

Interactive Learning for Identifying Relevant Tweets to Support Real-time Situational Awareness

Luke S. Snyder, Yi-Shan Lin, Morteza Karimzadeh, Dan Goldwasser, and David S. Ebert, *Fellow, IEEE*

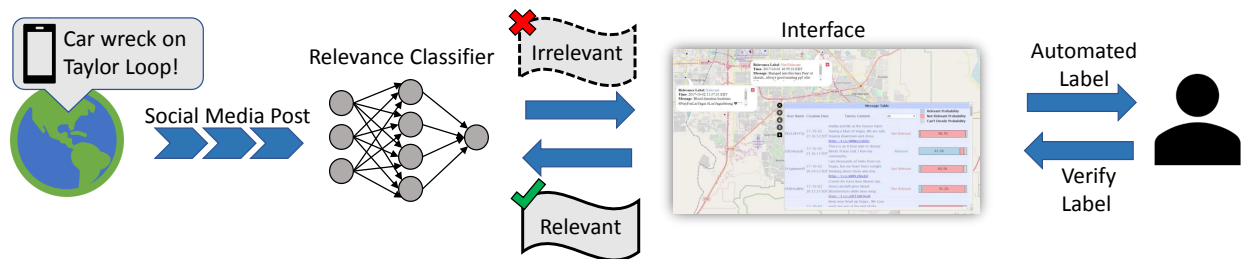


Fig. 1. Our interactive learning framework allows users to train text relevance classifiers in real-time to improve situational awareness. In this example, a real-time tweet regarding a car accident is incorrectly classified as “Irrelevant”. Through the SMART 2.0 interface, the user can view its label and correct it to “Relevant”, thereby retraining and improving the classifier for incoming streaming data.

Abstract—Various domain users are increasingly leveraging real-time social media data to gain rapid situational awareness. However, due to the high noise in the deluge of data, effectively determining semantically relevant information can be difficult, further complicated by the changing definition of relevancy by each end user for different events. The majority of existing methods for short text relevance classification fail to incorporate users’ knowledge into the classification process. Existing methods that incorporate interactive user feedback focus on historical datasets. Therefore, classifiers cannot be interactively retrained for specific events or user-dependent needs in real-time. This limits real-time situational awareness, as streaming data that is incorrectly classified cannot be corrected immediately, permitting the possibility for important incoming data to be incorrectly classified as well. We present a novel interactive learning framework to improve the classification process in which the user iteratively corrects the relevancy of tweets in real-time to train the classification model on-the-fly for immediate predictive improvements. We computationally evaluate our classification model adapted to learn at interactive rates. Our results show that our approach outperforms state-of-the-art machine learning models. In addition, we integrate our framework with the extended Social Media Analytics and Reporting Toolkit (SMART) 2.0 system, allowing the use of our interactive learning framework within a visual analytics system tailored for real-time situational awareness. To demonstrate our framework’s effectiveness, we provide domain expert feedback from first responders who used the extended SMART 2.0 system.

Index Terms—Interactive machine learning, human-computer interaction, social media analytics, emergency/disaster management, situational awareness

1 INTRODUCTION

Social media data has been used extensively in a variety of applications and research endeavors due to its ability to provide useful information on the public’s opinions and behavior. Analysts in various domains are increasingly using social media to gain rapid situational awareness. For instance, first responders are leveraging Twitter data to obtain actionable information for crisis response and prevention (see [29] for an extensive list of literature on this subject). However, the vast amounts of unstructured text make the identification of relevant information nontrivial, limiting situational awareness. This issue is further compounded by changes in topics of interest (to end users) over time, since the computational models built to determine relevant information for one event or one user group may not apply to other events or other user groups due to variations in diction, word structure, or user expectations.

Several classification approaches have been developed to identify relevant and irrelevant social media information, such as clustering [5, 6], keyword matching [45], and term-vector similarity [12]. However,

to the best of our knowledge, no existing work in this area includes interactive learning with real-time data, focusing instead on improving the machine learning algorithms themselves [5, 16, 23, 31–33, 38, 45, 46, 54] or interactively training on archived datasets [9, 18]. Continuing on our example of first responders, a pre-trained classifier may not fulfill first responders’ varying needs, since one first responder may be interested in monitoring road closures, and another one might be interested in identifying disinformation and misinformation on social media in order to take counter-action. Ultimately, first responders’ definition of relevancy will depend on the situation at hand, which can vary over time. Interactively training classifiers through iterative user labeling can alleviate this problem.

In this paper, we present a novel interactive framework in which the user iteratively (re)labels the relevancy of streaming social media data to adaptively train the underlying model to match their needs for improved situational awareness. We compare three different types of neural networks in terms of classification performance and computational efficiency for real-time learning. Furthermore, we optimize and computationally evaluate the selected models by simulating the real-time user feedback on several crisis-related datasets. Our results show that our interactive model outperforms state-of-the-art machine learning-based classification models.

To incorporate our evaluated models into a working application, we extend an existing visual analytics system tailored for situational awareness called the Social Media Analytics and Reporting Toolkit (SMART) [52, 53], which has been successfully used by many first responder groups in the United States. SMART allows users to interactively explore trending topics on social media through integrated topic modeling and spatial, temporal, and textual visualizations. We call the

• Luke S. Snyder, Yi-Shan Lin, Dan Goldwasser, and David S. Ebert are with Purdue University.

Email: {snyder238, lin670, dgoldwas, eberrd}@purdue.edu

• Morteza Karimzadeh is with the University of Colorado Boulder (formerly at Purdue University). Email: karimzadeh@colorado.edu

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

newly extended system SMART 2.0, which incorporates our interactive learning framework to address the needs raised by the aforementioned first responder users and reduce noise in the incoming stream of data.

Finally, we present domain-expert feedback on the usefulness of our approach as experienced by multiple first responders who used SMART 2.0 for crisis-related use cases. In addition, we include two usage scenarios of the system to illustrate its application to real-life situations. Overall, the major contributions of this paper are as follows:

1. We present a novel interactive learning framework for classification of streaming text data.
2. We compare three different types of neural networks in terms of performance and computational efficiency, and tune the models for learning at interactive rates. We further computationally evaluate the selected model on several disaster-related datasets.
3. We integrate our models in SMART 2.0, a visual analytics application for situational awareness, and present user feedback obtained from domain experts using the system for crisis events.

In the remainder of the paper, we discuss related work in section 2, the design of the framework and model in section 3, SMART 2.0 in section 4, evaluation of our framework in section 5, discussion and future work in section 6, and concluding remarks in section 7.

2 RELATED WORK

2.1 Short Text Classification

Researchers have presented many techniques to classify text documents into categories such as sentiment or topics [17, 35, 37, 49]. However, classifying short text, e.g. social media posts, is more challenging due to the lack of contextual information and loose adherence to standard grammar. To tackle the brevity of short text, auxiliary resources such as external corpora [10] or knowledge bases [20], or methods such as term frequency-inverse document frequency (TF-IDF) [16], have been proposed for improving classification.

Representing words as n -dimensional vectors (i.e. word embedding) has become increasingly prevalent, since vectors can be used as inputs to machine learning models for finding semantic similarities [47, 50]. In particular, Google’s Word2Vec [30] has been employed extensively in classification tasks [4, 26, 27, 32, 33] due to its impressive ability in capturing linguistic regularities and semantics. For instance, words frequently used together are likely to be closer in the Word2Vec vector space than words that are not, and vector operations reveal meaningful semantics (e.g., the vector “King” – “Man” + “Woman” is close to the vector “Queen” [30]). Since pre-trained Word2Vec models encode embeddings learned from larger web corpora, they have been increasingly used in short text classification tasks [4, 32, 33, 44].

Neural networks have generated state-of-the-art results in recent years for text classification problems [30, 32, 33, 43] and have also been used with Word2Vec [32, 33, 44]. Neural networks are well-suited for online learning processes in which training data is supplied iteratively since they can learn adaptively from new data [32, 33]. Nguyen et al. [33] presented a convolutional neural network with Word2Vec that outperformed non-neural classifiers, and Nguyen et al. [32] proposed a new online learning classification algorithm for deep neural networks utilizing the log-loss and gradient of sequential training batches. Their methods were evaluated with disaster-related datasets. However, these methods were not adapted to user-guided learning in which time constraints are essential and the provided batches may be small. In particular, the online learning method designed by Nguyen et al. [32] was evaluated with batch sizes of 200. In our work, we assume the user needs to train with flexibly interactive amounts of data (10-20 samples) to view immediate predictive improvements for situational awareness.

Classification for situational awareness. Utilizing real-time social media data for situational awareness (and crisis prevention in particular) is a heavily researched topic [16, 23, 31–33, 38, 45, 54]. However, identifying situationally-relevant information is nontrivial due to the high noise-to-signal ratio. Karimi et al. [23] found that classification methods, such as Support Vector Machine and multinomial Naïve

Bayes, can identify disaster-related tweets, although generic features such as hashtag count and tweet length are preferable so that the model does not learn relevancy only for a specific disaster. Researchers have used clustering [5, 6, 46] or enhanced keyword matching [45] to detect relevant crisis and event information, and provided human-annotated Twitter corpora that can be used to train word embedding models [21].

