Title: Developing an automatic pipeline for analyzing chatter about health services from social media: A case study for Medicaid

Authors:

*Yuan-Chi Yang, PhD[1]
Mohammed Ali Al-Garadi, PhD[1]
Whitney Hogg-Bremer, BS[1]
Jane M. Zhu, MD, MPP, MSHP[2]
David Grande, MD, MPA[3]
Abeed Sarker, PhD[1]

[1]Department of Biomedical Informatics, School of Medicine, Emory University, Atlanta, GA, United States;
[2]Division of General Internal Medicine and Geriatrics, Oregon Health & Science University, Portland, OR;

[3]Division of General Internal Medicine and Geriatrics, School of Medicine, University of Pennsylvania, Philadelphia, PA;
*Corresponding author

Postal address: 101 Woodruff Circle, 4th Floor East, Atlanta, GA 30322

Email: yuan-chi.yang@emory.edu

Phone: 404-727-6123

Word count:

Abstract: 249

Body: 3968

Keywords:

Medicaid, Social Media, Consumer Feedback, Machine Learning, Natural Language Processing

Abstract:

Objective

Social media can be an effective but challenging resource for conducting close-to-real-time assessments of consumers' perceptions about health services. Our objective was to develop and evaluate an automatic pipeline, involving natural language processing and machine learning, for automatically characterizing user-posted Twitter data about Medicaid.

Material and Methods

We collected Twitter data via the public API using Medicaid-related keywords (Corpus-1), and the website's search option using agency-specific handles (Corpus-2). We manually labeled a sample of tweets into five pre-determined categories or *other*, and artificially increased the number of training posts from specific low-frequency categories. We trained and evaluated several supervised learning algorithms using manually-labeled data, and applied the best-performing classifier to collected tweets for post-classification analyses assessing the utility of our methods.

Results

We collected 628,411 and 27,377 tweets for Corpus-1 and -2, respectively. We manually annotated 9,571 (Corpus-1: 8,180; Corpus-2: 1,391) tweets, using 7,923 (82.8%) for training and 1,648 (17.2%) for evaluation. A BERT-based (bidirectional encoder representations from transformers) classifier obtained the highest accuracies (83.9%, Corpus-1; 86.4%, Corpus-2), outperforming the second-best classifier (SVMs: 79.6%; 76.4%). Post-classification analyses revealed differing inter-corpora distributions of tweet categories, with *political* (63%) and *consumer-feedback* (43%) tweets being most frequent for Corpus-1 and -2, respectively.

Discussion and Conclusion

The broad and variable content of Medicaid-related tweets necessitates automatic categorization to identify topic-relevant posts. Our proposed pipeline presents a feasible solution for automatic categorization, and can be deployed/generalized for health service programs other than Medicaid. Annotated data and methods are available for future studies (LINK_TO_BE_AVAILABLE).

BACKGROUND

Consumers' perspectives and feedback are crucial for improving products or services. Over the last two decades, widespread adoption and use of the Internet has led to its utilization as a major platform for collecting targeted consumer feedback. Businesses often allow consumers to rate specific products and/or services, and also provide detailed comments or reviews, and this has become a key feature of e-commerce platforms. For example, consumer-generated reviews and ratings of products play an important role in differentiation on Amazon, which currently has a global presence.[1,2] There are also companies, such as Yelp, that focus specifically on enabling crowdsourcing consumer feedback.[3-6] Similarly, as social media has become the primary platform of communication for many people, many companies have started maintaining and communicating via social media accounts, often enabling direct communications, both private and public, with consumers. Not only do consumers provide comments or seek assistance through those social media accounts, but they also often engage in discussions about products or services within their own social networks. Consequently, such consumer-generated chatter is often utilized to assess their perceptions of specific topics, which may range from products or services to social programs, legislations and politicians.

Social media is a rich resource for obtaining perspectives on public health, since it enables the collection of large amounts of data directly and in real-time. It is commonly used for sentiment analysis—a field of study that analyzes opinions, sentiments, attitudes and emotions from written language. Sentiment analysis research involving social media data has covered a wide range of topics, events, individuals, issues, services, products, and organizations.[7,8] The use of social media has not, however, been limited to sentiment analysis in open domains. Over recent years, research within the broader medical domain have embraced social media, and it is currently being utilized for conducting real-time public health surveillance, including for topics such as influenza surveillance, pharmacovigilance and toxicovigilance.[9-11] Meanwhile, similar to corporate businesses in the United States (US), health service providers such as local health departments and hospitals have also started adopting social media specifically as a consumer-facing communication channel.[12,13] Prior studies in this space have investigated how the social media data linked to such health services accounts reflects the consumers' perspectives about them. The simplest studies have focused on utilizing structured or numeric information, such as likes or ratings, associated with the accounts belonging to hospitals or nursing homes, and these metrics had been compared against traditional quality reports and ratings.[14-16] Building on the advances in open-domain natural language processing (NLP), some studies within the broader health domain have attempted to use unstructured data, including postings related to patient experiences about hospitals, to infer consumer sentiments[17,18] or extract topics that summarize content.[19]

