Social Media Image Analysis for Public Health

Venkata Rama Kiran Garimella
Aalto University
Helsinki, Finland
kiran.garimella@aalto.fi

Abdulrahman Alfayad
Carnegie Mellon University
Doha, Qatar
afayad@qatar.cmu.edu

Ingmar Weber
Qatar Computing Research Institute
Doha, Qatar
iweber@qf.org.qa

ABSTRACT

Several projects have shown the feasibility to use *textual* social media data to track public health concerns, such as temporal influenza patterns or geographical obesity patterns. In this paper, we look at whether geo-tagged *images* from Instagram also provide a viable data source. Especially for "lifestyle" diseases, such as obesity, drinking or smoking, images of social gatherings could provide information that is not necessarily shared in, say, tweets. In this study, we explore whether (i) tags provided by the users and (ii) annotations obtained via automatic image tagging are indeed valuable for studying public health. We find that both user-provided and machine-generated tags provide information that can be used to infer a county's health statistics. Whereas for most statistics user-provided tags are better features, for predicting excessive drinking machine-generated tags such as "liquid" and "glass" yield better models. This hints at the potential of using machine-generated tags to study substance abuse.

INTRODUCTION

The annual cost of "lifestyle diseases" in the U.S. is estimated at up to a *trillion* U.S. dollars. In order to target these diseases and induce required lifestyle changes it is useful to first gain a better understanding of the phenomenon through (i) identifying the population groups most affected, (ii) understanding the interaction between environment, population and other variables, and (iii) tracking changes over time to monitor the effect of any campaign or interventions.

Over the last couple of years, social media has emerged as a viable data source to use for large-scale public health studies. Using social media seems particularly appropriate for studying lifestyle diseases, such as obesity, as more and more of people's life and social interaction is publicly shared online.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
CHI'16, May 07 - 12, 2016, San Jose, CA, USA
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3362-7/16/05 ... $15.00
DOI: http://dx.doi.org/10.1145/2858036.2858234

One type of social media data that, to the best of our knowledge, has not been tapped into for large-scale public health studies is *images*. More than 1.8 billion images are shared online every day [1], (compared to 500M tweets per day), making rich media content a prolific data source.

In addition to its sheer size, social media image data might contain complementary data that is not shared in textual status updates. For example, after or even during a social gathering it is often customary to share images taken with a smart phone, rather than post a textual summary of the event. Images shared might or might not be annotated by the user. Until recently, automatically processing the *content* of images was beyond the realm of technical feasibility. Due to advances in deep learning [9], it has become possible to automatically understand objects in images and generate textual descriptions. In this paper, we look at whether the *content* of images shared on Instagram provides health-related information, and, in particular, whether machine-generated tags add value over user-provided ones.

The use of automatic image annotation tools holds a range of potential advantages one of which is language agnosticism. Here, "language" refers both to English vs. say, Spanish but it also refers to different and inconsistent conventions of what users might label as, say, "#healthy". Another advantage is the inference of "hidden" labels that an image owner would not assign. Studies have shown that a large fraction of young social media users have posted a picture depicting substance use [15]. Such images about alcohol or drug intake may not be explicitly labeled as, say, "#alcohol" and hence difficult to track using text-based methods. The insights obtained from this study could be of use in designing systems for understanding and monitoring substance abuse and other stigmatized behaviors.

We find that both user-provided and machine-generated labels for geo-tagged Instagram images hold valuable information in predicting county-level health statistics. Though for most variables user-provided labels outperform machine-generated ones, the opposite holds for "Excessive Drinking".

[1] `http://read.bi/1ouGzgQ`, last accessed on July 13, 2015.

RELATED WORK

Social media data for public health analysis Traditional methods to track public health metrics are based on pooling data from doctors, pharmacies and other sources which comes with certain drawbacks related to latency and cost. For tracking lifestyle diseases, where latency is not an issue, survey-based approaches have difficulties in collecting a "holistic" view of a patient that includes many facets of their lifestyle.

