Why is an Event Affective? Classifying Affective Events based on Human Needs

Haibo Ding and **Tianyu Jiang** and **Ellen Riloff**
School of Computing
University of Utah
Salt Lake City, UT 84112
{hbding, tianyu, riloff}@cs.utah.edu

Abstract

Affective events are events that impact people in positive or negative ways. When people discuss an event, people understand not only the affective polarity but also the reason for the event being positive or negative. In this paper, we aim to categorize affective events based on the reasons why events are affective. We propose that an event is affective to people often because the event describes or indicates the satisfaction or violation of certain kind of human needs. For example, the event "I broke my leg" affects people negatively because the need to be physically healthy is violated. "I play computer games" has a positive affect on people because the need to have fun is probably satisfied. To categorize affective events in narrative human language, we define seven common human need categories and introduce a new data set of randomly sampled affective events with manual human need annotations. In addition, we explored two types of methods: a LIWC lexicon based method and supervised classifiers to automatically categorize affective event expressions with respect to human needs. Experiments show that these methods achieved moderate performance on this task.

Introduction

Affective events are activities or states that are typically positive (desirable) or negative (undesirable). For example, "I went to Disneyland" and "I got engaged" are usually positive events. We could often infer that people who experience these events would usually have positive affective states. On the other hand, people who experience events such as "I got fired" and "I broke my leg" would often have negative affective states. Knowing the affective polarities of events plays an important role in various human language understanding tasks such as opinion analysis (Deng, Wiebe, and Choi 2014), sarcasm recognition (Riloff et al. 2013), and narrative text understanding (e.g., plot units (Goyal, Riloff, and Daumé III 2013; Lehnert 1981)).

However, when people discuss events, people understand not only the affective polarities of events but also the reason why an event is positive or negative. For example, though both "go to Disneyland" and "get engaged" are positive, the reasons are different. People feel happy about "go to Disneyland" because it's a fun/entertaining activity. Whereas the reason why "get engaged" is desirable for people is mainly about a social relationship. Events are also negatively affective due to various reasons. For example, "got fired" is undesirable because people would not have financial income. However, "broke a leg" is negative mainly because it causes damage to one's body. Understanding the reasons for events being affective would enable us to obtain deep understanding of narrative stories (Goyal, Riloff, and Daumé III 2013; Lehnert 1981). For example, by analyzing the affective polarities and the corresponding reasons we could infer the protagonist's motivations and plans (Schank and Abelson 1977) in a narrative story. In addition, automatically recognizing the reasons for events being affective could help improve the effectiveness of automatic psychotherapy methods by identifying the potential causes for certain mental issues.

However, this problem has not received much attention in the field of Natural Language Processing (NLP). In our work, we aim to categorize affective events based on their human needs, which is inspired by the intuition that human beings' actions or behaviors are often motivated by basic human needs. An event is positive often because its implied human needs are satisfied. On the contrary, an event is undesirable for people usually because human needs are violated. Inspired by previous human need theories, i.e., Maslow's Hierarchy of Needs (Maslow et al. 1970) and Fundamental Human Needs (Max-Neef, Elizalde, and Hopenhayn 1991), we propose seven common categories of human needs to categorize affective events: *Physiological Needs*, *Health Needs*, *Leisure Needs*, *Social Needs*, *Financial Needs*, *Cognition Needs*, *Freedom Needs*. We also define another two categories: *Emotions/Sentiments/Opinions* to capture explicit emotions and opinions, and *None of the Above* for events that can not be categorized into the human need categories.

We conduct a manual annotation study on a set of affective (positive or negative) events, which were manually annotated with affective polarity by (Ding and Riloff 2018). In our study, annotators are asked to annotate each affective event using the defined human need categories. Our analysis shows that the majority of the affective events could be manually categorized into the seven human need categories with good inter-annotator agreements (Cohen's $\kappa \geq$.65). The resulting annotated events establish a new data set on this task of understanding the reasons for events being affective. We plan to make this new data set freely available.

We formalize the problem of recognizing the reasons for events being affective as a multi-class classification task which is to categorize an affective event into one of the human need categories. With the newly annotated data, we create two types of methods including a LIWC (Pennebaker, Booth, and Francis 2007) based method and supervised classifiers to automatically categorize affective events. For supervised classifiers, we experiment with SVM and logistic regression models using Ngram and word embedding features. Our experiments show that the logistic regression classifier with word embeddings achieved the best performance (54.8 average F1) on our data.