Extracting knowledge from social media data is notoriously difficult for NLP methods due to factors such as the presence of misspellings, colloquial expressions, lack of context, and noise. These problems are exacerbated for health-related data due to the complexities of domain-specific terminologies, the lack of expert knowledge among

common social media users, and the uniqueness of health-related topics. Consequently, there is considerably less research exploiting social media free text data for health-related tasks. Past studies closely related to ours have focused on analyzing sentiments towards attributes of health insurance plans[20] and social media users' responses to public announcements about health policies.[21] However, to the best of our knowledge, there have been no close-to-real-time automatic pipeline that provides comprehensive data collection and analysis on social media chatter about health services and insurance coverage provided by large public insurers such as Medicaid and Medicare. This is perhaps because social media chatter associated with these topics cover more diverse perspectives, which increase the complexities posed by them to NLP methods. For example, for a medical entity such as a hospital, it is easy, from a relative perspective, to characterize the chatter in terms of user sentiment (*e.g.*, how the user feels about the services provided by the hospital). However, chatter associated with an entity like Medicaid contains discussions about politics and legislations, academic research, statistics and factual information, consumer feedback, and so on. Chatter related to politics will be different, in terms of content, compared to chatter related to consumer feedback. In fact, sentiment may also mean different things for these two broad categories of chatter—negative sentiment in political chatter may represent a user's emotions associated with a political decision about the health service (*e.g.*, changes in policies related to insurance coverage or covered benefits within Medicare or Medicaid), rather than the service itself.

Due to these additional complexities, there is a need to identify the broad categories of information in such social media chatter before they can be used for targeted analyses, such as sentiment analyses. These complexities, combined with the promise of social media data and the lack of past research in this specialized area, served as the primary motivation for the study described in this paper. We chose Medicaid as our target health service because it is the single largest public insurance program in the country[22] and contains large volumes of related chatter on social media.

The specific objectives of this paper are to:

- Assess if a social media platform, specifically Twitter, contains sufficient volumes of chatter about health services so that it can be used to conduct largescale analyses, using Medicaid as our target service.
- Develop and discuss a data-centric pipeline, involving NLP and machine learning, for automatically collecting, categorizing and analyzing Twitter chatter associated with Medicaid.
- Describe the manual annotation of a Twitter-Medicaid dataset, and its composition.
- Describe supervised classification strategies for automatically classifying Medicaid-related tweets into broad categories, and evaluating the performances of several machine learning models, with particular emphasis on tweets that potentially represent consumer feedback.
- Conduct post-classification content analyses to verify the potential utility of our data-centric pipeline.

The main contributions of this paper are as follows:

- We present the methods and results of collecting Medicaid-related Twitter data, analyzing a sample of the data manually, and developing an annotation guideline suitable for preparing a large dataset for training classification algorithms.
- We present details of automatic supervised classification experiments, including methods, results and evaluations, and provide suggestions about how to further improve the performance.
- We discuss the post-classification analyses of the collected data, including data distribution and content analyses.
- We make the NLP and machine learning scripts in this work publicly available, along with the labeled training dataset and a larger set of unlabeled Medicaid-related data.

MATERIALS AND METHODS

Data Collection

To develop our models for the analyses of data related to Medicaid from Twitter, we collected two sets of publicly available data from the network, which we label as Corpus-1 and Corpus-2. Corpus-1 contains tweets mentioning the term '*medicaid,*' or Medicaid agency (MA) and managed care organization (MCO, an organization that provides Medicaid-related health services under contracts from the agency) names that are branded and thus easily distinguishable on Twitter (e.g., *Medi-cal*: California's Medicaid program; and *TennCare*, Tennessee's Medicaid program). These tweets were collected via Twitter's public streaming API (application programming interface) from May 1, 2018 to October 31, 2019. As it is reported that misspellings appear frequently on social media platforms,[23] particularly Twitter, we used an automatic spelling variant generator to generate common misspellings for '*medicaid*' to increase retrieval rate.[24] The full list of keywords, including the misspellings, are shown in Table 1 in the Supplementary Material. We then identified and removed tweets whose contents were not directly related to Medicaid and repeated/duplicated contents/tweets (*e.g.*, fund-raising or political campaign). To focus on tweets expressing personal opinion, we also removed retweets, deemed as duplicates of the original tweets. The final dataset consisted of 628,411 tweets for Corpus-1.

While most of the chatter regarding Medicaid posted by consumers only included the term '*medicaid*' (or its variants), some directly tagged or mentioned relevant Twitter handles associated with MAs or the MCOs (e.g., '@organization_name'). Corpus-2 is composed of such tweets, and the MA and MCO Twitter handles were identified in a previous study.[25] The full list of the handles used in data collection is presented as Table 2 in the Supplementary Material. These tweets were retrieved by targeted searching '(*e.g.*, *to:organization_name*)' on Twitter. These tweets were posted between December 12, 2008 and the time of search (January 9, 2020). We filtered the tweets using the same approaches that were used for Corpus-1. In total, there are 27,337 tweets in the corpus.

