

Long Nights, Rainy Days, and Misspent Youth: Automatically Extracting and Categorizing Occasions Associated with Consumer Products

David B. Bracewell

Oculus360

16301 Quorum Dr. Suite 100B

Addison, TX, USA

dbracewell@oculus360.us

Abstract

One way in which marketers gain insights about consumers is by identifying the occasions in which consumers use their products and which are invoked by their products. Identifying occasions helps in consumer segmentation, answering why consumers purchase a product, and where and when they use it. Additionally, the types of occasions a consumer participates in and the social settings surrounding those occasions provide insights into the consumer's personality and sociocultural self. Insights such as these are required for understanding consumer behavior, which marketers need to better design and sell their products. In this paper, we describe a methodology for extracting and categorizing occasions from product reviews, product descriptions, and forum posts. We examine using a maximum entropy markov model (MEMM) and a linear chain conditional random field (CRF) for extraction and find the CRF results in a 72.4% F1-measure. Extracted occasions are categorized as one of six high-level types (*Celebratory*, *Special*, *Seasonal*, *Temporal*, *Weather-Related*, and *Other*) using a support vector machine with an 88.5% macro-averaged F1-measure.

1 Introduction

Social media provides an outlet for consumers to discuss, praise, chastise, and recommend products and services. These consumer generated reviews and commentaries provide marketers insight into the who, what, when, where, why, and how (i.e. the six W's) surrounding the procurement and usage of

their products. One way in which marketers answer the six W's is by identifying the occasions, particular times or events, in which their products are used or with which consumers associate their products. These occasions may be routine, e.g. "work" or "at the office", seasonal/weather related, e.g. "rainy day" or "winter", special, e.g. "birthday" or "Christmas", or time related, e.g. "on the run" or "early morning." More than just answering the six W's, occasions also provide a marketer insight into the personality, social status, social circle, and behavior of consumers.

Marketers traditionally rely upon surveys and ethnographic studies in order to gain insights about consumers. The results of these surveys and studies are: (1) consumer segments; (2) when and where the respondents are likely to purchase or use a product; (3) whether they are likely to use the product alone or with others; (4) whether or not the respondents like the product; and (5) are the respondents likely to purchase the product again. These surveys and studies are costly and limited to a much smaller sample size than is obtainable via online reviews and social media. However, current computational approaches to gaining consumer insights typically are limited to the volume and trend of positive and negative comments, reviews, tweets, etc. (Pang et al., 2002; Dini and Mazzini, 2002; Smith et al., 2012; Socher et al., 2013).

Research from the fields of consumer and social psychology, dialogue processing, and affective must be incorporated into computational systems in order for them to replace surveys as a marketer's source of consumer insights. Drawing on these fields of re-

search facilitates an understanding of the attitudes, behaviors, and personal and sociocultural qualities of consumers. Critical to the success of such a computational system is the automated extraction of the occasions in which consumers use a product or with which they associate a product. These occasions and their implicatures provide answers to the six W's and are a basis for understanding a consumer's personal and sociocultural self.

In this paper, we present a methodology for automatically extracting and categorizing occasions in product reviews, product descriptions and forum posts. The extraction of occasions is cast as a sequence labeling problem using the standard BIO encoding. Extracted occasions are categorized as one of six high-level types, *Celebratory*, *Special*, *Seasonal*, *Temporal*, *Weather-Related*, and *Other*, based on common occasions marketers seek to capture in surveys.

2 Related Work

The most related area of research to the extraction of occasions is event extraction. Event extraction deals not only with the extraction of events, but also with the extraction of the entities participating in the events, and other attributes of the event, such as the time (Moschitti et al., 2013), location (Speriosu et al., 2010), and modality (Bracewell et al., 2014). Despite the advances in the extraction of events, the definition of an "event" is ill-defined and changes based on the problem being solved. The Automated Content Extraction (ACE) program defines events using a limited set of types (ACE, 2005). TimeML defines events as "situations that happen or occur" and mainly focuses on the duration properties of the event (Pustejovsky et al., 2003). Instead of precisely defining what an event is, Monahan and Brunson (2014) identify the qualities representative of events.

