Context-Aware Experience Extraction from Online Health Forums

Yunzhong Liu
Computer Science & Eng
Arizona State University
Tempe, AZ 85287, USA
Email: liuyz@asu.edu

Yi Chen
School of Management
New Jersey Institute of Technology
Newark, NJ 07103, USA
Email: yi.chen@njit.edu

Jiliang Tang
Yahoo! Labs
Yahoo! Inc
Sunnyvale, CA, USA
Email: jlt@yahoo-inc.com

Huan Liu
Computer Science & Eng
Arizona State University
Tempe, AZ 85287, USA
Email: huan.liu@asu.edu

***Abstract*—Online health forums provide a large repository for patients, caregivers, and researchers to seek valuable information. The extraction of patient-reported personal health experience from the forums has many important applications. For example, medical researchers can discover trustable knowledge from the extracted experience. Patients can search for peers with similar experience and connect with them. In this paper, we model the extraction of patient experience as a classification problem: classifying each sentence in a forum post as containing patient experience or not containing patient experience. We propose to exploit the sentence context information for such experience extraction task, and classify the context information into global context and local context. A unified Context-Aware expeRience Extraction (CARE) framework is proposed to incorporate these two types of context information. Our experimental results show that the global context can significantly improve the experience extraction accuracy, while the local context can also improve the performance when less labeled data is available.**

I. Introduction

Online health forums, such as MedHelp [1], PatientsLikeMe [2], WebMD [3], and Healthboards message boards [4], provide a large repository for patients, caregivers, and researchers to seek valuable information. According to a recent survey, 72% of Internet users looked for online health information [1]. 18% of Internet users or 23% of those who have chronic conditions would like to connect with people with similar concerns [2].

Recently, the vast majority of research has been focused on online health forums to show its importance in knowledge discovery, e.g., discovering adverse drug reaction (ADR) knowledge [3], [4]. ADR has become a leading cause of death in the United States [5]. The online health forums provide a large-scale and timely fashion to obtain the ADR knowledge based on the aggregated evidences reported by patients or their caregivers. However, the health forum information can be very noisy. Some information includes patient-reported personal health experience, such as "Took Keppra for a week. I felt very nauseous", which provides new ADR evidence. The other information can be hearsay, such as "You should not take Keppra as I heard that it could cause birth defects." Such unverified hearsay, if used for ADR discovery, can be very misleading. To discover ADR knowledge based on the aggregated evidences, it is critical for us to extract patient-reported first-hand health experience and differentiate it from the other less-convincing forum information.

The online health forums are also commonly used for patients to seek valuable information and provide mutual support [6], [7]. Lots of patients would like to seek other patients with similar experience. With those peers with similar experience, they can form a self-help group to support and encourage each other. In this case, patient experience extraction can help them find the right people to connect.

Some existing work on experience extraction from the Web includes [8], [9]. In [9], the experience is extracted from a MedHelp health forum post by classifying each sentence as personal experience or hearsay. They used the bag of words (BOW), bigram, and Part of Speech (POS) features extracted from each individual sentence for such classification task. In [8], experience-containing sentences are detected from Weblogs. They formulated the problem as a classification task with various linguistic features extracted from each sentence. Section VI has a more detailed discussion of the related work.

The existing work classifies a sentence as experience-containing sentence or not experience-containing sentence independently, and does not consider the context of a sentence. However, the context information can be very important for accurately classifying a sentence. First, an individual sentence can be only part of a complete experience description, and it is difficult to identify whether it describes a patient's experience or not by the sentence itself. The extensive use of informal language, like incomplete sentences, make it even more difficult to classify each sentence separately. Consider this short sentence "*ER+/PR+ HER2-*." It is a breast cancer type, and has nothing to do with patient experience. Now let us check its local context: "*A little about me. Found knot in July ... Found 0.09mm invasive area. No lymph node involvement. ER+/PR+ HER2-. Onco score of 20 1/2. Now on tamoxifen. A survivor.*" With its context sentences, we can recognize that breast cancer type as part of the patient's personal experience. Second, the same sentences can be either part of experience or not part of experience given different global context, e.g.,