Tweet Contents and Manual Annotations

To better understand the contents of the tweets posted by users, and to develop methods to automatically categorize the posts, we first manually inspected the contents of the posts and identified commonly occurring themes from a random sample of tweets. In line with our extensive past work on similar topics, we used the grounded theory approach[26]—we first analyzed a random sample of tweets to identify recurring topics and then grouped the topics into broader categories/themes. We decided to settle for 5 broad categories—(i) academic, (ii) consumer feedback, (iii) information/outreach, (iv) news, and (v) political opinion/advocacy. Tweets that could not be categorized as any of these were labeled as *other*. Descriptions are given as the follows:

- **Academic (a):** Tweets related to research about Medicaid. These include tweets by persons or organizations with academic affiliations or think tanks, or any tweet relating to education, scholarship, and thought including (links to) journal publications and reports.
- **Consumer Feedback (c):** Tweets related to consumers' experiences or questions related to Medicaid services, coverage, benefits, or health issues. The tweets are typically from Medicaid consumers or family members of consumers, and can also include discussions with others.
- **Information/Outreach (i):** Tweets directed at consumers and beneficiaries of Medicaid to convey information including agency services, programs, events, enrollment, eligibility criteria, etc. The tweets containing information about general health or public health reminders are also included.
- **News (n):** News and announcements–including any tweets from a news agency/organization. The tweets from Medicaid agencies/plans are excluded.
- **Political opinion / advocacy (p):** Comments, personal opinions, and feedback about politics related to Medicaid.
- **Other (o):** Tweets that are not relevant, typically the noise that isn't captured by the initial screening.

Two trained annotators performed a first round of annotations (for the data in Corpus-1) in multiple iterations, developed annotation guidelines, and resolved ambiguities via discussion. Following the completion of this round of annotations, the annotation disagreements were resolved by AS and WH. We found the class distribution to be very imbalanced, with most of the tweets annotated as class *n*, *p*, and *o*, while only a small portion were in class *a*, *c*, and *i* (Table 1). To understand how this imbalanced distribution affected the classifier performances on the smaller classes, particularly the *c* class, we performed preliminary automatic classification experiments using three classifiers: Naïve Bayes (NB), Support Vector Machine (SVM) and Random Forest (RF). We split the data into training (80%) and test (20%) sets, and found the best performance on class *c* to be low for all the classifiers, with best F_1-score = 0.3 (SVM). Tweets belonging to the consumer feedback (*c*) class was of particular importance to our overarching project objectives, so we devised two strategies for improving performance for this class—the first involving additional annotations of targeted tweets from the same dataset, and the second focusing on collecting an additional dataset (Corpus-2, as described earlier).

For the first strategy, we conducted another round of annotation of tweets from Corpus-1 to increase the number of tweets for the *c* class. Due to the very low numbers of *c* class

tweets in the original dataset, we realized that it would not be feasible to annotate sufficient numbers of these tweets by drawing random samples because of budgetary and other constraints. Therefore, rather than randomly drawing tweets for the next round of annotations, which would again lead to finding a small number of tweets belonging to category *c*, we attempted to artificially increase the number of tweets for this category. We achieved this by running our above-described *weak* classifier on a larger set of unlabeled tweets and only picking tweets classified as class *c* by the SVM classifier. This significantly increased the number of class *c* tweets in the data to be annotated. The new set of annotated data were then added into the training set, and the data distribution is presented in Table 1.

We followed the same annotation strategy for Corpus-2 (i.e., annotating tweets classified as *c* by the classifier trained on the previously annotated data). An outline of the overall annotation process is presented in Figure 1. While we tried to decrease the class imbalance in the training sets of the two corpora, to ensure that our evaluations represented the classifier performances on real-world distributions of the data, we did not artificially balance these. Therefore, the test sets represented class distributions that we would expect if we employed the classifier on real-time streaming data.

