FROM LANGUAGE TO LOCATION USING MULTIPLE INSTANCE NEURAL NETWORKS

by

Sneha Nagpaul A Thesis Submitted to the Graduate Faculty of George Mason University In Partial fulfillment of The Requirements for the Degree of Master of Science Computer Science

Committee:

Date:	Spring 2018 George Mason University Fairfax, VA
	Dr. Kenneth S. Ball, Dean, Volgenau School of Engineering
	Dr. Sanjeev Setia, Chairman, Department of Computer Science
	Dr. Aditya Johri, Committee Member
	Dr. Jessica Lin, Committee Member
	Dr. Huzefa Rangwala, Thesis Director

From Language to Location using Multiple Instance Neural Networks

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science at George Mason University

By

Sneha Nagpaul Bachelor of Science University of Delhi, 2010

Director: Dr. Huzefa Rangwala, Professor Department of Computer Science

> Spring 2018 George Mason University Fairfax, VA

 $\begin{array}{c} \mbox{Copyright} \textcircled{C} \mbox{ 2018 by Sneha Nagpaul} \\ \mbox{ All Rights Reserved} \end{array}$

Dedication

To, the Cat.

Acknowledgments

Foremost, I would like to thank to my advisor and thesis director, Dr. Huzefa Rangwala for his motivation, enthusiasm, and immense knowledge. This project could not have been completed without his patience and guidance during the process. I am particularly appreciative of the focus he provided to this work in a rapidly evolving research environment.

In addition to my advisor, I would like to thank the rest of my thesis committee: Dr. Jessica Lin and Dr. Aditya Johri for their encouragement, time and insightful comments. Dr. Lin's feedback helped polish the research. Dr. Johri pointed out various applications, hence putting this work in a fresh perspective.

I would like to extend my gratitude to Dr. Vadim Sokolov, Assistant Professor with the Operations Research Department at George Mason, for sharing his insights on deep learning while entertaining research ideas and providing notes.

I would also like to acknowledge Mohammad Arifur Rahman, who helped me navigate the murky waters of multiple instance learning research. A special thanks to Kyle Jackson for reading the draft and providing notes on language.

Finally, I would like to thank my family and friends.

Table of Contents

			Pag	je
List	t of T	ables .	· · · · · · · · · · · · · · · · · · ·	ii
List	t of F	igures .	vi	ii
Abs	stract			х
1	Intr	oductio	n	0
		1.0.1	Contributions	1
		1.0.2	Thesis Outline	2
2	Bac	kground	l and Related Work	3
	2.1	Multip	le Instance Learning	3
		2.1.1	Single Instance Learning (SIL)	5
		2.1.2	Support Vector Machines for MIL	5
		2.1.3	Group Instance Cost Function (GICF)	6
	2.2	Geogra	aphic Information Retrieval	7
		2.2.1	Traditional Methods	7
		2.2.2	Language Modeling	8
		2.2.3	Neural Networks for GIR	8
	2.3	Neural	Networks and Deep Learning Primer	8
		2.3.1	Multilayer Neural Networks and Back-propagation	9
		2.3.2	Convolutional Neural Networks	0
		2.3.3	Distributed Representations of Words 1	1
		2.3.4	Recurrent Neural Networks	2
3	Prop	posed N	ſethod	3
	3.1	Proble	m Statement	3
	3.2	Metho	d Description	3
	3.3	Metho	d Characteristics	4
4	Res	ults		8
	4.1	Datase	ets	8
	4.2	Experi	mental Setup	9
	4.3	Explor	ing the Hyper-parameter space 2	1
	4.4	Perform	mance comparison with MIL methods	5

	4.5	Runnir	ng Time Comparison	27
		4.5.1	Running time vs Accuracy	28
		4.5.2	Running time vs F measure	30
	4.6	Illustra	ative Examples	32
		4.6.1	New York City	34
		4.6.2	San Francisco	38
	4.7	Insight	s from results	42
5	Disc	ussion a	and Conclusion	44
	5.1	Future	Work	44
		5.1.1	Transfer Learning	45
		5.1.2	Multiple Classification	45
		5.1.3	Other instance level models	45
		5.1.4	Exotic Aggregation Functions	46
		5.1.5	Other Applications	46
Bib	liogra	aphy		47

List of Tables

Table		Page
4.1	Datasets	20
4.2	Initial high level results on datasets	21
4.3	Hyper-Parameter Accuracy	22
4.4	Hyper-Parameter F-score	23
4.5	Hyper-Parameter Running Time	23
4.6	Accuracy	26
4.7	F measure	26
4.8	Running Time	28

List of Figures

Figure		Page
2.1	Multiple Instance Learning Framework	5
3.1	Multiple Instance Learning for Twitter	13
3.2	Model Architecture	15
3.3	Instance Level Model	16
3.4	Aggregation Model	17
4.1	Datasets Map	19
4.2	Hyperparameter Comparison	24
4.3	Accuracy vs Running Time by City	29
4.4	F measure vs Running Time by City	31
4.5	Word Cloud - NYC	32
4.6	Word Cloud - SF	33
4.7	Positive Example in New York City : User 1	35
4.8	Positive Example in New York City : User 2	36
4.9	Negative Example in New York City: User 3	37
4.10	Positive Example in San Francisco : User 1	38
4.11	Positive Example in San Francisco : User 3	39
4.12	Positive Example in San Francisco : User 3	40
4.13	Negative Example in San Francisco :User 4	41

Abstract

FROM LANGUAGE TO LOCATION USING MULTIPLE INSTANCE NEURAL NET-WORKS

Sneha Nagpaul

George Mason University, 2018

Thesis Director: Dr. Huzefa Rangwala

Given the deluge of data caused by crowd generated content from social media websites, the complexity of extracting information from has increased manifold. An important characteristic of such text is its original location which can in turn be used to respond to emergencies such as floods and crimes[1]. The patterns discovered by such geolocation of social media related unstructured text can also be used commercially for targeted advertising and recommender systems[2][3][4]. This work deals with geolocating short texts from social media that are labeled with a user's information. However, instead of locating the user who can be viewed as a collection of these texts, it focuses on locating each such text, here a tweet. For this task, the problem is described within the multiple instance learning framework and a novel approach using neural networks is designed which trains a tweet level classifier using only user location labels. The model outperforms the state of the art in multiple instance learning and provides significant scalability and speedup compared to existing methods. Exceeding the Bag of Words models prevalent in prior geolocation research, the intuitive tweet level neural network classifier discovers complex features such as grammar and identifies name places without feature engineering[5][6].

Chapter 1: Introduction

Due to the growth in content generated by social media platforms, it is becoming increasingly complex to distill information from this potential resource. A simple and useful characteristic of the these data is the location of the content generator at a given time. This information can be leveraged to encode events that are actionable for authorities such as natural calamities and crimes [7][1]. Additionally, learning the location specific patterns in this language can be used for commercial purposes such as targeted advertising, resource allocation and recommender systems [2][3].However, only a small percentage of the generated content is geo-tagged and thus there exists a need to automate the process of inferring location from the content itself.

It is often the case in real world scenarios that labeled information is available at an aggregated level for a set of instances; however, knowledge distillation is required at a more granular level. An example of a nefarious use of these methods could be if the outcome of an election in an area is known, the probability of each of the voters' choices could be determined; thus, threatening their privacy [8]. However, a use case that furthers social good can be found in medicine where given the genetic makeup and the occurrence of a disease(label), individual genes can be held responsible and targeted for treatment and diagnosis. In the language processing community, researchers have worked on the problem that if the sentiment of a product or movie review is known, how this sentiment will be propagated to the sentence or the word level [9]. The method that is proposed here can be viewed as an application of the multiple instance framework. In this work, given discourse from a certain region (aggregate level), the method will endeavor to transfer this regional label to a language substructure like a phrase (disaggregate level).

Prior work on predicting location by using short social media text tries to classify substantially lengthy content using indicative words into regional groups. In particular, research has focused on location labels available for a user active on Twitter [5][10] using a body of their tweets. This work identifies words pertaining to a certain region by using a Bag of words model to process language; and thus distorts phrasal information in the language.

Viewing this problem within the multiple instance learning framework can help to solve the geolocation problem at a granular level while preserving the structure of the text. In this formulation each tweet can be regarded as an instance and each user can be regarded as a group of tweets, or a bag. The multiple instance learning model then propagates the label at the user level to each tweet which is most indicative of the users location.

Prior work in multiple instance learning makes strong assumptions about the membership of instances inside a bag in order to train the classifier based on the bag label [11]. More recently, researchers have tried to relax this assumption by assuming that bag level labels are simply an aggregation of instance level predictions [9]. However, even these relaxed methods rely on a similarity measure that requires training of deep neural networks in order to learn meaningful embeddings for classification. This process can be subjective and specific and, hence, can be viewed as de-facto feature engineering. A common issue with prior approaches is the use of the similarity kernel which makes the methods hard to scale as the size of the dataset grows.

In this work, the problem is set up as finding patterns of language that represent a certain geographic location when aggregate level information is provided and needs to be transferred to each instance to locate an individual tweet.

1.0.1 Contributions

In this work, a novel framework for Multiple Instance Learning is proposed using neural networks which addresses some of the gaps from prior research in multiple instance learning and geographic information retrieval. Thus, the proposed method makes the following contributions:

1. Provides a framework for transferring bag labels to instance level predictions using a representation learning framework which can accommodate any input directly without

the need to encode with embeddings trained from another neural network as in [9] or kernel. This is a major contribution of the work since it eliminates the need for feature engineering which is often subjective and case specific. The neural network learns the latent space directly for the unlabeled data using the aggregated information labels.