[1] http://www.medhelp.org
[2] http://www.patientslikeme.com
[3] http://exchanges.webmd.com
[4] http://www.healthboards.com/boards

978-1-4673-9548-9/15 $31.00 © 2015 IEEE
DOI 10.1109/ICHI.2015.11

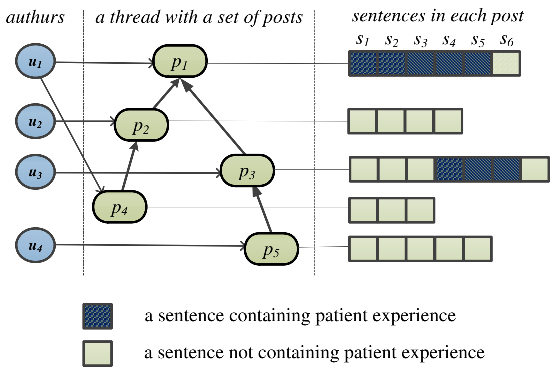

Fig. 1: The Context of Each Sentence.

different post authors. Consider a post with two sentences, "Keppra could cause birth defects. It is really not a safe drug." Such sentences are more likely to be part of patient experience if they are from a post written by a patient who reported the drug usage experience in an earlier post.

In this paper, we classify a sentence in a health forum post as containing patient experience or not containing patient experience with its context information. We differentiate two types of context information: global context and local context. We then propose a Context-Aware expeRience Extraction (CARE) framework to incorporate these two types of context information, and evaluate its effectiveness with experiments.

II. Problem Statement

We extract the patient-reported first-hand personal health experience, or patient experience in short, from online health forums. Experience is a collection of events and/or activities from which an individual or group may gather knowledge, opinions and/or skills [5]. In [8], experience is defined as knowledge embedded in a collection of activities or events which an individual or group has actually undergone. It can be subjective as in opinions as well as objective. In this work, we define patient experience as the health-related events and/or activities undergone by a specific patient and the corresponding opinions on those events and/or activities perceived by that patient. For example, "Took Keppra for a week. I felt very nauseous" is patient experience. "Keppra could cause birth defects. It is really not a safe drug" can be part of patient experience if they are based on a past event that the patient has actually taken the drug. On health forums, a patient is sometimes represented by her/his caregiver, e.g., the patient's mother. As they are closely related, the caregiver's description or corresponding opinions are also counted as patient's first-hand experience.

Fig. 1 shows an abstracted forum thread. It has 5 posts from 4 different authors. p_1 is the first post in the thread, which is replied by p_2 and p_3, denoted by the solid arrows. Some posts, such as p_1 and p_3, include sentences containing patient experience, which are represented by the shadowed grids. The other sentences do not contain any patient experience, which are represented by the unshadowed grids.

[5]http://en.wikipedia.org/wiki/Experience_(disambiguation)

Our goal is to classify each sentence in a forum post as a sentence containing patient experience or not containing patient experience. Suppose we have a set of threads collected from an online health forum. Each thread consists of a sequence of posts. Each post consists of a sequence of sentences, and is written by a forum user. We extract a set of sentences $\mathcal{S} = \{s_1, s_2, ..., s_l\}$ from the threads. Each sentence s_i has its author context u_i, e.g., the author's profile, post context p_i, e.g., the neighboring sentences in the same post, and thread context t_i, e.g., the post position in the thread.

We model the problem as a classification task. Let $Y = \{y_1, y_2, ..., y_l\}$ be the corresponding labels, where $y_i \in \{-1, 1\}$, $y_i = 1$ indicates that s_i is a sentence containing patient experience, and $y_i = -1$ indicates that s_i is a sentence not containing patient experience. Each sentence $s_i \in \mathcal{S}$ has its context information $\langle u_i, p_i, t_i \rangle$. We aim to learn a classifier from labeled sentences with their content, context and labels, which can predict the labels for the unlabeled sentences.