Figure 1: Flow chart for the entire annotation process involving multiple rounds

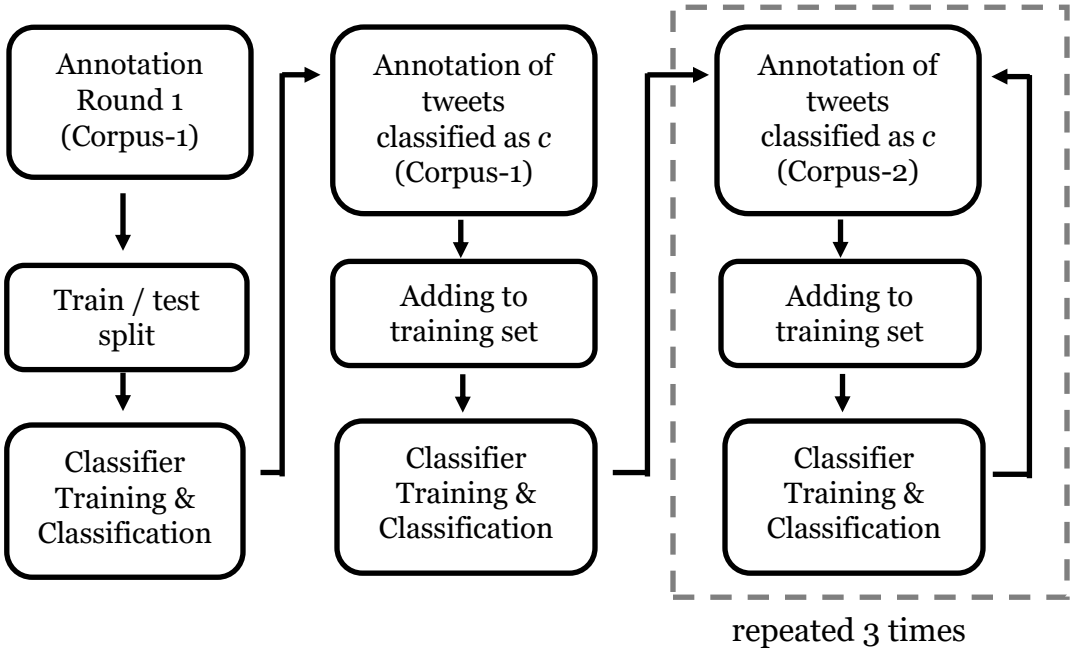

Table 1: Distribution for annotated data in the first round of annotations (row 2 & 3), and the final data sets (Corpus-1 for row 4 and 6; Corpus-2 for row 5 and 7)

Category	**a**	**c**	**i**	**n**	**p**	**o**	**Total**
Training set (1st round)	61	158	198	1288	3613	477	5795 (80%)

Test set (1st round)	35	37	49	317	897	114	1449 (20%)
Training set (Corpus-1)	83	356	429	1299	3710	854	6731
Training set (Corpus-2)	3	779	60	31	37	282	1192
Test set (Corpus-1)	35	37	49	317	897	114	1449
Test set (Corpus-2)	3	87	17	6	7	79	199
Total	124	1259	555	1653	4651	1329	9571

Classification

We experimented with five traditional classification algorithms, including Gaussian NB,[27,28] SVM,[29,30] RF,[31] k-nearest neighbor (KNN),[28] and shallow neural networks (NN), and two advanced classification algorithms, bi-directional long short-term memory (BLSTM)[32,33] and Bidirectional Encoder Representations from Transformers (BERT)[34,35]. We merged the training sets from the two corpora for training the traditional classifiers. Although the origin and distributions of tweets in the two corpora were different, we decided to combine them since our past research suggests that multi-corpus training, or distant supervision, leads to performance improvements for social media text classification.[36] For the advanced classifiers, the training data was further split into training (80%) and validation (20%) sets. The tweets were pre-processed before training or testing by lowercasing and anonymizing URLs and user names. For the traditional classifiers, the non-English characters were further removed, and each word was stemmed by the Porter stemmer. The features were 3000 most frequent n-grams (contiguous sequences of words with n ranging from 1 to 3), and 972 word clusters (a generalized representations of words learnt from medication-related chatter collected on Twitter).[37,38] For the advanced classifiers, each word or character sequence was replaced with a dense vector, and the vectors were then fed into the relevant algorithms for training. The performance on test set from Corpus-1 and Corpus-2 are shown in Table 2.

The hyperparameters for the traditional classifiers were as follows: 8 nearest neighbors for KNN, 3 hidden layers [numbers of units are (32,16,8)] for NN, 1000 estimators (trees) for RF, radial basis function kernel for SVM. The number of the estimators for RF, the number of nearest neighbors used in KNN, and the class weights for SVM were optimized for accuracy and F1-score on class c, using the training data (Corpus-1). We used Twitter GloVe word embeddings for the BLSTM[39] classifier, where each word is converted to 200-dimensional vector. BLSTM was then trained with 40 epochs and dropout regularization and the best model was selected through accuracy on validation data. We chose RoBERTa-large for BERT algorithms,[35] trained with 3 epochs.

Post-classification Analyses

To assess the utility of out classification approaches and gain understanding of the data, we used the best-performing classifier (the classifier based on BERT) to label all

collected unlabeled data and compute the data distribution. We then performed content analysis using the term-frequency-inverse-document-frequency method (TFIDF),[40] focusing on the tweets in Corpus-1 that contained the term '*medicaid*.' Our intent was to qualitatively assess that the classifier was capable of distinguishing tweets based on contents that are manually verifiable. For all content analyses, the text was first pre-processed by lowercasing, removing URLs, user names, non-English characters, stopwords, and any word with less than four characters.

RESULTS

Annotation and Class Distributions in test sets

We annotated a total of 8,180 tweets from Corpus-1 and 1,391 tweets from Corpus-2. We obtained substantial inter-annotator agreement (Cohen's $\kappa = 0.734$)[41,42] over 892 double-annotated tweets. The test data sets were randomly selected from the corpora and, therefore, can be considered a sample of the collected data. For Corpus-1, the test data contained 1449 tweets, among which the *political discussion* (class p) was the dominant class (62%) followed by *news* (class n; 22%), while *consumer feedbacks* (class c) made up less than 3% of the tweets. In contrast, consumer feedback comprised about 44% in Corpus-2 and roughly 40% of the tweets could not be categorized, most of which were part of conversations and could not be understood without full context.