Related Work

Recently, there has been growing interest in recognizing the affective polarities of events which are typically objective, factual statements. For example, Goyal, Riloff, and Daumé (2013) developed a bootstrapped learning method to learn *patient polarity verbs*, which impart affective polarities to their patients. Connotation frames (Rashkin, Singh, and Choi 2016) were recently proposed to incorporate the connotative polarities for a verb's arguments from the writer's and other event entity's perspectives. Another group of researchers have studied +/- effect events (Deng, Choi, and Wiebe 2013; Choi and Wiebe 2014) which they previously called benefactive/malefactive events. Their work mainly focused on inferring implicit opinions through implicature rules (Deng and Wiebe 2014; 2015). In addition, Ding and Riloff (2016) designed an event context graph model to harvest affective events. Lately Ding and Riloff (2018) developed a semantic consistency model to recognize a large set of affective events using three types of semantic relations. They also studied the prevalence of affective events in a set of randomly sampled events. In our work, we used their annotated affective events for human need analysis. Recently, some researchers also demonstrated that automatically acquired patterns could benefit the recognition of first-person related affective sentences (Reed et al. 2017). Previous work only predicts affective polarities of events. Instead, our work aims at analyzing the reason for events being affective.

Our human need categories are inspired by two previous theories. The first is Maslow's hierarchy of needs (Maslow et al. 1970) which was developed to study people's motivations and personalities. The second is Fundamental Human Needs (Max-Neef, Elizalde, and Hopenhayn 1991) which was developed to help communities identify their strengths and weaknesses. Inspired by these theories, we provide detailed definitions and examples for seven common human need categories to understand the reasons for events being affective in human language.

Our human need categories are also related to the concept of "goals", which has been proposed to analyze narrative stories (Schank and Abelson 1977). They proposed a taxonomy of 7 types of goals including Satisfaction Goals, Enjoyment Goals, Achievement Goals, Preservation Goals, Crisis Goals, Instrumental Goals, and Delta Goals (Schank and Abelson 1977). Goals could be very specific to a character in a particular narrative story, for example "terrorists have a goal to kill people". However, human needs are usually common sense that are universal, shared by most people (Max-Neef, Elizalde, and Hopenhayn 1991). In addition, our work is also related to the research on wish detection (Goldberg et al. 2009) and desire fulfillment understanding (Chaturvedi, Goldwasser, and Daumé III 2016).

Human Need Categories

By comprehending a narrative story, we understand not only the affective polarity of the events in it, but also the reason for the events being affective (positive or negative). With the affective polarity, we understand how the story characters are affected by different events. Knowing the reasons enables us to understand why they are affected in that way. In this section, we define seven categories of human needs in addition to two other categories. Detailed definitions and examples are explained as follows.

Physiological Needs

Physiological needs are the basic needs that must be satisfied to maintain our body's basic functions. For example, we need to eat food to grow, our body needs sleep, etc. An event in this category is affective often because it describes the satisfaction or violation of a physiological need. For example, "I am hungry" is negative because the need to have food is not satisfied.

More specifically, the physiological needs include (1) the need to be able to breathe beneficial or pleasant air, and to avoid unpleasant air; (2) the need to avoid hunger, to avoid unpleasant food, and to eat or obtain pleasing food; (3) the need to avoid thirst, to avoid unpleasant beverages, and to drink or obtain pleasing beverages; (4) the need to sleep, regularly and comfortably; (5) the need to maintain warmth of the human body, to not be too hot or too cold; (6) the need to have or obtain shelter (i.e., a place to live or stay) and to avoid unpleasant shelters. If an event is affective, and describes or indicates the satisfaction or violation of these needs, then it belongs to this category.

Examples

- "*I have not eaten for 2 days*" is negative because the need to have food is violated.
- "*I woke up at 2am*" is negative because the speaker violated the need of having enough sleep or sleeping soundly.
- "*I ate cake*" is positive because the need of having enjoyable food is satisfied.
- "*I bought a house*" is positive because the need of owning a shelter is satisfied.

Physical Health and Safety Needs

In our daily lives, many events will harm or endanger our physical health and safety. If we experience these events we would be affected negatively. However, some other events will improve our physical health or safety conditions, which will affect us in positive ways when we experience them. These events are affective to us because they satisfy or violate the needs to be physically healthy and safe. Affective events in this category could be related to health problems, body injuries, exercise, etc.

Examples
- "*My head hurts*" is negative because the need to be physically healthy is violated.
- "*I do exercise*" is positive because exercise is associated with good health.
- "*I was kidnapped*" is negative because the speaker is concerned about his/her physical health and safety.

