Predictive Models for Food Code Violations

Jim Hartman, R.S., M.S.

Retired Public Health Sanitarian, 196 S. Grant Ave., Columbus, OH 43215, USA

SUMMARY

Many health departments use hypothetical "risk classification" concepts, often including establishments' use of "risky" food processing procedures, to prioritize inspections. In this paper, descriptions of some strategies for scheduling inspections based on this sort of perceived risk will be followed by a review of more fine-tuned classification models based on inspection history and other factors. A strategy for scheduling inspections based on detection of sporadic cases of foodborne illness will then be outlined. The advent of social media, machine learning and portable digital technology (smart phones) have begun to revolutionize this field, although some of the investigative techniques used in earlier work have not yet been applied to the newer approaches. The goals of this review are to bring some of these strategies together and offer recommendations for advancing them.

OVERVIEW

Through much of the latter part of the twentieth century, public health practitioners searched for ways to assess the risk of a restaurant causing an outbreak of foodborne disease. The goal of predictive modeling in this context is to use risk for scheduling and focusing inspection programs. Because of the difficulty of measuring risk directly, easier-to-measure characteristics or performance measures have been proposed.

Because we often hear the words "analog" and "digital," it may be instructive to ponder exactly what they mean. An analog clock has hands that sweep around a dial to indicate the passage of time. However, this movement "is not time itself: it's a representation or analogy of time" *(28)*. In a similar way, the presence of certain features is not risk itself, but can represent it. A more complicated computer model, to be described later, uses concepts from earlier models using categories, but without employing any analogies to risk. Other models to be described really are digital: they use the presence or absence of illness, similar to 1s and 0s. Of course, merely being digital does not imply being more accurate than analog methods.

All these methods have the potential to benefit restaurant patrons by identifying food safety problems earlier than they would otherwise be found, thus preventing many cases of foodborne illness. In addition, because some of these illnesses are due to infectious agents, the general public, including people who rarely or never eat out, could benefit. Application of predictive modeling could help public health agencies in charge of conducting inspections by helping them achieve their public health goals even if they miss some of their mandated inspections. Finally, restaurants themselves stand to benefit in that their food preparation practices could be improved and their legal liability diminished.

However, much of the work in these areas has appeared in information technology journals rather than in food safety or public health publications. Furthermore, information technologists have not applied food safety concepts such as HACCP. This article is an attempt to spread the word about these methods to a broader audience and, it is hoped, improve them.

EARLY MODELS

The U.S. Food and Drug Administration (FDA) promulgated the United States' first "Ordinance regulating eating and drinking establishments," a mimeographed document, in 1935 *(25)*. The 1976 *Food Service Sanitation Manual (24)* recommended semi-annual inspections. The idea of adjusting the inspection frequencies of food service operations (restaurants) and retail food establishments (markets, etc.) according to some theory of their risk of causing illness has been evolving ever since.

Kaplan and EI-Ahraf *(14)* were among the first researchers to tabulate data on reported outbreaks of foodborne illness for the purpose of estimating risk according to the type of establishment involved. Ten years later, Irwin et al. *(11)* set up a case-control study to examine the violations reported on the last routine inspection report for the restaurant that had caused each of 28 outbreaks. Controls were matched to them by routine inspection date. The best predictor of which food services would later cause an outbreak was found to be "any improper food protection practice" *(11)*. Bryan *(3)* suggested basing inspection frequencies on the presence of foods often implicated as vehicles of foodborne illness, risky food processing steps, and average daily patronage. Briley and Klaus *(2)* used ideas from Bryan, as well as Kaplan and EI-Ahraf's, and added a statistic based on the average score from the last five inspections. Wodi and Mill *(27)* also used a predicted risk score based on the last two inspection scores and critical items violated.

Columbus models

Two predictive models have been proposed at the Columbus, Ohio, Health Department (now Columbus Public Health).

One model had some of the predictor variables in the earlier models already cited but also had some important differences, including extensive use of existing inspection records. The other model went in a different direction, without having any analogies to risk.

Classification and Regression Trees (CART)

A project in Columbus, Ohio, proposed using Classification and Regression Trees (CART) *(1)* software with food safety inspection data provided by the City to develop a new predictive model. Early applications of such a "tree-structured approach" to data analysis *(1)* included a project in which an airplane flying in large circles around ships of six different structural types produced continuous fluctuating radar signals that were used to classify the ships. Another application was analysis of mass spectra of airborne contaminants to characterize specific pollutants. Both of these applications are actually examples of artificial intelligence: machine learning without specification of the variables to use. The class intervals of the continuous random variables were selected by the computer software.

