Information Systems 54 (2015) 147–155

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/infosys

Mining citizen emotions to estimate the urgency of urban issues

Christian Masdeval [a,*], Adriano Veloso [b]

[a] Dataprev, Brazil
[b] Universidade Federal de Minas Gerais, Brazil

ARTICLE INFO

Article history:
Received 4 December 2014
Received in revised form
30 May 2015
Accepted 22 June 2015
Available online 26 June 2015

Keywords:
Emotion classification
Networked citizenship
Crowdsourcing

ABSTRACT

Crowdsourcing technology offers exciting possibilities for local governments. Specifically, citizens are increasingly taking part in reporting and discussing issues related to their neighborhood and problems they encounter on a daily basis, such as overflowing trash-bins, broken footpaths and lifts, illegal graffiti, and potholes. Pervasive citizen participation enables local governments to respond more efficiently to these urban issues. This interaction between citizens and municipalities is largely promoted by civic engagement platforms, such as See-Click-Fix, FixMyStreet, CitySourced, and OpenIDEO, which allow citizens to report urban issues by entering free text describing what needs to be done, fixed or changed. In order to develop appropriate action plans and priorities, government officials need to figure out how urgent are the reported issues. In this paper we propose to estimate the urgency of urban issues by mining different emotions that are implicit in the text describing the issue. More specifically, a reported issue is first categorized according to the emotions expressed in it, and then the corresponding emotion scores are combined in order to produce a final urgency level for the reported issue. Our experiments use the SeeClickFix hackathon data and diverse emotion classification algorithms. They indicate that (i) emotions can be categorized efficiently with supervised learning algorithms, and (ii) the use of citizen emotions leads to accurate urgency estimates. Further, using additional features such as the type of issue or its author leads to no further accuracy gains.

1. Introduction

City maintenance is extremely expensive, since it involves monitoring and fixing a variety of complex issues related to public safety, environmental problems, and quality of life. In particular, monitoring urban issues (e.g., potholes, damaged street signs, graffiti, street light issues, damaged trees, park maintenance) usually requires a large number of employees working on a permanent basis. Alternatively, crowdsourcing can promote a participatory way of monitoring urban environments, and thus municipality's resources can be redirected to effectively fixing the reported issues. As citizens are increasingly equipped with smartphones, they are also increasingly able to perform pervasive crowdsourcing at urban level, that is, to report urban issues with GPS location directly to the city's appropriate department. This model of crowdsourced city enables not only better resource allocation, but also fine-grained monitoring capabilities.

Early examples of crowdsourcing platforms for monitoring urban issues include See-Click-Fix,[1] FixMyStreet,[2]

* Corresponding author.
E-mail addresses: christiancleber@dataprev.gov.br (C. Masdeval), adrianov@dcc.ufmg.br (A. Veloso).

[1] seeclickfix.com
[2] www.fixmystreet.com

http://dx.doi.org/10.1016/j.is.2015.06.008
0306-4379/

CitySourced,[3] and OpenIDEO.[4] These platforms enables citizens to report issues by entering a description of what needs to be done, fixed or changed. As an example, Fig. 1 shows a report extracted from seeclickfix.com/issues/1273509-turns-on-and-off-all-night. Citizens may also vote for specific issues, thus endorsing the wish that the problem is solved. While the number of votes an issue receives ultimately reflects its urgency to be solved, acquiring a significant number of votes may take several days or weeks, leading to ineffective and late responses. In order to become more responsive, government officials must be able to prioritize more urgent issues, and thus estimating the number of votes an issue will receive may help officials to better meet the needs and concerns of the citizens.

We here are interested in estimating the urgency of a reported issue by the number of votes it receives, and we investigate the extent to which the textual description by itself determines the urgency of the reported issue. We observed that issues are usually described in an opinative way, but with different viewpoints and inclinations. Thus, we hypothesize that emotions[5] that are implicit in the textual description of an issue (i.e., fear, distress, shame) may be a good evidence of its urgency. We propose algorithms for categorizing the reported issues according to the emotions expressed in their textual description, and then we exploit the corresponding emotion scores in order to estimate the number of votes the reported issue will receive. Since such estimates are made based solely on the text used to describe the issue, government officials may have an immediate view of the urgency of the issues, enabling them to prioritize solving problems that are more urgent according to the citizens.

Experiments using real data obtained from the SeeClickFix hackathon (www.meetup.com/software/events/138126482) demonstrated the effectiveness of exploiting citizen emotions for the sake of estimating the urgency of urban issues. Specifically, we evaluated a set of emotion classification algorithms and conclude that estimates based on emotions implicit in the textual description are significantly more accurate than estimates obtained directly from the text. Also, using additional features related to the author of the issue and to the type of the issue leads to no further significant gains.

2. Related work

2.1. Participatory sensing and crowdsourcing in urban spaces

The revolution in communication and crowdsourcing [4,5] technologies has been changing not only the daily lives of people but also the interactions between governments and citizens. In recent years, many efforts have been made in order to understand these interactions. Kanhere [13] provides a comprehensive overview of these efforts.

For Chun et al. [8], the final stage of transformation of "open government", the so-called Government 2.0, implies that information should flow not only from the government to the citizens but also from citizens to the government and among citizens. It will require principles, functions and technological enablers to lead to a transformative and participatory model. Participation encourages the public engagement by increasing opportunities for the public to participate in policy making. They also define Web 2.0 as a collection of social media through which individuals are active participants in creating, organizing, editing, combining, sharing, commenting, and rating Web content as well as forming a social network through interacting and linking to each other and state that the required functions of Government 2.0 can be achieved by adopting the Web 2.0 technologies.