Leisure and Aesthetic Needs

Another type of need that most people share is the need to have fun, to be relaxed, to have leisure time. For example, people would often feel happy when they "play games". But for some other events, people would be affected negatively because they prevent people from having leisure time. For example, "work on holidays" is negative because people expect to have fun or leisure time on holidays. In addition, people have needs to appreciate and enjoy the beauty of certain things because people will often feel relaxed, joyful, or peaceful when these needs are satisfied. Since leisure needs and aesthetic needs are closely related we combine them together as a single category.

Leisure and aesthetic needs mainly include: (1) the need to have entertaining or fun activities, to avoid the lack of fun or entertaining activities; (2) the need to have leisure, to avoid too much work because it detracts from leisure time; (3) the need to have an enjoyable, pleasant environment; (4) the need to pursue and appreciate the beauty of nature, art, music and other aesthetically beautiful things. If an affective event describes or indicates the satisfaction or violation of these needs, then it belongs to this category.

Examples
- "*I play computer games*" is positive because it describes a fun activity.
- "*I work on Christmas Day* " is negative because it is typically unenjoyable and detracts from leisure time.
- "*room is noisy* " is negative because the environment is undesirable
- "*I saw a rainbow*" is positive because the need to appreciate beauty is satisfied.

Social, Self-Worth, and Self-Esteem Needs

According to Aristotle, human beings by nature are "social animals", which means that people naturally have needs to have close relationships with family and friends, and also enjoyable relations with the general public. Events indicating the satisfaction of having good family or friend relations make people feel good, otherwise they would impact people negatively. As a member of society, people also have needs to have and improve self-worth and self-esteem, and to be respected by others.

In our work, we group these social needs together as a single human need category. Specifically, these needs include: (1) the need to have family, to have close family relations, to avoid damaging family relations; (2) the need to have friendships, to avoid damaging friendships; (3) the need to maintain pleasant social relations with the general public, to avoid conflicts and arguments; (4) the need to maintain socially and culturally acceptable behaviors; (5) the need to realize and improve one's self-worth, to be recognized by others; (6) the need to maintain and improve self-esteem or dignity.

Examples
- "*My mom visited me*" is positive because a family relationship is maintained.
- "*I have many friends*" is positive because the friendship need is satisfied.
- "*Nobody talks to me*" is negative because social relations with others are not good.
- "*They mock me*" is negative because the speaker's self-esteem/dignity is hurt.

Finances, Possessions, and Job Needs

In our daily lives, many affective events involve earning money, having well-paid jobs, obtaining useful or valuable possessions. When these events happen to us, we would usually have positive feelings. However, if we experienced affective events like losing money, getting fired from a job, or losing some valuable possessions, we would usually be affected negatively. More specifically, this category includes: (1) the need to obtain and protect financial income; (2) the need to acquire possessions and maintain good condition of one's possessions; (3) the need to have a job and satisfying work because jobs are usually reliable sources to obtain financial income. We define that if the possession in an event is more directly related to another type of need, we prefer to categorize the event in that human need category. For example, "I bought steaks" mainly satisfies the physiological needs because the purpose of "steaks" is for eating.

Examples
- "*I got a lot of money*" is positive because the need to have financial income is satisfied.
- "*I bought a computer*" is positive because the need to obtain useful tools is satisfied.
- "*I lost my wallet*" is negative because the need to protect possessions is violated.
- "*I got fired*" is negative because the need to have a job has been violated.

Cognition and Education Needs

Another group of human needs are cognition and education needs, which motivate people to learn skills, obtain information, understand meanings, improve cognitive abilities, etc. When such needs are satisfied (e.g., "learned new skills") people would often feel good, otherwise people are affected negatively (e.g., "I did not get the Master degree"). More specifically, this group of needs includes: (1) the need to obtain skills, information, and knowledge, to receive education, and to improve one's intelligence; (2) the need to mentally process information correctly (e.g., remembering, calculation), and to have good cognitive abilities.

Examples
- "*I learned to mow the lawn*" is positive because the speaker learned a skill.
- "*I graduated*" is positive because the need to receive education is satisfied.

◦ "*I overestimated the number*" is negative because the speaker did not process information correctly.

Freedom of Movement and Accessibility Needs

There are many affective events that describe our movement or accessibility situations. When we can not move freely, or access something in a timely manner, then we are affected negatively. Specifically, these type of needs include: (1) the need to move or change positions freely; (2) the need to access things or services in a timely manner.

Examples

◦ "*I have been waiting for 5 hours*" is negative because the need to access things in a timely manner is not satisfied.

