MEDINFO 2021: One World, One Health – Global Partnership for Digital Innovation
P. Otero et al. (Eds.)

doi:10.3233/SHTI220149

Understanding Public Perceptions of Measles from Twitter Using Multi-Task Convolutional Neural Networks

Samuel Wang[a*], Jingcheng Du[a*], Lu Tang[b], Cui Tao [a]

[a] School of Biomedical Informatics, The University of Texas Health Science Center at Houston, Houston, TX, USA
[b] Department of Communication, Texas A&M University, College station, TX, USA
** co-first authors*

Abstract

Measles is a highly contagious cause of febrile illness typically seen in young children. Recent years have witnessed the resurgence of measles cases in the United States. Prompt understanding of public perceptions of measles will allow public health agencies to respond appropriately promptly. We proposed a multi-task Convolutional Neural Network (MT-CNN) model to classify measles-related tweets in terms of three characteristics: Type of Message (6 subclasses), Emotion Expressed (6 subclasses), and Attitude towards Vaccination (3 subclasses). A gold standard corpus that contains 2,997 tweets with annotation in these dimensions was manually curated. A variety of conventional machine learning and deep learning models were evaluated as baseline models. The MT-CNN model performed better than other baseline conventional machine learning and the signal-task CNN models, and was then applied to predict unlabeled measles-related Twitter discussions that were crawled from 2007 to 2019, and the trends of public perceptions were analyzed along three dimensions.

Keywords:

measles, deep learning, natural language processing, social media, public opinion

Introduction

Measles is a highly contagious virus and remains a leading cause of vaccine-preventable illness, claiming around 100,000 lives each year worldwide.[1] Measles was once declared eliminated in the United States in 2000 thanks to a highly effective vaccination program. However, the U.S. and other developed countries have experienced a resurgence of measles outbreaks in the last few years due to a decline in vaccination rate.[2] The recent surge in measles cases prompted the World Health Organization to declare vaccine hesitancy as one of the ten most significant threats to global health.[3]

Causes of vaccine refusal and delay are complicated and vary from individual to individual, including fear of adverse reactions, concerns over efficacy, etc. The rise of the Internet and social media allows for the rapid dissemination of anti-vaccine information, which creates more barriers in vaccine promotion. A timely understanding of public perceptions (e.g., what they are discussing related to measles, how they feel, and what they think about the measles vaccine) will allow public health agencies to respond appropriately and timely. Traditional surveys are useful in understanding public perceptions; however, they can be costly and time-consuming.[4] Due to the ubiquitous presence of the Internet, mobile devices, and advances in machine learning and natural language processing (NLP), social media has increasingly become an important venue to understand health-related public perceptions.

Twitter is a promising resource for the analysis of health-related public perceptions due to its public data, popularity, and convenient data access [5,6]. However, some distinctive characteristics of Twitter messages make them difficult to analyze, such as short length of text, frequent use of informal languages (e.g., hashtags, abbreviations). These create obstacles for NLP approaches and typically require substantial time and effort on feature engineering for conventional machine learning algorithms to process the embedded information to further analyses. Deep learning recently achieved state-of-the-art performance in many NLP tasks,[7] including the analysis of social media contents (e.g., Twitter).[8–10] Compared to conventional machine learning algorithms, deep learning can save great effort on features engineering.[7]

Multi-task (MT) learning is a machine learning strategy that aims to improve performance on specific tasks by learning shared-representations across related tasks.[11] Understanding of the public perceptions of measles along multiple dimensions requires learning from multiple similar tasks simultaneously, which provides a good opportunity for MT learning. MT can improve the performance of the deep neural network by reducing the risk of overfitting on a specific task.[11] Using parameters sharing, MT forces the deep neural network to find representations that are generalizable to capture all of the related tasks, which makes it less likely to overfit on a specific task.

The purpose of this study is to propose an MT deep learning framework to map Twitter contents related measles to multiple dimensions. The contributions of this study are three folds: 1) a manually annotated measles Twitter dataset in three dimensions, including Emotions Expressed, Attitudes towards Vaccination, and Type of Message; 2) an evaluation of MT deep learning and conventional machine learning and deep learning algorithms on the dataset; 3) the analysis of trends of public perceptions of measles from Twitter using additional crawled data.