Research in real-time event detection has benefited from the wide spread acceptance and adoption of social media. Sites like Twitter and Facebook act like social sensors facilitating the real time detection of disasters (Sakaki et al., 2010; Vieweg et al., 2010) and local events (Boettcher and Lee, 2012; Lee and Sumiya, 2010). Relying on the real-time nature of Twitter and the volume of tweets around unusual or significant events, Sakaki et al. (2010) construct a

real-time earthquake detection system using twitter users as sensors. Lee and Sumiya (2010) use Twitter to determine unusual local events happening in a given geographic area based on the regularity of tweets against the normal behavior of twitter users in the area.

The dialogue that takes place over social media makes it possible to find and extract life and social events for such purposes as detecting online bullies (Dinakar et al., 2011) and suicide prevention (Jashinsky et al., 2014). Li et al. (2014) target specific replies on Twitter containing manifestations of speech acts, namely congratulations/condolences, to extract major life events, e.g. marriage, using a distant-supervised approach. In addition to the detection of events, work has been done on identifying the social implicatures of dialogue which is in response to a set of events (e.g. Wikipedia page edit) or which may lead to a series events (e.g change in leadership) (Bracewell et al., 2011; Bracewell et al., 2012; Tomlinson et al., 2012).

Broader related research on mining consumer insights is found in the fields of consumer psychology and affective computing. Consumer psychology studies how thoughts, feelings, and perceptions influence the way individuals buy, use, and relate to products, services, and brands. Drawing from other areas in psychology, e.g. social psychology, consumer psychologists formalize the cognitive system of consumers using a categorical representation of products, services, brands and other marketing entities (Loken et al., 2008). Supported by Rosch's (1973) work on prototype theory, Loken and Ward (1990) find a link between the prototypicality of a product and consumers' affect toward it.

A critical component to understanding consumers' affect toward a product is identifying the brands, products, and attributes (or aspects) of the product consumers are mentioning. Wiegand and Klakow (2014) examine separating types from brands, e.g. "soda" vs "coke", using a ranking-based approach which alleviates the need for labeled data. Putthividhya and Hu (2011) use a named entity recognition system to extract product attributes from listing titles on eBay. They focus on extracting brand, style, size, and color within the clothing and shoes categories. Stoica et al. (2007) describe a WordNet-based approach to constructing hierarchi-

cal facets relating to aspects associated with a domain or product. Yu et al. (2011) present a domain-assisted approach to constructing aspect hierarchies.

Aspect-based sentiment analysis (Pontiki et al., 2014) merges affect and information extraction seeking to determine the sentiment toward aspects of products, e.g. the consumer sentiment toward the screen of a TV or the food at a restaurant. Approaches to aspect term identification range from standard BIO encoding (Chernyshevich, 2014; Toh and Wang, 2014) to rule-based approaches (Porcia et al., 2014). Techniques for aspect polarity detection include machine learning based techniques that integrate multiple sentiment lexicons (Wagner et al., 2014) to grammar based approaches (Brun et al., 2014).

More general than aspect-based sentiment analysis is sentic computing (Cambria and Hussain, 2012). Sentic computing synthesizes common-sense computing, linguistics, and psychology to infer both affective and semantic information about concepts. Cambria et al. (2014) show how SenticNet, a semantic and affective resource, can detect topics and determine polarity in patient opinions.

Another area of research relevant to consumer insights is around the identification of needs and wants on social media. Kanayama and Nasukawa (2008) examine the needs and wants of consumers using syntactic patterns to analyze the demand for products. Ramanand et al. (2010) examine the identification of wishes in reviews and surveys in which consumers make suggestions for improvements and show their intentions to purchase/use a product.

3 Modeling Occasions for Consumer Insights

Occasions are particular times or events and range from the everyday, such as waking up and going to bed, to the special, such as birthdays and weddings. While every occasion is of importance, those surrounding products are of the most use to marketers for gaining insights into consumer behavior. Thus, in this paper we focus only on occasions which are related to a product. More specifically we restrict the definition of an occasion to:

Times or happenings in which a product is used or with which a product is associated.