Recent studies have shown that a large scale, real time, non-intrusive monitoring can be done using social media to get aggregate statistics about the health and well being of a population [7, 23, 11]. Twitter in particular has been widely used in studies on public health [17, 21, 16, 12], due to its vast amount of data and the ease of availability of data.

Culotta [5] and Abbar et al. [1] used Twitter in conjunction with psychometric lexicons such as LIWC and PERMA to predict county-level health statistics such as obesity, teen pregnancy and diabetes. Paul et al. [18] make use of Twitter data to identify health related topics and use these to characterize the discussion of health online. Mejova et al. [14] use Foursquare and Instagram images to study food consumption patterns in the US, and find a correlation between obesity and fast food restaurants.

Social media has also enabled large-scale studies linking lifestyle and health data at an *individual* level. For example, Sadilek et al. [22], build a classifier to identify the health of a user based on their Twitter usage. Many more "hidden", yet important conditions such as depression [6, 4, 3], sleep problems [13], eating disorders [26], and substance use [20] have been studied using social media data. Our study is different from the ones discussed above in that we propose the use of image data to study public health.

Abdullah, et al. [2] use smile recognition from images posted on social media to study and quantify the overall societal happiness. Andalibi et al. [3] study depression related images on Instagram and "establish[ed] the importance of visual imagery as a vehicle for expressing aspects of depression". In our work, we study if the use of image recognition techniques helps in understanding a broader range of health related issues.

Image recognition, Deep learning Almost all the methods above rely on textual content though images and other rich multimedia form a major chunk of content being generated and shared in social media. Automatic image annotation has greatly improved over the last couple of years, owing to the development in deep learning [8]. Object recognition [27] and image tagging [9] have become possible because of these new developments, e.g. Karpathy et al. [9] use deep learning to produce descriptions of images, which compete with (and sometimes beat) human generated labels. A few studies already make use of these advances to identify [10] and study [25] food consumption from pictures.

DATA COLLECTION

The social media data used in our analysis consists of geo-referenced Instagram images (mostly taken at restaurants, related to food) obtained by Mejova et al. [14]. From this data, we are using the November subset consisting of 18M images. The choice of a smaller subset was randomly made to cater to the hard rate-limit restrictions imposed by Imagga and the choice for food related images was to maintain topical focus.

Each image is associated with various types of meta data, such as geographic location (latitude, longitude), user-provided hashtags, time stamp and comments. Using the (latitude, longitude), we mapped an image to a U.S. county (represented by a unique FIPS code) using an API provided by Federal Communication Commission.We then considered the top 100 counties in terms of image count and sampled 2,000 images uniformly at random from each of these 100 counties, due to computational restrictions imposed by the image annotation API.

We make use of two types of meta data obtained from images in this study: (i) machine-generated tags from a commercial service, and (ii) user-provided textual tags (hashtags).

Machine-Generated Tags

Each image was tagged using a state-of-the-art commercial image recognition tool from Imagga.com.[2] Figure 1 shows an example image along with the tags and their confidence. For 189,293 of our 200,000 images we managed to obtain at least one tag via Imagga. For the remaining 10,707 images the Instagram image was no longer publicly available, either due to deletion, changes in privacy settings, or account suspension by Instagram.[3]

Imagga API returns, for each image a set of tags along with a 'confidence' score. We only considered those tags with a confidence score above 20%. We experimented with different threshold values in the range of [5,60]. Based on the macro-averaged correlation for the nine different indices, we selected a global value of 20%. Though different threshold values might lead to slightly better results for a given index, we preferred not to explore this route to avoid the risk of overfitting. Tags appearing in less than 10 distinct counties were also ignored, leaving us with 1,750 distinct tags. Using this tag set, we created a representation of the tags used in each of the 100 counties by counting how many images were assigned a particular tag by Imagga. Imagga is not a restaurant- or food-specific annotation tool and its model is trained on general images. The ten most commonly returned machine annotations for the images in our data set were 'food', 'delicious','cocktail','people','gold'

[2] `http://docs.imagga.com/#auto-tagging`

[3] The Imagga tags dataset is available for download at `https://users.ics.aalto.fi/kiran/imagga_tags_data`