Nazer et al. [31] developed a system to detect requests for help by utilizing both tweet context (e.g., geotag) and content (e.g., URLs). Rudra et al. [38] designed a novel classification-summarization framework to classify disaster-related tweets, and then summarize the tweets by exploiting linguistic properties typical of disaster tweets (e.g., combinations of situational and non-situational information). Zoppi et al. [54] provided a relevance labeling strategy for crisis management that computed data relevance as a function of the data’s integrity (e.g., are the geo-coordinates incorrect?), statistical properties (e.g., can we select a subset of the data that are geographically close?), and clustering (e.g., what groups are present in the data?). Toriumi et al. [46] clustered tweets based on their retweet count in real-time to extract important topics and classify tweets accordingly.

The methods discussed so far, however, lack user interactivity. In particular, these classification methods are inflexible to user-dependent needs that change over time as new situations and events occur. As such, their practical use for real-time situational awareness is limited.

2.2 Visual Analytics and Interactive Learning for Situational Awareness

Researchers have presented a number of visual analytics (VA) solutions for situational awareness. Diakopoulos et al. [12] developed Vox Civitas, a VA application for journalistic analysis and user-guided filtering using social media content. Vox Civitas filters out unrelated data by automatically computing time-dependent term-vector similarities. Twit-Info [28] aggregates streamed Twitter data and automatically discovers events from activity peaks in real-time. The authors assign relevance to a tweet by counting its number of event-related keywords. Pezanowski et al. [36] designed the geovisual analytics system SensePlace3 to provide situational awareness by leveraging geographical information and place-time-theme indexing with string-based queries for exploring datasets. SensePlace3 primarily relies on TF-IDF for tweet retrieval in response to user queries. However, these tools do not employ machine learning for relevance classification and do not integrate user feedback to improve their underlying models or algorithms.

Visual analytics has also been increasingly used to improve various machine learning processes, such as feature selection [13], attribute weighting [48], and labeling [8, 9, 18], and even understanding the models themselves [22, 39, 42]. Sacha et al. [40] proposed a framework to discuss the various forms of human interaction with machine learning models in visual analytics systems and theorized that VA tools could increase knowledge and usability of machine learning components. Endert et al. [15] designed a system that classifies archived documents through user-guided semantic interactions (e.g., moving a document to another group) that improve the underlying model. Our work is based on the same idea in that we intend to improve model performance through user feedback, but with real-time social media data.

Heimerl et al. [18] analyzed three separate methods for user-guided classification of a set of archived text documents: the *basic* method, which does not employ sophisticated visuals; the *visual* method, which visually represents the labeled and unlabeled documents for user exploration; and the *user-driven* method, which provides the user with full control over the labeling process. The first two methods employ active learning, in which the model selects a data sample to be labeled by the user that most effectively helps it distinguish relevant from irrelevant data. This contrasts with the user deciding which instances they wish to label. The authors did not find any statistically significant differences in terms of F_1 score between the methods in their user study. Bosch et al. [9] developed ScatterBlogs2, a VA application that provides user-guided learning of filter classifiers on historical social media messages to support situational awareness. These two works are perhaps the most similar to ours, yet differ in two fundamental ways. First, they do not provide interactive learning in real-time, which strains the user, as

they are required to visit historical data for additional training. Second, they do not employ neural networks, which are better suited for online learning environments, such as social media streaming, in which training data is supplied sequentially over time [32, 33]. It is important to note that Bosch et al. [9] allow the user to adjust a filter’s focus (i.e., how precise the classification is) in real-time if it misses relevant data or does not sufficiently filter out irrelevant data. However, this could indicate that the model has not properly learned the distinction between relevant and irrelevant data. Since training can only be completed with historical posts, the user is unable to update the model immediately with the streamed data, limiting situational awareness. Our approach not only solves this issue by allowing the user to immediately train the model for improvement, but also provides the user with the ability to create classifiers on-the-fly to accommodate their real-time needs.

3 INTERACTIVE LEARNING FRAMEWORK

Our framework for interactively learning relevant social media posts in real-time consists of two primary components. The first is a formalized set of design goals necessary to effectively facilitate situational awareness in real-time through user interactivity. The second is a detailed underlying model that is adapted to user-guided training with real-time streaming data. In section 4, we discuss our implementation of the framework that realizes the design goals.

3.1 Design Goals

The framework’s design goals were iteratively defined through discussions with domain experts such as first responders who frequently use visual analytic social media applications for real-time situational awareness. In general, these experts found it necessary for the interactive framework to incorporate user feedback in a timely manner, as well as account for time and situation-dependent user needs. With their feedback, the following specific design goals were established:

- DG1 Filter and view relevant data:** Filtering data by relevancy removes noisy data, allowing the user to more quickly find data that may require immediate attention or contain important information. The ability to view the relevant data itself is equally important for determining the urgency and content of relevant data.
- DG2 Correct incorrect classifications:** Since classifiers may provide incorrect results, especially during the early stages of training, it is necessary for the user to be able to correct the label in real-time. This both improves the model’s performance and lowers the likelihood that incoming streamed data will be incorrectly classified and missed.
- DG3 Create new classifiers in real-time:** The needs of the user can change dramatically over time and vary across users themselves. As an example, one user may wish to train a classifier to find data related to a specific hurricane event to expedite identification of people in desperate need of assistance. However, another user may wish to find data related to safety in general, not just a hurricane. As such, they should each be able to create and train their own classifiers in real-time specific to their needs at the time.
- DG4 Minimize model training time:** Although it is important to design a high-performing model, time constraints are equally important. Specifically, when the model is trained by user feedback, the user should not have to wait for several minutes for the model to be retrained and relabel data. Previously streamed data labels may update with retraining, allowing the user to potentially find important information that they had not seen before. As such, it is necessary to provide these updated results as quickly as possible for real-time situational awareness.

3.2 Workflow

Fig. 2 shows the three primary components of our framework’s workflow applied to streaming tweets (however, the framework can be generalized to other kinds of text). First, as tweets are streamed in real-time,

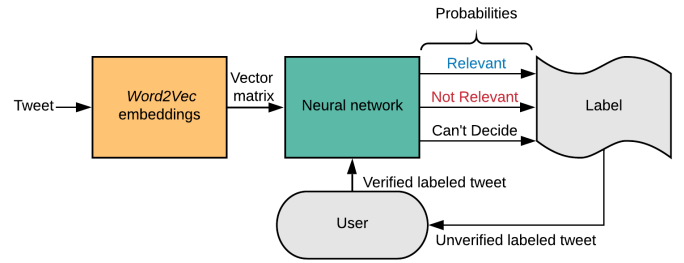


Fig. 2. High-level workflow of our framework with three main components: tweet vectorization, tweet classification, and user feedback.

they are vectorized using a word embedding model. Second, the vectorized tweets are provided as inputs to the neural network classifier (discussed in next section), which outputs a set of probabilities from the activation function of the tweet’s predicted relevancy and assigns an unverified relevance label. Third, the labeled tweet is relayed to the user through the user interface. If the user identifies tweets with incorrect labels, they can correct the label for the system to retrain and improve the model for relevance predictions.

3.3 Interactive Model Details

In the following subsections, we elaborate on the underlying representations and models used to support our interactive learning framework. We design, optimize, and evaluate our approach with the key assumption that classifiers are trained (from scratch) in real-time using user-provided labels for streaming text. We simulate this process by adding training examples in small batches of 10 and evaluating against testing data, as explained below. All simulations were completed on a server with 128 GB RAM, 32 TB of disk storage, and 2 Intel(R) Xeon(R) E5-2640 v4 CPUs at 2.40GHz.

3.3.1 Model Candidates

Selecting the underlying model for our framework was a key task, as it must be efficiently trainable with a continual stream of user-labeled data (DG4). As discussed in Section 2, neural networks are a natural choice for online learning scenarios in which training data is supplied sequentially over time [32, 33]. In addition, neural networks have generated impressive results with Word2Vec [30] embeddings [32, 33, 44]. Therefore, we employ a neural network as our classifier to determine text relevance based on real-time training examples provided by the user. To convert the text into vector inputs (of our neural network), we use word embeddings generated by Google’s Word2Vec skip-gram model [3, 30], which contains 3 million 300-dimensional word vectors pre-trained (and therefore, capturing word embeddings) on a subset of the Google News dataset with approximately 1 billion words.