Occasions matching this definition are in bold font in the following examples:

1. “They are GREAT to take along to a **party** if you’re serving crackers and cheese.”
2. “I bought these for my **vacation** and they did not disappoint.”
3. “Boy, do these take me back to those **misspent days of my foolish youth.**”

In the first example the occasion is a party relating to where the reviewer used the product. From this example we can infer that the occasion of use is social, i.e. involved more than just the reviewer, and most probably is informal. Furthermore, we learn that the reviewer believes the product is well suited for party occasions. Given further context about the kind of party, e.g. kids or work, would lead to further insights about the individual, such as if they have children, their age, their occupation, and their marital status. In the second review the occasion (“vacation”) is the reason for the reviewer to purchase the product and the answer to when the reviewer used it. Moreover, from the review we can infer that the use of the product was a positive experience for the reviewer. The third review is an example of how a product can be associated with an occasion, which in this case is a memory of the reviewer’s youth. Marketers use these type of occasions to connect with consumers at a subconscious and emotional level.

While occasions are closely related to events, not all fit nicely within the ACE and TimeML definitions. For example, take the following:

1. “These boots really kept me warm during the **winter.**”
2. “Every time I smell a freshly baked apple pie it **brings me back to my childhood.**”

In the first example the product is a pair of boots and the occasion of use is the winter. Within an event framework winter would not be identified as an event, but as a temporal attribute possibly of a “keep warm”. In the second example the occasion is “brings me back to my childhood” and is associated with the product (“apple pie”) by the reviewer. The event in the sentence is a “baking” event with

Occasion Type	Definition
Celebratory	Occasions meant to celebrate an event, person, or group of people (e.g. parties and award ceremonies)
Special	Occasions which have significant importance to an individual or group of individuals (e.g. holidays and life events)
Seasonal	Occasions related to the seasons of the year. (e.g. winter)
Weather-Related	Occasions strongly associated with the weather and/or temperature. (e.g. hot days and rainy nights)
Temporal	Occasions tied to a specific time (e.g. 9 to 5, late night, and last year)
Other	Occasions which do not fit in the other categories (e.g. a shopping spree, at the beach)

Figure 1: The six high-level occasion types used to categorize occasion mentions.

the apple pie being the item baked. The occasion is tangential to the event and most likely would not be associated with it by an event extraction system. However, this type of occasion provides evidence of a strong connection between the product and a specific time or event that is nostalgic for the consumer and is invaluable for marketers when crafting their marketing strategy.

Often occasions are associated with special events, such as ceremonies and celebrations. However, as with event types, there are a number of different types of occasions. We define six high-level types, listed in Figure 1, which are based on common occasions marketers use to segment consumers.

Celebratory occasions, which include parties and festivals, are social occasions and inform to the group with which the consumer belongs. An example of a celebratory occasion is:

“I wore it a couple weeks ago to a **party** and felt festive yet as comfy as if I was wearing loungewear.”

Some celebrations are due to special occasions. Special occasions are those which have significant meaning to the consumer, such as holidays and religious observances. The following review excerpt contains mentions of two special occasions:

“I recommend these for your **engagement party** or **rehearsal dinner**.”

Temporal and seasonal occasions relate to the time in which a product is used or associated. An example of a seasonal occasion is :

“A quintessential style to take you **between seasons**.”

The following excerpt from a product description contains two suggested temporal occasions of use:

“Just the right size for your **day-to-day life**, but elegant enough for **evening**.”

Weather-related occasions relate to the weather, e.g. rain and snow, or temperature, e.g. hot and 98 degrees. Two examples of weather-related occasions are seen in:

“The tea is great hot for **chilly nights** and iced for **hot days**.”

Finally, we define an other type for occasions that do not neatly fit in one of the previous five categories. An example of an occasion that is marked as other is:

“Taking a look at the latest summer fashion makes me want to **lie on the beach**.”

While there are a multitude of additional occasions types that are definable, we limit the categories to the six presented above in this paper.

4 Data Collection and Annotation

We collect 26,208 sentences from 1,000 product reviews, 500 product descriptions, and 800 forum posts discussing fashion and food related products for annotation. An iterative annotation process is used wherein during each iteration automated machine annotation is performed followed by manual correction. During the initial iteration automated machine annotations are produced using a gazetteer and successive iterations use a machine learning model. Manual correction of the machine annotations involves: (1) removing incorrect occasions; (2) adding missed occasions; and (3) fixing boundaries of partially correct occasions. Due to project constraints all manual correction is performed by one annotator. In the future, we hope to employ multiple annotators.

The initial iteration of the annotation process is performed on 7,000 randomly selected sentences. The gazetteer used during the initial iteration is semi-automatically constructed using WordNet (Miller, 1995). The full hyponym tree and all derivationally related forms for *social event*, *time period*, and the first noun sense of *activity* are extracted to construct the gazetteer.