Kavanaugh et al. [14] make an exploratory study of using traditional social media content as Twitter, Facebook, Flickr and YouTube, to detect, in real time, spikes in activity related to issues concerning public safety. Analyzing information from multiple social media sources should be possible to identify convergence situations (meaningful patterns and trends) and it will help, among other things, in treat crisis situations, from the routine (e.g., traffic, weather crises) to the critical (e.g., earthquakes, floods). This is difficult once a lot of noise should be filtered to make the information useful and reliable. To accomplish its goal, they employed mining techniques covering multiple media types (i.e., text, audio, image, and video). Also, they developed tools to recognize events and to help the visualization of the "big picture" of social media activity and content, and changes in both over time.

Handte et al. [10] proposed methods for crowd density estimation for improving public transportation. Specifically, the authors proposed approaches to estimate the number of passengers in a vehicle. Artikis et al. [1] presented a system for heterogeneous stream processing and crowdsourcing supporting intelligent urban traffic management, and Schnitzler et al. [21] provide an overview of an intelligent urban traffic management system, including approaches for dealing with complex events such as congestions. Litou et al. [15] proposed an approach for emergency notification using online social networks. The proposed approach selects the most efficient routes to maximize the information reach.

2.2. Sentiment analysis and emotion classification

Sentiment analysis methods are typically divided into two broad categories: those that are based on lexical approaches and those that are based on machine learning algorithms. One advantage of learning-based methods is their ability to adapt and create trained models for specific purposes and contexts. In contrast, lexical-based methods make use of a predefined list of words, where each word is associated with a specific sentiment. Next we will focus on learning-based methods, since this is the approach we will follow in this paper.

[3] www.citysourced.com

[4] openideo.com

[5] Although there is not an exact definition for the concept of emotion, most agree that emotions are reactions to events deemed relevant to the needs, goals, or concerns of an individual.

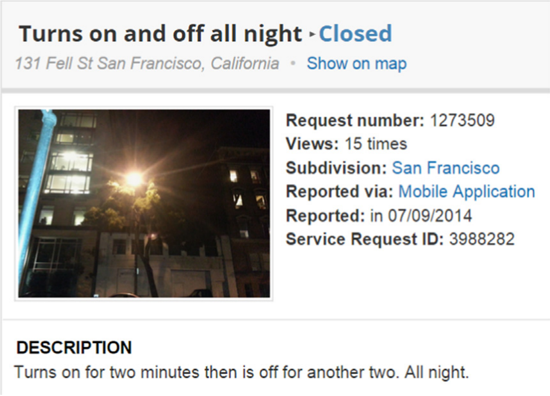

Fig. 1. A reported issue. The description indicates a lighting issue.

Multiple variants of machine learning algorithms, such as SVMs and Naive Bayes, have been used for sentiment classification. Pang et al. [19] used such algorithms to investigate the effectiveness of classification of documents by overall sentiment. Experiments demonstrated that machine learning algorithms are better than human produced baseline for sentiment analysis on movie review data. Features based on unigrams and bigrams are used for classification. Learning algorithms included Naive Bayes, Maximum Entropy, and SVMs. Interestingly, the authors suggested that the evaluated learning algorithms are better than human baselines for sentiment classification. Chaovalit and Zhou [7] investigated movie review mining using machine learning algorithms and semantic orientation. Specifically, supervised classification algorithms are used in the proposed approach to classify the movie review.

Zhu et al. [27,28] proposed aspect based opinion polling from free form textual customers reviews. The aspect related terms used for aspect identification were learned using multi-aspect bootstrapping. Bikel and Sorensen [3] implemented a Subsequence Kernel based Voted Perceptron and compared its performance with standard SVMs. It is observed that as the number of true positives increases, the increase in the number of false positives is much less in Subsequence Kernel based voted Perceptrons compared to SVM.

Manually labeling training data can be time-consuming. To reduce the labeling effort, opinion words can be utilized in the training procedure. Tan et al. [23] used opinion words to label a portion of informative examples and then learn a new supervised classifier based on labeled ones. A similar approach is also used in Qiu et al. [20]. In addition, opinion words can be utilized to increase the sentiment classification accuracy. Melville et al. [18] proposed a framework to incorporate lexical knowledge in supervised learning to enhance accuracy. Ensembles have also been evaluated for the sake of sentiment analysis. An ensemble works by combining the outputs of several base classification models to form an integrated output. Xia et al. [26] proposed different ensembles and made a comparative study of their effectiveness for sentiment classification. Gomide et al. [9] estimate the overall online sentiment expressed by the Brazilian population during the Dengue season, and exploit the obtained sentiment scores to anticipate outbreaks. As a result, Dengue incidence has been reduced by 20% in Brazil.