◦ "*I was stuck in my car*" is negative because the need to move freely is violated.

Emotions, Sentiments, and Opinions

There are also many affective events or state expressions in human language that do not belong to any human need categories, but only describe people's affective states such as sentiments, emotions, or opinions. We define a separate category to group these affective events together, which includes (1) events that directly describe experiencers' sentiments, emotions, feelings, or physical expressions of emotions; and (2) events that express an opinion towards an object. We also define that if an event both expresses a sentiment/emotion and is also related to a previous human need category, we prefer to categorize it into the corresponding human need category. For example, "I love my family" not only expresses a positive emotion, but also indicates a good social relationship. Therefore, we categorize it into the *Social Needs* category.

Examples

◦ (a) "*I am happy*" is positive because it describes a positive internal emotion state.

◦ (b) "*Canadians are good*" is positive because it describes a positive opinion.

None of the Above

There are some affective events that do not belong to any of previous human need categories or the *Emotions* category. We categorize all of these affective events in the *None of the Above* category which includes, but is not limited to (1) events or situations that are too general or abstract to be assigned to any of the other categories; or (2) the reason why the event is positive or negative falls into a different category than the ones listed above.

Examples

◦ "*I had a problem*" is negative, but we do not know the specific reason why.

◦ "*I got a mistake*" is negative but we do not know what the mistake is.

Affective Event Data Set

The goal of our research is to study the reasons for events being affective. Since affective events are the objects of study, we used the affective event data set created by (Ding and Riloff 2018) with affective polarity annotations. We augmented the affective event data set with human need category annotations. In this section, we first introduce the personal story corpus and representation method used for extracting events by (Ding and Riloff 2018), and then the manual annotations for affective polarity.

Personal Story Corpus

The events in the affective event data set were extracted from a personal story corpus, which was obtained by applying a personal story classifier (Gordon and Swanson 2009) on 177 million blog posts. All of the blog posts are from the ICWSM 2009 and 2011 Spinn3r data sets (Burton, Java, and Soboroff 2009; Burton, Kasch, and Soboroff 2011). After removing the near-duplicate stories and those with no first person mentions to make sure that the extracted stories are related to the first person (i.e., the author), the final personal story corpus contains 1,383,425 personal story blogs.

Event Representation and Extraction

Each event is represented using a frame-like event structure to capture the semantic meanings of various events. Each event representation contains four components: ⟨**Agent**, **Predicate**, **Theme**, **PP**⟩ which denote the agent, predicate, theme, and prepositional phrase (PP) of an event respectively. Event structures are extracted from parsed sentences using Stanford dependency relations (De Marneffe and Manning 2008). To extract events that are specific enough to distinguish between events with different meanings, an event structure is required to have a Predicate component, and also must have an Agent or a Theme (i.e., at least one of them).[1] The PP component is optional.

The **Predicate** is a simple verb phrase, which typically corresponds to an action or state. The Predicate component is required to contain a verb, and may also include a particle, infinitive verb, negation term "not", if they appear. For example, the Predicate could be "eat" or "not want to take off". The **Agent** component is composed of an **event entity** which could be a named entity, compound noun, or pronoun. The term "Theme" is used loosely to allow an event entity or adjective to fill the **Theme** component (e.g., "I am brilliant"). The event structure also includes a prepositional phrase (**PP**) formed with a preposition and an event entity, which can be essential to distinguish between dramatically different event types. For example, "I go to beach" is a very different kind of event than "I go to prison". However, without PP component the event representation will be "I go". Although multiple PPs are common and can be important, only a single PP is kept in the event representation to prevent the representation from becoming overly specific. If multiple PPs are attached to the Predicate, only the closest one is kept.

In addition, each event structure is normalized in two ways. First, all words in the event structures are lemmatized. Second, active and passive voice event expressions are normalized. For example, "I was killed by him" and "he killed me" are both represented by the structure: "⟨he, kill, me, -⟩".

[1] For "be", "have" verbs, both an Agent and Theme are required.

Physiological	Health	Leisure	Social	Finance	Cognition	Freedom	Emotion	None
19 (4%)	52 (10%)	75 (14%)	108 (20%)	29 (5%)	26 (5%)	7 (1%)	128 (24%)	98 (18%)

Table 1: Distribution of Human Need Categories (each cell shows the frequency and percentage).