This graduate student project *(8)* used the aforementioned automated nonparametric statistical methodology in 1992 in an attempt to show that the outcome of the most recent inspection of a food service operation could be predicted through use of the preceding series of inspections between 1986 and 1990. The outcomes of interest were time-temperature violations or inspection failure (a score below 90 out of 100 and/or a "critical violation," one capable of causing an outbreak of foodborne illness). This was similar to the previous examples in that a variety of independent variables *(Table 1)* were available, but the computer selected the ones with the most predictive power.

The results of one classification tree are shown in *Fig. 1*. In a "learning sample" of 1,000 full-menu restaurants, 528 had failed a standard inspection at least once in 5 years. CART made its first split with the question "Was the standard deviation of scores in the previous year above 1.95?" 398 of the 573 establishments for which the answer was "yes" had failed at least once. The bar at the right end of the box indicates that it is a "terminal node," and it was not split further. Of the 427 for which the first answer was "no," 159 establishments that did not have an extra inspection included 86 failures and could not be split further. A split of the 268 establishments that had had an extra inspection identified 41 failures by asking whether the average interval between inspections was 241 days or more. Finally, of those 41, 36 had had 2 or 3 extra inspections in the previous year.

The results of the analyses generally make sense. Restaurants with a variable score are not under effective control and therefore could be expected to have problems. If not getting an extra inspection predicts failure, the inspections are generally doing what they're supposed to be doing. Also, perhaps going without an inspection for 241 days has a protective

TABLE 1. Potential predictor variables; adapted from references *8* and *22*

AVAILABLE TO CART IN COLUMBUS	USED IN CHICAGO FORECASTING MODEL
Average number of days between inspections, prior year	Length of time since last inspection
Standard deviation of scores, prior year	Three-day average high temperature
Number of extra inspections, prior year	Nearby garbage and sanitation complaints
Days since last regular inspection	
Average duration in minutes, prior year	Has a tobacco or alcohol consumption
Average income in zip code	Nearby burglaries
Type of operation (food service or retail food establishment)	Type of facility
Vending location or regular	Length of time operation has been operating
Frequency of the outcome of interest, prior year	
Any food safety violation, prior year	Had a previous critical or serious violation
Commercial versus non-commercial	
Any violation, any year	
Number of previous violations	
Purpose of the index inspection	Inspector assigned

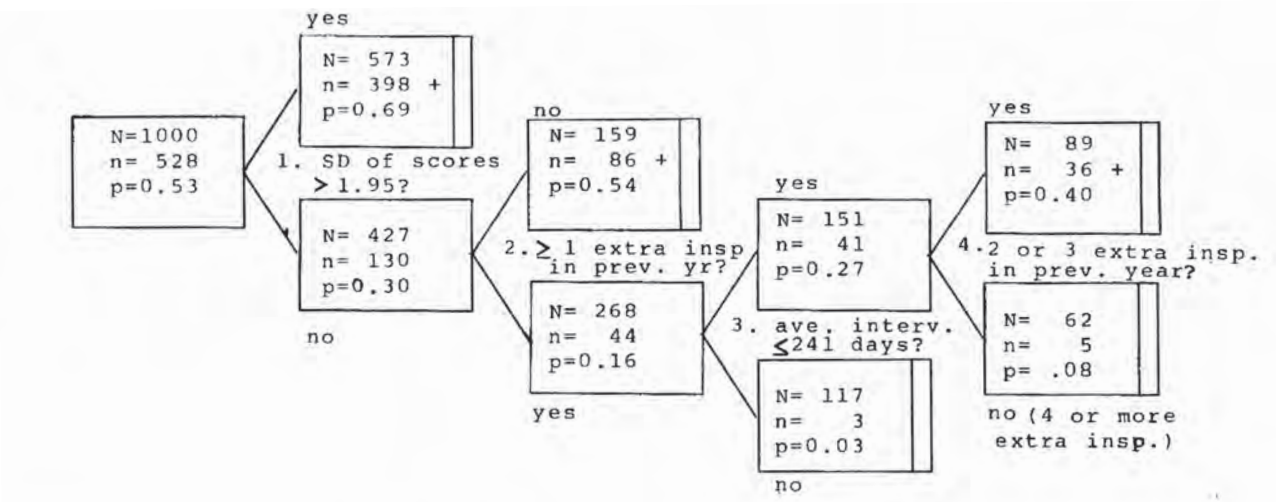

FIGURE 1. Inspection failure rates in a learning sample of 1,000 full-menu restaurants as a function of risk markers identified by CART.

Questions 1 through 4 about the risk markers classified 528 establishments that failed at least once, and others that never failed in 5 years, into groups with higher and lower failure rates (p).

Figure 1. Example of CART output *(8).*

effect in that if sanitarians know which restaurants will get along without an inspection for longer than recommended, they are likely to skip them.

