

#iranElection: Quantifying Online Activism

Devin Gaffney

Visiting Researcher, Rensselaer Polytechnic Institute

110 8th Street

Troy, NY, 12180

1-(503)-319-2931

DGaffney@Bennington.edu

ABSTRACT

In a retrospective analysis of the Iran national election protests, the meteoric rise of the Twitter platform as a purported tool of dissidence has become a subject of debate and importance in the burgeoning field of online activism. By using methods of data collection novel to Web 2.0 social media applications, can a finer granularity be achieved in directly measuring the impact of the internet on politics and society?

Keywords

Twitter, Online Activism, Iran Election 2009

1. INTRODUCTION

In a June 2009 interview with TED, media analyst and internet pundit Clay Shirky announced “this is it. The big one[1],” in reference to the apparent meteoric rise of social media technologies following Mir Hossein Mousavi’s repudiation of the Iranian national election results. As the story goes, when the Iranian government enacted a media blackout, nationals, news services, and the international audience turned towards non-traditional means of communication, most notably Twitter. In the immediate days following the blackout, articles with headlines such as CNN’s “Tear gas and Twitter: Iranians take their protests online” presumed the importance and touted the capabilities of services such as Twitter in allowing dissidents communicate not only to the international audience and news services, but amongst themselves in order to protest and organize effectively [2].

Beyond traditional media, the US government was also drawn to the apocryphal Twitter-Iran connection; on June 16, the State Department requested that Twitter postpone updates to the service by “highlight[ing] to them that this was an important form of communication” both in terms of external information exchange as well as internal organization [3].

As illustrated by Evgeny Morozov, “[i]t is easy to see why so many pundits accepted this narrative: they had seen something similar before. The exultant hordes of attractive, obstreperous young people, armed with fax machines and an occasional Xerox copier, taking on the brutal dictators” - in Morozov’s work, Shirky’s tweet analogizing Tehran 2009 to Leipzig 1989 is conjured, and indeed, another similar event seems to bolster such a comparison: in Seattle 1999, it was the cash-strapped, liberal, technologically savvy, globally aware citizenry - the self-described “electro-hippies” - that proved to leverage information and communications technologies (ICTs) for their own protests. Why would it not be assumed that a similarly liberal (or more accurately, pro-Western), technologically savvy and globalized portion of the population, unhappy with the results of the election,

would turn to the unconventional online activist platform in the face of a government crackdown?

Sandor Vegh’s “Classifying Forms of Online Activism” provides a useful and spartan framework with which to analyze cases of online activism [4]. In his study on the World Bank, Vegh identifies three distinct dimensions of online activism: awareness/advocacy, or the generation of sympathetic information, organization/ mobilization, or the planning and deliberation as a result of sympathetic information, and action/reaction, or the result of the planning and deliberation [4].

Although these dimensions necessarily assume a temporal progression of one after another, it is important to understand that in actuality many instances of each dimension could be occurring at once and at different speeds and magnitudes. In researching a given instance of online activism, we can determine both the role as well as the efficacy of the instance by evaluating it’s focus and resultant impact in each of Vegh’s three dimensions.

For Vegh, the case for cyber-protest against the World Bank, and subsequently, its adherence to his categorization scheme, is made from an analysis of specifically curated websites. This near-anthropological methodology is popular in internet studies [5][6][7][8], and is fairly straightforward: by involving oneself in a given instance of online activism, curating a catalog of websites representative of the overall environment in which the particular case occurs, then examining the qualitative nature of the content and utility of these sites, it is possible to identify particularly interesting occurrences, outline general themes, and, as in the case of Van Aelst and Walgrave (2004), illustrate a link topology within the catalog. We could characterize this method as a Web 1.0 method, that is, a method without using any of the data creation/collection techniques that typify Web 2.0.

This type of evaluation, however, requires laborious amounts of primary data collection, knowledge of both the pre-existing environment as well as the boundaries of that environment, and risks omissions of perhaps influential, yet unreferenced or under-represented, sources of primary data. Traditionally, internet studies has been plagued with the problem of a lack of discrete data, or more accurately, a lack of knowledge about the entire ecosystem which provides the discretion of that data. Simply put, it’s hard to quantify a moving target, particularly when the only methods that can be employed are cumbersome and require technical expertise in collecting the data.

When talking about web-based studies, it is technically impossible to ensure that “all” related data sources are included in any particular study, or all relevant communications are included. Ultimately, this leads to the problem of not being able to accurately portray the environment and context in which a given instance of online activism occurs. Although no particular methodology may completely mitigate this, some may do better than others, dependent on the subject. As such, this paper is

largely about the feasibility of a new form of methodology in web science that leverages machine-readable data that has become a hallmark of Web 2.0 sites. The Iran election is, for our purposes, illustrative of this new methodology, used as a concrete example of the ability, but is used primarily in service of explaining the potentially novel results this methodology can yield.