III. Context Information for Patient Experience Extraction

In this section, we introduce two types of context information: global context and local context. The global context is the context information shared by the whole post, while the local context is the context information within a post.

A. Global Context

The global context of a sentence is the context information of the post containing that sentence, which is used to differentiate the posts that are more likely to contain patient experience from the other posts. It can be extracted from the post author context or the thread context of a post.

As different post authors may have different preferences about sharing personal experience, we can take the author context or profile into consideration when classifying a sentence posted by that author. For example, a patient tends to share their personal experience, while a doctor rarely does that. We can thus use the author type, whether the author is a doctor or not as an author context feature. More author context features, such as the previous posts from the same author, the number of replies posted by the author, and the number of badges the author has won, can also be considered.

The thread context of a post can also help extract the patient experience sentences. For example, the post position, in particular, whether it is the first post in a thread, can be used as a thread context feature. The first post tends to share more personal experience, while the following posts tend to contain some suggestions and comments from the other forum users. We can further consider whether a post is from the thread initiator, i.e., the author of the first post in a thread. If a post is from the thread initiator, it is more likely to share some personal experience.

B. Local Context

The local context refers to the context information within a post. An example post from a health forum is shown in Fig. 2. Each sentence is preceded by its sentence sequence number.

> ``(1) ***I*** *take Sotalol.* (2) *It is ok in preventing ventricular arrhymias.* (3) *especially on high dosage.* (4) *After 18 months of assumption* ***my ICD recorded*** *only a few Vtachs, while when* ***I was*** *on Coreg* ***I had*** *at least one Vtach a week.* (5) *Your bpm and bp seem just ok.* (6) *You want them lower?* (7) *You can...* "

Fig. 2: An Example Post with a Sequence of Sentences.

In Fig. 2, the first 4 sentences are patient experience sentences, while sentence (5)-(7) are not. We notice that sentence (1) and (4) can be identified as patient experience sentences much easier than sentence (2) and (3). The reason is that both sentence (1) and (4) contain some distinguishable features. For example, as shown in the bold font, sentence (1) uses the pronoun "I" as the subject and sentence (4) further uses the verbs in their past tense. Such sentence-level features are useful for the patient experience identification. On the other hand, it is difficult to identify (2) and (3) when we only consider themselves; however, we can identify them as patient experience with their context sentences (1) and (4).

Aforementioned analysis suggests that the label of a sentence tends to be consistent with the labels of its local context sentences. Such local context may be leveraged to improve the patient experience sentence extraction performance. More specifically, we make use of local context based on the following two observations[6].

- *Observation 1: The adjacent sentences tend to have the same label.*
- *Observation 2: The sentences in the same post tend to have the same label.*

The first observation follows the fact that the experience or non-experience sentences are usually contiguous. A sentence used to describe the patient's experience is more likely to be followed by another sentence that continues the experience description. Even if a post includes both experience and non-experience sentences, the experience sentences or non-experience sentences are usually clustered together. For example, a forum user can make some comments or suggestions in the first part of a post as a reply to a preceding post, and then describe some of her own experience in the second part. Or, she can describe her own experience first, and then give some suggestions in the following set of sentences, as shown in Fig. 2.

The second observation follows the fact that the label of an individual sentence usually depends on the labels of the other sentences in the same post. If most of the other sentences in that post are patient experience sentences, then it is more likely that the individual sentence is also a patient experience sentence. Otherwise, if none of the other sentences are patient experience sentences, then it is less likely one individual sentence is a patient experience sentence.