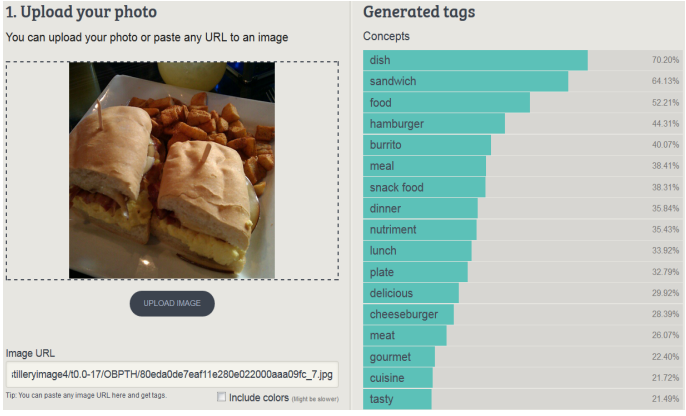

Figure 1. Example output for the Imagga Auto-Tagging Demo for one of the Instagram images in our data set. Each tag is associated with a confidence score (shown on the right).

, 'healthy', 'adult', 'caucasian', 'person', 'kitchen'. We observe that the tags contain references to food, restaurants and people. In general, the Imagga generated tags include a wide range of information from identifying types of food ('spaghetti'), drinks ('cappuccino'), objects ('stringed instrument'), people ('male'), and actions ('drunkard').

User-Provided Tags

Many images are provided with textual tags by the user uploading the image. For example, the image shown in Figure 1 comes with the tags #breakfastsandwich, #brunch, and #amazingpotatoes. Note that images might or might not have textual tags assigned by the user (in our dataset, 49% of the images have at least one tag). To have an equal size comparison to the machine-annotated data set, we sampled 200k images with at least one user-provided tag, 2,000 images for each of the 100 counties. Only user-provided tags appearing in at least 10 distinct counties were kept, leaving us with 5,865 distinct tags.

Demographic variables

As a baseline, we also obtained a list of demographic variables from the County Health Rankings website.[4] For each of the 100 counties we used same the demographic features that were used in Culotta [5], representing information related to (i) age distribution, (ii) racial distribution, and (iii) income. The ground-truth offline data on public health was obtained from the same County Health Rankings site. This data contains a wealth of health-related variables ranging from "Teen births" to "Diabetic monitoring". Given that our image data was obtained from *food* locations, we decided to focus our analysis on the following nine variables: (i) % *Smokers*, (ii) % *Obese*, (iii) *Food environment Index*, (iv) % *Physically inactive*, (v) % *Excessive drinking*, (vi) %*Alcohol – impaired driving deaths*, (vii) % *Diabetics*, (viii) %*Food insecure*, and (ix) *% Limited access to healthy food*.

[4] `http://www.countyhealthrankings.org/rankings/data`

METHODS

To test whether social media image data helps in studying public health, we perform regression to predict each of the nine health-related variables (dependent variables) using the image meta data features described above (independent variables). Given the high-dimensionality of our feature vectors relative to the number of validation points (N=100 counties), we use ridge regression to avoid over fitting (with a smoothing parameter of α=0.1 as in [5]). We perform 10 fold cross-validation, training a model on 90 counties and testing it on the remaining 10 to evaluate the accuracy of the regression.

For each county, the number of images containing a tag was used to create a county x tags matrix (size 100 x 1,750 for Imagga, 100 x 5,865 for user tags). We normalized each row (county) in this matrix using Euclidean norm in order to remove differences in image counts across counties. We also tried out various extensions of the text tags using external dictionaries such as LIWC [19] and PERMA [24]. We do not report results for these here as they did not lead to improvements, neither applied to human-generated nor to machine-generated tags.

To measure the accuracy, we used (i) Pearson's r - Correlation between the actual and predicted values (larger values are better), and (ii) Symmetric Mean Absolute Percentage Error (SMAPE, smaller values are better).

RESULTS

Table 1 summarizes the prediction performance for different combinations of (i) user-provided tags as features ("U"), (ii) machine-generated tags as features ("I"), and (iii) county-level demographics ("D"). We observe that using user-provided tags in conjunction with demographic features helps significantly improve the performance over just using demographic features, both in terms of Pearson's r and SMAPE.