Examples of occasions identified using the gazetteer are as follows:

1. “Darling **cocktail party** or *date night* dress .”
2. “We only stayed at the **party** an hour because my shoes were killing my feet.”

In the examples listed above, occasions in bold font are correctly identified by the gazetteer and left as is, underlined occasions are incorrectly identified by the gazetteer and removed, and occasions in italic font are not in the gazetteer and added during manual correction. After manual correction (involving the previously three mentioned steps) of the initial 7,000 sentences, 4,500 are randomly selected and held out as test data, and 500 are held out as a development set for occasion extraction. The remaining 2,000 sentences are used as training data for the machine learning model in the second iteration.

The second and successive iterations work on batches of 500 sentences. At each iteration a machine learning model is trained and then used to extract occasions in the new batch of sentences. During each iteration we switch the model we train between the two described in Section 5. We alternate models to ensure we do not bias toward one model and because each model is likely to find something the other did not. The machine identified annotations are manually corrected and added to the set of training data for the next iteration. This process is repeated until all sentences are annotated.

2,393 occasions are annotated across the 26,208 sentences making up the corpus. This an average of 1 occasion every 11 sentences. There is approximately 1 occasion per product review and forum post and 1 occasion every 3 product descriptions.

The next step in the annotation process is to assign a type to each of the 2,393 annotated occasions. We use WordNet to assign an initial type and manually correct the assigned labels. We construct a map-

ping between WordNet senses and occasion types by starting with a set of twelve seeds, listed in Figure 2. The full hyponym tree and all derivationally related forms of each seed sense are extracted and mapped to the seed’s associated occasion type.

WordNet Sense	Occasion Type
party#N#4	Celebratory
celebration#N#1	Celebratory
season#N#2	Seasonal
temperature#N#1	Weather-Related
day#N#1	Temporal
day#N#2	Special
valentine#N#1	Special
gift#N#1	Special
anniversary#N#1	Special
birthday#N#1	Special
special#A#3	Special
New Year#N#1	Special

Figure 2: Seed senses for mapping from WordNet senses to occasions types. Where the sense is described in *lemma#POS#sense* number form.

WordNet lemmas found in a given occasion annotation are examined in right-to-left order. All senses for a lemma are considered in order of sense number. Assignment is performed greedily with the type of the first sense found in the mapping being assigned to the occasion. The *Other* type is assigned if no mapping is found.

Type	Count
Celebratory	107
Seasonal	525
Special	336
Temporal	263
Weather-Related	48
Other	1,114

Table 1: The number of occasions annotated for the six high-level types.

After automatic type assignment is complete the types are manually corrected. Most types are easily determined by an annotator. However, the celebratory and special types do have an overlap, e.g. birthday party. Annotators are told to assign the category of special instead of celebratory when the celebration is associated with a life event (e.g. birthday and engagement parties). The breakdown of the number occasions of each type is shown in Table 1.

5 Computational Methodology and Experimental Results

We divide the extraction and categorization of occasions into two different tasks. We found in preliminary experiments that this division produces better results than jointly performing the two tasks. The rest of this section details the models and results for each task.

5.1 Automatically Extracting Occasions

We model the extraction of occasions using the standard BIO encoding. Words in a sentence are labeled as *B-Occasion*, *I-Occasion*, or *Other* depending on if the word begins an occasion phrase, is within an occasion phrase, or is outside of an occasion phrase respectively. We experiment using a maximum entropy markov model (MEMM) (McCallum et al., 2000) and a linear chain conditional random field (CRF) (Lafferty et al., 2001) to perform extraction. We use an in-house implementation of MEMMs, which uses the LibLinear library (Fan et al., 2008), and CRFsuite (Okazaki, 2007) for the CRF implementation. Parameters are tuned using a grid search to maximize the F1-measure over the 500 sentence development set. The optimal parameters for the MEMM are $C = 3$ and the optimal parameters for the CRF are $C1 = 0$ and $C2 = 2$.

The feature templates used for the extraction of occasions are listed in Figure 3. The features consist of surface, syntactic, and semantic information about the word and its context. Syntactic information is in the form of part-of-speech information and semantic information is in the form of WordNet super sense, i.e. lexicographer filenames (note that all possible super senses are for a word, i.e. no sense disambiguation is performed). These features, with the exception of the WordNet-based feature, are commonly used in other sequence labeling tasks, such as shallow parsing and named entity recognition. We eliminate all features that occur only once in our training set.

