HARVARD | BUSINESS | SCHOOL

# Using Text Analysis to Target Government Inspections: Evidence from Restaurant Hygiene Inspections and Online Reviews

Jun Seok Kang Polina Kuznetsova
Yejin Choi Michael Luca

## Working Paper

**14-007**

**July 12, 2013**

# Using Text Analysis to Target Government Inspections: Evidence from Restaurant Hygiene Inspections and Online Reviews

## — Working Paper —

**Jun Seok Kang** **Polina Kuznetsova** **Yejin Choi**
Department of Computer Science
Stony Brook University
Stony Brook, NY 11794-4400
{junkang,pkuznetsova,ychoi}
@cs.stonybrook.edu

**Michael Luca**
Harvard Business School
Soldiers Field Road
Boston, MA 02163
mluca@hbs.edu

### Abstract

Restaurant hygiene inspections are often cited as a success story of public disclosure. Hygiene grades influence customer decisions and serve as an accountability system for restaurants. However, cities (which are responsible for inspections) have limited resources to dispatch inspectors, which in turn limits the number of inspections that can be performed. We argue that NLP can be used to improve the effectiveness of inspections by allowing cities to target restaurants that are most likely to have a hygiene violation. In this work, we report the first empirical study demonstrating the utility of review analysis for predicting restaurant inspection results.

## 1 Introduction

Public health inspection records help customers to be wary of restaurants that have violated health codes. In some counties and cities, e.g., LA, NYC, it is required for restaurants to post their inspection grades at their premises, which have shown to affect the revenue of the business substantially (e.g., Jin and Leslie (2005), Henson et al. (2006)), thereby motivating restaurants to improve their sanitary practice. Other studies have reported correlation between the frequency of unannounced inspections per year, and the average violation scores, confirming the regulatory role of inspections in improving the hygiene quality of the restaurants and decreasing food-borne illness risks (e.g., Jin and Leslie (2003), Jin and Leslie (2009), Filion and Powell (2009), NYC-DoHMH (2012)).

However, one practical challenge in the current inspection system is that the department of health has only limited resources to dispatch inspectors, leaving out a large number of restaurants with unknown hygiene grades. We postulate that online reviews written by the very citizens who have visited those restaurants can serve as a proxy for predicting the likely outcome of the health inspection of any given restaurant. Such a prediction model can complement the current inspection system by enlightening the department of health to make a more informed decision when allocating inspectors, and by guiding customers when choosing restaurants.

Our work shares a similar spirit with recently emerging studies that explore social media analysis for public health surveillance, in particular, monitoring influenza or food-poisoning outbreaks from microblogs (e.g., Aramaki et al. (2011), Sadilek et al. (2012b), Sadilek et al. (2012a), Lamb et al. (2013), Dredze et al. (2013), von Etter et al. (2010)). However, no prior work has examined the utility of review analysis as a predictive tool for accessing hygiene of restaurants, perhaps because the connection is not entirely conspicuous: after all, customers are neither familiar with inspection codes, nor have the full access to the kitchen, nor have been asked to report on the hygiene aspects of their experience.

In this work, we report the first empirical study demonstrating the utility of review analysis for predicting health inspections, achieving over 82% accuracy in discriminating severe offenders from places with no violation, and find predictive cues in reviews that correlate with the inspection results.