Classification Results

The F1-scores and the accuracies of the classifiers on the test sets are presented in Table 2, including confidence intervals estimated using bootstrapping, while the precisions and the recalls are given in Table 3 in the Supplementary Materials due to space limit. For the test set from Corpus-1, the classifiers showed high performance for class *p*, but relatively lower for class *c*. This was expected based on the large imbalance described earlier. Among all the traditional classifiers experimented, SVM performed the best, with F_1-score of 0.53 on class *c* and 0.89 on class p, and overall accuracy 79.6%. RF has a comparable accuracy score (78.7%) and F_1-score on class *p* (0.87) but the F_1-score on class c was only 0.35. In contrast, we found that the advanced classifiers in general had higher overall accuracy, but only BERT had higher F_1-score bother on both class *p* (0.91) and class *c* (0.63), while BLSTM showed similar F_1-score on *c* (0.53).

For the test set from Corpus-2, the classifiers performed well on the class *c* while the overall accuracies were comparable to those on test set from Corpus-1. The classifiers performed relatively poorly on class *p*, due to its low frequency. SVM still performed the best among the traditional classification algorithms experimented, with accuracy of 76.4% and F_1-score of 0.84 on class *c*. On the other hand, BERT performed better than all the other classifiers, with accuracy of 86.4% and class *c* F_1-score of 0.92.

Table 2: Classification performances of the classifiers on the test sets of Corpus-1 & 2. The 95% confidence intervals are given in parenthesis.

Classifier	test set (Corpus-1)	F_1-score (.XX)						Accuracy (%)
		a	c	i	n	p	o	
	NB	11	17 (11 – 24)	24	57	71	34	55.7 (53.1 – 58.2)
	SVM	37	**53 (40 – 64)**	41	75	**89**	47	**79.6 (77.5 – 81.6)**

	RF	5	35 (17 – 51)	21	75	87	46	78.7 (76.6 – 80.8)
	KNN	0	27 (12 – 41)	9	55	66	26	52.9 (50.3 – 55.3)
	NN	17	40 (26 – 53)	23	71	86	46	73.8 (71.6 – 76.1)
	BLSTM	0	53 (39 – 66)	31	75	88	52	79.0 (76.8 – 81.0)
	BERT	53	**63 (50 – 74)**	53	82	**91**	58	**83.9 (82.0 – 85.8)**
	test set (Corpus-2)	**F_1-score (.XX)**						**Accuracy (%)**
		a	**c**	**i**	**n**	**p**	**o**	
	NB	0	74 (66 – 80)	22	21	0	19	45.2 (38.2 – 52.3)
	SVM	50	**84 (78 – 90)**	50	46	0	78	**76.4 (70.4 – 81.9)**
	RF	0	81 (74 – 87)	0	44	10	73	68.8 (62.3 – 75.4)
	KNN	0	22 (11 – 33)	0	0	0	60	45.2 (38.2 – 52.3)
	NN	50	74 (67 – 81)	47	20	6	72	63.8 (57.3 – 70.4)
	BLSTM	0	77 (69 – 83)	33	77	23	70	65.3 (58.8 – 71.9)
	BERT	80	**92 (87 – 96)**	84	50	42	90	**86.4 (81.4 – 91.0)**

Error analysis

We conducted a brief analysis of the errors made by the BERT-based classifier. We first calculated the confusion matrix on both test sets (

Table 3). For Corpus-1, we highlighted that the classifier mostly misclassified political tweets as news or consumer feedback, where the latter, along with the low volume of consumer feedback, contributes to the low precision on the *c* class. This is not surprising because users sometimes commented and discussed policy with personal experience and some news content was related to opinions about the policy. Similarly, tweets labeled as consumer feedback or news were most frequently misclassified as political discussion. We also observed that some news tweets were misclassified as the *i* class because information is frequently spreading through news media, and that the uncategorized tweets were most frequently misclassified as political content, possibly because the tweets were about political discussion, but not directly about Medicaid policy.

For Corpus-2, the dominating classes were consumer feedback (*c*) and uncategorized tweets (*o*), and these two were most frequently misclassified as political tweets (*p*) or as each other. The reason that they were misclassified as political tweets is the same as in Corpus-1. We suspect they were misclassified as each other because tweets sometimes lacked context, making their meanings ambiguous and hard for the machine to understand. For example, the tweet *'<organization_name> RUDE RUDE RUDE RUDE RUDE RUDE,'* though ambiguous, might be understood as that the customer service representative was rude, and thus, we categorized it as consumer feedback. However, the machine learning algorithms are not capable of deciphering such implicit contexts—that 'rude' is usually associated with attitude, and, in tweets directed to the agency's handle, it is likely about a customer service representative. Similarly, the tweet *'<organization_name> that's the number I called,'* mostly likely belongs to a conversation between a customer and a representative but the lack of information other than 'number' renders it to the *o* class. However, the machine learning are unlikely to be able to capture this understanding.