We have verified the above two observations based on statistical analysis. We crawled the data from the MedHelp Heart Disease forum, and randomly selected a subset of threads as our data set. Two annotators labeled the same set of sentences in 264 posts from 85 threads and resolved all the label disagreements. Among the 264 posts, 105 posts, 41.6% of all the posts, do not contain any patient experience sentences, which indicates that the experience sentences are not evenly distributed in all the posts. The total number of sentences we get from the 264 posts is 2007 by using Stanford NLP for sentence splitting [10]. By removing the sentences with less than 3 tokens, e.g., a sentence containing only punctuation marks, and the question sentences, which are considered as non-experience, we get 1481 sentences with 759 sentences labeled as experience sentences.

[6]We may make other observations about local context; for example, sentences within a distance or containing co-referent entities tend to have the same label, and we would leave it as our future work.

TABLE I. Statistical Analysis of Label Pairs

	#pairs	#pairs with the same label	ratio
Any two sentences	2191880	1095884	0.5
Two adjacent sentences	2458	2126	0.865
Two sentences in the same post	13220	10774	0.815

As shown in Table I, two random sentences have the same label with the probability 0.5. Two adjacent sentences (from the same post) have the same label with the probability 0.865. Two sentences from the same post have the same label with probability 0.815. The statistics show that the adjacent sentences or the sentences from the same post indeed tend to share the same label. These observations are used to model local context in the problem of patient experience extraction.

IV. The Proposed Framework: CARE

In this section, we introduce how to incorporate context information for patient experience extraction. We first introduce the basic classifier of our framework. Then we present how to model the global context and local context, which results in the proposed CARE framework.

A. The Basic Classifier

We use a support vector machine (SVM) as the basic model for the proposed framework CARE to classify each sentence as containing patient experience or not containing patient experience due to its superior performance in many real-world applications [11]. Let $f = \{f_1, f_2, ..., f_{k_1}\}$ be the set of k_1 sentence-level features extracted from the content of each sentence. Let $Z = \{z_1, z_2, ..., z_l\} \in \mathbb{R}^{k_1 \times l}$ be the feature matrix representation for l labeled sentences $\{s_1, s_2, ..., s_l\}$. Each feature vector $z_i, 1 \leq i \leq l$, contains all the k_1 feature values extracted from sentence s_i. We also have $Y = \{y_1, y_2, ..., y_l\}$ as the corresponding labels, where $y_i = 1$ indicates sentence s_i is a patient experience sentence, and $y_i = -1$ if it is not. Given Z and Y, we learn a standard soft-margin SVM by solving the following optimization problem.

$$\begin{aligned} \min_{w,b,\xi} \quad & \frac{1}{2} w^T w + C_e \sum_{i=1}^{l} \xi_i \qquad (1) \\ s.t. \quad & y_i(w^T z_i + b) \geq 1 - \xi_i, \\ & \xi_i \geq 0, i = 1, \ldots, l, \end{aligned}$$

where w is the feature weight vector, b is the bias, and $w^T z_i + b$ is the predicted value for the label. ξ_i is the soft-margin slack variable for z_i, which allows the noise in the training sample. C_e controls the total error from training samples.

B. Exploiting Global Context

The global context, such as author context and thread context, can be considered as features of each sentence, which suggests that we can exploit global context via feature engineering. Let $g = \{g_1, g_2, ..., g_{k_2}\}$ be the set of k_2 features extracted from global context. We expand the sentence-level feature set f with the context feature set g, and let $F = f \bigcup g$ be the feature set of $k = k_1 + k_2$ features for each individual sentence. Let $X = \{x_1, x_2, ..., x_l\} \in \mathbb{R}^{k\times l}$ be the new feature matrix representation for l labeled sentences $\{s_1, s_2, ..., s_l\}$ with the feature set F. Each feature vector, $x_i, 1 \leq i \leq l$, contains k_1 sentence-level features and k_2 features from its global context.