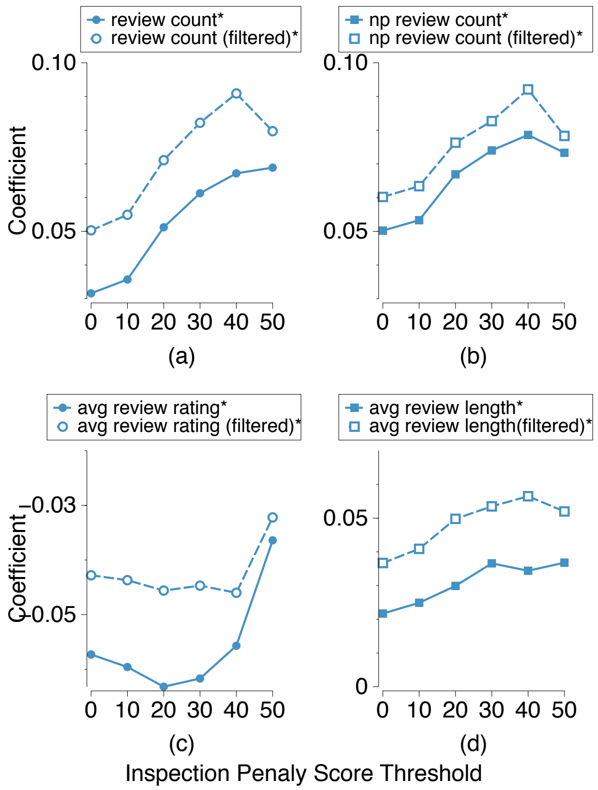


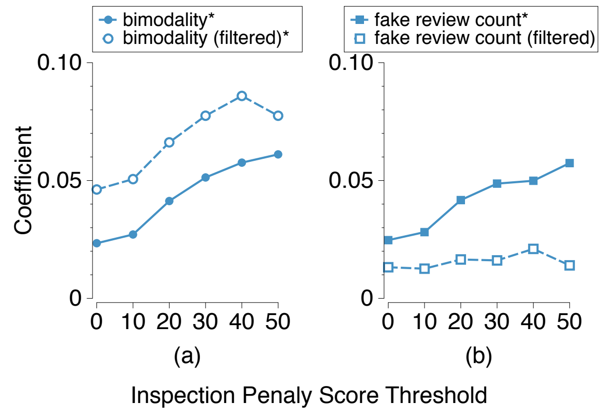


Figure 2: Spearman's coefficients of factors & inspection penalty scores. '*': statistically significant ($p \leq 0.05$)

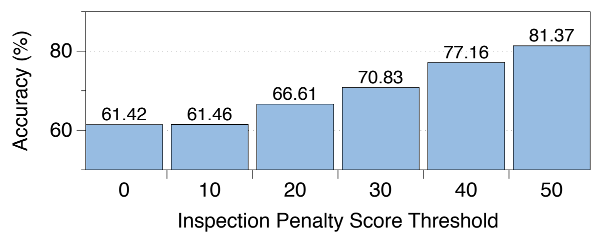


Figure 1: Spearman's coefficients of factors & inspection penalty scores. '*': statistically significant ($p \leq 0.05$)

Figure 3: Trend of penalty score thresholds & accuracies.

## 2 Data

We scraped entire reviews written for restaurants in Seattle from Yelp over the period of 2006 to 2013. The inspection records of Seattle is publicly available at `www.datakc.org`. More than 50% of the restaurants listed under Yelp do not have inspection records, implying the limited coverage of inspections.[1] After integrating reviews with inspection records, we obtained about 13k inspections over 1,756 restaurants with 152k reviews. For each restaurant, there are typically several inspection records. We define an "inspection period" of each inspection record as the period of time starting from the day after the previous inspection to the day of the current inspection.[2] Each inspection period corresponds to an instance in the training or test set. We merge all reviews within an inspection period into one document when creating the feature vector.

[1]We convert street addresses into canonical forms when matching restaurants between Yelp and inspection database.

[2]If there is no previous inspection, then the period stretches to the past 6 months in time.

Note that non-zero penalty scores may not necessarily indicate alarming hygiene issues. For example, violating codes such as *"proper labeling"* or *"proper consumer advisory posted for raw or undercooked foods"* seem relatively minor, and unlikely to be noted and mentioned by reviewers. Therefore, we focus on restaurants with severe violations, as they are exactly the set of restaurants that inspectors and customers need to pay the most attention to. To define restaurants with "severe violations" we experiment with a varying threshold $t$, such that restaurants with score $\geq t$ are labeled as *"unhygienic"*.[3]

## 3 Correlates of Inspection Penalty Scores

We examine correlation between penalty scores and several statistics of reviews:

**I. Volume of Reviews:**

- *count of all reviews*
- *average length of all reviews*

[3]For restaurants with *"hygienic"* labels, we only consider those without violation, as there are enough number of such restaurants to keep balanced distribution between two classes.