Methods

We collected two corpora of measles-related tweets: one corpus purchased from DiscoverText.com, which was used to develop an annotated Twitter corpus for model training and evaluation, and the other corpus was crawled using TweetScraper.[12] The second corpus was used to analyze the trends of public perceptions of measles.

Twitter Data Annotation Scheme

An annotation scheme containing three dimensions was developed, including Attitude towards Vaccination, Emotion Expressed, and Type of Message. The Emotion Expressed and Type of Message were adapted based on Chew and Eysenbach [13,14] and the Attitude towards Vaccination was created inductively. Dimensions and sub-classes are provided in Table 1. Attitude Towards Vaccination has three dimensions: pro(pro-vaccination), against(anti-vaccination), and not applicable. Type of Message has six categories: news update, resource, personal experience, personal opinion and interest, question, and others. Emotion Expressed also contained six categories: alarm/concern, reassurance, anger, humor/sarcasm, neutral, and not applicable.

Twitter corpus collection and annotation

We used the keyword "measles" to collect two corpuses of English Twitter text. The first corpus was collected from DiscoverText.com from December 1, 2014 to February 2, 2015 and included 1,154,156 tweets. The gold standard corpus (tweets with the annotation), consisting of 2,997 tweets, were chosen using a systematic sampling method from this dataset. The tweets were annotated along these three dimensions: Emotion Expressed, Attitude towards Vaccination, and Type of Message, respectively. Two independent coders annotated these three variables with satisfactory intercoder reliability (Cohen's Kappa: .78, .72, and .80, respectively). We further crawled Twitter for measles-related content using TweetScraper[12]. A total of 1,917,032 tweets posted between Jan. 18, 2007 and Apr. 23,2019 were collected. This unlabeled corpus was then used for prediction and trends analysis purposes.

Multi-task Convolutional Neural Network (MT-CNN)

We proposed a multi-task convolutional neural network (MT-CNN) to classify measles-related tweets along three dimensions. MT learning has achieved successes in many machine learning applications. Compared to single-task learning, MT typically optimizes more than one loss functions. According to Ruder[15], there are two major types of MT learning for deep neural networks, including 1) soft parameter sharing, where each task has its own model, and the distance between the parameters of the model is regularized to be similar; and 2) hard parameter sharing, in which lower hidden layers are shared across tasks and each task has its own task-specific layers.

In this study, we adopted the hard parameter sharing, which is the most commonly used approach in MT deep learning. The architecture of the proposed framework can be seen in Figure 1. The MT-CNN adopts convolutional neural networks as lower layers, which have been widely used for text classification tasks and have achieved superior performance compared to conventional machine learning algorithms.[16,17] All three tasks share the same weights in the convolutional neural networks, including embedding layer, convolutional layer, and max-pooling layers. On top of the max-pooling layers, each task has its own task-specific fully connected layers. For one particular target task, the other two tasks were served as auxiliary tasks and the training was achieved by minimizing the weighted sum of losses from every task. The weights for loss were hyper-parameters for each target task and were searched on the validation set of the target task. Al shared layers and task-specific layers were trained and finetuned during the training process.

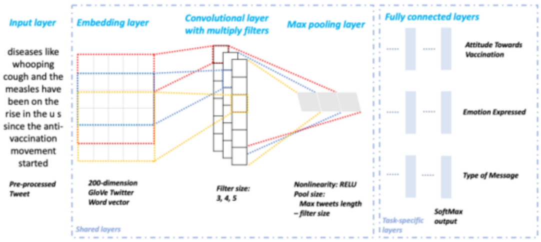

Figure 1. The architecture of MT-CNN model

Experiments and Evaluation

A pretrained GloVe Twitter embedding from Stanford was used to initialize the weights in the embedding layer.[18] As the embedding in dimension 200 achieved the best performance on our validation set, we chose dimension 200 as the size of embedding. We applied padding to the tweets and the max length of tweets was set at 60. All the tweets with fewer than 60 tokens were padded to 60 tokens with zeros. The CNN has three convolutional layers with filter sizes 3, 4, and 5. Each filter size has 128 filters. The layers are then flattened, and a dropout is applied at the rate of 0.5. The algorithms were implemented using Keras framework, with Tensorflow as the core. Accuracy, precision, recall, and F-1 Score were calculated for these algorithms. For comparison, we further evaluated the CNN model with single task setting.