2. WEB 1.0's "MANUAL CURATION" VERSUS WEB 2.0's "AUTOMATIC COLLECTION"

Essentially, we can call the Web 1.0 nearly anthropological methodology of becoming a part of an online community/familiarizing oneself with the network of websites/users/groups, then analyzing some subset of that a method of "manual curation." In the Web 1.0 environment, this system is reasonable: publishing one's own content generally took the form of a personal/group website, and conversations took place in forums or via e-mail. In both cases, one could have automated the collection and analysis of such data, but a simple system of manual detection/selection/analysis seemed to be sufficient to make a point.

Indeed, in many rights, the manual curation method is still preferable for certain situations and contexts; in trying to understand the support network for patients of a particular disease, it would likely make sense to employ this method. This does not, however, mean that it should be the only methodology, or even be the predominant methodology for studying the internet; if the methodology should reflect the case being studied, then this can't be applied to something like the Iran election, and indeed, it is either too cumbersome or not wide-reaching enough to cover many topics of interest.

In Web 1.0 studies such as Smith and Smythe's analysis of the "Battle of Seattle"¹, emphasis was placed on groups, websites, and entire entities, which abstracted the subject to a manageable level, but didn't allow for a very fine granularity. Similarly, most early papers seem distanced from the subject being discussed, and as a result, cannot derive specific conclusions or propose clear models for exactly how the internet augments political and societal processes. In Web 2.0, data has become largely machine-accessible; Twitter is a perfect example. By the very nature of the tagging system, we can quickly identify exact communication transmissions that are of interest, and analyze them accordingly. As such, the level of abstraction from our actual subject does not need to exist; we can collect the actual piece-by-piece conversation, and analyze how actors use the technology and how successful that use is on an actor-to-actor level, not a website-by-website basis.

In other words, whereas Web 1.0's manual curation method allows for an abstracted look at communication, Web 2.0's automatic collection method allows for direct analysis of that communication. In a situation such as the Iran election, one is obviously drawn to understand the case through the social networking sites that were so emblematic of it; if this is the case, then we must leverage the existing tools, namely the Twitter API,

in order to collect more data and therefore have a broader base of understanding what actually happened, who specifically made it happen, and why it went as far as it did. The Web 1.0 method in many respects and for many cases no longer applies; people simply have shifted to the web-application as communications medium over the past decade. These applications employ APIs, RSS feeds, and tagging systems to organize their information; by using automatic collection techniques, we can collect more exact data, more of it, and faster. To continue using the old system would, in certain cases such as this, be foolish.

3. OVERVIEW OF THE SOCIAL NETWORK COLLECTION METHODOLOGY

With this comparison of manual curation/automatic collection in mind, and with the new approaches at hand, it is important to talk about the scope of the current work, as well as the specific methodology in collecting the data. In attempting to understand the role of Twitter in the Iran election, the most straightforward approach was to collect tweets carrying the most seemingly popular "hashtag," in this case, #IranElection, and store the basic metadata associated with the tweet: the time it was posted, the user who posted it and their basic information such as the number of friends and followers, their stated location, and so forth².

Immediately upon realizing the research potential of a data set consisting of a representative amount of "#IranElection tweets," a simple scraping program was written in Ruby to query the Twitter Search API and store the results. The scraping began on June 16, and was able to capture Tweets posted after midnight (Iran local time) the day of the election. This program was run consistently for two weeks, then was sporadically turned on with new features to collect previous tweets missed in initial queries, verify that certain sets were being collected, and so forth. By October 24, the Program had collected 766,263 Tweets across 73,693 Users.

After converting the raw data into a database, analysis was conducted on an array of characteristics. These characteristics fell into three different categories: histogram-based analysis of the Users and Tweets, network-based analysis of the re-tweets³, tentatively called "re-tweet influence maps", and simple language based analysis such as tag clouds.

The histogram-based analysis was rather simple: within the program, two object groupings were created: GraphPoints and Graphs. GraphPoints would be any row of a particular histogram, where the GraphPoint.tag would specify the Graph it belonged to, the GraphPoint.label would represent the x-axis data (time posted, username, number of friends, etc), and the GraphPoint.frequency would represent the value of that row. For instance, with User account creation dates, the GraphPoint.tag would specify this was a "User account creation graph," the GraphPoint.label would specify the date a User joined, and GraphPoint.frequency would specify the number of Users who joined on that date. A Graph's given GraphPoints were then converted into CSV data, which is a widely accepted table data format.

¹ The "Battle of Seattle" is a colloquial term for the events that unfolded in November 1999 in Seattle; the WTO scheduled a round of talks in the city, and were met with significant protests that ultimately stalled the talks; this case has been largely cited as one of the first clear cases of online activism as an efficacious tool for political leverage.

³ A re-tweet, was defined strictly: a re-tweet was counted if it followed the case-insensitive syntax of "rt @{user_name}." If multiple re-tweets were detected in a given message, it was treated as two re-tweets. For example "rt @{user_1} rt @{user_2} test message" would count as two re-tweets, one referencing user_1 and another referencing user_2.

