

Predicting Future Personal Life Events on Twitter via Recurrent Neural Networks

Maryam Khodabakhsh^a, Mohsen Kahani^{a,*}, Ebrahim Bagheri^b

^a*Ferdowsi University of Mashhad*

^b*Department of Electrical and Computer Engineering, Ryerson University*

Abstract

Social network users publicly share a wide variety of information with their followers and the general public ranging from their opinions, sentiments and personal life activities. There has already been significant advance in analyzing the shared information from both micro (individual user) and macro (community level) perspectives, giving access to actionable insight about user and community behaviors. The identification of personal life events from user's profiles is a challenging yet important task, which if done appropriately, would facilitate more accurate identification of users' preferences, interests and attitudes. For instance, a user who has just *broken his phone*, is likely to be upset and also be looking to purchase a new phone. While there is work that identifies tweets that include mentions of personal life events, our work in this paper goes beyond the state of the art by predicting a future personal life event that a user will be posting about on Twitter solely based on the past tweets. We propose two architectures based on *recurrent neural networks*, namely the classification and generation architectures, that determine the future personal life event of a user. We evaluate our work based on a gold standard Twitter life event dataset and compare our work with the state of the art baseline technique for life event detection. While presenting performance measures, we also discuss the limitations of our work in this paper.

Keywords: Life Event Prediction; Recurrent Neural Networks; Social Networks; Twitter

1. Introduction

Social microblogging platforms, such as Twitter, are accessible information sharing environments where individuals, governments and even multi-national corporations share their thoughts, views, sentiments and recent news. The live stream of content that is generated and shared on these platforms serve as a valuable source of information for performing macro and micro-level analyses ranging from the understanding of individual users' interests and preferences [1, 2] to detecting and monitoring the formation and evolution of user communities [3, 4, 5] as well as identifying important trending topics [6, 7] and how they can enable descriptive and prescriptive analytics [8]. Researchers have already extensively explored the possibility of using

*Corresponding Author

Email addresses: maryamkhodabakhsh@mail.um.ac.ir (Maryam Khodabakhsh), kahani@um.ac.ir (Mohsen Kahani), bagheri@ryerson.ca (Ebrahim Bagheri)

Preprint submitted to Intelligent Information Systems

July 19, 2018

socially shared data to build predictive models. For instance, work has been done on using social data to offer macro-scale predictions on dependent variables such as stock market prices [9, 10, 11, 12, 13], movie sales [14, 15], disease outbreak [16, 17] and election results [18, 19, 20], just to name a few. There has also been interesting studies showing that with even scant amount of information from individual users on the social network, it is possible to perform micro-level predictions such as detecting psychopathy [21, 22], anxiety and depression [23, 24].

One of the main focal research areas on social network data is to derive such *actionable insights* in order to facilitate decision making. To this end, detecting and monitoring events on the social network has been of interest as trends and events on social networks have shown to reflect many of the characteristics of the events of the real world [25, 26]. Most of the techniques that focus on the detection of trending topics and events on social networks are based primarily on the hypothesis that trending social topics or events can be characterized by the change in observation frequency. In other words, if a certain topic is frequently mentioned in a certain time period, while the topic had been dormant in the past time periods, it can be considered to be a reflection of an event happening in the real world or a trending issue. Therefore, the success of such *event detection* models relies on the global monitoring of social content across, at least, a subsection of the broad social network.

More recently, researchers have become interested in studying a more nuanced form of events that are specific to a given user, known as *personal life events*. Personal life events, also known as life events, relate to an event or incident that happened for an individual social network user that was then publicly reported by this user on the social network. Life events can include getting engaged, moving to a new house, buying a new phone, and graduating from college, among others. If detected, personal life events can provide a good basis for making recommendations to users. For instance, it would be possible to recommend new phone products on the market to a user who has reported a personal life event related to breaking his/her phone. As another example, it would be possible to recommend relevant career opportunities for a user who is engaged with the life event of changing jobs.

The detection of life events has a different set of challenges compared to the detection of trending topics and global events as these events are only mentioned locally by the user him/herself and at most reacted to by some close social connections. Therefore, methods that rely on changes in content frequency would not be suitable for this purpose. For instance, a user would tweet *'I just got married thanks monbebes for tuning into our wedding ceremony'*; reporting that she/he has gotten married. This tweet was liked three times and re-tweeted only once; therefore, it is quite difficult to determine its significance and importance. Furthermore, social network users use special jargon to report life events that are quite hard to pick out and identify. As an example, a Twitter user posts *'i don't miss being alone!'* and posts a picture of his engagement party. It is difficult to determine this life event given the brevity and rather cryptic nature of the message. Moreover, there are many messages posted on Twitter that are similar to life events but are in fact information sharing or advertising content. A user posting *'That I'll get hoodwinked into a 2nd marriage.'* is not reporting a life event but rather expressing his/her opinion about a topic that has a similar representation to a life event. Finally, the high *class imbalance* of reported personal life events compared to non-life event content makes the detection of life event messages a difficult task. Existing work [27, 28, 29, 30] already extensively explore various manually curated features such as syntactic, semantic, sentiment as well as behavioral features to build classification models that can determine whether one or a collection of tweets are discussing a personal life event.

The objective of this paper is to move beyond the state of the art by attempting to predict *if*

and what personal life event a given user will discuss in a future time interval by solely relying on and processing the user’s historical tweets. In other words, in contrast to earlier work, this paper presents several methods for anticipating whether a given user will report on or mention a certain personal life event in her future tweets by considering previous tweets. For instance, by considering a past tweet from a Twitter user, such as *‘I do not understand why planes don’t board from the back first. That would save so much time and frustration. Is that not logical?!’*, our work would determine that the user will be reporting on a *Travel* related life event in the future; however, without having access to the future tweet that mentions *‘Left Chicago today & moved into new NYC apartment’*. From an objectives perspective, the main differentiating aspect of our work from the literature is that earlier work always benefit from the tweet that has reported the personal life event as one of the inputs to their classification model; however, our work performs predictive analytics based on past tweets assuming that the tweet that will be reporting the personal life event has not yet been posted by the user.

In this paper, we view the problem of predicting a user’s future life event as a sequence generation process where the objective is to process a sequence of past tweets and generate a future sequence that would determine the future personal life event of the user through our proposed variants of recurrent neural networks. Succinctly, the contributions of this paper are as follows:

- We propose two variants of a classification architecture based on recurrent neural networks that consider a sequence of input life events and a sequence of tweets, respectively and directly generate an output label that represents a future life event;
- In addition, we introduce a sequence generation architecture based on recurrent neural networks that considers past tweets and their life event labels *in tandem* to generate an output sequence resembling a future tweet, which would then be used to determine the future life event;
- Our proposed work is evaluated on a gold standard dataset that is systematically collected from Twitter and has real-world characteristics such as having a high class imbalance. We further discuss examples of where our methods fail to identify future life events and elaborate on the reasons for them.

The rest of this paper is organized as follows: In the next section, we classify the related work in three subsections covering pertinent work in event prediction, life event detection and deep learning techniques applied on social content. In Section 3, a high-level overview of our proposed approach is presented, which is then followed by an in-depth presentation of the details of the work in Section 4. The details of our experiments, gold standard dataset, measurement metrics, baselines and our findings are presented in Section 5. The next section is dedicated to error analysis and the paper is finally concluded in Section 7.

2. Related Work

The prior related work to this paper can be classified broadly into three categories, namely event prediction, personal life event detection, and deep learning methods for social content. We present work in these areas in the following subsections.

2.1. Event Prediction

Users' behavior on social networks is often a reflection of the events and emotions that they experience in the real world. As such researchers have been interested to use social content in order to model user behavior and make prediction on that basis. A wide range of prediction models based on Twitter content have already been developed that span several domains including traffic [31], election [18, 19, 20], healthcare [32, 16, 17], and the stock market [9, 10, 11, 12, 13]. For instance, Tsakalidis et al. [19] relied on Twitter content to predict the outcome of elections in the EU (Germany, Netherlands and Greece). Their model included eleven features based on Twitter, such as the number of times that different parties were mentioned on Twitter. In addition, they used sentiment analysis to assign a sentiment value to each tweet. They found that the use of Twitter features leads to statistically significant better results compared to just using poll information. Cameron et al. [20] found that the size of candidates network, such as the number of followers on Twitter and friends on Facebook, is not a good predictor of election results.

Mao et al. [11] studied whether the daily number of tweets can predict the S&P 500 stock indicators while the work in [12] employed sentiment features for predicting stock market trends. Nguyen et al. [9] proposed a model for predicting the stock market by capturing company and topic specific sentiments. They proposed a novel feature called the topic-sentiment which represents the sentiments towards specific company topics. Nofer et al. [13] considered the number of Twitter followers in their analysis in addition to sentiments because they hypothesized that the importance of every tweet depends on the number of users recognizing the original message.

Liu et al. [14] used bag of words, mentions, presence of a URL, emoticons, tweet length and trigger word features to predict box-office revenues for movies prior to their release. Choudhery et al. [15] presented a social data mining system that takes into account the number of tweets per day, as well as percentages of positive and negative tweets (based on sentiment analysis on those tweets) to build a polynomial regression model to predict the expected box office revenue. Based on the features of the social graph and social text content, Lassen et al. [33] developed a linear regression model that transforms iPhone-related tweets into an accurate prediction of the quarterly iPhone sales with an average error close to the established prediction models from investment banks. Radosavljevic et al. [34] developed a model called *Goalr* that predicts the World Cup outcomes and the relative strength of each country based solely on Tumblr blog posts. They used both team mentions and player mentions in their analysis. In all above methods, feature extraction from social content is a necessary task for building the prediction model.