This model was never deployed, with no reason ever offered. However, it had never been requested by the City of Columbus, and no city time was used for its development. The Ohio State University provided the statistical software for initial processing of the data as well as the CART software. At that time (1992), the City of Columbus may not have had the necessary data processing resources; it had not yet even computerized its complaint system *(8)*. The model is mentioned here because it seemed to work, has potential as a predictive model, and is in general agreement with the first Chicago model, described later; as *Table 1* shows, its input was remarkably similar. The CART model also had some of the predictor variables in the earlier models previously cited, but with some important differences: it could be run separately for different kinds of food operation (e.g., markets and food service operations); it could predict various outcomes (such as inspection failures or time-temperature violations, specifically); and it was self-updating, in the sense that after the results of inspections it had scheduled were entered, CART could be re-run to predict the next batch of inspection results.

Surveillance system using a weight-loss website

The Columbus Health Department (now Columbus Public Health) applied in 1995 and 1997 for grants from the U.S. Department of Agriculture Cooperative State Research, Education, and Extension Service National Research Initiative Competitive Grants Program (NRICGP) to fund a surveillance system for sporadic foodborne illness, using a weight-loss website. After health department management priorities changed, the Scioto Valley Health Systems Agency (SVHSA) applied for the same grant for the same purpose in 2000. SVHSA was a consortium of 15 county and city health departments, including the Columbus Health Department. The original proposal *(9)* was modified after each grant application and has also been described in a short article in *Food Protection Trends (10)*.

The basic idea was to identify potential sporadic cases of foodborne disease originating with restaurants among users of a weight-loss website that would allow people with weight problems or other conditions requiring dietary monitoring to log on and enter a daily food history. It was expected that a small number would experience symptoms of a foodborne disease later and that they would be able to report them on the site. This would minimize problems of recall bias, etc. The Health Department would investigate commercial sources of food with methods usually used in outbreak investigations. The proposal included use of ATP bioluminometers to focus environmental testing and immunomagnetic separation, followed by PCR, to expedite microbiological testing. Offering a bounty for stool samples was proposed to obtain clinical specimens from squeamish clients. By using a prospective rather than a retrospective design, the system would show the actual risk associated with specific foods and processing errors and determine how much deviation from control is necessary to cause illness. For example, we know the percentage of outbreaks caused by improper cooling, but we do not know the percentage of improperly cooled foods that cause illness. What is the probability that bean soup cooled in a huge pot for eight hours will cause illness? To answer this question, we would have to investigate meals eaten by people who did and people who did not get sick. Many (or most) inspections in the jurisdiction would be targeted to investigate restaurants or other sources of foods identified by the web site. A data-processing system capable of tabulating and comparing food processes used by restaurants involved with cases and those involved with controls would be necessary.

For example, suppose the case had cramps and diarrhea and a stool test revealed the presence of *Bacillus cereus*. As *Table 2* shows *(4)*, attention would focus on foods eaten 7–12 hours earlier. *Table 3* *(26)* suggests focusing attention on starchy foods that may have been held at incubating temperatures. Under the theory that improper hot holding of starchy food caused the illness, either such a food would not appear in the food history of any controls matched to this case, or the food would have been held at a safe temperature.

NRICGP reviewers rejected each proposal. The 1996 review panel summary, for example, suggested that "the proposed work may exceed the capabilities of the [Principal Investigator] … the proposed testing using PCR methods may not be readily available and may be extremely costly … [and a] study design limited to a small case-control approach to investigate foodborne exposures may provide a clearer focus for future studies." This proposal is cited here because it used a unique method for identifying cases, and the case follow-up methods appear to be applicable to later work.

RECENT MODELS

For many decades, it has been possible to call in complaints to the health department by telephone; now it is also easy to file complaints via government websites. These methods, however, are notoriously subject to underreporting. In the past few years, papers have appeared about the use of social media to monitor disease – so-called infodemiology and infoveillance *(18)*. Oldyrod et al. *(18)* did a computer-assisted search of 5 databases and reviewed 5,239 papers discussing use of these methods for disease surveillance and found the papers focused mostly on influenza-like illness; only 10 discussed use of consumer-generated data for foodborne illness surveillance.

Harrison et al. *(7)* used custom software to analyze 294,000 New York City restaurant reviews on the business review website Yelp (www.yelp.com) and considered 468 of them to have been reportable as foodborne illness outbreaks, although only 3% of those had actually been reported to the New York City Department of Health and Mental Hygiene. Of 129 reviewers describing 2 or more sick persons or scombroid poisoning or a neurological illness, 102 refused to be interviewed by investigators.

Kang et al. *(13)* also used Yelp reviews and claimed to be the first to compare them with official restaurant inspection results. They "scraped" Seattle restaurant reviews from 2006 to 2013 and found that over half the restaurants had no corresponding inspection records. They claim 82% accuracy in predicting the inspection results where these results existed. Their work generated tables similar to Table 4. Interestingly, they found that restaurants that reviewers called "pretentious" were hygienic, whereas when reviews called restaurants "cheap" or singled out specific food ingredients, e.g., noodles, "one can extrapolate that the overall experience probably was not glorious."

