SOTICS 2014

The Fourth International Conference on Social Eco-Informatics

ISBN: 978-1-61208-372-8

October 12 - 16, 2014

Nice, France

SOTICS 2014 Editors

Pascal Lorenz, University of Haute-Alsace, France

Petre Dini, Concordia University, Canada / China Space Agency Center, China

SOTICS 2014

Forward

The Fourth International Conference on Social Eco-Informatics (SOTICS 2014), held between October 12 - 16, 2014 in Nice, France, was an event on social eco-informatics, bridging different social and informatics concepts by considering digital domains, social metrics, social applications, services, and challenges.

We take here the opportunity to warmly thank all the members of the SOTICS 2014 technical program committee, as well as all of the reviewers. We also kindly thank all the authors who dedicated much of their time and effort to contribute to SOTICS 2014. We truly believe that, thanks to all these efforts, the final conference program consisted of top quality contributions.

We also gratefully thank the members of the SOTICS 2014 organizing committee for their help in handling the logistics and for their work that made this professional meeting a success.

We hope that SOTICS 2014 was a successful international forum for the exchange of ideas and results between academia and industry and to promote further progress in the area of social eco-informatics. We also hope that Nice, France provided a pleasant environment during the conference and everyone saved some time to enjoy the charm of the city.

SOTICS 2014 Chairs

SOTICS General Chairs

Philip Davis, Bournemouth and Poole College - Bournemouth, UK
David Newell, Bournemouth University - Bournemouth, UK

SOTICS Advisory Chairs

Petre Dini, Concordia University, Canada & IARIA, USA
Krzysztof Juszczyszyn, Wrocław University of Technology, Poland
Ulrike Lucke, University of Potsdam, Germany
Abdulrahman Yarali, Murray State University, USA

SOTICS Special Area Chairs on Social Networks

Feng Gao, Microsoft Corporation, USA
Lynne Hall, University of Sunderland, UK

SOTICS Special Area Chairs on eGovernment

Jennifer Watkins, Los Alamos National Laboratory, USA

SOTICS Special Area Chairs on Digital Cities

Tom Erickson, IBM T. J. Watson Research Center, USA

SOTICS Special Area Chairs on Digital Libraries

Paolo Garza, Politecnico di Milano, Italy

SOTICS Special Area Chairs on Media

Claus Atzenbeck, Hof University, Germany / Aalborg University, Denmark

SOTICS Publicity Chairs

Nima Dokoohaki, Royal Institute of Technology (KTH) - Stockholm, Sweden
Christine Langeron, Hubert Curien Laboratory, Lyon University, France
Sandra Sendra Compte, Polytechnic University of Valencia, Spain

SOTICS 2014

Committee

SOTICS General Chairs

Philip Davis, Bournemouth and Poole College - Bournemouth, UK
David Newell, Bournemouth University - Bournemouth, UK

SOTICS Advisory Chairs

Petre Dini, Concordia University, Canada & IARIA, USA
Krzysztof Juszczyszyn, Wrocław University of Technology, Poland
Ulrike Lucke, University of Potsdam, Germany
Abdulrahman Yarali, Murray State University, USA

SOTICS Special Area Chairs on Social Networks

Feng Gao, Microsoft Corporation, USA
Lynne Hall, University of Sunderland, UK

SOTICS Special Area Chairs on eGovernment

Jennifer Watkins, Los Alamos National Laboratory, USA

SOTICS Special Area Chairs on Digital Cities

Tom Erickson, IBM T. J. Watson Research Center, USA

SOTICS Special Area Chairs on Digital Libraries

Paolo Garza, Politecnico di Milano, Italy

SOTICS Special Area Chairs on Media

Claus Atzenbeck, Hof University, Germany / Aalborg University, Denmark

SOTICS Publicity Chairs

Nima Dokoohaki, Royal Institute of Technology (KTH) - Stockholm, Sweden
Christine Langeron, Hubert Curien Laboratory, Lyon University, France
Sandra Sendra Compte, Polytechnic University of Valencia, Spain

SOTICS 2014 Technical Program Committee

Chris Geiger, Duesseldorf University of Applied Sciences, Germany
Christos K. Georgiadis, University of Macedonia, Greece
Lilia Georgieva, Heriot-Watt University, UK
Apostolos Gkamas, University Ecclesiastical Academy of Vella of Ioannina, Greece
Bogdan Gliwa, AGH University of Science and Technology, Poland
William I. Grosky, University of Michigan-Dearborn, USA
Adrian Groza, Technical University of Cluj-Napoca, Romania
Richard Gunstone, Bournemouth University, UK
Panos Hahamis, University of Westminster – London, UK
Lynne Hall, University of Sunderland, UK
R. Gy Hashim, Universiti Teknologi MARA (UiTM) - Shah Alam, Malaysia
Patrick Hoefler, Know-Center GmbH, Austria
Hana Horak, University of Zagreb, Croatia
Yuh-Jong Hu, National Chengchi University, Taiwan
Yun Huang, Northwestern University, USA
Omar S. Hujran, Princess Sumaya University for Technology – Amman, Jordan
Darko Huljenic, Ericsson Nikola Tesla d. d., Croatia
Colette Hume, University of Sunderland, UK
Sergio Ilarri, University of Zaragoza, Spain
Carlos E. Jiménez, IEEE Technology Management Council Spain, Spain
Maria João Simões, Universidade da Beira Interior - Covilhã, Portugal
Dima Kagan, Ben Gurion University of the Negev, Israel
Nina Kahnwald, University of Siegen, Germany
Michail Kalogiannakis, University of Crete, Greece
Georgios Kapogiannis, Coventry University, UK
Charalampos Karagiannidis, University of Thessaly, Greece
Jarmo Kärki, National Institute for Health and Welfare Information Structures and Classifications Unit, Finland
Zinayida Kensche (Petrushyna), RWTH Aachen University, Germany
Christian Kittl, evolaris next level GmbH, Germany
Dave Kocsis, University of Nebraska at Omaha, USA
Dimitris Kotzinos, TEI of Serres, Greece
Tewart Kowalski, Stockholm University, Sweden
Peter Kraker, Know-Center, Austria
Dimosthenis Kyriazis, University of Piraeus, Greece
Renaud Lambiotte, FUNDP - Namur, Belgium
Shou-de Lin, National Taiwan University, Taiwan
Xumin Liu, Rochester Institute of Technology, USA
Jinhu Lu, RMIT University, Australia
Sultana Lubna Alam, University of Canberra, Australia
Richard Lucas, University of Canberra, Australia
Paul Lukowicz, German Research Center for Artificial Intelligence (DFKI) - Kaiserslautern, Germany
Lorenzo Magnani, University of Pavia, Italy

Rudi Vansnick, Internet Society Belgium, Belgium
Nikos Vrakas, University of Piraeus, Greece
Fang Wang, UMASS - Amherst, USA | Nankai University - Tianjin, China
Liqiang Wang, University of Wyoming, USA
Jiang Wei, Missouri University of Science and Technology - Rolla, USA
Huadong Xia, Microstrategy Inc. / Network Dynamics and Simulation Science Laboratory - Virginia Tech, USA
Levent Yilmaz, Auburn University, USA
Eiko Yoneki, University of Cambridge, UK
Ryan Zammit, Middlesex University, UK
Weining Zhang, University of Texas at San Antonio, USA
Wenbing Zhao, Cleveland State University, USA
Xingquan Zhu, Florida Atlantic University, USA

Table of Contents

An Audiobooks-based Approach for Creating a Speech Corpus for Acoustic Models

Salvatore Michele Biondi, Vincenzo Catania, Ylenia Cilano, Raffaele Di Natale, Antonio Rosario Intilisano, Giuseppe Monteleone, Daniela Panno
Dipartimento di Ingegneria Elettrica Elettronica e Informatica
University of Catania
Catania, Italy
{salvo.biondi, vincenzo.catania, ylenia.cilano, raffaele.dinatale, aintilis, giuseppe.monteleone, daniela.panno}@dieei.unict.it

***Abstract*—The limited availability of a valid speech corpus is one of the major problems affecting the design of speech recognition acoustic models. As a matter of fact, large amounts of manually-transcribed data is necessary in order to build a valid acoustic model. Nevertheless, obtaining large datasets is generally both time- and resource- consuming as it requires a continuous supervision of the entire building process. Speech corpora can be used to generate an acoustic model, however a large part of these are not suitable or freely available. This paper aims at showing the use of audiobooks as databases for creating speech corpora. An automatic algorithm that processes audiobooks for building speech corpora is proposed. This method allows to replace traditional manual transcription of audio- recordings and to automatically obtain a phonetic dictionary. An Italian acoustic and linguistic model was generated as use-case to test the effectiveness of the proposed procedure.**

Keywords-Automatic Speech Recognition; Audio Databases; Audiobook; Speech processing.

I. INTRODUCTION

Automatic Speech Recognition (ASR) maps an acoustic signal into a sequence of phonemes or words (discrete entities). A typical ASR process performs the decoding of input speech data by means of signal processing techniques and acoustic or language models. Acoustic models refer to the representation of acoustics and phonetics, while language models describe the words that are recognized by the ASR process.

The complexity of natural language makes its use for human – computer interaction a difficult task to perform [1]. For example, a specific context may influence the generation of phonemes or speech signal can vary significantly according to speaker sex, style, speed and can be affected by noise.

Speech recognition engines require statistical representation of each of the distinct sounds that makes up a word (Acoustic model). This model is created by a training process that compiles recordings and their transcriptions into a statistical representation of the sounds for each word. Therefore, large amounts of transcribed recordings are generally required.

Obtaining these data is an expensive and time-consuming task: several hours of recordings are necessary, recorded in a quiet environment and by different voices. Finding all these data with the relative transcription is not trivial.

An alternative to build from scratch a database with audio recordings and manual transcriptions is the use of audiobooks.

This paper is organized as follows: the prior works are described in Section II, Section III depicts the novel approach, Section IV the procedure of training acoustic models, Section V describes the algorithm developed to obtain the phonetic transcriptions, Section VI shows the results; finally, in Section VII, we draw conclusions and possible future works.

II. RELATED WORK

Different approaches are available to generate an audio corpora database to be used for creating an acoustic model: a first approach [2] that uses speakers with different age and gender for providing recordings related to a predefined text; a second approach that uses available audio sources (radio broadcast, news broadcasts) and the corresponding text transcriptions [3][4]. However, the first method is accurate but time- and resource-consuming, while the second one is easy to develop but it often provides incomplete results because most of its audio data sources have no corresponding accurate word transcriptions.

For example, a complete commercial recording database of Italian language records already exists and is called APASCI [5]. This is an Italian speech database designed for researching on acoustic modeling. The process to create a database using these features from scratch requires much time.

Another available solution is Voxforge [6], a free project that collects corpora of spoken speeches in several languages thanks to the collaboration of users who provide their voices. All audio files are available under the General Public License (GPL) and allow obtaining acoustic models for many speech recognition engines such as Carnegie Mellon University (CMU) Sphinx [7] and Hidden Markov Model Toolkit (HTK) [8]. Sphinx does not provide an Italian acoustic model, thus, it has to be created with a suitable audio database. Beside Voxforge, no Italian free audio databases, which work with the aforementioned engine, were created.

At the time of writing this paper Voxforge contained about 11 hours of recordings for Italian language. Such a small amount of records does not allow to obtain a good acoustic model using Sphinx, because to create a new model 50 hours of audio recordings of 200 different speakers are required [9]. So, the need of a specific audio database for creating an Italian language acoustic model with Sphinx has arisen.