Byrd et al. [17] collected a set of tweets based on influenza related keywords, classified them based on their sentiment characteristics and identified Twitter users who are affected by influenza in selected cities. Eichstaedt et al. [32] used language expressed on Twitter to show that community-level psychological features correlate with age-adjusted mortality from atherosclerotic heart disease (AHD). They identified that the language pattern showing negative emotions are risk factors for AHD whereas as positive emotions act as a protective factor against AHD.

2.2. Personal Life Event Detection

Most of the earlier work that have focused on the detection of life events from Twitter employ either an individual tweet [35, 27, 36, 28] or a stream of conversations [37, 38]. One of the widely used approaches for detecting life events is based on the analysis of individual tweets with the objective to process the content of the given tweet using natural language processing (NLP) techniques to explore whether a mention of a certain life event can be detected or not. However, this task can be very challenging due to noisy, ambiguous and the short length of tweets. The use

of conversations; however, can improve the performance of life event detection methods. Briefly, conversations can not only help identify events with a higher *precision*, but can also be a way to infer additional information that might not have been present in the original tweet. To this end, Choudhury et al. [29] used the collection of a users tweets within a specific time interval in order to detect whether a life event has been reported by the user in that time or not. However, in their work, they do not detect the type of the life event and only resort to detecting whether a life event has been mentioned.

Regardless of whether one tweet or a collection of tweets are selected, the primary focus of the work in the literature has been to extract strong discriminative features from tweets in order to identify life events. One of the most common features is the bag of words feature that models tweets as a vector whose entries are nonzero if the corresponding terms appear in the tweet. Di Eugenio et al. [35] and Dickinson et al. [27] have explored various features and found that n-grams (an extension of the bag of words approach) show the highest performance. However, this method suffers from the *curse of dimensionality* when the vocabulary size is large. This is particularly the case for Twitter content due to the large number of slangs, acronyms, abbreviations, and misspellings that are predominantly observed in tweets. Furthermore, the temporal ordering of words and the semantic and syntactic features of the text, e.g., named entities and part of speech tags, are overlooked in this approach. Several researchers have explored the impact of linguistic and structural features such as part of speech tagging [35, 30] and reported that this feature is not sufficient alone because few users actually observe grammatical structure on Twitter given the informal communication style.

Another group of features that have shown to be useful are based on the so called named entity vectors, which attempt to extract information by answering the 4W questions: who, what, when, and where. To this end, named entity extraction [27] and semantic role labeling [35] techniques have been used to generate *semantic* features. Furthermore, given the fact that personal life events can impact a user's feelings and emotions such as becoming glad, upset, and restless, among others, the use of sentiment features have also been considered [36] and shown to be strong indicative features for personal life event detection. From a different perspective, attention features [29] define how the content posted by a user are noticed by other users and are measured in terms of replies and retweets, and reflect how many times the user is addressed/talked about by other users within a given time interval. The motivation for using these features is based on the simple logic that important events are bound to generate more attention and activity within the immediate personal network of an individual [27, 36, 29]. Dickinson et al. [27] refer to attention features as *interaction* features and rather than just considering the number of retweets, favorites or replies, they consider who are the users performing these actions and their interaction patterns towards the poster of the tweet. They found almost no effect for the interaction features for the purpose of life event detection when applied on Twitter content and compared to other features such as n-grams, sentiments, and emojis.

Finally, while many of the introduced features have good performance on the precision metric, they do not perform too well with respect to *recall*. To address this issue, the work in [39] adopts a word vector representation of tweets that is obtained from neural word embedding techniques. These neural word embedding based features incorporate both syntactic and semantic aspects of the content and therefore show to be effective features for identifying personal life events from the perspective of both recall and precision.

2.3. Deep Learning on Social Content

Unlike work on life event detection and event prediction on Twitter, our work in this paper focuses on predicting future life events prior to them being reported by the user. We employ recurrent neural networks (RNN) for this purpose. In tasks that involve sequential inputs, such as speech and language, recurrent neural networks process each element of the input sequence in the order observed in the input and codify the observed past information in the sequence as a vector in the hidden states, which can be considered to be a form of a feature. RNNs are in essence a folded feed-forward network. RNNs are suitable for predicting the next characters in text [40], the next word in a sequence [41] or more complex tasks such as translation [42, 43]. After reading a sentence word by word, an encoder network is used to learn the thought expressed in that sentence in its final hidden state, referred to as the *thought vector*. The final hidden state is then used as the initial hidden state of a jointly trained decoder that would produce a probability distribution for the first word in the output sequence. Once a particular first word is chosen from this distribution and provided as input to the decoder, it will then output a probability distribution for the second word of the output sequence and so on until a termination symbol is observed.

In theory, RNNs can make use of information in arbitrarily long sequences, but in practice they are limited to looking back only a few steps. Long Short-Term Memory (LSTM) was introduced to overcome this problem by augmenting the network with a memory cell which remembers inputs for a long time [44]. LSTM networks or related forms of gated units are also currently used for encoder and decoder networks that perform well in machine translation [42, 43], sentiment analysis [45, 46], text classification [47, 48, 49], generation of sequence models [50].

Beyond RNNs, other deep learning methods have also shown promising results on various tasks on Twitter content. Tang et al. [51] extended the C&W model and developed three neural networks to learn sentiment-specific word embeddings from Twitter content. The tweets containing positive and negative emotions were used as training data. Severyn et al. [52] used a three-step process to train their deep learning model for sentiment analysis of tweets at both message and phrase levels. Word embeddings were initialized using a neural language model which is trained on a large unsupervised collection of tweets, and a convolutional neural network is used to further refine the embeddings on a large distant supervised corpus. Limsopatham et al. [53] presented a model based on bidirectional LSTMs to automatically learn orthographic features without requiring feature engineering for named entity recognition in Twitter messages. Their model consists of three main components: (1) orthographic sentence generator, (2) word representations as input vectors, and (3) a bidirectional LSTM. Li et al. [54] proposed an attention-based LSTM model, which incorporates topic modeling into the LSTM architecture through an attention mechanism in order to recommend hashtags for tweets.

A deep learning framework for rumor debunking on Twitter and Weibo is introduced in [55]. The framework learns RNN (tanh, GRU, LSTM) models by utilizing the variation of aggregated information across different time intervals related to each event. Lin et al. [56] believed that individuals psychological stress is recognizable via social media. They defined two sets of attributes: 1) low-level content attributes from a single tweet, including text, images and social interactions, and 2) user-scope statistical attributes such as behavioral attributes, social engagement and linguistic style attributes from users weekly tweet postings. To combine content attributes with statistical attributes, they designed a convolutional neural network (CNN) with cross autoencoders to generate user-scope content attributes from low-level content attributes. Finally, they proposed a deep neural network (DNN) model to incorporate the two types of user-scope attributes to detect users psychological stress.

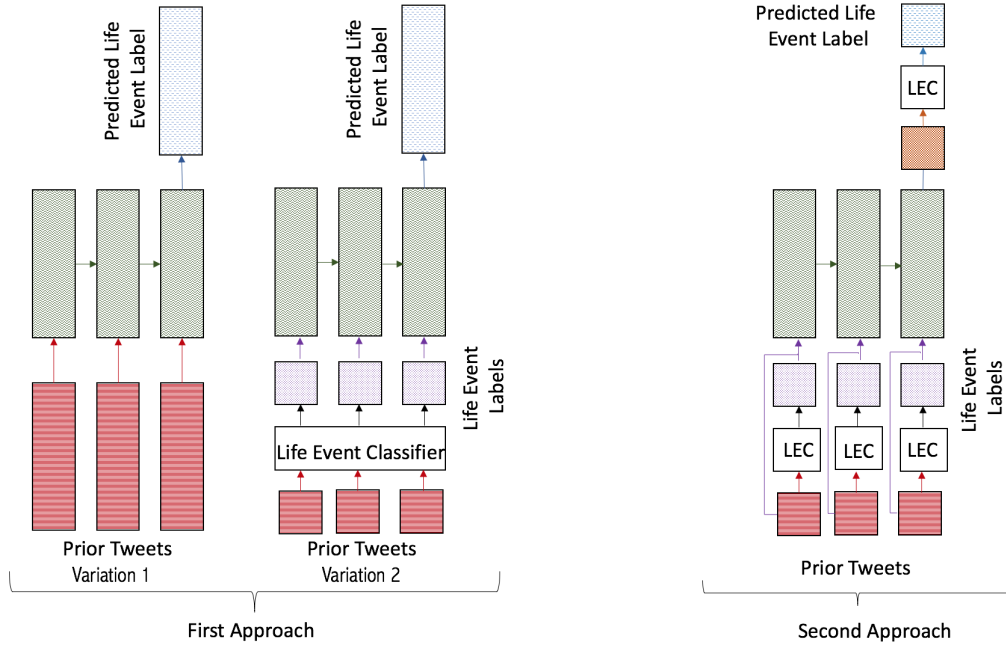


Figure 1: Overview of the proposed Seq2Seq models for life event prediction (LEC stands for Life Event Classifier).

3. Approach Overview

In our work, we propose three variations of recurrent encoder-decoders for predicting future personal life events based on user’s past tweets. We employ Sequence-to-Sequence (Seq2Seq) models to perform (i) direct life event label prediction; and, (ii) indirect life event prediction through the generation of future tweets. In the first approach, we predict future life event labels directly based on a *many-to-one* Seq2Seq architecture. The output label of the sequence to sequence model is the future life event that the user might report on her Twitter timeline. We present two variations of the many-to-one Seq2Seq architecture. In the first variation, a set of tweets are directly fed into the encoding stage and predicted life event label is produced in the decoding stage. The second variation; however, works with a collection of past life event labels from that user to predict the future life event. Therefore, the input to the decoder stage is a sequence of past personal life events reported by the user and the output is a predicted personal life event in the future. Now, given the second variation requires a set of input life events, we generate such labels by using state of the art methods that classify a given tweet into certain life event classes. As such, a set of past tweets of the user are passed through such classifier that would produce one life event label per tweet. These life event labels are then passed onto the many-to-one Seq2Seq model as input to the encoder stage.