5.1.1 Results

Performance is measured using the CoNLL precision, recall, and F1-measure and the percentage instance error in which an occasion is correct if and only if it exactly matches a gold standard annota-

Current word	w_i	$\& t_i$
Current word & POS	w_i, p_i	$\& t_i$
Previous word & POS	w_{i-1}, p_{i-1}	$\& t_i$
Word two back & POS	w_{i-2}, p_{i-2}	$\& t_i$
Next word & POS	w_{i+1}, p_{i+1}	$\& t_i$
Word two ahead & POS	w_{i+2}, p_{i+2}	$\& t_i$
Bigram word	w_{i-2}, w_{i-1}	$\& t_i$
	w_{i-1}, w_i	$\& t_i$
	w_i, w_{i+1}	$\& t_i$
	w_{i+1}, w_{i+2}	$\& t_i$
Bigram word & POS	$w_{i-2}, p_{i-2}, w_{i-1}, p_{i-1}$	$\& t_i$
	$w_{i-1}, p_{i-1}, w_i, p_i$	$\& t_i$
	$w_i, p_i, w_{i+1}, p_{i+1}$	$\& t_i$
	$w_{i+1}, p_{i+1}, w_{i+2}, p_{i+2}$	$\& t_i$
Trigram word	w_{i-2}, w_{i-1}, w_i	$\& t_i$
	w_i, w_{i+1}, w_{i+2}	$\& t_i$
Current POS	p_i	$\& t_i$
Previous POS	p_{i-1}	$\& t_i$
POS two back	p_{i-2}	$\& t_i$
Next POS	p_{i+1}	$\& t_i$
POS two ahead	p_{i+2}	$\& t_i$
Bigram POS	p_{i-2}, p_{i-1}	$\& t_i$
	p_{i-1}, p_i	$\& t_i$
	p_i, p_{i+1}	$\& t_i$
	p_{i+1}, p_{i+2}	$\& t_i$
Current word is punct.	$isPunctuation(w_i)$	$\& t_i$
Current word is digit	$isDigit(w_i)$	$\& t_i$
Current word is letter	$isLetter(w_i)$	$\& t_i$
Current word is upper	$isUppercase(w_i)$	$\& t_i$
Current word is lower	$isLowercase(w_i)$	$\& t_i$
WordNet super sense	$ss_{ij} \forall sense(w_i)$	$\& t_i$

Figure 3: Feature templates used for extracting occasions. w_1, \dots, w_n are the words in the sentence and w_i the current word. p_1, \dots, p_n is the part-of-speech sequence for the sentence and p_i is the part-of-speech for the current word w_i . $sense(w_i)$ returns all possible senses for the current word, w_i , and ss_{ij} is the super sense associated with sense j . t_i is the tag assigned to the i 'th word.

tion. Results for the MEMM and CRF are listed in Table 2. As is seen in the table, the CRF model outperforms the MEMM with an increase in precision of 4.5%, recall of 20.6%, and F1-measure of 16.5%. Additionally, the CRF has an approximately 57% decrease in instance error rate.

Model	P	R	F1	Err
MEMM	79.2%	43.2%	55.9%	4.9%
CRF	83.7%	63.8%	72.4%	2.8%

Table 2: CoNLL Precision, Recall, F1-measure, and percentage instance Error results for extracting occasions.

Table 3 lists the precision, recall, and F1-measure

by length of the occasion in words. The performance of the MEMM degrades as the length of the occasion increases whereas the performance of the CRF is consistent across the varying lengths. One explanation of why the CRF performs better than the MEMM is the label-bias problem. Label-bias is a known weakness of MEMMs, which CRFs address, in which contextual information is lost around low-entropy transitions due to the use of a per-state (vs single) exponential model (Lafferty et al., 2001).

Model	Length	P	R	F1
MEMM	1	79.8%	55.6%	65.6%
	2	84.6%	41.8%	55.9%
	3	75.0%	25.5%	38.1%
	4	76.9%	35.7%	48.8%
	5+	40.0%	66.7%	11.4%
CRF	1	83.9%	52.8%	64.8%
	2	80.8%	74.6%	77.6%
	3	89.2%	70.2%	78.6%
	4	85.7%	85.7%	85.7%
	5+	80.8%	70.0%	75.0%

Table 3: CoNLL Precision, Recall and F1-measure by length of occasion in words.

