Development and Evaluation of Tagalog Linguistic Inquiry and Word Count (LIWC) Dictionaries for Negative and Positive Emotion

Amanda L. Andrei

23 December 2014

Introduction

As the use of online and social media increases globally, the need for sentiment analysis tools in multiple languages is critical in order to understand and analyze the vast amount of data that may contain users' feelings, perceptions, and beliefs. Users from different countries convey their messages in various languages, which may convey different sentiments and cultural connotations. Developing non-English sentiment analysis tools can ensure that data is not lost due to language. A proof-of-concept Tagalog Linguistic Inquiry and Word Count (LIWC) dictionary for positive and negative emotion was developed for use in analyzing mixed language Twitter data from the Philippines and evaluated against human-annotated sentiment for Twitter, referred to as groundtruth.

Background

Linguistic Inquiry and Word Count (LIWC)

The LIWC program (Pennebaker, 2001; Pennebaker, 2007) is a social psychology tool that is increasingly being used for content analysis of social media, particularly Twitter (Servi & Elson, 2012; Gunn & Lester, 2012). Originally created in English and spanning over 4,000 words and stems, LIWC uses "dictionaries" of words which correspond to various domains of linguistic processes, psychological processes, personal concerns, and spoken categories. For the purposes of this study, the most salient dictionaries for sentiment analysis were Positive Emotion and Negative Emotion (contains sub-dictionaries for Anger, Anxiety, and Sadness).

Translating LIWC into other languages may reveal insights into cross-cultural psychology (Hayeri et al., 2010). The dictionaries currently exist in twelve different languages, although not all of the languages have been verified and validated to the same extent as the original English dictionary (LIWC.net). Additionally, the creators of the Spanish dictionary (Ramirez-Esparza et al., 2007) and the German dictionary (Wolf et al., 2008) fitted the non-English lexicons to the English grammatical scheme, and therefore excluded significant sociolinguistic aspects of the languages, such as the T-V distinction (from the Latin *tu* and *vos*, respectively the singular and plural forms of "you"), where use of different pronouns indicates different roles of respect, power, kinship, and other sociocultural norms (see Brown and Gilman, 1960). The Arabic dictionary team (Hayeri et al., 2010) chose to translate the English LIWC into Arabic as well as create a dictionary based on an Arabic grammatical scheme and found high consistency between the two dictionaries; however, the Arabic dictionary does not appear to have been evaluated against human annotated data. Finally, an evaluation of the Brazilian Portuguese LIWC negative and positive dictionaries against two other sentiment lexicons revealed that the Brazilian Portuguese LIWC positive dictionary performs well, while the negative dictionary does not perform as well as the other lexicons (Balage Filho et al., 2013).

Our work started with development of a Tagalog LIWC dictionary for disaster (Andrei, et al. 2014); in this study we describe the development of positive and negative emotion dictionaries in Tagalog and their inclusion in LIWC list (LIWC.net).

Philippines

With a complex and politicized history of multilingualism (Gonzalez 1998; Nical et al. 2004), most Filipinos use multiple dialects and languages (primarily English and Filipino, which is heavily derived from Tagalog), both in everyday speech and online. In fact, despite the fact that standardized and annotated corpora for Tagalog do not exist, the language maintains a significant presence on the web (Zuraw, 2006). Furthermore, the Filipino diaspora has extended the number of Tagalog-speakers beyond just the Philippines, with significant communities in the United States, the Middle East, and other parts of Europe and Asia (Philippine Statistics Authority, 2013). Many of these Filipinos use social media to stay in touch with family and friends back in the Philippines.

Additionally, other Philippine groups and organizations use social media in mixed languages to broadcast messages. For instance, the Philippine government and weather station disseminate public service announcements during natural disasters via Twitter and Facebook. Extremist groups such as the Moro National Liberation Front also tweet messages and post on Facebook to rally followers and denounce the Philippine government. From a policy perspective, both of these social media use cases—disaster management and extremist behavior—are prime candidates for mixed language sentiment analysis.