- 2. The model is trained using gradient descent through an end-to-end trainable neural network. Since the model training can work with batches, the proposed method scales up easily in terms of memory requirements and provides a significant computational speedup.
- 3. Relaxes assumptions about instance level memberships in bags, while providing a flexible setting to train an instance level classifier. This increases applicability by providing an intuitive architecture for finding instance level classifier which has the flexibility to be replaced with any trainable neural network architecture depending on the use case. This flexibility is also extended to how the instance level labels are aggregated to form bag level labels.

The proposed model tests itself on a novel application for a geolocation task that classifies tweets as pertaining to a particular location and thus discovers language indicative of a region.

1.0.2 Thesis Outline

The rest of this work is organized into 5 sections. Section 2 provides background and related research in fields of multiple instance learning and geolocating text while pointing out to their contributions and gaps. Section 3 describes the method used to address the shortcomings of existing research and goes into detail on model architecture and training methods. Section 4 details the dataset, experimental setting and results and provides illustrative examples from the model. Section 5 concludes with a discussion and suggests future directions.

Chapter 2: Background and Related Work

Given this setup of the geolocation problem in this work, we discuss prior literature in Multiple Instance learning which deals with transferring of labels from a broader to a finer level and Geographic Information Retrieval which is used to map language to location in general. Finally, a brief review of neural networks is provided that introduces the capability of representation learning for discovering non-linear relationships in data without feature engineering.

2.1 Multiple Instance Learning

Multiple Instance Learning refers to a class of learning methods that work with labeled information at an aggregate level while trying to distill information at a disaggregate level. This natural framework is prevalant in a variety of problems and has been employed in solving content based image retrieval and classification [12], privacy related applications [13], categorising texts [11] and recognising objects [14].

This framework was traditionally employed to explore the feature space for a supervised learning problem, where many contenders of feature sets constituted the individual instances for the original instance (the bag) but only one of those feature vectors was responsible for the bag level labels [15]. Thus, the aggregation function was an OR relationship, wherein negative bags had only negative instances and positive bags had at least one positive instance [15]. This stringent assumption fails to generalize as soon as the use case moves away from exploring potential features for a certain supervised learning problem. Additionally, the main goal of this formulation was to infer bag level labels as opposed to exploring granularity [15]. In the context of locating a person from a tweet, someone belonging to a certain location might travel and tweet about the destination. It is also possible that they might relocate and absorb the linguistic style of the new region. Hence, the OR formulation is inappropriate for the location use case as all tweets may not be indicative of region. There might be a general language pattern to account for or travel which may result a negative bag to have positive instances.

In order to relax this strong assumption, researchers have proposed other ways to link instance level labels to higher level labels. A work by Xu and Frank posits that there is equal contribution from all the constituents of a bag which are considered independent [16]. Zhou et al.[17] make the assumption that within a negative bag all instances are negative and then use a semi supervised approach to classify further. In another generalisation of the original aggregation, Weidmann et al.[18] classify bags based on the presence of a collection of types of instances where a type is a concept which can be viewed as an atleast-one-of-each formulation instead of an at-least-one-formulation. This would be akin to engineering features and looking for their presence in the user. As a hypothetical example, if they are health conscious and eat orgranic food and reference technology in their tweets then we might want to classify them as a belonging to San Francisco. This formulation, though intuitive, would require separate treatment of different datasets and would stand on subjective assumptions made by the modeler as opposed to discovering patterns present within the data.

A more pertinent body of work tries to shift focus from bag level to instance level predictions. In particular, Liu et al. [19], utilize this framework to identify key instances within a bag i.e. those that are most relevant to the positively labeled bags using nearest neighbor heuristics subject to various aggregation schemes. While this formulation is particularly fitting to the geolocation use case, the combinatorial complexity for the nearest neighbor based metrics and checking for various ways instances in the positive bags might have voted towards the bag label make this approach unscalable.



Figure 2.1: Multiple Instance Learning Framework - In the traditional setting, the negative bags(left) would only contain negative instances and the positive bags(right) would contain at least one positive instance. This is the implication of the OR function constraint on the instance level to bag level label transfer[15].

2.1.1 Single Instance Learning (SIL)

A naive way to tackle the problem is Single Instance Learning (SIL) which assumes that every instance of a bag has the bags label [20]. This assumption is also equivalent to an AND transfer relationship. This approach seems to match the framework of the problem under consideration as it can be considered reasonable to assume that if one tweet is from New York, the rest are as well. However, it doesn't take into account general phrases across various places and thus could confuse the classifier due to additional noise in the dataset [20].

2.1.2 Support Vector Machines for MIL

A more recent body of work harnesses the Support Vector Machine (SVM) framework while maintaining the instance to bag label transfer relationship as the traditional OR [11] [15]. In their work Adrews et al. [11] propose MISVM models which train the parameters using all possible combinations of instance and bag level labels to find the model that maximizes the soft margin criterion. A kernel based formulation can be seen in Gartner et al. [21], where vanilla support vector machines are deployed to work on the bag level directly. In computer vision related applications, researchers have attempted to find regions of interest within an image by connection through instance level labels [22][23]. Sparse multi- instance learning approaches are mentioned in [24] and transductive SVM's [25] which explore viable low density regions for placing the decision boundaries.

Thus, due to the presence of a kernel and testing for many possibilities, the complexity of support vector machine based methods becomes intractable, rendering them unscalable in terms of memory and computational resource needs. These drawbacks limit the use of MISVM to relatively smaller datasets.

2.1.3 Group Instance Cost Function (GICF)

More recently, neural networks are being leveraged in order to address the rigidity of aggregation functions[9]. The Group Instance Cost Function (GICF) method uses an adaptive cost function that tries to find the parameters optimized for instance and bag level costs. They achieve this using a two part approach wherein one part of the loss function assigns labels that justify the instance-bag relationship and the other part assigns similar labels to similar instances across bags. To achieve this GICF uses a similarity measure that determines differences among instances and bags.

GICF's use of a similarity measure makes it a kernel based method like SVMs which causes similar scalability issues. Another perspective on the similarity kernel is that there are certain requirements for feature engineering before the model can process the bags and instances. The generated feature vectors are then passed on to a training method with the two-tier cost function to find a classifier at the bag level. Following this, the instances can be treated by the classifier as a bag of one to be classified. While this method takes into account modern deep learning architectures, it still requires designing of the similarity measure which can be subjective and hence harder to generalize.

While GICF is the closest in terms of architecture to the work described in this paper, there are a few more things to consider while comparing the proposed method to it; particularly the problems these methods have originally tackled. GICF has been tested on sentiment data wherein the instances are sentences in a review and heavily share context. In case of a bag of tweets, even though the tweets come from the same person, they have no requirement to have been originated at the same time and share any context as is the case with a paragraph of a review.

MISVM and related works use strong assumptions for instance membership and aggregation functions. The work that relaxed these assumptions is GICF but requires hard to train embedding vectors and similarity measures. Old works have focused on context heavy instance bag combinations such as reviews. The proposed method will try to complement the prior research by applying the framework within the constraints of geolocation of twitter based dataset which consists of tweet instances that share the source but not the context necessarily. Thus relaxing assumptions of membership and aggregation enough for the data to be able to provide those patterns to the model.

2.2 Geographic Information Retrieval

Geographic Information Retrieval refers to a class of methods that deal with mapping language to location [26]. General tasks include geolocating text, essentially predicting the location for a given text and dialectology which deals with studying language patterns of a certain geographical area. This work extends GIR literature as it endeavours to predict the location of tweets.

2.2.1 Traditional Methods

The common theme in traditional GIR is an auxiliary dataset, called gazetteer, that maps words to locations statically [26][27]. In addition to this gazetteer lookup, heuristic based feature engineering methods such as term frequency analysis and pinpointing the positional attributes of words are leveraged to accomplish this task. These methods were traditionally used on regional applications due to the non-scalable nature of the gazetteer [28]. For application at a larger scale to online datasets, a naive implementation of place name detection followed by disambiguation was adopted to keep the size of the auxiliary dataset tractable [29].

2.2.2 Language Modeling

A scalable class of methods that modelled languages emerged as datasets became bigger and more intractable towards the early 2010s [30]. Unsupervised learning methods such as topic modelling were initially explored which are harder to scale [27]. To overcome the gazetteer disadvantage and achieve larger scalability, supervised learning methods were used on the datasets that are labelled with user level information [30]. In the context of this paper, they deal with user level labels and require a string of tweets to geolocate on the bag level.

2.2.3 Neural Networks for GIR

With the emergence of deep learning and neural networks, recent work has focussed on modelling language features with neural networks. Liu et al. [31] work at locating users instead of language using feedforward neural networks. Another recent work, uses bag of words representations which distorts structural language features of user documents for a neural network to classify and find indicative words[5].

A shortcoming of these methods is that they require a considerable amount of information even to geolocate at a bag level - they do not locate a tweet, they locate a user. Another issue is the extensive use of the bag of words model which distorts structural information in language that could be indicative of region. This work addresses both of these problems by leveraging representation learning at a fine grained level with minimal preprocessing. By generalising the design of the neural network, milNN also provides a framework to vary architecture choices based on expected complexity without feature engineering.