Examples where the CRF and MEMM extract an occasion correctly are:

1. “Just what you need for a **hot summer day!**”
2. “We (my son and I) purchased this gift set for my wife on **Valentines day.**”
3. “It’s the perfect size to take me from a **day at work** to a **night out for drinks with friends.**”

In the first example, the occasion (“hot summer day”) is noun phrase representing the reviewer’s belief of a good time to use the product. In the second example the occasion is a holiday (“Valentines day”). The final example contains two occasion mentions that represent a time range, in the form of *from time₁ to time₂*.

5.2 Automatically Categorizing Occasions

Once an occasion is extracted it is categorized as one of the previously defined six types. We examine the effectiveness of categorizing occasions given only the occasion and no context. This task is an example of a short-text classification problem (Sriram et al., 2010). To solve this task we use a multi-class

support vector machine as implemented in the Lib-Linear library (Fan et al., 2008). We use the default values for the C and ϵ parameters.

Three features are used for determining the type of an occasion. The first is the standard bag of words with words normalized to lowercase. The second feature is the WordNet super senses of all possible senses found in the occasion. The super senses for adjectives and adverbs in WordNet are not as well defined as they are for nouns and verbs. Because of this, we use the super sense for the associated noun sense using the derivationally related form relation for adjectives and the pertainym (adverb to adjective) and derivationally related form (adjective to noun) relations for adverbs. The final feature is the SUMO concepts (Benzmüller and Pease, 2012) associated with all WordNet senses in the occasion.

5.2.1 Results

Table 4 lists the 10-fold cross-validation results for determining the type of a given occasion. As is seen in the table, F1-measures range from 71.9% for weather-related to 96.7% for seasonal.

Type	P	R	F1
Celebratory	92.6%	84.7%	88.5%
Seasonal	95.9%	97.5%	96.7%
Special	96.5%	93.8%	95.1%
Temporal	80.6%	88.6%	84.4%
Weather-Related	78.0%	66.7%	71.9%
Other	94.5%	93.8%	94.2%
Macro-avg	89.7%	87.5%	88.5%
Micro-avg	93.1%	93.1%	93.1%

Table 4: 10-fold cross-validation Precision, Recall, and F1-measure for categorizing occasions as *Celebratory*, *Special*, *Seasonal*, *Temporal*, *Weather-Related*, or *Other*.

Examples of errors in type assignment are shown in Figure 4. The errors in the first two examples happen due to “spring” and “time” being highly associated with seasonal and temporal occasions respectively. In the third example, the system assigns the type other whereas the true type is special. While the act of “shooting photos” is itself not special the type of photos (“engagement”) in the example does make it special. In the fourth example the occasion “upcoming year” is assigned special by the system most likely due to its similarity to the variations of the “new year” special occasions in the corpus. The fi-

Occasion	Gold	System
1.) “new spring semester”	Temporal	Seasonal
2.) “spend time with the one you love”	Other	Temporal
3.) “shooting your engagement photos”	Special	Other
4.) “upcoming year”	Temporal	Special
5.) “Halloween party”	Special	Celebratory

Figure 4: Examples of errors in the assignment of types to occasions.

nal error is a common example of confusion dealing with celebrations taking place as part of a special occasion. The gold standard annotations label these as special occasions whereas the system mostly identifies them as celebratory.

6 Conclusion

In this paper we introduce a methodology for extracting and categorizing occasions in which a product is used or with which a product is associated. We focus primarily on product descriptions, product reviews, and forum posts which are comments or reviews about a product. Occasions are categorized as one of six types: *Celebratory*, *Special*, *Seasonal*, *Temporal*, *Weather-Related*, and *Other*. Extraction and categorization are treated as separate tasks with extraction casted as a BIO encoded sequence labeling problem and categorization as a short text classification problem.

We examine the use of a MEMM and CRF for extracting occasions and find that the CRF model outperforms the MEMM. Categorization is cast as six-class classification problem with a support vector machine used to predict the best type. Categorization results in a macro-averaged F1-measure of 88.5%.