Methodology

Creating positive and negative emotion Tagalog LIWC dictionaries

The first factor to consider when creating non-English LIWC dictionaries is whether or not to translate the original English LIWC into the non-English version, or to start with a non-English corpus and categorize it based on the LIWC structure. If the purpose is to create a tool that can be used for general texts, or to see how accurate the LIWC structure is, then it may be best to start with translating from the English into the target language and to start with closed class vocabulary categories, such as pronouns. However, the purpose of these Tagalog LIWC dictionaries were to measure sentiment across a specific kind of text (tweets), and therefore it was more appropriate to begin with a corpus of tweets and other online text and filter it for Tagalog words.

The three main phases in creating the dictionaries were: obtaining a corpus of words, verifying translations, and classifying words as positive or negative. As the original focus of this project was violent extremism, the original corpus used to create the LIWC dictionaries consisted of a list of the most frequently used terms from a variety of Philippine blogs, online news articles, and Twitter posts about violent extremism, containing domain-specific vocabulary such as *NPA* (National People's Army), or *nakubkob* (besieged). The list totaled 18,254 unigrams in English, Tagalog (Filipino), Cebuano, Indonesian, and Spanish, as well as punctuation, Twitter handles, hashtags, and other Internet slang and acronyms.

The list was run through Google Translate for Filipino and then manually verified by a heritage speaker with intermediate knowledge of Tagalog. The resulting verification produced a filtered list of 1,510 Tagalog words, which excluded potential Tagalog Internet slang and acronyms. The rationale behind this filtering was because the English LIWC dictionary as of yet does not have a separate dictionary for Internet slang, and also because capturing the rapidly evolving cyber-language in the Philippines was beyond the scope of this project.

In the original English LIWC, verifying dictionaries involves at least three judges independently assigning a word to categories ("judges" are renamed as "lexicographers" to more accurately capture their role). In the development of the Tagalog dictionary, four native Tagalog-speaking lexicographers were available to verify translation and categorize words as positive or negative. To ensure that lexicographers received an equal workload and to minimize fatigue, the main Tagalog corpus was split into two sets of 755 words each. Three lexicographers reviewed each set of words to verify translation and indicated if a word was Positive Emotion, Negative Emotion, Neither, or Both. They also recorded any alternative spellings, meanings, or notes about words (such as dialectal or slang use). Scoring a word as Positive Emotion or Negative Emotion required at least two lexicographers to independently tag a word with that emotion. Outcomes are summarized in Table 1. When lexicographers disagreed on a word, the word was deemed ambiguous and stored in a list to be re-evaluated in the future.

Lexicographer 1	Lexicographer 2	Lexicographer 3	Outcome
Positive	Positive	Positive	Positive
Positive	Positive	Negative/Neither/Both	Positive
Positive	Negative	Neither/Both	Evaluate
Both	Both	Both	Evaluate
Neither	Both	Both	Evaluate
Neither	Neither	Both	Neither
Neither	Neither	Neither	Neither
Negative	Negative	Negative	Negative
Negative	Negative	Positive/Neither/Both	Negative

Table 1: Scoring for lexicographers to classify a word as Positive, Negative, Neither, or Both.

Creating groundtruth from tweets

Due to unfolding events during the course of this research, the topic of the project shifted from violent extremism to disaster management, and so the groundtruth set of tweets was derived from a larger set of tweets about super typhoon Yolanda in the Philippines. This initial set of tweets consists of approximately 1.2 million tweets originating from the Philippines from November 3-18, 2013 that were harvested based on if hashtag or keywords related to the disaster (both in English and Tagalog), such as *typhoon*, *supertyphoon, Yolanda*, *bagyo, #yolandaph*, etc. (see Appendix for complete list). The preliminary groundtruth set was created by searching for tweets that contained negative or positive words from the Tagalog dictionaries. A random sample of twenty tweets per day was taken for each day, from both dictionaries.