These methods do not seem to depart radically from earlier methods, and they have a critical drawback: they are time- and resource-consuming. However, they also have an important advantage: they supplement traditional reporting systems.

Chicago's food safety inspection forecasting model

In 2002, Chicago Mayor Rahm Emanuel issued an executive order making "Open Government Data" a priority and requiring all city agencies to make "all appropriate data sets" available for public use *(5)*. The city won a grant from Bloomberg Philanthropies in 2013 to develop innovations they could share with other cities. A working group, including the data science team at Allstate Insurance, chose the idea of a predictive model for food code violations partly because the city had 15,000 food service operations but only 36 food sanitarians *(17)*, which comes out to 2.4 hours per inspection, including travel time, if each location received an average of 2 inspections per year. The model they developed was tested over 2 months in 2014 and went into use in 2015 *(17)*.

The right side of *Table 1* shows the 9 predictor variables Chicago uses *(22)*. Chicago tested the model by using it to generate inspection schedules and compared those inspection results to those resulting from inspection assignments generated in the ordinary way. Managers did not inform sanitarians which lists were generated by the model. The result was that 55% of the inspections scheduled in the usual way found critical violations, whereas 69% of the inspections scheduled by the model resulted in critical violations. Also, the model identified the problem restaurants (of which there were 35) 7.5 days earlier than the normal scheduling method would have.

Chicago published their model as "open source code" so that it could be used and possibly improved on in other jurisdictions. But as of early 2016, only one other jurisdiction had used it – Montgomery County, Maryland, near Washington, D.C. They also found 27% more violations in the first month than would have been found with the normal scheduling and found them 7.5 days earlier *(22, 23)*.

TABLE 2. Classification of acute enteric diseases by symptoms, incubation periods, and agents (adapted from references *4* and *9*)

Disease group	Predominating or initial symptoms	Incubation (hours)	Likely etiologic agents
Upper GI	Nausea, vomiting	< 1	Heavy metals
		1–6	*Bacillus cereus*[a]
			Staph. aureus
		7–12	Mushrooms[b]
Sore throat & respiratory	Sore throat, fever, nausea	< 1	Lye
		24–72	*Strep. pyogenes*
Lower GI	Cramps, diarrhea	7–12	*B. cereus*
			Cl. perfringens
		13–72	*Campylobacter*
			Pathogenic *E. coli*
			Salmonella
			Shigella
			Norwalk virus
		> 72	Other viruses
			Entamoeba
			Giardia
Neurological	Vision problems, tingling, paralysis	< 1	Insecticides
		1–6	Ciguatera fish
		18–36	*Cl. botulinum*
		> 72	Mercury[c]
Generalized infection	Fever, chills, aches	> 72	*Listeria*
			Salmonella Typhi
			Hepatitis A
			Toxoplasma
Allergic	Facial flushing, itching	< 1	Scombroid fish
			Monosodium glutamate
		1–6	Hypervitaminosis A[d]

[a]exo-enterotoxigenic strains
[b]endo-enterotoxigenic strains
[c]grain fungicides; meat of animals fed contaminated grain
[d]due to consumption of liver and kidneys of animals from cold regions

Smart phone applications

The models to be described in the following sections used the global positioning system (GPS) data of smart phone users who had their phones set to share their location data (for navigation purposes, for example) to detect users' restaurant visits by combining the location data with official health department restaurant licensing records and Google Maps, Google Places API, etc. Each visit that is within 50 meters of a food venue is automatically "snapped" to the nearest one as determined by the Google application. "Cell phones determine their location through a combination of satellite GPS, WiFi access point fingerprinting, and cell-tower triangularization… Location accuracy typically ranges from 9 meters to 50 meters and is highest in areas with many cell towers and Wi-Fi access points. In such cases, even indoor localization (for example, within a mall) is accurate *(21)*."

TABLE 3. Classification of acute foodborne disease outbreaks by preparation method, significant ingredient, agent, and contributing factors, New York State, 1980–1991. Adapted from references *9* and *26*