In selecting the specific neural network model type, we experimented with the well-known Convolutional Neural Network (CNN) [25], Long-Short Term Memory (LSTM) Neural Network [19], and Recurrent Neural Network (RNN) [14] since they have performed well in various text classification tasks [51]. Hybrid architectures, such as recurrent convolutional neural networks [24], have also been proposed in recent years, but have not been made available in well-supported libraries. Therefore, we did not consider them in this paper, since our goal was to also support a well-tested SMART 2.0 system for end users.

Our CNN model contains the traditional convolutional and max-pooling layers before activation [51]. Specifically, we apply a 1-dimensional convolutional layer, 1-dimensional max-pooling layer, flatten the output, and then activate it with softmax and a dense layer. The filter and kernel sizes of the convolutional layer are optimized during the hyperparameter stage (explained in Section 3.3.4). We use Hochreiter’s LSTM [19] and the traditional RNN [14] architectures as provided by Keras [2]. The LSTM and RNN hidden layer each contain 300 hidden neurons and use softmax activation.

3.3.2 Design

As mentioned before, to enable the use of neural networks for classifying text, we convert the unstructured text (of the tweets) into vectors ready for consumption by the neural network. When using Word2Vec vectors as features for classification, a common approach is to convert each word in the sentence to its vector, average the word vectors in the sentence, and then use the resulting feature vector for model training [4, 43]. However, averaging the vectors results in the loss of syntactic information, which can negatively impact classification results [27]. As an example, the two sentences “Only Mary will attend the ceremony.” and “Mary will only attend the ceremony.” would generate identical averaged sentence vectors since they contain the same set of words, but they differ in meaning. Therefore, to capture both semantic and syntactic information, we represent a sentence as a matrix where each row i is a 300-dimensional Word2Vec vector corresponding to word i in the original sentence.

The input to the neural network consists of the matrix representing the sentence (as described above) and the output consists of the classification labels for the input sentence (Fig. 2). Specifically, we allow a tweet to be (1) Relevant, (2) Not Relevant, or (3) Can’t Decide. The label with the highest probability from the activation function corresponds to the final label given to it. The “Can’t Decide” label indicates that the tweet may or may not be relevant depending on the context. This is useful if the user finds a social media post such as “Remembering when Hurricane Irma destroyed my home...” that may not directly relate to the current event, but may be semantically relevant, and the user does not want to mark such cases as “Not Relevant”. This gives the user more flexibility to accommodate their needs since the definition of relevancy will depend on both the user and the situation.

3.3.3 Corpus for Model Selection and Optimization

To experiment with different neural network model types and optimize the selected model, we used a disaster-related corpus annotated on the crowd-sourcing platform, Figure Eight [1]. The dataset contains 10,876 tweets related to different types of disaster events, such as hurricanes and automobile accidents. The data was collected using keywords such as “ablate” or “quarantine”, and therefore, covers a wide variety of disaster-related topics. Our main motivation for using this open dataset is its size (as well as topical relevance), enabling the optimization of hyperparameters and comparison of various models. In the corpus, each tweet is manually labeled by Figure Eight’s workers as “Relevant”, “Not Relevant”, or “Can’t Decide”, and the distribution of labels is unbalanced. Specifically, there are 4,673 “Relevant” instances, 6,187 “Not Relevant” instances, and 16 “Can’t Decide” instances. This dataset has been used in other tweet classification research projects [45]. However, the researchers of that study remove the tweets with the “Can’t Decide” label to improve training data quality. As explained in the previous section, we find the “Can’t Decide” option useful for users to apply to cases with insufficient context for relevance determination. We randomly shuffle the data and divide the dataset into 80% training, 10% validation, and 10% testing sets.

It is important to note that we only use the Figure Eight dataset to optimize the hyperparameters and provide an initial evaluation of the model by simulating the provision of labels in real-time by the user. Since each tweet in the dataset contains true labels that were manually assigned by humans, it allows us to evaluate the model performance by comparing the model’s predictions to the true labels after each training iteration. Our proposed approach as well as its integration within the SMART 2.0 system, however, allows for the creation of the models from scratch (with no prior training) (DG3), leveraging real-time labels provided by users on streaming data for training.

3.3.4 Optimization

In order to experiment with the different neural network model types, we ran several training simulations with random combinations of hyperparameters (i.e., random grid search) to see which model converged to the best F_1 score. The F_1 score is a metric widely used to evaluate the quality and performance of machine learning models and neural networks [41]. It is computed as the harmonic mean of *precision*

(the proportion of true positive predictions compared to the total number of positive predictions) and *recall* (the proportion of true positive predictions compared to the overall number of positive instances) : $F_1 = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$. The F_1 score provides a balanced measure, combining these two performance aspects. It is therefore more informative compared to other metrics such as accuracy, especially when the training and testing sets are imbalanced [11], as in our case.

A central part of our approach to the training, validation, and verification of learning models is simulating the interactivity of visual analytics for real-time data, i.e. for use cases in which training data does not exist prior to user interaction. We assume the user (re)labels the incoming stream of data and therefore iteratively trains a model, which consequently meets their real-time needs. To replicate this process, we computationally evaluate the model’s performance (as if it is successively trained by user-labeled data) by iteratively training the model with 10 new samples from the training dataset. We average the F_1 score obtained from each of these iterations and use the resulting number to measure the model’s performance. In addition, we introduce a new variable, *window size*, for our training iterations. Specifically, due to the considerably small amount of training data provided by the user, we found that an appropriately small number of epochs (one forward and one backward pass over the training data in the model) was necessary to reduce performance degradation from initial overfitting. However, we also found that increasing the number of epochs could lead to higher F_1 scores as more data was provided. Thus, we use a sliding window of 110 samples that includes the (successively provided) new training data (10 samples) as well as the most recently used training data (100 samples) to both account for small amounts of training samples and increase the number of total training epochs for a given sample.

We use the validation data to optimize the hyperparameters for each of the CNN, LSTM, and RNN models. Specifically, after each training iteration with 10 new samples, we evaluate the neural network’s F_1 score on the validation set to view its simulated performance as if it was trained by gradual user labeling. After identifying the optimal hyperparameters for each of the CNN, LSTM, and RNN models, we evaluate their performance on the testing set.

Table 1 demonstrates the results from our validation stage. Specifically, it lists the average F_1 score obtained during each training simulation along with the total CPU time required to complete the simulation (accumulated with each training and evaluation iteration). Although in many applications, F_1 score alone is sufficient to evaluate machine learning models, it is not for ours. To see why, note that the LSTM model yields an F_1 score of 0.75, the highest of any hyperparameter combination. However, the LSTM model (with the highest F_1 score) takes approximately 4,242 seconds to complete training, whereas the CNN model (with the highest F_1 score) only takes 504 seconds. Thus, the LSTM model takes roughly eight times longer to simulate than the CNN model, but does not improve its F_1 score by a significant amount (LSTM: 0.75 vs. CNN: 0.74). In the context of interactive learning, we wish to balance the training/CPU time and performance such that the model both performs well and retrains in a short amount of time for rapid improvement (DG4). Therefore, it is necessary to consider both the CPU time and average F_1 score. With these optimization standards in mind, we chose the hyperparameters that yielded the highest F_1 scores for each model since the other hyperparameter combinations generated lower F_1 scores and higher or comparable CPU times. The selected combinations correspond to rows 1, 4, and 7 in Table 1 with the average F_1 scores in bold.

The testing process is identical to the validation process: after the model is trained with 10 new samples, its performance is measured by computing the average F_1 score on the testing set (using the optimized hyperparameters from the validation stage). Our results are summarized in Table 2. We found that the LSTM model yielded the highest F_1 score of 0.75. The CNN and RNN models achieved a 0.73 and 0.70 F_1 score, respectively. Based on these results and the previously discussed optimization standards, we selected the optimized CNN model for our classifier. In particular, the CNN simulation not only yielded a competitive average F_1 score of 0.73, but also achieved this score 6 to 8 times more quickly than the LSTM or RNN (Fig. 3), which is

Table 1. Average precision, recall, F_1 score, and CPU time for the top three performing hyperparameter combinations on each of the CNN, LSTM, and RNN models. Bold numbers correspond to the highest F_1 scores and lowest CPU times for each of the three model types. We report the recall, precision, and F_1 score to four decimal places (when necessary) to distinguish the average F_1 scores.

Model	Learning Rate	Batch Size	Epochs	Dropout	Recurrent Dropout	Filter Size	Kernel Size	Optimizer	Average Precision	Average Recall	Average F_1 score	CPU Time (sec)
CNN	0.0079	10	1	–	–	16	2	Adam	0.75	0.73	0.74	503.82
CNN	0.01	50	2	–	–	16	2	Adagrad	0.73	0.71	0.72	522.47
CNN	0.0063	10	3	–	–	16	2	Adam	0.73	0.71	0.72	553.43
LSTM	0.0002	10	10	0.4	0.2	–	–	Adam	0.7597	0.7475	0.7534	4241.97
LSTM	0.0002	20	8	0.2	0.6	–	–	Adam	0.7597	0.7468	0.7530	4100.37
LSTM	0.0006	100	12	0.6	0.6	–	–	Adam	0.7559	0.7431	0.7493	4209.37
RNN	0.0001	10	7	0.0	0.2	–	–	Adam	0.7037	0.6957	0.6996	3069.81
RNN	0.0001	20	5	0.0	0.0	–	–	Adam	0.7028	0.6921	0.6973	2805.52
RNN	0.0001	100	12	0.0	0.2	–	–	Adam	0.70	0.69	0.69	3160.35

Table 2. Testing results with the optimal hyperparameter combinations for the CNN, LSTM, and RNN models. The bold numbers correspond to the highest F_1 score and lowest CPU time among the three models.