The preliminary set was given to three native Tagalog-speaking annotators to mark on a Likert scale of "Strongly Negative," "Somewhat Negative," "Neutral," "Somewhat Positive," and "Strongly Positive." The tweets were then relabeled as "positive," "negative," or "neutral" depending on the majority score of the lexicographers. If a tweet was split across polarity, the tweet was labeled as 'unknown.' Tweets annotated with unknown were removed from the set, for a final groundtruth set of 575 tweets.

Since the original LIWC does not have a neutral dictionary, the tweets were classified in two different ways based on a different definition of 'neutral.' The unscored set of tweets defines positive as all tweets containing only positive words, negative as containing only negative words, and neutral containing both positive and negative words. This method of classification is not entirely true to the LIWC methodology of outputting percentages of categories, and so the same set was reclassified based on a sentiment score S, where

$$S = \frac{n_{pos} - n_{neg}}{N}$$

and n_{pos} = number of positive words, n_{neg} = number of negative words, and N = total number of words in the tweet. If $S > 0$, the tweet was classified as positive, if $S = 0$ the tweet was classified as neutral, and if $S < 0$, the tweet was classified as negative.

Therefore, the scored set of tweets contains a slightly greater amount of positive and negative tagged tweets than the unscored set. The set also accounts for the proportion of negative to positive words within each tweet. This method takes the LIWC program one step further—LIWC provides a percentage of negative words and positive words per document, but does not subtract them from each other to show proportion. Since this scoring method focuses specifically on emotion, it allows us to immediately rate the tweet as overall positive or negative.

Results

Tagalog positive and negative dictionaries were successfully created by closely following the methodology for the English LIWC. However, as this replication did not account for Tagalog morphology and used a web-based corpus on a specific topic, the dictionaries are not as extensive as the original English dictionaries and operate on simple word-matching to find sentiment within a document.

By evaluating the Tagalog dictionaries on precision and recall, it is also clear that the dictionaries can be further improved to detect sentiment. Furthermore, conducting this evaluation on these dictionaries is one of the first natural language processing evaluations on a LIWC dictionary (including English and other languages) and suggests a framework for evaluating LIWC dictionaries in the future.

Final dictionaries

The final numbers for the dictionaries are summarized in Table 2. Only 40.86% of the corpus was judged as associated with sentiment. More words were judged negative than positive. Examples of entries for the Tagalog dictionary are given in Table 3. Note, these numbers are total unique terms, and the English dictionary actually matches more words since it includes wildcards after some terms, therefore allowing a unique term to match multiple forms of the word variants. Considering that Tagalog has a rich morphology that includes prefixes, infixes, suffixes, and reduplication, the original LIWC methodology of including a wildcard after a root word is insufficient for analyzing the Tagalog output. Examples of the English dictionary are given in Table 4.

Dictionary	Total unique terms		% of Tagalog corpus (N = 1510)
	Tagalog	English	
Positive emotion	273	408	18.08%
Negative emotion	344	499	22.78%
Total	617	907	40.86%

Table 2: Final number of terms in the Tagalog dictionary, with English dictionary as comparison. Note that the English dictionary includes wildcards for some word stems, while the Tagalog dictionary is strictly word-matching.

Positive		Negative	
Tagalog	English translation	Tagalog	English translation
asa	hope	bagyo	Storm
asang	hope	baha	flood
awa	mercy	bahala	care
aya	governess	bakakon	liar
ayos	order	bakbakan	scrimmage

Table 3: Examples of words from the positive and negative Tagalog LIWC dictionaries.

In the Tagalog dictionaries, not all morphological forms of a word are listed. For instance, the word *asa* is listed in direct case (*asa*) and indirect case (*asang*), but *bagyo* exists only in the direct case. This is due to the fact that other inflections of the word were not present in the original corpus of high frequency words.

Positive	Negative
accept	fail*
accepta*	fake
accepted	fatal*
accepting	fatigu*
accepts	fault*

Table 4: Examples of words from the positive and negative English LIWC dictionaries.