III. A NOVEL APPROACH

As described above, the acoustic model is a building block of an ASR system and it can manage a specific language only if an acoustic model for that language is available.

In this paper, a novel approach for creating an acoustic and linguistic model for Sphinx is described. It provides a method to obtain in a simple and fast way many transcribed recordings from audiobooks. Audiobooks are a valid and reliable source to create acoustic models because they are recorded in an echoic chamber by different voices and the text of the audiobook is available. Audiobooks are a good and free statistical basis in terms of cost / time, but the problem is that, in general, the training for creating an acoustic model requires small audio files in terms of duration (for example the optimal length for Sphinx is not less than 5 seconds and not more than 30 seconds [9]). Because audiobooks provides audio files that are too long for Sphinx, a method to split audiobooks and to associate to each part obtained the corresponding transcription is necessary. The tool HTK was used for this purpose. To solve the problem of the creation of acoustic models, we need to provide the Italian phonetic transcription of each word in audio files. An algorithm that creates phonetic transcriptions has been developed For the training of the Italian acoustic model we used SphinxTrain [9] (it is the tool used for the training of an acoustic model for Sphinx).

This method can be applied for the creation of acoustic models in several languages.

IV. TRAINING BY AUDIOBOOK

Audiobooks are usually available like a single long audio file with the corresponding text transcription. For our research, we created an automatic method to develop a training set for SphinxTrain from audiobooks. So, this method splits a single audiobook in several little audio files with the corresponding text transcription.

These audio files make the Audio Database required from sphinx to extract statistic from the speech.

HTK toolkit [10] was used to split the audiobooks. HTK [11] requires a little acoustic model to perform the work described above. This acoustic model was developed by using VoxForge training through a set of collected transcribed speech corpora. During tests, VoxForge held a database with about 11 hours recording with related text transcriptions. Usually, an audio database with 11 hours of

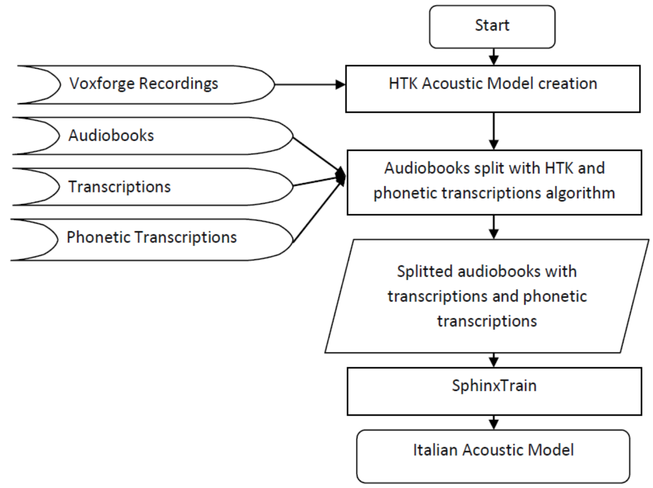

Figure. 1. Procedure of training

recording is not enough to create an efficient acoustic model, but it was sufficient to use in our automatic split method through HTK.

Our test shows that this database is a valid choice for our task. In fact, we use this first model to split the selected audiobooks because it allows to align the audio to its transcription.

Next, the phonetic transcriptions of the words must be provided: for this purpose, the algorithm described in Section 5 has been developed. Finally, audio recordings obtained by the split audiobooks and their phonetic transcriptions are input for SphinxTrain that produces an Italian acoustic model. Figure 1 shows the whole procedure.

V. PHONETIC DICTIONARY GENERATION

The training function must be configured with the sound units, the corresponding transcriptions and the phonetic dictionary. It maps every word into a sequence of sound units (phonetic transcription). To derive the sequence of sound units, the phonetic transcriptions are associated with each word:

```
ABACO      a b a k o
ABACHI     a b a k I
```

In the example above, the word ABACO is the Italian for "Abacus", while ABACHI is its plural. At their right, we show their corresponding phonetic translation.

Initially, it uses an existing dictionary that consists of a customized version of the Festival [12] lexicon (brackets and number were removed). It includes more than 440,000 words, but it is not complete enough to perform the phonetic transcriptions of all the words present in the split audio files. Accents are important for Italian language because the meaning of the word may change based on their position.

So, an algorithm to derive the missing phonetic transcriptions is necessary. The developed algorithm is able to derive two different phonetic transcriptions. The input of the algorithm consists of a list of words not found in the Festival lexicon and for each word it performs the following steps:

- [optional] an online search of the word in an Italian dictionary-database [13]: if the word is found in the dictionary the position of the accent is obtained. In fact, if the word is a lemma (or entry-word) missing in the phonetic dictionary, the algorithm generates the phonetic transcription with the correct accent. This step is optional since calling the web service takes a considerable time. Furthermore, grandidizionari.it service does not contain all the necessary words, so:
- If the accent was not obtained, the lemma from which the word is derived is looked up in the dictionary of Morph-IT! [14], that is a lexicon of inflected forms with their lemma and morphological features. Output of this module is a tuple composed by:

1. Form
2. Lemma
3. Features

- At this point, the substring with which the word ends is compared with a list of available desinences written according to the following format:

```
vowel_or_consonant+desinence,category,
part_of_speech,phonetical_transcription
```

-`vowel_or_consonant` indicates if "`desinence`" must be preceded by a vowel, a consonant, or both. This parameter is optional and can take 3 values: V (vowel), C (consonant), VC (vowel plus consonant).

- `desinence` is the analyzed desinence, that is the string compared with the substring with which the word ends. The desinences have been obtained from [15]. For each desinence, its inflected versions are obtained by analyzing the `category`.

- `category` is used to get the inflections of each desinence. A letter represents each category. There is a list of categories in which each category is associated with some characteristics. Each characteristic is written according to the following format:

$$\texttt{category: } s_{1,1}\text{-}pt_{1,1} > si_{1,1}\text{-}pti_{1,1}, \ldots, si_{1,n}\text{-}pti_{1,n}; \ \ldots \ ; s_{k,1}\text{-}pt_{k,1} > si_{k,1}\text{-}pti_{k,1}, \ldots, si_{k,n}\text{-}pti_{k,n};$$

(where s=substring, pt=phonetic transcription, i=inflected).

In general, if a desinence belongs to a certain category and it ends in $s_{k,1}$, to obtain the inflected form, $s_{k,1}$ is removed from the desinence and it is replaced with $si_{k,1}$, $pt_{k,1}$ is removed from its phonetic transcription and it is replaced with $pti_{k,1}$. This procedure is applied for all eventual n inflections.

For example, consider the "D" category that is written as:

```
D:gia-dZ i! a>gie-dZ i! e;cia-tS i! a>cie-tS i! e;
```

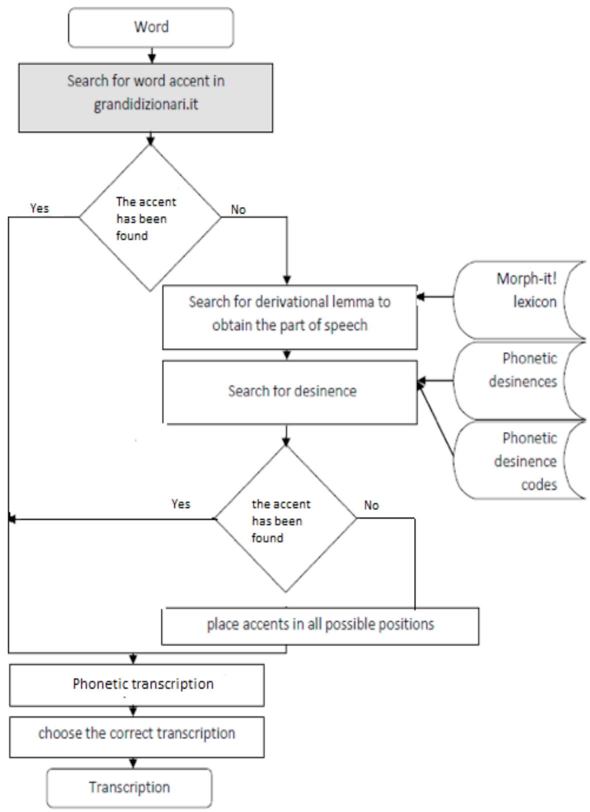

Figure. 2. Algorithm for phonetic transcriptions, in gray the optional step

This means that if a desinence belongs to the category D, if it ends in "-gia" and the correspondent phonetic transcription ends in "-dZ i! a", the inflected form is obtained by removing "-gia" and adding "-gie " in the desinence and by eliminating " dZ-i! a" and adding "-dZ i! e" in the phonetic transcription.

There is also the "I" category for desinences that do not have inflected forms: `I:;`

- `part_of_speech` indicates the part of speech and is used because some equal desinences endings have different accents emphasis depending on the part of speech. The used parts of speech are: V for verbs, N for nouns and adjectives, A for other;

- `phonetical_`transcription is the phonetic transcription of the desinence.

If no desinence has been identified and the word does not have an accent, a set of strings is created. This set

 ISBN: 978-1-61208-372-8

includes all possible inflected forms for the analyzed word. Transcriptions are created for each of these forms and the user can choose the correct one. The entire procedure is shown in Figure 2.

For the Italian language, the Italian phonemes are used [16][17] with the following simplification: close-mid front unrounded vowels and open-mid front unrounded vowels are mapped as a single phoneme “e”. The same simplification is applied to the close-mid back rounded vowels and the open-mid back rounded vowels, mapped in a single phoneme “o”.

The algorithm provides the phonetic transcription of a word both with accents and without (setting by the software), and the user can choose which version to take into account. The word to be phonetically transcribed is analyzed character by character. Optionally, the software cannot take into account the position of accent and therefore does not generate a phonetic transcription that includes the accent, skipping the entire branch "no" in Figure 2 of the first if statement (including grandidizionari.it service).

VI. EXPERIMENTAL RESULTS

The goal of this test was not to validate the Sphinx acoustic model we developed for the Italian language, since no other models were available for comparison. Instead, the purpose of the test was to demonstrate that the approach based on audiobooks can deliver a valid acoustic model. We expect that increasing the size of the training set should provide higher performance.

A first acoustic model with HTK has been created by using Italian VoxForge database. The splitting algorithm has been applied to the free audiobooks [18]. We obtained about 34 hours of split recordings, 56% of them being spoken by female voices. Audiobooks were read by six different male voices and two different female voices.

The automatic phonetic generation gave the same results of querying the “grandidizionari.it” web service. In addition, we obtained phonetic translation for words not provided by the online dictionary.

Finally, a linguistic model was obtained using the transcriptions of audiobooks. To test the performance of the acoustic model two different speakers (a male and a female) pronounced a list of 40 words taken from the vocabulary of the linguistic model. Table 1 shows the results:

TABLE I. RESULTS

	speaker 1 (female)	speaker 2 (male)
recognized words %	70%	77.50%

The acoustic model was trained with two female and six men voices. Although female voices are 56% compared to the total, the male voice is better recognized than the female one. The results show that the variety of voices affects the quality of the result [19]. In addition, we noted that audio recordings that did not match with corresponding words are recognized as very similar words from the phonetic point of view (for example, the word “velocemente” is recognized as “velatamente”).

VII. CONCLUSIONS AND FUTURE WORK

This paper shows an automatic method to obtain recordings, transcriptions and phonetic transcriptions in order to create an acoustic and linguistic model. Our goal was to investigate the possibility of using a set of free audiobooks for generating a dataset as a complete database of ad hoc audio recordings. Then, an Italian acoustic model has been created by SphinxTrain, actually not available in [20]. Currently, the performed tests are based on approximately 30 hours of recordings.