Network-based analysis was a bit trickier, as a system had to be developed that could adequately represent different network mappings of re-tweets based on different granularities. A network map of every re-tweet that occurred in the entire data set would inform us as to who the most influential users were, as well as the general flow of re-tweets from those influential users to “hubs” of people who tend to re-tweet frequently, and then throughout the network. What that particular map would not tell us, though, is the number of messages that occurred within a particular time-space; if we wanted to know the moment at which a particular User became a highly re-tweeted person, it would not be immediately apparent. For this reason, sub network-graphs were developed against day-to-day data, hour-to-hour data, and minute-to-minute data. In this way, for example, we could look at every re-tweet that occurred within one particular minute, then shift further in time to one minute later and see the differences in that map.

From a technical standpoint, every re-tweet was treated as an atomic network: two nodes with a directed edge. The object was defined as a Retweet, where the Retweet.retweeted_user represented the original user being referenced, the Retweet.retweeting_user represented the user referencing them, and the Retweet.edge_id was the Twitter-internal id of the Tweet in which the re-tweet occurred. Retweet objects were then collected into RetweetCollectors, which grouped Retweet objects by different time metrics. For example, RetweetCollector.minute, RetweetCollector.hour, and RetweetCollector.date would specify the Retweet objects to belong to that particular RetweetCollector. Within the network map generation code, queries would filter based on the minute, hour, and date settings to collect subsets and super-sets of the Retweet to view the data at varying granularities. Retweet data was then converted to GraphML files, which is a widely accepted file format for network data. When the networks are actually interpreted, a high out-degree for any particular user’s node corresponds to a high number of other user nodes re-tweeting that user’s content. Conversely, a high in-degree for any particular user’s node corresponds to a high number of re-tweets created by that particular user.


This approach allows us to actually treat the network of influence as one sample study. In previous Web 1.0-era manual curation methods, if this type of analysis were to be conducted, it would have been infinitely trickier, for a multitude of reasons. With Web 2.0 applications and norms, however, the task of generating such information is rather straightforward; once the data is pulled from the API, we can create simple algorithms to leverage it and create such analytical tools. By converting the entire data set into a re-tweet influence map we see the signal, which any researcher could find with enough effort, but we also are able to look at the noise and determine if there is anything valuable in that as well. This noise also contextualizes the signal, and allows us to see if the user base tends to uniformly follow the signal, or actually operates tangential to the signal as would be expected.

It should be stated that this is not the first time such an analysis has been proposed. Indeed, the Web Ecology Project’s laconic analysis of 2,042,166 tweets containing relevant Iran election search terms revealed interesting and novel results that could not have been captured in any traditional manual curation method [9]. By analyzing the data at a finer granularity and separating the process into three categories of analysis has, however, yielded novel results that a cursory study such as the Web Ecology Project’s article did not. As such, it is important to share this deeper analysis in the hopes that it will allow for new insight.

Finally, as a tangential note to the findings, it should be said that the methodology above has been re-factored into a general-purpose utility that uses a distributed-computing approach to manage and collect data requests from researchers with similar data needs in their own studies on the impact of social networks, specifically Twitter, on politics and society. Without going into too much detail, a user-friendly web interface will soon allow researchers to initialize new data requests or “scrapes,” select different methods and depths of analysis, and retrieve analytical findings as well as raw data sets for more specific questions they may have. In short, by abstracting the process used to collect this particular data set, we can generalize a toolkit for conducting these types of studies with relative ease. As work is ongoing, it is premature to release specific details concerning the functionality as of now, but the source code is licensed under creative commons and is freely available as the project continues⁴.

4. FINDINGS: HISTOGRAM-BASED ANALYSIS

Figure 1: Tweet volume per day. The x-axis represents a given date, the y axis is the number of tweets in the data set that occurred on the given date.



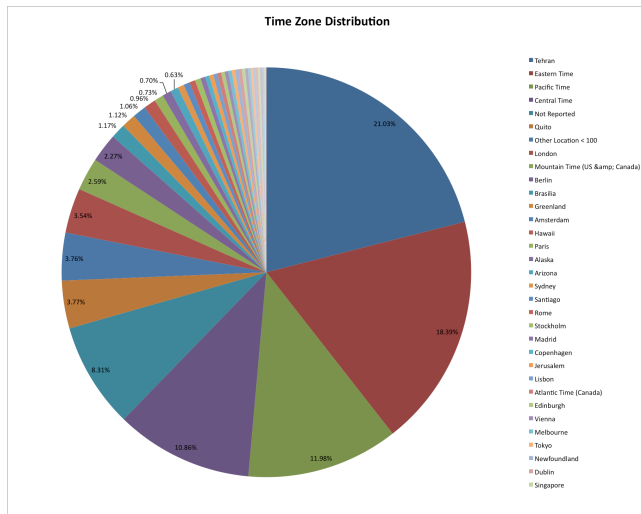
First, a high-level overview of the data is appropriate. As Figure 1 shows, unsurprisingly, the vast majority of Tweets collected in the data set were posted within the few weeks following the election itself. There are interesting outliers, notably the slight spike in traffic on September 18th, 2009, which corresponded with Quds day. Quds day, a national holiday proclaimed by Ayatollah Khomeini in 1979 in order to “demonstrat[e] the solidarity of Muslims world-wide, [and] announce their support for the legitimate rights of the Muslim people” [10], has become an institutionalized propagandistic tool to both bolster political support from muslim masses both domestic and abroad as well as assert the predominance of the Islamic Republic over its neighbors in middle eastern affairs. In the 2009 Quds day demonstrations, a number of election protests were held throughout the country as well as abroad [11]. Another interesting point is the apparent drop in traffic on June 30th, 2009. With this graph alone, and from reading a random sampling of messages from that particular date, a possible cause of this is it is not immediately clear.