Six competitive conventional machine learning models were evaluated as baseline models, including support vector machines (SVM), Extra Trees (ET), Random Forest, Logistic Regression, AdaBoost, and GradientBoosting. We used the term frequency-inverse document frequency vectorizer (TF-IDF) as the feature. In our pre-experiment, other commonly used text classification features, including bag-of-words (BOW) and mean embedding (averaged word vectors for all words in a text) feature, were evaluated. TFIDF provided a relatively strong and robust performance on three tasks. Python scikit-learn package was used to implement these algorithms.

The gold standard corpus was divided into training, validation, and testing sets with a proportion of 7:1:2. The models were then trained using the training dataset, tuned using the hyperparameters on the validation set, and their performance was evaluated on the testing set. Considering random effects on the training of machine learning and deep learning models, for each model and setting, we run the model 5 times and report the average metrics.

Trends of measles public perceptions

The best-performing model (i.e., MT-CNN) was used to predict the unlabeled data in our crawled Twitter corpus. The MT-CNN model was repeated five times, and the final predictions for the unlabeled tweets were based on the majority voting of the five predictions. The prevalence of each class was then calculated for each of the three dimensions (Emotion Expressed, Attitude towards Vaccination, and Type of Message) by taking the ratio of the number of tweets classified to that class to the total number of tweets, using month as time week. This prevalence for each week was calculated from 2007 to 2019.

Dimension	Classes	Definition	Example
Attitude towards Vaccination	Pro(pro-vaccination)	The tweet supports vaccination	Vaccinating your kids is the best way to prevent the spread of measles!
	Against(anti-vaccination)	The tweet is against vaccination	It turns out that the measles outbreak was actually caused by measles vaccine and not the unvaccinated kids.
	Not Applicable	The tweet does not express any opinion about vaccination.	We have seen a spike for measle cases over the past couple of months.
Type of Message	News/Update	The tweet contains measles-related news and updates of the outbreak.	Measles outbreak in Samoa have killed 50 children.
	Resource	The tweet provides information that helps people understand the nature of measles, how to prevent it or available treatment methods.	Measles does long-term damage to your immune system. The best way to fight measles is MMR vaccination.
	Personal Experience and interest	Twitter user mentions a direct (personal) or indirect (e.g., friend, family, co-worker) experience with measles or the social/economic effects of measles on them.	I'm stuck at home with the measles.
	Personal Opinion	Twitter user posts their opinion of the measles or expresses a need for or discovery of information. General measles chatter or commentary, including jokes and parodies.	Unvaccinated people are the reason for the Disneyland measles outbreak.
	Question	The tweet is a genuine question about measles with no answer.	Should we allow unvaccinated kids to go to school after measles outbreak?
	Others	The tweet is in a foreign language, contains only a URL, is unrelated to measles, or is an advertisement.	The stars sure are bright tonight.
Emotion Expressed	Alarm/Concern	Tweet expresses measles-related fear, anxiety, worry, or sadness for self or others.	Vaccinating your kids is the best way to prevent the spread of measles!
	Reassurance	**Relief:** Tweet expresses joy, happiness, or a sense of peace; **Downplayed risk**: Tweet attempts to de-emphasize the potential risk of measles or bring it into perspective. May also express a lack of concern or disinterest.	It turns out that the measles outbreak was actually caused by measles vaccine and not the unvaccinated kids.
	Anger	Tweet expresses anger, annoyance, scorn, or volatile contempt.	We have seen a spike for measle cases over the past couple of months.
	Humor/Sarcasm	The tweet is comedic or sarcastic.	Only anti-vaxxers will get this: measles
	Neutral	The message does not have an obvious tone.	Measles does long-term damage to your immune system. The best way to fight measles is MMR vaccination.
	Not Applicable	When the tweet is in a foreign language, contains only a URL or is unrelated to measles.	I'm stuck at home with the measles.

Table 1. The architecture of MT-CNN model

Results

MT-CNN Performance

The comparison of different classification algorithms is shown in Table 2. As shown, the MT-CNN model achieved the best accuracy in the classifications of Emotion Expressed and Type of Message, with an accuracy of 0.6992 and 0.7353 respectively, and second-best accuracy in classifying Attitude towards Vaccination with an accuracy of 0.8057. Furthermore, the MT-CNN model achieved higher accuracy than the single-task CNN model on all tasks. The specific metrics for each class for the three dimensions for the MT-CNN model are shown in Table 3. The MT-CNN achieved relative higher scores on the classes of Attitude towards Vaccine than that of the other dimensions, with F-score scores ranging from 0.57 to 0.85. For some classes, the scores can be very low (e.g., F-score of Question in Type of Message is 0), which could largely be due to insufficient sample sizes in our corpus.