C. Exploiting Local Context

Different from the global context, it is difficult to model local context via feature engineering, hence we model the local context as labeling constraints between sentences within a post. Our previous analysis suggests that two adjacent sentences or two sentences from the same post tend to share the same label. We define a matrix $A^N \in \mathbb{R}^{l\times l}$ to encode the adjacent-sentence relationship where $A^N_{ij} = 1$ if sentences s_i and s_j are adjacent sentences and $A^N_{ij} = 0$ otherwise. Similarly, we can introduce a matrix $A^P \in \mathbb{R}^{l\times l}$ to encode the same-post relationship where $A^P_{ij} = 1$ if sentences s_i and s_j are from the same post and $A^P_{ij} = 0$ otherwise. The following derivations are based on the sentence-sentence label consistency matrix $A \in \mathbb{R}^{l\times l}$. We can obtain A from either A^N, A^P or their combination as $A = A^N + \lambda A^P$, where λ controls the weight of different local context in the model. In this paper, we focus on studying whether the context can improve the performance. We simply combine two different local contexts with equal weight $\lambda = 1$ to construct a sentence-sentence label consistency matrix. More study will be conducted to investigate the best way to combine them.

To model local context for patient experience extraction, the basic idea is to minimize the label difference between two sentences if they are two adjacent sentences or from the same post. This idea can be mathematically formulated as solving the following minimization problem:

$$\min_{w} \quad \frac{1}{2}\sum_{i,j} A_{ij}(w^T x_i - w^T x_j)^2 = w^T X L X^T w, \qquad (2)$$

where $X \in \mathbb{R}^{k\times l}$ is the matrix representation of sentences in the training set. $L = D - A$ is a Laplacian matrix where D is a diagonal matrix with $D_{ii} = \sum_{j=1}^{l} A_{ij}$.

D. The Optimization Problem for CARE

With model components for global and local context, the proposed patient experience extraction framework CARE is to solve the following optimization problem:

$$\begin{aligned} \min_{w,b,\xi} \quad & \frac{1}{2}w^T w + C_e \sum_{i=1}^{l} \xi_i + C_r w^T X L X^T w \qquad (3)\\ s.t. \quad & y_i(w^T x_i + b) \geq 1 - \xi_i,\\ & \xi_i \geq 0, i = 1, \ldots, l, \end{aligned}$$

where C_r is the parameter to control the contribution from local context. Note that the global context has been incorporated into the expanded feature vector x_i.

We can easily collect a large amount of unlabeled sentences from online health forums. The proposed framework CARE can be extended to incorporate unlabeled sentences. Let $\mathcal{S}^U$ be a set of m unlabeled sentences. We use $X^U \in \mathbb{R}^{k\times m}$ to denote the matrix representation of $\mathcal{S}^U$ and $X' = [X, X^U] \in \mathbb{R}^{k\times(l+m)}$ to denote the matrix representation of labeled and unlabeled sentences. Similarly, we extract features to model global context and construct the sentence-sentence label consistency matrix $A' \in \mathbb{R}^{(l+m)\times(l+m)}$ to model local context. The Laplacian matrix L' is built based on A' for both labeled and unlabeled data. The formulation extending CARE to incorporate unlabeled data is then updated by replacing X and L with X' and L' in Eqs. (3). The optimization problem can be solved by the primal Laplacian SVM solver [12], [13].

V. Experiments

In this section, we conduct experiments to answer the following two questions - (1) can the proposed framework CARE help patient experience extraction? and (2) how does contextual information affect the performance of CARE? We begin by introducing experimental settings.

A. Experimental Settings

We use the same data set as described in Section III-B. In that data set, we have 1481 labeled sentences in 264 posts from 85 threads crawled from the MedHelp forum.