One of the differences between the Tagalog and English dictionaries is that the English dictionary uses a simple rule for morphology with use of a wildcard. Therefore, words such as fail, failed, failing, failure, etc. are all tagged as negative words, increasing the number of words that LIWC can categorize beyond the simple word-matching mechanism of the Tagalog LIWC dictionaries. However, some inflections of words are entered as unique terms, as seen in the various forms of accept. It is not entirely clear from the LIWC 2007 manual how they determined why certain inflections are included and other excluded, e.g., why accept* is not entry on its own.

Evaluating the dictionaries

Both the positive and negative dictionaries were evaluated based on the metrics of precision, recall, and F-measure as defined below. While LIWC does not have a neutral dictionary, the same metrics were also used to evaluate the neutral tags.

$$precision = \frac{\#\ of\ correct\ polarity\ tags\ given\ by\ system}{total\ \#\ of\ polarity\ tags\ given\ by\ system}$$

$$recall = \frac{\#\ of\ correct\ polarity\ tags\ given\ by\ system}{total\ \#\ of\ possible\ correct\ polarity\ tags\ based\ on\ groundtruth}$$

$$F\ measure = \frac{2\ \times precision\ \times recall}{precision + recall}$$

The tool had higher precision when tagging negative tweets, but lower precision for identifying positive tweets. Additionally, reclassifying the tweets based on sentiment score increased the recall, but lowered its precision. The low scores for the neutral tweets are due to the fact that LIWC has no neutral dictionary, and therefore there is no direct comparison between the sentiment scores and LIWC. Metrics are summarized in Table 5.

Tweet Classification	Label	Precision	Recall	F-measure	# of correct tags given by system	# of tags given by groundtruth
Unscored	Positive	25.67%	49.48%	33.44%	48	97
	Negative	75.14%	38.50%	50.92%	139	361
	Neutral	9.09%	17.78%	12.03%	16	90
Scored	Positive	23.89%	55.67%	33.44%	54	97
	Negative	70.80%	44.32%	54.51%	160	361
	Neutral	11.38%	15.56%	13.15%	14	90

Table 5: Comparison of tweets with manually annotated groundtruth and Tagalog LIWC tool.

Discussion

The following discussion of the dictionaries is based on qualitatively examining patterns from the scored tweets. Tweet examples have been cleaned so as to remove Twitter handles, links, and emojis[1]. Since emojis convert into strings of characters in the files, they are removed; however, emoticons remain intact.

[1] Emojis are pictographs originating in Japan that are "more fleshed-out versions of emoticons" as well as "a pantheon of objects, activities, and events" (Lebduska, 2014).

Negative dictionary

The negative dictionary accurately detected negative emotion words such as *nakakatakot* (scared) and *kawawa* (pitiful). Swear words were also an accurate measure of a tweet being negative. The dictionary also correctly matched tweets with words related to death, such as *patay* (dead; death) and *namatay* (died; casualty) to negative sentiment tweets. Examples of tweets are provided in Table 6.

Tweet	Groundtruth	Tool	S
Huhu ***nakakatakot*** *yung bagyong parating :/* [crying sounds] this coming storm is **scary** :/	NEG	NEG	-0.167
Sa tala ng Marabut Municipal Gov't, 24 na ang ***namatay*** *sa bayan dahil sa Bagyong #YolandaPH* Of note from Marabut Municipal Gov't, 24 **casualties** due to Typhoon #YolandaPH	NEG	NEG	-0.063

Table 6: Examples of correctly tagged negative tweets. Even though the second tweet does not state emotion (and should be labeled as a fact), the lexicographers rated it as negative emotion. See Discussion of creating groundtruth for more information.

One reason the negative dictionary may have performed well on precision is due to the fact that one of the dictionary entries is *bagyo*. Since the set of tweets analyzed is about a typhoon, it is unsurprising that the Tagalog word for typhoon (*bagyo*) would appear so many times since a natural disaster event typically has negative connotations (see Bankoff 2002 for the cultural, historical, and political discourse around disasters). Additionally, the lexicographers classified *bagyo* as a negative word when creating the dictionaries which would lead to higher precision when the Tagalog LIWC negative dictionary matches words.