**II. Sentiment of Reviews:**

- *average review rating*
- *count of negative ($\leq 3$) reviews*

**III. Deceptiveness of Reviews:** Restaurants with unhappy customers might be more motivated to hire fake reviewers.[4]

- *bimodal distribution of review ratings*
  The work of Feng et al. (2012) has shown that the shape of the distribution of opinions, sharpe bimodal distributions in particular, can be a telltale sign of deceptive reviewing activities.[5]
- *volume of fake reviews*
  We also explore the use of deception classifier to measure the amount of fake reviews.[6]

**Filtering Reviews:** When computing above statistics over the set of reviews corresponding to each restaurant, we also consider removing a subset of reviews that might be dubious or just noise. In particular, we remove reviews that are too far away (delta $\geq 2$) from the average review rating. Another filtering rule can be removing all reviews that are classified as deceptive by the deception classifier explained above. For brevity, we only show results based on the first filtering rule, as we did not find notable differences in different filtering strategies.

**Results:** Fig 1 and 2 show Spearman's rank correlation coefficient with respect to the statistics listed above, with and without filtering, computed at different threshold cutoffs $\in \{0, 10, 20, 30, 40, 50\}$ of inspection scores. Although coefficients are not strong,[7] they are mostly statistically significant with $p \leq 0.05$ (marked with '*'), and show interesting contrastive trends as highlighted below.

[4]It is also possible that some of the most assiduous restaurants that abide by health codes strictly might also stay diligent in soliciting fake positive reviews.

[5]We approximate this by computing the variance of review ratings. Naturally the distribution of online review rating is bimodal, rather than unimodal (Hu et al., 2009).

[6]For this purpose we collected a set of fake reviews and truthful reviews (250 reviews for each class) in the restaurant domain, following the data collection method introduced by Ott et al. (2011). Based on unigram and bigram features, we were able to achieve 79.2% accuracy with 10 fold cross validation.

[7]Spearman's coefficient assumes monotonic correlation. We suspect that the actual correlation of these factors and inspection scores are not entirely monotonic.

| Features | Acc. | MSE | SCC |
|---|---|---|---|
| - | *50.00 | 0.500 | - |
| review count | *50.00 | 0.489 | 0.0005 |
| np review count | *52.94 | 0.522 | 0.0017 |
| cuisine | *66.18 | 0.227 | 0.1530 |
| zip code | *67.32 | 0.209 | 0.1669 |
| avrg. rating | *57.52 | 0.248 | 0.0091 |
| inspection history | *72.22 | 0.202 | 0.1961 |
| unigram | 78.43 | 0.461 | 0.1027 |
| bigram | *76.63 | 0.476 | 0.0523 |
| unigram + bigram | 82.68 | 0.442 | 0.0979 |
| all | 81.37 | 0.190 | 0.2642 |

Table 1: Feature Compositions & Respective Accuracies, Respective Mean Squared Errors(MSE) & Squared Correlation Coefficients (SCC), np=non-positive

In Fig 1, as expected, average review rating is negatively correlated with the inspection penalty scores. Interestingly, all three statistics corresponding to the volume of customer reviews are positively correlated with inspection penalty, and if the reviews are filtered, then the correlation gets stronger, suggesting the existence of deceptive reviews. Also notice that correlation is generally stronger when higher cutoffs are used (x-axis), as expected. Fig 2 looks at the relation between deception level and the inspection scores more directly. As suspected, restaurants with high penalty scores show increased level of deceptive reviews. Although correlation coefficients are insightful, those alone are not informative enough to be used as a predictive tool, hence we explore content-based classification next.