2.3 Neural Networks and Deep Learning Primer

Machine learning has conventionally been a process where researchers try to transform raw input data into an internal representation that is appropriate for a learning subsystem to consume in order to identify or classify patterns of interest [32]. Representation learning is a class of methods that deal with automatically finding these internal representations in order to perform machine learning tasks. Deep Learning is one kind of representation learning method that deal with finding these features through a series of nonlinear transformations on minimally processed input to discover representation that make the data suitable for tasks like classification and detection [32].

A key feature of these methods is that there is minimal intervention with the data before being fed to the model and there is no requirement for explicit feature engineering from experts in a particular field. For example, a linguist need not be involved to work on language related tasks. Hence, if the method was initially developed for english, it would just as easily generalize to other formal languages such as chinese or french and informal languages, such as those used on Twitter and Facebook and vice versa.

Another major advantage of deep learning over any methods that use feature engineering is that if the usage changes in the future, the model will simply update itself based on the newer pattern. This idea is known as concept drift, where the data inherently changes and older patterns become obsolete [32]. In the case where a conventional machine learning approach is being used, a new feature vector would need to be designed using considerable subjectivity and time of a domain expert in order to accommodate such changes manually.

2.3.1 Multilayer Neural Networks and Back-propagation

A simple yet powerful application can be found in the popular feed forward neural networks or multi layer perceptron (MLP) which map a fixed size input and to a fixed size output. The intermediate layers are a series of nonlinear activations on affine transformations of the inputs which are all trainable [32]. These intermediate weights are learned using backpropagation which is an application of the chain rule for composite functions wherein the error is propagated backwards from the point where a label is available for a particular training instance and a loss can be calculated and adjustments can be made at initial levels all the way to the first layer where the input was fed into the network [33]. The direction of change of the intermediate weights is determined by calculating the gradient which is derivative of the nonlinear transformation and the adjustment is made in the opposite direction of the gradient.

While deep learning methods have been widely accepted by the computer vision community, this review will focus on how language related applications have benefited from the neural network paradigm with exception of a few illustrative examples.

Due to the implications of the universal approximation theorem [34], a neural network with a single hidden layer is capable of generalising any function given enough data and time. While this theoretical results makes basic neural networks attractive, it is often faster to explore exotic architectures to accomplish these tasks in reasonable time with data restrictions. Hence, the rest of this section will discuss convolutional neural networks and recurrent neural networks.

2.3.2 Convolutional Neural Networks

Convolutional Neural Networks are a class of deep learning models that deal with spatial correlations in the input [35][36]. The models involve sliding a filter on some structured data that may have compositional structure and subsequently discovering motifs that may be present throughout the raw structure. An example of such a structure is an image that has two balls at opposite corners and a filter would detect the ball having strong local similarity yet being able to occur in another part of the picture. This filter is a mathematical convolution and was traditionally hand designed as a part of a feature engineering step (eg. Edge detection). However, with neural networks, these filters are learned as trainable weights which are shared within each feature map which is essentially the output of applying the filter to the previous layer. The non linearity can be introduced as an activation on the tensor generated as the feature map. While this formulation is intuitive and popular in vision related applications, more recently they have shown to be useful with language and sequence related tasks[37]. In these language related applications, the convolution can be seen as a one dimensional window being slid over language substructures looking for

patterns.

2.3.3 Distributed Representations of Words

When shifting focus from the vision application to lanaugage the input space changes considerably. Now, as opposed to an image which is a tensor of numbers, symbols like words and letters require a distinct point of view for being tackled by a neural network. It is helpful to understand this handling with the example of an experiment where given a sequence of words the next word is the target for prediction as described in [38]. The input to this model is a one hot encoded series of words that belong to the sequence. At the first layer, every word in the series generates a pattern of activations or a word vector. This vector is then mapped by subsequent layers to the output which is a large vector that includes the probability of each of the words being next. While training, this model's output is the one hot encoded vector for the predicted word. While backpropagating from those errors, a distributed representation is learned by the activations in the input layer. LeCun et al.[32] talk of them as multiple micro-rules for symbols that the network automatically learns. This can also be viewed as 'intuitive' inference as opposed to logical inference that would stem from rigid grammatical rules.

One such example of these distributed representation is the Word2Vec model[39]. In the vector space, the word for king and queen appear closer whereas they would be far apart in a lexicon. Thus moving away from rigid representations to those that share meaning beyond fixes rules that need to manufactured by domain experts. Vector representations appear widely in natural language processing literature to accomplish a multitude of tasks [40][41][42][43][44][45][46]. Representations learned during training of a certain network can be indicative of how the language/symbols relate to the task. For example, the earlier experiment of predicting the next word provides spatial set up of language in the word vectors. Another perspective for looking at these vectors (or embeddings) is as a language model where the vector constitutes a particular word's role and usage instructions in that model's space. Older work goes as far as N-gram type modeling which is a generalisation of

a bag of words model where N = 1. Those models do not have the capabilities of modeling relationships between words like the distributed representations do.

2.3.4 Recurrent Neural Networks

While MLPs and CNNs are successful neural network models, they exhibit rigidity in the input they can take in that the input to these feed forward neural networks must be of a constant length. Recurrent Neural Networks provide more flexibility to this input structure as they process the input sequence by the element and maintain an internal state with the information they have encountered [32]. A modification of RNNs build a longer context into the state vectors called Long Short Term Memory Models [47]. LSTMs consist of gated units that chose to remember or forget information at each time step. Thus, they are able to maintain states with older information than vanilla RNNs. These methods have been extensively employed in machine translation [45] along with processing other sequential inputs.

In contrast to RNN's which can be viewed as a deep feed forward neural networks with all layers sharing the weights [32], the method proposed in this paper also has sharing of weights, but in the width of the network rather than the depth.

Chapter 3: Proposed Method

3.1 Problem Statement

Given a user U_i , a binary location label $y_i \in \{0,1\}$ where 1 denotes that the user is from a particular city and 0 denotes otherwise, and a collection of unlabeled tweets t_{ij} , j=1,2,...n associated with this user, the task is then to devise two functions, 1) $g(t) \rightarrow y$ that labels a new user, U as belonging to a certain city and 2) $f(t) \rightarrow y$ which labels individual tweets as belonging to the city under consideration by using only the U_i training labels.

3.2 Method Description



Figure 3.1: Multiple Instance Learning for Twitter- This figure shows how the Muliple Instance Learning framework is mapped to the current problem. Each user is a bag and is labeled negative or positive depending on whether they belong to a particular city. There instances are not assumed to belonging to certain classes given user/bag level labels but the aggregation function assumed here is and average.

Hence, as Figure 3.1 describes, a user is a bag labelled with location as being positive to belonging to a city or negative which the tweets form the instances in the multiple instance learning setting which are devoid of labels at the training stage. Thus, an instance level model is designed with trainable neural network components as though labels were available. This model is them simply plugged into the architecture described in Figure 3.2 along with an association function that can aggregate these labels to the bag level for loss discovery and training through gradient descent.

For a treatment of the problem as formulated here, a fully trainable neural network architecture is proposed in this work and is called milNN. The model's architecture is described in Figures 3.2, 3.3 and 3.4.

Instance Level Classifier: (Fig. 3.3) The tweet classifier is an embedding layer followed by a hidden layer that maps to a classifier. The model is regularized using dropout after the hidden layer.

Bag Level classifier: (Fig. 3.4) This instance level classifier is then applied to N tweets and the results are averaged to get the bag level labels. This 'average' can be substituted by any assumption given the complexity of the relationship.

Loss function and Training Method: (Fig. 3.2) The bag level loss is binary cross entropy loss (Equation. 3.1) and it is back-propagated using Adam Optimizer [48] to learn the weights of the network. Thus, the instance level classifier gets trained as a result of the bag level losses being back-propagated.

$$\mathcal{L}(\hat{y}, y) = -\frac{1}{N} \sum_{i}^{N} [y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i)]$$
(3.1)

3.3 Method Characteristics

Due to being fully embedded in the neural network and representation learning paradigm, milNN doesnt have any explicit feature engineering requirements. Hence, it negates the



Figure 3.2: Model Architecture - This general figure provides the framework for multiple instance learning with neural network where instance level models feed their predictions to a bag level aggregation layer and losses are subsequently back-propagated through the network to learn the weights. In effect, all instance level models share weights and are actually a single model being learned.

immeasurably time consuming and subjective involvement of a domain expert. Additionally, it is equipped to handle any changes that might occur organically in the data as it veers away from a rigid logic based paradigm.

As another point of improvement over kernel based methods prevalent in prior research, milNN learns distributed representations which are implicit to a similarity measure in this case. Moreover, it doesn't require high compute and memory restrictions and can be learned using stochastic gradient descent which is easy to parallelize.

Additionally, as described in the problem formulation there are no assumptions on the



Figure 3.3: Instance Level Model - Consists of an embedding layer that learns a vector for each word in the tweet which feeds into a fully connected layer before being classified as belonging to the city or not. Dropout is added to regularize before the sigmoid activation.

membership proportions within the bag. The aggregation assumption of computer the average is also not particularly stringent as the sigmoid layer doesn't provide exact labels for the instances, but rather a probability which is averages across all tweet classifications for a user level label.

The architecture is also flexible and the model described here can be seen as one example of the possibilities. Figure 3.2 describes a high level abstraction of the general idea of weights being shared in a wide sense and all the constituent parts can be replaced to accommodate the needs of the problem. This flexibility can be viewed as a use case level customization as opposed to a dataset level customization which would be a feature engineering approach that milNN avoids.



