classification Using Massive Examples Extracted from the Web*. COLING.
- J. H. Turner. 2007. *Human Emotions: A sociological theory*. New York : Routledge, 2007.
- J. H. Turner. 2000. *On the origins of human emotions: A sociological inquiry into the evolution of human affect*. Stanford, CA: Stanford University Press.
- J. H. Turner. 1996. The Evolution of Emotions in Humans: A Darwinian–Durkheimian Analysis. *Journal for the theory of social behaviour* 26:1-34
- L. Xu, H. Lin, J. ZHAO. 2008. Construction and Analysis of Emotional Corpus. *JOURNAL OF CHINESE INFORMATION PROCESSING*.
- X. Y. Xu, and J. H. Tao. 2003. The study of affective categorization in Chinese. The 1st Chinese Conference on Affective Computing and Intelligent Interaction. Beijing, China.
- A. Wierzbicka, 1996. *Semantics: Primes and Universals*. Oxford: Oxford University Press.