3. Urgency estimates based on citizen emotions

Our main objective in this paper is to investigate if the emotions associated with a reported issue (i.e., the distribution of emotions) are indicative of its urgency to be solved. Basic emotions are thought to be somewhat universal, and several scales have been developed for capturing emotions reported by people (typically less than 10 emotions are suggested). Since urban issues are expected to be associated with negative emotions, we decided to use the PANAS negative scale [24], which includes the following emotions: distress, disillusion, shame, hostility, irritability, anxiety, jitters, fright, and fear. Specifically, our basic assumption is that we can fit a model which estimates the urgency of the reported issues based on the different emotions implicit on the corresponding textual descriptions. Thus, we divided the task of estimating the urgency of urban issues based on citizen emotions into two separate steps:

- *Emotion classification*: The input for the emotion classification problem is a training set (referred to as $\mathcal{D}$) which consists of a set of records of the form $\langle d, S \rangle$, where d is the textual description of a reported issue (represented as a set of terms), and S is a set of emotions associated with the reported issue (i.e., an arbitrary issue may evoke more than one emotion). The training set is used to build functions relating textual patterns in the descriptions of the issues to the emotions associated with them. A set of future issues (referred to as $\mathcal{T}$) consists of records of the form $\langle d, ? \rangle$ for which only the terms in the description d are known, while the emotions associated with t are unknown. Classification models obtained from $\mathcal{D}$ are used to score the emotions associated with each reported issue in $\mathcal{T}$. The result is a set $\{x_1, x_2, \ldots, x_9\}$, where each x_i is the probability associated with a specific emotion i in the PANAS negative scale. Fig. 2 illustrates how different emotions may be distributed over two reported issues. The objective of an emotion classification algorithm is to approximate as best as possible the true distribution of emotions for each reported issue.
- *Urgency estimate*: The input for the urgency estimate problem is a training set (referred to as $\mathcal{D}'$) which consists of a set of records of the form $\langle E, v \rangle$, where E is the emotion distribution of a reported issue (represented as a vector of nine emotion probabilities $\{x_1, x_2, \ldots, x_9\}$), and v is the number of votes the reported issue has received. The training set is used to build a function relating the emotion distribution of an issue to the number of votes it received. A set of future issues (referred to as $\mathcal{T}'$) consists of records of the form $\langle E, ? \rangle$ for which only the emotion estimates E are known, while the number of votes v associated with the corresponding issue is unknown. Regression

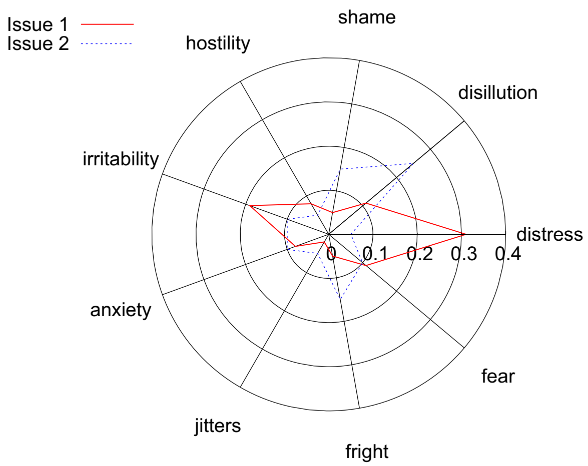

Fig. 2. Distribution of emotions obtained from the textual descriptions of two illustrative issues. Emotion classification aims at approximating as best as possible these distributions. Our hypothesis is that the urgency of an issue can be efficiently estimated by the emotion distribution extracted from the corresponding textual descriptions.

models obtained from $\mathcal{D}'$ are used to estimate the number of votes for each reported issue in $\mathcal{T}'$.

Next, we discuss the features and algorithms used to solve the emotion classification problem. Then, we discuss algorithms for estimating the urgency of reported issues based on the emotions associated with them.

3.1. Feature engineering and emotion classification algorithms

As most supervised learning applications, the main task of emotion classification is to engineer an effective set of features [16]. In this work we consider features that can be derived from textual descriptions of urban issues:

- Individual words and word bi-grams in the textual description of the reported issue.
- The author of the textual description.
- The type of the issue being described.
- The number of terms composing the textual description.

Given the above features, any existing supervised learning algorithm can be applied to emotion classification [16]. Specifically, we investigate three algorithms: Naive Bayes, Support Vector Machines, and Associative Classifiers [19,12,9,22,17], which are briefly described next:

- *Naive Bayes* (*NB*): This algorithm computes the probability of an emotion s given a reported issue d as

$$p(s|d) = \frac{p(s) \times p(d|s)}{p(d)}.$$

To estimate the term $p(d|s)$, the algorithm decomposes it by assuming that features are conditionally independent given s:

$$p(d|s) = \frac{p(s) \times (\prod_i p(f_i|s))}{p(d)}.$$

where each f_i is a feature of d.

- *Support Vector Machines* (*SVM*): This algorithm finds a hyperplane that separates reported issues according to emotions expressed in them. The separation, or margin, is as large as possible. The algorithm returns a decision function $h(d)$, so that the probability of an emotion s given a reported issue d is given as

$$p(s|d) = \frac{1}{1 + e^{ah(d)+b}}.$$

where a and b are estimated by minimizing the negative log-likelihood function in $\mathcal{D}$.

- *Associative Classifier* (*AC*): This algorithm receives an arbitrary reported issue d and extracts association rules $\{X \rightarrow s\}$ from $\mathcal{D}$, such that $X \subseteq d$. These rules have a confidence factor, denoted as θ, which corresponds to a weak estimate of $p(s|d)$. The final estimate of $p(s|d)$ is given as

$$p(s|d) = \frac{\sum_i^r \theta_i}{r}$$

where r is the total number of rules $\{X \rightarrow s\}$ for which $X \subseteq d$.

3.2. Fitting urgency estimates

We experiment with two regression algorithms, which produce as output a real value that represents the urgency of a reported issue.