% [(a)]	Significant ingredient	% [(b)]	Agent	% [(c)]	Contributing factors	% [(d)]
14%	**Eaten raw or lightly cooked**					
	Shellfish	95%	General viral	59%	Unapproved source	60%
			Norwalk virus	31	Eating raw meat	53
			Hepatitis A	2	Contaminated ingredient	49
6%	**Solid masses of potentially hazardous food**					
	Starchy food	58%	***B. cereus***	79%	Improper hot holding	42%
					Improper cooling	23
	Beef	21	*C. perfringens*	72	Improper cooling	50
					Improper reheating	39
	Egg	16	*Salmonella*	100	Inadequate cooking	100
					Contam. ingredient	49
5%	**Cook/serve foods**					
	Egg	31%	*Salmonella*	95%	Inadequate cooking	82%
	Poultry	21	*Salmonella*	33	Inadequate cooking	33
	Beef	20	*Salmonella*	36	Inadequate cooking	29
5%	**Natural toxicant**					
	Finfish	83%	Scombroid toxin	99%	Inadequate refrigeration	79%
5%	**Roasted meats and poultry**					
	Poultry	41%	***Salmonella***	52%	Inadequate cooking	41%
	Beef	41	*C. perfringens*	34	Improper hot hold	28
	Pork	16	*C. perfringens*	27	Improper cooling	45
1%	**Salads with one or more cooked ingredients**					
	Poultry	35%	***Salmonella***	50%	Inadequate refrigeration	50%
1%	**Liquid or semi-solid mixtures of potentially hazardous foods**					
	Poultry	33%	***Salmonella***	50%	Inadequate refrigeration	50%
1%	**Chemical contamination**					
	Beverages	42%	Heavy metals	58%	Added poison	33%

[a]Percentage of 1,528 reported outbreaks involving food with given method of preparation
[b]Percent of reported outbreaks for given significant ingredient in above category
[c]Percent of reported outbreaks for the specific agent in a significant ingredient category
[d]Percentage of outbreaks where specific contributing factor was reported in significant ingredient category
Smaller categories are not represented.

nEmesis and Foodborne Chicago using Twitter

In 2013, researchers at the University of Rochester in Rochester, New York, published the results of a project they called nEmesis, using Twitter to identify possible sporadic foodborne illness in New York City *(15, 19)*. At that time, Twitter was new enough that they felt it necessary to explain:

> Twitter is a widely used online social network and a particularly popular source of data for its real-time nature and open access… Twitter users post message updates (tweets) up to 140 characters long… [As of 2001] 13% of online adults use[d] Twitter, most of them daily and often via a [smart] phone… These mobile users often attach their current GPS location to each tweet… *(19)*.

The researchers detected users' restaurant visits preceding the onset of a suspected foodborne illness as already described. The users had to actively send at least one tweet from the location. The researchers used what they called human-guided machine learning to sort through each tweet for each user through the 72-hour period after a restaurant visit to identify "sick tweets" consistent with foodborne illness. The process involved hiring a panel of people to rate components of tweets, which were called "features," according to the correlation of the panel's rating of the features to the sender's having a condition consistent with a foodborne illness. *Table 4* shows their rating table.

Almost one-third of messages indicating foodborne illness could be traced to a restaurant. The "health score" calculated for each restaurant, based on the proportion of customers who got sick shortly after their visit, correlated well with worse official inspection scores ($r = 0.30$, $P = 0.0006$).

Next, nEmesis was deployed in Las Vegas, Nevada *(21)*. This time, nEmesis picked 71 restaurants to be inspected, and the standard protocol was used to pick a paired control restaurant matched in location, size, cuisine and permit type. Inspectors were not told which was which.

Results were similar. More control restaurants passed inspection with 0 or 1 demerits. Inspections scheduled via nEmesis resulted in significantly more demerits, 9 versus 6 per inspection ($P = 0.019$).

Foodborne Chicago, launched in 2013, also used Twitter, but used only the keyword "food poisoning" and its wildcard variants to identify tweets *(6)* and also used human-guided machine learning *(12)*. It differed from nEmesis in that health department staff contacted individuals who had sent "sick tweets" and directed them to fill out a form at the Chicago 311 non-emergency complaint system. However, complaints received via Twitter and complaints that people filed at the Foodborne Chicago site on their own were lumped together.

Of the 133 health inspections prompted by Foodborne Chicago between March 2013 and January 2014, 20.3% found at least 1 critical violation, compared to 16.4% of inspections prompted by complaints outside of Foodborne Chicago during that period. Of the people who filed Foodborne Chicago complaints, 9.8% reported seeking medical attention.

"Foodborne Chicago is currently offline and in the process of being updated" *(12)*. nEmesis, Foodborne Chicago, and FINDER (below) are all separate *(12)*.

As of 2014, Foodborne Chicago was being shared with Boston and New York City *(6)*.

FINDER

FINDER (Foodborne Illness Detector in Real Time) started as a project developed by Google and the Harvard T. H. Chan School of Public Health *(20)*. Unlike in the case of the Twitter applications, individuals did not have to do anything at a restaurant except have their smart phones set to share their location data, as already described. In this case, however, users remained anonymous, but the entire sequence of locations each one visited during the 3 days prior to the user performing a Google search of web pages about foodborne illness were included; such pages included sources such as Wikipedia articles about foodborne illness and the CDC website on foodborne illness. FINDER classified the web searchers as "sick" or "not sick" by using what they refer to as their "web search model," the content of which is not revealed in detail in published reports.