Table 3: BERT Classifier's confusion matrix on test set.

test set (Corpus-1)		Predicted Value					
		a	c	i	n	p	o
True Value	a	17	0	3	9	6	0
	c	0	27	1	0	7	2
	i	1	0	30	9	5	4
	n	8	0	17	273	18	1
	p	1	18	8	46	816	8
	o	2	4	6	10	39	53
test set (Corpus-2)		Predicted Value					
		a	c	i	n	p	o
True Value	a	2	0	0	1	0	0
	c	0	80	0	0	4	3
	i	0	0	13	3	0	1
	n	0	1	1	4	0	0
	p	0	1	0	1	4	1
	o	0	5	0	1	4	69

Table 4: Examples of Misclassified tweets by BERT Classifier on Corpus-1 and Corpus-2

	Tweets	True class (prediction)	Comments
test set (Corpus -1)	for me to get the help i need, i need to be diagnosed and get on disability so i can get off marketplace and on medicaid. my current psychiatrist has been against this for the 20 years i've seen him. i need to get set up with a mental health center. /2	*p* (*c*)	Discussion about Medicaid policy with personal experience
	the cost: gov. newsom (idiot) demanding $260 million for medicaid expansion to illegal immigrants | sean hannity <URL>	*p* (*n*)	Opinion on Medicaid policy presented as a news title
	<USERNAME> i'm on medicaid so i can't be obligated to pay, but that probably just means they won't let me have it.	*c* (*p*)	Customer's discussion about Medicaid coverage. It might be misclassified due to similarity to the

			Medicaid policy discussion
	enrollment numbers have been steadily declining and current, publicly available department goals indicate that the objective is to reform medicaid further due to a lack of funds. and into this state of public health care the general assembly is considering torpedoing funding.	*n* (*p*)	News about possible Medicaid policy reformation
	2nd hearing on gov. dewine's state operating #ohbudget is underway. there will be testimony from a number of cabinet directors on medicaid, health & human services agencies. watch live online at <URL> read testimony at <URL> under 'hearings' <URL>	*n (i)*	News about information related to Medicaid
	why is it okay for tricare, medicaid, and medicare to use private doctors but not the va? people who want to restrict the options of vets care are obviously in someones pocket & want to hide mass systemic failures from civilian eyes.	*o* (*p*)	Political discussion about VA insurance rather than Medicaid
test set (Corpus -2)	Tweets	True class (prediction)	Comments
	<organization_name> RUDE RUDE RUDE RUDE RUDE RUDE	*c* (*o*)	Most likely comment that the customer's service representative is rude but hard to pick up by algorithm.
	<organization_name> that's the number I called	*o* (*c*)	Classified as others because of lack of information other than 'number,' but the algorithm might recognized that it could be a conversation between a customer and a representative.

Post-classification analyses—data distribution

We applied the best-performing classifier to label both corpora. The class distribution obtained is shown in Figure 2. We found that the majority of tweets in Corpus-1 were news (class *n*, 23%) and political discussion (class *p*, 63%), while consumer feedbacks (class *c*) only made up 4%, consistent with the data distribution of the test set of Corpus-1. The data distribution indicates that this corpus suits to analysis of chatters regarding to political discussion or news. For Corpus-2, the majority of the tweets were labeled as consumer feedback (class *c*, 43%) and uncategorized (class *o*, 37%), also consistent with the data distribution in the test set.

Figure 2: Post-classification class distributions among to two corpora, as per the automatically classified tweets.

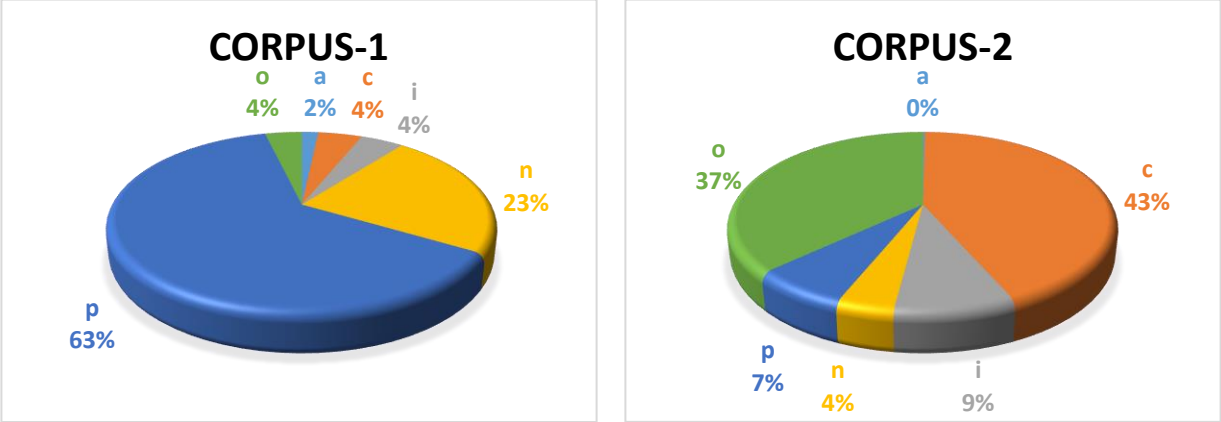

Post-classification Analyses: content of each class in Corpus-1

We now briefly summarize the findings from content analyses on the tweets in Corpus-1 that contain the terms associated with '*medicaid*,' in order to understand, from a high-level perspective, the contents within each category. The 10 highest-ranking bigrams and trigrams detected by TFIDF method are listed in Table 4 in the Supplementary Materials.[40] Not surprisingly, the Academic tweets (class *a*) are dominated by terms starting with *'study …'* and terms indicating research finding. Similarly, the *i* class contain terms related to *'service,' 'care,'* … etc, consistent with information outreach. For the *n* class, we found that many tweets were about news on medicaid work requirements in Kentucky and Arkansas (blocked by federal judge on March 27, 2019). In addition, *'social security'* and *'Trump …'* are also highly-ranked among the *n* and *p* classes. For the tweets belonging to the *c* class, some of the high-ranking terms were shared with other classes (e.g. *'… insurance,' 'social security,'* or *'… care'*) while some were specific to this class (*'make much'* or *'doesn cover'*) and potentially indicated comments about Medicaid income cap and coverage.