⁴ To download the source code and learn more about this portion of the current work, go to <http://github.com/DGaffney/TwitterGrab>

The data at a high level does seem to confirm popularly assumed notions, however, and this quantitative proof allows a level of certainty to be applied to the argument that would not be possible with manual curation: most traffic occurred in the immediate aftermath of the election results, and rapidly decreased. It is likely that a possible cause of the astronomic June 21, 2009 spike was a result of a “Twitter storm”: as User’s tag Tweets with a particular tag, those Users followers become aware of the term (as they likely view their tweets in the aggregate), which in turn pushes them to Tweet or Retweet, and so forth. As such, we can isolate much of our study to the few weeks immediately following the election

Figure 2 shows the breakdown of self-reported locations. Of particular importance in this data set is the fact that “further complexity comes from those outside Iran changing the location on their *Twitter* profiles to Tehran or time zone to +3:30

Figure 2: Breakdown of self-reported locations. It should be noted that one form of slacktivism⁵ employed was changing one’s location information to Tehran in order to “confuse” the Basij and assorted pro-government forces.



in an attempt to create cover for those tweeting in Iran” [12]. Indeed, it is hard to tell how many people changed their information precisely because so many may have. Conversely, it is more likely users *within* Tehran could have changed their own location, as their motivation to not be found by pro-government supporters would be much stronger. For these complicated reasons, the location-breakdown can’t be fully reliable. Additionally since this particular form of slacktivism has now become popularized in this instance, it’s likely that this will remain to be a problem for future geographically specific case studies. In and of itself, however, we can identify this activity as action/reaction, and then discuss to what extent this is actually a useful form of online activism. Whereas Morozov may rightly conclude that slacktivism has no utility, with the actual data sets


⁵ The term “Slacktivism,” or “slacker activism,” has come into popular use to describe activities which require little from the “activist,” yet still provide the feeling that one has done something to help a given cause. In our case, the changing of ones location may give the sense that they are confusing the authorities, but it’s not a proven successful tactic, and it requires little from the user, so it is viewed as a largely slacktivist activity.

available, we can determine what would constitute efficacious online activism, and conclusively label an act like this slacktivism if it does indeed seem to have no impact proportional to the level of attention devoted to it.

At the same time, we can now assume some things in the aggregate: aside from the Tehran anomaly, the vast majority of self-selected locations are firmly located in economically developed, generally European or North American, locations. From this, we can begin to shape a general demographic: people who would be likely to use Twitter already, or a population already well-connected to the internet (and to the social network Twitter tends to attract), and likely urban (due to the frequency of major city self-selection) likely contribute the majority of content, if not the majority of the user-base.

The next chart, Figure 3, is perhaps the most intriguing initial finding for this Tweet data. Burns and Eltham argue that a “surge in users meant that the network was available to significant numbers of people for the first time, allowing these users to mobilise the social media platform to attempt to influence international events” [13]. This “surge” seems to be corroborated by the data; from June 12th to June 25th, or the first two weeks following the election, 11,384, or 16.8%, of the users in the data set, joined Twitter with new accounts. Within those days, the number of account creation increased an order of magnitude, and clearly shows a correlation with the actual content of the conversation. From this, we can safely reason that these were not simply pre-existing users that switched conversation to the hot-button political issue of the month, but were actively joining the network in order to participate. Similarly, we can safely reason that in this case, online activism was not necessarily a primary method of activism until traditional media was broken down; the spike in account creation began the day reporters started being arrested, as the data shows [14].

Figure 3: Account creation histogram. The x-axis represents the day the particular User account was registered, the y-axis represents the number of users registering that day



Julian Bajkowski’s article cites Al-Jazeera’s Head of New Media, Moeed Ahmad, as only attributing 60 user accounts to Tehran proper [15]. This number, according to Ahmad, dropped to only 6 “active” accounts when communications were cut. In other words, if this surge in account creation was not caused by Tehranians, who else could have caused it? A safe assumption would be that the new users had a strong opinion on the matter, and wanted to become involved in the conversation, but this does not provide us

with their demographics or their actual efficacy/role in altering the outcome of the Iranian election. Similarly, the fact that this is a geographically specific situation does not necessarily mean that the interested parties themselves are in that geographically specific region; it could be that the vast Iranian diaspora of southern California is partially responsible for this surge in account creation, for example.