Figure 3.4: Aggregation Model - This model takes instance level predictions and averages them arithmetically to provide the bag level prediction which can be compared to labels and provide a loss that can be back-propagated through the network.

Chapter 4: Results and Evaluation

The following sections describe the dataset, the experimental setup with the hyperparameters of the model and training details. It then goes on to describe the Accuracy, F-measure and Running time comparisons of the various MIL architectures on the datasets to contrast with milNN. Finally, illustrative examples of the milNNs performance are shown as relating to the most indicative datasets.

4.1 Datasets

The datasets were created by reverse geocoding lat/long information from Twitter North America dataset [49]. These latitude and longitude readings were recorded when the user signed up location and subsequently tweets were recorded for this user. For use in this work, the top cities in the data were split into 15 datasets of equal number of positive and negative samples from the reverse geocoded dataset [49]. The negative samples for each city were randomly selected from the rest of the dataset after stratified sampling from other cities. Table 4.1 describes the number of training and testing instances in the respective datasets. The New York City data is the biggest with 19000 total samples. Chicago, San Francisco, Philadelphia, Washington DC and Toronto have more than 5000 examples and hence the results related to these datasets carry more weight. The rest of the datasets also have at least 2100 total users.

It is evident that there are some differentiating relationships that exist at the bag level at least by observing the results of the Multinomial Naive Bayes Classifier which averages at 70.3% accuracy. The highest accuracies being achieved on San Francisco, New York City and Philadelphia out of the bigger datasets.

It is also interesting to observe the results of a Support Vector Classifier on the dataset



Figure 4.1: Datasets map - This figure indicates the cities under consideration in this work. A color legend is included which will remain constant throughout this work for indicating cities in figures that follow.

to see preliminary existence of nonlinear relationships that might be present in the dataset. San Francisco and New York datasets perform better on this initial litmus test of non linearity in language use. However, the Philadelphia data has a borderline classification accuracy of about 50% using the Support vector classifier.

4.2 Experimental Setup

Instance level model: The tweets are preprocessed by changing URLs, @mentions, and hashtags to a generic word for each. This choice was tested on datasets and while the accuracy was comparable, the instance level results were drastically different. For instance the model for San Francisco predicted 0.99633 for the words "San Francisco" under the

Table 4.1: Datasets - Among the bigger datasets are NYC and SF. Cities like Boston and San Diego rank lower on the amount of tweets found in the dataset. It should be noted that only half of the total are positive examples. The negative examples were compiled randomly after stratified sampling from other cities.

City	Train	Test	Total
Atlanta,GA	3531	883	4414
Austin,TX	2332	583	2915
Baltimore,MD	2159	541	2700
Boston,MA	1911	478	2389
Chicago,IL	6628	1658	8286
New Orleans,LA	2072	520	2592
New York City,NY	15200	3800	19000
Paradise,NV	2475	620	3095
Philadelphia,PA	4633	1159	5792
San Diego,CA	1960	492	2452
San Francisco,CA	6168	1542	7710
Seattle,WA	2680	670	3350
Toronto,Canada	4029	1008	5037
Washington, D.C.	4584	1148	5732
Weehawken,NJ	1756	440	2196

chosen model and 0.06842 under the model containing hashtags and mentions. Thus, this preprocessing method was chosen so as to isolate language features from graph features of the dataset. Subsequently the tweets are tokenized and the top 5000 vocabulary words are henceforth considered in the modelling process. The tweet is then padded to a 20 word maximum (16 being the average for a tweet with 140 characters) and then fed through an embedding layer with 32 dimensions which is randomly initialized. Following this, there is a single hidden layer with 100 nodes that process the various language level relationships and feed the relu activations to the sigmoid layer for classification after adding a dropout of 25% for regularization.

Bag level model: N is set to 10 and the outputs of the instance level models are averaged at a higher layer for the bag level output. At this level binary cross entropy loss is calculated using the bag level labels and back-propagated using the Adam optimizer. A batch size of 256 bags at a time is chosen and trained for 200 epochs with a learning rate of 0.0001. An early stopping condition is included which breaks out of training when the loss of the epoch

Table 4.2: Initial high level results on datasets: As a way to explore the capabilities of the data, basic bag of words classifiers are explored to demonstrate existence of features within the dataset. Multinomial Naive Bayes picks up frequency based capabilities, while Support Vector Classifier tries to model non-linearities at the user level. While these methods show high performance in some cases, they do not train any instance level classifiers.

City	MNB	SVC
Atlanta	0.7022	0.5005
Austin	0.7238	0.5026
Baltimore	0.7338	0.5009
Boston	0.6276	0.5000
Chicago	0.6701	0.5723
New Orleans	0.6731	0.5981
New York City	0.7076	0.5795
Paradise	0.6903	0.5016
Philadelphia	0.7092	0.5038
San Diego	0.6626	0.5671
San Francisco	0.7607	0.5797
Seattle	0.7388	0.6104
Toronto	0.7639	0.5009
Washington, D.C.	0.6760	0.5400
Weehawken	0.7114	0.5636

converges and waits for 5 iterations to confirm the convergence.

4.3 Exploring the Hyper-parameter space

As the hyper-parameter space in neural networks can be large, here we only consider the implications changing of embedding dimensions and number of nodes in the dense layer of the neural network. The batch size, epochs, stopping conditions, inputs, optimizers and general architecture remain the same through these tests. The tests in this section were conducted by extracting a 20% validation set from the training data.

For exploring the model architecture, choices for embedding space and number of dense nodes was considered as shown in figure 4.2. While these choices reacted differently on the 15 datasets, a general conclusion to be made was that the accuracy wasn't particularly sensitive to these choices. However, the running time¹ was affected as the complexity in

¹MacBook Pro (2016)- Processor: 2.9 GHz Intel Core i7 ; Memory: 16 GB 2133 MHz LPDDR3

Table 4.3: Hyper-Parameter Accuracy(dense nodes/embedding dimension): This table makes a case of higher complexity as the accuracy scores are bolder towards bigger embeddings and more nodes in the dense layers

City	50/32	100/10	100/32	100/50	150/32	200/32
Atlanta	0.6239	0.6169	0.6268	0.6197	0.6154	0.6204
Austin	0.6781	0.6642	0.6738	0.6899	0.6770	0.6835
Baltimore	0.6520	0.6206	0.6381	0.5000	0.6427	0.6473
Boston	0.6872	0.6649	0.6715	0.6780	0.6911	0.6793
Chicago	0.6377	0.6494	0.6460	0.6460	0.6411	0.6419
Houston	0.6010	0.5884	0.6225	0.6288	0.6338	0.6376
New Orleans	0.6630	0.6534	0.6667	0.6618	0.6679	0.6775
New York City	0.6956	0.6961	0.7030	0.6933	0.6992	0.6982
Paradise	0.6141	0.6000	0.6222	0.6141	0.6162	0.6172
Philadelphia	0.6625	0.6490	0.6398	0.6523	0.6436	0.6506
San Diego	0.6237	0.6301	0.6531	0.6620	0.6569	0.6441
San Francisco	0.7599	0.7534	0.7591	0.7616	0.7506	0.7612
Seattle	0.7239	0.7099	0.7192	0.7155	0.7155	0.7127
Toronto	0.7398	0.7354	0.7335	0.7522	0.7335	0.7410
Washington, D.C.	0.6599	0.6507	0.6545	0.6572	0.6490	0.6556
Weehawken	0.6538	0.6439	0.6496	0.6524	0.6524	0.6553

the model grew. As expected, the simplest model runs the fastest. However, a trade-off needs to be considered between adding complexity in terms of dense layer nodes (and/or embedding dimensions) and running time add on. As per figure 4.2 it was particularly fruitful to move from 10 embedding dimensions to 32. However, not so much from 32 to 50. Additionally, the dense layer nodes seem to perform close to each other, perhaps due to the regularization mechanism of dropout applied in the model which essentially reduces all these models to an ensemble [50].

Thus, in order to be able to model complexity, the embedding dimension was chosen to be 32, and 100 dense nodes were picked as the model seemed to have considerable slowing down that couldn't compensate with improvement in accuracy when moved to the choices of 50 and 200 respectively.

Table 4.4: Hyper-Parameter F-score(dense nodes/embedding dimension): Even though the model is not particularly sensitive to the hyper-parameter changes here(all f-scores are within a few decimal points of each other), the higher f-scores occur towards the right of the table where the model is more complex

City	50/32	100/10	100/32	100/50	150/32	200/32
Atlanta	0.6242	0.6235	0.6312	0.6324	0.6091	0.6369
Austin	0.6725	0.6652	0.6800	0.6929	0.6753	0.6831
Baltimore	0.6725	0.6554	0.6329	0.6667	0.6419	0.6530
Boston	0.6948	0.6605	0.6693	0.6737	0.6974	0.6910
Chicago	0.6069	0.6197	0.6184	0.6130	0.6188	0.5963
Houston	0.6156	0.6110	0.6340	0.6278	0.6514	0.6479
New Orleans	0.6883	0.6742	0.6920	0.6796	0.6864	0.6942
New York City	0.6827	0.7012	0.6894	0.6824	0.6880	0.6908
Paradise	0.5901	0.5696	0.5978	0.5892	0.6000	0.6097
Philadelphia	0.6712	0.6625	0.6473	0.6646	0.6552	0.6564
San Diego	0.5564	0.5760	0.5988	0.6097	0.6084	0.5951
San Francisco	0.7564	0.7546	0.7556	0.7578	0.7476	0.7603
Seattle	0.7228	0.6989	0.7114	0.7120	0.7064	0.7027
Toronto	0.7380	0.7377	0.7340	0.7545	0.7360	0.7405
Washington, D.C.	0.6276	0.6309	0.6292	0.6275	0.6220	0.6376
Weehawken	0.6313	0.6006	0.6382	0.6422	0.6199	0.6409

Table 4.5: Hyper-Parameter Running Time(dense nodes/embedding dimension) : The simple model runs the fastest.