Public Perception Analysis

From the crawled tweets dataset from 2007 to 2019, there was a very high peak in February 2015 due to the measles outbreak in Disneyland in Orange County, California.[19] Another peak was found in February 2019, when there were 50 confirmed cases of measles identified in Clark County, Washington.[20] The trends (i.e., weekly prevalence) fluctuated significantly in all three dimensions (Attitude towards Vaccination, Emotion Expressed, and Type of Message). In terms of Attitude towards Vaccination, we observed a slight increase of Pro-vaccination tweets. There was a sharp increase in Anti-vaccination tweets near the end of 2014 while Pro-vaccination tweets dropped significantly during the same time. In general, there have always been more Pro-vaccination tweets than Anti-vaccination tweets. Among Emotion Expressed, Anger has been on the rise since 2016 and 2017. Alarm/Concern has remained relatively stable since 2007, though there was a significant spike near the end of 2014 and early 2015, which could be due to the measles outbreak at Disneyland, California.[19] The prevalence of Humor/Sarcasm dipped to a near all-time low at the same time the prevalence of Anger has increased to a near all-time high. Among the Types of messages, the prevalence of resources has been declining since a sudden spike near the end of 2016. Generally, when there is a decreasing trend in the prevalence of resources, there is an opposite positive trend of news updates. In September 2010, there was a spike for Pro-vaccination tweets, which coincided with a measles outbreak in Zimbabwe that claimed several children in the country.[21] During the same time, many tweets report declining measles vaccination rates, despite increased overall vaccination, which could contribute to the overall increase in Pro-vaccination sentiment during that time.

Similarly, in July 2012, there was an increase of Humor/sarcasm in tweets, with one particular tweet, comparing love to measles, being repeatedly retweeted on the site, which may have contributed to the spike.

	SVM	ExtraTrees	Random Forest	Logistic Regression	AdaBoost	Gradient Boosting	CNN	MT-CNN Model
Attitude Towards Vaccination	0.7873	**0.8107**	0.7889	0.7705	0.7889	0.7923	0.8023	0.8057
Emotion Expressed	0.6951	0.6901	0.6801	0.6600	0.6767	0.6884	0.6908	**0.6992**
Type of Message	0.7052	0.6918	0.6868	0.6951	0.7102	0.6968	0.7313	**0.7353**

Table 2. Performance comparison (measured in accuracy) of machine learning and deep learning algorithms

		Precision	Recall	F-Score	Count
Attitude Towards Vaccination	Pro-vaccination	0.7623	0.7525	0.7570	1077(35.94%)
	Against	0.8424	0.4353	0.5732	186(6.21%)
	Not Applicable	0.8311	0.8776	0.8536	1734(57.86%)
Emotion Expressed	Alarm/Concern	0.5862	0.3135	0.4066	400(13.35%)
	Reassurance	0.4467	0.0636	0.1097	125(4.17%)
	Anger	0.5754	0.3423	0.4236	258(8.61%)
	Humor/Sarcasm	0.6917	0.5333	0.6009	399(13.31%)
	Neutral	0.7238	0.9249	0.8117	1778(59.33%)
	Not Applicable	0.4667	0.1500	0.2255	37(1.23%)
Type of Message	News Update	0.8171	0.8532	0.8342	1242(41.44%)
	Resource	0.6019	0.6202	0.6088	583(19.45%)
	Personal Experience and Interest	0.7000	0.1043	0.1812	76(2.54%)
	Personal Opinion	0.7174	0.7943	0.7535	1008(33.63%)
	Question	0.0000	0.0000	0.0000	49(1.63%)
	Others	0.5667	0.0800	0.1368	39(1.30%)

Table 3. Detailed metrics of each class from three dimensions for MT-CNN model

Discussion

This study demonstrates the superiority of deep learning models over conventional machine learning algorithms in classifying measles-related tweets. The use of multi-task setting for deep learning model can further improve the performance by learning more generalizable features from multiple tasks. Our best performing algorithm by far is the MT-CNN model, combining aspects of deep learning with multi-tasking. This model is superior to the CNN model on all three tasks. Multi-task learning models can help alleviate many of the shortcomings of other machine and deep learning models, especially overfitting. The superiority of the MT-CNN model can further advance performance on other Twitter-classification tasks. We also find that conventional machine learning algorithms provide strong baseline performances on many tasks. These algorithms can still be good options for Twitter text classification tasks.