In order to obtain an enhanced Italian acoustic model, a selection of at least 50 hours of audiobooks recordings is required. The audiobooks recordings must be selected by different voices, ages, genders of the speaker or the topic of the audiobook. In total, about 200 different speakers are needed.

ACKNOWLEDGEMENTS

Authors were in part supported by the Sicily Region grants PROGETTO POR 4.1.1.1: "Rammar Sistema Cibemetico programmabile d’interfacce a interazione verbale".

REFERENCES

[1] M. Forsberg, "Why is speech recognition difficult?" Technical Report from Department of Computing Science Chalmers University of Technology, Sweden, 2003.

[2] VoxForge, http://www.voxforge.org (last visited March 5, 2014).

[3] L. Lamel, J. Gauvain, and G. Adda, "Lightly supervised and unsupervised acoustic model training", Computer Speech & Language, Spoken Language Processing Group, CNRS-LIMSI, Orsay, France, Volume 16 Issue 1, 2002, pp. 115-129.

[4] C. Cieri, D. Graff, and M. Liberman, “The TDT-2 Text and Speech Corpus” in In Proceedings of Darpa Broadcast News Workshop, 1999, pp. 57–60.

[5] APASCI, http://catalog.elra.info/product_info.php?products_id=168 (last visited February 25, 2014).

[6] VoxForge repository, http://www.repository.voxforge1.org/downloads/it/Trunk/Audio/ (last visited March 5, 2014).

[7] K.-F. Lee, H.-W. Hon, and R. Reddy, "An overview of the SPHINX speech recognition system, "IEEE Transactions on Acoustics, Speech and Signal Processing, Vol. 38, No. 1, January, 1990, pp. 35 - 45.

[8] HTK Documentation, http://htk.eng.cam.ac.uk/docs/docs.shtml (last visited February 25, 2014).
[9] Training Acoustic Model for CMUSphinx, http://cmusphinx.sourceforge.net/wiki/tutorialam (last visited March 3, 2014).
[10] HTK Toolkit and tutorial, http://www.voxforge.org/home/dev/autoaudioseg (last visited March 3, 2014).
[11] HTK JuliusTutorial, http://www.voxforge.org/home/dev/acousticmodels/linux/create/htkjulius/tutorial (last visited March 4, 2014).
[12] Festival website, http://festvox.org/ (last visited March 3, 2014).
[13] Grandizionari online service, http://www.grandidizionari.it/Dizionario_Italiano.aspx?idD=1 (last visited March 3, 2014).
[14] E. Zanchetta and M. Baroni, "Morph-it! A free corpus-based morphological resource for the Italian language" , proceedings of Corpus Linguistics 2005, University of Birmingham, Birmingham, UK.
[15] Desinences list, http://www.dipionline.it/guida/docs/DiPI_Terminazioni_Desinenze.pdf (last visited March 3, 2014).
[16] F. Albano Leoni and P. Maturi, “Manuale di fonetica” ,2002, Roma, Carocci.
[17] P. Maturi, “I suoni delle lingue, i suoni dell'italiano”, 2006, Bologna, Il Mulino.
[18] Free audiobooks repository from librivox, http://librivox.org/ (last visited March 3, 2014).
[19] M. Gerosa, D. Giuliani, and S. Narayanan: "Acoustic analysis and automatic recognition of spontaneous children’s speech", In Interspeech-2006, paper 1082-Wed2CaP.9.
[20] SPHINX free language model repository, http://sourceforge.net/projects/cmusphinx/files/Acoustic%20and%20Language%20Models/ (last visited March 3, 2014).

Electronic Word-of-Mouth Spread in Twitter as a Function of Message Sentiment

Benny Bornfeld, Sheizaf Rafaeli, Daphne Ruth Raban
Faculty of Management
University of Haifa
Haifa, Israel
bennyb@ruppin.ac.il, sheizaf@rafaeli.net, draban@univ.haifa.ac.il

***Abstract* - Which is more viral, positive or negative electronic word-of-mouth? Can you tell which of the following tweets will spread more in Twitter: "saw the movie … last night, must see it" or "saw the movie … last night, avoid at any cost"? This study is about electronic Word-Of-Mouth spread as a factor of its sentiment. Some theories support a negative bias, while others support a positive bias. Some suggest that both biases are possible, depending on the product type. This paper presents the main theories, related studies and the results of quantitative research based on movie related tweets containing sentiment polarity. Due to the dual nature of Twitter, as mass medium and social network, the research focuses on Twitter's social sub network which contains ordinary users having a small to medium number of followers. The main findings are that tweets with positive sentiment polarity spread 15-20 percent more than tweets containing negative sentiment polarity.**

Keywords-Information flow; Information Dissemination; Sentiment Analysis; electronic Word Of Mouth; Twitter.

I. INTRODUCTION

Word-of-mouth (WOM) is known to have a strong influence on the user's purchase decision. In addition to marketing aspects, message spread or virality is important for intellectual, learning and political reasons. Several books on this topic were recently published; amongst them are "Going Viral" [1] and "Memes in Digital Culture" [2].

The recognition of the importance of WOM in the two-step flow theory dates back to Lazarsfeld and Katz [3]. Electronic Word-Of-Mouth (eWOM) is an important product-related message spreading mechanism. The Internet-based WOM, eWOM, travels fast and can potentially reach very large audiences. One of the most salient Internet services today is Twitter. Twitter is a powerful platform for spreading many kinds of messages, including eWOM. Some eWOM messages carry a negative sentiment polarity and others carry positive or neutral polarity. Does message sentiment polarity influence the extent of message spread? The old marketer's belief that "bad is stronger than good" dominated the pre-Internet WOM era. Is this negative bias still dominant in the Internet social networks of today? Or is the spirit of Facebook's only “Likes” and no “dislikes” catching in the eWOM communication?

Looking at the question of eWOM spread and influence as a function of the eWOM sentiment polarity, there are several theories and evidence, elaborated in the next sections, which provides support for both directions. Those contradicting directions were the trigger for this study, hoping to make a contribution to this open question.

The rest of the paper is structured as follows: In the theoretical background, we present differing views by the Negativity Bias and Product Moderator Theory, followed by an overview of the data source, Twitter. In the related work section, we discuss and summarize related work on this and similar topics. We continue with presenting the research hypotheses and the method followed by a discussion of the finding. We finalize with conclusions and future work section.

II. THEORY

Looking at the question of WOM spread and influence as a function of the WOM sentiment polarity, several theories and evidence provide support for both directions. Some theories postulate that negative is more influential and some claim that positive is more influential. Others state that both directions are valid and the effect depends on the type of the product. In principal, message dissemination depends on two factors: the message value and the messenger’s preference. The information value theories presented here are (1) the negativity bias and (2) the product moderator theory which distinguishes the bias according to different product types.

A. *Negativity bias*

When examining the literature and theory related to the WOM and sentiment polarity, there is strong evidence for the negativity bias. Several articles [4][5][6] show evidence in support of the WOM negativity bias for products and services. They report that negative WOM is twice to 4 times stronger than positive WOM. These sources fit the statement "Bad is stronger than good", which is the title of an article [7] . They state that "The greater power of bad events over good ones is found in everyday events, major life events (e.g., trauma), close relationship outcomes, social network patterns, interpersonal interactions, and learning processes. Bad emotions, bad parents, and bad feedback have more

 ISBN: 978-1-61208-372-8

impact than good ones, and bad information is processed more thoroughly than good". Cheung and Thadani [8] provided a mapping of studies showing the prevalence of the negativity bias over a broad range of areas from information processing, memories, feedback, emotion, marital status, WOM, impression formation, choice, value and frames and customer satisfaction. Baumeister et al. [7] present a social evolutionary argument that this cross-area prevalent phenomenon shows that it is a human adaptive mechanism. They postulate that "The relative strength of bad over good is an adaptive response of the human organism to its physical and social environment. In view of how pervasive the relative strength of bad is, it seems unlikely that this pattern is maladaptive". They explain why it is more adaptive in the following manner -"bad events signal a need for change, whereas good ones do not. If satisfaction and pleasure were permanent, there might be little incentive to continue seeking further benefits or advances. The ephemeral nature of good feelings may therefore stimulate progress (which is adaptive). If bad feelings wore off, however, people might repeat their mistakes, so genuine progress would best be served by having the effects of bad events linger for a relatively long time." More evidence and possible explanations to the negativity bias is summarized in the work of Rozin and Royzman [9].

Another argument in support of the negative bias is the rarity argument. This argument claims that since negative information is rarer, it is highly informative by definition. In order to support the rarity argument, we first need to describe and validate the positivity dominance in languages. Rozin and Royzman [9] provide a summary of evidence based on the work of Osgood [10], showing that positive adjectives are used more frequently. They present a study done on 17 languages, which validated the positivity bias. Quantitative analysis studies by Blenn et al. [11] and Asur and Huberman [12] observed that there are more positive than negative tweets. The biased feedback features provided by social networks platforms like Facebook's “Like” and Google's “Plus One” contribute to the overall positive polarity in social networks. Back to the rarity argument, given the positivity bias in language, negative information is rarer and therefore more informative.

The reliability argument: In many online systems, such as recommender systems, the anonymity of the writer makes the user suspicious as to the credibility of the information, especially towards positive information. Lam et al. [13] studied the credibility issue and gave the following example: “consider a dishonest seller on eBay who wishes to increase his feedback score. He could create a large number of identities and use them to leave himself positive feedback.” According to attribution theory [40], the reader may attribute the positive information to the reviewer self-serving or other non-product-related reasons, leading him to discount positive information.

B. Product type moderator theory

Trying to settle the contradicting findings in different studies, several studies postulated and provided evidence that the effect is subject to the product type. Two contemporary studies, of “hedonic vs. utilitarian” and “promotion vs. prevention consumption” products, relate to the same basic factor of product type. Zhang et al. [14] draw on regulatory focus theory and propose that: "the consumption goals that consumers associate with the reviewed product moderate the effect of review valence on persuasiveness". Higgins et al. [15] who phrased the regulatory focus theory provide the following description of their theory: "A promotion focus would involve a state of eagerness to attain advancement and gains where as prevention focus would involve a state of vigilance to assure safety and non-losses".

Based on Attribution Theory, Sen and Lerman [16] examined the usefulness of published consumer reviews for the reader. They claim that: "trust that the reviewer’s opinions are based on external (product, or other related aspect) and not internal (subjective, or reviewer related) reasons will determine the review’s usefulness to the reader." Attribution theory examines whether the reader attributes the reviewer’s opinions to product related motivations, or believes that they are motivated by self-serving or other non-product-related reasons. The authors find that: "compared with the utilitarian case, readers of negative hedonic product reviews are more likely to attribute the negative opinions expressed, to the reviewer's internal (or non-product related) reasons; and therefore are less likely to find the negative reviews useful". Referring to the previously mentioned study, hedonic products map to product with promotion consumption goals and utilitarian products map to products with prevention consumption goal.

C. The Messenger’s preferences

Messages spread only if the messengers decide to pass them. In WOM communication, the messengers have the choice to pass the message or not. One factor on their pass-or-not decision would be the message value, which was discussed in the previous paragraph. Another important factor is their subjective preference. In his seminal book, "The Presentation of Self in Everyday Life", Goffman writes about how people present themselves in a way they want to be perceived by others. Berger and Milkman [17] stated that “Consumers often share content for self-presentation purposes [18] or to communicate identity, and consequently positive content may be shared more because it reflects positively on the sender. Most people would prefer to be known as someone who shares upbeat stories or makes others feel better rather than someone who shares things that make people angry or upset”. People often relate the message to the messenger. This is the perception in the roots of the phrase "shooting the messenger". Relating this theory to the question of eWOM dissemination in social networks, people will prefer to repeat positive eWOM in order to present themselves in a positive way.