City	50/32	100/10	100/32	100/50	150/32	200/32
Atlanta	20.9	13.1	27.6	27.4	25.3	26.9
Austin	14.5	9.5	19.0	23.4	20.1	22.7
Baltimore	12.6	7.9	16.3	3.5	19.4	21.3
Boston	11.6	7.0	14.2	19.4	16.8	18.9
Chicago	28.6	22.8	29.8	37.9	37.6	34.6
Houston	11.4	7.4	14.4	19.7	20.4	21.7
NewOrleans	11.8	8.0	15.1	20.8	19.9	24.4
NewYorkCity	47.8	31.0	49.7	463.2	51.1	57.2
Paradise	14.0	10.1	18.1	29.7	20.9	20.5
Philadelphia	25.8	16.2	32.7	40.0	48.0	44.7
SanDiego	11.5	7.3	14.3	20.1	26.2	21.7
SanFrancisco	40.3	23.2	35.2	57.7	52.7	47.2
Seattle	20.0	10.0	19.5	29.4	26.2	30.8
Toronto	28.8	16.9	29.1	47.3	36.9	44.6
WashingtonDC	29.2	16.9	26.2	34.9	33.1	38.4
Weehawken	10.8	7.5	13.0	18.9	15.5	21.0



Figure 4.2: Hyperparameter Comparison - In the above visual, size is the embedding size, color is number of nodes in the dense layer, cities are labeled near the running time vs accuracy point of each run. The running time is on the x-axis and accuracy is on the y-axis. Here the upper left corner of the graph is better where the accuracy is high and running time is low. Another characteristic that is observable is how the results of a particular model might cluster together to demonstrate the robustness of the model given various different datasets. On this count the lower embedding dimension (faster) models don't impress as accuracy remains low on many datasets and results seem more scattered. While the higher embedding dimensions remedy the accuracy, they increase the running time.It is, however, useful to notice that the difference in accuracy is insignificant compared to running times as the models get more complex.

4.4 Performance comparison with MIL methods

To judge the performance of milNN, it is compared with prior work in MIL namely SIL, MISVM and GICF. Here, the metrics of interest - Accuracy and F measure are compared and the running time comparison is given detailed treatment in the next section with this context.

As per the results in Table 4.2 the classifier has a higher accuracy on the bigger San Francisco and New York datasets. It outperforms all other methods on 14 of the 15 datasets considered. For the exception, the Paradise, NV dataset is one of the smaller ones and does better with the GICF approach. However, milNN outperforms GICF in the F score calculation of results (Table 4.3).

milNN outperforms all other methods on 10 out of the 15 datasets under consideration for the F measure. It loses to SIL on the Paradise and Washington DC datasets. Of these, the DC dataset is bigger and on observation the proposed model comes within the third decimal point while using a fraction of the time and still outperforms GICF by a considerable margin, as it does for the Paradise data. The method loses to GICF on the Atlanta, Baltimore and Philadelphia datasets; however, GICF takes longer to run and can't perform as well on the accuracy metric either. Additionally, it is worth noting that F scores are not as consequential as accuracy given the balanced nature of the dataset.

City	MISVM	SIL	GICF	milNN
Atlanta	0.4990	0.5780	0.6025	0.6602
Austin	0.4970	0.6070	0.6501	0.7015
Baltimore	0.4990	0.5180	0.6248	0.6858
Boston	0.5000	0.5460	0.5774	0.6276
Chicago	0.5000	0.5760	0.6429	0.6502
New Orleans	0.5000	0.5230	0.6365	0.6962
New York City	0.5000	0.5740	0.6476	0.7024
Paradise	0.4980	0.6060	0.6629	0.6565
Philadelphia	0.4960	0.5190	0.6195	0.6644
San Diego	0.5000	0.6300	0.6504	0.6850
San Francisco	0.5000	0.6300	0.7322	0.7542
Seattle	0.5000	0.6210	0.6970	0.7269
Toronto	0.4990	0.6390	0.6895	0.7520
Washington, D.C.	0.5000	0.5740	0.6298	0.6437
Weehawken	0.5000	0.6160	0.6727	0.7000

Table 4.6: Accuracy - milNN performs higher 14 of the 15 datasets under consideration.

 Table 4.7: F measure - The proposed model outperforms on 10 of the 15 datasets under consideration.

City	MISVM	SIL	GICF	milNN
Atlanta	0.0000	0.6900	0.6982	0.6568
Austin	0.0000	0.6650	0.6291	0.6848
Baltimore	0.0000	0.6560	0.6957	0.6816
Boston	0.0000	0.5820	0.5960	0.6130
Chicago	0.0000	0.5180	0.5163	0.6420
New Orleans	0.0000	0.6690	0.6976	0.7041
New York City	0.0000	0.6230	0.6349	0.6988
Paradise	0.0000	0.6510	0.6365	0.6468
Philadelphia	0.0000	0.6620	0.7006	0.6830
San Diego	0.0000	0.6080	0.6325	0.6548
San Francisco	0.0000	0.6980	0.7398	0.7603
Seattle	0.0000	0.6840	0.7112	0.7215
Toronto	0.0000	0.6810	0.7056	0.7485
Washington, D.C.	0.0000	0.6190	0.4910	0.6108
Weehawken	0.0000	0.6590	0.6588	0.6827

4.5 Running Time Comparison

milNNs running time comparison provided a 15 out of 15 improvement in comparison to the other methods. This underlines the guaranteed contribution of this work in terms of speed and optimization. Another place the proposed model saves time is in its negation of feature engineering and minimal preprocessing. However, these times can be quite arbitrary based on what feature vector the domain experts (linguists in the current context) decide to come up with and have not been documented here.

The GICF and milNN experiments were run on a consumer laptop² and the MISVM and SIL algorithms required a cluster with higher memory requirements³. Thus, the first two methods have running times that are the upper bounds and would be considerably lower on faster machines used on MISVM and SIL. Conversely, the MISVM and SIL running times can be thought of as being at least this high on consumer computers.

The computational complexity could be observed in action for MISVM and SIL running into hours of running on sampled instances. MISVM, even though theoretically superior for MIL than SIL, learns nothing from the hours of training which is reflected in the 50% accuracy across the datasets and zero F-score despite the longest running times. Perhaps this can be attributed to the strong membership assumptions and OR aggregation function which doesn't apply to the current use case. Due to these results, MISVM is not considered any more comparisons henceforth.

SIL does considerably well as the simplicity generalises to the user tweet relationship of the dataset for geolocation. Even though, SIL finishes training faster than MISVM, the method is still considerably slower that the neural network methods (GICF and milNN). The method proposed in this paper, milNN, provides an average 450x speedup over MISVM, over 70x speedup over SIL, and 4x speedup over GICF. While these numbers provide a high level overview of the applicability of milNN, they are not weighted according to the size of the datasets under consideration. The next two subsections consider this problem by

²MacBook Pro (2016)- Processor: 2.9 GHz Intel Core i7 ; Memory: 16 GB 2133 MHz LPDDR3

³Argo Research Cluster - Processor: 64 Core AMD Opteron; Memory: 512 GB(50 used)

City	MISVM ³	SIL 3	SIL 2	GICF 2	milNN^2
Atlanta	38,799	3,031	-	106	38
Austin	21,304	$3,\!072$	-	81	34
Baltimore	27,069	3,030	-	69	35
Boston	19,166	2,724	-	60	28
Chicago	22,935	$2,\!866$	-	288	45
New Orleans	33,043	$2,\!826$	-	76	52
New York City	10,096	$3,\!448$	over 12 hours	822	85
Paradise	12,891	$2,\!825$	-	149	40
Philadelphia	14,453	$3,\!320$	-	212	52
San Diego	22,547	$3,\!108$	-	67	57
San Francisco	12,562	$3,\!883$	-	222	64
Seattle	17,088	3,738	-	143	41
Toronto	17,861	$3,\!847$	-	131	48
Washington, D.C.	21,026	2,790	-	193	42
Weehawken	9,460	$1,\!894$	-	63	29

Table 4.8: Running Time - Consistent and significantly better performance is demonstrated by the proposed method (milNN)

looking at these run-times and comparing them with metrics by each dataset.

4.5.1 Running time vs Accuracy

While comparing running times to accuracy, in Figures 4.4 and 4.5, it is evident that milNN outperforms all other methods. Even on the datasets where accuracy was lower than others(Paradise), it is noticeable that the difference in accuracy is not as significant and the speed up provided by milNN. Thus the gain in performance outweighs the insignificant loss in accuracy for this particular dataset.

In Figure 4.3, it can be seen that the upper left corner of the graph is the more optimized area where accuracy is high and running times are low.

Another interesting observation to be made in Figure 4.3 is that milNNs accuracy vs running time scores are clustered together suggesting consistency and robustness of the approach when faced with various datasets with different feature sets. Conversely, GICF demonstrates a spread out scatter plot, indicating inconsistent performance across datasets.