## 4 Content-based Prediction

We examine the utility of the following features:

**Features based on customers' opinion:**

1. Aggregated opinion: average review rating
2. Content of the reviews: unigram, bigram

**Features based on restaurant's metadata:**

3. Cuisine: e.g., Thai, Italian, as listed under Yelp
4. Location: first 5 digits of zip code
5. Inspection History: a boolean feature ("hygienic" or "unhygienic"), a numerical feature (previous penalty score rescaled $\in [0, 1]$), a numeric feature (average penalty score over all previous inspections)
6. Review Count
7. Non-positive Review Count

| |
|---|
| **Hygienic** gross, mess, sticky, smell, restroom, dirty |
| **Basic Ingredients**: beef, pork, noodle, egg, soy, ramen, pho, |
| **Cuisines** Vietnamese, Dim Sum, Thai, Mexican, Japanese, Chinese, American, Pizza, Sushi, Indian, Italian, Asian |
| **Sentiment**: cheap, never, |
| **Service & Atmosphere** cash, worth, district, delivery, think, really, thing, parking, always, usually, definitely<br>- door: "The wait is always out the *door* when I actually want to go there",<br>- sticker: "I had *sticker* shock when I saw the prices.",<br>- student: "heap, large portions and tasty = the perfect *student* food!",<br>- the size: "i was pretty astonished at *the size* of all the plates for the money.",<br>- was dry: "The beef *was dry*, the sweet soy and anise-like sauce was TOO salty (almost inedible).",<br>- pool: "There are *pool* tables, TV airing soccer games from around the globe and of course - great drinks!" |

Table 2: Lexical Cues & Examples - Unhygienic (dirty)

**Classification Results** We use liblinear's SVM (Fan et al., 2008) with L1 regularization and 10 fold cross validation. We filter reviews that are farther than 2 from the average rating. We also run Support Vector Regression (SVR) using lib linear. Fig 3 shows the results. As we increase the threshold, the accuracy also goes up in most cases. Table 1 shows feature ablation at threshold $t = 50$, and '*' denotes statistically significant (p$\leq$0.05) difference over the performance with all features based on student t-test.

We find that metadata information of restaurants such as location and cuisine alone show good predictive power, both above 66%, which are significantly higher than the expected accuracy of random guessing (50%). Aggregated opinion, which is the average review rating during the corresponding inspection period, also performs substantially better than chance (57.52%), but that alone is not a strong predictor. Interestingly, the inspection history feature alone is highly informative, reaching accuracy unto 72%, suggesting that past performance is a good predictor of future performance. Textual content of the reviews (unigram+bigram) turn out to be the most effective features, reaching upto 82.68% accuracy. Lastly, when all the features are combined together, the performance decreases slightly to 81.37%, perhaps because n-gram features perform drastically better than all others.

| |
|---|
| **Hygienic**: |
| **Cooking Method & Garnish**: brew, frosting, grill, crush, crust, taco, burrito, toast<br>**Healthy or Fancier Ingredients**: celery, calamity, wine, broccoli, salad, flatbread, olive, pesto |
| **Cuisines** : Breakfast, Fish & Chips, Fast Food, German, Diner, Belgian, European, Sandwiches, Vegetarian |
| **Whom & When**: date, weekend, our, husband, evening, night |
| **Sentiment**: lovely, yummy, generous, friendly, great, nice |
| **Service & Atmosphere**: selection, attitude, atmosphere, ambiance, pretentious |

Table 3: Lexical Cues & Examples - Hygienic (clean)

### 4.1 Insightful Cues

Table 2 and 3 shows representative lexical cues for each class with example sentences excerpted from actual reviews when context can be helpful.

**Hygiene:** Interestingly, hygiene related words are overwhelmingly negative, e.g., *"gross", "mess", "sticky"*. What this suggests is that reviewers do complain when the restaurants are noticeably dirty, but do not seem to feel the need to complement on cleanliness as often. Instead, they seem to focus on other positive aspects of their experience, e.g., details of food, atmosphere, and their social occasions.

**Service and Atmosphere:** Discriminative features reveal that it is not just the hygiene related words that are predictive of the inspection results of restaurants. It turns out that there are other qualities of restaurants, such as service and atmosphere, that also correlate with the likely outcome of inspections. For example, when reviewers feel the need to talk about *"door", "student", "sticker"*, or *"the size"* (see Table 2 and 3), one can extrapolate that the overall experience probably was not glorious. In contrast, words such as *"selection", "atmosphere", "ambiance"* are predictive of hygienic restaurants, even including those with slightly negative connotation such as *"attitude"* or *"pretentious"*.