To study events related to the experiencer (i.e., the blogger), only the events that satisfy at least one of the following criteria are extracted. (1) The event has a first person reference (e.g., "I", "my"). (2) The event mentions a family member (e.g., "mom"). These events are related to the blogger because the affective state of the blogger usually parallels that of family members (e.g., "mom is sick" is undesirable for both mom and the blogger). 92 family member terms were manually compiled to recognize family members in events. (3) The event does not mention any other people[2]. In this case, the event pertains to the blogger (e.g., "the computer died"). Events that only mention other people are not extracted because they may be describing someone else's experience, not the blogger's.

To create the event structures, the personal story corpus was first preprocessed using StanfordCoreNLP (Manning et al. 2014) for POS and NER tagging, and SyntaxNet (Andor et al. 2016) for parsing. Then event structures were extracted from the parsed sentences. This resulted in 19,794,187 unique events. After filtering events with frequency < 5, the *event data set* totally contains 571,424 unique events.

Affective Polarity Annotations

To obtain manual affective polarity annotations, a set of 1,500 randomly sampled events were given to three human annotators, who were asked to manually assign each event with a polarity label from four polarity categories including Positive, Negative, Neutral, and Mixed (which could be positive or negative depending on different meanings). The pairwise inter-annotator agreemnts (Cohen's kappa κ) were κ=.76 κ=.70, and κ=.69. Then the majority label was assigned to each event as the gold standard polarity label. Only one event was labeled as Mixed, so the Mixed class was abandoned. In addition, there are 9 events that received three different labels from the annotators, which were also discarded. This resulted in a gold standard data set of 1,490 events, among which 295 (20%) events are labeled as positive, 264 (18%) are negative, and 931 (62%) are neutral. Through this process, 559 (38%) events were annotated as affective, which were used for the human needs analysis described below.

Annotation and Analysis of Human Needs

We hypothesized that the reasons for events being affective are often explained using human needs. However, two questions remain to be answered. First, how many affective events could be classified into our seven human need categories? Second, can human annotators consistently agree on the appropriate human need category to each affective event?

To answer these two questions, we added human needs annotations to the affective event data set which contains 559 affective events. We asked three human annotators to assign the most appropriate category label to each affective event. We measured their pairwise inter-annotator agreements using Cohen's kappa, which were κ=.69, κ=.66 and κ=.65. This demonstrates that human annotators could achieve good agreements ($\kappa \geq$.65) on this task. We then assigned the majority category label to each event as the gold category label. We found that 17 affective events were assigned with three different labels, which were discarded because annotators did not agree on these cases. We ended up with a gold standard data set of 542 affective events with human need category labels. Some affective events are shown in Table 2.

Positive Events	Human Need
⟨ I, take, advantage, of breakfast⟩	Physiological
⟨ ear, be, better, - ⟩	Health
⟨ I, watch, Hellboy II, -⟩	Leisure
⟨ we, get, marry, -⟩	Social
⟨ I, get, my new laptop, -⟩	Finance
⟨ my memory, be, vivid, -⟩	Cognition
⟨ my heart, feel, happy, -⟩	Emotion
⟨ we, be, legal, -⟩	None
Negative Events	**Human Need**
⟨ I, grow, hungry, - ⟩	Physiological
⟨ my face, look, pale, - ⟩	Health
⟨ -, rain out, game, -⟩	Leisure
⟨ girl, laugh, -, at me⟩	Social
⟨ house phone, not work, -, -⟩	Finance
⟨ I, lose, attention, -⟩	Cognition
⟨ I, be, scared, -⟩	Emotion
⟨ it, not work, -, for me⟩	None

Table 2: Affective Event Examples with Human Needs.

The distribution of human need categories is shown in Table 1. Table 1 shows that the majority (58%) of the affective events could be categorized into the seven human need categories. There are 24% affective events that are simply sentiment or emotion expressions such as "I'm happy", "I hate it", etc. Table 1 also shows that 18% of affective events cannot be categorized into any previous category. We conducted a further analysis on these affective events, and found three main reasons for events belonging to the None category. First, some events describe very abstract positive or negative situations where we know the polarity but not the specific reason (e.g., "we have trouble", "it does not work for me"). Second, some meanings are not clear which may be caused by parsing or extraction errors (e.g., "bit beats up", "I fell things"). Third, some events are affective because of other

[2]An entity is identified as "other people" if it's second or third person pronoun, a PERSON Named Entity, or nominal person mention based on WordNet (e.g., "plumber").

nuanced human need categories. For example, "I'm powerless" is negative because we have the need to be strong, to have power and authority.

The distribution in Table 1 also shows that the Freedom category (i.e., Freedom of Movement and Accessibility) only contains 7 instances. We concluded that this class is too small to provide sufficient evaluation data. So, we merged the Freedom category into the None category in our subsequent study and experiments.