We did not know the compositions of the two datasets we had collected *a priori*. Thus, the results of our classification experiments provided us very important knowledge about which type of Twitter data to use when conducting targeted studies about Health services in general, or Medicaid here. For example, when studying consumer feedback, it is best to use data from Corpus-2 (*i.e.*, tweets containing Twitter handles of the MA or MCO); for studying public perceptions of political decisions, Corpus-1 would be more

useful. Detailed content analyses of the tweets in each category, such as their temporal and geolocation-specific distributions are likely to reveal more relevant information. However, such analyses is outside the scope of the current study, and we plan to build on the NLP pipeline described in this paper to conduct more thorough content analyses in the future.

DISCUSSION AND CONCLUSION

We have developed a social media mining pipeline, involving NLP and machine learning, for continuously collecting and categorizing Twitter chatter about the Medicaid program. Our study demonstrates that it is possible to collect data about a large, complex health services and coverage program like Medicaid, using Twitter to obtain close-to-real-time knowledge about consumer perceptions and opinions. The automatic classification of streaming data is crucial, specifically for smaller classes, such as *c*, for studying targeted topics.

Moving forward, we believe that the automatic classification may be improved via modifying the annotation guidelines to further eliminate the ambiguities between classes. Using a multi-label classification model (*i.e.*, multiple labels for a single tweet) may also lead to better categorization, particularly for the borderline tweets. It is also possible to further improve the performance via training the classifiers for each corpus separately, or even constructing ensemble classifiers. Also, the content can be analyzed using other text mining techniques, such as topic modeling or sentiment analysis, or detailed, manual analysis on selected samples.

Our analysis can inform public health researchers on how to use public discussion about health programs and services like Medicaid. Similarly, our pipeline can be deployed by research groups or Medicaid agencies for continuous, on-going research on evolution of the public opinions on social media (e.g. the impact of certain policy changes or rulings). We also note that, though this work focuses on Medicaid, our methods and open-source code can readily be applied to other health services.

FUNDING

Research reported in this publication was supported by Robert Wood Johnson Foundation (RWJF) under award number 76158 (JMZ, DG). The content is solely the responsibility of the authors and does not necessarily represent the official views of the RWJF.

AUTHER CONTRIBUTIONS

YY conducted and directed the machine learning experiments, evaluations and data analyses, with assistance from MAA, AS and WH. YY, AS, MAA and WH contributed to the data collection, annotation and analyses. JMZ and DG provided their expertise in preparing the annotation guidelines and categories, and helped formulate the overarching objectives of the project. AS, JMZ and DG provided supervision for various aspects of the study. YY drafted the manuscript and all authors contributed to the final manuscript.

CONFLICT OF INTEREST

None declared

ACKNOWLEDGEMENTS

The authors thank the support from the Robert Wood Johnson Foundation.

REFERENCES

1. Chen P-Y, Wu S-y, Yoon J. The impact of online recommendations and consumer feedback on sales. *ICIS 2004 Proceedings.* 2004:58.
2. Mudambi SM, Schuff D. Research note: What makes a helpful online review? A study of customer reviews on Amazon. com. *MIS quarterly.* 2010:185-200.
3. Hu M, Liu B. Mining and summarizing customer reviews. Paper presented at: Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining2004.
4. Akay A, Dragomir A, Erlandsson B-E. A novel data-mining approach leveraging social media to monitor consumer opinion of sitagliptin. *IEEE journal of biomedical and health informatics.* 2013;19(1):389-396.
5. Lim Y-s, Van Der Heide B. Evaluating the wisdom of strangers: The perceived credibility of online consumer reviews on Yelp. *Journal of Computer-Mediated Communication.* 2015;20(1):67-82.
6. Luca M. Reviews, reputation, and revenue: The case of Yelp. com. *Com (March 15, 2016) Harvard Business School NOM Unit Working Paper.* 2016(12-016).
7. Yue L, Chen W, Li X, Zuo W, Yin M. A survey of sentiment analysis in social media. *Knowledge and Information Systems.* 2018:1-47.
8. Liu B. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies.* 2012;5(1):1-167.
9. Broniatowski DA, Paul MJ, Dredze M. National and Local Influenza Surveillance through Twitter: An Analysis of the 2012-2013 Influenza Epidemic. *Plos One.* 2013;8(12).
10. Sarker A, O'Connor K, Ginn R, et al. Social media mining for toxicovigilance: automatic monitoring of prescription medication abuse from Twitter. *Drug safety.* 2016;39(3):231-240.
11. O'Connor K, Pimpalkhute P, Nikfarjam A, Ginn R, Smith KL, Gonzalez G. Pharmacovigilance on twitter? Mining tweets for adverse drug reactions. Paper presented at: AMIA annual symposium proceedings2014.
12. Griffis HM, Kilaru AS, Werner RM, et al. Use of Social Media Across US Hospitals: Descriptive Analysis of Adoption and Utilization. *Journal of Medical Internet Research.* 2014;16(11):e264.
13. Harris JK, Mueller NL, Snider D. Social media adoption in local health departments nationwide. *Am J Public Health.* 2013;103(9):1700-1707.
14. Glover M, Khalilzadeh O, Choy G, Prabhakar AM, Pandharipande PV, Gazelle GS. Hospital Evaluations by Social Media: A Comparative Analysis of Facebook Ratings among Performance Outliers. *Journal of General Internal Medicine.* 2015;30(10):1440-1446.
15. Campbell L, Li Y. Are Facebook user ratings associated with hospital cost, quality and patient satisfaction? A cross-sectional analysis of hospitals in New York State. *BMJ Qual Saf.* 2018;27(2):119-129.
16. Hefele JG, Li Y, Campbell L, Barooah A, Wang J. Nursing home Facebook reviews: who has them, and how do they relate to other measures of quality and experience? *BMJ Qual Saf.* 2018;27(2):130-139.