In the second approach, we employ a *many-to-many* Seq2Seq model that operates over the combination of the inputs of the two earlier variations of the first approach, namely the tweets and the past personal life events, in the encoding stage. The model then generates a word sequence in the decoding stage that estimates the content of a future tweet. The idea of the second approach is that it might be possible to predict a user’s future tweet based on the personal life events they

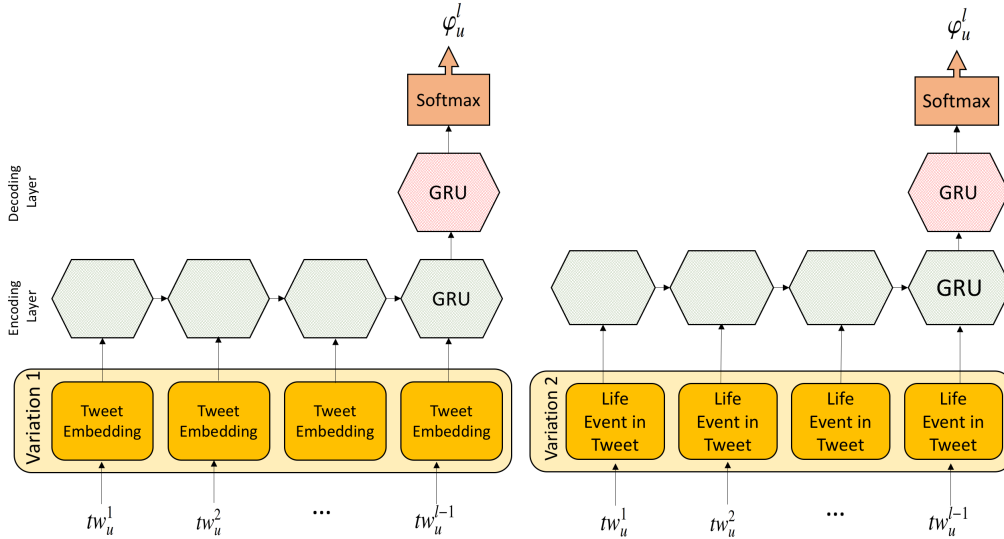


Figure 2: The architecture for the many-to-one Seq2Seq model for both variations: (a) based on past posted tweets, and (b) based on past observed life events.

have reported in the past as well as the tweets they have posted. We generate the user’s future tweet through a many-to-many recurrent model, the output of which is then classified into one of the personal life events using state of the art life event classification methods. Figure 1 shows the graphical representation of the two proposed approaches for life event prediction.

4. Proposed Approach

We provide the details of our two proposed approaches for life event prediction based on a sequence-to-sequence architecture. The Seq2Seq architecture is trained for classification in the first approach and for sequence generation in the second.

4.1. Many-to-one Classification Architecture

In the first approach, a Seq2Seq architecture is used to directly generate a personal life event label for a user given her past shared content. The most intuitive approach for predicting future behavior based on social content is to take the tweets that the user has posted into account to predict what the user will post in the future. For instance, a person who is getting married is likely to post tweets talking about preparing for the wedding day or a person changing jobs might post about how she is looking for a new job in the days or months leading to her finding the job. Therefore, our main hypothesis is that a user’s past tweets can be seen as potential indicators for a future life event.

In the *first variation* of the many-to-one Seq2Seq architecture, a sequence of tweets from the timeline of a user is taken into consideration to predict the future personal life event that the user will mention in her timeline. To this end, we need to work with an input sequence consisting of l tweets of user u , denoted as $tw_u = [tw_u^1, tw_u^2, \dots, tw_u^l]$. In order for the sequence tw_u to be used as input, a representation form for each tweet needs to be adopted. One representation

model for a tweet is to represent it as a bag of words, where each tweet is represented as a vector with v dimensions where v is the number of unique words in the tweet corpus. This method suffers from the curse of dimensionality, which is aggravated within the context of Twitter given the wide range of slangs, acronyms and abbreviations. As an alternative, Vosoughi et al. [57] proposed a model for generating general-purpose vector representation of tweets. The model learns tweet embeddings using a character-level CNN-LSTM encoder-decoder. Dhingra et al. [58] also adopted a similar strategy and proposed a Bi-directional Gated Recurrent Unit (GRU) neural network for finding vector space representations of tweets by learning complex, non-local dependencies in character sequences. We adopt a similar strategy to [58] and compute a vector representation for each tweet to be used as input. Based on the vector representations of the past tweets and an observed personal life event reported by the user at time interval l , denoted by φ_u^l , the objective of the model is to maximize the selection probability of φ_u^l given tw_u :

$$\arg \max_u P(\varphi_u^l | tw_u^1, tw_u^2, \dots, tw_u^{l-1}) \quad (1)$$

4.1.1. Encoding

In the proposed architecture, the encoder converts the input tweet vectors into a set of high dimensional hidden representations $h = [h_1, \dots, h_{l-1}]$ where h_i is the hidden state of the Seq2Seq architecture at time i . In other words, at time i , the encoder reads tw_u^i and updates its hidden state h_i as follows:

$$h_i = RNN(h_{i-1}, tw_u^i) \quad (2)$$

where $h_0 = 0$ and the RNN applied in the encoder is often either a Long Short-Term Memory (LSTM) [44] or Gated Recurrent Units (GRU) [59], which are both capable of learning long-term dependencies and do not suffer from the vanishing gradient problem. Similar to our proposed interval encoding mechanism shown in Equation 2, there have been other work that employ comparable approaches. For instance, Tang et al. [45] used an LSTM to produce sentence vectors, which are then applied within the context of a bi-directional gated recurrent neural network to compose the sentence vectors into document vectors. Likewise, Yang et al. [60] proposed a hierarchical attention network for document classification, which employs GRU for encoding both words and sentences. Similarly in our encoders, we use GRUs primarily because they have similar performance to LSTMs but are more computationally efficient.

4.1.2. Decoding

In our architecture, the decoder is a single-layer GRU that is responsible for generating the output predicted life event based on the encoded input in the form of a vector in the embedding space. Given the output sequence only consists of one token in our first approach, i.e., the predicted personal life event, the decoder would predict φ_u^l based on h_{l-1} :

$$\begin{aligned} P(\varphi_u^l = le | tw_u^1, tw_u^2, \dots, tw_u^{l-1}) &= P(\varphi_u^l = le | h_{l-1}) \\ &= \text{softmax}(d^1, \varphi_u^l = le) \end{aligned} \quad (3)$$

where le is some life event such as changing jobs or getting married and $d^1 \in R^d$ is the hidden state of the GRU decoder:

$$d^1 = GRU(d^0, \varphi_u^1) \quad (4)$$

It should be noted that d^0 is initialized based on the hidden state of the interval encoding h_{l-1} based on Equation 2.

4.1.3. Model Training

In order to train the model, we maximize Equation 1 whose parameters are learned by maximizing the log-likelihood of the condition. Back propagation through time (BPTT) [61] is used for computing the gradient of the objective function. The overall scheme for the first approach is depicted in Figure 2, where the encoder computes the hidden representation for all φ_u^l s, where $l > i \geq 1$ through GRU. This hidden representation is then used to compute the probability of φ_u^l .

The *second variation* of the same many-to-one Seq2Seq architecture, as shown in Figure 2.b, is based on a more dense input representation where the input to the encoder is the sequence of personal life events that have been reported by the user in the past. As discussed in the related works, there are already work in the literature that can classify a tweet into one of several classes of life events depending on the content of that tweet. For instance, such methods are able to label a tweet such as ‘Just booked my flight to Cali, a vacation sounds nice’ as one that is related to *travel* and another tweet such as ‘So excited to start my new job i’ve practically bought an entirely new wardrobe’ as a tweet that reports on the *job change* life event. Based on such personal life event classifiers, we are able to produce a sequence of life events, including the none event, which refers to cases when no life events are mentioned in the tweet, for each user. In our work, we adopt the method proposed in [39] for detecting personal life events. This method proposes to collect a set of words to *discriminatively* express each life event. The word vector representation of these words are then computed using word embedding methods, i.e., the Skipgram model [62]. Based on the vector representation of the life event words, it is possible to calculate the distance between each life event and a given tweet, also represented as the vector of its constituent words using the *word mover’s distance* measure [63]. Now, the advantage of this model is that it not only identifies the life event mentioned in a tweet, but also provides a vector representation for each of the life events based on the centroid of its constituting discriminative words. As such, the corresponding vector representation to each life event is used as input to the encoder in the Seq2Seq architecture when that life event has been detected in the previous tweets. In summary, we label a tweet with a life event based on the closeness of its vector representation to the vector representation of the life events and chose the vector representation of the selected life event as the representation of the tweet to be input into the encoder of the Seq2Seq model. The rest of the details of the architecture are similar to the first variation.