To understand the performance of the models in more detail, we decided to zoom in on the four variables with the largest absolute improvement in Pearson's r over the demographics-only baseline: physically active (U+D), excessive drinking (I), alcohol impaired driving deaths (U), and limited access to healthy food (U).

Figure 2 shows scatter plots for these four health statistics. Each dot corresponds to a county with its actual health value on the x-axis and the predicted value, using the best-performing model, on the y-axis. As it might be of interest to domain experts, we labeled some of the outliers on the scatter plot with the name of the area.

According to our data, Arlington, VA "looks" healthier than it actually is, whereas the opposite holds for Hamilton, TN. In future work, it would be interesting to see what differences exist between such counties.

To understand *why* a particular model works, we performed a feature analysis for the same set of four statistics. In each case we computed the top 10 features with

Table 1. Prediction performance for nine health statistics across 100 counties in a 10-fold cross validation setting. Feature sets used were (U)ser tags, (I)magga tags, (D)emographics and combinations of these. Statistical significant improvement over the demographics-only baseline was tested using a Fisher r-z transform paired test between the two correlations * = .05, ** = .01, * = .001.**

	Pearson's r						SMAPE					
	U	I	D	U+D	I+D	U+I+D	U	I	D	U+D	I+D	U+I+D
Smokers	0.55	0.53	0.72	0.73	**0.75**	0.73	8.1	8.7	7.1	6.8	**6.7**	6.9
Obese	0.42	0.48	0.81	**0.84***	0.82	0.84	6.5	6.2	4	**3.7**	3.9	3.7
Food Env. Index	0.54	0.45	0.87	**0.91*****	0.87	0.90	4.2	4.3	2.4	**2**	2.4	2
Physically active	0.46	0.58	0.70	**0.79***	0.74	0.79	7.7	7.4	5.9	**4.8**	5.6	4.9
Excessive Drinking	0.22	**0.49****	0.38	0.33	0.48	0.38	6.2	**5.2**	6.1	6.3	5.7	6
Alcohol Impaired	**0.51****	0.26	0.25	0.43	0.28	0.44	**7.9**	8.9	9.1	8.5	8.9	8.2
Diabetic	0.34	0.40	0.77	**0.83****	0.78	0.82	7.7	7.2	5.1	**4.5**	4.9	4.6
Food insecure	0.37	0.38	0.84	0.86	0.84	**0.87**	7.8	7.2	4.4	**4.2**	4.5	4.3
Limited access	**0.61***	0.40	0.46	0.58	0.48	0.58	**19.7**	22.6	21.5	20.5	20.9	20.8

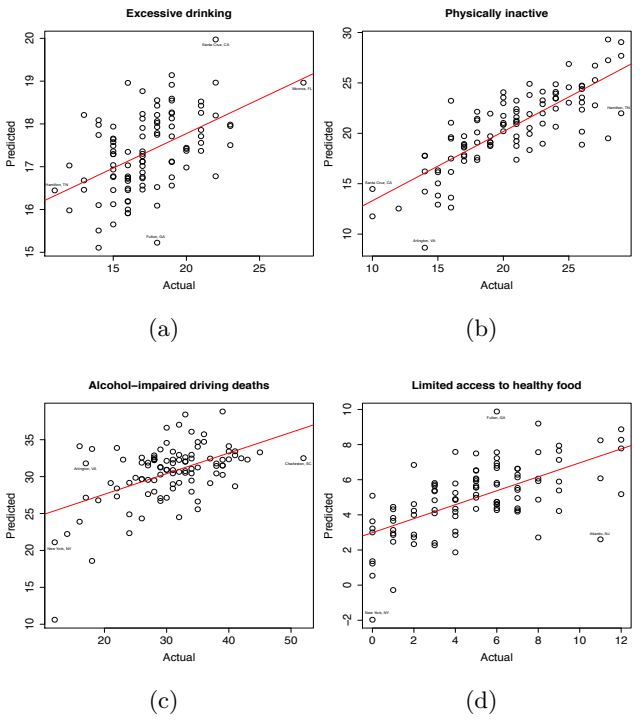

Figure 2. Predicted vs. actual health statistics value using (a) Imagga tags for predicting excessive drinking, (b) User-provided tags + demographics for predicting physically inactive, (c) User-provided tags for predicting alcohol-impaired driving deaths, and (d) User-provided tags for predicting limited access to healthy food.

the highest (absolute) value of correlation. Figure 3 shows the results.