A1 \ A2	Phy	Hlth	Leis	Socl	Fnc	Cog	Emo	None	#Tot
Phy	15	3	0	0	0	0	1	1	20
Hlth	2	36	0	5	0	0	2	3	48
Leis	2	1	60	9	4	1	5	6	88
Socl	0	5	1	85	0	0	1	6	98
Fnc	2	0	0	2	22	0	0	2	28
Cog	0	1	1	4	0	20	3	3	32
Emo	0	3	2	9	2	3	90	12	121
None	2	3	4	13	5	2	22	73	124
#Tot	23	52	68	127	33	26	124	106	559

Figure 1: Confusions between Two Annotators.

Figure 1 is the annotation confusion matrix [3] between two human annotators (A1 and A2) whose agreement is κ=0.66. Each row denotes the number of affective events that were annotated with a row label (e.g., "Phy") by annotator A1. Each column denotes the number of affective events that were annotated with a column label (e.g., "Hlth") by annotator A2. For example, the cell ("Phy", "Hlth") denotes that there are 3 affective events that were labeled "Phy" by A1 but "Hlth" by A2. The #Tot denotes the sum of each row or column. We notice that human annotators often confused Emotion and None. In addition, human annotators had disagreements about Social and Leisure.

Automatic Methods for Affective Events Categorization based on Human Needs

Categorizing affective events according to their human needs is a new natural language understanding task. We propose two types of methods to assess the difficulty of this task. First, we designed a rule-based method to infer the human need category of an affective event using the LIWC lexicon. Second, we created several supervised classification models to predict the human needs. Details of these two types of methods are presented below.

LIWC Lexicon based Method

The LIWC lexicon (Pennebaker, Booth, and Francis 2007) is a dictionary of words associated with various lexical categories. Besides pronominal and emotion categories, it also contains cognition and psychology categories of words, which are closely related to our task. To effectively explore this lexicon, we first manually analyzed the categories in LIWC, and built a mapping from LIWC categories to our human need categories, which is shown in Table 3. To predict the human need category of an event, we designed a voting based system which first looks up the LIWC category of each word in an event, maps it to our human need category, and then uses the majority category[4] across all words in the event expression as the final human need category. If none of the words in an event are contained in LIWC, or their categories cannot be mapped to our categories, then we assign a None label to that event.

LIWC		**Human Need**	**LIWC**		**Human Need**
Ingest	$\rightarrow$	Physiological	Social	$\rightarrow$	Social
Health	$\rightarrow$	Health	Work	$\rightarrow$	Finance
Body	$\rightarrow$	Health	Money	$\rightarrow$	Finance
Death	$\rightarrow$	Health	Insight	$\rightarrow$	Cognition
Leisure	$\rightarrow$	Leisure	Inhib	$\rightarrow$	Cognition
Affect	$\rightarrow$	Emotion			

Table 3: Mapping from LIWC to Human Need Categories.

Supervised Classifiers

Our task of categorizing affective events based on human needs is a multi-class classification task. We propose to experiment with two types of multi-class classification strategies. First, we trained a one-vs.-rest (ovr) binary classifier for each category. For prediction, if an instance is labeled with multiple labels, we select the most confident one as its category label. In our experiments, we used two base classifiers: the logistic regression classifier (**Logit**$^{\mathbf{ovr}}$) from scikit-learn (Pedregosa et al. 2011) and the linear SVM classifier (**SVM**$^{\mathbf{ovr}}$) from LIBSVM(Chang and Lin 2011). In addition, instead of decomposing a multi-class classification problem into multiple binary classifiation tasks, we explored two classification models that train a single multi-class classifier directly. The first one is a multinomial logistic regression model (**Logit**$^{\mathbf{multi}}$) which is also called softmax classifier. The second one is a multi-class SVM (**SVM**$^{\mathbf{multi}}$) (Crammer and Singer 2001). We used the implementation from the Liblinear library (Fan et al. 2008).

For our supervised classifiers, we experimented with two types of features: Ngram features and event embedding features as described below.

Ngram Features We used the lemmatized words in an event structure as its Unigram features.

Event Embedding Features We also evaluated the effectiveness of event embeddings on our task. For each event, we compute its embedding as the average of its words' embeddings. In our experiments, we used the 200 dimension word embeddings which were pre-trained on 27 billion tweets using GloVe (Pennington, Socher, and Manning 2014).