Figure 4.3: Accuracy vs Running Time by City - The results of the proposed model (+) plus signs are consistently faster with $GICF(\blacksquare)$ squares being less reliable given diverse requirements of datasets.

4.5.2 Running time vs F measure

4.4 indicates a scatterplot that is particularly interesting as the method couldn't outperform every old method on this metric.

GICF points are towards the right and lower side of the this scatterplot. However, they are both more scattered in terms of results with values below 0.6. This effect can be observed most predominantly the Chicago and Washington DC datasets. This questions the variability in these methods which is most probably an effect of feature engineering requirements of different datasets.

milNN clearly wins as the results for all datasets remain clustered in the top left hand corner of the graph which is the best quadrant to be in with high F scores and low running times. All F scores are also greater than 0.6 and the scatterplot is fairly clustered together which speaks to the robustness of the method.

Thus, after the comparison were made by extensive experiments on 15 datasets, the proposed model demonstrated superior scalability and running time performance while also being robust to the changing needs of various datasets.



Figure 4.4: F measure vs Running Time by City - milNN's results are consistently high even in diverse datasets and the + signs cluster in the best area of the graph where running times are low and metric is high.

4.6 Illustrative Examples



Figure 4.5: Word Cloud for NYC - It can be seen that the the word cloud created has a range of location names indicating the NYC users' proclivity to tweet their location.

The following sections deal with the biggest and highest performing language geolocation datasets - New York City and San Francisco.

The word cloud visual is created using all the test instance tweets that were classified to be over 0.95 by the instance level classifier. Location entities were subsequently extracted from these tweets using StandfordNER [51] and then weighted into a word cloud. High level comparison of the SF word cloud to the NYC word cloud indicates that NYC tweets mention various locations throughout the city more than SF tweets do. At the very least, the models learned by these datasets tend to pick out more place name features for NYC than SF. Hence, one would also expect there to be more language level features to be discovered in the SF dataset given the high accuracy of the method on these tweets.

The structure of the following analysis is to go over examples of users that belong



Figure 4.6: Word Cloud for SF - In contrast, the San Francisco word profile doesn't consist of as many place names as New York City and the model was able to determine higher level language features from this data.

to the city to analyse language patterns discovered by the model. Additionally the antipatterns are explored by analysing tweets from users that do not belong to New York and SF respectively in order to discover how these cities don't tweet. The tweets being classified as mostly certain are colored green (prediction over 0.95), the undecided tweets are colored orange (prediction between 0.4 and 0.6) and those that the models found definitely not belonging to the particular city are colored red (prediction below 0.4). They will be referred to as green, orange and red tweets respectively in the following analysis

4.6.1 New York City

User 1 We can clearly see that the model can figure out place names by itself that relate to a particular location. It found Brooklyn, NYC, New York. In the second case, it actually classified the tweet certainly as from New York because of multiple mentions of the name of the city.

Another indicative example can be seen in the mention of bags being lost by United Airlines. It is noteworthy that during the preprocessing phase, the hashtags and @mentions were abstracted away from the model. This might be indicative of people thanking the airlines for losing their baggage nationwide. Thus the model is not quite certain where the tweet is from.

In this classic case where the bag level label doesn't hold, the user is traveling to Dallas and Chicago. The model caught this fact even though it did not have the label at the tweet level.

User 2 This example is particularly indicative of someone in New York. The model clearly catches the easy location names like Lincoln Tunnel and New York City with near certainty. The more interesting nuances can be noticed in the tweets when the "view is beautiful" (high probability) vs when "the beautiful view overlooks the ocean" (low probability). The model realizes that even though NYC has views they don't overlook the ocean. Another interesting thing to note about this ocean tweet is the mention of the Hilton hotel. In the undecided

New York City - User 1

1.0 - " I'm at Fashion's Night Out NYC (NYC, New York) w/ 158 others http://t.co/r2XDXEp "

0.99496824 - " I'm at Brooklyn Bridge w/ 2 others [pic]: http://t.co/ ry1BAJZo "

0.43607074 - " Piece of crap airline! Thanks for losing my bag! @united @UnitedAirlines #united #unitedairlines "

0.014939007 - "I'm at Sfuzzi (2533 Mckinney Ave, Routh St, Dallas) w/ 7 others http://t.co/BnxbYtSr "

0.0078560486 - " I'm at Public House (400 N State St, at Kinzie, Chicago) w/ 7 others http://t.co/q9tiqW2m "

Figure 4.7: New York Tweets: User 1 - For demonstration of location related entities in high and low probability tweets.

orange tweet we can see how pervasive Hilton hotels are and they might not necessarily be in any one place. However, with the mention of the ocean, the model decreases the probability that the tweet is from New York.

Another cultural indication is seen in the last red tweet which signifies how the activity of standing in line at Walmart is not common to New Yorkers as there are no Walmart stores in city.

User 3 This is an interesting example as the user is not from New York City, but Arlington, NY. The example is particularly illustrative as the user is near New York City and frequently travels there. This is indicated in travel related tweets that are marked green by the model despite the bag level label being not New York.

The model remains undecided about the cinema tweet mentioning The Twilight Saga due to the movie being popular nationwide.

New York City - User 2

0.99559683 - "Beautiful view... good food... great music... romantic husband = perfect evening (@ Le Kaveka Restaurant & Bungalows) http:// t.co/sE4EjSGV '

0.99972242 - " A lovely fall brunch with @Accarrino (@ Anthony David's w/ @accarrino) [pic]: http://t.co/4wzMBut6 "

0.99978334 - " I'm at Lincoln Tunnel (New York City) w/ 3 others http:// t.co/EY5B7SB5"

0.49836314 - "I just became the mayor of Hilton Moorea: Toatea Crepes & Bar on @foursquare! http://t.co/jopZTuuE"

0.2133007 - "Beautiful breakfast overlooking the ocean (@ Hilton Moorea: Arii Vahine Restaurant) [pic]: http://t.co/277WG6QF"

0.11722157 - "Standing in line to return what we bought last night. Efficiency! Walmart is out of cash. Waiting 15 min for refund. http://t.co/ VIXay1dv"

Figure 4.8: New York Tweets : User 2 - For demonstrating nuanced place name recognition with better context.

There are standard location related red examples that were captured at the end.

Not New York City - User 3 (Arlington, NY)

0.99989402 - "Back in New Yawk Citay (@ Grand Central Terminal w/ 28 others) http://t.co/2CBNSMtJ"

0.99894804 - "Back to Vassar on the 2:45 Metro North... Snow fall was pretty while it lasted (@ Grand Central Terminal) [pic]: http://t.co/ MRE8JqMk"

0.50644404 - "@paradisetaylor and I on date night ï\x98\x8a (@ Regal Columbiana Grande Stadium 14 for The Twilight Saga: Breaking Dawn -...)http://t.co/54ywGeq9"

0.018137755 - "I just ousted @aashim_91 as the mayor of College Center -Vassar College on @foursquare! http://t.co/qSIjibJA"

0.3946189 - "I just became the mayor of Matthew's bean on @foursquare! http://t.co/jLE2je5g "

Figure 4.9: Negative Example in New York City: User 3 - This user is from Arlington, NY which is near the city and yet the model can distinguish when the user is tweeting from within the city.

4.6.2 San Francisco

San Francisco was the best performing model and the dataset was bigger indicating the popularity of Twitter in SF. Thus, there were many interesting observations to be made from the illustrative examples that went above and beyond catching direct location mentions in individual tweets. The following examples illustrate, besides the obvious name place related classifications, a proclivity to stick to english grammar even while being restricted to 140 characters by the people in this city.

San Francisco - User 1

0.9999975- " I'm at Alcatraz (Alcatraz Island, San Francisco Bay, San Francisco) w/ 6 others http://t.co/47YWmX9p "

0.9999856 - " I'm at Chinatown Gate (500 Bush St, at Grant Ave, San Francisco) http://t.co/rb49RnFa "

0.5055542 - " @matthewharkin @phillo haha, now I'm worried ")

0.025480814 - " I'm at Tiffany & Co. (210 N Rodeo Dr., Beverly Hills) http:// t.co/YppqS7ix ")

Figure 4.10: Positive Example in San Francisco : User 1 - Place name recognition related example for the SF instance level model

User 1 These are standard location names being caught by the model and are being rightly classified. The undecided tweet looks generic enough to not have any indicative pattern here. Hence, the model looks promising at this base level example.

User 2 However, sure enough, there is a mention of technology being labelled as green for SF by the model. This is a particularly indicative of an SF tweets considering the model did not have the explicit @mentions while prediction.

This users tweets also demonstrate the high level language features the model is capable

San Francisco - User 2

0.99910492 - "Can't wait for @BankSimple, @usbank is such a joke from a technology / ease-of-use perspective."

0.54251802 - "My whole morning **has** been devoted to banking. Not done yet. Living the life."

0.3358801 - "My whole morning **had** been devoted to banking. Not done yet. Living the life."

Figure 4.11: Positive Example in San Francisco : User 2 - Demonstration of higher language level features being learned by the model as it recognizes the better grammar choices in the undecided tweet and increases the probability of belonging to SF.

of dealing with thus filling the gap of all prior bag of words related geolocation research. The general tweet of wasting a morning at the bank is classified as though it could be from anywhere and the model remains undecided. However, it quickly lowers the probability of the tweet belonging to SF as soon as it notices the awkward grammar in the red instance which is only a one character difference from the orange tweet. No locations have been mentioned in either of these tweets. They don't contain any hashtags or mentions either.