Among the machine-derived tags most predictive of excessive drinking (Fig. 3 (a)), "liquid","beverage" and "meeting" (c.f. "party") and others make intuitive sense. These tags are especially interesting since such tags might not be part of human annotations. For the user-provided tags predictive of physical inactivity (Fig. 3 (b)), the hashtag #fatgirlstatus makes sense. It is used in a self-ironic manner to refer to food obsession and related topics.In addition to these intuitive features, there are also many that require further investigation. As the goal of this study was to test the general feasibility of

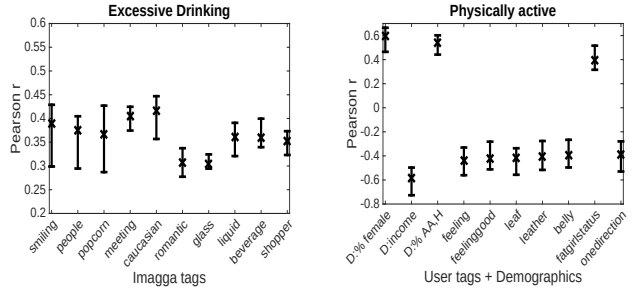

Figure 3. Feature analysis, showing the top ten features with the highest absolute value of correlation, for Imagga tags for predicting excessive drinking, (b) User-provided tags + demographics for predicting physically inactive. In (b), features with "D:" indicate demographic features, "D: AA,H" indicates demographic features pertaining to African Americans and Hispanics. Error bars indicate 95% confidence interval.

using *images* for public health studies, rather than to fully explain the social dynamics behind one particular health issue, we leave this for future work.

CONCLUSIONS

In this paper, we studied if social media image data can be used to understand public health. We used both user-provided as well as machine-generated tags to see if they can predict county-level health statistics. Our study relies on recent advances in deep learning for image analysis, a very active field that is undergoing constant advances. We show that both user-provided and machine-generated tags provide information that improves over models using only demographic information. Whereas in most cases, models based on user-provided tags outperform machine-generated ones, for statistics on excessive drinking the opposite holds. Machine-generated tags such as "glass", "liquid" and "beverage" (Figure 3) are all found to be good predictors, and such tags might not be part of human annotations. Machine-generated tags might prove of particular value in the case of stigmatized behaviors such as substance abuse, where explicit hints are unlikely to be added by the image owner. This could also be used in designing interfaces that could identify and help users in need of support.