FINDER operated in Las Vegas, Nevada between May and August 2016 and in Chicago between November 2016 and March 2017. The actual number of individuals tracked in this way is not known. However, over 15,000 Google searches were rated for the probability that their query was related to foodborne illness.

> Every morning, each city was provided with a list of restaurants in their jurisdiction that were automatically identified by FINDER. The health department in each city would then dispatch inspectors (who were unaware of whether or not the inspection was prompted by FINDER) to conduct inspections at those restaurants to determine whether there were health code violations. In addition to FINDER-initiated inspections, the health departments continued with their usual inspection protocols.

The result was that 52.3% of restaurants identified by FINDER had serious health code violations, compared with 22.7% for other restaurants. Also, because FINDER aggregates data from numerous individuals who ate at the same location, the researchers were able to determine that the restaurant most likely to have caused a person's illness was the one they visited most recently only 62% of the time; 19.4% of the time it was the 2nd most recently visited; 11.5% of the time it was the 3rd most recently visited, and 7.2% of the time it was the 4th or even an earlier restaurant, based on the "relative signal strength" for each restaurant. Finally, it was found that the odds of being identified as unsafe by FINDER were higher in restaurants with lower a priori risk levels assigned at licensing.

TABLE 4. Top 20 most significant negatively and positively weighted features related to foodborne illness as used by nEmesis. Adapted from reference *20*

Positive features		Negative features	
Feature	Weight	Feature	Weight
stomach	1.7633	think I'm sick	-0.8411
stomachache	1.2447	I feel sooo	-0.7156
nausea	1.0935	f*** I'm	-0.6393
tummy	1.0718	@mention sick to	-0.6212
#upsetstomach	0.9423	sick of being	-0.6022
nauseated	0.8702	ughhh cramps	-0.5909
upset	0.8213	cramp	-0.5867
nautious [sic]	0.7024	so sick omg	-0.5749
ache	0.7006	tired of	-0.5410
being sick man	0.6859	cold	-0.5122
diarrhea	0.6789	burn sucks	-0.5085
vomit	0.6719	course I'm sick	-0.5014
@mention I'm getting	0.6424	if I'm	-0.4988
#tummyache	0.6422	is sick	-0.4934
#stomachache	0.6408	so sick and	-4907
i've never been	0.6353	omg I am	-0.462
threw up	0.6291	@link	-0.4744
i'm sick great	0.6204	@mention sick	-0.4704
poisoning	0.5879	if	-0.4695
feel better tomorrow	0.5463	I feel better	-0.4670

CONCLUSION AND RECOMMENDATIONS

It is not clear from available published descriptions what kinds of inspections cities used in restaurants identified through nEmesis, Foodborne Chicago and FINDER. Apparently, they were all standard inspections. It may be possible to improve them by implementing the following suggestions:

TRY TO ESTIMATE INCUBATION PERIODS FOR NEMESIS AND FINDER

Because Foodborne Chicago links Twitter reporting of symptoms to contact with health department staff, it appears to be possible to ask respondents the time of onset of symptoms and compare that to the time of the initial tweet. This may provide estimates of incubation periods in nEmesis, because it might indicate the lag between getting sick and tweeting about it. Perhaps the estimates could be used to calibrate the FINDER results as well: the lag between onset of symptoms and the sending of tweets may approximate that between onset and web searching about symptoms.

TRY TO DETERMINE SYMPTOMS

Foodborne Chicago staff could also ask about symptoms while they communicate with cases.

Researchers employing nEmesis used the "features" listed in *Table 4* to classify the senders as cases; this table also seems to identify symptoms.

The "web search model" in FINDER" might also suggest symptoms. Obviously, if cases searched for "diarrhea," it's not hard to imagine what they might be going through.

INVESTIGATE CASES AND CONTROLS WITH METHODS NORMALLY USED TO INVESTIGATE FOODBORNE ILLNESS OUTBREAKS

The methods suggested for the surveillance system for potential sporadic cases of foodborne disease from restaurants among users of the weight-loss website would be better than standard inspections for cases and controls identified using the newer methods. *Tables 2 and 3* would help focus the inspections.

Knowing the symptoms and incubation periods would point to likely foods and mishandling errors, which would help focus corrective action and enforcement.

USE FINDER TO IDENTIFY CONTROLS

Restaurants that multiple people visited without ever following up with web searches related to foodborne illness might make better controls than randomly selected restaurants.