The analysis of trends based on the three dimensions we predicted (Attitude towards Vaccination, Emotion Expressed, and Type of Message) can provide a better understanding of public perception towards the measles on Twitter and how this perception changes over time. These trends can also help us better understand the effectiveness of vaccination promotion strategies. Understanding trends over time on platforms like Twitter would allow public health professions to improve vaccination promotion programs.

However, there are still limitations to our study. The first limitation is the assumption that all predicted labels are true labels. Because it is almost impossible for machine learning model to achieve 100% accuracy in prediction, some of the predicted labels may be inaccurate, which could lead to information bias due to the misclassification rates. The second limitation is the relatively small size of our gold standard corpus. The 2,997 tweets we trained our models on may not be representative of the nearly 2 million unlabeled tweets we eventually used to identify trends. Though it would take more manual labeling and significantly more processing time, training the models on more tweets could prevent overfitting and be more representative of the tweets eventually labeled. The potential shift in data distribution between labeled and unlabeled tweets would lead to bias to prediction as well. As future work, we will evaluate training strategies such as domain-adversarial training to reduce such bias.[22]

The classification accuracy of deep learning models on our datasets can be further improved. This study specifically looked at the CNN deep learning model, mainly due to its generalizability, robust performance and fast training time (i.e., fewer trainable parameters compared with other deep learning models). Future studies could evaluate more state-of-the-art algorithms on relevant tasks. The multi-task learning framework is generalizable and can be easily adopted for other deep learning models (e.g., replace the shared layers in Figure 1 with lower layers of other deep learning models).

Conclusions

In this study, we developed machine learning and deep learning models to automatically label measles-related tweets to understand public perceptions. We found that deep learning models performed better than conventional machine learning algorithms, and multi-task learning further improved the accuracy of deep learning algorithms in our datasets. Time series analysis MT-CNN model labeled tweets indicated the evolving trends in public perception towards measles disease and the measles vaccine. This study could allow for the design of future vaccine promotion strategies by revealing how populations change their views over time. The development of an effective multi-task learning model can also be applied to other public health problems if similar tasks exist for a target task, opening avenues for using accurate predictions to adjust current programs to target the public more.

Acknowledgements

SW, JD and CT have received funding support from National Institutes of Health under Award Numbers R01LM011829 and R01AI130460.

References

[1] Vaccine Knowledge, Measles | Vaccine Knowledge, (2019) Measles. https://vk.ovg.ox.ac.uk/vk/measles (accessed May 11, 2020).

[2] Decline in measles vaccination is causing a preventable global resurgence of the disease | National Institutes of Health (NIH). https://www.nih.gov/news-events/news-releases/decline-measles-vaccination-causing-preventable-global-resurgence-disease (accessed February 23, 2020).

[3] World Health Organization, Ten threats to global health in 2019. https://www.who.int/emergencies/ten-threats-to-global-health-in-2019, World Heal. Organ. (2019) 1–18. https://www.who.int/news-room/spotlight/ten-threats-to-global-health-in-2019%0Ahttps://www.who.int/news-room/spotlight/ten-threats-to-global-health-in-2019%0Ahttps://www.who.int/news-room/feature-stories/ten-threats-to-global-health-in-2019%0Ahttps://www. (accessed May 11, 2020).

[4] M. Dredze, D.A. Broniatowski, M.C. Smith, K.M. Hilyard, Understanding Vaccine Refusal: Why We Need Social Media Now, Am. J. Prev. Med. 50 (2016) 550–552. https://doi.org/10.1016/j.amepre.2015.10.002.

[5] J. Du, R.M. Cunningham, Y. Xiang, F. Li, Y. Jia, J.A. Boom, S. Myneni, J. Bian, C. Luo, Y. Chen, C. Tao, Leveraging deep learning to understand health beliefs about the Human Papillomavirus Vaccine from social media, Npj Digit. Med. 2 (2019) 27. https://doi.org/10.1038/s41746-019-0102-4.