Model	Average Precision	Average Recall	Average F_1 score	CPU Time (sec)
CNN	0.74	0.73	0.73	501.10
LSTM	0.76	0.74	0.75	4211.01
RNN	0.70	0.69	0.70	3085.59

significant in terms of responding to user feedback in a timely manner.

The optimized CNN model yielded 0.74 and 0.73 average precision and recall scores respectively (Table 2, row 1). This model performance may be due to the initial lack of sufficient training data and difficulty in classifying certain tweets. For instance, after examining the testing dataset, we found that many misclassified tweets were extremely short (e.g., the tweet “screams internally” was misclassified as “Relevant”) or contained complex disaster-related diction (e.g., the tweet “emergency dispatchers in boone county in the hot seat” was misclassified as “Relevant”). However, as we demonstrate in the next section, our model still outperforms state-of-the-art learning models on tweet datasets.

It is worth noting that we do not save the trained model from the validation or testing stages for evaluation in the next stage (or for use with SMART 2.0). We only save the optimized hyperparameters. This is because we assume that users start training a new model (for any event or topic they choose) by labeling the incoming stream of tweets.

In this section, we optimized the model on a sufficiently large dataset that contained tweets related to several kinds of disasters. In the next section, we evaluate the model on datasets containing tweets on specific events, which is representative of cases for situational awareness.

3.3.5 Evaluation

To further demonstrate the optimized CNN model’s performance, we computationally evaluated it on wildfire, bombing, and train crash datasets from CrisisLexT26 [34], each of which contain approximately 1,000 tweets collected during 2012 and 2013 from 26 major crisis situations labeled by relevance. We apply a similar process to evaluate our optimized CNN model on these datasets as we did with the Figure Eight [1] dataset. Specifically, we split the data into 50% training and 50% testing sets (to replicate the experimental setting of To et al. [45], against which we will compare our results), train the model by supplying 10 tweets from the training set at a time (to simulate user labeling of streaming data), evaluate the resulting model on the entire testing set, and then average the F_1 scores for each evaluation.

We summarize our results in Table 3 and graph the model’s performance for retraining with 10 new incoming tweets in Fig. 4, 5, and 6.

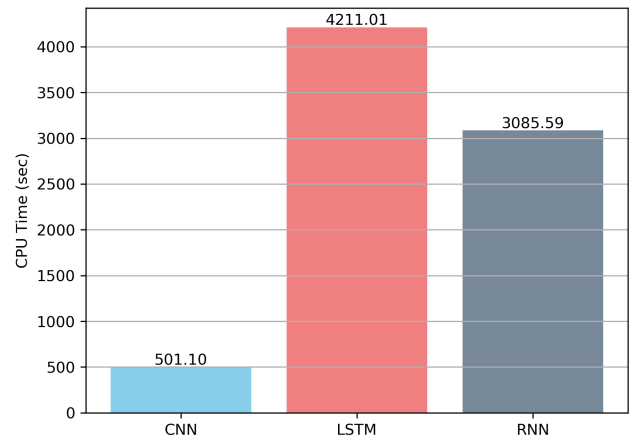


Fig. 3. The total CPU time required for each model to complete the testing simulation. The CNN model is noticeably faster than both the LSTM and RNN models.

In addition, we report the average CPU times to train the model during a single iteration (10 tweets) with each dataset in Table 3. Since the datasets vary slightly in size, we only compute the averages from the first 45 iterations since the smallest dataset (Boston Bombings) required 45 iterations to complete the simulation. We found that per-iteration training was fast and approximately 0.5 seconds with each dataset, which meets our timing demands (DG4).

We obtained 0.71, 0.64, and 0.88 F_1 scores for the Colorado wildfires, Boston bombings, and NY train crash datasets, respectively. Interestingly, the variance of the F_1 scores over the datasets is significant. The textual data in the Boston bombings dataset, which yielded the lowest average F_1 score, was not as easy to separate into the different relevance categories by the model compared with the other two datasets. However, the F_1 score does eventually converge towards a higher value similar to the other datasets, indicating the potential presence of outliers during the first few training iterations. In addition, we found that the simulations converged to the average F_1 scores after training with approximately 190–230 tweets, depending on the dataset, meaning that users need to label 190–230 tweets to achieve the reported F_1 scores. However, the CrisisLexT26 datasets also correspond to specific events, such as wildfires. As such, we surmise that interactively training the model on specific, well-defined events will reduce the amount of training data needed to achieve satisfactory results than with generic constraints on relevance (e.g., a classifier about safety in general).

Finally, we compare our results with the learning-based algorithm employed by To et al. [45], who also evaluated their model’s performance with CrisisLexT26 datasets. In particular, their learning-based approach used Word2Vec, TF-IDF, latent semantic indexing, and logis-

Table 3. Average precision, recall, and F_1 score for three CrisisLexT26 [34] datasets.

Category	Average Precision	Average Recall	Average F_1 score	Average CPU Time (sec)
Colorado wildfires	0.72	0.71	0.71	0.49
Boston bombings	0.64	0.65	0.64	0.50
NY train crash	0.86	0.90	0.88	0.49

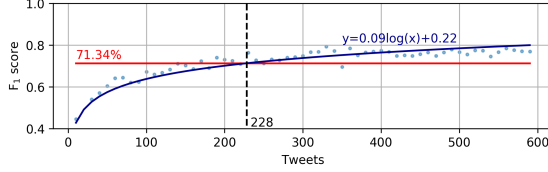


Fig. 4. Optimized CNN F_1 score per training iteration of 10 tweets with the Colorado wildfires dataset (Table 3). The F_1 scores are logarithmically fitted and intersect with the average F_1 score (0.7134) at 228 tweets.

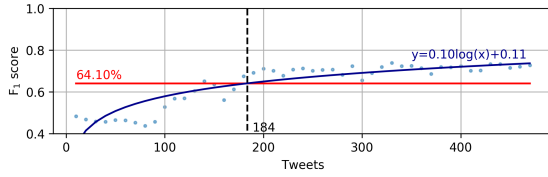


Fig. 5. Optimized CNN F_1 score per training iteration of 10 tweets with the Boston bombings dataset (Table 3). The F_1 scores are logarithmically fitted and intersect with the average F_1 score (0.6410) at 184 tweets.

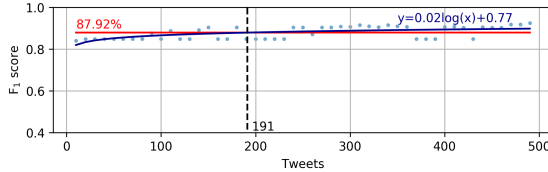


Fig. 6. Optimized CNN F_1 score per training iteration of 10 tweets with the NY train crash dataset (Table 3). The F_1 scores are logarithmically fitted and intersect with the average F_1 score (0.8792) at 191 tweets.

tic regression for classifying data as relevant or irrelevant. The authors of that study split the dataset into two equal parts: one for training and one for testing. They trained the model once (as opposed to our iterative approach) and evaluated on the testing set. Their algorithm was able to yield high precision scores between 0.85–0.95, compared to our scores of 0.64–0.86. However, their recall scores were approximately 0.22–0.45, considerably lower than our recall scores of 0.65–0.90. Therefore, our approach outperforms the learning-based model presented by [45], in terms of the overall F_1 score: our interactive approach achieves F_1 scores of 0.64–0.88 (depending on the dataset) compared to 0.45–0.64 by [45]. The authors also presented a matching-based approach that achieved a much higher F_1 score of 0.54–0.92, which is comparable to ours. However, they generate the set of hashtags to be used for matching by scanning all of the tweets in the dataset. Since we assume the data is streamed in real-time, and therefore, not available altogether, we use an iterative learning approach.

4 SMART 2.0

4.1 SMART

The Social Media Analytics and Reporting Toolkit (SMART) [52,53] is a visual analytics application designed to support real-time situational

awareness for first responders, journalists, government officials, and special interest groups. SMART obtains real-time publicly available geo-tagged data from the Twitter streaming API. The user is able to explore the trending and abnormal topics on various integrated visualizations, including spatial topic model visualization and temporal views. The tweet time chart and theme river visuals convey the temporal distributions of topics if the user wishes to determine how the content of streamed social data has changed over time.