- *Ordinary Least Squares* (*OLS*): The first algorithm is an ordinary least squares multivariate linear regression model. It estimates the number of votes v^* an issue will receive as a linear function of nine predictor variables $x_1, x_2, \ldots, x_9$ (i.e., $v^* = \beta_0 + \beta_1 \times x_1 + \beta_2 \times x_2 + \ldots + \beta_9 \times x_9$). Each x_i is the probability associated with a specific emotion. Model parameters $\beta_0, \beta_1, \ldots, \beta_9$ are determined by the minimization of the least squared errors [11] in the training set $\mathcal{D}'$.
- *Support Vector Regression* (*SVR*): We also consider the more sophisticated Support Vector Regression algorithm [2], a state-of-the-art regression algorithm. Unlike the OLS model, SVR does not consider errors that are within a certain distance of the true value (within the margin). It also allows the use of different kernel functions, which help solving a larger set of problems, compared to linear regression. We use both linear and radial basis function (RBF) kernels, available in the LIBSVM package [6], as the latter handles non-linear relationships.

4. Experimental section

In this section we empirically analyze the performance of the algorithms for emotion classification and urgency estimation. We employ p@k (precision at the first k

predicted emotions) as the basic evaluation measure in order to assess emotion classification performance. We employ RMSE (Root Mean Square Error) and MAP (Mean Average Precision) in order to assess regression and ranking performance. RMSE evaluates how close are v_i and v_i^*, which are respectively the actual and the estimated number of votes received by issue $i \in \mathcal{T}'$. MAP evaluates how two ranked lists correlate to each other. That is, we sort the reported issues using v_i and v_i^* and then compare how the obtained ranked lists correlate. Next we present relevant information about the dataset used in our experiments, and then we discuss our experimental methodology and the results obtained.

4.1. Dataset

Our experiments use the SeeClickFix hackathon data, which consist of 3-1-1 reported issues from four cities (Oakland, Richmond, New Haven, Chicago) covering the time period from 1-1-2012 to 4-1-2013. SeeClickFix.com is a web-based service designed to help citizens report non-emergency issues in their neighborhood. Submissions can be made via a web interface, by iPhone, Blackberry and Android reporting apps and Facebook application. The data come with several attributes, including the textual description of the issue, the author and type of the issue, and the number of votes it has received. We randomly selected 30,000 reports, which were then labeled by at least 3 human annotators, who were asked to label each issue with all the emotions that they thought to be implicit in the textual description. A total of 2818 reported issues do not express emotions, and in this case all probability values are set to 0. Thus the final dataset we used in our experiments is composed of 30,000 reported issues along with the emotions that are implicit in them.

Fig. 3 (Left) shows the word frequency distribution, while Fig. 3 (Right) shows the cumulative distribution function for the number of votes associated with reports. Most of the reports (about 72%) received only one vote, and less than 1% of the reports received more than 10 votes. Fig. 4 shows the average number of words for each type of issue. Some types of issue are usually associated with long textual descriptions (> 60 words), while other types of issue are usually associated with short textual descriptions (< 20 words). Fig. 5 shows the frequency of each type of issue. Specifically, issues related to trash, trees and potholes correspond to more than 30% of all issues. On the other hand, issues related to lost and found items, or to public art are rarely reported by citizens. Figs. 6 and 7 show the total and the average number of votes depending on the type of issue, respectively. Fig. 8 shows the relationship between urgency and frequency for different types of issue. Not always the most frequent issue is the most voted one. Issues related to hydrant occupy the 7th position in term of frequency (2209 occurrences), but only the 13th position in the list of the most voted (2796 votes). This suggests that despite many citizens register issues about hydrant, few citizens (or at least not in the same proportion) actually consider it a relevant problem. The inverse situation may also happen: there are 106 reported issues related to drug dealing and this type of issue occupies the 21th position in the list of occurrences. But in the list of votes, it is in the 17th position, revealing the natural concern citizens have about drug dealing. Finally, Fig. 9 shows how different emotions are distributed over the reported issues. More than 40% of the reported issues show some irritability from the author of the issue. Disillusion is also an emotion that is frequently shown by citizens when reporting issues. Shame and hostility are the emotions with less occurrences. A reported issue has, on average, 2.38 emotions assigned to it.

4.2. Experimental methodology

First, we conducted 5-fold cross validation for evaluating the emotion classification step. Thus, the original dataset with the textual description (and possible additional features) was arranged into five folds, including training and test sets (i.e., $\mathcal{D}$ and $\mathcal{T}$). After classifying the emotions for all reported issues, another dataset is formed where each record consists of the emotion distribution for the corresponding reported issue. Then, a second round of 5-fold cross-validation was performed in order to evaluate the urgency estimate step. Again, the dataset with emotion distributions was arranged into five folds, including training and test sets (i.e., $\mathcal{D}$ and $\mathcal{T}'$). For both emotion classification and urgency estimate steps, the results reported are the average of the five runs, and we used the Wilcoxon signed-rank test [25] for determining if the difference in performance was statically meaningful. All

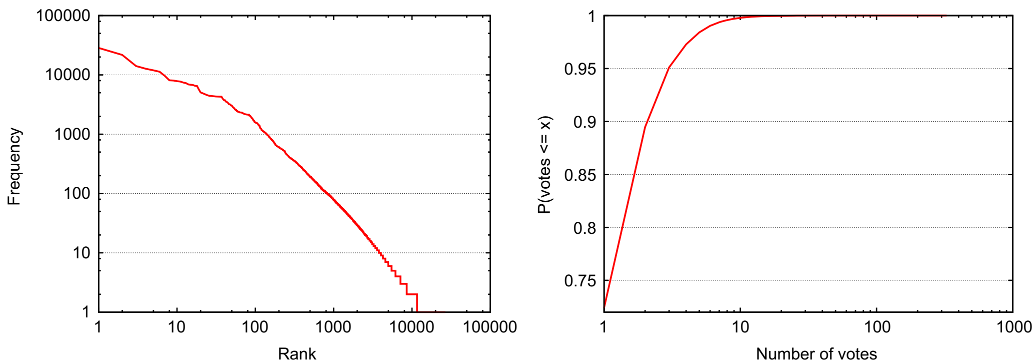

Fig. 3. Left – word frequency distribution. Right – cumulative distribution function for the number of votes.