D. Theory summary

Message dissemination depends on the information value and the messenger preference. Regarding the information value, there are more theories and evidence in support of a negative bias. The prevalence of the negativity bias across many areas, the adaptive argument, the rarity argument and the reliability issue associated mainly with positive eWOM lead to the hypothesis that negative messages spread more. On the other hand, the reliability is less of an issue in Twitter because the follower knows who sent the message. The regulatory focus theory, present the moderator role of product type. It predicts a positive bias for hedonic promotion consumption goal products. Since movies are definitely a hedonic, promotion consumption goal product it predicts that negative information will be "discounted". Regarding the messenger preference, the need to present oneself in a positive light contributes to a positive bias.

E. The research framework

The research question is: Does the eWOM message sentiment polarity influence the extent of message dissemination?

We studied the dissemination of eWOM in Twitter as a function of message sentiment polarity. This study examined the spread of tweets via the retweet mechanism in the movies domain. Hence, the presentation and discussion of related work focus on studies which share one or more of this research characteristic.

F. Twitter

Twitter is currently the most popular microblogging service. Microblogging is a broadcast medium in the form of blogging, which differs from typical blogging by its small content length. Twitter enables its users to send text-based posts called tweets to their followers. A tweet length is up to 140 characters. Users may subscribe to other users' tweets - this is known as “following” and subscribers are known as “followers”. Tweets are publicly visible by default. Twitter carries hundreds of millions of tweets per day.

Like social network sites, profiles are connected through an underlying articulated network, but these connections are directed rather than undirected. The number of user’s followers varies from zero to millions. Among the many ordinary users which have up to hundreds of followers there are some highly followed users. Those highly followed celebrities, politicians, news channels, corporations and others use this network as a mass medium communication channel. Participants have different strategies for deciding who they follow - some follow thousands, while others follow few. Some follow only those that they know personally, while others follow celebrities and strangers that they find interesting. In the following section, we further discuss if Twitter can be considered to be a social network.

Several social conventions were introduced and then embraced by the service users’ community itself. The most notable conventions are:

1. Reply: a way to reply to, or to mention another user. Syntax: @user (e.g., @barackobama)
2. Hashtag: a way to indicate the message topic. Syntax: #topic (e.g., #iranelections)
3. Retweet: forwarding others messages to your followers. Syntax: RT @user (e.g., RT @ladygaga). This spreading mechanism plays a major role in this research method and in many other studies and services. We elaborate on its role when answering question 5 in the next section.

III. RELATED WORK

Related work covered in this section presents studies which asked similar questions or methods. This review aims to bridge theory, methods and their findings.

The related work is divided into the following research questions:

1. From studies related to valence effect on product eWOM: What is the effect of positive versus negative product eWOM on message dissemination and influence in Twitter and in recommender systems?
2. Moving from product related information to the neighboring non-product information dissemination, several studies researched the following question: What is the effect of positive versus negative polarity on non-product related message dissemination and influence in Twitter and other media channels?
3. To show the practical relevance of this question, several studies asked: Does eWOM polarity influence product sales?
4. To validate the selection of Twitter as the data source, several studies examined the question: Is Twitter a social network?
5. Due to the key role of the retweet mechanism, several studies explored the question: What are the roles and characteristics of the retweet mechanism?

1. Spread and influence of positive and negative product eWOM in Twitter and in recommender systems.

The question of dissemination and influence of messages within online social network was explored from three different angles: the nodes, the arcs, and the substance.

1. The networks' node angle is focusing on the people and asking: who is influential?
2. The network structure angle examines the ties (arcs) between people and exploring how the networks structures affect the spreading.

3. The content angle examines spread as a function of the message content attributes, such as topic, emotion and sentiment (this is this research angle)

Barbagallo et al. [19] studied tweets containing tourist information about the city of Milan in Italy. They found that negative posts are retweeted more. Sen et al. [16] researched online reviews and found support for the negative bias. In addition, they observed that: "in the case of hedonic products however, readers were more likely to discount than value the negative reviews. Readers found 72% of reviews “not helpful” as compared to 28% being “helpful”".Jansen et al. [20] Twitter-based eWOM research covered products from several industries. They found that on average 50% of the tweets were positive and 33% were critical of the company or product. Zhang et al. [14] conducted an experiment in which they measured the reaction to positive and negative Amazon product reviews. The reviews covered two types of products: a promotion consumption goal product and a prevention consumption goal product. In accordance with Attribution Theory, they found that: "For products associated with promotion consumption goals, consumers show a positivity bias, whereby they rate positive reviews as more persuasive than negative ones. Conversely, consumers show a negativity bias for products associated with prevention consumption goals". In the preventive consumption product, the experiment participants were suspicious towards positive reviews. One common perception is that some of those reviews might be written by non subjective reviewers, such as the product seller. On the other hand, negative reviews for promotion consumption product were attributed to the reviewer’s subjective perspective. With regard to Attribution Theory, there is a difference between classic recommender systems and Twitter. In recommender systems, the reader has no information/acquaintance with the review writer. Therefore, she derives the attributes from cues in the review content. In Twitter, the reader is presumed to be familiar with the person she is following who sent the tweet.

2. What is the effect of positive versus negative content on non-product related message (e.g., news) dissemination and influence in Twitter and other media types?

Several scholars examined the dissemination of non product related content, such as news, articles and phatic communication. Berger and Milkman [17] studied the spread of NY Times articles by email. They found that the spread is related to activation "Content that evokes either positive (awe) or negative (anger or anxiety) emotions characterized by activation (i.e., high arousal) is more viral. Content that evokes deactivating emotion (sadness) is less viral". Stefan and Dang-Xuan [21] studied German politics related tweets and found that emotionally charged Twitter messages (positive or negative) tend to be retweeted more often and more quickly compared to neutral ones. Hansen et al. [22] Twitter based research found that negative news and positive phatic communication are more viral. They proposed that "if you want to be cited: Sweet talk your friends or serve bad news to the public". Somewhat contradicting results were presented by Bakshy et al. [23] . They found that tweets containing URLs linking to positive stories where more dominant in the top retweeted list. Thelwall et al. [24] studied tweets peaks around large events. They observed that "popular events are normally associated with increases in negative sentiment strength and some evidence that peaks of interest in events have stronger positive sentiment than the time before the peak". The rise in both positive and negative sentiment is at the expense of neutral tweets. Another interesting observation by this research supports the writers subjectivity claim: "a surprisingly small average change in sentiment associated with popular events (typically 1% and only 6% for Tiger Woods’ confessions) is consistent with events affording posters opportunities to satisfy pre-existing personal goals more often than eliciting instinctive reactions".

3. How does eWOM polarity influence sales?

The very practical question regarding the relation between eWOM and sales was addressed by several studies. Many of them chose to focus on the movies industry. Some studies took the challenge of solving the eWOM and sales chicken and egg question, using time series analysis. Some studies aimed at finding sales predictions based on eWOM characteristics, such as influencers, overall chatter and message sentiment. Some of those studies are based on recommender systems while others are Twitter based.

Liu et al. [25] found that positive Twitter WOM increases movie sales while negative WOM decreases it. They divided the tweets to pre-consumption (e.g., I want to watch the movie) and post consumption (e.g., the movie was….). They found that the strongest effect on movie sales comes from pre-consumption tweets where the authors express their intention to watch. Asur and Huberman [12] presented evidence that although eWOM volume is the main predictor for movie sale, sentiments extracted from Twitter can be utilized to improve the forecasting power of social media. Addressing the chicken and egg problem [26] conducted a study based on reviews from recommender web sites. They state that WOM is both a precursor to and an outcome of retail sales and that WOM polarity significantly influences the WOM volume. Those studies, showing the relation between sentiment, WOM and sales provide ground for tying the terms of influence and spread.

4. Is Twitter a social network?

Twitter's popularity, the buzz around it, the open nature of its communication and the opportunity it provides for computational social science research has made it a fertile ground for scientific research. Twitter combines characteristics of both mass media, broadcasting news and advertisement and characteristics of social network with relations and interaction between the users. The question of the nature of Twitter and how its users perceive it has

 ISBN: 978-1-61208-372-8

implications to this research and others which explore phenomena within social networks.

The 2010 article "What is Twitter, a Social Network or a News Media?" by Kwak et al. [27] addressed this basic question by conducting a large scale quantitative analysis. They showed that Twitter is not a classic social network: "In its follower-following topology analysis I have found a non-power-law follower distribution, a short effective diameter, and low reciprocity, which all mark a deviation from known characteristics of human social networks [28].Among reciprocated users we observe some level of homophily". Huberman et al. [28] studied the interaction between users on Twitter and came to the following conclusion: "most of the links declared within Twitter were meaningless from an interaction point of view. Thus the need to find the hidden social network; the one that matters when trying to rely on word of mouth to spread an idea, a belief, or a trend".

Other studies took a more qualitative approach. Gruzd et al. [29] conducted a case study on Barry Wellman's twitter followers and friends and stated: "there is a possibility that Twitter can form the basis of interlinked personal communities—and even of a sense of community. The analysis of Barry's Twitter network shows that it is a basis for a real community, even though Twitter was not designed to support the development of online communities". Boyd et al. [30] examined tweets and retweets and found that "Spreading tweets is not simply to get messages out to new audiences, but also to validate and engage with others". Marwick [31] conducted a research by tweeting questions to Boyd's Twitter followers and analyzing their answers. Among them was a question aimed at understanding to whom are they tweeting. They found that some users "imagined their audience as people they already knew, conceptualizing Twitter as a social space where they could communicate with pre-existing friends".

A comparison to another "quasi social" network may shed some light on this question. Social questions and answers sites, such as Yahoo! Answers also possess a dual nature. Golbeck [32] showed that this service fully meets the Golbeck's accessibility, relationship and support criteria for a web-based social network. Furthermore, a study by Rechavi and Rafaeli [33] showed that within this service there are actually two interdependent networks, a social and an informational network.

5. *What are the roles and characteristics of the retweet mechanism?*

Twitters' retweet feature receives a lot of attention in Twitter-based studies. Some studies provide descriptive data concerning retweet probability. Some studies correlate it against other characteristics, such as number of followers, number of friends, tweeting rate, mentions, hashtags, urls, etc. Other studies try to predict the retweet probability based on the user and content characteristics. Besides its obvious role in message spreading, several researchers claim that the retweet is an important indicator of influence in Twitter.

In a world where endless amount of information is flowing through social networks, competing for the user's attention, the message sender has to overcome the basic passivity of the message receiver. Based on retweets, Romero et al. [34] propose an algorithm to determine the influence and passivity of users based on their information forwarding activity (retweets).They see the retweet as an action performed by the retweeterer. Driving the user to take an action indicates influence. Cha et al. [35] compared three possible measures of influence, indegree, retweets and mentions. They argue that: "it is more influential to have an active audience who retweets or mentions the user. Retweets are driven by the content value of a tweet, while mentions are driven by the name value of the user.". Zaman et al. [36] built a model for retweet probability based on the tweeterer and the tweet content. Suh et al. [37] conducted a large scale study of retweets and found that "URLs and hashtags have strong relationships with retweetability. Amongst contextual features, the number of followers and followees as well as the age of the account seem to affect retweetability, while, interestingly, the number of past tweets does not predict retweetability of a user's tweet".