SMART uses string matching-based classifiers to visualize relevant data. Specifically, the user can either (a) select pre-defined filters, such as *Safety* or *Weather* (Fig. 7(c)), each using a series of related keywords for inclusion and exclusion of tweets in the subsequent topic-modeling (Fig. 7(f)) and (geo)visualizations (Fig. 7(b)), or (b) create their own filters by supplying keywords, and intersect or union multiple filters according to their needs. However, keyword-based matching is insufficient for finding relevant information as it fails to accurately capture semantic relevance and therefore effectively filter out noisy data. As an example, if the user were to apply the *Safety* classifier, it would be possible for the tweets “My house is on fire!” and “I was just fired from my job.” to pass through the filter since they both include the keyword *fire*. However, the latter is unrelated to the intended semantic context of *Safety* and thus dilutes the filter’s quality.

To address this problem, we integrate our interactive learning framework (the focus of this paper) in the existing SMART application [52,53] and seek domain expert feedback on the use of these models. We call the resulting extended application SMART 2.0. SMART 2.0 allows users to define string matching-based keyword filters (similar to SMART), but adds the ability for users to then iteratively refine and train the newly integrated models by labeling the filtered data as semantically relevant or not. In addition, the SMART 2.0 interface includes interactive visuals to facilitate user exploration, filtering, and refinement of relevant data (Fig. 7).

As with the model simulations in Section 3.3, SMART 2.0’s underlying models are trained with successive batches of 10 user-labeled tweets. In cases where model predictions conflict with user labels, user labels override the model’s since they represent the ground truth. In addition, users should not need to manually relabel the same data multiple times. Although conflicts might indicate that the model is not sufficiently trained, the model trains with the same data during several successive iterations (as discussed in Section 3.3.4), so conflicts might be resolved after future iterations.

4.2 SMART 2.0 Interface

The extensions to SMART 2.0’s user interface, compared with SMART, concern the new interactive visuals that allow users to iteratively train machine learning models, utilize model predictions for rapid relevancy identification, and understand a model’s reliability. The SMART 2.0 interface (Fig. 7) extends the interactive features of SMART for relevance identification in three primary ways:

1. Extending the tweet table (containing a tweet’s creation date and text) by including the predicted relevance label, relevance label probabilities, label modification, model training performance, and relevance filtering.
2. Extending the interactive map containing the geo-tagged tweets whose relevancy can be individually inspected or modified.
3. Altering the content of existing SMART views (e.g., topic models and spatial topic lenses) using either all data or only relevant data (as identified by the model and corrected by the user).

4.2.1 Table

The SMART 2.0 table (Fig. 7(g)–(j)) is extended from SMART in that it not only provides a tweet’s creation date and text, but also provides the predicted relevance label (Fig. 7(i)) and the probabilities of a tweet belonging to any of the relevance classes (Fig. 7(j)) (DG1).

In particular, the relevance of a tweet can be “Relevant”, “Not Relevant”, or “Can’t Decide”. The “Relevant” label is colored blue, the “Not Relevant” label red, and the “Can’t Decide” label gray to visually

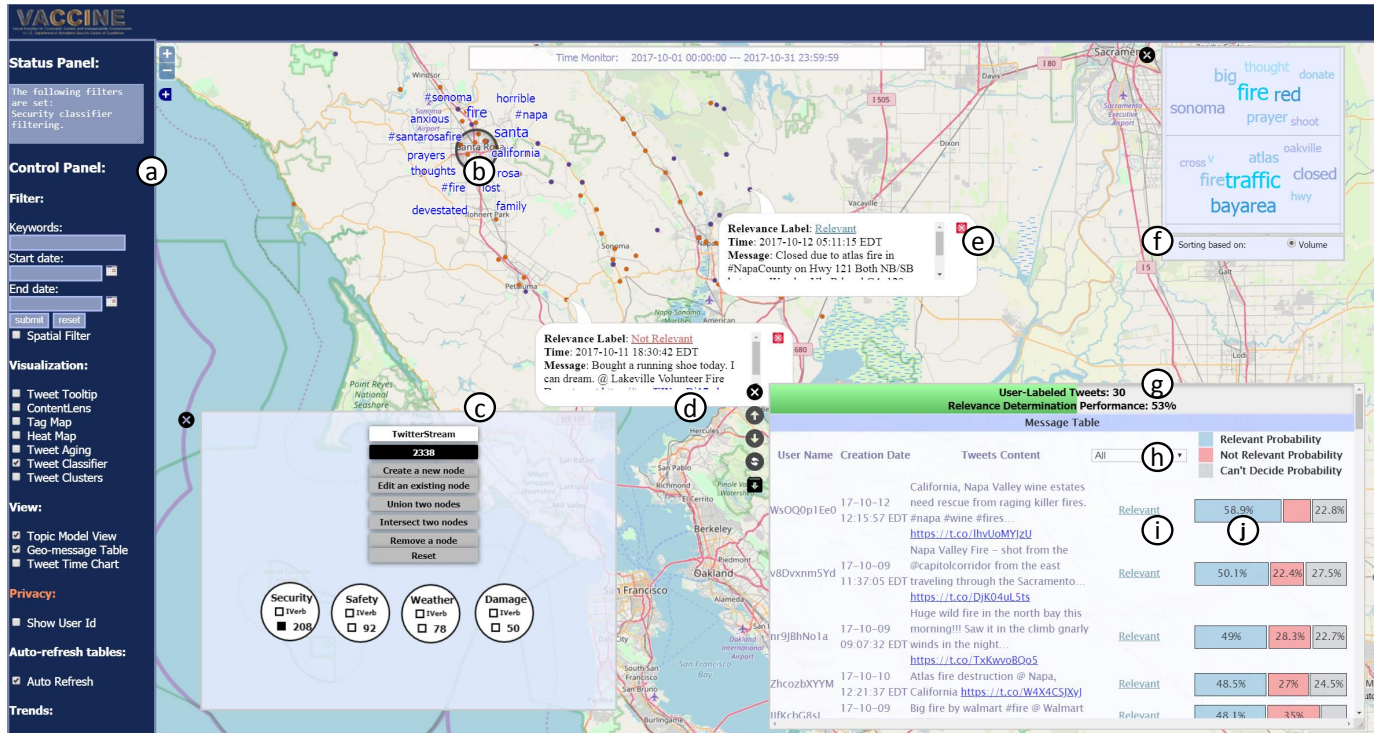


Fig. 7. SMART 2.0 overview: (a) The control panel provides several filters, visualizations, and views. (b) The content lens visualization provides the most frequently used words within a selected area. (c) The tweet classifier visualization provides keyword-based filters to help reduce noisy data. (d)(e) Clicking a tweet on the map with the tweet tooltip visualization displays the tweet's time, message, and relevance label. (f) The topic-modeling view, based on Latent Dirichlet Allocation, extracts trending topics and the most frequently used words associated with each topic among tweets with specified relevancy. (g)–(j) The message table aggregates the tweets for efficient exploration with (g) the model's estimated classification performance (F_1 score), (h) a drop down box to filter data by their relevance labels, (i) color-coded relevance labels that can be changed by clicking on the label itself, and (j) associated relevance probabilities. Tweet map symbols are colored orange and purple to distinguish Twitter data from Instagram-linked tweets, respectively, since the latter contains potentially useful images for situational awareness.

separate tweets with different relevance. SMART's preexisting blue color scheme motivated us to use the blue, red, and gray diverging coloring for relevancy in order to maintain visual appeal and harmony.

Users can directly click on relevance labels to correct the classifier's prediction (DG2). For instance, if a tweet is incorrectly marked "Relevant", clicking the label will change it to "Not Relevant" or "Can't Decide", depending on the label the user wishes to assign. Further, a drop down box is included at the top of the relevance label column (Fig. 7(h)), which provides the option to filter out data that does not have a specified relevancy (DG1). For example, by selecting "Relevant" from the drop down box, the table will remove tweets with labels "Not Relevant" and "Can't Decide" from all views and visualizations in SMART, including geovisualizations and temporal views.

The table also displays the *degree* (or confidence) of a tweet's relevancy. In specific, the probabilities of a tweet being "Relevant", "Not Relevant", or "Can't Decide" are represented as a horizontal segmented bar graph and sized proportional to their respective percentages (Fig. 7(j)). In addition, the user can sort tweets based on relevancy probability in ascending or descending order.

We provide the relevance probabilities and associated sorting actions as a supplementary relevance filtering mechanism (DG1). In particular, it is possible for tweets to be classified as "Relevant" by the model, for example, but with low confidence. The probability filtering allows the user to specifically view high-confidence relevant data and therefore further reduce potentially noisy data.