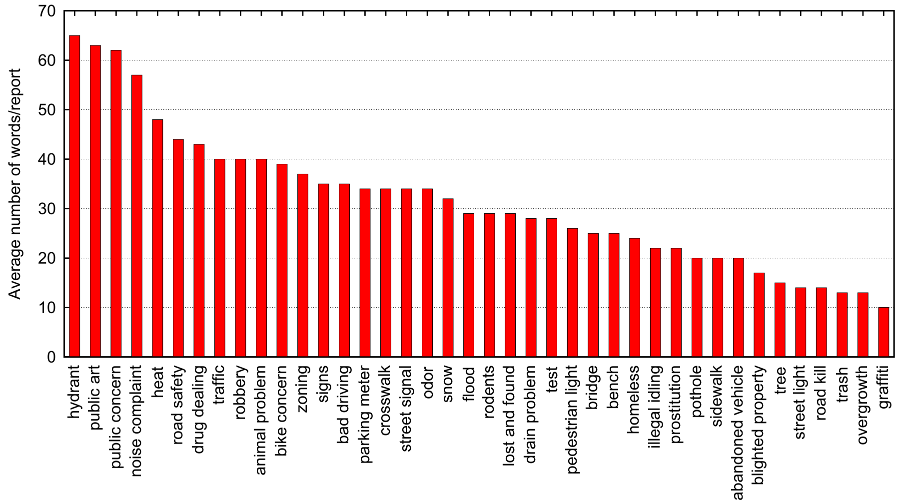

Fig. 4. Average number of words for each type of issue.

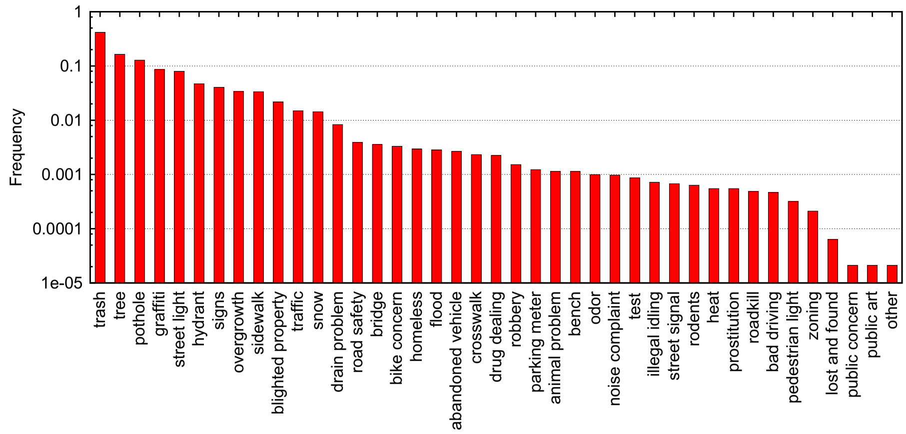

Fig. 5. Frequency of each type of issue.

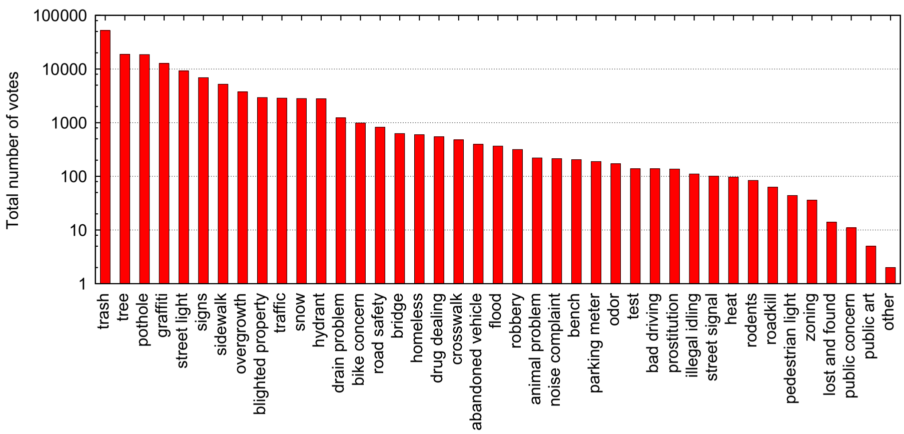

Fig. 6. Total number of votes for each type of issue.