**Whom and When:** If reviewers talk about details of their social occasions, it seems to be a good sign.

**The way food items are described:** Another interesting aspect of discriminative words are the way food items are described by reviewers. In general, mentions of basic ingredients of dishes, e.g., *"noodle", "egg", "soy"* do not seem like a good sign. In contrast, words that help describing the way dish is prepared or decorated, e.g., *"grill", "toast", "frosting", "bento box" "sugar"* (as in "sugar coated") are good signs of satisfied customers.

**Cuisines:** Finally, cuisines have clear correlations with inspection outcome, as shown in Table 2 and 3.

**Conclusion:** We have reported the first study demonstrating the promise of review analysis for predicting health inspections, introducing a task that has potentially significant societal benefits, while being relevant to much research in NLP for opinion analysis based on customer reviews.

## References


Eiji Aramaki, Sachiko Maskawa, and Mizuki Morita. 2011. Twitter catches the flu: Detecting influenza epidemics using twitter. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1576, Edinburgh, Scotland, UK., July. Association for Computational Linguistics.

Mark Dredze, Michael J. Paul, Shane Bergsma, and Hieu Tran. 2013. Carmen: A twitter geolocation system with applications to public health. In *AAAI Workshop on Expanding the Boundaries of Health Informatics Using AI (HIAI)*.

Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. Liblinear: A library for large linear classification. *The Journal of Machine Learning Research*, 9:1871–1874.

Song Feng, Longfei Xing, Anupam Gogar, and Yejin Choi. 2012. Distributional footprints of deceptive product reviews. In *ICWSM*.

Katie Filion and Douglas A Powell. 2009. The use of restaurant inspection disclosure systems as a means of communicating food safety information. *Journal of Foodservice*, 20(6):287–297.

Spencer Henson, Shannon Majowicz, Oliver Masakure, Paul Sockett, Anria Johnes, Robert Hart, Debora Carr, and Lewinda Knowles. 2006. Consumer assessment of the safety of restaurants: The role of inspection notices and other information cues. *Journal of Food Safety*, 26(4):275–301.

Nan Hu, Jie Zhang, and Paul A. Pavlou. 2009. Overcoming the j-shaped distribution of product reviews. *Commun. ACM*, 52(10):144–147, October.

Ginger Zhe Jin and Phillip Leslie. 2003. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2):409–451.

Ginger Zhe Jin and Phillip Leslie. 2005. The case in support of restaurant hygiene grade cards.

Ginger Zhe Jin and Phillip Leslie. 2009. Reputational incentives for restaurant hygiene. *American Economic Journal: Microeconomics*, pages 237–267.

Alex Lamb, Michael J. Paul, and Mark Dredze. 2013. Separating fact from fear: Tracking flu infections on twitter. In *the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.

NYC-DoHMH. 2012. Restaurant grading in new york city at 18 months. *New York City Department of Health and Mental Hygiene*.

Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T. Hancock. 2011. Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 309–319, Portland, Oregon, USA, June. Association for Computational Linguistics.

Adam Sadilek, Henry Kautz, and Vincent Silenzio. 2012a. Predicting disease transmission from geotagged micro-blog data. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*.

Adam Sadilek, Henry A. Kautz, and Vincent Silenzio. 2012b. Modeling spread of disease from social interactions. In John G. Breslin, Nicole B. Ellison, James G. Shanahan, and Zeynep Tufekci, editors, *ICWSM*. The AAAI Press.

Peter von Etter, Silja Huttunen, Arto Vihavainen, Matti Vuorinen, and Roman Yangarber. 2010. Assessment of utility in web mining for the domain of public health. In *Proceedings of the NAACL HLT 2010 Second Louhi Workshop on Text and Data Mining of Health Documents*, pages 29–37, Los Angeles, California, USA, June. Association for Computational Linguistics.