Figure 4: simplified curve illustrating general traffic trends for Twitter as of November 11th, 2009. Set at Central European Time, the data in this graph is offset against the data in our own graph by one hour (Our data is set at UTC). Note the significant dip in the morning (or the American late evening).

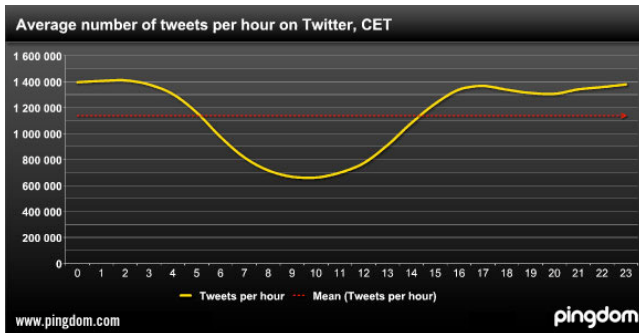
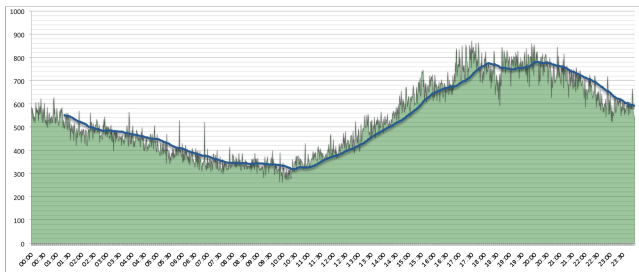


Figure 5: Traffic for data set as compared against general traffic trends; note that the curve, beyond the one hour offset, is very similar to general traffic, suggesting that a representative user base employed the #IranElection hashtag



Figures 4 and 5 show the general traffic trends on Twitter as of late 2009 against the data in this data set. The first graph, from Pingdom.com, a data aggregation service, shows the general traffic from a sample data set collected from October 21, 2009 to November 11, 2009 [16]. Some of the results are clear: The traffic is in phase with the average North American’s sleep and wake schedule, and traffic from other nations does not adequately mitigate that impact. From this, we may draw two conclusions: the impact of North American influence is so considerable to any Twitter data set so that it may be exceedingly rare to see any significant phase shifting, and that any change from the “average” phase shift is likely due to either a new user base or a major event that occurs at a specific time frame.

5. FINDINGS: NETWORK-BASED ANALYSIS

The network analysis is more difficult work, as it currently requires manual analysis after the actual networks are generated. With the current data set, the RetweetCollector process generated 61,838 GraphML files: 135 day-to-day graphs, 24 combined hour graphs, 3,143 hour-to-date graphs, and 58,536 minute-to-minute graphs. In general, it is difficult to work with large network data,

and further work is necessary to streamline the analytical process on the information. In its current implementation, however, interesting results can be yielded.

Figure 6: GUESS visualization generated from #IranElection-hour_date-2009.06.20_17.graphml, or June 20, 2009, at 17, or 5pm UTC (1pm EST, 11:30pm IRST). The listed users, from top to bottom, are persiankiwi, mousavi1388, tedchris, and stopahmadi. Persiankiwi, tedchris, and stopahmadi are well-known “hub” accounts for tweets, and mousavi1388 is Mir Hossein Mousavi’s official Twitter presence.



For analysis, GraphML files are loaded into Network Workbench [17]. Within the environment, the GUESS visualization is initialized, which includes a Python-based interpreter for querying specific nodes and edges. As Figure 6 shows, the tool is invaluable in quickly identifying the most influential “hubs” of communication for a given timeframe. In this particular set, June 20th, 5pm UTC, we can quickly identify persiankiwi, tedchris, stopahmadi, and mousavi1388 as high-degree nodes. Persiankiwi, tedchris, and stopahmadi are well-known “hub” accounts that were influential and pervasive throughout the Twitter re-tweet conversation, and mousavi1388 is Mir Hossein Mousavi’s official Twitter presence. By using basic querying, we can identify the specific Tweets being re-tweeted; for instance, one of the highly re-tweeted messages from mousavi1388 by another user was "RT @mousavi1388 I am prepared For martyrdom, go on strike if I am arrested #IranElection⁶".

In Vegh’s classification, a tweet such as this would likely be considered one of mobilization; by calling for actions to be taken, Mousavi is leveraging the communications platform to quickly and efficiently direct the justifications and basis for action. If, in the process of analyzing the data set, large amounts of this type of signaling is present, then it is reasonable to assume that this case of online activism had a basis of organization/mobilization; the reasons for why it was or was not efficacious becomes the debatable subject.