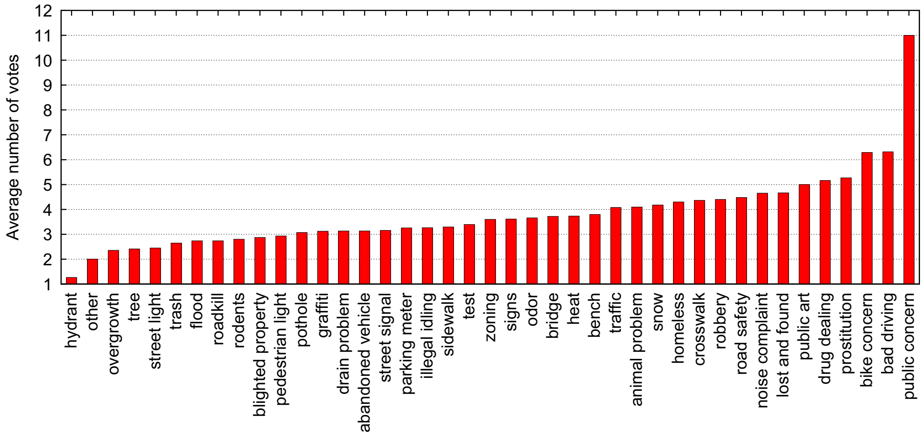

Fig. 7. Average number of votes for each type of issue.

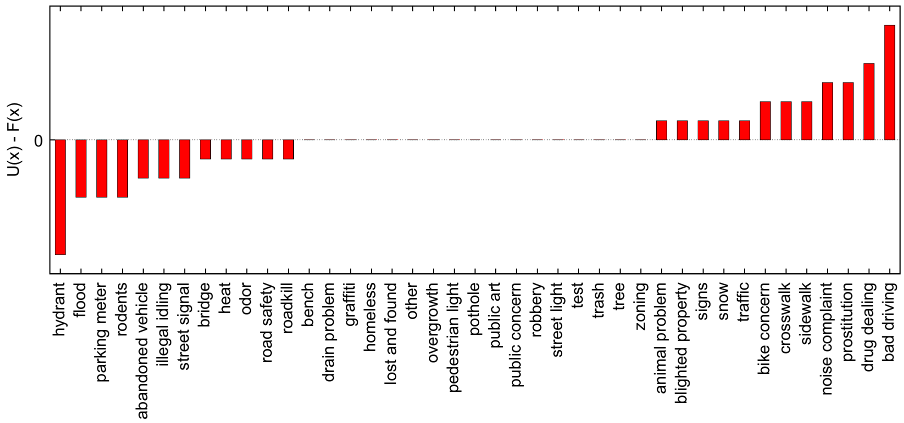

Fig. 8. Relationship between urgency and frequency for different types of issue. $U(x)$ gives the position of the type of issue x in terms of the average number of votes. $F(x)$ gives the position of the type of issue x in terms of its frequency.

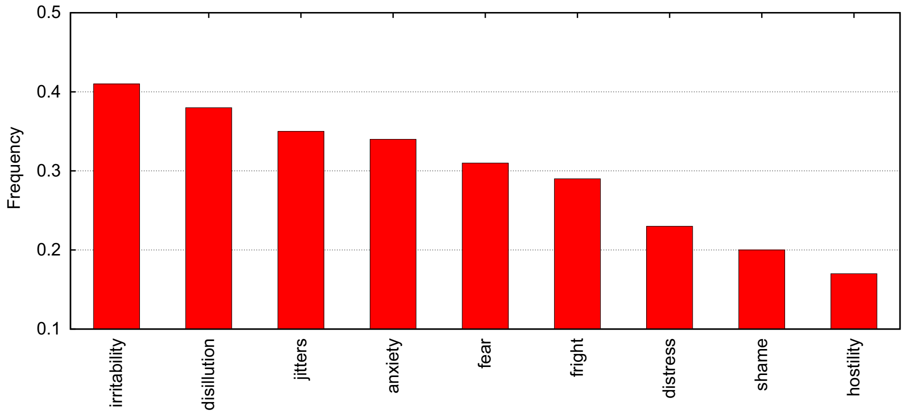

Fig. 9. Frequency of emotions. Note that frequencies do not sum up to 1 because the same reported issue may be associated with multiple emotions.

parameters used correspond to the ones that lead to the best performance.

4.3. Emotion classification

We next evaluate the emotion classification performance associated with different supervised learning algorithms, namely Naive Bayes (NB), Support Vector Machines (SVM), and Associative Classifier (AC). Fig. 10 shows p@k and r@k numbers for different values of k. Specifically, for each reported issue in the test set $\mathcal{T}$, we separate the k emotions associated with the highest scores, and compared this with the true set of emotions associated with the issue. All three algorithms show

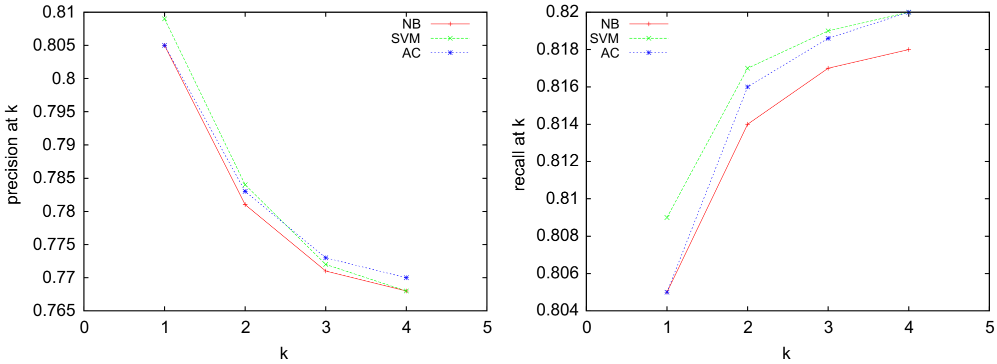

Fig. 10. Precision and recall numbers for different emotion classification algorithms.