[6] L. Sinnenberg, A.M. Buttenheim, K. Padrez, C. Mancheno, L. Ungar, R.M. Merchant, Twitter as a tool for health research: A systematic review, Am. J. Public Health. 107 (2017) e1–e8. https://doi.org/10.2105/AJPH.2016.303512.

[7] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature. 521 (2015) 436–444.

[8] P. Nakov, A. Ritter, S. Rosenthal, F. Sebastiani, V. Stoyanov, SemEval-2016 task 4: Sentiment analysis in Twitter, in: Proc. SemEval, 2016: pp. 1–18.

[9] S. Mohammad, F. Bravo-Marquez, M. Salameh, S. Kiritchenko, SemEval-2018 Task 1: Affect in Tweets, in: Proc. 12th Int. Work. Semant. Eval., 2018: pp. 1–17. https://doi.org/10.18653/v1/s18-1001.

[10] S. Rosenthal, N. Farra, P. Nakov, SemEval-2017 task 4: Sentiment analysis in twitter, in: ArXiv, 2019: pp. 1–18. https://doi.org/10.18653/v1/s17-2088.

[11] S. Ruder, An overview of multi-task learning in deep neural networks, ArXiv Prepr. ArXiv1706.05098. (2017).

[12] TweetScraper. https://github.com/jonbakerfish/TweetScraper.

[13] C. Chew, G. Eysenbach, Pandemics in the age of Twitter: Content analysis of tweets during the 2009 H1N1 outbreak, PLoS One. 5 (2010) 361–367. https://doi.org/10.1371/journal.pone.0014118.

[14] J. Du, L. Tang, Y. Xiang, D. Zhi, J. Xu, H.Y. Song, C. Tao, Public perception analysis of tweets during the 2015 measles outbreak: Comparative study using convolutional neural network models, J. Med. Internet Res. 20 (2018). https://doi.org/10.2196/jmir.9413.

[15] S. Ruder, An Overview of Multi-Task Learning in Deep Neural Networks∗, ArXiv. (2017). http://arxiv.org/abs/1706.05098.

[16] Z. Jianqiang, G. Xiaolin, Z. Xuejun, Deep Convolution Neural Networks for Twitter Sentiment Analysis, IEEE Access. 6 (2018) 23253–23260. https://doi.org/10.1109/ACCESS.2017.2776930.

[17] Y. Kim, Convolutional neural networks for sentence classification, EMNLP 2014 - 2014 Conf. Empir. Methods Nat. Lang. Process. Proc. Conf. (2014) 1746–1751. https://doi.org/10.3115/v1/d14-1181.

[18] J. Pennington, R. Socher, C.D. Manning, GloVe: Global vectors for word representation, EMNLP 2014 - 2014 Conf. Empir. Methods Nat. Lang. Process. Proc. Conf. (2014) 1532–1543. https://doi.org/10.3115/v1/d14-1162.

[19] California Department of Public Health, California Measles Surveillance Update, (2015) 1–2. https://www.cdph.ca.gov/Programs/CID/DCDC/CDPH Document Library/Immunization/IMM-MeaslesUpdate2015-04-17.pdf (accessed May 30, 2020).

[20] Measles Outbreak Highlights the Importance of Vaccination - Copy | February 2019 Blogs, (n.d.). https://www.ihs.gov/newsroom/ihs-blog/february2019/measles-outbreak-highlights-the-importance-of-vaccination/ (accessed May 30, 2020).

[21] WHO | Zimbabwe Measles Immunization and Child Health Days campaign 2010, (n.d.). http://www.who.int/hac/crises/zwe/sitreps/24may_2june2010/en/ (accessed April 10, 2020).

[22] Y. Ganin, V. Lempitsky, Unsupervised domain adaptation by backpropagation, 32nd Int. Conf. Mach. Learn. ICML 2015. 2 (2015) 1180–1189. http://arxiv.org/abs/1409.7495 (accessed May 20, 2019).

Address for correspondence

Cui Tao, Ph.D., FACMI

7000 Fannin, Suite 600, Houston, TX 77030, USA

Cui.Tao@uth.tmc.edu

Phone: 713-500-3981 Fax: 713-500-3929