However, as mentioned in the Results section about the Final dictionaries, the morphology of words was not completely addressed, which may have had an impact on the performance of the negative dictionary. For instance, this version of the dictionary included several different forms of *patay* (death), which increased its precision when tagging negative sentiment words, but the dictionaries missed words such as *pinakalakas* (strongest), which uses the prefix *pinaka-* (superlative comparison) and the root *lakas* (strength).[2] Another feature about this version of the dictionary is that it does not search for these wildcards and instead searches the tweets only for the words in the dictionary. This condition produces results such as classifying tweets with *bagyo* (storm) as negative, but not catching tweets with *bagyong* (where –ng is an indirect case marker), which may have produced a greater recall for negative tweets as well as increased the number of neutral tweets. Examples of these tweets are given in Table 7.

Tweet	Groundtruth	Tool	S
Bagyong Yolanda; Pinakamalakas na bagyong dumaan sa buong mundo ngayong taon.-Balita Typhoon Yolanda; the strongest storm to pass across the whole world this year. –News	NEG	POS	0.091

[2] Note that *lakas* is an entry in the positive dictionary word, but was often used in a negative sense in these tweets.

Si bagyong #YolandaPH ang pinakamalakas na bagyo sa buong mundo sa taong ito. typhoon #YolandaPH is the strongest storm in the whole world in this year	NEG	NTR	0

Table 7: Examples of tweets incorrectly tagged that contain instances of *bagyong* and *pinakamalakas*, words which were not included in the tool which would have changed the tool rating.

Despite using three lexicographers, the negative emotion dictionary contained some words that are not negative. For instance, *lang* (only, just) often tagged a tweet as negative when groundtruth marked it as positive or neutral. Use of *lang* may have led lexicographers to tag tweets as negative, simply because the topic of the tweets was a negative event. This points to a training issue. Examples of misclassification are in Table 8.

Tweet	Groundtruth	Tool	S
*Bagyong Yolanda singing: Ikot Ikot **lang**! Ikot ikot ikot **lang**! Ikot ikot **lang**! Ikot ikot ikot!*[3] Typhoon Yolanda singing: **Just** round and round! **Just** round and round! **Just** round and round! Round and round and round!	NTR	NEG	-0.176
*Grabee kung mayaman **lang** ako mag - dodonate ako para sa mga nasalanta ng bagyong #YOLANDA! :(((* Yikes if **only** I'm rich, I will donate to those who were devastated due to typhoon #YOLANDA! :(((	POS	NTR	0

Table 8: Examples of incorrectly tagged tweets based on the miscategorized negative dictionary word *lang*.

Positive dictionary

The lower precision in the positive dictionary output is primarily due to sentiment words having different orientations within a certain domain. Because LIWC dictionaries are assessed on a word level, a word that may be initially classified as one sentiment may have the opposite sentiment meaning within a certain domain. For instance, *biktima* (victim) was judged as a negative sentiment word—however, this word would often co-occur with the positive word *tulong* (help; also included other inflections, e.g., *tumulong*, helping). Based on the reclassification of the tweets, this condition usually created a neutral tag (S = 0), but lexicographers often marked these types of tweets as positive, since the main message of the tweets were appeals for help or statements that victims were receiving help.

Malakas (strong) and *lakas* are other examples of sentiment words with different orientations. Both terms were classified as positive in the LIWC dictionary, yet often occurred in conjunction with words like *bagyo, ulan* (rain), and *hangin* (wind), which produced statements often judged as negative, such as "The wind is very strong." This condition led to lower precision in the positive output from the tool.

However, as with the negative dictionary, the positive dictionary was more accurate when detecting positive emotion words that were also deemed positive by groundtruth. For instance, words that

[3] Note: This is a lyric of a Filipino pop song, "Ikot Ikot" by Sarah Geronimo

conveyed gratitude and hope, such as *bangon* (rise up) and *salamat* (thank you), often matched positive groundtruth tweets.