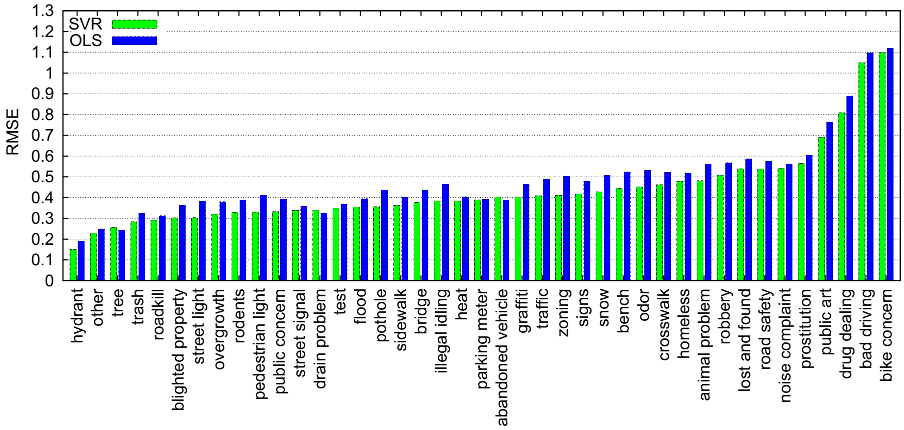

Fig. 11. RMSE numbers associated with different types of issue.

very similar performance. Precision numbers range from ≈0.805 to ≈0.770, while recall numbers range from ≈805 to ≈820, as the value of *k* increases. This suggests that supervised learning is an effective method for capturing emotions that are implicit in the reported issues.

We also evaluate possible precision improvements due to the inclusion of additional features, such as the author of the issue, the type of the issue, or the number of words in the textual description of the issue. The type of the issue showed to be the best additional feature, but still, it was unable to provide significant improvements when compared with the scenario where only the textual description is given to the emotion classification algorithm. Employing the length of the issue as additional feature (as well as the author), provided almost no improvement in terms of precision. From now on we will proceed with experiments that employ only the words in the textual description as features for emotion classification.

4.4. *Urgency estimate*

The urgency of an issue is given by the number of votes it has received. Now, we turn our attention to investigate how effectively we can estimate the number of votes an issue will receive based solely on the emotions extracted from its textual description (after the emotion classification step). Fig. 11 shows RMSE numbers for the different types of issue. Lower numbers are obtained for issues that usually receives less votes (e.g., hydrant and tree). Also, issues that usually receives more votes are also those with higher RMSE numbers (e.g., drug dealing and bad driving). Finally, SVR shows a significant superior performance when compared with OLS.

Table 1
MAP numbers for SVR and OLS algorithms.

Algorithm		MAP
SVR	Emotion distribution	0.833
OLS	Emotion distribution	0.791
SVR	Textual description	0.566
OLS	Textual description	0.526

Our last experiment is devoted to investigate the ranking performance of SVR and OLS algorithms. Both algorithms were evaluated under two scenarios: (i) the regression model was built from the distribution of emotions (after the emotion classification step), and (ii) the regression model was built directly from the textual description. Table 1 shows MAP numbers for SVR and OLS under these two evaluation scenarios. Again, SVR showed a significantly better performance when compared with OLS. Finally, regression models built from the distribution of emotions showed to be extremely better

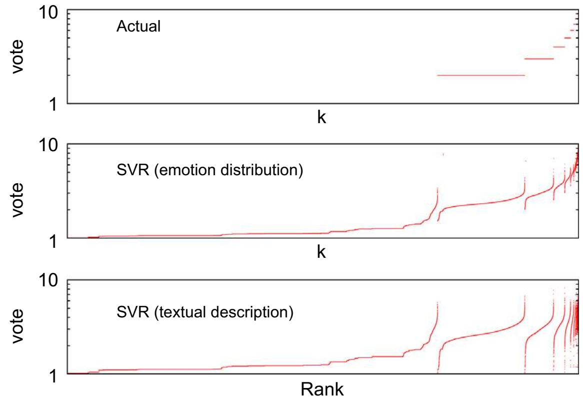

Fig. 12. Top– issues ranked according to their actual urgencies (i.e., number of votes). Middle and bottom – urgencies as predicted using emotion distribution and textual description.

than the regression models built directly from the textual descriptions. Fig. 12 shows the issues ranked by their actual urgencies. The figure also shows the predicted urgency of these issues, using either emotion distribution or textual description.

5. Conclusions

Continuous advances in technology are helping governments to develop innovative approaches to serving the public, making it easier to access government services, communicate with government officials, and make valuable government information readily available to the public. In this paper we investigate the important problem of anticipating the urgency of urban issues based solely on how citizens describe such issues. Often, citizens demonstrate emotions such as irritability, fear, or distress, while describing urban issues. We hypothesize that these emotions are indicative of the urgency of the reported issues, and we proposed a two-step approach to produce urgency estimates. The first step, emotion classification, is devoted to categorize an issue according to the emotions implicitly expressed in the textual description of the issue. The second step is devoted to estimate the urgency of the issue based solely on the emotions categorized during the first step. We conducted a systematic set of experiments using data obtained from the SeeClickFix hackathon, which demonstrated the effectiveness of the proposed approach to estimate the urgency of urban issues based on citizen emotions.

Acknowledgments

This work was partially sponsored by the Brazilian National Institute of Science and Technology for the Web (Grant MCT/CNPq 573871/2008-6) and AV's grants and scholarships from CNPq.