4.2. Many-to-many Generation Architecture

The main idea of the second approach is to predict a future personal life event by trying to predict the content of a user’s future tweet. Similar to [64], we hypothesize that there are two information pieces that can assist with the prediction of a future tweet, namely the past tweets and the user’s past tweet topics. In the context of our work and given we are interested in personal life events, we customize these two information pieces as (i) user’s past tweets, and (ii) user’s past reported life events; therefore, we exploit users’ tweets and the life events that they reported as prior knowledge for estimating the next tweet. In order to incorporate these information pieces into our model, we need to first create a uniform representation as the two information sources are not compatible because past tweets are a set of words while the reported life events are merely a label. We address this issue by converting the reported life event labels into a set of representative words for the life event. The representative words used here are those derived based on [39] and used in the second variation of the first approach. Based on these words, depending on the life event that a tweet is reporting, it is augmented with a set of words representing that life event. It

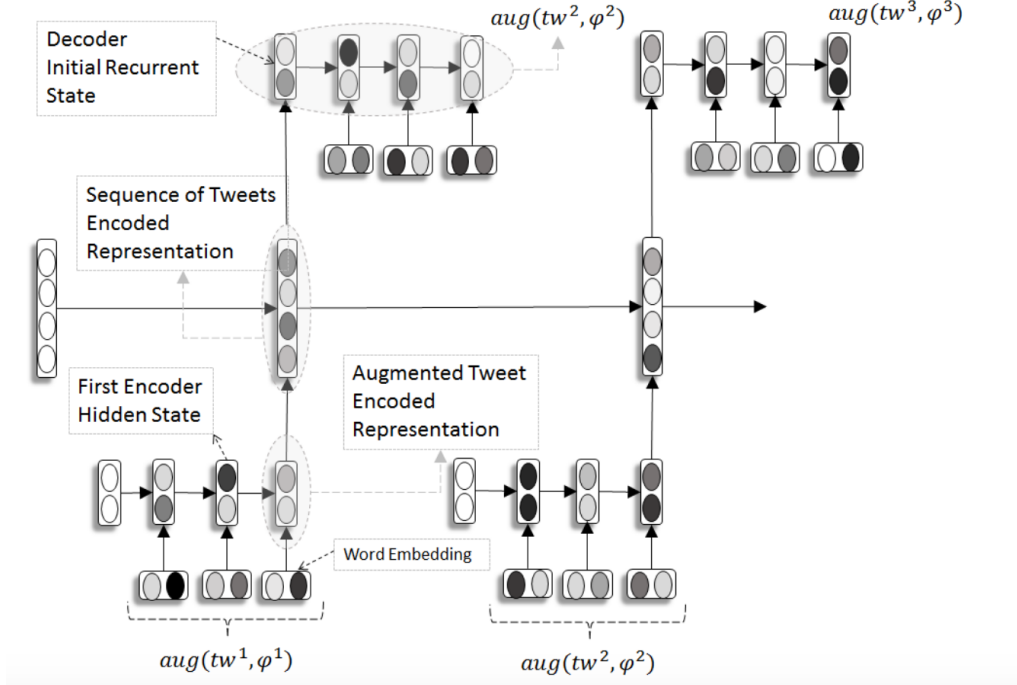


Figure 3: The architecture for the many-to-many Seq2Seq model for augmented tweet prediction (the different color circles point to the fact that vector elements can have different values).

should be noted that the reported life event for each tweet is automatically determined and those tweets that are determined not to be reporting any personal life events are not augmented with any additional words. Now, based on the vector representation of the augmented tweets as well as an augmented future tweet, i.e. the sequence of words in a future tweet plus the representative words of the life event that tweet is representing, the objective is to maximize the generation probability of the augmented future tweet as follows:

$$\arg \max_u P\left(\text{aug}(tw_u^l, \varphi_u^l) | \text{aug}(tw_u^1, \varphi_u^1), \text{aug}(tw_u^2, \varphi_u^2), \dots, \text{aug}(tw_u^{l-1}, \varphi_u^{l-1})\right) \quad (5)$$

where tw_u^l is the user's l^{th} tweet, and φ_u^l is the reported life event in that tweet. Also, $\text{aug}()$ is a function that augments the words within the input tweet with the relevant representative life event words. This function returns a vector whose elements are the tweet words and life event words sequentially. In other words, the words in the tweet are concatenated with the representative words of the life event, forming one sequence of words consisting of the tweet words first followed by life event words. The overall architecture of the Seq2Seq model for the second approach is shown in Figure 3. This many-to-many Seq2Seq architecture predicts the next tweet given the tweets already posted by the user. The history of past posted tweets is considered as a sequence at two levels: a sequence of words for each tweet and a sequence of tweets. The many-to-many Seq2Seq architecture models this hierarchy of sequences using two RNNs: one at the word level and one at the tweet level. In the sequence of tweets, the first encoder RNN maps each tweet to a tweet vector. The tweet vector is the hidden state obtained

after the last token of the tweet has been processed. The higher-level encoder RNN keeps track of past tweets by processing each tweet vector iteratively. After processing a tweet, the hidden state of the second encoder RNN represents a summary of the sequences up to and including the tweet, which is used to predict the next tweet. The next tweet prediction is performed by means of a decoder RNN, which takes the hidden state of the second encoder RNN and produces a probability distribution over the tokens in the next tweet. The decoder RNN is similar to an RNN language model but with the important difference that the prediction is conditioned on the hidden state of the second encoder RNN. It can be seen as a tweet generation module.

4.2.1. Encoding

The encoding process of our architecture consists of two encoding processes: (1) encoding of the augmented tweets, and (2) the encoding of the sequence of augmented tweets. In the first encoding step, we employ a recurrent neural network architecture for encoding the augmented tweets, which produces a fixed-length vector, after reading the last word of an augmented tweet. The RNN processes the tweet words and the life event words sequentially from the augmented tweet and updates its hidden state as follows:

$$h_i^1 = RNN_E(h_{i-1}^1, aug(tw_u^k, \varphi_u^k)[i]) \quad (6)$$

The hidden state (h_i^1) is expressed with a superscript to denote that it is related to the first encoding step and $aug(tw_u^k, \varphi_u^k)[i]$ is the i^{th} word of the augmented tweet tw_u^k . It is important to mention that our work has resemblance to the work by Semeniuta et al [65], which is based on the the variational autoencoder framework for generating tweets. Its scoring components are a convolutional encoder and a deconvolutional decoder combined with an LSTM recurrent layer. Similarly, Oak et al. [66] also employed LSTMs at both character and word levels to generate realistic-looking tweets with the same statistical properties as the actual data. Furthermore, Shang et al [67] have introduced a Neural Responding Machine (NRM), based on the general encoder-decoder framework, in order to generate response for Weibo (a popular Twitter-like microblogging service in China) conversations. They used GRU for both the encoder and decoder. Also, Serban et al [68] proposed the multi-resolution recurrent neural network for generatively modeling sequential data at multiple levels of abstraction. They used GRUs and applied it to dialog response generation for Twitter conversations. One of the shortcomings of GRUs is that they are only able to make use of previous context and have restrictions as the future input information cannot be reached from the current state. Bidirectional RNNs (BRNNs) [69] address this issue by processing data in both directions with two separate hidden layers, which are stacked on top of each other. In our work, we use Bidirectional GRUs, to represent RNN_E , which compute a *forward* hidden layer by iterating through the input from beginning to end and a *backward* hidden layer by iterating through the input sequence from the end to the beginning. These two hidden layers are concatenated into a single vector, denoted by h_i^1 . Therefore, h_i^1 represents the encoded version of the augmented tweet up to word i .

The second encoding step is focused on reading the encoded version of the augmented tweets sequentially and encoding this sequence into a fixed-length vector. The RNN processes the encoded representation of each augmented tweet and updates the encoding as follows:

$$h_i^2 = RNN_S(h_{i-1}^2, h_i^1) \quad (7)$$

Likewise, RNN_S is modeled as a GRU.

4.2.2. Decoding

In this architecture, the decoder is another GRU, which is trained to generate the next augmented tweet given the previous augmented tweets. In other words, the GRU decoder processes the hidden state of the encoded augmented tweets as input and produces a probability distribution over the tokens in the next time interval as output and is computed as follows:

$$P(\text{aug}(tw_u^l, \varphi_u^l) | \underbrace{\text{aug}(tw_u^1, \varphi_u^1), \dots, \text{aug}(tw_u^{l-1}, \varphi_u^{l-1})}_{\Psi}) = \prod_{i=1}^N P(w_i | w_1, \dots, w_{i-1}, \Psi) \quad (8)$$

where N is the length of tw_u^l and w_i is its i^{th} word. The probability of word w_i given previous words in the tweet and the previous augmented tweets can be estimated as follows:

$$\prod_{i=1}^N P(w_i = v | w_1, \dots, w_{i-1}, \Psi) = \text{softmax}(d^{l-1}, w_i = O_v) \quad (9)$$

where O_v is the encoded embedding of word v and $d^{l-1} \in R^d$ is the hidden state of the RNN decoder:

$$d^j = \text{RNN}_D(d^{j-1}, w_j) \quad (10)$$

Here, RNN_D is a GRU. It should be noted that we use beam search [70] for in the decoding process to approximate the most probable augmented tweet given the past augmented tweets.

4.2.3. Model Training

Let us assume that a time interval I consists of l augmented tweets and that a generative model parameterizes a probability distribution P with parameters θ over the set of possible intervals of arbitrary length. The probability of an interval I can be defined as follows:

$$\begin{aligned} P_{\theta}(\text{aug}(tw_u^1, \varphi_u^1), \dots, \text{aug}(tw_u^l, \varphi_u^l)) &= \prod_{t=1}^l P_{\theta}(\text{aug}(tw_u^t, \varphi_u^t) | \underbrace{\text{aug}(tw_u^1, \varphi_u^1), \dots, \text{aug}(tw_u^{t-1}, \varphi_u^{t-1})}_{\Upsilon}) \\ &= \prod_{t=1}^l \prod_{i=1}^{N_t} P(w_i^t | w_1^t, \dots, w_{i-1}^t, \Upsilon) \end{aligned} \quad (11)$$

where N_t is the length of the t^{th} augmented tweet and w_i^t is its i^{th} word.