Boyd et al. [30] studied retweeting as conversational practice and claim that: "While retweeting can simply be seen as the act of copying and rebroadcasting, the practice contributes to a conversational ecology in which conversations are composed of a public interplay of voices that give rise to an emotional sense of shared conversational context. For this reason, some of the most visible Twitter participants retweet others and look to be retweeted. This includes users of all kinds, but notably marketers, celebrities and politicians". A research on celebrities influence in Twitter [38] defined influence in the following manner: "the ability to, through one's own behavior on Twitter, promote activity and pass information to others". He found that retweet-based influence is the most significant type of influence.

A. Related work summary

Most product related WOM studies report a negative bias. Some studies derive from the regulatory focus theory and show the moderator role of the product type, leading to a positive bias in hedonic promotion consumption goal products. Studies on non-product content dissemination, such as news, present contradicting results. Some studies focus on the emotions the message arouses and the message affordance. Several studies support the two ways relationship between the movie's Tweets dissemination and its box office success. Several studies pondered over the nature of Twitter and found that it exhibits also social network characteristics. Studies that focused on the role of Twitter's retweet feature found that it plays an important role in the social sub-network and that retweet is an indicator of influence.

In light of the theory and related work, there were three motivations for this research:

 ISBN: 978-1-61208-372-8

1. The contradicting evidence from different studies.
2. The tension between theories supporting positive versus negative bias.
3. Modest availability of evidence based on updated high volume data collected from online social networks.

In order to be consistent with the view of Twitter as a social network, the focus of the research is on WOM flow between ordinary users and not WOM originating from the highly followed users.

Related studies show that this research stands on solid ground when choosing Twitter as the data collection field, examining the retweet dissemination and choosing to study the movies domain.

IV. RESEARCH HYPOTHESES

Consistent with other researches which used movies and other products tweets and the language positivity bias previously described -

H1: There are more positive than negative tweets in movie-related Twitter messages.

Following the theory and studies that show support for more spread of positive polarity messages in hedonic, promotion consumption goal -

H2: Positive polarity movie tweets will spread more, in number of retweets and audience size, than negative polarity movie tweets.

V. METHOD

The method is based on measuring the message dissemination (the dependent variable) as a function of the message sentiment polarity (the independent variable). The basic categorical values in sentiment polarity are positive, neutral and negative. The retweet mechanism drives dissemination. The collected data contains tweets about movies which came out between the end of 2011 and the beginning of 2012. The focus was on new movies since they were more tweeted about. The rationale of choosing movies was discussed in the related work section. Due to the dual nature of Twitter as mass medium and social network, the research focuses on Twitter's social sub network which contains ordinary users having medium number of followers. The cutoff number below which we considered a user to be an ordinary user was having 1000 followers. This number was chosen due to its being the round number above Dunbar's number for social group size and wanting to address a significant portion of Twitter users. 70% of twitter users have 50-1000 followers (Figure 4). The reason there is a lower limit of 50 followers is that there was almost no retweet activity for tweets sent by users with fewer than 50 followers.

A. *Research process description*

This section describes the research process and results, discussing the rationale of the different steps, challenges and results. The main steps were:

1. Data collection
2. Data cleaning
3. General tweets and retweets statistics
4. Followers analysis
5. Sentiment classification
6. Manual sentiment classification
7. Naïve Bayes classification endeavour
8. Study of tweets dissemination as a factor of the sentiment

1) *Data Collection*

About half a million movie related tweets were collected during 4 months using a service by a company called GNIP [39] (see Figure 1). GNIP provides full access to tweets which was not available directly from Twitter in the time this data was collected.

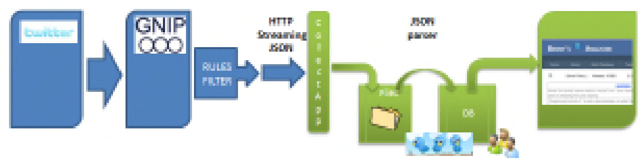

Figure 1. Tweets collection architecture.

2) *Data Cleaning*

After the tweets collection step, where ~500,000 tweets containing movie names were collected, the large data set was analyzed using an application that we've developed for this purpose, called Twitter Analyzer. The application main features are presenting, sorting and filtering all tweet's and user's fields. It also calculates and aggregates number of retweets and exposure.

Manual inspection of the tweets using the Twitter Analyzer indicated that many of the tweets do not contain WOM content. The first step was cleaning the data set in order to get a higher percentage of WOM content. The cleaning process first step was removing several movies related tweets which contained a large number of non WOM content. The second step was using a white list for filtering tweets which contained words indicating that the user had seen the movie (was, were, went to, saw, have seen, had seen, watched, is). After the cleaning process, the clean data set contained 21,000 tweets. Eventually we verified that our data set contained low level of spam in it (less than 3%).

3) *General tweets and retweets statistics*

For the second hypothesis, we're interested in the retweet mechanism and followers count. The relative share of retweets to the overall traffic is described in Table 1. It is

 ISBN: 978-1-61208-372-8

based on the initial data set of 108,000 tweets. The data shows that ~9% (7/78) of the tweets were retweeted and that 22% of the total number of tweets is due to retweets.

4) Followers

Analysis of the number of followers showed that the distribution is according to a power law and 70% of the users have 50 – 1000 followers (see Figure 2). The power law distribution of number of followers is the reason that the 100 most retweeted tweets (~1%) are responsible for about half the retweets in the data set.

TABLE I :MESSAGE TYPE DISTRIBUTION IN THE INITIAL DATASET.

Retweets	22%
Original tweets that were retweeted	7%
Tweets that were not retweeted	71%
Total	100%

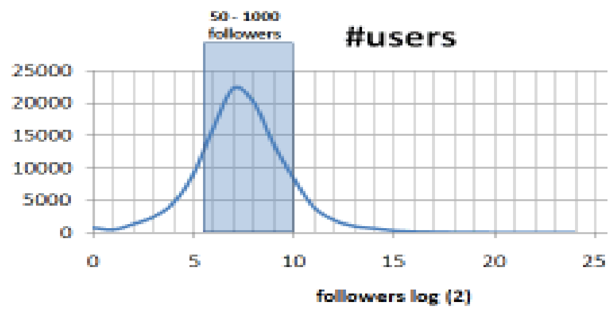

Figure 2: Distribution of users according to the number of followers (log 2).

5) Sentiment Classification

The following categories were defined (*):

- Positive (e.g., "… WAS SOOOOOO GOOD!!!! ")
- Negative (e.g., "… is the first movie where i actually fell off to sleep...#flop")
- Neutral (e.g., "Just saw … and now I think Im Keke Palmer singing all these slow jams lmao")

(*) Two other categories were used to classify the tweets, "Not Relevant" and "Before" (e.g., "going to see the movie"). Due to their relatively low number, they were joined with the Neutral category in order to simplify the analysis.

The following classification guidelines were set:

- In cases where the tweet contained both positive and negative content (e.g., "the movie was too long but interesting"), it was classified as neutral.
- The sentiment classification refers to the movie and not to the general sentiment of the tweet. For example, the following tweet: "had a great time with my friend but the movie been boring" was classified as negative.

6) Manual classification

Two human coders classified 8,600 tweets according to the categories described above. The 8,600 tweets were sampled randomly from the clean data set and contained all the messages that were retweeted and part of the messages that were not retweeted. The human coders' inter-classification-agreement rate was ~84% (7195/8600). Those 7195 classified tweets (Table 2) were used for the dissemination analysis. This data set is called the classified set. The main result of this stage is that the ratio between negative and positive eWOM is 0.18 (7.3/40.4).

7) Naïve Bayes classification endeavour

Having a large tweets data set, our goal was to use the manual classification to train a Naïve Bayes classifier. The low percentage of negative tweets (~7%) led to a relatively high classification error rate which made it unsuitable for usage as a reliable classifier for the larger data set. Unbalanced data sets are a known issue with naïve bayes classifier.

8) Tweets dissemination as a factor of the sentiment

The overall ratio: The overall ratio between negative and positive retweet count was 0.18 (table 3), which is the same as the ratio between negative and positive tweets.

Ordinary users' retweet ratio: A closer examination of retweets count for ordinary users (50 – 1000 followers) showed that positive retweets are retweeted about 15% more on average (Figure 3). Further, neutral tweets get retweeted more times (Table 3) than both positive and negative tweets, this is due to a lot of retweets of neutral tweets that were tweeted by highly followed users.

TABLE II: MANUAL SENTIMENT CLASSIFICATION DISTRIBUTION RESULTS.

Retweets	Negative	Positive	Neutral	Total
Quantity	527	2907	3761	7195
Percent	7.3%	40.4%	52.3%	100%

TABLE III: SUM OF THE TIMES THE TWEETS GOT RETWEETED.

Retweets	Negative	Positive	Neutral	Total
Count(*)	173	976	5365	6514(**)
%	2.7%	15%	82.3%	100%

(*) After removing the highest record in every category
(**) Retweets from the classified data were counted from the large data set

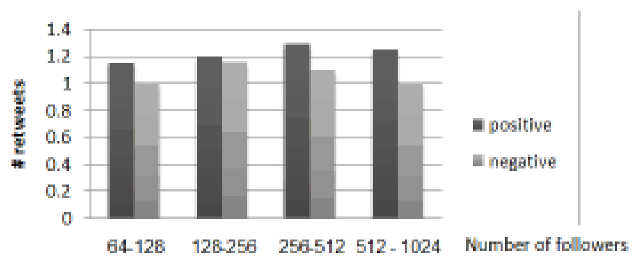

Figure 3: Average number of the times a tweet was retweeted for ordinary users.

TABLE IV: TOTAL EXPOSURE OF TWEETS.

Tweets	Negative	Positive	Neutral	Total
Count(*)	233349	1289641	4867790	6390780
%	3.6%	20%	76.4%	100%

Exposure ratio: Counting the number of users who received the tweet, the negative/positive ratio for the total exposure was again 0.18 (3.6% / 20% in Table 4). Consistent with the previous observation, neutral tweets get the highest exposure.

Ordinary users' exposure ratio: A closer look at the distribution of exposure count of users who received (exposed to) the retweet showed a significant difference of ~15%-20% more positive tweets in the range of 100-600 exposure level (Figure 4).

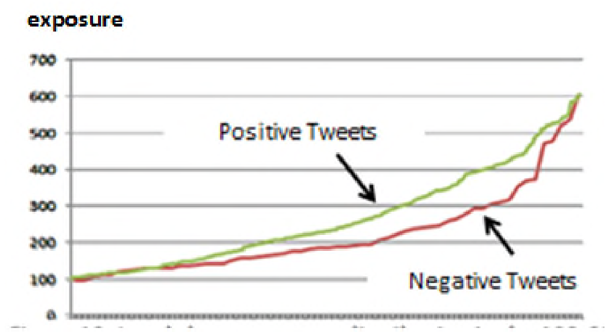

Figure 4: Retweet exposure distribution in the 100-600 exposure range.

9) Dissemination as a factor of sentiment - results summary

The negative/positive ratio of ~0.18 was consistent over the full classified tweet set. This ratio, showing positive dominance over negative in the original tweets (Table 2) is consistent with other studies, such as Blenn et al. [11] and Asur and Huberman [12]. When focusing on ordinary users, in the 100-1000 followers range, there is a positive bias: positive tweets are retweeted more times (Figure 3) and positive retweets exposure is higher (Figure 4).