Investigating control restaurants paired with case restaurants, as has been described for nEmesis, might finally begin to resolve one of the questions the surveillance system using the weight-loss website was hoping to answer: What is the actual risk of illness from a specific food code violation? The control restaurants would have to be inspected anyway. Some of them might make errors similar to the case restaurants, but not cause illness. This scenario might set up an enforcement quandary, but discovering what kept the control restaurant's errors from causing illness might also suggest new control measures.

CONSIDER USING CART OR AN ANALOG WITH CHICAGO'S FOOD SAFETY INSPECTION FORECASTING MODEL

Unlike the Model, CART updates its predictions with each run.

DEVELOP A CLEARINGHOUSE LISTING USERS OF THESE NEW METHODS

There seems to be no sort of central repository for information about application of these methods. Developers in Chicago have endeavored to make their models available to others by publishing their code. When other locations using their methods are known, they could be identified if a repository were available. Not only is it unclear where these models are in use, it is also sometimes unclear when use of a model was started and stopped. It may be helpful for various health departments to collaborate as described in this section.

ENCOURAGE REVIEW WEBSITES TO INCLUDE A LINK TO THE REVIEWER'S LOCAL HEALTH DEPARTMENT'S REPORTING SYSTEM

CONSIDER PRIVACY ISSUES

If FINDER were to detect an individual who visits the same restaurant on multiple consecutive days and gets identified as sick during the period of being tracked, the person may be a sick food employee. In that case, it may be appropriate to break the individual's anonymity and notify the restaurant.

There may be other issues. Regarding FINDER, for example, anonymity is assured by following the "Google Privacy Policy and Terms of Service" *(20)* and other procedures. However, a recent article in *The New York Times (16)* suggests that "de-identifying" people or substituting fake values may not be adequate: "In 2006, individuals were identified from the web-browsing histories of three million Germans, data that had been purchased from a vendor."

ACKNOWLEDGMENTS

The author acknowledges the work of Kathyrn Madden, R.S., to scan Reference 8 to prepare a digital copy.

***Author for correspondence:**
Phone: +1 614.589.8994
Email: jimh117@earthlink.net

REFERENCES

1. Breiman, L., J. H. Friedman, R. A. Olshen, and C. J. Stone. 1984. Classification and Regression Trees. Wadsworth, Pacific Grove, CA.
2. Briley, R. T., and E. F. Klaus. 1985. Using risk assessment as a method of determining inspection frequencies. *Dairy Food Sanit.* 5:468–474.
3. Bryan, F. L. 1992. Foodborne illness risk assessment of food-service establishments in a community. *J. Food Prot.* 45:93–100.
4. Bryan, F. L., C. A. Bartleson, O. D. Cook, J. J. Guzewich, D. Maxson, R. C. Swanson, E. C. D. Todd, and L. Wisniewski. 1999. *Procedures to Investigate Foodborne Illness – 5th Edition.* International Association of Milk, Food and Environmental Sanitarians, Inc. (IAMFES). Des Moines, IA.
5. Emanuel, R. 2012. Open data executive order, Pub. L. No. (No. 2012-2). Available at https://www.cityofchicago.org/city/en/narr/foia/open_data_executiveorder.html. Accessed August 5, 2019.
6. Harris, J. K., R. Mansour, B. Choucair, J. Olson, C. Nissen, and J. Bhatt. 2014. Health department use of social media to identify foodborne illness — Chicago, Illinois, 2013–2014. *Mort. Morb. Wkly. Rep.* 63:681–685. Available at https://www.cdc.gov/ mmwr/ preview/ mmwrhtml/mm6332a1.htm. Accessed 17 July 2019.
7. Harrison, C., M. Jorder, H. Stern, F. Stavinsky, V. Reddy, H. Hanson, H. Waechter, L. Lowe, L. Gravano, and S. Balter. 2014. Using online reviews by restaurant patrons to identify unreported cases of foodborne illness — New York City, 2012–2013. *Mort. Morb. Wkly. Rep.* 63(20):441–445. Available at https://www.cdc.gov/MMWr/preview/ mmwrhtml/mm6320a1. htm. Accessed 28 September 2019.
8. Hartman, J. 1992. "A computer model to identify risk markers for food service sanitation violations associated with outbreaks of foodborne illness." Unpublished master's thesis, The Ohio State University. Available at: http://www.jimh117.com. Accessed 22 August 2019.
9. Hartman, J. 2007. Web-based system for monitoring diet and physical activity and for active surveillance of enteric diseases. Available at http://www.jimh117.com. Accessed 22 August 2019.
10. Hartman, J. 2009. A call for a new surveillance system for sporadic foodborne illness using a weight-loss website. *Food Prot. Trends* 29:7:472, 465. Available at https://www.foodprotection.org/publications/food-protection-trends/archive/2009-07-a-call-for-a-new-surveillance-system-for-sporadic-foodborne-illness-using-a-weight-loss-web-/. Accessed 30 July 2019.
11. Irwin, K., J. Ballard, J. Grendon, and J. Kobayashi. 1989. Results of routine restaurant inspections can predict outbreaks of foodborne illness: The Seattle-King County experience. *Amer. J. Publ. Hlth.* 79:586–590. Available at https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1349498/. Accessed 1 July 2019.
12. Ivanova, E. 19 July 2019. Re: article, Elena. Ivanova@cityofchicago.org. Available from the author at jimh117@earthlink.net.