The parameters θ can be learned by maximizing the log-likelihood of an interval I . We employ back propagation through time (BPTT) [61] for computing the gradient of the objective function as defined in Equation 11. Based this equation, it is possible to generate a future augmented tweet based on its relevance to the past augmented tweets. In our work, the recurrent neural networks RNN_E , RNN_S and RNN_D are GRUs whose parameters are learnt in this way. In light of the fact that the sequence of words produced by the decoder in the output is intended to represent a user's future tweet, we employ it to predict the next life event of the user. We use the same method presented in [39], denoted as Life Event Classifier (LEC) in Figure 1, to classify the generated sequence into one of the life events. The determined life event class for the generated tweet would represent the predicted life event in this model.

5. Experiments

We evaluate the different variations of our sequence to sequence model, and compare it against the state of the art baseline methods. The objective of our experiments is to measure the effectiveness of our proposed work for predicting future personal life events solely based on the previous tweets of the user.

Table 1: Specification of the dataset with the distribution of life events.

Life Event	Number of tweets	Percentage
Broken device	5,768	8.69%
Device upgrade	2,822	4.25%
Moving	4,001	6.03%
New job	4,001	6.03%
Travel	3,480	5.24%
Wedding	4,015	6.05%
Negative samples (no life event)	42,275	63.7%

5.1. Experimental Setup

The primary dataset used in our experiments is composed of 66,362 labeled tweets collected and annotated manually by twenty experts¹. The labeled tweets cover six personal life events as shown in Table 1. The dataset also consists of tweets that are not about any life events, which are used as negative samples. As shown in Table 1, we have preserved the significant *class imbalance* that is prevalent in life events on Twitter in this dataset. Given most tweets do not necessarily talk about a life event on Twitter, in our dataset, over 63% of the tweets are related to no life events.

For each of the labeled tweets in the dataset, we also extracted tweets by the same user from one week prior to the given tweet through the Twitter RESTful API in order to have access to historical tweets. We pre-processed the dataset as follows. First, we removed URLs and reposting marks such as ‘RT’ and ‘//’. Then, we filtered stop words. Finally, we used the Stanford CoreNLP tool set for lemmatizing all the tweets.

In all models, back propagation through time [61] and Adam [71] are used for computing gradient of the objective function and optimization, respectively. All our models were trained on NVIDIA Tesla K80 GPU provided through SHARCNET². Our implementation is done using the open-source Python library Theano [72]. The hidden state dimension for the encoders and decoders is set to 512 in both variations of the first approach. Also the size of the embeddings is set to 100. The training stops if the likelihood of the validation set does not improve for 10 consecutive iterations (early stop). The learning rate was set to 0.001. The dimension of hidden state of the RNN in the second approach is set to 1,000. We initialize the word vectors using a pre-trained word embedding with size 400 obtained by executing the Skip-gram methods [62] on a 10 million English tweet corpus collected through the Twitter RESTful API. This model is trained with a learning rate of 0.0001 and with mini-batch containing 25 training samples. The size of beam is 1. In the test phase, we consider variant lengths of 4, 6 and 8 because the average number of words in a tweet is close to 6 based on an experiment that we did on our tweet corpus.

¹Available upon request.

²<https://www.sharcnet.ca/>

5.2. Baselines

While there are no prior work on predicting future personal life events, as mentioned in the related work, there are work that determine whether a life event is mentioned in a tweet or a collection of tweets. One of the most related works in the literature is proposed by Li et al. [30], which is a pipelined system for detecting major life events. This work first identifies the life event category the tweet is referring to and subsequently identifies whether the tweets is a self report of the life event or it is referring to the life event for a third party. As suggested by the authors, in order to replicate the baseline, we extract features including bag of words, named entity mentions, dictionary and window for tweet modeling. Bag of words and NER features are the sequence of words in the tweet and named entities, respectively. A dictionary of the top-40 words for each life event category is built (automatically inferred by a topic model). The dictionary feature value is the terms probability generated by the corresponding event. If a dictionary term exists in the tweet, left and right context words within a window of 3 words and their part-of-speech tags are extracted as the window feature. Named entity tags are assigned using Ritter et al.’s Twitter NER system [73] and Part-of-Speech tags are assigned based on Twitter POS package [74]. As suggested by the authors [30], we train a 7-class maximum entropy classifier with all these features as the experiments by Li et al. reported this classifier to have the best performance. The other work that we use as baseline is the life event detection method proposed by Choudhury et al [36]. In their work, the authors propose to extract a set of features including hashtags, emoticons, named entities, and sentiments to build classifiers that can find life event mentions. We use these features as well to build a second baseline using a maximum entropy model.

Given the objective of our evaluation is to investigate the performance of the proposed method for life event prediction, we use the standard information retrieval metrics including *precision*, how many of the life events detected by our method were correct, *recall*, how many life events in the ground truth set were retrieved by our method and F-Score, which is the harmonic mean of precision and recall.

5.3. Results

In order to perform the evaluations, we split the life event dataset introduced in Table 1 into three parts, namely training, validation and testing, each of which contained 80%, 10%, and 10% of the original dataset, respectively. The results of the experiments are shown in Table 2 for both variations of the first approach (classification architecture) as well as the second approach (generation architecture). The results are compared against the strongest baselines available in the literature reported by Li et al. [30] and Choudhury et al [36]. Given our objective is to predict the personal life event of a future time interval, we executed Li et al.’s method using both the last tweet prior to the future tweet as well as the past one week of tweets leading to the future tweet of interest. However, given the nature of the features by Choudhury et al., we report the findings based on building classifiers for each individual feature as well as a classifier where all features are included using the Maximum Entropy classifier. The results are reported in Table 2. As seen in the table, the precision of both variations of the baseline by Li et al is quite poor and in fact close to random when keeping in mind that we are performing a 7-class prediction task and therefore a random classifier would perform at a 14% precision rate compared to the current precision of the baseline which is 13.3%. It should be noted that the work by Li et al is considered to be a strong baseline as it performs at precision rates of $\sim 75\%$ when the life event is mentioned in the tweet that is being classified; however, as observed in our results, it is not

Table 2: The performance of the proposed classification and generation methods.

Method		Precision	Recall	F-Score
Classification Architecture	Variation 1	0.234	0.211	0.222
	Variation 2	0.272	0.328	0.297
Sequence Length				
Generation Architecture	100	0.289	0.363	0.322
	8	0.357	0.358	0.358
	6	0.377	0.359	0.368
	4	0.394	0.369	0.381
Baseline by Li et al [30]	Latest Tweet	0.13	0.361	0.191
	All Tweets from Past Week	0.133	0.364	0.195
Baseline by Choudhury et al [36] based on Maximum Entropy Classification	Hashtags	0.129	0.359	0.19
	Emoticons	0.13	0.36	0.191
	Named Entities	0.13	0.36	0.191
	Sentiments	0.133	0.365	0.195
	All Features	0.19	0.358	0.248

able to show competitive performance for performing future life event prediction. Among other conclusions, this significant difference in classification precision shows that the task of predicting the personal life event of a future task is quite difficult and cannot be performed using existing life event classification techniques. Furthermore, we observe that the work by Choudhury et al [36] shows quite similar performance to the work by Li et al. [30] when individual features are used to train the classifier. The performance of these classifiers are weaker than an average baseline, which would have a precision of 14%. However, the performance of the classifier based on all of the features is better than the other baselines in terms of precision at 0.19. However, this performance is still weaker than all variations of our generation architecture and weaker than the second variation of the classification architecture.

Now, in terms of the comparative performance of the two variations of the first approach, which is primarily based on a many-to-one Seq2Seq classification architecture, one can see that the second variation shows stronger performance. While the first variation works directly with the tweets, the second variation only works based on a set of life event labels that have been observed in the past. In other words, the input sequence to the second variation is a sequence of personal life events that the user has tweeted about (including a *none* label). Our observation has been that reducing symbol space of the Seq2Seq model has improved the performance of the model in terms of both precision and recall. The main reason for this could be that the first variation needs to model any sequence, which could consist of many different tokens spanning possibly beyond the number of English words, as tweets could include abbreviations and slang words, and therefore, creating an accurate representation for such a large symbol space would be challenging. In contrast the input sequence space of the second variation consists of only seven symbols each representing one of the life events or the *none* life event label.

While the second variation shows better performance compared to the first variation, it is important to point out that it is limited to the performance of the underlying tweet life event classification algorithm. In other words, given the input to this model is a sequence of life

event labels, and tweets themselves do not have the life event labels with them and hence the life event label for each tweet is automatically determined using a classification technique, the performance of the second variation is dependent on the performance of the employed life event classifier. Therefore, it is possible to get better performance using the second variation if more accurate life event classifications per tweet were available.

In the second approach, which unlike the first approach, is based on a sequence generation architecture, an augmented tweet is predicted, which is assumed to be the rough estimated representation of a future personal life event-related tweet by the user. As seen in Table 2, we have tested the second approach on varying lengths of the output sequence. The output sequences that were generated had a length of 4, 6, 8, and 100 tokens. We were interested to see whether both smaller and larger sequence lengths had any significant impact on the performance of the prediction model. The first important observation is that the predictions made based on the second architecture, regardless of the length of the output sequence, have both higher precision and recall rates compared to the classification architecture. This can be attributed to the fact that the generation architecture employs an augmented tweet as input, which consists of the tweet itself and a specific set of words for the related life event. This way, the set of words related to the life event drives the encoding and decoding process and hence leads to better prediction performance. It should be noted that more so than the second variation of the classification approach, the performance of the generation architecture is even more dependent on the performance of the employed life event classifier as it is used for both generating the augmented tweets as well as classifying the generated output sequence.