The table provides a performance bar that encodes the estimated performance (F_1 score) of the underlying learning model (Fig. 7(g)), as well as the number of user-labeled tweets, to inform the user of the model reliability. Since labeled testing data is not available to evaluate the model for real-time training (because we assume the user may train on any type of event data and has their own specifications

for relevancy), the model's performance can only be estimated. Based on our evaluations in Section 3.3.5 with datasets typical of situational awareness scenarios (Table 3), the Colorado wildfires dataset generated the F_1 score (0.71) closest to the average of the three datasets (0.74). Therefore, we use the Colorado wildfires dataset's logarithmic trendline $y = 0.09 \log_e(x) + 0.22$ (Fig. 4) to approximate the model's F_1 score as a function of the number of user-labeled tweets.

4.2.2 Map

The SMART 2.0 map is extended from SMART in that it includes a tweet's relevance label (which can be modified) in addition to its text and creation date (Fig. 7(d)(e)). Through the *Tweet Tooltip*, the user can directly click on tweet symbols on the map to view their text and associated relevancy (DG1). In addition, the user can correct the classified relevance label (DG2) by clicking on the label itself. Map inspection can allow the user to view and investigate potential geographical relevancy trends. For example, during crisis events, relevant tweets might be closely grouped on the map, so it may be more beneficial for the user to view predicted relevance from the map itself.

The interactions between the table and map are synchronized. If the user relabels data on the map, the associated new label will also be updated in the table, and vice versa. In addition, selecting a relevancy filter from the drop down box in the table filters the tweets on the map.

4.2.3 Integration with Existing Visualizations

Many of SMART's original visualizations, such as the topic-model views, spatial topic lenses, and temporal views help users make sense of spatiotemporal text data. Therefore, we integrated all of these views in SMART 2.0 with the relevance extensions.

Users have the option to view only relevant or all the data (including irrelevant tweets) in various visualizations in case the interactive

classifiers are not yet trained to desirable accuracies since, as we show in Section 3.3.5, classifiers typically require around 200 user-labeled tweets to achieve F_1 scores of 0.70–0.80. If they choose to view only relevant tweets, any relevance filtering action also updates the data used by other visuals. For example, the topic-modeling view (Fig. 7(f)) extracts the top 10 topics from the tweets and displays the most frequently used words for each topic. If the user filters out irrelevant tweets, the topic-modeling view will only be applied to the remaining relevant tweets. It is important to note that the majority of visualizations in SMART 2.0 require a minimum number of tweets in order to render. When filtered relevant data is scarce, visualizations do not populate, in which case users can individually inspect tweets. For instance, the topic-modeling view requires at least 10 tweets to extract topics.

Overall, SMART 2.0’s suite of visualization tools can be used in combination with relevance interactions to further understand trends and important spatiotemporal characteristics of relevant data.

5 USER EXPERIENCE

In this section, we provide usage scenarios and feedback from domain experts that demonstrate our framework’s effectiveness and usability.

5.1 Usage Scenario 1

Alice is an emergency dispatcher interested in identifying people in need for help or hazardous locations during a hurricane. She uses SMART 2.0 to find any related social media posts near the affected area. She adds a new filter *Hurricane* and provides an appropriate set of filter keywords such as “hurricane”, “help”, “blocked”, and “trapped”.

After applying the *Hurricane* filter, she explores the filtered tweets in the table and finds a tweet labeled “Relevant” that says “Does anyone know how to get help setting up my TV?”. Since the tweet is unrelated to a hurricane, she relabels it as “Not Relevant”. After further browsing the table, she finds a tweet that says “The road near Taylor Loop is blocked from a broken tree.”, but it is labeled as “Not Relevant”. Since the tweet contains actionable information, she relabels it as “Relevant”. After labeling several more tweets for model training and noticing that the model predicts correctly, she decides to only view “Relevant” tweets and sort them by most relevant. She promptly identifies a tweet posted only a few minutes ago marked as highly relevant. It reads “Car just crashed into tree blocking road near Taylor Loop!”. Alice immediately notifies first responders of the location to provide assistance.

By using SMART 2.0, Alice is able to identify important, relevant data more quickly through interactively training the model to remove noise and then filtering by relevance.

5.2 Usage Scenario 2

To demonstrate the generalizability of our framework to other domains, we applied our interactive framework in real-time during the Purdue vs. Virginia 2019 March Madness game in the Kentucky area. We assumed the role of a journalist who wanted to follow public discourse on the game by identifying the relevant tweets. We first constructed a *Sports* filter, which included keywords such as “Purdue”, “game”, “score”, and “#MarchMadness”. We then interacted with the streaming data by iteratively labeling the relevancy of tweets (from scratch) and found that the system correctly classified incoming data after roughly 80 training samples (Fig. 8). We noticed that the time intervals between successive trainings increased, indicating that it was more difficult to find incorrectly labeled data towards the end and that the model gradually learned from user feedback. In particular, the interval between the first and second training iterations was 2 minutes, whereas the interval between the final two was 4 minutes.

5.3 Domain Expert Feedback

We piloted SMART 2.0 with two groups of first responders, each containing two individuals, who frequently use SMART during events for situational awareness in their operations. Both groups participated in separate 1-hour long sessions via conference call in which they iteratively trained a classifier from scratch and applied relevance filtering and visualizations to assess the implemented framework. They received a tutorial of SMART 2.0 30 minutes beforehand and were provided

with web access to the system to complete the session. For both groups, we simulated the real-time use of SMART 2.0 by feeding in a stream of historical data on events (previously collected). For the first group, the system presented unlabeled tweets from the Las Vegas shooting on October 1, 2017 in the Las Vegas area. For the second group, we used unlabeled tweets from the October 2017 Northern California wildfires. We used historical event datasets to ensure the existence of sufficient training relevant samples for a situational awareness scenario.

The domain experts in the first group applied the *Safety*, *Damage*, and *Security* filters during the iterative training process, resulting in 317 tweets. They trained on the same underlying model for all three filters, as they considered them semantically related. In total, after relabeling approximately 200 tweets, they indicated that they could trust the model to predict accurately and were pleased that the tweets they had not seen before from the *Security* filter were labeled correctly. Their definition for relevancy was tweets containing actionable information. For instance, they marked tweets containing information about road closures, blood drive locations, or death counts as relevant. They labeled data with general comments regarding the shooting, such as “I hope everyone is safe now...terrible shooting...”, as irrelevant since they did not provide actionable information. Interestingly, they also marked tweets that would influence public opinion (and therefore may cause action) such as those from bots or trolls as relevant since they still contained actionable information.

The domain experts from the second group followed a similar process in which they applied the *Safety*, *Damage*, and *Security* filters, resulting in 445 tweets, and trained a learning model for relevance. They found that after training on roughly 67 tweets, the model satisfactorily predicted relevancy. As with the first group, these domain experts labeled tweets as relevant if they contained actionable information.

The domain experts from both groups found SMART 2.0 to be easy to use and effective in identifying important data. For instance, they discovered relevant, actionable information after training the model: specific blood drive locations to aid shooting victims. Notably, the users mentioned that they felt less the need to relabel data as they progressed since the system provided more correct labels. They were pleased that they had the option to view only relevant data, but could see all of the data regardless of relevancy to avoid potentially missing important misclassified data, and that the model was responsive to user training. In addition, they found the relevance percentage bars to be helpful in determining the tweets that were potentially the most relevant.

One concern the domain experts had was that SMART 2.0 does not indicate the number of tweets that are predicted as relevant. They felt this extension could help them infer the occurrence of events or potential crises. For example, the number of relevant tweets for the *Safety* classifier would likely increase significantly during a widespread disaster. We plan to introduce this feature in the next development cycle. However, we have added a visualization of estimated model performance in SMART 2.0 (Fig. 7(g)) to help users ascertain the reliability of a model’s relevancy predictions.

Overall, the feedback from the domain experts was positive and helpful, indicating the system’s practicality and usefulness in facilitating real-time situational awareness. In addition, they have asked to use SMART 2.0 in their emergency operations center.

6 DISCUSSION AND FUTURE WORK

Our interactive learning framework and SMART 2.0 integration were developed with the user in mind, influencing all of our design, computational evaluation, and implementation choices. Our user-centered model and SMART 2.0 application contribute to both the machine learning and visual analytics communities. We bridge the two fields by demonstrating how models can be interactively trained and evaluated while keeping the user in mind, and used to facilitate situational awareness for real-life, practical use.