Interestingly enough, we can also begin to identify different groups of users: re-tweet-ed’s, re-tweet-ing’s, and other. In this particular case, we are identifying persiankiwi, tedchris, stopahmadi, and mousavi1388 as re-tweet-ed’s, with an out-degree (number of messages that re-tweet their content) of 50, 86, 68, and 114, respectively. The emblematic example of the re-

⁶ The source for this tweet is located at <http://twitter.com/mousavi1388/status/2254485463>. The re-tweeting user, who has since privatized their account, could not be seen.

tweet-ing type is the user edward, which re-tweeted 2,262 tweets from other accounts. In the scope of this work, it will be interesting to see if users can gain importance by essentially becoming their own RSS feed, curating a collection of tweets. In the case of edward, despite the frequent re-tweeting, the account only managed to be re-tweeted 83 times over the data set (solidly in the long tail), zero of which occurred at this particular hour. The other accounts are as yet uncategorized: those who are not necessarily influential in either spreading or creating content, but are responsible for the majority of the traffic. Understanding this group will aid in understanding whether or not a “Twitter revolution” argument is even tenable in Iran, which Morozov argues it is not.

Without context, these figures may be of little use, but in the wider spectrum of the whole data-set, we can begin to understand what this implies; for persiankiwi, this particular hour’s 50 re-tweets represents a small fraction of the account’s total re-tweets, 5,917. From this context, we can begin to talk about order of importance for a particular Tweet or timeframe, and identify the relative significance of specific users and their content, which allows us to say with a fair degree of certainty what the users’ effects were in using the social media platform.

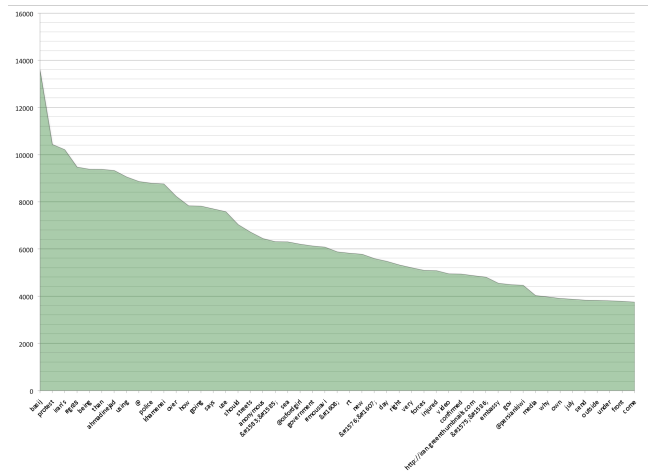
Currently, this method of investigating “re-tweets of interest” via querying is sufficient for attempting to understand the qualitative nature of tweets for a given time. In analyzing a representative amount of tweets in this form, we can generate a rough qualitative understanding of both the demographics as well as the content of the messages, and approach a case study of online activism from a radically novel approach with perhaps surprising results. In the aggregate, however, more research into this network-based analysis and extensions on the current analytical toolset needs to be conducted to properly streamline both the querying and identification process that is generally used in determining “what matters.”

Furthermore, this is still at best a cursory analysis. As stated in the abstract, future work should focus on the difficult question of natural language processing. As argued in Verspoor et. al., by developing an ontology, the data set can be annotated with the proper terms and definitions, and “serve as input to a learning algorithm that aims to generalize from the original examples by determining commonalities among them through their linguistic properties” [18]. By doing so, progress could be made beyond simple “tag cloud” structures and actually hint at overall structure of the conversation. As noted by Morozov, this is a profound challenge, as “by its very design Twitter only adds to the noise: it’s simply impossible to pack much context into its 140 characters” [19], and without that context, NLP would be exceedingly difficult.

6. FINDINGS: LANGUAGE-BASED ANALYSIS

NLP is clearly the preferred method of understanding what is actually being said, and what the messages are actually intending. As research continues, this will be an area of intense exploration. In cursory language-based analysis, however, initial findings can still be particularly useful.

Figure 7: Chart of top 50 words used over all Tweets in data set. The spike at the top term, basij, is worth noting as a possible point of interest.



If we exclude stop words (such as “is”, “it”, “he”, “she”, etc...), we can begin to analyze the qualitative nature of all Tweets through a simple word frequency qualitative analysis. Although it may not clearly demonstrate a particular leaning to any of Vegh’s three areas of online activism classification, it does begin to illustrate exactly what was being said. For instance, by and large, the top word is “basij”, or the Basij paramilitary organization that are, translated, the “volunteers” that are vehemently, sometimes violently, opposed to anti-government protests.