User 3 The model would be remiss if it didn't capture the engineering and technology culture in the Bay area particularly as it relates to Apple Inc. This user's green tweets show that "Engineers love free food!" is certainly something that is indicative of San Francisco. The Mac and iOS mentions in the other two tweets also classify high. However, the really interesting tweet is the one that was classified red even though it mentions Apple by name. It is clear that the user is visiting the Cupertino campus of the company as there was a name tag provided. Hence, the model picked up on being away from San Francisco. This is an example of travel related discrepancies caught by the model without any place being mentioned.

San Francisco - User 3

0.97942388 - "Engineers love free food! #IDF2011 http://t.co/rThzEeKO http://t.co/oyYqSmYo"

0.99747145 - "@hashimwaheed the left one is USB serial into Mac, the other is normal iPhone USB into Mac"

0.9944582 - "iOS5 beta expires today! "limited-edition b7b" redsn0w lets you sync data+ pics: OSX http://t.co/EbVEGO0t Win http://t.co/vC5PK2Eg"

0.0067595979 - "@alexeheath my host at Apple surprised me with that visitor's name tag...I had expected it to be my real name :)"

Figure 4.12: Positive Example in San Francisco : User 3 - Technology related jargon is modeled high in SF social media as expected while qualifying these mentions with context.

User 4 This negative example is a clear indication of the SF population sticking to proper grammar while tweeting as every tweet by this person has low probability and not one of them has the correct sentence structure or spelling. The example that is particularly indicative of this is where the user mentions "technical issues" in a badly formed sentence and still doesn't receive any substantial probability for belonging to San Francisco.

Not San Francisco - User 4 (Louisiana, Arabi)

0.16304019 - "Follow the OG triple OG @thad4mayor to ensure that he don't steal ur wallet when he see you in the streets...<>jtfo"

0.23261635 - "@jbdachamp u show me no luv :("

0.0011654327 - "Nap time"

0.011714808 - "somethins gotta give"

0.10918618- "@Cree_Oh_Lay_CO how ya been?"

0.00020607341 - "@jbdachamp and u won't lol",

0.0094076423 - "@jbdachamp I was MIA 4 a min due 2 **technical** issues but now I'm baaaaack lol"

0.010172283 - "da best part is that the downs dont last always"

0.20761815 - "I luv fridays :)"

0.064976566 - "TGIF"

0.073392898 - "@Cree_Oh_Lay_CO we're great :)"

0.18137941 - "@thad4mayor "our" hmmm lol"

Figure 4.13: Negative Example in San Francisco :User 4 - This user belongs to Louisiana and has poor grammar choices in tweets. This is picked up by the model even when the word 'technical' is mentioned which should have rated higher for being from SF.

4.7 Insights from results

Thus, it was discovered after running extensive experiments on datasets from 15 locations that even at it's worst, the proposed architecture proved itself to run faster than the state of the art. It outperformed older methods on most datasets on various metrics and saved time by eliminating feature engineering and kernel design.

Subsequently, it demonstrated ability to capture different features in datasets that had considerably different needs with New York City tweets being location heavy and San Francisco Tweets adhering to strict grammar even within the 140 character limit. The same general model was able to cater to the needs of both datasets and successfully discovered patterns that were indicative of the local language style.

Prior works have handled the missing instance level labels in various ways to evaluate the instance level classification. Such labels might be intuitive where the use case is sentiment analysis [9], but they can be quite subjective for the geolocation modeling case as the architecture doesn't assume a particular characteristic associated with the cities under consideration. For example, if a subset of the data were to be labeled where locations are mentioned by name, it would only test the instance level model on being able to catch that particular pattern and would fail on instances where other characteristics may be present in the data like superior grammar usage. Another method for instance level classification judging is presented in [7], where instances are triaged and classification on a bag level is subsequently evaluated. If the bag level model trains better with the new selected instances. then the instance level model is deemed to have discovered the underlying patterns in the instances. While this approach might suit a paragraph based dataset (news, reviews), it falls short for a collection of tweets as they share little context among each other in the bag. Additionally, the neural network is constructed to model non-linear relationships in the tweet level data which might not be caught at the bag level by a classifier. For example, if this model of evaluation was to be considered and a Mulitnomial Naive Bayes classifier was trained after triaging the tweets, instances that differ by a word would be considered to bring the same value even though one might be grammatically incorrect due to the wrong word placement(eg. fig 4.11).

Chapter 5: Discussion and Conclusion

To help deal with the information deluge caused by content generation on social media platforms, this work focused on geo-locating tweets while working from user(set of tweets) level location labels. This was accomplished by providing a framework consisting of an end to end trainable neural network architecture for Multiple Instance Learning that helped transfer user level location labels to individual instances of tweets. The method was subsequently applied to datasets from 15 cities and was compared with state of the art approaches as well as various architecture designs. The model outperformed earlier methods and proved to be faster across all 15 datasets while eliminating the need for feature engineering by learning representations. This was subsequently demonstrated using illustrative examples where the model could identify place names and grammatical structure from short text.

Since the results heavily depend on the datasets, the models would need to retrain as the social media data formats change. For example, tweets have recently been increased from 140 character to 280. While this change can be accommodated by changing the words per tweet assumption during tokenizing tweets, it would require an updated dataset for training to catch the newer language patterns that might have emerged with the increased character limit. Many other changes have occurred with twitter data since the publication of the dataset in 2012 [49] used to construct the demonstrations in this work, such as urls, photos, videos, quotes and @replies ceasing to be included in the 140 character limit [52].

5.1 Future Work

As the architecture is modular and various pieces can be substituted to cater to particular dataset characteristics, many improvements can be made depending on the use case and instance level relationships in the datasets.

5.1.1 Transfer Learning

Any twitter related embeddings can serve as a starting point for the embedding based classifier [53][39]. It would then be equipped with some idea of what tweets are like and will retrain (or fine-tune) to serve the needs of the particular use case of geo-location.

The transfer learning approach will also generalize to other language based datasets with their particular targets like sentiment analysis.

5.1.2 Multiple Classification

Instead of 15 datasets, if a dataset with multiple locations were to be treated, a simple change from sigmoid to softmax layer classification will equip the model to deal with multiple locations at once. This approach would require some exploration as the model might be confused by similar place names in various cities(like 'Main Streets') or high level language patterns that are pervasive throughout the country. However, it might also be able to contrast language patterns better given exact labels for each instance as opposed to City X/not City X type labels.

5.1.3 Other instance level models

Depending on the complexity of the instance level data, a sequence model might be appropriate for classification given the success of Recurrent Neural Networks with sequential data [32]. This could be used to accommodate a variable number of instances per bag as the current Multi Layer Perceptron model is limited to accepting a fixed number of inputs.

In fact, the instance level classifier can be substituted by any trainable architecture including CNNs, ResNet blocks or other exotic architectures that might emerge, depending on the requirements of the dataset.

5.1.4 Exotic Aggregation Functions

While this work uses a basic average function to aggregate the instance level predictions to the bag level, this relationship might not be as straightforward for other use cases particularly if the instances are not relatively independent like tweets in each bag.

However, in this case it would be important to be careful so as to not interfere with the instance level model extracting information as a non linear relationship would reduce the instance level classifications to a spatial feature set generator.

5.1.5 Other Applications

Given the generality of the architecture, it can be leveraged to solve other Multiple Instance Learning problems in medicine and genetic research. The embeddings model as employed in milNN, would translate easily to any tag related feature set like language words that can be embedded into a space and modeled. Since the model doesn't use any pre-trained embeddings like Word2Vec and starts with randomly initialized embeddings, it intuitively translates to other use cases that don't emerge from natural language applications. The embeddings that would be learned by such a model could be invaluable in exploratory data analysis of entities and their latent space. Bibliography

Bibliography

- [1] Z. Ashktorab, C. D. Brown, M. Nandi, and A. Culotta, "Tweedr: Mining twitter to inform disaster response," in *ISCRAM*, 2014.
- [2] K. Lim and A. Datta, A topological approach for detecting Twitter communities with common interests. Springer Berlin Heidelberg, 2013, pp. 23–43.
- [3] S.-S. Ho, M. Lieberman, P. Wang, and H. Samet, "Mining future spatiotemporal events and their sentiment from online news articles for location-aware recommendation system," in *Proceedings of the First ACM SIGSPATIAL International Workshop on Mobile Geographic Information Systems*, ser. MobiGIS '12. New York, NY, USA: ACM, 2012, pp. 25–32. [Online]. Available: http://doi.acm.org/10.1145/2442810.2442816
- [4] K. W.-T. Leung, D. L. Lee, and W.-C. Lee, "Personalized web search with location preferences," in 2010 IEEE 26th International Conference on Data Engineering (ICDE 2010). IEEE, 2010. [Online]. Available: https://doi.org/10.1109/icde.2010.5447911
- [5] A. Rahimi, T. Cohn, and T. Baldwin, "A neural model for user geolocation and lexical dialectology," *CoRR*, vol. abs/1704.04008, 2017. [Online]. Available: http://arxiv.org/abs/1704.04008
- [6] J. Eisenstein, B. O'Connor, N. A. Smith, and E. P. Xing, "A latent variable model for geographic lexical variation," in *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, ser. EMNLP '10. Stroudsburg, PA, USA: Association for Computational Linguistics, 2010, pp. 1277–1287. [Online]. Available: http://dl.acm.org/citation.cfm?id=1870658.1870782
- [7] Y. Ning, S. Muthiah, H. Rangwala, and N. Ramakrishnan, "Modeling precursors for event forecasting via nested multi-instance learning," in *Proceedings of the 22Nd* ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ser. KDD '16. New York, NY, USA: ACM, 2016, pp. 1095–1104. [Online]. Available: http://doi.acm.org/10.1145/2939672.2939802
- [8] D. Kotzias, M. Denil, P. Blunsom, and N. de Freitas, "Deep multiinstance transfer learning," *CoRR*, vol. abs/1411.3128, 2014. [Online]. Available: http://arxiv.org/abs/1411.3128
- D. Kotzias, M. Denil, N. de Freitas, and P. Smyth, "From group to individual labels using deep features," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '15. New York, NY, USA: ACM, 2015, pp. 597–606. [Online]. Available: http://doi.acm.org/10.1145/2783258.2783380