Tweet	Groundtruth	Tool	S
*Sana naman yung milyong-milyong **tulong** ay makakarating talaga dun sa mga **biktima** ng Bagyong Yolanda!!!! pleaaase lang!!!* I hope that the millions of Yolanda aid really reach the real victims of Typhoon Yolanda!!!! Please!!!	POS	NEG	-0.059
*Sendong survivors sa Iligan City **tumutulong** na rin sa mga **biktima** ng bagyong yolanda sa Tacloban City* Sendong survivors in Iligan City are also helping the victims of typhoon yolanda in Tacloban City	POS	NTR	0
*Maraming **salamat** po! Ito ay para sa mga nabiktima nang bagyong Yolanda!* Thank you very much! This is for the victims of typhoon Yolanda!	POS	POS	0.067
*1sang Bayan tayo'y Babangon..walang maiiwan sa pag #**Bangon** #TaclobanLeyteSamar!* One country, we will rebuild… no one will be left behind as we #Rebuild #TaclobanLeyteSamar	POS	POS	0.222

Table 9: Examples of incorrectly and correctly tagged positive tweets.

Groundtruth

Although groundtruth was created based on the majority of annotators' assessments, the groundtruth could be made more rigorous. Since annotators were given open-ended instructions on how to judge polarity, there may be some tweets that could arguably change orientations. For instance, the tweets in Table 7 tagged as negative could also be interpreted as neutral facts (which in the case of the second tweet, would render the tool output as correct). The first tweet in Table 8, quoting a pop song lyric, might even be considered as positive, since it is expressing a humorous outlook (albeit this outlook could also be considered dark, cynical, or tasteless). The second tweet in Table 9 about victims helping each other can be considered a fact that the lexicographers marked as positive but under differing guidelines, may be classified as having neutral sentiment.

Improvements

Creation of dictionaries

Several improvements could be made to this methodology to make the dictionaries more robust or to develop a better analytic capability for sentiment analysis in Tagalog. For instance, classifying words into dictionaries was primarily successful with emotional words, but cases such as *lang* suggest that the final draft of the dictionary should be reviewed again by the lexicographers. If a manual-based approach were desired for developing sentiment dictionaries, lexicographers could classify sets of adjectives or adverbs, since these words are typically found to be related to opinions. This approach would be difficult to execute if the main developers of the dictionaries were not native speakers of the language and would need to include an additional step of either finding extensive lists of words to judge, or enlisting lexicographers to brainstorm the words before judging them. A crowdsourcing service, such as the Amazon Mechanical Turk (see www.mturk.com), may be useful in this case, as the service sets up a system which allows requestors to submit human intelligence tasks (such as translating or tagging words) to a vast and diverse group of people willing to work on-demand, thereby saving time and costs while improving reliability.

While the brainstorming method is more in keeping with the original English LIWC methodology, such lists could result in words that are not used in Twitter (e.g., users do not use such terms because they are too long). In addition, it may be best to start with Twitter words to develop Twitter based emotion dictionaries. Alternatively, dictionary-based approaches could automate this process and reduce time and costs. This would involve developing or obtaining a small set of sentiment words, using an algorithm to search an online dictionary (such as www.tagalog-dictionary.com) for the synonyms and antonyms of the seed words, and adding the seed words to the original set and iterating over the online dictionary until word collection ceases (see Hu and Liu 2004). Liu 2012 lists other more sophisticated computational and mathematical models for developing sentiment lexicons in various languages.

It is also worth noting that the English LIWC dictionaries are limited by the growing numbers of Internet-speak and slang. Although developing a lexicon of Tagalog Internet slang was beyond the scope of this project, recommended next steps for analyzing Twitter data would be to create such a lexicon.