SMART 2.0 currently only collects English tweets, although supporting non-English languages (one at a time) with our current design (e.g., Spanish only) is straightforward since Word2Vec embeddings can be independently trained on a corpus in the target language [7]. Extending our system to support multilingual tweets would be a powerful asset,

xhdUCgkm6L	19-03-31 09:47:45 EDT	GURU Sports @ Lawrenceburg, Kentucky https://t.co/Eaj3BXWVAd	Not Relevant	32.1%	65.5%
uZWVby932r	19-03-31 07:11:58 EDT	@Emceegreg But with sports, most people are on board. I feel weird because I'm not. Like 1% can fly, and many hate... https://t.co/GrF6xvAzX2	Not Relevant		68.9%
ahpZ9WznKu	19-03-30 21:00:33 EDT	@boilerball ☐☐ #boilerup #boilermakers #purdue #marchmadness #elite8 @ Louisville, Kentucky https://t.co/CgWYqyp1aK	Relevant	61%	38.9%
kNqZZwQf05	19-03-30 20:55:52 EDT	#marchmadness has hit the 'Ville! Purdue looks to have the fam support tonight @ KFC Yum! Center https://t.co/M66YT1Xvev	Relevant	77.6%	21.6%
VOQLyHC5Ic	19-03-30 20:09:23 EDT	#BoilerUp @boilerball #CarsenEdwards #MarchMadness @ KFC Yum! Center https://t.co/gcmSxtwPuc	Relevant	55.2%	44.7%
jb81Y4KEP3	19-03-30 19:58:07 EDT	It's raining like cats and dogs outside! However, it's nice and cozy inside for Brax's Birthday Video Game Truck 🎮... https://t.co/aClyK9O39R	Not Relevant		82.7%
E2OrPlmFI9	19-03-30 19:06:33 EDT	View from the office tonight. #MarchMadness (at @KFC Yum! Center in Louisville, KY) https://t.co/YZ4E8CYrB2 https://t.co/TMe8PSAs04	Relevant	79.7%	

Fig. 8. Tweets with correctly predicted relevancy from the Purdue vs. Virginia 2019 March Madness game after the user (re)labels 80 tweets.

especially for multilingual users, in amplifying situational awareness by leveraging relevancy of tweets issued in different languages. However, the multilingual model performance evaluation and testing is an open area for research. In addition, determining the specifics of how a single relevance classifier might be trained with multilingual tweets requires careful attention. For instance, training iterations with Spanish tweets should also affect the relevancy of semantically-related English tweets. Since similar words in different languages likely have different vector representations (embeddings), multilingual mappings must be learned or training must be performed differently, such as with parallel corpora [7]. Multilingual support also requires changes in SMART 2.0's language-dependent visualizations, such as the topic-modeling view (Fig. 7(f)). Translation to a unified language or extracting topics separately for each language are two potential solutions.

The scalability of our framework is a natural concern, especially since SMART and many deployed real-time visual analytics applications contain multiple users who require responsive interfaces while monitoring crisis events. We deliberately designed the framework architecture with scalability in mind. As mentioned in Section 3.3.4, we selected the model and optimal hyperparameters based on training/CPU time in an effort to maximize the model's computational speed. Further, SMART 2.0 filters and views at most 800-900 tweets at a time, although user-specified filtering (typical in situational awareness scenarios) reduces the data to only a few hundred tweets, as demonstrated in Section 5.3. It takes only 2-3 seconds to calculate and retrieve their relevance labels over the network from the server where the model resides, and per-iteration training is fast, as established in Section 3.3.5.

Training a model during a particular event, such as a disaster, can be straightforward due to potentially larger amounts of relevant data. However, social media data during periods without major events are likely to contain very few, if any, relevant tweets. As such, if the user *only* trains the model on irrelevant data, it will poorly predict relevant data since it has only be taught what is irrelevant. Although the user can improve the model through training during a real-life disaster, they are required to know when and where the disaster occurs to begin training. This can be problematic if the user wishes to rely on relevancy predictions to detect hazardous situations.

To accommodate time periods in which relevant data is scarce, we plan to introduce an interactive feature in which users can provide example tweets or external resources for specific relevance labels. For instance, if the user would like to train a *Hurricane* classifier before a hurricane event, they could provide a relevant text such as "I'm stranded by this hurricane. Please help!". The model could then detect relevant tweets once the hurricane begins as opposed to requiring user training during the event. In addition, we plan to provide the user with the

option to visit specific historical data to train existing classifiers, as done by Bosch et al. [9].

Our interactive learning performed well on target datasets (i.e., wild-fire, bombing, and crash) as explained in Section 3.3.5. Specifically, it required the users to label approximately 200 tweets to achieve acceptable F_1 scores. However, tweets are short, and therefore, more research is required to investigate the suitability of our approach for "general" classifiers, such as ones that learn to classify relevant data to "safety". As safety can be affected by many events or situations, the model may need additional training that is typical of a targeted dataset.

Finally, although we rigorously optimize and evaluate our machine learning model, the hyperparameter combinations were only tuned with the Figure Eight dataset [1]. Since optimal hyperparameters can depend on the dataset itself, it is possible our model may not be optimally tuned for different datasets, even though that optimization may be negligible from a user standpoint. We did use other datasets in our model evaluation to show the satisfactory resulting performance (Section 3.3.5). Given that the Figure Eight dataset classifies generic events as relevant or irrelevant, as opposed to specific events, we expect that our model performs well on many different event types.

7 CONCLUSION

We presented a novel interactive framework in which users iteratively (re)train neural network models with streaming text data in real-time to improve the process of finding relevant information. We optimized and evaluated a machine learning model with various datasets related to situational awareness and adapted the model to learn at interactive rates. According to evaluation results, our model outperforms state-of-the-art learning models used in similar classification tasks. Finally, we integrated our framework with the SMART application and extended it to SMART 2.0, allowing users to interactively explore, identify, and refine tweet relevancy to support real-time situational awareness. Our discussions with multiple first responders who use SMART 2.0 indicated positive feedback and user experience. In particular, their assessments demonstrated that our interactive framework significantly improved the time-consuming process of finding crucial information during real-time events.

8 ACKNOWLEDGEMENTS

This material is based upon work funded by the U.S. Department of Homeland Security (DHS) VACCINE Center under Award No. 2009-ST-061-CI0003, DHS Cooperative Agreement No. 2014-ST-061-ML0001, and DHS Science and Technology Directorate Award No. 70RSAT18CB0000004.