Top 50 Word Frequencies over all Tweets

basij	13604
protest	10437
iran's	10215
#gr88	9476
being	9385
ahmadinejad	9330
using	9057
@	8864
police	8792
khamenei	8763
over	8235
how	7836
going	7823
says	7704
use	7586
should	7038
streets	6718
anonymous	6447
در	6315
sea	6308
@oxfordgirl	6209
government	6133
#mousavi	6085
و	5884
new	5783
به	5598

day	5482
right	5330
very	5210
forces	5100
injured	5087
video	4954
confirmed	4944
http://iran.greenthumbnails.com	4874
از	4812
embassy	4553
gov	4494
@persiankiwi	4468
media	4033
why	3973
own	3909
july	3876
send	3836
outside	3827
under	3809
front	3789
come	3753
kill	3708
gas	3683
!	3676
taking	3619

Interestingly, Neda⁷, which seemed to be a large topic of discussion, is much lower on the list, and does not appear in the top 50. We could then begin to surmise that perhaps the impact of her death may not have been as significant a bolstering for online activists on Twitter as would have been thought; perhaps the fact that it is not as high on the list speaks volumes towards who is actually speaking on the social network.

7. CONCLUSION

Currently, it is difficult to say with any certainty what the role of Twitter was in the Iran election; further analysis and close inspection of certain aspects of the data set is needed. What is clear, however, is that a methodology as outlined above has the ability to produce new insights into the study of online activism. Previous approaches, specifically the “manual curation” method, are useful, but may not appropriately capture the context in which online communication occurs, and tend to lean towards predominately qualitative analysis, which is, while useful, not the only approach that should be used.

Additionally, as the mass amount of online communication now occurs on machine readable platforms, new methodologies should be attempted or adopted in order to possibly capitalize on this fortuitous change. By embracing the new methods employed by Web 2.0 technology, and by leveraging the rapid prototyping languages available for use in such endeavors, it is hard to imagine why not to include this into consideration when conducting research.

By combining Vegh’s framework and this new methodology, interesting conclusions can be gleaned; we can quickly identify the influential users, and more importantly, the influential tweets, using the mass data collection and automatic analysis method, and then look at their content to determine which of Vegh’s three categories their content falls in. By using histogram-based analysis, we can identify at which point the users became most active - if we see that many users joined the network immediately following the election, we can then posit that they are likely not going to fall into the awareness/advocacy category; if they were trying to advocate, they may have been involved early on, advocating for Mousavi well before the post-election fallout. By using even rudimentary tag clouds, we can identify the key terms used in Tweets; by looking at different tag clouds over time, we can perhaps even see terms reflect a general shift from awareness/advocacy towards organization/mobilization, and eventually action/reaction. What is clear, however, is that the evidence so far suggests a demographic of non-iranians generating awareness about the situation; by posting/circulating/retweeting on-the-ground content, and by participating in the conversation, Sandor Vegh’s awareness category is clearly satisfied; more work is needed for the other two categories, however.

This does not necessarily propose that revolutionary insights will be gained; they are merely novel due to the granularity at which those insights can be investigated. By capturing a sample of the network, we can look at the entire environment in which an instance of online activism occurs, and query that data set methodically to gain a better signal from the noise. Although Twitter is a notoriously noisy platform, there remains crucial

⁷ Neda is a reference to Neda Agha Soltan, a 27-year-old woman whose death was captured by video and broadcast by youtube, and then by traditional media. Her death, widely viewed as a result of basij militia, was a rallying cry for protesters [20].

questions for this particular case that need attention: specifically, if this isn’t a “Twitter revolution,” what role does it play? Likely, the results of a further inquiry to the data will show that the main utility of the social media platform was not organization/mobilization or action/reaction, but in the category of awareness/advocacy, particularly with respect to the international audience. “The world is watching you” is a powerful phrase that has been echoed by numerous protestors in trying to achieve their goals; by shedding light via the transmission of imagery and video, Twitter likely plays a more important role than some commentators have given it. In the end, the primary goal of constructing the internet was to create a communications system resilient to attack; investigating how it handles under pressure is of importance to computer scientists, political scientists, and anyone interested in knowing what technologies allow or do not allow significant impacts in politics and society.

8. TO BE STUDIED

Although this methodological approach clearly shows promise, there is a significant question left unanswered in this paper primarily due to its complicated nature: ethics. As of now, there seems to be some controversy over whether or not it is ethical to collect tweets and associated user information. In conducting this research, the ethics discussion was guided by Twitter’s terms of service. Simply put, Twitter informs its users that “What you say on Twitter may be viewed all around the world instantly. You are what you Tweet!” In more legalistic terms, Twitter clearly states in the first paragraph that “The Content you submit, post, or display will be able to be viewed by other users of the Services and through third party services and websites. You should only provide Content that you are comfortable sharing with others under these Terms.” Additionally, Twitter provides robust privacy features that disallow any data collection via the API for users who have made their accounts private.

With this in mind, it is still a controversial subject, however, as Morozov clearly points out the political and possibly life-threatening conditions under which this particular event took place: “As it happens, both Twitter and Facebook give Iran’s secret services superb platforms for gathering open source intelligence about the future revolutionaries, revealing how they are connected to each other. These details are now being shared voluntarily, without any external pressure. Once regimes used torture to get this kind of data; now it’s freely available on Facebook.” Understanding the ethical ramifications of such a study, and whether or not publicly viewable data necessarily constitutes some allowance for researchers to conduct such a study on it is a crucial question, and would require a paper just as long to begin addressing. For this reason, it is something worth looking at both beyond and in complement to this work.