B. Results summary

General characteristics and statistics of the retweet mechanism: Retweets constitute 7% of all tweets, and counting the repeats they amount to ~29%. The retweets are power law distributed. Some tweets are retweeted in high numbers, 1% of the tweets that were retweeted were responsible for 50% of the total number of retweets.

Followers' distribution: 70% of the users have 50 to 1000 followers. This is the group that is referred to as ordinary users in our analysis. There is a lower boundary of 50 followers below which there are almost no retweets.

Tweets dissemination:

- Full classified data set (7195 tweets):
  - Positive tweets outnumber negative tweets by a ratio of 5.5 (1 / 0.18).
  - Neutral tweets are the most retweeted.
- Ordinary users subset (50 – 1000 followers) (4627 tweets):
  - There was a positive bias of (15%-20%) in dissemination measurements

1) Limitations

The generalization power of these findings is somewhat limited due to the focus on one product type in a specific network. Nevertheless, Twitter is a very big network and movie related tweets are popular (see Asur and Huberman [12]).

VI. Discussion

Our findings support H1 by showing that there are five times more positive than negative sentiment polarity tweets. This is consistent with the language positivity bias. An alternative explanation is that since the dataset contains tweets about movies, most people enjoy the movie they see and avoid going to movies which they will not like by reading reviews and getting recommendations from friends. Regarding the main hypothesis, this study provides support for H2 which predicts that positive sentiment polarity tweets spread more than negative tweets in the social sub-network of Twitter. While the dissemination results of the full classified set showed no preference to positive or negative polarity, a closer look at the ordinary users showed a positive bias of about 15%-20%. These results support the regulatory focus theory and messenger preference, both predicting a positive bias / negative discount, for promotion consumption goal product, such as movies.

The limited explanatory of power of 15-20% suggests that there are other significant factors that affect messages dissemination, such as content, attributes, structure and user characteristics. Some of them are described in the studies referenced in this paper.

The justification of using Twitter as our data field relies on the existence of a social sub-network for ordinary users. Marwick [31] and Gruzd et al. [29] claimed that Twitter has social network aspects in addition to mass medium characteristics. We extricated the social sub-network by restricting the analysis to ordinary users, those with 50-1000 followers. This is a novel approach, following Liu et al. [22] who also split their Twitter users by follower count. Empirically, there are two findings which support the dual approach view. With a cutoff parameter of 600-1000 followers we found significant differences in sentiment dissemination and in tweets length between the two groups.

A repeated finding shows high exposure and number of retweets of neutral tweets. This can be explained in light of the two sub networks approach. Most of the neutral tweets are tweeted by highly followed users. Those highly followed users are often channels of information.

VII. Conclusion and Future Work

Positive tweets get the stage on the social sub network of Twitter with the topic of movies. Furthermore, the positivity bias hypothesis for hedonic promotion consumption goal products received support in these data.

Future work: A complimentary focus on the same question can study the results for eWOM spread on different products. It may be interesting to compare a promotion consumption goal product with a prevention consumption goal one. Message dissemination of the same product (movies) can be studied in other social networks and compared to the results presented here.

ACKNOWLEDGMENT

This research was supported by the I-CORE Program of the Planning and Budgeting Committee and the Israel Science Foundation (grant no. 1716/12) and with affiliation and support of the Internet Research Center.

REFERENCES

[1] K. Nahon and J. Hemsley, Going Viral, Polity, 2013, .

[2] L. Shifman, Memes in Digital Culture, MIT Press, 2013, .

[3] P. F. Lazarsfeld and E. Katz, "Personal influence," Glencoe, Free Presse, 1955.

[4] J. Arndt, "Role of product-related conversations in the diffusion of a new product," J.Market.Res., 1967, pp. 291-295.

[5] M. L. Richins, "A multivariate analysis of responses to dissatisfaction," Journal of the Academy of Marketing Science, vol. 15, 1987, pp. 24-31.

[6] C. W. L. Hart, J. L. Heskett and W. E. Sasser, "The profitable art of service recovery," Harv.Bus.Rev., vol. 68, 1990, pp. 148-156.

[7] R. F. Baumeister, E. Bratslavsky, C. Finkenauer and K.D. Vohs, "Bad is stronger than good." Review of General Psychology, vol. 5, 2001, pp. 323.

[8] C. M. K. Cheung and D. R. Thadani, "The Effectiveness of Electronic Word-of-Mouth Communication: A Literature Analysis," 23rd Bled eConference eTrust: Implications for the Individual, Enterprises and Society, Bled, Slovenia, 2010.

[9] P. Rozin and E. B. Royzman, "Negativity bias, negativity dominance, and contagion," Personality and Social Psychology Review, vol. 5, 2001, pp. 296-320.

[10] C. E. Osgood, Cross-cultural universals of affective meaning, University of Illinois Press, 1975, .

[11] N. Blenn, K. Charalampidou and C. Doerr, "Context-Sensitive sentiment classification of short colloquial text," NETWORKING 2012, 2012, pp. 97-108.

[12] S. Asur and B. A. Huberman, "Predicting the future with social media," Arxiv Preprint arXiv:1003.5699, 2010.

[13] S. Lam, D. Frankowski and J. Riedl, "Do you trust your recommendations? An exploration of security and privacy issues in recommender systems," Emerging Trends in Information and Communication Security, 2006, pp. 14-29.

[14] J. Q. Zhang, G. Craciun and D. Shin, "When does electronic word-of-mouth matter? A study of consumer product reviews," Journal of Business Research, vol. 63, 2010, pp. 1336-1341.

[15] E.T. Higgins, J. Shah and R. Friedman, "Emotional responses to goal attainment: strength of regulatory focus as moderator." J.Pers.Soc.Psychol., vol. 72, 1997, pp. 515.

[16] S. Sen and D. Lerman, "Why are you telling me this? An examination into negative consumer reviews on the Web," Journal of Interactive Marketing, vol. 21, 2007, pp. 76-94.

[17] J. Berger and K. Milkman, "Social transmission, emotion, and the virality of online content," Wharton Research Paper, 2010.

[18] A. Wojnicki and D. Godes, "Word-of-mouth as self-enhancement," HBS Marketing Research Paper no.06-01, 2008.

[19] D. Barbagallo et al, "An empirical study on the relationship between Twitter sentiment and influence in the tourism domain." in Information and communication technologies in tourism 2012, Helsingborg, Sweden, January 25-27, 2012. 2012.

[20] B. J. Jansen, M. Zhang, K. Sobel and A. Chowdury, "Twitter power: Tweets as electronic word of mouth," J.Am.Soc.Inf.Sci.Technol., vol. 60, 2009, pp. 2169-2188.

[21] S. Stieglitz and L. Dang-Xuan, "Political Communication and Influence through Microblogging--An Empirical Analysis of Sentiment in Twitter Messages and Retweet Behavior," in System Science (HICSS), 2012 45th Hawaii International Conference on, 2012, pp. 3500-3509.

[22] L. K. Hansen, A. Arvidsson, F. A. Nielsen, E. Colleoni and M. Etter, "Good friends, bad news-affect and virality in twitter," Future Information Technology, 2011, pp. 34-43.

[23] E. Bakshy, J. M. Hofman, W. A. Mason and D. J. Watts, "Everyone's an influencer: quantifying influence on twitter," in Proceedings of the fourth ACM international conference on Web search and data mining, 2011, pp. 65-74.

[24] M. Thelwall, K. Buckley and G. Paltoglou, "Sentiment in Twitter events," J.Am.Soc.Inf.Sci.Technol., vol. 62, 2011, pp. 406-418.

[25] Y. Liu, H. Rui and A. Whinston, "Whose and What Chatter Matters? The Impact of Tweets on Movie Sales Framework," Working Papers, 2011.

[26] W. Duan, B. Gu and A.B. Whinston, "The dynamics of online word-of-mouth and product sales—An empirical investigation of the movie industry," J.Retail., vol. 84, 2008, pp. 233-242.

[27] H. Kwak, C. Lee, H. Park and S. Moon, "What is Twitter, a social network or a news media?" in Proceedings of the 19th international conference on World wide web, 2010, pp. 591-600.

[28] B. Huberman, D. Romero and F. Wu, "Social networks that matter: Twitter under the microscope," 2008.

[29] A. Gruzd, B. Wellman and Y. Takhteyev, "Imagining twitter as an imagined community," Am.Behav.Sci., vol. 55, 2011, pp. 1294-1318.

[30] D. Boyd, S. Golder and G. Lotan, "Tweet, tweet, retweet: Conversational aspects of retweeting on twitter," in System Sciences (HICSS), 2010 43rd Hawaii International Conference on, 2010, pp. 1-10.

[31] A. E. Marwick, "I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience," New Media and Society, vol. 13, 2011, pp. 114-133.

[32] J. Golbeck, "The dynamics of web-based social networks: Membership, relationships, and change," First Monday, vol. 12, 2007.

[33] A. Rechavi and S. Rafaeli, "Knowledge and Social Networks in Yahoo! Answers," in System Science (HICSS), 2012 45th Hawaii International Conference on, 2012, pp. 781-789.

[34] D. Romero, W. Galuba, S. Asur and B. Huberman, "Influence and passivity in social media," Machine Learning and Knowledge Discovery in Databases, 2011, pp. 18-33.

[35] M. Cha, H. Haddadi, F. Benevenuto and K.P. Gummadi, "Measuring user influence in twitter: The million follower fallacy," in 4th International AAAI Conference on Weblogs and Social Media (ICWSM), 2010, pp. 10-17.

[36] T. R. Zaman, R. Herbrich, J. Van Gael and D. Stern, "Predicting information spreading in twitter," in Workshop on Computational Social Science and the Wisdom of Crowds, NIPS, 2010.

[37] B. Suh, L. Hong, P. Pirolli and E.H. Chi, "Want to be retweeted? Large scale analytics on factors impacting retweet in Twitter network," in Social Computing (SocialCom), 2010 IEEE Second International Conference on, 2010, pp. 177-184.

[38] E. T. R. Rosenman, "Retweets—but Not Just Retweets: Quantifying and Predicting Influence on Twitter," 2012.

[39] http://gnip.com (last visited June 2014)

[40] B. Weiner, "An attributional theory of achievement motivation and emotion", Psychological review, 92(4), 548, 1985.

Supporting Social Discovery of Appropriate Candidates for Promotion to Administrative Status in Wikipedia

Mark Zachry
Human Centered Design & Engineering
University of Washington
Seattle, WA, USA
zachry@uw.edu

David W. McDonald
Human Centered Design & Engineering
University of Washington
Seattle, WA, USA
dwmc@uw.edu

***Abstract*—Online systems that rely on their participants to provide oversight and governance face the challenge of assigning administrative rights to appropriate people in the community of users. This study examines how system tools in one such system, Wikipedia, support and constrain the social processes of administrative rights granting. Reporting on the results of a survey of community participants in the English Wikipedia "Request for Adminship" process, the study offers an analysis of system tools in support of a high-stakes social process. It discovers that the existing tools offering basic counts of user actions in the system are not generally perceived as valuable; instead, members of the community develop nuanced, individualized means of assessing administrative candidates. The study offers insights into the implications of tool design for self-governing online systems.**

Keywords—Wikipedia; administrators; management; collaboration; system tools.

I. INTRODUCTION

In self-governing, social contributor systems, community participants need processes and tools that are well integrated to support their administrative work. Whether the community is setting policy, disciplining those engaged in disruptive behavior, or selecting who among the community will have access to special system tools, integrated processes and tools are essential for facilitating sustainable, communally valued practices [1]. This integrated process-tool complex serves as a kind of bedrock on which a relatively predictable set of recurring actions can occur. At the same time, because a social contributor system is a dynamic entity made up of an evolving set of people and embedded in a larger, changing social context, it must necessarily mutate to remain viable. This need for relative stability and openness to change presents a challenge for those who aspire to design viable, socially attractive contributor systems.