Evaluation

In this section, we evaluate the methods described above. For supervised classifiers, the results are averages based on

[3] The abbreviations for each category are Physiological (Phy), Health (Hlth), Leisure (Leis), Social (Socl), Finance (Fnc), Cognition (Cog), Emotion (Emo).

[4] For ties, we remove a component one by one in the order of Agent, PP, Theme until we obtain a majority label.

3-fold cross-validation. Since LIWC based method is unsupervised, we did not use the training sets. We evaluated the LIWC based method on each of the test sets across 3-folds, and the average results are reported.

		Precision	Recall	F1
LIWC		47.7	39.0	38.6
Supervised Classifiers				
Feature	Method			
Unigram	Logit^{ovr}	33.6	28.7	27.3
	Logit^{multi}	40.4	31.0	30.5
	SVM^{multi}	50.4	40.5	42.2
	SVM^{ovr}	52.3	43.1	44.8
EventEmb	SVM^{multi}	50.0	49.9	49.3
	SVM^{ovr}	51.3	50.7	50.5
	Logit^{multi}	61.7	**51.7**	54.5
	Logit^{ovr}	**64.2**	**51.7**	**54.8**

Table 4: Affective Events Categorization Results.

Table 4 shows the average precision, recall, and F1 scores. The first row denotes the performance of LIWC, which shows that using LIWC we can recognize the human needs of affective events with 39% recall and 47.7% precision. One issue is that some words in LIWC categories do not perfectly correspond to our human need categories. For example, "abandon" and "damage" are categorized in the Affect category in LIWC, but they actually do not belong to our Emotion category because they do not express sentiments or emotions directly, though they imply sentiments.

The following rows show the performance of classifiers with Unigram and event embedding features. We notice that event embedding features achieve much better performance than Unigram features regardless of the classification models, and one-vs.-rest based methods obtain better performance than multi-class models in most cases. In our experiments, we also tried Bigram features and non-linear kernels for SVM, but they all performed worse. We also trained classifiers using both Unigram and event embedding features, but we did not obtain better performance. The reason could be that the classifiers may overfit the training data when using more features.

Our best performance is achieved by the **Logit^{ovr}** using the event embedding features which obtains 54.8% average F1 score on our data set. Table 5 shows the precision (Pre), recall (Rec), and F1 for each category by the best system. The results show that our best system achieved $\geq$ 60% F1 on four categories, and 40% to 49% on other three categories, which indicates that this task is difficult. Table 5 also shows that we achieved higher precision than recall for most of the human need categories (e.g., Health, Finance, Cognition), which indicates that we could improve the overall performance by improving the recognition coverage in future work.

Figure 2 shows the confusions between the predictions of the best method **Logit^{ovr}** and gold annotations. Each cell shows the sum of confusions over the 3-folds of cross-validation. Similar to humans, the system often confuses Emotion and None, and Emotion and Social. The system also has difficulty distinguishing Social from Health and Leisure.

Physiological			**Health**			**Leisure**			**Social**		
Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1
82	57	67	65	40	49	62	59	60	61	72	66
Finance			**Cognition**			**Emotion**			**None**		
Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1
61	31	40	75	31	42	60	75	66	47	49	48

Table 5: Precision, Recall, F1 of the Best Logit^{ovr} Method.

Pred. \ Gold	Phy	Hlth	Leis	Socl	Fnc	Cog	Emo	None	#Tot
Phy	11	0	0	0	0	0	0	2	13
Hlth	1	21	1	0	1	1	5	4	34
Leis	1	1	44	6	2	2	3	12	71
Socl	1	8	6	78	3	4	14	14	128
Fnc	1	1	2	0	9	0	1	2	16
Cog	0	0	0	1	0	8	1	2	12
Emo	2	11	10	16	4	3	96	18	160
None	2	10	12	7	10	8	8	51	108
#Tot	19	52	75	108	29	26	128	105	542

Figure 2: Confusions between Predictions and Gold Labels.

Conclusion

In this work, we studied the reason for events being affective. We proposed that an event was affective mainly because of the satisfaction or violation of certain kinds of human needs. We defined a set of human need categories to explain the affect of events. We also manually added manual annotations of human need categories to a previous collection of affective events. We demonstrated that the majority of affective events could be categorized into our human need categories, and human annotators can identify the human need category for an affective event with good annotation agreement. We plan to make the manually annotated events freely available to encourage future research in this direction.

In addition, we formalized the problem of recognizing the reason for an event being affective as a multi-class classification task. We evaluated two types of methods: rule-based system using LIWC, and supervised classifiers. Our experimental results showed that these methods achieved moderate performance on our data set. More future research work is needed to improve the performance of categorizing affective events based on their human needs.