From the perspective of the generated sequence length, it can be observed that while it does not significantly impact the performance of the recall metric, it does have a substantial impact on precision. Longer generated sequences have a lower precision and hence less effective for predicting a future personal life event. This can be potentially explained at least by two reasons. First, a study in 2012³ showed that while tweets can be as long as 140 characters, in practice they are 40 characters long on average. This is approximately 4-6 words in length per tweet. This is also in line with our own observation in our Twitter gold standard dataset. Therefore, generating output sequences of the same length as the average tweet would prevent the generation of possibly irrelevant tokens in the output. Second, given the output sequence generated at the decoder relies on *beam* search, the longer the sequence is, the more likely it will be that the generated sequence would have lower cohesion [70]. We observed that the output sequence of length four provides the best performance in terms of both precision and recall.

5.4. Bootstrap Aggregating

Several researchers have already shown that ensemble meta-algorithms can improve the performance of recurrent neural networks. More specifically, the bootstrap aggregating (*bagging*) method can lead to improved classification performance without the need to change the model being trained and by only learning ensembles over different subsets of the training data. Bagging operates by generating n training subsets by uniformly sampling data points with replacement from the original training set, which would then be used for training n similar models trained on different data subsets. Breiman [75] has shown that bagging can improve the performance of *unstable* methods such as neural networks that have been used in this paper. On this basis, we applied the bagging approach on both the baselines and the various proposed architectures in this paper.

³<https://goo.gl/ohBcD8>

Table 3: The performance of the proposed classification and generation methods after bagging.

Method		Precision	Recall	F-Score
Classification Architecture	Variation 1	0.262	0.330	0.292
	Variation 2	0.278	0.338	0.305
Sequence Length				
Generation Architecture	4	0.519	0.373	0.434
Baseline by Li et al [30]	Latest Tweet	0.247	0.264	0.255
	All Tweets from Past Week	0.278	0.305	0.291
Baseline by Choudhury et al [36] based on Maximum Entropy Classification	Hashtags	0.129	0.359	0.19
	Emoticons	0.129	0.359	0.191
	Named Entities	0.129	0.359	0.191
	Sentiments	0.223	0.256	0.195
	All Features	0.19	0.358	0.248

To apply the bagging approach, we trained several models with different number of epoches ranging from 200 to 400 with an interval of 50 epoches in between while also systematically applying the training dataset splits. Given the results in Table 2 showed that the generation architecture shows its best performance with a sequence length of 4, we used this sequence length in the models trained in the bagging approach, as well. The multiple predicted labels were gathered and used to get an aggregated predicted label based on a *plurality vote*. It should be noted that we experimented with both weighted and unweighted bagging methods and the difference was not noticeable to the third decimal point. The results of the bagging approach are reported in Table 3.

There are two main observations from the obtained results. The first is that all models have been improved as a result of the bagging process on both precision and f-score measures. This is most noticeable on the generation architecture where the precision improved as a result of bagging from 0.394 to 0.519, which is an improvement of 31.7%. The second observation is that the bagging process does not have a significant impact on recall as the improvements are only marginal. This observation can be broadened by comparing the recall rates across the different methods as well as before and after the bagging process. In all cases, while the recall rate of the generation model both after and before bagging outperforms the other methods but still the differences are not substantial. This shows that the models have very similar *retrieval* capacity with recall values in the range of 0.33 to 0.373, while the generation architecture has a strong *classification* ability shown by the precision of 0.519.

It should be noted that the results of the life event prediction should be understood within the context of the work on personal life prediction and the extremely difficult, highly noisy and class imbalanced dataset that it offers. We would like to point out that the strongest baseline proposed by Li et al [30] offers an f-score of 0.54 when working directly on the tweet of interest and performing a tweet classification task. In contrast, we attempt to predict the future life event without having access or working with the content of the tweet of interest and solely rely on past tweets to predict the future life event and achieve an f-score of 0.434, which is meaningful and strong when compared to the results obtained in [30].

In the next section, we will present tweets from our dataset for both correctly labeled and

incorrectly labeled instances and discuss the predominant reasons for the incorrect classifications.

6. Error Analysis and Discussion

Beyond the information retrieval metrics that report the performance of the proposed methods, it is also important to understand the cases when the models succeed in identifying the correct personal life event as well as when and why the models do not correctly identify the life events. To this end, we have reviewed the result of the better performing generation architecture. We found that the model performs very well for cases when a logical progression between the tweets of the same user can be observed. In other words, for those twitter users who consistently and continuously discuss their personal life events on Twitter, the model is able to show good performance. This is primarily due to the fact that the Seq2Seq model builds an internal representation by encoding a sequence of tweets and their life event labels (*augmented tweets*), which is then used to predict a future tweet. The more cohesive the set of tweets in this sequence are, the more consistent the encoded internal representation would be and hence the generated sequence through decoding is more likely to represent the future tweet and its life event. However, if the past sequence of tweets has arbitrary order, then the internal representation would be convoluted based on the inputs from the unrelated tweets in the sequence.

An example of a case where the model is successful in predicting the personal life event is when a user reports ‘*Smashed my iPhone, okay fairs*’ and subsequently posts ‘*first one I’ve ever smashed!!! Dan do apple fix your phone or give you a replacement??*’. For such a consistent sequence of tweets, the model generates the sequence: ‘note drop phone break’, which is easily labeled as the ‘Broken Device’ label. The first two rows of Table 4 provides two more examples of how coherent flows of tweets lead to correct classification.

Now in terms of the incorrectly predicted personal life event labels, we identified three primary reasons for the incorrect predictions. The first and primary cause (Error Type I) for the errors was related to the cases when the user was reporting a life event related to another person. In such cases, it was possible that the reported life event would be determined to be a personal life event and hence labeled incorrectly. For instance, for a user who tweeted ‘*OMG in shock with the news I just received. You never know when your time is up*’ and subsequently ‘*Mom going house hunting in Florida*’, the method incorrectly predicted the ‘Move’ life event having mistakenly labeled the second tweet as a *move* related tweet and hence using the incorrect life event in the augmented tweet and hence negatively influencing the sequence generation process. It should be noted that the performance of the Life Event Classifier [39] used in our experiments leads to such errors. While this model does implicitly consider self-reports in the process, it does not explicitly do so from a linguistics perspective. There are other techniques in the literature [30] that have a pipeline that determines whether the user who is reporting the life event is in fact directly involved in the life event or not, e.g., through checking whether the subject of the tweet is a *first person singular*, but given those models do not have 100% precision, still such instances can happen.

The second reason for the errors (Error Type II) was due to cases when a user was providing information on a topic that had similar components to a life event. For instance, for a user who tweeted ‘*If you was a Passport where would you hide?*’ and then later tweets ‘*I lost my passport, got a new one. Lost new one, found old one. Can I use it??*’, the user is in fact communicating some information about his passport to his followers. However, given the tweets have similar components to a travel tweet, the model would predict a ‘Travel’ life event for this user. The third reason for the errors (Error Type III) in the prediction was related to a sudden change of

Table 4: Sample Tweets and the Predicted Labels.

	Past Tweets	Actual Label	Predicted Label
Correct Predictions	Smashed my iPhone, okay fairs you'll be back! Hopefully!,Xx	Broken Phone Move	Broken Phone Move
Error Type I	Got to bring this beautiful girl to my sister's wedding last night [url] My Friend From Bible School Is Pregnant	No Life Event No Life Event	Wedding Wedding
Error Type II	OMG my grandma lol [url] g @user You can find rent, floor plans and all other info on the property here: [url]	No Life Event No Life Event	Wedding Move
Error Type III	@user @user No appointments at passport office so I have to pay £200 for an external companies help Exploring the trails and the falls. #WebstersFalls #DiscoverOntario #ExploreCanada #hamont [url]	Broken Phone Broken Phone New Job	Travel Travel Travel

subject in the users tweets. In these cases, the user is discussing some topic and suddenly changes the topic and discusses a new topic. For such cases, the proposed sequence generation approach cannot build a coherent internal encoded representation and hence has difficulty in predicting the correct life event. For instance, a user has tweeted ‘@user love ya mucho’ and then immediately afterwards posts ‘*unpacking from one trip and packing for another*’. In this case, the proposed approach having encoded the word *love* mistakenly predicts the ‘Wedding’ life event. As such Error Type III is essentially related to cases where the subject matter changes abruptly and hence the proposed approach in this paper is very likely (or at least expected) to predict the wrong personal life event given it cannot predict abrupt changes. Table 4 provides further examples for these three types of systematic errors that the proposed approach suffers from.

In addition to the above three types of errors, we also observed some other cases that a human oracle might also have problem in predicting if the actual future tweet was not available. For instance, for the following two tweets: ‘*Man I’m mad I broke my phone. That s[.]t killing me*’ and ‘*Cant wait until Friday to get my new phone*’, our proposed approach predicted a ‘Broken Device’, which is in fact correct; however, the user went on to tweet about a new phone and hence a ‘Device Upgrade’ would have been the more appropriate label, which was not correctly predicted by our approach. It seems that this type of error is similar to Error Type III in the sense that the user moves beyond the immediate scope of the previous tweet onto a different life event; however, it differs from Type III errors in that the transition to the new life event is logical.