- [10] A. Rahimi, T. Baldwin, and T. Cohn, "Continuous representation of location for geolocation and lexical dialectology using mixture density networks," *CoRR*, vol. abs/1708.04358, 2017. [Online]. Available: http://arxiv.org/abs/1708.04358
- [11] S. Andrews, T. Hofmann, and I. Tsochantaridis, "Multiple instance learning with generalized support vector machines," in *Eighteenth National Conference on Artificial Intelligence*. Menlo Park, CA, USA: American Association for Artificial Intelligence, 2002, pp. 943–944. [Online]. Available: http://dl.acm.org/citation.cfm?id=777092.777234
- [12] O. Maron and T. Lozano-Pérez, "A framework for multiple-instance learning," in Proceedings of the 1997 Conference on Advances in Neural Information Processing Systems 10, ser. NIPS '97. Cambridge, MA, USA: MIT Press, 1998, pp. 570–576.
 [Online]. Available: http://dl.acm.org/citation.cfm?id=302528.302753
- [13] D. Kifer, "Attacks on privacy and definetti's theorem," in Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data, ser. SIGMOD '09. New York, NY, USA: ACM, 2009, pp. 127–138. [Online]. Available: http://doi.acm.org/10.1145/1559845.1559861
- [14] H. Kuck and N. de Freitas, "Learning about individuals from group statistics," arXiv preprint arXiv:1207.1393, 2012.
- [15] T. G. Dietterich, R. H. Lathrop, and T. Lozano-Perez, "Solving the multiple instance problem with axis-parallel rectangles," *Artificial Intelligence*, vol. 89, no. 1-2, pp. 31–71, 1997. [Online]. Available: http://lis.csail.mit.edu/pubs/tlp/multiple-instanceaij.pdf
- [16] X. Xu and E. Frank, "Logistic regression and boosting for labeled bags of instances," in Proc. of the PacificAsia Conf. on Knowledge Discovery and Data Mining. Springer-Verlag, 2004, pp. 272–281.
- [17] Z.-H. Zhou and J.-M. Xu, "On the relation between multi-instance learning and semi-supervised learning," in *Proceedings of the 24th International Conference on Machine Learning*, ser. ICML '07. New York, NY, USA: ACM, 2007, pp. 1167–1174.
 [Online]. Available: http://doi.acm.org/10.1145/1273496.1273643
- [18] N. Weidmann, E. Frank, and B. Pfahringer, "A two-level learning method for generalized multi-instance problems," in *In Proceedings of the Fourteenth European Conference* on Machine Learning, 2003.
- [19] G. Liu, J. Wu, and Z.-H. Zhou, "Key instance detection in multi-instance learning," in *Proceedings of the Asian Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, S. C. H. Hoi and W. Buntine, Eds., vol. 25. Singapore Management University, Singapore: PMLR, 04–06 Nov 2012, pp. 253–268. [Online]. Available: http://proceedings.mlr.press/v25/liu12b.html
- [20] S. Ray and M. Craven, "Supervised versus multiple instance learning: An empirical comparison," in *Proceedings of the 22Nd International Conference on Machine Learning*, ser. ICML '05. New York, NY, USA: ACM, 2005, pp. 697–704. [Online]. Available: http://doi.acm.org/10.1145/1102351.1102439

- [21] T. Grtner, P. A. Flach, A. Kowalczyk, and A. J. Smola, "Multi-instance kernels," in In Proc. 19th International Conf. on Machine Learning. Morgan Kaufmann, 2002, pp. 179–186.
- [22] Y.-F. L. J.-H. Hu and Y. J. Z.-H. Zhou, "Towards discovering what patterns trigger what labels," 2012.
- [23] Y.-F. Li, J. T. Kwok, I. W. Tsang, and Z.-H. Zhou, A Convex Method for Locating Regions of Interest with Multi-instance Learning. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 15–30.
- [24] R. C. Bunescu and R. J. Mooney, "Multiple instance learning for sparse positive bags," in *Proceedings of the 24th International Conference on Machine Learning*, ser. ICML '07. New York, NY, USA: ACM, 2007, pp. 105–112. [Online]. Available: http://doi.acm.org/10.1145/1273496.1273510
- [25] O. Chapelle and A. Zien, "Semi-supervised classification by low density separation," 01 2005.
- [26] X. Zheng, J. Han, and A. Sun, "A survey of location prediction on twitter," CoRR, vol. abs/1705.03172, 2017. [Online]. Available: http://arxiv.org/abs/1705.03172
- [27] F. Melo and B. Martins, "Automated geocoding of textual documents: A survey of current approaches," *Transactions in GIS*, vol. 21, no. 1, pp. 3–38, jun 2016. [Online]. Available: https://doi.org/10.1111/tgis.12212
- [28] A. Woodruff and C. Plaunt, "Gipsy: Geo-referenced information processing system," vol. 45, 05 1994.
- [29] E. Amitay, N. Har'El, R. Sivan, and A. Soffer, "Web-a-where: Geotagging web content," in *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '04. New York, NY, USA: ACM, 2004, pp. 273–280. [Online]. Available: http://doi.acm.org/10.1145/1008992.1009040
- [30] Z. Cheng, J. Caverlee, and K. Lee, "You are where you tweet: a content-based approach to geo-locating twitter users," in *CIKM*, 2010.
- [31] J. Liu and D. Inkpen, "Estimating user location in social media with stacked denoising auto-encoders," in VS@HLT-NAACL, 2015.
- [32] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," vol. 521, pp. 436–44, 05 2015.
- [33] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *nature*, vol. 323, no. 6088, pp. 533–536, 1986.
- [34] V. Kůrková, "Kolmogorov's theorem and multilayer neural networks," Neural Networks, vol. 5, no. 3, pp. 501–506, Mar. 1992. [Online]. Available: http://dx.doi.org/10.1016/0893-6080(92)90012-8

- [35] Y. LeCun, B. E. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. E. Hubbard, and L. D. Jackel, "Handwritten digit recognition with a back-propagation network," in Advances in neural information processing systems, 1990, pp. 396–404.
- [36] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [37] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," arXiv preprint arXiv:1612.08083, 2016.
- [38] Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin, "A neural probabilistic language model," *Journal of machine learning research*, vol. 3, no. Feb, pp. 1137–1155, 2003.
- [39] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," *CoRR*, vol. abs/1310.4546, 2013. [Online]. Available: http://arxiv.org/abs/1310.4546
- [40] R. Socher, C. C. Lin, C. Manning, and A. Y. Ng, "Parsing natural scenes and natural language with recursive neural networks," in *Proceedings of the 28th international* conference on machine learning (ICML-11), 2011, pp. 129–136.
- [41] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in Advances in neural information processing systems, 2014, pp. 3104–3112.
- [42] 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, 2013, pp. 3111–3119.
- [43] H. Schwenk, "Continuous space language models," Computer Speech & Language, vol. 21, no. 3, pp. 492–518, 2007.
- [44] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa, "Natural language processing (almost) from scratch," *Journal of Machine Learning Research*, vol. 12, no. Aug, pp. 2493–2537, 2011.
- [45] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," arXiv preprint arXiv:1409.0473, 2014.
- [46] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," arXiv preprint arXiv:1406.1078, 2014.
- [47] 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.
- [48] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," CoRR, vol. abs/1412.6980, 2014. [Online]. Available: http://arxiv.org/abs/1412.6980

- [49] S. Roller, M. Speriosu, S. Rallapalli, B. Wing, and J. Baldridge, "Supervised text-based geolocation using language models on an adaptive grid," in *Proceedings of* the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, ser. EMNLP-CoNLL '12. Stroudsburg, PA, USA: Association for Computational Linguistics, 2012, pp. 1500–1510. [Online]. Available: http://dl.acm.org/citation.cfm?id=2390948.2391120
- [50] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014. [Online]. Available: http://jmlr.org/papers/v15/srivastava14a.html
- [51] J. R. Finkel, T. Grenager, and C. Manning, "Incorporating non-local information into information extraction systems by gibbs sampling," in *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, ser. ACL '05. Stroudsburg, PA, USA: Association for Computational Linguistics, 2005, pp. 363–370. [Online]. Available: https://doi.org/10.3115/1219840.1219885
- [52] A. Pardes, "A brief history of the ever-expanding tweet," Sep 2017. [Online]. Available: https://www.wired.com/story/a-brief-history-of-the-ever-expanding-tweet/
- [53] F. Godin, B. Vandersmissen, W. De Neve, and R. Van de Walle, "Multimedia lab@ acl w-nut ner shared task: named entity recognition for twitter microposts using distributed word representations," ACL-IJCNLP, vol. 2015, pp. 146–153, 2015.