REFERENCES

1. Sofiane Abbar, Yelena Mejova, and Ingmar Weber. 2015. You Tweet What You Eat: Studying Food Consumption Through Twitter. In *CHI*. 3197–3206.
2. Saeed Abdullah, Elizabeth L Murnane, Jean MR Costa, and Tanzeem Choudhury. 2015. Collective Smile: Measuring Societal Happiness from Geolocated Images. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 361–374.
3. Nazanin Andalibi, Pinar Ozturk, and Andrea Forte. 2015. Depression-related Imagery on Instagram. In *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing*. ACM, 231–234.
4. Sairam Balani and Munmun De Choudhury. 2015. Detecting and Characterizing Mental Health Related Self-Disclosure in Social Media. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1373–1378.
5. Aron Culotta. 2014. Estimating County Health Statistics with Twitter. In *CHI*. 1335–1344.
6. Munmun De Choudhury, Michael Gamon, Scott Counts, and Eric Horvitz. 2013. Predicting Depression via Social Media.. In *ICWSM*.
7. Mark Dredze. 2012. How Social Media Will Change Public Health. *IEEE Intelligent Systems* 27, 4 (2012), 81–84.
8. Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural computation* 18, 7 (2006), 1527–1554.
9. Andrej Karpathy and Li Fei-Fei. 2014. Deep visual-semantic alignments for generating image descriptions. *arXiv preprint arXiv:1412.2306* (2014).
10. Yoshiyuki Kawano and Keiji Yanai. 2015. FoodCam: A real-time food recognition system on a smartphone. *Multimedia Tools and Applications* 74, 14 (2015), 5263–5287.
11. Patty Kostkova. 2013. A Roadmap to Integrated Digital Public Health Surveillance: The Vision and the Challenges. In *WWW*. 687–694.
12. Patty Kostkova. 2015. Public Health. In *Twitter: A Digital Socioscope*, Yelena Mejova, Ingmar Weber, and Michael Macy (Eds.). Cambridge University Press, 111–130.
13. D.J. McIver, J.B. Hawkins, R. Chunara, A.K. Chatterjee, A. Bhandari, T.P. Fitzgerald, S.H. Jain, and J.S. Brownstein. 2015. Characterizing Sleep Issues Using Twitter. 17 (2015), e140.
14. Yelena Mejova, Hamed Haddadi, Anastasios Noulas, and Ingmar Weber. 2015. #FoodPorn: Obesity Patterns in Culinary Interactions. In *DigitalHealth*.
15. Elizabeth M Morgan, Chareen Snelson, and Patt Elison-Bowers. 2010. Image and video disclosure of substance use on social media websites. *Computers in Human Behavior* 26, 6 (2010), 1405–1411.
16. Jon Parker, Yifang Wei, Andrew Yates, Ophir Frieder, and Nazli Goharian. 2013. A Framework for Detecting Public Health Trends with Twitter. In *ASONAM*. 556–563.
17. Michael J. Paul and Mark Dredze. 2011. You Are What You Tweet: Analyzing Twitter for Public Health. In *ICWSM*. 265–272.
18. Michael J. Paul and Mark Dredze. 2014. Discovering Health Topics in Social Media Using Topic Models. *PLOS ONE* 9 (2014), e103408. Issue 8.
19. J. Pennebaker, J. Francis, and R. Booth. 2007. Linguistic inquiry and word count: LIWC 2001. (2007). `http://homepage.psy.utexas.edu/homepage/faculty/pennebaker/reprints/liwc2007_operatormanual.pdf`.
20. Salimian P.K., Chunara R., and Weitzman E.R. 2014. Averting the perfect storm: addressing youth substance use risk from social media use. *Pediatric Annals* 43 (2014), e242. Issue 10.
21. Kyle W Prier, Matthew S Smith, Christophe Giraud-Carrier, and Carl L Hanson. 2011. Identifying health-related topics on twitter. In *Social computing, behavioral-cultural modeling and prediction*. Springer, 18–25.
22. Adam Sadilek and Henry Kautz. 2013. Modeling the Impact of Lifestyle on Health at Scale. In *WSDM*. 637–646.
23. Abeed Sarker, Rachel E. Ginn, Azadeh Nikfarjam, Karen O'Connor, Karen Smith, Swetha Jayaraman, Tejaswi Upadhaya, and Graciela Gonzalez. 2015. Utilizing social media data for pharmacovigilance: A review. *Journal of Biomedical Informatics* 54 (2015), 202–212.
24. Martin Seligman. 2012. *Flourish: A Visionary New Understanding of Happiness and Well-being*. Atria Books.
25. Kyoko Sudo, Kazuhiko Murasaki, Jun Shimamura, and Yukinobu Taniguchi. 2014. Estimating Nutritional Value from Food Images Based on Semantic Segmentation. In *UbiComp*. 571–576.
26. Morgan Walker, Laura Thornton, Munmun De Choudhury, Jaime Teevan, Cynthia M. Bulik, Cheri A. Levinson, and Stephanie Zerwas. 2015. Facebook Use and Disordered Eating in College-Aged Women. *Journal of Adolescent Health* 57 (2015), 157–163.
27. Ren Wu, Shengen Yan, Yi Shan, Qingqing Dang, and Gang Sun. 2015. Deep image: Scaling up image recognition. *arXiv preprint arXiv:1501.02876* (2015).