Creation of groundtruth

Creating the groundtruth could also be revised. Annotators received open-ended guidelines on a five-point Likert scale that included an option for 'neutral,' a concept that is not included in the original or Tagalog LIWC dictionary. While this option was given to give annotators more freedom in categorizing, the scale could be revised in a variety of ways: removing neutral, adding an option for mixed sentiment, or modifying sentiment intensity. A subjectivity analysis could also be performed on the tweets before they are given to annotators, which would produce guidelines on how to judge a fact versus an opinion, or how to classify a negative versus positive fact. In examples such as quoting the pop lyric song, assessing humor or jokes within the tweets would require more subject matter expertise in Filipino culture and current events.

Additionally, another reason for low precision and recall could be due to different domains between the original corpus of words and the groundtruth. Since the LIWC tool is a "bag-of-words approach" instead of a more sophisticated natural language processing technique, many words from the original corpus were related to violent extremism instead of disasters. Creating groundtruth based on tweets about violent extremism may yield different scores and could be an area for future experimentation.

Application of dictionaries

While revising the dictionary development and groundtruth creation has potential to improve the performance of the tool, the tool itself has several issues worth addressing, namely morphology and domain application.

As mentioned in the results, Tagalog has a complex morphology that does not conform to the original LIWC methodology of adding a wildcard after a root. As mentioned previously, examples such as *pinakamalakas* and *bagyong* were words that were not included in the dictionaries and may have changed the scores of the tools. A way to address this issue would be to process the tweets through a morphological analyzer prior to sentiment analysis, and then populate the sentiment lexicon with the root words. Several morphological analyzers for Tagalog exist (Roxas and Mula, 2008; Fortes-Galvan and Roxas, 2006; Bonus, 2003).[4]

With respect to domain issues, Liu (2012) states, "although finding domain-specific sentiment words and their orientations are useful, it is insufficient in practice," and explains a method of including the aspect with the sentiment word to create an opinion context. In this case, one way to improve the accuracy of the Tagalog sentiment analysis would be to create opinion contexts for a disaster, e.g., pairing *ulan* and *malakas* together. Developing rules for intra-sentential and inter-sentential sentiment consistency would also help, such as considering if a statement has adversative expressions such as *pero* (but). For example, if a tweet states "*Ulan malakas, pero walang takot kami*" (Strong rain, but we have no fear), such a statement could be classified as a positive sentiment despite the negative opinion context. It is unlikely that a bag-of-words approach could identify this, which may mean tools other than LIWC are needed.

Conclusions

Developing the Tagalog LIWC positive and negative dictionaries as proof-of-concept was an interdisciplinary process that revealed some of the challenges to consider when creating and using a sentiment lexicon for analyzing Twitter. While the dictionaries were successfully completed based on the LIWC methodology, Tagalog language sentiment analysis would benefit from more sophisticated natural language processing techniques and tools (e.g., better equipped to handle morphological changes) in order to analyze social media data more accurately. Researchers should also carefully consider the guidelines and requirements for creating groundtruth for their data. Finally, it is critical to evaluate these tools before testing them on other data—while the Tagalog dictionaries allow for some interesting sociolinguistic and psycholinguistic perspectives on emotion and sentiment during Typhoon Yolanda, they should be refined and tested more before being applied authoritatively.

Acknowledgements

This project was funded by the Early Career Research Project (ECRP) at MITRE. Many thanks also to the Tagalog lexicographers: Maria Casipe, Risa Mayan, Cheryl Anne Lee-Monahan, Lysa Olimpo, and Carl Rosario. Special thanks also to Alison Dingwall, Theresa Dillon, Karine Megerdoomian, Beth Elson, Jennifer Mathieu, Tod Levitt, Geoff Emmer, Tom Kiley, and Scot Lunsford.

[4] However, upon contacting the authors, it was difficult to obtain use of these analyzers.

Appendix

Search terms used to collect data.
#bagyo
#bangon
#bangonpilipinas
#bangonvisayas
#haiyan
#reliefph
#supertyphoon
#typhoon
#yolanda
#yolandaph
bagyo
bagyong
bangon
haiyan
supertyphoon
typhoon
yolanda
yolandaph

References

Amazon Mechanical Turk. (2014). Retrieved 16 Dec 2014 from https://www.mturk.com/

Andrei, A., Dingwall, A., Dillon, T., & Mathieu, J. (2014). Developing a Tagalog linguistic inquiry and word count (LIWC) 'disaster' dictionary for understanding mixed language social media: A work-in-progress paper. *EACL 2014,* 91.