9. REFERENCES

- [1] Anderson, Chris. "Q&A with Clay Shirky on Twitter and Iran." TED Blog. 16 Jun 2009. TED, Web. 18 Jan 2010. <http://blog.ted.com/2009/06/qa_with_clay_sh.php>.
- [2] Nasr, Octavia. "Tear gas and Twitter: Iranians take their protests online." CNN.com 15 Jun 2009: n. pag. Web. 24 Feb 2010. <<http://www.cnn.com/2009/WORLD/meast/06/14/iran.protests.twitter/index.html>>.

- [3] Fleming, Sue. "U.S. State Department speaks to Twitter over Iran." Reuters 16 Jun 2009: n. pag. Web. 24 Feb 2010. <<http://www.reuters.com/article/idUSWBT01137420090616>>.
- [4] Vegh, Sandor. 2003. "Classifying Forms of Online Activism: The Case of Cyberprotests against the World Bank." Pp. 71-96 in *Cyberactivism*, edited by Martha McCaughey and Michael D. Ayers. New York, NY: Routledge.
- [5] Smythe, Elizabeth, and Peter J. Smith. "New Technologies and Networks of Resistance." *Cyber-Diplomacy: Managing Foreign Policy in the Twenty-First Century*. (2002): 48-82. Print.
- [6] Van Aelst, Peter, and Stefaan Walgrave. "New media, new movements? The role of the internet in shaping the 'anti-globalization' movement." *Cyberprotest: New media, citizens, and social movements*. (2004): 97-122. Print.
- [7] Pini, Barbara, Kerry Brown, and Josephine Previte. "Politics and identity in cyberspace: a case study of Australian Women in Agriculture online." *Cyberprotest: New media, citizens, and social movements*. (2004): 259-275. Print.
- [8] Edwards, Arthur. "The Dutch women's movement online: internet and the organizational infrastructure of a social movement." *Cyberprotest: New media, citizens, and social movements*. (2004):183-206. Print.
- [9] Hwang, Tim, et. al. "The Iranian Election on Twitter: The first eighteen Days." Web Ecology Project (2009): n. pag. Web. 24 Feb 2010. <<http://www.webecologyproject.org/2009/06/iran-election-on-twitter/>>.
- [10] "Imam's message announcing Quds Day from Sahifa-yi Nur , Vol. 8, p. 229, dated 7 August 1979," *Islamic Republic of Iran Broadcasting*, accessed Oct. 5, 2007, <http://www.irib.ir/worldservice/imam/palestin_E/10.htm>
- [11] Worth, Robert F. "Despite Warning, Thousands Rally in Iran ." *The New York Times*. 18 Sep 2009. New York Times, Web. 4 Mar 2010. <http://www.nytimes.com/2009/09/19/world/middleeast/19iran.html?_r=1&bl>.
- [12] Fisher, Ali. "Bullets with Butterfly Wings; Tweets, Protest Networks, and the Iranian Election." *Media, Power, and Politics in the Digital Age*, forthcoming. (2010): Print.
- [13] Burns, Alex and Eltham, Ben (2009) *Twitter Free Iran: an Evaluation of Twitter's Role in Public Diplomacy and Information Operations in Iran's 2009 Election Crisis*. In: *Communications Policy & Research Forum 2009*, 19th-20th November 2009, University of Technology, Sydney.
- [14] "Post-election clampdown in Iran ." *BBC News* 15 Jun 2009: n. pag. Web. 4 Mar 2010. <<http://news.bbc.co.uk/2/hi/8100310.stm>>.
- [15] Bajkowski, Julian. "Al Jazeera offers reality check for the Twittiverse." *MIS Australia* 22 Feb 2010: n. pag. Web. 25 Feb 2010. <<http://www.misaustralia.com/viewer.aspx?EDP://1266801386432§ion=newsletter>>.
- [16] "In-depth study of Twitter: How much we tweet, and when." *Royal Pingdom* 13 Nov 2009: n. pag. Web. 25 Feb 2010. <<http://royal.pingdom.com/2009/11/13/in-depth-study-of-twitter-how-much-we-tweet-and-when/>>.
- [17] NWB Team. (2006). *Network Workbench Tool*. Indiana University, Northeastern University, and University of Michigan, <<http://nwb.slis.indiana.edu>>.
- [18] Verspoor, K., Sanfilippo, A., Elmore, M., MacKerrow, E. (2006). "Deploying Natural Language Processing for Social Science Analysis." In: *proceedings 2006 Chicago Colloquium on Digital Humanities and Computer Science*.
- [19] Morozov, Evgeny. "Iran: Downside to the "Twitter Revolution"." *Dissent*. (2009): 10-14. Print.
- [20] "A death in Tehran." *Frontline*. Public Broadcasting Service: 17 Nov 2009. Television.