We randomly divide all the posts into two subsets of equal size $\mathcal{A}$ and $\mathcal{B}$. All the sentences from $\mathcal{A}$ are used for training, and the others from $\mathcal{B}$ are used for testing. We always fix $\mathcal{B}$ as the testing set; while choose $\alpha\%$ of $\mathcal{A}$ as labeled data and the remaining $1 - \alpha\%$ as unlabeled data. We vary α as $\{10, 25, 50, 100\}$ in this work. We draw 10 random splits and report the average classification accuracy. Here the accuracy is defined as the ratio of the number of correctly classified sentences to the total number of sentences in the testing set.

We use two global context features: **CtxA** and **CtxT**, and one regularization term for local context: **Reg**. CtxA is extracted from the author context: whether the post author is a doctor or not. There are many doctors who help answer questions and give suggestions to patients in the MedHelp forums. Their profile webpages usually indicate that they are doctors. We crawled all the profile webpages for those forum users who are a post author in our data set, and extract such

TABLE II. Performance Comparison

	Labeled Data	10%	25%	50%	100%
Sentence-level features	BOW	0.721	0.758	0.77	0.785
	Bigram	0.531	0.642	0.703	0.740
	BOW + POS	0.728	0.77	0.779	0.795
	N-gram(Weka)	0.719	0.755	0.758	0.766
	SentLing	0.777	0.797	0.809	0.819
CARE		0.823	0.842	0.848	0.854

context feature. CtxT is extracted from the thread context of a post: whether the post is the first post in the thread. We use the regularization defined in Section IV-C.

B. Comparison Methods

We compare our work with representative existing work, which also uses SVM as the basic classifier. Each of the existing work uses a different set of sentence-level features. The parameters of SVM for all comparison methods are determined by cross validation.

1) Bag of Words, Bigram, POS tags: N-gram and POS tag features are commonly used for text classification. [9] has used bag of words (**BOW**, also known as unigram), **Bigram**, and **BOW + POS** features to distinguish personal experience from hearsay. We use their approaches as baselines. We also use Weka [14] to generate an n-gram combination of 1000 most important unigram, bigram, or trigram features based on the TF/IDF transformation for each sentence. We denote such feature set as **N-gram(Weka)**.

2) Linguistic features: Most of the existing work for the experience extraction and classification task uses a set of linguistic features extracted from each sentence [8], [15]. We follow them and use three sets of linguistic features, totally 21 features, as our sentence-level features. We denote them as **SentLing**.

- *Pronouns*: As shown in [15], the presence of pronouns is very useful in identifying personal experience revealing sentences. We follow the approaches in [15] to design 8 features by extracting pronouns in 8 pronoun categories.
- *Tense*: As shown in [8], [15], the verb tense in a sentence is very helpful in identifying personal experience. Most of the experience descriptions are using the past tense or present tense. We follow [8] to design 6 tense features.
- *Modality*: As shown in the same existing work as above, the modality expresses the possibility of some activities or events, and thus it can be used to differentiate experience from non-experience. We follow [8] to design 7 modality features by checking the modality verb in a sentence.

C. Performance Comparison for Patient Experience Extraction

The comparison results are shown in Table II. Note that we use features in SentLing as the feature set for the content of sentences for the proposed framework **CARE**. We made the following observations:

- BOW and BOW + POS always outperform Bigram. The feature sets of Bigram are much larger than BOW and BOW + POS, which will degrade the performance of SVM because of curse of dimensionality.
- SentLing obtains better performance than the other compared methods, which suggests the importance of linguistic features in patient experience extraction.
- The proposed framework CARE can significantly improve the performance of patient experience extraction. For example, compared with the best performance of comparison methods, CARE improves the performance by 6.0% with 10% labeled data and 4.3% with 100% labeled data. The major reason is that the proposed framework captures both global and local context. In the following subsection, we will investigate the impact of context information on CARE.