Wikipedia, one of the most widely hailed social contributor systems in the world, is a flagship effort of the Wikimedia movement. As a multi-language encyclopedia notable for its relevance to the representation of human knowledge, Wikipedia is also notable for its reliance on interested people to make it work—including such key contributions as writing its content, supplying images and sounds, ensuring editorial integrity, and adding functional code to improve the underlying system itself. This reliance on contributors further extends to the selection of individuals who will be granted access to administrative tools that enable and sustain the system. In Wikipedia, such individuals are designated as administrators, or admins. Wikipedia administrators have access to tools and features that help them execute their maintenance responsibilities, such as the ability to block user accounts, to restore deleted pages, and to hide page revisions from standard users. (For a comprehensive list of administrator tools and capabilities in Wikipedia, see details at [2].)

A. The Administrator Promotion Process in Wikipedia

Within Wikipedia, Requests for Adminship (RfA) are handled via a formalized process by which editors determine who will become administrators. All registered editors in the community can potentially be involved by nominating users whom they see fit for the administrator role. Subsequently, all registered editors can potentially comment on and render an opinion about nominees to help decide who is a suitable candidate for an administrator role. This open process is exposed to the community with all transactions appearing on a wiki page. Prior research about the viability of this community-driven process [3] has identified a degree of urgency in developing a better understanding of the administrator promotion process because of the system's need for more administrative help.

When participating in the RfA deliberation process, Wikipedia editors may use whatever means the community will allow to make their assessment. A resource available to them is a set of guidelines published on the wiki, which provides some sense of appropriate means for evaluating whether candidates are experienced, active, responsible, interactive, and can be trusted to uphold Wikipedia policies. These and other characteristics are listed in the Wikipedia Guide to RfA, which is available on the wiki at [4]. Further, for each case presented to the community for consideration, a set of nominee-relevant exploration tools is presented via links on the nomination page.

The RfA guidelines and the tools presented to participants in the process—like many other aspects of Wikipedia—have largely emerged from the joint efforts of editors working together. The guidelines, consequently, do not stipulate official or mandatory requirements to becoming an administrator. Instead, the guidelines offer community-derived general advice for nominees and nominators as well as a description of the nomination process. Likewise, the

 ISBN: 978-1-61208-372-8

nominee-specific tools presented with each RfA case are offered by members of the community invested in the process, but open to the involvement of others. Reflecting the priorities of those community members who have invested work in the process, both these guidelines and the nominee exploration tools embed a set of values. The guidelines and tools are tightly connected to characteristics and factors that RfA contributors hope to see (and hope not to see) in a candidate. This process of human defined values becoming embedded in system code is a phenomenon of interest to researchers in Science and Technology Studies generally, but within Wikipedia specifically, it has been studied by such researchers as [5].

Our study advances this line of inquiry by investigating the relationships between the values embedded in the guidelines and tools and the actual work of editors participating in the RfA process. Through this investigation, we seek to discover the relationship between values in structural elements of the system (guidelines and tools) and the work that is conducted with and around them. Our work has potential value, then, for not only understanding a mechanism of self-governance in Wikipedia, but potentially also for understanding other systems that rely on their communities of users to develop community-governed policies to support collaborative action [6]. Researchers who are motivated to understand the dynamics of online elective processes like RfA [7], for example, would potentially benefit from our study.

Prior work examining the RfA process within Wikipedia has addressed several concerns, most of which are indirectly related to the present study. For example, Burke and Kraut [8] propose models of behavior based on the outcome of nomination cases. Another study [9] employs social network analysis techniques to examine the effects of relationships among people on their decision patterns when participating in RfA. A third study [10] examines the editing histories of administrators and analyzes them in relationship to voting patterns in the RfA process. A final study [11] uses interviews to identify sensemaking practices of participants in the RfA process and to consider the design of a visualization tool to support such work. Our work complements these previous studies, examining the thoughts and experiences of RfA participants and connecting them to the general affordances of the tools readily available to support community deliberation about administrator candidates.

In the paper that follows, we describe our means of understanding the social, tool-mediated process by which Wikipedians select community members to be elevated to administrative status (see Section II), examine the values that such members use to guide their opinions (see Section III), consider how the tools support such opinion development (see Section IV), and then, conclude with suggestions about designing for such contexts (see Section V).

II. METHOD

Our study employs survey and system analysis techniques to identify the relationship between community practices and values and the primary tools present in the context of RfA work. It is worth noting that our work focuses specifically on practices, values, and tools in the English Wikipedia, which is the largest instance of Wikipedia. Additional investigations of RfA in other language wikis (e.g., [10]) would be appropriate to determine if our results correspond to relationships and tool use in other language wikis.

A. Design of the Survey

We designed and offered an online survey of editors to study the RfA practices and values of active Wikipedians. Participants were solicited from within Wikipedia itself on pages and related forums on which editors involved in the RfA process were likely to notice the invitation (e.g., Wikipedia pages, such as the RfA talk page, the Village Pump Miscellaneous page, and various related Internet Relay Chat (IRC) channels). Additionally, Kudpung, a Wikipedia Online Ambassador and administrator, re-posted our invitation on the RfA Reform page within Wikipedia (see [12]), noting that WMF founder Jimbo Wales had recently expressed a firm desire to see more data driven contributions to proposals for RfA reform.

The survey asked participants a mix of questions, including demographic characteristics, opinions about the process, frequency with which they used tools from the RfA toolbox, and what characteristics they value most when evaluating an administrator candidate.

The survey was open for one month, during which 61 Wikipedians responded.

B. Data Analysis

To analyze the responses of our participants, we considered the characteristics they identified as important in their consideration of RfA nominees and the tools present in the process that they reported using. We then used their expressions about their deliberation practices to understand how they considered these characteristics and tools.

1) Valued Characteristics

Survey participants provided information about the characteristics they value in nominees in response to questions about (1) the normal evaluation practices, and (2) descriptions of what they look for when assessing a nominee. We then analyzed these responses to determine the frequency with which nominee characteristics were identified as valued. To categorize these responses, we used Wikipedia's Guide to RfA to group characteristics users hope to see in a candidate. As shown in Table II, we created one subcategory that lists the exact characteristics described in Wikipedia's Guide to the RfA, and one subcategory that lists additional characteristics participants mentioned in their responses that closely relate to those described in the guide. We made a count every time a participant mentioned one of the listed characteristics in their responses. In addition to these characteristics, we created a separate table (Table III) that lists characteristics participants collectively mentioned that do not relate to the categories in the guide to the RfA.

2) Tools

We analyzed tools presented to RfA participants in the RfA toolbox (see Figure 1 for an example) to see if there were direct or indirect relationships between the functionalities of the tools and characteristics users would like to see in an administrator candidate. For every characteristic (presented later in Table IV), we determined if there are tools that could assist in researching if a candidate possesses such a characteristic.

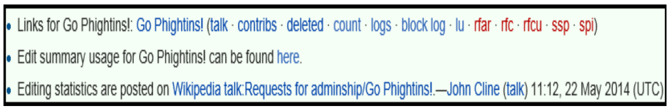

Figure 1. A set of tools presented to editors considering an RfA case. In this instance, the nominee is Go Phightins! An active version of this tool may be accessed here [13].

In these automatically generated tool boxes, participants in the RfA process are presented with links to the candidate's user page and user talk page, as well as links that yield basic data about the candidate's prior system actions. These count links include a log of diffs committed by the editor over time, a count of edits over time, a list of instances in which the candidate's account was blocked administratively, a count of the number of times the editor was involved in an arbitration process, and a count of the number of times the candidate's account was flagged for investigation based on a request of another editor. Additionally, much of this type of data can be accessed via an "editing statistics" link to a talk page associated with the candidate's nomination.

III. RESULTS

Our 61 participants included three types of Wikipedians:

- 31 registered editors
- 29 administrators
- 1 bureaucrat

Registered editors include those users who have created an account within Wikipedia (as opposed to those who edit anonymously). Administrators are editors who have been selected to have elevated systems permissions. Bureaucrats are those with system permissions to add and remove administrators. At the time of the survey, there were approximately 22,100,000 registered editor accounts, 14,000 administrators, and 35 bureaucrat accounts. (For more detail about the overall population of Wikipedians and their unique roles in the work and maintenance of the system, see [14]).

All participants in our survey expressed opinions about the RfA process, and nearly all were regular contributors to such deliberations (Table I). Among the survey participants, 20% reported that they participated in nearly every RfA case.

TABLE I. NUMBER OF RFA CASES THAT SURVEY PARTICIPANTS REPORT CONTRIBUTING TO ON A MONTHLY BASIS

Not monthly	3 (~5%)
Fewer than 2 RfAs per month	26 (~43%)
Two or more RfAs per month	20 (~33%)
Nearly every RfA	12 (~20%)

Survey participants were asked an open question about what they looked for when they considered a nominee: "In a few sentences, please describe what characteristics of an RfA candidate are most important to you as you evaluate him or her and why." Participants offered responses of varying lengths. All responses were analyzed to extract a list of characteristics, and all characteristics were included in the final list. This list was then transformed into a regular list of common characteristics and counts of occurrence, as indicated in Table II. Some of these categories matched the characteristics mentioned in the RfA Guidelines and others were identified as closely associated.

TABLE II. GUIDELINE-BASED NOMINEE CHARACTERISTICS THAT SURVEY PARTICIPANTS REPORT AS IMPORTANT IN THEIR CONSIDERATIONS

Editor characteristics identified in the Guide to RfA	**# of mentions**	**Closely related characteristics mentioned by participants**	**# of mentions**
strong edit history	22	has substantial contributions	27
observes consensus; follows policy	22		
clean block log, good editing behavior	16	no serious concerns about editing history	4
varied/diverse experience	14	experience; breadth of interest, tenure	8
helpful, polite, evident in talk pages, interacts well	13	civility, friendliness, patience, good attitude, work ethic	26
constructive use of edit summaries	5		
trustworthy, reliable, uses admin rights carefully	4	stand by decisions, sound judgment, admits mistakes	10
helps with chores; does admin work	4	good use of tools; good work	4
high quality articles	3	good content	12

 ISBN: 978-1-61208-372-8

Beyond those characteristics that matched or were closely aligned with characteristics from the RfA guidelines, our participants also identified others (Table III). These unique responses include such things as the nominee having an insider perspective (e.g., being “clued” in) and personal histories of interaction (e.g., knowing something about the individual because we have encountered each other while editing).

TABLE III. ADDITIONAL NOMINEE CHARACTERISTICS THAT SURVEY PARTICIPANTS IDENTIFIED AS IMPORTANT IN THEIR CONSIDERATIONS

Unique Responses	# of Mentions
maturity	15
deals well with conflict	7
knowledgeable, intelligent	9
"clued" in / cluefulness	7
Wikipedia-experienced, familiarity, length of experience	7
personal history of interaction / familiar with candidate's name	8

Given the characteristics identified by our survey participants, we then considered how closely they were aligned with the tools presented in RfA cases. As expected, some of the nominee characteristics that the community is encouraged to consider are discoverable in the toolbox (see Table IV). Notably, some of the characteristics are supported by multiple tools. For example, eight of the tools in the toolbox would allow an RfA participant to make an assessment about the strength of a nominee’s editing history. When considered in this framework, it became evident that the tools offered to the community are based on editor characteristics that are relatively simple and easy to represent as a count number. Notable then, too, is the fact that not all the suggested character attributes can be discovered with a tool. These characteristics are things that are not readily countable (e.g., trustworthiness and helpfulness). Additionally, none of the unique (non-guideline based) characteristics identified as important by our participants were discoverable via an offered tool.