In the future, we plan to identify the relation between products/attributes and occasions and between two occasions. We envision product-occasion relations to include usage and procurement and relations between two occasions to include standard event relations, such as causation. We also plan to increase the amount of training data including multiple new product domains. With the addition of new training data we will also expand upon the current set of six occasions types. In particular, we

will examine the use of topic models, such as Latent Dirichlet Allocation, to split the “Other” category into multiple topically relevant ones. We posit that while there exists a set of core occasion categories the vast majority are domain-dependent.

References

- ACE. 2005. The ace 2005 (ace05) evaluation plan. <http://www.nist.gov/speech/tests/ace/ace05/doc/ace05-evalplan.v3.pdf>.
- Christoph Benz Müller and Adam Pease. 2012. Higher-order aspects and context in SUMO. *Journal of Web Semantics (Special Issue on Reasoning with context in the Semantic Web)*, 12-13:104–117.
- Alexander Boettcher and Dongman Lee. 2012. Eventradar: A real-time local event detection scheme using twitter stream. In *GreenCom’12*, pages 358–367.
- David B Bracewell, Marc Tomlinson, Ying Shi, Jeremy Bensley, and Mary Draper. 2011. Who’s playing well with others: Determining collegiality in text. In *Semantic Computing (ICSC), 2011 Fifth IEEE International Conference on*, pages 21–26. IEEE.
- David B Bracewell, Marc T Tomlinson, Mary Brunson, Jesse Plymale, Jiajun Bracewell, et al. 2012. Annotation of adversarial and collegial social actions in discourse. In *Proceedings of the Sixth Linguistic Annotation Workshop*, pages 184–192. Association for Computational Linguistics.
- David B Bracewell, David Hinote, and Sean Monahan. 2014. The author perspective model for classifying deontic modality in events. In *The Twenty-Seventh International Flairs Conference*.
- Caroline Brun, Diana Nicoleta Popa, and Claude Roux. 2014. Xrce: Hybrid classification for aspect-based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 838–842, Dublin, Ireland, August. Association for Computational Linguistics and Dublin City University.
- Erik Cambria and Amir Hussain. 2012. *Sentic computing*. Springer.
- Erik Cambria, Daniel Olsher, and Dheeraj Rajagopal. 2014. Senticnet 3: a common and common-sense knowledge base for cognition-driven sentiment analysis. In *Twenty-eighth AAAI conference on artificial intelligence*.
- Maryna Chernyshevich. 2014. Ihs r&d belarus: Cross-domain extraction of product features using crf. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 309–313, Dublin, Ireland, August. Association for Computational Linguistics and Dublin City University.