17. Hawkins JB, Brownstein JS, Tuli G, et al. Measuring patient-perceived quality of care in US hospitals using Twitter. *Bmj Quality & Safety.* 2016;25(6):404-413.
18. Rastegar-Mojarad M, Ye Z, Wall D, Murali N, Lin S. Collecting and Analyzing Patient Experiences of Health Care From Social Media. *JMIR Res Protoc.* 2015;4(3):e78.
19. Ranard BL, Werner RM, Antanavicius T, et al. Yelp Reviews Of Hospital Care Can Supplement And Inform Traditional Surveys Of The Patient Experience Of Care. *Health Affairs.* 2016;35(4):697-705.
20. van den Broek-Altenburg EM, Atherly AJ. Using social media to identify consumers' sentiments towards attributes of health insurance during enrollment season. *Applied Sciences.* 2019;9(10):2035.
21. Hatchard JL, Neto JQF, Vasilakis C, Evans-Reeves KA. Tweeting about public health policy: Social media response to the UK Government's announcement of a Parliamentary vote on draft standardised packaging regulations. *Plos One.* 2019;14(2).
22. Medicaid enrollment changes following the ACA. https://www.macpac.gov/subtopic/medicaid-enrollment-changes-following-the-aca/. Accessed March 11, 2020.
23. Han B, Cook P, Baldwin T. Lexical Normalization for Social Media Text. *Acm Transactions on Intelligent Systems and Technology.* 2013;4(1):1-27.
24. Sarker A, Gonzalez-Hernandez G. An unsupervised and customizable misspelling generator for mining noisy health-related text sources. *J Biomed Inform.* 2018;88:98-107.
25. Zhu J, Sarker A, Gollust S, Merchant R, Grande D. *Characteristics of Twitter Use by State Medicaid Programs in the U.S (Preprint).* 2020.
26. Martin PY, Turner BA. Grounded Theory and Organizational Research. *Journal of Applied Behavioral Science.* 1986;22(2):141-157.
27. Rish I. An empirical study of the naive Bayes classifier. Paper presented at: IJCAI 2001 workshop on empirical methods in artificial intelligence2001.
28. Cover TM, Hart PE. Nearest Neighbor Pattern Classification. *Ieee Transactions on Information Theory.* 1967;13(1):21-+.
29. Chang CC, Lin CJ. LIBSVM: A Library for Support Vector Machines. *Acm Transactions on Intelligent Systems and Technology.* 2011;2(3):1-27.
30. Platt J. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers.* 1999;10(3):61-74.
31. Ho TK. Random decision forests. Paper presented at: Proceedings of 3rd international conference on document analysis and recognition1995.
32. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput.* 1997;9(8):1735-1780.
33. Schuster M, Paliwal KK. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing.* 1997;45(11):2673-2681.
34. Devlin J, Chang M-W, Lee K, Toutanova K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:181004805.* 2018.
35. Liu Y, Ott M, Goyal N, et al. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:190711692.* 2019.
36. Sarker A, Gonzalez G. Portable automatic text classification for adverse drug reaction detection via multi-corpus training. *Journal of biomedical informatics.* 2015;53:196-207.
37. Sarker A, O'Connor K, Ginn R, et al. Social Media Mining for Toxicovigilance: Automatic Monitoring of Prescription Medication Abuse from Twitter. *Drug Saf.* 2016;39(3):231-240.
38. Nikfarjam A, Sarker A. O Connor K, Ginn R, Gonzalez G. Pharmacovigilance from social media: Mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *J Am Med Inf Assoc.* 2015;22:671-681.

39. Pennington J, Socher R, Manning CD. Glove: Global vectors for word representation. Paper presented at: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)2014.
40. Beel J, Gipp B, Langer S, Breitinger C. Research-paper recommender systems: a literature survey. *International Journal on Digital Libraries.* 2016;17(4):305-338.
41. Viera AJ, Garrett JM. Understanding interobserver agreement: the kappa statistic. *Fam Med.* 2005;37(5):360-363.
42. Cohen J. Weighted kappa: nominal scale agreement with provision for scaled disagreement or partial credit. *Psychol Bull.* 1968;70(4):213-220.