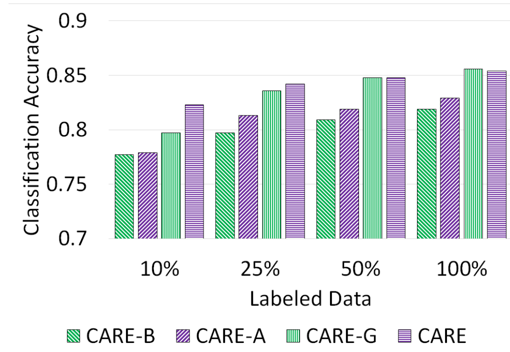

Fig. 3: The Impact of Context Information.

We also noticed that such a classification task is very difficult, mainly due to the usage of informal language and the complexity of natural language patterns. More advanced features, such as deep syntactic and semantic analysis of a sentence, combined with a larger training set should be helpful for an improved performance in the future.

D. Impact of Context Information

In this subsection, we investigate the effects of context information on the performance of CARE by systematically defining its variants.

- **CARE-B**: it only uses the basic classifier SVM with the SentLing features without any context information;
- **CARE-A**: it combines SentLing features and only CtxA features from global context;
- **CARE-G**: it combines SentLing features, and both CtxA and CtxT features from global context;

Parameters in these variants are determined via cross-validation and the comparison results are shown in Figure 3. First, **CARE-A** outperforms **CARE-B**, which supports the importance of CtxA features from global context especially when more labeled data is available. Second, **CARE-G** improves the performance by combining CtxA and CtxT features from global context, which suggests that CtxA and CtxT contain complementary information. Third, **CARE** can further improve the performance by incorporating local context into **CARE-G** when less labeled data is available. For example, the accuracy increases from 0.797 to 0.823 with 10% labeled training data. These observations suggest that both global and local context are useful for the performance improvement for patient experience extraction.

VI. Related Work

Some related work about extracting experience from the Web includes [8], [16], [15], [9], [17]. [8] has been focused on domain-independent and objective experience mining from Weblogs. They proposed a set of sentence-level linguistic features to classify a sentence. [9] classifies a sentence from the MedHelp online health forums into either personal experience or hearsay. [15] is focused on the detection of experience revealing sentences in product reviews. Similar to [8], each sentence is considered independently. None of them use the context information to improve the sentence classification performance. In an explorative study, [16] investigates some features for automatic detection of reports of experiences with products from online forums. Although no further experiments are provided, these explored features can also be incorporated into our framework. [17] classifies a tweet instead of a sentence in our work. The tweet context, such as the twitter's profile, has been taken into consideration in their classification model, which is similar to our global context features.

Some work about experience knowledge mining or sentence classification is also related to our work. [18] is mainly focused on the factuality analysis of the event mentions described in a sentence, e.g., whether the event indeed took place. Although their proposed techniques are orthogonal to ours, they can be used to extract more sentence-level features to improve our patient experience extraction accuracy. [19], [20] focus on more fine-grained experience knowledge extraction, including the extraction of events, entities, and relationships. [21] classifies a sentence from online medical forums into three categories: physical examination/symptoms, medications, and others. They use the post position feature and also use the conditional random fields (CRF) model to capture the dependency between sentences. However, their experimental results show that CRF-based model performs even worse than a SVM-based model without considering the dependency between sentences.

VII. Conclusions and Future Work

In this paper, we propose context-aware experience extraction from online health forums. We classify the context information into two types: global context and local context, and incorporate them into a unified framework CARE. The experimental results confirm its effectiveness.

There are several interesting directions for further investigations. First, we will investigate more advanced linguistic features, context features, and constraints. For example, we can extract features from the sentence dependency parsing graph and the post-level statistics like sentiment strength. Second, sentences in online health forums are usually unlabeled and it is time and effort consuming to obtain labeled data; hence we would like to investigate how to extract patient experience in unsupervised scenarios. Finally, we will evaluate the proposed CARE framework with different data sets from various health forums.

Acknowledgments

This material is based on work partially supported by NSF CAREER Award IIS-1322406, Google, and Leir Charitable Foundations.