13. Kang, J. S., P. Kuznetsova, M. Luca, and Y. Choi. 18–21 October 2013. Where not to eat? Improving public policy by predicting hygiene inspections using online reviews. Proceedings of the Conference on Empirical Methods in Natural Language Processing, Seattle, WA, 22:58–77. Available at https://www.aclweb.org/anthology/D13-1150. Accessed 18 September 2019.

14. Kaplan, O. B., and A. El-Ahraf. 1979. Relative risk ratios of foodborne illness in food service establishments: an aid in deployment of environmental health manpower. *J. Food Prot.* 42:446–447.

15. Kautz, H. 23 June 2013. "There's a fly in my tweets." *New York Times* (New York Edition) p. SR12. Available at: https://www.nytimes.com/2013/06/23/opinion/ sunday/theres-a-fly-in-my-tweets.html?_r=0. Accessed 19 June 2019.

16. Kolata, G. 23 July 2019. Can data be fully anonymous? New algorithm can still identify you. T*he New York Times International.* Available at https://www.nytimes.com/2019/07/ 23/health/data-privacy-protection.html. Accessed August 19, 2019.

17. McBride, K., G. Aavik, M. Toots, T. Kalvet, and R. Krimmer. 2018. How does open government data driven co-creation occur? Six factors and a 'perfect storm'; insights from Chicago's food inspection forecasting model. *Government Information Quarterly* 36:88–97. Available at https://doi.org/10.1016/j.giq.2018.11.006. Accessed August 5, 2019.

18. Oldroyd, R. A., M. A. Morris, and M. Birkin. 2018. Identifying methods for monitoring foodborne illness: review of existing public health surveillance techniques. *JMIR Publ. Hlth. Surveill.* 2018. 4(2):e57. Available at https://publichealth.jmir.org/2018/2/e57. Accessed 18 September 2019.

19. Sadilek A., S. Brennan, H. Kautz, and V. Silenzio. 2013. nEmesis: Which restaurants should you avoid today? p. 138–146. *In* Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing, AAAI Press, Palo Alto, CA. Available at https://www.aaai.org/ocs/index.php/HCOMP/HCOMP13/paper/viewFile/7475/7413. Accessed August 8, 2019.

20. Sadilek, A., S. Caty, L. DiPrete, R. Mansour, T. Schenk Jr., M. Bergtholdt, A. Jha, P. Ramaswami, and E. Gabrilovich. 2018. Machine-learned epidemiology: real-time detection of foodborne illness at scale. *NPJ Digital Med.* 36:1–7. Available at: https://www.nature.com/articles/s41746-018-0045-1. Accessed 17 July 2019.

21. Sadilek, A., H. Kautz, L. DiPrete, B. Labus, E. Portman, J. Teitel, and V. Silenzio. 2017. Deploying nEmesis: preventing foodborne illness by data mining social media. *AI Magazine* 38:37–48. Available at https://aaai.org/ojs/index.php/aimagazine/ article/view/2711. Accessed 17 July 2019.

22. Schenk, T. 2017. Food inspection forecasting — City of Chicago. Available at https://chicago.github.io/food-inspections-evaluation/. Accessed August 7, 2019.

23. Spector, J. 2016. Chicago is predicting food safety violations. Why aren't other cities? *CityLab.* Available at https://www.citylab.com/solutions/2016/01/chicago-is-predicting-food-safety-violations-why-arent-other-cities/422511/. Accessed August 7, 2019.

24. U.S. Department of Health, Education, and Welfare. 1978. Food service sanitation manual, including a model food service sanitation ordinance. Washington, D.C.: U.S. Government Printing Office.

25. U.S. Food and Drug Administration. 2017. Food Code. U.S. Department of Health and Human Services. Available at https://www.fda.gov/media/110822/download. Accessed 1 July 2019.

26. Weingold, S. E., J. J. Guzewich, and J. K. Fudala. 1994. Use of foodborne disease data for HACCP risk assessment. *J. Food Prot.* 57:820–830.

27. Wodi, B. E., and R. A. Mill. 1985. A priority system model for sanitation management in food service establishments. *Amer. J. Publ. Hlth.* 75:1398–1401.

28. Woodward, C. 16 April 2019. Analog and digital. Available at: https://www.explainthatstuff.com/analog-and-digital.html. Accessed August 10, 2019.