Based on these three types of errors, we believe that the work in this paper can be extended in the following directions:

1. The generation architecture proposed in this paper depends on a method to accurately determine the life events of each individual tweet. While there has already been work in the literature that perform life event classification at tweet level, as used in this paper, the performance can still be improved. One of the direction of our future work will focus on learning life event representations in tandem with the Seq2Seq architecture. In other words, instead of performing a pre-processing step to extract the life events for each tweet, we will use an *autoencoder* architecture to learn dense representations for life events. This will have two advantages: 1) it will remove dependence on the life event classifier, and 2) the life event representations will be learnt based on the available training set and will hence be customized for the prediction task.
2. Based on our observation that there are many tweets that have similar components to personal life events but are not personal reports of a life event, we believe that the inclusion of a binary classifier that would determine whether each individual past tweet is a personal report or not could potentially serve as an indicative feature of whether the user will be engaging in a personal report or not. The hypothesis for our future work is that if a user has engaged in a pattern of self reports in the past tweets then it is more likely to do so in the future as well. This binary classifier can either be used independently to act as a filter or in tandem with the recurrent neural network in the same way that the life event classifier was incorporated in our proposed approach in this paper.
3. Finally, the identified type III errors seem to be related to the fact that recurrent neural networks are more dependent on short-range dependencies as opposed to longer-range dependencies [76]. As such, in our model, the latest tweets would have the highest impact on the generated output sequence that would subsequently determine the predicted personal life event. There has been some research that already proposed *hacks* such as reversing the order of the input to bring longer-range dependencies into the short-range [77]; however,

a more systematic approach would be to use *attention models* [78]. In our future work, we are interested in employing attention models in our work so as to determine which portions of the past tweets or which phrases therein should be considered more heavily within the recurrent neural network. This can potentially address type III errors by lowering attention on less relevant content to the personal life event prediction objective.

7. Concluding Remarks

In this paper, we have addressed the problem of predicting future mentions of personal life events on social microblogging platforms, with specific attention to Twitter content. While the literature is abundant with methods that are able to identify and monitor events on social networks at a macro-scale, the detection of personal life events, which are user-specific and have an insignificant social resonance is yet to be fully explored. Within the area of personal life events, a few researchers have shown that it is possible to perform extensive feature engineering to train machine learning classifiers to classify a given tweet into a certain class of life events. Our focus in this paper has been to go beyond the state of the art and predict if and what personal life events will be mentioned in a future and yet unobserved tweet of a given user. To this end, we have proposed two main architectures based on recurrent neural networks that either directly or indirectly predict future personal life event mentions. We have shown in our experiments that this is a non-trivial task and our work is able to provide reasonable performance on a gold standard dataset despite the highly class-imbalanced nature of the personal life event data. We have further explored the errors observed in the obtained results and classified the errors into three primary error type classes, based on which, we have proposed three avenues for future work. In addition, we have only considered six life event types in our work in this paper that would only constitute a subset of all possible personal life event types that can be observed on Twitter. The main reason for this has been the high cost of curating and labeling life event related tweets from Twitter. Future work will need to focus on a larger and broader set of life event types that would show generalizability of this work beyond the six main life event types that have been covered in this paper.

References

- [1] F. Zarrinkalam, H. Fani, E. Bagheri, and M. Kahani, "Inferring implicit topical interests on twitter," in *Advances in Information Retrieval: 38th European Conference on IR Research, ECIR 2016, Padua, Italy, March 20–23, 2016. Proceedings* (N. Ferro, F. Crestani, M.-F. Moens, J. Mothe, F. Silvestri, G. M. Di Nunzio, C. Hauff, and G. Silvello, eds.), (Cham), pp. 479–491, Springer International Publishing, 2016.
- [2] Z. Zhao, H. Lu, D. Cai, X. He, and Y. Zhuang, "User preference learning for online social recommendation," *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 9, pp. 2522–2534, 2016.
- [3] H. Fani, F. Zarrinkalam, E. Bagheri, and W. Du, "Time-sensitive topic-based communities on twitter," in *Canadian Conference on Artificial Intelligence*, pp. 192–204, Springer, 2016.
- [4] H. Fani, E. Bagheri, and W. Du, "Temporally like-minded user community identification through neural embeddings," in *26th ACM International Conference on Information and Knowledge Management (CIKM)*, 2017.
- [5] H. Fani, E. Bagheri, F. Zarrinkalam, X. Zhao, and W. Du, "Finding diachronic like-minded users," *Computational Intelligence*, 2017.
- [6] A. Madani, O. Boussaid, and D. E. Zegour, "Real-time trending topics detection and description from twitter content," *Social Network Analysis and Mining*, vol. 5, p. 59, Oct 2015.
- [7] S. B. Kaleel and A. Abhari, "Cluster-discovery of twitter messages for event detection and trending," *Journal of Computational Science*, vol. 6, pp. 47–57, 2015.
- [8] W. Fan and M. D. Gordon, "The power of social media analytics," *Communications of the ACM*, vol. 57, no. 6, pp. 74–81, 2014.

- [9] T. H. Nguyen, K. Shirai, and J. Velcin, "Sentiment analysis on social media for stock movement prediction," *Expert Systems with Applications*, vol. 42, no. 24, pp. 9603–9611, 2015.
- [10] T. H. Nguyen and K. Shirai, "Topic modeling based sentiment analysis on social media for stock market prediction," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1354–1364, Association for Computational Linguistics, 2015.
- [11] Y. Mao, W. Wei, B. Wang, and B. Liu, "Correlating s&p 500 stocks with twitter data," in *Proceedings of the First ACM International Workshop on Hot Topics on Interdisciplinary Social Networks Research, HotSocial '12*, (New York, NY, USA), pp. 69–72, ACM, 2012.
- [12] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *Journal of computational science*, vol. 2, no. 1, pp. 1–8, 2011.
- [13] M. Nofer and O. Hinz, "Using twitter to predict the stock market," *Business & Information Systems Engineering*, vol. 57, no. 4, pp. 229–242, 2015.
- [14] T. Liu, X. Ding, Y. Chen, H. Chen, and M. Guo, "Predicting movie box-office revenues by exploiting large-scale social media content," *Multimedia Tools and Applications*, vol. 75, no. 3, pp. 1509–1528, 2016.
- [15] D. Choudhery and C. K. Leung, "Social media mining: Prediction of box office revenue," in *Proceedings of the 21st International Database Engineering & Applications Symposium, IDEAS 2017*, (New York, NY, USA), pp. 20–29, ACM, 2017.
- [16] H. Woo, Y. Cho, E. Shim, J.-K. Lee, C.-G. Lee, and S. H. Kim, "Estimating influenza outbreaks using both search engine query data and social media data in south korea," *Journal of medical Internet research*, vol. 18, no. 7, 2016.
- [17] K. Byrd, A. Mansurov, and O. Baysal, "Mining twitter data for influenza detection and surveillance," in *Proceedings of the International Workshop on Software Engineering in Healthcare Systems, SEHS '16*, (New York, NY, USA), pp. 43–49, ACM, 2016.
- [18] F. Franch, "(wisdom of the crowds) 2: 2010 uk election prediction with social media," *Journal of Information Technology & Politics*, vol. 10, no. 1, pp. 57–71, 2013.
- [19] A. Tsakalidis, S. Papadopoulos, A. I. Cristea, and Y. Kompatsiaris, "Predicting elections for multiple countries using twitter and polls," *IEEE Intelligent Systems*, vol. 30, no. 2, pp. 10–17, 2015.
- [20] M. P. Cameron, P. Barrett, and B. Stewardson, "Can social media predict election results? evidence from new zealand," *Journal of Political Marketing*, vol. 15, no. 4, pp. 416–432, 2016.
- [21] R. Wald, T. M. Khoshgoftaar, A. Napolitano, and C. Sumner, "Using twitter content to predict psychopathy," in *Machine Learning and Applications (ICMLA), 2012 11th International Conference on*, vol. 2, pp. 394–401, IEEE, 2012.
- [22] D. Preotiuc-Pietro, J. Eichstaedt, G. Park, M. Sap, L. Smith, V. Tobolsky, H. A. Schwartz, and L. Ungar, "The role of personality, age and gender in tweeting about mental illnesses," in *NAACL HLT*, vol. 2015, p. 21, 2015.
- [23] M. De Choudhury, M. Gamon, S. Counts, and E. Horvitz, "Predicting depression via social media.," *ICWSM*, vol. 13, pp. 1–10, 2013.
- [24] G. Coppersmith, M. Dredze, C. Harman, and K. Hollingshead, "From adhd to sad: Analyzing the language of mental health on twitter through self-reported diagnoses.," in *CLPsych@ HLT-NAACL*, pp. 1–10, 2015.
- [25] S. Unankard, X. Li, and M. A. Sharaf, "Emerging event detection in social networks with location sensitivity," *World Wide Web*, vol. 18, no. 5, pp. 1393–1417, 2015.
- [26] G. Paltoglou, "Sentiment-based event detection in twitter," *Journal of the Association for Information Science and Technology*, vol. 67, no. 7, pp. 1576–1587, 2016.
- [27] T. Dickinson, M. Fernandez, L. A. Thomas, P. Mulholland, P. Briggs, and H. Alani, "Identifying prominent life events on twitter," in *Proceedings of the 8th International Conference on Knowledge Capture, K-CAP 2015*, (New York, NY, USA), pp. 4:1–4:8, ACM, 2015.
- [28] P. Cavalin, M. Gatti de Baysar, and C. Pinhanez, "Towards personalized offers by means of life event detection on social media and entity matching," in *HT (Doctoral Consortium/Late-breaking Results/Workshops)*, vol. 1210, 01 2014.
- [29] S. Choudhury and H. Alani, "Detecting presence of personal events in twitter streams," in *International Conference on Social Informatics*, pp. 157–166, Springer, Springer International Publishing, 2014.
- [30] J. L. Li, A. Ritter, C. Cardie, and E. H. Hovy, "Major life event extraction from twitter based on congratulations/condolences speech acts," in *EMNLP*, 2014.
- [31] M. Ni, Q. He, and J. Gao, "Using social media to predict traffic flow under special event conditions," in *The 93rd Annual Meeting of Transportation Research Board*, 2014.
- [32] J. C. Eichstaedt, H. A. Schwartz, M. L. Kern, G. Park, D. R. Labarthe, R. M. Merchant, S. Jha, M. Agrawal, L. A. Dziurzynski, M. Sap, *et al.*, "Psychological language on twitter predicts county-level heart disease mortality," *Psychological science*, vol. 26, no. 2, pp. 159–169, 2015.
- [33] N. B. Lassen, R. Madsen, and R. Vatrappu, "Predicting iphone sales from iphone tweets," in *Proceedings of the 2014 IEEE 18th International Enterprise Distributed Object Computing Conference, EDOC '14*, (Washington,