REFERENCES

- [1] Figure eight: Disasters on social media. <https://www.figure-eight.com/data-for-everyone/>.
- [2] Keras. <https://keras.io>.
- [3] Word2vec. <https://code.google.com/archive/p/word2vec>.
- [4] O. Abdelwahab and A. Elmaghraby. UofL at SemEval-2016 task 4: Multi domain word2vec for twitter sentiment classification. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pp. 164–170, 2016. doi: 10.18653/v1/s16-1024
- [5] Z. Ashktorab, C. Brown, M. Nandi, and A. Culotta. Tweedr: Mining twitter to inform disaster response. In *Proceedings of the 11th International ISCRAM Conference*, pp. 354–358, 2014.
- [6] H. Becker, M. Naaman, and L. Gravano. Beyond trending topics: Real-world event identification on twitter. In *Proceedings of the 5th International Conference on Weblogs and Social Media*, pp. 438–441, 2011.
- [7] A. Bérard, C. Servan, O. Pietquin, and L. Besacier. MultiVec: a multilingual and multilevel representation learning toolkit for NLP. In *The 10th edition of the Language Resources and Evaluation Conference (LREC)*, 2016.
- [8] J. Bernard, M. Hutter, M. Zeppelzauer, D. Fellner, and M. Sedlmair. Comparing visual-interactive labeling with active learning: An experimental study. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):298–308, 2018. doi: 10.1109/tvcg.2017.2744818
- [9] H. Bosch, D. Thom, F. Heimerl, E. Puttmann, S. Koch, R. Kruger, M. Worner, and T. Ertl. Scatterblogs2: Real-time monitoring of microblog messages through user-guided filtering. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2022–2031, 2013. doi: 10.1109/tvcg.2013.186
- [10] M. Chen, X. Jin, and D. Shen. Short text classification improved by learning multi-granularity topics. In *22nd International Joint Conference on Artificial Intelligence*, pp. 1776–1781, 2011.
- [11] A. Dal Pozzolo, O. Caelen, R. A. Johnson, and G. Bontempi. Calibrating probability with undersampling for unbalanced classification. In *IEEE Symposium Series on Computational Intelligence*, pp. 159–166, 2015. doi: 10.1109/ssci.2015.33
- [12] N. Diakopoulos, M. Naaman, and F. Kivran-Swaine. Diamonds in the rough: Social media visual analytics for journalistic inquiry. In *IEEE Symposium on Visual Analytics Science and Technology*, pp. 115–122, 2010. doi: 10.1109/vast.2010.5652922
- [13] J. G. Dy and C. E. Brodley. Visualization and interactive feature selection for unsupervised data. In *Proceedings of the 6th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 360–364, 2000. doi: 10.1145/347090.347168
- [14] J. L. Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990. doi: 10.1016/0364-0213(90)90002-e
- [15] A. Endert, P. Fiaux, and C. North. Semantic interaction for visual text analytics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pp. 473–482, 2012. doi: 10.1145/2207676.2207741
- [16] S. Ghosh and M. S. Desarkar. Class specific TF-IDF boosting for short-text classification: Application to short-texts generated during disasters. In *Companion of the The Web Conference 2018 (WWW '18)*, pp. 1629–1637, 2018. doi: 10.1145/3184558.3191621
- [17] A. Go, R. Bhayani, and L. Huang. Twitter sentiment classification using distant supervision. *CS224N Project Report*, Stanford, 1(12), 2009.
- [18] F. Heimerl, S. Koch, H. Bosch, and T. Ertl. Visual classifier training for text document retrieval. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2839–2848, 2012. doi: 10.1109/tvcg.2012.277
- [19] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735
- [20] X. Hu, N. Sun, C. Zhang, and T.-S. Chua. Exploiting internal and external semantics for the clustering of short texts using world knowledge. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, CIKM '09, pp. 919–928, 2009. doi: 10.1145/1645953.1646071
- [21] M. Imran, P. Mitra, and C. Castillo. Twitter as a lifeline: Human-annotated twitter corpora for NLP of crisis-related messages. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC)*, 2016.
- [22] M. Kahng, P. Andrews, A. Kalro, and D. Chau. ActiVis: Visual exploration of industry-scale deep neural network models. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):88–97, 2018. doi: 10.1109/tvcg.2017.2744718
- [23] S. Karimi, J. Yin, and C. Paris. Classifying microblogs for disasters. In *Proceedings of the 18th Australasian Document Computing Symposium*, ADCS '13, pp. 26–33, 2013. doi: 10.1145/2537734.2537737
- [24] S. Lai, L. Xu, K. Liu, and J. Zhao. Recurrent convolutional neural networks for text classification. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, vol. 3, pp. 2267–2273, 2015.
- [25] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791
- [26] J. Lilleberg, Y. Zhu, and Y. Zhang. Support vector machines and word2vec for text classification with semantic features. In *IEEE 14th International Conference on Cognitive Informatics & Cognitive Computing (ICCICC)*, pp. 136–140, 2015. doi: 10.1109/icci-cc.2015.7259377
- [27] W. Ling, C. Dyer, A. Black, and I. Trancoso. Two too simple adaptations of word2vec for syntax problems. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 1299–1304, 2015. doi: 10.3115/v1/n15-1142
- [28] A. Marcus, M. Bernstein, O. Badar, D. Karger, S. Madden, and R. Miller. TwitInfo: aggregating and visualizing microblogs for event exploration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pp. 227–236, 2011. doi: 10.1145/1978942.1978975
- [29] M. Martínez-Rojas, M. D. C. Pardo-Ferreira, and J. C. Rubio-Romero. Twitter as a tool for the management and analysis of emergency situations: A systematic literature review. *International Journal of Information Management*, 43:196–208, 2018. doi: 10.1016/j.ijinfomgt.2018.07.008
- [30] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pp. 3111–3119, 2013.
- [31] T. H. Nazer, F. Morstatter, H. Dani, and H. Liu. Finding requests in social media for disaster relief. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 1410–1413, 2016. doi: 10.1109/asonam.2016.7752432
- [32] D. T. Nguyen, S. Joty, M. Imran, H. Sajjad, and P. Mitra. Applications of online deep learning for crisis response using social media information. *arXiv preprint arXiv:1610.01030*, 2016.
- [33] D. T. Nguyen, K. A. A. Mannai, S. Joty, H. Sajjad, M. Imran, and P. Mitra. Rapid classification of crisis-related data on social networks using convolutional neural networks. *arXiv preprint arXiv:1608.03902*, 2016.
- [34] A. Olteanu, S. Vieweg, and C. Castillo. What to expect when the unexpected happens: Social media communications across crises. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, pp. 994–1009, 2015. doi: 10.1145/2675133.2675242
- [35] M. Pennacchiotti and A.-M. Popescu. A machine learning approach to twitter user classification. In *5th International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2011.
- [36] S. Pezanowski, A. M. Maceachren, A. Savelyev, and A. C. Robinson. SensePlace3: A geovisual framework to analyze place-time-attribute information in social media. *Cartography and Geographic Information Science*, 45(5):420–437, 2018. doi: 10.1080/15230406.2017.1370391
- [37] X.-H. Phan, L.-M. Nguyen, and S. Horiguchi. Learning to classify short and sparse text & web with hidden topics from large-scale data collections. In *Proceedings of the 17th International Conference on World Wide Web*, pp. 91–99, 2008. doi: 10.1145/1367497.1367510
- [38] K. Rudra, S. Ghosh, N. Ganguly, P. Goyal, and S. Ghosh. Extracting situational information from microblogs during disaster events: a classification-summarization approach. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*, vol. 19-23-of CIKM '15, pp. 583–592, 2015. doi: 10.1145/2806416.2806485
- [39] D. Sacha, M. Kraus, D. A. Keim, and M. Chen. VIS4ML: An ontology for visual analytics assisted machine learning. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):385–395, 2019. doi: 10.1109/tvcg.2018.2864838
- [40] D. Sacha, M. Sedlmair, L. Zhang, J. Aldo Lee, D. Weiskopf, S. North, and D. Keim. Human-centered machine learning through interactive visualization: Review and open challenges. In *Proceedings of the 24th European Symposium on Artificial Neural Networks (ESANN)*, pp. 641–646, 2016.
- [41] M. Sokolova, N. Japkowicz, and S. Szpakowicz. Beyond accuracy, F-score and ROC: A family of discriminant measures for performance evaluation.

AI 2006: Advances in Artificial Intelligence, pp. 1015–2021, 2006. doi: 10.1007/11941439_114

- [42] J. Talbot, B. Lee, A. Kapoor, and D. Tan. EnsembleMatrix: interactive visualization to support machine learning with multiple classifiers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, pp. 1283–1292, 2009. doi: 10.1145/1518701.1518895
- [43] D. Tang, B. Qin, and T. Liu. Document modeling with gated recurrent neural network for sentiment classification. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1422–1432. Association for Computational Linguistics (ACL), 2015. doi: 10.18653/v1/d15-1167
- [44] D. Tang, F. Wei, N. Yang, M. Zhou, T. Liu, and B. Qin. Learning sentiment-specific word embedding for twitter sentiment classification. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, vol. 1, pp. 1555–1565. Association for Computational Linguistics (ACL), 2014. doi: 10.3115/v1/p14-1146
- [45] H. To, S. Agrawal, and C. Shahabi. On identifying disaster-related tweets: Matching-based or learning-based? In *IEEE 3rd International Conference on Multimedia Big Data (BigMM)*, 2017. doi: 10.1109/bigmm.2017.82
- [46] F. Toriumi and S. Baba. Real-time tweet classification in disaster situation. In *Proceedings of the 25th International Conference Companion on World Wide Web*, pp. 117–118, 2016. doi: 10.1145/2872518.2889365
- [47] L. Trieu, H. Tran, and M.-T. Tran. News classification from social media using twitter-based doc2vec model and automatic query expansion. In *Proceedings of the 8th International Symposium on information and communication technology*, vol. 2017- of *SoICT 2017*, pp. 460–467, 2017. doi: 10.1145/3155133.3155206
- [48] E. Wall, S. Das, R. Chawla, B. Kalidindi, E. T. Brown, and A. Endert. Podium: Ranking data using mixed-initiative visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):288–297, 2018. doi: 10.1109/tvcg.2017.2745078
- [49] T. Wilson, J. Wiebe, and P. Hoffmann. Recognizing contextual polarity: An exploration of features for phrase-level sentiment analysis. *Computational Linguistics*, 35(3):399–433, 2009. doi: 10.1162/coli.08-012-r1-06-90
- [50] P. C. Wong, E. G. Hetzler, C. Posse, M. A. Whiting, S. L. Havre, N. O. Cramer, A. Shah, M. Singhal, A. E. Turner, and J. Thomas. IN-SPIRE InfoVis 2004 contest entry. In *Proceedings of the IEEE Symposium on Information Visualization*, 2004. doi: 10.1109/infvis.2004.37
- [51] W. Yin, K. Kann, M. Yu, and H. Schütze. Comparative study of CNN and RNN for natural language processing. *arXiv preprint arXiv:1702.01923*, 2017.
- [52] J. Zhang, S. Afzal, D. Breunig, J. Xia, J. Zhao, I. Sheeley, J. Christopher, D. S. Ebert, C. Guo, S. Xu, J. Yu, Q. Wang, C. Wang, Z. Qian, and Y. Chen. Real-time identification and monitoring of abnormal events based on microblog and emergency call data using SMART. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 393–394, 2014. doi: 10.1109/vast.2014.7042582
- [53] J. Zhang, J. Chae, C. Surakitbanharn, and D. S. Ebert. SMART: Social media analytics and reporting toolkit. In *The IEEE Workshop on Visualization in Practice 2017*, pp. 1–5, 2007.
- [54] T. Zoppi, A. Ceccarelli, F. Lo Piccolo, P. Lollini, G. Giunta, V. Morreale, and A. Bondavalli. Labelling relevant events to support the crisis management operator. *Journal of Software: Evolution and Process*, 30(3), 2018. doi: 10.1002/smr.1874