References

[1] A. Artikis, M. Weidlich, F. Schnitzler, I. Boutsis, T. Liebig, N. Piatkowski, C. Bockermann, K. Morik, V. Kalogeraki, J. Marecek, A. Gal, S. Mannor, D. Gunopulos, D. Kinane, Heterogeneous stream processing and crowdsourcing for urban traffic management, in: EDBT (2014) 712–723..
[2] D. Basak, S. Pal, D. Patranabis, Support vector regression, Neural Inf. Process.—Lett. Rev. 11 (10) (2007) 203–224.
[3] D. Bikel, J. Sorensen, If we want your opinion, in: ICSC, 2007, pp. 493–500.
[4] I. Boutsis, V. Kalogeraki, Crowdsourcing under real-time constraints, in: IPDPS, 2013, 753–764.
[5] I. Boutsis, V. Kalogeraki, On task assignment for real-time reliable crowdsourcing, in: ICDCS, 2014, pp. 1–10.
[6] C. Chang, C. Lin, LIBSVM: A library for support vector machines, ACM Trans. Intell. Syst. Technol. 2 (3) (2011) 27.
[7] P. Chaovalit, L. Zhou, Movie review mining: a comparison between supervised and unsupervised classification approaches, in: HICSS, 2005.
[8] S.A. Chun, S. Shulman, R. Sandoval, E. Hovy, Government 2.0: making connections between citizens, data and government, Inf. Polity 15 (1) (2010) 1–9.
[9] J. Gomide, A. Veloso, W. Meira Jr., V. Almeida, F. Benevenuto, F. Ferraz, M. Teixeira, Dengue surveillance based on a computational model of spatio-temporal locality of twitter, in: Web Science, 2011, p. 3.
[10] M. Handte, M. Iqbal, S. Wagner, W. Apolinarski, P. Marrón, E. Navarro, S. Martinez, S. Barthelemy, M. Fernández, Crowd density estimation for public transport vehicles, in: EDBT/ICDT Workshops, 2014, pp. 315–322.
[11] T. Hastie, R. Tibshirani, J. Friedman, The Elements of Statistical Learning: Data Mining, Inference and Prediction, Springer, New York, NY, USA, 2008.
[12] T. Joachims, A statistical learning model of text classification for support vector machines, in: SIGIR, 2001, pp. 128–136.
[13] S. Kanhere, Participatory sensing: crowdsourcing data from mobile smartphones in urban spaces, in: ICMDM, 2011, pp. 3–6.
[14] A.L. Kavanaugh, E.A. Fox, S.D. Sheetz, S. Yang, L.T. Li, D.J. Shoemaker, A. Natsev, L. Xie, Social media use by government: from the routine to the critical, Gov. Inf. Q. 29 (4) (2012) 480–491.
[15] I. Litou, I. Boutsis, V. Kalogeraki, Efficient dissemination of emergency information using a social network, in: EDBT/ICDT Workshops, 2014, pp. 347–354.
[16] B. Liu, L. Zhang, A survey of opinion mining and sentiment analysis, in: Mining Text Data, 2012, pp. 415–463.
[17] R. Lourenco Jr., A. Veloso, A. Pereira, W. Meira Jr., R. Ferreira, S. Parthasarathy, Economically-efficient sentiment stream analysis, in: SIGIR, 2014, pp. 637–646.
[18] P. Melville, W. Gryc, R. Lawrence, Sentiment analysis of blogs by combining lexical knowledge with text, KDD, In, 2009, 1275–1284.
[19] B. Pang, L. Lee, S. Vaithyanathan, Thumbs up? Sentiment classification using machine learning techniques, CoRR cs.CL/0205070, 2002.
[20] L. Qiu, W. Zhang, C. Hu, K. Zhao, SELC: a self-supervised model for sentiment classification, in: CIKM, 2009, pp. 929–936.
[21] F. Schnitzler, A. Artikis, M.Weidlich, I. Boutsis, T. Liebig, N. Piatkowski, C. Bockermann, K. Morik, V. Kalogeraki, J. Marecek, A. Gal, S. Mannor, D. Kinane, D. Gunopulos, Heterogeneous stream processing and crowdsourcing for traffic monitoring: Highlights, in: ECML/PKDD, 2014, pp. 520–523.
[22] I.S. Silva, J. Gomide, A. Veloso, W. Meira Jr., R. Ferreira, Effective sentiment stream analysis with self-augmenting training and demand-driven projection, in: SIGIR, 2011, pp. 475–484.
[23] S. Tan, Y. Wang, X. Cheng, Combining learn-based and lexicon-based techniques for sentiment detection without using labeled examples, in: SIGIR, 2008, pp. 743–744.
[24] D. Watson, L. Clark, A. Tellegen, Development and validation of brief measures of positive and negative affect: the panas scales, J. Personal. Soc. Psychol. 54 (6) (1988) 1063–1070.
[25] F. Wilcoxon, Individual comparisons by ranking methods, Biometrics 1 (1945) 80–93.
[26] R. Xia, C. Zong, S. Li, Ensemble of feature sets and classification algorithms for sentiment classification, Inf. Sci. 181 (6) (2011) 1138–1152.
[27] J. Zhu, H. Wang, B. Tsou, M. Zhu, Multi-aspect opinion polling from textual reviews, in: CIKM, 2009, pp. 1799–1802.
[28] J. Zhu, H. Wang, M. Zhu, B. Tsou, M. Ma, Aspect-based opinion polling from customer reviews, Trans. Affect. Comput. 2 (1) (2011) 37–49.