Balage Filho, P. P., Pardo, T. A., & Aluısio, S. M. (2013). An evaluation of the Brazilian Portuguese LIWC dictionary for sentiment analysis. *Proceedings of the 9th Brazilian Symposium in Information and Human Language Technology,* 215-219.

Bankoff, G. (2002). Cultures of disaster: Society and natural hazard in the Philippines. RoutledgeCurzon: New York, NY.

Bonus, D. E. (2003). The Tagalog stemming algorithms (TagSA). *Proceedings of the Natural Language Processing Research Symposium*.

Brown, R. & Gilman, A. (1960). "The pronouns of power and solidarity." Style in Language. MIT Press: Cambridge, MA. 253-76.

Fortes-Galvan, F., & Roxas, R. A. (2006). Constraint-based morphological analyzer for concatenative and non-concatenative morphology. *Proceedings of the Pacific Asia Conference on Language, Information, and Computation,*

Gonzalez, A. (1998). The language planning situation in the Philippines. *Journal of Multilingual and Multicultural Development, 19*(5), 487-525.

Gunn, J. F., & Lester, D. (2012). Twitter postings and suicide: An analysis of the postings of a fatal suicide in the 24 hours prior to death. *Present Tense*, 27(16), 42.

Hayeri, N., Chung, C., & Pennebaker, J. (2010). The development of linguistic inquiry and word count (LIWC) for Arabic texts. Austin, TX: LIWC.Net.

Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2004),*

Linguistic inquiry and word count. Retrieved 4 Aug 2014 from http://liwc.net/

Liu, B. (2012). Sentiment analysis and opinion mining (synthesis lectures on human language technologies). Morgan & Claypool Publishers.

Nical, I., Smolicz, J. J., & Secombe, M. J. (2004). Rural students and the Philippine bilingual education program on the island of Leyte. In J. W. Tollefson, & A. Tsui (Eds.), *Medium of instruction policies - which agenda? whose agenda?* (pp. 153-176). Mahwah, New Jersey: Lawrence Erlbaum Associates.

Pennebaker, J. W., Chung, C. K., Ireland, M., Gonzales, A., & Booth, R. J. (2007). The development and psychometric properties of LIWC2007. *Austin, TX, LIWC.Net,*

Pennebaker, J. W., Francis, M. E., & Booth, R. J. (2001). Linguistic inquiry and word count: LIWC 2001. *Mahwah: Lawrence Erlbaum Associates, 71*, 2001.

Philippine Statistics Authority. (2014, May). Statistical tables on overseas contract workers (OCW): 2013. Retrieved, 4 Aug 2014 from http://www.census.gov.ph/content/statistical-tables-overseas-contract-workers-ocw-2013

Ramírez-Esparza, N., Pennebaker, J. W., García, F. A., & Suriá Martínez, R. (2007). La psicología del uso de las palabras: Un programa de computadora que analiza textos en español.

Roxas, R. R., & Mula, G. T. (2008). A morphological analyzer for Filipino verbs. *22nd Pacific Asia Conference on Language, Information and Computation,* 467-473.

Servi, L., & Elson, S. (2012). A mathematical approach to identifying and forecasting shifts in the mood of social media users. The MITRE Corporation.

Wolf, M., Horn, A. B., Mehl, M. R., Haug, S., Pennebaker, J. W., & Kordy, H. (2008). Computer-aided quantitative text analysis: Equivalence and reliability of the German adaptation of the linguistic inquiry and word count. *Diagnostica, 54*(2), 85-98.

Zuraw, K. (2006). Using the web as a phonological corpus: A case study from Tagalog. *Proceedings of the 2nd International Workshop on Web as Corpus,* 59-66.