References

[1] S. Fox and M. Duggan, "Health online 2013," in *Pew Research Report*, 2013.

[2] S. Fox, "Peer-to-peer healthcare," in *Pew Research Centers Internet and American Life Project*, 2011.

[3] C. C. Yang, H. Yang, and L. Jiang, "Postmarketing drug safety surveillance using publicly available health consumer contributed content in social media," *ACM Transactions on Management Information Systems*, vol. 5, no. 1, April 2014.

[4] X. Liu, J. Liu, and H. Chen, "Identifying adverse drug events from health social media: A case study on heart disease discussion forums," in *Proceedings of International Conference for Smart Health*, 2014.

[5] B. W. Chee, R. Berlin, and B. Schatz, "Predicting adverse drug events from personal health messages," *AMIA Proceedings*, 2011.

[6] Y.-C. Wang, R. Kraut, and J. M. Levine, "To stay or leave? the relationship of emotional and informational support to commitment in online heath support groups," in *ACM 2012 Conference on Computer Supported Cooperative Work (CSCW)*, 2012.

[7] V. V. Vydiswaran, Y. Liu, K. Zheng, and Q. Mei, "User-created groups in health forums: What makes them special?" in *Eighth International AAAI Conference on Weblogs and Social Media (ICWSM'14)*, 2014.

[8] K. Park, Y. Jeong, and S. Myaeng, "Detecting experiences from weblogs," in *Proceedings of ACL*, 2010.

[9] X. Liu and H. Chen, "Azdrugminer: An information extraction system for mining patient-reported adverse drug events in online patient forums," in *Proceedings of International Conference for Smart Health*, 2013.

[10] C. D. Manning, M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky, "The Stanford CoreNLP natural language processing toolkit," in *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2014, pp. 55–60. [Online]. Available: http://www.aclweb.org/anthology/P/P14/P14-5010

[11] N. Cristianini and J. Shawe-Taylor, *An introduction to support vector machines and other kernel-based learning methods*. Cambridge university press, 2000.

[12] M. Belkin, P. Niyogi, and V. Sindhwani., "On manifold regularization," in *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics (AISTAT 2005)*, 2005, pp. 17–24.

[13] S. Melacci and M. Belkin, "Laplacian Support Vector Machines Trained in the Primal," *Journal of Machine Learning Research*, vol. 12, pp. 1149–1184, March 2011.

[14] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The weka data mining software: An update," *SIGKDD Explorations*, vol. 11, 2009.

[15] Q. Nguyen, "Detecting experience revealing sentences in product reviews," Master's thesis, the University of Amsterdam, 2012.

[16] V. Jijkoun, M. de Rijke, W. Weerkamp, P. Ackermans, and G. Geleijnse, "Mining user experiences from online forums: An exploration," in *Proceedings of NAACL HLT Workshop on Computational Linguistics in a World of Social Media*, 2010.

[17] W. Murphy, "Using supervised learning to identify descriptions of personal experiences related to chronic disease on social media," Master's thesis, The Pennsylvania State University, 2014.

[18] S. Abe, K. Inui, K. Hara, H. Morita, C. Sao, M. Eguchi, A. Sumita, K. Murakami, and S. Matsuyoshi, "Mining personal experiences and opinions from web documents," *Web Intelligence and Agent Systems: An International Journal*, vol. 9, 2011.

[19] T. Kurashima, K. Fujimura, and H. Okuda, "Discovering association rules on experiences from large-scale blog entries," in *Proceedings of ECIR*, 2009.

[20] S. Myaeng, Y. Jung, and Y. Jeong, "Experiential knowledge mining," *Foundations and Trends in Web Science*, vol. 4, no. 1, 2012.

[21] P. Sondhi, M. Gupta, C. Zhai, and J. Hockenmaier, "Shallow information extraction from medical forum data," in *Proceedings of International Conference on Computational Linguistics (COLING)*, 2010.