- Karthik Dinakar, Roi Reichart, and Henry Lieberman. 2011. Modeling the detection of textual cyberbullying. In *The Social Mobile Web*.
- Luca Dini and Giampaolo Mazzini. 2002. Opinion classification through information extraction. In *Proceedings of the Conference on Data Mining Methods and Databases for Engineering, Finance and Other Fields (Data Mining)*, pages 299–310.
- Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874.
- Jared Jashinsky, Scott H Burton, Carl L Hanson, Josh West, Christophe Giraud-Carrier, Michael D Barnes, and Trenton Argyle. 2014. Tracking suicide risk factors through twitter in the us. *Crisis*.
- Hiroshi Kanayama and Tetsuya Nasukawa. 2008. Textual demand analysis: Detection of users’ wants and needs from opinions. In *Proceedings of the 22Nd International Conference on Computational Linguistics - Volume 1, COLING ’08*, pages 409–416. Association for Computational Linguistics.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML ’01*, pages 282–289, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Ryong Lee and Kazutoshi Sumiya. 2010. Measuring geographical regularities of crowd behaviors for twitter-based geo-social event detection. In *Proceedings of the 2nd ACM SIGSPATIAL international workshop on location based social networks*, pages 1–10. ACM.
- Jiwei Li, Alan Ritter, Claire Cardie, and Eduard Hovy. 2014. Major life event extraction from twitter based on congratulations/condolences speech acts. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1997–2007, Doha, Qatar, October. Association for Computational Linguistics.
- Barbara Loken and James Ward. 1990. Alternative approaches to understanding the determinants of typicality. *Journal of Consumer Research*, pages 111–126.
- Barbara Loken, Lawrence W Barsalou, and Christopher Joiner. 2008. Categorization theory and research in consumer psychology. *Handbook of consumer psychology*, pages 133–65.
- Andrew McCallum, Dayne Freitag, and Fernando CN Pereira. 2000. Maximum entropy markov models for information extraction and segmentation. In *ICML*, pages 591–598.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Sean Monahan and Mary Brunson. 2014. Qualities of eventiveness. In *Proceedings of the Second Workshop on EVENTS: Definition, Detection, Coreference, and Representation*, pages 59–67, Baltimore, Maryland, USA, June. Association for Computational Linguistics.
- Alessandro Moschitti, Siddharth Patwardhan, and Chris Welty. 2013. Long-distance time-event relation extraction. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 1330–1338. Asian Federation of Natural Language Processing.
- Naoaki Okazaki. 2007. Crfsuite: a fast implementation of conditional random fields (crfs).
- Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment classification using machine learning techniques. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 79–86.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. Semeval-2014 task 4: Aspect based sentiment analysis. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 27–35, Dublin, Ireland, August. Association for Computational Linguistics and Dublin City University.
- Soujanya Poria, Erik Cambria, Lun-Wei Ku, Chen Gui, and Alexander Gelbukh. 2014. A rule-based approach to aspect extraction from product reviews. *SocialNLP 2014*, page 28.
- James Pustejovsky, José Castaño, Robert Ingria, Roser Saurí, Robert Gaizauskas, Andrea Setzer, and Graham Katz. 2003. Timeml: Robust specification of event and temporal expressions in text. In *in Fifth International Workshop on Computational Semantics (IWCS-5)*.
- Duangmanee (Pew) Putthividhya and Junling Hu. 2011. Bootstrapped named entity recognition for product attribute extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP ’11*, pages 1557–1567. Association for Computational Linguistics.
- J. Ramanand, Krishna Bhavsar, and Niranjan Pedanekar. 2010. Wishful thinking: Finding suggestions and ’buy’ wishes from product reviews. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, CAAGET ’10*, pages 54–61. Association for Computational Linguistics.
- Eleanor H Rosch. 1973. Natural categories. *Cognitive psychology*, 4(3):328–350.
- Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake shakes twitter users: real-time event

- detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, pages 851–860. ACM.
- Andrew N Smith, Eileen Fischer, and Chen Yongjian. 2012. How does brand-related user-generated content differ across youtube, facebook, and twitter? *Journal of Interactive Marketing*, 26(2):102–113.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA, October. Association for Computational Linguistics.
- Michael Speriosu, Travis Brown, Taesun Moon, Jason Baldrige, and Katrin Erk. 2010. Connecting language and geography with region-topic models. *Models of Spatial Language Interpretation at Spatial Cognition 2010 (COSLI-2010)*, 2010.
- Bharath Sriram, Dave Fuhry, Engin Demir, Hakan Ferhatosmanoglu, and Murat Demirbas. 2010. Short text classification in twitter to improve information filtering. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 841–842. ACM.
- Emilia Stoica, Marti A Hearst, and Megan Richardson. 2007. Automating creation of hierarchical faceted metadata structures. In *HLT-NAACL*, pages 244–251.
- Zhiqiang Toh and Wenting Wang. 2014. Dlirec: Aspect term extraction and term polarity classification system. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 235–240, Dublin, Ireland, August. Association for Computational Linguistics and Dublin City University.
- Marc T Tomlinson, David B Bracewell, Mary Draper, Zewar Almissour, Ying Shi, and Jeremy Bensley. 2012. Pursing power in arabic on-line discussion forums. In *LREC*, pages 1359–1364.
- Sarah Vieweg, Amanda L Hughes, Kate Starbird, and Leysia Palen. 2010. Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1079–1088. ACM.
- Joachim Wagner, Piyush Arora, Santiago Cortes, Utsab Barman, Dasha Bogdanova, Jennifer Foster, and Lamia Tounsi. 2014. Dcu: Aspect-based polarity classification for semeval task 4. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 223–229, Dublin, Ireland, August. Association for Computational Linguistics and Dublin City University.
- Michael Wiegand and Dietrich Klakow. 2014. Separating brands from types: an investigation of different features for the food domain. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2291–2302, Dublin, Ireland, August. Dublin City University and Association for Computational Linguistics.
- Jianxing Yu, Zheng-Jun Zha, Meng Wang, Kai Wang, and Tat-Seng Chua. 2011. Domain-assisted product aspect hierarchy generation: towards hierarchical organization of unstructured consumer reviews. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 140–150. Association for Computational Linguistics.