- DC, USA), pp. 81–90, IEEE Computer Society, 2014.
- [34] V. Radosavljevic, M. Grbovic, N. Djuric, and N. Bhamidipati, “Large-scale world cup 2014 outcome prediction based on tumblr posts,” in *KDD Workshop on Large-Scale Sports Analytics*, 2014.
- [35] B. Di Eugenio, N. Green, and R. Subba, “Detecting life events in feeds from twitter,” in *Semantic Computing (ICSC), 2013 IEEE Seventh International Conference on*, pp. 274–277, Ieee, 2013.
- [36] S. Choudhury and H. Alani, “Personal life event detection from social media,” in *Doctoral Consortium and Workshop Proceedings of the 25th ACM Hypertext and Social Media Conference*, vol. 1210, 01 2014.
- [37] P. R. Cavalin, L. G. Moyano, and P. P. Miranda, “A multiple classifier system for classifying life events on social media,” in *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on*, pp. 1332–1335, IEEE, 2015.
- [38] L. G. Moyano, P. R. Cavalin, and P. P. Miranda, “Life event detection using conversations from social media,” in *Brazilian Workshop on Social Network Analysis and Mining*, 2015.
- [39] M. Khodabakhsh, M. Kahani, E. Bagheri, and Z. Noorian, “Detecting life events from twitter based on temporal semantic features,” *Knowledge-based Systems Journal*, 2017, *second revision submitted*, http://1s3.rnet.ryerson.ca/wiki/images/4/48/Life_Event_Detection.pdf.
- [40] I. Sutskever, J. Martens, and G. Hinton, “Generating text with recurrent neural networks,” in *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML’11, (USA), pp. 1017–1024, Omnipress, 2011.
- [41] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *Proceedings of the 26th International Conference on Neural Information Processing Systems*, NIPS’13, (USA), pp. 3111–3119, Curran Associates Inc., 2013.
- [42] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Proceedings of the 27th International Conference on Neural Information Processing Systems*, NIPS’14, (Cambridge, MA, USA), pp. 3104–3112, MIT Press, 2014.
- [43] K. Cho, B. van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning phrase representations using rnn encoder–decoder for statistical machine translation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, (Doha, Qatar), pp. 1724–1734, Association for Computational Linguistics, Oct. 2014.
- [44] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Comput.*, vol. 9, pp. 1735–1780, Nov. 1997.
- [45] D. Tang, B. Qin, and T. Liu, “Document modeling with gated recurrent neural network for sentiment classification,” in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, (Lisbon, Portugal), pp. 1422–1432, Association for Computational Linguistics, September 2015.
- [46] X. Glorot, A. Bordes, and Y. Bengio, “Domain adaptation for large-scale sentiment classification: A deep learning approach,” in *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML’11, (USA), pp. 513–520, Omnipress, 2011.
- [47] X. Zhang, J. Zhao, and Y. LeCun, “Character-level convolutional networks for text classification,” in *Proceedings of the 28th International Conference on Neural Information Processing Systems*, NIPS’15, (Cambridge, MA, USA), pp. 649–657, MIT Press, 2015.
- [48] J. Y. Lee and F. Deroncourt, “Sequential short-text classification with recurrent and convolutional neural networks,” in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics*, 2016.
- [49] Y. Kim, “Convolutional neural networks for sentence classification,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 2014.
- [50] A. Graves, “Generating sequences with recurrent neural networks,” in *ICML Representation Learning Workshop*, 2012.
- [51] 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 (Volume 1: Long Papers)*, (Baltimore, Maryland), pp. 1555–1565, Association for Computational Linguistics, June 2014.
- [52] A. Severyn and A. Moschitti, “Twitter sentiment analysis with deep convolutional neural networks,” in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’15, (New York, NY, USA), pp. 959–962, ACM, 2015.
- [53] N. Limsopatham and N. Collier, “Bidirectional lstm for named entity recognition in twitter messages,” in *Proceedings of the 2nd Workshop on Noisy User-generated Text*, pp. 145–152, 2016.
- [54] Y. Li, T. Liu, J. Jiang, and L. Zhang, “Hashtag recommendation with topical attention-based lstm,” in *Proceedings of the 26th International Conference on Computational Linguistics*, pp. 943–952, Coling, 2016.
- [55] J. Ma, W. Gao, P. Mitra, S. Kwon, B. J. Jansen, K.-F. Wong, and M. Cha, “Detecting rumors from microblogs with recurrent neural networks,” in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI’16, pp. 3818–3824, AAAI Press, 2016.

- [56] H. Lin, J. Jia, Q. Guo, Y. Xue, Q. Li, J. Huang, L. Cai, and L. Feng, "User-level psychological stress detection from social media using deep neural network," in *Proceedings of the 22Nd ACM International Conference on Multimedia*, MM '14, (New York, NY, USA), pp. 507–516, ACM, 2014.
- [57] S. Vosoughi, P. Vijayaraghavan, and D. Roy, "Tweet2vec: Learning tweet embeddings using character-level cnn-lstm encoder-decoder," in *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '16, (New York, NY, USA), pp. 1041–1044, ACM, 2016.
- [58] B. Dhingra, Z. Zhou, D. Fitzpatrick, M. Muehl, and W. W. Cohen, "Tweet2vec: Character-based distributed representations for social media," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, vol. 2, pp. 269–274, 05 2016.
- [59] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," in *Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation (SSST-8)*, 2014.
- [60] Z. Yang, D. Yang, C. Dyer, X. He, A. J. Smola, and E. H. Hovy, "Hierarchical attention networks for document classification," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 1480–1489, 2016.
- [61] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.
- [62] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," in *Proceedings of Workshop at ICLR*, 2013.
- [63] M. J. Kusner, Y. Sun, N. I. Kolkin, and K. Q. Weinberger, "From word embeddings to document distances," in *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37, ICML'15*, pp. 957–966, JMLR.org, 2015.
- [64] A. Sordani, Y. Bengio, H. Vahabi, C. Lioma, J. Grue Simonsen, and J.-Y. Nie, "A hierarchical recurrent encoder-decoder for generative context-aware query suggestion," in *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, CIKM '15*, (New York, NY, USA), pp. 553–562, ACM, 2015.
- [65] S. Semeniuta, A. Severyn, and E. Barth, "A hybrid convolutional variational autoencoder for text generation," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pp. 638–648, 2017.
- [66] M. Oak, A. Behera, T. Thomas, C. O. Alm, E. Prud'hommeaux, C. Homan, and R. W. Ptucha, "Generating clinically relevant texts: A case study on life-changing events," in *CLPsych@HLT-NAACL*, pp. 85–94, 2016.
- [67] L. Shang, Z. Lu, and H. Li, "Neural responding machine for short-text conversation," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing*, pp. 1577–1586, The Association for Computer Linguistics, 2015.
- [68] I. V. Serban, T. Klinger, G. Tesauro, K. Talamadupula, B. Zhou, Y. Bengio, and A. C. Courville, "Multiresolution recurrent neural networks: An application to dialogue response generation," in *AAAI*, pp. 3288–3294, 2017.
- [69] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.
- [70] A. Graves, "Sequence transduction with recurrent neural networks," in *International Conference of Machine Learning (ICML) 2012 Workshop on Representation Learning*, 2012.
- [71] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *the 3rd International Conference for Learning Representations*, 2015.
- [72] F. Bastien, P. Lamblin, R. Pascanu, J. Bergstra, I. Goodfellow, A. Bergeron, N. Bouchard, D. Warde-Farley, and Y. Bengio, "Theano: new features and speed improvements," in *Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop*, 2012.
- [73] A. Ritter, S. Clark, O. Etzioni, *et al.*, "Named entity recognition in tweets: an experimental study," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pp. 1524–1534, Association for Computational Linguistics, 2011.
- [74] O. Owoputi, C. Dyer, K. Gimpel, N. Schneider, and N. A. Smith, "Improved part-of-speech tagging for online conversational text with word clusters," in *In Proceedings of NAACL*, Association for Computational Linguistics, 2013.
- [75] L. Breiman, "Bagging predictors," *Mach. Learn.*, vol. 24, pp. 123–140, Aug. 1996.
- [76] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE transactions on neural networks*, vol. 5, no. 2, pp. 157–166, 1994.
- [77] X. Shi, K. Knight, and D. Yuret, "Why neural translations are the right length," in *EMNLP*, pp. 2278–2282, 2016.
- [78] K. Cho, A. Courville, and Y. Bengio, "Describing multimedia content using attention-based encoder-decoder networks," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 1875–1886, 2015.