The lack of tool support for reasoning about a nominee’s characteristics, presented in a way that the community values, may indicate an opportunity for design enhancements to the RfA toolbox. Such opportunity, however, merits careful consideration. Additional input from our survey participants (discussed in the following section) offers a more complex picture of how tools are connected to the practices of RfA participants thinking through the characteristics they value and how an individual nominee exhibits those characteristics.

TABLE IV. NOMINEE CHARACTERISTICS IDENTIFIED IN THE RFA GUIDELINES THAT ARE SUPPORTED BY TOOLS PRESENTED TO PARTICIPANTS IN THE DELIBERATIVE CONTEXT PAGE

Guideline Characteristic	# of Tools Supporting
Strong edit history	8
Constructive and frequent use of edit summaries	3
Has a clean block log, good editing behavior	2
User interaction (helpful, polite, evident in talk pages, interacts well)	2
Observes consensus (ensures neutrality and verifiability) knows and follows policy and guidelines	1
Varied/diverse experience	0
Trustworthy, reliable, uses admin rights carefully	0
Helps with chores; participating in admin work	0
High quality of articles (articles featured, "good articles")	0

IV. DISCUSSION

Wikipedians are invested in the process of selecting high-quality administrators from their ranks. As the process is now realized in the community, many tools (e.g., guidelines, essays, link-based count logs) have been developed to support this work. As representative voices from the community, our survey participants provide a sense of the nuanced and largely invisible aspects of the process-tool complex involved. As one participant explains,

“The reality is that adminship is oriented to communal trust and confidence, not percentages and numbers, and each user will have their own way to assess candidates' readiness for the role.”

As expressed here, among all the participants, there is some disbelief that tools based on automated counts can represent all the qualities of interactions and edits that are worth consideration. The implicit implication of these comments is that humans must judge candidates based on their reading/sensemaking about candidate behavior.

A. Patterns of Consideration Used

Although we see evidence among participant responses that each user has “their own way” of thinking about nominees, we note that these ways can broadly be organized into two categories based on the intensity of consideration.

Some participants describe a "systematic" investigation approach to considering nominees. These people explore a wide range of participation examples by methodically looking through the places where the candidate has edited. These people also look systematically at the exchanges the candidate has had with others. They specifically consider the user's page and corresponding user_talk page (the primary on-wiki spaces used for direct communication between two or more editors) and those pages' revision histories. These systematic users explore deeply and would seem to be spending considerable time in their exploration.

Other survey participants describe a comparatively shallow form of exploration. These users make general comments pointing out that there are other users who are likely to do a deep exploration and that any glaring flaws would likely be found by those others. Thus, these shallow explorers feel that they only need to look at a small number of issues which they care about. In these cases, there is not one thing that they consider as a group, but in general they are not looking widely. They tend to mention using the RfA comments posted by previous participants in the deliberation to help make their determination, before looking at one or two other items to render their opinion.

B. Considerations Beyond Those Represented in Tools

Regardless of how intensely and by what means the survey participants considered nominees, factors they considered extended beyond those that could be explored through existing tools. These factors were diverse, as suggested in the quotes that follow.

1) Breadth of Knowledge

Participants cared about how a nominee demonstrated potential for thinking broadly. All nominees are specialists in certain topical content areas and/or types of work within Wikipedia, but a good nominee needed to show potential beyond those areas of specialization. A participant explained they wanted to see a "breadth and depth of knowledge of policies and guidelines, not simply being able to regurgitate sentences from policy pages." Another participant expanded on this idea: "For me, a user must be competent and well aware of things that are going on on the wiki, not only in the area in which they wish to work (although this is most important)."

A third participant offered a slightly different take on this expectation for breadth of knowledge. In particular, this participant wanted nominees to demonstrate an appreciation for the diverse work of Wikipedians. Such an appreciation is deemed necessary because of the power that a confirmed nominee would hold upon becoming an administrator:

"Candidates don't have to be prolific content contributors, but they must show that they know all the content basics and can work effectively with experienced content editors. The 'janitor' metaphor is misleading: a janitor can't tear up your homework and ban you from school, but an admin can."

2) Bravery

Paired with this concern about desiring admins who appreciated a breadth of experiences was a desire to identify nominees who would not shrink away from the necessary hard decisions an administrator must perform. In short, the community values an administrator who will step up and do difficult things. As a survey participant explained, an ideal nominee is "not afraid to step on a few toes to do their job as an admin."

3) Social Adeptness

A valued consideration among survey participants was the sociability of nominees. Sociability included many dimensions. For example, it is related to social trust: "Probably the main question I ask overall is 'Can I trust the candidate?'" Sociability is also related to moderate behavior in the face of heated debate. A survey participant explains, "A calmness and maturity that means the editor is generally a moderating influence in contentious or heated circumstances. This does not just apply to editors who can become uncivil in conflict, but also those who are naive and lack tact and sensitivity." For some survey participants, social adeptness was something that could best (and perhaps only) be judged by prior intimate interactions: "Generally I comment only on the RfAs of users I've interacted with, positively or negatively, and base my position on that."

4) Fitness as a Representative

Finally, survey participants weighed how well a candidate would represent the overall community. A survey participant explained that when considering a nominee, they thought about "whether the candidate will be a good spokesperson; admins are on a pedestal to some degree; they have the illusion of authority. As such, they should represent the project well to new editors, etc." Related to this, there was concern for finding good representatives of the diversity of the community: "For Wikipedia to be 'the free encyclopedia that anyone can edit,' all demographics should be welcomed."

V. CONCLUSION AND FUTURE WORK

As this study demonstrates, the means by which users of a large contributor system engage in the difficult work of self-governance are complex [15]. To ensure that the best possible users are elevated to the status of administrators, people create and participate in an integrated tool-process set of activities through which they endeavor to discover administrator candidates who exhibit characteristics they value. As the operators of such a contributor system consider options to better support the work of the community, they would ideally consider such integrated tool-process activities.

Design considerations for future versions of such systems should include a rich model of the socio-technical aspects of sustainable community-driven collaborative work. Our study offers some initial basis for thinking through such design work in the future.

 ISBN: 978-1-61208-372-8

Following the results of our survey, the design of the RfA experience could be modified to better align tools with the needs of the community. This redesign would include the expansion of the tool set to be more than a collection of basic counters. It would be desirable, for example, to offer tools that could be configured by users to represent more complex patterns of counts that correspond to the kinds of patterns that many RfA participants seek out. Such configurability is technically feasible, drawing on the techniques of end user configuration used in other social media sites. Beyond supporting configurations of basic counts, it would be possible to allow users to weight such counts, which addresses some of the nuanced judgments that survey participants reported as part of their typical RfA work.

Beyond supporting more complex, configurable views of counts, a redesigned system might also include new measures of candidate behaviors that are not currently available. Our examination of how existing tools support valued administrator characteristics as expressed in the RfA Guidelines yielded a set of characteristics not explorable based on simple edit counts. With the development of more sophisticated data mining and interpretation techniques, however, it seems feasible that some insight into these valued characteristics might be possible. For example, it would not be too difficult to classify previous user contributions to expose how frequently a candidate has engaged in actions that the community values as "chores" or has conducted work that is perceived as high quality. Additionally, textual analysis software could be used to characterize how individuals participate on talk pages which could help others to understand the social adeptness of a candidate with whom they had not yet interacted.

Drawing on the expressed needs of the community (both through this survey and beyond), it seems reasonable that the community of editors in Wikipedia could be better supported in the social discovery of valued administrators to conduct important and much needed work in the system.

ACKNOWLEDGMENTS

This work is supported in part by the National Science Foundation under Grants IIS-0811210 and IIS-1162114. The authors thank Anna DelaMerced for her contributions to conducting the survey.

REFERENCES

[1] A. Schroeder and C. Wagner, "Governance of open content creation: A conceptualization and analysis of control and guiding mechanisms in the open content domain," in J. Am. Soc. Inf. Sci. Technol. 63, 10 (October 2012), pp. 1947-1959, doi:10.1002/asi.22657 http://dx.doi.org/10.1002/asi.22657.

[2] Wikipedia:Administrators. Available from http://en.wikipedia.org/wiki/Wikipedia:Administrators

[3] M. Burke and R. Kraut, "Taking up the Mop: Identifying Future Wikipedia Administrators," in CHI '08 Extended Abstracts on Human Factors in Computing Systems, ACM, 2008, pp. 3441-3446, doi:10.1145/1358628.1358871.

[4] Wikipedia:Guide to Requests for Adminship. Available from http://en.wikipedia.org/wiki/Wikipedia:Guide_to_requests_for_adminship

[5] C. Müller-Birn, L. Dobusch, and J. D. Herbsleb, "Work-to-rule: The Emergence of Algorithmic Governance in Wikipedia," Proc. of the 6th International Conference on Communities and Technologies (C&T '13), ACM, 2013, pp. 80-89, doi:10.1145/2482991.248299.

[6] T. Kriplean, I. Beschastnikh, D. W. McDonald, and S. A. Golder, "Community, Consensus, Coercion, Control: CS*W or How Policy Mediates Mass Participation," Proc. 2007 International ACM Conference on Supporting Group Work (GROUP), ACM, 2007, pp. 167-176. doi:10.1145/1316624.1316648.

[7] J. B. Lee, G. Cabunducan, F. G. Cabarle, R. Castillo, and J. A. Malinao, "Uncovering the Social Dynamics of Online Elections," in J. UCS 18, no. 4, 2012, pp. 487-505.

[8] M. Burke and R. Kraut, "Mopping Up: Modeling Wikipedia Promotion Decisions," Proc. 2008 ACM Conference on Computer Supported Cooperative Work (CSCW), ACM, 2008, pp. 27-36.

[9] G. Cabunducan, R. Castillo, and J. B. Lee, "Voting Behavior Analysis in the Election of Wikipedia Admins," Proc. Advances in Social Networks Analysis and Mining (ASONAM), 2011, pp. 25-27, doi:10.1109/ASONAM.2011.42.

[10] P. Turek, J. Spychala, A. Wierzbicki, and P. Gackowski, "Social Mechanism of Granting Trust Basing on Polish Wikipedia Requests for Adminship," in Lecture Notes in Computer Science, vol. 6984, Springer Berlin Heidelberg, 2011, pp. 212-225.

[11] K. Derthick, P. Tsao, T. Kriplean, A. Borning, M. Zachry, and D. W. McDonald, "Collaborative Sensemaking during Admin Permission Granting in Wikipedia," in A.A. Ozok and P. Zaphiris (Eds.): Online Communities, HCII 2011, LNCS 6778, 2011, pp. 100–109.

[12] Wikipedia:RfA Reform 2012. Available from http://en.wikipedia.org/wiki/Wikipedia:RfA_reform_2012

[13] Wikipedia:Requests for adminship/Go Phightins!. Available from https://en.wikipedia.org/wiki/Wikipedia:Requests_for_adminship/Go_Phightins!

[14] Wikipedia:Wikipedians. Available from http://en.wikipedia.org/wiki/Wikipedia:Wikipedians

[15] I. Beschastnikh, T. Kriplean, and D. W. McDonald, "Wikipedian Self-governance in Action: Motivating the Policy Lens," Proc. Int. Conf. on Weblogs and Social Media, 2008.

 ISBN: 978-1-61208-372-8