Acknowledgements

This material is based in part upon work supported by the National Science Foundation under Grant Number IIS-1619394. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Andor, D.; Alberti, C.; Weiss, D.; Severyn, A.; Presta, A.; Ganchev, K.; Petrov, S.; and Collins, M. 2016. Globally normalized transition-based neural networks. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.

Burton, K.; Java, A.; and Soboroff, I. 2009. The ICWSM 2009 Spinn3r dataset. In *Third Annual Conference on Weblogs and Social Media*.

Burton, K.; Kasch, N.; and Soboroff, I. 2011. The ICWSM 2011 Spinn3r dataset. In *Proceedings of the Annual Conference on Weblogs and Social Media*.

Chang, C.-C., and Lin, C.-J. 2011. LIBSVM: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)* 2(3):27.

Chaturvedi, S.; Goldwasser, D.; and Daumé III, H. 2016. Ask, and shall you receive? understanding desire fulfillment in natural language text. In *Processings of the AAAI Conference on Artificial Intelligence*.

Choi, Y., and Wiebe, J. 2014. +/-EffectWordNet: Sense-level lexicon acquisition for opinion inference. In *Proceedings of the 2014 Conference on Empirical Methods on Natural Language Processing*.

Crammer, K., and Singer, Y. 2001. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of machine learning research* 2(Dec):265–292.

De Marneffe, M.-C., and Manning, C. D. 2008. Stanford typed dependencies manual. Technical report, Stanford University.

Deng, L., and Wiebe, J. 2014. Sentiment propagation via implicature constraints. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*.

Deng, L., and Wiebe, J. 2015. Joint prediction for entity/event-level sentiment analysis using probabilistic soft logic models. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.

Deng, L.; Choi, Y.; and Wiebe, J. 2013. Benefactive/malefactive event and writer attitude annotation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*.

Deng, L.; Wiebe, J.; and Choi, Y. 2014. Joint inference and disambiguation of implicit sentiments via implicature constraints. In *Proceedings of the 25th International Conference on Computational Linguistics*.

Ding, H., and Riloff, E. 2016. Acquiring knowledge of affective events from blogs using label propagation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Ding, H., and Riloff, E. 2018. Weakly supervised induction of affective events by optimizing semantic consistency. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Fan, R.-E.; Chang, K.-W.; Hsieh, C.-J.; Wang, X.-R.; and Lin, C.-J. 2008. Liblinear: A library for large linear classification. *Journal of machine learning research* 9(Aug):1871–1874.

Goldberg, A. B.; Fillmore, N.; Andrzejewski, D.; Xu, Z.; Gibson, B.; and Zhu, X. 2009. May all your wishes come true: A study of wishes and how to recognize them. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*.

Gordon, A., and Swanson, R. 2009. Identifying personal stories in millions of weblog entries. In *Third International Conference on Weblogs and Social Media, Data Challenge Workshop*.

Goyal, A.; Riloff, E.; and Daumé III, H. 2013. A Computational Model for Plot Units. *Computational Intelligence* 29(3):466–488.

Lehnert, W. G. 1981. Plot units and narrative summarization. *Cognitive Science* 5(4):293–331.

Manning, C. D.; Surdeanu, M.; Bauer, J.; Finkel, J.; Bethard, S. J.; and McClosky, D. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics*.

Maslow, A. H.; Frager, R.; Fadiman, J.; McReynolds, C.; and Cox, R. 1970. *Motivation and personality*, volume 2. Harper & Row New York.

Max-Neef, M.; Elizalde, A.; and Hopenhayn, M. 1991. *Human Scale Development: Conception, Application and Further Reflections*. The Apex Press.

Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.

Pennebaker, J. W.; Booth, R. J.; and Francis, M. E. 2007. Linguistic inquiry and word count: LIWC2007. *Austin, TX: liwc.net*.

Pennington, J.; Socher, R.; and Manning, C. D. 2014. GloVe: Global vectors for word representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Rashkin, H.; Singh, S.; and Choi, Y. 2016. Connotation Frames: A data-driven investigation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.

Reed, L.; Wu, J.; Oraby, S.; Anand, P.; and Walker, M. A. 2017. Learning lexico-functional patterns for first-person affect. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.

Riloff, E.; Qadir, A.; Surve, P.; De Silva, L.; Gilbert, N.; and Huang, R. 2013. Sarcasm as contrast between a positive sentiment and negative situation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*.

Schank, R. C., and Abelson, R. P. 1977. *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. Hillsdale, NJ: L. Erlbaum.